
Revisiting the SHINE high-contrast imaging exoplanet survey using the Regime Switching Model

Mariam Sabalbal



*A thesis submitted in partial
fulfillment for the degree of
Doctor of Science*

University of Liège

2026

STAR Institute
Planetary & Stellar system Imaging Laboratory (PSILab)

Supervisor

Prof Olivier Absil

Jury Members

Dr Markus Bonse

Dr Faustine Cantalloube

Dr Valentin Christiaens

Dr Carl-Henrik Dahlqvist

Prof Michaël De Becker

Dr Gilles Orban De Xivry

Declaration

This work has made use of the High Contrast Data Centre, jointly operated by OSUG/IPAG (Grenoble), PYTHEAS/LAM/CeSAM (Marseille), OCA/Lagrange (Nice), Observatoire de Paris/LESIA (Paris), and Observatoire de Lyon/CRAL, and supported by a grant from Labex OSUG@2020 (Investissements d'avenir – ANR10 LABX56).

This project received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 819155).

Computational resources were provided by the Consortium des Équipements de Calcul Intensif (CÉCI), funded by the Fonds de la Recherche Scientifique de Belgique (F.R.S.-FNRS) under Grant No. 2.5020.11 and the Walloon Region.

ChatGPT (OpenAI) was used during the preparation of this thesis for language editing, proofreading, improving the clarity and readability of the manuscript, and enhancing the presentation of selected figures. All suggestions generated by ChatGPT were critically reviewed, modified where appropriate, and validated by the author.

Copyright ©2026 by Mariam Sabalbal, Université de Liège, Place du 20 Août, 7, B-4000 Liège, Belgium.

All right reserved. No part of this publication may be reproduced in any form or by any means, without the prior written permission of the author and supervisor.

Abstract

High-contrast imaging has become a major technique for the detection and characterization of young giant exoplanets, probing a region of parameter space largely inaccessible to other detection methods. Advances in adaptive optics, coronagraphy, and post-processing techniques have enabled the discovery of a growing population of directly imaged companions. However, the sensitivity of current ground-based observations remains strongly limited by residual stellar noise, particularly at small angular separations where planets are expected to be more common. Consequently, the scientific return of many archival imaging surveys remains fundamentally limited by the performance of the post-processing and detection frameworks used to analyze them.

This thesis investigates the application of the Regime Switching Model (RSM), a recent high-performance detection framework, to the F150 sample of the SHINE survey with the aim of improving the detection sensitivity at small angular separations.

We applied the RSM algorithm to more than 200 observations from the F150 sample of the SHINE survey, deriving detection maps and detection limits. A clustering approach based on environmental parameters was used to investigate the influence of observing conditions on the post-processed noise distribution, and two radial thresholding strategies based on lognormal noise modeling and F1-score maximisation were developed and evaluated. The detected signals were then assessed using multi-epoch observations, color-magnitude diagrams, and additional diagnostics to evaluate their astrophysical nature and distinguish genuine sources from false positives. Finally, we applied the RSM framework to Keck observations of ϵ Eridani to search for ϵ Eridani b and derive constraints on its properties.

The application of RSM to the SHINE F150 sample yields contrast improvements of approximately a factor of two overall and up to a factor of five at small angular separations compared with standard PCA processing. A total of 89 signals are recovered, including four background-star candidates and no new planetary candidate. The analysis reveals distinct observational regimes characterized by different noise properties and detection thresholds, with threshold variations of up to an order of magnitude between observing conditions. Both the lognormal and F1-score thresholding approaches provide consistent detection thresholds, with the lognormal approach offering a simple, generalizable framework applicable across different observing conditions. The analysis also reveals

a substantial population of false positives, corresponding to a false-positive rate of approximately 18%. Finally, the application of RSM to archival Keck observations of the ϵ Eridani system provides additional constraints on the proposed companion compared with previous analyses based on Keck data.

Overall, this work demonstrates that RSM is an effective detection framework for improving the sensitivity of high-contrast imaging observations, particularly at small angular separations. It also shows that robust threshold calibration requires accounting for variations in the observing conditions, as the post-processed noise properties differ significantly between observational regimes. Finally, the substantial false-positive population identified in this work highlights that both the adopted thresholding strategy and the classical PSF-subtraction stage remain major challenges for the reliable application of RSM. These findings lay the foundation for further improving the framework and extending its application to the re-analysis of archival surveys and future high-contrast imaging observations.

Acknowledgements

Almost five years ago, I started this PhD believing that research was simply a continuation of university learning, only focused on a single topic. That belief lasted until I found myself alone in my office, with only a few papers to read and the subject of my doctorate ahead of me. It was indeed very different from anything I had expected. Research comes with no measure of progress, no syllabus to follow, and often no clear beginning or end. Since then, this PhD has been a journey of highs and lows. Difficult as it was at times, it has been one of the most enriching chapters of my life. It taught me that no one is born knowing how to do high-contrast imaging, that accepting what I do not know is the first step toward understanding, and that learning is built slowly, little by little, through patience, mistakes, and perseverance. Although this dissertation bears my name, it is far from being the result of my efforts alone. It is the product of the guidance, support, friendship, and kindness of many people, to whom I owe my deepest gratitude.

First and foremost, I would like to express my deepest gratitude to my supervisor, Olivier. I often realize how fortunate I have been to work under your supervision. You made my PhD remarkably smooth in ways that many doctoral students unfortunately never get to experience. Your door was always open, our conversations were always easy, and I always felt heard and understood. No matter how busy you were, you generously made time whenever I needed guidance. Scientifically, I learned a great deal from you, especially from the different perspectives you consistently brought to every discussion. Even more importantly, I learned from the moments when you simply said, "I don't know." Seeing someone with your experience acknowledge uncertainty taught me that not having all the answers is actually an essential part of doing research. During the many moments when I felt that my work was going nowhere, you never seemed to lose confidence in the process. Your confidence gave me the strength to keep going, even when I could not yet see any sign that the work was moving forward.

I would also like to extend my gratitude to Faustine. Knowingly or unknowingly, your openness in sharing your experience with imposter syndrome during a lunch at the Sagan Workshop helped me navigate my own struggles during the final stages of my PhD. Knowing that these feelings are shared by so many researchers allowed me to better understand them and to face them with a different perspective. To everyone in PSILab and AGO, thank you for your wonderful company throughout these years. A special thank

you goes to Malavika and Maxime for their constant emotional support. I would also like to thank Gilles, Agustina, Carles, Patricia, Iremsu, Ludo, Guillaume, Jyotirmay, Jahanvi, Valentin, Sandrine, Carl-Henrik, Antoine and Irene. Your conversations, laughter, and lunch breaks made this journey one that I will always remember with warmth.

To Belgium, my second home, thank you for being the place where I never felt like a foreigner. To everyone who feels concerned by these words, Belgian or, like me, Belgian at heart, thank you for making this country what it is. Thank you for your kindness, for the countless everyday interactions that made me feel welcomed, and for the genuine sense of belonging I have found here.

To my closest friends, Jana, Malavika, and Maitha, thank you for always being there. Every difficult moment, every doubt, and every setback eventually found its way to one of you, and I always left our conversations with a little more clarity. Your unconditional support and your willingness to always make time for me have carried me through this journey and played a tremendous role in bringing me to this point. I would also like to extend my gratitude to Muna. Your perspective, wisdom, and guidance have had a profound impact on me.

To my family, thank you for giving me the roots that made this journey possible. To my father, Amer, thank you for always believing in me, for encouraging me to pursue my ambitions, and for supporting every decision that led me here. To my mother, Sawsan, whose dedication to mathematics shaped me from an early age, thank you for planting the seeds that eventually grew into my love for mathematics and physics. Thank you for supporting my decision to leave, despite how difficult it must have been for you. To my brother, Ahmad, thank you for being someone I could rely on. And to my little sister, Nour, this journey came at the cost of time that I wish I could have spent with you. I left when you were still an adolescent, and now you are about to get married. Thank you for understanding the distance that this journey required.

Finally, I would like to thank Lebanon, my small, beautiful, and wounded country. You gave me the foundation upon which everything else was built. I am deeply grateful to the school teachers who first inspired my love for mathematics and physics: Mme Racha Barakeh, Mme Maysa Akel, and Mme Rayan Fakhri, whose encouragement planted in me the idea that I could one day pursue higher studies. I am equally grateful to Prof. Adnan Naja and Prof. Ahmad Osman for their continuous dedication, and to Dr. Hani

Hamzeh and Dr. Roy Aad for the sacrifices they made in support of the students of the Lebanese University. I hope that one day I can return and give back what these years have given me, and offer future students the same opportunities that were once offered to me.

Contents

Declaration	iii
Abstract	iv
Acknowledgements	vi
Acronyms	xii
1 Introduction	1
1.1 Planets through history	2
1.2 Studying exoplanets	3
1.2.1 Exoplanets through indirect detection	4
1.2.2 Exoplanets through direct imaging	9
1.3 Ground-based high-contrast imaging	11
1.3.1 Enabling technologies in HCI	11
1.3.2 The anatomy of a high contrast image	16
1.3.3 Observing strategies	20
1.3.4 PSF subtraction techniques	22
1.4 Detection techniques in HCI	25
1.4.1 Signal-to-noise ratio (S/N)	26
1.4.2 Forward models	28
1.4.3 Other detection techniques	30
1.4.4 Performance of detection techniques in HCI	32
1.4.5 Observables in HCI	33
1.5 Direct imaging surveys and their scientific contributions	37
1.5.1 First generation instruments surveys	38
1.5.2 Second generation instruments surveys	40
1.5.3 What did we learn?	44
1.6 Scope and outline of the dissertation	49
2 Clustering the SHINE survey	52
2.1 Dataset description	53
2.1.1 HC-DC pre-processing	53

CONTENTS	x
2.1.2 Selection of datasets	55
2.2 Choice of relevant parameters	56
2.3 Clustering	60
2.3.1 Effectiveness of clustering	60
2.3.2 Results of clustering	60
2.3.3 Noise correlation test	62
2.4 Improving the clustering	64
3 Regime Switching Model algorithm	67
3.1 RSM explained	68
3.1.1 RSM concept	68
3.1.2 RSM applied to HCI	69
3.1.3 Updates to RSM	73
3.1.4 RSM performance	73
3.2 Parameters in RSM	75
3.2.1 PSF subtraction parameters	76
3.2.2 RSM parameters	77
3.2.3 Parametrisation in the thesis	81
3.3 Achieving better RSM performance	82
3.3.1 Combination of different parameter ranges	84
3.3.2 Combination of different nights	84
4 Detection threshold in RSM	86
4.1 Noise behavior in RSM	87
4.1.1 Fitting a lognormal distribution to the RSM noise	88
4.1.2 Threshold using maximum F1 score	89
4.1.3 Comparison and reliability of the detection thresholds	92
4.2 Sensitivity limits	94
4.2.1 Influence of threshold and PSF subtraction techniques	95
4.2.2 Optimal usage of RSM	97
4.2.3 Reliability of the contrast curves	98
5 RSM application to SHINE F150 sample	102
5.1 Contrast limits improvements	103
5.2 Signals in RSM maps	104
5.2.1 Comparison of SHINE survey findings with other algorithms	106
5.2.2 Identification of new point sources	108

CONTENTS	xi
5.3 Exoplanet candidate: π Ara	118
5.3.1 Context	118
5.3.2 RSM application on π Ara	119
5.3.3 Nature of the detected signals	121
5.4 Discussion	123
6 RSM applied to Epsilon Eridani with Keck	130
6.1 Introduction	131
6.2 Observations and pre-processing	133
6.3 RSM applied to ϵ Eridani	134
6.3.1 RSM application	134
6.3.2 Improvements to the RSM analysis	135
6.4 Discussion	140
6.5 Summary	142
7 Conclusions & perspectives	144
7.1 Conclusions	144
7.2 Perspectives	146
7.2.1 Future developments of the RSM framework	146
7.2.2 Future scientific applications	147
Appendices	
A F150 observations	149
B Variation of optimal RSM parameters across cluster centers and individual datasets	161
C Comparison with the Dahlqvist contrast curve definition	164
D Comparing RSM noise distribution models	166
E Maximum F1 score threshold generalized across a cluster	169
Bibliography	171

Acronyms

ADI	Angular Differential Imaging
AO	Adaptive Optics
APCA	Annular Principal Component Analysis
ASDI	Angular and Spectral Differential Imaging
AU	Astronomical Unit
CMD	Color Magnitude Diagram
DM	Deformable Mirror
EIDC	Exoplanet Imaging Data Challenge
ELT	Extremely Large Telescope
FM KLIP	Forward-Modelled Karhunen–Loève Image Projection
FM LOCI	Forward-Modelled Locally Optimized Combination of Images
FOV	Field of View
FP	False Positive
FPR	False Positive Rate
FWHM	Full Width at Half Maximum
HC-DC	High Contrast Data Center
HCI	High-Contrast Imaging
IRDIS	Infra-Red Dual-beam Imager and Spectrograph
IWA	Inner Working Angle
JWST	James Webb Space Telescope
LLSG	Local Low-rank plus Sparse plus Gaussian-noise decomposition
LOCI	Locally Optimized Combination of Images
LWE	Low Wind Effect
NA SODINN	Noise-Adaptive Supervised exOplanet Detection via Direct Imaging with Neural Networks
NIRC2	Near-Infrared Camera 2
NIRCam	Near Infrared Camera (JWST)
NMF	Non-negative Matrix Factorization
PA	Parallactic Angles
PACO	PAch COvariance
PSF	Point Spread Function

RDI	Reference Differential Imaging
RSM	Regime Switching Model
RTC	Real-Time Computer
RV	Radial Velocity
S/N	Signal-to-Noise Ratio
SDI	Spectral Differential Imaging
SHINE	SpHere INfrared survey for Exoplanets
SPHERE	Spectro-Polarimetric High-contrast Exoplanet REsearch
SR	Strehl Ratio
TP	True Positive
TPR	True Positive Rate
VIP	Vortex Image Processing
VLT	Very Large Telescope
WDH	Wind Driven Halo
WFS	WaveFront Sensor

Chapter 1

Introduction

Abstract. This introductory chapter provides an overview of exoplanet detection, with a particular focus on the direct imaging of exoplanets and the technological and methodological developments that have enabled the detection of giant planets. The chapter begins in Section 1.1 with the historical efforts to understand planets and planetary motion. Building on this foundation, Section 1.2 introduces the modern techniques used to detect exoplanets, including indirect methods (Sect. 1.2.1) and direct imaging (Sect. 1.2.2). Section 1.3 focuses on the direct imaging of giant planets with ground-based telescopes, detailing the enabling technologies for stellar-light suppression and aberration correction in Sect. 1.3.1, the resulting anatomy of a high-contrast image in Sect. 1.3.2, the observing strategies used to mitigate quasi-static noise in Sect. 1.3.3, and the PSF-subtraction techniques employed to reveal planetary signals in Sect. 1.3.4. Section 1.4 discusses how potential signals are identified and assessed. This includes the signal-to-noise ratio in Sect. 1.4.1, forward modeling in Sect. 1.4.2, alternative detection techniques in Sect. 1.4.3, performance assessments in Sect. 1.4.4, and the scientific outcomes enabled by these methods in Sect. 1.4.5. Section 1.5 is devoted to high-contrast imaging surveys, presenting an overview of the first generation instruments surveys in Sect. 1.5.1, second generation instruments surveys in Sect. 1.5.2 and their scientific results in Sect. 1.5.3. Finally, Section 1.6 discusses the motivations for this PhD and defines its scope, objectives, and main contributions.

1.1 Planets through history

The observation of planets in the night sky marked one of the earliest foundations of astronomy. Ancient observers noticed that a small number of bright objects did not remain fixed relative to the background stars, but instead changed position over time. Because of this apparent motion, these objects were referred to as planets, from the Greek word *planētai*, meaning wanderers. Their motions were carefully tracked and became central to early astronomical observations, calendars, and cosmological interpretations.

Early descriptions of planetary motion were developed within a geocentric framework by civilizations including Babylonian, Greek, Indian, and Islamic astronomers. These models progressively improved the prediction of planetary trajectories through increasingly sophisticated geometrical constructions such as epicycles. A major transition occurred during the 16th and 17th centuries with the heliocentric model proposed by Copernicus, Kepler's laws of planetary motion, and Newton's law of gravitation, which established the foundations of celestial mechanics.

The invention of the telescope further transformed planetary science. Observations by Galileo Galilei in the early 17th century, including the phases of Venus and the moons of Jupiter, provided strong support for the heliocentric model and expanded the known architecture of the Solar System. Later developments such as spectroscopy, astrometry, and photometry enabled the characterization of planetary atmospheres, orbital properties, and physical compositions.

These advances naturally led to the question of whether planets could exist around stars other than the Sun. Over the past decades, numerous observational techniques have been developed to address this question. The first confirmed detections of exoplanets were achieved in 1992 around a pulsar using pulsar-timing measurements (Wolszczan & Frail, 1992), followed in 1995 by the discovery of 51 Pegasi b around a main-sequence star using the radial velocity method (Mayor & Queloz, 1995). Since then, thousands of exoplanets have been discovered through methods including radial velocity, transits, direct imaging, microlensing, and astrometry (a general review in Perryman, 2018). Figure 1.1 presents the number of confirmed exoplanets as a function of the detection method.

The diversity of the detected planetary systems has profoundly transformed our understanding of planet formation and evolution, revealing architectures and physical properties very different from those observed in the Solar System (Winn & Fabrycky, 2015).




Figure 1.1: Number of confirmed exoplanets discovered since 1992, categorized by detection method. Image produced by NASA Exoplanet Archive.

This growing diversity has also motivated the establishment of a definition for what constitutes a planet. A planet is commonly defined as a body that is nearly spherical due to its self-gravity, has cleared the neighborhood of its orbit, and has a mass below approximately $13M_J$, the threshold above which deuterium fusion can occur and objects are generally classified as brown dwarfs. However, this definition remains debated, and alternative definitions based on the planet-to-host mass ratio (Grether & Lineweaver, 2006) or the formation mechanism of the object (Bowler et al., 2020) have been proposed.

1.2 Studying exoplanets

Because of the small angular separation between exoplanets and their host stars, combined with the large brightness ratio between them, directly resolving exoplanets has long remained impossible. Consequently, early exoplanet detections relied primarily on *indirect methods*, which infer the presence of a planet through its gravitational or photo-




Figure 1.2: Distribution of confirmed exoplanets as a function of orbital separation and planetary mass, categorized by detection method. Image produced by NASA Exoplanet Archive.

metric effects on the host star. These techniques have played a foundational role in the discovery and characterization of thousands of exoplanets spanning a broad diversity of planetary systems. Figure 1.2 illustrates the distribution of planetary masses and orbital separations revealed by the different detection techniques.

1.2.1 Exoplanets through indirect detection

Pulsar timing

Pulsars are highly magnetized, rapidly rotating neutron stars that emit narrow beams of radio waves. As these beams sweep across Earth, they are observed as pulses arriving at extremely regular intervals, ranging from milliseconds to seconds. A planet orbiting a pulsar induces a small reflex motion through its gravitational interaction with the host star. This motion is detected as periodic variations in the arrival times of the pulses, which arrive slightly earlier when the pulsar moves toward Earth and slightly later when it moves away.

Pulsar timing is extraordinarily precise and can detect companions with masses as low as that of the Moon. This technique led to the first confirmed discovery of exoplanets in 1992 by Aleksander Wolszczan and Dale Frail (Wolszczan & Frail, 1992). They identified two Earth-mass planets orbiting the pulsar PSR B1257+12, with orbital periods of approximately 66 and 98 days.

Nevertheless, this method is limited by the rarity of pulsars suitable for high-precision timing, as well as by the unusual and still poorly understood processes through which planets may form in such extreme environments.

Radial Velocity Method

The radial velocity method detects exoplanets by measuring periodic variations in a star's radial velocity caused by the gravitational pull of an orbiting planet. While the star exerts a gravitational force on the planet, the planet also exerts a force on the star, much smaller relatively, causing it to move in a small orbit around the system's center of mass. This motion induces periodic variations in the star's radial velocity, which result in Doppler shifts in the star's spectrum.

A major milestone for this technique was reached in 1995 with the Nobel-winning discovery of 51 Pegasi b by Michel Mayor and Didier Queloz (Mayor & Queloz, 1995), marking the first confirmed detection of an exoplanet orbiting a main-sequence star. The companion, characterized by a mass comparable to that of Jupiter and an orbital period of only 4.2 days, revealed the existence of a new class of planets later referred to as *hot Jupiters*, namely giant gaseous planets orbiting extremely close to their host stars.

By analyzing the radial velocity curve, astronomers can infer key planetary properties, including the minimum mass, orbital period, and orbital eccentricity. The method is especially effective for detecting massive close-in planets. However, advances in instrumentation, such as the HARPS spectrograph (Mayor et al., 2003) and the ultra-precise ESPRESSO spectrograph (Pepe et al., 2014), have significantly improved the sensitivity of the technique, enabling the detection of lower-mass planets down to the rocky-planet regions and planets orbiting at larger separations from their host stars.

Transit Method

A transit occurs when a planet passes in front of its host star along the observer's line of sight, producing a small but measurable decrease in the star's apparent brightness. These periodic dimmings, detected through photometric monitoring, reveal the presence of an orbiting planet.

By analyzing transit light curves, astronomers can determine several planetary properties, including the planet's radius, orbital period, and orbital inclination. When combined with radial velocity measurements, transit observations also make it possible to estimate the planet's mean density, providing insights into its internal composition. In addition, spectroscopic observations obtained during transits and secondary eclipses enable the study of exoplanetary atmospheres, revealing information about their chemical composition, temperature structure, and the possible presence of clouds or hazes.

The transit method has been instrumental in the discovery of thousands of exoplanets, particularly through space missions such as Kepler (Borucki et al., 2010) and TESS (Ricker et al., 2015). These missions have been groundbreaking in their ability to detect a large fraction of the exoplanets known today, including small Earth-sized planets, and have profoundly improved our understanding of the diversity of planetary systems and exoplanet populations. Ground-based surveys such as TRAPPIST (Gillon et al., 2011) also played a major role in the field through the discovery of the TRAPPIST-1 system (Gillon et al., 2017), one of the most remarkable known planetary systems, composed of seven Earth-sized planets orbiting an ultracool dwarf star. More recently, facilities such as James Webb Space Telescope (Gardner et al., 2006) have enabled the atmospheric characterization of transiting exoplanets through high-precision spectroscopy.

Gravitational Microlensing

Gravitational microlensing occurs when the light from a distant background star is bent and magnified by the gravitational field of a foreground object, typically a star that may host a planet, passing between the source and the observer. If a planet is present, it can induce a brief, distinctive perturbation in the resulting light curve, particularly when it lies near the region of maximum magnification known as the Einstein ring. The characteristics of this perturbation depend on the planet's mass and its projected separation from the host star.

Microlensing is especially powerful for detecting low-mass planets, planets on wide orbits, and even free-floating planets that do not orbit any star. Moreover, it can probe planetary systems at great distances, including those in the Galactic bulge, which are often inaccessible to other detection methods. However, microlensing events are rare, and non-repeating, which makes follow-up observations and confirmations difficult. Despite these challenges, the method offers a unique window into the galactic population of exoplanets, particularly those in regimes underexplored by other techniques.

Astrometry

Astrometry detects exoplanets by precisely measuring the subtle positional shifts of a star in the plane of the sky, caused by the gravitational pull of an orbiting planet. Unlike the radial velocity method, which measures changes in a star's velocity along our line of sight, astrometry tracks the star's motion perpendicular to the line of sight, offering a complementary perspective.

By monitoring these minute displacements, astronomers can determine the planet's orbital parameters, including period, eccentricity, and inclination, and crucially, can directly derive the planet's true mass, rather than just a minimum mass estimate. This makes astrometry a powerful tool for constructing a complete dynamical picture of exoplanetary systems.

The method is especially well-suited for detecting massive planets on wide orbits around nearby stars. While technically demanding due to the extreme precision required, recent and upcoming astrometric facilities, most notably the *Gaia* mission (Gaia Collaboration et al., 2016), are significantly advancing the potential of astrometry for exoplanet discovery and orbital characterization.

What Have We Learned?

Indirect detection techniques have profoundly expanded our understanding of planetary systems, revealing a diversity that far exceeds what was once imagined based on the Solar System alone. One of the most transformative findings is the wide variety of exoplanet types, including Earth-sized planets, super-Earths (1–2 Earth radii), mini-Neptunes (2–4 Earth radii), gas giants, and ice giants. Several of these populations, particularly super-Earths and mini-Neptunes, have no counterparts in the Solar System, indicating that planet formation can produce a much broader range of planetary architectures and physical properties than previously anticipated. A comprehensive review of exoplanet demographics is provided by Cloutier (2026).

Among the most surprising discoveries are hot Jupiters. These objects, such as 51 Pegasi b, challenge classical models of planet formation, which predict that giant planets should form beyond the snow line. A comprehensive review of the proposed formation and migration mechanisms for hot Jupiters is provided by Dawson & Johnson (2018). Similarly, the detection of planets with extremely eccentric orbits (e.g., HD 80606 b (Naef et al., 2001) and HD 28254 b (Wu & Murray, 2003), with $e \approx 0.95$) points toward past dynamical interactions, such as planet–planet scattering or Kozai-Lidov cycles induced by distant stellar companions. The discovery of free-floating planets, which are not gravitationally bound to a host star, further suggests that planet ejection may be a common outcome of early dynamical evolution and gravitational interactions within planetary systems.

Statistical analyses, especially from the Kepler mission, indicate that planets are extremely common throughout the Galaxy. On average, there is at least one planet per star in the Milky Way (Howard et al., 2012). Small planets dominate the population: Earth-sized and mini-Neptune planets are the most common, with occurrence rates above 30–50% around both Sun-like and M-dwarf stars (Howard et al., 2012; Dressing & Charbonneau, 2015). By contrast, hot Jupiters are relatively rare, with occurrence rates below 1% (Wright et al., 2012).

Finally, detection methods have enabled the identification of planets located within the habitable zone, the range of orbital distances around a star where a planet could sustain liquid water at its surface. Kepler has discovered dozens of such candidates, particularly around M-dwarf stars (Kopparapu et al., 2013). Ongoing and upcoming facilities are expected to greatly expand the characterization of exoplanet atmospheres. JWST is exploring a broad diversity of exoplanetary atmospheres over a wide range of masses, temperatures, and compositions, while ARIEL will target a diverse sample of approximately 1000 extrasolar planets through simultaneous observations at visible and infrared wavelengths.

Although *indirect detection* methods have been highly successful in advancing the field of exoplanets, they remain intrinsically biased toward planets orbiting close to their host stars and generally provide only indirect constraints on planetary properties. Direct imaging is complementary to these techniques because of its sensitivity to young giant planets at wide orbital separations, as illustrated in Fig. 1.2, and its ability to directly collect photons emitted or reflected by the planet itself. This enables detailed atmospheric characterization, astrometric monitoring, and the study of planet-disk interactions, thereby opening a unique window onto the formation and evolution of planetary systems.

1.2.2 Exoplanets through direct imaging

Direct imaging of exoplanets is the detection of a planet as a point source of light spatially separated from its host star. This method is also known in the field as **high-contrast imaging (HCI)**, where the contrast is defined as the apparent flux ratio (F_p/F_*) between the planet and its host star. High contrast refers to a small flux (typically $F_p/F_* < 10^{-3}$) ratio, corresponding to faint planetary objects, while low contrast refers to a larger ratio, corresponding to a brighter planetary signal (or the realm of binary stars / brown dwarfs). The achieved contrast between the planetary signal and its host star has been a key metric of progress in the field, with the ultimate goal of reaching contrasts similar to those in the Solar System, such as a present-day (4.5 Gyr) Jupiter–Sun contrast of 10^{-9} , or an Earth–Sun contrast of 10^{-10} in the optical.

The light detected from a planet can originate either from stellar light reflected by the planet toward the observer or from the planet's own thermal emission, arising predominantly from the energy acquired during its formation. The brightness of a star depends primarily on its mass, which determines the nuclear fusion processes occurring in its interior and therefore its luminosity. In contrast, planets do not sustain nuclear fusion (mass $< 13M_J$) and their luminosity mainly originates from the gradual release of this accumulated energy. As a result, planetary brightness depends primarily on the planet's mass and age: it increases with mass and decreases as the planet evolves and cools over time. Other important factors affecting the brightness of the planetary signal include its orbital separation from the host star, its formation scenario (e.g., cold-start or hot-start models), and its atmospheric composition. In addition to these dependencies, the planet/star contrast also strongly depends on the observing wavelength. Planetary thermal emission peaks in the infrared, making this wavelength range particularly favorable for the detection of young giant planets. The near-infrared, in particular, represents an optimal compromise between achievable contrast and angular resolution for current high-contrast imaging instruments.

The first direct imaging detection of a planetary-mass companion was reported by Chauvin et al. (2005) in April 2004, corresponding to the discovery of a young, $\sim 5M_J$, self-luminous companion at a projected separation of 41 au from a $\sim 25 M_J$ brown dwarf. The relatively late success of direct imaging compared to other detection methods is due to the major challenges faced by this technique, including the need to achieve very high contrast at small angular separations from the host star, where planetary signals are typically buried under diffraction patterns. Direct detections became possible in the early 2000s through the combination of several major technological developments, including

adaptive optics systems for correcting atmospheric turbulence, coronagraphs for attenuating stellar light, large aperture 8-10 m class telescopes capable of resolving small angular separations, and advanced post-processing techniques designed to suppress residual instrumental aberrations.

With the development of these technologies, the high-contrast imaging field has focused on observing young (10–100 Myr) nearby ($d < 100$ pc) stars in the infrared using 8–10 m class ground-based telescopes (VLT, Keck, Gemini, Subaru, LBT) as well as dedicated high-contrast instruments such as NACO (Lenzen et al., 2003), LMIRCam (Skrutskie et al., 2010), GPI (Macintosh et al., 2014), SCExAO (Jovanovic et al., 2015), NIRC2 (Serabyn et al., 2017) and SPHERE (Beuzit et al., 2019), in order to search for **young Jupiter-like planets** at wide orbital separations. Ongoing technological efforts aim to reach higher contrast levels in order to detect terrestrial planets. Future space missions such as the Roman Space Telescope and the Habitable World Observatory (HWO) in the optical, as well as extremely large ground-based telescopes such as the ELT, operating in the infrared, are expected to significantly extend the detection capabilities of direct imaging, enabling the detection of lower-mass planets at smaller orbital separations.

Direct imaging has been complementary to indirect detection methods by probing wide-orbit planets that cannot be accessed by most indirect techniques. The existence of planets at wide orbital separations challenges the traditional formation pathway that explains close-in planets through core accretion and suggests that alternative formation mechanisms, such as gravitational instability, may also play an important role at large separations. Direct imaging has also enabled the detection of circumstellar disks in scattered light, providing valuable insight into disk–planet interactions. In addition, direct imaging allows the characterization of planets by analyzing the light originating from the planets themselves, enabling atmospheric properties such as temperature, chemical composition, and cloud structure to be probed.

Direct imaging provides measurements of both the location of the planet (**astrometry**) and its brightness (**photometry**). Astrometric measurements enable the identification and characterization of the orbital motion of the planet when multiple observational epochs are available. Photometric measurements, when combined with evolutionary models, allow estimates of the planetary mass. In addition, in the case of non-detections, detection limits combined with evolutionary models can be used to derive occurrence-rate constraints for specific classes of planets.

1.3 Ground-based high-contrast imaging

Imaging a planet orbiting a star represents a major technological and instrumental challenge. This section presents the main components and techniques of HCI, which enable contrasts on the order of 10^{-6} required for the detection of massive exoplanets. Section 1.3.1 introduces the enabling technologies, including large-aperture telescopes, adaptive optics systems, and coronagraphs. Section 1.3.2 describes the general anatomy of an HCI observation. Section 1.3.3 presents the observing strategies developed to mitigate residual instrumental aberrations, while Section 1.3.4 discusses the PSF-subtraction techniques used to remove the remaining stellar light in the recorded images.

1.3.1 Enabling technologies in HCI

Angular resolution and the telescope aperture

For distant astronomical sources such as stars, the wavefront reaching the telescope can be approximated as planar across the aperture (in the absence of wavefront perturbations). When this wavefront passes through a finite aperture, diffraction redistributes the incoming light into a structured intensity pattern, called the **point spread function (PSF)** (Racine, 1996), rather than forming a perfect point-source image. For a circular aperture, this diffraction pattern is described by the Airy function shown in Figure 1.3. The angular extent of the central lobe of the PSF is given by $\theta \approx 1.22\lambda/D$, where λ is the observing wavelength and D is the telescope diameter. The **full width at half maximum (FWHM)** of the PSF core is approximately $1.04\lambda/D$.

In a star–planet system, the planetary companion is located at a slightly different angular position on the sky and therefore produces a second PSF in the focal plane that is shifted with respect to the stellar PSF. The ability to distinguish these two signals depends on the angular width of the PSF, which scales with λ/D . A larger telescope aperture therefore produces a narrower PSF and improves the ability to spatially separate the exoplanet signal from the stellar diffraction pattern. This is particularly critical for detecting planets located at very small angular separations from their host stars. For example, accessing a projected separation of about 2 au for a star located at 50 pc at a wavelength of $\lambda = 1.65\mu\text{m}$ (H-band) requires a telescope diameter of at least 8 m. This requirement motivates the use of current 8–10 m class telescopes and the development of next-generation 30–40 m class facilities such as the Extremely Large Telescope (ELT), which are essential for improving access to planets at very small angular separations from their host stars.




Figure 1.3: *Left:* Point spread function of a circular aperture showing the Airy-disk and the three first Airy rings. Image from Pedrotti et al. (2017). *Right:* Radial profile of the same PSF showing the full width at half maximum (FWHM) measure in red. Image from Majewski lecture notes (University of Virginia).

Atmospheric turbulence and adaptive optics

In real ground-based observations, the image of a star is not a diffraction-limited PSF but a blurred image affected by atmospheric turbulence. As starlight propagates through the Earth's atmosphere, variations in temperature and wind velocity across different atmospheric layers introduce phase distortions in the incoming wavefront. The spatial scale over which the wavefront can be considered approximately planar is characterized by the coherence length r_0 . Over apertures larger than r_0 , the wavefront is no longer planar and optical aberrations are introduced. One example of such aberrations is tip-tilt, which modifies the apparent direction of the incoming light. As a result, different coherent patches of size r_0 across the telescope aperture produce slightly displaced images in the focal plane, leading to a broadened stellar PSF. In addition to these spatial distortions, atmospheric turbulence also evolves rapidly in time. This temporal variability is quantified by the coherence time τ_0 , which describes the characteristic timescale over which the wavefront changes significantly.

These atmospheric effects strongly degrade the PSF observed in ground-based images and produce what is known as a seeing-limited PSF, as shown in Fig. 1.4 (left). In this regime, the angular resolution is determined by the atmospheric coherence length r_0 rather than by the telescope diameter, and is approximately given by $\theta \approx \lambda/r_0$. For the example considered previously, this corresponds to resolving spatial scales of about 50 au instead of 2 au, illustrating the strong limitations that atmospheric turbulence imposes on high-contrast imaging observations.




Figure 1.4: Adaptive optics corrections in H band at VLT/SPHERE. *Left:* Long exposure of a point source without AO correction. *Right:* AO corrected image in nominal seeing conditions. Taken from Petit et al. (2016)

To overcome these limitations, high-contrast imaging relies on adaptive optics (AO), which corrects aberrations at small spatial scales and restores the incoming wavefront toward a planar shape while operating at high speed to follow its rapid temporal variations. The AO system, as shown in Fig. 1.5, is composed of three main components: a deformable mirror (DM), a wavefront sensor (WFS), and a real-time computer (RTC). The DM is relying on a grid of actuators (N_{act}) with a spatial sampling comparable to the coherence length r_0 , allowing it to dynamically reshape the mirror surface to counteract distortions in the incoming wavefront. The WFS measures how the current wavefront deviates from a planar reference, accounting for the continuous temporal evolution of atmospheric turbulence, and can rely on different technologies such as the Shack–Hartmann wavefront sensor (Shack & Platt, 1971) or the Pyramid wavefront sensor (Ragazzoni & Farinato, 1999). The RTC rapidly processes this information to compute the necessary corrections, which are then applied to the DM in a closed feedback loop operating on millisecond timescales. This correction must be faster than the atmospheric coherence time τ_0 (typically 3ms) so that the AO system can track the evolving distortions and compensate for the wavefront before it becomes fully decorrelated.

The impact of AO correction is clearly visible in the PSF shown in Fig. 1.4 (right), where the starlight is concentrated into a much sharper core. The performance of an AO system is typically quantified using the Strehl ratio, defined as the ratio between the observed peak intensity of a star and the peak intensity expected for a theoretical diffraction-limited PSF. Although AO correction never reaches a strehl ratio of 100%, the level of correction is nevertheless crucial and has enabled the detection of bright companions even without the use of coronagraphy, such as the brown dwarf companion HR 7672B (Liu et al.,




Figure 1.5: Illustration of a simplified adaptive optics system. Image taken from Skaf et al. (2022).

2002). First-generation AO systems such as VLT/NACO (Fusco et al., 2004), Keck AO (Wizinowich et al., 2000), and Subaru/CIAO (Tamura et al., 1998) achieved Strehl ratios of around 50-60%, while second-generation systems, known as extreme adaptive optics, such as VLT/SPHERE (Fusco et al., 2006), Gemini/GPI (Macintosh et al., 2014), and Subaru/SCEAO (Jovanovic et al., 2015), can reach values of 90% or higher under optimal conditions. These improvements have had a major impact on the development of high-contrast imaging techniques (see, e.g., (Chauvin, 2018)).

High exoplanet-star contrast and coronagraphy

After restoring the Airy pattern of the PSF with adaptive optics, the main challenge in high-contrast imaging remains the large brightness difference between the stellar PSF and the planetary PSF. In practice, the observation does not appear as two distinct PSFs located next to each other; instead, the planetary signal is buried within the diffraction pattern of the stellar PSF, especially at small angular separations from the star, which is precisely the region of greatest interest for HCI. This effect is illustrated in Figure 1.6, where a companion with a contrast of 10^{-3} has a peak intensity lower than the diffraction rings of the stellar PSF, making it difficult to distinguish from the surrounding stellar diffraction structure.




Figure 1.6: Illustration of a simulated stellar PSF with a planetary companion of contrast 10^{-3} located at a separation of $2\lambda/D$. *Left:* horizontal cut through the stellar and planetary PSFs. *Right:* two-dimensional representation of the stellar and planetary PSFs.

To overcome this limitation, Bernard Lyot developed the coronagraph originally to study the solar corona (Lyot, 1939). The modern-day Lyot coronagraph is an optical system designed to suppress on-axis stellar light while transmitting off-axis light from nearby sources such as planetary companions. As illustrated in Figure 1.7, it consists of an opaque mask placed in the focal plane to block the central core of the stellar PSF, followed by a pupil-plane stop, known as the Lyot stop, which removes diffracted stellar light redistributed near the edges of the pupil. The off-axis planetary signal remains largely unaffected by these elements and can therefore be detected more easily in the final image. As shown in the panel (f) of Figure 1.7, the diffraction pattern is attenuated compared to the case shown in panel (b).

The performance of a coronagraph is typically characterized using several metrics, including the raw contrast, which quantifies the level of residual stellar light suppression, and the inner working angle (IWA), which defines the smallest angular separation at which a companion can be detected. In the case of the classical Lyot coronagraph, the achievable raw contrast remains limited and the IWA is typically about $3-4\lambda/D$, as illustrated in panel (c), defined by the size of the opaque mask. Modern coronagraphs, such as the Four-Quadrant Phase Mask (4QPM, Rouan et al., 2000) and the vortex phase mask (Mawet et al., 2005), suppress stellar light through destructive interference rather than simple amplitude blocking. These designs provide significantly improved raw contrast and enable smaller inner working angles, typically reaching separations of about




Figure 1.7: The Lyot coronagraph. A telescope with a circular unobstructed pupil (a) points at a star along the optical axis of the telescope and instrument. The telescope pupil is reimaged through the instrument optics to a focal plane (b), showing an Airy core surrounded by diffraction rings. A focal plane mask (FPM) then blocks the central region of the stellar image (c). Another set of optics reimages c to the pupil plane in panel d. The removal of the Airy core at the FPM redistributes light into a broad ring around the edge of the reimaged telescope pupil. The Lyot stop [more generally referred to as a pupil plane mask (PPM)] blocks this starlight (e), at the cost of reduced off-axis throughput and decreased angular resolution. A final set of relay optics then forms the science camera focal plane (SCFP) (f). Image and caption taken from Kenworthy & Haffert (2025)

$1-2\lambda/D$, and are now widely used in high-contrast imaging instruments installed on several large ground-based telescopes, including the VLT, Subaru, Keck and LBT. A detailed review of these coronagraphic techniques may be found in Kenworthy & Haffert (2025).

1.3.2 The anatomy of a high contrast image

After coronagraphic suppression of the stellar core and diffraction pattern, reaching raw contrast levels typically of the order of 10^{-3} , the image is still not completely free of residual structures. Instead, several additional features become visible, including instrumental diffraction patterns, residual atmospheric aberrations not fully corrected by adaptive optics, and quasi-static instrumental aberrations. Figure 1.8 highlights the different




Figure 1.8: Example SPHERE high-contrast image. White dashed circles highlight the DM control radius, the spider diffraction patterns, the four satellite spots, and the Arago spot at the center. A representative speckle at an arbitrary position is indicated in blue.

structures typically present in a high-contrast image. These features can be broadly divided into two categories: static structures appearing at predictable locations and non static structures appearing at random locations. The static features, highlighted in white in Figure 1.8, include the following:

- Arago spot at the center: produced by diffraction around the focal-plane coronagraphic mask.
- Adaptive optics control radius: angular extent of the DM-corrected region in the image, located at approximately $N_{\text{act}}\lambda/(2D)$. Due to the finite number of actuators in the DM, only spatial frequencies up to a cutoff value can be corrected. Beyond this radius, the wavefront remains uncorrected and the image becomes seeing-limited.
- Diffraction from the telescope pupil and spider vanes create structures that are not fully blocked by the Lyot stop, producing residual diffraction patterns aligned with the spider directions.
- Satellite spots: each satellite spot is a replica of the stellar PSF, created on purpose by injecting a specific shape on the DM, used to determine the stellar position behind the coronagraph and to provide photometric calibration.

These structures are well understood and can be accounted for in the PSF modeling described later. A more detailed description of high-contrast imaging features in SPHERE observations is provided in Cantalloube et al. (2019).




Figure 1.9: Illustration of the simulation of the wind driven halo (WDH) (left), and the low wind effect LWE in the right. Images taken from (Cantalloube et al., 2019)

In contrast, non static randomly distributed features, referred to as speckles, can be divided into two main categories:

- *Atmospheric residual speckles:* since AO correction is never perfect (the Strehl ratio does not reach 100%), residual wavefront errors remain. These residuals evolve rapidly on millisecond timescales and tend to average out over long exposures (leaving a smooth halo).
- *Quasi-static speckles:* these originate from instrumental and optical aberrations, including non-common path aberrations arising in the optical path between the wavefront sensor and the science detector, which are therefore not corrected by the adaptive optics control loop (see Figure 1.5). These speckles evolve slowly, typically on timescales of minutes to hours, and constitute a major limitation for the detection of exoplanets because they have a morphology and angular scale similar to those of planetary signals (of order λ/D). Further discussion of their temporal behavior is given by Hinkley et al. (2007).

In addition to the fundamental limitations of HCI, specific atmospheric and instrumental effects occurring under particular observing conditions can further degrade the achievable raw contrast, the quality of the final post-processed images, and the resulting contrast curves (see Cantalloube et al., 2019). These include the wind-driven halo (WDH), which arises under high wind conditions, and the low-wind effect (LWE), which occurs under low wind conditions. Both effects are illustrated in Figure 1.9.

Wind-Driven Halo

The wind-driven halo (WDH, Cantalloube et al., 2020; Cantalloube et al., 2018) appears as a radial extension of the PSF along a preferred direction in the image, within the adaptive optics corrected region, as shown in Figure 1.9a. This effect occurs when atmospheric turbulence evolves faster than the AO system can compensate for. Because the AO correction operates in a closed control loop, it introduces a temporal delay, known as the servo-lag error (about 3 ms in SPHERE). When the atmospheric perturbations evolve on timescales shorter than this delay, residual phase errors remain correlated along the wind direction, producing the characteristic elongated halo structure. This situation typically occurs during strong high-altitude jet stream conditions, which induce rapid wavefront variations and short coherence times.

More challengingly, because the WDH is aligned with the wind direction in pupil coordinates, it rotates in the detector frame during angular differential imaging (ADI) (later explained in Section 1.3.3) observations as the parallactic angle changes. As a result, this structure behaves differently from static instrumental speckles and is not fully captured by standard PSF subtraction techniques, leaving residual signatures in the final post-processed images. This represents a significant limitation for high-contrast imaging performance.

Low Wind Effect

The low wind effect (LWE, Milli et al., 2018) is characterized by the appearance of secondary lobes around the PSF in non-coronagraphic focal-plane images, typically forming two to four asymmetric PSF structures. This phenomenon arises from temperature differences between the telescope spiders and the surrounding air, which induce local refractive index variations across the pupil through thermal exchanges and introduce phase discontinuities between different pupil regions. As a result, the PSF becomes distorted and multiple secondary lobes appear in the image, as shown in Figure 1.9b. This effect occurs primarily under low wind conditions, which increases the duration of thermal exchanges between spiders and the air flow, typically for wind speeds below about 3 m s^{-1} . It was especially common in SPHERE observations obtained between 2015 and August 2017.

Observations affected by the LWE experience significant contrast degradation. Milli et al. (2018) reported a degradation of the raw contrast by up to a factor of 50 at separations of $0.1''$ between LWE-affected and unaffected images. To mitigate this effect, a dedicated low-emissivity coating was applied to the telescope spiders to reduce temperature

differences with the surrounding air, effectively lowering the wind-speed threshold for the onset of the LWE from about 3 m s^{-1} to approximately 1 m s^{-1} . However, since a large fraction of archival SPHERE exoplanet survey observations were obtained between 2015 and 2017, around 15% of the observations acquired between 2015 and 2018 remain affected by the LWE for wind speeds below 4 m s^{-1} .

Quasi-static speckles, wind driven halo and low wind effect constitute some of the major limitations in high-contrast imaging observations (Marois et al., 2003; Cantalloube et al., 2019) and motivate the development of dedicated observing strategies and post-processing techniques designed to differentiate between speckle noise and planetary signals.

1.3.3 Observing strategies

To address the limitation of quasi-static speckles, HCI observations rely on differential imaging techniques, in which sequences of coronagraphic images of a single target are acquired under varying observational conditions to model the speckle field while preserving the signal of potential companions. In each method, a reference PSF model can be constructed from the image sequence and subtracted from the data, after which the frames are combined to enhance the companion signal while reducing residual speckle noise.

Angular Differential Imaging

Angular Differential Imaging (ADI, Marois et al., 2006) relies on observing in pupil-tracking mode, in which the instrument derotator is operated such that the telescope pupil remains fixed with respect to the instrument. As a result, the field of view, including any off-axis companions, rotates with respect to the detector, while instrumental aberrations and quasi-static speckles remain fixed in the focal plane.

The ADI sequence therefore consists of multiple frames in which the astrophysical signal appears at different detector positions following the evolution of the parallactic angle. A reference PSF can be constructed from the image sequence through the median combination of all the image sequence (see Section 1.3.4 for more advanced PSF model construction). Each frame is then subtracted from this reference, derotated according to its parallactic angle, and combined into a final image in which residual speckle noise is attenuated and the planetary signal is enhanced (see Figure 1.10).




Figure 1.10: Illustration of the ADI sequence. Image taken from Follette (2023).

ADI is the most widely used observing strategy in high-contrast imaging, and most of the classical and advanced post-processing algorithms are based on this approach (see Sections 1.3.4 and 1.4). One limitation of ADI is that its performance depends on the amount of field rotation achieved during the sequence. At small angular separations, limited rotation causes partial inclusion of the companion signal in the PSF model, leading to signal attenuation known as *self-subtraction*. Further limitations of ADI are discussed in Ruane et al. (2019).

Spectral Differential Imaging

Spectral Differential Imaging (SDI, Sparks & Ford, 2002) relies on the chromatic scaling of the stellar PSF and speckle field with wavelength, which follows the λ/D dependence of diffraction. By rescaling images obtained at different wavelengths, the speckle pattern can be modeled and subtracted while preserving the signal of potential companions, whose angular separation remains fixed on the sky.

SDI is commonly used in HCI and is often combined with ADI, known as ASDI, to exploit both spectral and angular diversity and improve the detectability of faint companions. It is implemented in instruments such as SPHERE and GPI through the use of integral field spectrographs, where the simultaneous spectral information can also be used to characterize the atmospheric properties of detected companions.

Reference Differential Imaging

Reference differential imaging (RDI, Lafrenière et al., 2009) consists of constructing a model of the stellar PSF using observations of reference stars that do not host detectable companions. This reference PSF is then subtracted from the science target in order to suppress residual speckles. For this approach to be effective, the reference frames should be obtained with similar stellar brightness, wavelength range, instrumental configuration, and adaptive optics correction quality than the science target.

Although RDI performs better than ADI at small angular separations, in ground-based observations the temporal variability of atmospheric conditions and instrumental aberrations makes it difficult to obtain a reference PSF that perfectly matches the science target. This limitation has been partially mitigated for SPHERE by constructing large PSF libraries from archival observations (Xie et al., 2022). RDI is generally more effective for space-based observations, where the PSF remains more stable in time.

Other observing strategies include Polarimetric Differential Imaging (PDI, Kuhn et al., 2001) and the more recent Coherent Differential Imaging (CDI, Potier et al., 2025), both of which are particularly relevant for the detection of circumstellar disks.

1.3.4 PSF subtraction techniques

The ADI technique presented in Sect. 1.3.3 constitutes the baseline PSF subtraction method for ADI, also referred to as "median subtraction". In this approach, the median combination of all frames in step B of Fig. 1.10 models the static component of the HCI data shown in Fig. 1.8. Atmospheric speckles evolve over time and are only correlated over a limited number of frames, while quasi-static speckles, although long-lived, vary in intensity and position throughout the sequence. A better model for the PSF is one that also accounts for the speckles in the sequence, which can lead to more efficient

noise suppression in the final post-processed image. This motivates the development of PSF subtraction techniques based on constructing a reference PSF that is sufficiently correlated with each frame in the sequence. We present here several PSF subtraction methods used in HCI.

LOCI

Locally Optimized Combination of Images (LOCI, Lafrenière et al., 2007) relies on constructing a linear combination of reference images in localized regions throughout the image, in which each reference frame contributes with an optimized weight. These weights are determined by minimizing the residual intensity within each region through a least-squares approach. Several variations of LOCI have since been developed (Currie et al., 2012; Marois et al., 2014; Thompson & Marois, 2021).

PCA

Principal Component Analysis (PCA/KLIP, Amara & Quanz, 2012; Soummer et al., 2012) decomposes the image sequence into an orthogonal set of basis images, called principal components. Each component represents a structure of the stellar PSF and speckle field, with the first components describing the dominant features showing the highest variance, while higher-order components describe progressively weaker features. A PSF model is then reconstructed using a selected set of principal components and subtracted from each frame of the sequence.

The number of components is one of the most important parameters of the method: using too few components leads to an incomplete reconstruction of the PSF and leaves significant residual speckle structures after subtraction, whereas using too many components can include part of the planetary signal in the PSF model, leading to attenuation of the companion signal in the final image. Another important parameter is the minimum field rotation required to ensure sufficient displacement of a potential companion between frames, typically on the order of λ/D , which depends on the angular separation from the star (the linear displacement scales as $r\Delta\theta$). Absil et al. (2013) introduced annular PCA to account for this radial dependence, as well as for the different behavior of the speckle noise and subtraction requirements at different separations.

LLSG

Local Low-rank plus Sparse plus Gaussian-noise decomposition (LLSG, Gomez Gonzalez et al., 2016) separates the image sequence into three components through a local decomposition applied to small regions of the frames. The low-rank component accounts for the dominant stellar PSF structures and quasi-static speckle field, the sparse component isolates compact astrophysical signals such as potential planetary companions, and the Gaussian-noise component models the remaining small-scale residual noise not captured by the low-rank approximation.

NMF

Non-negative Matrix Factorization (NMF, Ren et al., 2018) provides a low-rank approximation of the image sequence similar to PCA, but enforces non-negative basis components and coefficients. This constraint helps reduce the *over-subtraction* of astrophysical signals during PSF reconstruction, i.e., the attenuation of the planetary signal caused by its correlation in shape and intensity with speckle noise, all following the PSF core shape. A more detailed explanation of this effect is given by Pueyo (2016); Bonse et al. (2025).

4S

Signal Safe Speckle Subtraction (4S, Bonse et al., 2025) is a PSF subtraction technique based on PCA that aims to mitigate both *over-subtraction* and *self-subtraction*. The method relies on a linear noise model designed to prevent the algorithm from learning the structure of the PSF core, thereby limiting the attenuation of the planetary signal. In addition, it employs a signal-invariant loss function that prevents the inclusion of the planetary signal in the PSF model, further preserving its flux during the subtraction process.

More recently, the field of post-processing has started moving toward physically modeling the speckle structures from wavefront-sensing telemetry rather than relying solely on stellar-light suppression in the science images, as explored for instance through differentiable rendering approaches (Feng et al., 2025). This emerging direction aims to incorporate information from the wavefront aberrations responsible for speckle formation, allowing a more accurate reconstruction of the stellar PSF and an improved separation between speckles and planetary signals.




Figure 1.11: The four pillars in HCI for the imaging of the four giant planets around the HR8799 star. Images from Faustine Cantalloube

All these PSF-subtraction methods produce a final residual frame whose pixel values remain expressed in flux units, such that any potential planetary signal is directly related to its flux. Candidate detections may initially be identified through visual inspection of the residual images. However, robust detection requires the use of statistical tests to characterize the noise properties and evaluate the significance of a signal relative to the underlying noise distribution. The topic of detection in HCI is discussed in the next section (Section 1.4).

The combination of observing strategies and PSF subtraction techniques, together with the enabling technologies in HCI has enabled contrasts on the order of 10^{-6} , sufficient to directly image young giant exoplanets in the infrared (see Figure 1.11). This has led to some of the most significant discoveries of directly imaged planetary systems, including HR 8799 bcde (Marois et al., 2008b, 2010), β Pictoris b (Lagrange et al., 2009b), and PDS 70 b (Keppler et al., 2018), which have significantly advanced our understanding of giant exoplanets.

1.4 Detection techniques in HCI

The question of detection is central across many scientific fields, including astronomy. In HCI, this question is particularly important, as it determines whether a signal corresponds to an astrophysical source, representing a true unique discovery and providing valuable insights into the field, or arises from *another* residual noise.

The most widely used metric in astronomy is the signal-to-noise ratio, introduced in Section 1.4.1, which generally represents the final step in ADI processing following the PSF subtraction techniques described in Section 1.3.4. Alternative approaches for detection are presented in Sections 1.4.2 and 1.4.3. The performance of these methods is discussed in Section 1.4.4. Finally, the scientific implications of detections and non-detections in HCI are presented in Section 1.4.5.

1.4.1 Signal-to-noise ratio (S/N)

The signal-to-noise ratio (S/N) is a widely used metric for assessing detections in astronomy. It quantifies how strongly a measured signal deviates from the noise statistics, typically by normalizing the signal intensity to the standard deviation of the noise.

The topic of detection can be framed as a statistical hypothesis test between two competing hypotheses: the null hypothesis H_0 , in which the observation is explained by noise alone, and the alternative hypothesis H_1 , in which the observation is better explained by a signal-plus-noise model. Assuming that the noise follows a Gaussian distribution and that the noise and planetary signal measurements are constructed from the mean pixel values measured within apertures, the noise sample is defined as $X = \{\bar{X}_1, \dots, \bar{X}_n\}$, with mean μ_X and standard deviation σ_X , while the planetary signal sample is defined as $Y = \{\bar{Y}_1\}$, with mean μ_Y . The signal-to-noise ratio is then defined as

$$S/N = \frac{\mu_Y - \mu_X}{\sigma_X} \quad (1.1)$$

As illustrated in the middle panels of Fig. 1.12, the confidence in a detection increases as the measured signal deviates further from the noise distribution under H_0 . In other words, a signal located farther in the tail of the noise distribution is more likely to be explained by H_1 than by H_0 .

In HCI, the noise distribution is not straightforward. Analyses by Soummer & Aime (2004); Fitzgerald & Graham (2006) showed that the intensity distribution in AO-corrected images is described by a modified Rician distribution. This distribution arises from the combination of a coherent (static) component, associated with the diffraction-limited PSF (e.g., Airy pattern), and an incoherent (random) component due to atmospheric speckles. After coronagraphy, AO-corrected images still exhibit static features (as shown in Fig. 1.8), as well as quasi-static and atmospheric speckles. The static and quasi-static contributions are also well described by a modified Rician distribution (Soummer et al.,

2007). In contrast, the atmospheric speckle component, when averaged over a large number of independent realizations, tends toward a Gaussian distribution as predicted by the central limit theorem. Observing strategies and post-processing techniques aim to suppress and decorrelate quasi-static speckles. As a result, the residual noise in the final processed frames is often assumed to be approximately Gaussian ("whitened") (Marois et al., 2008a), an assumption that has been widely adopted in the literature (Mawet et al., 2014).

Regarding the construction of the sample used to estimate the noise distribution, HCI typically considers resolution elements of size λ/D rather than individual pixels, in order to reduce spatial correlations and improve the independence of the samples. In addition, the noise statistics are estimated locally in annuli to account for their radial dependence across the image. Mawet et al. (2014) define the samples using apertures of diameter λ/D placed within an annulus, as illustrated in the left panel of Fig. 1.12, where each sample corresponds to the mean intensity of the pixels within the aperture. More recently, Bonse et al. (2023) adopted a similar aperture-based approach but considered only the central pixel of each aperture as the representative value, further reducing correlations between samples. At small angular separations, the number of independent resolution elements becomes limited (e.g., as few as 6 at $1 \lambda/D$), placing the problem in the regime of small-sample statistics. To account for this, Mawet et al. (2014) propose the use of the Student's t -distribution, with the test statistic defined as

$$T = \frac{\mu_Y - \mu_X}{\sigma_X \sqrt{1 + 1/n}}. \quad (1.2)$$

This expression is similar to the signal-to-noise ratio defined in Eq. 1.1, with the additional factor $\sqrt{1 + 1/n}$ accounting for the limited number of independent samples. In the limit of large n , this term approaches unity, and the statistic converges to the standard S/N.

To quantify the risk of incorrectly rejecting H_0 in favor of H_1 in the absence of a planetary signal, the false positive fraction (FPF) is defined as

$$\text{FPF} = \int_{T_{\text{obs}}}^{\infty} p(T = t | H_0) dt, \quad (1.3)$$

where T_{obs} is the value of T for the observation and $p(T = t | H_0)$ is the probability density function of the test statistic under the null hypothesis. For example, the case $T_{\text{obs}} = 2.28$ (as illustrated in the right panel of Fig. 1.12) corresponds to $\text{FPF} = 0.0243$, meaning that 2.43% of noise realizations would be incorrectly classified as a detection.




Figure 1.12: The three main steps of the hypothesis testing framework introduced by (Mawet et al., 2014). Figure taken from (Bonse et al., 2023)

A commonly adopted detection threshold in astronomy is 5σ , which corresponds to $\text{FPF} = 2.87 \times 10^{-7}$, i.e., an extremely low probability of false positives. The strategy presented here remains the most widely used detection metric after post-processing, mainly due to its simplicity and accessibility. However, several studies have highlighted limitations of the S/N framework, as the residual noise in the final processed frames is observed to deviate from the Gaussian assumption. Pairet et al. (2019) showed that the tail of the residual noise distribution is better described by a Laplacian distribution than by a Gaussian. More recent studies further confirmed the non-Gaussianity of the residual noise (Bonse et al., 2023; Cantero et al., 2023). In particular, Cantero et al. (2023) identified different statistical regimes within the final residual frames, with some regions being approximately Gaussian while others exhibit clear non-Gaussian behavior.

Today, it is widely acknowledged that the residual noise in HCI is not strictly Gaussian, motivating the exploration of alternative approaches. Methods aiming at modeling the planetary signal are discussed in Section 1.4.2, while alternative detection frameworks that do not rely directly on the signal-to-noise ratio are presented in Section 1.4.3.

1.4.2 Forward models

Forward models in HCI provide an alternative framework for the detection and characterization of planetary signals while better preserving their morphology and flux. Unlike classical PSF subtraction techniques (Section 1.3.4), which primarily focus on suppressing the stellar light and speckle noise before identifying planetary signals in the residuals, forward-modeling techniques explicitly model the expected planetary signal and search for its signature in either preprocessed or post-processed frames, i.e., with or without prior stellar light subtraction.

ANDROMEDA

The ANgular DiFFeRential OptiMal Exoplanet Detection Algorithm (ANDROMEDA, Cantalloube et al., 2015) is an ADI-based forward-modeling algorithm that searches for the expected signature of a planetary companion using maximum likelihood estimation. ANDROMEDA relies on pairwise subtraction, where the reference frame associated with each science frame is selected to be close in time while satisfying an exclusion criterion. This ensures that the two frames remain sufficiently correlated, while allowing the planetary signal to move enough between the frames. If present, the planetary signal appears in the differential image as a characteristic positive-negative dipole resulting from the subtraction of the two rotated signals. The algorithm then searches for this expected signature using maximum likelihood estimation. ANDROMEDA assumes a modified Rician distribution at large angular separations and a Student's t -distribution at small angular separations.

FMMF

The Forward Model Matched Filter algorithm (FMMF, Ruffio et al., 2017) is a KLIP-based forward-modeling technique that searches for the distorted planetary signal after KLIP PSF-subtraction rather than the original off-axis PSF. In particular, the algorithm accounts for the effects of *self-subtraction* introduced by the PSF subtraction process, including negative lobes and signal attenuation, using the analytical forward model developed by Pueyo (2016). FMMF then applies a matched-filter framework based on maximum likelihood estimation, similar to the approach used in Cantalloube et al. (2015), to search for the expected processed planetary signature in the residual images. FMMF assumes a Gaussian distribution for the S/N statistic.

PACO

The PATch COvariances algorithm (PACO, Flasseur et al., 2018) is an ADI-based forward-modeling technique that searches for exoplanet signals by modeling the spatial correlations of the noise within local patches of the image sequence. PACO models the noise using a multivariate Gaussian distribution and searches for a planetary signal represented by an off-axis PSF following the expected circular trajectory induced by field rotation throughout the ADI sequence. Unlike the previous techniques mentioned, PACO does not explicitly suppress the stellar light. Instead, it relies on a generalized likelihood ratio test to discriminate between the hypotheses of noise alone and signal-

plus-noise. Several extensions of PACO have since been developed, including Flasseur et al. (2020a), which improves robustness against corrupted frames, and Flasseur et al. (2020b), which extends the framework to combined ADI and SDI observations (PACO-ASDI). PACO assumes a Gaussian distribution for the S/N statistic.

All of these methods produce both a flux estimate and an associated S/N ratio map.

1.4.3 Other detection techniques

The methods presented in this section explore alternative approaches for assessing detections without relying on the S/N method. These techniques generally do not introduce new models for either the noise or the planetary signal, but instead rely on one of the PSF-subtraction methods described in Section 1.3.4, most commonly PCA, while proposing alternative strategies for signal vetting and detection assessment.

STIM map

The Standardized Trajectory Intensity Mean map (STIM; Pairet et al., 2019) exploits the temporal information contained in the residual ADI sequence (step D in Figure 1.10) rather than the spatial statistics at a given angular separation, as commonly done in classical S/N ratio maps (see Figure 1.12). For each pixel, the method follows the corresponding ADI trajectory and computes a detection statistic defined as the ratio between the temporal mean of the residual intensities and their temporal standard deviation. The method is motivated by the theoretical result that, under the null hypothesis, the probability of obtaining large STIM values decreases exponentially with the number of frames. The detection threshold is estimated empirically by processing the sequence with *opposite parallactic angles*, which suppresses any real astrophysical signal while preserving the statistical properties of the residual speckle noise.

RSM

The Regime Switching Model (RSM, Dahlqvist et al., 2020, 2021b,a) uses the spatial and temporal information contained in the residual ADI sequence (step D in Fig. 1.10) to classify the data into either a noise regime or a planetary regime through a regime-switching framework inspired by econometrics (Hamilton, 1988). The noise distribution

is modeled using a weighted combination of Gaussian and Laplacian distributions, while the planetary model is represented either by an off-axis PSF or by a forward-modeled PSF. The likelihood associated with each regime is computed for every pixel, from which the probability of belonging to either the planetary or noise regime is derived.

The final output of the algorithm is a probability detection map with values ranging between 0 and 1. As the main detection framework used throughout this thesis, the RSM algorithm will be described in greater detail in Chapter 3. In addition, the definition of detection thresholds for RSM constitutes an important contribution of this thesis and is explored in Chapter 4.

SODINN/NA-SODINN

Supervised exoplanet detection via Direct Imaging with deep Neural Networks (SODINN, Gomez Gonzalez et al., 2018) is a supervised deep-learning algorithm that uses convolutional neural networks to classify image patches into either planetary signal class or noise class. The method operates on 3D patches extracted from the final ADI output (step E in Fig. 1.10), with a spatial size of approximately twice the FWHM. The third dimension corresponds to the evolution of the residual intensity in the patch for different numbers of principal components used in the PCA. The training set is constructed using injected off-axis PSF as positive examples and real residual noise patches from the processed observations as negative examples. After training, the network assigns a confidence score corresponding to the probability that a given patch belongs to the planetary-signal class, producing a detection map with values between 0 and 1.

Building upon this framework, Cantero et al. (2023) introduced Noise-Adaptive SODINN (NA-SODINN), designed to account for the different statistical properties of the residual noise across the final ADI frame, being non Gaussian close to the star and approximately Gaussian at larger angular separations. NA-SODINN performs the full training and classification procedure independently for each noise regime in the image. A detection threshold based on F1-score maximization is adopted for this method in the forthcoming work of Cantero et al. (2026), following our work described in Chapter 4.

1.4.4 Performance of detection techniques in HCI

The previous sections presented a number of post-processing and detection techniques commonly used in HCI. However, the number of available methods has grown significantly in recent years, and the present review does not aim to provide an exhaustive overview of all existing approaches. A broader review of current detection and post-processing techniques is provided in the *talk* by Dr. Faustine Cantalloube during the Sagan Workshop.

To compare the performance of these techniques on a fair basis, Cantalloube et al. (2021) organized the Exoplanet Imaging Data Challenge (EIDC), in which several PSF-subtraction techniques (PCA, LOCI, PACO and FMMF) combined with various detection methods (RSM, SODINN, STIM) participated. The goal of the challenge was to provide a standardized framework for comparison by testing all algorithms on the same datasets, acquired with various instruments, under various observing strategies and observing conditions. The algorithms were also evaluated using the same metrics, including the F1-score, the area under the curve (AUC) of the true positive rate (TPR), and the AUC of the false discovery rate (FDR).

Such a comparison framework is particularly important, as different algorithms are often tested on different datasets and assessed using different evaluation metrics, making direct comparisons difficult. The results of the first phase of the challenge, focused on detection performance, are presented in Fig. 1.13 for the ADI-based algorithms. The challenge highlighted particularly strong performances from methods such as RSM, PACO, and FMMF, with RSM showing the smallest AUC for FDR. These results have not yet been updated to include the performance of more recent techniques, such as NA-SODINN, which demonstrated improved performance compared to SODINN, and 4S, which was able to recover AF Lep b in archival VLT/NACO data more than a decade before its official discovery (Bonse et al., 2025).




Figure 1.13: Ranking of the different families of ADI-based algorithms for the EIDC challenge. Left column: Ranking based on the F1-score. Middle column: Ranking based on the AUC of the TPR. Right column: Ranking based on the AUC of the FDR. From top to bottom, the 5 different families of algorithms: PSF-subtraction techniques combined with S/N maps (red), PSF-subtraction techniques combined with advanced detection maps (orange), inverse problem approaches (blue) and supervised machine learning detection algorithms (green). The light, medium and dark colors correspond to the three VLT/SPHERE-IRDIS, Keck/NIRC2, and LBT/LMIRCam data sets respectively. Image taken from Cantalloube et al. (2021).

1.4.5 Observables in HCI

Producing residual frames in which the stellar light is sufficiently attenuated and the planetary signal can be reliably assessed provides highly valuable information for HCI analysis, regardless of whether a detection is achieved. In the case of a detection, dedicated characterization tools can provide accurate estimates of both the flux and the astrometry of the companion, while non detections define the detection limits achievable by a given instrument and post-processing method, commonly expressed through contrast curves.

Astrometry and Photometry

Once a signal is confidently detected above a given detection threshold, this does not automatically imply a planetary origin. The detected signal may still correspond to a false positive produced by residual speckle noise or limitations of the detection technique. In the case of a genuine astrophysical source, exoplanets are also not the only objects capable of producing detectable signals in the field of view. Background stars, particularly for targets located close to the Galactic plane, can exhibit faint apparent magnitudes that become detectable in HCI observations and may therefore mimic planetary companions.

Differentiating between a planetary companion and a background star is commonly achieved through monitoring the astrometric evolution of the candidate over at least two observational epochs. A bound planetary companion follows the apparent motion of the host star in terms of parallax and proper motion (possibly with some measurable orbital motion), while a background object does not. Tools such as `backtracks` (Balmer et al., 2025b) are commonly used to model the expected trajectory of background stars relative to the target star. Photometric information also provides important constraints on the nature of the detected signal. Planetary companions often exhibit characteristic colors linked to molecular absorption features in their atmospheres when observed at multiple wavelengths, while background stars generally follow different spectral energy distributions. False positives can additionally be identified through follow-up observations. A candidate that is not recovered at subsequent epochs is unlikely to be astrophysical in origin. Additional diagnostics can also be obtained from the chromatic behavior of the signal: residual speckles scale with wavelength due to diffraction, whereas genuine astrophysical sources preserve their physical position independently of wavelength.

As accurate astrometric and photometric measurements are central to determining the nature of a detected signal, several tools have been developed in HCI to estimate the flux and position of candidate companions while accounting for the biases introduced by post-processing techniques, such as positional offsets and flux attenuation. One widely used approach is the NEGative Fake Companion technique (NEGFC, Lagrange et al., 2010; Wertz et al., 2017), which relies on injecting a negative PSF template at the location of the candidate companion. The flux and position are then iteratively adjusted to identify the set of parameters that provides the best cancellation of the signal in the residual image, typically quantified through the minimization of a residual metric such as the sum of absolute residual intensities.

Once a planetary signal is confirmed, astrometric measurements obtained over multiple observational epochs enable the estimation of the orbital elements of the companion. Tools such as *orbitize!* (Blunt et al., 2020, 2024) explore the orbital parameter space using the measured angular separations and position angles of the candidate companion at different epochs in order to constrain its orbital architecture. These measurements provide valuable information about the dynamical properties of planetary systems. For instance, astrometric measurements of the HR 8799 system revealed the possible long-term dynamical stability associated with three of the four planets being nearly coplanar (Goździewski & Migaszewski, 2020). Similarly, the inferred periastron and apastron distances of HD 95086 b suggested that the companion did not possess a sufficiently large eccentricity to explain the truncation of the outer edge of the debris disk (Rameau et al., 2016). Photometric measurements also provide key physical constraints on the detected companion. When combined with accurate stellar age estimates and evolutionary models, the measured flux can be converted into an estimate of the planetary mass.

Contrast curves

In the case of non-detection, HCI observations still provide valuable information about the detection limits achievable around a given target using a specific instrument and post-processing pipeline. These sensitivity limits are commonly expressed through contrast curves, defined as the minimum detectable flux ratio between a companion and the host star as a function of angular separation.

In classical post-processing techniques, the sensitivity limit is directly linked to the adopted detection threshold, for instance a 5σ threshold under the assumption of gaussian noise. Because post-processing techniques attenuate the planetary signal through effects such as self-subtraction, the recovered flux is lower than the true injected flux. To account for this attenuation, a throughput correction is introduced, typically estimated by injecting synthetic companions into the data and measuring the ratio between the recovered and injected fluxes. The contrast at a given separation is then commonly estimated as

$$\text{Contrast} = \frac{5\sigma}{F_{\star}} \times \frac{1}{\text{Throughput}}, \quad (1.4)$$

where $F_{\star}$ is the stellar flux. Contrast curves generally improve with increasing angular separation. At small separations, the residual speckle noise dominates and strongly limits the achievable contrast, while at larger separations the noise level decreases and the contrast curve becomes flatter.

In classical post-processing frameworks, the flux limit associated with a given detection threshold corresponds to placing the mean of the planetary-signal distribution at that threshold. Under this definition, approximately 50% of the companions injected at the sensitivity limit would be recovered above the threshold, while the remaining 50% would remain undetected. In this sense, the classical 5σ contrast curve corresponds to a 50% completeness level. Other completeness definitions have also been adopted in the literature, such as the 95% completeness criterion introduced by Wahhaj et al. (2013a), which provides a more conservative and robust estimate of the detection capability.

To avoid the dependance of contrast curves on Gaussian assumptions, alternative approaches have been proposed to avoid assuming a specific underlying distribution for the residual noise. For instance, Jensen-Clem et al. (2018); Dahlqvist et al. (2021b) introduced methods based directly on controlling the number of true positives and false positives accepted in the frame. The approach of Jensen-Clem et al. (2018) relies on fixing an acceptable number of false positives per separation and iteratively injecting companions until this limit is reached. In addition, their work represents completeness as a two-dimensional map spanning values between 0 and 1. Dahlqvist et al. (2021b) implemented a practical version of this concept within the RSM framework through an iterative injection and recovery procedure converging toward the requested completeness level.

The importance of these detection limits is that, when combined with evolutionary models, they can be converted into mass detection limits, allowing one to constrain the presence of planets of a given mass around a star. This is particularly important for systems such as ϵ Eridani b, which has been observed with multiple HCI instruments over the years. The achieved contrast limits have progressively constrained the physical properties compatible with the observations, such as excluding the detectability of the companion under specific age assumptions (Mawet et al., 2019) or at certain effective temperatures (Llop-Sayson et al., 2025). More generally, mass detection limits derived across large surveys provide key constraints on the occurrence rate and population statistics of exoplanets. This topic is discussed further in Section 1.5.

1.5 Direct imaging surveys and their scientific contributions

The major HCI surveys were motivated by several scientific goals. By observing large samples of stars, these surveys aimed to maximize the probability of detecting giant planets and to perform demographic studies of planetary populations at wide orbital separations, a region of the planet mass vs semi-major axis parameter space that is largely inaccessible to indirect detection techniques. Beyond the detections themselves, direct imaging provides constraints on the masses, orbital architectures, and atmospheric characteristics of giant planets. Combined with the demographic information derived from both detections and non-detections, these observations provide important constraints on the formation and evolution of giant planets at large separations.

Target selection was tailored to these objectives. Members of young moving groups were favored because they are likely to host young planets that are sufficiently luminous for direct imaging, while their ages can be constrained relatively well, enabling more reliable planetary mass estimates through evolutionary models. Surveys also targeted young star-forming regions to search for protoplanets, stars hosting debris disks to investigate planet–disk interactions, intermediate- and high-mass stars because of their higher giant-planet occurrence rates, and stars exhibiting astrometric accelerations, which offer a higher probability of companion detection than blind surveys.

These large observing programs have contributed to the discovery of a substantial fraction of the directly imaged exoplanets known today. Modern surveys typically encompass several hundred targets and achieve deep contrast limits, enabling increasingly precise constraints on the distribution of giant planets at wide orbital separations. General reviews of direct-imaging surveys are provided by Bowler (2016); Currie et al. (2023a).

HCI surveys can be broadly divided into two generations, reflecting the technological developments that enabled them. In this section, we distinguish between (i) first-generation surveys conducted on large ground-based telescopes equipped with adaptive optics, with or without coronagraphs, and (ii) second-generation surveys employing extreme adaptive optics systems. Particular emphasis is placed on the SHINE survey, which constitutes the main focus of this thesis and is therefore discussed in greater detail.

1.5.1 First generation instruments surveys

The first direct imaging surveys began with the space-based HST/NICMOS survey of young A–M stars (Lowrance et al., 2005), which led to the detection of only a few brown dwarf companions before HCI efforts expanded to large ground-based telescopes. The first generation of surveys benefited from major advances in HCI instrumentation and data analysis, including the use of 8m class telescopes, adaptive optics, differential observing strategies, PSF-subtraction techniques, and, in some surveys, coronagraphy. These programs were specifically designed for the direct detection of giant exoplanets. Typical survey samples contained 50–300 stars and achieved contrast limits of approximately $\Delta\text{mag} \sim 10\text{--}11$ at $0.5''$. Examples of these surveys include:

NaCo survey of young nearby austral stars

This program (Chauvin et al., 2010) utilized NaCo (Rousset et al., 2003) at the VLT between 2002 and 2007, together with Lyot coronagraphy in the H and K_S bands and PSF-subtraction techniques. The target sample consisted of 88 nearby young stars, primarily G- to M-type dwarfs younger than 100 Myr and located within 100 pc. The most important discovery from this instrument was the detection of 2M1207 b, the first directly imaged planetary-mass companion, with an estimated mass of $5 \pm 2M_J$, orbiting the $\sim 25M_J$ brown dwarf 2M1207-3932 in the ~ 10 Myr TW Hydrae Association (Chauvin et al., 2005).

Gemini NICI Planet-finding Campaign

This campaign (Liu et al., 2010; Biller et al., 2013; Nielsen et al., 2013; Wahhaj et al., 2013b) used the Near-Infrared Coronagraphic Imager (NICI) adaptive-optics instrument on the 8.1 m Gemini South telescope. Observations combined a Lyot coronagraph with single-channel H-band imaging using ADI, as well as simultaneous dual-channel imaging in the CH_4S ($1.578\mu\text{m}$) and CH_4L ($1.652\mu\text{m}$) filters using both ADI and SDI. The survey targeted 230 young stars spanning a wide range of spectral types.

This survey discovered brown-dwarf companions, including the highly eccentric companion PZ Tel B in the β Pictoris moving group (Biller et al., 2010). Although no planetary-mass companion was detected, the survey provided important constraints on the occurrence rate of giant planets at wide separations. Using a sample of 80 members of nearby young moving groups, Biller et al. (2013) constrained the frequency of $1\text{--}20M_J$ planets between 10 and 150 au to be below 6–18% at the 95.4% confidence level. For a sample

of 70 B- and A-type stars, Nielsen et al. (2013) found that the frequency of planets more massive than $4M_J$ between 59 and 460 au is below 20% at 95% confidence. Similarly, from observations of 57 debris-disk stars, Wahhaj et al. (2013a) showed that fewer than 13% host planets more massive than $5M_J$ beyond 80 au at the 95% confidence level.

International Deep Planet Search

The International Deep Planet Survey (IDPS) (Vigan et al., 2012; Galicher et al., 2016) was a large imaging survey carried out with NaCo at the VLT, NIRC2 at Keck, NICI at Gemini South, and NIRI at Gemini North, targeting nearly 300 young A–M stars over a period of 14 years. The target stars were typically younger than 300 Myr and located at distances of approximately 10–80 pc.

One of the main results of this collaboration was the breakthrough discovery of the HR 8799 planetary system (Marois et al., 2008b, 2010), the first directly imaged multi-planet system, which became a cornerstone of HCI studies (discussed further in Section 1.5.3). The statistical analysis of the full survey sample by Galicher et al. (2016) concluded that $1.05^{+2.80}_{-0.70}\%$ of stars host at least one giant planet with masses between 0.5 and $14M_J$ and orbital separations between 20 and 300 au.

NaCo Large Program

The NaCo Large Program (NaCo-LP) survey (Desidera et al., 2015; Chauvin et al., 2015; Reggiani et al., 2016; Vigan et al., 2017) used the NaCo instrument at the VLT, combined with a Lyot coronagraph and ADI observations, to target 86 young, bright, primarily FGK stars in the H band between 2009 and 2013.

The statistical analysis presented by Chauvin et al. (2015) placed upper limits on the occurrence rate of giant planets over a wide range of orbital separations. In particular, fewer than $\sim 15\%$ of Sun-like stars were found to host planets more massive than $5M_J$ between 100 and 500 au, while fewer than $\sim 10\%$ host planets more massive than $10M_J$ between 50 and 500 au. Using the complete survey sample of 199 FGK stars, Vigan et al. (2017) derived a substellar companion frequency of 0.75–5.70% for companions with masses between 0.5 and $75M_J$ and orbital separations between 20 and 300 au.

These surveys reached unprecedented sensitivities, enabling access to the planetary-mass regime and leading to the first landmark discoveries in HCI. The demographic analyses derived from these surveys consistently indicated that giant planets at wide orbital separations are intrinsically rare. However, their sensitivity remained insufficient to place strong constraints on the dominant planet formation pathways. This motivated the development of second-generation instruments and larger survey programs with substantially improved sensitivities and sample sizes.

1.5.2 Second generation instruments surveys

This generation of surveys benefited from the development of extreme adaptive optics (XAO) systems, employing high-order wavefront correction with large numbers of actuators and deformable mirrors. Combined with optimized coronagraphs and post-processing techniques, these instruments routinely achieved contrast limits of approximately $\Delta\text{mag} \sim 12\text{--}13$ at $0''.5$. The largest surveys of this generation are the Gemini Planet Imager Exoplanet Survey and the SPHERE High-Contrast Imaging Survey for Exoplanets, which targeted samples of approximately 400–600 stars.

Gemini Planet Imager Exoplanet Survey

The Gemini Planet Imager Exoplanet Survey (GPIS; Nielsen et al., 2019b) is a 600-star survey conducted with the Gemini Planet Imager (Macintosh et al., 2006) at Gemini South. The instrument combines extreme adaptive optics, an apodized Lyot coronagraph, and advanced post-processing techniques to image giant planets at small inner working angles. The survey targeted young nearby FGKM stars within 100 pc, with the goals of measuring the occurrence rate of giant planets, investigating its dependence on stellar properties, and constraining the atmospheres and orbital properties of substellar companions.

This survey led to several discoveries, among which 51 Eridani b (Macintosh et al., 2015) is one of the most notable. This young (~ 20 Myr) giant planet, located at a separation of 13 au, has an estimated mass of only $\sim 2M_J$ assuming hot-start evolutionary models, probing a region of parameter space that was previously inaccessible to direct imaging surveys. On the demographic side, Nielsen et al. (2019b) probed giant planets down to separations of approximately 10 au and measured an occurrence rate of $9^{+5}_{-4}\%$ for planets with masses between 5 and $13M_J$ and separations between 10 and 100 au. One important result from these measurements was the identification of a strong correlation

between the occurrence rate of giant planets at wide separations and stellar mass. In contrast, the occurrence rate of brown-dwarf companions does not exhibit the same trend, suggesting different formation pathways for the two populations. The demographic implications of this study are discussed in more detail in Section 1.5.3.

SpHere INfrared survey for Exoplanets

The SpHere INfrared survey for Exoplanets (SHINE; Chauvin et al., 2017a) used the Spectro-Polarimetric High-contrast Exoplanet REsearch (SPHERE; Beuzit et al., 2019) instrument installed at the Very Large Telescope (VLT) in Chile. SPHERE combines the SAXO extreme adaptive-optics system (Fusco et al., 2006), capable of delivering Strehl ratios of up to $\sim 90\%$ in the H band, with apodized Lyot coronagraphs (Guerri et al., 2011) and differential-imaging techniques such as ADI and SDI. Observations in SHINE were acquired using both near-infrared subsystems of SPHERE: the Integral Field Spectrograph (IFS; Claudi et al., 2008), operating in the YJ and YJH spectral modes with a field of view of $1.77'' \times 1.77''$, and the InfraRed Dual-band Imager and Spectrograph (IRDIS; Dohlen et al., 2008), operating in dual-band imaging mode in the H and K bands with a field of view of $11'' \times 11''$.

The full SHINE survey targeted nearly 400 nearby young stars. The sample has a median age of 45Myr (90% between 11 and 450Myr), a median stellar mass of $1.15M_{\odot}$ (90% between 0.57 and $2.37M_{\odot}$), and a median distance of 48pc (90% between 11 and 137pc). A detailed description of the sample properties is presented in Desidera et al. (2021). The main scientific objectives of the SHINE survey are (i) to identify new planetary-mass and brown-dwarf companions, (ii) to study the architecture of planetary systems at wide orbital separations, (iii) to investigate the connection between the presence of planets and circumstellar disks, (iv) to determine the occurrence rate of giant planets with semi-major axes beyond 5 au, and (v) to explore the dependence of planetary companion frequency and properties on stellar mass over the range $0.5\text{--}3M_{\odot}$.

Preliminary statistical results of the SHINE survey were presented for the F150 sample, a subset of 150 stars representative of the full SHINE dataset. Desidera et al. (2021) describes the sample properties, Langlois et al. (2021) presents the detection performance obtained using classical post-processing techniques such as PCA and TLOCI, and Vigan et al. (2021) reports occurrence rate constraints derived from these observations. The SHINE survey led to the discovery of the planetary-mass companion HIP 65426 b (Chauvin et al., 2017b) and the brown dwarf companion HIP 64892 B (Cheetham et al., 2018). SPHERE observations also contributed to the detection of PDS 70 b (Keppler

et al., 2018; Müller et al., 2018), although this object was identified outside the SHINE survey itself, and more recently to the detection of HD 143811 b in combination with GPI (Squicciarini, V. et al., 2025). In addition, the SHINE dataset provided improved characterization and refined orbital constraints for several previously known planetary systems including 51 Eri b (Samland et al., 2017; Brown-Sevilla et al., 2023), HR 8799 (Apai et al., 2016; Zurlo et al., 2016), and β Pictoris (Lagrange et al., 2020). Several circumstellar disks were also detected in SHINE observations despite not being the primary science goal of the survey (e.g. Lagrange et al., 2016; Feldt et al., 2017; Perrot et al., 2019). The demographic analysis of Vigan et al. (2021) confirmed the correlation between giant-planet occurrence rate and host-star mass, with occurrence rates of $23.0^{+13.5}_{-9.7}\%$, $5.8^{+4.7}_{-2.8}\%$, and $12.6^{+12.9}_{-7.1}\%$ for BA-, FGK-, and M-type stars, respectively. This study also provided important constraints on planet formation pathways, which are discussed in Section 1.5.3.

Since the present thesis focuses on the detection performance of SHINE observations, we briefly summarize the results of Langlois et al. (2021), who reported typical 5σ contrast limits of 13 mag at 200 mas and 14.2 mag at 800 mas with IFS, and 11.8 mag at 200 mas, 13.1 mag at 800 mas, and 15.8 mag at $3''$ with IRDIS in the H band. These observations led to the detection of 1483 candidate signals with IRDIS (more than 95% obtained with the H2–H3 filter pair). Among these signals, seventeen were confirmed as substellar companions, 1176 were identified as background contaminants through follow-up observations, and 307 remained single-epoch candidates without confirmation. Because most detections were obtained using the H2–H3 filter pair (more than 95 %), which probes the methane absorption feature in the H band, the H2–H3 color was used to identify candidate companions as explained in (Bonnefoy et al., 2018). Methane-bearing companions typically show negative H2–H3 colors, while late-L objects may exhibit slightly positive values. In contrast, most background stars cluster around $H2-H3 \approx 0$. This is well illustrated in Fig. 1.14. Single-epoch candidates were therefore placed in the color-magnitude diagram (CMD) in order to evaluate their likelihood of being substellar companions. Sources located within the gray exclusion region shown in Fig. 1.14 were identified as likely background contaminants.

Several advanced post-processing techniques were applied within the SHINE survey to further improve detection performance, particularly at small angular separations. The ASDI-based PACO algorithm was applied to a subset of 24 datasets from the F150 sample (Chomez et al., 2023), demonstrating contrast improvements of up to a factor of five at small separations. The algorithm was subsequently applied to the full SHINE




Figure 1.14: Color-magnitude diagram showing the absolute H2 magnitude as a function of the H2–H3 color for SHINE candidate detections reported by Langlois et al. (2021), together with the photometry of confirmed companions and comparison sequences of M-, L-, T-, and Y-type objects. The gray region indicates the area of the diagram statistically dominated by background contaminants. Figure taken from Langlois et al. (2021).

survey of nearly 400 stars (Chomez et al., 2025), constituting the fourth paper of the SHINE survey. More than 3500 astrophysical signals were detected, among which eight were identified as promising companion candidates. That analysis also made use of SnapSHINE observations, consisting of short follow-up sequences obtained to provide additional epochs for candidate characterization. The work presented in this thesis similarly applies the advanced RSM algorithm to the F150 sample, with the results presented in Chapter 5. More recently, the NA-SODINN algorithm was also applied to the F150 sample, recovering several potential companion candidates (Cantero et al., 2026).




Figure 1.15: Non-exhaustive selection of substellar companions discovered through HCI. References are given in the same order as the images shown: Bailey et al. (2014); Marois et al. (2010); Kuzuhara et al. (2013); Currie et al. (2023b); Kalas et al. (2008); Haffert et al. (2019); Lagrange et al. (2009b); Carson et al. (2013); Currie et al. (2020); Macintosh et al. (2015); Chauvin et al. (2017b); Kraus et al. (2014); Rameau et al. (2013).

1.5.3 What did we learn?

Emblematic systems

HCI surveys have uncovered several benchmark planetary systems that have shaped our understanding of young giant planets and their environments. These discoveries have provided direct insights into planetary architectures, atmospheric properties, planet–disk interactions, and the process of planet formation itself, often probing regimes inaccessible to other detection techniques. Figure 1.15 presents some of these HCI detections. Below, we briefly review a few systems that have had a particularly strong impact on the field.

- **β Pic**

The A6V star β Pictoris is the eponymous member of the β Pictoris moving group, with an estimated age of ~ 20 Myr (Mamajek & Bell, 2014). It hosts an edge-on debris disk, the first debris disk resolved around a main-sequence star (Smith & Terrile, 1984). Subsequent observations revealed a warp in the inner disk, which was interpreted as the signature of gravitational perturbations from a giant planet (Mouillet et al., 1997). The later discovery of β Pic b (Lagrange et al., 2009b, 2010), orbiting at a semi-major axis of approximately 9 au, provided strong support for this interpretation, making the β Pictoris system one of the clearest observational demonstrations of planet–disk interactions and a benchmark for studies of how giant planets sculpt debris disks through long-term gravitational perturbations (see Hughes et al., 2018, for a review).

Radial velocity observations with HARPS later revealed a second giant planet, β Pic c (Lagrange et al., 2019), which was subsequently recovered through interferometric observations with VLTI/GRAVITY (Nowak et al., 2020). Recent observations with JWST/NIRSpec (Gibbs et al., 2026) and, independently, with JWST/NIRCam, VLT/SPHERE, and VLT/ERIS (Sutcliffe et al., 2026) led to the discovery of a third giant planet, β Pic d, at a projected separation of ~ 26 au, whose orbit is consistent with sculpting the inner edge of the debris disk. The system is one of the few directly imaged planetary systems for which direct imaging, radial velocity, absolute astrometry, and interferometric measurements can be combined, enabling precise orbital constraints and dynamical mass determinations for both planets (Lagrange et al., 2020; Brandt et al., 2021).

- **HR 8799**

HR 8799 is a nearby (~ 40 pc) young A-type star with an estimated age of ~ 40 Myr (Baines et al., 2012). HR 8799 was the first directly imaged multi-planet system, with the planets b, c, and d discovered in 2008 (Marois et al., 2008b) at projected separations of approximately 68, 38, and 24 au, respectively, followed by the discovery of planet e at a separation of ~ 15 au in 2010 (Marois et al., 2010). The presence of several massive ($> 5 M_J$) planets on wide orbits raised questions regarding their formation and long-term stability, which is now thought to be maintained by a near 1:2:4:8 resonant chain (Fabrycky & Murray-Clay, 2010; Konopacky et al., 2016). This resonant architecture suggests a migration-driven formation history (Goździewski & Migaszewski, 2014), although the exact formation pathway remains uncertain.

Atmospheric studies of the four planets revealed that young giant planets possess

cloudy, low-gravity atmospheres that differ significantly from those of field brown dwarfs (Barman et al., 2011; Currie et al., 2011; Skemer et al., 2012). As the first directly imaged multi-planet system, HR 8799 provides a unique laboratory for comparative exoplanetology, where atmospheric diversity can be studied among coeval giant planets orbiting the same host star. Such comparisons have highlighted the importance of cloud properties, chemistry, and formation history in shaping the observed spectra (e.g. Nasedkin et al., 2024). Consequently, the system has been observed with nearly all major HCI instruments (Lafrenière et al., 2009; Fukagawa et al., 2009; Metchev et al., 2009; Soummer et al., 2011; Currie et al., 2011; Greenbaum et al., 2018; Zurlo et al., 2022; Boccaletti et al., 2024), as well as interferometric and spectroscopic facilities (Wang et al., 2021; Nasedkin et al., 2024). These observations enabled detailed atmospheric abundance measurements, particularly of the C/O ratio, opening the possibility of linking atmospheric composition to giant-planet formation pathways (Konopacky et al., 2013; Xuan et al., 2026).

- **PDS 70**

PDS 70 is a young (~ 5 Myr), $0.87M_{\odot}$ T Tauri star hosting a transition disk with a large central cavity of radius ~ 70 au (Hashimoto et al., 2012). The morphology of the disk suggested the presence of a planetary companion within the cavity, either still in formation or already formed, a hypothesis later confirmed by the discovery of the giant planet PDS 70 b within the gap (Keppler et al., 2018), followed by the detection of PDS 70 c (Haffert et al., 2019). Located at orbital separations of approximately 22 and 34 au, respectively, the planets constitute the first confirmed direct detections of protoplanets embedded within their natal protoplanetary disk. Evidence for circumplanetary material around PDS 70 b and c (Christiaens et al., 2019; Keppler et al., 2019; Benisty et al., 2021), together with accretion signatures traced through $H\alpha$ emission (Wagner et al., 2018), indicates ongoing accretion and supports the interpretation that the planets are still forming. As a result, the PDS 70 system has become a benchmark laboratory for studying planet formation within evolving protoplanetary disks (for review, Andrews, 2020).

Recently, JWST has begun extending the parameter space of HCI toward lower-mass and colder planets. This includes the direct imaging of the previously known radial velocity and astrometric planet Eps Ind A b, a cold (~ 275 K) and mature (~ 3.5 Gyr) giant planet (Matthews et al., 2024), as well as the discovery of a sub-Jovian planet embedded in the debris disk of TWA 7 (Lagrange et al., 2025). Although JWST is not designed as a survey facility, several targeted surveys are exploiting its HCI capabilities to search for and characterize planetary companions around carefully selected targets.

Occurrence rate and demographics

Planet occurrence rates and their dependence on stellar and planetary properties, such as host-star mass, metallicity, planet mass, and orbital separation, provide important constraints on planet formation and evolution models. In direct imaging surveys, occurrence rates are derived by combining detections and non detections with the survey sensitivity. Contrast limits are converted into planet-mass detection limits using evolutionary models and stellar age estimates. Monte Carlo simulations of synthetic planet populations spanning a range of masses, semi-major axes, and orbital parameters are then used to evaluate the survey completeness, i.e. the probability of detecting a planet if it was present. The resulting completeness estimates allow the underlying planet population and occurrence rates to be inferred from the survey results (see review by Bowler & Nielsen, 2018).

Following the discovery of the first directly imaged exoplanets (Marois et al., 2008b; Lagrange et al., 2010) in first-generation HCI surveys, the statistical properties of wide-orbit giant planets began to be measured. These surveys were typically sensitive to giant planets with masses down to $\sim 5 M_J$ at projected separations of around 50 au, but their yields remained low, generally resulting in only 0–2 new planetary detections per survey. It was also noted that most directly imaged planets were found around intermediate- and high-mass stars, despite HCI observations generally achieving more favorable mass detection limits around lower-mass stars owing to their lower luminosities. This trend was later confirmed by second-generation surveys, which revealed a higher occurrence rate of giant planets around stars more massive than $\sim 1.5 M_\odot$.

Compared to RV demographic studies, which show that the occurrence rate of giant planets increases with both planet mass and orbital separation up to ~ 2.5 au (Cumming et al., 2008), Nielsen et al. (2019a) found an opposite trend at wider separations, with the occurrence rate of giant planets beyond 10 au decreasing toward higher planet masses and larger orbital separations. Combined with the RV results of Cumming et al. (2008),

these findings suggest that the giant-planet population peaks at orbital separations of roughly 3–10 au, near the H_2O snow line, with occurrence rates declining toward both smaller and larger separations. This picture was subsequently reinforced by RV surveys sensitive to giant planets beyond the H_2O snow line (Fernandes et al., 2019; Fulton et al., 2021). Nielsen et al. (2019a) also compared the demographics of giant planets and brown-dwarf companions. While the occurrence rate of giant planets decreases with increasing companion mass and orbital separation, the brown-dwarf population exhibits the opposite behavior, with occurrence rates increasing toward higher masses and wider separations. These contrasting trends support the interpretation that giant planets and brown dwarfs arise from distinct underlying populations and likely form through different formation mechanisms.

The two leading formation scenarios for giant planets are core accretion and gravitational instability. The core accretion scenario is a bottom up framework where a solid core forms through the accumulation of planetesimals and pebbles (Pollack et al., 1996) before the rapid accretion of gas that builds gas giant planets (Venturini et al., 2015). The gravitational instability scenario is a top down framework where a sufficiently massive and rapidly cooling disk becomes unstable to its own self-gravity in its outer part and fragments into bound clumps that may evolve into giant planets (Boss, 1998). These two scenarios form different planet populations; core accretion is expected to produce a population concentrated toward lower masses and smaller orbital separations, while gravitational instability preferentially forms massive companions at wide orbital separations, which makes demographics study informative about the possible dominant formation mechanism. Vigan et al. (2017) showed that the planet population predicted by core accretion models largely lies outside the sensitivity limits of the NACO-LP survey, whereas a large fraction of the population predicted by gravitational instability models falls within the survey detection space. The small number of detected companions therefore suggests that, if giant planets can form through gravitational instability, this formation pathway must itself be intrinsically rare. With the improved sensitivities of the SHINE and GPIES surveys, part of the parameter space populated by core accretion planets became accessible. Figure 1.16 compares synthetic planet populations generated from core accretion models (Emsenhuber et al., 2021) and gravitational instability models (Forgan et al., 2018) with the detections and survey sensitivities of SHINE (Vigan et al., 2021). Analyses by Vigan et al. (2021) and Nielsen et al. (2019a) suggest that both formation mechanisms contribute to the observed population of wide-orbit giant planets, with core accretion being the dominant formation pathway.




Figure 1.16: Comparison of the SHINE survey sensitivity for 77 FGK stars with simulated planet populations derived from the core accretion (CA) and gravitational instability (GI) formation models in the mass semi-major axis plane. Figure taken from Vigan et al. (2021).

1.6 Scope and outline of the dissertation

The field of exoplanet detection has undergone major developments in the XXIst century, and has emerged as both a scientific and technological achievement enabled by the development of large telescopes, extreme adaptive optics systems, coronagraphy, and increasingly sophisticated post-processing techniques based on a diversity of observing strategies, noise modeling, and planetary signal extraction. Through large observational surveys, direct imaging has provided important constraints on giant exoplanets at wide orbital separations and young evolutionary stages. Despite these advances, the performance of post-processing algorithms remains a major limiting factor in high-contrast imaging, particularly at small angular separations where residual speckle noise strongly affects detection sensitivity. This limitation is highlighted in Langlois et al. (2021), where classical processing techniques showed only limited improvement over the raw contrast close to the inner working angle. As a consequence, the close-in parameter space for young planets remains only weakly explored by direct imaging.

In this context, the main goal of this thesis focuses on improving the detection limits of the archival SHINE survey through the application of the Regime Switching Model, a high-performance detection framework that demonstrated strong performance in the EIDC. This work primarily addresses the first scientific objective of the SHINE survey, namely

the detection of new planetary-mass and substellar companions, while also providing updated detection limits that can contribute to the survey's fourth objective of placing tighter constraints on the occurrence rate of giant planets at wide orbital separations. Any new detection is bound to valuable information on planetary atmospheres, formation conditions, and system dynamics, considering the small number of directly imaged planets. This is even more important at shorter orbital separations, where planets are expected to be more frequent and where detections offer greater opportunities for complementary follow-up with indirect detection techniques, enabling more robust constraints on planetary masses and properties. The application of advanced post-processing techniques to the SHINE survey has also been explored in recent studies, notably with PACO (Chomez et al., 2023, 2025) and, more recently, NA-SODINN (Cantero et al., 2026). The analysis presented in this thesis was initiated prior to our awareness of the PACO application to the SHINE dataset. Following the publication of the PACO results, the scientific interest of this work was further reinforced by the fact that the two approaches rely on different statistical assumptions and detection mechanisms. As a consequence, they may exhibit different sensitivities near the contrast limit and therefore identify candidate signals that remain undetected by other techniques. This work focuses on the F150 sample of the SHINE survey, which was the only publicly available subset of the survey at the beginning of this PhD. The complete SHINE dataset has only recently become publicly available.

Reaching these objectives also requires understanding how the heterogeneous nature of the data in a large survey influences the performance of the detection framework. This thesis therefore investigates the different noise behaviors observed across the F150 sample as a consequence of varying atmospheric and observational conditions, and studies their impact on the detection thresholds derived from the RSM. Different approaches for determining these thresholds are explored, both assuming an underlying statistical distribution and without making explicit assumptions on the noise distribution through the maximization of the F1 score. The latter approach is not specific to the RSM framework and can be extended to other detection algorithms.

The rest of the present manuscript is structured as follows. Chapter 2 presents the selection of the datasets and their classification according to differences in noise behavior associated with atmospheric and observational parameters. Chapter 3 investigates the RSM framework and its different parameters in order to identify optimal configurations for HCI analyses. Chapter 4 investigates different thresholding strategies for detection, including approaches based on fitting statistical distributions across observational groups, as well as methods based on the maximization of the F1-score without explicit assump-

tions on the underlying noise distribution. Chapter 5 focuses on the resulting contrast improvements obtained through the application of the RSM, the recovery and assessment of candidate signals identified in the RSM maps through multi-epoch verification and photometric consistency analysis, as well as the investigation of false-positive detections within the RSM framework. Chapter 6 presents the application of the RSM framework to the ϵ Eridani system using Keck HCI observations. Finally, Chapter 7 summarizes the main conclusions of this work and discusses its broader implications for future HCI surveys and instruments.

Chapter 2

Clustering the SHINE survey

Abstract. The SHINE survey is one of the largest HCI surveys to date and provides important constraints on the occurrence rate of massive exoplanets. Targeting approximately 400 stars, the SHINE survey encompasses a broad diversity of observing conditions during the acquisition of its datasets, some of which strongly affect the performance of HCI observations. In this chapter, a clustering approach based on the k-means algorithm is employed to classify observations into meaningful groups reflecting distinct noise behaviors. This classification is motivated by the high computational cost of the RSM: rather than optimizing the RSM parameters for every individual target, the optimization is performed only for the representative center of each cluster and subsequently generalized to all observations belonging to that cluster. The resulting clustering identifies several characteristic observational regimes, including conditions dominated by the wind-driven halo and the low-wind effect. These regimes, characterized by different residual noise properties, provide insight into the impact of observational conditions on the statistical behavior of the noise.

This chapter is organized as follows. Section 2.1 describes the dataset sample, including the HC-DC reduction presented in Sect. 2.1.1 and the dataset selection procedure adopted in this work in Sect. 2.1.2. Section 2.2 presents the selection of the observational parameters used in the clustering analysis. Section 2.3 describes the application of the k-means clustering algorithm to the F150 sample and presents the corresponding results reported in Sabalbal et al. (2026). Finally, Sect. 2.4 discusses possible improvements to the clustering procedure.

2.1 Dataset description

The present analysis focuses on IRDIS observations obtained in the H2–H3 filter pair and within a field of view of $2''.2$ in diameter, where the RSM algorithm provides its best detection performance. The datasets analyzed here correspond to the F150 sample, comprising 150 stars observed in 343 observation sequences, later filtered down to 213 datasets after quality selection.

This section describes the datasets acquired for the F150 subsample of the SHINE survey. It is divided into two parts: Sect. 2.1.1 describes the HC-DC pre-processing and the corresponding data products used as input for this study, and Sect. 2.1.2 presents the criteria adopted for the selection of datasets included in the analysis.

2.1.1 HC-DC pre-processing

The observations were pre-processed by the High-Contrast Data Center (HC-DC; Delorme et al., 2017; Galicher et al., 2018; Beuzit et al., 2019; Maire et al., 2016). The HC-DC reduction process is described in detail by Delorme et al. (2017) and delivers several levels of processed data products. In this work, we use Level 1 data, which correspond to calibrated science-ready frames prior to speckle subtraction and high-contrast post-processing. These products include the coronagraphic science frames, off-axis PSF reference frames, parallactic angle information required for angular differential imaging, and additional metadata describing the environmental observing conditions and instrumental performance.

Calibrations

Cosmetic calibrations (dark subtraction, flat-field correction, and bad-pixel correction) follow the SPHERE Data Reduction and Handling (DRH) pipeline (Pavlov et al., 2008), provided by the SPHERE consortium to ESO. The centering of the star behind the coronagraph is performed using satellite spots generated from waffle patterns, or alternatively using unsaturated off-axis exposures obtained without the coronagraph.

The astrometric calibration of the IRDIS datasets is performed using on-sky calibration measurements combined with fixed instrumental corrections applied in pupil-tracking mode. In particular, the HC-DC apply the correction of 135.99° corresponding to the instrument pupil offset, together with the on-sky calibration derived by Maire et al. (2016). This results in a True North correction of $-1.75 \pm 0.08^\circ$ and a plate scale of $12.255 \pm 0.009 \text{ mas pixel}^{-1}$ for IRDIS.

HC-DC Products

The HC-DC reduction for IRDIS produces master reduced data cubes together with several auxiliary products required for the analysis:

- `SCIENCE_PARA_ROTATION_CUBE-rotnth.fits`: a one-dimensional array containing the parallactic angle associated with each science frame, used for angular differential imaging processing.
- `REDUCED_MASTER_CUBE-center_im.fits`: a four-dimensional master science cube. The first dimension corresponds to the two spectral channels (H2 and H3), the second dimension contains the temporal sequence of frames, and the last two dimensions correspond to the spatial axes of the coronagraphic images.
- `SCIENCE_PSF_MASTER_CUBE-median_unsat.fits`: a four-dimensional cube containing unsaturated PSF reference images used for photometric calibration and PSF modeling. The first dimension corresponds to the two IRDIS wavelengths, the second dimension typically contains two frames (before and after the science sequence), and the last two dimensions correspond to the spatial axes of the PSF images.

Additional useful products for this study include:

- `PROFILES-all_contr_curve_1sigma.fits`: raw contrast curves derived from the HC-DC reduction.
- `IRD_ATMO_CONDITIONS-seeing.txt`: seeing and coherence time measurements obtained with the Differential Image Motion Monitor (DIMM).
- `FRAME_SELECTION_VECTOR-frame_selection_vector.fits`: frame-selection quality flags associated with both coronagraphic science frames and off-axis PSF exposures.

In addition to these products, the data headers provide atmospheric and AO telemetry information from SPARTA, including seeing estimates, coherence time, Strehl ratio, and wind speed measurements. In this work, however, we consider atmospheric parameters such as seeing, wind speed, and coherence time derived primarily from differential image motion monitor (DIMM, Sarazin & Roddier, 1990) measurements, following the recommendations of Milli et al. (2017). When DIMM seeing measurements are unavailable, they are inferred from SPARTA seeing values using the empirical relationship established by Milli et al. (2017). The Strehl ratio is estimated empirically from the unsaturated PSF exposures.

2.1.2 Selection of datasets

To ensure a robust sample of stars observed under conditions ranging from fair to excellent, I applied a series of selection criteria based on the quality of both coronagraphic and non-coronagraphic data, as well as atmospheric conditions. Observations not meeting these criteria were excluded from the analysis.

Non-Coronagraphic Images: The following criteria were used to discard poor-quality observations based on the off-axis PSF:

- The star was located outside the coronagraphic mask.
- The off-axis PSF exhibited strong secondary lobes, indicative of the low-wind effect (see Milli et al., 2018).
- The full width at half maximum (FWHM) exceeded 5.5 pixels.

Coronagraphic Data: For the coronagraphic sequences, I combined several quality assessment tools to identify and reject unsuitable data:

- Frame quality diagnostics from the HC-DC pipeline.
- Frame correlation analysis using the VIP package (Gomez Gonzalez et al., 2017; Christiaens et al., 2023).
- Visual inspection.

Cubes in which the majority of frames were corrupted or misaligned were systematically excluded.

Atmospheric Conditions: Observations were rejected if any of the following atmospheric criteria were satisfied:

- Seeing exceeded $2''$.
- Strehl ratio was below 0.5.
- Less than 40 frames were available in the ADI sequence.
- The raw contrast at 500 mas was nonphysical (negative values or greater than unity).

These filtering steps resulted in a total of 213 observations for 150 stars. These constitute the data sample used in this paper. The star names, along with the corresponding descriptive information and parameters, are all summarized in Table A.1 see (Appendix A).

2.2 Choice of relevant parameters

The datasets considered in this work were obtained under a wide range of observing conditions, which directly affect the residual speckle noise properties and therefore the performance of PSF-subtraction algorithms and detection methods such as RSM. The optimization of the RSM parameters in Dahlqvist et al. (2021a) showed that different observing conditions can lead to different optimal processing parameters. For this purpose, we adopt the same approach as used in Dahlqvist et al. (2022) for the SHARDDS survey, where datasets were grouped according to their observing conditions to improve the robustness and efficiency of the post-processing optimization.

Figure 2.1 presents the distributions of key atmospheric, instrumental, and observational parameters for the selected datasets, illustrating the variability in conditions such as seeing, Strehl ratio, raw contrast, and stellar magnitude. The figure also highlights the diversity in wind speed, spanning both low- and high-wind regimes, as well as variations in coherence time. This broad range of observing conditions is directly reflected in the diversity of residual noise structures present in the coronagraphic observations.

Below, we summarize the main parameters typically considered in HCI analyses, together with the interdependencies revealed by the correlation matrix shown in Fig. 2.2:

- **Seeing:** describes the blurring of astronomical images caused by atmospheric turbulence. Practically, it is quantified by the angular FWHM of the atmospheric PSF, commonly approximated by $0.976\lambda/r_0$, and expressed in arcseconds. Lower seeing values generally correspond to improved AO correction and sharper PSFs, whereas higher seeing values lead to stronger residual wavefront errors and increased speckle noise.
- **Stellar Magnitude:** The apparent brightness of the target star directly affects AO performance. Brighter stars enable more efficient wavefront sensing and improved correction quality, whereas fainter stars generally lead to reduced AO performance and increased residual noise.
- **Strehl Ratio:** A direct indicator of AO correction quality, with values approaching unity corresponding to near diffraction limited performance.
- **DIT:** Detector integration time, primarily determined by the target brightness and the observing strategy. In general, fainter stars require longer integration times to achieve a sufficient S/N. The DIT is therefore naturally correlated with stellar magnitude and influences the S/N of individual frames.




Figure 2.1: Distributions of atmospheric, instrumental, and observational parameters for the 213 selected observations in the F150 sample. These parameters include seeing, Strehl ratio, stellar magnitude, number of frames in the ADI sequence, wind speed, coherence time, raw contrast at 500 mas, ambient and primary mirror temperatures, and total parallactic angle rotation.

- **Number of Frames:** For a given observing sequence, the total integration time is proportional to the product of the DIT and the number of frames. Consequently, longer DITs result in fewer frames. The number of frames determines the temporal sampling of quasi-static speckle evolution, with larger frame counts providing a finer sampling that can improve the performance of post-processing algorithms such as the RSM.
- **Wind Speed:** A key parameter influencing noise structures in coronagraphic observations. Low wind speeds can trigger the low-wind effect, producing characteristic secondary lobes in the PSF, while high wind speeds can generate extended AO residual halos. Both regimes degrade achievable contrast (see Sect. 1.3.2 for description of both of these effects).
- **Coherence Time** (τ_0): Represents the characteristic timescale over which atmospheric wavefront distortions remain correlated. Larger coherence times generally improve AO correction efficiency, although they may coincide with low-wind conditions that introduce additional PSF artifacts. Short coherence times correspond to rapidly evolving turbulence and reduced correction performance.
- **Raw Contrast:** Provides an empirical estimate of the achievable contrast at a reference angular separation (here 500 mas).
- **Total Parallax Angle Rotation:** Measures the field rotation during the observing sequence and is a critical parameter for ADI performance. Insufficient rotation increases self-subtraction of astrophysical signals and reduces detection sensitivity.

The correlation matrix shown in Fig. 2.2 reveals several expected relationships between these parameters:

- A strong correlation between the DIT and the number of frames reflects their inverse proportionality, together with the relatively similar total integration times adopted across the observations (typically less than one hour).
- Correlations between seeing, coherence time, and wind speed are consistent with their common dependence on atmospheric turbulence conditions.
- Stellar magnitude correlates with the Strehl ratio, the DIT, and inversely with the number of frames, reflecting their shared dependence on target brightness and AO performance.
- Raw contrast correlates with several atmospheric and instrumental parameters, confirming its role as a global indicator of observation quality.




Figure 2.2: Correlation matrix between the atmospheric, instrumental, and observational parameters of the F150 sample. These parameters include the total parallactic angle rotation, seeing, wind speed, coherence time, raw contrast at 500 mas, Strehl ratio, exposure time, ambient temperature, primary mirror M1 temperature, stellar magnitude, the temperature difference between ambient air and M1, and the number of frames in the ADI sequence. Additionally, two clustering-related quantities are shown: the cluster assignment of each dataset and the distance of each dataset to the center of its assigned cluster.

Considering the strong correlation between exposure time, number of frames, and stellar magnitude, we retain only the number of frames in the final parameter set, as it has a direct impact on the performance of the RSM used in this study.

In conclusion, the clustering analysis is performed using the following parameters: seeing, raw contrast, Strehl ratio, wind speed, coherence time, total parallactic angle rotation, and number of frames. This selection follows the methodology proposed by Dahlqvist et al. (2022) and provides a compact yet informative description of the observing conditions relevant to high-contrast imaging performance.

2.3 Clustering

We adopted the clustering method outlined in Dahlqvist et al. (2022), grouping observations with similar characteristics using k-means clustering. This approach not only helps identify distinct classes of SPHERE observations but also reduces computational load when optimizing PSF subtraction techniques and RSM parameters. Observations were grouped into six clusters. The number of clusters was determined by maximizing the silhouette score, a metric used to evaluate the quality of clustering. The k-means algorithm automatically selected a representative observation for each cluster. The parameters for PSF subtraction and the RSM algorithm can then be optimized for these representative data sets, and generalized to other observations within the same cluster. To ensure unbiased parameter selection, we made sure that none of the cluster centers contained datasets with bright circumstellar signals.

2.3.1 Effectiveness of clustering

To visualize the clustering effectiveness based on selected parameters, we applied Principal Component Analysis (PCA) for dimensionality reduction, enabling clear separation of clusters in the latent space. We projected the environmental datasets onto their first three principal components. Figure 2.3 (top) shows that most clusters are well-separated, although clusters 2 and 6 overlap. This overlap is resolved when examining the third principal component, as shown in Plot 2.3 (bottom).

2.3.2 Results of clustering

Figure 2.4 displays a 200×200 pixel crop of representative individual frames from different clusters. These images illustrate the varying noise behavior across clusters, with some frames showing secondary lobes at close separations from the star indicative of the low-wind effect (Fig. 2.4a), while others exhibit strong elongated (aka butterfly) patterns characteristic of the wind driven halo (Fig. 2.4e).

For a physically grounded interpretation of the nature of the various clusters, Fig. 2.5 presents their visualization based on key environmental parameters, revealing distinct traits for each cluster:

- cluster 1 (Fig. 2.4a, 11% of observations) is associated with low wind speed (Fig. 2.5b) and high coherence time (Fig. 2.5a), suggesting the presence of the low-wind effect;




Figure 2.3: Visualization of clusters using principal component analysis. *Top* shows the first vs second component projection. *Bottom* shows second vs third component projection.

- cluster 2 (Fig. 2.4b, 48% of observations) represents observations under generally fair to good conditions;
- cluster 3 (Fig. 2.4c, 8% of observations) is characterized by a large number of frames (Fig. 2.5d);
- cluster 4 (Fig. 2.4d, 3% of observations) is defined by low raw contrast, low coherence time, and high seeing (Figs. 2.5a, 2.5c);
- cluster 5 (Fig. 2.4e, 15% of observations) shows high wind speed, high seeing, and low coherence time (Figs. 2.5a, 2.5b), which makes it susceptible to strong wind driven halo;




Figure 2.4: Image representation of single frames from various observations across different clusters, labeled (a) to (f) corresponding to clusters 1 to 6 presented in log scale.

- cluster 6 (Fig. 2.4f, 15% of observations) is associated with large parallactic angle ranges (Fig. 2.5d).

This successful clustering of data into meaningful categories will allow us to later analyze noise distributions and detection limits based on well-defined observational conditions, enabling a more targeted approach to interpreting results.

2.3.3 Noise correlation test

To assess how well the chosen cluster centers reflect the noise behavior of other datasets within the same cluster, we performed a pixel correlation test. This test computes the correlation between the selected cluster center and the other datasets in the cluster. Additionally, we compare this correlation to that of the dataset exhibiting the highest correlation with other datasets within the cluster.

To study the correlation between datasets, we adopted the approach used in Juillard et al. (2024), using a correlation metric from the VIP package based on the Spearman rank correlation coefficient. This coefficient was chosen to avoid the normality assumption required by the Pearson coefficient and to circumvent the dependence of the Structural Similarity Index (SSIM) on the image dynamic range. For each dataset, we identified the frame within the science cube that exhibits the highest correlation with other




Figure 2.5: Projection of cluster distributions across parameters: seeing, coherence time, wind speed, raw contrast, Strehl ratio, total parallactic angles, and number of frames.

frames, designating it as the representative frame of that dataset. We then computed the pixel correlation between this representative frame and the representative frames of all other datasets in the cluster. This process was repeated for all datasets within one cluster as well as for all clusters. The dataset exhibiting the highest correlation with the other datasets in its cluster is expected to share the most prominent noise features and thus be the most representative. We then compared the Spearman correlation of the selected cluster centers, chosen based on environmental conditions, to that of the most correlated dataset in each cluster. The results are shown in Figure 2.6.

The high correlation values for the cluster centers confirm that these selections reliably represent their respective clusters. This correlation metric proves less effective for clusters affected by strong wind-driven halo effects (clusters 4 and 5). This limitation arises from the rotation of wind-driven speckles with the parallactic angle (Cantalloube




Figure 2.6: Spearman rank correlation comparison between cluster centers and most correlated datasets in each cluster.

et al., 2020), which introduces additional complexity. In such cases, high correlation coefficients tend to identify datasets with less pronounced wind effects rather than true feature similarity, reducing the effectiveness of correlation as a metric in these specific scenarios.

2.4 Improving the clustering

This section is motivated by limitations identified in the initial clustering results, particularly in the identification of datasets affected by the LWE. In the clustering shown in Fig. 2.5, some datasets exhibiting LWE-like signatures were not grouped within the cluster expected to represent these conditions (cluster 1). In addition, this cluster does not systematically show the low Strehl ratios typically expected from observations affected by this effect.

In the initial clustering, the LWE-dominated group was primarily driven by two parameters: coherence time and wind speed. However, a more physically motivated description of this effect should also include the temperature difference between the ambient air and the telescope structure, in particular the spiders, as discussed in Milli et al. (2018). The metadata do not provide measurements of the spiders temperatures used in Milli et al. (2018) to quantify the LWE. However, they do include measurements of the ambient temperature and the primary mirror M1 temperature, which does not correspond exactly to the spider temperature used in that study.

Nevertheless, the distributions of these temperatures shown in Fig. 2.1 indicate that the M1 temperature is generally slightly lower than the ambient temperature. We therefore use the temperature difference between these quantities as an approximate proxy for the thermal gradients relevant to the LWE. This temperature difference shows stronger correlations with the Strehl ratio, coherence time, and wind speed than either temperature considered individually, suggesting that it captures part of the physical conditions associated with the LWE.

Including this temperature difference as an additional parameter in the clustering improves the identification of datasets affected by the LWE. As shown in Fig. 2.7, the main physical classes identified in the initial clustering are preserved, while the datasets affected by the LWE are more consistently grouped together.

In the improved clustering shown in Fig. 2.7, cluster 4 is characterized by large temperature differences combined with low wind speed and systematically reduced Strehl ratios, which are consistent with typical LWE conditions. Another group (cluster 3) also shows moderately large temperature differences and low wind speed, but retains relatively high Strehl ratios. This group includes observations obtained after August 2017, following the implementation of mitigation measures on the telescope spiders, which reduced the impact of the LWE despite similar environmental conditions.




Figure 2.7: Projection of the dataset clusters in the parameter space defined by seeing, coherence time, wind speed, raw contrast at 500 mas, Strehl ratio, total parallactic angle rotation, number of frames, and the temperature difference between ambient air and the primary mirror M1.

Chapter 3

Regime Switching Model algorithm

Abstract. The RSM algorithm is one of the most advanced detection methods in HCI, relying on the modeling of astrophysical signals in post-processed frames through planetary and noise hypotheses within a regime-switching framework. Its performance depends on a number of parameters that can significantly affect both the sensitivity and the reliability of detections. In this chapter, these parameters are systematically investigated in order to quantify their impact on the resulting contrast curves. The analysis demonstrates that parameter selection can lead to substantial differences in performance, with contrast improvements of up to a factor of five between unfavorable and optimal configurations.

The chapter is divided as follows. Section 3.1 presents the RSM algorithm, including its concept in Sect. 3.1.1, its application to HCI in Sect. 3.1.2, its updates in Sect. 3.1.3, and its performance in Sect. 3.1.4. Section 3.2 presents the parametrization of the RSM algorithm, including the PSF-subtraction parameters discussed in Sect. 3.2.1, the internal RSM parameters presented in Sect. 3.2.2, and the parametrization adopted in this thesis in Sect. 3.2.3. Finally, methods for improving the detection performance of the RSM framework are presented in Section 3.3, including the exploration of different parameter ranges in Sect. 3.3.1 and the combination of consecutive observing nights in Sect. 3.3.2.

3.1 RSM explained

3.1.1 RSM concept

The Regime Switching Model (RSM), also referred to as a Markov Regime Switching Model (Hamilton, 1988), is a well-known statistical framework in economics used to describe time series whose behavior changes depending on the underlying conditions in which the system evolves, such as periods of expansion and recession. These underlying conditions are called regimes (or states). A regime represents the physical phase occurring in reality that produces changes in the observed time series. Since the regime itself is not directly observable, it must be inferred from the statistical properties of the data.

Although the distinction between such phases may appear straightforward, financial time series typically exhibit strong short-timescale fluctuations, and not every increase necessarily indicates the onset of an expansion phase. The RSM provides a framework to distinguish between these phases through several components. First, the model assumes a statistical description for each regime and evaluates the likelihood that the observed data are explained by each of these models. Second, it incorporates transition probabilities between regimes based on realistic assumptions about how systems evolve over time. For example, a rapid transition from recession to expansion is generally unlikely, and this is reflected by assigning a low transition probability between these two regimes. Finally, the model includes a short-memory dependence between successive time steps, meaning that the probability of the current regime depends partly on the regime probability inferred at the previous step.

Dahlqvist et al. (2020) identified a similarity between financial time series and the temporal evolution of intensity values in residual derotated frames obtained in HCI. In the case of quasi-static speckles, the staticity of the quasi-static signal is lost during the derotation process, which leads to fluctuations in intensity values between low and high levels at a given spatial location when viewed over time. In contrast, an astrophysical signal remains temporally consistent and therefore produces persistently higher intensity values over time. Dahlqvist et al. (2020) showed that the RSM can infer these two regimes: noise regime and planetary regime, from the distinct signatures they leave in the time series. As a result, the RSM can exploit temporal information that is typically not fully utilized in classical post-processing approaches, thereby improving the detection of faint companions.

3.1.2 RSM applied to HCI

Theory of RSM

As discussed in Chapter 1, the final derotated ADI frames remain dominated by residual noise structures that evolve radially. Building on this, Dahlqvist et al. (2020) exploited the spatial-temporal statistics of the derotated ADI residuals, together with a planetary signal model, to construct the RSM framework capable of inferring the state of each spatial-temporal position as either speckle dominated or planetary signal dominated. In the noise regime, the spatial-temporal intensity series of pixels in a given annulus can be modeled as the sum of a constant mean residual level μ and a space and time dependent fluctuation term $\varepsilon_{r,t}$ describing the local residual variability, where r denotes the spatial variable (R being the total number of pixels in the annulus) and t denotes the temporal variable (T being the total number of frames).

In the planetary regime, an additional contribution corresponding to an astrophysical signal is included. This signal is modeled using a scaled PSF βP , where P represents the PSF model and β denotes the flux scaling factor. The true state of each pixel is not treated as strictly belonging to one regime or the other, but instead as a probabilistic quantity that accounts for both states, given by $K_{s(r,t)}$, where the sum of the probabilities over the two regimes equals 1. The following equation describes this model:

$$x_{r,t} = \mu + \beta K_{s(r,t)} P + \varepsilon_{s(r,t)} = \begin{cases} \mu + \varepsilon_{s(r,t)} & \text{for pure noise regime} \\ \mu + \beta P + \varepsilon_{s(r,t)}, & \text{for pure planetary regime} \end{cases} \quad (3.1)$$

where $x_{r,t}$ represents the spatial-temporal series of pixels at indices (r,t) , and $K_{s(r,t)}$ is the probabilistic term at each (r,t) position for each state s .

Since the planetary signal exhibits a two-dimensional PSF morphology with a characteristic width comparable to the FWHM, the statistical analysis is performed on small spatial patches rather than on individual pixels. These patches, referred to as crops, have a spatial size of $\theta \times \theta$ pixels and each constitute one realization of the spatio-temporal dataset $x_{r,t}$. The spatio-temporal realizations $x_{r,t}$ are therefore obtained by extracting such crops at different temporal frames and spatial positions along a given annulus. Consequently, $x_{r,t}$ has dimensions of $\theta \times \theta \times R \times T$.

Dahlqvist et al. (2020) showed that the residual term $\varepsilon_{s(r,t)}$, which corresponds to the observed signal $x_{r,t}$ after subtracting the mean level μ , follows approximately a Gaussian distribution at large separations and a Laplacian distribution at small separations using VLT/NACO data. The likelihood model therefore evaluates how well the residual term $x_{r,t} - \mu$ (or equivalently $x_{r,t} - \mu - \beta P$ under the planetary hypothesis) follows the assumed noise distribution. This likelihood function $\eta_{s(r,t)}$ for state s at indices (r,t) , assuming a Gaussian distribution, is given by

$$\eta_{s(r,t)} = \sum_n \frac{\theta^2}{\theta^2} \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\frac{\left(x_{(r,t)}^n - K_{s(r,t)} \beta P^n - \mu \right)^2}{2\sigma^2} \right], \quad (3.2)$$

where σ denotes the standard deviation.

As a reminder, the probability $\xi_{s,(r,t)}$ for a pixel (r,t) of belonging to regime s does not depend only on this likelihood, but also on the probability of being in regime s at the previous step, denoted as $\xi_{q,(r,t)-1}$, and on the transition probabilities between regimes given by the matrix $p_{q,s}$. This probability is computed as

$$\xi_{s(r,t)} = \sum_{q=0}^1 \frac{\eta_{s(r,t)} p_{q,s} \xi_{q,(r,t)-1}}{\sum_{q=0}^1 \sum_{s=0}^1 \eta_{s(r,t)} p_{q,s} \xi_{q,(r,t)-1}} \quad (3.3)$$

where the denominator represents the normalization factor ensuring that the posterior probabilities sum to unity.

The transition probabilities between regimes are described by the matrix $p_{q,s}$, which includes four components: remaining in the noise regime ($p_{0,0}$), remaining in the planetary regime ($p_{1,1}$), transitioning from noise to planetary regime ($p_{0,1}$), and transitioning from planetary to noise regime ($p_{1,0}$). Based on the realistic assumption that planetary signals are rare within a given annulus, Dahlqvist et al. (2020) assumed that at most one planetary signal is present per annulus. Under this assumption, the transition matrix becomes

$$p_{q,s} = \begin{pmatrix} p_{0,0} = 1 - 1/(TR) & p_{1,0} = 1/T \\ p_{0,1} = 1/(TR) & p_{1,1} = 1 - 1/T \end{pmatrix}. \quad (3.4)$$

The recursive evaluation of $\xi_{s,(r,t)}$ for all pixels (r,t) produces probabilities for each regime at each spatial-temporal position. The RSM detection map is constructed from the posterior probability of belonging to the planetary regime.

RSM in practice

The RSM analysis can be divided into three main steps: (i) transforming the ADI sequence into post-processed spatial-temporal series, (ii) modeling these series and computing the probability of belonging to each of the two regimes, and (iii) finally deriving the RSM detection map.

1. Input:

Starting from an ADI sequence, the RSM algorithm employs a variety of classical post-processing algorithms such as PCA, NMF, LOCI, and LLSG, as presented in Dahlqvist et al. (2020), as well as forward-model-based approaches such as FM-LOCI and FM-KLIP, as presented in Dahlqvist et al. (2021b). After applying PSF subtraction and derotating the frames to a common angle, as illustrated in Fig. 1.10, the RSM algorithm uses the cube obtained after step D instead of combining all frames into a single image as seen in step E.

Using this cube of derotated frames, the RSM spans the data both radially and temporally. A set of annuli is defined between a minimum radius corresponding to the IWA and a maximum radius defining the outer boundary of the field of view. For example, this range can extend from 10 pixels to 100 pixels in separation. The annuli are defined at integer separations (e.g., 10, 11, ..., 100).

For each annulus, starting from a first pixel located at position r within the annulus, a spatial patch of size $\theta \times \theta$ centered on this pixel is extracted. This patch is then stacked temporally across the T frames, producing a temporal realization of size $\theta \times \theta \times T$ associated with the pixel at position r . The same procedure is repeated for the next pixel within the annulus, and the corresponding temporal realization is appended after that of the previous pixel. This process is performed sequentially for all pixels belonging to the annulus, thereby constructing a complete spatial-temporal series $x_{r,t}$ associated with that annulus.

This procedure is then repeated independently for each annulus across the considered radial range.

2. RSM:

This step corresponds to the estimation of the regime-switching probabilities for each state, as described in the theoretical framework presented in the previous section. Several parameters must be estimated for this purpose, including the assumed probability distribution of the intensity values and its two components, namely the mean μ and the variance σ , as well as the planetary model P , the planetary flux β , and the transition probability $p_{q,s}$.

The choice of the probability distribution is discussed in more detail in section 3.2.2 and is fully reviewed in Dahlqvist et al. (2020). For each $\theta \times \theta$ spatial-temporal series, the pixel values are fitted to the assumed distribution, from which the mean and variance are estimated. The PSF model P is derived either from the off-axis PSF or from the forward-modeled PSF, while the planetary flux β is defined as a multiplicative factor δ of the standard deviation, given by $\beta = \delta\sigma$, where $\delta = 5$ corresponds to searching for a 5σ planetary signal. Finally, the transition probability $p_{q,s}$ can be inferred directly from the number of pixels in the annulus R and the number of frames T , as shown in Equation 3.4.

Once these parameters are estimated, the probability $\xi_{s(r,t)}$ of belonging to each state is computed for every position (r,t) following Equation 3.3.

3. Output:

The scientific objective is to estimate the confidence that a pixel located at position r belongs to the planetary regime, quantified by $\xi_{1(r)}$. The RSM algorithm produces both state probabilities $\xi_{0(r,t)}$ and $\xi_{1(r,t)}$ at each spatial-temporal position (r,t) . In its final step, only the probability associated with the planetary regime ($s = 1$) is retained, forming a cube of RSM scores across all frames T for each spatial position r within the annulus. The mean score over the T frames is then computed to obtain a single probability value per spatial position r . This results in a 2D detection map based on the RSM probability associated with the planetary state.

While an intuitive interpretation would suggest that a high RSM score reflects high confidence in the presence of a planetary signal and a low score reflects low confidence, the RSM framework corresponds to an imbalanced classification problem in which this interpretation does not strictly hold from a statistical perspective. This behavior arises from the assumption imposed by the algorithm that

at most one planetary signal is present within a given annulus, which introduces a strong imbalance between noise-dominated and planetary-signal classifications. As a result, a dedicated detection threshold must be defined, which is explored in Chapter 4.

3.1.3 Updates to RSM

Dahlqvist et al. (2021b,a) presented several important updates to the RSM algorithm that significantly improved its robustness and detection performance. First, Dahlqvist et al. (2021b) introduced the use of forward-modeled PSFs within KLIP and LOCI in order to account for distortions affecting the planetary signal during PSF subtraction, such as over-subtraction and self-subtraction (Pueyo, 2016). The same work also proposed the forward–backward RSM approach, which refines the estimation of regime probabilities by allowing them to depend not only on the probability of the previous step, as described earlier, but also on the probability of the subsequent step. In addition, it investigated the possibility of combining several PSF-subtraction techniques to construct a single spatial-temporal series, thereby improving the sensitivity to faint companions.

A second major improvement was introduced in Dahlqvist et al. (2021a) through the development of the auto-RSM framework, which optimizes both the PSF-subtraction parameters and the internal RSM parameters in order to maximise detection performance. All these parameters are described in more detail in Sect.3.2.

3.1.4 RSM performance

The RSM algorithm has demonstrated improved performance compared to classical detection techniques in Dahlqvist et al. (2020) and Dahlqvist et al. (2021b). These studies showed that RSM provides better detection sensitivity than S/N maps derived from classical algorithms such as PCA, NMF, LLSG, and LOCI. This improvement was demonstrated both through visual inspection of the recovery of injected synthetic companions and through quantitative evaluation using ROC curves. The detection sensitivity was further improved when combining several PSF-subtraction techniques, as shown in Dahlqvist et al. (2021b), and through the optimization of the RSM parameters in Dahlqvist et al. (2021a).

This performance was further confirmed in the EIDC, as discussed in Sect. 1.4.4. RSM ranked among the most advanced detection techniques and demonstrated clear improvements over both classical and other advanced ADI-based approaches. In particular, RSM achieved one of the highest F1 scores while maintaining one of the lowest false-positive rates, highlighting its strong detection capability.

The performance metric used in the studies of Dahlqvist et al. (2020) and Dahlqvist et al. (2021b) was primarily based on TPR and FPR. In this work, we instead adopt the RSM completeness curve described in Dahlqvist et al. (2021b), as the main performance metric. At a given angular separation, the completeness represents the faintest planet-to-star flux ratio for which a specified fraction of injected synthetic companions is recovered above a defined detection threshold. For example, a completeness level of 50% corresponds to recovering half of the injected synthetic companions above the adopted detection threshold, which is the convention typically used to define 5σ contrast curves. For these contrast curves, the detection threshold is set to 5σ , corresponding to a false alarm probability of approximately 3×10^{-7} under the assumption of Gaussian noise. In the RSM framework, however, the final probability is a nonlinear function of the input signal, preventing the direct application of such a threshold. In Dahlqvist et al. (2021b), the completeness formalism was therefore adapted to the RSM framework by defining the detection threshold as the RSM probability associated with the brightest speckle in the field of view. Alternative definitions of the detection threshold, explored in this thesis, are presented in Chap. 4. A comparison between the 5σ contrast curve and the 50% completeness curve for PCA is shown in Figure 3.1, together with the corresponding RSM completeness curve.

Instead of defining the detection threshold using the brightest false positive in the entire field of view, we define the detection threshold annulus-wise based on the brightest speckle within each annulus. This allows us to investigate the behavior of the algorithm as a function of angular separation, since the brightest speckle in the full field of view typically corresponds to a false positive located in the innermost annulus (near the IWA). We acknowledge that this approach implies accepting a higher number of false positives than would be used in a final detection analysis; however, it is adopted here as an intermediate step to better characterize the radial behavior of the algorithm. A more rigorous definition of detection thresholds is presented in Chapter 4.




Figure 3.1: Comparison of the HCI adopted 5σ PCA contrast curve (blue), the 50% completeness contrast for PCA (red), and the 50% completeness contrast for RSM (green).

3.2 Parameters in RSM

The RSM algorithm relies on several parameters, including those related to the ADI post-processing techniques, the internal RSM parameters, and the choice of the optimal combination of PSF-subtraction methods. Proper tuning of these parameters is essential for improving the subtraction of speckles and enhancing the detectability of planetary signals. In particular, a key requirement for the RSM to perform well is the construction of spatial-temporal series that accurately represent the statistical behavior of both speckles and potential planetary signals, together with models that properly describe these two classes.

In this section, we present the different parameters used by the RSM algorithm and discuss their influence on the overall performance of the method.

3.2.1 PSF subtraction parameters

The PSF-subtraction techniques implemented within the RSM framework are APCA (Absil et al., 2013), LOCI (Lafrenière et al., 2007), NMF (Ren et al., 2018), LLSG (Gomez Gonzalez et al., 2016), and forward-model implementations of KLIP and LOCI (Pueyo, 2016; Dahlgqvist et al., 2021b). The parameters associated with these PSF-subtraction techniques include:

- the number of principal components for APCA, NMF, LLSG (rank), and FM-KLIP;
- the minimum field rotation for APCA, FM-KLIP, LOCI, and FM-LOCI;
- the error-minimization tolerance for LOCI and FM-LOCI;
- the number of azimuthal segments for APCA, LLSG.

Careful tuning of these parameters is essential. For instance, selecting too few principal components may result in insufficient speckle suppression, whereas an overly aggressive parameterization may lead to self-subtraction and partial removal of the planetary signal. Choosing appropriate parameter values is therefore critical for maximizing the detection performance of the RSM.

In the RSM framework, the optimization of these parameters relies on the concept of inverted parallactic angles, which blur any planetary signal while preserving the temporal correlation and statistical properties of the noise. The optimal parameter values are those that minimize the 5σ contrast achieved by the corresponding PSF-subtraction technique when applied to data processed with inverted parallactic angles. The auto-RSM procedure explores a range of values for each parameter and evaluates the corresponding planet-to-star flux ratio. For discrete parameters (e.g., the number of principal components or the number of azimuthal segments), the optimization is performed by testing each value within the selected range. For continuous parameters, such as the tolerance or the minimum field rotation, the optimization relies on a Bayesian framework or alternatively on a particle swarm optimization approach (Dahlgqvist et al., 2022) to determine the optimal values.

The parameter ranges were tested for different instruments, and recommended values were provided for SPHERE, NIRC2, and LMIRCam in Dahlgqvist et al. (2021a).

3.2.2 RSM parameters

Several parameters in the RSM algorithm are left variable and must be optimized for each dataset because of their importance and their impact on the final RSM map. These parameters include the batch crop size θ used for the planetary model, the planetary flux β , the assumed probability distribution describing the noise properties, the regions selected in the de-rotated cube to estimate these noise properties, and the optimal combination of PSF-subtraction techniques.

In the auto-RSM framework, these parameters are optimized using the following procedure: a synthetic planetary signal with a contrast corresponding to 5σ in a final PCA-processed frame is injected at a representative location within the annulus, and the ratio between the maximum RSM score measured within a 1 FWHM region around the injection location and the maximum RSM score measured elsewhere in the annulus is computed. This ratio is then maximised in order to determine the optimal set of RSM parameters. The main RSM parameters and their impact on the detection performance are detailed below:

Crop size θ

The batch crop size corresponds to the pixel region used for the planetary model. For off-axis PSF models, the crop size is typically chosen to be approximately one FWHM in diameter. For forward models, the crop size requires an increase (approximately 1.5 FWHM) in order to account for the negative lobes of the PSF. This increase was shown to improve performance at small separations, as demonstrated in Dahlqvist et al. (2021b).

At larger separations, however, the RSM performance can decrease when using a large crop size. This behavior results from the fact that the apparent displacement of the planetary signal increases linearly with separation for a fixed parallactic rotation angle. As a result, the negative lobes become less significant at large separations, while increasing the crop size introduces additional noise around the PSF into the planetary model rather than relevant signal information, which can lead to partial overfitting of the local noise structure.

Region considered for the computation of noise properties

The RSM algorithm includes a parameter defining the region used to estimate the noise properties, from which the noise mean μ and the standard deviation σ are computed. The default configuration in RSM uses both spatial and temporal information and is referred to as the *spatial-temporal (ST)* estimation. In this case, all pixels within a given annulus and across all frames are used to estimate a single mean and standard deviation for that annulus.

Other estimation regions were explored by Dahlqvist et al. (2021a), including the *frame-based (FR)* estimation, in which the noise properties are computed separately at each frame using all pixels within the annulus. Another approach is the *segment with mask (SM)* estimation, which consists of using the information within a segment of approximately 3 FWHM around each pixel located at separation r , while excluding the central 1 FWHM region around that pixel, and considering this segment across all frames. Additional estimation strategies were also proposed in Sect. 3.2.2 of Dahlqvist et al. (2021a), but were discarded in practice. This region selection parameter is referred to as the "var" parameter in the RSM implementation.

We found no significant difference in performance between these estimation strategies, as shown in Fig. 3.2b, with all configurations converging to similar completeness curves at all separations.

The choice of the probability distribution used in the RSM framework was initially considered as an additional parameter, with several candidates including Gaussian, Laplacian, Huber loss based distributions, and a weighted combination of Gaussian and Laplacian distributions. In practice, this parameter was later fixed to a weighted combination of Gaussian and Laplacian distributions, which was found to provide the best overall performance.

Planetary flux β (or the multiplicative factor δ)

The planetary flux assumed for the planetary model is an important parameter in the RSM framework. The default approach is to define the planetary flux as a multiplicative factor δ of the standard deviation σ , and to optimize the value of δ , with δ typically taken as an integer (e.g. 1, 2, ..., 5, ...). A high value of δ corresponds to searching for signals located further in the tail of the noise distribution, and therefore to a higher




Figure 3.2: Completeness variability for selected RSM parameters, together with a comparison to their optimal values. *Left:* multiplicative planetary flux factor δ spanning values between 1 (1σ) and 7. *Right:* different regions used for estimating the noise properties: ST (spatial-temporal), FR (frame-based), and SM (segment-with-mask).

confidence level in distinguishing a planetary signal from noise. In contrast, low values of δ correspond to testing fainter signals, but also make the separation between the noise regime and the planetary regime more difficult for the RSM algorithm, with $\delta = 0$ corresponding exactly to the noise regime.

This behavior is illustrated in Fig. 3.2a, which shows a degradation in performance for low values of δ . In this work, β and δ are used interchangeably, since alternative definitions of β were explored in Sect. 3.2.3 of Dahlqvist et al. (2021a) but were later discarded.

PSF subtraction method combination parameter

The RSM is one of the few algorithms that can combine multiple PSF-subtraction techniques simultaneously. This capability relies on the idea that different PSF-subtraction techniques produce residual noise with distinct statistical properties. As a result, combining several techniques creates a longer spatial-temporal series with increased variability in the noise regime, while the planetary signal remains temporally consistent. This allows the RSM algorithm to accumulate stronger evidence for the planetary regime over time. However, combining techniques that produce highly redundant residual noise can instead degrade performance, as repeated noise structures may artificially reinforce




Figure 3.3: Completeness variability for 15 different combinations of PSF-subtraction techniques based on APCA, NMF, LOCI, and LLSG. The curve corresponding to the optimal combination selected by the auto-RSM is shown in black.

speckle features in the probability map. Dahlqvist et al. (2021b) discussed the difficulty of identifying a single combination of methods that provides optimal performance in all cases. The selection of an optimal PSF-subtraction combination for a given ADI dataset was further addressed in Dahlqvist et al. (2021a).

We investigated the effect of this parameter using completeness curves for 15 different combinations of PSF-subtraction techniques based on APCA, NMF, LOCI, and LLSG (Fig. 3.3). The completeness curves show strong variability between individual PSF-subtraction techniques and combinations of multiple techniques, in agreement with the results presented in Dahlqvist et al. (2021b). These results confirm that the RSM strongly benefits from the diversity of residual noise structures produced by different PSF-subtraction techniques, and that its performance generally improves when several techniques are combined.

The curve corresponding to the optimal combination selected by the auto-RSM is also shown in black. However, this curve does not provide the best completeness limits in our analysis. This discrepancy is discussed Sect. 3.2.3.

All these parameters can be optimized using two different approaches: annular mode and full-frame mode. The annular mode consists of optimizing the parameters independently for each annulus, which accounts for the radial variation of the noise properties with separation from the star and for the corresponding changes in the optimal configuration

of the RSM algorithm. In contrast, the full-frame mode determines a single set of parameters valid for all separations. This is achieved by identifying the parameter set that minimises a normalised contrast metric over a selected range of separations, typically chosen to prioritise performance at small angular separations (see Sect. 3.4 in Dahlgvist et al., 2021a).

3.2.3 Parametrisation in the thesis

The parametrisation adopted in this thesis follows the reduced configuration introduced for the SHARDDS survey (Dahlgvist et al., 2022), in which only the parameters with the strongest impact on the RSM detection performance are retained and optimized. In particular, we use the full-frame optimization mode and restrict the analysis to the minimal set of parameters described in Sect. 3.2. The RSM maps are computed using the forward-backward formulation (Dahlgvist et al., 2021b), together with a Bayesian optimization framework.

Regarding the choice of PSF-subtraction techniques, the LLSG method was excluded from the analysis. The residual maps obtained with LLSG showed non-uniform behavior across separations, which translated into irregular 5σ contrast curves (outside RSM) and unstable optimization results. This lack of consistency propagates into the RSM parameter optimization and ultimately degrades its performance. The analysis therefore relies on APCA (Absil et al., 2013), LOCI (Lafrenière et al., 2007), NMF (Ren et al., 2018), and forward-model implementations of KLIP and LOCI (Pueyo, 2016; Dahlgvist et al., 2021b).

The RSM maps are evaluated over a radial range extending from $0''.11$ (9 pixels) to $1''.1$ (90 pixels). Forward-model techniques are restricted to separations smaller than 30 pixels ($\sim 0''.35$), where the impact of their negative lobes is most significant and where they provide the largest improvement over classical subtraction methods.

The parameter ranges adopted for the PSF-subtraction methods follow those recommended for SPHERE observations in Dahlgvist et al. (2021a). The only modification concerns the number of principal components used in PCA, which was restricted to the interval $[5, 25]$, since larger values produced excessive subtraction when applied across different datasets.

The optimization of the combination of PSF-subtraction techniques during the early stages of this thesis was performed using the particle swarm optimization framework. However, an unidentified limitation at that time led to suboptimal selections of the method combinations (see the black curve in Fig. 3.3). For this reason, we explored all possible combinations of PSF-subtraction techniques for all datasets in the F150 sample, and the configurations providing the best performance were identified empirically, as described in Chapter 4. Although the issue affecting the selection of the optimal combinations has since been resolved, we retained this exhaustive exploration in order to ensure a consistent choice of combinations across all datasets.

As for the optimization strategy adopted for the RSM parameters, we rely on the similarities in the noise properties observed between the different members of each cluster, as demonstrated in Sect. 2.3.3, together with the similarities in the optimal PSF-subtraction parameters reported in Sect. 4.4 of Dahlqvist et al. (2021a). Following the approach introduced in Dahlqvist et al. (2022), the optimization of the RSM parameters was therefore performed only for the cluster centers. The parameters selected for the cluster centers are listed in Table 3.1. We observe that the RSM consistently favors a crop size θ of approximately 1 FWHM for off-axis PSF models and ~ 1.5 FWHM for forward-model PSFs. Similarly, the multiplicative factor δ is consistently found to be close to $\delta = 5$. In contrast, the optimal choice of the noise-estimation region shows larger variability, although it has little impact on the final RSM performance, as illustrated in Fig. 3.2b.

These results indicate that the internal RSM parameters can be fixed without significantly affecting performance for SPHERE datasets, and that the variability required to account for differences in noise properties between ADI sequences is primarily driven by the PSF-subtraction parameters. This behavior is further illustrated in Appendix B.

3.3 Achieving better RSM performance

In this section, we explore several strategies that can be used to further improve the detection limits achieved with the RSM algorithm. The approaches presented here are computationally demanding and are therefore not adopted for the SHINE survey analysis. Instead, they are used to investigate the extent to which the performance of RSM can be improved under less restrictive computational constraints.

Table 3.1: Optimal parameters for PSF subtraction techniques and RSM for each cluster center.

Parameters	Center 1	Center 2	Center 3	Center 4	Center 5	Center 6
APCA components	20	14	18	8	7	12
APCA segments	4	3	3	3	3	4
APCA FOV rotation	0.42	0.29	0.26	0.31	0.5	0.76
NMF components	9	10	10	22	11	6
LOCI tolerance	0.0038	0.0016	0.0012	0.0057	0.0096	0.002
LOCI FOV rotation	0.49	0.14	0.48	0.1	0.13	0.73
FM KLIP components	9	19	8	12	5	5
FM KLIP FOV rotation	0.51	0.27	0.27	0.34	0.25	0.81
FM LOCI tolerance	0.0042	0.0074	0.0072	0.0073	0.0099	0.0016
FM LOCI FOV rotation	0.11	0.13	0.13	0.11	0.2	0.61
APCA δ	5	5	5	5	5	5
NMF δ	5	5	5	5	5	5
LOCI δ	5	5	5	5	5	5
FM KLIP δ	5	5	5	5	5	5
FM LOCI δ	5	5	5	5	5	5
APCA crop size	4	4	4	4	4	4
NMF crop size	4	4	4	4	4	4
LOCI crop size	4	4	4	4	4	4
FM KLIP crop size	7	7	7	7	7	7
FM LOCI crop size	7	7	7	7	7	7
APCA variance	SM	ST	SM	SM	SM	FR
NMF variance	SM	SM	SM	SM	SM	SM
LOCI variance	FR	ST	ST	SM	SM	ST
FM KLIP variance	FR	ST	ST	SM	ST	SM
FM LOCI variance	SM	SM	ST	ST	ST	ST

3.3.1 Combination of different parameter ranges

As shown previously, the RSM benefits from the use of multiple PSF-subtraction techniques, since the diversity of residual noise structures across different methods enriches the spatial-temporal series while allowing a real astrophysical signal to accumulate more coherently. Here, we explore a similar strategy by applying the same PSF-subtraction technique multiple times using different parameter ranges. For example, the same PSF-subtraction technique can be included twice in the RSM framework using two separate parameter intervals, such as $[5, 25]$ and $[26, 45]$ for the number of principal components, allowing the algorithm to optimize over a broader parameter space.

This effect is illustrated in Figure 3.4, which shows improved performance overall when employing two parameter ranges per technique instead of a single range. The curves shown in Figure 3.4 represent the highest contrast achieved at each separation among all tested combinations of PSF-subtraction techniques. This representation therefore illustrates the maximum performance that can be achieved with these configurations, although such performance may be difficult to reproduce systematically in practice.

3.3.2 Combination of different nights

Another way to improve the performance of the RSM framework is to combine observing sequences acquired over several consecutive nights. This approach increases the effective time series while ensuring that potential astrophysical signals do not move by more than one resolution element between observations. The benefits of this strategy are illustrated in Chapter 6, where the combination of approximately three consecutive high-quality observing nights in the search for ϵ Eridani leads to an improvement in contrast performance by factors of two to five compared to individual observations processed with RSM.

However, it was also observed that combining nights with significantly different observational conditions can deteriorate the overall performance of the RSM analysis. In particular, the inclusion of poorer quality observations may reduce the final performance to a level comparable to that of the poorer dataset. This situation occurs in several SHINE observations, where consecutive observing nights often correspond to datasets that were re-observed due to unfavorable observing conditions. The presence of such low-quality datasets motivated the selection procedure presented in Section 2.1.2, where only observations satisfying specific quality criteria were retained for the analysis.




Figure 3.4: Completeness curves corresponding to the highest contrast achieved at each separation for different combinations of PSF-subtraction techniques. The blue curve corresponds to the best-performing combinations using APCA, NMF, and LOCI with a single parameter range. The orange curve includes the addition of forward-model approaches. The green curve corresponds to the case where all techniques are applied twice using two distinct parameter ranges for each parameter.

Chapter 4

Detection threshold in RSM

Abstract. The detection maps produced by the RSM framework span values ranging from 0 to 1. While higher values naturally indicate higher confidence levels, the RSM framework suffers from class imbalance effects where lower values may still correspond to significant detections. This issue is common among several advanced detection techniques lacking well-defined detection thresholds. In this chapter, detection thresholds are investigated using two approaches: statistical distribution fitting of the RSM maps within each observational group and angular separation, and threshold estimation through the maximization of the F1-score without explicit assumptions on the underlying noise distribution, an approach that can also be generalized to other detection techniques.

This chapter is organized as follows. Section 4.1 describes the statistical behavior of the noise in the RSM framework, including a first threshold estimation method based on log-normal distribution fitting in Sect. 4.1.1, and then threshold estimation through F1-score maximization in Sect. 4.1.2, and the comparison and reliability assessment of the two approaches in Sect. 4.1.3. Section 4.2 investigates the resulting sensitivity limits, including the influence of the threshold choice and PSF-subtraction technique in Sect. 4.2.1, the optimal combination of PSF-subtraction methods within the RSM framework in Sect. 4.2.2, and the reliability of the resulting contrast curves in Sect. 4.2.3. These sections are directly extracted from Sabalbal et al. (2026).

4.1 Noise behavior in RSM

In HCI, detection limits are commonly set at a 5σ level, corresponding to a false alarm probability (FAP) of 3×10^{-7} based on a Gaussian noise assumption in processed frames (Mawet et al., 2014). However, this approach may not accurately reflect the true noise characteristics in HCI datasets. Alternative approaches based on, e.g., Laplacian distributions (Pairet et al., 2019; Dahlgvist et al., 2020) have been proposed to address this limitation, especially at small separations. Moreover, methods such as those in Jensen-Clem et al. (2018), Bonse et al. (2023), and Daglayan et al. (2024) avoid assuming any specific noise model, instead using adaptive techniques to derive thresholds directly from data, reducing the reliance on potentially mismatched distributions.

In the context of RSM, fitting a theoretical distribution to the noise in the final RSM maps (hereafter referred to as RSM noise) is challenging due to the nonlinear dependency of the final RSM score on the likelihood functions and transition probabilities, leading to the distortion of any pre-assumed distribution. Dahlgvist et al. (2021b) defines the detection threshold for the contrast curve as the brightest FP in the final detection map. As noted in Appendix C, this threshold is broadly consistent with those defined in this section. However, it lacks robustness in the presence of astrophysical signals and is not designed to detect them. An alternative approach explored by Dahlgvist et al. (2022), based on reversed parallactic angles, was also considered as a potential threshold applicable to both companion detection and contrast curves. However, it posed challenges by revealing a high number of false positives, ultimately limiting its utility as a reliable detection criterion.

In this study, we investigate a new, consistent threshold applicable both for generating contrast curves and identifying candidate companions in RSM. Here, we examine two methods for defining a detection threshold, either (i) by fitting an empirical distribution to independent RSM noise realizations in the final RSM map (Sect. 4.1.1), or (ii) by balancing false positives and true positives while maximizing the F1 score, without assuming a specific RSM noise distribution (Sect. 4.1.2). We also explore how varying observing conditions influence the RSM noise distribution and the resulting threshold. Finally, we evaluate the strengths and limitations of each threshold within the context of RSM.




Figure 4.1: Quantile-quantile plot of the RSM noise in Cluster 1 at $0.25''$ (left) and $1''$ (right) using APCA. The plot compares the RSM noise distribution to the lognormal distribution, excluding values above the 3×10^{-7} FAP of the fit. The coefficient of determination (R^2) is also displayed, indicating a good fit.

4.1.1 Fitting a lognormal distribution to the RSM noise

In this method, we take advantage of the large amount of data in the SHINE F150 survey to build noise distribution samples at each angular separation. Following Bonse et al. (2023), we treat independent pixel realizations, i.e., the central pixels of non-overlapping resolution elements, as distribution samples. Noise analysis at a given separation across the entire survey shows multiple RSM noise distributions, leading us to analyze realizations per cluster and per combination of PSF subtraction techniques in RSM, at each considered angular separation.

The final RSM maps contain only positive values, leading to a noise distribution that is not centered on zero and exhibits positive skewness. After testing multiple distributions as shown in Appendix D, we selected the lognormal distribution for its minimal parameter requirements and its consistency with the shape of the RSM noise distribution across separations and clusters. To define a detection threshold, we set a 3×10^{-7} FAP for the lognormal distribution, aligning with the conventional 5σ threshold in Gaussian statistics. To prevent bias when fitting the distribution, we exclude scores above 0.1 from the distribution samples, as these correspond to bright signals.

Figure 4.1 presents quantile-quantile (Q-Q) plots for cluster 1 at small (left) and large (right) separations using RSM with APCA to these plots. We compare the logarithm of independent sample values from all observations in cluster 1 at $0.25''$ and $1''$ with the corresponding fitted distribution values, considering only those below the 3×10^{-7} threshold. The Q-Q plots exhibit an approximately linear trend in both cases, indicating a

strong fit, even though the extremities of the distribution show some outliers as it is often the case in Q-Q plots. Moreover, the coefficient of determination (R^2), representing the Pearson correlation between quantiles, is close to 1, further confirming the goodness of fit.


In Figure 4.2, we show how varying observing conditions impact the RSM noise distribution and, consequently, the chosen detection threshold. Figure 4.2a presents the lognormal fits for different clusters at $0''.25$ using RSM with APCA. Cluster 3, featuring a large number of frames per target, shows a narrower distribution and noise tail, while cluster 4, with mostly poor-quality observations, shows a broader tail and wider distribution, indicating a higher detection threshold. This plot highlights the varying RSM noise behaviors across observations under different conditions. Figure 4.2b illustrates the variations of the 3×10^{-7} FAP threshold across clusters and separations, highlighting the advantage of setting separate thresholds per separation to accommodate distinct noise properties. It also emphasizes the benefit of adjusting thresholds for observations under varying conditions. This approach, consistent with predictions, results in lower thresholds for Cluster 3 and higher ones for Cluster 4 and 5. This analysis underscores the importance of considering observing conditions when setting thresholds, as applying a single threshold across all datasets could lead to significant underestimations or overestimations of the detection thresholds, as evidenced by the almost two orders of magnitude difference in thresholds between the best and worst clusters.

4.1.2 Threshold using maximum F1 score


In this section, we explore the possibility of defining a detection threshold without assuming a specific RSM noise distribution, an approach particularly relevant for samples containing a limited number of observations, where robust noise histograms cannot be built, especially for small angular separations. Following Daglayan et al. (2024), we define detection thresholds by empirically maximizing the F1 score at each angular separation through fake companion injections. The F1 score is defined as follows:

$$\text{F1 score} = \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}}, \quad (4.1)$$

with TP the true positives, FP the false positives, and FN the false negatives. Maximizing the F1 score ensures a good balance between true positive rate and false positive rate.



(a)



(b)

Figure 4.2: *Top.* Lognormal distributions obtained by fitting RSM noise histograms at $0''.25$ for various observing conditions (clusters) using RSM with APCA. *Bottom.* Detection thresholds derived from the lognormal distributions at 3×10^{-7} FAP, as a function of angular separation and of observing conditions.




Figure 4.3: Variation of the TPR, FPR, and F1 score for the cluster 1 center at $0''.7$. The dashed purple line marks the threshold that maximizes the fitted F1 score curve (dark blue curve), balancing high true positive (orange curve) recovery with minimal false positives (dark red curve).

To compute the false positives, all independent pixel realizations at each separation are considered as devoid of circumstellar signal (pure noise). Varying the threshold, we calculate the number of aperture values exceeding the threshold, which represent the number of false positives in the considered annulus. The false positive rate (FPR) is then defined as the ratio of these false positives to the total number of independent realizations. For true positives, we follow Mawet et al. (2014) by first estimating noise levels in the residual cube produced by APCA. We then inject synthetic point sources with fluxes between two and three σ of the measured noise. An injected source is considered recovered if its RSM score surpasses the considered detection threshold, and the number of such recoveries within the annulus defines the true positives. The optimal threshold is determined by maximizing the F1 score, ensuring a balance between high true positive recovery and minimal false positives.

Figure 4.3 shows the evolution of TPR, FPR, and the F1 score for varying thresholds, applied to the cluster 1 center (corresponding to the BD-15 705 dataset) using RSM with APCA at $0''.7$ separation. The figure highlights the fact that the optimal F1 score favors minimal FPR and high TPR, thereby effectively distinguishing noise from potential true




Figure 4.4: Comparison of detection thresholds with false positives and injected companions for the cluster 1 center at an angular separation of $0''.6$ and showing the RSM scores of injected fake companions (gold dots), false positives (red crosses), the maximum F1 score threshold (dashed blue line), and the lognormal threshold (solid blue line).

positives. This metric is highly dataset-specific as it relies on the false positive rate and the true positive rate, the latter also depending on the flux of injected signals, which is influenced by the noise distribution. Consequently, this approach captures the unique characteristics of each dataset and aligns the threshold closely with the specific noise properties of the data. This technique, however, has several limitations. First, while it minimizes the false positive rate, it does not explicitly control the acceptable number of false positives per annulus. Additionally, it assumes all independent realizations are false positives, which may not hold if a real signal is present. In such cases, the maximum F1 score metric tends to exclude the true signal by minimizing false positives. One solution is to remove bright signals by injecting negative fake companions to avoid bias. Finally, this approach requires noise computation, fake companion injections, and a complete RSM processing for each separation and for each fake companion injection, making it computationally expensive.

4.1.3 Comparison and reliability of the detection thresholds

To evaluate the reliability of detection thresholds, we examine their proximity from false positives for the cluster center of cluster 1. Figure 4.4 compares the RSM scores of fake companions injected at an angular separation of $0''.6$, across flux levels ranging from 0 to 8σ , with the false positives detected at that separation and the two associated detection

thresholds. For each injection, we computed the number of independent apertures within the considered annulus whose mean values exceeded the recovered flux, treating them as false positives. Both the maximum F1 score threshold and the lognormal threshold lie above the false positives, with the F1-score threshold closely aligning with the 2.5σ injected fake companion, near the noise-dominated region. The lognormal threshold, slightly higher at approximately 3σ , indicates a more conservative detection criterion. Figure 4.5 shows the median detection thresholds across all clusters, with each lognormal threshold computed individually for each cluster considering all its members, and the maximum F1 score thresholds computed for the cluster center datasets. The dip in thresholds below $0''.4$ in the bottom-right panel of Fig. 4.5 corresponds to the influence of forward-model techniques (FM KLIP and FM LOCI), which are active within this separation range. This comparison highlights the differences between the two methods, with the lognormal threshold being generally more conservative. The larger variability in the detection threshold derived from the maximum F1 score across separations reflects its sensitivity to outliers in different annuli.

While it arguably provides the most relevant thresholds for a given dataset, the maximum F1 score is computationally demanding to apply on each dataset of the survey, and too specific to generalize across an entire cluster (see more about that in Appendix E). Conversely, the detection thresholds derived from a lognormal distribution are more conservative, computationally efficient, and reflect the noise behavior under similar observing conditions. However, they rely on a predefined distribution, large sample sizes, and do not capture dataset-specific nuances. For the present study, focusing on a large survey, we adopt the threshold derived from the 3×10^{-7} FAP of the lognormal distribution, and derive it at each separation, for each combination of PSF subtraction technique, and for each cluster. We recommend using the F1 score threshold for smaller samples, while noting its limitations described in Sect. 4.1.2.




Figure 4.5: Median detection thresholds for the six different clusters using the 3×10^{-7} FAP under lognormal distribution (orange), compared with the threshold derived from the maximum F1 score on the cluster centers (blue), for different combinations of PSF subtraction techniques as described in the plot titles.

4.2 Sensitivity limits

In this study, we adopt the 50% completeness curve to assess the sensitivity limits to point-like sources. Such curves indicate the flux at which 50% of injected signals are detected above the selected threshold, here the 3×10^{-7} FAP from the lognormal distribution, which corresponds to the 5σ contrast curve when assuming Gaussian noise. Here, we analyze the effects of threshold selection and PSF subtraction technique combinations on contrast curves, define the optimal RSM contrast curve, and evaluate its reliability.




Figure 4.6: Contrast comparison at 50% completeness using thresholds based on the maximum F1 score (blue), the lognormal RSM noise assumption (orange), and the 5σ APCA contrast curve (black) across different PSF subtraction techniques applied to the center of cluster 1.

4.2.1 Influence of threshold and PSF subtraction techniques

To evaluate how the chosen threshold impacts contrast curves, Fig. 4.6 presents a comparison for the center of cluster 1. It includes contrast curves derived using the maximum F1 score threshold (blue) and the lognormal threshold (orange) across four combinations of PSF subtraction techniques: APCA, APCA-NMF, APCA-NMF-LOCI, and APCA-NMF-LOCI-FMKLIP-FMLOCI. A standard 5σ contrast curve based on full-frame PCA with five principal components is shown in black for comparison. The RSM contrast curves show a good level of consistency between the two ways to define the detection threshold. This is especially the case when multi-technique combinations are used (see bottom panels in Fig. 4.6), while the lognormal threshold leads to slightly worse sensitivity than the F1-score one when RSM is performed with APCA only (Fig. 4.6 top left) due to




Figure 4.7: Effect of different combinations of PSF subtraction techniques in RSM on the fitted lognormal distributions, detection thresholds, and contrast curves. *Top left:* Difference of lognormal fit across different PSF subtraction techniques with colors ranging from dark violet for multi combinations to yellow for singular combinations. *Top right:* Distribution of the 3×10^{-7} FAP under lognormal distribution for the different combinations across all separations. *Bottom left:* Contrast curve delivered by the different combinations with 50% completeness.


the higher threshold in that case. Consistency was also found between the lognormal contrast curves derived here and the contrast curves based on the appearance of a first false positive proposed in Dahlqvist et al. (2022) for the SHARDDS survey (see Appendix C).

The choice of which PSF subtraction techniques to combine is an important aspect of the RSM algorithm, which is not straightforward to optimize. Instead of relying on a single combination identified as optimal by the RSM framework, we propose to systematically evaluate the performance of various combinations of PSF subtraction techniques. We analyze the differences in noise distributions in the final frames, the corresponding thresholds, and the contrast curves achievable with each combination. Figure 4.7 examines these effects across the 28 possible combinations of PSF subtraction techniques,

ranging from single PSF subtraction techniques (in yellow) to multi-combinations (in dark violet). Figure 4.7 (top left) displays the median lognormal fits for each combination, clearly showing narrower noise distributions for multi-combinations and broader distributions for single techniques, indicating improved performance with combined techniques. Figure 4.7 (top right) presents the thresholds derived from these fits across all combinations and separations. Noise thresholds, on average, range from approximately 10^{-4} for multi-combinations to nearly ten for LOCI – exceeding RSM’s maximum score. This renders LOCI non-physical and unreliable when used alone, a consequence of its high false positive rate in RSM for many datasets. Figure 4.7 (bottom left) illustrates the contrast curves for the various combinations, showing an improvement by a factor of three for multi-combinations. Contrast values with lognormal thresholds exceeding 0.1 in RSM were excluded from the bottom-left panel due to difficulties in achieving convergence. This further demonstrates that incorporating multiple PSF subtraction techniques in RSM significantly enhances detection limits. Moreover, it highlights RSM’s ability to leverage the diverse noise patterns of different PSF subtraction methods, effectively reducing the misclassification of speckles as planetary signals. RSM’s efficiency in integrating additional information – either by combining different PSF subtraction techniques or by increasing the number of frames – is evident through narrower noise distributions, improved thresholds, and enhanced detection limits (see also Figure 4.2a).

4.2.2 Optimal usage of RSM

Although combining multiple techniques yields the best RSM performance, the optimal contrast at each separation does not always result from the same combination. Figure 4.8 presents a histogram illustrating the contributions of various RSM combinations to the best contrast values across all separations, based on all datasets. We identify the best contrast at each separation, shown as the red curve in Fig. 4.9. Figure 4.8 indicates that approximately 26 combinations out of the 28 possible combinations of PSF subtraction techniques contribute to the red curve of Fig. 4.9, with two combinations standing out due to their significantly higher contributions compared to others: APCA-NMF-LOCI-FMKLIP, and APCA-LOCI-FMKLIP. Since tracking candidates across 26 different RSM maps is complex, we opted to use only the two combinations with the highest contributions. The final contrast curve was then computed by taking the median of their contrast



4.2.3 Reliability of the contrast curves

Evaluating the adopted contrast curve metric through fake companion injections helps verify the consistency of its core assumptions, including the applied noise threshold, the definition of the median contrast curve, and the expected 50% recovery rate of injected signals. We inject fake point sources at the predicted sensitivity limit using the VIP package on a cluster center and examine their recovery across different RSM maps. Figure 4.10 compares the recovery of injected signals using RSM maps against the signal-to-noise ratio (S/N) map produced by PCA with five principal components, for the center of cluster 1. The flux levels of the injected companions have been chosen to be just above the expected RSM sensitivity limits as illustrated in Fig. 4.11, which also compares the detection limits of RSM with the full-frame PCA 5σ contrast curve. The injections are performed at multiple separations, with one injection at each separation and varying position angles.




Figure 4.9: Optimal contrast curve achieved by different combinations of PSF subtraction techniques in RSM. Gray curves show the median contrast performance for each combination, computed across all datasets in cluster 1. The red curve highlights the best achievable performance at each separation, given by the median of the minimum contrast values across all combinations. The black curve shows the median of the maximum contrast values delivered by the two best-performing combinations: APCA-NMF-LOCI-FMKLIP and APCA-LOCI-FMKLIP.

As expected, a significant fraction of the injections are successfully detected above the threshold in RSM maps, achieving approximately 50% recovery consistent with the definition of the contrast curve. Conversely, the same injections all appear well below the 5σ level in the S/N PCA map. This highlights RSM's ability to deliver deeper detection limits and reliable contrast curves, surpassing conventional post-processing methods in HCI.




Figure 4.10: Comparison of recovered fake point sources in cluster center 1 using RSM maps with different PSF subtraction technique combinations: (a) APCA-NMF-LOCI-FMKLIP and (b) APCA-LOCI-FMKLIP, compared to the signal-to-noise ratio PCA map with five principal components (c). The RSM maps shown in this figure are thresholded (RSM map minus the separation-dependent threshold). Injected signals recovered above the detection threshold are marked with yellow squares, while unrecovered signals appear in red squares.




Figure 4.11: Injected fake point sources compared to the optimal RSM contrast curve (black) and PCA five-component contrast curve (red) for cluster center 1

Chapter 5

RSM application to SHINE F150 sample

Abstract. Following the optimization of the RSM framework and the calibration of detection thresholds, this chapter presents its application to the SHINE F150 sample. The analysis first quantifies the improvements in detection performance achieved with RSM through comparisons with the classical PCA reductions used in Langlois et al. (2021). The recovered signals are then compared with the PACO detections, benefiting from the larger PACO analysis of the full SHINE survey and the availability of SnapSHINE follow-up observations. Particular attention is given to signals detected only with RSM, for which we investigate their reliability and possible nature using follow-up observations when available, color–magnitude information, wavelength-scaling behavior, and estimates of the local background contamination. These analyses allow us to identify promising companion candidates while quantifying the false-positive rate associated with the RSM framework.

In this chapter, we present the contrast improvements achieved for the SHINE F150 sample following the application of the RSM method in Section 5.1. Section 5.2 presents the detected RSM signals, compares them with previous SHINE results obtained using other algorithms (Section 5.2.1), and describes the identification of new point sources (Section 5.2.2), together with an assessment of the reliability of these detections. Section 5.3 focuses on the investigation of the candidate companion around π Ara. Finally, Section 5.4 discusses the overall performance of RSM within the SHINE survey. Sections 5.1 and 5.2 (only 5.2.1) are based on results presented in Sabalbal et al. (2026) while the rest of the sections include updated analyses and new results.




Figure 5.1: Comparison of RSM and full-frame PCA median contrast curves for the F150 SHINE sample (dashed bold lines). The illustration provides a visualization of all RSM contrast curves (shaded blue) and PCA contrast curves (shaded pink), highlighting RSM's significant improvement in detection limits at both small and large separations.

5.1 Contrast limits improvements

In this section, we present the detection limits for the SHINE F150 sample derived from 213 observations using the lognormal fitting for the detection threshold and the 50% completeness curve described in Chapter 4. The F150 sample was originally reduced using classical post-processing techniques such as PCA and TLOCI in Langlois et al. (2021), while the complete SHINE survey was subsequently reprocessed with the advanced PACO algorithm (Chomez et al., 2025). A review of these previous results is provided in Sect. 1.5.2.

Figure 5.1 displays the envelope of all RSM contrast curves (shaded blue) together with full-frame PCA-ADI curves (shaded pink) computed with five principal components, as described in Langlois et al. (2021). For consistency within the Langlois et al. (2021) analysis, we do not use a high-pass filter at large separations, although we acknowledge that it could deliver better results. The PCA-ADI curves may show slight differences with those published in Langlois et al. (2021), reflecting variations in the data reduction process. Nevertheless, our PCA-ADI results align well with those reported in Chomez et al. (2023), as both are based on the same reduction pipeline. The PCA contrast

curves display outliers due to bright companions and background star signals. The median contrast curves indicate RSM's substantial improvement, achieving a factor of two enhancement at 800 mas and a factor of four to five at 135 mas compared to the five-component PCA results.

This improvement is further highlighted in Fig. 5.2, where the magnitude difference between RSM and full-frame PCA sensitivity limits shows a median improvement of 1.8 mag at $0.''135$, emphasizing RSM's superior detection at small separations and its ability to distinguish planetary signals amidst stellar residual noise. The improvement closely aligns with that achieved using PACO (Chomez et al., 2023), with both methods delivering an improvement of approximately a factor of five at small separations and a factor of about two overall, compared to full-frame PCA. These results emphasize the consistent advancements in sensitivity limits across different post-processing techniques. Figure 5.2 shows that a few datasets exhibit slightly better performance with PCA than with RSM. This occurs primarily for observations located near the boundaries of their respective clusters, where the adopted RSM parameters may deviate from their true local optima, or in cases where the generalized noise model across the cluster is more restrictive than the dataset's actual RSM noise properties. These exceptions are rare and do not affect the overall improvement trends observed across the survey, although they do highlight certain limitations of parameter and noise generalization for a small subset of datasets.

This study also allows us to evaluate RSM's performance under varying observational conditions. Figure 5.3 highlights the improvements across distinct clusters in the SHINE survey, demonstrating consistent enhancements across different datasets. This underscores RSM's robustness in handling diverse observational scenarios and its ability to improve point source detection across the entire survey.

5.2 Signals in RSM maps

The application of the RSM algorithm to the SHINE F150 sample resulted in the identification of 90 signals across 60 different observations, while 153 observations from the F150 sample showed no detectable signal within the $1.''1$ FOV. A signal is considered significant when it lies above the detection threshold in all four RSM detection maps, considering the two selected algorithm combinations and the parallel observations performed by SPHERE-IRDIS in the H2 and H3 filters. A dedicated tool was developed to




Figure 5.2: Comparison of RSM and full-frame PCA improvements in magnitude at separations of $0''.135$, $0''.2$, $0''.4$, and $0''.8$. The median difference is shown as a dashed line, while histograms represent the full dataset.

cross-match detections between maps within half a FWHM, ensuring consistent identification of point sources across filter combinations. The photometric and astrometric measurements were derived using the negative fake companion algorithm (NEGFC, Wertz et al., 2017) implemented in the VIP package.

Among the detected signals, 49 were already recovered by Chomez et al. (2025), while 41 were not identified with PACO and required a dedicated analysis to distinguish genuine astrophysical sources from false positives. This counting of false positives does not include some artifacts that tend to appear in RSM maps in the presence of bright companions, at the same angular distance as those companions. This section presents a detailed comparison between the signals detected with RSM and PACO in Sect. 5.2.1, followed by an analysis of the new point sources recovered exclusively with RSM in Sect. 5.2.2 using proper motion and color-magnitude diagram tests.




Figure 5.3: Contrast performance improvements of RSM across various observing conditions in the SHINE survey clusters: median contrast curves are shown in red for RSM and black for PCA, demonstrating RSM's superior efficiency overall.

5.2.1 Comparison of SHINE survey findings with other algorithms

Given the comparable performance of RSM and PACO, a direct comparison of their detected signals provides valuable insight into the relative sensitivity of both methods, the reliability of faint detections, and potential missed detections on either side. To ensure a fair comparison between PACO and RSM, only the common observations in the archival point-source detections reported in Chomez et al. (2025) were considered. These correspond to common H2/H3-band observations covering fields of view between $0''.11$ and $1''.1$. Within this sample, PACO reported no detection in 181 observations. Among these, 31 revealed a detection in RSM, while 150 are classified as non-detections in both methods. These 31 observations contain 41 signal detections that are examined in more detail in Sect. 5.2.2. In addition to these 181 observations classified as non-detections, PACO recovered 52 signals in the remaining 32 observations, including 24 confirmed planets (observed at different epochs), 24 background stars, and four ambiguous detections. Here, we compare these 52 PACO detections with our RSM results.




Figure 5.4: Contrast distribution and separation of the 52 PACO-detected signals and their recovery by RSM. Dark blue points represent signals recovered in all four RSM maps, other colors correspond to sources recovered in only some RSM maps.

Out of the 52 PACO detections, 49 were also recovered by RSM based on our four-map detection criterion. Figure 5.4 presents the contrast values of each PACO-detected signal, with colors indicating the recovery status in RSM; dark blue points represent signals fully recovered in four maps, while other colors correspond to unrecovered sources. As shown in the figure, most signals detected by PACO were also recovered by RSM, with only a few exceptions.

The unrecovered signals correspond to the confirmed planet HD 95086 b (in two observations), and a background star around HD 151726, although these sources were recovered in at least one of the RSM maps. The primary reason for this discrepancy lies in the different detection criteria employed by the two algorithms. The RSM method requires recovery above the detection threshold in both H2 and H3 filters, whereas PACO derives a single S/N that combines information from both spectral channels. Consequently, sources with marginal S/N in one band but a stronger signal in the other may be recovered by PACO but not by RSM.

This behavior is clearly illustrated by the case of HD 95086 b, observed on 2015-05-05 and 2015-05-11, with PACO S/N values of 6.8 and 5.0, respectively. Desgrange et al. (2022) showed that the planet-star contrast for this target is twice as large in H3 compared to H2 (see their Fig. 5), consistent with the Chomez et al. (2025) database, where the H2 contrast errors are nearly as large as the measured contrast values themselves.

5.2.2 Identification of new point sources

The application of the RSM algorithm to the F150 SHINE sample revealed 40 newly detected point sources across 31 different observations for 29 stars. Among these targets, ten stars have additional epochs within the F150 sample, while the remaining 19 correspond to single-epoch observations.

Multiple epochs

Multi-epoch data enable the differentiation between false positives and genuine astrophysical sources through proper motion tests. Out of the ten targets that have additional epochs, two exhibit detections in more than one epoch, while the remaining eight show detections in only a single epoch despite the availability of multiple observations. Regarding the multi-epoch detections, RSM successfully recovered a candidate companion around PZ Tel in an observation from 2014, which was not detected by PACO, unlike other epochs where the signal was identified by both methods. RSM also recovered a background star around HD 164249 identified through the proper motion test between the 2015 and 2016 epochs. These signals were not reported by Chomez et al. (2025) but were previously mentioned by Langlois et al. (2021).

For detections found in only one epoch among several available observations, both the expected position of the source within the FOV and the contrast curves of the non-detection epochs were examined. This analysis is illustrated in Figure 5.5, where we show the contrast curves for all multi-epoch targets except PZ Tel. The figure includes the contrast curve of the detection epoch together with those of the other epochs, the measured contrast of the detected point source, and the expected position of the detection assuming background-star motion. By comparing these positions with the corresponding contrast curves, we assess whether the source should have been detectable in the other epochs.

We find that detections around HD 164249 (2015 epoch), ϵ Hyi, CD-31 16041, HD 104125, and ϵ Pav do not present corresponding signals in the FOV of the other observations, despite their generally better contrast performance. Assuming a background-star hypothesis, the expected positions of these signals would lie above the contrast limits of the other epochs, implying that they should have been detected if they were real astrophysical sources. We therefore classify these detections as false positives, even though the hypothesis of a significant orbital motion between epochs cannot be formally ruled out.

For the remaining targets, the non-detections can be explained because the expected position of the source falls either at separations with poorer contrast (e.g., ASAS J205428+0906.1, V* CC Phe, and CD-51 10924) or within the inner working angle, as in the case of HD 128819. We therefore retain these signals for further analysis.

Individual epochs

Verifying the nature of a signal for single-epoch targets is more challenging. In these cases, following Chomez et al. (2025), a CMD analysis (introduced previously in Chapter 1.5) is used to assess whether each detection is more likely to be consistent with a planetary companion or a background star. The 19 targets that had no additional epochs within the SHINE F150 survey generated a total of 26 significant signals in our RSM maps. To these 26 signals, we add the 4 ambiguous candidates identified in the previous paragraph, and perform a CMD analysis to evaluate the astrophysical relevance of each signal.

The CMD is shown in Fig. 5.6 and includes all detections reported in Chomez et al. (2025), supplemented by the new RSM candidates. The corresponding astrometric and photometric measurements of these candidates are summarized in Table 5.1.

As discussed by Chomez et al. (2025), background stars occupy a well-defined region in the diagram. We therefore adopt their background-star region, referred to as the exclusion area in Fig. 5.6. Many of the RSM candidates fall within this exclusion region, suggesting that they are likely background stars. Most of the remaining signals lie near the $\Delta\text{mag} \sim 0$ region, with two signals (π Ara and CD-51 10924) located in the methane-rich planetary region of the CMD.




Figure 5.5: Contrast curves for observations with at least one detected signal, overlaid with the measured contrasts of the detections. Circles indicate sources recovered with RSM, while triangles show the expected location of the source assuming a background object with constant contrast. The gray dashed area indicates the inner working angle area.




Figure 5.6: *Updated* Color-magnitude diagram showing the new RSM detections in the F150 SHINE sample, plotted alongside the signals reported in Chomez et al. (2025) for confirmed exoplanets, background stars identified through proper motion, and background stars classified via CMD analysis. The gray shaded area indicates the exclusion region for background stars defined by Chomez et al. (2025). This figure is an updated version of the CMD presented by Sabalbal et al. (2026), using photometry derived with the NEGFC technique in this version instead of the RSM characterization tool.

Table 5.1: Astrometric and photometric measurements for the 30 identified RSM candidates. The table presents the SIMBAD target name, the observing date, the separation, the position angle, and the contrast and magnitudes measured in the H2 and H3 bands for each detection.

Star	Date	Band	Separation [mas]	Position angle [°]	Contrast	Δmag
* 30 Ser	2015-06-05	H2	723.75	214.28	6.4×10^{-6}	12.98
		H3	759.61	214.81	5.3×10^{-6}	13.19
* alf Sgr	2018-05-08	H2	622.29	8.15	1.3×10^{-5}	12.25
		H3	655.43	7.84	1×10^{-5}	12.47
* alf Sgr	2018-05-08	H2	658.48	352.56	5.7×10^{-6}	13.11
		H3	658.48	352.56	5.7×10^{-6}	13.11
* mu. PsA	2016-10-14	H2	416.86	196.99	2.8×10^{-6}	13.87
		H3	438.81	196.69	2.7×10^{-6}	13.92
* pi. Ara	2016-06-11	H2	883.07	213.92	3.4×10^{-6}	13.67
		H3	939.33	215.35	4.7×10^{-6}	13.31
* pi. Ara	2016-06-11	H2	311.09	194.18	1.2×10^{-5}	12.26
		H3	334.16	194.78	1.2×10^{-5}	12.26
* pi. Ara	2016-06-11	H2	895.37	114.38	7.6×10^{-6}	12.79
		H3	925.49	113.59	1.9×10^{-6}	14.31
* pi. Ara	2016-06-11	H2	894.59	356.14	4.2×10^{-6}	13.43
		H3	934.61	357.10	3.6×10^{-6}	13.60
ASAS J193304+0345.7	2018-06-19	H2	147.57	355.47	1×10^{-2}	4.97
		H3	143.82	366.64	9.7×10^{-3}	5.03
ASAS J205428+0906.1	2018-05-13	H2	72.04	-0.00	1.4×10^{-4}	9.60
		H3	102.34	0.00	9×10^{-5}	10.12
CD-51 10924	2016-05-31	H2	681.39	325.04	3.2×10^{-6}	13.73
		H3	735.43	324.70	7.9×10^{-7}	15.25
CD-57 1054	2016-01-03	H2	898.40	188.28	1×10^{-5}	12.46
		H3	927.36	187.76	7×10^{-6}	12.89
CD-57 1054	2016-01-03	H2	901.07	217.38	5.4×10^{-6}	13.17
		H3	943.31	216.95	3.8×10^{-6}	13.55

Continued

Table 5.1 – continued

Star	Date	Band	Separation [mas]	Position angle [°]	Contrast	Δmag
CD-81 812	2016-06-28	H2	816.15	153.99	5.5×10^{-6}	13.14
		H3	845.52	153.84	3×10^{-6}	13.80
HD 987	2016-10-15	H2	341.64	312.36	1.9×10^{-5}	11.81
		H3	365.96	311.30	1.2×10^{-5}	12.30
HD 1466	2015-10-26	H2	746.56	187.82	7.3×10^{-6}	12.85
		H3	678.41	189.21	8.7×10^{-6}	12.65
HD 1466	2015-10-26	H2	419.59	10.36	1.8×10^{-5}	11.84
		H3	433.50	9.37	2.5×10^{-5}	11.50
HD 74576	2016-03-28	H2	788.32	122.63	1.4×10^{-6}	14.61
		H3	825.14	123.26	1.2×10^{-6}	14.80
HD 101615	2016-06-02	H2	158.45	44.93	4.4×10^{-5}	10.90
		H3	167.21	41.23	1.8×10^{-5}	11.86
HD 113457	2016-04-17	H2	1127.42	359.04	8.7×10^{-5}	10.15
		H3	1127.64	359.08	9.2×10^{-5}	10.09
HD 118100	2016-06-27	H2	1025.69	18.20	3.1×10^{-6}	13.76
		H3	1072.80	18.31	2.9×10^{-6}	13.86
HD 128819	2016-04-17	H2	111.25	276.43	1.1×10^{-3}	7.43
		H3	109.25	278.07	1.1×10^{-3}	7.41
HD 152391	2015-06-05	H2	954.08	68.38	1.6×10^{-6}	14.49
		H3	1004.81	68.33	1.3×10^{-6}	14.68
HD 317617	2015-05-31	H2	658.94	145.62	1.9×10^{-5}	11.80
		H3	659.81	145.40	1.9×10^{-5}	11.81
PM J19560-3207	2015-05-04	H2	597.39	278.88	6.6×10^{-6}	12.95
		H3	609.88	278.32	7.4×10^{-6}	12.82
PM J19560-3207	2015-05-04	H2	561.02	103.42	6.6×10^{-6}	12.95
		H3	561.02	103.42	6.6×10^{-6}	12.95
Smethells 20	2015-05-16	H2	477.27	245.03	1.3×10^{-5}	12.22
		H3	486.96	244.61	9.9×10^{-6}	12.51

Continued

Table 5.1 – continued

Star	Date	Band	Separation [mas]	Position angle [°]	Contrast	Δmag
V* CC Phe	2017-09-29	H2	430.40	355.54	4×10^{-6}	13.49
		H3	468.25	354.94	1.7×10^{-6}	14.45
V* KU Lib	2016-06-01	H2	642.27	35.82	1.7×10^{-6}	14.41
		H3	676.64	36.14	2.3×10^{-6}	14.08
V* V1005 Ori	2015-09-28	H2	1000.88	120.14	4.7×10^{-6}	13.33
		H3	1053.18	120.17	4.2×10^{-6}	13.45

Reliability of RSM signals

The photometry of the RSM signals does not show strong discrepancies between the H2 and H3 magnitudes that would allow an automatic classification as either real astrophysical objects or false positives. Instead, the photometric properties of most signals are broadly compatible with both a background-star and an exoplanet interpretation, except for a few cases (e.g., HD 101615). However, given the similar performance of RSM and PACO, we do not expect such a large number of new candidates that were not reported by (Chomez et al., 2025). We therefore explore additional approaches to assess whether these detections correspond to real astrophysical sources.

We use the Besançon Galactic model (Robin & Creze, 1986) to infer, based on the position of each target, the expected number of background stars within a 2.4 arcsec FOV, considering H-band magnitudes ranging from 0 to +25, consistent with the RSM limits shown in Fig. 5.6. The cumulative counts of the detected signals and those predicted by the model are shown in Fig. 5.7. To assess whether the observed number of detections is compatible with the background-star population predicted by the Galactic model, we perform a statistical hypothesis test. The null hypothesis assumes that all detected sources originate from the expected background population, and the corresponding p-value quantifies the probability of obtaining a discrepancy at least as large as the observed one under this assumption. The figure indicates an expected total of only one background star across the 23 targets, compared to 30 detected signals, of which 28 are photometrically compatible with background stars. Assuming a Poisson distribution, the probability that these detections are explained solely by a background-star origin is extremely low, with a p-value of $\sim 1.33 \times 10^{-25}$. This indicates that the ensemble of detections cannot be explained by background contamination alone and suggests that the majority of these signals are likely false positives.




Figure 5.7: Cumulative number of detected signals compared with the expected number of background stars predicted by the Besançon Galactic model.

we repeated the same statistical test for each target individually. The resulting p-values are shown in Fig. 5.8. We find three stars with relatively high p-values: HD 113457, ASAS J193304+0345.7 and CD-51 10924. The remaining signals have p-values ranging between 0.02 and 0.1, with values close to zero in some cases. Since this analysis is performed on individual targets, we do not consider p-values around 0.05 to be statistically significant in this context, as they do not remain robust when interpreted collectively. This suggests that all detections, except for possibly the three targets mentioned above, are most likely false positives.

We performed a second consistency check to assess whether the detected signals are compatible with an astrophysical origin. This test relies on the chromatic behavior of speckles: residual speckles produced by diffraction scale approximately with wavelength ($\sim \lambda/D$), causing their apparent separation from the central star to change between spectral channels. In contrast, a genuine astrophysical source remains at a fixed angular separation. We therefore compared the measured displacement of each signal between the H2 and H3 bands with the displacement expected from speckle




Figure 5.8: Comparison between the detected number of point-like sources (blue columns), the number of background stars expected from the Besançon Galactic model (pink columns), and the p-values associated with each target (black dashed columns).

wavelength scaling. The expected relative scaling between the H2 and H3 bands is computed as $(\lambda_{H3} - \lambda_{H2})/\lambda_{\text{mean}}$, corresponding to a scaling of about 2%. Most of the detected signals exhibit a displacement between H2 and H3 consistent with this value, as shown in Fig. 5.9, suggesting that a large fraction of these detections are residual speckles and therefore likely false positives. However, a subset of detections exhibits little or no measurable displacement, or a displacement smaller than expected from speckle wavelength scaling, namely those around HD 128819, alf Sgr, CD-81 812, HD 317617, HD 1466, ASAS J193304+0345.7, π Ara, Smethells 20, HD 113457, PM J19560–3207 (two signals) and CD–57 1054. These detections are therefore not fully consistent with the expected chromatic behavior of speckles and cannot be identified as false positives based on this diagnostic alone.

In summary, the two approaches used to assess the nature of the signals yield broadly consistent results. The background-star expected-count analysis indicates that three out of the 30 detections are likely explainable as background stars. The H2-H3 scaling analysis, on the other hand, suggests that approximately twelve out of the 30 detections exhibit astrometric measurements not compatible with speckles chromatic behavior. We




Figure 5.9: Change in separation between the H2 and H3 bands for all RSM detections (blue columns), compared with the expected H2-H3 wavelength scaling indicated by the dashed gray line.

note that ASAS J193304+0345.7 and HD 113457 are consistent with a background-star interpretation in both approaches. In addition, HD 128819 and HD 317617 show small displacements between bands, together with background-star p-values around 0.08-0.09. This suggests that these four detections can be classified as newly identified background stars, while the remaining signals with background star photometry are likely false positives.

Regarding the planetary-like signals, namely π Ara and CD-57 1054, we find that the signal associated with π Ara is also recovered by NA-SODINN (Cantero et al., 2026) at 888 mas and 114° in both the H2 and H3 bands, consistent with the positions reported here. Owing to the availability of an additional epoch in the SHINE survey (not included in the F150 sample), we therefore decided to follow up this target, although the behavior of this signal (as well as that of CD-57 1054) scales with wavelength, as shown in Fig. 5.9.

5.3 Exoplanet candidate: π Ara

5.3.1 Context

π Ara is an A5 IV/V star with an estimated age of 600 ± 220 Myr, located at a distance of ~ 41 pc. The system hosts a debris disk that was spatially resolved in the far-infrared using dual-band imaging with the *Herschel*/PACS instrument at $100 \mu\text{m}$ (Morales et al., 2013). Subsequent analyses (Mittal et al., 2015; Morales et al., 2016; Lazzoni et al., 2018) interpreted the system’s spectral energy distribution using a two-belt debris disk model with an inner warm belt located at ~ 9.6 AU ($\sim 0''.22$), and an outer cold belt located at ~ 98 AU ($\sim 2''.4$), as described in (Lazzoni et al., 2018).

The presence of two spatially separated dust belts is reminiscent of the Solar System structure, with an asteroid-belt-like inner component and a Kuiper-belt-like outer component. Such a configuration is generally interpreted as evidence for dynamical sculpting by one or more planetary companions that maintain the dust-depleted region between the belts. Using dynamical stability arguments, Lazzoni et al. (2018) explored the range of planetary masses capable of sustaining the observed gap. They showed that a single planet on a circular orbit would require a mass exceeding $\sim 50 M_J$, while substantially lower masses are allowed if the orbit is eccentric ($\sim 11.8 M_J$). Alternatively, multi-planet configurations remain viable, with three planets on circular orbits permitted down to masses of $\sim 5.1 M_J$.

π Ara has been targeted by several exoplanet imaging surveys, including observations with GPI (Nielsen et al., 2019a; Squicciarini et al., 2025) and the SHINE/SPHERE survey (Langlois et al., 2021; Chomez et al., 2025). No detections were reported within the FOV, even after applying advanced post-processing algorithms such as PACO on both instruments (Squicciarini et al., 2025; Chomez et al., 2025). Astrometric measurements based on the cross-calibration of Hipparcos and Gaia EDR3 (Brandt, 2021) showed no evidence of significant astrometric acceleration. Radial-velocity monitoring, primarily with HARPS (Rupprecht et al., 2004), revealed substantial variability with amplitudes reaching up to $\sim 612.4 \text{ m s}^{-1}$ (Lagrange et al., 2009a; Díaz et al., 2011; de Bruijne & Eilers, 2012; Saffe et al., 2021; Grandjean et al., 2023). However, these variations have consistently been attributed to stellar pulsations, which are common among A-type stars.

Within the SHINE survey, Langlois et al. (2021) and Chomez et al. (2025) identified nine background stars within the full SPHERE FOV but no planetary companion. In this work, several signals are recovered by the RSM within the $2''.4$ FOV, with one signal lying outside the exclusion area in the CMD and therefore suggesting a potential planetary origin. Motivated by this possibility, we revisited the π Ara observations obtained in 2016 and additionally analyzed a second epoch acquired in 2019, which was not included in the F150 sample.

5.3.2 RSM application on π Ara

The initial RSM analysis reported four signals around π Ara in the 2016 observation. We therefore re-optimized the PSF subtraction and RSM parameters for this dataset to improve accuracy and to minimize potential mismatches between the parameters derived from the cluster center and those appropriate for this target. The same optimization procedure was applied to an additional observation of π Ara obtained in 2019. The adopted parameters are listed in Table 5.2.

The new optimization recovered three of the four previously reported signals, labelled a, b, and c in the 2016 dataset, and additionally revealed several new detections in the 2019 observation (d–g), all shown in Fig. 5.10. The astrometry and photometry of these signals are summarized in Table 5.3.

As shown in the figure, the 2019 observation also exhibits several candidate signals. No strict detection threshold was imposed at this stage, as we aimed to explore all potential cases, including signals such as d and e, which are recovered in H2 but not necessarily in H3 (as in the case of e). This behavior is of particular interest because the candidate signal c detected in 2016 is brighter in the H2 band than in H3, a characteristic expected for a planetary signal. However, none of the candidates detected in the 2019 observation shows a photometric behavior similar to signal c.

Furthermore, assuming a companion located near the inner belt at ~ 9.8 AU, corresponding to a minimum orbital period of about 24 years, the three-year time baseline between the two observations would imply only a small positional shift (approximately 60 mas). Such a displacement is not observed for any of the signals recovered in the 2019 data.

Table 5.2: Optimal parameters for PSF subtraction techniques and RSM for the two observation of Pi Ara in 2016 and 2019, taken in the both bands H2 and H3.

Parameters	2016 H2	2016 H3	2019 H2	2019 H3
APCA components	21	20	37	16
APCA segments	3	2	4	2
APCA FOV rotation	0.59	0.39	0.66	0.57
NMF components	15	20	8	10
LOCI tolerance	0.0034	0.0065	0.0061	0.0072
LOCI FOV rotation	0.6	0.57	0.1	0.13
FM KLIP components	16	17	5	6
FM KLIP FOV rotation	0.7	0.55	0.26	0.53
FM LOCI tolerance	0.0045	0.0042	0.0085	0.009
FM LOCI FOV rotation	0.53	0.49	0.1	0.23
APCA delta	5	5	5	5
NMF delta	5	5	5	5
LOCI delta	5	5	5	5
FM KLIP delta	5	5	5	5
FM LOCI delta	5	5	5	5
APCA crop size	4	4	4	4
NMF crop size	4	4	4	4
LOCI crop size	4	4	4	4
FM KLIP crop size	7	7	7	7
FM LOCI crop size	7	7	7	7
APCA variance	SM	ST	SM	ST
NMF variance	ST	FR	ST	ST
LOCI variance	ST	ST	SM	ST
FM KLIP variance	ST	ST	ST	ST
FM LOCI variance	SM	ST	ST	SM




Figure 5.10: Final RSM maps (left and middle panels) in comparison to the S/N PCA (right) for π Ara taken in 2016 (top) and 2019 (bottom) in both bands H2 (left) and H3 (middle).

5.3.3 Nature of the detected signals

We investigate here whether the detected signals around π Ara could be associated with background stars. A proper-motion test was performed for all candidates, as presented in Fig 5.11. This analysis reveals no compatibility between the detected signals and a background-star scenario. This conclusion is further supported by the very low number of expected background stars predicted by the Besançon Galactic model, previously shown in Fig. 5.8, which yields an expectation value of only 0.11 background stars.

In addition, the behavior of the signals between the H2 and H3 bands follows the characteristic scaling of false positives, with a relative displacement of approximately 2%, consistent with the expected speckle-noise behavior.

These observations were also processed with NA-SODINN, leading to the recovery of signal *b* in the 2016 epoch and signals *d* and *f* in the 2019 epoch. In contrast, PACO analyses reported by Squicciarini et al. (2025); Chomez et al. (2025) did not recover any of these signals.




Figure 5.11: Proper-motion test comparing the signals detected in the 2016 and 2019 observations of π Ara. The three signals identified in 2016 are fixed as reference points, and for each subplot the expected positions in 2019 are compared with the signals recovered in the 2019 data. The analysis is performed using the backtracks code (Balmer et al., 2025a).

Table 5.3: Photometry and astrometry of the detected signals in π Ara in both observations taken in 2016 and 2019.

Signal	Date	Band	Separation [mas]	Position angle [°]	Contrast	Δmag
a	2016-06-10	H2	328.48	193.86	1.2×10^{-5}	12.30
		H3	298.41	196.76	1.1×10^{-5}	12.35
b	2016-06-10	H2	895.73	213.96	5.1×10^{-6}	13.23
		H3	930.08	213.49	6.4×10^{-7}	15.48
c	2016-06-10	H2	895.37	114.38	7.6×10^{-6}	12.79
		H3	925.49	113.59	1.9×10^{-6}	14.31
d	2019-05-19	H2	742.26	85.63	4.5×10^{-6}	13.36
		H3	779.70	85.63	3.6×10^{-6}	13.61
e	2019-05-19	H2	1039.02	344.43	2.1×10^{-6}	14.20
		H3	1076.52	344.23	1.7×10^{-6}	14.41
f	2019-05-19	H2	504.26	82.10	6.3×10^{-6}	13.01
		H3	532.14	82.58	5.7×10^{-6}	13.11
g	2019-05-19	H2	892.40	170.04	2.6×10^{-6}	13.97
		H3	920.11	170.48	2.4×10^{-6}	14.06

Taking all these elements together, including the chromatic scaling of signals consistent with speckle behavior, the absence of recovery at the expected positions in the second epoch, and the non-detection of any of these signals with PACO analyses in both SHINE and GPI datasets, we conclude that the detected signals are most likely false positives. Regarding RSM and NA-SODINN, while both methods rely on different models for the planetary signal and residual noise during the detection stage, they both depend on similar classical PSF-subtraction algorithms for the modeling and subtraction of the stellar PSF, a step on which PACO does not rely. The recovery of these signals may therefore reflect limitations of classical PSF-subtraction techniques in accurately modeling the stellar PSF.

5.4 Discussion

This work was initially motivated by the sensitivity limitations close to the IWA of the standard S/N-based detection techniques used in Langlois et al. (2021), and by the potential improvement that the RSM framework could provide at these separations. This potential was well illustrated with an enhancement of a factor of four to five in contrast at the IWA and approximately a factor of two overall. These results are broadly consistent with the recent results reported for PACO-ASDI in Chomez et al. (2025).

Out of the 213 observations in the SHINE F150 sample, 153 showed no signal in the RSM analysis, while the remaining 60 observations exhibited candidate signals. Owing to the PACO analysis of the full SHINE survey, which benefits from additional observations at multiple epochs for several targets, together with the use of SnapSHINE follow-up observations that could not be incorporated into the RSM framework because they typically contain too few frames for reliable RSM analysis, we adopted the signal classifications and assessments reported by Chomez et al. (2025) as a reference for the interpretation of the candidate signals. A summary comparison between the PACO and RSM results, restricted to the same observations, filters, and FOV, is presented in Figure 5.12. Both methods showed no signal in 150 observations, while 49 signals were recovered by both approaches across 29 observations. PACO additionally reported three signals that were not recovered using the adopted RSM detection threshold; their non-recovery was discussed previously in Sect. 5.2.1. Conversely, the RSM analysis identified in the remaining 31 observations one known planetary signal corresponding to PZ Tel, two background stars previously reported by Langlois et al. (2021) but not by Chomez et al. (2025), and four additional background-star candidates. However, the analysis also revealed approximately 33 false positives, corresponding to a false positive rate of 18%, whereas the true positive rate, computed using the PACO classifications as the reference, reached 94%. The high true positive rate demonstrates the strong detection capability of the RSM framework, although the elevated false positive rate suggests a potential issue in the adopted thresholding strategy.

To visualize this behavior, we plot all signals recovered by RSM, including signals common with PACO, non-common detections, and false positives, as a function of their attributed RSM score in Figure 5.13. We find that most detections common with PACO, referred to here as true positives, are recovered at high RSM confidence values close to 1. A smaller number of true positives are recovered at lower RSM scores, slightly above the adopted detection threshold; however, their recovery occurs at the expense of a significantly larger number of false positives. This elevated false positive rate is directly linked to the underestimation of the noise tails highlighted in Figure 4.1.

However, the limitations of RSM are not solely related to the threshold estimation. The method also recovers false positives with relatively high confidence scores (between 0.1 and 1). This behavior is illustrated by the example of π Ara discussed in Sect. 5.3.2, where several false positives were recovered in both observations and were also identified by NA-SODINN. Although NA-SODINN relies on a deep-learning classifier, both detection methods operate on residual images produced by classical PSF-subtraction




Figure 5.12: Summary of the detection statistics for the F150 sample, comparing the outcomes of RSM and PACO across four cases: no signal in either method, signal in PACO only, signal recovered by both methods, and signal recovered by RSM only. The dashed gray bars indicate the number of observations in each category, with the corresponding counts labeled above them. The colored stacked bars represent the number of detected signals, with colors denoting their classification: confirmed planets, background stars, ambiguous sources, false positives. This figure presents an updated version of the statistics reported by Sabalbal et al. (2026), incorporating the revised candidate classifications obtained in this work.

techniques such as PCA. The recovery of the same false-positive signals by both methods therefore suggests that at least part of these high-confidence false positives may originate from limitations of the PSF-subtraction stage, rather than from the RSM classification itself.

We explore in Figure 5.14 the behavior of the RSM detections as a function of atmospheric and observational parameters. No individual parameter appears to be clearly associated with the production of false positives. However, a larger number of false positives is observed for lower total rotation angles, shown in panel (b), as well as for datasets with a lower number of frames (panel f), larger temperature differences between the ambient air and M1 (panel h), lower coherence time (panel e) and lower




Figure 5.13: Distribution of detected signals as a function of their corresponding RSM score and separation. The figure shows astrophysical signals detected by both RSM and PACO (blue), astrophysical signals detected only by RSM (green), and false positives detected by RSM (red).

raw contrast values (panel c), the latter also being highlighted by Cluster 5 in panel (j). The apparent reduction of false positives at high frame numbers is likely affected by a selection bias, since observations with the highest number of frames predominantly correspond to known planetary detections observed with SPHERE.

Interestingly, parameters commonly associated with high-contrast imaging quality, including the seeing, Strehl ratio, wind speed, and even the euclidean distance from the cluster center, which could indicate an inadequate association between a dataset and its assigned cluster, potentially leading to suboptimal PSF-subtraction parameters, do not exhibit any clear trend with the occurrence of false positives.




Figure 5.14: Distribution of detected signals as a function of their corresponding RSM score and various observational and atmospheric parameters: seeing, total field rotation, raw contrast, wind speed, coherence time, number of frames, Strehl ratio, temperature difference between the ambient air and M1, stellar magnitude, corresponding cluster number, and distance to the cluster center. The figure shows astrophysical signals detected by both RSM and PACO (blue), astrophysical signals detected only by RSM (green), and false positives detected by RSM (red).




Figure 5.14: (continued)




Figure 5.14: (continued)

Chapter 6

RSM applied to Epsilon Eridani with Keck

Abstract. The ϵ Eridani system represents one of the most important targets for direct imaging studies of nearby planetary systems. Its proximity and the presence of a giant planet with a mass comparable to that of Jupiter make it a valuable benchmark for understanding giant planets in solar system-like architectures. Despite extensive observations with many of the major HCI instruments, ϵ Eridani b, originally identified through radial-velocity measurements, has remained undetected in direct imaging. In this chapter, we investigate the application of the RSM framework to archival Keck observations of ϵ Eridani in an effort to further improve the detection limits achieved for this system. This analysis was conducted prior to the publication of the latest JWST observations, which reached significantly deeper detection limits and provided stronger constraints on the properties of the planet.

This chapter is organized as follows. Section 6.1 presents an overview of the ϵ Eridani system. Section 6.2 describes the Keck observations used in this work. Section 6.3 presents the application of the RSM framework and the resulting detection limits. Finally, Section 6.4 discusses the results of the analysis and compares them with the most recent constraints obtained with JWST.

6.1 Introduction

Epsilon Eridani is a young K2V main-sequence star (Keenan & McNeil, 1989) located at a distance of 3.22 pc from the Sun (Gaia Collaboration, 2020), making it one of the closest Sun-like stars to Earth. The ϵ Eridani system hosts a debris disk composed of a main outer ring extending from approximately 35 to 90 AU, an inner belt at ~ 3 AU, and an intermediate belt near ~ 20 AU (Backman et al., 2009). A Jupiter-like planet has been confirmed between the two inner belts at a semi-major axis of about 3.4 AU (Hatzes et al., 2000), with an orbital period of approximately 7.4 yr (Mawet et al., 2019). Several studies have suggested the presence of additional planets inferred from disk morphology, including a Jovian-mass planet near 20 AU proposed to explain the cleared region between the intermediate and outer belts (Backman et al., 2009), and another candidate inferred from the narrow width of the outer belt, suggesting a planet at ~ 48 AU (Booth et al., 2017).

The age of ϵ Eridani has long been debated. Early estimates of approximately 200 Myr were based on kinematic arguments associating the star with the Ursa Major moving group and adopting the age inferred for that population (Fuhrmann, 2004). Using chromospheric activity diagnostics, (Mamajek & Hillenbrand, 2008) derived a significantly older age in the range of 400–800 Myr. As a result, a broad range of ages has been adopted across the literature. More recent studies based on isochrone fitting, such as (Sahlholdt et al., 2019), suggest that the system is likely older, with an age in the ~ 400 –900 Myr range. The most recent gyrochronology study by Sanghi et al. (2026) further supports an older age, estimating approximately 1.1 Gyr.

The ϵ Eridani system is of particular interest due to its resemblance to the Solar System, with multiple belts potentially shaped by planets at comparable relative separations. Its young age makes it a valuable laboratory for studying planet formation pathways and the early dynamical evolution of planetary systems, offering insight into the past evolution of the Solar System. In addition, its main-sequence lifetime provides long-term dynamical stability, favorable for habitability studies. Finally, its close proximity enables high spatial resolution observations, making ϵ Eridani an exceptional target for detailed investigations of planet–disk interactions.




Figure 6.1: Constraints on the expected on-sky position of ϵ Eridani b during the 2017 and 2019 Keck observing epochs. Figure adapted from (Thompson et al., 2025).

Epsilon Eridani b was first reported by (Hatzes et al., 2000) based on nearly 20 years of radial velocity monitoring and was subsequently confirmed by (Benedict et al., 2006) through a combination of radial velocity measurements and astrometry. Early analyzes suggested a giant planet with an orbital period of approximately 7 years, a minimum mass of $0.86 M_J$, and a relatively high eccentricity ($e \sim 0.6$). Since then, and over the following two decades, ϵ Eridani has remained a prime target for numerous surveys and facilities.

Multiple direct-imaging campaigns—including observations with Spitzer (Marengo et al., 2006, 2009; Janson et al., 2015), Gemini (Wahhaj et al., 2013b; Nielsen et al., 2019b), NACO (Janson et al., 2007; Launhardt et al., 2020), Keck (Mawet et al., 2019; Llop-Sayson et al., 2021), SPHERE (Tschudi et al., 2024), and most recently JWST/NIRCam (Llop-Sayson et al., 2025)—have not resulted in a direct detection of the planet. Nevertheless, progressively tighter constraints on the orbital parameters of ϵ Eridani b have been obtained through the combined analysis of radial velocity, astrometric, and direct-imaging data (Mawet et al., 2019; Llop-Sayson et al., 2021). The most recent analysis by (Thompson et al., 2025) favors a planet with an orbital period of 7.32 years, a mass of approximately $1 M_J$, and a low eccentricity of $e \simeq 0 - 0.09$. This study also places constraints on the expected on-sky position of the planet at different epochs, as shown in Figure 6.1.

The deepest contrast limits for this planet have been achieved with ground-based facilities and, more recently, with space-based observations from JWST, reaching contrasts of $\sim 7 \times 10^{-6}$ at $1''$ in the Keck M band (Llop-Sayson et al., 2021) and $\sim 1 \times 10^{-6}$ with JWST/NIRCam $F444W$ band (Llop-Sayson et al., 2025).

In this work, we aimed to push the detection limits of Keck/NIRC2 observations of ϵ Eridani in the M_s band using the RSM algorithm. This study was conducted prior to the publication of the JWST/NIRCam results. Nevertheless, we present here the improvements in contrast limits achieved with Keck data through the application of the RSM framework.

6.2 Observations and pre-processing

This study is using the deepest available ground-based observations of ϵ Eridani obtained with the near-infrared camera NIRC2 mounted on the 10 m Keck Telescope at the W. M. Keck Observatory, following the installation of its vector vortex coronagraph (Serabyn et al., 2017). The observations were carried out in pupil-tracking mode, enabling ADI observations. The NIRC2 narrow camera configuration was used throughout, providing a plate scale of $9.942 \text{ mas pixel}^{-1}$.

The data were acquired at two different epochs, in 2017 and 2019, using the Keck adaptive optics system operating in both Shack-Hartmann and pyramid wavefront sensor modes. In total, nine individual observing sequences were obtained. The details of each observation are summarized in Table 6.1. Observing conditions were generally good, with median DIMM seeing values ranging from $0''.52$ to $0''.97$, and Strehl ratios reaching up to 90%. All observations were pre-reduced using the standard Keck pipeline, which includes flat-fielding, bad-pixel correction, sky subtraction, and image registration, following the procedure described in (Xuan et al., 2018). Additional details on the observations can be found in (Mawet et al., 2019; Llop-Sayson et al., 2021).

The dataset was subsequently post-processed by (Llop-Sayson et al., 2021) using principal component analysis (PCA) combined with a matched filter (Ruffio et al., 2017), adopting 30 principal components. In their analysis, observations obtained within the same year were combined to produce a single detection map, leading to a significant improvement in the achieved contrast limits. The resulting contrast limits at angular separations of $0''.4$ and $1''$ are reported in Table 6.1.

Table 6.1: Summary of observations using Keck/NIRC2 Ms band including the contrast limits reached by (Llop-Sayson et al., 2021). Table is derived from (Llop-Sayson et al., 2021).

Epoch	Exposure Time (s)	# frames	total PA	5σ 0.''4	5σ 1.''0	WFS
2017-01-09	0.5	210	88.6	4.8e-05	1.1e-05	SH
2017-01-10	0.5	260	100.2	4.9e-05	1.1e-05	SH
2017-01-11	0.5	160	69.5	6.3e-05	1.0e-05	SH
Combined 2017				3.3 e-05	7.0e-06	
2019-10-20	0.25	173	84.9	5.1e-05	2.0e-05	Pyramid
2019-10-21	0.25	195	82.0	3.9e-05	1.8e-05	Pyramid
2019-10-22	0.25	239	89.3	5.4e-05	2.0e-05	Pyramid
2019-11-04	0.25	222	68.2	6.3e-05	1.8e-05	Pyramid
2019-11-05	0.25	217	68.0	4.6e-05	1.3e-05	Pyramid
2019-12-08	0.3	217	103.7	1.0e-04	2.4e-05	SH
Combined 2019				2.3 e-05	7.9e-06	

In this work, we adopt a similar multi-epoch combination strategy, but employ the RSM algorithm with the aim of further pushing the achievable detection limits.

6.3 RSM applied to ϵ Eridani

In this section, we present the deepest detection limits obtained for the ϵ Eridani dataset within the RSM framework. We first describe the application of RSM to the individual observations (Sect. 6.3.1) and then explore strategies to further improve these limits by exploiting the multi-epoch nature of the data (Sect. 6.3.2).

6.3.1 RSM application

We begin by applying the RSM algorithm to each of the nine observations individually. The analysis is performed over a field of view ranging from 30 to 125 pixels, corresponding to angular separations of approximately 0.''3 to 1.''25. An annulus width of 10 pixels is adopted for the RSM computation.

Within the RSM framework, we consider annular PCA (APCA) and LOCI as PSF subtraction techniques. Compared to the analysis of the F150 SHINE sample, the NMF method is excluded due to unstable behavior observed across different observations. We also discard forward-modeling approaches based on KLIP and LOCI, as they in-

introduce a significant computational cost without yielding measurable improvements in performance. As a result, we focus on three RSM configurations: RSM with annular PCA (RSM APCA), RSM with LOCI (RSM LOCI), and a combined approach using both APCA and LOCI (RSM APCA-LOCI).

The final RSM maps are highly sensitive to the tuning of both the PSF subtraction and RSM parameters. For this reason, the APCA and LOCI parameters, as well as the RSM parameters are optimized for each observing sequence. The full set of adopted parameters is listed in Table 6.2.

To quantify the detection limits achieved in this analysis, we adopt the 50% completeness criterion introduced by Dahlgvist et al. (2022), in which the detection threshold is defined as the highest false-positive value. The limited number of observations available in this dataset makes it difficult to robustly apply thresholding approaches based on assumptions about the underlying noise distribution, such as those developed in Chapter 4. In addition, the Keck datasets are substantially larger than the SPHERE observations analyzed in this thesis, typically containing 100–200 frames. Combined with the use of multiple epochs and post-processing techniques, this makes the application of the maximum F1-score threshold computationally expensive. Since Appendix C shows that the different thresholding criteria yield very similar contrast limits in the absence of bright signals, we adopt the 50% completeness criterion with the highest false-positive threshold. This choice provides a robust estimate of the detection sensitivity while remaining broadly comparable to the standard 5σ detection limits.

The resulting contrast curves for each observation are presented in Fig. 6.2 and summarized in Fig. 6.3. For comparison, we also show signal-to-noise ratio (SNR) PCA results computed with 30 principal components using the same completeness formalism. An improvement in the completeness limits is already evident, reaching sensitivity levels comparable to those reported by Llop-Sayson et al. (2021).

6.3.2 Improvements to the RSM analysis

In this section, we investigate potential improvements to the detection sensitivity by exploiting the availability of multiple observing nights, as well as prior expectations on the planet location derived from orbital constraints (Fig. 6.1).




Figure 6.2: 50% completeness curves for nine individual observations across 2017 and 2019 corresponding to the optimized version of RSM APCA (blue), RSM LOCI (orange), RSM APCA-LOCI (green), in comparison to the SNR PCA with 30 principal components.

Combining observations

Motivated by the improvement reported by Llop-Sayson et al. (2021) through the combination of multiple observing epochs, we adopt a similar strategy by combining observations obtained on consecutive nights into a single data cube. The resulting extended time series is then analyzed within the RSM framework to produce a single final RSM map. We combine all observations obtained in 2017 into a single dataset, and similarly merge the observations taken in October and November 2019. The observation acquired in December 2019 is not combined with the other 2019 data, as the expected on-sky displacement of the planet would place it outside a common resolution element relative to the October and November epochs.

The resulting completeness curves obtained from these combined datasets are shown in Fig. 6.3. They are compared to the limits achieved for individual observations discussed in the previous section, as well as to PCA results obtained with 30 principal components. A factor of two to three improvement in contrast is achieved when comparing the RSM combined observations to the PCA combined observations for the 2017 dataset. The slightly degraded detection limits obtained with combined PCA compared to Llop-Sayson et al. (2021) comes from the fact that our PCA implementation does not include the matched-filter step used in that study. Nevertheless, the RSM completeness curve reaches a level of $\sim 5 \times 10^{-6}$ at $1''$ in the 2017 observations, which is deeper than the limits previously reported by Llop-Sayson et al. (2021), as shown in Table 6.1.

We present in Fig. 6.4 the RSM maps corresponding to the deepest contrast limits achieved in this study. The maximum RSM scores in both maps remain low overall, indicating the absence of any statistically significant detection.

Location-focused RSM

Given the prior constraints on the expected on-sky position of ε Eridani b at different epochs, we explored the use of azimuthal partition within the RSM framework to better isolate a potential planetary signal from unrelated false positives located at the same radial separation. In the RSM framework, azimuthal partition consists of dividing an annulus into several azimuthal segments, such that each segment is analyzed independently across the temporal sequence. This approach aims to suppress bright false positives elsewhere in the annulus that could otherwise dominate the detection threshold.




Figure 6.3: 50% completeness curves for the ϵ Eridani Keck/NIRC2 observations. Individual-night results obtained with PCA are shown as filled light-blue regions, while optimized RSM results for individual nights are shown as filled orange regions. The combined multi-night PCA and RSM results are displayed as the blue and red lines, respectively. *Left:* observations obtained in 2017. *Right:* observations obtained in 2019.

We tested azimuthal partitions consisting of 2, 4, 8, and 16 segments, corresponding to azimuthal widths of at least 4 FWHM at a separation of $1''$. These configurations were compared to the optimal number of segments obtained during the RSM parameter optimization, as shown in Fig. 6.5. Overall, we observe similar performance across the different azimuthal partitions. For the 2017 dataset, slightly improved contrasts at $1''$ are obtained for the 2 and 4 segment configurations, whereas the 8 and 16 segment configurations lead to degraded performance. This behavior can be understood as a trade-off intrinsic to azimuthal partitions. While partition can effectively exclude bright false positives located elsewhere in the annulus, increasing the number of segments reduces the length of the temporal sequence available to the RSM at a given location. This reduction limits the model's ability to robustly distinguish between noise and a genuine planetary signal, ultimately leading to degraded detection performance.




Figure 6.4: RSM maps corresponding to the combined observations obtained in 2017 and in 2019.

6.4 Discussion

The results obtained in this work correspond to the deepest detection limits achieved to date with ground-based Keck/NIRC2 observations of ϵ Eridani. The contrast limit reached, $\sim 5 \times 10^{-6}$, represents an improvement by a factor of ~ 1.4 relative to the deepest Keck results reported by (Llop-Sayson et al., 2021). Assuming a planet mass of $\sim 1 M_J$, these limits translate into sensitivity to ϵ Eridani b at an age of up to ~ 430 Myr, compared to ~ 350 Myr for the results of (Llop-Sayson et al., 2021), as shown in Fig. 6.6. This figure illustrates the expected contrast of ϵ Eridani b for a planet mass of $\sim 1 M_J$, based on the models of (Linder et al., 2019), using the available M_p (NACO) band as a proxy for the Keck M_s band, together with the expected temporal evolution of the planetary contrast.

Recent JWST/NIRCam observations reach substantially deeper limits, as shown in Fig. 6.6, probing contrasts consistent with planetary ages of ~ 1000 – 1200 Myr. The recent gyrochronology analysis by Sanghi et al. (2026), suggesting an age of $\sim 1.1 \pm 0.1$ Gyr for ϵ Eridani, naturally explains the non-detection of the planet reported by (Llop-Sayson et al., 2025). The same study further argues that the lack of detection may imply an enhanced atmospheric metallicity ($[M/H] \approx +0.5$ dex) and/or the presence of water clouds.

We also explore the relative performance of Keck and JWST not only at $\sim 1''$, corresponding to the expected position of ϵ Eridani b, but across all separations. Figure 6.7 shows this comparison, using the combined 2017 contrast curve from this work as the deepest Keck result and the JWST limits in the F444W filter. At $1''$, JWST improves the achievable contrast by roughly a factor of ~ 10 compared to Keck, highlighting the




Figure 6.5: 50% completeness curves for the combined ε Eridani observations, comparing the optimal RSM configuration (black) with fixed azimuthal partitions of 2 (blue), 4 (orange), 8 (green), and 16 (red). *Right:* observations obtained in 2017. *Left:* observations obtained in 2019.

superior performance of space-based observations in the background-limited regime. At smaller separations, however, the performances become more comparable. In particular, the Keck M_s -band data presented in this work reach contrast levels similar to JWST/F444W at separations of $0''.35 - 0''.45$. This implies that, for signals expected at these angular separations, Keck observations combined with advanced post-processing can achieve sensitivities comparable to JWST on bright targets like ε Eridani. Moreover, similar contrast levels in the M_s and F444W bands do not imply identical detection limits. The slightly longer wavelength of the M_s band corresponds to a slightly higher planetary brightness, as predicted by the evolutionary models of Linder et al. (2019), potentially enabling the detection of slightly lower-mass planets. This effect is illustrated in Figure 6.6, where the contrast required to detect a planet of a given mass is systematically higher in F444W than in the M_s band. As a result, the Keck and JWST observations provide comparable performance at small separations, with the Keck M_s data potentially offering slightly better sensitivity in this regime.




Figure 6.6: Temporal evolution of the expected contrast of ϵ Eridani b based on models from (Linder et al., 2019), assuming a mass of $\sim 1 M_J$ with $[M/H] = 0$ and 0.4 , and including both cloudy and cloud-free atmospheric models. The model predictions are compared to contrast limits obtained with Keck and JWST. *Left:* M -band contrast limits from Keck. *Right:* F444W band contrast limits from JWST.

6.5 Summary

In summary, the RSM analysis of the Keck/NIRC2 observations of ϵ Eridani does not yield any statistically significant detection, a result fully consistent with the non-detection reported by Llop-Sayson et al. (2025) using JWST/NIRC2 data.

The detection limits obtained in this work through the application of RSM and the combination of multiple observing nights show an improvement by a factor of ~ 1.4 compared to the latest published Keck results. We also find an improvement of approximately a factor of 2 relative to individual-night RSM results, highlighting the enhanced performance of the RSM framework when multi-epoch data are available. These improvements provide new Keck constraints on the ϵ Eridani system, reaching sensitivity to ~ 430 Myr ϵ Eridani b, compared to limits below ~ 400 Myr in previous studies.

Finally, we show that at large separations JWST clearly outperforms Keck, reaching contrasts roughly an order of magnitude deeper. At smaller separations, however, the performance becomes more comparable, with the Keck M_s observations showing similar than JWST around $0.35\text{--}0''.45$.




Figure 6.7: Comparison of contrast limits for ϵ Eridani obtained with JWST (F210M and F444W) and M_s Keck observations from this work.

Chapter 7

Conclusions & perspectives

7.1 Conclusions

This thesis was motivated by the limited sensitivity of HCI observations at small angular separations, where residual stellar noise remains a major obstacle to the detection of planetary companions. To address this challenge, the RSM was applied to the archival F150 sample of the SHINE survey with the objective of improving detection performance and searching for previously undetected companion candidates. The application of RSM to the F150 sample yielded contrast improvements of up to a factor of five at small angular separations and approximately a factor of two overall, enabling the recovery of around 89 point-like sources, including four background-star candidates and no new planetary candidate. Similar improvements were obtained for the ϵ Eridani system, where the resulting detection limits provided additional constraints on the age of the proposed companion based on the Keck observations. These results demonstrate that the RSM framework can significantly improve the sensitivity of archival HCI observations and increase their scientific return.

The clustering of datasets based on observational parameters revealed groups of observations sharing common HCI features, including wind-driven halo and low-wind effects. This result suggests that observational metadata can successfully identify physically meaningful observing regimes and provides a useful framework for studying their impact on HCI observations. These groups were also found to exhibit distinct noise properties in the final RSM maps, with different lognormal distributions leading to significant variations in the derived detection thresholds. In particular, observations with a larger number of frames generally led to lower thresholds, whereas observations affected by degraded observing conditions required higher thresholds. These results demonstrate that detection thresholds should be adapted to the characteristics of each observational regime rather than assume a universal value across an entire survey.

The assumption of an underlying noise distribution was not the only approach explored for determining detection thresholds. A distribution-free method based on the maximization of the F1 score was also investigated and was found to yield results broadly consistent with those obtained using the lognormal distribution. However, the resulting thresholds were found to be highly sensitive to the specific false positives present in a given dataset, making them difficult to generalize across different datasets within the same observational group. In addition, their application within the RSM framework proved computationally expensive when performed on a dataset-by-dataset basis. Nevertheless, the method provides a valuable alternative when no assumptions can be made about the underlying noise distribution and can be applied to other detection algorithms, as recently demonstrated for NA-SODINN (Cantero et al., 2026).

The analysis of the RSM detections in the F150 sample revealed a larger number of signals than expected from the comparison with PACO, despite the similar overall detection performance of the two methods. While more than 95% of the expected astrophysical signals were recovered, an additional 33 signals were identified, corresponding to approximately 18% of the total detections. Although most of these sources occupy regions of the color–magnitude diagram typically associated with background stars, neither the expected background-star density nor their wavelength-scaling behavior supports such an interpretation for all of these signals. As a result, only four new signals were classified as likely background stars, while the majority are likely false positives. The origin of these false positives appears to be twofold. First, low-RSM-score false positives are associated with an underestimation of the detection threshold derived from the fitted lognormal distribution. This result highlights the risks of relying on empirical distributions without a theoretical foundation, as inaccuracies in the tail directly translate into inaccurate false-positive probabilities. In this respect, the problem is analogous to the limitations of the Gaussian assumption commonly adopted in HCI post-processing (Pairet et al., 2019). The present work provides a direct quantification of the impact that such threshold underestimations can have on the resulting detections. Second, the presence of high-RSM-score false positives suggests limitations that are intrinsic to the current implementation of the framework. The results may indicate that the use of classical post-processing techniques within RSM can generate high-confidence detections that are not associated with astrophysical sources. Together, these findings demonstrate that both the underlying post-processing technique and the adopted detection threshold strategy

have a significant impact on the reliability of RSM detections. The reduction of false positive detections therefore remains a key challenge for the reliable application of the RSM framework, particularly given the substantial observational resources required for follow-up observations.

7.2 Perspectives

The results of this thesis suggest several directions for future developments of the RSM framework, as well as scientific applications of the improved detection capabilities demonstrated in this work.

7.2.1 Future developments of the RSM framework

Characterization of the noise in the residual spatio-temporal series with the SHINE survey

The statistical distribution of the residual spatio-temporal noise was investigated by Dahlqvist et al. (2020) on VLT/NACO datasets, who found that it can be described by a Gaussian distribution, a Laplacian distribution, or a mixture of both. Extending this analysis to the full F150 sample would provide a more comprehensive assessment of the validity of these models across a wide range of observing conditions in SPHERE. Such an analysis could reveal limitations of the currently assumed distributions or identify more appropriate statistical models, thereby improving the likelihood estimation within the RSM framework and ultimately enhancing the algorithm's detection performance.

Validation of the ADI+SDI implementation of the RSM

Although this thesis focuses on the ADI implementation of the RSM, the framework already includes an ADI+SDI implementation that has not yet been systematically tested on scientific datasets. Recent studies have demonstrated the benefits of combining angular and spectral diversity, with ADI+SDI approaches achieving deeper detection limits than ADI alone. In particular, the ASDI implementation of PACO (Flasseur et al., 2020b) contributed to the recovery of three faint astrophysical signals close to the contrast limit in the SHINE survey (Chomez et al., 2025) that were not recovered by the

ADI implementation of the RSM. Validating and applying the ADI+SDI implementation of the RSM therefore represents a promising avenue for further improving its detection performance and assessing its ability to recover faint companions beyond the capabilities of the current ADI implementation.

Integration of advanced post-processing techniques

Several advanced post-processing methods have recently been proposed to overcome some of the limitations of classical techniques, including the AMAT algorithm (Daglayan et al., 2024) and 4S (Bonse et al., 2025). These methods have demonstrated improved performance compared to the PCA and LOCI implementations currently used within RSM. Their integration into the framework could improve the sensitivity of the algorithm and help reduce the population of high-RSM-score false positives identified in this thesis.

Computational acceleration of the RSM framework

Despite the efforts made to reduce its computational cost through parallelization, the RSM framework remains computationally demanding. Further optimization would facilitate its adoption within the HCI community and reduce its dependence on access to high-performance computing resources. In addition, faster implementations would enable the routine use of computationally intensive threshold-determination methods, such as F1-score maximization, which could help address the threshold underestimation responsible for part of the low-RSM-score false-positive population identified in this work.

7.2.2 Future scientific applications

Application to the full SHINE survey

A natural continuation of this work would be the application of the RSM framework to the full SHINE survey. Extending the analysis beyond the F150 sample would provide homogeneous RSM-derived detection limits across the survey and allow the investigation of additional observations with the framework. The resulting detection limits could subsequently be incorporated into occurrence-rate studies, although such analyses have already been performed for the survey using PACO (Chomez et al., 2025).

Application to other archival HCI surveys

The scientific potential of reprocessing archival HCI datasets with modern detection frameworks has been highlighted by several recent studies. Examples include the recovery of AF Lep b by Bonse et al. (2025) and the discovery of the planetary companion HD 143811 b in archival GPI and SPHERE observations (Squicciarini, V. et al., 2025). While major surveys have already benefited from re-analysis efforts, large archives such as those obtained with Keck/NIRC2, the BEAST survey with SPHERE (Janson et al., 2021), and the LEECH survey (Skemer et al., 2014) remain valuable targets for modern detection frameworks. The application of RSM to these datasets may improve existing detection limits and potentially reveal previously undetected astrophysical sources.

Application of RSM to future HCI facilities

The next generation of HCI facilities, notably the Nancy Grace Roman Space Telescope and the Extremely Large Telescope (ELT), will significantly expand the discovery space for directly imaged exoplanets. The Roman Space Telescope will enable the detection of mature giant planets at contrasts inaccessible from the ground. Although Roman observations will primarily rely on dedicated observing strategies such as roll subtraction, deep observations with a large number of exposures may also benefit from the RSM framework to further improve the achievable detection limits. The ELT, equipped with instruments such as MICADO, will provide unprecedented angular resolution and sensitivity. RSM could become part of the suite of advanced post-processing techniques used to further enhance the sensitivity of these observations. The experience gained in this thesis from applying RSM to the SHINE survey, together with the identification of the framework's current limitations, provides a strong foundation for further developing RSM in preparation for next-generation HCI observations.

Appendix A

F150 observations

Table A.1 lists all the observations included in this thesis, excluding those identified as low quality. For each dataset, the corresponding observational parameters are provided, including the total parallactic angle range, seeing, wind speed, Strehl ratio, coherence time, number of frames, and the raw contrast at 500 mas. The table also reports the cluster assigned to each dataset through the clustering analysis, together with its corresponding cluster center.

Table A.1: Name of the star, observation date, total parallactic angle range (Δ PA), seeing (S), wind speed (WS), Strehl ratio (SR), coherence time (CT), number of frames (F), raw contrast at 500 mas (RC), and associated cluster number (C) (defined in Sect. 2.2) are provided for each F150 target. Observations marked with shaded areas represent cluster centers.

Star	Date	Δ PA ($^\circ$)	S ($''$)	WS (m/s)	SR	CT (ms)	F	RC	C
* bet Pic	2016-09-16	38.71	0.42	3.53	0.7	1.72e-02	464	5.30e-05	1
* bet Pic	2015-11-28	5.97	0.94	1.27	0.93	1.70e-02	80	4.65e-08	1
* del Crv	2018-01-25	94.44	0.59	1.13	0.79	1.05e-02	72	6.68e-05	1
* e Vir	2018-04-13	30.25	0.34	1.7	0.94	1.05e-02	90	2.72e-05	1
* q Vel	2018-02-27	55.55	0.69	4.75	0.93	7.90e-03	140	2.50e-05	1
ASAS J205428+0906.1	2016-09-17	37.47	0.62	0.75	0.54	2.96e-02	75	1.87e-04	1
BD+02 1729	2017-02-13	38.4	0.75	4.75	0.8	1.06e-02	80	8.44e-05	1
BD-15 705	2017-11-06	71.08	0.56	1.08	0.77	8.00e-03	48	5.51e-05	1
BD-21 1074A	2017-02-07	27.06	0.71	1.48	0.84	1.79e-02	72	4.10e-05	1
CD-31 16041	2018-04-17	84.52	0.88	4.38	0.8	1.50e-02	42	8.97e-05	1
CD-54 8168	2018-04-19	36.19	0.64	1.55	0.81	1.52e-02	50	8.20e-05	1
HD 105	2015-09-26	44.98	1.06	2.38	0.73	1.59e-02	72	1.48e-04	1
HD 61005	2015-02-03	93.04	0.66	1.55	0.77	3.55e-02	64	2.56e-04	1
HD 75505	2018-01-06	27.43	0.59	3.65	0.81	9.20e-03	60	9.32e-05	1
HD 90712	2016-01-03	74.77	0.67	0.38	0.8	6.50e-03	64	1.22e-04	1
HD 106048	2018-03-30	47.27	0.45	2.83	0.79	1.66e-02	60	9.76e-05	1
HD 107821	2015-02-04	37.18	0.74	3.35	0.87	1.20e-02	92	2.49e-05	1
HD 109573	2015-02-03	48.5	0.67	3.6	0.92	1.28e-02	112	3.34e-05	1
HD 128819	2017-03-19	67.61	0.63	1.63	0.86	7.50e-03	77	3.99e-05	1

Table A.1: continued.

Star	Date	Δ PA (°)	S (″)	WS (m/s)	SR	CT (ms)	F	RC	C
HD 159911	2018-04-12	72.72	0.35	2.25	0.78	6.90e-03	50	6.44e-05	1
V* PZ Tel	2016-09-16	29.04	0.56	1.98	0.6	1.50e-02	80	1.46e-04	1
V* V1005 Ori	2015-09-28	32.88	0.91	1.27	0.8	4.23e-02	64	5.84e-05	1
* 51 Eri	2016-12-13	45.03	0.84	2.92	0.66	6.70e-03	72	1.05e-04	2
* 51 Eri	2015-12-26	37.56	1.18	3.13	0.79	1.80e-03	256	2.01e-04	2
* L01 Pup	2016-03-30	47.7	0.93	4.75	0.9	2.30e-03	80	3.39e-05	2
* bet Pic	2016-10-14	53.82	0.75	9.38	0.67	3.20e-03	280	1.40e-04	2
* bet03 Tuc	2016-09-16	31.53	0.54	7.43	0.69	7.40e-03	80	3.45e-05	2
* e Vir	2017-02-10	31.19	0.78	7.83	0.84	2.60e-03	112	4.64e-05	2
* eps Hyi	2016-09-16	29.06	0.42	7.65	0.7	8.40e-03	160	2.63e-05	2
* eps Pav	2015-07-01	21.3	0.32	0.88	0.71	1.20e-03	128	1.48e-04	2
* eps Pav	2015-05-12	22.34	0.99	4.43	0.7	3.20e-03	256	4.45e-05	2
* eta Cha	2016-01-02	24.17	0.74	3.28	0.88	4.30e-03	80	6.32e-05	2
* eta Cha	2016-12-13	28.47	0.69	7.88	0.65	4.10e-03	90	8.20e-05	2
* eta Tel	2015-05-05	46.83	1.11	3.58	0.7	1.10e-03	168	5.44e-05	2
* eta Tuc	2015-10-27	25.16	1.16	2.95	0.82	1.30e-03	128	1.25e-04	2
* q Vel	2015-05-10	27.31	0.83	8.28	0.66	3.50e-03	128	5.48e-05	2
* rho Aql	2016-06-11	35.7	0.64	2.92	0.63	3.50e-03	96	1.53e-04	2
* rho Vir	2016-06-11	30.23	0.71	5.25	0.64	2.30e-03	72	1.33e-04	2

Table A.1: continued.

Star	Date	Δ PA (°)	S (″)	WS (m/s)	SR	CT (ms)	F	RC	C
1RXS J070659.0-535347	2017-02-11	38.97	0.52	10.2	0.69	1.26e-02	80	1.76e-04	2
ASAS J193304+0345.7	2018-06-19	34.68	0.88	9.1	0.64	4.30e-03	48	4.27e-04	2
ASAS J193304+0345.7	2016-05-23	35.93	0.58	8.3	0.52	7.10e-03	70	3.18e-04	2
ASAS J205428+0906.1	2018-05-13	22.67	0.88	7.35	0.6	4.00e-03	41	3.21e-04	2
BD+01 2447	2016-01-01	34.56	0.72	4.97	0.84	3.50e-03	64	7.97e-05	2
BD+02 1729	2015-12-01	32.85	1.0	6.1	0.73	4.90e-03	64	1.75e-04	2
BD+20 1790	2015-02-07	27.22	1.04	9.0	0.69	3.50e-03	76	9.39e-05	2
BD+20 1790	2017-02-10	27.3	0.55	4.3	0.69	7.00e-03	80	1.73e-04	2
BD-12 243	2015-12-02	50.97	1.1	2.38	0.72	9.40e-03	64	1.45e-04	2
BPS CS 22881-0003	2016-06-11	49.52	0.55	2.53	0.55	5.50e-03	49	3.94e-04	2
CD-40 2458	2016-12-13	63.99	0.59	5.9	0.65	6.10e-03	80	8.30e-05	2
CD-46 1064	2015-12-28	42.11	0.89	5.65	0.79	3.80e-03	64	8.09e-05	2
CD-46 3212	2018-01-26	43.06	0.93	8.78	0.81	3.80e-03	47	4.32e-04	2
CD-46 3212	2015-11-29	58.57	0.95	4.2	0.73	8.90e-03	80	1.66e-04	2
CD-48 2972	2015-11-30	36.79	1.28	5.4	0.65	5.20e-03	64	1.88e-04	2
CD-51 10924	2015-05-06	34.54	1.19	1.83	0.62	1.40e-03	256	1.10e-04	2
CD-52 381	2018-12-17	48.53	0.88	7.97	0.59	4.40e-03	58	4.29e-04	2
CD-52 9381	2015-05-09	34.25	0.72	5.03	0.67	2.20e-03	64	2.29e-04	2
CD-52 9381	2015-05-30	31.02	0.88	6.9	0.55	1.50e-03	64	1.91e-04	2
CD-54 7336	2015-06-03	30.79	0.99	5.68	0.8	5.60e-03	64	8.35e-05	2
CD-54 8168	2015-06-27	39.3	0.58	6.0	0.69	3.50e-03	80	8.30e-05	2

Table A.1: continued.

Star	Date	Δ PA (°)	S (″)	WS (m/s)	SR	CT (ms)	F	RC	C
CD-57 1054	2016-01-03	28.93	0.75	1.85	0.74	3.10e-03	64	1.07e-04	2
CD-61 1439	2016-01-02	26.06	0.9	1.58	0.76	1.50e-03	64	1.35e-04	2
CPD-72 2713	2016-06-27	42.72	0.92	5.95	0.64	4.10e-03	115	2.11e-04	2
EUVE J0625-60.0	2015-02-07	33.13	0.71	4.08	0.66	7.80e-03	76	1.03e-04	2
HD 377	2016-10-14	35.75	0.63	6.03	0.68	2.00e-03	80	8.86e-05	2
HD 987	2016-10-15	29.68	0.71	7.65	0.63	2.30e-03	90	2.90e-04	2
HD 1466	2015-10-26	25.0	1.07	1.77	0.6	1.10e-03	64	3.92e-04	2
HD 8558	2016-10-14	28.16	0.56	6.03	0.63	3.10e-03	64	2.95e-04	2
HD 8558	2015-10-28	28.22	1.02	1.7	0.6	1.90e-03	64	1.76e-04	2
HD 15115	2015-10-26	29.64	0.82	3.2	0.69	1.90e-03	64	2.35e-04	2
HD 24636	2016-01-03	25.51	0.69	4.03	0.81	3.50e-03	80	1.02e-04	2
HD 35650	2015-12-28	56.67	0.82	3.45	0.81	2.90e-03	64	6.16e-05	2
HD 44627	2015-02-06	28.28	1.08	3.25	0.68	5.70e-03	64	9.15e-05	2
HD 49855	2015-12-28	21.38	0.84	8.03	0.71	2.50e-03	64	1.19e-04	2
HD 51797	2016-01-03	39.53	0.55	1.2	0.77	4.90e-03	64	1.15e-04	2
HD 55279	2017-02-08	32.24	0.6	4.53	0.73	6.90e-03	78	1.59e-04	2
HD 74576	2016-03-28	42.58	1.08	8.35	0.87	3.60e-03	64	5.05e-05	2
HD 75505	2015-02-06	30.42	1.2	7.33	0.63	2.30e-03	84	2.48e-04	2
HD 84075	2017-02-11	24.41	0.51	11.0	0.68	4.00e-03	80	1.07e-04	2
HD 95086	2015-05-11	22.48	1.08	2.4	0.61	4.20e-03	64	9.38e-05	2
HD 95086	2015-05-05	22.44	0.74	5.5	0.66	2.70e-03	64	3.60e-04	2

Table A.1: continued.

Star	Date	Δ PA (°)	S (″)	WS (m/s)	SR	CT (ms)	F	RC	C
HD 96064	2017-02-09	49.86	0.45	5.88	0.83	4.00e-03	58	4.96e-05	2
HD 100453	2016-01-21	7.63	0.91	4.28	0.81	4.30e-03	176	8.00e-05	2
HD 101615	2016-06-02	28.88	0.53	6.03	0.68	4.20e-03	48	1.21e-04	2
HD 103599	2015-06-02	30.96	1.17	3.53	0.82	2.40e-03	64	1.47e-04	2
HD 104125	2015-06-03	28.55	1.17	5.8	0.88	2.60e-03	64	1.81e-04	2
HD 113457	2016-04-17	30.42	0.83	9.43	0.75	3.40e-03	80	1.21e-04	2
HD 115470	2016-04-01	44.7	0.84	3.45	0.83	2.30e-03	64	3.85e-05	2
HD 115600	2015-06-04	27.32	1.15	3.9	0.75	2.40e-03	64	1.44e-04	2
HD 116434	2016-06-26	42.18	0.69	2.23	0.66	3.10e-03	80	2.57e-04	2
HD 116434	2017-02-07	44.27	0.58	8.8	0.84	5.80e-03	80	5.23e-05	2
HD 116434	2016-05-31	34.22	0.84	7.35	0.59	2.10e-03	64	1.87e-04	2
HD 118100	2016-06-27	52.04	0.8	11.05	0.64	2.80e-03	64	1.42e-04	2
HD 118588	2015-05-15	33.65	0.56	6.13	0.71	8.60e-03	48	3.39e-05	2
HD 118697	2016-06-02	38.29	0.44	3.33	0.61	3.40e-03	80	2.10e-04	2
HD 122705	2016-03-30	27.73	0.86	5.18	0.77	3.40e-03	40	9.06e-05	2
HD 143488	2016-06-11	67.97	0.76	3.38	0.65	5.60e-03	65	2.30e-04	2
HD 143488	2015-05-09	36.43	0.9	8.13	0.67	1.90e-03	48	2.58e-04	2
HD 159911	2015-05-12	65.29	0.83	2.6	0.69	2.70e-03	64	5.11e-05	2
HD 159911	2016-04-16	54.59	0.54	7.45	0.76	6.80e-03	80	6.21e-05	2
HD 160305	2015-05-13	34.6	0.94	2.1	0.68	4.50e-03	64	4.20e-05	2
HD 164249	2018-04-11	31.95	0.52	7.75	0.97	5.50e-03	40	2.30e-04	2

Table A.1: continued.

Star	Date	Δ PA (°)	S (″)	WS (m/s)	SR	CT (ms)	F	RC	C
HD 172555	2015-05-11	19.3	0.79	8.13	0.69	3.90e-03	192	3.59e-05	2
HD 175726	2016-09-16	38.4	0.53	6.5	0.68	4.20e-03	65	5.36e-05	2
HD 175726	2016-06-12	31.37	0.88	7.33	0.66	3.20e-03	64	1.36e-04	2
HD 189245	2015-05-14	16.18	0.75	3.48	0.72	6.50e-03	128	2.47e-05	2
HD 189285	2016-05-29	58.98	0.64	5.08	0.55	2.50e-03	72	4.78e-04	2
HD 189285	2018-05-07	50.27	0.77	7.55	0.68	2.60e-03	50	1.15e-04	2
HD 189285	2015-06-03	40.87	0.98	4.53	0.81	4.40e-03	64	7.10e-05	2
HD 202917	2015-05-31	49.56	1.25	3.6	0.67	1.10e-03	86	2.43e-04	2
HD 202917	2016-05-31	32.68	0.83	9.93	0.52	2.70e-03	64	2.54e-04	2
HD 218396	2018-06-19	34.4	0.63	5.03	0.86	7.40e-03	64	1.23e-04	2
HD 218396	2016-11-18	16.84	0.98	5.58	0.65	4.90e-03	64	2.18e-04	2
HD 218860	2015-09-30	43.75	0.68	1.85	0.76	3.00e-03	64	6.56e-05	2
HD 219482	2015-10-01	25.64	0.76	4.03	0.81	2.00e-03	240	8.21e-05	2
HD 326277	2016-03-30	33.62	1.18	6.1	0.7	2.20e-03	45	1.08e-04	2
LP 745-70	2015-05-11	53.75	0.57	7.0	0.6	5.20e-03	64	8.64e-05	2
SCR J0017-6645	2016-09-17	29.37	0.63	1.95	0.51	1.49e-02	76	1.51e-04	2
V* CC Phe	2015-09-30	46.35	0.98	3.1	0.74	3.00e-03	160	1.03e-04	2
V* CC Phe	2017-09-29	37.0	0.43	2.8	0.68	5.20e-03	40	1.35e-04	2
V* CT Tuc	2015-11-29	25.78	1.57	4.33	0.65	5.50e-03	64	2.65e-04	2
V* NZ Lup	2016-05-31	46.64	0.92	9.07	0.61	2.70e-03	64	2.39e-04	2
V* PZ Tel	2014-07-15	16.27	0.92	5.88	0.66	2.10e-03	141	1.69e-04	2

Table A.1: continued.

Star	Date	Δ PA (°)	S (″)	WS (m/s)	SR	CT (ms)	F	RC	C
V* PZ Tel	2017-05-18	33.15	0.81	4.72	0.64	2.90e-03	60	3.20e-04	2
V* PZ Tel	2015-05-06	11.41	1.43	1.77	0.61	2.10e-03	82	2.94e-04	2
V* V1358 Ori	2017-02-09	36.47	0.63	6.08	0.83	5.00e-03	56	6.38e-05	2
V* V1358 Ori	2015-10-28	48.54	1.01	4.3	0.78	2.10e-03	80	9.28e-05	2
V* WW PsA	2015-10-01	54.73	0.65	4.4	0.56	2.20e-03	128	2.15e-04	2
* alf PsA	2015-06-30	8.7	0.94	4.7	0.82	1.30e-03	380	1.45e-04	3
* bet Leo	2015-05-31	14.37	1.11	6.15	0.79	1.10e-03	800	1.16e-04	3
* bet Pic	2016-04-16	18.06	0.79	5.97	0.8	6.20e-03	462	9.38e-05	3
* bet Pic	2015-12-26	37.33	1.44	4.93	0.88	2.60e-03	490	1.23e-04	3
* bet Pic	2015-10-01	39.95	1.42	5.83	0.76	1.00e-03	704	2.04e-04	3
* bet Pic	2016-01-20	29.72	1.14	8.35	0.9	2.80e-03	660	1.00e-04	3
* bet Pic	2018-09-12	26.12	0.65	1.52	0.87	5.50e-03	768	1.27e-04	3
* bet03 Tuc	2015-09-28	28.04	0.9	3.88	0.87	8.00e-03	512	9.95e-05	3
* e Vir	2015-06-03	28.86	1.3	2.05	0.9	4.60e-03	256	7.15e-05	3
* eps Hyi	2015-10-01	34.62	1.06	2.28	0.76	1.40e-03	608	1.55e-04	3
* kap Psc	2015-09-25	41.34	1.32	6.43	0.82	4.20e-03	300	1.23e-04	3
* phi Eri	2015-09-26	36.71	0.9	3.6	0.88	1.15e-02	512	5.80e-05	3
* q Vel	2018-01-05	44.4	0.55	1.33	0.9	1.00e-02	675	7.46e-05	3
* rho Aql	2015-09-28	25.22	1.51	2.13	0.88	8.80e-03	256	9.48e-05	3

Table A.1: continued.

Star	Date	Δ PA (°)	S (″)	WS (m/s)	SR	CT (ms)	F	RC	C
* zet Aql	2015-06-28	29.88	0.41	6.83	0.79	1.20e-03	560	1.62e-04	3
HD 100546	2015-05-29	34.73	0.84	8.32	0.75	1.60e-03	384	2.16e-04	3
V* HN Peg	2015-09-30	25.69	0.81	6.78	0.89	3.20e-03	256	2.84e-05	3
V* TW PsA	2015-06-27	100.73	0.91	2.8	0.82	2.10e-03	512	6.84e-05	3
* 30 Mon	2016-01-19	40.43	1.8	7.75	0.86	1.30e-03	448	1.79e-03	4
* pi. Ara	2016-06-11	25.92	0.62	3.2	0.65	3.30e-03	95	9.03e-04	4
CD-52 381	2016-01-17	14.32	1.77	0.82	0.72	2.60e-03	40	6.41e-04	4
CD-81 812	2016-06-28	16.63	1.83	13.6	0.56	1.40e-03	56	1.11e-03	4
Smethells 20	2015-05-16	25.88	1.32	13.1	0.61	1.60e-03	64	1.02e-03	4
V* BO Mic	2015-06-04	61.52	1.25	7.2	0.77	2.60e-03	64	1.36e-03	4
V* CE Ant	2015-05-09	74.0	1.31	6.38	0.56	2.10e-03	60	7.05e-04	4
* 18 Boo	2016-06-12	25.97	1.03	15.78	0.65	1.90e-03	128	1.55e-04	5
* 30 Ser	2015-06-05	39.42	1.54	11.53	0.84	1.40e-03	64	2.72e-04	5
* L01 Pup	2015-12-29	34.71	1.96	14.8	0.75	7.00e-04	128	3.21e-04	5
* alf CrA	2015-06-06	49.74	1.29	12.82	0.89	2.10e-03	256	7.67e-05	5
* alf PsA	2015-09-24	48.23	1.55	15.13	0.71	1.20e-03	512	3.03e-04	5
* alf Sgr	2018-05-08	26.66	1.04	12.13	0.76	1.60e-03	70	2.03e-04	5

Table A.1: continued.

Star	Date	Δ PA (°)	S (″)	WS (m/s)	SR	CT (ms)	F	RC	C
* bet Cir	2015-03-30	25.86	1.68	9.25	0.66	2.30e-03	224	1.09e-04	5
* chi Cap	2015-09-29	80.57	1.91	9.9	0.76	2.10e-03	256	5.29e-04	5
* e Vir	2016-03-29	31.24	1.1	10.2	0.89	1.40e-03	104	6.49e-05	5
* eta Tuc	2016-06-28	27.33	1.08	14.53	0.67	1.90e-03	240	8.57e-05	5
ASAS J205428+0906.1	2015-06-06	25.88	1.52	12.55	0.67	1.50e-03	64	2.07e-04	5
BD-03 4778	2015-05-15	40.35	1.69	11.3	0.67	1.80e-03	64	1.51e-04	5
CD-51 10924	2016-05-31	37.87	1.05	11.78	0.67	1.90e-03	70	1.94e-04	5
HD 25284	2015-09-30	38.75	1.34	6.85	0.71	1.30e-03	128	7.72e-05	5
HD 30447	2015-12-29	16.8	1.43	10.68	0.66	1.30e-03	96	1.78e-04	5
HD 96064	2016-04-21	42.87	1.57	10.4	0.63	1.40e-03	64	2.45e-04	5
HD 96338	2015-02-06	39.24	1.7	9.43	0.68	2.10e-03	64	1.59e-04	5
HD 104125	2016-03-30	28.4	1.34	6.95	0.78	2.50e-03	52	2.80e-04	5
HD 107301	2015-06-05	22.12	1.41	11.18	0.85	2.00e-03	60	1.22e-04	5
HD 107301	2016-01-19	20.52	1.85	9.05	0.67	1.50e-03	128	5.05e-04	5
HD 107821	2019-03-10	28.02	0.99	12.98	0.68	2.30e-03	48	1.26e-04	5
HD 109536	2016-04-21	55.0	1.63	10.6	0.72	1.70e-03	216	3.17e-04	5
HD 128819	2016-04-17	49.99	0.91	12.45	0.75	3.20e-03	64	1.23e-04	5
HD 138204	2016-03-29	55.93	1.33	13.73	0.81	1.90e-03	64	8.65e-05	5
HD 148055	2016-03-29	54.28	1.67	11.88	0.59	1.20e-03	64	2.55e-04	5
HD 151798	2016-06-12	58.94	1.11	10.35	0.65	2.80e-03	64	2.44e-04	5
HD 152391	2015-06-05	35.04	1.66	14.05	0.83	1.50e-03	64	-1.58e-06	5

Table A.1: continued.

Star	Date	Δ PA (°)	S (″)	WS (m/s)	SR	CT (ms)	F	RC	C
HD 156097	2016-06-28	94.11	1.52	15.68	0.6	2.10e-03	56	3.93e-04	5
HD 164249	2016-04-17	37.29	1.72	13.73	0.63	2.00e-03	61	5.72e-04	5
HD 164249	2019-09-08	40.61	1.1	12.75	0.68	2.20e-03	51	1.92e-04	5
HD 164249	2015-05-10	34.32	1.87	2.63	0.69	1.20e-03	56	1.69e-04	5
HD 218396	2017-06-15	19.28	0.92	11.45	0.92	2.70e-03	144	1.50e-04	5
HD 224228	2015-10-25	59.29	1.48	12.03	0.64	1.30e-03	64	1.99e-04	5
NAME Proxima Centauri	2016-03-28	25.58	1.95	9.95	0.7	1.50e-03	200	1.24e-04	5
* alf Sgr	2016-05-21	71.49	1.26	7.72	0.86	2.10e-03	176	8.80e-05	6
* del Crv B	2016-06-11	99.26	0.84	14.07	0.66	3.10e-03	80	1.99e-04	6
* mu. PsA	2016-10-14	84.12	0.74	5.3	0.69	3.30e-03	80	1.22e-04	6
* zet Lep	2017-02-11	83.53	0.63	6.68	0.91	6.90e-03	288	3.43e-05	6
ASAS J200136-3313.0	2015-09-30	87.66	1.0	6.75	0.71	2.90e-03	64	1.72e-04	6
CD-31 16041	2015-05-15	81.15	1.1	9.68	0.68	2.90e-03	64	6.03e-05	6
CD-35 1167	2016-01-02	67.81	0.72	4.7	0.78	4.10e-03	64	1.46e-04	6
CD-52 381	2015-09-25	81.54	0.9	8.18	0.58	4.40e-03	160	2.50e-04	6
HD 17925	2016-10-15	70.8	0.66	6.05	0.67	2.70e-03	160	1.50e-04	6
HD 19545	2016-10-13	145.16	1.4	9.23	0.68	1.60e-03	90	1.96e-04	6
HD 77825	2016-04-17	78.33	0.64	6.85	0.73	3.60e-03	57	1.43e-04	6
HD 131835	2015-05-14	72.66	0.59	5.47	0.7	9.00e-03	64	4.21e-05	6

Table A.1: continued.

Star	Date	Δ PA (°)	S (″)	WS (m/s)	SR	CT (ms)	F	RC	C
HD 142851	2015-06-03	91.84	1.03	7.68	0.83	3.20e-03	64	7.72e-05	6
HD 143567	2015-05-04	147.95	0.65	4.22	0.68	2.70e-03	60	1.21e-04	6
HD 143637	2017-03-20	115.97	0.77	4.65	0.76	3.10e-03	67	1.10e-04	6
HD 143637	2015-06-04	123.56	1.04	8.3	0.81	3.10e-03	68	1.12e-04	6
HD 147553	2015-05-30	88.85	0.72	7.68	0.71	1.60e-03	136	6.83e-04	6
HD 151726	2016-04-01	78.55	1.06	6.0	0.8	2.40e-03	80	6.35e-05	6
HD 151726	2017-05-18	64.81	0.92	11.6	0.67	2.30e-03	68	2.44e-04	6
HD 151726	2019-05-20	76.33	0.98	12.68	0.71	2.50e-03	57	1.32e-04	6
HD 159911	2019-04-15	67.51	0.41	3.73	0.77	5.60e-03	43	7.57e-05	6
HD 159911	2016-05-22	71.48	0.48	7.25	0.78	5.90e-03	70	6.03e-05	6
HD 168210	2017-06-03	115.43	0.44	6.22	0.76	5.00e-03	74	8.27e-05	6
HD 176367	2015-06-07	85.15	1.16	8.65	0.69	3.10e-03	80	9.41e-05	6
HD 201919	2015-06-03	92.85	0.98	5.75	0.83	4.10e-03	64	8.42e-05	6
HD 317617	2015-05-31	86.69	1.07	8.75	0.65	1.20e-03	64	2.10e-04	6
PM J19560-3207	2015-05-04	81.53	0.76	3.48	0.64	1.70e-03	64	3.27e-04	6
V* AS Col	2017-02-12	72.94	1.03	5.85	0.76	5.40e-03	96	1.60e-04	6
V* BX Ant	2015-05-05	73.22	0.79	2.55	0.61	1.70e-03	64	1.90e-04	6
V* BX Ant	2016-03-30	91.82	1.04	4.58	0.62	2.90e-03	60	1.35e-04	6
V* CE Ant	2017-02-07	90.94	0.65	8.0	0.75	4.20e-03	163	1.03e-04	6
V* KU Lib	2016-06-01	69.38	0.84	8.48	0.66	2.30e-03	96	2.14e-04	6
V* NZ Lup	2017-04-30	60.37	0.52	8.53	0.81	3.70e-03	80	5.34e-05	6

Appendix B

Variation of optimal RSM parameters across cluster centers and individual datasets

This section examines the impact of clustering on optimal parameter selection in RSM, emphasizing how different observing conditions influence the choice of PSF subtraction technique parameters and assessing the reliability of generalizing cluster center parameters to all datasets within a cluster. To quantify these effects, we measure the normalized distance between parameters. For numerical parameters, we compute their normalized distances, while for categorical parameters (e.g., the noise estimation methods), we use the dissimilarity index, a normalized measure of how different two RSM parameter sets are, with 0 indicating identical configurations and 1 indicating distinct ones. The mean normalized distance across all parameters is then calculated for each cluster center pair. Figure B.1 (top) presents these distances, revealing substantial variability in PSF subtraction parameters, with a median distance of 0.47, while RSM parameters exhibit lower variability, with a median distance of 0.1. This indicates significant differences in optimal parameter choices between cluster centers for PSF subtraction techniques. Overall, different PSF subtraction methods exhibit varying median distances across the clusters, with APCA having the lowest median distance of 0.33 and LOCI the highest at 0.68.

To assess the validity of generalizing parameters from a single representative dataset to an entire cluster, we conducted the same analysis within a small but diverse cluster (Cluster 1). Figure B.1 (bottom) illustrates the normalized distances between 20 datasets and the cluster center, showing a reduced median distance of 0.33 for PSF subtraction parameters. One dataset (HD 109573) failed to converge during optimization due to the presence of bright disk emission, explaining the use of 20 instead of 21 datasets.

B Variation of optimal RSM parameters across cluster centers and individual datasets

The slightly larger distance of 0.15 observed for RSM parameters primarily stems from variations in the noise estimation region parameter, which has minimal impact on the final RSM maps. In contrast, PSF subtraction parameters have a much stronger influence on the results. The reduction in parameter distance suggests that datasets with similar observing conditions tend to converge toward more comparable parameter choices.

Overall, this analysis underscores the substantial differences in optimal parameter selection between different cluster centers, highlighting the need to adapt parameter choices to specific datasets across a survey. At the same time, it demonstrates that datasets with similar observing conditions generally exhibit closer parameter distances. However, some datasets within a single cluster still show notable variations, reflecting dataset-specific noise properties. Despite these differences, we adopt a cluster-based parameter optimization approach in this study, as optimizing RSM parameters for each individual dataset remains computationally costly.





Figure B.1: Normalized distances for PSF subtraction (blue) and RSM parameters (orange) between cluster centers (top) and datasets within a single cluster (bottom).


Appendix C

Comparison with the Dahlqvist contrast curve definition

This section compares the contrast curve methodology used in Dahlqvist et al. (2022) for the SHARDDS survey with the approach applied in our study. Dahlqvist et al. (2022) define the detection threshold based on the brightest false positive within the field of view and considers 95% recovery of injected signals. In contrast, our approach employs a detection threshold corresponding to an FAP of 3×10^{-7} under a lognormal distribution at a given separation and evaluates contrast at the 50% injection recovery level. Figure C.1 presents these contrast curves for two datasets: one without a potential signal (Fig. C.1a) and another with a potential signal (Fig. C.1b). In Fig. C.1a, the contrast curves derived from different detection thresholds remain consistent at the same completeness levels. However, this trend does not hold in Fig. C.1b, where the contrast curve based on the lognormal detection threshold outperforms that derived from the first false positive at both completeness levels. In both figures, there is only a minor gap between the contrast levels required for 50% and 95% signal recovery. This result highlights a key limitation of the approach in Dahlqvist et al. (2022), where the contrast is influenced by the brightest signal in the field of view without determining whether it originates from noise or a potential candidate—a limitation that also applies to thresholds based on the maximum F1 score. Moreover, this method defines the detection threshold based on the entire field of view rather than at each separation, limiting its ability to assess contrast variations across different separations.



(a)



(b)

Figure C.1: Comparison of contrast curves for two datasets, one with no potential companion (a) and one with a potential companion (b). Contrasts are generated with 95% completeness using the first false positive threshold and lognormal threshold (in purple and green), and 50% completeness for both thresholds (in red and blue)

Appendix D

Comparing RSM noise distribution models

The distribution of independent RSM noise realizations, as discussed in Section 4.1.1, exhibits a near-normal profile when viewed on a logarithmic scale. Here, we investigate various normal-like distributions to determine the best fit for the logarithm of the datasets. Specifically, we compare the log of lognormal distribution (three parameters), the normal distribution (two parameters), the power-normal distribution (three parameters), the generalized normal distribution (three parameters), along with the Laplace distribution (two parameters). Figure D.1 presents the distribution of independent RSM noise realizations using APCA in the RSM context for cluster 1 at two separations, along with the fitted distributions. The histograms demonstrate a clear normal-like profile, with all tested distributions providing a reasonable fit, with the laplace being the least appropriate.

To assess the goodness of fit, we perform a Kolmogorov-Smirnov (KS) test and compute the p-values for each distribution at different separations. Table D.1 presents the results for Cluster 1, where RSM noise realizations were generated using APCA in the RSM detection map. The table shows that the generalized normal, the logarithm of the lognormal, and the power-normal distributions achieve p-values above 0.05, indicating a good

Table D.1: Kolmogorov-Smirnov test for different distributions fits for cluster 1 RSM noise realization samples: p-value at $0''.25$ and $1''$ using RSM APCA.

Distributions	sep = $0''.25$	sep = $1''$
Log of Lognorm	0.5	0.09
Norm	0.72	0.00
Generalized norm	0.43	0.71
Power norm	0.56	0.17
Laplace	0.04	0.00

fit. While the generalized normal and power-normal distributions yield high p-values, they do not consistently perform well across all clusters (e.g., failing for Cluster 6). Given its strong overall fit across different clusters and PSF subtraction technique combinations, we adopt the lognormal distribution as the preferred model.




Figure D.1: Histograms of independent RSM noise realizations (log scale) in the RSM maps of Cluster 1, obtained using APCA RSM, at $0''.25$ (top) and at $1''$ (bottom). The fitted distributions include normal (blue), generalized normal (orange), power-normal (green), Laplace (red), and log of lognormal (purple).

Appendix E

Maximum F1 score threshold generalized across a cluster

This section evaluates the generalization of thresholds determined using the maximum F1 score from a given dataset across other datasets with similar observing conditions. The maximum F1 score is highly sensitive to false positives, often influenced by the brightest speckle in the representative dataset rather than the overall noise profile. Generalizing such thresholds can lead to inconsistencies, as the presence of these speckles reflects specific environmental parameters and their interaction with adaptive optics. These factors are inherently unique to individual observations and cannot be reliably generalized, even under broadly similar observing conditions. Figure E.1 illustrates this issue by generalizing the maximum F1 score threshold from the cluster center BD-15 705 to the dataset HIP 92680 in cluster 1. The threshold derived from the cluster center (pink) shows a factor of ten difference compared to the self-computed maximum F1 score (cyan) and the lognormal threshold (dark blue). This significant discrepancy leads to the generalized threshold accepting false positives that would otherwise be excluded, underscoring the difficulties in applying overly tuned thresholds across datasets.




Figure E.1: Maximum F1 score threshold for the cluster center 1 BD-15 705 (in pink) applied on the HIP 92680 dataset, in comparison to the maximum F1 score computed by the dataset HIP 92680 itself denoted as "self" (in cyan) and the lognormal threshold (in dark blue) done by the cluster 1, at separation 0''55.

Bibliography

- Absil, O., Milli, J., Mawet, D., et al. 2013, *A&A*, 559, L12, doi: 10.1051/0004-6361/201322748
- Amara, A., & Quanz, S. P. 2012, *MNRAS*, 427, 948, doi: 10.1111/j.1365-2966.2012.21918.x
- Andrews, S. M. 2020, *ARA&A*, 58, 483, doi: 10.1146/annurev-astro-031220-010302
- Apai, D., Kasper, M., Skemer, A., et al. 2016, *The Astrophysical Journal*, 820, 40, doi: 10.3847/0004-637X/820/1/40
- Backman, D., Marengo, M., Stapelfeldt, K., et al. 2009, *ApJ*, 690, 1522, doi: 10.1088/0004-637X/690/2/1522
- Bailey, V., Meshkat, T., Reiter, M., et al. 2014, *ApJ*, 780, L4, doi: 10.1088/2041-8205/780/1/L4
- Baines, E. K., White, R. J., Huber, D., et al. 2012, *ApJ*, 761, 57, doi: 10.1088/0004-637X/761/1/57
- Balmer, W. O., Otten, G., & Stolker, T. 2025a, backtracks: a python package to compare relative astrometry with background helical motion, v0.6.0, Zenodo, doi: 10.5281/zenodo.14838370
- Balmer, W. O., Otten, G. P. P. L., & Stolker, T. 2025b, backtracks: a python package to compare relative astrometry with background helical motion: v0.6, v0.6, Zenodo, doi: 10.5281/zenodo.14838369
- Barman, T. S., Macintosh, B., Konopacky, Q. M., & Marois, C. 2011, *ApJ*, 733, 65, doi: 10.1088/0004-637X/733/1/65
- Benedict, G. F., McArthur, B. E., Gatewood, G., et al. 2006, *AJ*, 132, 2206, doi: 10.1086/508323
- Benisty, M., Bae, J., Facchini, S., et al. 2021, *ApJ*, 916, L2, doi: 10.3847/2041-8213/ac0f83
- Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, *A&A*, 631, A155, doi: 10.1051/0004-6361/201935251
- Biller, B. A., Liu, M. C., Wahhaj, Z., et al. 2010, *ApJ*, 720, L82, doi: 10.1088/2041-8205/720/1/L82
- . 2013, *ApJ*, 777, 160, doi: 10.1088/0004-637X/777/2/160

- Blunt, S., Wang, J. J., Angelo, I., et al. 2020, *AJ*, 159, 89, doi: 10.3847/1538-3881/ab6663
- Blunt, S., Wang, J., Hirsch, L., et al. 2024, *The Journal of Open Source Software*, 9, 6756, doi: 10.21105/joss.06756
- Boccaletti, A., Mâlin, M., Baudoz, P., et al. 2024, *A&A*, 686, A33, doi: 10.1051/0004-6361/202347912
- Bonnefoy, M., Perraut, K., Lagrange, A.-M., et al. 2018, *A&A*, 618, A63, doi: 10.1051/0004-6361/201832942
- Bonse, M. J., Garvin, E. O., Gebhard, T. D., et al. 2023, *AJ*, 166, 71, doi: 10.3847/1538-3881/acc93c
- Bonse, M. J., Gebhard, T. D., Dannert, F. A., et al. 2025, *AJ*, 169, 194, doi: 10.3847/1538-3881/adab79
- Booth, M., Dent, W. R. F., Jordán, A., et al. 2017, *MNRAS*, 469, 3200, doi: 10.1093/mnras/stx1072
- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Science*, 327, 977, doi: 10.1126/science.1185402
- Boss, A. P. 1998, *ApJ*, 503, 923, doi: 10.1086/306036
- Bowler, B. P. 2016, *PASP*, 128, 102001, doi: 10.1088/1538-3873/128/968/102001
- Bowler, B. P., Blunt, S. C., & Nielsen, E. L. 2020, *AJ*, 159, 63, doi: 10.3847/1538-3881/ab5b11
- Bowler, B. P., & Nielsen, E. L. 2018, in *Handbook of Exoplanets*, ed. H. J. Deeg & J. A. Belmonte, 155, doi: 10.1007/978-3-319-55333-7_155
- Brandt, G. M., Brandt, T. D., Dupuy, T. J., Li, Y., & Michalik, D. 2021, *AJ*, 161, 179, doi: 10.3847/1538-3881/abdc2e
- Brandt, T. D. 2021, *The Astrophysical Journal Supplement Series*, 254, 42, doi: 10.3847/1538-4365/abf93c
- Brown-Sevilla, S. B., Maire, A.-L., Mollière, P., et al. 2023, *A&A*, 673, A98, doi: 10.1051/0004-6361/202244826
- Cantalloube, F., Dohlen, K., Milli, J., Brandner, W., & Vigan, A. 2019, *The Messenger*, 176, 25, doi: 10.18727/0722-6691/5138
- Cantalloube, F., Mouillet, D., Mugnier, L. M., et al. 2015, *A&A*, 582, A89, doi: 10.1051/0004-6361/201425571
- Cantalloube, F., Por, E. H., Dohlen, K., et al. 2018, *Astronomy & Astrophysics*, 620, L10, doi: 10.1051/0004-6361/201834311
- Cantalloube, F., Farley, O. J. D., Milli, J., et al. 2020, *A&A*, 638, A98, doi: 10.1051/0004-6361/201937397

- Cantalloube, F., Gomez-Gonzalez, C., Absil, O., et al. 2021, arXiv e-prints, arXiv:2101.05080, doi: 10.48550/arXiv.2101.05080
- Cantero, C., Absil, O., Dahlqvist, C. H., & Van Droogenbroeck, M. 2023, *A&A*, 680, A86, doi: 10.1051/0004-6361/202346085
- Cantero, C., Sabalbal, M., Absil, O., et al. 2026, New substellar candidates identified through deep learning in the F150 sample of the large-scale SHINE direct imaging survey. <https://arxiv.org/abs/2605.23700>
- Carson, J., Thalmann, C., Janson, M., et al. 2013, *ApJ*, 763, L32, doi: 10.1088/2041-8205/763/2/L32
- Chauvin, G. 2018, in *Adaptive Optics Systems VI*, ed. L. M. Close, L. Schreiber, & D. Schmidt, Vol. 10703, International Society for Optics and Photonics (SPIE), 1070305, doi: 10.1117/12.2314273
- Chauvin, G., Lagrange, A.-M., Dumas, C., et al. 2005, *A&A*, 438, L25, doi: 10.1051/0004-6361:200500116
- Chauvin, G., Lagrange, A.-M., Bonavita, M., et al. 2010, *A&A*, 509, A52, doi: 10.1051/0004-6361/200911716
- Chauvin, G., Vigan, A., Bonnefoy, M., et al. 2015, *A&A*, 573, A127, doi: 10.1051/0004-6361/201423564
- Chauvin, G., Desidera, S., Lagrange, A. M., et al. 2017a, in *SF2A-2017: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, ed. C. Reyl  , P. Di Matteo, F. Herpin, E. Lagadec, A. Lan  on, Z. Meliani, & F. Royer, Di
- Chauvin, G., Desidera, S., Lagrange, A. M., et al. 2017b, *A&A*, 605, L9, doi: 10.1051/0004-6361/201731152
- Cheetham, A., Bonnefoy, M., Desidera, S., et al. 2018, *A&A*, 615, A160, doi: 10.1051/0004-6361/201832650
- Chomez, A., Lagrange, A.-M., Delorme, P., et al. 2023, *A&A*, 675, A205, doi: 10.1051/0004-6361/202245723
- Chomez, A., Delorme, P., Lagrange, A.-M., et al. 2025, *A&A*, 697, A99, doi: 10.1051/0004-6361/202451751
- Christiaens, V., Cantalloube, F., Casassus, S., et al. 2019, *ApJ*, 877, L33, doi: 10.3847/2041-8213/ab212b
- Christiaens, V., Gonzalez, C., Farkas, R., et al. 2023, *The Journal of Open Source Software*, 8, 4774, doi: 10.21105/joss.04774

- Claudi, R. U., Turatto, M., Gratton, R. G., et al. 2008, in *Ground-based and Airborne Instrumentation for Astronomy II*, ed. I. S. McLean & M. M. Casali, Vol. 7014, International Society for Optics and Photonics (SPIE), 70143E, doi: 10.1117/12.788366
- Cloutier, R. 2026, in *Encyclopedia of Astrophysics, Volume 1*, Vol. 1, 288–306, doi: 10.1016/B978-0-443-21439-4.00009-2
- Cumming, A., Butler, R. P., Marcy, G. W., et al. 2008, *PASP*, 120, 531, doi: 10.1086/588487
- Currie, T., Biller, B., Lagrange, A., et al. 2023a, in *Astronomical Society of the Pacific Conference Series*, Vol. 534, *Protostars and Planets VII*, ed. S. Inutsuka, Y. Aikawa, T. Muto, K. Tomida, & M. Tamura, 799, doi: 10.48550/arXiv.2205.05696
- Currie, T., Fukagawa, M., Thalmann, C., Matsumura, S., & Plavchan, P. 2012, *ApJ*, 755, L34, doi: 10.1088/2041-8205/755/2/L34
- Currie, T., Burrows, A., Itoh, Y., et al. 2011, *ApJ*, 729, 128, doi: 10.1088/0004-637X/729/2/128
- Currie, T., Brandt, T. D., Kuzuhara, M., et al. 2020, *ApJ*, 904, L25, doi: 10.3847/2041-8213/abc631
- Currie, T., Brandt, G. M., Brandt, T. D., et al. 2023b, *Science*, 380, 198, doi: 10.1126/science.abo6192
- Daglayan, Vary, Absil, et al. 2024, *An Alternating Minimization Algorithm with Trajectory for Direct Exoplanet Detection – The AMAT Algorithm*. <https://arxiv.org/abs/2410.06310>
- Daglayan, H., Vary, S., Absil, O., et al. 2024, *A&A*, 692, A126, doi: 10.1051/0004-6361/202451242
- Dahlqvist, C. H., Cantalloube, F., & Absil, O. 2020, *A&A*, 633, A95, doi: 10.1051/0004-6361/201936421
- . 2021a, *A&A*, 656, A54, doi: 10.1051/0004-6361/202141446
- Dahlqvist, C. H., Louppe, G., & Absil, O. 2021b, *A&A*, 646, A49, doi: 10.1051/0004-6361/202039597
- Dahlqvist, C. H., Milli, J., Absil, O., et al. 2022, *A&A*, 666, A33, doi: 10.1051/0004-6361/202244145
- Dawson, R. I., & Johnson, J. A. 2018, *Annual Review of Astronomy and Astrophysics*, 56, 175–221, doi: 10.1146/annurev-astro-081817-051853
- de Bruijne, J. H. J., & Eilers, A.-C. 2012, *A&A*, 546, A61, doi: 10.1051/0004-6361/201219219

- Delorme, P., Meunier, N., Albert, D., et al. 2017, in SF2A-2017: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics, ed. C. Reyl  , P. Di Matteo, F. Herpin, E. Lagadec, A. Lan  on, Z. Meliani, & F. Royer, 237, doi: 10.48550/arXiv.1712.06948
- Desgrange, C., Chauvin, G., Christiaens, V., et al. 2022, *A&A*, 664, A139, doi: 10.1051/0004-6361/202243097
- Desidera, S., Covino, E., Messina, S., et al. 2015, *A&A*, 573, A126, doi: 10.1051/0004-6361/201323168
- Desidera, S., Chauvin, G., Bonavita, M., et al. 2021, *A&A*, 651, A70, doi: 10.1051/0004-6361/202038806
- D  az, C. G., Gonz  lez, J. F., Levato, H., & Grosso, M. 2011, *A&A*, 531, A143, doi: 10.1051/0004-6361/201016386
- Dohlen, K., Langlois, M., Saisse, M., et al. 2008, in Ground-based and Airborne Instrumentation for Astronomy II, ed. I. S. McLean & M. M. Casali, Vol. 7014, International Society for Optics and Photonics (SPIE), 70143L, doi: 10.1117/12.789786
- Dressing, C. D., & Charbonneau, D. 2015, *ApJ*, 807, 45, doi: 10.1088/0004-637X/807/1/45
- Emsenhuber, A., Mordasini, C., Burn, R., et al. 2021, *A&A*, 656, A69, doi: 10.1051/0004-6361/202038553
- Fabrycky, D. C., & Murray-Clay, R. A. 2010, *ApJ*, 710, 1408, doi: 10.1088/0004-637X/710/2/1408
- Feldt, M., Olofsson, J., Boccaletti, A., et al. 2017, *A&A*, 601, A7, doi: 10.1051/0004-6361/201629261
- Feng, B. Y., Ferrer-Chvez, R., Levis, A., et al. 2025, *IEEE Transactions on Computational Imaging*, 11, 36, doi: 10.1109/TCI.2025.3525971
- Fernandes, R. B., Mulders, G. D., Pascucci, I., Mordasini, C., & Emsenhuber, A. 2019, *ApJ*, 874, 81, doi: 10.3847/1538-4357/ab0300
- Fitzgerald, M. P., & Graham, J. R. 2006, *ApJ*, 637, 541, doi: 10.1086/498339
- Flasseur, O., Denis, L., Thi  baut,   ., & Langlois, M. 2018, *A&A*, 618, A138, doi: 10.1051/0004-6361/201832745
- , 2020a, *A&A*, 634, A2, doi: 10.1051/0004-6361/201935859
- , 2020b, *A&A*, 637, A9, doi: 10.1051/0004-6361/201937239
- Follette, K. B. 2023, *PASP*, 135, 093001, doi: 10.1088/1538-3873/aceb31
- Forgan, D. H., Hall, C., Meru, F., & Rice, W. K. M. 2018, *MNRAS*, 474, 5036, doi: 10.1093/mnras/stx2870

- Fuhrmann, K. 2004, *Astronomische Nachrichten*, 325, 3, doi: 10.1002/asna.200310173
- Fukagawa, M., Itoh, Y., Tamura, M., et al. 2009, *ApJ*, 696, L1, doi: 10.1088/0004-637X/696/1/L1
- Fulton, B. J., Rosenthal, L. J., Hirsch, L. A., et al. 2021, *ApJS*, 255, 14, doi: 10.3847/1538-4365/abfcc1
- Fusco, T., Ageorges, N., Rousset, G., et al. 2004, in *Advancements in Adaptive Optics*, ed. D. B. Calia, B. L. Ellerbroek, & R. Ragazzoni, Vol. 5490, International Society for Optics and Photonics (SPIE), 118 – 129, doi: 10.1117/12.563009
- Fusco, T., Rousset, G., Sauvage, J.-F., et al. 2006, *Optics Express*, 14, 7515, doi: 10.1364/OE.14.007515
- Gaia Collaboration. 2020, *VizieR Online Data Catalog: Gaia EDR3* (Gaia Collaboration, 2020), *VizieR On-line Data Catalog: I/350*. Originally published in: 2021A&A...649A...1G, doi: 10.26093/cds/vizier.1350
- Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, *A&A*, 595, A1, doi: 10.1051/0004-6361/201629272
- Galicher, R., Marois, C., Macintosh, B., et al. 2016, *A&A*, 594, A63, doi: 10.1051/0004-6361/201527828
- Galicher, R., Boccaletti, A., Mesa, D., et al. 2018, *A&A*, 615, A92, doi: 10.1051/0004-6361/201832973
- Gardner, J. P., Mather, J. C., Clampin, M., et al. 2006, *Space Sci. Rev.*, 123, 485, doi: 10.1007/s11214-006-8315-7
- Gibbs, A., Ruffio, J.-B., Bidot, A., et al. 2026, *Discovery of an Exterior Third Planet Orbiting β Pictoris*. <https://arxiv.org/abs/2606.23789>
- Gillon, M., Jehin, E., Magain, P., et al. 2011, in *European Physical Journal Web of Conferences*, Vol. 11, *European Physical Journal Web of Conferences (EDP)*, 06002, doi: 10.1051/epjconf/20101106002
- Gillon, M., Triaud, A. H. M. J., Demory, B.-O., et al. 2017, *Nature*, 542, 456, doi: 10.1038/nature21360
- Gomez Gonzalez, C. A., Absil, O., Absil, P. A., et al. 2016, *A&A*, 589, A54, doi: 10.1051/0004-6361/201527387
- Gomez Gonzalez, C. A., Absil, O., & Van Droogenbroeck, M. 2018, *A&A*, 613, A71, doi: 10.1051/0004-6361/201731961
- Gomez Gonzalez, C. A., Wertz, O., Absil, O., et al. 2017, *AJ*, 154, 7, doi: 10.3847/1538-3881/aa73d7

- Goździewski, K., & Migaszewski, C. 2014, *MNRAS*, 440, 3140, doi: 10.1093/mnras/stu455
- . 2020, *ApJ*, 902, L40, doi: 10.3847/2041-8213/abb881
- Grandjean, A., Lagrange, A.-M., Meunier, N., et al. 2023, *A&A*, 669, A12, doi: 10.1051/0004-6361/202141235
- Greenbaum, A. Z., Pueyo, L., Ruffio, J.-B., et al. 2018, *AJ*, 155, 226, doi: 10.3847/1538-3881/aabcb8
- Grether, D., & Lineweaver, C. H. 2006, *ApJ*, 640, 1051, doi: 10.1086/500161
- Guerri, G., Daban, J.-B., Robbe-Dubois, S., et al. 2011, *Experimental Astronomy*, 30, 59, doi: 10.1007/s10686-011-9220-y
- Haffert, S. Y., Bohn, A. J., de Boer, J., et al. 2019, *Nature Astronomy*, 3, 749, doi: 10.1038/s41550-019-0780-5
- Hamilton, J. D. 1988, *Journal of Economic Dynamics and Control*, 12, 385, doi: None
- Hashimoto, J., Dong, R., Kudo, T., et al. 2012, *ApJ*, 758, L19, doi: 10.1088/2041-8205/758/1/L19
- Hatzes, A. P., Cochran, W. D., McArthur, B., et al. 2000, *ApJ*, 544, L145, doi: 10.1086/317319
- Hinkley, S., Oppenheimer, B. R., Soummer, R., et al. 2007, *ApJ*, 654, 633, doi: 10.1086/509063
- Howard, A. W., Marcy, G. W., Bryson, S. T., et al. 2012, *ApJS*, 201, 15, doi: 10.1088/0067-0049/201/2/15
- Hughes, A. M., Duchêne, G., & Matthews, B. C. 2018, *ARA&A*, 56, 541, doi: 10.1146/annurev-astro-081817-052035
- Janson, M., Quanz, S. P., Carson, J. C., et al. 2015, *A&A*, 574, A120, doi: 10.1051/0004-6361/201424944
- Janson, M., Brandner, W., Henning, T., et al. 2007, *AJ*, 133, 2442, doi: 10.1086/516632
- Janson, M., Squicciarini, V., Delorme, P., et al. 2021, *A&A*, 646, A164, doi: 10.1051/0004-6361/202039683
- Jensen-Clem, R., Mawet, D., Gomez Gonzalez, C. A., et al. 2018, *AJ*, 155, 19, doi: 10.3847/1538-3881/aa97e4
- Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, *PASP*, 127, 890, doi: 10.1086/682989
- Juillard, S., Christiaens, V., Absil, O., Stasevic, S., & Milli, J. 2024, *A&A*, 688, A185, doi: 10.1051/0004-6361/202449747
- Kalas, P., Graham, J. R., Chiang, E., et al. 2008, *Science*, 322, 1345, doi: 10.1126/science.1166609

- Keenan, P. C., & McNeil, R. C. 1989, *ApJS*, 71, 245, doi: 10.1086/191373
- Kenworthy, M. A., & Haffert, S. Y. 2025, *Annual Review of Astronomy and Astrophysics*, 63, 179, doi: <https://doi.org/10.1146/annurev-astro-021225-022840>
- Keppler, M., Benisty, M., Müller, A., et al. 2018, *A&A*, 617, A44, doi: 10.1051/0004-6361/201832957
- Keppler, M., Teague, R., Bae, J., et al. 2019, *A&A*, 625, A118, doi: 10.1051/0004-6361/201935034
- Konopacky, Q. M., Barman, T. S., Macintosh, B. A., & Marois, C. 2013, *Science*, 339, 1398, doi: 10.1126/science.1232003
- Konopacky, Q. M., Marois, C., Macintosh, B. A., et al. 2016, *AJ*, 152, 28, doi: 10.3847/0004-6256/152/2/28
- Kopparapu, R. K., Ramirez, R., Kasting, J. F., et al. 2013, *ApJ*, 765, 131, doi: 10.1088/0004-637X/765/2/131
- Kraus, A. L., Ireland, M. J., Cieza, L. A., et al. 2014, *ApJ*, 781, 20, doi: 10.1088/0004-637X/781/1/20
- Kuhn, J. R., Potter, D., & Parise, B. 2001, *ApJ*, 553, L189, doi: 10.1086/320686
- Kuzuhara, M., Tamura, M., Kudo, T., et al. 2013, *ApJ*, 774, 11, doi: 10.1088/0004-637X/774/1/11
- Lafrenière, D., Marois, C., Doyon, R., & Barman, T. 2009, *ApJ*, 694, L148, doi: 10.1088/0004-637X/694/2/L148
- Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, *ApJ*, 660, 770, doi: 10.1086/513180
- Lagrange, A.-M., Desort, M., Galland, F., Udry, S., & Mayor, M. 2009a, *A&A*, 495, 335, doi: 10.1051/0004-6361:200810105
- Lagrange, A.-M., Gratadour, D., Chauvin, G., et al. 2009b, *A&A*, 493, L21, doi: 10.1051/0004-6361:200811325
- Lagrange, A.-M., Bonnefoy, M., Chauvin, G., et al. 2010, *Science*, 329, 57, doi: 10.1126/science.1187187
- Lagrange, A. M., Langlois, M., Gratton, R., et al. 2016, *A&A*, 586, L8, doi: 10.1051/0004-6361/201527264
- Lagrange, A.-M., Meunier, N., Rubini, P., et al. 2019, *Nature Astronomy*, 3, 1135, doi: 10.1038/s41550-019-0857-1
- Lagrange, A. M., Rubini, P., Nowak, M., et al. 2020, *A&A*, 642, A18, doi: 10.1051/0004-6361/202038823
- Lagrange, A.-M., Wilkinson, C., Mâlin, M., et al. 2025, *Nature*, 642, 905, doi: 10.1038/s41586-025-09150-4

- Langlois, M., Gratton, R., Lagrange, A. M., et al. 2021, *A&A*, 651, A71, doi: 10.1051/0004-6361/202039753
- Launhardt, R., Henning, T., Quirrenbach, A., et al. 2020, *A&A*, 635, A162, doi: 10.1051/0004-6361/201937000
- Lazzoni, C., Desidera, S., Marzari, F., et al. 2018, *A&A*, 611, A43, doi: 10.1051/0004-6361/201731426
- Lenzen, R., Hartung, M., Brandner, W., et al. 2003, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 4841, *Instrument Design and Performance for Optical/Infrared Ground-based Telescopes*, ed. M. Iye & A. F. M. Moorwood, 944–952, doi: 10.1117/12.460044
- Linder, E. F., Mordasini, C., Mollière, P., et al. 2019, *A&A*, 623, A85, doi: 10.1051/0004-6361/201833873
- Liu, M. C., Fischer, D. A., Graham, J. R., et al. 2002, *ApJ*, 571, 519, doi: 10.1086/339845
- Liu, M. C., Wahhaj, Z., Biller, B. A., et al. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7736, *Adaptive Optics Systems II*, ed. B. L. Ellerbroek, M. Hart, N. Hubin, & P. L. Wizinowich, 77361K, doi: 10.1117/12.858358
- Llop-Sayson, J., Wang, J. J., Ruffio, J.-B., et al. 2021, *AJ*, 162, 181, doi: 10.3847/1538-3881/ac134a
- Llop-Sayson, J., Beichman, C., Bryden, G., et al. 2025, *AJ*, 170, 229, doi: 10.3847/1538-3881/adf727
- Lowrance, P. J., Becklin, E. E., Schneider, G., et al. 2005, *AJ*, 130, 1845, doi: 10.1086/432839
- Lyot, B. 1939, *Monthly Notices of the Royal Astronomical Society*, 99, 580, doi: 10.1093/mnras/99.8.580
- Macintosh, B., Graham, J., Palmer, D., et al. 2006, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 6272, *Advances in Adaptive Optics II*, ed. B. L. Ellerbroek & D. Bonaccini Calia, 62720L, doi: 10.1117/12.672430
- Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, *Proceedings of the National Academy of Science*, 111, 12661, doi: 10.1073/pnas.1304215111
- Macintosh, B., Graham, J. R., Barman, T., et al. 2015, *Science*, 350, 64, doi: 10.1126/science.aac5891

- Maire, A.-L., Langlois, M., Dohlen, K., et al. 2016, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9908, Ground-based and Airborne Instrumentation for Astronomy VI, ed. C. J. Evans, L. Simard, & H. Takami, 990834, doi: 10.1117/12.2233013
- Mamajek, E. E., & Bell, C. P. M. 2014, MNRAS, 445, 2169, doi: 10.1093/mnras/stu1894
- Mamajek, E. E., & Hillenbrand, L. A. 2008, ApJ, 687, 1264, doi: 10.1086/591785
- Marengo, M., Megeath, S. T., Fazio, G. G., et al. 2006, ApJ, 647, 1437, doi: 10.1086/505531
- Marengo, M., Stapelfeldt, K., Werner, M. W., et al. 2009, ApJ, 700, 1647, doi: 10.1088/0004-637X/700/2/1647
- Marois, C., Correia, C., Véran, J.-P., & Currie, T. 2014, in IAU Symposium, Vol. 299, Exploring the Formation and Evolution of Planetary Systems, ed. M. Booth, B. C. Matthews, & J. R. Graham, 48–49, doi: 10.1017/S1743921313007813
- Marois, C., Doyon, R., Nadeau, D., Racine, R., & Walker, G. A. H. 2003, in EAS Publications Series, Vol. 8, EAS Publications Series, ed. C. Aime & R. Soummer (EDP), 233–243, doi: 10.1051/eas:2003067
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, ApJ, 641, 556, doi: 10.1086/500401
- Marois, C., Lafrenière, D., Macintosh, B., & Doyon, R. 2008a, ApJ, 673, 647, doi: 10.1086/523839
- Marois, C., Macintosh, B., Barman, T., et al. 2008b, Science, 322, 1348, doi: 10.1126/science.1166585
- Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. 2010, Nature, 468, 1080, doi: 10.1038/nature09684
- Matthews, E. C., Carter, A. L., Pathak, P., et al. 2024, Nature, 633, 789, doi: 10.1038/s41586-024-07837-8
- Mawet, D., Riaud, P., Absil, O., & Surdej, J. 2005, ApJ, 633, 1191, doi: 10.1086/462409
- Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, ApJ, 792, 97, doi: 10.1088/0004-637X/792/2/97
- Mawet, D., Hirsch, L., Lee, E. J., et al. 2019, AJ, 157, 33, doi: 10.3847/1538-3881/aaef8a
- Mayor, M., & Queloz, D. 1995, Nature, 378, 355, doi: 10.1038/378355a0
- Mayor, M., Pepe, F., Queloz, D., et al. 2003, The Messenger, 114, 20
- Metchev, S., Marois, C., & Zuckerman, B. 2009, ApJ, 705, L204, doi: 10.1088/0004-637X/705/2/L204

- Milli, J., Mouillet, D., Fusco, T., et al. 2017, arXiv e-prints, arXiv:1710.05417, doi: 10.48550/arXiv.1710.05417
- Milli, J., Kasper, M., Bourget, P., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10703, Adaptive Optics Systems VI, ed. L. M. Close, L. Schreiber, & D. Schmidt, 107032A, doi: 10.1117/12.2311499
- Mittal, T., Chen, C. H., Jang-Condell, H., et al. 2015, ApJ, 798, 87, doi: 10.1088/0004-637X/798/2/87
- Morales, F. Y., Bryden, G., Werner, M. W., & Stapelfeldt, K. R. 2016, ApJ, 831, 97, doi: 10.3847/0004-637X/831/1/97
- Morales, F. Y., Werner, M. W., Bryden, G., et al. 2013, in American Astronomical Society Meeting Abstracts, Vol. 221, American Astronomical Society Meeting Abstracts #221, 144.05
- Mouillet, D., Larwood, J. D., Papaloizou, J. C. B., & Lagrange, A. M. 1997, MNRAS, 292, 896, doi: 10.1093/mnras/292.4.896
- Müller, A., Keppler, M., Henning, T., et al. 2018, A&A, 617, L2, doi: 10.1051/0004-6361/201833584
- Naef, D., Latham, D. W., Mayor, M., et al. 2001, A&A, 375, L27, doi: 10.1051/0004-6361:20010853
- Nasedkin, E., Mollière, P., Lacour, S., et al. 2024, A&A, 687, A298, doi: 10.1051/0004-6361/202449328
- Nielsen, E., De Rosa, R., Macintosh, B., et al. 2019a, in AAS/Division for Extreme Solar Systems Abstracts, Vol. 51, AAS/Division for Extreme Solar Systems Abstracts, 100.02
- Nielsen, E. L., Liu, M. C., Wahhaj, Z., et al. 2013, ApJ, 776, 4, doi: 10.1088/0004-637X/776/1/4
- Nielsen, E. L., De Rosa, R. J., Macintosh, B., et al. 2019b, AJ, 158, 13, doi: 10.3847/1538-3881/ab16e9
- Nowak, M., Lacour, S., Lagrange, A.-M., et al. 2020, A&A, 642, L2, doi: 10.1051/0004-6361/202039039
- Pairet, B., Cantalloube, F., Gomez Gonzalez, C. A., Absil, O., & Jacques, L. 2019, MNRAS, 487, 2262, doi: 10.1093/mnras/stz1350
- Pavlov, A., Möller-Nilsson, O., Feldt, M., et al. 2008, in Advanced Software and Control for Astronomy II, ed. A. Bridger & N. M. Radziwill, Vol. 7019, International Society for Optics and Photonics (SPIE), 701939, doi: 10.1117/12.789110
- Pedrotti, F. L., Pedrotti, L. M., & Pedrotti, L. S. 2017, Introduction to optics (Cambridge university press)

- Pepe, F., Molaro, P., Cristiani, S., et al. 2014, *Astronomische Nachrichten*, 335, 8, doi: 12.1002/asna.201312004
- Perrot, C., Thebault, P., Lagrange, A.-M., et al. 2019, *A&A*, 626, A95, doi: 10.1051/0004-6361/201834687
- Perryman, M. 2018, *The Exoplanet Handbook*
- Petit, C., Sauvage, J.-F., Costille, A., et al. 2016, *Journal of Astronomical Telescopes, Instruments, and Systems*, 2, 025003, doi: 10.1117/1.JATIS.2.2.025003
- Pollack, J. B., Hubickyj, O., Bodenheimer, P., et al. 1996, 124, 62, doi: 10.1006/icar.1996.0190
- Potier, A., Galicher, R., Baudoz, P., et al. 2025, *A&A*, 704, A98, doi: 10.1051/0004-6361/202556606
- Pueyo, L. 2016, *ApJ*, 824, 117, doi: 10.3847/0004-637X/824/2/117
- Racine, R. 1996, *Publications of the Astronomical Society of the Pacific*, 108, 699, doi: 10.1086/133788
- Ragazzoni, R., & Farinato, J. 1999, *Astronomy & Astrophysics*, 350, L23
- Rameau, J., Chauvin, G., Lagrange, A.-M., et al. 2013, *ApJ*, 772, L15, doi: 10.1088/2041-8205/772/2/L15
- Rameau, J., Nielsen, E. L., De Rosa, R. J., et al. 2016, *ApJ*, 822, L29, doi: 10.3847/2041-8205/822/2/L29
- Reggiani, M., Meyer, M. R., Chauvin, G., et al. 2016, *A&A*, 586, A147, doi: 10.1051/0004-6361/201525930
- Ren, B., Pueyo, L., Zhu, G. B., Debes, J., & Duchêne, G. 2018, *ApJ*, 852, 104, doi: 10.3847/1538-4357/aaa1f2
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 014003, doi: 10.1117/1.JATIS.1.1.014003
- Robin, A., & Creze, M. 1986, *A&A*, 157, 71
- Rouan, D., Riaud, P., Boccaletti, A., Clénet, Y., & Labeyrie, A. 2000, *PASP*, 112, 1479, doi: 10.1086/317707
- Rousset, G., Lacombe, F., Puget, P., et al. 2003, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 4839, *Adaptive Optical System Technologies II*, ed. P. L. Wizinowich & D. Bonaccini, 140–149, doi: 10.1117/12.459332
- Ruane, G., Ngo, H., Mawet, D., et al. 2019, *AJ*, 157, 118, doi: 10.3847/1538-3881/aafef2
- Ruffio, J.-B., Macintosh, B., Wang, J. J., et al. 2017, *ApJ*, 842, 14, doi: 10.3847/1538-4357/aa72dd

- Rupprecht, G., Pepe, F., Mayor, M., et al. 2004, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 5492, Ground-based Instrumentation for Astronomy, ed. A. F. M. Moorwood & M. Iye, 148–159, doi: 10.1117/12.551267
- Sabalbal, M., Absil, O., Dahlqvist, C.-H., & Delorme, P. 2026, *A&A*, 706, A275, doi: 10.1051/0004-6361/202555206
- Saffe, C., Miquelarena, P., Alacoria, J., et al. 2021, VizieR Online Data Catalog: Chemical analysis of early-type stars with planets (Saffe+, 2021), VizieR On-line Data Catalog: J/A+A/647/A49. Originally published in: 2021A&A...647A..49S, doi: 10.26093/cds/vizier.36470049
- Sahlholdt, C. L., Feltzing, S., Lindegren, L., & Church, R. P. 2019, *MNRAS*, 482, 895, doi: 10.1093/mnras/sty2732
- Samland, M., Mollière, P., Bonnefoy, M., et al. 2017, *A&A*, 603, A57, doi: 10.1051/0004-6361/201629767
- Sanghi, A., Mang, J., Llop-Sayson, J., et al. 2026, *AJ*, 171, 225, doi: 10.3847/1538-3881/ae4909
- Sarazin, M., & Roddier, F. 1990, *A&A*, 227, 294
- Serabyn, E., Huby, E., Matthews, K., et al. 2017, *AJ*, 153, 43, doi: 10.3847/1538-3881/153/1/43
- Shack, R. B., & Platt, B. C. 1971, *Journal of the Optical Society of America*, 61, 656, doi: 10.1364/JOSA.61.000656
- Skaf, N., Guyon, O., Gendron, É., et al. 2022, *A&A*, 659, A170, doi: 10.1051/0004-6361/202141514
- Skemer, A. J., Hinz, P. M., Esposito, S., et al. 2012, *ApJ*, 753, 14, doi: 10.1088/0004-637X/753/1/14
- Skemer, A. J., Hinz, P., Esposito, S., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9148, Adaptive Optics Systems IV, ed. E. Marchetti, L. M. Close, & J.-P. Vran, 91480L, doi: 10.1117/12.2057277
- Skrutskie, M. F., Jones, T., Hinz, P., et al. 2010, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III, ed. I. S. McLean, S. K. Ramsay, & H. Takami, 77353H, doi: 10.1117/12.857724
- Smith, B. A., & Terrile, R. J. 1984, *Science*, 226, 1421, doi: 10.1126/science.226.4681.1421
- Soummer, R., & Aime, C. 2004, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 5490, Advancements in Adaptive Optics, ed. D. Bonaccini Calia, B. L. Ellerbroek, & R. Ragazzoni, 495–503, doi: 10.1117/12.551985

- Soummer, R., Ferrari, A., Aime, C., & Jolissaint, L. 2007, *ApJ*, 669, 642, doi: 10.1086/520913
- Soummer, R., Hagan, J. B., Pueyo, L., et al. 2011, *ApJ*, 741, 55, doi: 10.1088/0004-637X/741/1/55
- Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJ*, 755, L28, doi: 10.1088/2041-8205/755/2/L28
- Sparks, W. B., & Ford, H. C. 2002, *ApJ*, 578, 543, doi: 10.1086/342401
- Squicciarini, V., Mazoyer, J., Lagrange, A.-M., et al. 2025, *A&A*, 693, A54, doi: 10.1051/0004-6361/202452310
- Squicciarini, V., Mazoyer, J., Wilkinson, C., et al. 2025, *A&A*, 702, L10, doi: 10.1051/0004-6361/202557104
- Sutlief, B. J., Bonse, M. J., Christiaens, V., et al. 2026, Direct Imaging Discovery of Giant Exoplanet β Pictoris d: A Decade-Long Game of Hide-and-Seek. <https://arxiv.org/abs/2606.23801>
- Tamura, M., Suto, H., Takami, H., et al. 1998, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 3354, *Infrared Astronomical Instrumentation*, ed. A. M. Fowler, 845–852
- Thompson, W., & Marois, C. 2021, *AJ*, 161, 236, doi: 10.3847/1538-3881/abee7d
- Thompson, W., Nielsen, E. L., Ruffio, J.-B., Blunt, S., & Marois, C. 2025, *AJ*, 170, 301, doi: 10.3847/1538-3881/ae0cbd
- Tschudi, C., Schmid, H. M., Nowak, M., et al. 2024, *A&A*, 687, A74, doi: 10.1051/0004-6361/202449442
- Venturini, J., Alibert, Y., Benz, W., & Ikoma, M. 2015, *A&A*, 576, A114, doi: 10.1051/0004-6361/201424008
- Vigan, A., Patience, J., Marois, C., et al. 2012, *A&A*, 544, A9, doi: 10.1051/0004-6361/201218991
- Vigan, A., Bonavita, M., Biller, B., et al. 2017, *A&A*, 603, A3, doi: 10.1051/0004-6361/201630133
- Vigan, A., Fontanive, C., Meyer, M., et al. 2021, *A&A*, 651, A72, doi: 10.1051/0004-6361/202038107
- Wagner, K., Follete, K. B., Close, L. M., et al. 2018, *ApJ*, 863, L8, doi: 10.3847/2041-8213/aad695
- Wahhaj, Z., Liu, M. C., Biller, B. A., et al. 2013a, *ApJ*, 779, 80, doi: 10.1088/0004-637X/779/1/80
- Wahhaj, Z., Liu, M. C., Nielsen, E. L., et al. 2013b, *ApJ*, 773, 179, doi: 10.1088/0004-637X/773/2/179

- Wang, J. J., Delorme, J.-R., Ruffio, J.-B., et al. 2021, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11823, Techniques and Instrumentation for Detection of Exoplanets X, ed. S. B. Shaklan & G. J. Ruane, 1182302, doi: 10.1117/12.2596484
- Wertz, O., Absil, O., Gómez González, C. A., et al. 2017, *A&A*, 598, A83, doi: 10.1051/0004-6361/201628730
- Winn, J. N., & Fabrycky, D. C. 2015, *Annual Review of Astronomy and Astrophysics*, 53, 409, doi: <https://doi.org/10.1146/annurev-astro-082214-122246>
- Wizinowich, P., Acton, D. S., Shelton, C., et al. 2000, *PASP*, 112, 315, doi: 10.1086/316543
- Wolszczan, A., & Frail, D. A. 1992, *Nature*, 355, 145, doi: 10.1038/355145a0
- Wright, J. T., Marcy, G. W., Howard, A. W., et al. 2012, *ApJ*, 753, 160, doi: 10.1088/0004-637X/753/2/160
- Wu, Y., & Murray, N. 2003, *ApJ*, 589, 605, doi: 10.1086/374598
- Xie, C., Choquet, E., Vigan, A., et al. 2022, *A&A*, 666, A32, doi: 10.1051/0004-6361/202243379
- Xuan, J. W., Ruffio, J.-B., Chachan, Y., et al. 2026, *ApJ*, 1000, 27, doi: 10.3847/1538-4357/ae448f
- Xuan, W. J., Mawet, D., Ngo, H., et al. 2018, *AJ*, 156, 156, doi: 10.3847/1538-3881/aadae6
- Zurlo, A., Vigan, A., Galicher, R., et al. 2016, *A&A*, 587, A57, doi: 10.1051/0004-6361/201526835
- Zurlo, A., Goździewski, K., Lazzoni, C., et al. 2022, *A&A*, 666, A133, doi: 10.1051/0004-6361/202243862