High-resolution spectroscopy for high-contrast imaging of exoplanets with ELT instruments

Sebastian Magnusson

Division of Astrophysics
Department of Physics

Lund University

2026-EXA253 FYSK04

Degree project of 15 higher education credits
June 2026

Supervisor: Adrien Simonnin
Co-supervisor: Jens Hoeijmakers

Division of Astrophysics
Department of Physics
Box 118
SE-221 00 Lund
Sweden

Abstract

Combining high-contrast imaging (HCI) with high-resolution spectroscopy (HRS) is one of the leading strategies for the next generation of direct-imaging instruments on the Extremely Large Telescope (ELT). HCI suppresses the host star's light spatially through adaptive optics and coronagraphy, while HRS suppresses it spectrally by exploiting the Doppler-shifted molecular lines in the planet's atmosphere. In principle, the two methods multiply, and a joint contrast gain of several orders of magnitude beyond either alone has been predicted. This is the design baseline of the ELT instruments ANDES and PCS.

This thesis develops an end-to-end simulation pipeline to quantify that gain for self-luminous giant exoplanets observed with ANDES on the ELT. The pipeline generates stellar and planetary spectra from physical models, propagates them through a simplified on-axis/off-axis differential observation including coronagraphic suppression and detector noise, and retrieves the planetary signal via cross-correlation against a noiseless template.

The pipeline is applied to four self-luminous companions with $T_{\mathrm{eff}} = 400 - 1600\,\mathrm{K}$ around a common early-F host star, across angular separations of $20 - 79\,\mathrm{mas}$, lower and upper halves of the $0.95 - 1.80\,\mu\mathrm{m}$ wavelength window, and resolving powers $R = 5 \times 10^4 - 1.9 \times 10^5$. At the reference configuration the cross-correlation deepens the accessible contrast by roughly three orders of magnitude in contrast relative to the raw coronagraphic contrast. Detection distances up to $\sim 340\,\mathrm{pc}$ are found for the hottest companion, and up to $\sim 15\,\mathrm{pc}$ for the coolest. The lower half of the wavelength window drives most of the detection for the cooler companions, while the resolving-power dependence is weak.

Populärvetenskaplig beskrivning

Någonstans där ute, i bana kring en avlägsen stjärna, kan det finnas en planet med hav, moln och en atmosfär fylld med de kemiska fingeravtrycken av liv. Vi vet inte än. men inom de kommande decennierna kanske vi får svaret.

Sedan 1995 har astronomer hittat fler än sex tusen planeter kring andra stjärnor. Nästan alla har dock upptäckts indirekt. Att faktiskt direkt avbilda en planet, och framför allt studera dess atmosfär, har visat sig betydligt svårare. Föreställ dig att från flera kilometers håll försöka se en eldfluga som flyger bredvid en fyr. Det är ungefär den utmaning som astronomer står inför. En ung Jupiter, en av de enklaste måltavlorna, är ungefär en miljon gånger ljussvagare än sin stjärna. En jordlik planet kan vara en miljard gånger svagare. Stjärnans ljus målar helt enkelt ut allt annat.

Mitt examensarbete utgår från en central fråga: hur kan vi övervinna denna enorma kontrast och faktiskt fånga ljuset från avlägsna planeter? En metod, *högkontrast avbildning*, blockerar stjärnans ljus med hjälp av adaptiv optik och koronagrafer, med andra ord, konstgjorda förmörkelser. En annan metod, *högupplöst spektroskopi*, separerar planet och stjärna i ljus: när planetens ljus delas upp i ett spektrum framträder fingeravtryck av molekyler så som vatten, koldioxid och metan, och eftersom planeten rör sig i bana kring sin stjärna förskjuts dessa fingeravtryck något i våglängd. Var för sig räcker metoderna inte hela vägen fram. Tillsammans kan de däremot göra det.

I projektet bygger jag en simulering av just en sådan kombinerad observation. Jag genererar realistiska planet- och stjärnspektra, lägger till instrumentbrus, och försöker återfinna planetsignalen i bruset. Genom att variera planetens temperatur, observationens våglängdsområde och instrumentets upplösning kartlägger jag hur djupt vi kan se. Arbetet är särskilt relevant för nästa generations Extremely Large Telescope (ELT) instrument ANDES och PCS som just är designade för denna kombination. Detta är ett steg på vägen mot att en dag kunna besvara en av de äldsta frågorna... är vi ensamma?

Contents

Chapter 1

Introduction

1.1 Background and Motivation

Since the discovery of the first planet orbiting a sun-like star in 1995 (Mayor & Queloz 1995), the catalog of confirmed exoplanets has grown to more than 6 200 (Christiansen et al. 2025). The overwhelming majority of these discoveries were obtained through indirect techniques, transit photometry and radial velocity surveys, which infer the presence of a planet from periodic variations in the host star's brightness or radial velocity. Direct imaging, in which photons emitted or reflected by the planet itself are spatially separated from the stellar light, has so far succeeded in fewer than 100 discoveries (Christiansen et al. 2025); this because most planets are too faint to be distinguished from the glare of their host star and thus remains restricted to young, self-luminous giants at relatively wide orbital separations (Chauvin 2024).

Even though direct detection has limited yield, it provides information no other method can. The recovered planetary spectrum encodes the chemical composition, temperature structure, formation history, and ultimately provides a way to assess the habitability of terrestrial worlds through spectral signatures of atoms and molecules (Snellen 2025). Pushing atmospheric characterization toward smaller, cooler, and ultimately Earth-like planets is therefore one of the central goals of the next decades. However, doing so requires reaching planet-to-star contrasts on the order of 10^{-8} in thermal emission and 10^{-9} in reflected light for terrestrial planets (Chauvin 2024), which are several orders of magnitude beyond current facilities capabilities.

Reaching such contrasts motivates the next generation of ground-based observatories, in particular the Extremely Large Telescope (ELT). Two of its instruments are central to this thesis: ANDES (ArmazoNes high Dispersion Echelle Spectrograph), a high-resolution spectrograph designed to combine with high-contrast imaging (Palle et al. 2025), and the planned second-generation PCS (Planetary Camera and Spectrograph), aimed at the contrast regime of terrestrial planets (Kasper et al. 2021). Both rely on combining high-contrast imaging with high-resolution spectroscopy, the technique this work investigates.

1.2 Aim and Research Questions

The primary aim of this project is to investigate the detectability of self-luminous giant exoplanets through the combination of high-contrast imaging (HCI) and high-resolution spectroscopy (HRS) in the context of future ELT instrumentation. In particular, the thesis focuses on evaluating the performance of the instrument ANDES.

- **What HRS contrast gain is achievable, when combined with HCI in the ELT regime?** Theoretical work (e.g., Snellen et al. 2015) predicts that multiplicative combination of HCI and HRS suppression can extend contrast limits by several orders of magnitude relative to imaging alone.

- **How does the choice of wavelength range affect detectability?** Different molecular carriers dominate at different temperatures, and their opacities are not uniformly distributed across the $0.95 - 1.80\,\mu$m window covered by ANDES.

- **Where is the, or is there a, trade-off between resolving power and throughput (the photon flux collected per spectral channel)?** A higher resolving power splits previously blended molecular bands into more individually resolved lines, but at fixed exposure time it also distributes the same photon budget across more spectral channels.

- **How does the planet's spectral content, set primarily by its effective temperature, shape the recovered signal?** Comparing detection thresholds at a common H-band photon budget against those obtained at the native surface flux...

Chapter 2

Theory

As of 2026, more than 6 200 exoplanets have been confirmed (Christiansen et al. 2025) and all discovered through a handful of complementary observational techniques. Figure 2.1 shows the current confirmed exoplanet population split by discovery method. Transit and radial velocities dominate the sample, while direct imaging only accounts for a small fraction of detections. Each technique is sensitive to a different range of planet masses, radii, orbital separations, and host-star types, and also yield a different subset of physical properties.

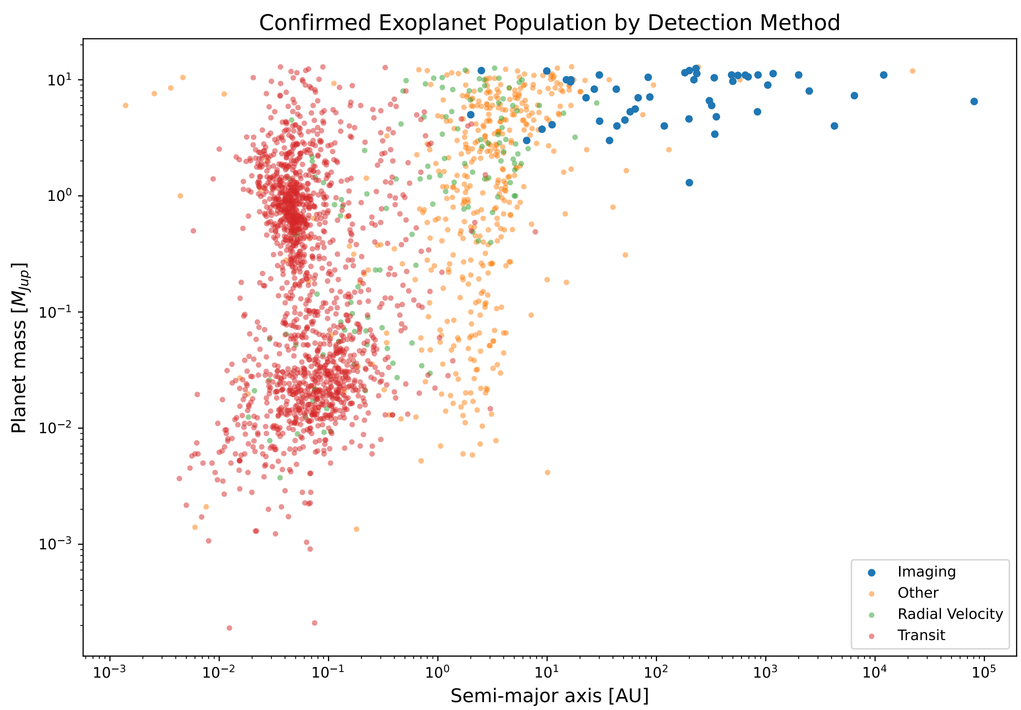

Figure 2.1: The confirmed exoplanet population by detection type (data from exoplanet.eu, 14 May 2026).

2.1 Detection Methods

2.1.1 Radial Velocity

The radial velocity (RV) method exploits the reflex motion of a star around the common center of mass in a star-planet system. As the star moves towards or away from the observer, its spectral lines are periodically Doppler shifted and a planet can be inferred from the resulting velocity variations (Wright & Gaudi 2013). For a planet of mass M_p on a circular orbit of period P and inclination i around a star of mass $M_* \gg M_\mathrm{p}$, the line-of-sight velocity semi-amplitude is

$$K = \left(\frac{2\pi G}{P}\right)^{1/3} \frac{M_\mathrm{p} \sin i}{M_*^{2/3}}.$$

This determines the maximum line-of-sight velocity reached by the star relative to the system's barycentric velocity (Endl 2014).

2.1.2 Transits

Transit photometry detects the small periodic dimming that occurs when a planet's orbit is aligned closely enough to the line of sight that it passes in front of its host star. Therefore, the method preferentially detects planets that are large relative to their host (Seager & Deming 2010; Endl 2014). Its dominance among the confirmed exoplanet population (Figure 2.1) is due to that a single telescope can monitor thousands of stars simultaneously, and the same transit repeats every orbit, making short-period planets easy to confirm (Seager & Deming 2010).

2.1.3 Direct Imaging

Direct imaging detects a planet by spatially separating its photons from those of the host star and recording them directly. Unlike indirect methods, it isolates the planetary signal and provides access to the planet's photometric and spectral properties, enabling atmospheric characterization (Chauvin 2024).

The technique is challenging for two main reasons. First, the contrast, defined as planet-star flux ratio, is large. A young, self-luminous giant emits roughly 10^{-6} of the flux of its host star, and mature Jupiters around 10^{-8} (Chauvin 2024). Secondly, the angular separations involved are small, a Jupiter-like planet at 5 AU seen from 50 pc away would have an angular separation of 100 mas and the habitable zone of a nearby (10 pc) Sun-like star spans only tens of milliarcseconds. Because the contrast and angular separation requirements both relax for young, hot planets on wide orbits, the directly imaged sample to date is strongly biased towards giant planets of a few to tens of Jupiter masses, see Figure 2.1 (Chauvin 2024; Christiansen et al. 2025).

A prototypical directly imaged system is HR 8799. Shown in Figure 2.2 is three of the four self-luminous giant planets found, with masses $7 - 10\,M_\mathrm{J}$ orbiting a young ($\sim 60\,\mathrm{Myr}$) A-type star at projected separations of 24, 38 and 68 AU (Marois et al. 2008).

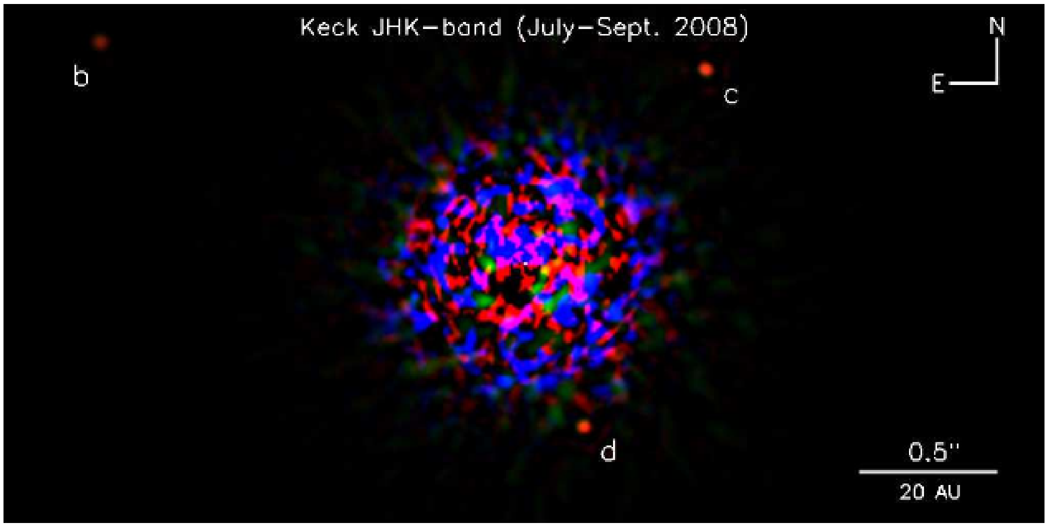

Figure 2.2: The HR 8799 system imaged with Keck II (J/H/Ks composite, ADI-reduced). Three self-luminous giant planets (b, c, d) are resolved around the central A-type host star at projected separations of 24, 38, and 68 AU. Taken from Marois et al. (2008), Figure 1 (bottom panel).

2.2 Observing Modes of Exoplanet Atmospheres

2.2.1 Transmission Spectroscopy

During a primary eclipse, when a planet passes in front of its host star, stellar light passes through the planet's atmosphere resulting in part of the light being absorbed depending on the atmospheric makeup. Differentiating between the in- and out-of-transit spectra isolates features resulting from the stellar light passing through the planet's atmosphere (Seager & Deming 2010; Madhusudhan 2019).

2.2.2 Reflected Light

A planet scatters part of the incident starlight back toward the observer; unlike transmission spectroscopy, where starlight passes through the atmosphere and picks up the planet's own absorption features, the reflected spectrum largely retains the shape of the stellar spectrum, affected by the planet's albedo and by scattering from e.g. clouds(Snellen 2025). The signal is intrinsically small, but the combination of HCI and HRS on the ELT is expected to make it accessible for cooler, more reflective terrestrial planets (Snellen 2025; Palle et al. 2025; Kasper et al. 2021).

2.2.3 Emission Spectroscopy

Emission spectroscopy targets the planet's own thermal radiation. The shape of the spectrum is set by the temperature-pressure profile of the photosphere together with the opacities of the atmospheric gas, see §2.3, (Snellen 2025; Madhusudhan 2019). For transiting planets, the secondary eclipse, when a planet passes behind the host star, provides a measurement by differencing the combined star and planet flux against the star alone while the planet is hidden (Seager & Deming 2010). For spatially resolved systems, direct imaging records the planetary flux without relying on a transit, and is the mode in which young self-luminous giant planets, still hot from the residual heat of formation, are characterized by current and upcoming instruments such as SPHERE, GPI, ANDES, and PCS (Chauvin 2024; Palle et al. 2025).

2.3 Modeling Exoplanet Atmospheres

Interpreting an observed planetary spectrum requires a forward model that, for a given set of parameters such as gravity, effective temperature, and elemental abundances, predicts the emergent spectrum, typically under simplifying assumptions such as a plane-parallel atmosphere in hydrostatic and local thermodynamic equilibrium (Fortney 2018; Madhusudhan 2019). For such a model, three components are needed: a pressure-temperature ($P{-}T$) profile of the atmosphere, a chemistry makeup providing the molecular and atomic abundances at each atmospheric layer, and a radiative transfer scheme. The specific grid and code adopted here are described in §3.2.

2.3.1 Radiative Transfer

Radiative transfer describes how the intensity I_λ of a beam is transformed by absorption and emission as it traverses a layer of gas (Ambaum 2021). For pure absorption, the Beer-Lambert law states that the intensity decreases exponentially along the path,

$$\mathrm{d}I_\lambda = -I_\lambda\,\mathrm{d}\delta_\lambda, \qquad \mathrm{d}\delta_\lambda = \tilde{n}\,\sigma_\lambda\,\mathrm{d}l, \tag{2.1}$$

where $\tilde{n}$ is the number density of absorbers, σ_λ their cross-section, $\mathrm{d}l$ the layer's infinitesimal geometric depth, and $\mathrm{d}\delta_\lambda$ the infinitesimal optical depth at wavelength λ (Ambaum 2021).

A real atmospheric layer does not only absorb but also emits. The thermal emission of a layer at temperature T is set by the Planck function $B_\lambda(T)$. Including both contributions gives the Schwarzschild equation, which describes the radiative transfer through a medium assuming local thermodynamic equilibrium,

$$\frac{\mathrm{d}I_\lambda}{\mathrm{d}\delta_\lambda} = I_{B,\lambda}(T) - I_\lambda, \tag{2.2}$$

with $I_{B,\lambda}(T) = \frac{B_\lambda(T)}{\pi}$ describing the blackbody intensity for temperature T (Ambaum 2021). Equation 2.2 is the equation any atmospheric radiative transfer code must solve given a P-T profile and an opacity structure. `petitRADTRANS` solves this in its line-by-line mode at native resolving power $R = 10^6$ (Mollière et al. 2019); §3.2 describes its subsequent use for this analysis.

2.3.2 Atmospheric Chemistry

The cross-section σ_λ in the Beer-Lambert law is set by the abundances of the atomic and molecular species present in each layer where the $P{-}T$ profile sets the local conditions. Under chemical equilibrium these abundances follow from minimizing the total Gibbs free energy, the thermodynamic potential whose minimum defines equilibrium at fixed pressure and temperature, of the gas at the local pressure and temperature; this is the approach adopted in the `ATMO 2020` grids used as input to the planetary spectra in this work (Madhusudhan 2019; Phillips et al. 2020).

The species set by the chemistry produces opacity in two ways. Line opacity from the rotational, vibrational, and electronic transitions of individual atoms and molecules. Continuum opacity is caused by Rayleigh scattering, the elastic scattering of photons off molecules, which is largely negligible in the near-infrared (Ambaum 2021), and by collision-induced absorption, in which colliding H_2 and He pairs briefly acquire an induced dipole and absorb across broad structureless bands (Terragni et al. 2025).

Figure 2.3 shows what this means in practice. The left panel gives the $P - T$ profile of a $1600\,K$ and a $400\,K$ self-luminous Jupiter at $log\,g = 4$: at fixed gravity the hotter atmosphere is offset toward higher temperature at every depth. The right panel shows the equilibrium volume mixing ratios, the number of molecules of a given species relative to the total number of gas molecules. H_2O is present at both temperatures, but the carbon carrier switches: at $1600\,K$ CO carries essentially all of the carbon, while at $400\,K$ it has converted to CH_4, and NH_3 becomes prominent.

The equilibrium assumption is an idealization. The most important caveat is that vertical mixing in real atmospheres can "quench" the CO - CH_4 conversions at altitudes where the chemistry is slower than the mixing, leaving more CO and less CH_4 and NH_3 than equilibrium predicts (Phillips et al. 2020).

2.3.3 Key Parameters Affecting Spectra

Effective temperature, T_{eff}. Sets the overall flux level through $L = 4\pi R^2 \sigma T_{\text{eff}}^4$. Through the equilibrium chemistry it also controls which carbon carrier is dominant: at high T_{eff} CO dominates as a carbon carrier while at low T_{eff} CH_4 takes over, with H_2O present at all temperatures (Fortney 2018).

Surface gravity, $\log g$. At fixed T_{eff}, higher gravity compresses the atmosphere into a thinner layer, so the photosphere sits at higher pressure. This broadens the alkali (Na, K) line profiles strengthen H_2 collision-induced absorption (Fortney 2018).

Elemental composition. Beyond temperature and gravity, the spectrum depends on what the atmosphere is made of: its overall metal content, conventionally parameterized by the metallicity [Fe/H] relative to the Sun, and the relative abundances of individual elements, most importantly the carbon-to-oxygen ratio C/O (Madhusudhan 2019). Higher

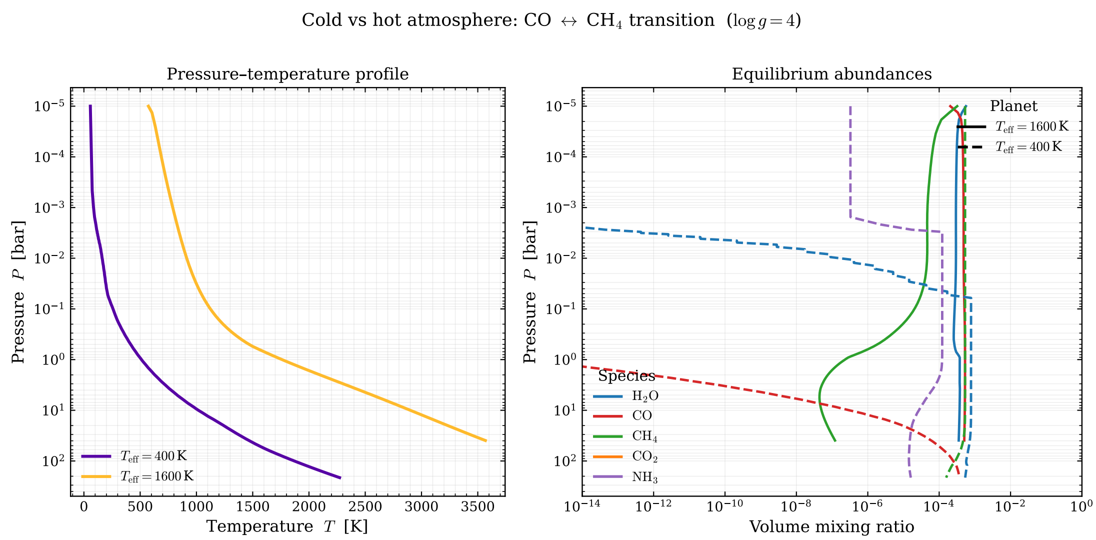

Figure 2.3: Pressure-temperature profiles (left) and equilibrium volume mixing ratios (right) for a self-luminous Jupiter at $T_{\rm eff} = 400\,K$ and $T_{\rm eff} = 1600\,K$, both at $log\,g = 4.0$. Profiles and abundances from the ATMO 2020 grid (Phillips et al. 2020).

metallicity deepens the molecular bands by raising the absorber abundances, while the C/O ratio controls which carbon and oxygen carriers dominate the chemistry.

2.4 High-Contrast Imaging

High-contrast imaging encompasses the instrumental and observational techniques that allow the light of a planet to be spatially separated from that of its host star. As established in §2.1.3, this must overcome both a large planet-to-star flux ratio and small angular separations, constraints that grow even more severe for smaller, terrestrial planets, e.g. in reflected light (Follette 2023; Chauvin 2024). Together this places the planetary signal in the wings of the stellar point-spread function, where it is overwhelmed by stellar photons and their associated noise. Modern HCI addresses this through three stages, each targeting a different aspect of the stellar light contaminating the planet position (Chauvin 2024).

Point spread function. The point spread function (PSF) describes how a point source, such as a star, appears on the detector after passing through the telescope and atmosphere (Davies & Kasper 2012; Kenworthy & Haffert 2025). For an ideal telescope above the atmosphere the PSF is the Airy diffraction pattern: a sharp central core surrounded by faint concentric rings, with an angular width set by the ratio of the observing wavelength to the telescope diameter, FWHM $\simeq 1.03\,\lambda/D$ (Howell 2006). For a ground-based telescope, atmospheric turbulence smears this pattern into a much broader, blurry blob; on a large telescope the result is a PSF up to two orders of magnitude wider than the diffraction-limited core (Davies & Kasper 2012; Chauvin 2024). Any planet hiding within the wings of this stellar PSF is then lost in the glare of its host star (Chauvin 2024).

Adaptive optics. Adaptive optics (AO) sharpens the PSF by removing the atmospheric distortion in real time. A wavefront sensor measures how the incoming wavefront has been deformed by the atmosphere, a fast computer calculates the inverse distortion, and a deformable mirror applies that correction to the beam before it reaches the science instrument (Davies & Kasper 2012; Chauvin 2024). A well-corrected AO image recovers a near-diffraction-limited core, but the correction is never perfect: the residual phase errors left in the wavefront form a granular, time-variable pattern of bright dots around the PSF known as speckles. Speckles look similar to faint companions and are the dominant limitation on how close to the star a planet can be detected (Kenworthy & Haffert 2025).

Coronagraph. Even with a sharp, diffraction-limited PSF, the central core and inner diffraction rings of the star contain enormously more flux than any nearby planet. A coronagraph is a set of optical elements that suppresses this on-axis stellar light while letting the off-axis planet light pass through. The most common design uses an opaque mask placed in the image plane to block the central core of the star, together with a second mask in a re-imaged pupil that blocks the diffracted light leaking around the first mask (Kenworthy & Haffert 2025).

2.5 High-Resolution Spectroscopy

High-resolution spectroscopy (HRS) offers a complementary approach to direct imaging by exploiting the information contained within a planets atmosphere. A planetary spectrum encodes atmospheric composition and temperature structure through absorption and emission features. The relevant instrument is a spectrograph: a dispersing element such as a diffraction grating that separates incoming light by wavelength, recording each wavelength at a different position on the detector (Howell 2006). The ability of a spectrograph to distinguish two nearby wavelength is quantified by its resolving power R, defined formally in §2.8.1. At high resolving powers $R \gtrsim 10^4$ molecular bands such as those of CO, H_2O, CH_4 break up into many individually resolved lines, each carrying information about the atmospheric composition and thermal structure (Snellen 2025), see Figure 2.5 for an illustration of this.

An advantage of HRS is that the planet spectrum is Doppler-shifted with respect to contaminating components. The orbital motion of the planet produce a radial velocity, while the stellar spectrum is comparatively stationary and telluric (atmospheric) absorption lines remain fixed in the observer frame. This behavior enables efficient suppression of host-star contamination and allows the planetary signal to be isolated even when its contribution is small. In practice, the planetary signal is commonly extracted using cross-correlation techniques (see §2.7) where the planetary spectra is cross correlated with a model template, coherently combining the contribution of many lines to boost detection significance (Snellen 2025).

2.6 Combining HCI and HRS

Both HCI and HRS tackle the contrast problem in two separate ways. HCI suppresses the star spatially by attenuating its flux at the planet position, while HRS suppresses it spectrally by filtering the residual stellar light using the planet's Doppler-shifted spectrum. Combining them to a joint technique (HRS+HCI) in which the effective contrast is the product of the two suppressions (Snellen et al. 2015). Schematically, an HCI contrast of 10^{-5} combined with an HRS contrast of 10^{-5} can in principle approach a joint contrast of 10^{-10}; in practice the achievable performance is set by noise sources (see §2.8). This multiplicative gain is illustrated in Figure 2.4

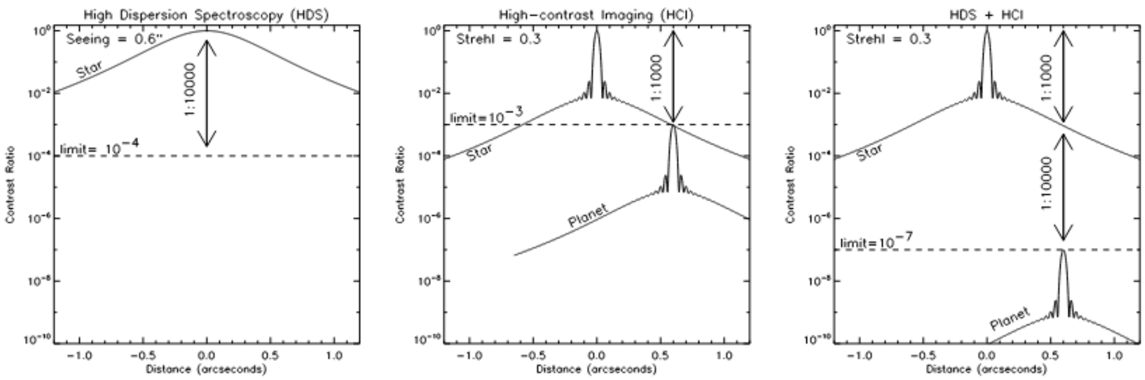

Figure 2.4: Left: Contrast limit of HRS only (labeled "HDS", high dispersion spectroscopy, in the original figure which is synonymous with HRS as used in this thesis). Middle: Contrast limit of HCI only. Right: theoretical gain by combining HRS+HCI. Taken from Snellen et al. (2015), Figure 1.

A key motivation for HRS+HCI is that it mitigates limitations inherent to each technique individually. By first attenuating the star locally at the planet position, the stellar photon noise is reduced and the subsequent high-resolution analysis becomes substantially more efficient (Snellen et al. 2015). The multiplicative contrast gain is a central motivation for next-generation ELT instruments that combine extreme adaptive optics, coronagraphy, and high-resolution spectroscopy.

2.7 Cross-Correlation Spectroscopy

The spectral suppression term that combines multiplicatively with HCI in §2.6 is in practice realized through cross-correlation spectroscopy. At sufficiently high resolving power, molecular bands separate into many narrow lines whose collective pattern uniquely identifies each molecule (Birkby 2018; Snellen 2025). Any individual line is buried in noise, but not when treated as a collective. The positions and relative strength of the lines are predicted by a forward model of the planet's atmosphere. The cross-correlation technique uses this model as a template, shifting it over a range of radial velocities and computing its similarity with the observed spectrum at each step. When the template matches the

Doppler-shifted planetary signal the line contributions add constructively, resulting in a peak in the cross-correlation function (CCF) at the corresponding velocity (Birkby 2018).

For two discrete sequences, the CCF is defined as

$$z[k] = \sum_l x_l y^*_{l-k}, \tag{2.3}$$

where x_l is the observed spectral flux, y_l is the model template, k is an integer lag, and y^* denotes the complex conjugate. For the real-valued spectra considered here the conjugation has no effect, and the operation reduces to a sum of pixel-wise products between the spectrum and a copy of the template shifted by k pixels. On a wavelength grid sampled at constant resolving power, an integer lag corresponds to a fixed velocity step $\Delta V = c/R$, so $z[k]$ can equivalently be written as a function of radial velocity v. This is the definition adopted in §3.5, where the CCF is computed with `scipy.signal.correlate` (Virtanen et al. 2020).

In practice the CCF serves two purposes. First, by co-adding the contributions of many lines it raises the planetary signal above the noise floor, enabling detection that would otherwise be impossible at the level of individual lines (see §2.8.2) (Birkby 2018). Secondly, the velocity at which the CCF peaks provides a direct measurement of the planet's radial velocity which can constrain orbital parameters and atmospheric dynamics (Snellen et al. 2010; Birkby 2018).

2.8 Instrumental Parameters and Noise

The detectability of a planetary signal in a high-resolution cross-correlation analysis depends on two properties of the instrument: the spectral resolving power, which sets how much of the line forest in a planetary spectrum can be resolved and used, and the detector noise budget, which sets the floor against which the recovered signal must be measured.

2.8.1 Spectral Resolution

The resolving power of a spectrograph is defined as

$$R = \frac{\lambda}{\Delta\lambda}, \tag{2.4}$$

where $\Delta\lambda$ is the smallest wavelength interval the instrument can separate. From the point of view of detectability, R enters the cross-correlation analysis in two ways. First, it sets the velocity sampling: a close-in planet whose Doppler shift moves its spectrum across more than ~ 10 pixels during a time series is well resolved, while at low R the lines either blend together or never separate from the stationary stellar and telluric features (Birkby 2018). Second, R controls both the depth and the number of individually resolvable spectral lines. Figure 2.5 shows the same hot-Jupiter model spectrum convolved to a range of resolving powers; many of the individual CO lines visible at $R = 10^5$ have disappeared by $R = 300$, and those that remain are substantially shallower. Because the cross-correlation co-adds the contribution of $N_{\rm lines}$ resolved lines, the planet signal-to-noise ratio scales as $\sqrt{N_{\rm lines}}$ (Birkby 2018).

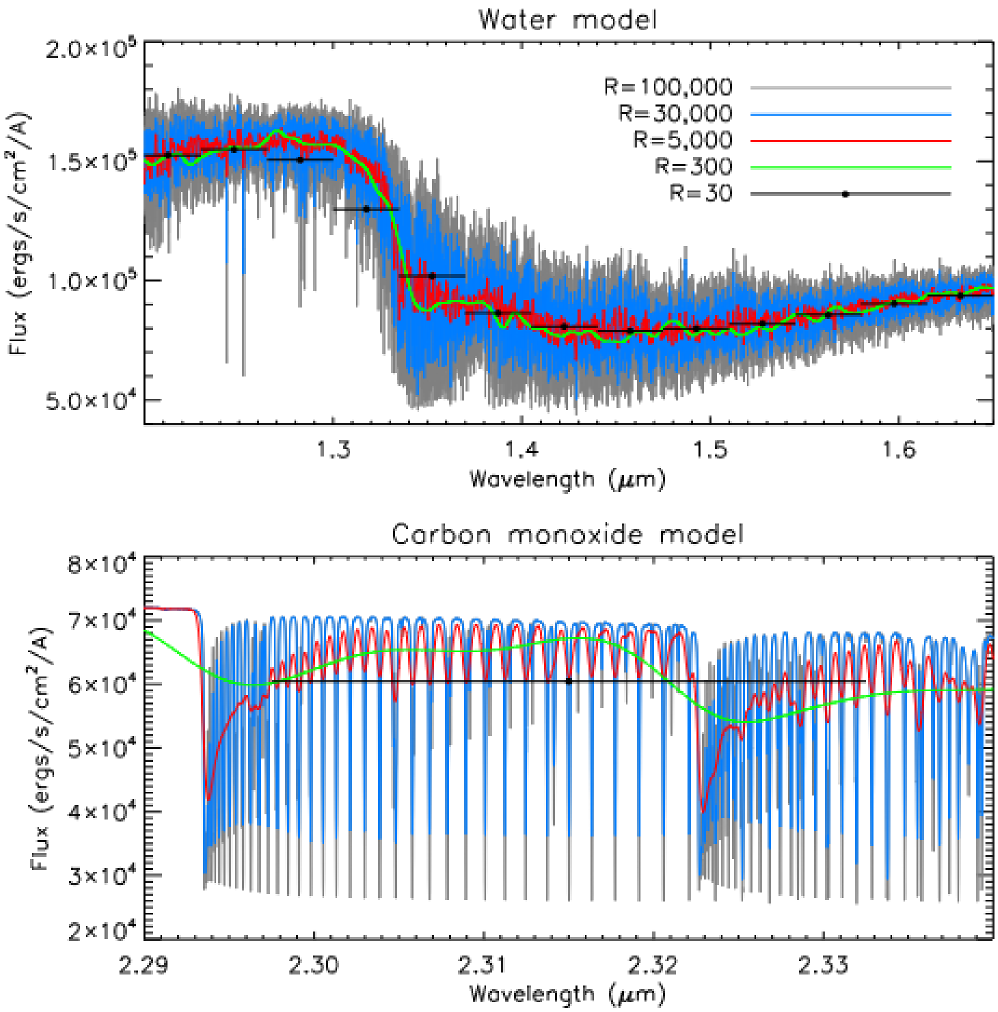

Figure 2.5: Hot Jupiter model spectrum of water and carbon monoxide convolved at a range of resolving powers. Taken from Birkby (2018), Figure 2.

2.8.2 Noise Sources

The noise budget on a charge-coupled device (CCD) measurement is set by

$$\frac{S}{N} = \frac{N_\star}{\sqrt{N_\star + n_{\rm pix}\left(N_S + N_D + N_R^2\right)}}, \tag{2.5}$$

where $N_\star$ is the total signal collected from the source (in electrons), $n_{\rm pix}$ is the number of pixels involved in the measurement, N_S is the sky background per pixel, N_D is the dark current per pixel, and N_R is the read noise per pixel. The first noise term, $\sqrt{N_\star}$, is the Poisson noise of the source itself (Howell 2006). The other three terms become important in the faint-source regime, which is relevant for direct detection of an exoplanet.

Photon noise. Photon noise describe the randomness inherent in photons arriving at the detector and follow Poisson statistics. This applies both to the planet signal itself and to the sky and stellar background that lands on the same detector pixels (Howell 2006).

Dark current. Thermal agitation in the detector, even in the absence of incident light, liberates electrons which are indistinguishable from photoelectrons on readout; their generation is a steep function of temperature and follow a Poisson distribution (Howell 2006).

Read noise. Read noise is the uncertainty in the conversion of charge in a pixel to a digital number. This is independent of exposure time and added to every pixel on every readout. This sets a noise floor that dominates in faint-source short exposures (Howell 2006).

Signal-to-noise of the planetary signal. For an exoplanet detection via cross-correlation, the relevant signal is not a single pixel but the coherent sum over $N_{\rm lines}$ resolved spectral lines. To first order, assuming the contaminating telluric features can be removed, the planetary signal-to-noise is

$$\mathrm{S/N}_p = \frac{S_p\sqrt{N_{\rm lines}}}{\sqrt{S_\star + \sigma^2_{\rm bg} + \sigma^2_{\rm read} + \sigma^2_{\rm dark}}}, \tag{2.6}$$

where S_p and $S_\star$ are the planet and stellar signal per resolution element. Here $\sigma^2_{\rm bg}$, $\sigma^2_{\rm read}$ and, $\sigma^2_{\rm dark}$ denote the variances of the same sky, dark-current, and read-noise contributions as N_S,N_D and, N_R in Equation 2.5 (Birkby 2018).

Chapter 3

Methodology

This chapter describes the simulation pipeline developed to assess the detectability of self-luminous giant exoplanets for HRS+HCI, covering end-to-end from physical stellar and planetary models to a detection significance. The pipeline takes the physical models introduced in Chapter 2, a stellar spectrum, a planetary spectrum, and a coronagraphic contrast curve, and propagates them through a simplified ANDES-like observation that includes photometric calibration, detector conversion, and noise, and retrieves the planetary signal through cross-correlation against a noiseless template.

3.1 Pipeline Overview

The pipeline takes stellar and planetary atmosphere models together with an instrument configuration and returns the detection significance of the recovered planetary signal. It consists of four stages: (i) generation of stellar and planetary spectra (§3.2), (ii) photometric calibration and detector conversion (§3.3), (iii) construction of synthetic observations including coronagraphic suppression and detector noise (§3.4), and (iv) recovery of the planetary signal through cross-correlation (§3.5).

3.2 Stellar and Planetary Spectra

Stellar spectra are drawn from the PHOENIX library (Husser et al. 2013) via the `PhoenixStarTable` interface in `petitRADTRANS`, which returns $F_\nu(\lambda)$ on the native PHOENIX grid; this is converted to $F_\lambda = F_\nu\, c/\lambda^2$ $[\mathrm{erg\,cm^{-2}\,s^{-1}\,\AA^{-1}}]$. Throughout this work the host star is an early F-type dwarf with $T_{\mathrm{eff},\star} = 7200\,\mathrm{K}$ and $H = 5.0$; surface gravity and metallicity follow the PHOENIX defaults at this temperature.

Planet emission spectra are generated with `petitRADTRANS` (Mollière et al. 2019) in line-by-line mode, preserving the individual molecular line profiles required for cross-correlation. $P{-}T$ profiles and equilibrium volume mixing ratios are read directly from the ATMO 2020 grid (Phillips et al. 2020); mixing ratios are converted to mass fractions using the local mean molar mass. Every species available in both ATMO 2020 and the `petitRADTRANS` database is included as a line opacity contributor, with H_2 and He as Rayleigh scatterers and H_2-H_2/H_2-He collision-induced absorption included as a continuum. The emergent flux $F^p_\lambda(\lambda)$ is computed at the planet surface over 0.95-1.80 μm, matching the reach of ANDES (Palle et al. 2025).

The four model companions studied here span the near-infrared sequence, relevant to direct imaging with ANDES; their parameters are summarized in Table 3.1 and their four most dominant molecular contributors at each temperature are documented in Appendix A (Figures A.1-A.4). They share the same host star, separation range, rotational velocity, and barycentric kinematics; only $T_{\rm eff,p}$ (and hence the realized H-band contrast) differs. Surface gravity is fixed at $\log g = 4.0$ for all four companions, a representative value for young, self-luminous giant planets adopted in comparable HRS+HCI simulation studies, (e.g. Houllé et al. 2021).

	Host star	**Companion A**	**Comp. B**	**Comp. C**	**Comp. D**
($T_{\rm eff}$) [K]	7200	400	800	1200	1600
Surface gravity ($\log g$)	4.0	4.0	4.0	4.0	4.0
Angular separation [mas]	-	20.0-78.9			
H band magnitude	5.0	variable ($5.0 + \Delta m_H$)			
Radial velocity [$\mathrm{km\,s^{-1}}$]	12.6	$12.6 + 8.0$			
Rotational velocity [$\mathrm{km\,s^{-1}}$]	-	20			

Table 3.1: Default parameters for the host star and its four simulated companions. Values listed once across the companion columns are common to all four.

3.3 Photometric calibration and detector conversion

Both spectra are placed on a common observational basis through a single calibration chain: a log-uniform wavelength grid at constant resolving power R spanning 0.95–1.80 μm ($\Delta\lambda_i = \lambda_i/R$, default $R = 10^5$); convolution with a Gaussian instrumental line-spread function of FWHM $\lambda_{\rm med}/R$, applied on the native oversampled grids before resampling; rotational broadening of the planet spectrum with $v \sin i = 20\,\mathrm{km\,s^{-1}}$ through an analytic kernel; Doppler shifts of $v_{\rm tot} = v_{\rm orb} + v_{\rm sys}$ for the planet and $v_{\rm sys}$ for the star, performed by rescaling the wavelength axis by $1+v/c$ followed by log-space interpolation; and photometric anchoring of the star such that its 2MASS H-band integral equals the Vega-calibrated photon flux corresponding to $H = 5.0$.

For the planet, two flux modes are supported. In *magnitude mode* (default) the planet flux is rescaled so that its H-band integral relative to the star equals $10^{-\Delta m_H/2.5}$; Δm_H then parameterizes the contrast. In *native flux mode* the `petitRADTRANS` surface flux is geometrically diluted to the observer through $(R_p/d)^2$ with R_p and d the planet radius and distance to the planet, respectively, expressed in a common length unit; this mode preserves the planet's absolute brightness and is used in §4.2 to convert detection thresholds into maximum-detectable distances.

Photon fluxes are converted to detector electrons per spectral channel via

$$N_e^{\star/p}(\lambda) = F_{\rm phot}^{\star/p}(\lambda)\, A\, \eta\, \mathrm{QE}(\lambda)\, \mathrm{DIT}, \tag{3.1}$$

with $A = 980\,\mathrm{m}^2$ (ELT collecting area), $\eta = 0.07$ (representative end-to-end ANDES transmission), and QE(λ) the wavelength-dependent detector quantum efficiency. A noiseless, zero-velocity copy of the planet spectrum is propagated through the same chain in parallel and stored as the cross-correlation template (§3.5).

3.4 Synthetic Observation

The coronagraphic contrast curve $C(\lambda;\rho)$ used here is a simulated ANDES reference profile (private communication, ANDES Consortium), interpolated onto the simulation wavelength grid at the requested angular separation ρ. Figure C.1 (top panel) shows the resulting contrast curve at 24 mas.

The synthetic detector frames mirror the on-axis/off-axis subtraction scheme used in real HRS+HCI observations,

$$s_1(\lambda) = N_e^\star(\lambda)\, C(\lambda;\rho) + N_e^p(\lambda) \quad \text{(on-axis: planet + stellar leakage)}, \tag{3.2}$$
$$s_2(\lambda) = N_e^\star(\lambda)\, C(\lambda;\rho) \quad \text{(off-axis: stellar reference)}, \tag{3.3}$$
$$s_3(\lambda) = \tilde{s}_1(\lambda) - \tilde{s}_2(\lambda) \quad \text{(residual)}, \tag{3.4}$$

where $\tilde{s}_i$ denotes the noisy realization of s_i. Each $\tilde{s}_i$ is corrupted by independent Poisson photon noise $\tilde{s}_i \sim \mathcal{P}(s_i)$, Poisson dark current $\mathcal{P}(d \cdot \mathrm{DIT})$ with $d = 0.0056\, e^- \,\mathrm{pix}^{-1}\,\mathrm{s}^{-1}$, and Gaussian read noise $\mathcal{N}(0, \sigma_R^2)$ with $\sigma_R = 4.5\, e^-$ (private communication, ANDES Consortium). The two realizations use independent random seeds. See Figure C.1 for the resulting spectra following this process.

3.5 Signal Retrieval via Cross-Correlation

The residual frame s_3 is cross-correlated against a noiseless planet template $T(\lambda)$, obtained by propagating the planet model through the same calibration chain with $v_{\rm orb} = v_{\rm sys} = 0$, no rotational velocity, and no detector noise. Equation 2.3 is computed over $|v| \leq 500\,\mathrm{km\,s^{-1}}$; on the log-uniform grid one integer pixel lag corresponds to $\Delta v = c/R \simeq 3\,\mathrm{km\,s^{-1}}$ at $R = 10^5$.

To isolate the planet contribution from correlations between the template and any residual stellar or detector structure in s_3, a differential CCF is formed:

$$\mathrm{CCF}_{\rm planet}(v) = \mathrm{CCF}(s_3, T)\,(v) - \mathrm{CCF}(s_3 - N_e^p, T)\,(v), \tag{3.5}$$

where the second term is the CCF of the same noisy frame with the noiseless planet electrons subtracted; it captures everything that correlates with T in the absence of planet light. The detection significance is the ratio of the peak of $\mathrm{CCF}_{\rm planet}$ at the injected velocity $v_{\rm inj} = v_{\rm orb} + v_{\rm sys}$ to the standard deviation of the noise-only CCF baseline,

$$(S/N)_{\rm planet} = \frac{\mathrm{CCF}_{\rm planet}(v_{\rm inj})}{\sigma[\mathrm{CCF}(s_3 - N_e^p, T)]}. \tag{3.6}$$

A signal is reported as detected when $(S/N)_{\rm planet} \geq 5$ (Houllé et al. 2021; Snellen 2025).

Chapter 4

Results and Discussion

4.1 Estimated Theoretical Gain

Figure 4.1 shows the 5σ detection threshold returned by the pipeline for the four model companions A-D around the reference host star over the full 0.95-1.80 μm range at $R = 10^5$. The dashed orange curve gives the raw coronagraphic contrast at each separation, expressed in magnitude; it represents the floor that would be reached by an imaging-only observation, before any spectral information is used.

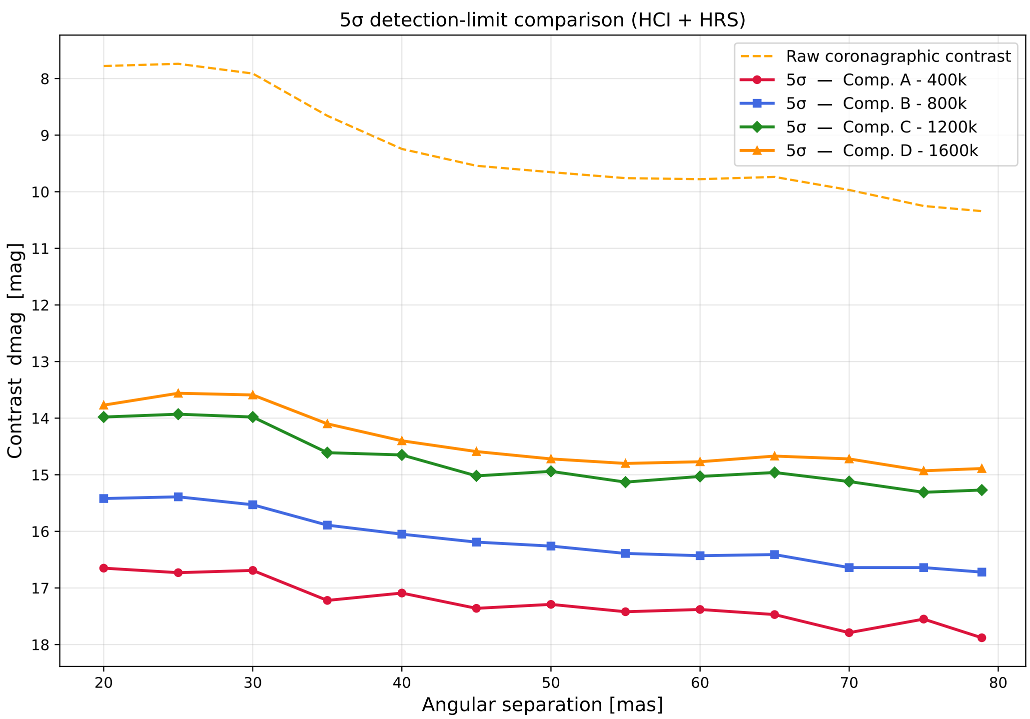

Figure 4.1: 5σ detection limits achieved by the simulated HRS+HCI pipeline for the four model companions A-D around an $H = 5.0$, $T_{\mathrm{eff},\star} = 7200\,K$ host star. The solid curve gives the faintest planet-to-star contrast at which the differential CCF reaches $S/N \geq 5$, as a function of angular separation. The dashed orange curve shows the raw coronagraphic contrast for reference.

At a separation of 40 mas, the 5σ thresholds reach $\Delta m_H \approx 17.2$ for companion A, 15.9 for B, 14.6 for C, and 14.1 for D, corresponding to contrasts of 1.3×10^{-7}, 4.3×10^{-7}, 1.4×10^{-6}, and 2.3×10^{-6} respectively. For reference, the raw coronagraphic contrast at the same separation reaches $\sim 2.1 \times 10^{-4}$.

The cross-correlation analysis therefore deepens the accessible contrast by ~ 8 magnitudes (≈ 3 orders of magnitude in flux ratio) relative to the imaging-only floor for the coolest companion, and by ~ 5 magnitudes for the hottest one. This is in line with the order-of-magnitude argument of Snellen et al. (2015), that HRS+HCI can in principle approach a contrast given by the product of the two suppression mechanism. Additionally, it also agrees with the PCS design specification of an HRS contrast gain of at least $1 : 10^4$ (Kasper et al. 2021).

The ordering of the four curves at fixed dmag is set by the spectral content of the planet rather than its brightness, since all four companions share the same H-band photon budget in its magnitude mode. What differs is the absolute range of each planet's spectrum across the $0.95 - 1.80\,\mu$m window. Figure B.1 (upper panel) shows that the raw surface emission of the 400 K companion spans many orders of magnitude across the wavelength range, with the flux concentrated into bright opacity windows around 1.07, 1.27, and $1.55\,\mu$m (Figure A.1), while the 1600 K spectrum varies by less than an order of magnitude over the same range and has a smoother continuum (Figure A.4). The H-band photometric anchoring then amplifies this disparity. After all four spectra are forced to share the same H-band integral (Figure B.1 (middle panel)), the relative structure of the raw spectra is preserved in absolute electrons per channel: the 400 K per-channel count fans out over roughly seven orders of magnitude, while the 1600 K count varies by roughly one. Anchoring at H therefore concentrates the photon budget of a cool companion into the narrow opacity windows where its emission actually emerges, deepening the per-line contrast that the CCF coherently sums over (§3.5).

The small improvement of every curve with separation reflects the corresponding improvement of the coronagraphic contrast profile in Equation 3.2, as the separation increases, the stellar leakage decreases.

4.2 Detection Horizons at Native Flux

The H-band anchored picture in §4.1 provides an estimation of the contrast gain from HRS+HCI but not on actual detectability of physical systems, due to the planet brightness being decoupled from its size and distance §3.3. In native flux mode the planet is instead set by its geometrically diluted surface emission. The host star remains anchored at $H = 5.0$, and at each separation the system distance d is swept to find the maximum distance at which the CCF still reaches 5σ.

Figure 4.2 reorders the four companions relative to Figure 4.1: the hottest planet is now detectable to the largest distance, and the coolest is restricted to the immediate solar neighborhood. Companion D (1600 K) remains 5σ-detectable out to $\sim$200-340 pc across the probed separation range, Companion C (1200 K) to $\sim$ 100-200 pc, Companion B (800 K) to $\sim$70-130 pc and Companion A (400 K) only to $\sim$10-15 pc.

Due to the intrinsic brightness scaling as T_{eff}^4 the 1600 K planet outshines the 400 K one by roughly 2 orders of magnitude in total flux. The scaling effects, which favored companion A

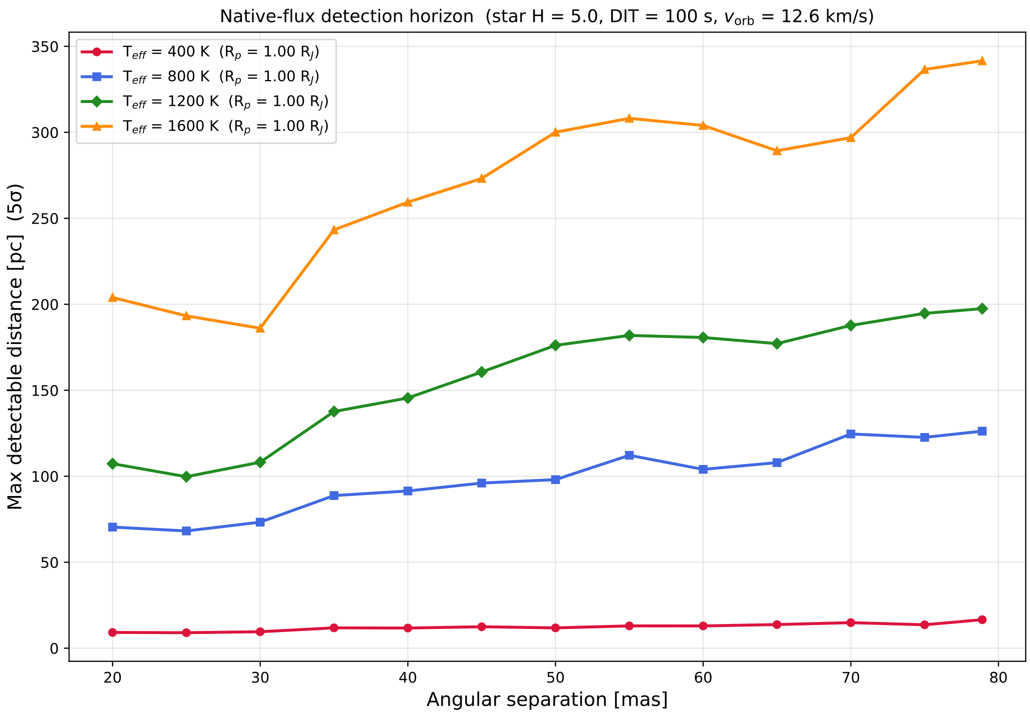

Figure 4.2: Native-flux detection horizon: maximum distance d_{max} at which each model companion is recovered at 5σ in the differential CCF, as a function of angular separation. The host star is held at $H = 5.0$, $T_{\mathrm{eff},\star} = 7200\,K$ and the planet flux is set by `petitRADTRANS` surface emission evaluated at the observer. For reference: current limitations for a similar stellar type have found planet at ~ 110 pc at separation ~ 800 mas (Chauvin et al. 2017; Ray et al. 2025). This would be located far right of the plot.

in §4.1, cannot compensate for the lost photons when the source brightness is set physically rather than parametrically. For reference, current contrast-limited direct-imaging have detected companions of similar effective temperature at ~ 100 pc and ~ 800 mas with similar host star properties (Chauvin et al. 2017; Ray et al. 2025), well to the right of the angular separation range plotted in Figure 4.2.

The irregular shape of the curves in Figure 4.2 follows the adopted coronagraphic contrast profile: the raw coronagraphic contrast shown in Figure 4.1 (dashed curve) is comparatively shallow and nearly flat between 20 and 30 mas before deepening steadily from ~30 mas onward, and this same behavior propagates directly into the native-flux detection horizon, producing the depressed values at the smallest separations in Figure 4.2. Furthermore, since the residual stellar leakage entering Equation 3.4 is set by this contrast curve, the recovered S/N, and therefore the maximum detectable distance, inherit the same structure. Because the coronagraphic contrast also varies with wavelength (Figure C.1), the precise separations at which local peaks and valleys occur are additionally sensitive to which wavelengths dominate the recovered S/N at each separation.

4.3 Wavelength-Range Dependence

§4.1 is averaged over the full 0.95-1.80 μm window. To investigate how much of the detection comes from each part of the range, the same pipeline run on the two halves of the window: 0.95-1.40 μm and 1.40-1.80 μm. The slicing is applied to the H-band anchoring, so the absolute electron budget per channel is unchanged; only the wavelength range entering the CCF is restricted.

4.3.1 H-band Anchored

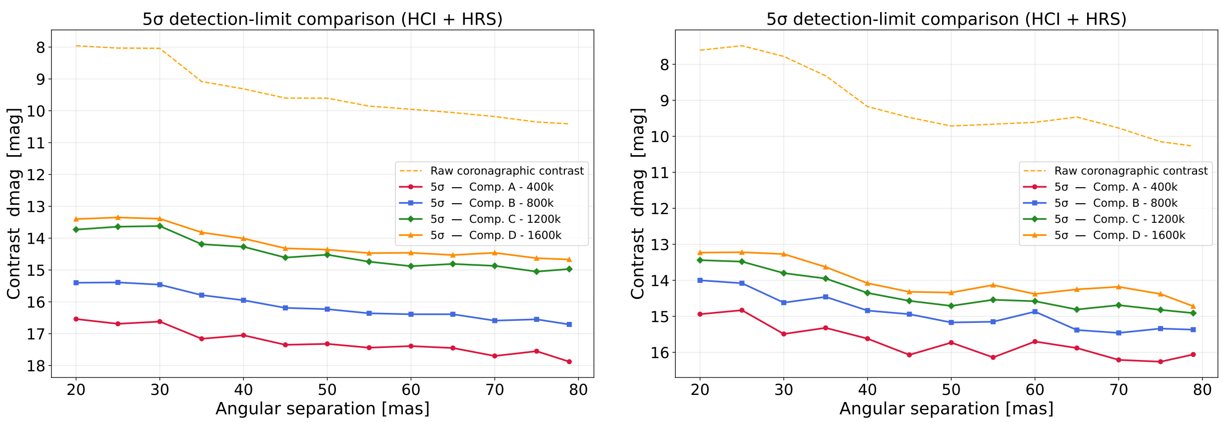

(a) 0.95-1.40 μm (b) 1.40-1.80 μm

Figure 4.3: 5σ detection limits for the four model companions, restricted to two halves of the full wavelength range. Same axes and reference contrast as Figure 4.1. The H-band anchoring is preserved across panels; only the wavelength range entering the CCF differs.

The two panels of Figure 4.3 show that the lower half, 0.95-1.40 μm drives most of the detection. For all four companions the 5σ curves in panel (a) are deeper than in panel (b), with a larger gap for companion A and B compared to C and D.

The effects discussed in §4.1 reappear here, now sub-band by sub-band. Companions A and B have most of their emission concentrated in narrow opacity windows around 1.07, 1.27, and 1.55 μm (Figures A.1, A.2), and the H-band anchoring places its reference integral inside the upper sub-band, where the spectrum sits in a single broad continuum window. Two of the three cool-companion emission peaks fall in the lower half, and they are also where the deepest H_2O, NH_3 and CH_4 absorption features sit; the lower sub-band therefore carries most of the depth that the CCF can coherently sum over, while the upper half contributes a smoother, shallower spectrum. For the hotter companions C and D (Figures A.3, A.4) the underlying spectrum is comparatively smooth across the whole $0.95 - 1.80$ μmrange, the anchoring has less range to amplify, and the two halves contribute more equally, hence the much smaller gap between panels.

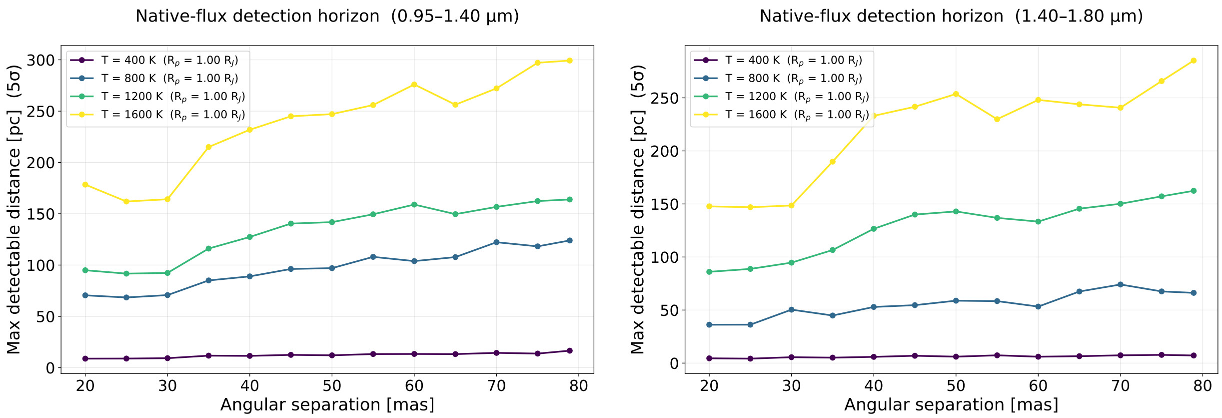

(a) 0.95-1.40 μm (b) 1.40-1.80 μm

Figure 4.4: Native-flux 5σ detection horizons for the four model companions, restricted to two halves of the full wavelength range. Same axes and assumptions as Figure 4.2

4.3.2 Native Flux

When used in native flux mode (Figure 4.4), the lower range yields a larger detection horizon for every companion, but the margin decreases noticeably with temperature. For companion A ($400\,K$) and B ($800\,K$) the lower range exceeds the higher by roughly a factor of two at all separations, while for companion C ($1200\,K$) and D ($1600\,K$) the two halves produce more comparable horizons.

Native flux mode removes the H-band anchoring, so the amplification that was concentrating the cool companions' photon budget into the lower sub-band no longer applies, each planet now emits whatever absolute flux its surface produces (Figure B.1). Only where the planetary spectrum deposits photons and where the line contrast can convert them into CCF significance matter. For the cool companions, most of the emerging flux still falls in the same narrow opacity windows at 1.07 and 1.27 μm, so the same features as in §4.3.1 also drive the detection from the lower half. For the hot companions, the underlying spectrum is broader and increasingly weighted toward the upper range, and the line depth is not receiving the photon budget benefit as in §4.3 and the two halves contribute comparably. Again, the irregular shape of the curves follows from the same reasoning as discussed in §4.2.

4.4 Resolving Power Dependence

The pipeline is re-run at $R \in [5\times10^4, 1.9\times10^5]$, holding the DIT, contrast curve, and wavelength range fixed at the §4.1 values. Figure 4.5 shows the 5σ detection threshold Δm_H as a function of R for each companion, at four representative angular separations (20, 40, 60, 79 mas).

The cooler companions A and B show no clear systematic trend; across the probed range of R the curves differ within $\sim 0.2 - 0.6$ mag of one another, with variance comparable to

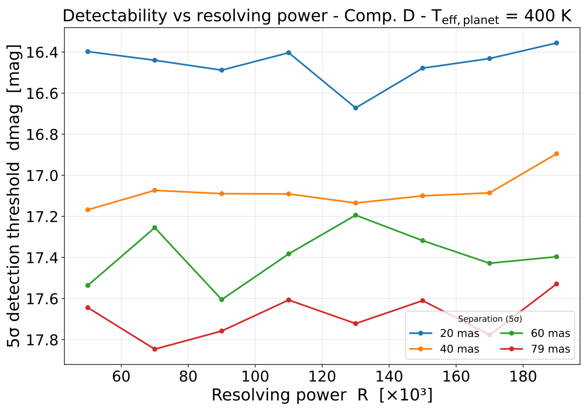

(a) Companion A ($T_{\mathrm{eff,p}} = 400$ K).

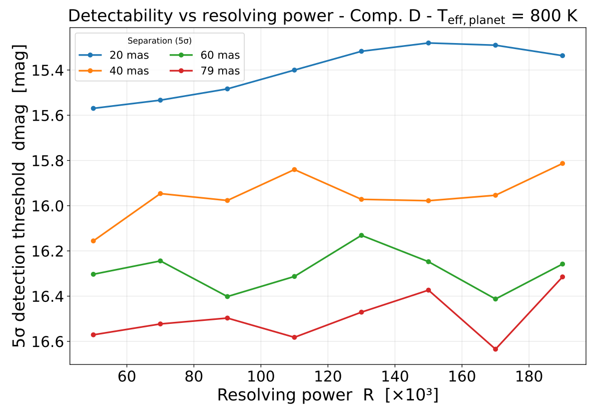

(b) Companion B ($T_{\mathrm{eff,p}} = 800$ K).

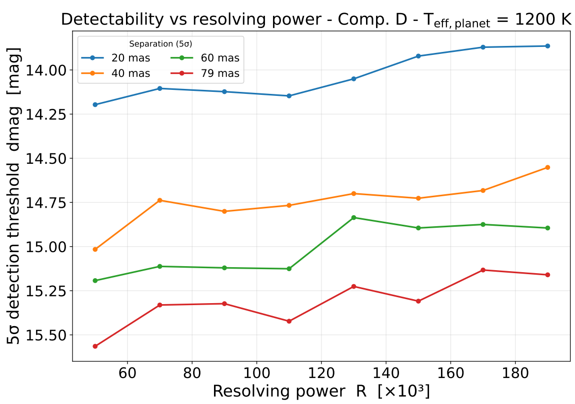

(c) Companion C ($T_{\mathrm{eff,p}} = 1200$ K).

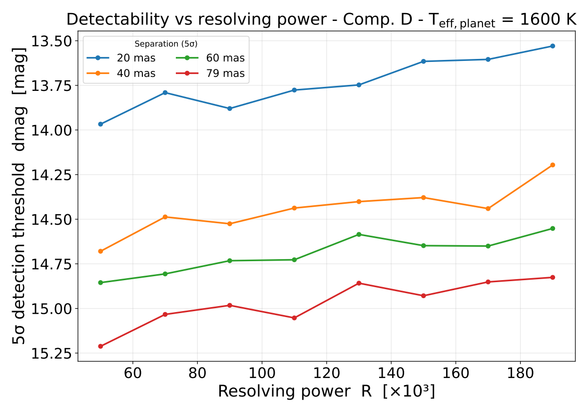

(d) Companion D ($T_{\mathrm{eff,p}} = 1600$ K).

Figure 4.5: 5σ detection threshold Δm_H as a function of resolving power, at four representative angular separations, for each of the four model companions. All other pipeline parameters are held at the reference values (0.95-1.80 μm, $DIT = 100\,s$, magnitude-anchored mode).

the spread between independent noise draws (see below). The hotter companions C and D show a moderate preference for the lower end of the R range, drifting by ~ 0.5 mag over the range.

Increasing R splits previously blended molecular bands into more individually resolved lines and preserves their depth, which raises N_{lines} and, all else being equal, should improve the CCF S/N (§2.7). However, at fixed DIT and throughput a higher R distributes the same photon budget across more spectral channels, so each channel collects fewer planet electrons while its read-noise and dark-current contribution stay fixed (§2.8.2). The CCF gain from splitting and deepening lines therefore competes against a per-channel signal that drops with R, and is dependent on how much spectral structure the planet contributes and on how close the per-channel signal sits to the noise floor.

One caveat is worth noting: the apparent point-to-point variation in Figure 4.5 does not reflect a stochastic variability of the pipeline; every point is fully deterministic for the same input parameters and random seed for the noise. Because each R value requires its own wavelength grid the same seeds map to an independent noise realization, and the variance between adjacent R values measures the dispersion of those independent draws rather than any intrinsic randomness of the method. Without averaging over multiple seeds it is therefore not possible to conclude whether the inconclusive shapes for companion A and B hide a trend, or whether the magnitude drift for C and D reflects a real trend rather than a lucky/unlucky draw.

Chapter 5

Conclusion and Outlook

This thesis investigated the detectability of self-luminous giant exoplanets through the combination of high-contrast imaging and high-resolution spectroscopy in the context of future ELT instruments, particularly the ANDES spectrograph. An end-to-end simulation pipeline was developed to generate physically motivated stellar and planetary spectra, propagate them through a simplified HRS+HCI observation that includes coronagraphic suppression and detector noise, and retrieve the planetary signal via cross-correlation against a noiseless template. The pipeline was applied to a grid of four self-luminous companions with $T_{\text{eff}} = 400,\ 800,\ 1200,$ and $1600\,K$ orbiting a common early F-type host star, and the 5σ detection threshold was mapped as a function of angular separation, wavelength range, and resolving power.

Summary of findings

The cross-correlation analysis was found to deepen the accessible contrast by approximately five to eight magnitudes (roughly three orders of magnitude in contrast) relative to the raw coronagraphic floor §4.1, with the largest gain obtained for the coolest companion. This is broadly consistent with the multiplicative-gain argument of Snellen et al. (2015).

When the four companions were instead compared at their native surface fluxes §4.2, the $1600\,K$ planet remained 5σ-detectable out to $\sim 200 - 340\,\text{pc}$ across the probed separation range, while the $400\,K$ planet was restricted to $\sim 10 - 15\,\text{pc}$.

The wavelength-range study §4.3 showed that the lower half of the $0.95 - 1.8\,\mu\text{m}$ window ($0.95 - 1.4\,\mu\text{m}$) drives most of the detection for all four companions when the H-band anchoring is preserved, with a particularly large gap for the cool companions whose emission is concentrated in narrow opacity windows around 1.07 and $1.27\,\mu\text{m}$. In native-flux mode, where the anchoring no longer redistributes the photon budget, this preference for the lower range persists for the cooler companions but weakens for the hotter ones, whose smoother emission spreads more uniformly across the full range.

The resolving power study §4.4 proved less conclusive. The two coolest companions showed no clear systematic trend over $R \in [5 \times 10^4,\ 1.9 \times 10^5]$ with the dispersion between points comparable to the variance expected from independent noise realizations. The two hotter companions showed a moderate ($\sim 0.5\,\text{mag}$ drift toward better detectability at the lower end of the R range. However, without averaging over multiple noise seeds, it is not possible to confirm whether this drift reflects a real trend or a particular sequence of noise draws.

Limitations

The results reported above should be read as upper-limit estimations rather than predictions of on-sky performance. The pipeline neglects telluric absorption and the residual variability of those features after correction, both of which are important in real ground observations. The companion grid samples only four effective temperatures at a single surface gravity, metallicity, and C/O ratio, and a single host star; the influence of e.g. clouds, disequilibrium chemistry, and rotational broadening beyond $v\,sin\,i = 20\,\mathrm{km\,s^{-1}}$ is not explored. Moreover, each pipeline run uses a single random seed for the noise, so quantitative statements about point-to-point variation, particularly for the resolving power study, would benefit from averaging over many independent realizations.

A further idealization concerns the cross-correlation template itself (§3.5). The template used is generated from the same planetary model which produce the injected signal, thus the recovered S/N corresponds to a perfectly matched filter and should be read as an upper limit. In practice, the atmospheric parameters of an exoplanet, for instance, its chemistry, $T-P$ profile, and C/O ratio, are not known a priori. The signal must therefore be searched for using candidate templates, none of which will generally provide a perfect match to the true spectrum. Consequently, template mismatch is expected to reduce the recovered cross-correlation signal and detection significance relative to the values reported here.

Outlook

Several extensions naturally follow from this work. The most immediate is to repeat each configuration with many noise seeds so that the results can be reported with statistical uncertainties rather than as single deterministic draws. Furthermore, bringing the simulated observations closer to real ground-based HRS+HCI by e.g. including tellurics, AO-residual speckles, and their associated variability, would likely lessen the contrast gains reported here. Also, probing a larger set of exoplanetary atmospheres, even beyond the equilibrium chemistry considered here, would test how sensitive the temperature driven trends of §4.1 and §4.3 are to changes in the underlying chemistry. To this extent, injecting specific molecules and biomarkers into the planetary template, would allow their individual detectability to be quantified and would extend the pipeline from a parametric contrast study toward a tool for assessing the spectroscopic characterization of habitable atmospheres. Finally, the same pipeline can in principle be adapted to simulate reflected light, instead of emissions, and thus shift the focus from young giants to cooler, terrestrial planets, which in turn could simulate the PCS regime. This is also the setting where biomarker injection becomes most relevant.

Bibliography

Ambaum, M. H. P. 2021, in Thermal Physics of the Atmosphere (Second Edition), second edition edn., ed. P. D. Williams, Developments in Weather and Climate Science (Elsevier), 205–222

Birkby, J. L. 2018, arXiv e-prints, arXiv:1806.04617

Chauvin, G. 2024, Comptes Rendus Physique, 24, 139

Chauvin, G., Desidera, S., Lagrange, A.-M., et al. 2017, A&A, 605, L9

Christiansen, J. L., McElroy, D. L., Harbut, M., et al. 2025, The Planetary Science Journal, 6, 186

Davies, R. & Kasper, M. 2012, ARA&A, 50, 305

Endl, M. 2014, in Encyclopedia of the Solar System (Third Edition), third edition edn., ed. T. Spohn, D. Breuer, & T. V. Johnson (Boston: Elsevier), 957–977

Follette, K. B. 2023, Publications of the Astronomical Society of the Pacific, 135, 093001

Fortney, J. J. 2018, arXiv e-prints, arXiv:1804.08149

Houllé, M., Vigan, A., Carlotti, A., et al. 2021, A&A, 652, A67

Howell, S. B. 2006, Handbook of CCD Astronomy, 2nd edn., Cambridge Observing Handbooks for Research Astronomers (Cambridge University Press)

Husser, T.-O., Wende-von Berg, S., Dreizler, S., et al. 2013, A&A, 553, A6

Kasper, M., Cerpa Urra, N., Pathak, P., et al. 2021, The Messenger, 182, 38

Kenworthy, M. A. & Haffert, S. Y. 2025, ARA&A, 63, 179

Madhusudhan, N. 2019, ARA&A, 57, 617

Marois, C., Macintosh, B., Barman, T., et al. 2008, Science, 322, 1348

Mayor, M. & Queloz, D. 1995, Nature, 378, 355

Mollière, P., Wardenier, J. P., van Boekel, R., et al. 2019, A&A, 627, A67

Palle, E., Biazzo, K., Bolmont, E., et al. 2025, Experimental Astronomy, 59, 29

Phillips, M. W., Tremblin, P., Baraffe, I., et al. 2020, A&A, 637, A38

Ray, S., Sallum, S., Hinkley, S., et al. 2025, ApJ, 983, L25

Seager, S. & Deming, D. 2010, ARA&A, 48, 631

Snellen, I., de Kok, R., Birkby, J. L., et al. 2015, A&A, 576, A59

Snellen, I. A. G. 2025, ARA&A, 63, 83

Snellen, I. A. G., de Kok, R. J., de Mooij, E. J. W., & Albrecht, S. 2010, Nature, 465, 1049

Terragni, J., Gordon, I., Adkins, E., et al. 2025, Journal of Quantitative Spectroscopy and Radiative Transfer, 347, 109631

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261

Wright, J. T. & Gaudi, B. S. 2013, in Planets, Stars and Stellar Systems. Volume 3: Solar and Stellar Planetary Systems, ed. T. D. Oswalt, L. M. French, & P. Kalas, 489

Appendix A

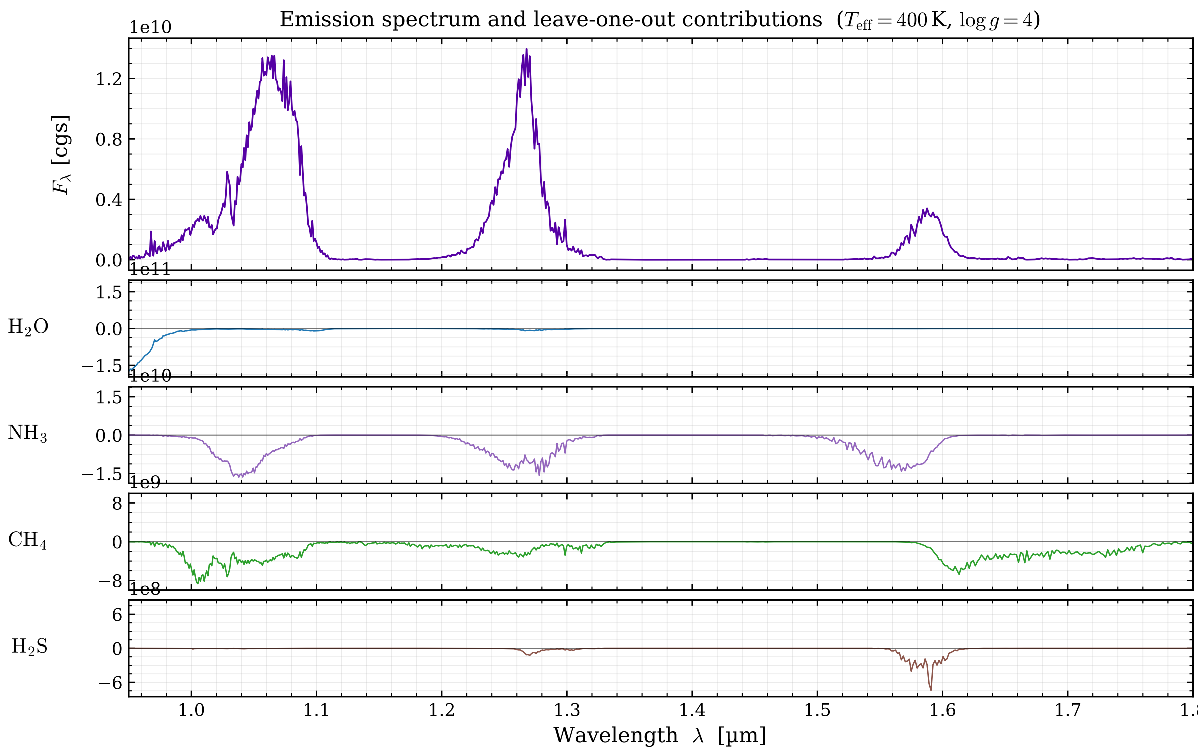

Figure A.1: Emission spectrum across $0.95 - 1.80\,\mu$m (top panel) decomposed into per-species leave-one-out contributions ($F_{\rm all} - F_{\rm not,x}$) for the four most dominant line contributors in companion A.

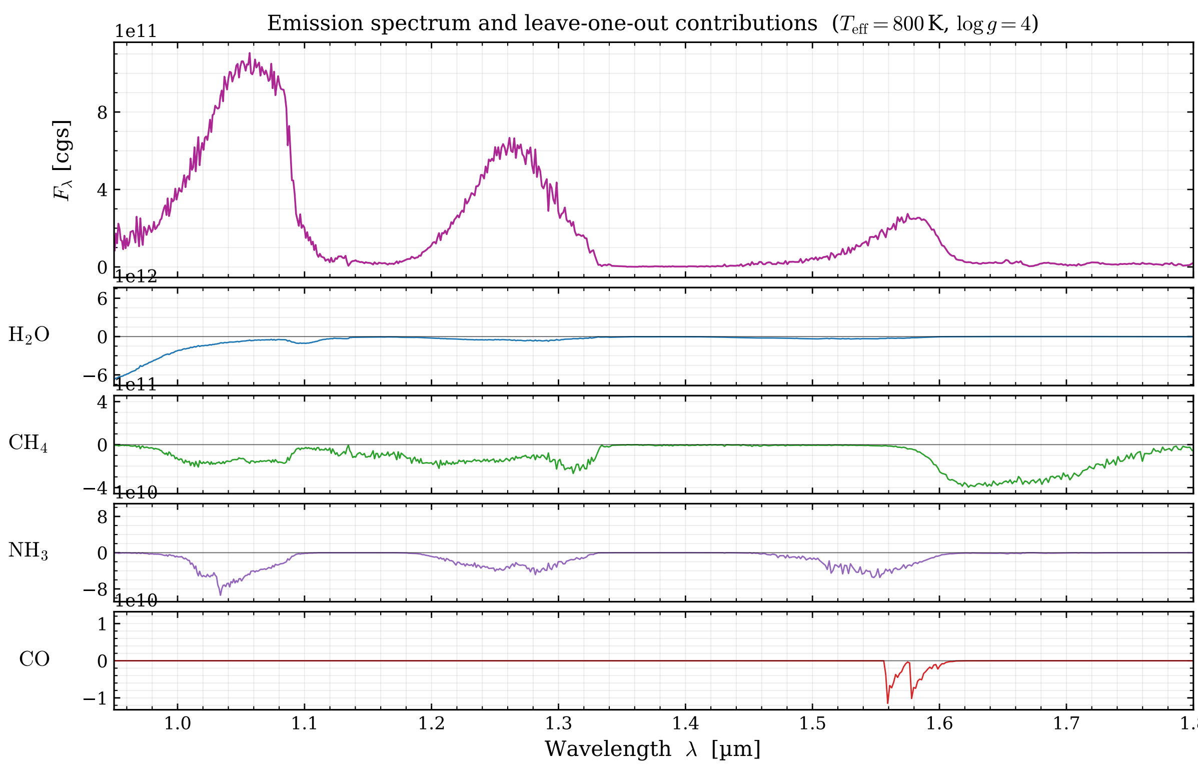

Figure A.2: Emission spectrum across $0.95 - 1.80\,\mu$m (top panel) decomposed into per-species leave-one-out contributions ($F_{\rm all} - F_{\rm not,x}$) for the four most dominant line contributors in companion B.

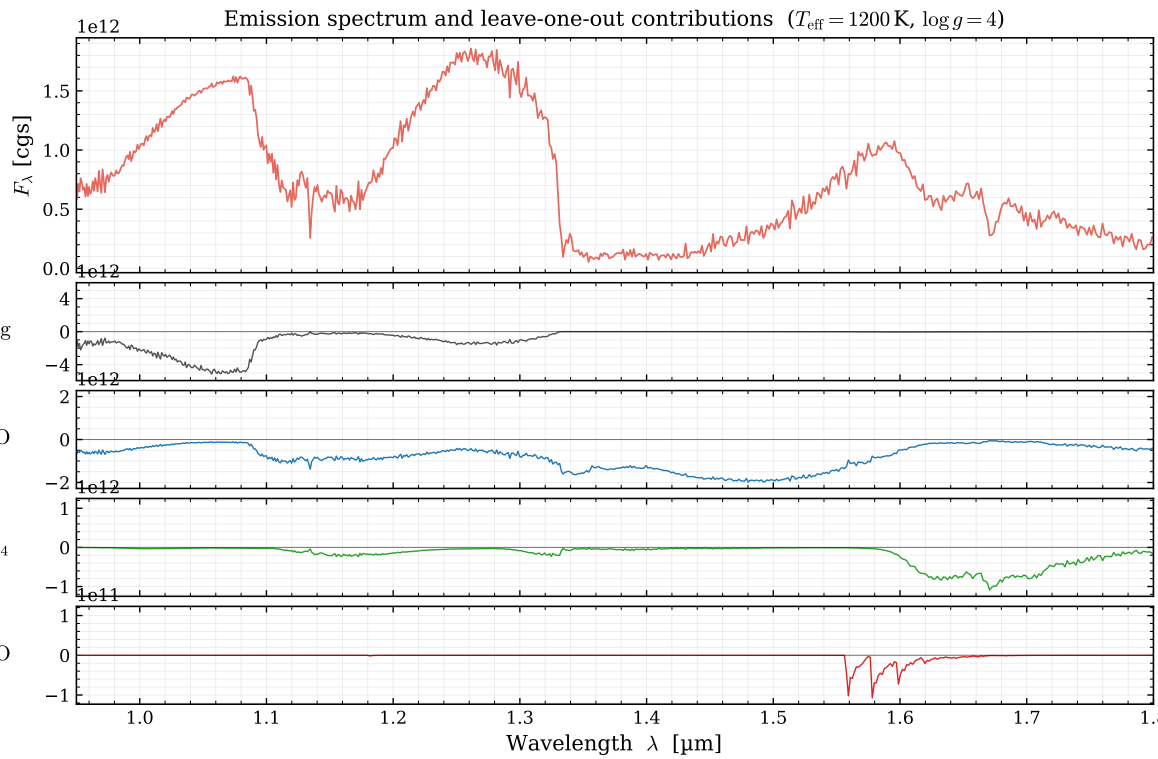

Figure A.3: Emission spectrum across $0.95 - 1.80\,\mu$m (top panel) decomposed into per-species leave-one-out contributions ($F_{\rm all} - F_{\rm not,x}$) for the four most dominant line contributors in companion C.

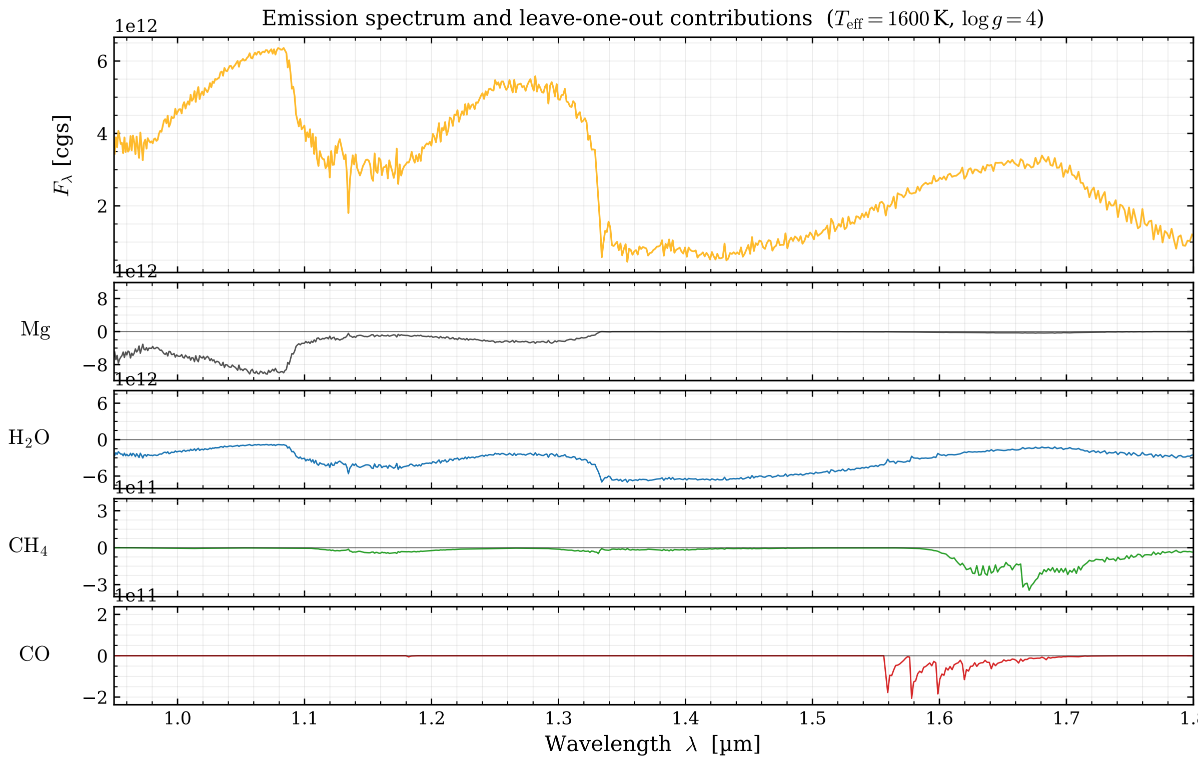

Figure A.4: Emission spectrum across $0.95 - 1.80\,\mu$m (top panel) decomposed into per-species leave-one-out contributions ($F_{\rm all} - F_{\rm not,x}$) for the four most dominant line contributors in companion D.

Appendix B

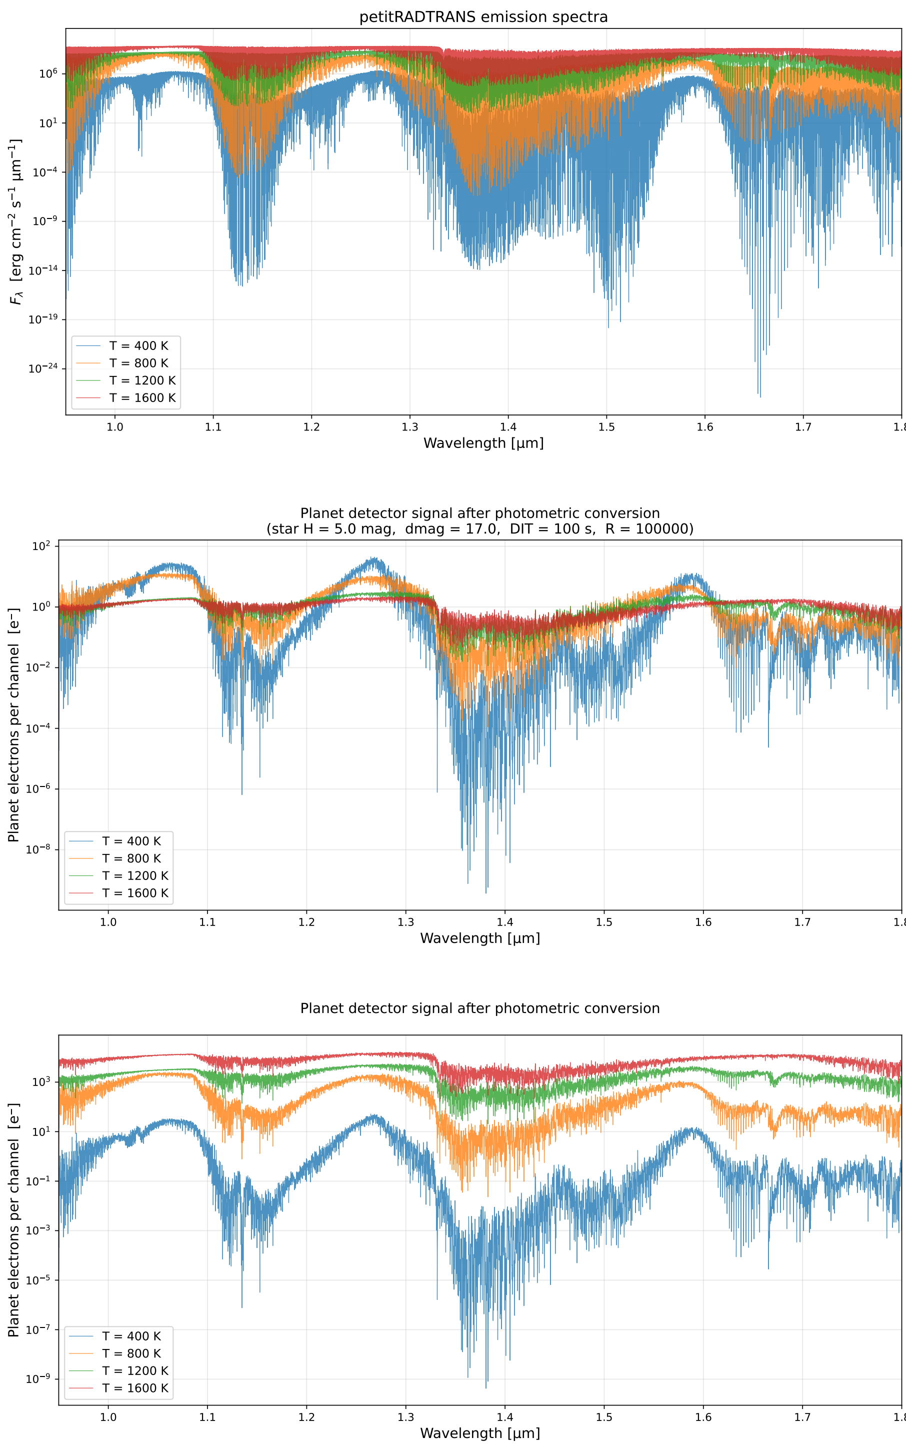

Figure B.1: Planetary spectrum at three stages of the §3.3 pipeline for the four model companions. **(upper)** Raw `petitRADTRANS` surface emission F_λ. **(middle)** Detector electrons per channel after H-band anchoring at $\Delta m_H = 17.0$ (magnitude mode, DIT $= 100\,$s, $R = 10^5$); all four spectra share a common H-band integral. **(lower)** Detector electrons per channel in native-flux mode, where the surface flux is geometrically diluted by $(R_p/d)^2$ without photometric anchoring, preserving the T_{eff}^4 brightness ordering.

Appendix C

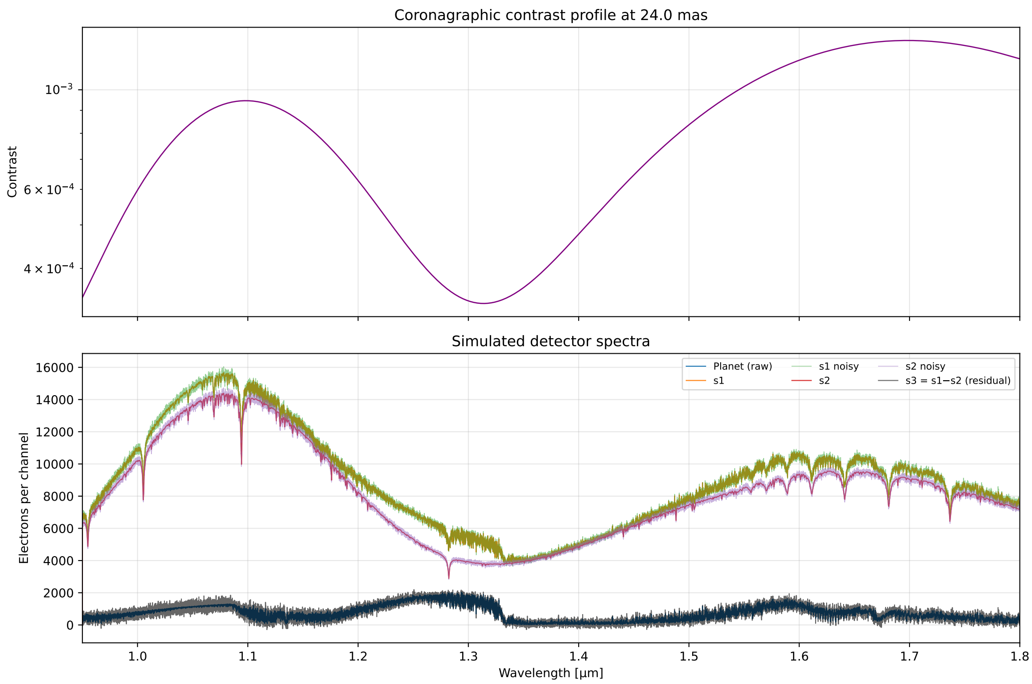

Figure C.1: Resulting spectrum following §3.4 for companion C at $\Delta m_H = 10.0$ at a separation of 24 mas including the coronagraphic contrast cure at the same separation.