

VLT/CRIRES+ observations of warm Neptune WASP-107 b

Challenges in detecting molecules with ground-based transmission spectroscopy of cooler and cloudy exoplanets

L. Boldt-Christmas^{1,*}, A. D. Rains^{1,2}, N. Piskunov¹, L. Nortmann³, F. Lesjak^{3,4}, D. Cont^{5,6},
O. Kochukhov¹, A. Hahlin¹, A. Lavail⁷, T. Marquart¹, U. Heiter¹, M. Rengel⁸, D. Shulyak⁹, F. Yan¹⁰,
A. Hatzes¹¹, E. Nagel³, A. Reiners³, and U. Seemann¹²

¹ Observational Astrophysics, Department of Physics and Astronomy, Uppsala University, Sweden

² Instituto de Astrofísica, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, 782-0436 Macul, Santiago, Chile

³ Institut für Astrophysik und Geophysik, Georg-August-Universität, Friedrich-Hund-Platz 1, 37077 Göttingen, Germany

⁴ Leibniz Institute for Astrophysics Potsdam (AIP), An der Sternwarte 16, 14482 Potsdam, Germany

⁵ Universitäts-Sternwarte, Ludwig-Maximilians-Universität München, Scheinerstrasse 1, 81679 München, Germany

⁶ Exzellenzcluster Origins, Boltzmannstrasse 2, 85748 Garching bei München, Germany

⁷ Namzitu Astro, 31130 Quint-Fonsegrives, France

⁸ Max-Planck-Institut für Sonnensystemforschung, Justus-von-Liebig-Weg 3, 37077 Göttingen, Germany

⁹ Instituto de Astrofísica de Andalucía – CSIC, c/ Glorieta de la Astronomía s/n, 18008 Granada, Spain

¹⁰ Department of Astronomy, University of Science and Technology of China, Hefei 230026, PR China

¹¹ Thüringer Landessternwarte Tautenburg, Sternwarte 5, 07778 Tautenburg, Germany

¹² European Southern Observatory, Karl-Schwarzschild-Str. 2, 85748 Garching bei München, Germany

Received 31 July 2025 / Accepted 24 April 2026

ABSTRACT

Context. The atmospheres of transiting exoplanets can be studied spectroscopically using space-based or ground-based observations. Each has its own set of strengths and weaknesses, so there are benefits to both approaches. This is especially true for more challenging targets such as cooler, smaller exoplanets whose atmospheres most likely contain many molecular species and cloud decks.

Aims. We aim to study the atmosphere of the warm Neptune-like exoplanet WASP-107 b ($T_{\text{eq}} \approx 740$ K). Several molecular species have been detected in this exoplanet in studies using the space-based JWST, and we aim to confirm and expand upon these detections by using the ground-based VLT and evaluating how well our findings agree with previously retrieved atmospheric parameters.

Methods. We observed two transits of WASP-107 b with VLT/CRIRES+ and created cross-correlation templates of the target atmosphere based on results from previous studies. We created different templates to investigate the impact of varying volume mixing ratios of species and the inclusion or exclusion of clouds. Considering this target's observational challenges, we created simulated observations prior to evaluating our real data in order to assess our expected detection significances with the cross-correlation technique.

Results. We report cross-correlation signals of two molecular species, CO ($\sim 6\sigma$) and H₂O ($\sim 4.5\sigma$), in our $K_p - v$ maps. This supports previous space-based detections and demonstrates, for the first time, the capability of VLT/CRIRES+ to detect species in targets cooler than hot Jupiters using transmission spectroscopy. We show that our analysis is sensitive to the inclusion of clouds but less so to different volume mixing ratios. Interestingly, our signal maximum deviates notably but consistently from its expected location in the K_p direction of our maps, and we speculate on the possible reasons for this effect. We demonstrate that the error budget for these relatively cool exoplanets is severely reduced in comparison to hotter exoplanets and emphasise the need for further work in the context of high-resolution spectroscopy.

Key words. methods: observational – methods: statistical – techniques: spectroscopic – planets and satellites: atmospheres – planets and satellites: composition

1. Introduction

As an exoplanet passes in front of its host star, its chemical composition may be studied by observing the fraction of starlight that is transmitted through the planet's upper atmosphere as a function of wavelength. This simple premise is what enables the field of transmission spectroscopy, and it has facilitated the growth of exoplanetary characterisation from a hypothetical possibility into a thriving interdisciplinary research topic. Today, out of the nearly 6000 exoplanets¹ that have been confirmed, over

250 planets have been chemically characterised by the methods of transmission, emission, and/or reflectance spectroscopy².

Spectroscopic studies of exoplanetary atmospheres can be broadly divided into two different branches: space-based observations of high photometric precision but lower spectral resolution and ground-based observations of higher spectral resolution that suffer from Earth's atmospheric effects. Each of these two avenues come with their own benefits and drawbacks, and as such, it is difficult to argue that either approach is superior to the other. Ground-based spectrographs can generally be of higher

* Corresponding author: linn.boldt-christmas@physics.uu.se

¹ NASA Exoplanet Archive – <https://exoplanetarchive.ipac.caltech.edu/>

² IAC ExoAtmospheres database – <https://research.iac.es/proyecto/exoatmospheres/>

resolving power, R , enabling them to measure strong contrasts between the cores and the wings of individual spectral lines. Meanwhile, instruments of lower R can only detect the combined effects of groups of lines that result in an increase or decrease (emission or absorption) of local flux. Observations of high resolution therefore provide spectra of greater detail and allow for robust identification of chemical species. In recent years, ground-based transmission spectroscopy has been crucial for the study of detailed atmospheric effects such as winds, jet streams, and vertical stratification (e.g. [Ehrenreich et al. 2020](#); [Seidel et al. 2020](#); [Gandhi et al. 2022](#); [Prinoth et al. 2022](#); [Lesjak et al. 2023](#); [Cont et al. 2024](#); [Nortmann et al. 2025](#); [Seidel et al. 2025](#)).

In contrast, instruments of higher spectral resolution currently remain prohibitively complex to use in space. For space-based telescopes such as JWST, the onboard spectrographs have a much lower spectral resolution ($R \lesssim 3000$), meaning such instruments are not able to resolve individual spectral lines. Nonetheless, space-based telescopes and their spectrographs maintain a number of notable benefits, including broader wavelength coverage as well as a lack of turbulent distortions, telluric absorption and emission, and daytime interruptions. The field of space-based atmospheric characterisation is more active than ever, and a small selection of recent successes include [JWST Transiting Exoplanet Community Early Release Science Team \(2023\)](#), [Tsai et al. \(2023\)](#), [Smith et al. \(2024\)](#), [Mukherjee et al. \(2025\)](#), [Teske et al. \(2025\)](#), and many more.

Thanks to these fundamental differences, space-based and ground-based transmission spectroscopy bring forth distinct but complementary contributions to the general research goal of exoplanet atmospheric characterisation. In the pursuit of recovering an atmospheric transmission spectrum of a planet and deriving as many details as possible about its physical conditions, both types of observations fulfil an important role that cannot be achieved by the other. This becomes increasingly important as the exoplanet community continues to move further towards the study of planets that are cooler and more chemically complex than the evergreen case studies of hot and ultra-hot Jupiters. With benefits from both types of data, in an ideal situation, one could capitalise on this by observing the same target with both a low-resolution space-based instrument as well as a high-resolution ground-based instrument in the same wavelength regime, ensuring that one may support the detections of the other.

It is this ideal situation, and the comparison of conclusions derived from the two observing techniques, that is explored in this paper. We analyse transit observations from two different nights of the exoplanet WASP-107 b ([Anderson et al. 2017](#)), observed using the ground-based Cryogenic Infrared Echelle Spectrograph (CRIRES⁺) on the Very Large Telescope (VLT) of the Paranal Observatory, and we compare these high-resolution observations of $R \approx 140\,000$ covering K-band wavelengths $\sim 2.0\text{--}2.5\ \mu\text{m}$ with the analysis results of space-based observations of the same target from a growing gallery of JWST observations. Previous studies have already yielded a number of molecular detections in WASP-107 b, and this work has been able to confirm some of these results at a wavelength range that no previous ground-based studies of this target have explored.

We use the high-resolution cross-correlation spectroscopy (HRCCS) technique to report the signals of two molecular species, CO at $\sim 6\sigma$ and H₂O at $\sim 4.5\sigma$, and that HRCCS using a multi-species template of the atmosphere as retrieved by a previous JWST study ([Welbanks et al. 2024](#)) produces a $\sim 6\sigma$ cross-correlation peak. Considering only ground-based

transmission spectroscopy studies, these detections are the first ever to be made for a target of $T_{\text{eq}} < 800\ \text{K}$, highlighting both the capability of instruments such as VLT/CRIRES⁺ and the challenges that will face observers as they proceed towards studying targets of even lower temperatures.

As the exoplanet community continues to reap the fruits of JWST data, the number of exoplanetary atmosphere studies using space-based observations will continue to grow. At the same time, the Extremely Large Telescope (ELT) nears completion, and it will usher in a new age for ground-based observations very soon. In this context, there is a need for assessing the extent to which the two types of observations can be combined, as whatever parameters we obtain from either type of data and incorporate into our HRCCS template will affect our final result. To explore this, we used simulated observations together with real VLT/CRIRES⁺ data to further explore (i) how sensitive our HRCCS detection significance is to the changes in parameters from space-based retrievals by cross-correlating templates derived from two different JWST retrievals and (ii) how the uncertainties in other parameters of our system can affect our results, exploring possible effects of cooler, smaller, and cloudier planets such as WASP-107 b.

Considering these properties means that a system such as that of WASP-107 b becomes more challenging to analyse using HRCCS and that there are many instances where the properties of that analysis are not fully understood or where the independence between conflating effects cannot be fully justified. As a result, this paper is above all an attempt to study the consequences of complexity and cross-talk between various components of the problem for the case of a relatively cool target than a nominal hot Jupiter based on good quality data and multiple transits. For cooler planets, the interplay between the dominant features in the transmission spectrum, the underlying stellar (esp. cool dwarf) spectrum, and the telluric contamination becomes a very difficult task for more simplistic algorithms, and likewise subtler effects such as the presence of clouds in the planetary atmosphere or minor errors in the wavelength solution may result in a major reduction in detection significance due to appearance of ‘systematic features’ in the $K_p - v$ maps. Thus the overarching question of this paper is to what extent may we fully account for artefacts in our cross-correlation maps and bring the characterisation of cooler exoplanets such as WASP-107 b to the same level of significance that was achieved for hotter Jupiters with our traditional methods. This paper uses the traditional HRCCS methodology, but as we demonstrate, non-traditional targets clearly require non-traditional methodology, which is yet to be fully developed. As such, our hope is that this paper will serve more as an exploration of such a methodology rather than a reporting of chemical detections and/or non-detections only using the tools available to us today.

In Sect. 2, we provide context for our study by exploring previous atmospheric studies of this target that have been conducted since its discovery. In Sect. 3, we provide the details of how our observational data were obtained and prepared for interpretation. In Sect. 4, we describe our methods for removing the stellar and telluric signal and our cross-correlation analysis. In Sect. 5, we use simulated observations to establish some context, explore the anticipated impact of certain parameters, and set expectations for what can be considered a reasonable detection significance for this target and instrument. In Sect. 6, we present the analysis of our VLT/CRIRES⁺ observations and investigate possible interpretations of the $K_p - v$ plots. Our results are presented and discussed in Sect. 7 before the final conclusions in Sect. 8.

Table 1. Host star and planetary parameters for WASP-107 b.

Star (WASP-107)		
Parameter	Symbol/units	Value
Mass ^a	$M_*(M_\odot)$	$0.683^{+0.017}_{-0.016}$
Radius ^a	$R_*(R_\odot)$	0.67 ± 0.02
Age ^a	(Gyr)	3.4 ± 0.7
Distance ^b	(pc)	64.401 ± 0.1078
Spectral type ^c		K7V
Luminosity ^a	$L_*(L_\odot)$	0.132 ± 0.003
Magnitude (K band) ^d	m_{K_s}	8.637 ± 0.023
Effective temperature ^a	$T_{\text{eff}} \text{ (K)}$	4425 ± 70
Metallicity ^a	(Fe/H)	0.02 ± 0.09
Surface gravity ($\log g$) ^a	$\log_{10} \text{ (cm/s}^2\text{)}$	4.633 ± 0.012
Carbon-to-oxygen ratio ^e	C/O	0.50 ± 0.10
$v \sin i^f$	(km/s)	$0.507^{+0.072}_{-0.086}$
Planet (WASP-107 b)		
Parameter	Symbol/units	Value
Mass ^a	$M_p (M_J)$	0.0960 ± 0.0050
Radius ^g	$R_p (R_J)$	0.94 ± 0.02
Density ^a	$\rho \text{ (g/cm}^3\text{)}$	$0.134^{+0.015}_{-0.013}$
Eq. temperature ^h	$T_{\text{eq}} \text{ (K)}$	736 ± 17
Int. temperature ⁱ	$T_{\text{int}} \text{ (K)}$	460 ± 40
Orbital period ^a	$P \text{ (days)}$	5.7214742
Orbital eccentricity ^a	e	0.06 ± 0.04
Orbital inclination ^g	$i \text{ (deg)}$	89.56 ± 0.08
Sky-projected inclination ^j	$ \lambda \text{ (deg)}$	118^{+38}_{-19}
Planet RV semi-amplitude ^k	$K_p \text{ (km/s)}$	105.2 ± 2.5
Argument of periastron ^l	$\omega \text{ (deg)}$	-2.3 ± 6.1
Semi-major axis ^m	$a \text{ (au)}$	0.055 ± 0.001
Transit depth ^m	$\delta \text{ (%)}$	2.17 ± 0.02
Transit duration ^f	$t_{14} \text{ (hours)}$	2.7528 ± 0.0072

Notes. ^aPiaulet et al. (2021), ^bGaia EDR3 (Gaia Collaboration 2021), ^cDressing et al. (2019), ^d2MASS All-Sky Catalog (Cutri et al. 2003), ^eHejazi et al. (2023), ^fBourrier et al. (2023), ^gKokori et al. (2023), ^hMočnik et al. (2017), ⁱSing et al. (2024), ^jRubenzahl et al. (2021), ^kGuilluy et al. (2024), ^lMurphy et al. (2024), ^mAnderson et al. (2017).

2. Previous studies

2.1. Discovery and earlier observational studies

WASP-107 b is a warm ($T_{\text{eq}} < 800 \text{ K}$) Neptune-like exoplanet with a radius of $0.924 R_J$ and mass of $0.096 M_J$ (Močnik et al. 2017; Piaulet et al. 2021). Its host star, WASP-107, is a cool dwarf of solar metallicity located 65 pc away in the constellation Virgo. It is of spectral type K7V with an effective temperature of 4358 K, a radius of $0.656 R_\odot$, and a mass of $0.696 M_\odot$ (Anderson et al. 2017; Dressing et al. 2019) with near-solar abundances and a carbon-to-oxygen ratio (C/O) of 0.50, compared to solar C/O = 0.54 (Hejazi et al. 2023). For a full overview of the stellar and planetary parameters of this system, see Table 1.

The matter of this planet's atmospheric conditions was identified as a point of possible scientific interest already at the time of its discovery by Anderson et al. (2017). In that work, its mass was established to be $0.12 M_J$, which placed the planet in a transition region between gas giants where the planetary mass is less than half that of Saturn ($0.5 M_{\text{H}} = 0.15 M_J$) but above twice that of Neptune ($2 M_{\text{N}} = 0.11 M_J$). This is a regime where our Solar System has no analogues, especially not at short

orbital distances, so speculations about the possible climate of WASP-107 b were difficult to justify at that time. It was noted that the planet could be characterised more accurately if one was to determine, for example, its atmospheric metallicity, considering this is a parameter that varies between the ice giants (higher metallicity) and gas giants (lower) in our Solar System.

Due to the potential scientific value of studying its atmosphere and the observational suitability of the target, WASP-107 b was suggested in the discovery paper as an excellent candidate for future transmission spectroscopy studies of its atmospheric composition. The observational advantages of WASP-107 b are twofold: (i) its large atmospheric scale height, and (ii) its small, bright host star. WASP-107 b has a notably large radius for its mass, and was found in Piaulet et al. (2021) to have an even lower density (0.134 g/cm^3) than previously measured, placing its mass at $30.5 M_\oplus$ or $0.096 M_J$ (which makes it even more exotic to our Solar System, at $\sim 10\%$ of Jupiter's mass and $\sim 90\%$ of its radius). As such, WASP-107 b is a very 'puffy' planet with a significantly inflated atmosphere – i.e. larger scale height – which is beneficial for transit studies. Furthermore, the host star is particularly bright in the infrared (K-band magnitude of 8.637 mag; see Table 1), making it highly suitable for obtaining high signal-to-noise (S/N) observations in the infrared and near-infrared regime, and thus for atmospheric spectroscopy studies as molecular lines appear at these wavelengths.

The first two atmospheric studies of WASP-107 b were published in May 2018, both using space-based observational data of transits from the *Hubble* Space Telescope (HST). The work by Kreidberg et al. (2018) used observations with HST/WFC3 (with the G141 grism, which covers the wavelength range of 1.1–1.7 μm) of a single WASP-107 b transit to detect water in its upper atmosphere through atmospheric retrievals. This work constrained atmospheric metallicity to an upper limit of $\times 30$ solar metallicity, and noted a depletion of methane. The work by Spake et al. (2018) also used HST/WFC3 observations (this time with the G102 grism of wavelength range 0.8–1.1 μm) to detect a post-transit tail at 10 833 Å (He I triplet), suggesting that this planet's atmosphere was highly extended. In this work, the authors also noted the possibility of studying this tail via observations of this triplet using high-resolution spectrographs in the infrared, as had been done at the time for, for example, WASP-69 b by Nortmann et al. (2018) and HAT-P-11 b by Allart et al. (2018) using the CARMENES instrument ($R \approx 80\,000$) on the 3.5 m telescope at the Calar Alto Observatory.

The first ground-based, high resolution study of WASP-107 b was published by Allart et al. (2019), which confirmed the target's extended helium atmosphere using Calar Alto Observatory/CARMENES to confirm the space-based HST/WFC3 findings whose helium feature was poorly resolved. The data obtained during these observations were also used in a study by Kesseli et al. (2020) who used archival CARMENES data to search for the molecule FeH in a number of exoplanetary atmospheres, including that of WASP-107 b, but no FeH was found in this target. Two subsequent follow-up surveys were made using more ground-based observations, both studies using Keck II/NIRSPEC. Kirk et al. (2020) confirmed that the signature of the escaping helium saw no significant temporal variation in the two years since the previous result, while Spake et al. (2021) managed to obtain significant post-transit coverage that confirmed the tail's length to be the equivalent of ~ 7 planet radii, corresponding to approximately twice the planet's Roche lobe radius. More recently, studies have confirmed and further characterised WASP-107 b's escaping helium tail with both ground-based observations – namely Guilluy et al. (2024),

this time using the GIARPS (GIANO-B + HARPS-N) observing mode of the Telescopio Nazionale *Galileo* – and also with a very recent space-based observational study by JWST that included pre-transit coverage (Krishnamurthy et al. 2025).

In 2021, a radial velocity (RV) study by Piaulet et al. (2021) using an extensive dataset from Keck I/HIRES and archival data (overall spanning observations from 2011–2020) confirmed that the density of WASP-107 b was significantly lower than previously measured and also the detection of another more massive planet in the WASP-107 system (WASP-107 c) at a much longer, eccentric orbit. Also using Keck I/HIRES, following up TESS data as part of the TESS-Keck survey collaboration, a study by Rubenzahl et al. (2021) measured the misalignment of WASP-107 b’s orbit, i.e. its obliquity, based on observations of the target’s Rossiter-McLaughlin effect. A polar/retrograde orbit had already been suspected by Dai & Winn (2017) who constrained the likely obliquity to be in the range of 40–140° based on the finding of a lower number of starspot-crossings by WASP-107 b in *K2* data than expected. The TESS-Keck study successfully confirmed this anticipated obliquity, with a final calculation of sky-projected inclination being $|\lambda| = 118^\circ$. This result was further supported by Bourrier et al. (2023) who confirmed its polar, retrograde orbit by analysis of its Rossiter-McLaughlin effect, and was complemented by Dholakia et al. (2025) who placed an upper bound on the planet’s oblateness of $f < 0.23$.

2.2. Transmission studies using JWST

Considering the diversity of these previous studies, the benefit of observing a target using both space-based and ground-based telescopes is clear for many science cases. By repeatedly studying the same target across overlapping wavelength ranges with both types of observations, one hopes to reach more robust conclusions about physical properties of the atmosphere in question.

While all the transmission studies listed so far of WASP-107 b have remained within the wavelength range between ~ 0.8 – $1.7 \mu\text{m}$, more studies featuring space-based results further into the infrared have started to be published in the past few years. In particular, JWST has been able to study WASP-107 b using transmission spectroscopy across a redder wavelength range that had not previously been probed for this target. Dyrek et al. (2024) observed WASP-107 b with JWST on 19–20 January 2023 under the JWST MIRI GTO programme. The transit was observed using the MIRI spectrometer, which provides a spectral resolution of 30 to 100 across 4.61 – $11.83 \mu\text{m}$ (mid-infrared). Together with archival HST data from Kreidberg et al. (2018) of wavelength region 1.121 – $1.629 \mu\text{m}$, the authors of this work performed atmospheric retrievals whose best fits to these datasets resulted in detections of SO_2 ($\sim 9\sigma$), H_2O ($\sim 12\sigma$), and silicate clouds ($\sim 7\sigma$) alongside a non-detection of CH_4 and tentative detections of H_2S ($\sim 4\sigma$), NH_3 (~ 2 – 3σ), and CO (~ 2 – 3σ). From these retrievals, the team also obtained volume mixing ratios (VMRs) for which SO_2 came out several orders of magnitude higher than equilibrium chemistry predicts. The interpretation advanced in their paper is that the atmosphere of the planet must therefore be in chemical disequilibrium, caused by photochemical reactions through (i) photodissociation of H_2O in the uppermost atmospheric layers, generating atomic H and OH radicals that in turn create SO_2 by oxidising H_2S , and (ii) further introduction of more OH radicals through photodissociation of other molecules beyond H_2O in the lower atmospheric layers.

Another recent publication of this target also used a combination of HST and JWST observation to further confirm and

expand upon these detections. Welbanks et al. (2024) observed WASP-107 b over two transits on 14 January 2023 and 4 July 2023 as part of the MANATEE NIRC+MIRI GTO program using NIRC+M F322W2 covering 2.4 – $4.0 \mu\text{m}$ and F444W covering 3.9 – $5.0 \mu\text{m}$. Combining these datasets with the MIRI results of Dyrek et al. (2024) plus the archival HST WFC3 data from Kreidberg et al. (2018) and Spake et al. (2018) resulted in an impressive sequence of spectra covering a total wavelength region of 0.8 – $1.7 \mu\text{m}$ and 2.4 – $12.2 \mu\text{m}$. In their atmospheric retrieval, the best fit confirmed the Dyrek et al. (2024) detections (CO_2 at 27σ , H_2O at 18σ , and SO_2 at 8σ), as well as two tentative detections of CO and NH_3 (both at 5σ). This work also provided the first detection of CH_4 (at 8σ), found at wavelengths that had not been previously studied for this target, i.e. ~ 3.2 – $3.8 \mu\text{m}$. Furthermore, this work managed to constrain internal temperatures to $>345 \text{ K}$, suggesting that the notable inflation of this exoplanet may be explained by a Neptune-like internal structure of tidally induced heating – speculated to possibly be due to the planet’s non-circular orbit of $e = 0.06 \pm 0.04$ (Piaulet et al. 2021).

Other studies have recently been published using transmission spectra from JWST of this target to study two very relevant atmospheric factors: core interior and limb asymmetry. The first study by Sing et al. (2024) analysed a JWST-NIRSpec transmission spectrum using the G395H grating, which provides a wavelength range of 2.7 – $5.18 \mu\text{m}$, and again detected SO_2 , CH_4 , H_2O , CO_2 , and CO through retrievals. Once abundances for these species were established – with special attention to the confirmation of previously detected methane depletion – the authors were able to run a grid of forward models in order to investigate vertical mixing, metallicity, and temperature structure of WASP-107 b. They establish a relatively hot intrinsic temperature for WASP-107 b of $T_{\text{int}} = 460 \pm 40 \text{ K}$, which is presumably responsible for the planet’s inflated atmosphere; and they infer a core mass of $11.5 M_\oplus$ (i.e. an approximate third of the planet’s total mass) which is significantly higher than previously established upper limits.

The other study by Murphy et al. (2024) uses spectra taken at 2.5 – $4.0 \mu\text{m}$ with the JWST/NIRC+M F210M filter and F322W2 grism in order to study the morning and evening limbs of WASP-107 b with the ambition to search for potential asymmetries between the morning and evening terminators. For tidally locked exoplanets of equilibrium temperatures above $\sim 1200 \text{ K}$, atmospheric models predict that limbs should become heterogeneous due to day-to-night circulation (e.g. Kataria et al. 2016; Powell et al. 2019) while remaining more homogeneous for exoplanets of lower temperatures such as WASP-107 b ($T_{\text{eq}} = 736 \text{ K}$ from Table 1). However, this work indicates that the planet’s morning limb is cooler by approximately 100 K , resulting in a scale height difference between the limbs for the wavelengths studied. These findings were confirmed and expanded upon by a follow-up study led by the same team, Murphy et al. (2025), now using JWST data from all of JWST’s science instruments (NIRISS, NIRC+M, NIRSpec, and MIRI) and covering for the first time the entire range of ~ 1 – $12 \mu\text{m}$. In this study, there is further evidence of strong variation between its morning and evening limb, specifically in abundances of SO_2 and CO_2 , and that clouds appear to only form on the morning limb, leaving the evening limb clear.

These studies, together with many other modelling studies not discussed here (e.g. Schlawin et al. 2018; Millholland et al. 2020; Wang & Dai 2021; Khodachenko et al. 2021; Linssen et al. 2022), demonstrate that the formation history, atmospheric structure, and dynamics of WASP-107 b are not yet fully understood.

Establishing specific pathways for the formation of exoplanets – especially those without analogues in our Solar System, such as WASP-107 b – is a very active field of research that is currently making good progress in providing evidence that complicated formation histories may manifest in both an exoplanet’s orbit (see e.g. [Maire et al. 2023](#), and references therein) and interiors (see e.g. reviews by [Nettelmann & Valencia 2021](#); [Guillot et al. 2022](#); [Foley 2024](#)). The studies listed in this section therefore shine light on the uniqueness of this target, and it has become increasingly clear that WASP-107 b appears to be a highly unusual planet both atmospherically, compositionally, and dynamically, which is an important piece of context for the reader moving forward. Thus far, the vast majority of these results have not been confirmed at high spectral resolution through any ground-based observations, which served as a key motivator for this study.

3. Observations and data reduction

In this section the observations and the treatment of the data are described in detail. In Sect. 3.1, we describe VLT/CRIRES⁺ observations, in Sect. 3.2 we summarise the data reduction, Sect. 3.3 details the alignment of the wavelength scales, in Sect. 3.4 we discuss stability and variability, and Sect. 3.5 covers cleaning and pre-processing necessary before our exoplanet detection methodology outlined in the next section.

3.1. Observations

In this work, two primary transits of WASP-107 b were observed as part of ESO programme IDs 108.C-0267(D) and 110.C-4127(D) (PI: L. Nortmann). A third transit was scheduled to be observed in May 2023, but this night was lost due to bad weather and no observations were taken. The two successfully observed transits were obtained using the CRIRES⁺ instrument that is installed on UT3 of the VLT at Paranal Observatory, Chile. CRIRES⁺ is the upgrade project of the CRIRES instrument, which was previously in use at the VLT until 2014. It is a cross-dispersed echelle spectrograph operating in the near-infrared and mid-infrared regions of 0.95 μm to 5.3 μm (YJHKLM bands) at a nominal spectral resolution of $R \sim 100\,000$ with the 0.2 arcsecond slit ([Dorn et al. 2023](#)).

This work is the first ground-based study of WASP-107 b at the K-band wavelengths of ~ 2.0 – 2.5 μm , which is significantly redder than previous ground-based studies at ~ 1 μm (Y band). While previous studies have covered this target at several different wavelength ranges from both space and the ground, this is the only ground-based study so far to cover this particular wavelength regime that crucially straddles the relatively unexplored area between the shorter ~ 1 μm regime and the longer >3 μm regime that is populated by several important molecular spectral features. All details regarding the parameters of the observations used in this work can be found in Table 2.

The first observation, henceforth referred to as Night 1 (N1), was obtained on the night to 11 March 2022. The second observation, henceforth Night 2 (N2), was obtained on the night to 23 February 2023. For both nights, observations consisted of 34 in-transit and the remainder of out-of-transit exposures³.

³ Following recommendations supported by e.g. [Cheverall & Madhusudhan \(2024\)](#) and [Palle et al. \(2025\)](#), the number of out-of-transit exposures taken was intentionally as high as our observing schedule allowed.

Table 2. VLT/CRIRES+ observations used in this work.

Night 1		
Parameter	Symbol/units	Value
Date	YYYY-MM-DD	2022-03-11
Obs. start	UTC	03:54:23
Transit start	"	05:32
Transit mid-point	"	06:55
Transit end	"	08:17
Obs. end	"	09:30:11
No. exposures (in/out)		64 (34/30)
Exposure length	[s]	300
Avg. S/N per pixel		132
Airmass		1.19–1.03–1.59
Relative humidity	%	15–27
Seeing	"	0.32–1.04
Avg. resolution	R	143 000
Night 2		
Parameter	Symbol/units	Value
Date	YYYY-MM-DD	2023-02-23
Obs. start	UTC	04:55:26
Transit start	"	05:48
Transit mid-point	"	07:11
Transit end	"	08:33
Obs. end	"	09:16:53
No. exposures (in/out)		50 (34/16)
Exposure length	(s)	300
Avg. S/N per pixel		130
Airmass		1.20–1.03–1.21
Relative humidity	%	27–38
Seeing	"	0.58–1.26
Avg. resolution	R	125 300
General settings		
Parameter	Symbol/units	Value
Wavelength setting		K2148
Wavelength coverage	μm	1.972–2.452
No. echelle orders		6
Exposure time	(s)	300
Readout time	(s)	14
Nodding pattern		ABBA

Notes. Relative humidity and seeing are given by their minimum and maximum values of the night. Airmass is given by three values corresponding to the start of the observation, the minimum value (occurring close to mid-transit), and the end of the observation.

All exposures were 300 seconds long, taken in an ABBA nodding pattern, where the change between two exposures during the transit is 370 m/s (or approx. 25% of the planetary diameter). This exposure cadence resulted in an average S/N of 130 per exposure, selected intentionally to optimise the trade-off between good S/N per exposure (in long exposures) and minimal smearing due to the planetary movement (in short exposures). Excessive smearing has been proven to have a significant impact on the HRCCS analysis used in this work (see Sect. 4.2), so our observing strategy followed the recommendations of [Boldt-Christmas et al. \(2024\)](#) to maximise the detection probability.

Figure 1 shows the variations across N1 and N2 for air-mass, S/N per exposure, and conditions (seeing and water




Fig. 1. Observing conditions as a function of time for Night 1 (left) and Night 2 (right). The vertical grey lines indicate the extent of the transit event. *Top:* airmass. The blue and orange coloured points represent nodding positions A and B. *Middle:* as above but for median S/N per pixel. *Bottom:* weather conditions. Red crosses correspond to the seeing in arcseconds for each exposure, and the blue line to the precipitable water vapour (PWV), i.e. water vapour along our line of sight, throughout the night in millimetres.

vapour). Weather conditions were generally favourable for both nights, but conditions were overall better on N1 compared to N2. At good seeing conditions, the adaptive optics (AO) system of CRIRRES⁺ performs so well that it delivers a central peak that is smaller than the $0.2''$ slit. This is an effect known as ‘super-resolution’, which has occurred in observations from VLT/CRIRRES⁺ previously – see, for example, Yan et al. (2023), Cont et al. (2024, 2025), Lesjak et al. (2025), and Nortmann et al. (2025). At super-resolution, the resolving power R can be as high as 150 000 while still sampled by 2.3 detector pixels (where the nominal $0.2''$ slit is projected on 3.5 detector pixels). The downside of super-resolution is uncertainty about where the star is located across the slit, which introduces a shift between A and B spectra obtained at different positions along the slit as part of the nodding procedure. The shift mostly manifests as RV offset (as large as 1 km/s) that is corrected for when resampling spectra on the common wavelength scale (see Sect. 3.3).

The seeing conditions for N1 systematically stayed below $0.6''$ within the transit window. The AO produced a stable central peak of the point spread function (PSF) with constant width and shape. The first 63 spectra show PSF width of 2.1 ± 0.1 km/s or $R = 143\,000$ measured from the cross-dispersion profile constructed by the data reduction software and verified using the stellar CO lines around 2325 nm. For N2, the seeing fluctuated above $0.6''$ throughout the transit, leading to noticeably variable PSF, leading to a resolving power R between 120 000 and 135 000 or 2.5 and 2.2 km/s correspondingly. 300 m/s is certainly visible in the averaged line profile, so we compared the results of the analysis for N1 and N2 (Sect. 6.1 and Fig. 8). No systematic differences that could be attributed to the variable spectral resolution were found.

3.2. Data reduction

The data were reduced using the standard ESO data reduction system for CRIRRES⁺ i.e. the cr2res pipeline (version 1.6.7)⁴ and recipes were executed with EsoRex (version 3.13.6)⁵. Raw calibrations were reduced by the recipes cr2res_cal_dark, cr2res_cal_flat, and cr2res_cal_wave, using daytime calibrations available under ESO Program ID 60.A-9051(A). Respectively, these: (i) produce the master dark used during cr2res_cal_flat and cr2res_cal_wave; (ii) produce the master flat and adopted bad pixel mask; and (iii) determine the wavelength scale to be adopted for science exposures. Science frames were grouped into A/B nodding pairs (32 pairs on N1, and 25 pairs on N2) and reduced using the recipe cr2res_obs_nodding which performs A/B nodding subtraction to remove the sky background and influence of hot pixels, before producing 1D extracted spectra for each exposure – one A and one B. Note that we do not perform dark correction on our science frames as the A/B nodding subtraction serves the same purpose. Finally, we corrected our science spectra for the effect of the echelle blaze function using the pipeline-generated blaze file from cr2res_cal_flat and a Python script external to the pipeline⁶.

⁴ ESO CR2RES Pipeline: <https://www.eso.org/sci/software/pipelines/cr2res/cr2res-pipe-recipes.html>

⁵ ESO Recipe Execution: <https://www.eso.org/sci/software/cpl/esorex.html>

⁶ Python software for interfacing with the CRIRRES⁺ data reduction pipeline, post-processing and cleaning reduced spectra, running SYS-REM, performing cross-correlation, and simulating synthetic exoplanet transits per Sect. 5 can be found at <https://github.com/adrains/luciferase>. Our wavelength alignment algorithm is available as an IDL script in the folder scripts_reduction on the same repository.

At this stage, these spectra are not yet science ready for our science case. The combination of super-resolution and optical effects introduced by separate A/B nodding positions produce two separate sequences of spectra each night – one A sequence and one B sequence – that are not coherently aligned to either an absolute or relative wavelength scale, which must be corrected before we can look for the subtle effects of a planet atmosphere on our data. We discuss how we approach this in the next section.

3.3. Wavelength alignment

The alignment of A and B frames required in-depth investigation of the wavelength scale assigned to the spectral orders by the ESO pipeline. The calibration process is based on the calibration source illuminating the whole slit, unlike the science target image created by the ‘overperforming’ AO system under super-resolution conditions. The projection of the slit on the detector is tilted and curved with respect to the detector pixels, and even though the shape of the slit image changes a lot (by a few pixels) across the focal plane, the ESO pipeline is capable taking this into account. What remains is the wavelength change for each science spectrum due to re-positioning the target on the slit during nodding in super-resolution conditions, which requires interruption of the AO operation.

For K-band observations, the data product of the ESO pipeline delivers 18 spectral segments (6 spectral orders registered by 3 detectors, with gaps between detectors and no wavelength overlaps between orders). The obvious solution of cross-correlating science spectra does not work well for all spectral segments due to the different number and origin of spectral features (stellar vs telluric). Another problem was aligning spectra between the nights, as the change in barycentric velocity has more impact on the cross-correlation between spectral segments that have more pronounced stellar lines. An attempt to use *molecfit* (Smette et al. 2015) to produce a robust template (that would not include stellar lines and thus will be insensitive to barycentric correction) also failed as some of the segments are essentially free of telluric features.

In the end, we selected a single spectral order dominated by telluric features (order 27 covering 2045.5–2089.3 nm range) and one exposure (for each night) as a template. Exposure selection was based on visual inspection for instrumental defects and cosmic ray hits missed by the data reduction system, and on the highest S/N pointed at spectrum 32 for N1 and 23 for N2 (spectrum 87 in the two-night combined sequence). In-depth discussion of the optical design of CRIRES⁺ with Ernesto Oliva (priv. comm.) convinced us that the change expected due to shifts across the slit is equivalent to the change in RV of the target rather than a geometrical translation of the focal plane image. With that assumption, we have written a wavelength morphing tool that transforms a spectral segment to its nominal wavelength scale given the reference template.

The actual wavelength scale of a given segment was derived using a polynomial approximation to the RV applied to the nominal wavelength solution provided by the standard calibrations of CRIRES⁺ for each night:

$$w_k(x) = w(x) \cdot \sum_{m=0}^M a_m^k \cdot x^m, \quad (1)$$

where $w(x)$ is the standard wavelength solution from calibration for pixel number x and $w_k(x)$ is the morphing of that solution to phase k ; adopted polynomial order M is between 0 and 2




Fig. 2. Wavelength correction illustrated as RV shifts to the wavelength in the middle of a spectral order (vertical axis) determined for each spectrum. Red and blue asterisks show AA or BB nodding pairs of consecutive exposures where no interruption in AO operation occurred and thus we expected no change in the across-the-slit position. The differences in these pairs provide an alternative estimate of the precision. The mean differences are 80 m/s for N1 and 36 m/s for N2, which are to be compared to the estimates based on telluric line centres giving 91 and 53 m/s, correspondingly. The large shift between the nights was measured by cross-correlating the two templates and converting it to an extra RV for each night to provide a sense of scale for the reader. In reality, most of this shift is translation (not RV) due to the reproducibility of the CRIRES⁺ wavelength settings.

with index m ; and a_m^k are the unknown coefficients for phase k . Coefficients are found through an optimisation process that tries to maximise the cross-correlation between the reference phase and every other phase of the same transit. This means that we do not revert to an absolute reference frame and the final data may have an additional fake ‘Doppler shift’ of the reference spectrum on the order of a few hundred metres per second. This shift will be present as an extra offset from the expected systemic velocity in the final results.

Cross-correlation between the reference spectra from different transits was used to connect the two nights. This shift combines the RV change due to super-resolution and the error in echelle positioning. The latter results in a shift along the main dispersion rather than a RV shift. The reproducibility of the echelle positioning in CRIRES⁺ is nominally one tenth of a pixel or about 100 m/s, but we find that in our case, the combined effect of positioning and super-resolution is around 300 m/s. This shift does not affect the analysis as the two nights are treated independently.

The precision of this correction was estimated by trying 30 different phases from each night (the same spectral order) as references (900 combinations) and comparing the resulting corrections. For $M = 0$, the root mean square (RMS) of the correction spread within each night was found to be 119 m/s for N1 and 99 m/s for N2. Increasing to $M = 1$ reduced these numbers to 82 m/s and 67 m/s, but $M = 2$ did not add any improvement. The wavelength correction derived for the reference order was applied to all spectral orders at that phase. The corrections are illustrated in Fig. 2 as RV values for the central column of the middle detector. An alternative estimate of the correction precision for a single night was done by measuring the wavelength of the telluric line centres. The latter were determined by fitting a




Fig. 3. Per-phase differences (black lines) between observed WASP-107 b spectra and an inverse model of the mean nightly telluric absorption with a mean PSF profile scaled to the appropriate airmass. This spectral segment is totally dominated by telluric features and contains no significant stellar lines. Vertical offsets between each subsequent phase were added for visibility. A blue overplotted spectrum shows the mean telluric spectrum for each night at the lowest airmass. Note that substantial structure in the residuals is only visible for two phases on Night 2 (third from bottom, and topmost), which is indicative of not only the PSF changes in those phases but also the general stability of the PSF otherwise.

Gaussian to the line profiles. For $M = 2$, we found the RMS to be 91 m/s for N1 and 53 m/s for N2. These values match nicely the correction scatter obtained for consecutive AA or BB spectra in the ABBA nodding sequences as shown in Fig. 2.

A version of the code used for this correction can be found in the repository⁵. Our intention is to integrate this tool into the official CRRES⁺ pipeline in the future.

3.4. PSF stability and telluric spectrum variability

Variations in the PSF or components of the telluric spectrum (e.g. H₂O absorption) may have strong negative consequences for detrending analysis methods applied to our WASP-107 b data. One way to probe the severity of these effects is to study the differences between telluric lines as a function of time. The advantage of using tellurics to probe PSF stability is that these lines do not experience time-varying Doppler shifts in excess of 100 m/s, unlike stellar or exoplanetary lines, but they do change in strength and shape as a function of the airmass, AO performance, and changes in the partial pressures of various molecules in the Earth’s atmosphere (e.g. humidity). A model of telluric absorption that could be adopted for the airmass of each of our spectra and compared with the data was therefore needed, where any systematic features in the residuals will indicate changes in the PSF and/or chemical composition of the atmosphere. Ideally, this telluric spectrum should be constructed from the data itself so as to match the instrumental profile and lines present – something more challenging to accomplish with a telluric spectrum produced via spectral synthesis methods such as *molecfit*. Fortunately, as part of related project, we have developed a method to recover such a telluric spectrum from observations that we could use here.

The method to reconstruct telluric absorption from observations – transmission spectroscopy decomposition (TSD) – is described in detail by the upcoming paper Piskunov et al. (2025). Briefly, TSD formulates the analysis of high-resolution exoplanet transmission spectroscopy data as an inverse problem, and optimises for three spectral components (stellar, telluric, exoplanet) which can be used to best reconstruct the observed data. Each

of these three components has a unique velocity frame, and by adopting a priori these velocities at each observed phase, the method can ‘disentangle’ the components⁷. TSD adopts the assumption that each night of data has a single constant telluric transmission spectrum affected only by airmass. While in reality this assumption is not strictly true, especially for wavelength regions affected by H₂O absorption, the resulting nightly telluric spectrum can effectively be considered the mean telluric spectrum at the mean PSF for that night. By comparing this telluric spectrum to each phase of our observed WASP-107 b data, systematic differences in the residuals would indicate a departure from this mean value, and thus highlight problems due to a changing PSF.

Figure 3 presents such a comparison for the middle detector of order 27 (2061–2075 nm). This order contains exclusively strong telluric lines, making it optimal for detecting PSF or telluric opacity variations during a transit. For N1, we see no substantial systematic deviations from the horizontal line for any of the telluric features. For N2, in two phases (third from the bottom and topmost, both out of transit phases), the residual shows inverted ‘v’ shape in the line cores, which is indicative of additional broadening of the observed profile – consistent with worse seeing for these two phases. From this comparison, we conclude that our dataset is luckily not affected with strong PSF variations, barring two phases in N2 that both occur outside the transit window.

3.5. Post-processing

With wavelength scales aligned and unified, the final step before analysis is to continuum normalise and clean our spectra for obvious artefacts that might interfere with the detection of an exoplanet atmosphere. To continuum normalise our data, we applied a first order polynomial correction to each of our 18

⁷ This is analogous to how one might disentangle time-series double-lined eclipsing binary spectra into the three spectral components of Star A, Star B, and tellurics which together can be used to reconstruct a model of the observed data at each phase given the velocities for the binary components.

spectral segments with reference to synthetic stellar and telluric template spectra (for e.g. masking; see Sect. 5 for more information on our adopted MARCS and molecfit templates). Each night was treated separately, as was each A/B sequence. However, within a night, the slope of the correction was set constant for each spectral segment, with only the scale allowed to vary between exposures. In other words, we assumed that the shape of the spectrum is constant across exposures, but the system throughput may change. Referencing synthetic spectra allows for better stability, especially in segments with numerous telluric absorption features. While the continuum normalisation framework is not ultimately required for the next detrending analysis, we find the results to be more robust.

For data cleaning, we looked for and sigma clipped two kinds of artefacts. The first only occur on a single pixel at a single exposure – perhaps the result of uncorrected cosmic rays – and the second are systematic differences between the A and B sequences – perhaps the result of uncorrected detector artefacts at one nodding position but not the other. For the first, we compute the standard deviation of the pixel in time for one sequence, and compared it to the median spectrum of the other sequence, clipping where this exceeds our adopted threshold. For the second, we computed the standard deviations of the entire spectral segment for the two nodding sequences, and compared these with the difference within nodding pairs. The assumption here is that systematic differences should only occur in cases of detector artefacts present in only one nodding position, in which case we must discard the entire column since ~50% of the data is now considered systematically aberrant. For both cases, we adopted a value of 5.0σ for sigma clipping to remove the worst of the artefacts without risking destroying subtle signals introduced by the planet, and only interpolated clipped pixels when fewer than 5 pixels (~10% of the exposures on a single night) are clipped – otherwise masking the entire column. Further, we chose not to extrapolate clipped edge (in time) pixels as testing indicated this introduced spurious features observable in later analysis. The end result is spectra packaged into arrays of shape $[N_{\text{exp}}, N_{\text{spec}}, N_{\text{px}}]$ where N_{exp} is the total number of exposures observed per night, N_{spec} is the number of spectral segments equal to $N_{\text{order}} \times N_{\text{detector}}$, and N_{px} is the number of spectral pixels per detector.

4. Method

This section describes how the reduced data were analysed using the cross-correlation technique. Section 4.1 details the removal of stellar and telluric features via SYSREM, and Sect. 4.2 summarises the cross-correlation of SYSREM residuals with `petitRADTRANS` exoplanet template spectra.

4.1. Removal of stellurics (SYSREM)

Spectra obtained through ground-based observations of a transiting exoplanetary system consist of spectral features from three different sources: (i) the stellar spectrum from the host star; (ii) the telluric contamination from Earth’s atmosphere; and finally, (iii) a very minor fraction of spectral features that have been imprinted onto the stellar light by transmission through the exoplanetary atmosphere. In this particular science case, it is only this final exoplanetary component that is of interest as it is this component that carries information about the planet’s atmosphere. To study it, this spectrum must be isolated from the stellar and telluric contribution – or ‘stellurics’ as they will be referred to collectively henceforth.

During a single transit, stellar lines shift relative to tellurics primarily due to the rotation of Earth. This effect is of the order of a few hundred metres per second, i.e. small compared to our spectral resolution, which warrants the assumption that telluric lines remain aligned during the relatively short observing window of a transit. From this assumption, a common method for telluric removal at infrared wavelengths is a type of Principal Component Analysis (PCA) algorithm known as SYSREM. This algorithm was first developed by Tamuz et al. (2005) for cleaning time-series photometry and later applied to observations of exoplanets, with early examples including Birkby et al. (2013), Brogi et al. (2014), and Schwarz et al. (2015).

Conceptually, SYSREM capitalises on the fact that the Doppler shift of the exoplanet will range across the transit event from being relatively blueshifted (at the beginning of the transit) to being relatively redshifted (at the end) from the perspective of the observer by ~350 m/s every exposure. Comparatively, the telluric lines will not be notably shifted: the stellar lines will only shift marginally (~10 m/s every exposure) based on the star’s relative motion, while the telluric lines will not shift at all with respect to the observer. As such, telluric lines will be recorded on roughly the same spectral pixel throughout the observing window, whereas the exoplanetary lines will first fall on bluer pixels and then shift throughout the night across to redder pixels from exposure to exposure. Notably, this distinction between the star and planet velocity frames becomes less significant for a lower orbital velocity of the planet, even for a relatively small reduction in the RV semi-amplitude K_p ; in the case of a hot Jupiter ($K_p \gtrsim 120$ km/s), the large orbital velocity of the planet results in significant and spectrally resolved differences between the stellar and planetary velocity frames, but in the case of an Earth-mass planet ($K_p \sim 50$ km/s), they might be indistinguishable as this would result in a change across the transit that is smaller than the resolution element of the spectrograph. For WASP-107 b, $K_p = 105$ km/s (see Table 1).

In practice, an algorithm such as SYSREM can fit and remove systematic trends – namely, the telluric lines that are persistently present at the same spectral pixel throughout the observation. Effectively, SYSREM models the wavelength-dependent and time-dependent systematics (stellurics) before subtracting them, which should in theory leave behind only the non-systematics (planetary features) in the resulting residuals. SYSREM does this iteratively, meaning that the process of fitting a model and then removing it from the observation is repeated some number of times, at which point the spectrum should in theory be fully free of stellurics. This model is created by taking the product of two components: a spectrum S that varies with wavelength λ , and a variation A that varies with time t to account for time-dependent fluctuations such as airmass. The model, f , may then be expressed as

$$f(\lambda, t) = S(\lambda) \cdot A(t), \quad (2)$$

which is repeatedly fitted for each iteration to the data for each pixel, minimising the sum of the residuals squared each time. In the case of cross-correlation searches for species that are present in the host star of the observed system (e.g. CO), SYSREM may be better executed in the stellar rest frame (i.e. where the telluric lines would shift marginally and the stellar lines do not; see e.g. Nortmann et al. 2025). For a more detailed mathematical description of SYSREM, see Sect. 3.1 of Czesla et al. (2024).

The benefit of SYSREM is that this procedure removes stellurics even when they densely overlap with the spectral region of the exoplanetary features, as is the case for our K-band observations, where simply discarding affected pixels would result in

an unsustainably large loss of data. However, there is a challenge associated with using SYSREM in that the user does not have full insight into exactly how many iterations result in a spectrum that is not still contaminated by tellurics (i.e. too few iterations) nor taken so far that the planetary signal also begins to be removed (too many iterations). There have been investigations of how to determine the most appropriate number of iterations, such as [Cheverall et al. \(2023\)](#) or [Meech et al. \(2022\)](#), but some general limitations remain in that certain factors – such as fewer exposures, unfavourable barycentric separation, varying humidity affecting tellurics, low signal – will inevitably create obfuscation in this procedure, which impacts us adversely, as discussed in upcoming sections. For example, significant changes in an opacity component (e.g. H₂O) of the telluric spectrum, or of the PSF width, during a transit may lead to a complete failure of SYSREM – especially for slow-moving planets, which is why we controlled that this is not the case for our dataset (Sect. 3.5). All of these concerns is ultimately why, in Sect. 6, the results of a wider range of SYSREM iterations ($n=8$, i.e. 3–10, out of the total 0–15) are shown rather than selecting a single iteration. To account for such imperfect detrending, we weighted the residuals produced by SYSREM by inverse of the standard deviation of each spectral pixel as a function of orbital phase (exposure). This has the effect of down-weighting pixels still strongly variable in time (e.g. those most affected by H₂O telluric absorption), improving the effectiveness of our cross-correlation analysis in the following section.

4.2. Cross-correlation with synthetic templates

Once observations have been cleaned of telluric contamination, the SYSREM residuals can be searched for planetary signal; however, this signal is so extremely small that statistical methods, such as cross-correlation analysis, must be employed. The HRCCS analysis uses synthetic templates computed with radiative transfer packages and an assumed model of the planetary atmosphere. Cross-correlation was then used to establish the statistically significant similarity (or lack thereof) with the planetary transmission spectrum hidden in the noisy residuals of SYSREM, which was done over a range of RV shifts for each exposure. Mathematically, this required evaluating the normalised cross-correlation function (CCF):

$$\text{CCF}(v, t) = \frac{\sum_{\lambda} x_{\lambda}(t) \cdot T_{\lambda(1+v/c)}}{\sqrt{\sum_{\lambda} x_{\lambda}^2(t) \cdot \sum_{\lambda} T_{\lambda(1+v/c)}^2}}, \quad (3)$$

where x_{λ} are the weighted SYSREM residual for data taken at time t and T_{λ} is the template offset with RV v . While the exact expression for CCF varies somewhat between applications, the HRCCS approach is by now well-established and has become a standard tool in the exoplanet community, as it has been employed by many similar studies (see e.g. the recent review by [Snellen 2025](#), and references therein).

4.2.1. Generating templates with petitRADTRANS

The templates used for cross-correlation in this work were generated using the radiative transfer package *petitRADTRANS* ([Molliere et al. 2019](#)), version 2.7.6. This package models high-resolution transmission spectra (using the 1b1 or ‘line-by-line’ method) at $R = 10^6$ that we later convolve to $R = 140\,000$, slightly above the average CRIRES⁺ resolution across the two nights.

Table 3. Parameters used to generate the *petitRADTRANS* templates.

For all templates				
Parameter	Units	Value		
Wavelength range	μm	1.8–2.7		
Pressure range	bar	1–10 ^{−10}		
Eq. temperature ^W	T _{eq} (K)	738		
Int. temperature ^S	T _{int} (K)	460		
Planet radius ^W	R _p (R _J)	0.94		
Gravity ^W	g (cm/s ²)	10 ^{2.45} ≈ 282		
Reference pressure ^D	P ₀ (bar)	0.01		
Atmospheric IR opacity ^D	κ _{IR} (cm ² /g)	0.01		
Ratio of optical and IR opacity ^D	γ	0.4		
Non-metallic mass fractions ^D	H:He	0.72:0.28		
Rayleigh species		H ₂ , He		
Continuum opacities		H ₂ –H ₂ , H ₂ –He		
Convolved spec. resolution	R	140 000		
Per retrieval				
Parameter	Units	Line list	CHI	Aur
Mean mol. weight [*]	u		3.02	2.64
VMR per species: ^{**}	(dex)			
CH ₄		HITEMP (2020)	−5.8	−6.1
CO		HITEMP (2010)	−1.9	−3.0
CO ₂		HITEMP (2010)	−3.9	−4.4
H ₂ O		ExoMol (2018)	−2.1	−2.6
H ₂ S		HITRAN (2013)	−8.5	−8.6
NH ₃		ExoMol (2011)	−5.0	−5.1
[†] SO ₂			−5.2	−5.7

Notes. Superscripts *D*, *S*, *W* correspond to values taken from [Dyrek et al. \(2024\)](#), [Sing et al. \(2024\)](#), and [Welbanks et al. \(2024\)](#) respectively. ^{*}The mean molecular weight (MMW), denoted by μ and as calculated in Eq. (4), is given here for the atmosphere when considering the presence of all species. ^{**}The values r_V cited here for VMRs are given in dex, i.e. as $\log_{10}(r_V)$ for each species. These values are taken from the Extended Data Table 2 columns for CHIMERA (CHI) and Aurora (Aur) of [Welbanks et al. \(2024\)](#), as indicated, and were originally published with errors. However, we do not include these errors here, as our templates are generated without uncertainties. [†]The published VMR for SO₂ was accounted for in our generation of templates insofar as its weight was included in our calculation of MMW, but this species is not part of our analysis, as it does not have features present in our wavelength range.

Two sets of templates were generated for our analysis, and a complete list of the parameters for all templates can be found in Table 3. The spectra are generated using *petitRADTRANS*’s associated opacity data which uses a number of different line lists for the species modelled. Most parameters were adopted from the studies of [Dyrek et al. \(2024\)](#) and [Welbanks et al. \(2024\)](#) in an attempt to facilitate comparability, favouring values from [Welbanks et al. \(2024\)](#) where the option existed as this work features more data points (including the [Dyrek et al. \(2024\)](#) dataset) across a broader wavelength range and can thus be considered more robust. Our adopted values include the assumption of the so-called ‘Guillot’ temperature–pressure profile, i.e. a profile based on the often used Eq. (29) in [Guillot \(2010\)](#). We set the infrared atmospheric opacity to be $\kappa_{\text{IR}} = 0.01 \text{ cm}^2 \text{ g}^{-1}$, and the ratio between optical opacity and infrared opacity to be $\gamma = 0.4$. The equilibrium temperature and intrinsic temperature were set to $T_{\text{eq}} = 738 \text{ K}$ ([Welbanks et al. 2024](#)) and $T_{\text{int}} = 460 \text{ K}$ ([Sing et al. 2024](#)) respectively, and at the planetary

radius $R_p = 0.94 R_J$, the pressure is $P_0 = 0.01$ bar and gravity is $\log_{10}(g) = 2.45$, which is $\approx 282 \text{ cm/s}^2$ (Dyrek et al. 2024; Welbanks et al. 2024).

Our two sets of templates are based on the results of the two retrievals from Welbanks et al. (2024), and the intention is to investigate how well each type of template fares both in contrast to one another, and also when accounting for the possible presence of a cloud deck (see Sect. 4.2.2). The two sets of retrievals from Welbanks et al. (2024) use Aurora and CHIMERA, two independent inference frameworks (whose outputs are found in their Extended Data Table 2), from hereon referred to as ‘Aur’ and ‘CHI’, respectively. The differences in the retrieved VMRs between these two forward models range between tenths of dex to more than 1 dex. For example, the two values for $\log_{10}(\text{VMR})$ of H_2S are CHI = -8.5 and Aur = -8.6 while the values for $\log_{10}(\text{VMR})$ of CO are CHI = -1.9 and Aur = -3.0 . One purpose of this exercise is to understand how sensitive the HRCCS method is to differences in retrieved values – or, at least, to set a lower limit on the parameter differences – that can be detected with sufficient confidence, and how this combines with the impact of modelling a cloud deck presence.

We took these VMRs r_V from the outcome of the Welbanks et al. (2024) retrievals. The mean molecular weight of μ for each retrieval was calculated from each r_V of species i (either metals, Z, or primordial, p , i.e. H_2 or He). Explicitly, μ was calculated using the relationship

$$\mu = \sum_i (r_{Vi} \cdot m_i) = \sum_Z (r_{VZ} \cdot m_Z) + \sum_p (r_{Vp} \cdot m_p). \quad (4)$$

This equation was solved noting that $\sum r_{Vp} = 1 - \sum r_{VZ}$ and assuming H:He mass fractions (X) of 0.72:0.28 gives $m_p = 2 \times 0.72m_{\text{H}} + 0.28m_{\text{He}}$. Values for r_{VZ} are known from the retrievals, and m_i are basic atomic and molecular data from the NIST database⁸. This mean molecular weight μ is assumed to be constant throughout the atmospheric layer that we are observing, thereby neglecting potential stratification due to photochemistry at lower altitudes or evaporation at high altitudes, as sufficient mixing is expected at the altitudes we observed, even though these processes may vary significantly and are currently poorly understood (e.g. Miguel & Kaltenegger 2014; Lavvas & Koskinen 2017; Soni & Acharyya 2024).

All templates were generated for the wavelength range covering the CRIRES⁺ K band plus a margin of approx. 0.2 μm in either direction (i.e. generated across 1.8–2.7 μm) and convolved to match the resolution of CRIRES⁺. Templates were generated for individual species as well as a ‘global’ template that includes contributions from all species. We included all species from the retrievals of Welbanks et al. (2024) except for SO_2 , as this species is not expected to have features in the K band. In all templates, however, SO_2 is still accounted for in our calculations of the mean molecular weight.

4.2.2. Inclusion of cloud deck in templates

The presence of clouds, hazes, and/or other condensates (here all included in the broader term of ‘aerosols’) in an exoplanetary atmosphere significantly impacts its transmission spectrum and are thus very important to accurately include in one’s analysis. Generally, aerosols scatter light and therefore quench the spectral features i.e. decrease the amplitudes of present species’

spectral lines. In the best case, this will only discreetly manifest as increased scattering slopes or reduced line depths (Pinhas & Madhusudhan 2017), but in less favourable cases, this can result in spectra that are at worst completely featureless. This is because the effect of aerosols on a transmission spectrum can be hugely varied as the total impact will be different depending on several different variables: planetary chemistry, atmospheric scale height, thermal and wind gradients across the planet’s limb, wavelength regime being studied, and many more. This large parameter space is partially why the field of exoplanetary cloud studies is so active today, as there are always several effects working in tandem for any given target that create complex results that can be challenging to interpret. For a more detailed discussion on the impact of aerosols on transmission spectra, see Sect. 3.1 of Gao et al. (2021).

While this impact is therefore well-known, in high resolution (ground-based) spectroscopy studies of hot exoplanets, there are arguments that are often made for not including clouds in the generation of templates for cross-correlation. Generally speaking, the reasoning for doing so would be that any potential clouds, if at all present, would be (i) below the atmospheric layer being probed, and (ii) sufficiently flat across our (comparatively narrow) wavelength range that their only impact should be an effective overall dampening of our line strength, meaning their inclusion or exclusion will not significantly impact the results of our cross-correlation detections (unless doing a simultaneous analysis over very broad wavelength regions; see Pino et al. 2018).

However, in the case of WASP-107 b, the matter of exactly how its clouds may substantiate is an active research question where recent studies have shed significant light on the topic but unknowns still remain. Dyrek et al. (2024) describe a high-altitude cloud deck at $P = 10^{-5}$ bar (see their Extended Data Fig. 4), while Murphy et al. (2024) and Murphy et al. (2025) demonstrate that WASP-107 b’s cloud formation differ across its morning limb and evening limb. Coupling this with the knowledge of its ongoing photochemistry from Dyrek et al. (2024), Sing et al. (2024), and Welbanks et al. (2024), and the varied shortcomings in our ability to model atmospheres (as acknowledged by all of the above), it is clear that a full description of clouds in WASP-107 b is still a work in progress with many unrestrained parameters. This is also supported by Changeat et al. (2025) who demonstrated the importance of cloud modelling in retrievals for JWST data across a range of targets including WASP-107 b.

Therefore, we only generated templates that assume two extreme border cases to straddle the upper and lower boundaries of reality: a grey (i.e. uniform across all wavelengths) cloud deck at $P = 10^{-5}$ bar versus no cloud deck. The decision to not model more complex or iterative manifestations of clouds was effectively due to the aforementioned reasons of why clouds are often neglected in ground-based analysis: if cloud decks are present at pressures higher than $P \approx 10^{-3}$ bar (i.e. at lower altitudes), we find no discernible differences appear in our petitRADTRANS generated spectra compared to the clear spectra (the standard deviation s of respective transit depths between the two being $s < 0.005$), and only minor differences ($s < 0.01$) appear at $P \approx 10^{-4}$ bar. Only at $P = 10^{-5}$ bar do differences begin to manifest relatively noticeably ($s = 0.01$), and as such, lower altitude cloud decks are in our case effectively equivalent to the clear case. Thus, the on/off treatment of clouds was considered acceptable for this work; in subsequent sections, we demonstrate why this question may need to be revisited in the future.

⁸ National Institute of Standards and Technology (NIST) database: <https://webbook.nist.gov/chemistry/name-ser/>




Fig. 4. Simulated global (i.e. including all species) transmission spectrum of WASP-107 b, coloured by species contribution, based on parameters from the results of the Welbanks et al. (2024) CHIMERA forward model and as generated by `petitRADTRANS` at a resolving power of $R = 10^6$. The plot demonstrates the notable change of the spectrum’s shape when either excluding (*top*, showing a ‘clear’ atmosphere) or including (*bottom*, a ‘cloudy’ atmosphere) the presence of a grey cloud deck at $P = 10^{-5}$ bar. The grey areas denote wavelength ranges that fall outside of the spectral orders of the VLT/CRIRES⁺ K band of $\sim 2.0\text{--}2.5\ \mu\text{m}$ (1972–2452 nm).

The assumed chemical composition of WASP-107 b means that the spectral features of the transmission spectrum are not uniformly affected by clouds across the wavelength regime of the K band. This is illustrated by Fig. 4, which shows the transmission spectrum of WASP-107 b based on the results of the CHI-model from Welbanks et al. (2024). In the K band, significant quenching can be seen in the spectral lines found between $\sim 2.1\text{--}2.3\ \mu\text{m}$, for which the lines are only barely visible above the cloud deck, compared to the lines at $\sim 2.0\text{--}2.1\ \mu\text{m}$ and $\sim 2.3\text{--}2.5\ \mu\text{m}$ that are relatively speaking less quenched. Crucially, not only does this mean that a significant amount of signal is lost compared to the cloud-free case, but this means that the cloud deck is not merely reducing the line strength of the spectrum’s lines across the whole spectral window as can be assumed in cross-correlation studies – instead, it is effectively altering the shape of the spectrum as a whole, which will have a significant effect on our HRCCS results (see Sect. 5). Figure 4 also demonstrates the relatively small and/or hidden contribution of $\text{CH}_4 + \text{CO}_2 + \text{H}_2\text{S}$ compared to H_2O , CO , and NH_3 , which is why no attempt is made to detect these species individually in Sect. 6.

Despite previously referenced findings that suggest differences between morning and evening limbs of WASP-107 b, we use one template at a time to construct the CCFs for the whole transit i.e. do not analyse the limbs differently. This was largely motivated by the claim by Murphy et al. (2025) in the caption of their Fig. 1, which states that ‘the combination of [their]

morning and evening spectra are consistent with the panchromatic limb-combined spectrum [...] presented in Welbanks et al. (2024)’. Supported by tests from simulations, we consider this to be a second order effect as only a small number of exposures – nine per transit – would carry the imprint of the ‘morning’ and ‘evening’ spectrum. Furthermore, the discontinuity arising from a sudden change of templates is non-physical, and implementing a smooth transition between the templates would require further considerations, making a simple switch unrealistic.

4.2.3. Construction of $K_p - v$ diagrams

Once the templates are generated, these can be used to perform cross-correlation with the output of the stelluric removal i.e. the residuals containing the extracted planetary spectrum. As denoted by Eq. (3), the cross-correlation analysis is performed individually on each exposure of total N exposures before these are summed. As the amount by which the exoplanet has been Doppler shifted varies over time t , the amount of RV shift v at which the CCF is maximised will vary between each exposure. By plotting the CCF of each exposure in a heat map (where each exposures are plotted on the vertical axis, i.e. each exposure is denoted by a single row), the maximal cross-correlation value in S/N is indicated through brightness at the value of v where this CCF maximisation happens. Due to the movement of the planet, the first exposures (pre-transit) and last exposures (post-transit)

should contain no planetary signal whereas the exposures taken during transit should contain planetary signal along the slanted line of CCF peaks, following the changing RV of the planet. In our analysis, CCF plots are generated separately for exposures in each night (N1 or N2) and nodding position (A or B) to create four groups of CCF plots of N1-A, N1-B, N2-A, and N2-B. These plots are shown later in Sect. 6, confirming that no obvious artefacts or anomalous features are present.

While the information necessary to infer a detection is technically available at this point, no obvious planetary trail is visible, and the CCF for any individual phase remains too noisy to detect the planet. In order to boost this, we can collapse our set of $N_{\text{phase}} \times N_v$ CCFs in the phase dimension – i.e. collapse the data in the time dimension – by aligning the CCFs in the exoplanet rest frame and combining. At each phase there are three velocity components to consider: the time varying barycentric velocity (known), the constant systemic velocity of WASP-107 relative to the Solar System barycentre (v_{sys} , known), and the time varying projected exoplanet orbital velocity (not necessarily well known). To get around not perfectly knowing the exoplanet velocity, we repeatedly collapse our CCFs for some set of trial exoplanet RV semi-amplitude (K_p) values, and expect our signal to be maximised at the real K_p of the planet.

Shifting CCFs according to trial RVs for a range of K_p values (extending on both sides of the target's actual K_p , i.e. from $K_p = 0$ –400 km/s, and $v = \pm 200$ km/s) and co-adding all the exposures produces a so-called $K_p - v$ map that shows the CCF surface in K_p and v coordinates. A well-defined peak in the $K_p - v$ maps indicates that the template has matches in the SYSREM residuals, and if this peak occurs near the true K_p and v values⁹, we have a detection. Significance of detection is estimated by calculating the standard deviation (computed from the full map excluding the central region around $v = 0$ km/s by ± 30 km/s, i.e. the outer regions only), and comparing the maximum S/N of the map to the standard deviation, where a larger detection significance is implied by a larger deviation.

5. Analysis of simulated data

As the field of HRCCS of exoplanet atmospheres has been active for nearly two decades by now, a large sample of exoplanets have been successfully characterised in that time as citations throughout this paper show. For understandable reasons, the strongest detections have come from targets that are (i) close-in to their host star, meaning there are frequent opportunities for observations, large velocity excursions, and usually with circular orbits and thus tidal locking, (ii) have extended atmospheres, and are sufficiently hot such that their atmospheric species produce strong atomic spectral lines with a minimal likelihood of cloud formation due to dissociation. This has resulted in an over-representation of hot Jupiters and ultra-hot Jupiters in the catalogue of well-studied exoplanetary atmospheres, as can be shown by statistics by the IAC ExoAtmospheres database² and as summarised in Table 1 of Cont et al. (2024).

In the case of WASP-107 b, there are a number of factors that complicate our analysis in terms of how strong the anticipated cross-correlation signal should be. The target is not hot (below 800 K), which partly explains why its atmosphere is dominated by molecular species at these temperatures. Due to the fact that molecular species produce orders of magnitude more spectral

lines than atomic species – especially in the infrared wavelengths – their features are intrinsically more challenging to identify as their line lists are often less accurate and complete in comparison with atoms. For further discussions on the impact of molecular line lists on exoplanet studies, see, for example, Hoeijmakers et al. (2015) for titanium oxide; Hedges & Madhusudhan (2016) for water; Bowesman et al. (2021) for aluminium oxide; Tannock et al. (2022) for methane; and Rengel (2022) for an overview of atomic and molecular databases used in the (exo)planetary community. Furthermore, as described in Sect. 4.2.2, the presence of a cloud deck may be detrimental in our analysis; and as discussed, the detailed properties of the WASP-107 b clouds are largely unknown.

To understand these challenges, we developed a transit simulator to study and predict how some of these effects might substantiate and impact our interpretation of the real data. Investigating how our target may appear through the use of simulations is a crucial guiding tool in this new and under-explored regime of transmission spectroscopy of cool planets, as we do not expect to see the canonical $K_p - v$ maps that are more closely associated with past studies of hotter exoplanets.

Our simulated observations require six main ingredients: (1) a simulated stellar spectrum for WASP-107; (2) a model of Earth's atmospheric transmission incorporating key telluric species; (3) a synthetic transmission spectrum for WASP-107 b; (4) instrumental throughput and a wavelength scale for CRIRES⁺; (5) system parameters describing the star, planet, and their orbits; and (6) velocities (star, planet, barycentric), airmasses, and slit losses computed at each simulated epoch. To simplify this process and increase the explanatory utility of the simulations, we opt to simulate the same two transits as we actually observed down to the timing and number of exposures as described in Table 2. Our star, planet, and orbital parameters are taken from Table 1, and the velocities of each rest frame (star, planet, barycentric) are computed for each epoch using these system parameters. Our planet transmission spectra are modelled as per Table 3 (CHI-values), and our telluric transmission spectra are two molecfit models each computed and then fitted to the master spectrum from each night of our observed CRIRES⁺ data. Finally, our stellar spectrum is based on a custom 1D local thermodynamic equilibrium spherically symmetric model atmosphere computed using the latest version of the MARCS code (Gustafsson et al. 2008). The stellar spectrum is generated at $R \sim 100\,000$ adopting parameters from Piaulet et al. (2021) and using MARCS opacity tables, the solar abundance ratios of Grevesse et al. (2007), and VALD database line data (Ryabchikova et al. 2015). This stellar intensity spectrum was computed for 49 radial stellar disk positions (i.e. μ angles), which enables us to interpolate specific intensities directly without needing to invoke limb-darkening approximations. A full description of our simulation methodology is available in the methodology paper (Piskunov et al. 2025).

While we did our best to simulate our real CRIRES⁺ observations, there are several approximations made that prevent a 1:1 simulation of reality, but we do not perceive these as limiting the interpretability and utility of our simulations. Firstly, beyond adopted instrumental transfer functions for CRIRES⁺ as made available by ESO, we make no attempt to model the optics of the spectrograph, the pipeline reduction of data from raw 2D to extracted 1D spectra, nor effects such as correlated noise, hot pixels, nodding, cosmic rays, or super-resolution. Due to the low $v \sin i$ of WASP-107 (0.45 km/s determined by Rubenzahl et al. 2021) as compared to the motion of the planet and velocity resolution of CRIRES⁺, we neglect stellar rotation – and thus the

⁹ Note that the choice of x -axis zero point is somewhat arbitrary, with the two main conventions being for $v = 0$ km/s to be the Solar System barycentre or, as we adopt here, $v = v_{\text{sys}}$ i.e. the rest frame of WASP-107.




Fig. 5. Simulated $K_p - v$ plots used to test the impact of including or excluding cloud decks in four different test cases. *Top:* simulations of a ‘clear’ WASP-107 b that has no cloud deck present cross-correlated with a ‘clear’ template that excludes clouds (left) and with a ‘cloudy’ template that includes clouds (right). *Middle:* same as the top row but now simulating a ‘cloudy’ WASP-107 b with a cloud deck present at $P = 10^{-5}$ bar, shown with the same colour scale as the test above for comparison. *Bottom:* same as the middle row but now shown with its own colour scale for reference.

Rossiter-McLaughlin effect (Rossiter 1924; McLaughlin 1924) – when modelling. Despite evidence indicating that WASP-107 b has observable limb-like 3D structure variations that presumably result in complex line profiles visible at high resolution, we treated the planet atmosphere as 1D and constant with time. We modelled slit losses and thus varying ‘seeing’ conditions, but we did not model time-varying tellurics (except for the changing air-mass), for example, as the result of changing humidity conditions throughout a transit. Finally, we simulated WASP-107 b transits at $S/N \sim 130$ to match our real data assuming normal distribution of noise.

5.1. Note on fluctuations in detection significance

For the simulations presented in the following sections, the detection significance can be seen to vary at the $\sim 1\sigma$ level. This is expected behaviour, and is the result of random noise and slit losses applied to each spectral pixel and exposure respectively. These random excursions from the average have already been proven to exist in a previous simulation study aimed at finding the best possible observing strategy – for further discussion of this effect, see Sect. 4.3 of Boldt-Christmas et al. (2024).

5.2. Test 1: Inclusion versus exclusion of cloud deck

One of our first tests was to measure the impact of including or excluding a cloud deck in our cross-correlation templates, considering a hypothetical scenario where we do not know if our target in reality has a cloud deck present (cloudy) or not (clear). In these simulations, we demonstrate four cases:

1. Using a clear template with a clear atmosphere.
2. Using a clear template with a cloudy atmosphere.
3. Using a cloudy template with a clear atmosphere.

4. Using a cloudy template with a cloudy atmosphere.

In all four cases, no other parameters are changed between each simulation. In order to recreate the most common HRCCS scenario of searching for a singular species at a time, both the cloudy and clear cross-correlation templates include spectra from H_2O only. The resulting $K_p - v$ maps (combined N1+N2) are shown in Fig. 5, and demonstrate two important effects. Firstly, there are relative differences in detection significance between the cases where the planet’s actual physical characteristics (i.e. in this case, whether a cloud deck truly is present or not) matches or does not match the template’s characteristics. In both the case of a clear planet + clear template and cloudy planet + cloudy template, the detection significance is higher than that of a clear planet + cloudy template and cloudy planet + clear template respectively, i.e. higher when the template matches the atmosphere. It may not be surprising that one is rewarded if using a template that is truer to reality, but it may be surprising that the magnitude of this reward is so significant, with the ‘mismatched’ cases achieving only approximately 65% and 75% of the total possible detection significances respectively. Secondly, there is a very stark difference in detection significance between the cases of cloudy and clear planets, with the clear planet + clear template case resulting in a maximum S/N that is nearly five times higher than the cloudy planet + cloudy template case; even the clear but ‘mismatched’ case of clear planet + cloudy template fares better as its maximum S/N is over three times higher (even when recalling the expected fluctuations described in Sect. 5.1). This is in line with previous work that led to similar conclusions (Pino et al. 2018; Allart et al. 2020; Lafarga et al. 2023) for hotter targets.

This effect underlines a clear obstacle in studying exoplanets with cloudy atmospheres, which becomes increasingly relevant as we move towards characterising cooler (and thus cloudier) exoplanets. In the specific case of WASP-107 b, the presence of




Fig. 6. Simulated $K_p - v$ plots to test varying the argument of periastron ω in ten different test cases. *Top*: simulations of a ‘clear’ WASP-107 b that has no cloud deck present cross-correlated with a ‘clear’ template that excludes clouds for five different values of ω . *Middle*: same but now simulating a ‘cloudy’ WASP-107 b cross-correlated with a ‘cloudy’ template, where both the simulation and template include a cloud deck at $P = 10^{-5}$ bar. The data are shown using the same colour scale as the test above for comparison. *Bottom*: same as the middle row but now shown with its own colour scale for reference.

a cloud deck clearly reduces the detection significance from a strong detection to one barely above the noise floor even under these idealised circumstances, making the choice of ‘correct’ template even more important.

5.3. Test 2: varying the argument of periastron

The role of orbital configurations on exoplanet observations has been discussed extensively in the literature, which is understandable considering its potentially great impact on many systems. As described in the beginning of this section, the field of cross-correlation spectroscopy has been primarily raised on case studies of hotter, close-in planets that tend to have circular orbits – but as the field continues to attempt to characterise increasingly smaller and cooler exoplanets and wider orbits, the likelihood of a given target having a more complex orbit increases. As such, many feasibility and methodology studies have been recently produced in the context of a wide range of parameters – to name a few, [Prinoth et al. \(2024\)](#) for high eccentricity and the argument of periastron; [Cheverall & Madhusudhan \(2024\)](#) for low velocity planets; and [Hong et al. \(2025\)](#) for orbital periods.

For the next test, we study the impact of the argument of periastron, also known as the argument of periapsis, which is one of the orbital elements required to fully describe an elliptical orbit. The argument of periastron, ω , is the angle between the ascending node (i.e. the point where the orbital plane crosses the reference plane, in the rising direction) and the orbit’s closest point to the star (measured in the plane of the orbit). ω defines the orientation of the semimajor axis in the plane of the orbit relative to the observer – the ‘twist’ in the spherical coordinate direction – and therefore determines what segment of the orbit is actually captured during the primary eclipse and how the RV of the planet changes during transit.

Here, we want to test the impact of arguments of periastron values given the eccentricity suggested for WASP-107 b, which is $e = 0.06 \pm 0.04$ ([Piaulet et al. 2021](#)), in the context of the previous test on cloud inclusion. In this simulation, we model WASP-107 b using its known orbital parameters (including

eccentricity) across a range of ω values, but construct our $K_p - v$ map using the equation for a circular orbit in order to mimic the often true reality of not perfectly knowing orbital parameters. This is also relevant considering the well-known bias in orbital eccentricity determinations where the e of a purely circular orbit may come out as slightly positive due to random noise, meaning non-zero e values $< 3\sigma$ away from zero (such as in our case) may be less trustworthy ([Lucy & Sweeney 1971](#); [Shen & Turner 2008](#); [Zakamska et al. 2011](#)).

Figure 6 demonstrates the effect of varying ω across a full half-revolution of 0–180° for both clear simulations and cloudy simulations. Both cases are cross-correlated with their ‘matching’ respective templates of H_2O , i.e. the clear simulations are cross-correlated with a clear template and the cloudy simulations with a cloudy template. All other parameters between them remain the same, and are otherwise the same as in the first test.

For both the clear and cloudy case, $\omega = 90^\circ$ provides the best possible detection significance in the $K_p - v$ map. This is to be expected to be the case for a slightly eccentric orbit such as ours, as demonstrated by [Prinoth et al. \(2024\)](#) as this will align the observation with capturing the primary eclipse at a point during which the planet’s RV is maximally changing. The other angles, however, have a more varied impact on the clear and cloudy cases. For the clear simulation, it is notable that the detection significance is effectively halved at $\omega = 0^\circ$ and $\omega = 180^\circ$ compared to $\omega = 90^\circ$, and that the impact on detection significance is not immediately clear with only a single realisation for $\omega = 45, 135^\circ$ due to the reasons detailed regarding noise realisations in Sect. 5.1. Meanwhile, for the cloudy simulation, while $\omega = 90^\circ$ fares better than other angles, the variation between them is significantly less notable; at this point, the signal is already diminished to the point that the effect of cloud-quenching is more significant than the effect of ω .

One will also note the varying tilt of the signal in the $K_p - v$ map, which reflects the asymmetry of the planet’s RV change throughout the transit. This is a known phenomenon that creates a ‘stripe’ structure on the $K_p - v$ plot, where stripes tend to be vertical for $\omega = 90^\circ$ and slanted for other values. This effect was




Fig. 7. Simulated $K_p - v$ plots to test varying the argument of periastron ω and orbital eccentricity in six different test cases. *First row:* simulations of a ‘clear’ WASP-107 b that has no cloud deck present cross-correlated with a ‘clear’ template that excludes clouds for three different values of ω , all for a WASP-107 b on a fully circular orbit. *Second row:* same as the top row but now simulating a ‘cloudy’ WASP-107 b cross-correlated with a ‘cloudy’ template, where both the simulation and template include a cloud deck at $P = 10^{-5}$ bar. *Third row:* same as the first row but now for a WASP-107 b on its true (somewhat) eccentric orbit of $e = 0.06$. *Fourth row:* same as the second row but now for a WASP-107 b on its true orbit of $e = 0.06$.

previously described in [Basilicata et al. \(2024\)](#), who saw this in GIANO-B transmission spectra of the warm Neptune-like planet HAT-P-11 b (Kepler-3 b), which is even more eccentric at $e \approx 0.26$ and shows a pattern similar to our plots.

5.4. Test 3: Varying orbital eccentricity

In the case of WASP-107 b’s only slightly eccentric orbit where $e = 0.06$, it is fairly intuitive that $\omega = 90^\circ$ should give a maximised detection significance. To compare: for a less eccentric (i.e. circular) orbit, the planet’s orbital velocity v_{orb} is constant at all points across the orbit, so there will be no particular angle that will create a higher Δv_{orb} across the transit; for a more eccentric orbit, the angle at which Δv_{orb} is maximised during the transit window would be further away from $\omega = 90^\circ$ (with exact extent varying on eccentricity e). In order to demonstrate this intuition, our third test varies the argument of periastron ω once more, but in this case, also for the case of a fully circular orbit, both with and without clouds. The results are shown in Fig. 7.

In these simulations, it is demonstrated that (i) the tilt arising in the $K_p - v$ plots as mentioned in Sect. 5.1.2 and in [Basilicata et al. \(2024\)](#) can indeed be attributed to the eccentricity, as it is absent in the simulations for the circular orbit; (ii) within the noise fluctuations expected for a single realisation (see again Sect. 5.1.1), the detection significance remains broadly the same for all ω angles in the case of the clear, circular orbit, but the detection significance drops by almost half for the clear, eccentric orbit case. In the cloudy cases, mirroring the findings of the previous tests, the variation between the circular

and eccentric cases are almost negligible as the signal strengths are here severely quenched.

6. Analysis of real data

6.1. Cross-correlation analysis of real data

Once the two nights of data from VLT/CRIRES⁺ have been pipeline reduced, wavelength corrected and detrended, we begin the cross-correlation analysis using the technique and templates described in Sect. 4.2. Following this, there are 4×4 templates to use in our cross-correlation analysis: three templates contain the spectral contribution of the single species (with continuum) CO, H₂O, and NH₃ while a fourth template contains all six species and is referred to as the ‘global’ template. The three species included in the global template that do not have their own templates (CO₄, CO₂, and H₂S) were determined to have vanishingly small contributions at our wavelength range – see Fig. 4 – and thus no analysis was attempted for them individually. Each of the four templates were generated in four models: based on values of the two [Welbanks et al. \(2024\)](#) model frameworks Aurora and CHIMERA; and including or not a grey cloud deck at $P_{\text{bar}} = 10^{-5}$ bar. The templates without clouds are referred to as nC for ‘no cloud’ and yC those with as ‘yes cloud’, creating the four labels Aur+yC, Aur+nC, CHI+yC, and CHI+nC. All cross-correlations were done in the instrument (telluric) rest frame, except for CO, which was shifted to the stellar rest frame considering the presence of CO in the stellar spectrum of WASP-107 b.

Before creating the $K_p - v$ plots, we inspected the CCF maps that span the offset of $v = \pm 200$ km/s for N1-A, N1-B, N2-A,




Fig. 8. Cross-correlation function maps of the A and B frames for Night 1 (N1) and Night 2 (N2) from our VLT/CRIRES⁺ data. For each plot, the dotted white lines denote the start and end of the transit. The data are shown for SYSREM iterations 0, 5, 10, and 15 in order to demonstrate that all non-zero iterations are clear of any obvious contamination or artefacts. The maps show cross-correlation with the ‘global’ template from the CHI-model (see Table 3) that includes a cloud deck at $P = 10^{-5}$ bar.

and N2-B as described in Sect. 4.2.3. The data have insufficient S/N to visually show a planetary trace, so these plots are merely inspected to ensure that there are no obvious artefacts or clear contamination in the data. The CCF plots with global template CHI+yC for telluric frame SYSREM iterations $n = 0, 5, 10, 15$ are shown in Fig. 8, where the dotted white lines represent the beginning and end of the transits.

Each template was cross-correlated across a $K_p - v$ space in the range of $v = \pm 200$ km/s centred on the stellar RV, then shifted and co-added in steps of 1 km/s across $K_p = 0-400$ km/s. Thus, the expected position of the $K_p - v$ peak should be at $K_p = 105$ km/s (Table 1) and $v = 0$ km/s as we adopt the rest frame of the exoplanetary system. For each of the 16 templates, a $K_p - v$ plot was generated after each iteration of SYSREM, which we run for 15 iterations. Considering the first few SYSREM iterations clearly contain residual stellar signal, and considering that detection significances drop off towards ~ 12 iterations, we only plot SYSREM iterations 3–10 for each template. The motivation behind plotting this wide range of SYSREM iterations rather than selecting a single iteration is to demonstrate the consistency of the plots’ broader features across multiple iterations. The plots for the global template and the CO template are found in Fig. 9, and the plots for the H₂O template and NH₃ template are found in Fig. 10.

In these plots, two detection significances are shown in the top left corner of each iteration: one value for the maximum peak of S/N across the full $K_p - v$ map, and one value for the maximum peak of S/N closer to the expected value i.e. the maximum value within a narrow column centred on $v = 0$ km/s. These maxima are

denoted by ‘F’ for ‘full’ and ‘E’ for ‘expected’. The F-maximum is calculated from the entire K_p space, as described in Sect. 4.2.3, i.e. by subtracting the background map (the $K_p - v$ values outside the central column around ± 30 km/s from $v = 0$ km/s) from the full map, all divided by the standard deviation of the non-central region. For all species except NH₃, the F-maximum is generally between $3-6\sigma$, where the F-maximum of NH₃ is generally $< 3\sigma$, so NH₃ can thus already at this point be considered a non-detection. However, the F-maximum for all species consistently appears at an offset from the expected $K_p - v$ maximum position for all templates, and for all SYSREM iterations in the direction of larger K_p .

The E-maximum is then calculated in the same way as the F-maximum, but this value is instead the maximum from within a box of $N = 13$ km/s from the expected value in the $K_p - v$ map. This value of N was chosen by first calculating the average difference (across all SYSREM iterations and all parameter sets for that species) between the maximum found within boxes of different N -values and the previous F-maximum. By requiring the average difference between these values to be $< 1\sigma$, i.e. that both the F-maximum and the E-maximum coordinates are both within where a 1σ contour would be, we calculate the lowest value of N that fulfils this requirement. Excluding the non-detection NH₃ maps, we find this N -value to be 13 km/s, so the E-maximum is then the highest S/N found within this 26×26 km/s box, i.e. within 13 km/s from the expected $K_p - v$ value. Across all SYSREM iterations and for all parameter sets in Figs. 9 and 10, the F-maximum and E-maximum are shown in the top left of each $K_p - v$ plot.




Fig. 9. $K_p - v$ plots for SYSREM iterations 3-10 cross-correlating with the global template *top* and with the CO template *bottom*. In each plot, the dotted white line denotes the expected location of the peak, and the dashed yellow line denotes the actual location of the peak. In the top left, the F-value denotes the maximum value found across the ‘full’ map, and the E-value denotes the maximum value found within the ‘expected’ location, i.e. within 13 km/s of $K_p = 105$ km/s and $v = 0$ km/s.

The E-maximum is $\sim 6\sigma$ for the global template (maximum = 6.34σ , SYSREM = 8) and CO templates (maximum = 6.41σ , SYSREM = 5), and $\sim 4\sigma$ for the H₂O template (maximum = 4.53σ , SYSREM = 7). In all three cases, the average E-maximum is higher for the two parameter sets that do include clouds than those that do not, with the average difference in E-maxima for yC and nC respectively being: 0.90σ for the global template; 1.64σ for CO; and 0.19σ for H₂O.

The differences between the retrievals from the two forward models (Aur versus CHI) have a much smaller impact on the detection significance than the inclusion of clouds (yC versus nC). To separate the effects, we construct four $K_p - v$ difference maps after each SYSREM iteration from 3 to 10, and compare their average values (Δ Map) and the differences of F-maxima (Δ Max). The purpose of the difference maps is to isolate the effect of the clouds from the effect of forward models, so we compute the following differences: (Aur+yC) – (Aur+nC) and (CHI+yC) – (CHI+nC) for the clouds and (Aur+yC) – (CHI+yC) and (Aur+nC) – (CHI+nC) for models. We then take the absolute value of the average, μ , for differences at each SYSREM iterations 3–10, and calculate the standard deviation, s , in that set of differences. A table summarising these values can be found in Table 4.

Table 4. Differences between $K_p - v$ results from the four parameter-set templates for each species.

		Global		CO		H ₂ O	
		μ	s	μ	s	μ	s
Δ Map	(Aur+yC) – (CHI+yC)	0.01	0.40	0.03	0.24	0.03	0.18
	(Aur+nC) – (CHI+nC)	0.01	0.22	0.01	0.18	0.00	0.09
	(Aur+yC) – (Aur+nC)	0.05	1.04	0.05	0.70	0.11	0.96
	(CHI+yC) – (CHI+nC)	0.04	0.98	0.04	0.65	0.09	0.90
Δ Max	(Aur+yC) – (CHI+yC)	0.61	1.55	0.28	0.42	0.11	0.15
	(Aur+nC) – (CHI+nC)	0.10	0.48	0.25	0.28	0.00	0.09
	(Aur+yC) – (Aur+nC)	0.32	0.65	1.02	0.57	0.63	0.59
	(CHI+yC) – (CHI+nC)	0.84	1.61	0.99	0.52	0.52	0.41

Notes. The Δ Map is the subtraction of full $K_p - v$ maps, and Δ Max is the subtraction of F-maxima. For both Δ Map and Δ Max, the first two rows correspond to differences where the cloud inclusion is the same, and the last two rows correspond to differences where the forward models are the same. Values shown are the modulus average (μ) and standard deviation (s) across SYSREM iterations 3–10. All values are rounded to 2 d.p.




Fig. 10. $K_p - v$ plots for SYSREM iterations 3-10 cross-correlating with the H₂O template *top* and with the NH₃ template *bottom*. In each plot, the dotted white line denotes the expected location of the peak, and the dashed yellow line denotes the actual location of the peak. In the top left, the F-value denotes the maximum value found across the ‘full’ map, and the E-value denotes the maximum value found within the ‘expected’ location, i.e. within 13 km/s of $K_p = 105$ km/s and $v = 0$ km/s.

In this table, it can be seen that the differences of ΔMap have $\mu \sim 0.02$ and $s \sim 0.22$ for forward models (i.e. templates of the same cloud inclusion) but $\mu \sim 0.06$ and $s \sim 0.87$ for differences in cloud inclusions (i.e. templates of the same forward models). For ΔMax , the trend is the same, with $\mu \sim 0.23$ and $s \sim 0.50$ for differences in forward models but $\mu \sim 0.72$ and $s \sim 0.73$ for differences in cloud inclusions. Across the board, this implies that the differences between the inclusion or exclusion of a cloud deck are more impactful than a different VMR of a species, even when the ratio is varied by fairly notable quantities (for example, 1.1 dex in the case of CO). This is particularly notable considering that the two retrievals do account for their respective MMW, which differ by 0.38 u (see Table 3), and HRCCS being seemingly more sensitive to clouds than to MMW may indicate that high-resolution observations could be a way to resolve the metallicity-cloud degeneracy.

6.2. Investigation of $K_p - v$ maximum location deviation

For all templates, the location of the F-maximum along K_p varied across all SYSREM iterations and parameter sets from as low as ~ 150 km/s (i.e. ~ 45 km/s from the expected location of 105 km/s) to as high as ~ 400 km/s (the edge of $K_p - v$ map). In the v direction, no significant variation was found, with all species-parameter combinations in all SYSREM iterations generally deviating less than 5 km/s (most often only by ~ 1 km/s).

The F-maximum location is relatively consistent across all templates, indicating that its offset in K_p from the expected location comes from the data and our analysis. In the global template, the maximum mostly appears at ~ 150 km/s in the two nC templates, except some SYSREM iterations that have F-maxima at the edge of the map, but consistently at ~ 200 km/s in the two yC templates. In CO and H₂O, it is instead the clear models that produce a larger deviation, placing the F-maximum closer to ~ 300 km/s in all four nC templates (again with some SYSREM iterations showing the F-maxima at the map edges for CO) but consistently at ~ 200 km/s in the two yC templates (ranging from approx. ~ 180 km/s to ~ 250 km/s, similar to the global yC-templates).

Some of these commonalities can be seen in Figs. 9 and 10. Here, a characteristic ‘chimney smoke’ pattern is visible, where the more centralised peak is accompanied by higher S/N at effectively all K_p values above it, and appears to be broadening towards higher K_p values. This pattern varies from being more or less pronounced for different species and at different SYSREM iterations, but this region is most distinct in the nC templates of H₂O (which also show a larger difference in F-E than any other combination).

To explore whether this effect may be an artefact in the data (possibly from a specific detector, a specific nodding position, etc.), we experimented with excluding different combinations of data from Night 1, data from Night 2, data from A-frames,




Fig. 11. Real and simulated $K_p - v$ plots for SYSREM iterations 3–10 cross-correlating with the H₂O template including clouds. In each plot, the dotted white line denotes the expected location of the peak, and the dashed yellow line denotes the actual location of the peak. *First row:* real data $K_p - v$ map as shown in Fig. 10. *Second row:* simulated observation including all expected components (a ‘full’ simulation). *Third row:* simulated observation with an intentionally perturbed wavelength solution (‘bad λ ’). *Fourth row:* simulated observation excluding telluric contamination (‘no T’). *Fifth row:* simulated observation excluding a planetary atmosphere (‘no P’).

and data from B-frames. We also tried excluding specific spectral orders in several different permutations, but none of these exclusion tests produced any notable changes in the F-maxima location (only lowering the detection significances).

Instead, we revisited the simulations that we generated in Sect. 5, as we had now determined them to be in relatively good agreement with the real data results. A plot showing our new sets of $K_p - v$ plots can be found in Fig. 11. We first took our simulated observation that is as close as possible to our real WASP-107 b observation; this is referred to as the ‘full’ simulation (as it includes all three stellar, telluric, and planetary components).

Next, considering the challenges we faced in our wavelength alignment (Sect. 3.3), we took the same simulation but perturbed its wavelength scale by Doppler shifting and interpolating fluxes and uncertainties by RVs drawn from a normal distribution ($\sigma_{RV} = 100$ m/s) before creating new $K_p - v$ plots from this new dataset (labelled ‘bad λ ’). By doing the interpolation and Doppler shift after the simulation had been created, noise should be constant between this perturbed simulation and the original ‘full’ simulation.

We then created two new simulated observations: one that did not include any telluric lines but maintained all other parameters, and one that did not include planetary signal (bare rock) but maintained all other parameters. The plots from these simulations are labelled ‘no T’ and ‘no P’ for no tellurics and no planetary atmosphere respectively. Together, these two final simulations should test if any of the many manifestations, off-sets, and other sources of spurious signal detailed in the paper – arising from telluric contamination or elsewhere – could combine to create a signal that appears to be seemingly genuine but is in fact not.

Finally, we contrasted all these simulations to what was shown in the real data (‘real’). All datasets in Fig. 11, both real and simulated, are cross-correlated with the CHI+YC template for H₂O; and in the simulations, the simulated atmosphere is the CHI+YC global template.

In Fig. 11, it becomes clear that the ‘chimney smoke’ pattern is present in the full simulation including telluric lines, and that the simulation without tellurics shows no such feature. In the full simulation, it also becomes obvious that the effect is not actually an extension of the signal itself, but rather its own region of increased signal that happens to be located relatively near the true detection maximum in the expected location. In the real data, the small gap between the ‘smoke’ and the true detection becomes less pronounced or vanishes, possibly due to noise properties of the data. One also notices that the F-maximum (denoted by the dashed yellow lines) is closer to the expected location (the dotted white lines) in the no T-simulation than in the full simulation, but that the shift in K_p is smaller for both simulations compared to the real data.

Our bad λ -simulation shows that a poor wavelength alignment will indeed reduce our detection significance by at least ~ 0.5 – 1σ without changing its location. Interestingly enough, the chimney smoke effect also appears reduced here. This could be interpreted as the chimney smoke being in some way a manifestation of real signal, but it is not clear how this would combine with the finding above that removing telluric lines also reduces the presence of chimney smoke.

In our no P-simulation, we find that no signal of significant S/N – or chimney smoke effect – can be detected anywhere near the expected location, indicating that our signal found along $v = 0$ km/s is indeed genuine. The F-maximum of this map reaches towards, but is never at or above, 4.0σ – but importantly, that signal is clearly from a manifestation that is significantly removed from the expected location in v -space. In our $K_p - v$ map, the no P-case never results in any signal above 2.0σ within ± 40 km/s of $v = 0$ km/s for any of the SYSREM iterations shown. In contrast, for all our simulations containing planetary features (and for our real data), the signal always appears within approx. ± 5 km/s of v (and most often at 0–1 km/s). This significantly supports our interpretation that the signal in our real data is in fact genuine.

Out of concern for the spurious $>3\sigma$ signal appearing in our no P-simulation, we then created three more such simulations




Fig. 12. Simulated $K_p - v$ plots for SYSREM iterations 3–10 cross-correlating with the H_2O template including clouds. In each plot, the dotted white line denotes the expected location of the peak, and the dashed yellow line denotes the actual location of the peak. All simulated observations include tellurics. *Top row:* cross-correlation with a template with equilibrium temperature $T_{\text{eq}} = 1200$ K (‘hot’). *Middle row:* cross-correlation with a template with equilibrium temperature $T_{\text{eq}} = 738$ K (‘true’). *Bottom row:* cross-correlation with a template with equilibrium temperature $T_{\text{eq}} = 300$ K (‘cold’).

of a bare rock, i.e. only modelling a transiting sphere of radius R_p with no wavelength-dependent absorption. We find that with each new simulation, and thus each new noise realisation, the location of this artefact (and often-times multiple artefacts) changes location and strength at random, appearing minimally at 2.6σ and maximally at 3.9σ (average and median both 3.5) across the different bare rock simulations for $\text{SYSREM} = 3$ to $\text{SYSREM} = 10$. In terms of displacement in the v -space, artefacts appear between $\sim \pm 30$ – 120 km/s from $v = 0$, with their respective strengths varying across SYSREM iterations. This is to be expected considering that, as explored in detail in Boldt-Christmas et al. (2024) and as mentioned in Sect. 5.1, there is evidence of the precedent that different noise realisations can have major impacts. Nonetheless, the persistence of the planet signal at (roughly) the expected location in all other simulations, aligned with what we see in the real data, gives us confidence that the potential for spurious signals there do not give us less confidence in our planet detection.

Following these findings, we considered that the displaced chimney smoke effect could perhaps be related to the relative line strengths of the molecular bands in the H_2O template. This should differ for different values of equilibrium temperature by changing the level populations of water molecules and thus the relative strength of different groups of lines. We therefore created two more sets of $K_p - v$ maps from the full simulation, this time using new generations of the CHI+YC template that now changed T_{eq} from 738 K to $T_{\text{eq}} = 1200$ K and $T_{\text{eq}} = 300$ K (i.e. approx. ± 450 K from the initial T_{eq}). The result of cross-correlating the same simulated observation with these templates of different T_{eq} values (labelled ‘hot’ for 1200 K, ‘true’ for 738 K, and ‘cold’ for 300 K) are shown in Fig. 12. In this plot, it can be seen that an incorrect template clearly affects both detection significance and K_p of the F-maximum, where the expected $K_p - v$ maximum and the F-maximum are once again denoted by dotted white lines and dashed yellow lines respectively. Here, the effect is somewhat more pronounced and produces an F-maximum at higher K_p for the hotter template, and is less pronounced and at lower K_p for the cold template. The F-maxima for the hot template peak at earlier SYSREM iterations, and for the cold template at later ones, but both have lower significance than with the true template.

7. Results and discussion

In this section we summarise our results and provide some further context, hypotheses, and proposed interpretations of what

has been demonstrated in this study. The regime of relatively cool gas giants is clearly understudied compared to hot and ultra-hot Jupiters (as discussed in Sect. 5), and this is even more true for HRCCS analyses. In this study, a dedicated data reduction effort was required for handling super-resolution (Sect. 3.3), and the use of simulations informed our interpretation (Sects. 5 and 6.2), with the combination of both ultimately leading to our detection of cross-correlation signal. Only with multiple previous studies having repeatedly confirmed the presence of species such as H_2O , essentially guaranteeing its existence somewhere in the data, and with reassurances from our simulations that the detection should not be particularly high even under ideal circumstances, was there an incentive to explore this dataset more thoroughly. As such, one of the key results of this paper is a firm recommendation that authors of other similar, future studies try to reproduce this environment of strong guidance from simulations and previous literature if attempting to study targets whose signal strengths are anticipated to be relatively low for whatever reason (e.g. cooler, cloudier, smaller, or some combination).

Through this process, we arrive at the reporting of two individual species, CO and H_2O at detection significances of $\sigma \gtrsim 4.5$, in the atmosphere of WASP-107 b using VLT/CRIRES+ data near the expected $K_p - v$ location. This is in line with previous space-based studies that have also found these species to be present, and with these detections, we confirm that transmission HRCCS is indeed able to detect molecular species in warm planetary atmosphere regimes of $T_{\text{eq}} < 800$ K. While we are not able to identify any other individual species, we also confirm that HRCCS using templates including all six species based on the parameters of Welbanks et al. (2024) provides a signal of $\sigma \sim 6$, with the maxima found near the expected $K_p - v$ location but at an off-set in the K_p direction. Our results also indicate a non-detection of NH_3 despite this species being confirmed by previous studies; however, this is not particularly surprising considering its low amount of spectroscopic signal at our wavelength range.

It is important to underline the context in which these detections are claimed. Under other circumstances, for example, if WASP-107 b had been a less studied target for which no previous detections had been made, our detections would have lacked the necessary context to favour one interpretation over others. However, thanks to the wealth of existing space-based literature to build upon, these species have already been detected in this planet multiple times so their physical existence is known a priori. Combining this with our simulations demonstrating

that our detection significances are in line with expectations, we do consider these detections to be genuine, even if caution is advised against showing the same degree of confidence for lesser studied targets.

With the previous space-based literature providing many atmospheric parameters, we shifted our focus towards investigating the sensitivity of our detections to the properties of the template. In particular, we compared the impact of the two forward models from Welbanks et al. (2024) and the inclusion of clouds. We discovered that clouds play a much more significant role in HRCCS, while different forward models have only a marginal effect. This points to the importance of combining broadband low-resolution spectra from space with high-resolution ground-based observations to get the full picture, echoing the sentiments in Sect. 1 on the benefits of the two types of data.

Proving the success of HRCCS for cooler targets is crucial to the future of the field, considering that high-resolution spectroscopy is our best method for analysing atmospheric circulation such as winds and jet streams. In light of the great interest within the community regarding the dynamics of atmospheres of planets such as WASP-107 b, honing our ability to analyse exoplanetary atmospheric spectra at high spectral resolution for cooler targets is the key to future success in this research area. This study provides important insight to some of the obstacles and challenges that come with this novel and thus relatively unexplored regime, so the methodology and lessons learned from this work will hopefully be of use to such future studies.

The obstacle that we are unable to describe fully in this paper is the deviation of the F-maximum from the expected location (Sect. 6.2). Our interpretation is not that this is a physical offset where the planetary signal peaks at a significantly different K_p , but rather that this is an area of additional higher S/N above the expected location in the K_p direction of our $K_p - v$ plots. Future work will be required to fully understand and explain its origin, but considering its partial appearance in our simulations, it can be speculated that it must arise due to a combination of both intrinsic and extrinsic reasons. Below, we discuss some primary suspects based on findings and tests so far:

(i) *Poor choice of templates*: the simulations used to test the origin of the ‘chimney smoke’ artefact show that using an incorrect atmospheric model may create trails in the $K_p - v$ plots very similar to what we get from the real data (see Fig. 12). This could be an indicator of a problem arising from our choice of planetary atmospheric model used to generate our templates, from which there are many parameters that may be impactful beyond the equilibrium temperature tested in Sect. 6.2. This question certainly requires further investigation before we embark on more HRCCS analysis of cooler gas giants. This is especially true considering our very simplistic treatment of models, and the clear impact of template selection seen in the simulated tests in Sect. 5, and the fact that the clouds of WASP-107 b are still being researched. Considering our finding that the inclusion of clouds is arguably the parameter that has the most striking impact on our detection significance, it can be expected that more accurate cloud models may improve detection significance further, although the extent to which this is true will presumably vary for different targets and wavelength ranges.

(ii) *Residual tellurics*: our simulations demonstrated that the K_p artefact could also be caused by the presence of tellurics (Fig. 11). It is reasonable to assume that SYSREM residuals remain correlated with tellurics, causing a similar effect in the real data analysis. This is further supported by the fact that the artefact is most pronounced in the cross-correlation with the

H₂O nC templates (which are a close match to the water vapour lines in the Earth atmosphere) and by Gandhi et al. (2020) who also finds that H₂O detections with HRCCS should be diminished due to telluric absorption obscuring the signal for cloudy cases in particular. In reality, the artefact may be caused by the combination of erroneous template (based on overly simplistic atmospheric model) and SYSREM residuals. Stellaric contamination and the challenges of its removal is a well-known obstacle at infrared wavelengths (see e.g. Maguire et al. 2024, and discussions within), so it is not surprising that this may have been done imperfectly – or at least imperfectly to a degree that perhaps would have been acceptable or negligible in a hot Jupiter study, but not for this target where the error budget is significantly depleted as our detection is closer to the noise floor.

(iii) *SYSREM performance across spectral orders*: across our wavelength range, telluric contamination changes dramatically in strength and origin, where the middle spectral orders have relatively little tellurics. H₂O is present everywhere, but changes its strength dramatically, while CO₂ and CH₄ also contribute significantly to some spectral intervals. Considering these differences across wavelength, different spectral orders will almost certainly require different numbers of SYSREM iterations before they are fully detrended – but at each iteration, SYSREM will not have fully removed tellurics in certain spectral intervals, while it will start ‘eating’ into planetary signal in the others. In the future, a robust estimate of convergence for SYSREM should be developed and be ran on a subset of wavelength points that are expected to behave similarly (i.e. points mostly affected by only one species), as the amount of telluric absorption may change throughout a transit due to changes in humidity and temperature.

(iv) *Wavelength solution*: the wavelength solution for CRILES⁺ is not perfect, and even the fine tuning described in Sect. 3.3 may still have some (small) contribution to the artefact that we see in our results – although our test exploring this in Fig. 11 indicates that a poor wavelength solution should only impact our detection significance, not necessarily create offsets. Future work should investigate the impact of this on our analysis, both in the context of SYSREM performance and possible impact on evaluating CCF.

(v) *Physical effects*: it is worth noting that there are also physical reasons to expect a deviation of the maximum S/N in a $K_p - v$ plot. In Wardenier et al. (2023), the location of maxima in $K_p - v$ plots for hot Jupiters are explored, and they find that several atmospheric effects such as morning-to-evening limb variation and cloud decks can create offsets of up to ~20 km/s due to the combination of several effects arising from the 3D nature of the exoplanet atmosphere. In this work, they underline the fact that a K_p offset of a species merely reflects the rate of change of its Doppler shift in the planetary rest frame, and as such, ΔK_p can be quite large compared to Δv . The relatively small K_p of a smaller planet similar to WASP-107 b (compared to hot Jupiters) also creates a challenge as the lower the velocity of the planet, the harder it is to distinguish the planetary signal from stellaric lines, compounding the concern raised in point (iii) above. As such, it can be rationalised that at least a partial offset in the peak signal is to be expected, but further work will be required to differentiate this from other possible effects – especially considering that these effects should not be significant enough to explain an offset of our magnitude on their own.

For all of the above, there is a thread of commonality in that these are all aspects that will realistically exist in all high resolution transmission spectroscopy studies to some extent – but at relatively low S/N where each of these effects may only

impact the result by $\lesssim 2\sigma$, meaning they should be present but not decisive for many studies. The difference is that for this target and observation, only a relatively low maximum S/N can be anticipated even under ideal circumstances as demonstrated by our simulations. As such, for targets such as WASP-107 b, the error budget is too limited for these effects to be affordably overlooked, so they must be accounted for – at least qualitatively. This finding is also applicable to observations of hotter (or otherwise more favourable) targets in the search for species of fainter signals, where the same factors become relevant. The impact of a vast range of uncertainties in HRCCS is thoroughly explored in Savel et al. (2025), and the interplay between HRCCS detections and laboratory data (line lists) is examined in a recent review by Yurchenko et al. (2025), with many previous studies already acknowledging the many challenges associated with smaller, cooler, and/or cloudy planets in particular (e.g. Mollière & Snellen 2019; Hood et al. 2020; Finnerty et al. 2023; Dubey et al. 2025). In practice, a number of recent attempts at transmission HRCCS for smaller planets resulting in non-detections also shed light on the challenges involved, such as Grasser et al. (2024), Dash et al. (2024), and Parker et al. (2025).

It has also been found that detection significances in HRCCS are broadly speaking punished in the use of templates that do not match the reality of the target atmosphere. However, not all parameters are punished (or rewarded) equally, and this work finds that certain parameters such as VMRs and thus mean molecular weight are less impactful on the cross-correlation detection than, for example, correctly accounting for cloud inclusion and temperature T_{eq} . This should be considered in work on the metallicity-cloud degeneracy, as this finding implies that cross-correlating with the two types of template for a target where this degeneracy is unresolved may help us determine which is a better fit and therefore reality. These factors appear to be somewhat species-dependent, reflecting the diversity of possible spectral bands, and will thus also vary across exact wavelength ranges. Considering that T_{eq} is largely calculated from the radius of the host star, and considering that stellar radius is difficult to constrain for cooler stars such as M dwarfs (e.g. Shields et al. 2016; Parsons et al. 2018; Cassisi & Salaris 2019), this consideration is particularly important for planets around cooler stars. In transmission especially, this needs to be accounted for carefully against the concern of spot-crossings, which is inherently more of a concern in cooler stars and has already been found to be a concern for this target by Murphy et al. (2025).

8. Conclusions

In this study, we have characterised the atmosphere of the warm Neptune-like exoplanet WASP-107 b using two transit observations from VLT/CRIRES⁺ in the K band. We used cross-correlation to confirm the detection of two individual species, CO at $\sim 6\sigma$ and H₂O at $\sim 4.5\sigma$, in WASP-107 b within a reasonable distance (13 km/s) of its expected location in the $K_p - v$ detection map. We confirm that the global transmission spectrum as presented by Welbanks et al. (2024) can also be detected at $\sim 6\sigma$ within the same distance from the expected location. We also searched for NH₃ but did not detect it ($\sim 1-2\sigma$). Through the use of simulations, we demonstrated that these findings are in line with the expected detection significances for these species, for this target, and for these observations. These detections represent the first molecular detections made using HRCCS with transmission spectra for a target of $T_{\text{eq}} < 800$ K, thus showing the community moving towards the characterisation of smaller, cooler exoplanets.

Through both simulations and analysis of real data, we demonstrated that our HRCCS analysis is sensitive to the inclusion of a cloud deck in our cross-correlation templates and to the equilibrium temperature of the template, but it is not particularly sensitive to the exact VMR of different species from space-based retrievals. This is true even for the maximum difference in VMR between the two forward models we tested, which is 1.1 dex for CO. We also find that our maximum $K_p - v$ peak deviates from its expected K_p location, which we speculate arises from a combination of possible reasons, such as insufficient removal of the stellar and telluric signal and the model template equilibrium temperature. Our interpretation is that this is not a true offset but rather an artefact of our HRCCS analysis. Understanding its exact nature will require further work.

As the combination of new instrumentation and maturing methodology brings research further into the realm of characterising smaller and cooler exoplanets with high-resolution spectroscopy, the parameter space of possible exoplanetary atmospheres expands – and with it the parameter space of uncertainties. This diversification marks exciting new scientific frontiers, but associated with this reduced error budget are many sources of noise and uncertainty that graduate from being negligible to impactful. Only by mapping out these pitfalls may we avoid them, and only then will we continue our progression towards characterising even cooler and smaller exoplanet targets with high-resolution transmission spectroscopy.

Acknowledgements. We thank Bengt Edvardsson for computing our custom model atmosphere of the star WASP-107. L.B.-Ch., A.D.R., and N.P. acknowledge support by the Knut and Alice Wallenberg Foundation (grant 2018.0192). A.D.R. acknowledges support from ANID/Fondo ALMA 2024/No. 31240064. F.L. acknowledges the support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project number 31466515. D.C. is supported by the LMU-Munich Fraunhofer-Schwarzschild Fellowship and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2094 – 390783311. O.K. acknowledges support by the Swedish Research Council (grant agreement no. 2023-03667) and the Swedish National Space Agency. U.H. acknowledges support from the Swedish National Space Agency (SNSA/Rymdstyrelsen). M.R. acknowledges the support by the DFG priority program SPP 1992 “Exploring the Diversity of Extrasolar Planets” (DFG PR 36 24602/41). D.S. acknowledges financial support from the project PID2021-126365NB-C21(MCI/AEI/FEDER, UE) and from the Severo Ochoa grant CEX2021-001131-S funded by MCIN/AEI/10.13039/501100011033. E.N. acknowledges the support from the Deutsches Zentrum für Luft- und Raumfahrt (DLR, German Aerospace Center) – project number 50OP2502. Based on observations collected at the European Organisation for Astronomical Research in the Southern Hemisphere under ESO programme 108.C-0267(D) and 110.C-4127(D). CRIRES+ is an ESO upgrade project carried out by Thüringer Landessternwarte Tautenburg, Georg-August Universität Göttingen, and Uppsala University. The project is funded by the Federal Ministry of Education and Research (Germany) through Grants 05A11MG3, 05A14MG4, 05A17MG2 and the Knut and Alice Wallenberg Foundation. This work was co-funded by the European Union (ERC-CoG, EVAPORATOR, Grant agreement No. 101170037). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. This research has made use of the NASA Exoplanet Archive, which is operated by the California Institute of Technology, under contract with the National Aeronautics and Space Administration under the Exoplanet Exploration Program. This work has also made use of EsoRex (ESO CPL Development Team 2015) and the following Python packages: Astropy (Astropy Collaboration 2013), iPython (Perez & Granger 2007), Matplotlib (Hunter 2007), NumPy (Harris et al. 2020), Pandas (McKinney 2010; team 2023), PyAstronomy (Czesla et al. 2019), SciPy (Virtanen et al. 2020), and Seaborn (Waskom et al. 2017). Finally, we wish to thank the anonymous referee of this paper, whose questions and suggestions has significantly improved this manuscript.

References

- Allart, R., Bourrier, V., Lovis, C., et al. 2018, *Science*, **362**, 1384
- Allart, R., Bourrier, V., Lovis, C., et al. 2019, *A&A*, **623**, A58
- Allart, R., Pino, L., Lovis, C., et al. 2020, *A&A*, **644**, A155

- Anderson, D. R., Cameron, A. C., Delrez, L., et al. 2017, *A&A*, **604**, A110
- Astropy Collaboration (Robitaille, T. P., et al.) 2013, *A&A*, **558**, A33
- Basilicata, M., Giacobbe, P., Bonomo, A. S., et al. 2024, *A&A*, **686**, A127
- Birkby, J. L., de Kok, R. J., Brogi, M., et al. 2013, *MNRAS*, **436**, L35
- Boldt-Christmas, L., Lesjak, F., Wehrhahn, A., et al. 2024, *A&A*, **683**, A244
- Bourrier, V., Attia, M., Mallonn, N., et al. 2023, *A&A*, **669**, A63
- Bowesman, C. A., Shuai, M., Yurchenko, S. N., & Tennyson, J. 2021, *MNRAS*, **508**, 3181
- Brogi, M., de Kok, R. J., Birkby, J. L., Schwarz, H., & Snellen, I. A. G. 2014, *A&A*, **565**, A124
- Cassisi, S., & Salaris, M. 2019, *A&A*, **626**, A32
- Changeat, Q., Bardet, D., Chubb, K., et al. 2025, *A&A*, **699**, A219
- Cheverall, C. J., & Madhusudhan, N. 2024, *AJ*, **167**, 272
- Cheverall, C. J., Madhusudhan, N., & Holmberg, M. 2023, *MNRAS*, **522**, 661
- Cont, D., Nortmann, L., Yan, F., et al. 2024, *A&A*, **688**, A206
- Cont, D., Nortmann, L., Lesjak, F., et al. 2025, *A&A*, **698**, A31
- Cutri, R. M., Skrutskie, M. F., van Dyk, S., et al. 2003, *VizieR Online Data Catalog*, II/246
- Czesla, S., Schröter, S., Schneider, C. P., et al. 2019, PyA: Python astronomy-related packages [*record ascl:1906.010*]
- Czesla, S., Lampón, M., Cont, D., et al. 2024, *A&A*, **683**, A67
- Dai, F., & Winn, J. N. 2017, *AJ*, **153**, 205
- Dash, S., Brogi, M., Gandhi, S., et al. 2024, *MNRAS*, **530**, 3100
- Dholakia, S., Dholakia, S., & Pope, B. J. S. 2025, *ApJ*, **987**, 150
- Dorn, R. J., Bristow, P., Smoker, J. V., et al. 2023, *A&A*, **671**, A24
- Dressing, C. D., Hardegree-Ullman, K., Schlieder, J. E., et al. 2019, *AJ*, **158**, 87
- Dubey, D., Majumdar, L., Beichman, C., et al. 2025, *ApJS*, **278**, 19
- Dyrek, A., Min, M., Decin, L., et al. 2024, *Nature*, **625**, 51
- Ehrenreich, D., Lovis, C., Allart, R., et al. 2020, *Nature*, **580**, 597
- ESO CPL Development Team 2015, Astrophysics Source Code Library [*record ascl:1504.003*]
- Finnerty, L., Schofield, T., Sappay, B., et al. 2023, *AJ*, **166**, 31
- Foley, B. J. 2024, in *Exoplanets: Compositions, Mineralogy, and Evolution*, Chapter 15 (Reviews in Mineralogy and Geochemistry (RiMG) Volume 90), [*arXiv:2404.15433*] [*astro-ph, physics:physics*]
- Gaia Collaboration (Brown, A. G. A., et al.) 2021, *A&A*, **649**, A1
- Gandhi, S., Brogi, M., & Webb, R. K. 2020, *MNRAS*, **498**, 194
- Gandhi, S., Kesseli, A., Snellen, I., et al. 2022, *MNRAS*, **515**, 749
- Gao, P., Wakeford, H. R., Moran, S. E., & Parmentier, V. 2021, *J. Geophys. Res. (Planets)*, **126**, e06655
- Grasser, N., Snellen, I. A. G., Landman, R., Picos, D. G., & Gandhi, S. 2024, *A&A*, **688**, A191
- Grevesse, N., Asplund, M., & Sauval, A. J. 2007, *Space Sci. Rev.*, **130**, 105
- Guillot, T. 2010, *A&A*, **520**, A27
- Guillot, T., Fletcher, L. N., Helled, R., et al. 2022, Giant Planets from the Inside-Out [*arXiv:2205.04100*] [*astro-ph*]
- Guilluy, G., D'Arpa, M. C., Bonomo, A. S., et al. 2024, *The GAPS Programme at TNG: LV. A HeI survey of close-in giant planets hosted by M-K dwarf stars with GIANO-B*
- Gustafsson, B., Edvardsson, B., Eriksson, K., et al. 2008, *A&A*, **486**, 951
- Hargreaves, R. J., Gordon, I. E., Rey, M., et al. 2020, *ApJS*, **247**, 55
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, **585**, 357
- Hedges, C., & Madhusudhan, N. 2016, *MNRAS*, **458**, 1427
- Hejazi, N., Crossfield, I. J. M., Nordlander, T., et al. 2023, *ApJ*, **949**, 79
- Hoeijmakers, H. J., Kok, R. J. d., Snellen, I. a. G., et al. 2015, *A&A*, **575**, A20
- Hong, K. S., Finnerty, L., & Fitzgerald, M. P. 2025, Velocity shift and SNR limits for high-resolution spectroscopy of hot Jupiters using Keck/KPIC [*arXiv:2505.09781*] [*astro-ph*] version: 1
- Hood, C. E., Fortney, J. J., Line, M. R., et al. 2020, *AJ*, **160**, 198
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- JWST Transiting Exoplanet Community Early Release Science Team (Ahrer, E.-M., et al.) 2023, *Nature*, **614**, 649
- Kataria, T., Sing, D. K., Lewis, N. K., et al. 2016, *ApJ*, **821**, 9
- Kesseli, A. Y., Snellen, I. A. G., Alonso-Floriano, F. J., Mollière, P., & Serindag, D. B. 2020, *AJ*, **160**, 228
- Khodachenko, M. L., Shaikhislamov, I. F., Fossati, L., et al. 2021, *MNRAS*, **503**, L23
- Kirk, J., Alam, M. K., Lopez-Morales, M., & Zeng, L. 2020, *AJ*, **159**, 115
- Kokori, A., Tsaras, A., Edwards, B., et al. 2023, *ApJS*, **265**, 4
- Kreidberg, L., Line, M. R., Thorngren, D., Morley, C. V., & Stevenson, K. B. 2018, *ApJ*, **858**, L6
- Krishnamurthy, V., Carteret, Y., Piaulet-Ghorayeb, C., et al. 2025, Continuous helium absorption from the leading and trailing tails of WASP-107b, [*arXiv:2505.20588*] [*astro-ph*]
- Lafarga, M., Brogi, M., Gandhi, S., et al. 2023, *MNRAS*, **521**, 1233
- Lavvas, P., & Koskinen, T. 2017, *ApJ*, **847**, 32
- Lesjak, F., Nortmann, L., Yan, F., et al. 2023, *A&A*, **678**, A23
- Lesjak, F., Nortmann, L., Cont, D., et al. 2025, *A&A*, **693**, A72
- Linssen, D. C., Oklopčić, A., & MacLeod, M. 2022, *A&A*, **667**, A54
- Lucy, L. B., & Sweeney, M. A. 1971, *AJ*, **76**, 544
- Maguire, C., Sedaghati, E., Gibson, N. P., Smette, A., & Pino, L. 2024, *A&A*, **692**, A8
- Maire, A.-L., Delrez, L., Pozuelos, F. J., et al. 2023, *PASP*, **135**, 106001
- McKinney, W. 2010, in *SciPy Proceedings*, Austin, Texas, 56
- McLaughlin, D. B. 1924, *ApJ*, **60**, 22
- Meech, A., Aigrain, S., Brogi, M., & Birkby, J. L. 2022, *MNRAS*, **512**, 2604
- Miguel, Y., & Kaltenegger, L. 2014, *ApJ*, **780**, 166
- Millholland, S., Petigura, E., & Batygin, K. 2020, *ApJ*, **897**, 7
- Mollière, P., & Snellen, I. A. G. 2019, *A&A*, **622**, A139
- Mollière, P., Wardenier, J. P., Boekel, R. v., et al. 2019, *A&A*, **627**, A67
- Močnik, T., Hellier, C., Anderson, D. R., Clark, B. J. M., & Southworth, J. 2017, *MNRAS*, **469**, 1622
- Mukherjee, S., Schlawin, E., Bell, T. J., et al. 2025, *ApJ*, **982**, L39
- Murphy, M. M., Beatty, T. G., Schlawin, E., et al. 2024, *Nat. Astron.*, **1**
- Murphy, M. M., Beatty, T. G., Schlawin, E., et al. 2025, *AJ*, **170**, 61
- Nettelmann, N., & Valencia, D. 2021, in *ExoFrontiers: Big Questions in Exoplanetary Science* (IOP Publishing)
- Nortmann, L., Pallé, E., Salz, M., et al. 2018, *Science*, **362**, 1388
- Nortmann, L., Lesjak, F., Yan, F., et al. 2025, *A&A*, **693**, A213
- Palle, E., Biazzo, K., Bolmont, E., et al. 2025, *Exp. Astron.*, **59**
- Parker, L. T., Mendonça, J. M., Diamond-Lowe, H., et al. 2025, *MNRAS*, **538**, 3263
- Parsons, S. G., Gänsicke, B. T., Marsh, T. R., et al. 2018, *MNRAS*, **481**, 1083
- Perez, F., & Granger, B. E. 2007, *Comput. Sci. Eng.*, **9**, 21
- Piaulet, C., Benneke, B., Rubenzahl, R. A., et al. 2021, *AJ*, **161**, 70
- Pinhas, A., & Madhusudhan, N. 2017, *MNRAS*, **471**, A355
- Pino, L., Ehrenreich, D., Allart, R., et al. 2018, *A&A*, **619**, A3
- Piskunov, N., Rains, A. D., & Boldt-Christmas, L. 2025, TSD: An inverse problem approach for recovering the exoplanetary atmosphere transmission spectrum from high-resolution spectroscopy, [*arXiv:2509.12737*]
- Polyansky, O. L., Kyuberis, A. A., Zobov, N. F., et al. 2018, *MNRAS*, **480**, 2597
- Powell, D., Loudon, T., Kreidberg, L., et al. 2019, *ApJ*, **887**, 170
- Prineth, B., Hoeijmakers, H. J., Kitzmann, D., et al. 2022, *Nat. Astron.*, **6**, 449
- Prineth, B., Sedaghati, E., Seidel, J. V., et al. 2024, *AJ*, **168**, 133
- Rengel, M. 2022, *Proc. Int. Astron. Union*, **18**, 87
- Rossiter, R. A. 1924, *ApJ*, **60**, 15
- Rothman, L. S., Gordon, I. E., Barber, R. J., et al. 2010, *JQSRT*, **111**, 2139
- Rothman, L. S., Gordon, I. E., Babikov, Y., et al. 2013, *JQSRT*, **130**, 4
- Rubenzahl, R. A., Dai, F., Howard, A. W., et al. 2021, *AJ*, **161**, 119
- Ryabchikova, T., Piskunov, N., Kurucz, R. L., et al. 2015, *Phys. Scr.*, **90**, 054005
- Savel, A. B., Bedell, M., Kempton, E. M. R., et al. 2025, *AJ*, **169**, 135
- Schlawin, E., Greene, T. P., Line, M., Fortney, J. J., & Rieke, M. 2018, *AJ*, **156**, 40
- Schwarz, H., Brogi, M., de Kok, R., Birkby, J., & Snellen, I. 2015, *A&A*, **576**, A111
- Seidel, J. V., Ehrenreich, D., Pino, L., et al. 2020, *A&A*, **633**, A86
- Seidel, J. V., Prineth, B., Pino, L., et al. 2025, *Nature*, **639**, 902
- Shen, Y., & Turner, E. L. 2008, *ApJ*, **685**, 553
- Shields, A. L., Ballard, S., & Johnson, J. A. 2016, *Phys. Rep.*, **663**, 1
- Sing, D. K., Rustamkulov, Z., Thorngren, D. P., et al. 2024, *Nature*, **630**, 831
- Smette, A., Sana, H., Noll, S., et al. 2015, *A&A*, **576**, A77
- Smith, P. C. B., Line, M. R., Bean, J. L., et al. 2024, *AJ*, **167**, 110
- Snellen, I. 2025, Exoplanet atmospheres at high spectral resolution, [*arXiv:2505.08926*]
- Soni, V., & Acharyya, K. 2024, *ApJ*, **977**, 52
- Spake, J. J., Sing, D. K., Evans, T. M., et al. 2018, *Nature*, **557**, 68
- Spake, J. J., Oklopčić, A., & Hillenbrand, L. A. 2021, *AJ*, **162**, 284
- Tamuz, O., Mazeh, T., & Zucker, S. 2005, *MNRAS*, **356**, 1466
- Tannock, M. E., Metchev, S., Hood, C. E., et al. 2022, *MNRAS*, **514**, 3160
- team, T. p. d. 2023, pandas-dev/pandas: Pandas
- Teske, J., Batalha, N. E., Wallack, N. L., et al. 2025, *AJ*, **169**, 249
- Tsai, S.-M., Lee, E. K. H., Powell, D., et al. 2023, *Nature*, **617**, 483
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Methods*, **17**, 261
- Wang, L., & Dai, F. 2021, *ApJ*, **914**, 99
- Wardenier, J. P., Parmentier, V., Line, M. R., & Lee, E. K. H. 2023, *MNRAS*, **525**, 4942
- Waskom, M., Botvinnik, O., O'Kane, D., et al. 2017, Zenodo, aDS Bibcode: 2017zndo....883859W
- Welbanks, L., Bell, T. J., Beatty, T. G., et al. 2024, *Nature*, **630**, 836
- Yan, F., Nortmann, L., Reiners, A., et al. 2023, *A&A*, **672**, A107
- Yurchenko, S. N., Barber, R. J., & Tennyson, J. 2011, *MNRAS*, **413**, 1828
- Yurchenko, S. N., Tennyson, J., & Brogi, M. 2025, *Nat. Rev. Phys.*, **1**
- Zakamska, N. L., Pan, M., & Ford, E. B. 2011, *MNRAS*, **410**, 1895