The background of the cover is a dark, star-filled space. A bright, white star with a prominent lens flare is positioned in the upper right quadrant. A diagonal band of reddish-orange light, possibly a nebula or a light trail, stretches from the top right towards the bottom left. In the lower left foreground, a large, dark, spherical planet is shown, partially illuminated by the star, creating a subtle gradient from dark blue to a lighter, hazy edge.

Advanced Data Processing Techniques for Exoplanet Detection in High Contrast Images

Carl-Henrik Dahlqvist

Dissertation presented for the degree of
Doctor of Philosophy in Space Sciences

University of Liège, STAR Institute

Advanced Data Processing Techniques for Exoplanet Detection in High Contrast Images

Carl-Henrik Dahlqvist



UNIVERSITÉ DE LIÈGE

Faculté des Sciences

Space sciences, Technologies and Astrophysics Research (STAR)

Institute

Thèse présentée en vue de l'obtention du grade de docteur en Sciences

Wednesday 7th September, 2022

Cover image credit : artist's view of Beta Pictoris B, ESO L. Calçada/N. Risinger (skysurvey.org)
Back cover image credit : M97 and M108, Carl-Henrik Dahlqvist

Copyright ©2022 by Carl-Henrik Dahlqvist, Université de Liège, Faculté des Sciences, 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.

This work was supported by the Fonds de la Recherche Scientifique - FNRS under Grant n° F.4504.18
and by the European Research Council (ERC) under the European Union's Horizon 2020 research
and innovation program (grant agreement n° 819155).

Doctoral Committee

Supervisor

Dr. Olivier Absil, Université de Liège

Jury members

Dr. Faustine Cantalloube, Laboratoire d'Astrophysique de Marseille

Professor Gilles Louppe, Université de Liège

Professor Marc Van Droogenbroeck, Université de Liège

Dr. Gaël Chauvin, Institut de Planétologie et d'Astrophysique de Grenoble

Dr. Laurent Puyo , Space Telescope Science Institute

President

Professor Grégor Rauw, Université de Liège

Secretary

Dr. Valentin Christiaens, Université de Liège

In loving memory of my mother, Ghislaine Massin.

Contents

Abstract	v
Résumé	ix
Acknowledgments	xiii
1 Introduction	1
1.1 A brief history of exoplanetology	1
1.1.1 History	1
1.1.2 Exoplanet definition	2
1.1.3 Indirect detection methods	4
1.2 The power of direct exoplanet imaging	9
1.2.1 Astrometry	11
1.2.2 Photometry	12
1.2.3 Spectroscopy	13
1.2.4 Planet-disk interactions	14
1.2.5 Surveys and population statistics	15
1.2.6 Complementarity of indirect methods	16
1.3 Challenges of high-contrast imaging	16
1.3.1 Angular resolution and adaptive optics systems	17
1.3.2 Contrast and coronagraphy	18
1.3.3 Speckle noise	20
1.4 Observing strategies	21
1.4.1 Angular differential imaging	21
1.4.2 Spectral differential imaging	22
1.4.3 Reference-star differential imaging	23
1.5 Data processing techniques	23
1.5.1 PSF-subtraction techniques	24
1.5.2 Maximum likelihood techniques	28
1.5.3 Supervised machine learning techniques	30
1.6 Scope and outline of the dissertation	32

2	RSM detection map	35
2.1	Introduction	35
2.2	Regime-switching model	36
2.2.1	Building the time series	37
2.2.2	Model description	38
2.2.3	Definition of the model probabilities	39
2.2.4	Transition probabilities estimation	40
2.2.5	Likelihood function definition	41
2.2.6	Model estimation	41
2.3	Detection map estimation	42
2.3.1	Computation of de-rotated cubes of residuals	42
2.3.2	Probability distribution function	43
2.3.3	Intensity parameter	44
2.3.4	Planetary signal model	46
2.3.5	RSM detection map estimation	46
2.4	Performance assessment	47
2.4.1	Data	47
2.4.2	Detection maps	47
2.4.3	Receiver operating characteristic curves	48
2.5	Conclusion	55
3	Improving the RSM map algorithm	59
3.1	Introduction	59
3.2	Using forward models in RSM	60
3.2.1	KLIP-based forward modelling	60
3.2.2	LOCI-based forward modelling	63
3.2.3	Forward model RSM map: summary	65
3.3	Forward-model RSM performance assessment	65
3.3.1	Data sets	65
3.3.2	Results	66
3.4	Optimal PSF subtraction techniques selection	73
3.5	Forward-backward model	76
3.6	Contrast curve	81
3.7	Conclusion	84
4	Auto-RSM	85
4.1	Introduction	85
4.2	PSF-subtraction techniques optimisation	86
4.2.1	Definition of the loss function	87
4.2.2	Parameter selection via Bayesian optimisation	89

4.3	RSM map optimisation	93
4.3.1	Parameter selection for the RSM map	93
4.3.2	Optimal combination of the likelihood cubes	98
4.3.3	Practical implementation	100
4.4	Performance assessment	103
4.4.1	Description of the data sets	103
4.4.2	Performance metrics	104
4.4.3	Results	105
4.4.4	Commonalities in optimal parametrisations	114
4.5	Conclusion	118
5	SHARDDS Survey	121
5.1	Introduction	121
5.2	Survey description	122
5.3	Data reduction	128
5.3.1	Pre-processing and extraction of environmental data	128
5.3.2	Image processing	128
5.3.3	Clustering	129
5.3.4	High contrast image processing	132
5.4	Identification of planetary candidates	139
5.4.1	RSM NEGFC algorithm	140
5.4.2	Point-source characterisation	143
5.5	Contrast curves analysis	145
5.5.1	Influence of clustering	146
5.5.2	Influence of environmental parameters	147
5.6	Survey Sensitivity	150
5.6.1	Target detection probability	150
5.6.2	Occurrence rate	153
5.7	Conclusion	155
6	Conclusion	159
6.1	Summary	159
6.2	Perspectives	164
Appendices		
A	RSM detection map	181
A.1	Mathematical notations for the RSM map model	182
A.2	NACO β Pictoris	183

B	Improving the RSM map algorithm	185
B.1	Mathematical notations for the RSM KLIP FM and LOCI FM	186
B.2	Forward-backward model and β parameter estimation	187
B.3	Crop size comparison for FM LOCI RSM	188
B.4	PSF-subtraction techniques combinations	189
B.5	Contrast curve for the SPHERE data set using KLIP RSM	191
C	Auto-RSM	193
C.1	Mathematical notations for auto-RSM	194
C.2	Bayesian optimisation	195
C.3	Average contrast computation	196
C.4	Auto-S/N	201
C.4.1	Algorithm definition	201
C.4.2	Performance assessment	202
C.5	Auto-RSM pseudo-code	207
C.6	Description of the ADI sequences	210
C.7	Definition of the parameter ranges	211
C.8	Detection maps for auto-RSM parametrisations	213
C.9	Parametrisation for the full-frame auto-RSM	216
D	SHARDDS Survey	219
D.1	Parametrisation commonalities	219
D.2	RSM detection maps	223
D.3	Detection maps for planetary candidates	227
D.4	Determination of projected angular separation	228
D.5	Disks analysed in Pearce et al. (2022)	230
E	Contributions summary	231
E.1	Publications as first author	231
E.2	Publications as co-author	235
E.3	Conferences and seminars	235

Abstract

High contrast imaging (HCI) is one of the most challenging techniques for exoplanet detection, but also one of the most promising. The main difficulties encountered with HCI arise from the small angular separation between the host star and the potential exoplanets, the flux ratio between them, and the image degradation caused by the Earth's atmosphere. Adaptive optics and coronagraphic techniques are now widely used to improve the quality and the dynamic range of the images with dedicated instruments. However, despite the use of these cutting-edge technologies, the resulting images are still affected by residual aberrations. Under good observing conditions, the performance of HCI instruments is limited by aberrations arising in the optical train of the telescope and instrument, generating quasi-static speckles in the field of view. Different post-processing techniques along with observing strategies have been proposed in the last decade to deal with these quasi-static speckles, whose shape and intensity are similar to potential companions. This PhD thesis builds upon these recent advances, focusing mainly on the development of a new data processing technique to unveil fainter planetary signals from angular differential imaging (ADI) sequences, and to retrieve their observed properties.

Most post-processing techniques are based on the ADI observing strategy and perform a subtraction of a reference point spread function (PSF), which models the speckle field. Such techniques generally make use of signal-to-noise maps to infer the existence of planetary signals via thresholding. An alternative method to generate the final detection map based on a regime-switching model (RSM) is developed in the first part of this thesis. This approach considers a planetary regime and a speckle regime to describe, via a Markov chain, the evolution of the pixels intensity within cubes of residuals generated by one or multiple PSF-subtraction techniques. The short memory process used in the RSM algorithm allows quasi-static speckles to be treated more effectively. Using multiple PSF-subtraction techniques helps reducing further the residual speckle noise level, better discriminating planetary signals from residual speckles. The RSM map algorithm showed an overall better performance in the receiver operating characteristic space when compared with standard signal-to-noise ratio maps for several state-of-the-art ADI-based post-processing algorithms.

Building on the good results obtained with the RSM algorithm, several improve-

ments of the vanilla RSM map algorithm are then implemented. We started by considering two forward-model versions of the RSM map algorithm based on the LOCI and KLIP PSF-subtraction techniques, allowing to account for the planetary signal self-subtraction observed at short separations. We then addressed the question of optimally selecting the PSF subtraction techniques to optimise the overall performance of the RSM map. A new forward-backward approach is also implemented to take into account both past and future observations to compute the RSM map probabilities, leading to improved precision in terms of astrometry and lowering the background speckle noise. Performance analysis demonstrate the benefits of these improvements.

Following these developments, the RSM map algorithm can use up to seven PSF-subtraction techniques. The selection of the optimal parameters for these PSF-subtraction techniques as well as for the RSM map is therefore not straightforward, time consuming, and can be biased by assumptions made as to the underlying data set. We propose in the fourth chapter of this thesis a novel optimisation procedure that can be applied to each of the PSF-subtraction techniques alone, or to the entire RSM framework. This optimisation procedure, called auto-RSM, consists of three main steps: (i) definition of the optimal set of parameters for the PSF-subtraction techniques, (ii) optimisation of the RSM algorithm, and (iii) selection of the optimal set of PSF-subtraction techniques and ADI sequences used to generate the final RSM probability map. The optimisation procedure is applied to the data sets of the exoplanet imaging data challenge (EIDC). The results demonstrate the interest of the proposed optimisation procedure, with better performance metrics compared to the earlier version of RSM, as well as to other HCI data-processing techniques.

The auto-RSM framework is finally applied to the SHARDDS survey to bring an additional piece to the exoplanet puzzle, by contributing to the characterisation of planetary population via the estimation of occurrence rate maps. This survey gathers 55 main-sequence stars within 100 pc, known to host a high-infrared-excess debris disk, allowing us to potentially better understand the complex interactions between substellar companions and disks. A clustering approach is used to divide the set of targets into multiple subsets, in order to reduce the computation time by estimating a single optimal parametrisation for each considered subset. A new planetary characterisation algorithm, based on the RSM framework, is developed and tested successfully. We uncover the companion around HD206893, but do not detect any new companion around other stars. Planet detection and planet occurrence frequencies are nevertheless derived from the generated contrast curves and show a high sensitivity between 10 and 100 au for substellar companions with masses $>10M_J$.

Throughout the different chapters of this thesis, we have built a complex but highly efficient post-processing framework for ADI sequences, adding in each chapter many new features and simplifying its use. All these developments have been compiled into a python package, called PyRSM, which offers a parameter-free detection map com-

putation algorithm with a very low level of residual speckles. This package has largely increased in maturity thanks to the SHARDDS survey and has become a robust HCI post-processing pipeline, achieving good performance in terms of contrasts. PyRSM will hopefully be used for many more surveys and provide unprecedented detection limits, allowing the detection of many exoplanets with the next generation of telescopes and instruments.

Résumé

L'imagerie à haut contrast (HCI) est l'une des techniques les plus complexes pour détecter des exoplanètes, mais aussi l'une des plus prometteuses. Les principales difficultés rencontrées avec l'imagerie à haut contrast proviennent de la faible séparation angulaire entre l'étoile hôte et ses compagnons potentiels, de leur contraste et de la dégradation de l'image causée par l'atmosphère terrestre. L'optique adaptative et les techniques de coronagraphie ont permis d'améliorer la qualité des images via l'utilisation d'instruments dédiés. Malgré l'utilisation de ces technologies de pointe, les images obtenues sont néanmoins toujours affectées par des aberrations résiduelles apparaissant dans le train optique du télescope et de l'instrument, générant des aberrations quasi-statiques dans le champ de vue. Différentes techniques de traitement d'images ainsi que des stratégies d'observation ont été proposées au cours de la dernière décennie pour traiter ces aberrations quasi-statiques, dont la forme et l'intensité sont similaires à celles de compagnons potentiels. Cette thèse de doctorat s'appuie sur ces avancées récentes, afin de développer de nouvelles techniques de traitement d'images pour détecter des signaux plus faibles et les caractériser.

La plupart des techniques de traitement d'images sont basées sur la stratégie d'observation ADI (angular differential imaging) et consiste en une soustraction d'un modèle du champ d'aberrations, des images issues du télescope. Ces techniques utilisent généralement des cartes signal sur bruit pour déduire l'existence de signaux planétaires via la définition d'un seuil. Une méthode alternative pour générer la carte de détection finale via un modèle à changement de régime (Regime Switching Model, RSM) est développée dans la première partie de cette thèse. Cette approche considère un régime planétaire et un régime dominé par les aberrations pour décrire, via une chaîne de Markov, l'évolution de l'intensité des pixels dans les cubes de résidus générés par une ou plusieurs techniques de soustraction du champ d'aberrations. Ce processus à mémoire courte permet de traiter plus efficacement les aberrations quasi-statiques. L'utilisation de plusieurs techniques de soustraction du champ d'aberrations aide à réduire davantage le niveau de bruit résiduel. L'algorithme RSM démontre une meilleure performance globale en terme de courbe de fonction d'efficacité du récepteur (ROC) par rapport aux cartes de rapport signal à bruit générées par des algorithmes de post-traitement standard.

S'appuyant sur les bons résultats obtenus avec l'algorithme RSM, plusieurs améliorations de l'algorithme ont ensuite été implémentées. Nous avons commencé par inclure dans l'algorithme RSM, deux techniques de soustraction du champ d'aberrations permettant de tenir compte de la déformation du signal planétaire lors de la soustraction, observable en particulier à de courtes distance de l'étoile dans les séquences ADI. Nous avons ensuite abordé la question de la sélection optimale des techniques de soustraction du champ d'aberrations pour optimiser les performances globales de l'algorithme RSM. Une nouvelle approche prenant en compte les observations passées et futures dans un même jeu de données pour calculer les probabilités de l'algorithme RSM a également été implémenté, permettant d'améliorer la précision en termes d'astrométrie et réduire le bruit de fond.

L'algorithme RSM pouvant utiliser jusqu'à sept techniques de soustraction du champ d'aberrations, la sélection des paramètres optimaux pour ces techniques ainsi que pour l'algorithme RSM est complexe et peut être biaisée par des hypothèses formulées quant aux conditions d'observations. Nous proposons dans le quatrième chapitre de cette thèse, une nouvelle procédure d'optimisation qui peut être appliquée aux techniques de soustraction du champ d'aberrations seule, ou à l'algorithme RSM en entier. Cette procédure d'optimisation, appelée auto-RSM, consiste en trois étapes principales : (i) définition de l'ensemble des paramètres optimaux pour les techniques de soustraction du champ d'aberrations, (ii) optimisation de l'algorithme RSM, et (iii) sélection de l'ensemble des techniques de soustraction du champ d'aberrations et des séquences ADI à considérer lors de l'estimation de la carte de détection finale. La procédure d'optimisation est appliquée aux jeux de données de l'Exoplanet Imaging Data Challenge (EIDC). Les résultats démontrent l'intérêt de la procédure d'optimisation proposée, avec de meilleurs résultats en comparaison de la version précédente de l'algorithme RSM, et d'autres techniques de traitement de d'image.

L'approche auto-RSM est finalement appliqué aux jeux de données du relevé SHARDDS afin contribuer à la caractérisation des populations de planètes via l'estimation de leur taux d'occurrence. Ce jeu de données rassemble 55 étoiles situé à moins de 100pc de la Terre et connues pour héberger un disque de débris, nous permettant potentiellement de mieux comprendre les interactions complexes entre les compagnons et les disques stellaires. Une approche de partitionnement de données est utilisée pour diviser l'ensemble du jeu de données en plusieurs sous-ensembles, afin de réduire le temps de calcul en estimant une seule paramétrisation optimale pour chaque sous-ensemble considéré. Un nouvel algorithme de caractérisation des signaux planétaires, basé sur l'algorithme RSM, est développé et testé avec succès. Le compagnon sub-stellaire HD206893 B, identifié précédemment dans la littérature, est à nouveau détecté, mais aucun autre nouveau compagnon n'est détecté. Des cartes de fréquences de détection et de taux d'occurrence d'exoplanètes sont ensuite calculées sur base des courbes de contraste précédemment générées et montrent une sensibilité élevée entre 10 et 100 ua

pour des compagnons substellaires de masse $>10M_J$.

Tout au long des différents chapitres de cette thèse, nous avons construit un algorithme de traitement d'image à haut contrast complexe mais efficace, ajoutant à chaque chapitre de nouvelles fonctionnalités et simplifiant son utilisation via une optimisation automatique de ces paramètres. Tous ces développements ont été compilés dans un package python appelé PyRSM. Ce package a largement gagné en maturité grâce à l'analyse des données du relevé SHARDDS et est devenu un algorithme de traitement d'images robuste, réalisant de bonnes performances en termes de contrastes. Nous espérons que PyRSM sera utilisé pour de nombreux autres relevés et fournira des limites de détection sans précédent, permettant la détection de nombreuses exoplanètes avec la prochaine génération de télescopes et d'instruments.

Acknowledgments

Après avoir soutenu une première thèse en Sciences Economiques et de Gestion en septembre 2018, je me suis lancé une seconde fois dans cette belle expérience qu'est une thèse de doctorat. D'aucuns me traiteraient de fou, mais ce ne fut certainement pas la thèse de trop, tant j'ai apprécié ces quatre années de recherche passionnante! Je n'aurais jamais pu espérer, quelques années auparavant, pouvoir poursuivre des recherches dans un domaine si exaltant.

Cette thèse m'a permis d'acquérir de nombreuses connaissances en Sciences spatiales, qui seront certainement un atout pour la carrière dans le secteur spatial à laquelle je me destine. Elle m'a également permis d'affermir certaines qualités telles que la créativité, la rigueur, l'esprit critique, l'humilité. J'ai aussi pris confiance en moi en tant que scientifique. Ayant pu me concentrer pleinement sur ma recherche, à la différence de ma première thèse, j'ai pu aller plus loin dans la réflexion et l'analyse, et j'espère avoir su contribuer au domaine de l'imagerie à haut contraste avec les divers projets que j'ai eu la chance de publier.

Mais la réalisation d'une thèse ne se résume pas seulement au développement de qualités personnelles, il s'agit d'un partage d'idées entre différents acteurs, comme le promoteur, les co-auteurs, les lecteurs, les collègues, les personnes assistants aux séminaires ou aux conférences. Tous ces commentaires et suggestions sont nécessaires pour faire évoluer un travail de recherche. Une thèse de doctorat est donc surtout une entreprise collaborative. Il est important de mettre en évidence l'apport de tous ces acteurs à la réalisation de ce travail.

Je tiens donc tout d'abord à remercier mon promoteur ,Olivier Absil, d'avoir accepté de m'encadrer durant la réalisation de cette thèse, de m'avoir accueilli au sein d'un groupe de recherche exceptionnel. Je le remercie également pour son soutien sans faille tout au long de cette thèse et pour toute l'aide qu'il m'a apportée lors de la réalisation et soumission des différents chapitres. Je remercie également les membres de mon jury et de mon comité d'encadrement, Faustine Cantalloube, Gilles Louppe, Marc Van Droogenbroeck, Gaël Chauvin, Laurent Puyo, Grégor Rauw, et Valentin Christiaens qui m'ont fait profiter de leur expertise et de leurs précieux conseils. Je remercie également mes co-auteurs Julien Milli, Faustine Cantalloube, Markus Janson, Gayathri Viswanath avec lesquels j'ai pu développer des collaborations internationales

très enrichissantes. Je remercie l'Université de Liège qui m'a permis de réaliser cette thèse et de suivre une formation doctorale approfondie en Sciences spatiales. Je remercie enfin Nicole Massin pour sa gentillesse et son aide dans toutes mes démarches administratives.

Je souhaite ensuite remercier l'ensemble de mes proches pour leur soutien durant toutes ces années, mes collègues, Christian, Carles, Sandrine, Lorenzo, Rakesh, Gilles, Prashant, Jyotirmay, Maxime, Mariam, Malavika. J'ai une pensée émue pour Matthew qui nous a quitté dans des conditions tragiques en janvier 2022.

J'ai également une pensée particulière pour ma petite maman qui nous a quitté il y a déjà sept ans et n'aura pas pu voir l'aboutissement de ce changement de cap, de la Finance vers l'Astrophysique, qui a débuté peu de temps avant son départ. Je lui dédie ce travail en souvenir de tous les moments de bonheur que nous avons eu la chance de vivre ensemble mais également pour toutes les valeurs qu'elle m'a transmises et qui me guident encore aujourd'hui. Je remercie également le reste de ma famille, mon frère qui est toujours enthousiaste et intéressé par ce que je fais, mais également mon papa et mes deux soeurs qui m'ont toujours soutenu. Et comme on garde toujours le meilleur pour la fin, je remercie surtout Virginie qui est la condition de ma réussite. Car c'est elle qui me soutient le plus, subissant les longues journées de travail, mon stress et parfois ma mauvaise humeur. Elle a toujours su me rassurer et me faire voir les choses du bon côté dans les moments difficiles. Merci d'être là et promis plus de thèse! Mais ces remerciements ne seraient pas complets, sans mentionner mes deux chatons adorés, Hugo et Eléonore, qui illuminent et animent (parfois un peu trop) chacune de mes journées.

Chapter 1

Introduction

1.1 A brief history of exoplanetology

1.1.1 History

The possible existence of planets around other stars has fascinated mankind since ancient times. The Copernican revolution in the sixteenth century, which positioned the Sun at the center of our solar system, fostered the idea that stars in the sky could harbour planetary systems, as stated by the Italian philosopher Giordano Bruno. This idea was later shared by Isaac Newton. In the eighteenth century, he wrote in the "General Scholium" that concludes his Principia: "... if the fixed stars are the centres of similar systems, they will all be constructed according to a similar design and subject to the dominion of One." (Isaac et al. 1999). Although the first attempts to detect exoplanets are attributed to Huygens (1698), the first claims to have detected an exoplanet dates back to the nineteenth century with the pioneer works of William Jacob and Thomas See, who observed deviations in the orbits of the stars composing the 70 Ophiuchi binary star system. These early detections, that have since been refuted, were followed by numerous false positive claims in the twentieth century (e.g. van de Kamp 1969), which were attributable to the low sensitivity of the instruments.

Different new techniques were proposed in the second half of the twentieth century. In 1952, Struve (1952) proposed to use Doppler spectroscopy (radial-velocity) and the transit method to detect super-Jupiters orbiting their star much closer than Mercury. The first detection of an exoplanet is attributable to another technique, which relies on pulsar timing. In 1992, the radio astronomers Aleksander Wolszczan and Dale Frail announced the discovery of two exoplanets (PSR B1257+12 A and B) orbiting a pulsar (Wolszczan & Frail 1992). Due to the particularity of pulsars, which are extremely dense remnant of supernovae (10^{17} kg/m³) rapidly rotating and presenting a very intense magnetic field (as large as 10^{10} to 10^{12} Tesla), these planets do not belong to the main population of exoplanets. The same year, a suspected detection was made

in the binary star system Gamma Cephei Ab, using radial-velocity (RV) variations. However, Walker et al. (1992) first dismissed this hypothesis because of weak variations in the Ca II emission line index, which followed the same periodicity, and it took until 2002 to confirm the detection. The first confirmed detection of an exoplanet orbiting a main sequence (MS) star followed in 1995. Mayor & Queloz (1995) relied on the RV technique to detect 51 Pegasi b, an exoplanet with a mass similar to Jupiter orbiting very closely to its host star with an orbital period of only 4.23 days.

These early discoveries opened up a new field in astronomy, with a harvest of new detections made in the following years, mainly via the RV technique. The 2000's saw the development of additional techniques such as the successful transit method, microlensing, or high contrast imaging (HCI). The large number of detections made via the transit method in the last two decades owes to the improved accuracy and precision of photometric measurements, as well as the development of dedicated space missions. The Kepler space telescope, launched in 2009, monitored continuously approximately 150,000 main sequence stars. The space telescope discovered during its nine and a half years of operation more than 2600 exoplanets, which represents to date the largest number of confirmed detections for a single instrument (e.g. Borucki et al. 2011). As regards high contrast imaging, the first direct image of a sub-stellar companion was made in 2004 (Chauvin et al. 2004).

Since 1992, approximately 4980 exoplanet detections have been confirmed ¹ (as of February 2022), in the close vicinity (less than 10kpc) of our Sun, illustrating the fast development of this new field of astronomy. The discoveries made in the past 30 years changed profoundly our understanding of the formation of planetary systems and of the place of our own solar system in the Universe. The large diversity of observed planetary system architectures drives our attempts to develop a general formation framework. More observations are needed to better characterize planetary systems, but also to get access to the entire zoo of exoplanets, as a large fraction of the expected population of exoplanets is not yet within reach, such as Earth analogues potentially harbouring life. This calls for the development of new instruments and data processing techniques.

1.1.2 Exoplanet definition

Exoplanets (or extra-solar planets) being defined as planets orbiting stars other than our Sun, their definition first requires a proper definition for planets. This definition evolved with our understanding about the Universe and our place in it. The word planet comes from the Ancient Greek word 'planetes', which means wanderer. Every moving light that could be observed at that time in the sky, was considered as a planet, including our moon. This definition evolved ever since, along with Science, with

¹<http://exoplanet.eu/catalog/>

sometimes controversies. The most recent definition of the International Astronomical Union (IAU), which was adopted in 2006, stipulates that a planet should satisfy the following three conditions:

1. A planet must orbit a star.
2. A planet must have sufficient mass, i.e. enough self-gravity, to overcome rigid body force and form a nearly spherical shape.
3. A planet must have sufficient mass to clear the neighbourhood along its orbit.

This definition refers mainly to the planets of our own solar system, and the IAU has not yet adopted a clear definition for exoplanets. A separate proposal to extend the IAU definition to exoplanets has been issued by Margot (2015), who suggests to rely on a simple metric to determine whether a planet can clear its orbital zone during its host star lifetime.

The IAU definition of a planet focuses mainly on small planets and does not incorporate information about the transition region between giant planets and brown dwarfs. An IAU working group issued a statement in 2007, to distinguish exoplanets from brown dwarfs based on their mass (Boss et al. 2005), leading to the following additional rule:

4. A planet must have a mass below the limit of deuterium fusion.

Brown dwarfs are defined by the IAU as sub-stellar objects whose mass are high enough for deuterium fusion but too small for hydrogen-1 fusion. However, the commonly admitted mass limit of 13 Jupiter masses (M_J) for deuterium fusion² does not seem to correspond to a clear transition in terms of mass distribution of the population of companions detected via the RV technique (Udry 2010). Based on these statistical results and considering the existence of a dry region between 25 and 40 M_J in the mass distribution, Schneider et al. (2011) proposed to set the upper limit of exoplanets mass to 25 M_J . However, more recent observations question the existence of this brown dwarf desert (e.g. Duchene & Kraus 2013).

An additional concern about the definition of a planet incorporating rules 1-4, is the fact that it does not account for potential free-floating planets, whose existence has been confirmed by gravitational microlensing (Sumi et al. 2011) and direct imaging (Delorme et al. 2012). An alternative strategy to define the concept of exoplanets could consider the formation processes of planets and how these processes affect the observational properties of planets throughout their entire lifetime. As brown dwarfs share the same gravitational collapse mechanism as low-mass protostars, they could be distinguished from exoplanets in young systems, by looking at their multiplicity

²This limit considers a solar metallicity.

fraction, their mass accretion rates or the existence of circumsubstellar disks. For mature giant planets and brown dwarfs, recent studies have focused on metallicity and chemical composition instead of mass to disentangle them (e.g. Ma & Ge 2014; Maldonado & Villaver 2017).

For exoplanets, two planetary formation scenarios are currently considered. The core accretion scenario (Pollack et al. 1996) states that the planet core is built by the accretion of planetesimals formed from the dust grains constituting the protoplanetary disk. As the core grows, the accretion rate increases due to its higher gravitational pull, forming a rocky planet or a gas giant if the core becomes sufficiently massive. According to the gravitational instability scenario (Boss 2002), if the protoplanetary disk is sufficiently massive, local instabilities lead to the fragmentation of the disk into dense cores of several Jupiter masses, forming giant planets. Direct collapse through gravitational instabilities is thought to proceed on an orbital timescale (Boss 2000) while simulations of planet formation via core-accretion suggest a much slower formation process taking several millions years (Hubickyj et al. 2005). Core accretion is often associated to a cold start formation mechanism as shocks in the accretion disk reduces the entropy of the material falling into the protoplanet (Marley et al. 2007). In contrast, gas collapsing because of gravitational instabilities is thought to retain much of its entropy leading to a hot start formation mechanism. However, more recent core accretion models are also compatible with hot start accretion process (Berardo et al. 2017), while planets may form via gravitational instabilities with a range of initial entropies (Spiegel & Burrows 2012). The faster gravitational instability formation scenario could account for the giant planets detected at large separation from young stars by direct imaging. However, the giant planets originating from the protoplanetary disk fragmentation should share the same solar abundance ratio as their host star, which is not in line with the evidences suggesting that the upper layers of Jupiter and Saturn are enriched with heavy elements (Matsuo et al. 2007). These two formation scenarios do not fully reflect the large variety of exoplanets observed since 1992, and more observations are needed to take into account more complex formation processes involving complex interactions with the planetary environment or orbital migration.

1.1.3 Indirect detection methods

The last section highlighted the difficulties associated with the definition of exoplanets and emphasised the necessity of more observations to adapt and challenge formation models. These observations have been made possible by the development of complex detection methods. In this section, we summarize the various indirect methods that have been developed and used over the past 30 years to detect exoplanets, before moving on to direct imaging in Section 1.2. For each method, we briefly present their general principle and discuss their main advantages and drawbacks. As suggested

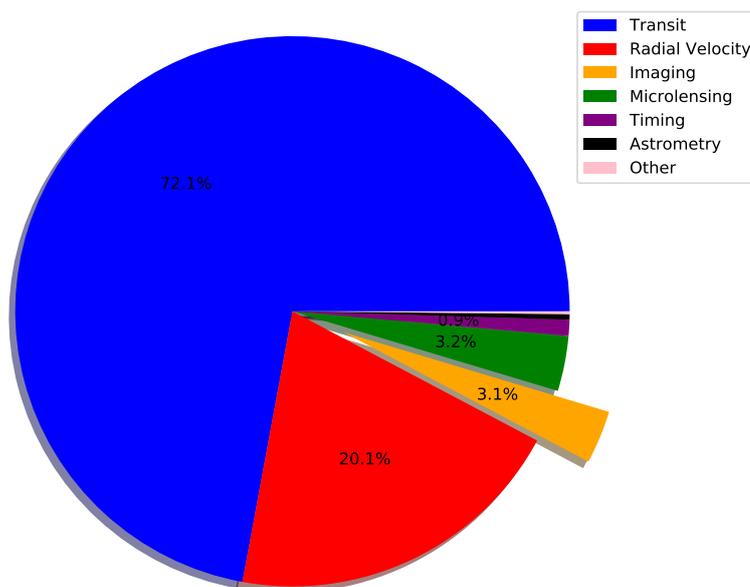


Figure 1.1: Percentage of confirmed detections attributable to the different detection methods from 1992 to February 2022. The detection methods are represented by different colors. The same color code is used in Figure 1.2. The raw data were extracted on the 2022-02-15 from the online Extrasolar Planets Encyclopaedia (Schneider et al. 2011).

by the term 'indirect', the methods presented in this section take advantage of the influence of exoplanets on their host star to detect and characterize them. In contrast with direct imaging presented in Section 1.2, no direct light from the planet is used for their characterization.

Throughout this section and the next one, the reader can refer to Figure 1.1 and 1.2 which illustrate respectively the fraction of exoplanet detections attributable to each method and the region of the mass/semi-major axis parameter space covered by the considered methods. As can be seen in Figure 1.2, each detection method is well suited for a given region of the physical parameters space, with most techniques being complementary. Nevertheless, a large fraction of the physical parameters space is not covered yet, as illustrated by the region including solar system planets.

Pulsar timing

As mentioned in Section 1.1.1, pulsar timing enabled the discovery of the first exoplanets. This approach relies on a frequency analysis of the time-periodic pulses emitted from the magnetic poles of the pulsar to unveil slight perturbations associated with the presence of a planet. Pulsar timing allows one to retrieve the mass of the detected companions as well as most of its orbital parameter (semi-major axis, eccentricity, inclination). Pulsars being significantly rarer than conventional stars in their main

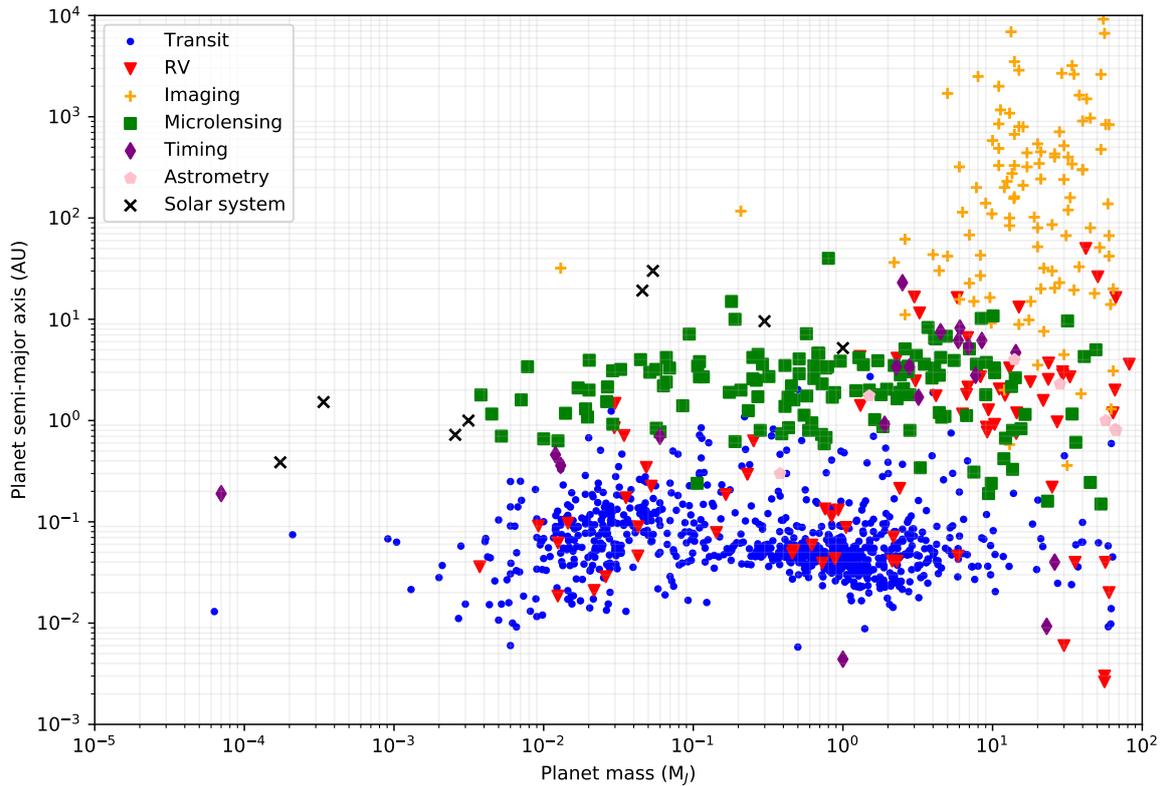


Figure 1.2: Semi-major axis of the confirmed exoplanets and brown dwarfs as a function of their mass, for the considered detection methods. The figure regroups 1331 out of the known 4978 exoplanets, since a large fraction of exoplanets is not yet characterized in terms of semi-major axis and mass. This concerns mainly the detections made by the RV techniques. The raw data were extracted on the 2022-02-15 from the online Extrasolar Planets Encyclopaedia (Schneider et al. 2011).

sequence, exoplanets detected via pulsar timing represent a marginal fraction of the number of detected exoplanets as illustrated in Figure 1.1. Pulsars being compact object, detected planetary companions are often close to the neutron star and/or massive. As no information about the origin of the detected exoplanet can be obtained (whether the exoplanets has been gravitationally captured, formed from the leftovers of the supernovae, or survived it), these detections do not provide much information about planet formation processes.

Radial velocities

The radial velocity detection method takes advantage of the revolution of an exoplanet's host star around the center of mass of the planet(s)/star system. If the exoplanet is massive enough and/or close enough to its parent star³, the center of mass

³The gravitational force of a companion decreases with the square of the distance to its host stars while the center of mass increases linearly with the distance.

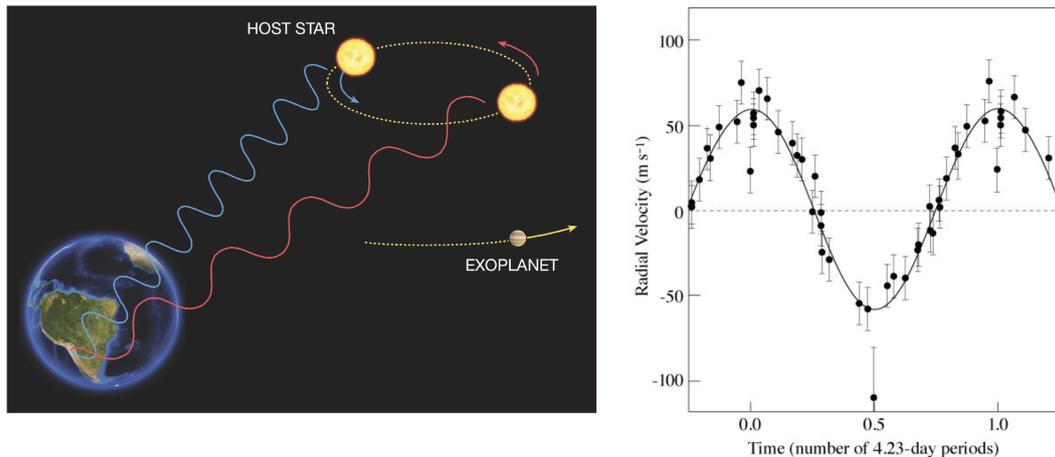


Figure 1.3: Graphical representation of the RV method (left), along with the evolution of the radial velocity curve of 51 Peg (right), the first exoplanet detected with the RV method. The dotted line represents the set of observations, while the solid line is the fit phased to a period of 4.23 days. (Credits: Mayor & Queloz (1995)).

shift becomes sufficiently large to allow a measurable periodic radial motion of the host star via Doppler spectroscopy (see Figure 1.3 for an illustration). The RV method gives access to a lower bound of the companion mass but not its true mass, as it depends on the inclination of the exoplanet orbital plane, which cannot be retrieved with this method alone. Detections through RV technique favour massive, edge-on and close-in companions (e.g. hot Jupiter) as they induce larger spectral displacements. Short orbital periods are also preferred, as they reduce the required observation timescale.

Astrometry

Besides the use of the Doppler effect to unveil radial velocity variations, a precise astrometry⁴ monitoring of the star can also lead to the detection of the periodic motion caused by an exoplanet. The GAIA space telescope dedicated to astrometry and launched in 2013, should increase significantly the number of exoplanets detected via astrometry in the next few years. Astrometry provides information about some orbital parameters of the exoplanet, including the inclination, and can therefore be used along with RV to define precisely the companion mass. The main limitation of astrometry pertains to the difficulty of obtaining precise measurements of the small displacement of stars under the gravitational influence of exoplanets. Space-based observations are required, as distortion caused by atmospheric turbulence affects dramatically the accuracy of measurements. Similarly to the RV method, massive exoplanets orbiting low mass stars are preferred. Astrometry has nevertheless a good complementarity with

⁴Astrometry consists in the precise measurement of the location of an object relative to reference background stars.

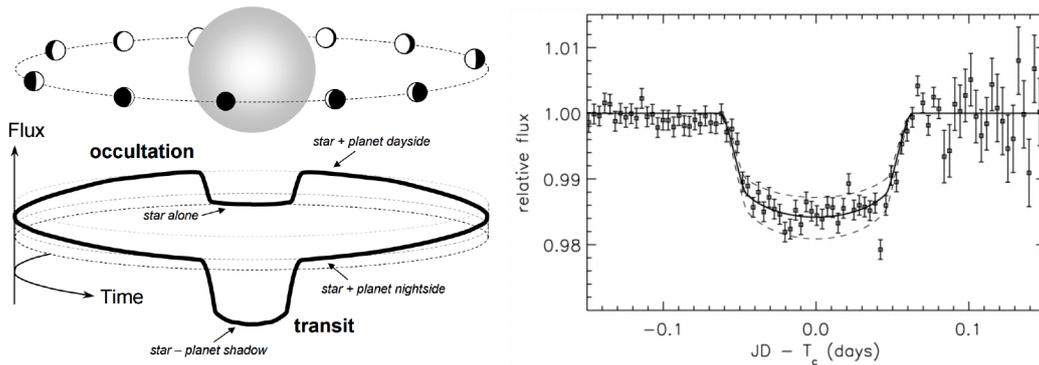


Figure 1.4: Graphical representation of the primary and secondary transits (left, credits: Perryman (2018)). Photometric time series of HD 209458 during a transit. The solid line represents the best fit model indicating a maximum reduction of 1.7% of the host star brightness during the transit (Credits: Charbonneau et al. (2000))

RV measurements as it favours pole-on system and is more sensitive to companions with large orbits.

Photometric Transit

The transit detection method has revolutionized the field of exoplanet astronomy with over 3400 confirmed transiting exoplanets, representing 72% of the discovered exoplanets. This method is based on the analysis of the evolution of the starlight intensity to track periodic eclipses (primary or secondary) caused by planets in a nearly edge-on configuration (see Figure 1.4). The idea behind the transit method dates back to O. Struve in 1952 but the first transit observation was made only in 1999 by Charbonneau et al. (2000). The NASA Kepler space observatory (Borucki et al. 2010) has been so far the most prolific instrument for the detection of exoplanets, monitoring simultaneously a hundred thousand stars in a region close to the solar system. Future missions, such as PLATO (PLANetary Transits and Oscillations of stars, expected to be launched in 2026, Rauer et al. 04 September 2014) will contribute to new discoveries by considering a larger set of stars.

The transit method gives directly access to the radius of the planetary companion by studying the transit depth, which is proportional to the star-planet surface ratio. It also provides information about the orbital period and the orientation of the exoplanet's orbit relative to the sky plane. Transit may be used along with RV to obtain the mass and eccentricity, which combined with the transit observations, give access to the density of the planet and therefore to its internal structure. Although the transit method is indirect, it can probe planetary atmospheres via spectroscopy, by comparing the spectrum during the transit with the spectrum of the star outside transit. The main limitation of the transit method owes to the inclination requirement to observe the eclipse. The probability of having a stellar system inclination match-

ing the Earth line of sight decreases with separation, favouring companions at short separation. Moreover, there is a clear preference for large companions compared to their host star, since they induce a deeper transit, easier to detect. A drawback of the transit method is its high false positive rate (see among others, Santerne et al. 2012; Morton et al. 2016), due to the large number of events that can produce transit signatures, such as brown dwarf eclipses, background eclipses, stellar spots, or diluted eclipsing binaries⁵.

Microlensing

The microlensing method exploits a galactic gravitational lens, i.e., a distribution of matter between a background light source and the observer, which curves the light rays emitted by the background light source, as they travel towards the observer. In the context of exoplanet detection, the background light source is a star and the distribution of mass acting as a magnifier is another star potentially surrounded by a planetary system. Such alignments being rare, only a few dozens of exoplanet detections have been made via microlensing since the first detection in 2000 (Bond et al. 2001). Interestingly, the method can also probe faint free-floating companions that are not accessible by any other technique, shedding light on the population of such objects (Sumi et al. 2011). In terms of discovery space, microlensing has an increased sensitivity to small exoplanets orbiting low-mass stars at large separations, just beyond the snow line (typically between 0.5 and 10 AU), and also multi-planet systems. Microlensing allows to recover the mass of the exoplanet relative to the mass of its host star as well as the projected semi-major axis. The main limitations of the microlensing method are the limited number of retrieved astrophysical parameters and its non reproducibility, considering the extremely low probability of follow-up alignment with another background star.

1.2 The power of direct exoplanet imaging

Although indirect detection methods have dominated exoplanet science in the past decades, recent instrumental developments and new data processing techniques allowed the fast development of direct imaging as a promising complementary detection method. As illustrated by Figure 1.2, direct imaging enables the detection of young and massive exoplanets at separations that are not yet covered by indirect methods (>5 au). It is bound to bridge the gap with indirect methods in terms of angular separation and companion mass with the advent of extremely large telescopes (Quanz et al. 2015). Moreover, as the method relies on the direct observation of photons emitted by

⁵When a background eclipsing binary is co-aligning with a target star, the eclipse is diluted by the foreground target star and can appear as the transit of a exoplanet.

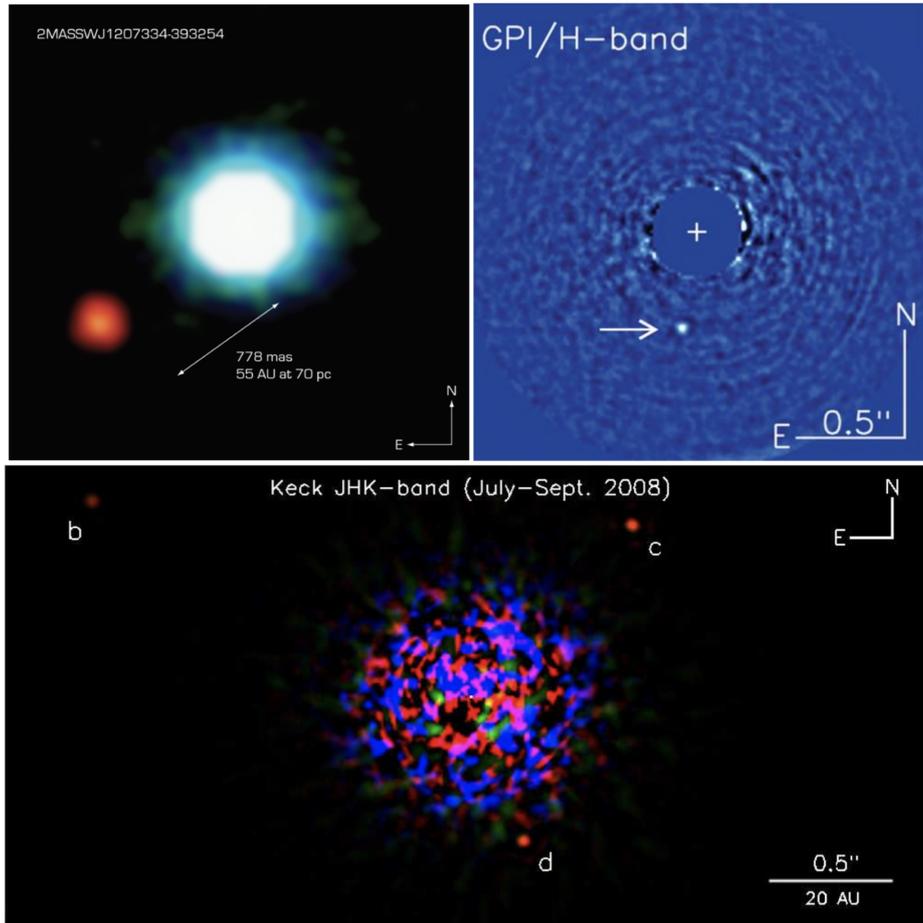


Figure 1.5: Exoplanet orbiting the 2MASSWJ 1207334-393254 brown dwarf (upper left panel, credits: Chauvin et al. 2004), Eri51 b (upper right panel, credits: Marois et al. 2008b), HR8799 b,c,d (bottom panel, credits: Macintosh et al. 2015)

the planets themselves, it gives access to a wider range of astrophysical parameters, putting new constraints on planet formation models.

The first exoplanet detected via direct imaging was a giant planet orbiting the 2MASSWJ 1207334-393254 brown dwarf situated in the TW Hydrae association (Chauvin et al. 2004). The low temperature of the brown dwarf and the relatively large angular separation (778 mas for a projected distance of 55 AU) eased the detection, the contrast between the two objects being only 10^{-2} . This discovery was followed by a few tens of exoplanets and brown dwarf companions detected via direct imaging, using both space based and ground based telescopes. Ground based telescopes are the most commonly used for direct imaging nowadays, thanks to their larger aperture and the development of advanced adaptive optics. Among these discoveries, we can point out the detection of HR8799 b,c,d,e using the Keck and Gemini-south telescopes (Marois et al. 2008b, 2010), the detection of Beta Pictoris b using the VLT (Lagrange et al. 2010), or more recently the discovery of Eri51 b with the Gemini-

south telescope (Macintosh et al. 2015). All these detections have been made possible by advanced techniques involving optics, instrumental design and data processing (see Pueyo 2018, for a review). These advanced techniques are designed to tackle the two main challenges of High Contrast Imaging (HCI), namely the contrast and the angular separation between the exoplanet and its host star. These state-of-the-art technologies will be discussed in Section 1.3. We consider first the physical parameters of detected exoplanets that can be inferred from HCI observations.

As already mentioned, direct imaging aims at detecting the photons emitted (mainly in the infrared) or reflected (mainly in the visible) by the exoplanet itself. The intensity of visible light from the host star reflected by the exoplanet depends on the planet albedo, which derives from its surface and atmosphere composition. As for the intensity of the photons emitted in the infrared, it depends on the planet surface temperature, which is affected by both the planet age and mass. Giant exoplanets as well as brown dwarfs contract with time and release gravitational potential energy as thermal emission, cooling down monotonically with age. Following evolution models, the temperature for giant exoplanets ranges from 500 to 3000 K during their first 10 Myr to cooler temperature below 500 K at 1 Gyr. Their thermal emission therefore peaks at near-to mid-infrared wavelength in their earlier stages, moving later to the mid-infrared regime at older ages. Particular attention has been paid to near-infrared emissions for HCI, which provides a good trade-off between the noisier mid-infrared bands and the more turbulent visible regime⁶. The contrast between the host star and hot young planets at these longer wavelengths is also lower than for visible light where the star black body emission is usually much higher compared to the reflected light. Although most HCI efforts focused on near-infrared H band, the Lp thermal infrared band ($3.78\mu\text{m}$) has led to many discoveries (e.g. β Pic b, HR8799 e or PDS 70 b, resp. Lagrange et al. 2010; Marois et al. 2010; Keppler, M. et al. 2018) due to its higher Strehl ratio, lower background star contamination, and the higher emission of young and intermediate giant exoplanets in these wavelengths.

1.2.1 Astrometry

Several key physical parameters of detected exoplanets may be inferred from HCI observations. Confirmed detections require usually at least two sets of observations taken at two different epochs to ensure that the planetary candidate is not an artefact due to residual noise, or a background star, especially for systems in the Galactic plane. The exoplanet orbital elements estimation requires however at least three sets of observations. A higher number of observations spanned over a long period are

⁶The higher level of noise in mid-infrared is explained by the temperature of the Earth atmosphere and the instrument itself, while atmospheric turbulence is more difficult to treat in the visible range due to the smaller coherence time (see Section 1.3.1).

often necessary, considering the slow revolving planets detected by HCI. The orbital periods of directly imaged exoplanets are so long, that no exoplanet has yet a closed orbit, i.e. astrometric measurements covering the whole orbit. The orbital elements are usually computed relying on Monte-Carlo Markov Chains or Large-Step Markov Chains (resp. MCMC and LSMC) estimation, which allows to explore the orbital parameter space, based on the estimated angular separation and angular position of the planetary candidate at the different epochs. Precise astrometry is therefore a key element of image processing techniques, as the precision of the orbital elements computation depends directly on the uncertainty associated with the estimation of the angular separation and angular position.

The definition of the orbital elements in a multi-planet system allows the simulation of potential dynamical configurations in order to infer the multi-planet system's long term stability, as well as the potential existence of additional planets. One of the best studied multi-planet system is HR 8799, for which the astrometric measurements are consistent with a stable configuration, with at least three of the four planets being co-planar, close to a 1:2:4:8 period ratio (Goździewski & Migaszewski 2020). Other systems such as β Pic have also been studied in details (Lagrange, A. M. et al. 2020). The geometry of the orbit in single-companion system may also be used to infer the possible existence of additional companions. For example, the determination of possible periastron/apastron distances for HD 95086 b by Rameau et al. (2016) suggested that the companion had not a sufficient eccentricity to explain the truncation of the outer edge of the debris disk. Orbit monitoring on single-companion systems can help differentiate perturbation due to an inner companion from non-zero eccentricity (e.g. Rameau et al. 2016; Maire et al. 2019; Lacour et al. 2021). Precise astrometry can also help investigate the formation scenario by studying the misalignment between orbital plane and stellar spin axes (Bryan et al. 2020), or by considering the orbital geometry differences existing between exoplanets and brown dwarfs (Bowler et al. 2020).

1.2.2 Photometry

Besides astrometry, precise photometry is also essential to properly characterize the detected planet. Relying on both the contrast between the companion and its host star and the estimated age of the host star, evolutionary models may be used to estimate the mass of the planet. The selected evolutionary models (often referred as hot start and cold start model) as well as the assumptions made on the physics of the atmosphere may affect drastically the estimated mass (e.g. Marley et al. 2007).

An accurate and precise estimation of the host star age is also of paramount importance, as a revision of stellar ages can affect drastically the interpretation of a detected signal, such as its classification between giant exoplanet and brown dwarf. As an illustration, a recent revision of the age of AB Pic from 30 Myr to 13 Myr (Chauvin

et al. 2005; Booth et al. 2021) led to a decrease in the estimated mass of AB Pic b from $14 M_J$ to $10 M_J$. A more striking example is GJ 504b which was first considered as an exoplanet with a mass of $4_{-1.0}^{+4.5} M_J$ for a star age of 160_{-60}^{+350} Myr estimated via gyrochronology and activity indicators (Kuzuhara et al. 2013). High-resolution optical spectrum analysis of the star led to an inferred isochronal age of $2.5_{-0.7}^{+1.0} Gyr$ (D’Orazi et al. 2017), implying a mass above $24 M_J$ for the companion and classifying therefore GJ 504b as a brown-dwarf. Most of the stars hosting a imaged planetary system belong to young moving groups or associations of stars (ages between 10-300 Myr) at close separation from the Earth ($\lesssim 100$ pc). Stars in such groups have been formed in the same region and share similar age and proper motions. Their age is estimated relying on Hertzsprung-Russell diagram, or Lithium abundance. For isolated stars, the estimation is more difficult and often based on stellar rotation and activity, using optical interferometry to resolve them.

1.2.3 Spectroscopy

An interesting additional feature of direct imaging, is its ability to probe multiple wavelengths via spectroscopy. To date, spectro-photometric measurements have been obtained for 41 giant planets or low mass brown dwarf companions via direct imaging and 136 jovians and super-Earths via the transit technique (Currie et al. 2022). While the transit method focuses on mature and highly radiated jovians and super-Earth, direct imaging allows the characterisation of young self-luminous giant exoplanets at wider distances, where the irradiation level is negligible. Multispectral observations of exoplanet can help constraining the surface temperature, the presence of clouds but also provide information about the surface gravity and the planetary radius when used along with advanced atmospheric models. High resolution spectroscopy can also unveil the chemical composition of the atmosphere and the presence of molecule such as water, methane or ammonia.

Models of exoplanet and brown dwarf atmospheres are more complex than those of stars, as they must deal with much lower temperature, lower surface gravity, the presence of clouds, disequilibrium chemistry due to vertical mixing or convection. Multiple 1D models have been developed and improved in the past decades. Their first objective was to provide spectra at the top of the atmosphere for multiple effective temperature and surface gravity, via the determination of atmospheric temperature profiles via radiative-convective flux balance. The temperature profile can then be used as boundary conditions for interior structure model providing information about temperature and surface gravity evolution over time (Chabrier & Baraffe 1997). Well known early models include the AMES-Cond models of Baraffe et al. (2003) or the model of Saumon & Marley (2008), which also incorporates a varying sedimentation efficiency in the upper atmosphere boundary to investigate the impact of clouds on the

evolution of brown dwarfs. Cloud models have first been developed by Allard et al. (2001) and Ackerman & Marley (2001) to account for the reddening of L dwarf spectra. Indeed, the presence of clouds is often associated to a hotter temperature profile, which can explain at least partially the redder colors of exoplanets presenting cloud coverage, like β Pic, HR 8799 bcde, or 51 Eri b (Currie et al. 2013; Rajan et al. 2017). Sedimentation efficiency is the main parameter characterising clouds in 1D models and defines the size of the cloud particle which should insure the equilibrium in the transport of condensible gases in the atmosphere. Other features of observed brown dwarf and giant exoplanets required the development of more complex models. High temperature transitions lines have been incorporated to describe more accurately atmosphere opacity (Phillips et al. 2020). Pressure broadened line shapes of alkali metals Na and K were introduced as it drives partly the shape of red-optical and near-infrared spectrum of brown dwarfs (Allard et al. 2016; Allard et al. 2019). Non-equilibrium chemistry has also been included to account for the vertical mixing observed in brown dwarfs (among others Saumon et al. 2003; Geballe et al. 2009; Leggett et al. 2017).

Quasi-periodic and high amplitude variation in the emission of brown dwarf and giant exoplanets indicates a dynamic and varying atmosphere that cannot be captured by global 1D models. Local 3D models have been developed to cope with the variability of cloud coverage and convective movements (Showman & Kaspi 2013; Tan & Showman 2017, 2021). All these different physics-based models, which are parametrized by planetary characteristics such as metallicity, temperature, surface gravity, or cloud coverage, provide spectra that can be compared with multispectral observations to infer the companion physical characteristics (see among others Baudino et al. 2015; Claudi et al. 2019; Cheetham et al. 2019; Phillips et al. 2020).

1.2.4 Planet-disk interactions

In our current understanding of planetary system formation, gas giant planets form in gas-rich protoplanetary disks that dissipate in a few million years, leaving behind one or several planets as well as belts of planetesimals that never managed to grow to full-sized planets. These belts are called debris disks, because the planetesimals collide between each other and generate small dust particles detectable either through their scattered light in optical regime or through their submillimetric thermal emission, creating an infrared excess above the stellar photosphere. Current far-infrared surveys can detect debris disks with an infrared excess above 10^{-6} , and have estimated that a debris disk must be present around at least 30% of A stars and 20% of FGK stars (e.g. Eiroa et al. 2013), but the real occurrence rate could be much higher (Pawellek et al. 2021). Those disks are a natural place to look for planets because planet formation succeeded to form at least large planetesimals in those systems. Meshkat et al. (2017) found indeed a tentative evidence that giant planets have a higher occurrence rate

around debris disc host. Currently forming or mature exoplanets could also explain the presence of rings, spirals, and gaps seen in many debris disks (Dong et al. 2015), stirring (Pearce et al. 2022), as well as sharp edges and pericenter offsets (Kalas et al. 2005). Modelling these perturbations may provide constraints on the masses of the exoplanets (Kalas et al. 2008; Lagrange et al. 2009a).

1.2.5 Surveys and population statistics

The multiple HCI surveys that have been done in the past two decades enabled most of the exoplanet detections via direct imaging, providing good contrast for a sample of more than 500 stars. Although the number of discovered exoplanets has been relatively small, the obtained contrast curves allowed the estimation of giant exoplanet population statistics. The most interesting population statistics is the inferred occurrence rate, more precisely the evolution of the occurrence rate depending on exoplanet or star properties such as mass, semi-major axis, or star metallicity. This key statistics allow to strongly constraints the different models of planetary formation and evolution.

Based on RV surveys, Cumming et al. (2008) found out the existence of a power law relationship between the occurrence rate of giant exoplanets and their semi-major axis up to 2.5 au. The first direct imaging surveys in the 2000's showed a low upper limit on the occurrence rate (below 10%) for giant exoplanets in wider orbit (beyond 10 au), indicating that the power law could not be extended beyond 2.5 au (Lafreniere et al. 2007b; Chauvin et al. 2010). With the increasing number of detected exoplanets and the advent of more powerful adaptive optics, more recent surveys have been able to provide estimates of the occurrence rate instead of upper limits. Higher achievable contrasts enable the potential detection of lower mass exoplanets (e.g., 50% detection probability for exoplanet with mass around $1 M_J$ at separations beyond 50 au for the SHINE survey, Vigan et al. 2021). Recent surveys show a high sensitivity for giant exoplanets with mass over $2 M_J$ between 10 and 200 au (Meshkat et al. 2017; Nielsen et al. 2019a; Vigan et al. 2021). The study of the obtained planetary occurrence rates allowed to derive (several) power laws for giant planet population between 10 and 100 au, with occurrence rates decreasing with planet mass and semi-major axis and increasing with stellar mass. Nielsen et al. (2019a) compared the obtained relationships for giant exoplanets and brown dwarfs and demonstrated that they were clearly two different populations, indicating different formation mechanisms. These studies along with more recent RV surveys (Fernandes et al. 2019; Fulton et al. 2021) beyond 2.5 au, suggest the existence of an occurrence rate peak for giant exoplanets near the snow line (3-10 au).

As mentioned in Section 1.1.2, two different formation scenarii have been proposed in the literature: the gravitational instability, which predicts high-mass exoplanets at wide separations, and the core accretion scenario, which predicts more low-mass

exoplanets at shorter separations with a dependence on the stellar mass. The results from the GPIES survey (Nielsen et al. 2019a) seem to demonstrate the consistency of the gravitational instability scenario with brown dwarf populations. However, the power law obtained for giant exoplanets at wide separations rather favours the core accretion scenario. Vigan et al. (2017) reached similar conclusion with a much higher contribution of the core accretion scenario than gravitational instability to explain giant exoplanet population. The eccentricity analysis made by Bowler et al. (2020) also suggests that brown dwarfs and giant exoplanets have different formation mechanisms involving for the first, binary-like formation, and for the second, core accretion within the proto-planetary disk.

1.2.6 Complementarity of indirect methods

A precise orbit determination via astrometric measurements of the planet position, combined with indirect techniques such as RV or absolute star astrometry may help further constraining companion mass (e.g. Mawet et al. 2019), without requiring the definition of the planet age and the selection of an evolutionary model. This will lead eventually to a set of benchmark planets, for which dynamical measurements are available together with large infrared spectral coverage. The first data releases of the Gaia survey are already revealing helpful for direct imaging analysis considering the length of the mission, which give access to long period orbits (Bonavita et al. 2021). Indirect techniques may also be helpful to fill the gaps when studying exoplanet population, direct imaging being limited to large and distant exoplanets. Indeed, considering multiple detection techniques help covering a larger fraction of the mass-distance parameter space as illustrated in Figure 1.2, but also in terms of star properties. RV studies combined with direct imaging can for instance help investigating the extent of giant exoplanets migration during their lifespan, by considering a wider range of separations (e.g. Lagrange, A. M. et al. 2020).

1.3 Challenges of high-contrast imaging

The direct detection of exoplanets is a very challenging task, considering the small angular distance between the companion and the host star (typically between 0.1 and a few arcseconds) and the very high contrast between them, which ranges from 10^{-3} (young massive planet emitting in the infrared) to 10^{-10} (Earth-like planet reflecting light from their host star). These elements explain why only giant planets orbiting at a large angular separation have been discovered so far, despite the technological breakthroughs of the past decade. In the case of ground-based telescopes, the use of adaptive optics along with coronagraphic devices has been the mainstream approach to reach both high contrast and high angular resolution.

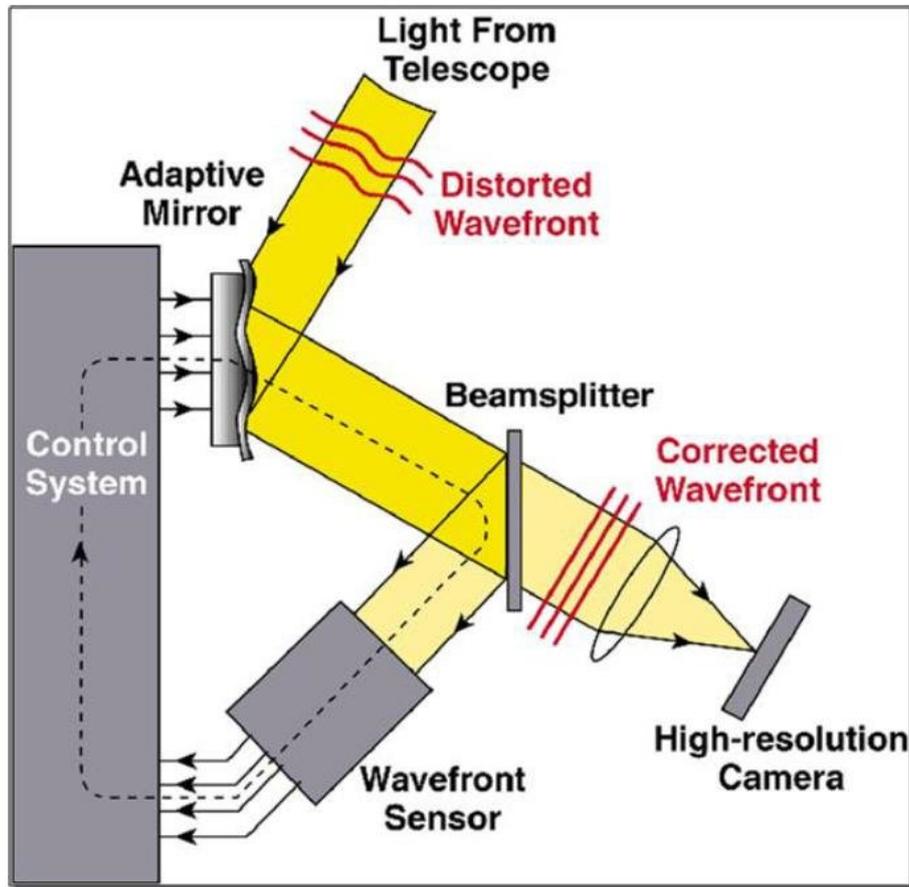


Figure 1.6: Simplified graphical representation of an adaptive optics system. Credits: Rimmele and Marino (2011).

1.3.1 Angular resolution and adaptive optics systems

The resolution achieved by a telescope is in theory inversely proportional to the diameter of its main aperture for a given wavelength. The image of a point-like source at large distance, such as a star, is a diffraction pattern called the point spread function (PSF), which consists for a circular aperture in an Airy function presenting a main peak surrounded by multiple lobes. In the diffraction-limited case, the resolving power of the telescope is given by the Rayleigh criterion, which imposes a minimum distance between two resolved objects of $1.22\lambda/D$, with D the diameter of the telescope primary mirror, and λ the considered wavelength. This distance corresponds to the separation between the PSF main peak and its first minimum.

However, for ground-based telescopes, atmospheric turbulence leads to time and space-varying phase shifts of the wavefront affecting the PSF sharpness. The expected sharp PSF core is replaced by a fuzzy blob or halo, decreasing the angular resolution limit. The angular resolution in the presence of atmospheric turbulence is given by the seeing $0.98\lambda/r_0$ (Dierickx 1992), with r_0 the Fried parameter providing an equivalent

telescope diameter in terms of achievable angular resolution. As the Fried parameter is scaling as $\lambda^{6/5}$, the angular resolution limit is less affected by atmospheric turbulence in the near-infrared compared to the visible⁷.

The angular resolution limit therefore becomes independent of the telescope aperture, when accounting for atmospheric turbulence. This could drastically limit the resolving power of the 8 m class telescopes that have been constructed throughout the 90s. The concept of adaptive optics has been proposed by Babcock (1953) to overcome, at least partially, the effect of turbulence by using deformable mirrors (DM) to correct the distorted wavefront. It was not until 1990 that adaptive optic was applied to astronomy at the La Silla Observatory (Rousset et al. 1990).

As illustrated in Figure 1.6, an adaptive optics system combines a wavefront sensor (WFS), a deformable mirror (DM) composed of hundreds or thousands of actuators (41x41 actuators for SAXO, in VLT/SPHERE, Beuzit et al. 2019) and a real time computer (RTC). The wavefront sensor, such as the Shack-Hartmann wavefront sensor (Shack & Platt 1971), the pyramid wavefront sensor (Ragazzoni & Farinato 1999) or the curvature wavefront sensor (Roddiier & Roddiier 1988), analyses the local deformation of the wavefront of a guide star (a star positioned near the astrophysical object of interest, or a laser guide star). This information about the local slope of the wavefront is sent to the RTC that reconstructs the entire wavefront. Based on the reconstructed wavefront, the RTC computes the correction to be applied to the DM, which consists in optical path delays or advances. These corrections are sent to the actuators which deform locally the mirror in order to cancel out the wavefront errors, flattening the wavefront and therefore sharpening the PSF. These steps are repeated in a closed loop mode at a high frequency (1.2 kHz for SAXO, in VLT/SPHERE, Beuzit et al. 2019).

The VLT/SPHERE instrument, which will provide most data sets used throughout this dissertation, relies on extreme AO (xAO), focusing on a small region of the sky in order to further improve the achievable Strehl ratio via a higher density DM and a fast RTC. As can be seen in Figure 1.7, the xAO installed in the VLT/SPHERE instrument can achieve a Strehl ratio of about 90 % in the H-band ($1.6\mu\text{m}$), meaning that around 90 % of the maximal intensity of the aberration-free PSF can be recovered for nominal turbulence conditions (to be compared to about 5 % without AO). A more complete review of the AO systems used in high contrast imaging can be found in Milli et al. (2016a).

1.3.2 Contrast and coronagraphy

Having dealt with the question of the angular resolution limit, we now turn to the large difference in terms of brightness between exoplanets and their host stars. The high

⁷This is explained by the smaller coherence time in the visible which leads to more severe phase shifts for a given optical path

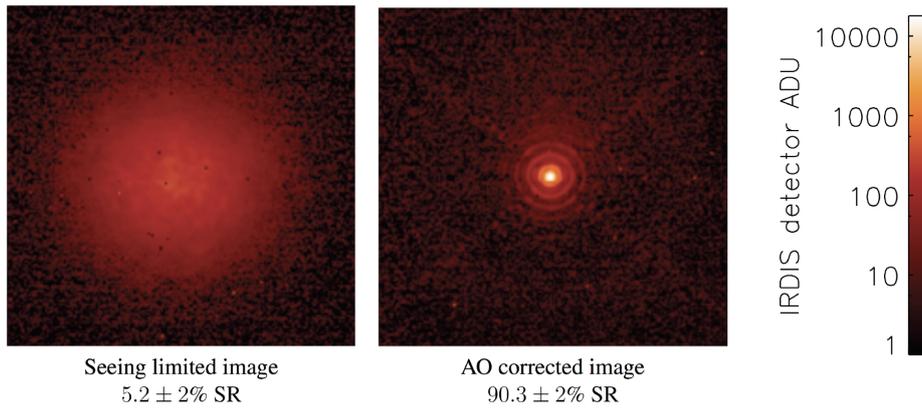


Figure 1.7: Example of a long exposure image of a star in H-band ($1.6\mu\text{m}$) taken with the VLT/SPHERE instrument in nominal turbulence seeing conditions (Seeing of 0.85 arc sec and wind speed of 12.5 m s^{-1}), without AO correction (left), where the size of the seeing pattern is around λ/r_0 ; and with AO correction (right) where the instrument resolution is closer to the diffraction limited resolution of λ/D (with $D = 8\text{m}$ for the VLT). Credits: (Sauvage et al. 2016).

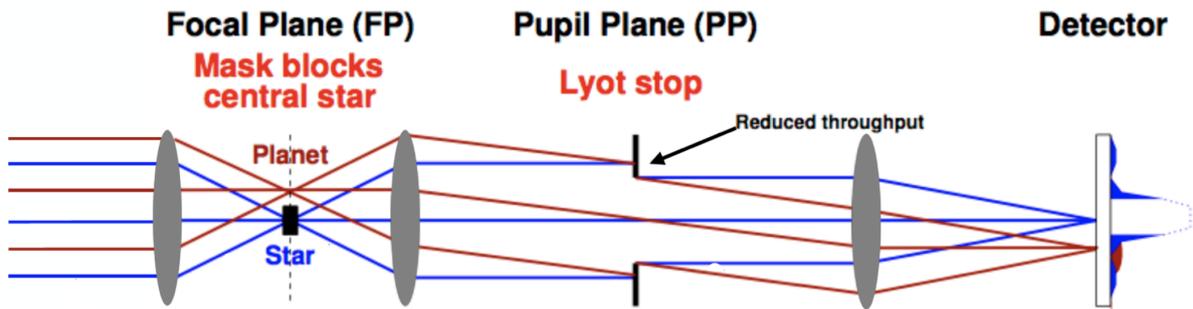


Figure 1.8: Optical layout of the classical Lyot coronagraph. Credits: M. Kenworthy (Leiden University).

contrast challenge is addressed by the use of coronagraphic techniques, which aim to (i) create a dark region around the host star to avoid detector saturation, and (ii) cancel the diffracted starlight in the off-axis region where planets may orbit. Coronagraphs can remove starlight by acting at different levels on the amplitude or the phase⁸, in the focal plane or in the pupil plane. Coronagraphy is a very active field and different strategies acting on these different levels have been investigated. The original concept dates back to Lyot (1939), and consists in an opaque disk blocking the starlight at the focal plane (Lyot mask), combined with a second annulus-shaped opaque mask (the Lyot stop) placed at the following pupil plane to remove most of the remaining diffracted starlight (Airy rings not blocked by the Lyot mask). The Lyot coronagraph acts on the amplitude of starlight at both focal and pupil plane.

⁸When acting on the phase, the coronagraph creates destructive interferences via phase shift to remove the starlight.

The main drawback of the classical Lyot coronagraph is the reduced throughput of planetary signal due to the Lyot stop, as well as the large inner working angle (IWA)⁹, and the non perfect cancellation of the diffracted starlight. More advanced designs have been proposed to cope with these limitations, taking into account additionally the spectral bandwidth at which the coronagraphic mask operates, as well as reducing the sensitivity of the set-up to potential instrumental aberrations. Among coronagraphic concepts operating in current state-of-the-art observatories, we may note for instance the vortex coronagraph (Mawet et al. 2005; Mawet et al. 2013), which creates destructive interferences via a phase-shift of the on-axis starlight, rather than blocking the starlight. This design improves significantly the throughput and reduces the IWA, however at the cost of a higher sensitivity to instrument misalignments (tip-tilt). We also note the apodized pupil Lyot coronagraph (APLC, Soummer 2004) which is the most commonly used coronagraphic device to reach high contrast at small IWA. The APLC is an hybrid coronagraph combining a Lyot coronagraph and a pupil plane apodizer, which helps further reduce the on-axis light intensity. Other designs such as apodized vortex coronagraphs are currently under development and foreseen for current and future ground-based instruments. A detailed review of most of these coronagraphic technologies may be found in Mawet et al. (2012).

High contrast imaging instruments are often equipped with multiple coronagraphic devices and multiple masks to cover a wide range of observational constraints, as well as a broader range of wavelengths. These coronagraphic devices are often coupled with a tip-tilt sensor to center the source and limit instrumental misalignments, further reducing the IWA. The VLT/SPHERE instrument is equipped with a classical Lyot coronagraph, a four-quadrant phase mask and an APLC, along with a differential tip-tilt sensor (DTTS, Baudoz et al. 2010).

1.3.3 Speckle noise

In spite of the AO system and the coronagraph, some stellar light always makes its way to the detector, in part due to residual aberrations coming from non corrected atmospheric turbulences, or to instrumental aberrations appearing after the beam splitter (also called non-common path aberrations, NCPA, see Figure 1.6), which can hinder the detectability of faint exoplanets. Part of these aberrations can be corrected via careful calibration. However, some of them are time dependent due to mechanical stress evolution, thermal variations during the observation, unfiltered vibrations, or imperfections in the moving parts of the instrument. These time-dependent aberrations lead to the appearance of quasi-static speckles, which are difficult to remove in post-processing because their variability is driven by the slow variation of the in-

⁹ The IWA is the smallest angle at which the required contrast is achieved with a throughput higher than 0.5.

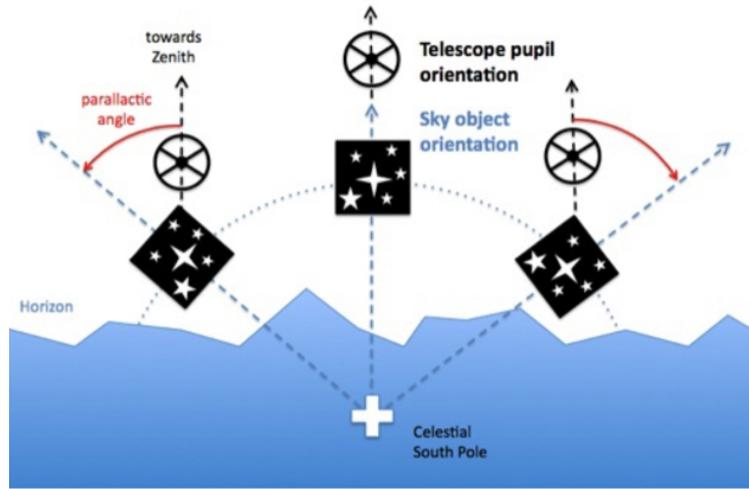


Figure 1.9: Graphical representation of the angular differential imaging observing strategy. Credits: Christian Thalmann.

strument mechanical characteristics on timescales of several minutes to several hours (Milli et al. 2016a; Goebel et al. 2018). Different image processing techniques, along with observing strategies have been proposed in the last decade to deal with these quasi-static speckles, whose shape and intensity (about 10^{-4} raw contrast to the star) are similar to potential companions.

1.4 Observing strategies

We start by considering the different observing strategies, i.e., how the data is acquired at the instrument level, before considering in the next section the various data processing techniques that have been built upon these observing strategies. All three observing strategies presented in this section refer to 'differential imaging', which relies on an observing sequence of a single target or multiple targets to define a model of stellar residuals and speckle noise, which are then removed from the observations. The modelling of the speckle field and its subtraction from the science images represents the corner stone of most post-processing techniques. The different observing strategies are designed to enable the construction of a model for the speckle field, affecting the planetary signal as little as possible during the subtraction process.

1.4.1 Angular differential imaging

Angular differential imaging (ADI, Liu 2004; Marois et al. 2006) is nowadays the most commonly used observing strategy to mitigate quasi-static speckles in HCI. This observing strategy consists in acquiring images in pupil-stabilised mode, i.e., with the

instrument derotator keeping the pupil orientation fixed with respect to the detector. The aim of this approach is to keep the quasi-static speckles fixed in the focal plane, so that they can easily be identified with respect to astrophysical objects rotating around the star along with the Earth rotation. Indeed, the instrumental aberrations and diffraction effects causing quasi-static speckles are mainly locked to the pupil. By using this angular diversity, a model of the speckle field, often referred to as the reference point spread function (reference PSF), may be built from the data. This reference PSF is then subtracted from the set of ADI images. The resulting subtracted frames are eventually aligned and combined to detect the signal of potential exoplanets or disks, which should not have suffered too much from the reference PSF subtraction, while the residual noise should average out incoherently. The main limitations of the ADI observing strategy are the significant observing time required to obtain a sufficient angular diversity (usually around one hour), as well as the planetary signal self-subtraction observed at close separation (see Chapter 3 for a discussion about self-subtraction and mitigation procedures). The main assumption behind the ADI approach is that speckles remain mostly static within the observing sequence. However, this assumption holds only partially, as small temporal changes in the wavefront quality due mainly to NCPA, can affect the speckle field modelling. ADI also has a limited sky coverage to achieve the required total parallactic angle rotation for an acceptable observing time.

1.4.2 Spectral differential imaging

Spectral differential imaging (SDI) seeks to decouple the planetary signal from the speckle field, based on their wavelength dependence rather than azimuthal rotation (ADI), overcoming one of the shortcomings of ADI by considering observations taken simultaneously. The initial concept was first proposed by Racine et al. (1999), who relied on two images taken in adjacent bands. Their idea was to select carefully the two wavelengths to insure the presence of the exoplanet in one image but limit its brightness in the other (the methane absorption line is commonly used, e.g. Lenzen et al. 2004; Marois et al. 2005). This allows to effectively suppress the speckle noise but avoid any self-cancelling of the planetary signal during the subtraction.

A more advanced multi-spectral SDI approach (requiring integral field spectrographs) proposed by Sparks & Ford (2002) relies on the fact that, at first order, the speckles expand radially in the field of view as the wavelength increases, while potential planetary signals remain fixed (although their PSF broadens slightly). Indeed, since speckles originate from diffraction of on-axis residual light from the host star, their spatial distribution scales with wavelength as for any diffraction effect, while their shape remains constant (in the absence of chromatic aberrations). It is therefore possible to rescale the images taken at different wavelengths to a common reference

wavelength, using the linear scaling of diffraction with wavelength from which we derive the scale $S(\lambda) = \lambda_{ref}/\lambda$. The speckles being aligned in the rescaled images, a model of the speckle field can be created and subtracted from them, before scaling back and mean-combining the resulting images. Because of the radial shift of the companion during the rescaling process, the self-subtraction of the planetary signal should be limited. SDI is often used along ADI to enjoy a larger diversity for the computation of the speckle field model (see among others, Zurlo et al. 2016; Flasseur et al. 2020; Kiefer et al. 2021).

1.4.3 Reference-star differential imaging

Reference-star differential imaging (RDI, Lafrenière et al. 2009; Soummer et al. 2011; Ruane et al. 2019) is an alternative approach to ADI and SDI that relies on observations of different stars to build a model of the speckle field. The speckle field model is then subtracted from science images after adapting its brightness. RDI tries to alleviate one of the main shortcomings of both ADI and SDI observing strategies, i.e. the over-subtraction observed especially at close separation (due lower parallactic angle rotation for ADI, and the less effective rescaling for the SDI).

However, the RDI observing strategy comes also with important constraints. As it relies on different stars to build the speckle field model, a high stability of the instrument is needed. RDI is therefore preferentially used with the Hubble space telescope (e.g., Golimowski et al. 2006; Schneider et al. 2008) which avoids the time-variable aberrations due to uncorrected atmospheric turbulences. It also requires a procedure to precisely determine the flux scaling factor used to match the magnitude of the stars and avoid the apparition of artefacts.

1.5 Data processing techniques

In the last decade, the field of HCI has been very active and a large number of data processing techniques have been developed to identify and characterise planetary candidates. Data processing represents one of the corner stones of high contrast imaging, along with instrumentation and observing strategies. The most common approach makes use of both angular differential imaging (ADI, Marois et al. 2006) observing strategy and reference PSF subtraction.

Data processing of HCI data sets is a multi-step procedure, which can be divided between pre-processing and post-processing of the FITS files produced by the telescope¹⁰. The pre-processing step involves the calibration of the images, with dark current subtraction, flat field correction (to treat possible non-uniformity inside the

¹⁰FITS files are the IAU standard for storing astronomical data (<https://fits.gsfc.nasa.gov/iaufwg>)

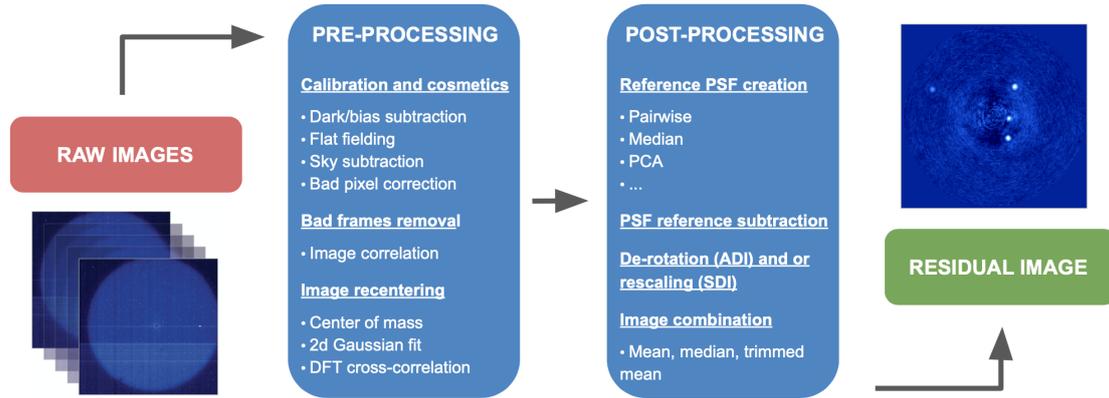


Figure 1.10: Representation of the typical data processing pipeline for PSF-subtraction techniques, from the set of raw images coming from the telescope until the final residual image.

field of view), subtraction of the thermal background radiation of the sky (important when the observations are done in the thermal infrared), and bad pixels removal (dead or hot pixels). Bad frames caused by star/coronagraph misalignments, bad observing conditions or AO correction errors are then identified by a correlation analysis of the images or by a pixel statistics analysis for a subregion of the field of view, and removed from the data set. The last pre-processing step is the recentering of all images to avoid any misalignment of the central star that could affect reference PSF modelling during the post-processing steps. This step is not mandatory for last generation instruments, such as VLT/SPHERE, which has real-time coronagraphic centering capabilities.

The post-processing of the pre-processed FITS files then aims to maximise the planetary signal to noise ratio (SNR), by tackling the noise resulting from NCPA, AO correction errors or residual diffracted light from the coronagraph. Post-processing techniques can be divided in three main categories: (i) speckle subtraction techniques, (ii) maximum likelihood techniques, and (iii) supervised machine learning techniques. We present here a non-exhaustive list of HCI post-processing techniques to illustrate these different approaches.

1.5.1 PSF-subtraction techniques

Speckle subtraction techniques, also called PSF-subtraction techniques, are the most commonly used approach in the HCI community, and follow the steps outlined in Figure 1.10: estimation of a model of the speckle field (also called reference PSF), subtraction of this model from each science image, de-rotation (ADI) and/or rescaling (SDI) of the PSF-subtracted images, and finally combination of these images to generate a so-called residual image. Several methods using this framework have been proposed to maximize the noise reduction by optimizing the estimation of the reference

PSF.

Median subtraction

The most simple approach to model the speckle field is the median combination of the entire set of calibrated science images. This simple approach was proposed along with the original ADI observing strategy by Marois et al. (2006). Because of its rotation in the set of science images, the planetary signal affects only marginally the estimation of the median. The median of the science images provides therefore a good first approximate of the speckle field. After the subtraction of the median, the residual noise is centred around zero, such that when all the PSF-subtracted images are de-rotated and combined, the residual noise should average out (the noise being mainly Gaussian from the central limit theorem, see for more details Marois et al. 2008a). More advanced versions of this approach involve a field rotation criterion to reduce self-subtraction of the planetary signal (Smart-ADI, Marois et al. (2006)) or the use of the centro-symmetry of the speckles to form the PSF model (IRS, Ren et al. (2012)). Despite these improvements, median combination does not perform well in the close vicinity of the host star, as self-subtraction is strong at these close separations and residual speckle noise remains strong even after subtraction. More advanced PSF-subtraction techniques have been proposed to mitigate these drawbacks by creating a model of the speckle field for each individual science image (see Figure 1.11 for an illustration of the main steps of these PSF subtraction techniques).

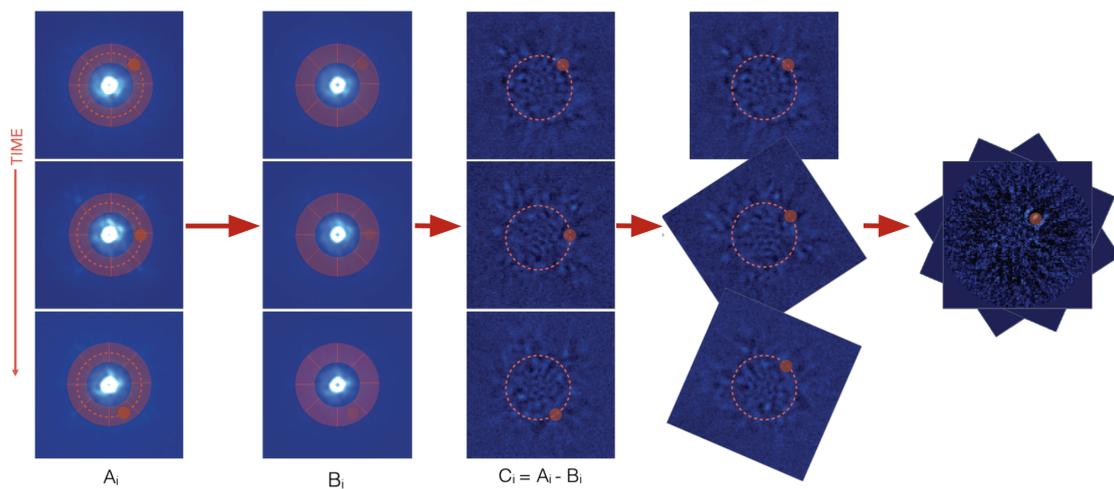


Figure 1.11: Graphical representation of the main steps of the post-processing framework of PSF-subtraction techniques, with from left to right, the calibrated science images (A), the frame-wise reference PSF (B), the residual images (C), the derotated residual images and final residual map. Adapted from C. Gomez Gonzalez (2018).

Locally Optimised Combination of Images (LOCI)

The locally optimized combination of images (LOCI, Lafreniere et al. 2007b) derives directly from the median combination approach as it also considers a combination of images to generate a model of the speckle field. However, unlike the median combination approach, LOCI computes a frame-wise reference PSF based on a linear combination of a set of reference images (a subset of the science images selected based on a field rotation criterion). Indeed, the correlation from one image to another should not be equivalent for every pair of frames and some frames should be favoured when estimating the frame-wise reference PSF. On top of being frame-wise, this approach treats independently sub-regions of the field of view to take into account local spatial correlation of the speckle noise. The coefficients of the linear combination are computed via a least square minimization of the residuals after subtraction of the speckle field. A region larger than the one for which the reference PSF is computed, is used for the least square minimization. This allows to limit the potential impact of a planetary signal during the coefficients estimation. Different flavour of LOCI have been proposed, such as damped LOCI (D-LOCI, Pueyo et al. 2012), template LOCI (T-LOCI, Marois et al. 2013), or adaptive LOCI (A-LOCI, Currie et al. 2012) among others. For a detailed presentation of the original approach, see Lafreniere et al. (2007b) or Section 3.2.2.

Principal Component Analysis (PCA)

PCA-based techniques (Soummer et al. 2012; Amara & Quanz 2012) rely on the decomposition of a set of reference images into their eigen-modes to construct the reference PSF. More precisely, the frame-wise reference PSF is defined as the projection of the considered image onto a lower-dimensional orthogonal basis (the principal components) computed from the reference images via PCA. The principal components are defined by a truncation of an orthogonal basis obtained using either eigen decomposition, a singular value decomposition (SVD), or a Karhunen-Loève transformation (KLIP, Soummer et al. 2012). The first few modes in the basis account for the larger fraction of the information content of the reference frames. The movement of a planet, when relying on ADI or SDI observing strategies, makes it a highly variable feature, observed only in the higher order modes. The main parameter of PCA-based PSF subtraction techniques is therefore the truncation order, i.e. the number of eigen modes kept for the reference PSF computation. The definition of the truncation order results from a trade-off between the quality of the speckle field representation (the amount of residuals) and the self-subtraction of the planetary signal contained in the high order modes. In the case of the RDI observing strategy, the principal components are learned from a reference star and used as basis for the projection of the science images.

As shown in Amara & Quanz (2012), full-frame PCA brings improvements both in

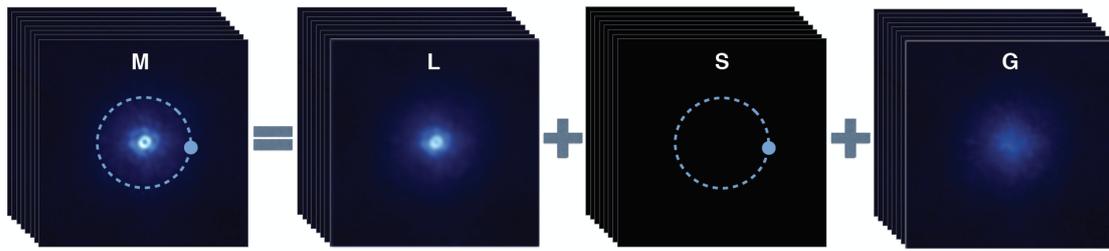


Figure 1.12: Representation of the different components of the local low-rank plus sparse plus Gaussian decomposition (Gonzalez et al. 2016).

terms of signal to noise ratio (SNR) and detectability of companions at close separations compared to LOCI. A more advanced annulus-wise PCA algorithm introduced by Absil et al. (2013), allows to treat more locally the speckle noise, further improving the SNR.

Non-negative Matrix Factorization (NMF)

The non-negative matrix factorization (NMF) algorithm shares the same low-rank matrix approximation approach as PCA. However, NMF relies on an additional non-negativity constraint on both the input matrix (cube of images) and the projection matrices. NMF aims to define a k -dimensional approximation of the input matrix in terms of non-negative factors W and H via the minimization of their Frobenius norm (Lee & Seung 1999). Ren et al. (2018) demonstrated its interest when dealing with extended structures such as debris disks.

LLSG

The Local Low-rank plus Sparse plus Gaussian-noise decomposition (LLSG, Gonzalez et al. 2016) derives again from the same low-rank matrix approximation family as PCA and NMF. LLSG applies a decomposition of the set of science images into low-rank, sparse and Gaussian components, adapting to HCI the approach proposed by Bouwmans & Zahzah (2014) in the context of computer vision. LLSG relies on the Go Decomposition robust PCA algorithm developed by Zhou & Tao (2011). The use of robust PCA decomposition, which considers multiple subspaces, enhances residual speckle noise suppression by considering an additional Gaussian components. The sparse term then mainly contains the planetary signal as demonstrated experimentally by Gonzalez et al. (2016). The LLSG algorithm improves the detectability of faint planets especially at close separations compared to full-frame PCA, but it is sensitive to outliers that can appear in the sparse term along with planetary signals.

1.5.2 Maximum likelihood techniques

We have so far presented different approaches to model the speckle field and subtract it from the set of science images to eventually unveil planetary signals. These PSF-subtraction techniques did not account for potential distortion of the planetary signal due to the subtraction of the speckle field model. The subtraction of the reference PSF or the decomposition in multiple components leads indeed to a self-subtraction of the planetary signal. Another family of post-processing algorithms replace or complement the reference PSF subtraction by a forward modelling of the planetary companion along with a maximum likelihood based estimation of the planetary flux. These techniques model and track the planetary signal based on the knowledge of the speckle noise statistics, the expected planetary movement (depending on the observing strategy), and the impact of the speckle field modelling. Based on this knowledge and on assumptions about the speckle noise distribution, these techniques rely on a maximum likelihood approach to provide an estimated contrast for every position within the field of view, alongside a SNR-based detection map.

ANDROMEDA

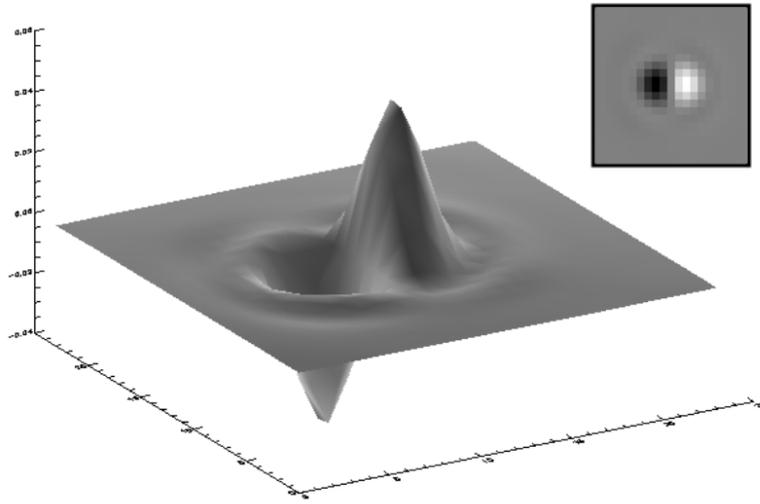


Figure 1.13: Planet signatures generated after performing the ADI pairwise subtraction (Cantalloube et al. 2015).

The ANgular DiffeRential OptiMal Exoplanet Detection Algorithm (ANDROMEDA, Mugnier et al. 2009; Cantalloube et al. 2015) relies on pairwise subtraction of adjacent images to reduce speckle noise, making the assumption that the speckle noise in adjacent images is highly correlated. The pairwise frame differencing ensures a sufficient stability of the speckle noise to allow an effective noise suppression, the autocorrelation

of images in HCI being often very high on short timescales ¹¹. It also allows to create a very specific signature for the planetary signal as can be seen from Figure 1.13. Based on the known signature of the planetary signal (forward model) after pairwise subtraction, ANDROMEDA uses a maximum likelihood approach to compute an estimated planetary flux for every position within the field of view. The residual noise after subtraction is assumed to follow Gaussian statistics, allowing the use of an analytical expression for the planetary flux estimation.

KLIP FMMF

One of the main drawbacks of the PCA-based PSF-subtraction techniques is the partial suppression of the planetary signal via self-subtraction (and over-subtraction, see Section 3.2 for more information). The KLIP forward-model matched filter approach developed by Pueyo (2016) and further improved by Ruffio et al. (2017), derives from the Karhunen-Loève transformation an analytical expression to compute a forward-modeled planetary template for the KLIP PSF-subtraction technique (Soummer et al. 2012). The forward-modeled planetary template incorporates, to the first order, the perturbations of the PSF due to the subtraction of the speckle field model estimated with KLIP (self-subtraction and over-subtraction). A Gaussian maximum likelihood approach similar to ANDROMEDA is then used alongside the forward-modeled planetary template to compute the flux at any point in the field of view.

PACO

The exoplanet detection algorithm based on PATCH COvariances (PACO, Flasseur et al. 2018) differs from the other techniques by the fact that it directly tries to model the speckle noise without any reference PSF-subtraction, avoiding self and over-subtraction of the planetary signal at short separations. Based on the set of science images, the PACO algorithm tries to infer both local spatial correlations and speckle noise statistics via a nonstationary multi-variate Gaussian model. In the case of the ADI observing strategy, it considers for each point within the field of view and each frame, the set of patches centred on the considered pixel in all the other science images to infer the speckle noise statistical properties. It then relies on a maximum likelihood estimator based on a multi-variate Gaussian model to compute the estimated flux using the off-axis PSF as a planetary template. The existence of a planetary candidate is then inferred at a predefined level via a generalized likelihood ratio test.

¹¹This pairwise subtraction can be seen as a PSF-subtraction technique using the second image as a reference PSF for the first one

TRAP

TRAP (Samland et al. 2017) tracks the planetary signal along the temporal axis and models locally the speckle field by considering specific areas that share similar noise statistics. A temporal model of the behaviour of the speckle pattern is estimated for each pixel, by considering a set of carefully chosen pixels that should share similar noise statistical properties but do not include the planetary signal. A model of transiting exoplanet is then fitted alongside the obtained temporal model describing the speckle noise to best fit the temporal evolution of the pixel intensity and determine the presence or absence of planetary signal. The TRAP algorithm aims to improve the reachable contrast at small angular separations ($<3\lambda/D$) compared to models based on spatial correlations between de-rotated images.

Half-sibling regression

Gebhard et al. (2022) proposed recently an image processing technique relying also on an inverse problem approach but without using likelihood maximization. Similarly to TRAP (Samland et al. 2017), the speckle noise associated to the times series of a given pixel is modelled based on a set of pixels situated in specific areas that should ensure some mutual information. Two regions of interest are considered, a circular region symmetrically across the origin from the studied pixel and a region around the pixel itself, but not causally linked to it. Half-sibling regression is then used to learn a model of the pixel speckle noise, based on the information contained in these two regions, and predict the systematic noise associated to the selected pixel. In order to avoid overfitting and self-subtraction, a model of the planet or a mask is applied on the time series. The position of the planet being not known, different models are estimated with the planet signal/mask positioned at different time step. These different noise models are then used to construct an hypothesis map, by estimating for each pixel the most probable temporal position for a planetary signal. A detection map is finally computed by confronting this hypothesis map with every possible planetary trajectories.

1.5.3 Supervised machine learning techniques

Another very promising research avenue for HCI consists in reformulating the exoplanet detection task as a supervised binary classification problem. This approach, which relies on machine learning techniques such as decision trees or deep neural networks, was first proposed by Gomez Gonzalez et al. (2018) in the context of direct imaging, with the Supervised exOplanet detection via Direct Imaging with deep Neural Networks (SODINN) algorithm. The main difficulty associated to the use of machine learning techniques pertains to the large number of labelled samples required for training. This difficulty is exacerbated with direct imaging, by the large imbal-

ance between positive and negative classes (resp. planets and speckle noise), due to the lack of confirmed exoplanet detections. Two different approaches have been proposed to generate the negative class, either by relying on standard machine learning data-augmentation techniques (Gomez Gonzalez et al. 2018) or by using a Generative Adversarial Network (Yip et al. 2020). This latter approach should help removing potential undetected exoplanets from the samples of speckle noises. Regarding the positive class, the samples are obtained by injecting, in negative class samples, synthetic exoplanets at different SNR levels, using an off-axis PSF of the host star as planetary model. The approaches proposed by Gomez Gonzalez et al. (2018) and Yip et al. (2020) apply then a neural network based on convolutional layers for pattern recognition to finally generate a detection map.

SODINN

The deep neural network in SODINN is fed by Multi-level Low-rank Approximation Residual (MLAR) samples estimated at every position within the field of view. The 3D MLAR samples are square residual patch sequences with a size of two FWHM, obtained by subtracting from ADI sequences reference PSFs computed at different ranks via PCA. The idea behind this multi-rank estimation is that a planetary signal is expected to show up consistently at different ranks, whereas residual noise patterns evolve with the number of principal components.

A multilevel information aggregation scheme is then applied to reduce drastically the information content of the 3D MLAR cubes. The proposed two-level aggregation scheme is composed of one layer of convolutional Long-Short Term Memory networks (LSTM, Shi et al. 2015) and a pooling layer. The resulting aggregated information is then sent to a fully connected layer featuring 128 hidden units followed by a rectified linear unit (Nair & Hinton 2010) activation. A sigmoid function is eventually used as output layer to allow a binary classification at a given threshold for each pixel (i.e., each MLAR sequence), providing the final detection map. The neural network is trained by considering two classes, one regrouping MLAR samples in which fake planetary signals have been injected, and one regrouping MLAR samples of speckle noise. The training and validation of the neural network requiring a large number of samples, standard machine learning data-augmentation techniques are used to increase the number of samples in the two classes.

Generative Adversarial Network and Convolutional Neural Network

Yip et al. (2020) take advantage of the presence of speckle noise in all direct images to train a Deep Convolutional Generative Adversarial Network (DCGAN) to learn a model of the speckle pattern and avoid possible contamination by undiscovered exoplanets, which could be wrongfully included in the negative class by standard data-

augmentation techniques. Multiple Convolutional Neural Network (CNN) classifiers are then trained and tested on negative and positive classes generated by the DCGAN based on images from the Hubble Space Telescope (HST) instrument NICMOS. Synthetic exoplanets are injected across a wide range of SNRs to test the ability of the CNNs to uncover planetary signal. The tested CNN classifiers exhibit a good performance in terms of Precision and Recall map and allow to re-confirm bright source detections in HST images.

1.6 Scope and outline of the dissertation

In this introduction, we have highlighted the many advantages of direct imaging and its complementarity with indirect methods, both in terms of retrieved physical properties and exoplanet population. With the future 20-40m class ground-based telescopes (e.g. ELT, TMT, or GMT) and the JWST space telescope, researchers will soon have access to an entire new population of medium to small-size planets via direct imaging. These bright perspectives, as well as the technological revolution of the past decades helped by facilities such as the VLT or the Keck observatory, contributed to the growing attention given to HCI by the scientific community. The various developments in adaptive optics systems, coronagraphic devices and advanced data processing techniques covered in this introduction, allowed to reach in recent years unprecedented contrast, below 10^{-6} .

This PhD thesis builds upon these recent advances, focusing on the development of new data processing techniques to make the most of current hardware technologies and prepare for the future. More specifically this thesis is devoted to the development of new post-processing techniques to unveil fainter planetary signals from both ADI and SDI calibrated sequences, and to retrieve their observed properties (position and flux).

The existence of a large number of different HCI post-processing techniques (as seen in Section 1.5) with their own peculiarities motivated my initial search for an algorithm that could take advantage of their strengths, while mitigating at least partly their drawbacks. I started with the simple observation that most PSF-subtraction techniques include a common final step to compute detection maps (SNR map) based on their PSF-subtracted images, relying mainly on the approach of Mawet et al. (2014). The main idea was therefore to start from the cubes of residuals generated by different post-processing techniques, and replace the estimation of the SNR map by a new aggregation scheme. Each PSF-subtraction technique treating differently the speckle field, their optimal combination should boost the strength of planetary signals, and further reduce the residual speckle noise.

Other approaches have been proposed to replace the SNR map estimation, such as the weighted angular differential imaging of Bottom et al. (2017), or the STIM

map of Pairet et al. (2019), which accounts for the non-Gaussianity of the residual speckle noise when computing the SNR map. However, these approaches are applied on a single PSF-subtraction technique at a time and do not consider the potential of combining several residuals map to enhance the achievable contrast.

The second chapter of this thesis explores the possibility of combining the information of several cubes of residuals (before their mean-combination, see Figure 1.10), using a Markov regime-switching model to generate a single probability detection map. Based on concatenated series of pixel-wise time sequences, the model estimates annulus-wise the probability associated with two regimes, the first one accounting for the residual noise and the second one for the planetary signal in addition to the residual noise. This first version of the so-called RSM map algorithm relies on three PSF-subtraction techniques: Annular PCA, LLSG and NMF. The performance of the algorithm is tested on data sets from two instruments, VLT/NACO and VLT/SPHERE, and the results indicate an overall better performance in the receiver operating characteristic (ROC) space when compared with standard SNR maps.

The third chapter is devoted to several improvements of the RSM map framework, first by considering, forward-modelled PSFs to account for the distortions resulting from reference PSF subtraction. Two forward-model versions of the RSM map algorithm are implemented, one based on the LOCI and the other on the KLIP PSF-subtraction techniques. We also investigate the added value of a forward-backward approach for the probability estimation. The forward-backward approach leads to a clear improvement of the astrometry precision at the cost of a reduced performance in terms of faint planetary signal detection. A new framework to generate contrast curves based on probability maps is also proposed in this third chapter.

Considering the improvements made in the third chapter, seven different PSF-subtraction techniques, including two forward-model versions, can be used within the RSM framework. As a result, the selection of the optimal parameters for these PSF-subtraction techniques as well as for the RSM map is not straightforward, time consuming, and can be biased by assumptions made on the underlying data set. We propose in chapter four, an unsupervised optimization framework for both the PSF-subtraction techniques and the RSM map algorithm, which allows a data-driven estimation of the optimal parameters. The optimisation procedure, called auto-RSM consists of three main steps, (i) the definition of the optimal set of parameters for the PSF-subtraction techniques using contrast as performance metrics, (ii) the optimisation of the RSM algorithm, and (iii) the selection of the optimal set of PSF-subtraction techniques and ADI sequences used to generate the final RSM probability map. The results obtained by applying this optimization framework on the data sets of the exoplanet imaging data challenge (EIDC, Cantalloube et al. 2020b) demonstrate clearly the interest of the proposed optimisation procedure, with better performance metrics compared to the earlier version of RSM, as well as to other HCI data processing techniques.

Chapter two to four allowed the development of a robust detection algorithm, unsupervised to a large extent thanks to the auto-RSM framework, and which has the ability of characterizing planetary candidates. In chapter five, we apply the auto-RSM framework to the SHARDDS survey to bring an additional piece to the exoplanet puzzle, by contributing to the characterisation of planetary population via the estimation of occurrence rate maps. The SHARDDS survey gathers observing sequences generated with the VLT/SPHERE instrument for 55 main-sequence stars within 100 pc, known to host a high-infrared-excess debris disk. The definition of constraints for planet population in terms of physical properties such as mass or semi-major axis is necessary to better understand planet formation and evolution, and try to develop a global formation model. Alongside the estimation of detection limits, a planetary characterisation algorithm based on the auto-RSM framework is developed to determine more precisely the astrometric and photometric properties of planetary candidates. This last chapter is also an opportunity to better characterise the auto-RSM framework by determining the degree of similarity existing between the optimal parametrisations obtained for different observing sequences. This should help us assess the necessity to consider multiple optimal parametrisations for large surveys and the potential impact in terms of reachable contrast. Finally, Chapter six summarises and concludes on the work done, providing as well future perspectives for potential improvements of the RSM framework and its future use in other surveys.

Chapter two to five are partial reproductions of articles that have been published in the *Astronomy and Astrophysics* journal (see appendix E for the articles abstracts and references). Their introduction and conclusion have been modified to allow for a smooth transition between the different chapters. Redundant definitions have been removed and more detailed descriptions have been included, especially in chapter five. The mathematical notation has also been modified to ensure its coherence throughout the thesis manuscript.

Chapter 2

RSM detection map for direct exoplanet detection in ADI sequences

2.1 Introduction

In this first chapter, we propose a novel approach dealing with the last step of the reference PSF-subtraction based post-processing framework (see Figure 1.10). Instead of averaging the set of de-rotated images obtained after the reference PSF subtraction and compute an S/N map, we propose to consider the entire set of residual frames and rely on a regime-switching algorithm to classify the pixels into two categories, regrouping either the planetary signals or the quasi-static speckles. The probability associated with the planetary regime then allows the creation of a detection map. The algorithm derives from the Markov regime-switching model first proposed by Hamilton (1988), which is widely applied to analyse economic and financial time series. The aim of our new detection algorithm is to treat more effectively the residual noise still observed in the cube of residuals provided by ADI methods, increasing our ability to disentangle faint signals from bright speckles. The flexibility of the algorithm allows the use of ADI cubes treated with most post-processing methods. The cubes of residuals obtained from the different post-processing methods may be used separately but also together, improving further the sensitivity of the detection algorithm to faint companions.

This chapter is organized as follows. In Section 2.2, we describe the new regime-switching model for the detection of exoplanets. Section 2.3 presents in detail the model estimation and the definition of the different parameters. The ability of our model to disentangle faint planetary signals from bright speckles is tested in Section 2.4, by injecting fake companions into two different data sets and by comparing the results with state-of-the-art ADI-based post-processing techniques, using receiver operating curves.

2.2 Regime-switching model

The proposed detection algorithm derives from the Markov-switching regressions introduced by Goldfeld & Quandt (1973) and Cosslett & Lee (1985), further improved by Hamilton (1988, 1994), who developed an iterative inference algorithm to estimate the model parameters, namely the Markov Regime-Switching Model (RSM). This approach is one of the most popular non-linear time series models in the econometric literature and many variants have been proposed. The aim of the RSM is to take into account possible dramatic changes in the behaviour of time series such as the transition between economic expansion and contraction in the case of financial time series. The regime-switching model relies on several linear equations to describe the different states in which a system described by a time series can be. The probability of being in a given state depends on both a pre-defined transition probability and on the ability of the different equations to describe properly the evolution of the time series. One of the model outcomes is the probability associated with the different regimes. The RSM provides for each element of the time series, the probability of being in the different regimes. Our detection map derives directly from these probabilities.

In the case of our RSM detection map, the time series is built from the de-rotated cube of residuals obtained after the reference PSF subtraction and de-rotation steps of the ADI sequence post-processing. Several cubes of residuals treated with different ADI PSF subtraction techniques may be stacked in the time axis to provide additional information and increase the ability of the model to detect faint companions. To allow for the detection of planetary signals, we rely on two different regimes to model the pixels' intensity in the de-rotated cube of residuals: a regime in which the residuals time series is described by speckle noise and a second regime with speckle noise plus a planetary signal. The planetary signal may be modelled as the measured off-axis PSF¹ or as a forward model of the off-axis PSF after the subtraction. We consider in this chapter the measured off-axis PSF for simplicity, although the algorithm may be easily adapted to a forward modelled off-axis PSF, as demonstrated in chapter 3.

The regime-switching model we propose here is a modified version of the original Markov-switching model, in which only one parameter is determined via a maximum log-likelihood estimation. We rely on the characteristics of the data set to define the other model parameters. Having presented the basic principles behind our regime switching model, we may now describe the detailed procedure for our RSM detection map computation.

¹For coronagraphic imaging, an off-axis non-coronagraphic image of the target is routinely acquired before and after the observing sequence. This PSF reference is used to calibrate the flux of the star and provide a model of the planetary signal for forward model-based algorithm. For non-coronagraphic imaging, this reference PSF is the unsaturated exposure.

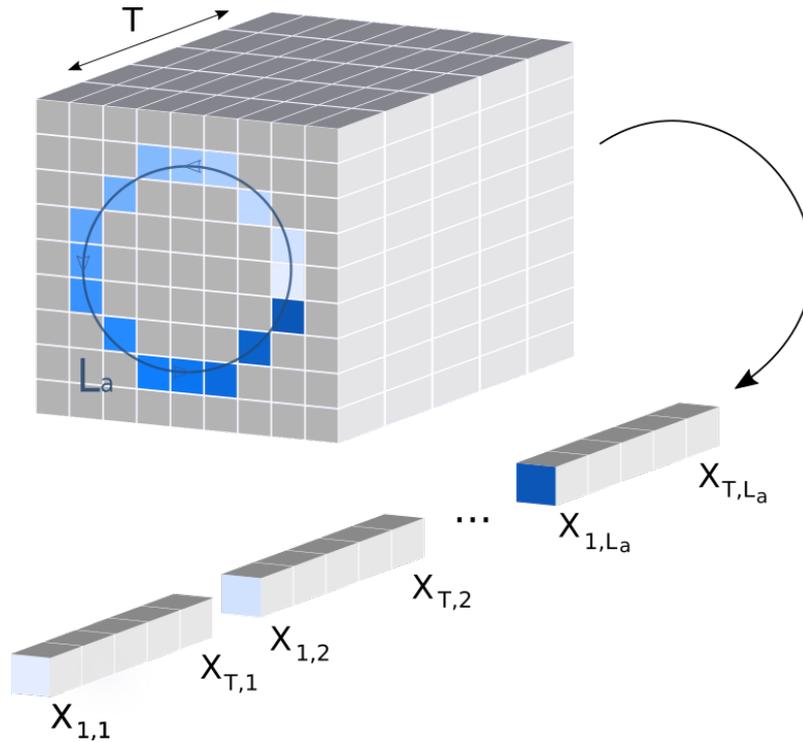


Figure 2.1: The residuals time series for a given annulus a is obtained by stacking the pixels values of the considered annulus along the time axis.

2.2.1 Building the time series

The first step of our estimation procedure is to build the time series that the regime switching model will try to model. As the noise properties are expected to evolve with radial distance, the regime switching model is applied annulus-wise. For each annulus a , a specific residuals time series X_{i_a} is built by vectorizing that part of the cube of residuals, indexed by i_a the flattened pixel number. The length of the time series X_{i_a} depends on the number of pixels in the considered annulus L_a but also on the number of frames in the original de-rotated cube of residuals T . We indeed take advantage of all the individual frames contained in the de-rotated cube of residuals instead of collapsing the cube as it is usually done when estimating an S/N map. As can be seen from Figure 2.1, the time series X_{i_a} is built by concatenating the set of T observations for every pixel contained in the annulus a , i.e. $X_{i_a} = \{X_{1,1}, X_{2,1}, \dots, X_{T,1}, X_{1,2}, \dots, X_{T,2}, \dots, X_{T,L_a}\}$ with $i_a \in \{1, \dots, T \times L_a\}$. The first subscript of X indicates the selected frame in the de-rotated cube of residuals, while the second one provides the position of the considered pixel in the selected annulus a . Both subscripts are replaced by a single index i_a to form the residuals time series that feeds the regime-switching model.

We consider first the time axis and then the spatial axis in order to stay, during the iterative process used to build the detection map, in the planetary regime during

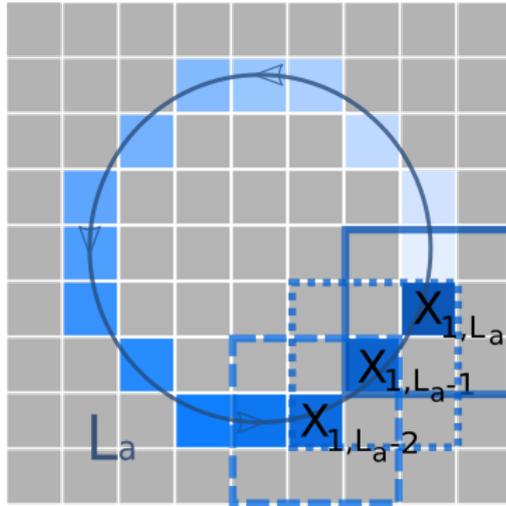


Figure 2.2: The time series \mathbf{X}_{i_a} is created by considering matrices of dimension $\theta \times \theta$ centred on every X_{i_a} in the cube of residuals. The figure shows the residuals matrices obtained from the first frame of the cube of residuals for the last three pixels of the annulus with θ equal to 3.

T steps instead of switching T times between both regimes when a planetary signal is present in a given annulus. Indeed, when travelling through the residuals time series, the planetary signal observed in a given pixel will act on the regime-switching model during T steps, allowing the probability to be in the planetary regime to build up thanks to the short term memory of the model. This helps to enhance the sensitivity of the algorithm to faint signal as it allows the probability to built up for a longer period.

2.2.2 Model description

The second step of the RSM detection map computation consists in defining the set of equations describing the residuals time series for the two considered regimes. In the first regime, the time series X_{i_a} is described by a residual noise following the statistics of the quasi-static speckle residuals contained in the annulus. In the second regime, the time series X_{i_a} is described by both the residual noise and the planetary signal model (off-axis PSF). The PSF being two-dimensional, we consider not only one pixel at a time but a batch of pixels in a square of size θ equal to the full width at half maximum (FWHM) of the PSF. In order to define the probability of observing a planetary signal at a given pixel x_{i_a} , we therefore need to consider a number of neighbouring pixels depending on the value of θ . As depicted in Figure 2.2, we define \mathbf{x}_{i_a} , the residuals matrices of dimension $\theta \times \theta$ centred on x_{i_a} , which will replace the time series X_{i_a} used so far. Larger values of θ may be considered in the case of a forward modelled off-axis PSF to take into account the signal self-subtraction, which, for instance, could create

negative wings in the azimuthal direction. Our regime switching model is therefore characterized by the following equations:

$$\mathbf{x}_{i_a} = \mu + \beta F_{i_a} \mathbf{m} + \boldsymbol{\varepsilon}_{s,i_a} = \begin{cases} \mu + \boldsymbol{\varepsilon}_{0,i_a} & \text{if } S_{i_a} = 0 \\ \mu + \beta \mathbf{m} + \boldsymbol{\varepsilon}_{1,i_a} & \text{if } S_{i_a} = 1 \end{cases} \quad (2.1)$$

where β provides the strength of the planetary signal, μ the mean of the quasi-static speckle residuals and $\boldsymbol{\varepsilon}_{s,i_a}$ their time and space varying part characterized by the quasi-static speckle residuals statistics (see Table A.1 for a summary of all the variables used in the regime switching model). \mathbf{m} is the model of the planetary signal, which is the normalized off-axis PSF in the FWHM region.

As can be seen from Eq. 2.1, there exist two possible states S_{i_a} , which are reflected in the value taken by the parameter F_{i_a} , with $F_{i_a} = 1$ in the case of a planetary signal detection and $F_{i_a} = 0$ in the other case. S_{i_a} is not directly observable, but we see its effect on the behaviour of \mathbf{x}_{i_a} via the realization F_{i_a} .

The parameter F_{i_a} is a realization of a two-states Markov chain allowing a short term memory. This implies that we only consider the state S_{i_a-1} in which the system was at index $i_a - 1$ to define the probability of being in a given state S_{i_a} for the current index i_a . The fact that the realization F_{i_a} is a probabilistic outcome implies that we cannot consider being in only one of the two regimes. We have instead a given probability of being in each of them. Our regime switching model tries to describe the behaviour of the time series \mathbf{x}_{i_a} via a probability weighted sum of the values generated by the equation describing each regime.

2.2.3 Definition of the model probabilities

The probability of \mathbf{x}_{i_a} being in a state or regime $S_{i_a} = s$ is characterized by the set of parameters of Eq. 2.1, i.e. \mathbf{m} the planetary signal model, and μ and β , the statistical properties of the residual noise $\boldsymbol{\varepsilon}_{s,i_a}$. We make here the simplifying assumption that the quasi-static speckles residuals $\boldsymbol{\varepsilon}_{s,i_a}$ may be characterized to a good level of precision by their mean μ and variance σ . We write the probability of observing \mathbf{x}_{i_a} in the state s at step i_a as follows:

$$\xi_{s,i_a} = P(S_{i_a} = s | \Omega_{i_a}, \mathbf{m}, \mu, \beta, \sigma), \quad (2.2)$$

where \mathbf{m} , μ , β , σ and $\Omega_{i_a} = \{\mathbf{x}_{i_a}, \mathbf{x}_{i_a-1}\}$ provide the parameters of the model.

This probability ξ_{s,i_a} is the key element of our RSM detection map as the map is constructed based on the value taken by ξ_{1,i_a} for every pixel of every annulus.

Indeed, ξ_{1,i_a} provides a detection probability for each pixel and each frame of the de-rotated cube of residuals. The final RSM detection map is created by averaging these probability along the time axis of the cube of residuals.

In the case of a two-state Markov chain, the computation of ξ_{s,i_a} necessitates the estimation of (i) the probability ξ_{q,i_a-1} of observing the system in the state q at step $i_a - 1$, (ii) the transition probability $p_{q,s}$ from state q to state s and (iii) the likelihood of observing \mathbf{x}_{i_a} in state s at step i_a , which we note η_{s,i_a} . The probability of being in a state s at index i_a can be computed as the normalized likelihood of being in state s at index i_a multiplied by the probability of having been in either of the two states at index $i_a - 1$ and by the transition probability $p_{q,s}$, which accounts for the short term memory of the algorithm. The expression of the state probability ξ_{s,i_a} is therefore given by the following expression (Hamilton 1988):

$$\xi_{s,i_a} = \sum_{q=0}^1 \frac{\eta_{s,i_a} p_{q,s} \xi_{q,i_a-1}}{f(\mathbf{x}_{i_a} | \Omega_{i_a-1}, \mathbf{m}, \mu, \beta, \sigma)}, \quad (2.3)$$

with the sum f of conditional densities for index i_a given by:

$$f(\mathbf{x}_{i_a} | \Omega_{i_a-1}, \mathbf{m}, \mu, \beta, \sigma) = \sum_{q=0}^1 \sum_{s=0}^1 \eta_{s,i_a} p_{q,s} \xi_{q,i_a-1}, \quad (2.4)$$

and the transition probabilities given by:

$$p_{q,s} = \text{P}(S_{i_a} = s | S_{i_a-1} = q), \quad (2.5)$$

with $q, s \in \{0, 1\}$. We consider the two possible states describing the system at index $i_a - 1$ via the sum over q . The function $f(\mathbf{x}_{i_a} | \Omega_{i_a-1}, \mathbf{m}, \mu, \beta, \sigma)$, which represents the numerator summed over the two possible states taken at index i_a , ensures that the sum of the probability ξ_{s,i_a} equals one for every index i_a .

2.2.4 Transition probabilities estimation

For our two-regime model, the transition probability $p_{q,s}$ regroups the probabilities of staying in either regimes along with the probabilities of switching to the other regime. The estimation of $p_{q,s}$ is relatively straightforward by imposing to the algorithm the potential existence of no more than one planetary signal per annulus. A number of planetary signal per annulus in the interval $]0, 1]$ may therefore be considered. Following our testing, a value of one companion per annulus must be privileged in the

case of faint companions as lower values decrease both the residual speckles and the companion intensities in our model. Considering the number of pixels L_a and the number of frames T , the parametrization of $p_{q,s}$ translates as follows in the case of one planetary signal per annulus:

$$p_{q,s} = \begin{pmatrix} p_{0,0} = 1 - 1/(T \times L_a) & p_{1,0} = 1/T \\ p_{0,1} = 1/(T \times L_a) & p_{1,1} = 1 - 1/T \end{pmatrix}. \quad (2.6)$$

2.2.5 Likelihood function definition

The determination of the likelihood is the key step of the model estimation. The challenge is to select the right probability distribution function to describe properly ε_{s,i_a} , the residual noise due to the quasi-static speckles. Indeed, the value taken by η_{s,i_a} depends directly on the position of the elements of \mathbf{x}_{i_a} , or the elements of $\mathbf{x}_{i_a} - \beta\mathbf{m}$, in the probability distribution of the quasi-static speckle residuals. Considering the small transition probabilities $p_{0,1}$, the probability of planetary signal detection ξ_{1,i_a} depends heavily on the value taken by η_{1,i_a} . The parametrization of the selected probability distribution function plays as well an important role.

Different probability distribution functions may be used. For the sake of clarity, we illustrate the likelihood function definition with a simple Gaussian distribution as it is done in Hamilton (1988). However, the next section will allow us to investigate the question of the optimal probability distribution function selection as different post-processing algorithms provide different noise distributions for different separations. The Gaussian distribution allows us to construct a likelihood function for state s at index i_a in the following manner:

$$\eta_{r,i_a} = \sum_n^{\theta^2} \frac{1}{\theta^2} \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\frac{(\mathbf{x}_{i_a}^n - F_{i_a}\beta\mathbf{m}^n - \mu)^2}{2\sigma^2} \right], \quad (2.7)$$

with n the index of the matrix elements for \mathbf{x}_{i_a} and \mathbf{m} . The sum over the matrix elements allows us to get only one value per considered $\theta \times \theta$ patch.

2.2.6 Model estimation

Since the estimation of ξ_{s,i_a} depends on its value at the previous step, we rely on an iterative procedure to estimate the entire set of ξ_{s,i_a} . This iterative procedure requires the definition of an initial condition for $\xi_{q,0}$. Assuming that the considered Markov chain is ergodic, we can simply set $\xi_{q,0} = P(S_t = q \mid \mathbf{m}, \mu, \beta, \sigma)$ equal to the unconditional probability $\xi_{q,0} = P(S_t = q)$. Following the approach proposed

by Hamilton (1994), the two initial probabilities $\xi_{0,0}$ and $\xi_{1,0}$ may be estimated via following system of equations:

$$\begin{cases} \xi_{0,0} = \xi_{0,0}p_{0,0} + \xi_{1,0}p_{1,0} \\ \xi_{1,0} = \xi_{1,0}p_{1,1} + \xi_{0,0}p_{0,1} \\ \xi_{1,0} + \xi_{0,0} = 1 \end{cases} \quad (2.8)$$

which translates in terms of matrices in:

$$\mathbf{A}\boldsymbol{\xi} = \boldsymbol{\psi} \quad (2.9)$$

with $\boldsymbol{\xi} = [\xi_{0,0}, \xi_{1,0}]$ the set of initial probabilities, $\boldsymbol{\psi} = [0, 0, 1]$ and \mathbf{A} given by:

$$\mathbf{A} = \begin{pmatrix} I_{2 \times 2} - P \\ 1 & 1 \end{pmatrix}, \quad (2.10)$$

with P the matrix of $p_{q,s}$, $I_{2 \times 2}$ a diagonal matrix of dimension 2×2 . Solving the system of equations (eq.2.8) to obtain the initial probabilities $\boldsymbol{\xi}$, is then equivalent to taking the third row of the matrix $(\mathbf{A}^t \mathbf{A})^{-1} \mathbf{A}^t$.

2.3 Detection map estimation

We propose in this section a procedure to produce a RSM detection map. The model we developed so far necessitates the computation of cubes of residuals along with the definition of several parameters: the probability distribution function of the quasi-static speckles residuals $\boldsymbol{\varepsilon}_{s,i\alpha}$ and its first two moments, the planetary signal model \mathbf{m} , the intensity parameter β , and the transition probability $p_{q,r}$. The transition probability $p_{q,r}$ is already defined in Section 2.2.4, we therefore consider the remaining three model parameters.

2.3.1 Computation of de-rotated cubes of residuals

The first step to create a RSM detection map, is the production of the de-rotated cubes of residuals for the selected ADI-based post-processing techniques feeding our regime-switching algorithm. As an illustration of the ability of our model to improve the detection when considering several methods at once, we consider in this chapter three different post-processing techniques: annular PCA, NMF and LLSG. For the two first approaches, the estimation of the cubes of residuals starts with the definition of a reference PSF. Annular PCA follows the PCA principles by computing the directions of maximal variance from the main matrix representing the ADI sequence, $\mathbf{M} \in \mathbb{R}^{n \times p}$,

with n the number of frames and p the number of pixels in the considered annulus. The determination of a reference PSF is done via the estimation of the eigenvectors \mathbf{V} of the matrix \mathbf{M} by taking \mathbf{V}_k , the first k components of \mathbf{V} . Annular PCA relies on a separate estimation for each annulus composing the original cube of data to take into account the radial evolution of the noise distribution. It allows to consider the local structure of the speckle noise instead of the entire frame. The cube of residuals is then obtained via the subtraction of the low rank matrix $\mathbf{M}\mathbf{V}_k^T\mathbf{V}_k$ from the initial ADI sequence \mathbf{M} .

As for Annular PCA, NMF can be understood as a low rank approximation, with an additional non-negativity condition. This method consists in the decomposition of a matrix into two factors of non-negative values via the minimization of the Frobenius norm:

$$\operatorname{argmin}_{\mathbf{W}, \mathbf{H}} \frac{1}{2} \|\mathbf{M} - \mathbf{W}\mathbf{H}\|_{FN}^2 = \frac{1}{2} \sum_{i,j} (M_{i,j} - WH_{i,j})^2, \quad (2.11)$$

where $\mathbf{W} \in \mathbb{R}^{n \times k}$ and $\mathbf{H} \in \mathbb{R}^{k \times p}$. The method allows the definition of a matrix $\mathbf{W}\mathbf{H}$ with rank k lower than the one of original matrix \mathbf{M} , keeping only the main components of \mathbf{M} . The matrix $\mathbf{W}\mathbf{H}$ provides a reference PSF for the entire set of frames representing the structure of the residual starlight. As for annular PCA, this matrix is subtracted from the original ADI sequence to obtain the cube of residuals, $\mathbf{M} - \mathbf{W}\mathbf{H}$.

Finally, the LLSG estimation is based on the decomposition of the cube intensities in three separate components, \mathbf{L} a low rank matrix, \mathbf{S} a sparse matrix expected to contain the potential planetary signal and \mathbf{G} the Gaussian part of the background noise. This explains partly why the distribution of the resulting residuals, observed in Figure 2.3, is far from being Gaussian, the Gaussian part of the noise having already been removed. More information about the algorithm may be found in Gonzalez et al. (2016). The cube of residuals is directly provided by \mathbf{S} .

2.3.2 Probability distribution function

We then move to the model parameters definition by first considering the selection of the probability distribution function describing the speckles residuals. Figure 2.3 (a-d) provides the distribution of the residuals for a VLT/NACO ADI sequence (see Section 2.4 for a description of the data set) obtained with respectively the annular PCA, the Non-negative Matrix Factorization (NMF) and the Local Low Rank plus Sparse plus Gaussian (LLSG) methods. We see from these graphs that the distribution of the residuals is either close to a Lapacian or to a Gaussian distribution depending on the selected post-processing techniques and on the angular separation. At small angular separations, the tails of the distributions of the residuals seem to be closer

to a Laplacian, while at larger separation they seem closer to a Gaussian, except for LLSG processing. This radial evolution is mainly due to the higher (relative) number of intense speckles, the lower number of pixels and the lower field rotation at small separation. Overall, the distribution of the residuals is close to a Gaussian for annular PCA and NMF, and close to a Laplacian for LLSG. This confirms partially the findings of Pairet et al. (2019), who demonstrated that the residuals were closer to a Laplacian than a Gaussian distribution, especially when looking at the tails of the distribution.

The results of Figure 2.3 illustrate the difficulty of defining the residuals distribution as there exists a dependence on both the separation and the post-processing technique along with differences between the tails and the core of the distribution. We propose therefore to consider both the Gaussian and Laplacian distributions in the performance assessment of Section 2.4.

The proposed regime-switching model provides a local detection probability as it considers one annulus at a time. The parameters of the residuals probability distribution should therefore be estimated locally. We have seen in the previous section that we considered not a single pixel at a time but a $\theta \times \theta$ matrix of pixels centred on the pixel of interest. We therefore estimate empirically the pixel-wise mean and variance of the residuals by considering an annulus with a width of θ pixels, centred on the selected annulus. The entire set of frames is used for the estimation of these two parameters. Although planetary signal may be included in the annulus, the effect of this signal on the estimation of the mean and variance is limited and decreases with angular separation.

2.3.3 Intensity parameter

For the estimation of the intensity parameter β , we propose to rely on the estimated variance of the pixel intensity in the annulus. We were inspired here by the signal to noise ratio maps that are usually created using the final frame provided by most of the PSF-subtraction techniques. We define the intensity parameter β as a multiple of the estimated variance σ :

$$\beta = \delta\sigma. \tag{2.12}$$

The β parameter is the only parameter we propose to estimate via a maximum log-likelihood. Several values of δ are tested in a given interval starting at $\delta = 1$, as $\delta = 0$ would imply a single regime model. The optimal δ in an annulus a is the one leading to the highest log-likelihood sum $\sum_{i_a}^{L_a \times T} \log [f(\mathbf{X}_{i_a} | \Omega_{i_a-1}, \mathbf{P}, \mu, \beta, \sigma)]$ in the considered interval.

Relying on this definition of β allows to get information about the position of the detected planetary signal inside the probability distribution of the residual speckles. A

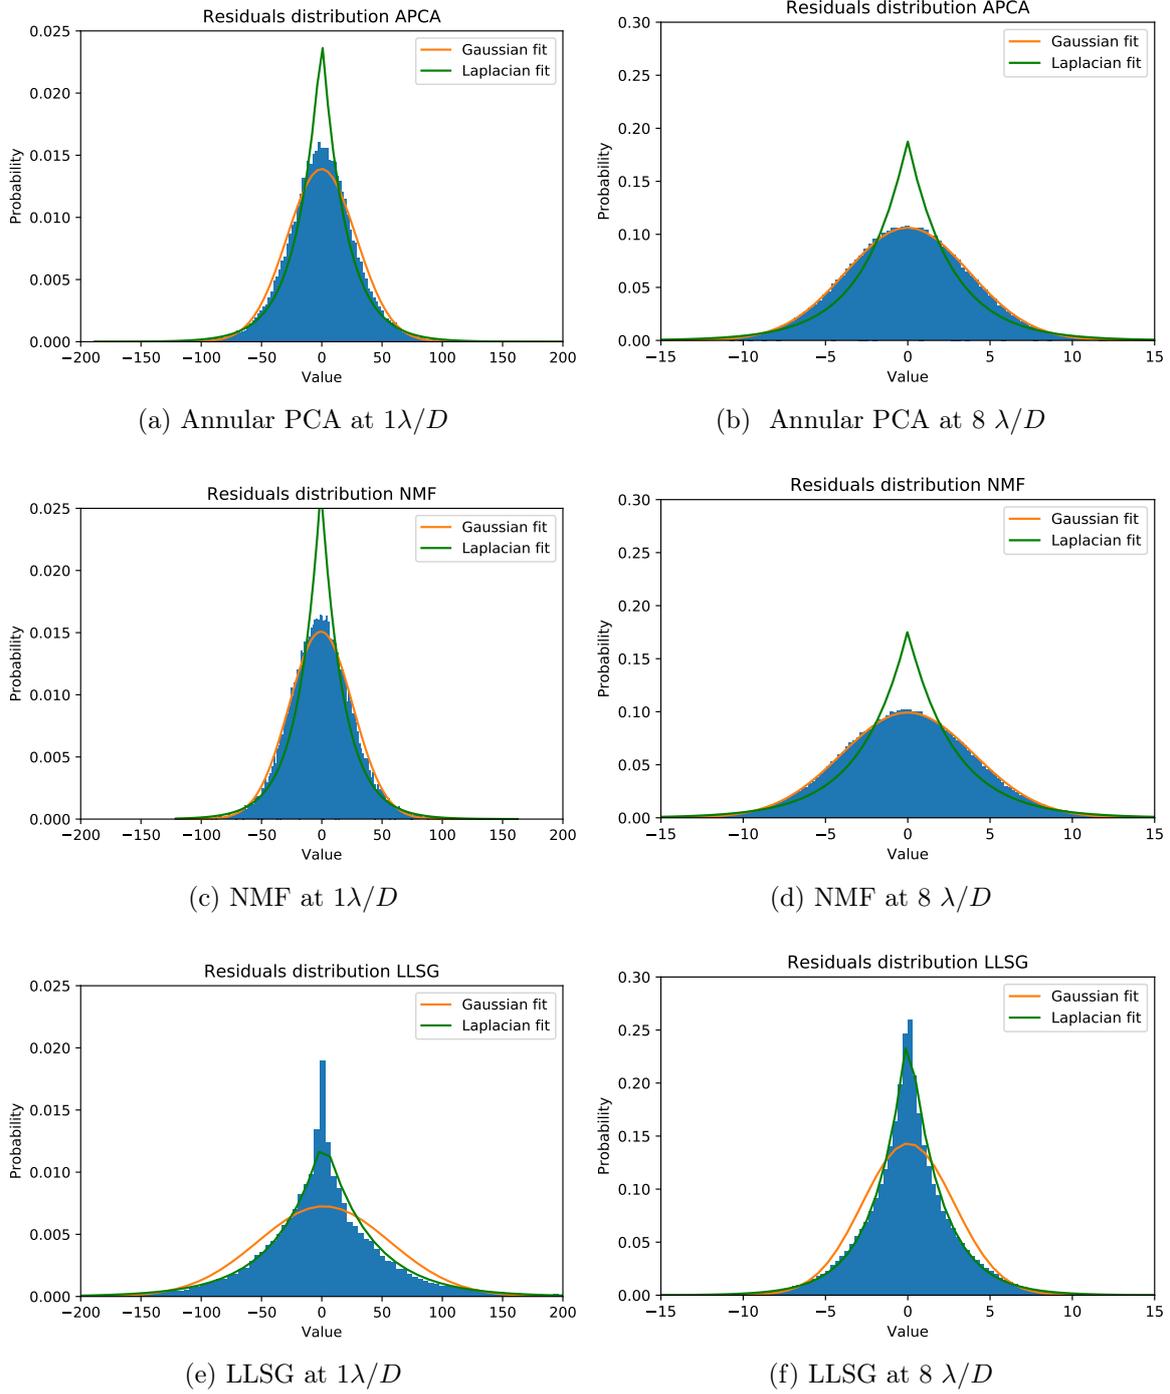


Figure 2.3: Distribution of the residuals for a VLT/NACO data set after PSF subtraction by Annular PCA (top), NMF (middle) and LLSG (bottom) along with a Gaussian (orange line) and Laplacian (green line) fit at small (left) and large separations (right), with respectively 20 components for the Annular PCA and the NMF and a rank of 5 for the LLSG.

higher δ implies that the detected signal is farther in the distribution tails, which indicates a higher level of confidence (which will generally translate as a higher probability in the RSM map) about the detected planet and a higher flux for a given noise distribution. However, the β parameter does not provide an estimation of the planetary flux as we are not using a forward model of the PSF for the planetary signal.

2.3.4 Planetary signal model

Using a forward model for the planetary signal would allow us to take into account the distortions (such as self-subtraction) created by ADI-based post-processing treatment when estimating cubes of residuals. Although a forward-modelled PSF should provide more accurate results, it should be noted that some PSF-subtraction techniques do not lend themselves to the analytical computation of a forward model (eg., LLSG, NMF). A more universal numerical way to compute a forward model is to compare the initial cube of residuals and the one in which a fake companion has been injected. Following this approach, we have tested numerical estimation of forward-modelled PSF for generating RSM detection maps but without managing to improve the algorithm accuracy compared to the use of measured off-axis PSF. We therefore decided to only consider measured off-axis PSFs in the rest of this chapter. However, a forward model variant of the proposed algorithm will be proposed in the next chapter, providing good results at short separations, at the expense of the computation time.

2.3.5 RSM detection map estimation

Now that we have defined the procedure to estimate the cubes of residuals feeding the RSM algorithm as well as the model parameters, we may summarize the main steps of the algorithm as follows:

1. Compute the residuals cubes for the selected ADI techniques and de-rotate all the resulting frames;
2. Define the separation to the star for the first and last annuli, respectively $a_{ini} = \text{FWHM}/2 + 1$ and $a_{fin} = (f_{size} - \text{FWHM})/2$ with f_{size} the size of the frame;
3. Define the series \mathbf{x}_{i_a} for the first annulus;
4. Estimate the mean and variance of the residuals inside the annulus separately for each residuals cube;
5. Using the iterative procedure described in Section 2.2, estimate ξ_{r,i_a} for each index i_a for the set of tested δ ;

6. Include the probability of planetary signal ξ_{1,i_a} providing the maximum likelihood in a three-dimensional matrix $\mathbf{U} \in \mathbb{R}^{L_a \times T}$;
7. Repeat steps 2 to 6 for the next annulus ($a + 1$) until a_{fin} is reached;
8. Average the detection probability contained in \mathbf{U} along the time axis to obtain the final RSM detection map.

The resulting detection map provides the averaged probability of observing a planetary signal in a given cube of data, along with the optimal β . The next section will explore the effectiveness of this new approach when applied to observational data sets.

2.4 Performance assessment

2.4.1 Data

We propose to use two ADI sequences acquired with two instruments of the Very Large Telescope (VLT): NACO and SPHERE. It will allow us to investigate the ability of our model to deal with the different noise profiles produced by these instruments.

The first data set focuses on β Pictoris and its planetary companion β Pictoris b. It was obtained in L' band in January 2013 with NACO in its AGPM coronagraphic mode (Absil et al. 2013). The ADI sequence is composed of 612 individual frames obtained by averaging 40 successive individual exposures, each frame providing an effective integration time of 8 s. The parallactic angle ranges from -15° to $+68^\circ$. We use every third frame to reduce the CPU time and cropped the central 101×101 pixels region to consider mainly the first arc-second.

The second data set is an ADI sequence on Eri51 produced by the SPHERE-IRDIS instrument, using an apodized pupil Lyot coronagraph (Samland et al. 2017). The sequence was taken in $K1$ band in September 2015 and regroups 194 frames with 16 s of integration time. The parallactic angle ranges from 297° to 339° . The data set was pre-processed using the SPHERE Data Center pipeline (for more details about the reduction see Delorme et al. 2017; Maire et al. 2019).

2.4.2 Detection maps

We start our analysis by considering the RSM detection map generated with the proposed algorithm, based on the residual cubes provided by annular PCA, NMF and LLSG, and compare it with the S/N map obtained with the same three post-processing algorithms. The post-processing as well as the S/N detection maps are generated for all three methods with the *VIP* package developed by Gomez Gonzalez et al. (2017) using the standard parametrization. Both Annular PCA and LLSG are performed

annulus-wise, with each annulus being divided into four segments in the case of LLSG. Other parametrizations are possible as the proposed approach works with any de-rotated cube of residuals. The three cubes of residuals obtained with the selected post-processing techniques are then stacked to create a single cube to feed the RSM. The variance and the mean of the residuals are estimated separately for each sub-cube as their noise profiles are specific, as demonstrated in the previous section when looking at the residual distributions.

Figure 2.4 displays the RSM detection map and the S/N maps obtained for the SPHERE-IRDIS Eri51 data set (see Figure A.1 for similar detection maps for the NACO β Pictoris data set). As can be seen, the difference in intensity between the planetary signal and the background speckles is much higher with our new approach than with the usual S/N maps. Eri51 b (contrast of $6.73 \times 10^{-6} \pm 9.02 \times 10^{-7}$ at a separation of 453.4 ± 4.6 mas, Samland et al. 2017; Maire et al. 2019) can be clearly identified on the lower left quadrant with RSM, annular PCA and LLSG, although we observe a higher number of false positives in the case of LLSG. The visual identification becomes more difficult when looking at the S/N map provided by NMF, which shows brighter wind-driven halo residuals.

To illustrate the computation of the RSM map, Figure 2.5 shows how the probability $\xi_{1,ia}$ builds up when getting closer to a planetary signal. It reports the RSM map probabilities along the radial axis crossing the peak value attributed to Eri51 b, along with the optimal δ for the respective annuli. The data includes 7 pixels \times 197 frames \times 3 ADI-based post-processing techniques and is centred on the pixel showing the highest probability. As can be seen, no signal may be found in the first 591 patches representing the first pixel. The probability builds up steadily for the next three pixels until reaching a peak probability of over 95 percent. The value of the optimal δ increases as well with a peak value of 4 reached at the fifth pixel, illustrating the displacement of the signal farther in the residuals distribution tail due to the increasing flux coming from the planetary candidate. We then observe a decrease of the probability and optimal δ , which eventually gets back to the background speckle noise level. The stacked cube of residuals encompasses the cubes of residuals generated first by the annular PCA, then by the NMF and finally by the LLSG. Looking at the sharp increase observed at the beginning of every pixel, we see that the strongest signal may be found in the annular PCA cube of residuals, confirming the visual analysis of the S/N maps. However the signal is still strong in the two other cubes of residuals to be able to maintain the high probability observed for the three central pixels.

2.4.3 Receiver operating characteristic curves

In order to explore in more detail the properties of our new approach and compare its performance with other state-of-the-art methods, we generate synthetic data sets

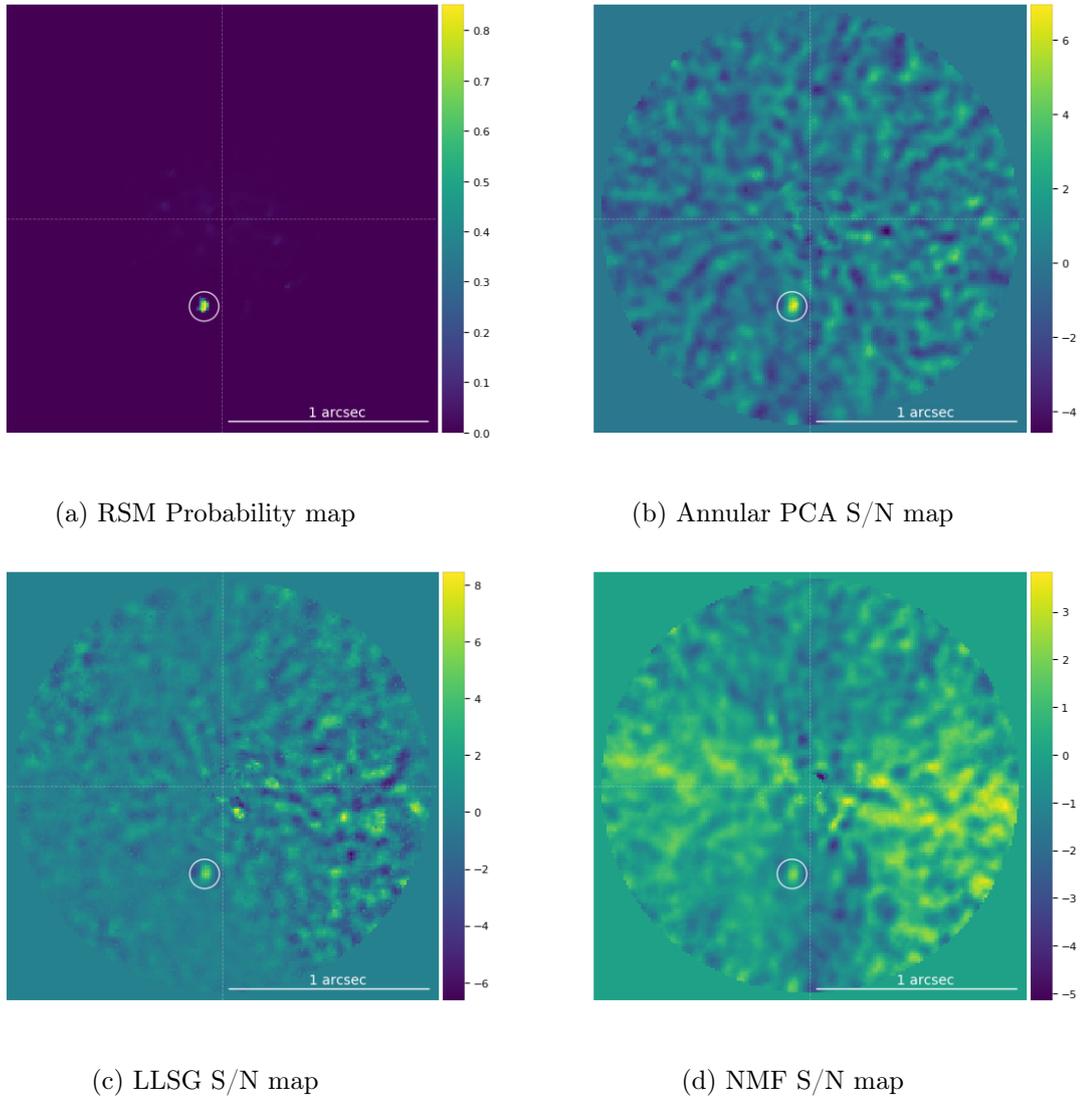


Figure 2.4: Probability map obtained for the SPHERE-IRDIS Eri51 data set, with the RSM using a Gaussian distribution along the S/N map generated with the cube of residuals obtained with Annular PCA, LLSG and NMF with respectively 20 components for the Annular PCA and the NMF and a rank of 7 for the LLSG. The colour scale indicates the probability for the RSM map and the signal-to-noise ratio for the three S/N maps (Mawet et al. 2014). The maps are centred on the star Eri51 while Eri51 b is identified by the white circle in the lower left quadrant.

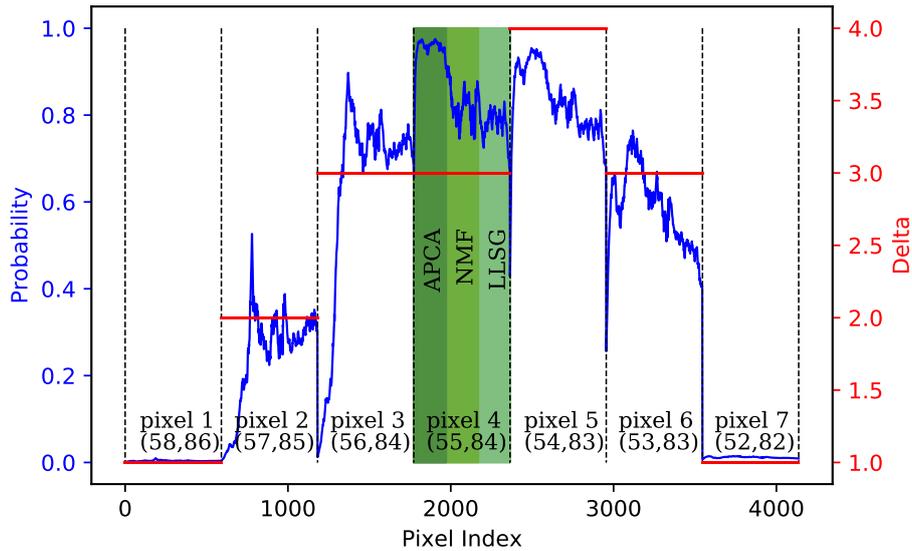


Figure 2.5: Evolution of the probability in the RSM detection map around the location where Eri51 is detected along with the optimal δ for the respective annuli.

based on the two ADI sequences presented in the previous subsection. We rely on the injection of fake companions in the initial ADI sequences, an approach widely accepted by the HCI community for generating synthetic data to assess the sensitivity of post-processing methods. Since the contrast that can be reached as well as the noise structure both depend on the angular separation, we consider three different annuli as described in Table 2.1. The comparison with the other methods is based on Receiver Operating Characteristic (ROC) curves, which are widely used to assess the performance of binary classifiers. In these curves, one axis provides the true positive rate and the other the false positive rate. When using ROC curves for performance assessment, the main proxy for the classifier performance is the area under the ROC curve: the better the classifier, the higher the area under the ROC curve, i.e. the higher the true positive rate for a given false positive rate. We replace the false positive rate by the number of false positive for the entire frame, averaged over the number of test data sets considered for a given separation as it is done in Gomez Gonzalez et al. (2018).

The fake companions are defined as the normalized off-axis PSF, generally measured by offsetting the target star from the coronagraph, multiplied by flux values from a predefined interval defined to challenge the set of tested methods. Five different flux values are tested for each separation with step size of 0.5 times the initial value. For each flux values, eight positions are tested to mitigate the impact of bright speckles or local minima. The resulting 40 test data sets are then used to estimate the ROC curves for each separation. The contrasts for the three selected separations are pro-

Table 2.1: Injected companions contrast range for the three considered separations.

	NACO (β Pic)	SPHERE (51 Eri)
Separation	Contrast	Contrast
$2 \lambda/D$	$3.3\text{-}8.2 \times 10^{-4}$	$1.0\text{-}2.6 \times 10^{-4}$
$4 \lambda/D$	$0.5\text{-}1.3 \times 10^{-4}$	$1.2\text{-}3.1 \times 10^{-5}$
$8 \lambda/D$	$1.3\text{-}3.3 \times 10^{-5}$	$2.1\text{-}5.2 \times 10^{-6}$

vided for the NACO and SPHERE data sets in Table 2.1. Before injecting the fake companion, we removed the known companions and some bright disk structures for the β Pictoris data set, using the negative fake companion technique (Lagrange et al. 2010). We consider as false positive a detected companion at any other location than the one selected for the fake companion injection.

The exoplanet detections for the annular PCA, the NMF and the LLSG method are based on S/N map generated using the procedure of Mawet et al. (2014). The detection of a true or false positive is done on the de-rotated median-combined individual frame by estimating for every pixel a signal-to-noise ratio. This estimation is done annulus-wise in order to take into account the evolution of the residuals distribution. The signal to noise ratio is calculated using the procedure developed by Mawet et al. (2014), by comparing the flux inside an aperture with a diameter of one FWHM centred on the considered pixel (i.e., 5 pixels for both data sets) with the flux of all the other apertures included in the annulus. This procedure implements a small-sample statistics correction, relying on a student t-test to determine the S/N. Increasing S/N or probability thresholds are applied to generate the different ROC curves for all the considered methods. Once the S/N map is computed, successive thresholds are applied onto the S/N map to create the ROC curves. For each threshold, the detection of the fake companion as well as the number of false positives are recorded and averaged over the entire set of synthetic data sets generated for the considered annulus to construct our false and true positive rates. We follow a similar procedure for the RSM detection map, replacing simply the S/N thresholds by successive percentage thresholds applied on the detection map.

The parameters of the different post-processing techniques have been selected to maximize the area under the ROC curves, i.e., maximize the true positive rate while minimizing the number of false positives. For annular PCA and NMF, the number of principal components used to construct the reference PSF have been set to 20 for both data sets. As for LLSG, we selected a rank value of 5 for the estimation of the matrix \mathbf{S} for the β NACO data set and 7 for the SPHERE-IRDIS data set. As regards the RSM, the mean and variance of the residuals distribution are again estimated annulus-wise. The fake companions injected into our simulations having relatively low flux values, we tested δ in the interval $[1, 5]$ and kept the one leading to the highest

total log-likelihood to generate the final RSM map.

As an illustration of the detection maps calculation for the ROC curves generation, Figure 2.6 shows the probability and S/N maps obtained by injecting fake companions with high contrast values at three different separations from the star Eri51 (2, 4 and 8 λ/D). As can be seen, apart from the signal injected at 8 λ/D which appears relatively clearly in the S/N for all three post-processing methods, the RSM map is the only map providing a clear detection for all three fake companions. A set of detection maps is shown in Figure A.2 for the NACO β Pictoris data set, leading to similar conclusions.

Influence of the probability distribution

We now turn to the estimation of the ROC curves which will provide more comprehensive results. We start by considering two different variants of RSM to investigate the choice of the probability distribution for the likelihood function definition. The two variants presented in Figure 2.7, use respectively the Gaussian and Laplacian distribution to construct the likelihood function appearing in η_{r,i_a} . The ROC curves are estimated for different separations as we have seen in the previous section that the probability distribution describing the residuals evolves with angular separation. As can be seen from Figure 2.7, the results of the two variants are very close in the case of the β Pictoris data set, while the distance between them becomes significant for the Eri51 data set. In both cases, the RSM model using the Laplacian distribution performs better for small separation while the Gaussian distribution leads to better results for larger separations.

These results confirm the findings made with Figure 2.3 and the importance of tails fit when selecting the optimal probability distribution. It demonstrates the interest of considering the residuals distribution evolution along the radial axis to optimally parametrize our model. We therefore propose to start the RSM detection map estimation with an analysis of the noise profile to select the right probability distribution for every separation. This additional step has been included in the RSM detection map *python* package that we have developed based on the model presented in this chapter². The function allows (i) to select one of the two distributions, (ii) to automatically select the best distribution based on a best fit approach, or (iii) to create a hybrid distribution consisting in a weighted sum of both distributions. This last possibility can be useful when facing asymmetrical probability distributions as the parameters of both distributions may be estimated separately based on a best fit approach.

²The RSM detection map *python* package is available on GitHub: <https://github.com/chdahlqvist/RSMmap>

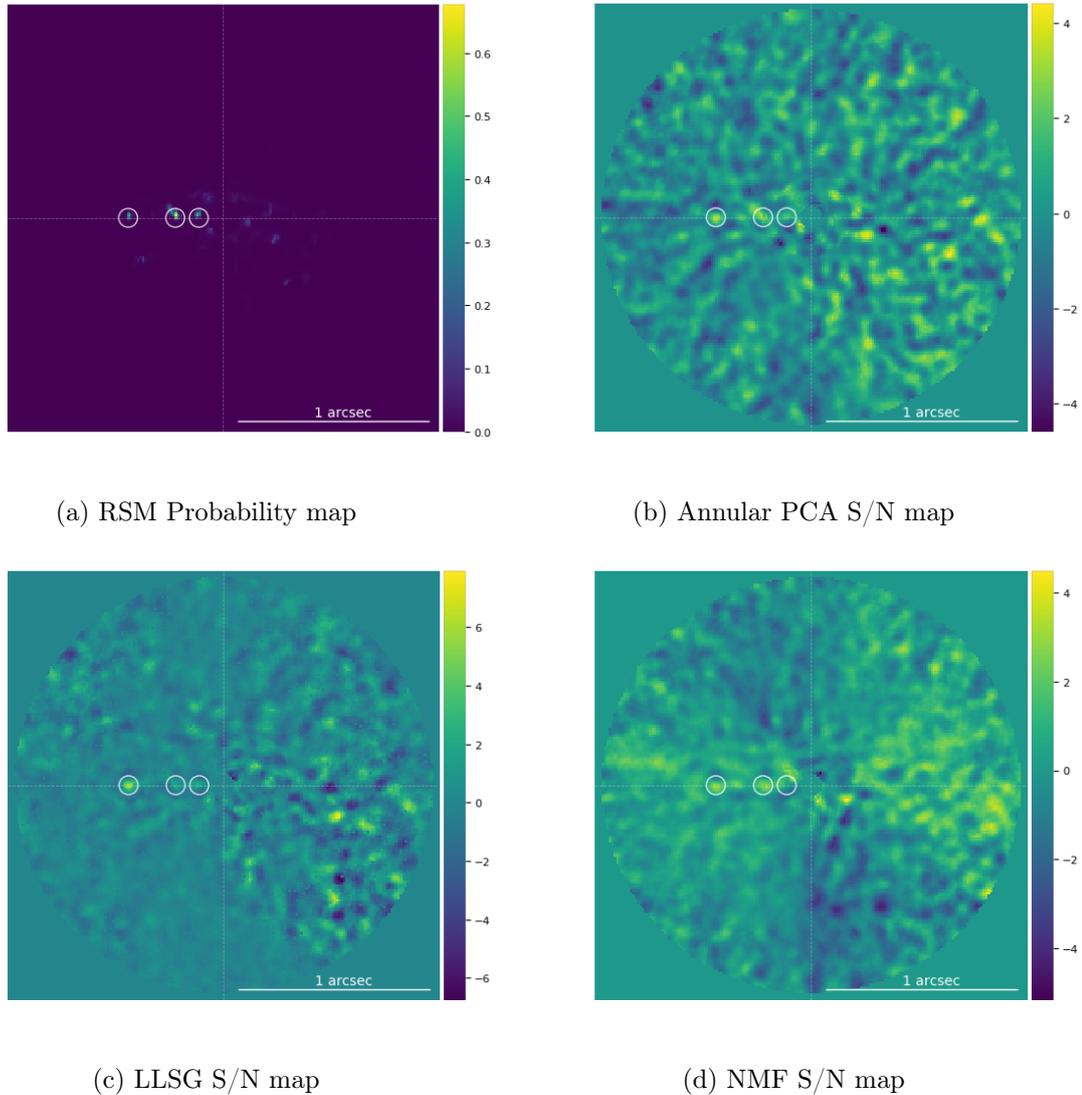


Figure 2.6: Detection map obtained after injecting three fake companions in the SPHERE-IRDIS Eri51 reference cube used for the ROC estimation, at a distance of 2, 4 and 8 λ/D with respectively a contrast of 1.0×10^{-4} , 1.2×10^{-5} and 3.7×10^{-6} . The colour scale indicates the probability for the RSM map and the signal-to-noise ratio for the three S/N maps. The maps are centred of the star Eri51 while the fake companions are identified by the white circles.

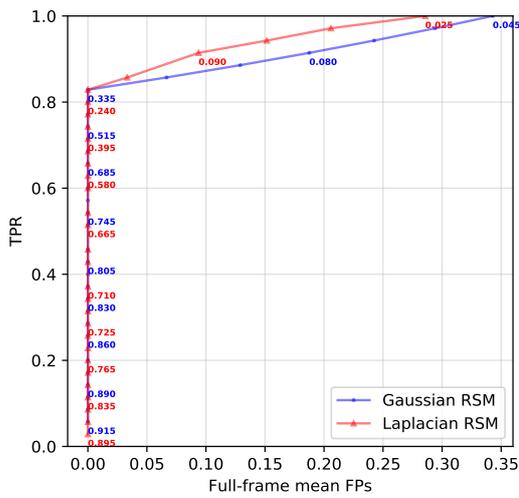
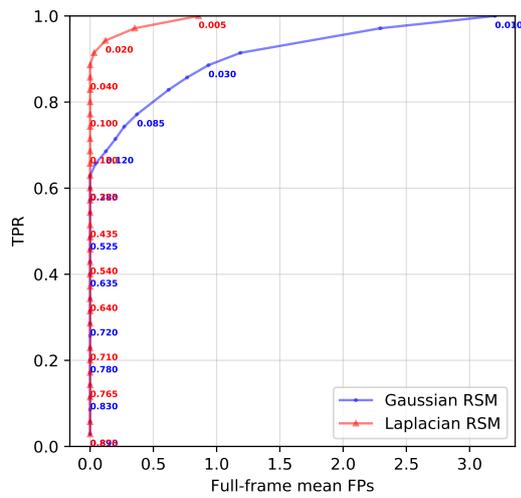
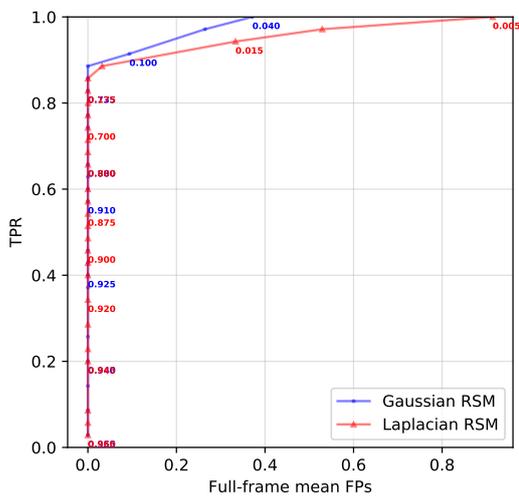
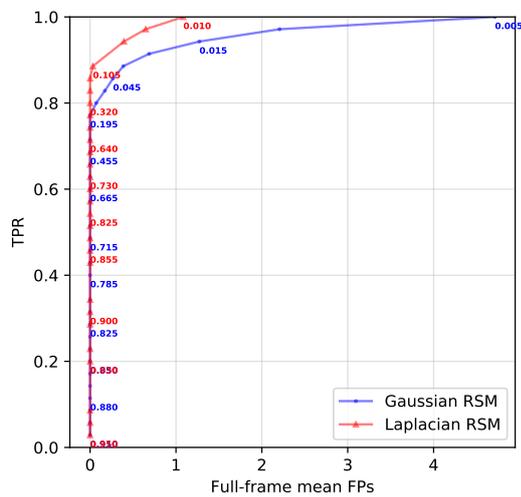
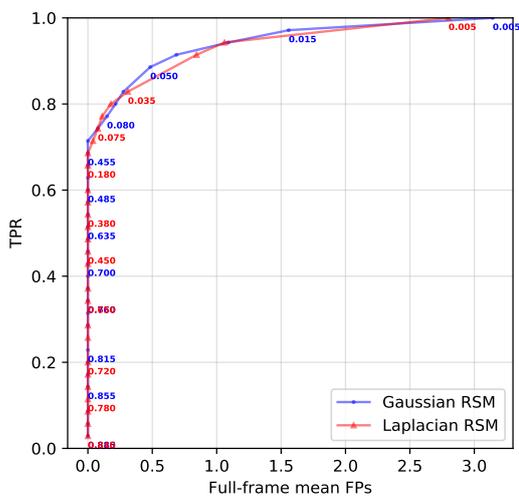
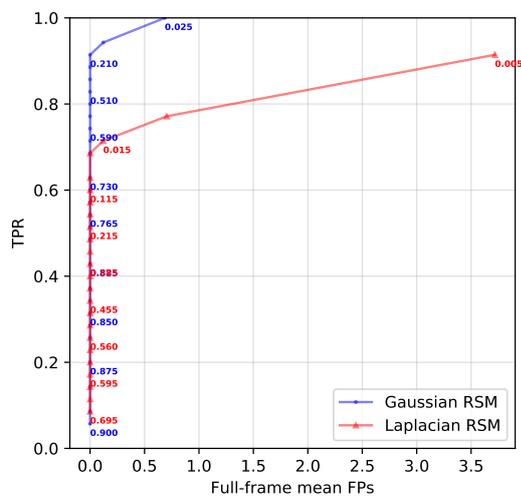
(a) Beta Pictoris at $2 \lambda/D$ (b) Eri51 at $2 \lambda/D$ (c) Beta Pictoris at $4 \lambda/D$ (d) Eri51 at $4 \lambda/D$ (e) Beta Pictoris at $8 \lambda/D$ (f) Eri51 at $8 \lambda/D$

Figure 2.7: ROC curves for the Beta Pictoris and Eri51 data sets, with the Regime Switching Model using respectively a Gaussian (blue) and Laplacian (red) distribution to construct the likelihood function.

Comparison with S/N-based detection

We now address the question of the performance of our algorithm compared to the three post-processing methods using S/N maps. Figure 2.8 reports, for the two data sets, the ROC curves of all four methods for the same separations as before. Considering the results presented in Figure 2.7, we have selected for each data set and each separation the distribution that provided the highest area under the ROC curve. The results demonstrate the interest of the new approach considering that the RSM performs better in every case. This may be explained by the ability of our model to be fed with multiple cubes of residuals, but also by its ability to focus only on relevant data thanks to the regime-switching feature. This allows our model to take advantage of the strength of the different post-processing methods used to produce the cubes of residuals. As speckles are not treated equally by these post-processing techniques, it is easier to remove them by taking into account several cubes of residuals. This ability to remove speckles is further improved by the memory of the RSM. Indeed, the dependence of ξ_{s,i_a} on the transition matrix $p_{q,s}$ and on the probabilities at step $i_a - 1$ (see Eq. 2.3) partly mitigates the effect of speckles on the detection map. Outliers caused by quasi-static speckles do not lead to a clear regime switch while, when facing a planetary signal, the detection probability builds up along the time axis as we have seen in Figure 2.5. The dependence on the past observation reduces significantly the noise in the final detection map.

Furthermore, the possibility of selecting the right probability distribution to describe the residuals allows us to describe more precisely the behaviour of these residual speckles, which is not possible with the S/N approach. The more significant improvements for the Eri51 data set may be explained by the lower level of noise inside this ADI sequence, which suggests that our model should perform better with last generation instruments.

2.5 Conclusion

We explored in this chapter the possibility of improving exoplanet detection using a Regime Switching Model deriving from the field of econometrics, with one regime representing the planetary signal in addition to the speckle noise and the other only the speckle noise. This novel approach allows the creation of probability maps based on cubes of residuals obtained with different PSF-subtraction techniques. The RSM algorithm can be associated with any PSF-subtraction techniques as it can be fed with different cubes of residuals separately or jointly. The short memory process at the heart of our RSM detection map allows quasi-static speckles to be treated more effectively when using several cubes of residuals provided by different post-processing algorithms and thereby to reach better detection performance.

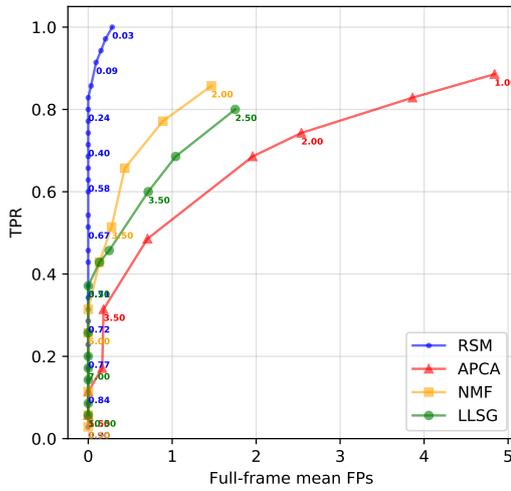
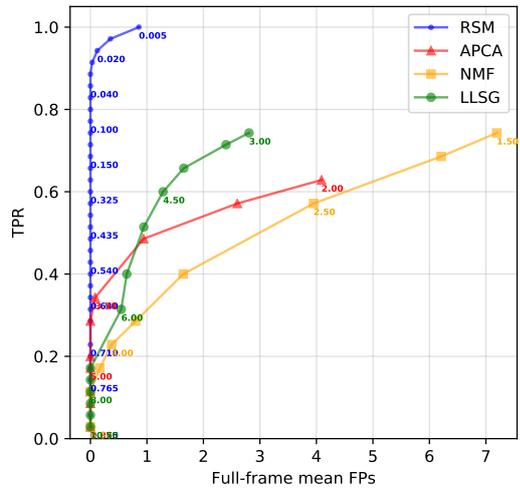
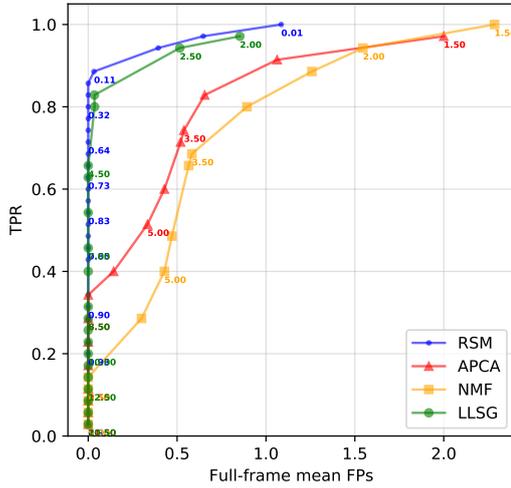
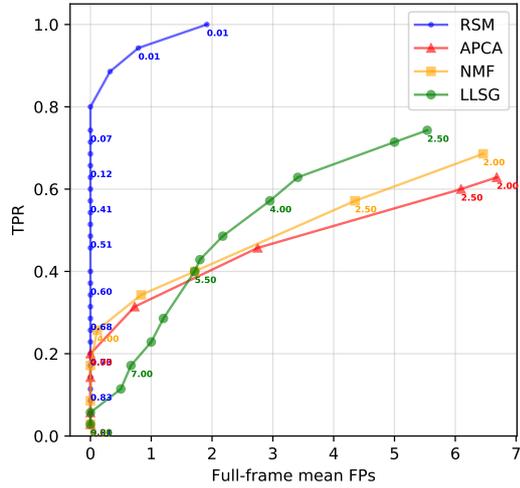
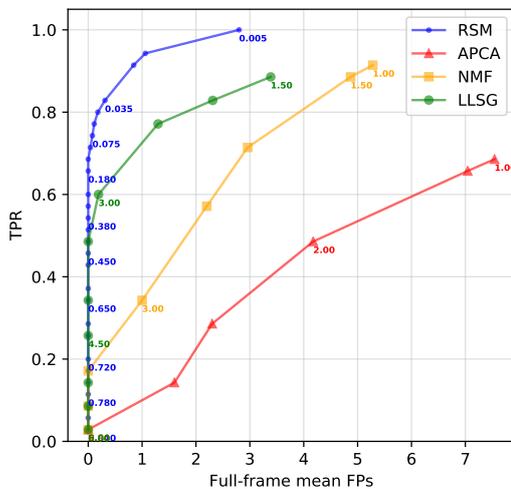
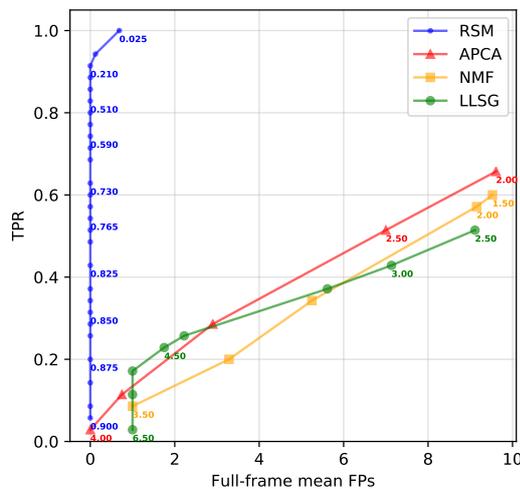
(a) Beta Pictoris at $2 \lambda/D$ (b) Eri51 at $2 \lambda/D$ (c) Beta Pictoris at $4 \lambda/D$ (d) Eri51 at $4 \lambda/D$ (e) Beta Pictoris at $8 \lambda/D$ (f) Eri51 at $8 \lambda/D$

Figure 2.8: ROC curves for the Beta Pictoris and Eri51 data sets, with in blue the Regime Switching Model using the likelihood function that provided the highest area under the ROC curve, in red, yellow and green the detection based on signal to noise ratio using respectively the annular PCA, the NMF and the LLSG as post-processing techniques for the estimation of the cube of residuals.

The RSM is easy to use as most of the parameters are estimated empirically. The only parameter that may need to be tuned is δ , which defines the strength of the signal coming from the planetary candidates. The model selects automatically this parameter via a maximum log-likelihood approach. However an upper value has to be defined for the interval.

We demonstrated the interest of our approach by injecting fake companions into two data sets provided by the VLT/NACO and VLT/SPHERE instruments. We compared the proposed RSM map with standard S/N maps obtained with three state-of-the-art methods, annular PCA, NMF and LLSG. The ROC curves demonstrate clearly the interest of our model as it outperforms all the other methods for the three angular separations we considered, and for both data sets. The results also confirm that the probability distribution of the residuals evolves with radial distance and that it should be taken into account in our model when defining the likelihood function used to estimate the probability of being in one of the two regimes. Indeed, the Laplacian distribution performs clearly better for close separations while the Gaussian one provides better results for larger angular distances. The possibility of optimally selecting the probability distribution based on the residual noise profile has been included in the RSM detection map framework.

Chapter 3

Improving the RSM map algorithm: PSF forward modelling and optimal selection of PSF subtraction techniques

3.1 Introduction

The goal of this third chapter is to further develop the RSM approach by considering a larger set of PSF subtraction techniques that includes the forward modelling of the point source. Indeed, as for ANDROMEDA and KLIP-FMMF, the RSM map relies on a matched filter to infer the existence of planetary candidates in residuals images. However, the initial version of the algorithm, presented in chapter 2, uses solely an off-axis PSF for the detection. Forward modelling could significantly improve the sensitivity of the algorithm to faint companions by taking into account the distortion generated by the speckle field subtraction. We propose a method that would rely on the KLIP forward model (KLIP-FM) developed by Pueyo (2016) as well as a forward model version of LOCI for investigating the added value of forward-modelled point sources. Section 3.2 is devoted to the development of the two forward-model versions of the RSM algorithm, while Section 3.3 provides a performance assessment of these versions.

As seen in chapter 2, the RSM map can accommodate several PSF subtraction techniques to generate a final probability map. That raises questions regarding the selection of the optimal set of techniques to reach the highest sensitivity as well as its dependence on the HCI instrument and on the radial distance. We compare, in Section 3.4, the performance of several set of techniques via ROC curves and investigate the impact of the considered instruments on this selection, considering three state-of-

the-art HCI instruments: NACO, SPHERE, and LMIRCam.

In Section 3.5, we propose an improved method for the probability estimation, relying on a forward-backward approach that allows the use of both past and future observations within the cube of residuals to generate the RSM map. The original RSM map uses a simple forward approach, which considers only past observations to build up the probabilities. We compare the performance of both approaches with standard S/N maps through the use of ROC curves. Finally, in Section 3.6, we present a new framework developed to compute contrast curves, as it is not possible when dealing with probability maps to rely on the standard procedure used for S/N maps.

3.2 Using forward models in RSM

The original RSM map relies on an off-axis PSF to model the planetary signal. A promising development of the current method would be to take into account, via a forward model, the effects of the PSF subtraction techniques on the planetary signal. Indeed, most PSF subtraction techniques lead to distortions of the planetary signal, such as over-subtraction and self-subtraction (Pueyo 2016). Over-subtraction is attributable to quasi-static speckles inside the set of reference frames, while self-subtraction is due to the presence of the planetary signal itself inside the same set of reference frames. The signature of self-subtraction is specific to planetary candidates, as quasi-static speckles coming from the optical train do not rotate with the field. Because of field rotation, the evolution of the reference frames composition leads to the appearance of a negative wing travelling in time from one side of the planet to the other in the azimuthal direction. The temporal motion of this negative wing should therefore help disentangling a planetary candidate from a bright speckle.

We investigate, in this section, two forward model versions of the RSM map relying on the KLIP and LOCI algorithms. Both algorithms can accommodate an analytical estimation of the forward-modelled PSF, which is not the case for other ADI-based techniques such as NMF and LLSG. This avoids the complex task of choosing the fake companion intensity when constructing a forward-modelled PSF empirically by comparing an initial cube of residuals and one in which a fake companion has been injected.

3.2.1 KLIP-based forward modelling

Karhunen-Loève image processing (KLIP) is a popular speckle subtraction technique first proposed by Soummer et al. (2012) and further improved by Pueyo (2016) who developed its forward model version. Similarly to PCA, KLIP estimates the reference PSF via a low-rank approximation of a reference library built to limit the impact of potential planetary signal on the speckle field estimation. For each frame of an ADI

sequence, the KLIP algorithm computes the directions of maximal variance from the reference library. It keeps the principal components up to a rank, K , that is smaller than the dimension of the reference library, discarding the higher order modes that should contain more of the planetary signal. The principal components are found via a decomposition of the covariance matrix of the mean-subtracted reference frames R via the Karhunen-Loève transform. They are given by:

$$\mathbf{Z}_K = \sqrt{\Lambda^{-1}} \mathbf{V}_K \mathbf{R}, \quad (3.1)$$

with $\Lambda = \text{diag}(\mu_1, \mu_2, \dots, \mu_k)^\top$ as the diagonal matrix with the eigenvalues of the image-to-image sample covariance matrix $\mathbf{R}\mathbf{R}^\top$ with \mathbf{R} the mean-subtracted reference library matrix and $\mathbf{V}_K = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k]$ as the respective eigenvectors up to the order $K \leq N_R$, N_R being the number of images in the reference library used to compute the reference PSF (see Table B.1 for a summary of all the variables used in the KLIP-based forward modelling). The reference PSF is then found by projecting the initial science image, \mathbf{i} , or a subsection of this science image, onto the selected principal components. The reference PSF is subtracted from the science image yielding the residual image \mathbf{x} as follows:

$$\mathbf{x} = \mathbf{i} - \mathbf{Z}_K^\top \mathbf{Z}_K \mathbf{i}. \quad (3.2)$$

The selection of the reference library is done via the definition of a minimal field of view (FOV) rotation between the science image, \mathbf{i} , and the set of selected reference images, \mathbf{R} . The minimal FOV rotation should be large enough to limit the distortion due to the planetary signal contained in the library (see Marois et al. 2010) but not too large so that it is possible to keep a sufficient correlation between the speckle field contained in the science image and the reference library. Pueyo (2016) proposed to model the distortion via an analytical expansion of the principal components to account for the presence of planetary signal inside the reference library. In the case of self-subtraction, the planetary signal appears in the principal components estimation via the covariance matrix and therefore the distortion is non linear as a result of the projection of \mathbf{m} onto the perturbed components $\Delta\mathbf{Z}_K$. In contrast, the distortion due to over-subtraction is linear in \mathbf{m} as it is defined as the projection of the planetary signal, \mathbf{m} , on the unperturbed components, \mathbf{Z}_K . The forward model of the planetary signal considers both type of subtractions as follows:

$$\mathbf{p} = \mathbf{m} - \mathbf{Z}_K^\top \mathbf{Z}_K \mathbf{m} - \left(\mathbf{Z}_K^\top \Delta\mathbf{Z}_K + (\mathbf{Z}_K^\top \Delta\mathbf{Z}_K)^\top \right) \frac{\mathbf{i}}{\beta}, \quad (3.3)$$

where \mathbf{m} represents the normalised planetary signal before reference PSF subtraction, typically the instrument off-axis PSF, and \mathbf{p} is the forward model of the planet after

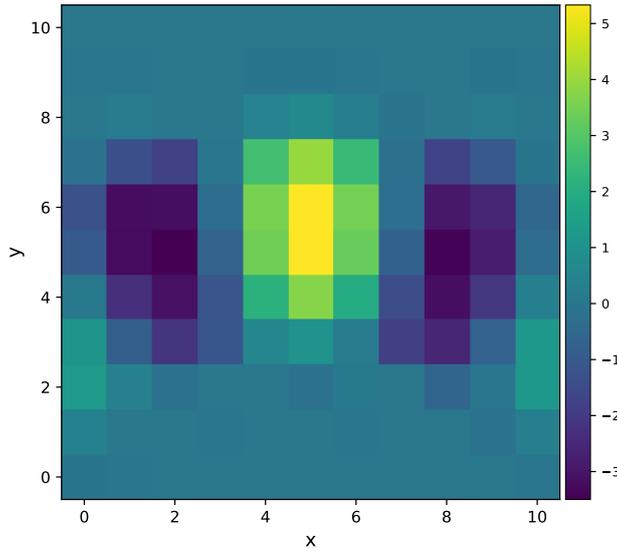


Figure 3.1: KLIP forward-modelled PSF for a NACO β Pictoris ADI sequence taken at a distance of $2\lambda/D$, cropped at two FWHM and summed along the time axis.

subtraction. The second term on the right provides the over-subtraction of the point source while the third term gives the self-subtraction due to rotation via $\Delta\mathbf{Z}_K$ (see Pueyo 2016, for the detailed derivation of this expression).

Having documented the estimation of the forward-modelled planetary signal and of the cube of residuals, we now consider how to include these elements in the RSM map framework. We rely on an annulus-wise estimation but in contrast with annular PCA or LLSG (Gomez Gonzalez et al. 2017), we do not estimate the speckle field for consecutive non-overlapping annuli. We estimate instead a specific speckle field for every radial distance, a . The self-subtraction wings appearing azimuthally, the brightest part of the planetary signal is contained in an annulus segment of width equal to one full width at half maximum (FWHM). The selected annulus with a width of one FWHM is centred on a and shifted by one pixel between each radial distance a instead of being shifted by one FWHM in the case of annular PCA or LLSG. This approach simplifies the forward model PSF estimation, provides a more accurate estimation of the speckle field, and avoids any non-linearities due to transitions between annuli.

The reference PSF and the forward-modelled PSF computation is done via Eq. 3.2 and Eq. 3.3, respectively. The resulting forward-modelled PSF is then derotated and cropped to form a set of patches, \mathbf{p}_{i_a} , where all the elements outside the selected annulus segment is set to zero as can be seen in Figure 3.1. Several crop sizes, from one to two FWHM, are tested in the next section. Once injected in the expression of the likelihood given in Eq. 2.7, this allows us to focus the model only on the region where strong intensity variations occur. The expression of the likelihood of being in

the state, s , for every patch, i_a , becomes (in the Gaussian case):

$$\eta_{s,i_a} = \sum_n \frac{\theta^2}{\theta^2} \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\frac{(\mathbf{x}_{i_a}^n - F_{i_a}\beta\mathbf{p}_{i_a}^n - \mu)^2}{2\sigma^2} \right], \quad (3.4)$$

where \mathbf{x}_{i_a} are the derotated and cropped patches obtained from the residual image, \mathbf{x} .

When relying on a forward-modelled PSF, the intensity parameter β directly provides an estimation of the planet luminosity, which may be helpful for characterizing the planetary candidate beyond its detection. Two methods are considered for estimating this intensity parameter. The first method is similar to the one proposed in the previous chapter, with β defined as a multiple of the estimated standard deviation of the pixel intensity in the annulus, $\beta = \delta\sigma$. In this case, the standard deviation σ is estimated empirically by considering all the frames and an annulus with a width of one FWHM centred on the annulus of interest, a . The δ parameter is defined via the maximisation of the total likelihood of the annulus. The second method relies on the definition of the intensity via a Gaussian maximum likelihood (see Cantalloube et al. 2015; Ruffio et al. 2017, for more details) before the computation of the RSM map itself, which allows the use of the following analytical form for the flux parameter β :

$$\tilde{\beta} = \frac{\sum_j^T \mathbf{i}_j^\top \mathbf{p}_j / \sigma_j}{\sum_j^T \mathbf{p}_j^\top \mathbf{p}_j / \sigma_j}, \quad (3.5)$$

with the standard deviation σ_j computed separately for each frame by considering an annulus with a width equal to one FWHM centred on a .

The main advantage of this second method is the simplicity of the intensity computation. It provides also a specific intensity for each pixel, which may help the algorithm to differentiate bright speckles from planetary candidates. The only drawback of this approach is that it makes the assumption of a Gaussian residuals distribution, which is not always the case especially near the host star. As for the parametrisation of the RSM likelihood function in this second case, the standard deviation used for the flux estimation is taken and the computation of the mean is also done frame-wise using the same procedure. Both approaches are investigated in Section 3.3.

3.2.2 LOCI-based forward modelling

The second PSF subtraction technique that we propose for investigating the RSM framework is a forward model version of LOCI, the locally optimised combination of images (Lafreniere et al. 2007b). Specifically, LOCI relies on a linear combination of reference images to model the speckle field in a given science image, \mathbf{i} . As for the KLIP algorithm, the definition of the reference library is based on the definition of a minimal FOV rotation between the frames composing the reference library and the

selected science image, that is, a minimal distance by which a potential point source in the science image would be displaced in the frames composing the reference library.

Besides this angular distance, LOCI relies on the definition of two different subsections within the ADI sequence. A first section, \mathbf{O}_K , is used for the computation of the linear combination factors, while a second smaller subsection \mathbf{R}_K is selected for the speckle field subtraction¹. The use of a larger section \mathbf{O}_K aims to reduce the weight of the potential planetary candidate in the estimation of the linear combination. Once the reference library is defined, the computation of the linear combination factors is simply done via the minimisation of the sum of squared residuals (Lafreniere et al. 2007b):

$$\epsilon^2 = \sum_{j=0}^{N_p} \left[\mathbf{o}_i^j - \sum_k^{N_R} c_k \mathbf{o}_k^j \right]^2, \quad (3.6)$$

with \mathbf{o}_i the section of the frame for which a model of the speckle field is computed via the factors c_k , and \mathbf{o}_k the section of the reference frame k (see Table B.1 for a summary of all the variables used in the LOCI-based forward modelling). The minimum of this last expression has an analytical form obtained by setting all the partial derivatives with respect to c_k equal to zero, which is equivalent to solving a simple system of linear equations of the form $\mathbf{A}\mathbf{x} = \mathbf{b}$ given by:

$$\sum_k^{N_R} c_k \left[\sum_{j=0}^{N_p} \mathbf{o}_l^j \mathbf{o}_k^j \right] = \left[\sum_{j=0}^{N_p} \mathbf{o}_l^j \mathbf{o}_i^j \right], \quad (3.7)$$

which holds $\forall l \in K$.

Once obtained, the factors, c_k , are multiplied by the subsection of the reference frames, \mathbf{r}_k , and subtracted from the subsection of the science image, \mathbf{i} , to get the residual, \mathbf{x} , as follows:

$$\mathbf{x} = \mathbf{i} - \sum_k^{N_R} c_k \mathbf{r}_k, \quad (3.8)$$

\mathbf{i} defined here as the subsection of the science image corresponding to the reference frames, \mathbf{r}_k .

The forward model of the planetary signal is easily computed using the same factors, c_k , and the planetary signal, \mathbf{m} :

$$\mathbf{p} = \mathbf{m}_i - \sum_k^{N_R} c_k \mathbf{m}_k. \quad (3.9)$$

¹We consider an annulus of three FWHM for \mathbf{O}_K and one FWHM for \mathbf{R}_K .

As in the case of the KLIP forward model, for both the residual image and the forward model PSF, an annulus with a width of one FWHM is used to focus on the region where the planetary signal is the most visible. The forward model PSF, \mathbf{p} , and the residuals images, \mathbf{x} , are again derotated and cropped to form the time series, \mathbf{x}_{i_a} and \mathbf{p}_{i_a} . The two methods used to estimate the flux parameter β are again considered.

3.2.3 Forward model RSM map: summary

We briefly summarise the main steps for the RSM map estimation when relying on the LOCI and KLIP forward-model approach as follows:

1. Compute the residuals for an annulus centred on a for each frame using the KLIP or LOCI procedure.
2. Compute the PSF forward model for every frame and every position within the annulus, a .
3. Derotate the resulting annuli and crop the forward-modelled PSF and science image to form the time series, \mathbf{x}_{i_a} and \mathbf{p}_{i_a} .
4. Estimate the mean and variance of the residuals for every frame, considering the annulus of width equal to one FWHM centred on a .
5. Using the iterative procedure described in Section 2.2, estimate, ξ_{1,i_a} for each index, i_a , using the forward model version of the likelihood (see Eq. 3.5).
6. Repeat steps 1 to 5 for every annulus.
7. Average the resulting probability matrix along the time axis to obtain the final RSM detection map.

3.3 Performance assessment of a forward-modelled RSM map

3.3.1 Data sets

We propose to rely on data sets provided by three different instruments to assess the performance of the two forward model versions of the RSM map. The two first ADI sequences are the same as the ones used in the first chapter and were acquired with two instruments of the Very Large Telescope (VLT), NACO, and SPHERE, while the third sequence was acquired with the LMIRCam instrument of the Large Binocular Telescope (LBT). This choice of data sets aims to investigate the behaviour of the algorithm when facing different noise profiles generated by a variety of instruments.

The first data set is an ADI sequence on β Pictoris and its planetary companion β Pictoris b obtained in L band in January 2013 with NACO in its AGPM coronagraphic mode (Absil et al. 2013). The ADI sequence is composed of 612 individual frames obtained by integrating 40 successive individual exposures of 200 ms. Every third frame was selected here to reduce the computation time, resulting in a final cube of 204 frames. The parallactic angle ranges from -15° to $+68^\circ$.

The second ADI sequence focuses on Eri51. It was obtained in K1 band in September 2015 with the SPHERE-IRDIS instrument, using an apodised pupil Lyot coronagraph (Samland et al. 2017). The data set regroups 194 frames pre-processed using the SPHERE Data Center pipeline (for more details about the reduction see Delorme et al. 2017; Maire et al. 2019). The integration time is 16 s and the parallactic angle ranges from 297° to 339° .

The last data set is an ADI sequence on HD183324 produced by the LMIRCam instrument of the LBT. The images were obtained in October 2018 in L' band without coronagraph using a single telescope. The pre-processed data set² contained 1394 frames with integration time of 109 ms. They were binned over 10 successive individual exposures to reduce the computation time, leading to 139 frames with an integration time of 1.09 s. The parallactic angles range from -13° to -39° . A region with a radius of one arcsecond is considered for all three data sets, which corresponds to around $16 \lambda/D$ for the SPHERE data set and $8 \lambda/D$ for the NACO and LMIRCam data sets.

3.3.2 Results

The performance assessment of the two forward model versions of the RSM map is done via the estimation of ROC curves. In contrast to the ROC curve usually used for assessing the performance of binary classifiers, the false positive rate (FPR) is replaced by the number of false positive (FP) for the entire frame, averaged over the number of test data sets used for the ROC curve computation (see previous chapter and Gomez Gonzalez et al. (2018)). Synthetic data sets are generated based on the three selected ADI sequences by injecting fake companions at two different angular separations to account for the radial evolution of the noise profile. The known companions and some bright disk structures for the β Pictoris data set were removed via the negative fake companion technique (Lagrange et al. 2010) prior to generating the synthetic data sets. The fake companions, which are simply defined as the normalised off-axis PSF, are injected at 16 different position angles with five different flux values for a given angular separation. This allows us to test the sensitivity of the forward model RSM map to different contrasts and mitigates the impact of local speckles on the estimation of the ROC curves. The contrasts used for the three ADI sequences are given in Table 3.1. The relatively low contrasts used for the LMIRCam data set

²Courtesy of Arianna Musso-Barcucci.

Table 3.1: Injected companions contrasts range for the two considered separations and the three ADI sequences.

	NACO	SPHERE	LMIRCam
Separation	Contrast	Contrast	Contrast
$2\lambda/D$	$3.3\text{-}8.2 \times 10^{-4}$	$1.0\text{-}2.6 \times 10^{-4}$	$3.4\text{-}8.6 \times 10^{-3}$
$8\lambda/D$	$1.3\text{-}3.3 \times 10^{-5}$	$2.1\text{-}5.2 \times 10^{-6}$	$3.4\text{-}8.6 \times 10^{-4}$

arise from the short integration time, the low number of frames after binning as well as the small angular rotation, all of which affect the performance of the PSF subtraction techniques. This provides an interesting way of exploring the algorithm performance in different HCI regimes.

We consider a true positive (TP) for a given threshold to be a peak value above the threshold in a circle with a diameter of one FWHM centred on the position of the injected fake companion. A value above the selected threshold at any other location is considered as a FP. In order to avoid double counting, we impose the condition that peak values outside the fake companion region should be separated by a minimal distance of one FWHM.

KLIP-FM RSM map

The forward model version of the KLIP algorithm was developed along a Gaussian matched filter to detect potential planetary candidates in the cube of residuals using the PSF forward model. We propose therefore to compare the performance of the forward model RSM map with the performance of the KLIP forward model matched filter (KLIP-FMMF) developed by Ruffio et al. (2017). We include additionally the original RSM map applied on the cube of residuals generated by the KLIP PSF subtraction techniques and the S/N map obtained with KLIP. In the last case, the S/N map is generated annulus-wise using the procedure of Mawet et al. (2014).

The parameters of the KLIP algorithm, namely, the number of principal components and minimum FOV rotation, were selected to optimise the ROC curves³ for the two considered angular separations. A single set of parameters was defined for each data set. The number of principal components was set to 20 for the SPHERE and NACO data sets while a value of 18 principal components was chosen for the LMIRCam data set. The FOV rotations expressed in terms of FWHM are respectively 0.5, 0.3, and 0.3.

We start by considering for all three data sets the impact of the selected crop size for the forward-modelled PSF used in the KLIP-FM RSM map. As can be seen from Figure 3.2, the larger crop sizes seem to outperform the crop size of one FWHM for

³We mean by optimizing the ROC curve, maximizing the true positive rate (TPR) while minimizing the number of FPs for the set of considered thresholds.

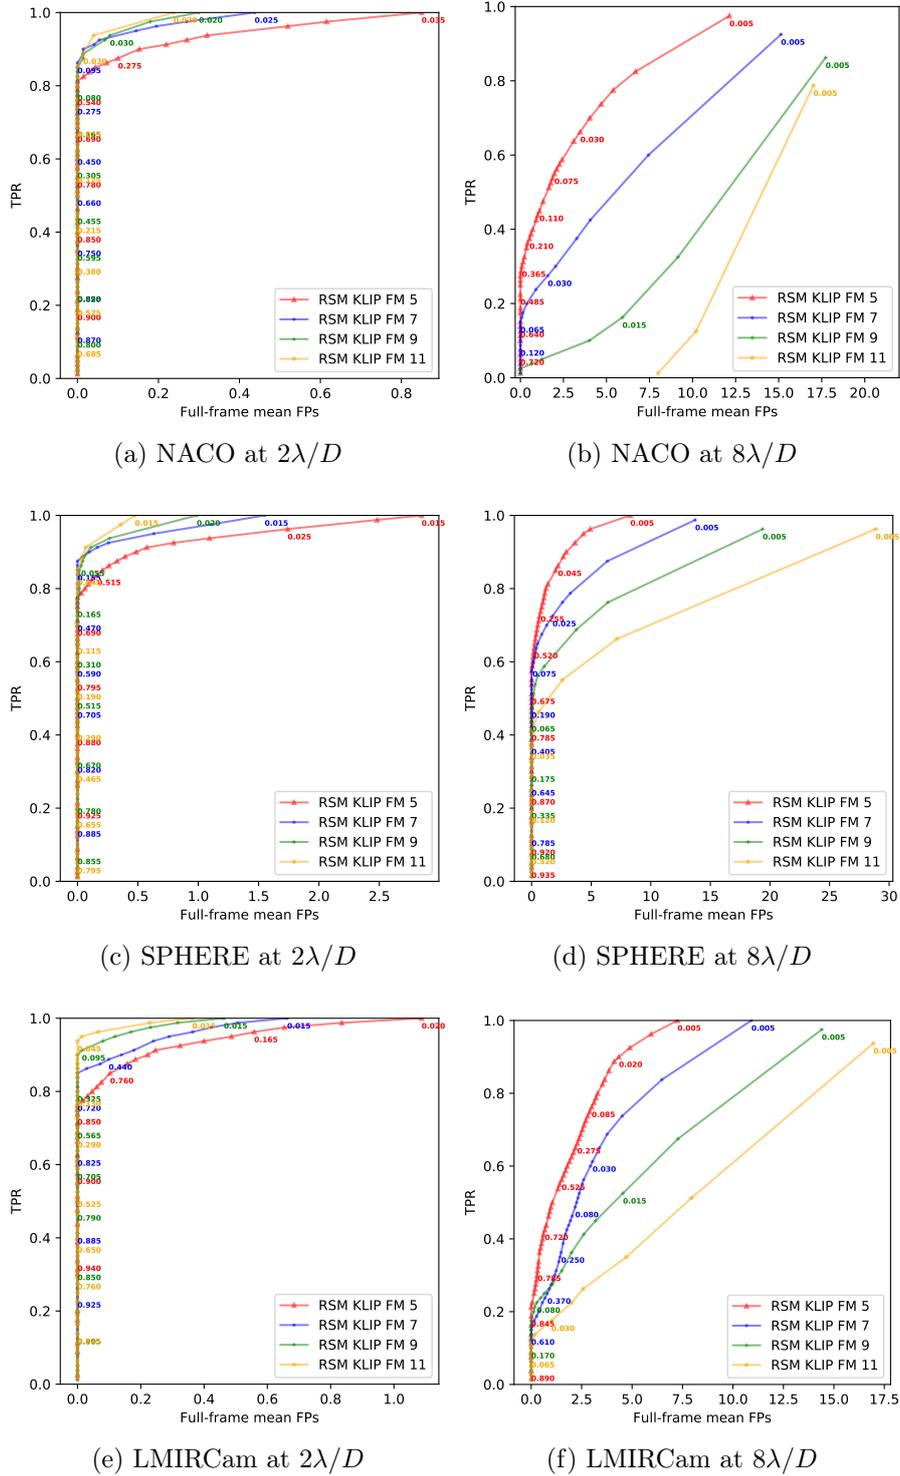


Figure 3.2: ROC curves for the NACO, SPHERE, and LMIRCam data sets, with the KLIP-FM RSM map using respectively a crop size for the forward modelled PSF of 5 (red), 7 (blue), 9 (green), 11 (orange) pixels ($\text{FWHM} \approx 5$ pixels for all three data sets).

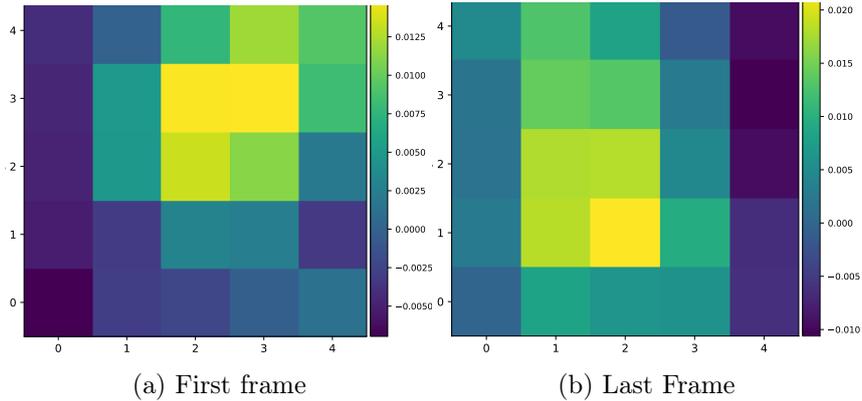


Figure 3.3: KLIP forward-modelled PSF for the NACO β Pictoris ADI sequence taken at a distance of $2\lambda/D$ with the same azimuthal orientation as in Figure 3.1 and cropped at one FWHM. The two images correspond respectively to the first and the last frame.

the small angular separation while the reverse is true for the large angular separation. However, we observe a much larger gap for the largest separation, especially in the case of the NACO data set. This may be explained by the reduced self-subtraction observed at large angular separations, the movement of potential astrophysical signals increasing linearly with the angular separation. This implies that the brightness of the negative lobes appearing on the sides of the main peak reduces with the angular separation. This makes the larger crop sizes unnecessary and more prone to speckle noise.

Considering the previous results, we select for the NACO data set a crop size of five pixels (one FWHM) and a crop size of seven pixels for the other two data sets. This provides a good performance trade-off between small and large angular separations. As can be seen from Figure 3.3, a crop size of one FWHM still captures part of the negative wing azimuthal translation. The results for the two versions of KLIP-FM RSM map, as well as KLIP-FMMF, KLIP RSM map and the KLIP using S/N map are given in Figure 3.4. We see from these plots that, at small separation, the KLIP-FM RSM map seems to slightly outperform KLIP-FMMF, while the reverse is true at large separation. The KLIP approach using S/N map has a higher ability to detect faint companions at large radial distances but it is no match to the other methods at small separations. The KLIP RSM map provides surprisingly good results, being often the closest to the KLIP S/N map for large separations and being relatively close to KLIP-FM RSM and KLIP-FMMF at $2\lambda/D$ from the host star. A combination of both KLIP RSM and KLIP-FM RSM could be interesting to keep the high sensitivity of KLIP-FM RSM at close separations while improving the sensitivity at larger radial distances. It seems also clear from Figure 3.4 that the version of the KLIP-FM RSM map using the Gaussian approximation for estimating the flux parameter β (Eq. 3.5)

outperforms the one relying on the maximum likelihood approach proposed in chapter 2, providing in all cases equivalent or better results. We also tested the maximum likelihood based approach and the Gaussian approximation with the KLIP RSM map with similar results (see Appendix B.2 for a comparison between the two approaches in the case of KLIP RSM), demonstrating the efficiency of this new way of estimating β on top of its faster estimation.

LOCI-FM RSM map

Turning to the RSM LOCI FM map, the tolerance level for the square-residuals minimisation and the minimum FOV rotation were also selected to provide the best overall performance. A tolerance of 9×10^{-3} was chosen for NACO and SPHERE and a tolerance of 1×10^{-2} for LMIRCam. The minimum FOV rotations are respectively, 0.6, 0.2 and 0.2 FWHM. The analysis of the ROC curves obtained with different crop sizes leads to similar conclusions to the case of KLIP-FM RSM. The crop size of one FWHM performs better, in a global sense, even though larger crop sizes lead to slightly better results at small angular separations. The ROC curves corresponding to the crop sizes performance comparison are presented in Appendix B.3. Regarding the performance of LOCI-FM RSM, the results in Figure 3.5 demonstrate again the interest of the Gaussian maximum likelihood to define the flux parameter β mainly for the largest separation. Both LOCI-FM RSM and LOCI RSM outperform clearly the LOCI S/N map for the $2\lambda/D$ angular separation, while the reverse is true for the $8\lambda/D$ angular separation. The ordering is similar to that of the KLIP case with LOCI-FM RSM leading at $2\lambda/D$ and LOCI RSM being closer to the LOCI S/N map at $8\lambda/D$, which again seems to favour a combination of both the LOCI and LOCI-FM to benefit from their respective strength. The search for an optimal mix between the different PSF subtraction techniques is investigated in the next section.

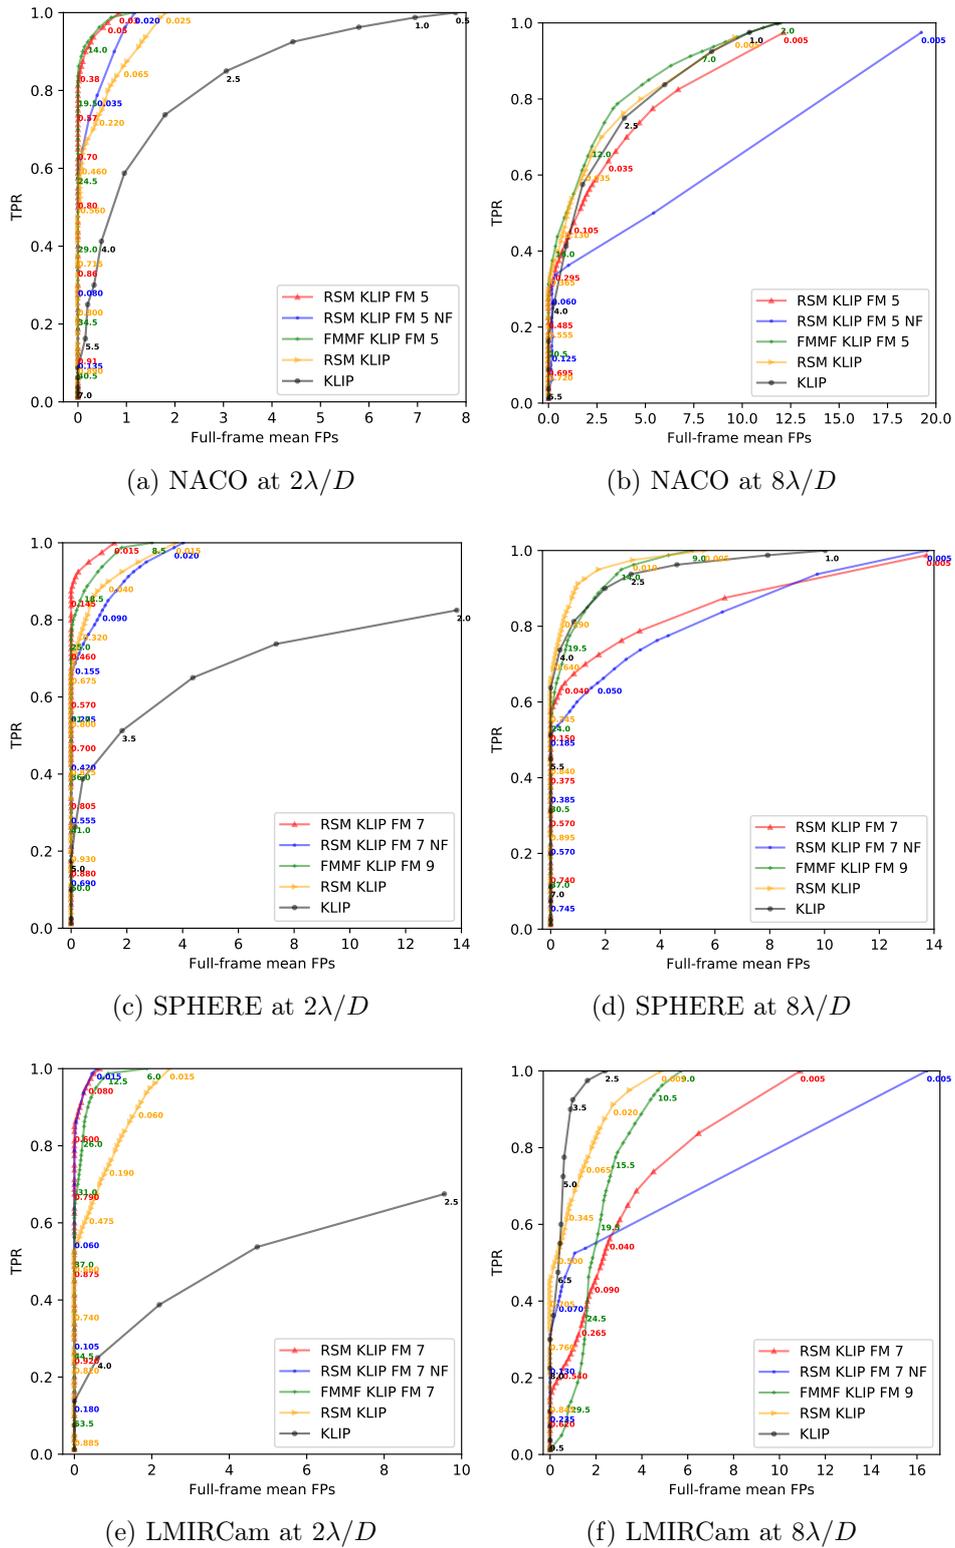


Figure 3.4: ROC curves for the NACO, SPHERE and LMIRCam data sets, with respectively, the KLIP-FM RSM map using the Gaussian maximum likelihood for the pre-optimisation of the flux parameter β (red), the KLIP-FM RSM map with no flux pre-optimisation (NF), which relies on the standard maximum likelihood used in the original RSM map for the estimation of flux parameter β (blue), the forward model matched filter KLIP-FMMF (green), the RSM map using KLIP (orange) and KLIP using the standard S/N map (black).

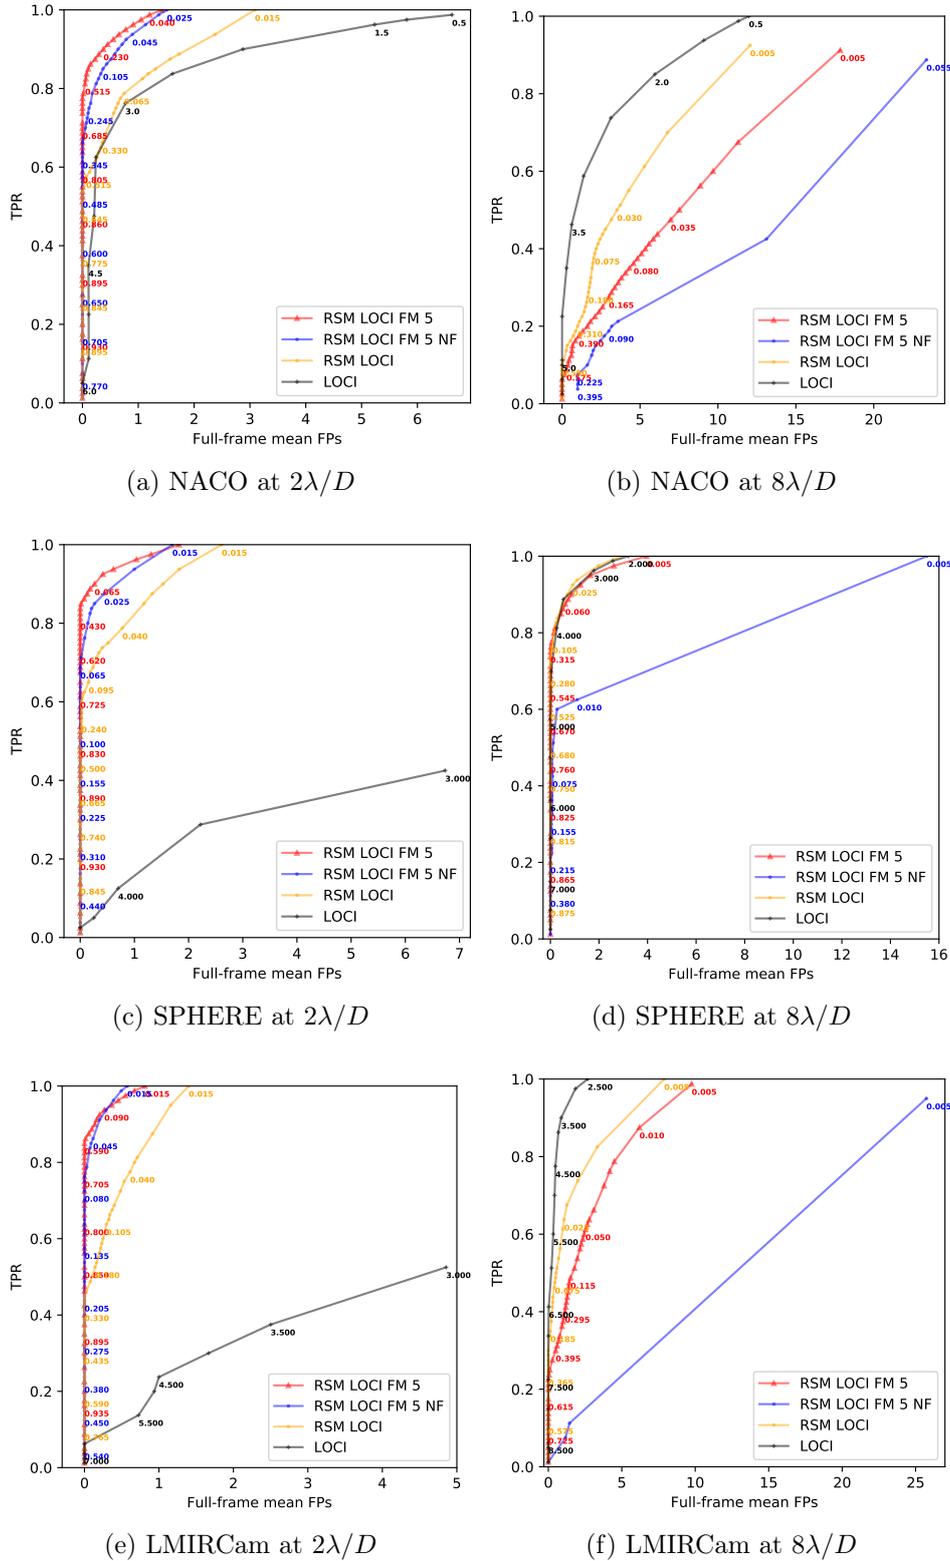


Figure 3.5: ROC curves for the NACO, SPHERE and LMIRCam data sets, with, respectively, the LOCI-FM RSM using the Gaussian maximum likelihood for the pre-optimisation of the flux parameter β (red), the LOCI-FM RSM map with no flux pre-optimisation (NF), which relies on the standard maximum likelihood used in the original RSM map for the estimation of flux parameter β (blue), the RSM map using LOCI (orange) and LOCI using the standard S/N map (black).

3.4 Optimal PSF subtraction techniques selection

Having demonstrated the added value of the forward model versions of the RSM map, at least at small angular separations, we are now left with five different PSF subtraction techniques (annular PCA, KLIP, NMF, LLSG, LOCI) plus two forward model versions to generate the RSM maps. Given that the Annular PCA and KLIP are relatively close in their definition, we decided to focus solely on KLIP, as preliminary results demonstrated their similarities in terms of performance and their non-complementarity. We address in this section the difficult question of optimally selecting these PSF subtraction techniques to optimise the overall performance of the resulting RSM maps. In particular, we investigate the dependence of the optimal combination on the instrument and radial distance. We rely again on ROC curves to assess the performance of the various combinations we considered.

In order to speed up the multiple RSM map estimations, we slightly modified the original RSM map procedure as presented in chapter 2, with, however, no impact on the final outcome of the algorithm. We divided the procedure into two separate steps, the first one being the estimation of the likelihood provided in Eq. 2.7 and the second one the estimation of the probability of being in the planetary regime given by Eq. 2.3. A separate likelihood cube is estimated for every considered PSF subtraction technique for the entire set of annuli. Some of these likelihood cubes are then stacked along the time axis depending on the selected combination. The probabilities are eventually estimated annulus-wise for every pixel of every frame and averaged along the time axis to generate the final probability map⁴. This allows us to estimate only once the likelihood cubes for the different PSF subtraction techniques, with the second step, which is also the fastest, as the only one to be repeated for each combination. The parametrisation of the underlying PSF subtraction techniques were selected to maximise the overall performance for each data set.

⁴This architecture is implemented in the PyRSM python package, which includes all the developments presented here, and is available on GitHub: <https://github.com/chdahlqvist/RSMmap>.

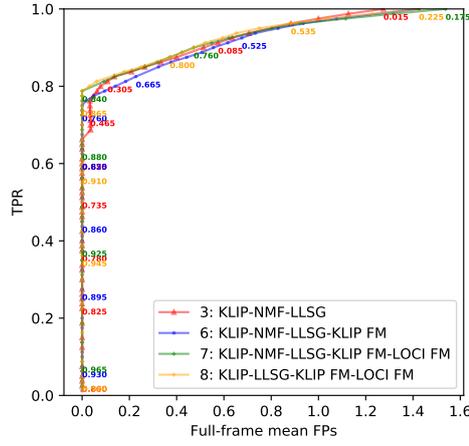
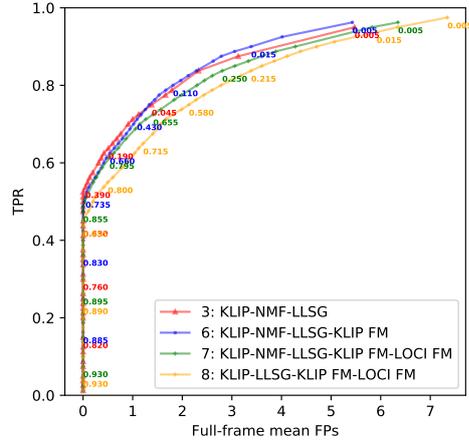
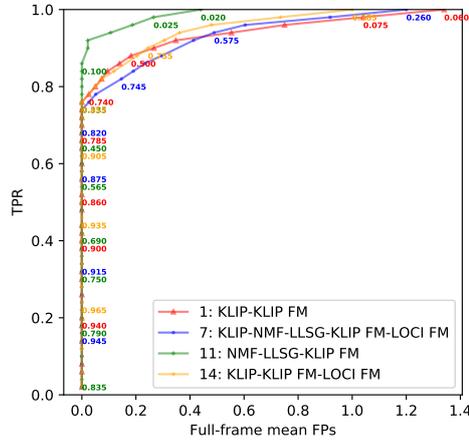
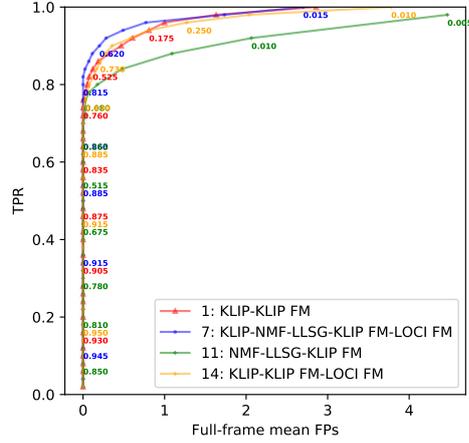
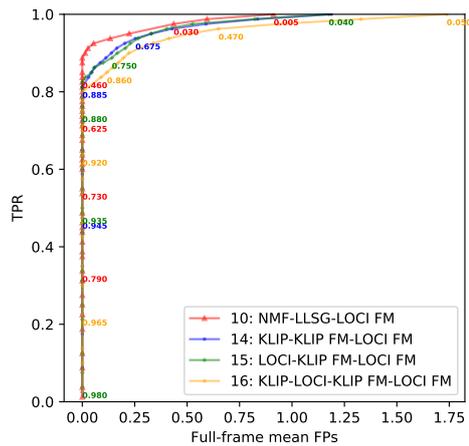
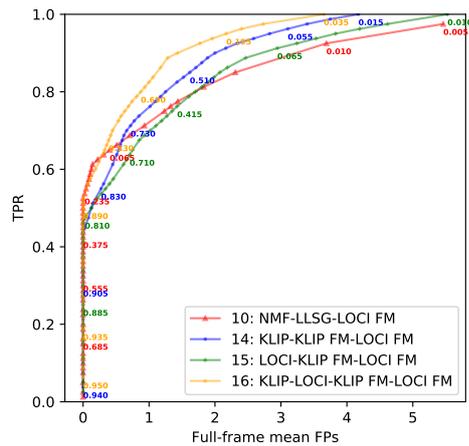
(a) NACO at $2\lambda/D$ (b) NACO at $8\lambda/D$ (c) SPHERE at $2\lambda/D$ (d) SPHERE at $8\lambda/D$ (e) LMIRCam at $2\lambda/D$ (f) LMIRCam at $8\lambda/D$

Figure 3.6: ROC curves for the NACO, SPHERE and LMIRCam data sets, with the four best combinations of PSF subtraction techniques used to generate the RSM map algorithm.

The ROC curve computation follows the same procedure as in the previous section with a region of one arcsecond considered for all three data sets and the curves being computed for the two same angular separations. The main characteristics of the ROC curves for the 16 selected combinations of PSF subtraction techniques may be found in Appendix B.4. The two parameters we used to select the best combinations, are the maximum TPR reached without any FP, and the average number of FPs inside the entire frame at TPR=1⁵. The first parameter is the most important one, as it gives clues about the highest contrast the algorithm can reach without any false detection. The second gives a measure of the number of bright background structures that have not been properly treated by both the PSF subtraction techniques and the RSM algorithm.

The results presented in appendix B.4⁶ show large differences in terms of performance between the considered combinations highlighting the importance of the PSF subtraction techniques selection. Based on the two metrics introduced in the previous paragraph, we selected the four best combinations for each data set. The ROC curves for the two considered angular separations are presented in Figure 3.6. When comparing these curves with the ones in Figures 3.4 and 3.5, we see that the improvement of the RSM map performance occurs mainly at larger separations when considering multiple PSF subtraction techniques. The ROC curves are indeed very close to the ones obtained with the KLIP-FM RSM and LOCI-FM RSM for the $2\lambda/D$ radial distance, while the gap is much wider for the $8\lambda/D$. Apart from the combinations 7 and 14 (see Figure B.3), which are selected for multiple data sets, the other combinations are specific to each data set. This seems to demonstrate that the selection of an optimal combination should be done at least on an instrument-specific basis. The definition of a single optimal combination for the entire set of annuli seems also difficult, as we often observe that higher performance at short separations goes hand in hand with lower performance at large separations. A last element to consider for the selection of the optimal combination is the threshold value for which the first false positive is observed, which should be as small as possible since large values imply the presence of bright structures in the probability map. We conclude from Figure 3.6 that in a global sense, the best combinations are the combinations 3, 11, and 10 for the NACO, SPHERE, and LMIRCam data sets, respectively. Looking at these three combinations, we see that they share a common structure, being composed of the LLSG and NMF PSF subtraction techniques, in addition to a LOCI- or KLIP-based PSF subtraction technique. The performance of this particular combination is probably due to

⁵The average number of FPs at TPR=1 is estimated by taking the highest threshold corresponding to a TPR of 1, or if a TPR of 1 cannot be reached, the smallest probability threshold we considered in our study, i.e., 0.5%.

⁶The ROC curves summary for the 16 selected combinations is presented in appendix B.4. via a table and a set of figures presenting the same results. Figure B.3 provides a graphical comparison between the combinations via bar charts while Table B.2 provides the detailed results.

the differences between these PSF subtraction techniques in terms of residuals noise profile. These differences should help to better average out the speckle noise via the RSM algorithm. This structure therefore appears to be an interesting starting point when studying a new data set. We further characterise these three combinations in Section 3.6 by estimating their contrast curve.

The results presented in this section demonstrate the dependence of the optimal set of PSF subtraction techniques on the instrument providing the ADI sequence, but also on the angular separation, although a common underlying structure could be seen. A larger set of ADI sequences would be needed to determine whether a single optimal set of PSF subtraction techniques may be identified for a given instrument or even multiple instruments, which could be very helpful when dealing with large surveys. We investigate this question in the next chapter, as well as the development of an automated procedure for the selection of the optimal set of PSF subtraction techniques to be considered when generating the final RSM map.

3.5 Forward-backward model

In this section, we discuss an additional improvement of the original RSM map by considering a forward-backward approach for the estimation of the probability, ξ_{1,i_a} . The current approach relies solely on past observations to construct the cube of probabilities while the entire cube of residuals is available for the estimation, that is, of both past and future observations. We propose therefore to replace the current forward approach by a forward-backward approach, which considers both past and future observations. This method computes two separate sets of probabilities, the forward probabilities as done in the original RSM framework:

$$\xi_{1,i_a}^f = \sum_{q=0}^1 \frac{\eta_{1,i_a} p_{q,1} \xi_{q,i_a-1}^f}{\sum_{q=0}^1 \sum_{s=0}^1 \eta_{s,i_a} p_{q,s} \xi_{q,i_a-1}^f}, \quad (3.10)$$

but also the backward probabilities, which rely on the probability estimated at index $i_a + 1$ instead of index $i_a - 1$ to compute the probability at current index i_a as:

$$\xi_{1,i_a}^b = \sum_{q=0}^1 \frac{\eta_{1,i_a} p_{q,1} \xi_{q,i_a+1}^b}{\sum_{q=0}^1 \sum_{s=0}^1 \eta_{s,i_a} p_{q,s} \xi_{q,i_a+1}^b}. \quad (3.11)$$

Once both sets have been estimated, the final probabilities are obtained by multiplying the two sequences of probabilities. A normalisation factor is applied, making sure that the total probability equals 1 for every index i_a . The final probabilities are

therefore given by:

$$\xi_{1,i_a} = \frac{\xi_{1,i_a}^f \xi_{1,i_a}^b}{\sum_{s=0}^1 \xi_{s,i_a}^f \xi_{s,i_a}^b} \quad (3.12)$$

Because the RSM features a short-term memory, the probability of being in the planetary regime builds up when we get closer to the planetary signal but with a small latency. As can be seen from Figure 3.7, this latency leads to a shift of the main peak towards the future for the forward approach and towards the past for the backward approach. When relying on the forward-backward approach these shifts cancel out and the main peak is centred on the true position of the planetary signal. The forward-backward approach should therefore allow to reach a higher precision in terms of exoplanet astrometry.

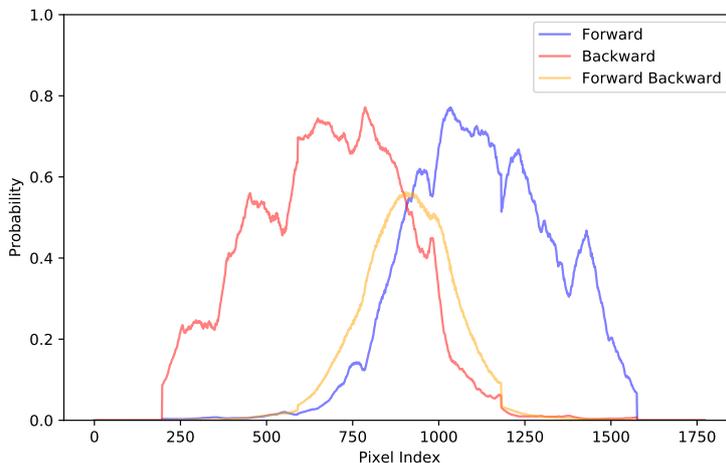


Figure 3.7: Evolution of the probabilities for the forward, backward, and forward-backward approaches using KLIP, around the location of a planetary candidate injected in the Eri51 data set (radial distance of $4\lambda/D$ with a contrast of 3.76×10^{-5}).

In order to investigate the ability of both approaches to derive an accurate astrometric measurement for the detected planetary signal, we propose performing a series of simulations based on the SPHERE data set. We study the evolution of the astrometric precision for a range of contrasts, considering again a radial distance of $2\lambda/D$ and $8\lambda/D$. As done in the previous sections, we base our simulation on synthetic data sets, on the basis of which we apply a KLIP RSM map using the forward and forward-backward version of the RSM algorithm. The negative fake companion (NEGFC) method (Lagrange et al. 2010; Marois et al. 2010; Wertz, O. et al. 2017) is also applied on the synthetic data sets, allowing for a comparison with a technique dedicated to astrometry⁷. For each radial distance, we inject fake companions at 16

⁷We relied on the function provided by the VIP package (Gomez Gonzalez et al. 2017) for the computation of the position via the NEGFC using a simplex (Nelder-Mead) optimisation.

different position angles. The set of considered contrasts are computed based on the KLIP RSM contrast curve, estimated using the approach proposed in the next section (see Figure B.4). We define two sets of contrasts ranging from one to six times the sensitivity limit at the considered radial distance in Figure B.4, with a step size of 0.5. In the case of the RSM map, the computation of the position is done by fitting a two-dimensional Gaussian to the detected planetary signal. The astrometric error bars for the three considered methods are computed as the root mean squared (rms) position error between the obtained position and the injected fake companion true position, averaged over the two axes. The rms is estimated over the 16 fake companions injected at each radial distance, for every contrast.

The results from Figure 3.8 demonstrate clearly the ability of the forward-backward approach to decrease the position error compared to the original forward approach. As can be seen from Figure 3.8, the RSM forward-backward approach performs better than the NEGFC approach at large radial distances and for high contrasts. However, for lower contrast, the RSM forward-backward approach reaches a noise floor around 4 mas, higher than the noise floor of the NEGFC approach, which is between 1.5 and 2 mas. This higher noise floor may be explained partly by the profile of the planetary signal in the RSM map. As can be seen from Eq. 2.3, the RSM approach response to a planetary signal is non linear and dependent on neighbouring pixels, leading potentially to asymmetries in the azimuthal direction, even in the forward-backward case. The algorithm architecture also leads to non-linearities along the radial axis because of the annulus-wise probabilities computation. Finally, as can be seen from Figure 3.7, the forward-backward approach reduces the amplitude of the planetary signal within the probability map. All these elements affect the Gaussian fit and therefore the astrometric precision that the RSM algorithm can reach. Nevertheless, as demonstrated by the results from Figure 3.8, the RSM forward-backward approach can reach a higher astrometric precision, especially at large radial distances and high contrasts. This is due to the better ability of the RSM algorithm to detect faint companions. It is also worth noting that the computation time is much lower when using the RSM map than with the NEGFC approach or the more advanced Markov Chain Monte Carlo version of NEGFC approach. The Gaussian fit of the RSM forward-backward signal provides therefore a good first estimate, especially for high contrast targets, which can then be used, for lower contrasts, as an initial position for more advanced astrometry techniques. A more advanced planet characterisation algorithm based on the RSM framework and using the NEGFC approach is developed in chapter 5. This new planet characterisation algorithm provides a more accurate estimation of the planet photometry and astrometry independently for the contrast (see 5.4.1).

Another advantage of the forward-backward approach lies in its ability to reduce the background speckle noise and smooth the probability curve, the noise being treated differently by the forward and backward components. Looking at Figure 3.9, we see

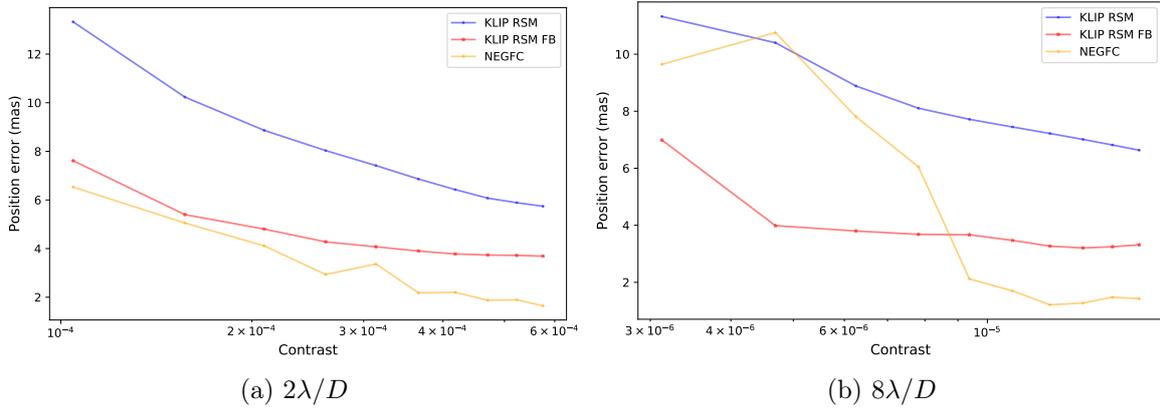


Figure 3.8: RMS position errors averaged over the two axes expressed in mas for the KLIP RSM map using the forward (blue) and forward-backward (red) versions of the RSM approach and for the NEGFC approach (orange). The two graphs show the dependence of the averaged RMS position error on the contrast for respectively a radial distance of $2\lambda/D$ and $8\lambda/D$.

that the level of the residual speckle noise has reduced drastically for the LMIRCam data set, the brightest speckle probability decreasing by around 40 percent. However this noise reduction comes along with a reduced brightness of the planetary signal. This reduced brightness is also illustrated in Figure 3.7, where the peak value obtained with the forward-backward approach is lower than the one obtained with the two other approaches, leading to a reduced detection threshold. This reduction of the planet signal strength does not impact the performance of the forward-backward approach in terms of ROC curve, with similar results for both the forward and forward-backward approaches. The forward-backward approach outperforms even slightly the forward approach at small separation (see Appendix B.2 for a comparison between the original RSM approach and the forward-backward version in the case of KLIP).

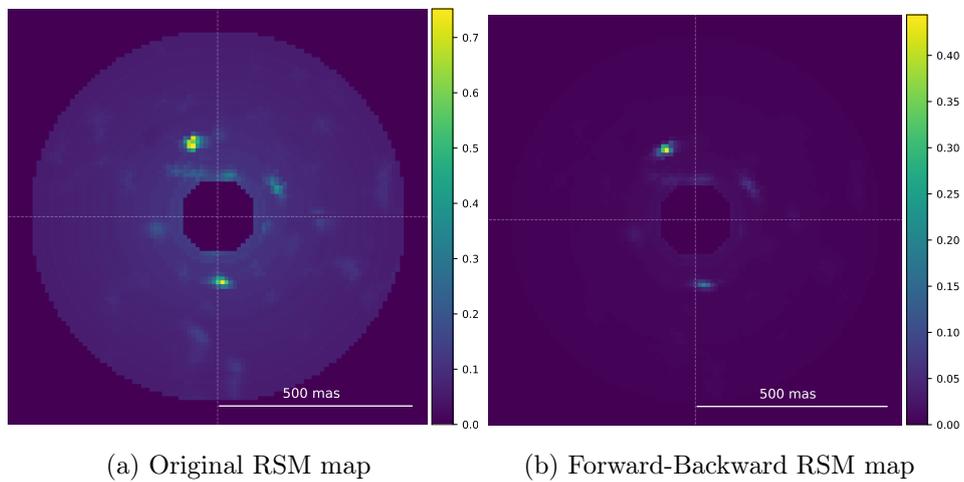


Figure 3.9: RSM map generated with forward and forward-backward approach for the LMIRCam data set using KLIP with 18 principal components, a FOV rotation expressed in terms of FWHM of 0.3 and a Gaussian distribution. A square root-based scale has been used to increase the background speckle noise brightness.

3.6 Contrast curve

This section is devoted to the estimation of contrast curves based on our RSM framework. When relying on probability map for exoplanet detection, we cannot use the traditional procedure to compute the contrast curve. This procedure selects the contrast corresponding to a TPR of 0.5 for a given probability of observing a FP (Jensen-Clem et al. 2017). A 5σ threshold is usually chosen, which corresponds to a 3×10^{-7} false alarm probability under a Gaussian noise hypothesis. In the case of the RSM map, the probability provided as an output by the algorithm is a non linear function of the underlying likelihoods, past observations, and the transfer matrix (see Eq. 2.3), which precludes us from defining a similar probability threshold for the FP. We therefore rely on a procedure similar to the one we used for the ROC curve computation. Considering a given flux, we inject fake companions at different position angles for a given radial distance. The resulting synthetic data sets allow for the computation of a threshold-dependent TPR. The next step is to select a convenient threshold. As it is not possible to reach a 5σ confidence empirically, we select the threshold that leads to the first detection of a FP within the entire frame (as done in Jensen-Clem et al. 2017, but considering the entire frame instead of the selected annulus). This threshold definition provides a direct link with the background noise level while avoiding the usual shortcomings of standard contrast curves, namely, the noise Gaussianity assumption and the definition of the region used to empirically estimate the first two moments of the noise.

We repeat the steps presented in the previous paragraphs on sets of fake companions injected with different flux values until a TPR of 0.5 is found when detecting the first FP. We use an iterative procedure based on linear interpolations to minimise the number of flux values to be considered before reaching the TPR of 0.5. The procedure starts with the estimation of the TPR for a selected pair of upper and lower flux values (step 1 in Figure 3.10, respectively below and above the TPR of 0.5 preferably). A linear interpolation then allows the determination of a best guess for the next flux value to be estimated (step 2). Once the TPR is estimated for this new flux value, a new linear interpolation is performed with this last flux value as upper bound (resp. lower bound) if the TPR is above 0.5 (resp. below 0.5) keeping the previous lower flux value (resp. upper flux value) (step 3). The procedure is repeated until a TPR $\in [0.45, 0.55]$ is found using a tolerance interval with a size of 0.1. These steps are summarised in Figure 3.10.

We estimated contrast curves for the three data sets, considering the original RSM map using the optimal combination found in Section 3.6, the forward-backward version of the algorithm, and the simple KLIP S/N map. We computed the contrast from a radial distance of $2\lambda/D$ up to $8\lambda/D$ using a step-size of one λ/D and then up to $16\lambda/D$ for the SPHERE data set with a step-size of two λ/D . The resulting contrast curves

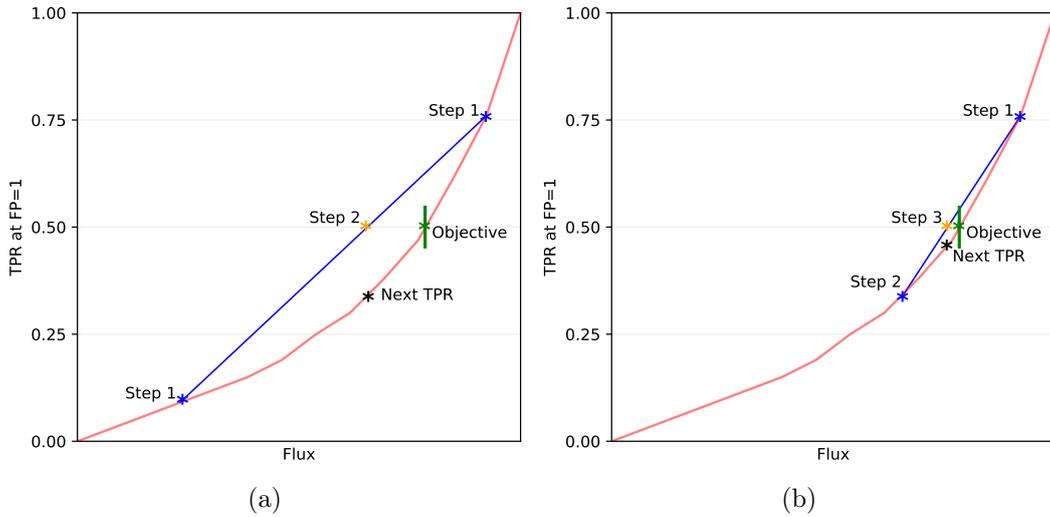
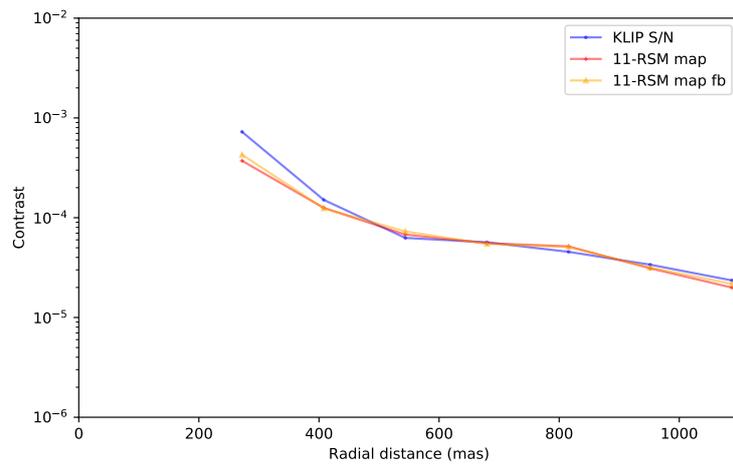
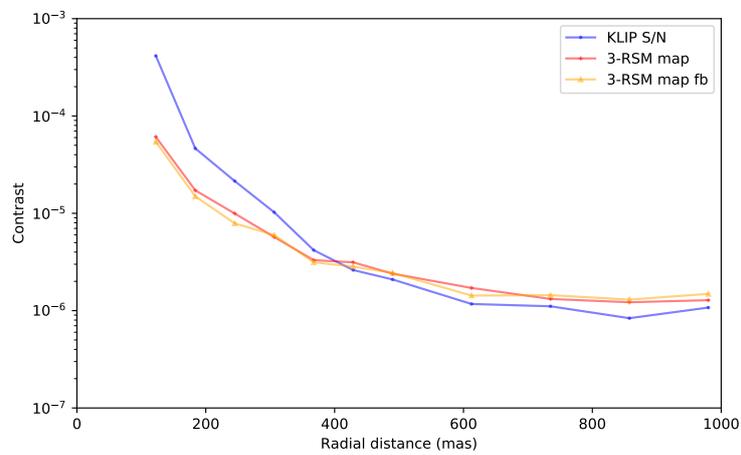


Figure 3.10: Procedure for the estimation of the contrast curve, the orange star corresponds to the linear interpolation, the black star to the corresponding TPR, the blue stars to the two previous estimations used for the linear interpolation and the green star to the flux corresponding to a TPR of 0.5. The red curve makes the link between the flux and the true positive rate in the case of a single false positive while the green line shows the tolerance's interval of $[0.45, 0.55]$.

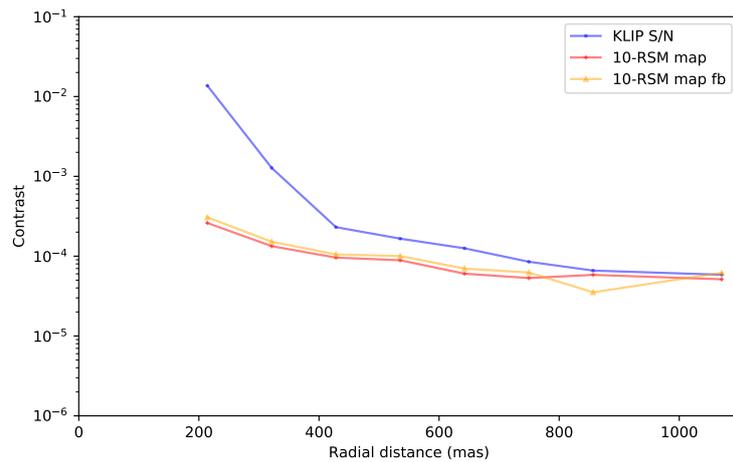
may be found in Figure 3.11. These contrast curves confirm the findings made with the ROC curves, the RSM map reaching in most cases a significantly higher contrast for small angular separations and providing comparable contrast at larger separations. The gap between the contrast curves of the KLIP S/N map and the RSM map is significantly larger for the LMIRCam data set, highlighting the interest of the RSM approach for this instrument. However, the lower number of frames and the smaller angular rotation for this particular data set may explain these results. A larger number of LMIRCam data sets is therefore needed to confirm these findings. We also find that the RSM map and the forward-backward version have a similar performance for all angular separations. Considering its higher precision and its lower background noise level, the forward-backward version seems to be a promising alternative to the standard RSM map.



(a) NACO



(b) SPHERE



(c) LMIRCam

Figure 3.11: Contrast curves for the three data sets using the optimal combination of PSF subtraction techniques obtained in Section 3.4 (resp. 3, 11, and 10) with the region $[2\lambda/D, 16\lambda/D]$ considered for the SPHERE data set and the region $[2\lambda/D, 8\lambda/D]$ considered for the other two data sets to get the contrast in the first arcsecond.

3.7 Conclusion

In order to improve the RSM map sensitivity to faint exoplanets, we implemented several improvements at different levels. We first considered a forward-model version of the RSM map for two different PSF subtraction techniques, KLIP and LOCI. The computation of forward-modelled PSFs allowed to take into account PSF distortions due to the speckle field subtraction. We demonstrated the interest of the forward model at small angular separations through the estimation of ROC curves. These ROC curves highlight the radial dependence of the optimal PSF crop size, with larger crop sizes leading to better results for small angular separations while the reverse is true for smaller crop sizes. This could be explained by the fact that larger PSF crop sizes better account for the self-subtraction patterns, whose intensity decreases with radial distance as the relative movement of astrophysical signals increases with the distance to the host star.

In a second step, we investigated the question of the optimal selection of the PSF subtraction techniques to be included in the estimation of the RSM map. Relying again on ROC curves, we demonstrated the importance of this selection with sometimes large differences between the performance of the considered combinations. The optimal combinations seems to depend on the instrument, but also on the angular separation.

We proposed then an improvement directly related to the way probabilities are estimated within the RSM algorithm by replacing the original forward approach by a forward-backward approach. This approach relies on both past and future observations to compute a final probability. This allows us to deal more efficiently with background noise as speckles are not treated in the same way when relying on a forward and backward approach. Another advantage of the forward-backward version of the RSM map is its ability to extract more precisely the planetary astrometry.

We finally implemented a new framework to compute contrast curves in the case of probability maps, which cannot be treated as S/N maps. We kept the TPR of 0.5 while removing the Gaussian based noise threshold definition, replacing it with a threshold based on the detection of the first FP. Using a procedure based on successive linear interpolations, we were able to compute the contrast curve for the original RSM map and the forward-backward version using the optimal combination obtained previously and compare their performance with the one of a simple KLIP S/N map. The results highlighted the ability of the RSM map to detect fainter companions at small angular separations (below 400 mas) and the overall similar results for larger separations. The contrast curves also confirmed the interest of the forward-backward approach as it provides similar contrast curves while reducing the speckle background noise and giving a higher precision in terms of astrometry.

Chapter 4

Auto-RSM: an automated parameter selection algorithm for the RSM map exoplanet detection algorithm

4.1 Introduction

The use of one or several PSF-subtraction techniques to generate a S/N map, or a probability map via the RSM algorithm, requires the definition of multiple parameters, specific to each method and varying potentially from one ADI sequence to another. The selected set of parameters can have dramatic effects on the final detection map, both in terms of noise and algorithm throughput. The selection of the optimal set of parameters is usually done manually, which requires time and can lead to bias, as the definition of the optimality of the set is driven by the ability of the user to analyse properly the generated detection maps. The complexity of the optimal parameter selection can be an obstacle to the use of some HCI data-processing techniques, and can lead to unreliable results. This makes it also difficult to compare properly the performance of HCI data-processing techniques, as their performance is parameter-driven to a large extent, and therefore depends on subjective choices made by the user. To mitigate these issues, we propose in the context of the RSM framework an optimisation procedure, called auto-RSM, to select automatically the best set of parameters for the PSF-subtraction techniques, as well as for the RSM algorithm itself. To our knowledge, such an extensive optimisation procedure has not yet been proposed in the HCI literature, although some earlier works have already addressed partly the question of the parameter optimisation. Approaches such as the S/N based optimisation of the number of components for PCA (Gomez Gonzalez et al. 2017), or the direct optimization of the nonlinear S/N function (Thompson & Marois 2021), focus on a single PSF-subtraction technique, whereas here we proposed a more generic framework applicable to most PSF-subtraction techniques. The proposed optimisation framework

can be divided into three main steps: (i) the selection of the optimal set of parameters for the different PSF-subtraction techniques (and ADI sequences) via Bayesian optimisation, (ii) the optimal parametrisation of the RSM algorithm, and (iii) the selection of the optimal set of PSF-subtraction techniques (and ADI sequences) to be considered for the computation of the final RSM probability map. This last step is motivated by the fact that, when relying on multiple PSF-subtraction techniques and multiple ADI sequences, some sequences may be noisier while some methods may better cope with the noise, independently from their parametrisation. Special attention should therefore be paid to the choice of the cubes of residuals generating the final probability map. The optimisation step for the PSF-subtraction technique parameters is not limited to auto-RSM, and can be performed separately if a S/N map is preferred to the RSM probability map. Besides the development of the auto-RSM, we therefore propose a variant of the algorithm adapted to the use of S/N maps instead of RSM probability maps. The auto-S/N relies on the first step of the auto-RSM algorithm, and on a modified selection framework allowing the optimal combination of multiple S/N maps. The performance of both optimisation procedures is assessed using the *exoplanet imaging data challenge* (EIDC Cantalloube et al. 2020b), which regroups ADI sequences generated by three state-of-the-art HCI instruments: SPHERE, NIRC2 and LMIRCam. The EIDC initiative aims to provide the tools, i.e., the data sets and performance metrics, to compare properly the various HCI data processing techniques that have been developed recently. The EIDC first phase, which ended in October 2020, regrouped and compared the results from 23 different submissions for the ADI sub-challenge (Cantalloube et al. 2020b), making it a great tool to assess the performance of a new approach. The results obtained with the different versions of the proposed optimisation procedure demonstrate its interest when compared to the results presented in Cantalloube et al. (2020b). The rest of the chapter is organised as follows. Section 4.2 describes the procedure used to optimise the parameters of the PSF-subtraction techniques. In Section 4.3, we introduce the RSM map framework and present the next two steps of the auto-RSM framework: the optimal RSM parameter selection, and the selection of the optimal set of cubes of residuals used to generate the final probability map. Section 4.4 is devoted to the performance assessment of the optimisation procedure along with the comparison of different versions of the optimisation procedure. Finally, Section 4.5 concludes on this work.

4.2 PSF-subtraction techniques optimisation

The proposed optimisation procedure relies on the concept of inverted parallactic angles, which has already been used in the HCI literature (e.g. Gomez Gonzalez et al. 2018; Pairet et al. 2019). The sign of the parallactic angles used to de-rotate the ADI

Table 4.1: Set of parameters selected for the optimisation of the six considered PSF-subtraction techniques.

Parameters	APCA	NMF	LLSG	LOCI	KLIP-FM	LOCI-FM
Number of principal components	X	X	X		X	
FOV minimal rotation	X			X	X	X
Number of azimuthal segments	X		X			
Tolerance for error minimisation				X		X

sequence is flipped, blurring any planetary signal while preserving the noise temporal correlation and statistical properties. The use of ADI sequences with flipped parallactic angles should allow us to avoid biases due to the contribution from potential planetary signals during the optimisation process. Although the inverted parallactic angles approach allows us to blur planetary signals, it is not immune to potential bright artefacts, which implies that particular attention needs to be paid to the elimination of outliers from the computed optimal parameters.

Following the development of chapter 3, the RSM map algorithm can accommodate up to six different PSF-subtraction techniques: LOCI, annular PCA, KLIP, NMF, LLSG, and forward-model versions of KLIP and LOCI. Each PSF-subtraction technique is characterised by its own set of parameters, which strongly affect the quality of the reference PSF modelling. Table 4.1 presents the parameters that we have identified as the most relevant for the optimisation of the six considered PSF-subtraction techniques. Other parameters, such as the annulus width, were tested during the auto-RSM development, but were discarded from the optimisation framework as their influence was found to be smaller or because of other practical considerations.

4.2.1 Definition of the loss function

Parameter optimisation requires the definition of a loss function f , which provides, for a given set of parameters \mathbf{p} , an outcome $f(\mathbf{p})$ that can be maximised or minimised. In the case of reference PSF modelling, the loss function should quantify the ability of the PSF-subtraction technique to remove the residual noise contained in the ADI sequence and to identify potential planetary companions. The definition of the achievable planet/star flux ratio or contrast, for a given detection significance, is therefore a good candidate to measure the PSF-subtraction technique performance. In the context of HCI, the contrast is defined as follows (Jensen-Clem et al. 2017):

$$\text{contrast} = \left(\frac{\text{factor} \times \text{noise}}{\text{stellar aperture photometry}} \right) \left(\frac{1}{\text{throughput}} \right). \quad (4.1)$$

The contrast is usually defined at a 5σ level which implies $\text{factor} = 5$ with $\text{noise} = \sigma$. As the parameter optimisation is done for a single ADI sequence at a time, the

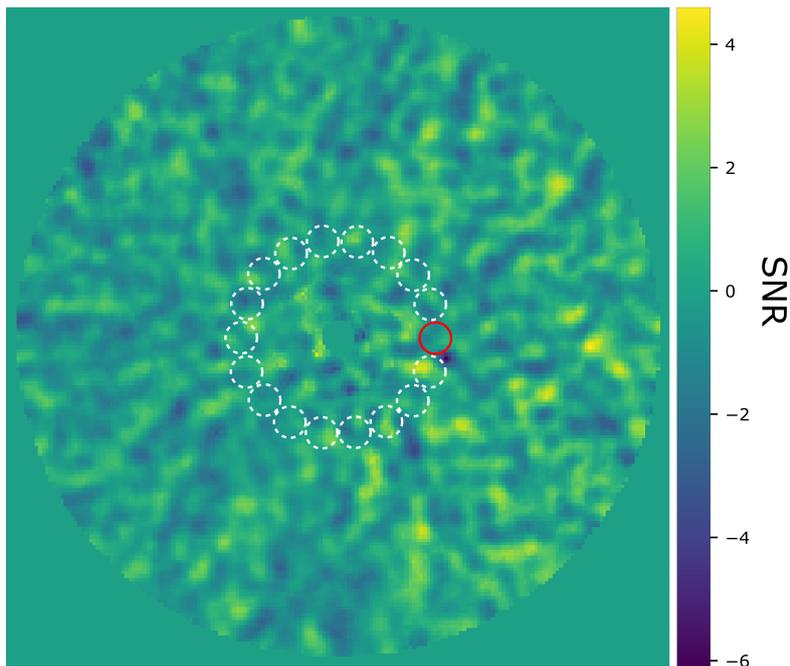


Figure 4.1: Estimation of the noise via the annulus-wise procedure proposed by Mawet et al. (2014). The dotted white circles indicate the apertures whose flux is used for the noise computation, while the red circular region is centred on the pixel for which the noise needs to be estimated.

stellar aperture photometry does not impact the optimisation process and is therefore irrelevant and set to 1. We rely on the procedure of Mawet et al. (2014) illustrated in Figure 4.1 to determine the noise annulus-wise. For a given annulus and for an aperture centred on the pixel of interest, with a diameter equal to the full width at half maximum (FWHM) of the PSF, the noise is computed by considering the standard deviation of the fluxes in all the non-overlapping apertures (one FWHM in diameter each) included in this annulus. The number of apertures being relatively small for small angular separations, the procedure implements a small statistics correction, relying on a Student t-test to correct the multiplicative factor for the noise.

The throughput quantifies the attenuation of the planetary signal due to reference PSF subtraction. In practice, the throughput is estimated by injecting a fake companion at a predefined position and computing the ratio between the injected aperture flux and the recovered aperture flux after the reference PSF subtraction. Contrast curves can be computed by averaging the sensitivity limit in terms of planet-to-star contrast obtained by injecting several fake companions at different position angles for a series of angular separations. Relying on several azimuthal positions and averaging the associated contrasts reduces the impact of the residual speckles on the estimated contrast. We follow this approach, but instead of injecting individual fake companions separately to compute the average contrast, we inject several fake companions at once,

which drastically reduces the computation time. We impose a minimum separation of one FWHM between the apertures containing the fake companions, and a maximum of eight fake companions per annulus, in order to limit potential cross-talk between the injected fake companions. This safety distance, as well as the small intensity of the injected fake companions¹, provides a good approximation of the average contrast (see Appendix C.3 for a comparison between sequential and multiple injections) while limiting the computation time, which is crucial here as parameter optimisation requires a large number of contrast estimations. The loss function computation may be summarised as follows:

1. Reference PSF estimation using the selected post-processing technique and set of parameters;
2. reference PSF subtraction from the original ADI sequence, de-rotation of the cube of residuals, and median combination of the resulting frames;
3. computation of the fluxes for the entire set of apertures within the selected annulus in the median-combined frame obtained in step 2, and estimation of the noise relying on a Student t-test;
4. injection of fake companions at the selected set of azimuths with flux value defined as five times the noise computed in step 3;
5. computation of the cube of residuals for the ADI sequence containing the fake companions and median combination;
6. computation of the throughputs by comparing the aperture flux of the injected companions to that of the retrieved companions after PSF subtraction (difference between final frame of step 5 and step 2);
7. estimation of the contrasts via Eq. 4.1 and computation of the average contrast.

4.2.2 Parameter selection via Bayesian optimisation

The NMF and LLSG PSF-subtraction techniques have integer parameters that are, in practice, restricted to a small range of possible values. One can therefore easily select their optimal parameters by going through their entire parameter space, and simply applying steps 1-7 to compute the contrast for each set of parameters. The optimal set of parameters is the one that minimises the contrast. However, for the other PSF-subtraction techniques, part of the parameter space is continuous, which

¹Following the methodology of Gomez Gonzalez et al. (2017), the intensity of the injected companions represents only a few percent of the pixel intensity within the ADI sequence for a given annulus, which limits the impact of the multiple injections on the estimation of the reference PSF.

prevents exploration of the entire parameter space. A more advanced minimisation algorithm is therefore needed. The derivatives and convexity properties of our loss function are unfortunately unknown. However, it is expected that our loss function, i.e. the function describing the evolution of the annulus-wise contrast in terms of the selected parameters, is non-convex and most probably non-linear. This implies that we cannot rely on mainstream minimisation approaches (e.g. Newton-Conjugate-Gradient algorithm or Nelder-Mead Simplex algorithm). In addition, evaluating the annulus-wise contrast is expensive, because of the numerous steps involved in its estimation. We therefore cannot simply rely on Monte Carlo simulation or random searches to explore the parameter space.

Considering all these constraints, we decided to rely on Bayesian optimisation to select the optimal set of parameters for the remaining PSF-subtraction techniques. Bayesian optimisation is a powerful strategy to limit the number of loss function evaluations needed to reach an extremum (see, e.g. Mockus et al. 1978; Jones et al. 1998). This strategy belongs to a class of algorithms called sequential model-based optimisation (SMBO) algorithms. This class of algorithms uses previous observations of the loss function to determine the position of the next point inside the parameter space to be evaluated. It is called Bayesian optimisation because it relies on Bayes' theorem to define the posterior probability of the loss function, on which the sampling strategy is based. Bayes' theorem states that the posterior probability associated with a model, given a set of observations, is proportional to the likelihood of the observations given the model, multiplied by the prior probability of the model:

$$P(f \mid \mathcal{O}_{1:t}) \propto P(\mathcal{O}_{1:t} \mid f)P(f), \quad (4.2)$$

where f is the loss function to optimise and $\mathcal{O}_{1:t} = \{\mathbf{p}_{1:t}, f(\mathbf{p}_{1:t})\}$ is the set of observations of the loss function, with $\mathbf{p}_{1:t}$ being the set of tested points (see Table C.1 for a summary of all the mathematical notions used throughout the chapter). In the case of Bayesian optimisation, we assume a Gaussian likelihood with noise, as follows:

$$P(\mathcal{O}_{1:t} \mid f) \sim \mathcal{N}(f(\mathbf{p}), \sigma_\epsilon^2), \quad (4.3)$$

where $\mathcal{O}_t = f(\mathbf{p}_t) + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$.

Regarding the prior distribution for our loss function, Mockus (1994) proposed relying on a Gaussian process (GP) prior as this induces a posterior distribution over the loss function that is analytically tractable². A GP is the generalisation of a Gaus-

²This implies that it is possible to update the posterior probability with the observations made

sian distribution to a function, replacing the distribution over random variables by a distribution over functions. A GP is fully characterised by its mean function $m(\mathbf{p})$, and its covariance function \mathbf{K} , where

$$f(\mathbf{p}_{1:t}) \sim \mathcal{GP}(m(\mathbf{p}_{1:t}), \mathbf{K}) \sim \mathcal{N}(m(\mathbf{p}_{1:t}), \mathbf{K}). \quad (4.4)$$

The GP process can be seen as a function that returns the mean and variance of a Gaussian distribution over the possible values of f at \mathbf{p} , instead of returning a scalar $f(\mathbf{p})$. We make the assumption that the prior mean is the zero function $m(\mathbf{p}) = 0$ and we select a commonly used covariance function, the squared exponential function:

$$[\mathbf{K}]_{i,j} = k(\mathbf{p}_i, \mathbf{p}_j) = \exp\left(-\frac{1}{2l^2} \|\mathbf{p}_i - \mathbf{p}_j\|^2\right), \quad (4.5)$$

where l is the length scale of the kernel.

Having documented the posterior probability computation for our loss function, we need to define a sampling strategy. Bayesian optimisation relies on an *acquisition function* to define how to sample the parameter space. This function is based on the current knowledge of the loss function, i.e. the posterior probability. The acquisition function is a function of the posterior distribution over the loss function f , which provides a performance metric for all new sets of parameters. The set of parameters with the highest performance is then chosen as the next point of the parameter space to be sampled. A popular acquisition function is the expected improvement (EI, Mockus et al. 1978; Jones et al. 1998) which is defined as follows:

$$\text{EI}(\mathbf{p}_{t+1}) = \mathbb{E}[\max\{0, f(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}})\}], \quad (4.6)$$

where \mathbb{E} is the expected value and $\hat{\mathbf{p}} = \operatorname{argmax}_{\mathbf{p}_i \in \mathbf{p}_{1:t}} f(\mathbf{p}_i)$ is the current optimal set of parameters. We see that in the case of Bayesian optimisation, we look for the maximum value of the loss function. As we are trying to minimise the contrast for a given set of parameters, we simply define our loss function $f(\mathbf{p})$ as the inverse of the contrast averaged over the selected set of azimuths (see Section 4.2.1).

An interesting feature of the EI is that it can be evaluated analytically under the GP model, yielding (see Appendix C.2 for more details about the derivation of these expressions)

with a new set of parameters. This will help us to create a continuous function to select the next point to sample in the parameter space.

$$\text{EI}(\mathbf{p}_{t+1}) = \begin{cases} (\mu(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}}))\Phi(Z) + \sigma(\mathbf{p}_{t+1})\phi(Z) & \text{if } \sigma(\mathbf{p}_{t+1}) > 0 \\ 0 & \text{if } \sigma(\mathbf{p}_{t+1}) = 0 \end{cases}, \quad (4.7)$$

where $\Phi(Z)$ and $\phi(Z)$ are respectively the cumulative distribution and probability density function of the Gaussian distribution, $\mu(\mathbf{p}_{t+1})$ and $\sigma(\mathbf{p}_{t+1})$ are the mean and variance of the Gaussian posterior distribution, and $Z = [\mu(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}})]/\sigma(\mathbf{p}_{t+1})$. We see from this last expression that the EI is high either when the expected value of the loss $\mu(\mathbf{p})$ is larger than the maximum value of the loss function $f(\hat{\mathbf{p}})$ or when the uncertainty $\sigma(\mathbf{p}_{t+1})$ around the selected set of parameters \mathbf{p}_{t+1} is high. The EI approach aims to minimise the number of function evaluations by performing a trade-off between exploitation and exploration at each step. The EI exploits the existing set of observations by favouring the region where the expected value of $f(\mathbf{p}_{t+1})$ is high, while it also explores unknown regions where the uncertainty associated with the loss function is high.

Bayesian optimisation starts with the initialisation of the posterior probability by estimating the loss function for several sets of parameters via random search in the parameter space. Once this initial population of observations is computed, the rest of the algorithm can be summarised as follows.

- Based on the GP model, use random search to find the \mathbf{p}_{t+1} that maximises the EI, $\mathbf{p}_{t+1} = \text{argmax} [\text{EI}(\mathbf{p}_{t+1})]$;
- compute the contrast for the new set of parameters \mathbf{p}_{t+1} ;
- update the posterior expectation of the contrast function using the GP model (see Appendix C.2);
- repeat the previous steps for a given number of iterations.

The number of random searches to compute the initial GP and the number of iterations for the Bayesian optimisation depend on the size of the parameter space associated with the considered PSF-subtraction techniques. A specific number of random searches and iterations are therefore selected for each PSF-subtraction technique. At the end of the Bayesian optimisation, the minimal average contrast for a given annulus a and PSF-subtraction technique m is stored in a matrix element $C_{a,m}$, along with the set of parameters \mathbf{p} in another matrix $P_{a,m}$.

This first step of the auto-RSM algorithm may be used outside the RSM framework, allowing the production of S/N maps based on the cubes of residuals generated by optimised PSF-subtraction techniques. A S/N-based version of the auto-RSM framework

called auto-S/N has been developed and is presented in Appendix C.4. Auto-S/N optimally combines S/N maps computed from the cubes of residuals generated by the optimised PSF-subtraction techniques, relying on the same greedy approach as for auto-RSM (see Section 4.3.2). The performance of auto-S/N is assessed in Appendix C.4.2 using the same metrics as for auto-RSM (see Section 4.4). The lower performance of auto-S/N implies that auto-RSM should be preferred, despite its longer computation time, although the two approaches can be complementary to some extent.

4.3 RSM map optimisation

4.3.1 Parameter selection for the RSM map

Following the optimisation of the PSF-subtraction techniques to be used in the RSM model (Section 4.2), the next step is to consider the parametrisation of the RSM algorithm itself. The use of the RSM algorithm requires the definition of four main parameters. These parameters are (i) the crop-size θ for the planetary model \mathbf{m} , (ii) the definition of the region of the cube of residuals considered for the computation of the noise properties, whose estimation can be done (iii) empirically or via best fit, and (iv) the method used to compute the intensity of the potential planetary candidate β . When defining the flux parameter β as a multiple of the noise standard deviation, an additional parameter δ has to be used to determine how far into the noise distribution tail we are looking for potential planetary candidates.

An optimal set of parameters for the RSM algorithm is computed separately for each PSF-subtraction technique, and is based on a performance metric computed using the generated RSM map. We do not rely on multiple simultaneous injections of fake companions at different azimuths, as done previously, as the RSM approach assumes a single planetary signal per annulus. Injecting the fake companions sequentially would largely increase the computation time. We therefore define, annulus-wise, a single median position in terms of noise intensity, common to all PSF-subtraction techniques. This allows a fair comparison between the PSF-subtraction techniques when selecting the best set of likelihood cubes to generate the final RSM map in the last step of the auto-RSM framework. The determination of this median position starts with derotation of the original ADI sequence and the median-combination of the resulting set of frames. We then compute the flux of every aperture contained in the selected annulus, each aperture centre being separated by a single pixel in contrast with the approach of Mawet et al. (2014), where the apertures centre are separated by one FWHM. We define the fake companion injection position as the centre of the aperture for which the flux is the median of all the apertures fluxes. We decided to compute this median-flux position in the original ADI sequence, as the median-flux position inside the PSF-subtracted final frame differs from one PSF subtraction technique to

the other, although a single common position is required for the final step of the auto-RSM algorithm. Regarding the contrast used for the optimisation of the RSM map parameters, for each PSF-subtraction technique, we select the average contrast $C_{a,m}$ obtained with the optimal set of parameters (see Table C.2 for a summary of the mathematical notation used to describe the auto-RSM framework). Here, we make the assumption that taking the median-flux position and the average contrast should provide a balanced optimised parametrisation that works for brighter as well as fainter planetary signals.

The performance metric used for the RSM algorithm optimisation is then defined as the peak probability in a circular aperture with a diameter of one FWHM centred on the position of the injected fake companion in the final RSM map divided by the maximum probability observed in the remaining part of the annulus of width equal to one FWHM. This allows us to account for potential bright speckles within the probability map as well as for the intensity of the planetary signal. Having defined the loss function used for the RSM parameter selection, we now consider the different parameters that should be optimised.

Crop size

The crop size θ is one of the parameters affecting the final probability map the most. This is especially true when relying on forward-model versions of the PSF-subtraction techniques, where a larger crop size should be considered to take advantage of the modelling of the negative side lobes appearing on either side of the planetary signal peak, which are due to self-subtraction associated with PSF-subtraction techniques. As seen in chapter 3, self-subtraction depends on the relative position of the planetary candidate compared to the host star, with stronger self-subtraction at small angular separations and almost no self-subtraction at large angular separations. Indeed, the apparent movement of the planetary candidate increases linearly with the distance to the host star as the parallactic angles remain fixed but the radius increases. Larger apparent movement between two frames goes along with reduced self-subtraction. This implies that the optimal crop size for forward-modelled PSF should decrease with angular separation, as the negative side lobes appearing on either side of the planetary signal peak are replaced by noise. The selection of the optimal crop size should account for this effect as well as the range of parallactic angles, which is specific to each data set and also affects self-subtraction patterns. For PSF-subtraction techniques relying on the off-axis PSF to model the planetary signal, we consider a smaller range of crop sizes, as we do not take into account the distortion due to reference PSF subtraction. A maximum size of one FWHM is considered when relying on off-axis PSFs compared to the two FWHMs used for forward-model PSF-subtraction techniques. The definition of a proper crop size is nevertheless still important, because considering the shape of

the PSF peak should help in disentangling planetary signals from speckle noise.

Parametrisation of the noise distribution

One of the corner stones of the RSM algorithm is the proper definition of the likelihood function associated with every patch contained in a given annulus. Four potential noise distribution functions are considered to compute these likelihoods, namely the Gaussian and the Laplacian distribution, the Huber loss (Pairet et al. 2019), and a hybrid distribution built as a weighted sum of Gaussian and Laplacian distributions (see chapter 2). The first two noise distribution functions require estimation of the noise mean and variance, whereas the other two require additional parameters. Selection of the optimal distribution is done automatically within the RSM algorithm via a best-fit approach. However, the estimation of the parameters characterising the residual noise distribution function necessitates proper definition of the set of pixels to be considered. Different approaches are tested in auto-RSM to determine the most relevant set of pixels inside the cube of residuals. We have selected five possible ways to evaluate the noise properties:

- Spatio-temporal estimation: The set of pixels incorporates the pixels inside the selected annulus³ for all the frames contained in the cube of residuals (see ‘Spatio-temporal’ in Figure 4.2). The distribution function parameters depend solely on the radial distance a (μ_a and σ_a^2).
- Frame-based estimation: The set of pixels incorporates the pixels of a given frame inside the selected annulus (see ‘Frame’ in Figure 4.2). The distribution function parameters depend on both the radial distance a and the time-frame t ($\mu_{a,t}$ and $\sigma_{a,t}^2$).
- Frame with mask-based estimation: The set of pixels incorporates the pixels of a given frame inside the selected annulus, apart from a region with a diameter of one FWHM centred on the pixels for which the likelihood is estimated (see ‘Frame with mask’ in Figure 4.2). The distribution function parameters depend on both the radial distance a and the pixels index i_a (μ_{a,i_a} and σ_{a,i_a}^2).
- Segment with mask-based estimation: The set of pixels incorporates the pixels of all frames inside a section (of length equal to three FWHMs) of the selected annulus, apart from a region with a diameter of one FWHM centred on the pixels for which the likelihood is estimated (see ‘Segment with mask’ in Figure 4.2). The distribution function parameters depend on both the radial distance a and the pixels index i_a (μ_{a,i_a} and σ_{a,i_a}^2).

³ By selected annulus, we are referring to the annulus of one FWHM in width centred on the radial distance of interest a .

- Temporal estimation: The last method is inspired by the approach developed in Flasseur et al. (2018). This approach relies on the cube of residuals before de-rotation. For a given patch inside the selected annulus, the pixels selected for computation of the distribution function parameters are the ones sharing the same position within the cube of residuals before de-rotation but taken at different times (see ‘Temporal’ in Figure 4.2). All the frames except for the frame containing the selected patch are therefore considered. The distribution function parameters depend on both the radial distance a and the pixels index i_a (μ_{a,i_a} and σ_{a,i_a}^2).

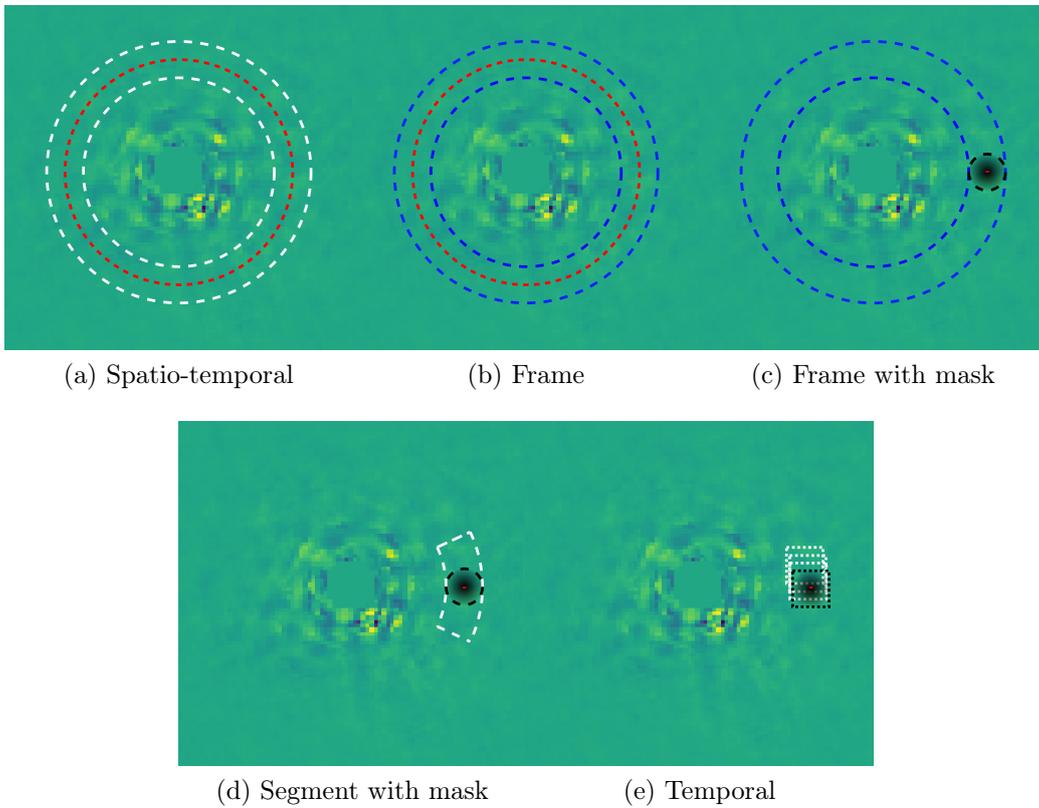


Figure 4.2: Graphical representation of the estimation of residual noise properties using the five proposed approaches. The red circle/point indicates the pixel for which the likelihood is estimated. White and blue circles encompass the set of pixels used for computation of the noise properties. White circles indicate that the entire set of frames from the derotated cube are used for the computation, while blue circles indicate that the estimation is done frame-wise. Black circles define a mask, i.e. pixels that are not considered in the estimation.

The use of these different methods allows us to investigate which part of the neighbourhood around the patch is relevant in order to correctly estimate the noise profile. This explains the wide variety of proposed methods both in terms of temporal and spatial position.

Depending on the region selected to compute the noise properties, a specific noise distribution function and parametrisation can be selected for a single patch, a single frame, or the entire set of frames and patches contained in the considered annulus. The estimation of noise distribution parameters can be done empirically or via best fit. The choice between empirical estimation and estimation via best fit represents an additional parameter to be considered during the RSM parameter optimisation.

Estimation of the planetary intensity

Two different methods were proposed to compute the planetary intensity parameter β in chapter 2 and 3. The first one relies on an additional parameter δ to define the expected position of the potential planetary signal intensity in the noise distribution. The intensity parameter β is defined as δ multiplied by the estimated noise standard deviation. A set of δ is tested and the optimal one is selected via maximisation of the total likelihood associated with a given angular distance (see Eq. 2.7). In the case of auto-RSM, this last step is removed and the optimal δ is selected during the auto-RSM optimisation process. Preliminary tests have shown that the optimisation of δ using the auto-RSM performance metric can significantly reduce the background noise in the RSM probability map compared to the total likelihood-based optimisation proposed in chapter 2, while leaving planetary signals almost unaffected for $\delta \leq 5$.

The second approach relies on Gaussian maximum likelihood to define a pixel-wise intensity. The estimation of the intensity parameter β via Gaussian maximum likelihood requires the computation of a frame-wise standard deviation. The expression for the pixel-wise intensity is

$$\tilde{\beta} = \frac{\sum_j^T \mathbf{i}_j^\top \mathbf{m}_j / \sigma_j}{\sum_j^T \mathbf{m}_j^\top \mathbf{m}_j / \sigma_j}, \quad (4.8)$$

with σ_j being the noise standard deviation, \mathbf{i}_j^\top the observed patch, and \mathbf{m}_j the planet model for frame j . Frame-wise computation of the standard deviation implies that the mean and variance computation described in the previous section should be performed via the frame-based estimation, the frame with mask-based estimation, or the temporal estimation, while for the other intensity computation methods, all five approaches can be used.

Sequential parameter optimisation

While optimisation of the PSF-subtraction techniques is done in a single step, optimisation of the RSM parameters is done partly sequentially. The estimation of the RSM map is indeed much slower than the estimation of the contrast. Depending on the region considered for the estimation of the noise properties, the computation time can

further increase, especially when relying on the frame with mask-based estimation or the temporal estimation, preventing optimisation of all the parameters in a single step. The selection of the optimal region to compute the noise properties is therefore treated separately. The selection of the optimal set of RSM parameters starts with computation of the RSM map performance metric for the two methods used to determine the intensity parameter β using the frame-based estimation of the noise properties⁴. For both methods, a separate performance metric is estimated for the selected range of crop sizes, but also for the selected range of δ for the method defining the intensity as a multiple of the noise standard deviation. The selection of the optimal value for all three parameters (i.e. δ , the crop size θ and the method to compute the intensity) is performed by comparing the obtained RSM performance metric. The faster computation of the RSM map when relying on the frame-based estimation of the noise properties allows optimisation of these three parameters in a single step. The next step involves the selection of the optimal region for estimation of the noise properties. Depending on the method selected to compute the intensity, a reduced set of regions may be considered. Optimisation of the RSM parameters ends by determining whether the noise properties are optimally computed empirically or via a best fit.

4.3.2 Optimal combination of the likelihood cubes

Having optimised the parameters of the PSF-subtraction techniques as well as the ones of the RSM algorithm, we are now left with a series of optimal cubes of likelihoods. One of the most interesting features of the RSM framework is its ability to use several cubes of residuals generated with different PSF-subtraction techniques to maximise the planetary signal, while minimising the residual speckle noise. The RSM algorithm takes advantage of the diversity of noise structures in the different cubes of residuals. This diversity is reflected in the noise probability distribution but also in the repartition of maxima and minima in the different speckle fields. By taking both aspects into account, the RSM algorithm is able to better average out the noise and improve the ratio between potential planetary signals and the residual speckle noise.

Despite optimisation of the parameters, some PSF-subtraction techniques may be less suited to generating a clean cube of residuals for some data sets. Redundancies in the information contained in several cubes of residuals may also degrade the performance of the RSM map by increasing the relative importance of some speckles. When dealing with several ADI sequences of the same object, some sequences can also be much noisier depending on the observing conditions. All these elements necessitate proper selection of the likelihood cubes used to generate the optimal final RSM map.

⁴The frame-based estimation of the noise properties has been selected as initial guess because it is shared by the two approaches used to compute the intensity parameter β , and is much faster than the frame with mask and the temporal estimations.

We propose the investigation of two possible approaches to select the set of likelihood cubes used for computation of the final probability map, a bottom-up approach and a top-down approach, making use of a greedy selection framework.

As the RSM algorithm relies on spatio-temporal series of likelihoods to compute annulus-wise probabilities (see Eq. 2.3), we start by defining the set of available series of likelihoods for a given radius a by $\mathbf{Y}^a = \{Y_{c,m}^a, \forall c \in [0, N_{sequence}], m \in [0, N_{technique}]\}$. The $Y_{c,m}^a$ time-series corresponds to the set of likelihoods η_{s,i_a} given in the Gaussian case by Eq. 2.7, generated for the cube c with the PSF-subtraction techniques m for all pixel indices i_a of the annulus a . This last step of auto-RSM is used to define a subset $\mathbf{Z}^a \subset \mathbf{Y}^a$ regrouping the series of likelihoods maximising the performance metric of the RSM probability map for annulus a . This selection step shares the same performance metric as the RSM parameters optimisation step. To compute the performance metric, a single fake companion injection is used for each annulus for the entire set of PSF-subtraction techniques⁵. The selected set of time-series of likelihoods \mathbf{Z}^a are then concatenated to form a single time-series per annulus and used to compute the probabilities via Eq. 2.3. The RSM performance metric, estimated based on these probabilities, allows us to select the optimal set \mathbf{Z}^a .

Bottom-up approach

When relying on a bottom-up approach, the iterative selection algorithm starts with an empty set \mathbf{Z}^a . At each iteration, the series of likelihoods $Y_{c,m}^a$ that leads to the highest performance metric increase is added to the set \mathbf{Z}^a . The procedure is repeated until no additional series of likelihoods leads to an increase in the performance metric. The bottom-up greedy selection algorithm can be summarised by the following steps.

- For each series of likelihood contained in \mathbf{Y}^a , compute the corresponding RSM map performance metric using the set of series of likelihoods $\mathbf{Z}^a \cup Y_{c,m}^a$ for annulus a .
- At each iteration, select the series of likelihoods providing the largest incremental performance metric increase and include the considered series of likelihoods Y_{c^*,m^*}^a in the set of selected series \mathbf{Z}^a . Remove from \mathbf{Y}^a the selected series Y_{c^*,m^*}^a , as well as any other series included in \mathbf{Y}^a that did not lead to an increase in the performance metric.
- Repeat the previous two steps until \mathbf{Y}^a is empty.

⁵For a given annulus a^* , the largest contrast in the set $C_{a^*,m}$ is used for the bottom-up approach and the smallest for the top-down, as they provide the best performance based on tests

Top-down approach

In contrast with the bottom-up approach, the top-down iterative selection algorithm starts with a set $\mathbf{Z}^a = \mathbf{Y}^a$ and relies on pruning steps to reduce the number of series of likelihoods included in \mathbf{Z}^a until an optimum is reached. The steps of the top-down greedy selection algorithm are the following.

- For each series of likelihood contained in \mathbf{Z}^a , compute the RSM map performance metric corresponding to the set of series of likelihoods $\mathbf{Z}^a \setminus Y_{c,m}^a$ for annulus a .
- At each iteration, select the series of likelihoods providing the largest incremental performance metric increase and remove the considered series of likelihoods Y_{c^*,m^*}^a from the set of selected series \mathbf{Z}^a .
- Repeat the two previous steps until no more incremental performance metric decrease can be observed.

Pseudo codes of both approaches are provided in Tables C.3 and C.4. The potential redundancies in the information contained in different cubes of likelihoods, as well as the iterative procedure used by the RSM algorithm to generate the final probability map, mean that the set of series of likelihoods are not truly independent, which prevents us from finding the global optimum while using a greedy approach. However, these bottom-up and top-down greedy selection algorithms provide a good approximation of the global optimum in a reasonable amount of time.

4.3.3 Practical implementation

After having presented the different steps of the proposed optimisation framework for the RSM map algorithm, these steps can now be merged into a single optimisation procedure, which is implemented in the PyRSM package⁶. Two different modes of this optimisation procedure are proposed: the full-frame mode and the annular mode. The two modes share a common structure but their output dependence on the angular separation is different. In the full-frame mode, there is no dependence between the optimal set of parameters and the angular separation to the host star, with a single set of parameters being used for every annulus. In the annular mode, the frames are divided into successive annuli of pre-defined width, and a set of optimal parameters is defined for each annulus. As the noise distribution and parameters evolve with the angular separation, this second mode accommodates possible evolutions of the optimal parametrisation with the angular separation to the host star. Figure 4.3 provides a graphical representation of the two different modes in the case of optimisation of the FOV minimum rotation used for the annular PCA estimation.

⁶<https://github.com/chdahlqvist/RSMmap>.

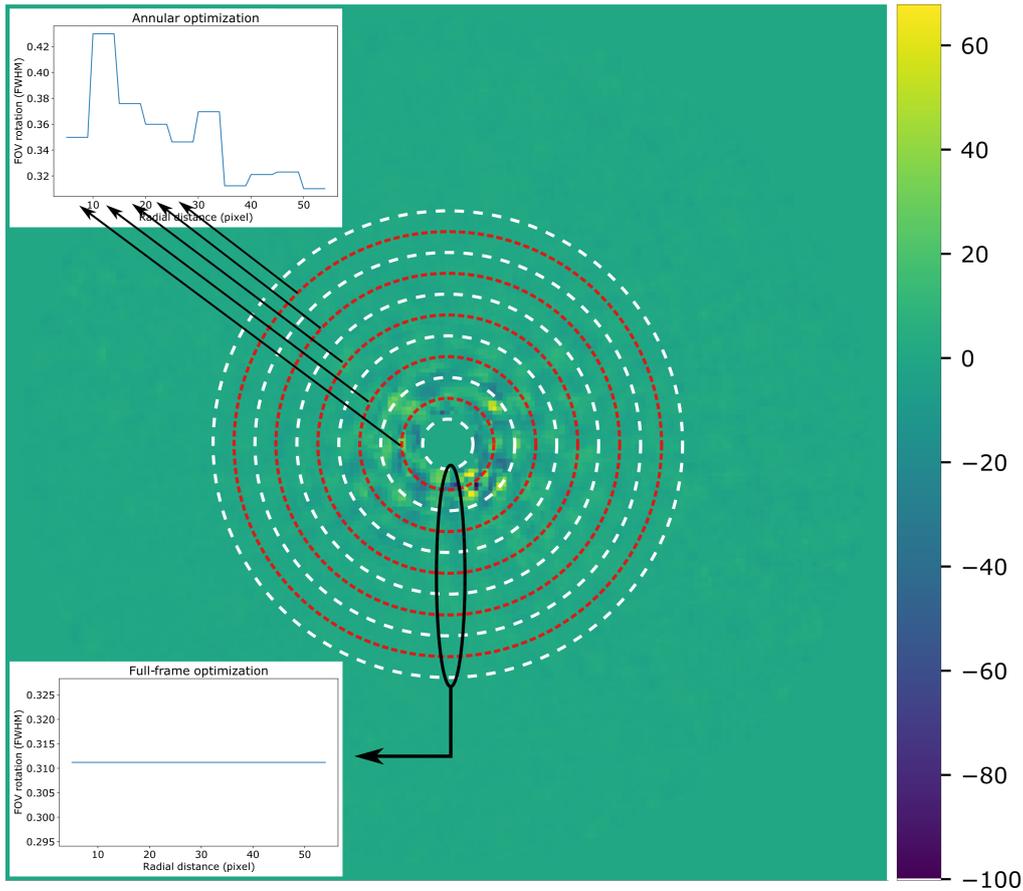


Figure 4.3: Graphical representation of the optimisation of the FOV minimal rotation for the annular PCA with the two modes of the auto-RSM algorithm for the SPHERE 1 data set of the EIDC. The full-frame version (illustrated in the bottom left corner) considers a set of annuli of width equal to one FWHM to provide a single set of optimal parameters. The annular version (top left) considers successive annuli of width equal to one FWHM to provide annulus-wise sets of optimal parameters.

As illustrated in Figure 4.3, in both cases the frames are divided into successive annuli of one FWHM⁷ in width. The red dotted circles represent the centres of the selected annuli on which the apertures for the optimisation of the PSF-subtraction techniques are centred, and for which probabilities are computed to optimise the parameters of the RSM algorithm and select the optimal set \mathbf{Z}^a . We do not consider all angular separations but a subset of them separated by one FWHM, as we expect a slow evolution of the parameters. This should give a good representation of the evolution of the parameters or a good overview in the case of the full-frame mode, while reducing the computation time⁸.

In the case of the full-frame mode, we consider a subset of the annuli of the annular

⁷A width equal to one FWHM often provides the best performance, but other widths can be used.

⁸A mode considering all angular separations has been tested and provides results close to the other modes while requiring a much longer computation time.

mode, with an increasing distance between the selected annuli as we move away from the host star. This allows us to increase the relative weight of small angular separations, the noisier region being located near the host star, and reduce the estimation time. The proposed annulus selection rule for the full-frame mode can be summarised as

$$\Delta a = \begin{cases} \text{FWHM} & \text{if } a \in [\text{FWHM}, 4 \text{ FWHM}] \\ 2 \text{ FWHM} & \text{if } a \in]4 \text{ FWHM}, 8 \text{ FWHM}] , \\ 4 \text{ FWHM} & \text{if } a \in]8 \text{ FWHM}, a_{max}] \end{cases} \quad (4.9)$$

where Δa is the separation between two successive annuli used in the optimisation procedure, and a_{max} is the largest annulus to be considered in the RSM map computation. Selection of the optimal parameter set for the PSF-subtraction techniques in the full-frame mode is achieved by comparing the normalised contrasts generated with the different tested parametrisations summed over the selected angular separations. We start by computing contrasts for a common set of parametrisations⁹ for each considered angular separation. For a given angular separation, the median of the obtained contrasts is then computed and used to normalise all the contrasts. The normalised contrasts are finally summed over the selected angular separations provided by Eq. 4.9 for each considered parametrisation. The optimal set of parameters is then the one that minimises the summed normalised contrast¹⁰. As the contrast decreases with the angular separation, the normalisation allows a proper summation of the contrasts generated at the different angular distances. A similar approach is used for optimisation of the RSM algorithm and selection of the optimal likelihoods, although no normalisation is required according to the definition of the performance metric. As regards the annular mode, no normalisation is required as the optimisation is done separately for each selected annulus.

The complete auto-RSM optimisation procedures for the two considered modes are summarised in Tables C.5 and C.6. As can be seen from both tables, the optimisation procedures can be divided into four main steps, (i) optimisation of the PSF-subtraction techniques, (ii) optimisation of the RSM algorithm, (iii) optimal combination of models and sequences, and (iv) computation of the final RSM probability map (respectively the `opti_model`, `opti_RSM`, `opti_combination`, and `opti_map` function of the `PyRSM` class). In both modes we include the estimation of a background noise threshold for every annulus, by taking, for each annulus, the maximum probability observed in

⁹We consider all the tested parametrisations for the NMF and LLSF and the parametrisations tested during the initialisation of the Gaussian process for the other PSF-subtraction techniques.

¹⁰The inverse of the normalised average contrast summed over the considered angular separation is used as loss function for the Bayesian optimisation.

the map generated with the reversed parallactic angles. Following subtraction of the angular separation-dependent thresholds, we set all negative probabilities to zero to generate the final map. The threshold subtraction should help to reduce the noise, especially near the host star where most residual speckles are observed. However, these thresholds should not be used as detection thresholds, as the noise statistics properties of the original ADI sequence are not exactly equivalent to the ADI sequence with sign-flipped parallactic angles.

Considering the existence of potential bright artefacts in the map generated with the reversed parallactic angles, we rely on a Hampel filter and a polynomial fit to smooth the radial evolution of the thresholds in the full-frame mode. As the parametrisation of the RSM algorithm has a large impact on planetary signals and background noise levels, we do not apply the threshold fit for the annular mode, as the RSM parametrisation evolves with the angular separations. However, we do apply a smoothing procedure for the parameters of the PSF-subtraction techniques by applying a moving average after a Hampel filter. This helps in smoothing potential discontinuities between annuli in the set of optimal parameters and provides a more consistent final probability map. As computation of the residual cubes is done annulus-wise¹¹, we need a single set of parameters¹² for a number of angular separations equal to the width of the annulus. However, the RSM map computation requires the definition of a set of parameters for every considered angular separation. A radial basis multiquadric function (RBF) is used to perform an interpolation (Hardy 1971) of the RSM optimal parameters for the annular mode to provide a set of parameters for each angular distance.

4.4 Performance assessment

4.4.1 Description of the data sets

As mentioned in the introduction, we base our performance analysis on the data set of the EIDC ADI subchallenge (Cantalloube et al. 2020b). This data set regroups nine ADI sequences, three for each considered HCI instrument, namely VLT/SPHERE-IRDIS (Beuzit et al. 2019), Keck/NIRC2 (Serabyn et al. 2017), and LBT/LMIRCam (Skrutskie et al. 2010). The ADI sequences were obtained in H2-band for SPHERE and Lp-band for the two other instruments. For each ADI sequence, a set of four fits files is provided: the temporal cube of images, the parallactic angles variation corrected from true north, a non-coronagraphic or non-saturated PSF of the instrument, and

¹¹An annular version of the NMF algorithm has been developed for the annular mode of the auto-RSM. The other PSF-subtraction techniques rely already on an annulus-wise estimation of the residuals.

¹²The set of parameters that has been optimised based on the set of apertures centred on the annulus

the pixel scale of the detector. The ADI sequences are pre-reduced using the dedicated pre-processing pipeline for the three instruments (more details about the reduction are provided by Cantalloube et al. 2020b).

As the LMIRCam ADI sequences regroup between 3219 and 4838 frames, we relied on moving averages to reduce this number to around 250 frames in order to limit the computation time. The reduction of the number of frames starts with the de-rotation of the original cube of images using the parallactic angle variations provided. A moving average is then applied on the de-rotated cubes along the time axis with a window and step size of 20 frames for the LMIRCam sequence 1 and 3, and 15 frames for the LMIRCam sequence 2. The same is done on the set of parallactic angle variations. The inverse reduced parallactic angle variations are then used to re-rotate the resulting ADI sequences. In addition to reducing the computation time, the moving average allows part of the noise to be averaged out in advance. More details about the nine ADI sequences are provided in Table C.7.

To assess the performance of the HCI data-processing techniques, fake companions were injected by the EIDC organisers using the VIP package (Gomez Gonzalez et al. 2017). Between 0 and 5 point sources were injected into each ADI sequence for a total of 20 planetary signals within the entire EIDC ADI subchallenge data set. These point sources were injected using the opposite parallactic angles, avoiding any interference with potential existing planetary signals while keeping the speckle noise statistics¹³. The separation, the azimuth, and the contrast of the injected fake companions were chosen randomly. The contrasts range between 2σ and 8σ based on a contrast curve computed with the regular annular PCA implemented in the VIP package (Gomez Gonzalez et al. 2017), which is referred to as the ‘baseline’ in the performance analysis presented in Cantalloube et al. (2020b). The detection maps of the baseline consist in S/N maps computed using the approach of Mawet et al. (2014). The detection maps generated with the baseline approach are used in our model comparison.

4.4.2 Performance metrics

The performance assessment of HCI data-processing techniques is done via the definition of a classification problem, counting detections and non-detections on a grid of FWHM-sized apertures applied to the detection maps. A true positive (TP) is defined as a value above the threshold provided by the user along with the S/N or probability maps within the FWHM aperture centred on the position of the injected fake companion. Any values above the provided threshold that are not in the set of apertures containing injected fake companions are considered as false positives (FPs).

¹³As the auto-RSM relies on reversed parallactic angles to optimise the model parameters, the optimisation is done on the original ADI sequences in the case of the EIDC data sets. However, the ADI sequences selected for the EIDC do not contain any known planetary candidates. The optimisation should therefore not be affected.

The false negatives (FNs) regroup all the non-detections at the position of injected fake companions, while the true negatives (TNs) are the non-detections at any other position. Different performance metrics are computed using these four categories:

- True positive rate: $TPR = \frac{TP}{TP+FN}$
- False positive rate: $FPR = \frac{FP}{FP+TN}$
- False discovery rate: $FDR = \frac{FP}{FP+TP}$
- F1 score: $F1 = \frac{2TP}{FP+FN+2TP}$

In addition to the F1 score computed at the pre-defined threshold, we follow the same approach as in Cantalloube et al. (2020b) and also consider the area under the curve (AUC) for the TPR, FPR, and FDR as a function of the threshold to classify the different versions of the proposed optimisation procedure. The AUCs of the TPR, FPR, and FDR are preferred to the values of these latter at the provided threshold, as this allows us to mitigate the arbitrariness of the threshold selection by considering their evolution for a range of thresholds. The AUC of the TPR should be as close as possible to 1 and the AUCs of the FPR and FDR as close as possible to zero. The F1 score being the harmonic mean of the recall and precision of the classification problem, it ranges between 0 and 1, with values close to 1 being favoured (perfect recall and precision).

4.4.3 Results

We have now all the elements to apply the auto-RSM optimisation procedure described in Sects. 4.2 and 4.3 to the nine selected ADI sequences. Only PSF-subtraction techniques relying on an off-axis PSF (and not on forward models) are considered during the optimisation procedure in order to reduce the computation time, considering the large set of ADI sequences on which the method is tested, and the numerous parametrisations of the auto-RSM algorithm that we consider. This also allows a fair comparison with the results of the RSM algorithm already used in Cantalloube et al. (2020b), which relied only on the PSF-subtraction techniques based on off-axis PSFs. We note however that the PyRSM Python package also accommodates the use of a forward-model version of KLIP and LOCI, where a parameter defines the maximum angular separation above which the forward model is no longer considered. This allows us to take advantage of the higher performance of forward-modeled PSF-subtraction techniques at small angular separations, while limiting their impact on the computation time at larger separations.

For the Bayesian optimisation of APCA and LOCI, the contrast is computed for respectively 80 and 60 points of the parameter space to initialise the Gaussian process,

while 60 iterations of the minimum-expectation Bayesian optimisation are used to determine the optimal set of parameters. The smaller number of points for the initialisation of LOCI comes from its smaller set of parameters compared to APCA. The number of points for the initialisation and the number of iterations have been chosen to ensure the convergence to the global optimum¹⁴. The ranges of possible values that have been selected to define the parameter space for the PSF-subtraction optimisation are shown in Table C.8. Most ADI sequences share a common range of possible values. However, differences may be found in the definition of the parameter space boundaries for the NIRC2 ADI sequences due to the reduced number of frames.

Full-frame and annular auto-RSM parametrisation

Regarding the parameters of the full-frame version of auto-RSM, the set of selected annuli is truncated at $a_{max} = 10\lambda/D$ to favour small angular separations during the optimisation, the region close to the host star being more noisy¹⁵. The order of the polynomial fit of the annular threshold is set to three in order to limit the impact of small artefacts appearing in the RSM map generated with inverted parallactic angles, while keeping the main characteristics of the angular evolution of the noise. The full-frame (FF) auto-RSM was tested with three different parametrisations, allowing the comparison between the bottom-up (BU) and top-down (TD) selection of the optimal cubes of likelihoods, as well as the comparison between the forward (F) and forward-backward (FB) approaches to compute the final probabilities.

The annular version of auto-RSM requires the definition of two additional parameters: the window sizes for the Hampel filter and the moving average used to smooth the PSF-subtraction parameters (see Table C.6). The window sizes are respectively equal to 3 and 5, and the window is centred on the angular distance for which the filtered value or the moving average is computed. Two different flavours of the annular auto-RSM were tested, one relying on the annular framework for the optimisation of the entire set of parameters (A), and one using the annular framework for the PSF-subtraction parametrisation and the full-frame framework for the RSM parametrisation and the selection of the optimal set of cubes of likelihoods. The hybrid approach mixing full-frame and annular frameworks (AFF) aims to reduce the angular variability of the background residual probabilities, which are mainly affected by the parametrisation of the RSM model.

¹⁴This parametrisation of the Bayesian optimisation algorithm ensures that the same set of optimal parameters is found when the algorithm is applied several times to the same ADI sequence.

¹⁵It also allows us to reduce the computation time, the larger angular separations being computationally more expensive.

Performance metric computation and model comparison

Having presented the five tested parametrisations of the full-frame and annular auto-RSM, we now turn to the estimation of the detection maps and the computation of the performance metrics, which will allow us to rank these parametrisations and compare them with both the original RSM algorithm and the baseline presented in Cantalloube et al. (2020b). All parametrisations of the auto-RSM were applied to the nine data sets of the EIDC. Figure 4.4 presents the detection maps generated with the full-frame auto-RSM using the bottom-up greedy algorithm to select the optimal set of cubes of likelihoods and the forward approach to compute the probabilities (auto-RSM FF_BU_F). The detection maps for all five parametrisations of the auto-RSM are provided in Appendix C.8. As can be seen from Figure 4.4, the contrast between detected targets and background residual probabilities is very high compared to standard S/N maps, demonstrating the ability of the proposed approach to easily disentangle planetary signals from residual speckles and ease the selection of a detection threshold. As an illustration, the ratio between the peak probability (or S/N) of the target and the mean of the background probabilities (or S/Ns) in the detection map of the SPHERE 1 data set is larger than 3000 for the auto-RSM FF_BU_F, and only 2 for the baseline¹⁶.

Following the EIDC procedure, a single threshold was selected for all data sets, for each parametrisation of the auto-RSM. This threshold allows estimation of the F1 score. As mentioned in Section 4.4.2, in addition to the F1 score, the AUCs of the TPR, FPR, and FDR are also computed. Figure 4.5 illustrates the computation of the different performance metrics for the nine data sets, relying on the detection maps generated with the auto-RSM FF_BU_F. The TPR, FPR, and FDR are computed for different threshold values ranging from zero to twice the selected threshold. The AUCs of the TPR, FPR, and FDR are computed in this interval. Apart from the NIRC2 data sets and LMIRCam-2, the AUC of the FDR is very small for the remaining data sets compared to the baseline. The AUC of the FPR is close to zero for all data sets, especially for the SPHERE data sets for which the AUC values are below the considered 0.001 limit.

¹⁶In the case of the baseline S/N map, the minimum S/N value has been added to the S/N map to have only positive values.

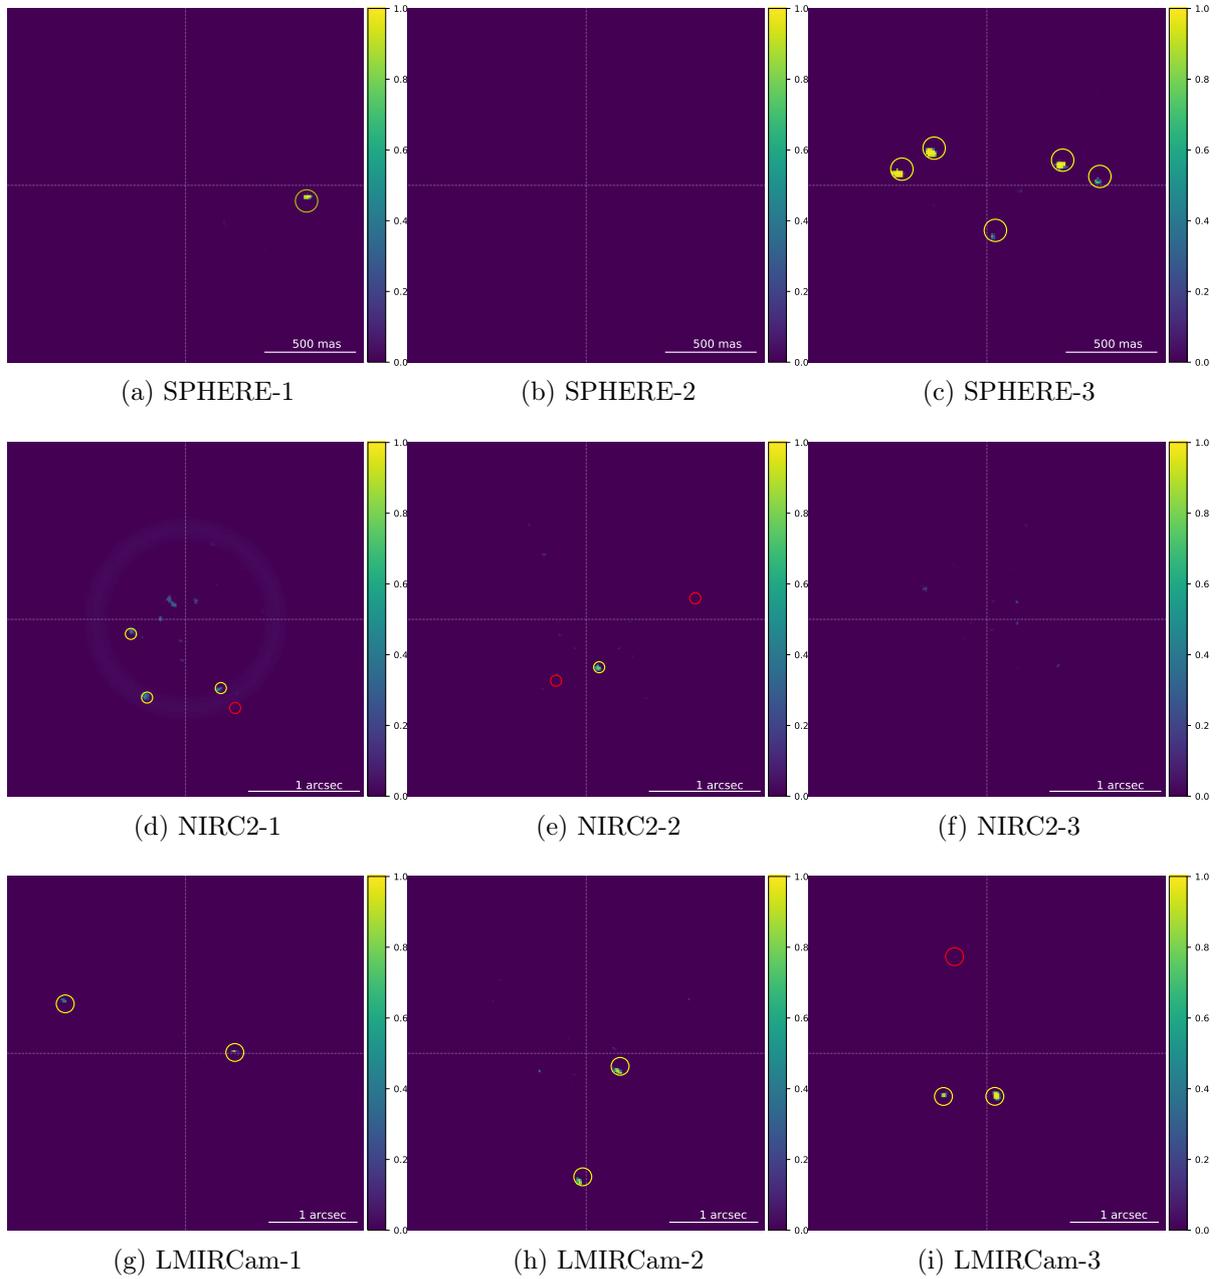


Figure 4.4: Detection maps corresponding to the nine data sets of the EIDC, generated with the full-frame version of auto-RSM using the bottom-up approach for the selection of the optimal set of cubes of likelihoods, as well as the forward approach for the computation of the probabilities. The yellow circles are centred on the true position of the detected targets (TP) and the red circles give the true positions of FNs.

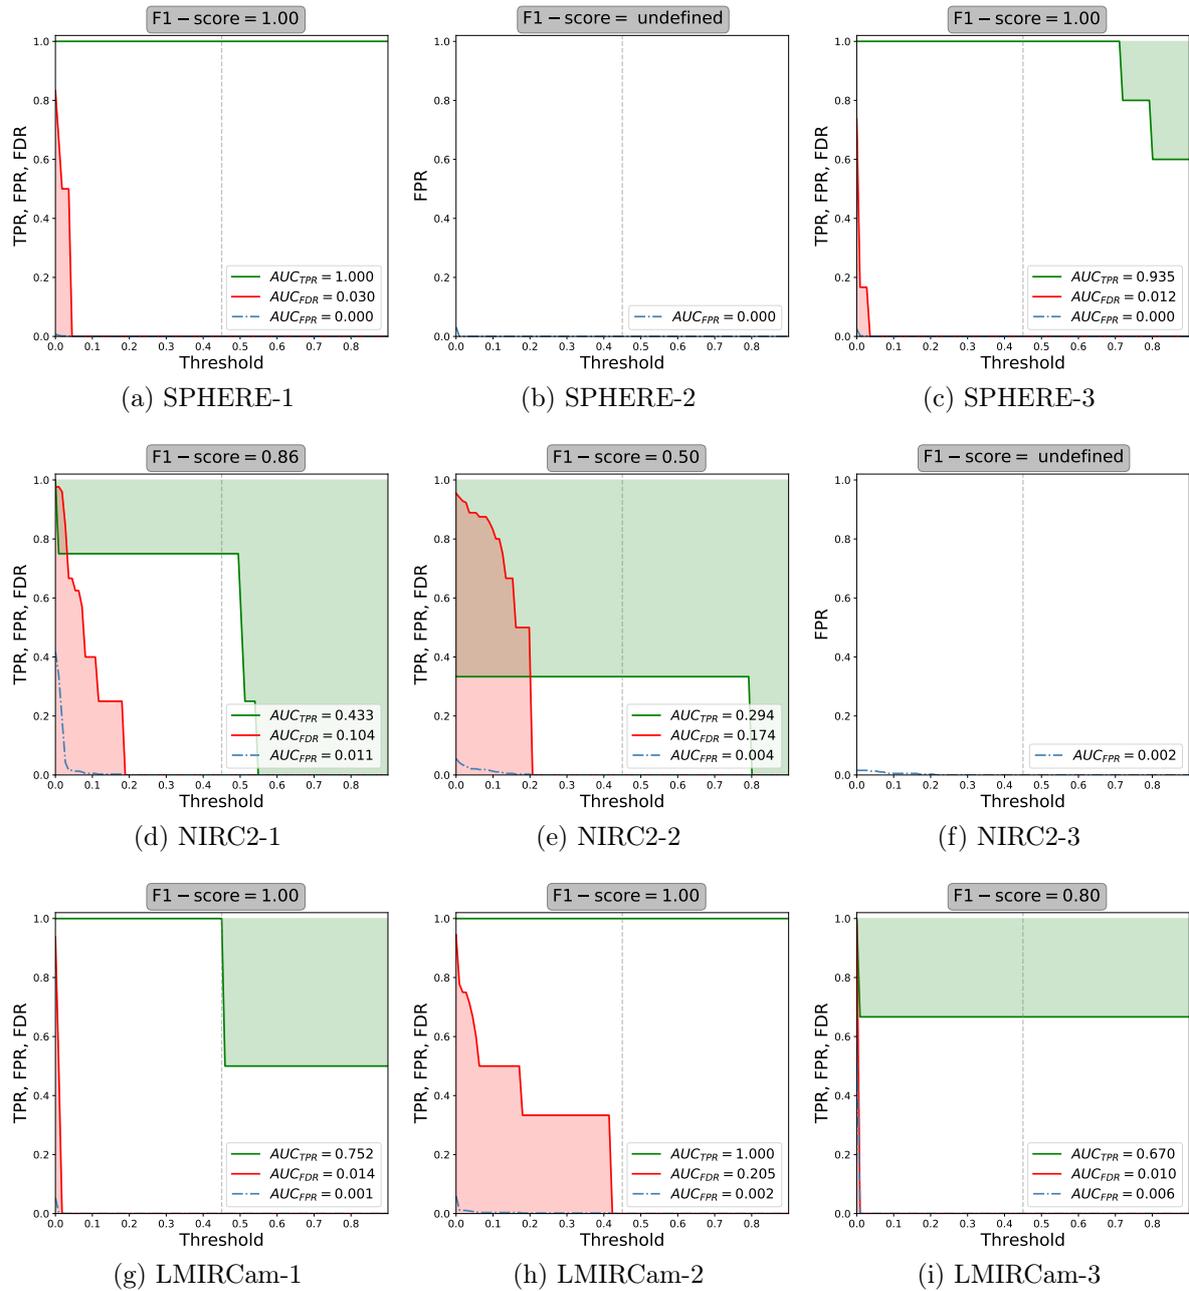


Figure 4.5: True positive rate (green), false discovery rate (red) and false positive rate (dash-dotted blue line) computed for a range of thresholds varying from zero to twice the selected threshold (represented by a dotted vertical line). These curves are computed for the nine data sets of the EIDC, relying on the detection maps estimated with the full-frame version of auto-RSM using the bottom-up approach for the selection of the optimal set of cubes of likelihoods, as well as the original forward approach for the computation of the probabilities. The green line representing the TPR should be as close as possible to 1 for the entire range of thresholds, while the red and dash-dotted blue line representing respectively the FDR and the FPR, should be as close as possible to zero.

Having illustrated the computation of the performance metrics for the different data sets, we now consider aggregated results to compare the performance of the five auto-RSM parametrisations with the baseline and RSM algorithm submission to the EIDC. The different rankings for the four considered performance metrics are shown in Figures 4.6 and 4.7. The light, medium, and dark colours correspond to the three instruments, with the VLT/SPHERE-IRDIS, Keck/NIRC2, and LBT/LMIRCam data sets, respectively. Figures 4.6 and 4.7 highlight the fact that the RSM-based approaches largely outperform the baseline with much higher F1 scores, a much larger AUC of the TPR, and much lower AUCs of the FDR and FPR. Regarding the five considered auto-RSM parametrisations, they all present a smaller F1 score compared to the RSM algorithm parametrised manually, except for the auto-RSM FF_BU_F, which performs slightly better. However, when considering the other performance metrics, the auto-RSM approach seems to perform better in most cases, especially when considering false positives. These results demonstrate the ability of the auto-RSM approach to better cope with residual speckle noise, while maintaining a high detection rate. This is a key element in reducing arbitrariness in the selection of the detection threshold. The selection of a detection threshold is indeed often a complex task, especially when relying on S/N maps, as the noise probability distribution is often non-Gaussian.

Looking in more detail at the five parametrisations of the auto-RSM, we see clearly that the auto-RSM FF_BU_F leads to the best performance metrics in most cases, and should therefore be favoured for detection when using the auto-RSM approach. The results for the annular and hybrid annular full-frame auto-RSM seem to demonstrate that considering the radial evolution of the optimal parameters does not lead to a significant improvement in performance. The slightly degraded performance of the annular mode can be explained by the fact that the auto-RSM optimisation relies on the inverted parallactic angle approach. The noise structure being similar but not equivalent when inverting the parallactic angles, the annular optimisation is more affected by local differences in the noise structure. These local differences in the noise structure prevent the algorithm from improving the overall performance. Considering the longer computation time required for the annular auto-RSM, and its performance, the full-frame version should clearly be preferred.

As regards the difference between the bottom-up and top-down approaches, the better results obtained with the bottom-up approach may be explained by its ability to select the cube of likelihoods in the right order. Indeed, the probability associated with a planetary signal increases along the temporal axis when computing the RSM detection map. This probability increases faster and stays high for longer when selecting first the cube of likelihoods providing the highest probability ratio between injected fake companion peak probability and background residual probabilities. Sorting the cubes of likelihoods in descending order of quality leads to a higher average probab-

ity for the planetary signal, while it should not affect the probability associated with residual noise.

In addition to the automated selection of the optimal parameters, the results obtained with the auto-RSM FF_BU_F show a clear performance improvement compared to the set of RSM detection maps originally submitted to the EIDC (Cantalloube et al. 2020b). We observe an overall reduction of 76% and 33% compared to the RSM submission for the AUC of the TPR and the AUC of the FDR, respectively, as well as an increase of 19% and 2% for the AUC of the TPR and the F1 score, respectively.

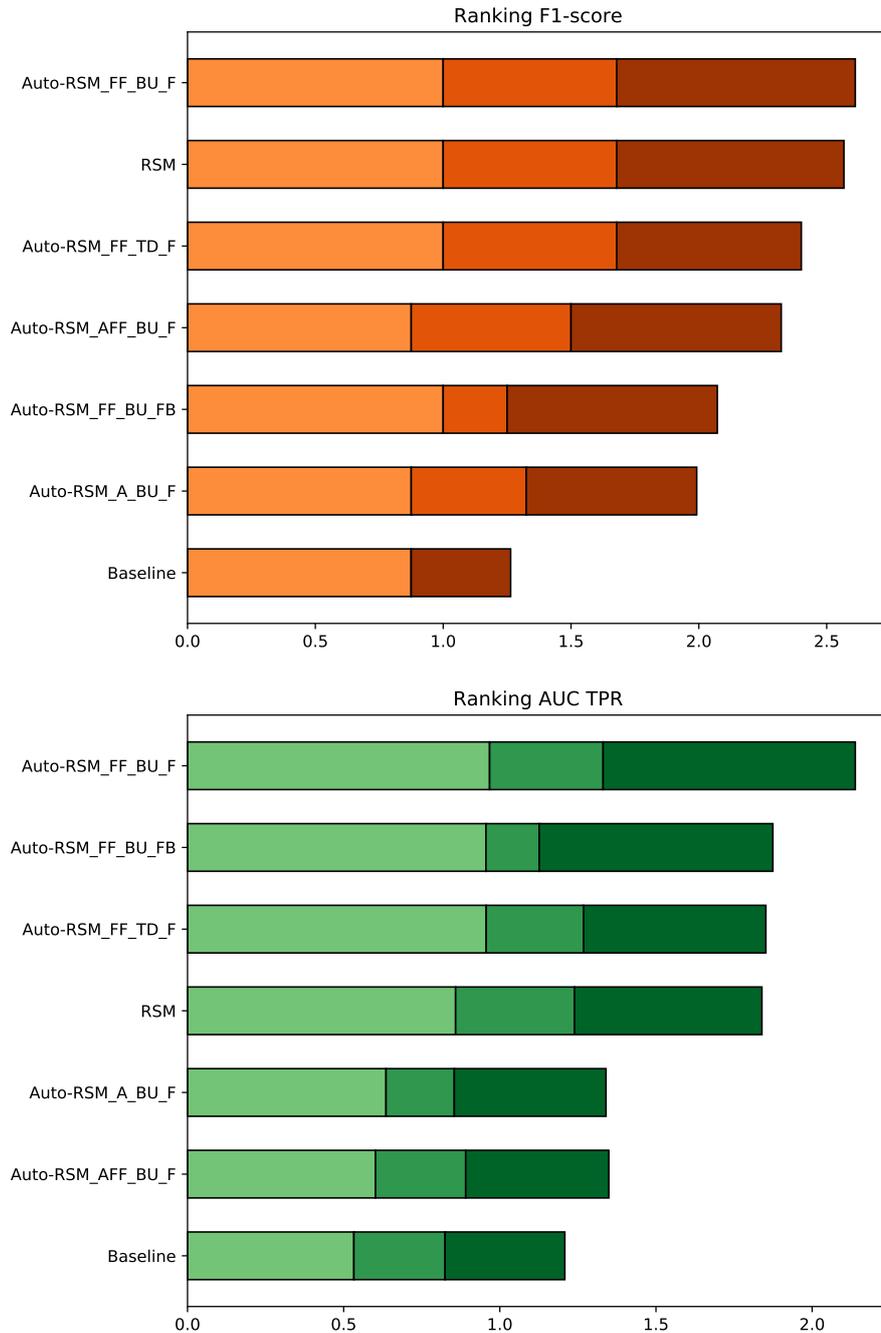


Figure 4.6: Ranking of the different parametrizations of the full-frame and annular versions of the auto-RSM along the original RSM and baseline presented in Cantalloube et al. (2020b). Figure (a) provides the ranking based on the F1 score obtained at the selected threshold. Figure (b) gives the ranking based on the AUC of the TPR. FF stands for full-frame, A for annular, AFF for annular full-frame (annular approach used to optimise the PSF-subtraction parameters and full-frame approach used for the RSM parameter optimisation and the selection of the optimal set of cubes of likelihoods), and BU, TD, F, and FB. The light, medium, and dark colours correspond to VLT/SPHERE-IRDIS, Keck/NIRC2, and LBT/LMIRCam data sets, respectively.

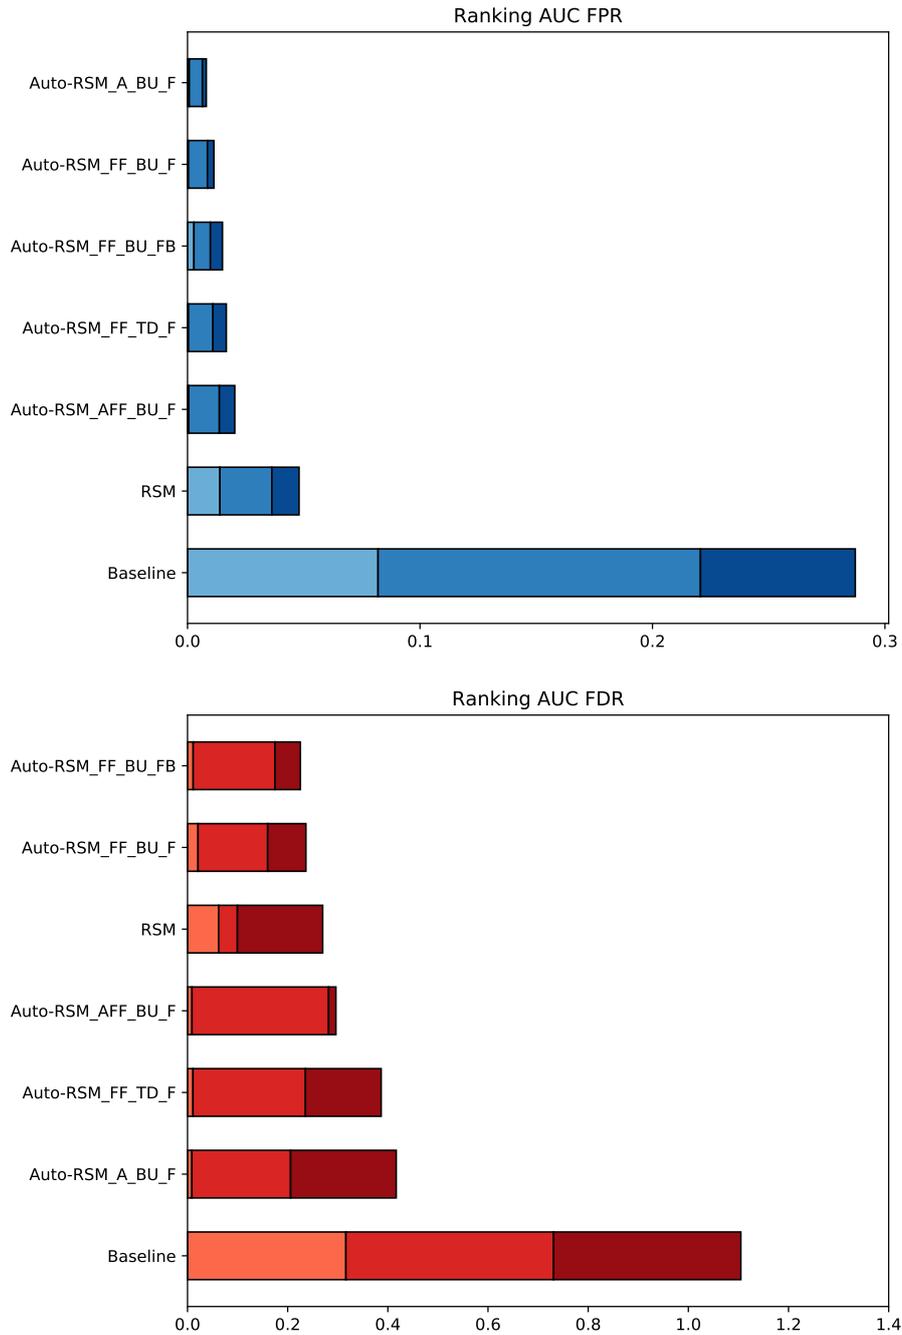


Figure 4.7: Ranking of the different parametrisations of the full-frame and annular versions of the auto-RSM along the original RSM and baseline presented in Cantalloube et al. (2020b). Figure (a) gives the ranking based on the AUC of the FPR, while Figure (b) provides the ranking based on the AUC of the FDR. FF stands for full-frame, A for annular, AFF for annular full-frame (annular approach used to optimise the PSF-subtraction parameters and full-frame approach used for the RSM parameter optimisation and the selection of the optimal set of cubes of likelihoods), and BU, TD, F, and FB. The light, medium, and dark colours correspond to VLT/SPHERE-IRDIS, Keck/NIRC2, and LBT/LMIRCam data sets, respectively.

4.4.4 Commonalities in optimal parametrisations

The proposed auto-RSM optimisation procedure is relatively time consuming even when relying on the full-frame mode, which may potentially preclude its use for very large surveys. In this section, we therefore investigate the possibility of using a smaller set of ADI sequences to generate an optimal parametrisation, which could then be applied to a larger set of ADI sequences. This requires a homogeneity in the sets of optimal parameters selected for different ADI sequences generated by a given HCI instrument, as surveys generally consider multiple observation sequences generated by a single HCI instrument. Sources of heterogeneities in the parametrisation of ADI sequences for a common instrument can originate from the number of frames, the observing conditions, the parallactic angle range, or the target position in the sky.

As the EIDC data set contains multiple ADI sequences generated with different instruments under different observing conditions and with different characteristics (see C.7 for the frames number, FOV rotation), it should allow us to estimate the homogeneity of the parametrisations for a common instrument, and potential heterogeneity between instruments. As the full-frame version of the optimisation algorithm provides the best performance, we rely on the set of parameters generated by this mode to conduct our analysis. We define a heterogeneity metric, which differs for the PSF-subtraction techniques and the RSM algorithm. We use the distance between parametrisations as a metric with which to gauge the performance of the PSF-subtraction techniques. This distance is defined as the difference between the optimal values of all the parameters (see Table C.9) for each pair of ADI sequences. For each parameter, we normalise the absolute value of the distance between the two considered ADI sequences with the mean of the two optimal values used to compute the distance. This allows proper comparison of the relative weight of the different parameters. For the RSM parametrisation, most parameters are non-numerical and we therefore replace the notion of distance by the notion of similarity. For a given pair of ADI sequences, as a metric for each parameter we use the percentage of dissimilarity within the entire set of PSF-subtraction techniques, i.e. the number of PSF-subtraction techniques for which the parameter values are different divided by the total number of PSF-subtraction techniques.

Figure 4.8 shows the cumulative normalised distances for every pair of ADI sequences within each instrument, along with the relative weight of all the parameters. The black line gives the mean distance computed based on the 36 possible pairs of ADI sequences. A cumulative distance below the back line indicates a higher homogeneity of the parameters for the considered pair of ADI sequences. As can be seen from Figure 4.8, the ADI sequences generated by the SPHERE and LMIRCam instruments seem to be characterised by a relatively homogeneous set of parameters, which implies that a common set of parameters could be defined and used for larger surveys. The larger heterogeneity for the NIRC2 samples, which seems to be mainly driven by the

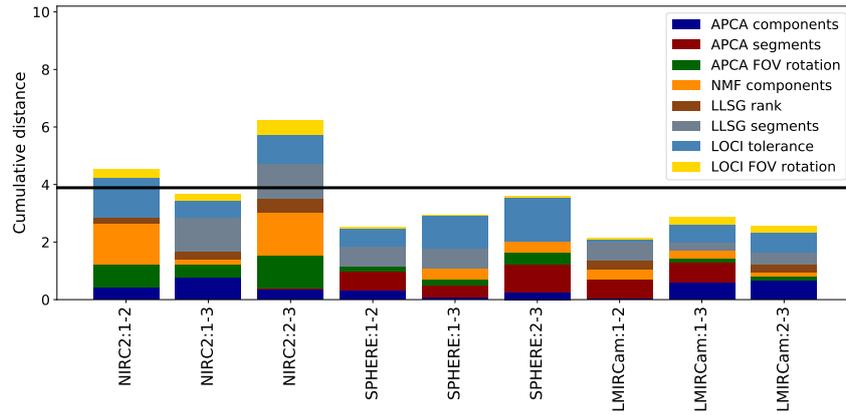


Figure 4.8: Normalised distances between PSF-subtraction-technique parameter sets for nine pairs of ADI sequences. The different coloured bars provide the contribution of the different parameters to the cumulative normalised distance. The considered pairs of ADI sequences are generated by the same instrument. The black horizontal line represents the normalised distances averaged over the 36 possible pairs of ADI sequences.

NIRC2-2 sequence, may be explained by the ADI sequence characteristics (see Table. C.9), or by differences in terms observing conditions.

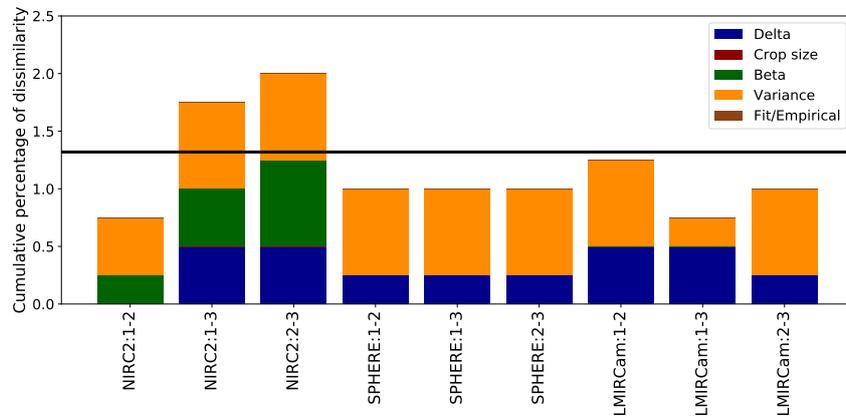


Figure 4.9: Percentage of dissimilarity between RSM parameter sets for nine pairs of ADI sequences. The different coloured bars provide the contribution of the different parameters to the cumulative dissimilarity. The selected pairs of ADI sequences are generated by the same instrument. The black horizontal line represents the percentage of dissimilarity averaged over the 36 possible pairs of ADI sequences.

Considering now the parametrisation of the RSM algorithm, the results presented in Figure 4.9, demonstrate again the larger parametric heterogeneity for the NIRC2 samples, with the NIRC2-3 sequence affecting the dissimilarity measures the most. This confirms the particularity of the NIRC2-3 sequence, which seems to have a different noise structure compared to the other ADI sequences. The crop size of 3 pixels

Table 4.2: Selected PSF-subtraction techniques for the computation of the final RSM-detection map for the nine ADI sequences in the case of the full-frame version of the auto-RSM procedure using the bottom-up approach.

ID/Selected model	APCA	NMF	LLSG	LOCI
SPHERE 1	X	X	X	X
SPHERE 2	X	X	X	X
SPHERE 3	X	X	X	X
NIRC2 1		X		X
NIRC2 2	X			
NIRC2 3		X	X	
LMIRCam 1			X	X
LMIRCam 2	X	X		X
LMIRCam 3	X		X	X

as well as the use of a best-fit approach to estimate the noise properties are common to all ADI sequences (see Table C.9). The heterogeneity is mainly driven by the definition of the region used for computation of the noise properties, which tends to demonstrate the advantage of considering multiple approaches for estimation of the noise properties. We finally consider the set of selected PSF-subtraction techniques used to generate the final RSM map when relying on the full-frame bottom-up auto-RSM. We see from Table 4.2 that the SPHERE ADI sequences share the same set of PSF-subtraction techniques while the set is different for the other ADI sequences.

In addition to the estimation of relative distances and dissimilarity measures, we also applied a K-means clustering algorithm to classify the nine ADI sequences into three clusters based on the set of parameters used for the PSF-subtraction techniques and the RSM algorithm, as well as the likelihood cubes selected for the detection map computation. Using these 32 parameters to characterise each ADI sequence¹⁷, the K-mean algorithm classified NIRC2-1 and NIRC2-3 into the first group, NIRC2-2 into the second group, and the remaining sequences into the third group. As expected, the NIRC2-2 is not in the same group as the other sequences generated with the NIRC2 instrument (see Figure 4.8 and Table 4.2). For the other sequences, we reach the same conclusion as before, apart from the fact that both SPHERE and LMIRCam data sets are regrouped into a single cluster whose centre is close to the SPHERE-1 ADI sequence. Increasing the number of clusters does not lead to clear separation between the SPHERE and LMIRCam instruments, demonstrating the similarity of the ADI sequences generated by both instruments.

We eventually tested the feasibility of using a single set of parameters for a given

¹⁷The categorical parameters such as the RSM parameters and the optimal set of likelihood cubes have been binarised. A standard normalisation, using the parameters mean and variance, has been applied on the data before the clustering algorithm. The K-means clustering algorithm relying on euclidean distance, a proper scaling of the parameters is necessary to avoid favouring some parameter.

instrument, allowing us to investigate the sensitivity of the detection map estimation to the parametrisation. We selected the optimal parametrisation of the SPHERE-1 data set as this parametrisation is the closest to the centre of the SPHERE-LMIRCam cluster, and estimated the detection maps for the SPHERE-2 and SPHERE-3 data sets. Despite the inhomogeneity of the SPHERE data sets in terms of observing conditions and selected wavelengths, the obtained detection maps differed only slightly from the one presented in Figure 4.4, with a reduced probability for one of the five targets in the SPHERE-3 data set but with a similar background noise level. As can be seen from Figure 4.10, the change in terms of AUC of the FPR and FDR is negligible, while the F1 score and the AUC of the TPR reduce a little when considering the 0.45 probability threshold but remain similar if the threshold is adapted.

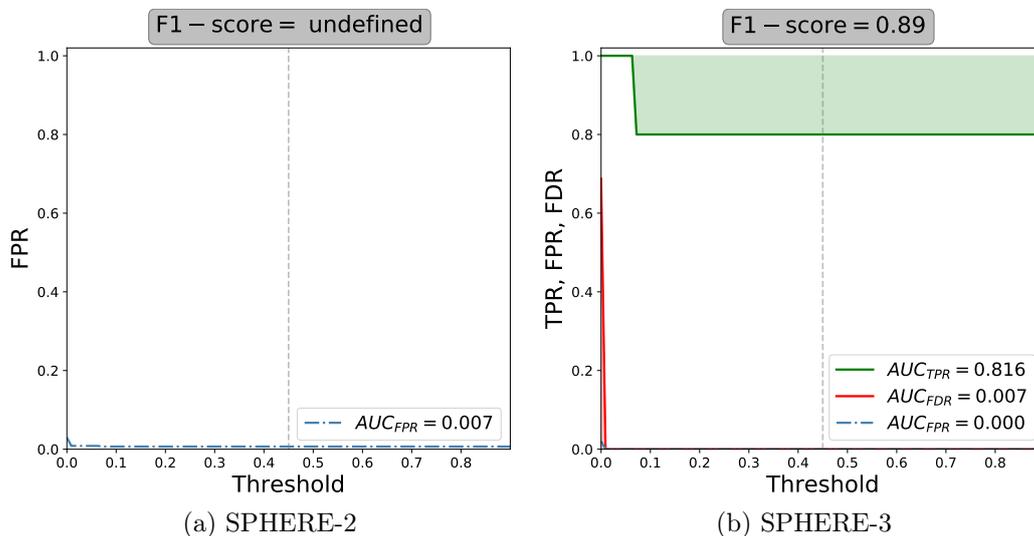


Figure 4.10: True positive rate (green), False discovery rate (in red), and False positive rate (dash-dotted blue line) computed for a range of thresholds varying from zero to twice the selected 0.45 detection threshold (represented by a dotted vertical line). These curves are computed for the SPHERE-2 and SPHERE-3 data sets of the EIDC, relying on the detection maps estimated with the SPHERE-1 optimal set of parameters along with the full-frame version of auto-RSM using the bottom-up approach .

Considering all these results, the use of a reduced number of parametrisations for large surveys seems feasible, especially for the SPHERE and LMIRCam instruments. As demonstrated by the ADI sequences generated by NIRC2, when large dissimilarities are observed in terms of background noise level, a more refined subdivision should be considered. The subdivision of large ADI sequence data sets in subgroups and the definition of representative ADI sequences for the computation of local optimal parametrisations will be investigated in the next chapter.

4.5 Conclusion

In this chapter, we presented a new automated optimisation framework for the RSM approach, called auto-RSM. The proposed automated parameter selection aims to reduce the complexity and sometimes arbitrariness of parameter selection when using HCI post-processing techniques and provide, to users, a simple framework to compute reliable detection maps. Based on a single or multiple ADI sequences, auto-RSM generates, after parameter optimisation, a single detection map with a high contrast between planetary candidates and residual speckles.

The proposed multi-step parameter optimisation framework can be divided into three main steps, (i) the selection of the optimal set of parameters for the considered PSF-subtraction techniques, (ii) the optimisation of the RSM approach parametrisation, and (iii) the selection of the optimal set of PSF-subtraction techniques and ADI sequences to be considered when generating the final detection map. The selection of the optimal set of parameters for the PSF-subtraction techniques is based on the minimisation of the mean contrast within the selected set of annuli, while the optimisation of the RSM approach and selection of the optimal set of cubes of likelihoods are based on the probability ratio between injected fake companions peak probability and background residual probabilities. Some PSF-subtraction techniques having a continuous parameters space, a Bayesian optimisation framework is proposed to explore the parameter space and select the optimal set of parameters. Two different versions of the auto-RSM algorithm are proposed, a full-frame version where a single set of parameters is selected for all angular separations, and an annular version where the set of optimal parameters evolves with radial distance. Different parametrisations of the full-frame and annular auto-RSM are tested to investigate the added value of different methods to select the optimal set of cubes of likelihoods or to compute the final probabilities.

The data sets of the EIDC and the performance assessment framework proposed in Cantalloube et al. (2020b) are used to compare the performance of the different versions and parametrisations of the auto-RSM. The performance assessment is performed via the computation of a data set-dependant F1-score at a predefined threshold, as well as the estimation of the AUC of the TPR, FPR and FDR. The auto-RSM results demonstrate the interest of the approach, as it provides, in most cases, better performance than the original RSM detection map submitted to the EIDC, while the original RSM map was already at, or close to the top of the ranking for all performance metrics in the EIDC. The full-frame auto-RSM using the bottom-up approach to select the optimal set of cubes of likelihood and the forward approach to compute the RSM probabilities provides overall the best performance in terms of detection. Considering the longer computation time and lower performance of the annular version, the full-frame auto-RSM should be preferred.

Auto-RSM being computationally expensive even when using the full-frame version, we investigate the possibility of using a common set of parameters for each instrument. We studied the commonalities existing between the parametrisations of the nine data sets of the EIDC, and found that the distance between the parametrisations for a common instrument is smaller than the distance with the parametrisations of other instruments. Potential differences in the noise characteristics of different data sets generated with a common instrument should nevertheless be taken into account, as illustrated by the NIRC2 data sets. However, the use of a limited number of parametrisations for large surveys seems possible and will be further investigated in the next chapter.

The auto-RSM framework is not limited to the RSM algorithm and the first step of the algorithm may be used separately to optimise the parametrisation of PSF-subtraction techniques and generate S/N maps. A S/N version of the proposed optimisation framework, called auto-S/N was developed. Despite the degraded performance compared to auto-RSM, auto-S/N is characterised by a reduced computation time and can sometimes be a good complement to auto-RSM. All these versions of the proposed optimisation framework are available in a single python package called PyRSM¹⁸. This python package offers a parameter free detection map computation algorithm with a very low level of residual speckles, especially for the auto-RSM, allowing a simple detection threshold selection.

¹⁸<https://github.com/chdahlqvist/RSMmap>

Chapter 5

SHARDDS Survey: Limits on Planet Occurrence Rates Based on Point Sources Analysis via the auto-RSM Framework

5.1 Introduction

Following the development of the auto-RSM framework, we propose in this last chapter to rely on this new post-processing pipeline to perform an in-depth analysis of the SHARDDS survey. The SPHERE High-Angular Resolution Debris Disks Survey (SHARDDS) gathers 55 main-sequence stars within 100 pc, known to host a high-infrared-excess debris disk. Those disks are a natural place to look for exoplanets because planet formation succeeded at least to form large planetesimals in those systems. This is one of the reasons why direct imaging surveys generally include many debris disk host stars, such as in the SPHERE-SHINE survey (Desidera et al. 2021) or the GPI-GPIES survey (Nielsen et al. 2019b). Meshkat et al. (2017) found indeed a tentative evidence that giant planets have a higher occurrence rate in debris disks hosts, and the first emblematic directly imaged planets were found in the massive debris disks system β Pic (Lagrange et al. 2009b) or HR 8799 (Marois et al. 2008). The SHARDDS survey already revealed debris disks resolved for the first time in scattered light: HD 114082 (Wahhaj et al. 2016, Engler et al. in prep.), 49 Ceti Choquet et al. (2017), HD 105 (Marshall et al. 2018) as well as a substellar companion (HD 206893 B) close to the deuterium burning limit (Milli et al. 2016b; Delorme et al. 2017; Romero et al. 2021).

In this context, our study aims to detect and characterise new exoplanets and brown dwarfs within these debris disks, allowing us to potentially better understand

the complex interactions between substellar companions and disks. We rely on the auto-RSM framework to compute detection maps and contrast curves. The auto-RSM framework being computationally expensive, a clustering approach is used to divide the set of targets into multiple subsets. For each subset, the cluster center is identified and the auto-RSM framework is applied onto it to provide the optimal parametrisation for the entire cluster.

Detection maps are then computed via the RSM approach, relying on these optimal parametrisations. The detection maps are used to identify potential planetary companions, and a new companion characterisation framework based on the RSM approach is introduced. This new framework is tested successfully, showing a higher astrometric and photometric precision for faint sources compared to standard approaches. We uncover the companion around HD206893, but do not detect any new companion around other stars.

The detection maps are then used to compute contrast curves. A correlation study between achievable contrasts and parameters characterising high contrast imaging sequences highlights the importance of the Strehl ratio, wind speed at a height of 30 meters and presence of wind-driven halo to define the quality of high contrast images. Planet detection and occurrence rate maps are eventually generated based on these contrast curves and on well-known evolutionary models, namely AMES-DUSTY and AMES-COND. These maps show, for the SHARDDS survey, a high sensitivity between 10 and 100 au for substellar companions with masses $>10M_J$.

The remainder of this chapter is organised as follows. Section 5.2 describes the target selection for the SHARDDS survey. In Section 5.3, we present our data reduction pipeline involving the definition of clusters along with cluster centres on which the auto-RSM optimisation procedure is applied. The computation of detection maps and contrast curves follows the estimation of the optimal parametrisations. Section 5.4 is devoted to the characterisation of potential planetary candidates. In Section 5.5, we consider the contrast curve as a performance metric and analyse the potential drivers of this performance. Section 6 focuses on the estimation of the planetary detection probability from which we derive an estimated planetary occurrence rate associated to the SHARDDS survey. Finally, Section 5.7 concludes this work.

5.2 Survey description

The SHARDDS survey was designed to image circumstellar disks around bright nearby stars (within 100 pc from the Earth) in the near-infrared using the VLT SPHERE instrument. The aim of the survey is to better understand the scarcity of debris disks detection in scattered light, by targeting disks without any scattered-light detection at the time of the survey design (2014), either because the target was not observed with

high-contrast instruments, or because the disk might be too compact and faint to be accessible with first-generation high-contrast instruments such as HST/NICMOS or VLT/NaCo, having poor performance below $0.5''$. The underlying goals are to characterise the disks architecture and properties, and statistically link these properties to the stellar age, spectral type, and potential presence of companions. This chapter contributes to the achievement of these objectives by applying the RSM detection algorithm on the data sets, to detect potential planetary candidates. The RSM detection algorithm was designed to unveil point-like sources and is therefore not fitted to detect extended features such as debris disks. The detection of companions can bring valuable information to better understand the secular interactions between debris disks, and companions and whether such interactions are always needed to explain particular signatures in disks such as azimuthal asymmetries, warps or sharp edges (see Mouillet et al. 1997; Lagrange et al. 2012; Lestrade & Thilliez 2015, for emblematic examples of signatures within debris disks attributable to a companion).

The SHARDDS survey includes 55 main-sequence stars visible from the Southern hemisphere, covering spectral types A-M and ages 10 Myr - 6 Gyr . This diverse sample of debris systems aims to provide a comprehensive view of planetary system properties and their time evolution. These stars were selected for the expected brightness of their disks (fractional luminosity above 10^{-4}) and because they were not yet resolved in scattered light. All stars that were not observable from Paranal with an airmass below 2, were excluded from the sample. The SPHERE-IRDIS instrument was used with the broad-band H filter ($\lambda = 1.625\mu m$, $\Delta\lambda = 0.290\mu m$), as well as an apodised Lyot coronagraph with a radius of 92 mas (N_ALC_YJH_S) to reach a high contrast in the innermost regions. The broad-band H filter was selected for its wide spectral band-pass allowing to collect more disk photons, but also because the performance of the extreme adaptive optics system improves at longer wavelengths and the dust from debris disks typically displays a red colour, while the thermal background is not as high as in the K band and does not dominate the noise budget at large separations. The observations were made in pupil-stabilised mode, using the Angular Differential Imaging observing strategy (ADI, Marois et al. 2008a). The targets were observed around meridian passage with a total execution time of one hour per target, which ensured a sufficient field rotation with about 40 minutes long coronagraphic images. The observations were grouped in two programs¹, 46 sources were imaged during P96 (1 October 2015 - 31 March 2016) and 9 during P97 (1 April 2016 - 30 September 2016). Due to adverse observing conditions, multiple observation sessions were required for some targets, leading to an actual data set of 73 ADI sequences. Tables 5.1-5.3 provides details on the set of targets, including the number of observation sequences acquired for each target (epoch). The distances, magnitudes and spectral types were taken from

¹Based on observations collected at the European Southern Observatory under ESO programmes 096.C-0388(A) and 097.C-0394(A)

the Hipparcos and GAIA catalogues (van Leeuwen 2007; Gaia Collaboration et al. 2016). The target Fomalhaut C, part of the SHARDDS sample, was excluded from our analysis as the observing conditions were very bad for all three epochs (data set published in Cronin-Coltsmann et al. 2021).

Table 5.1: Age, distance, spectral-type and magnitude distributions along with the number of ADI sequences for each SHARDDS target. For the definition of the star age multiple papers have been used: ¹ (Zuckerman & Song 2004), ² (Zuckerman et al. 2001), ³ (Chen et al. 2014), ⁴ (Moór et al. 2011), ⁶ (Rodriguez & Zuckerman 2012), ⁷ (Fernández et al. 2008), ⁸ (Rhee et al. 2007), ⁹ (Malo et al. 2013), ¹⁰ (Casagrande, L. et al. 2011), ¹¹ (Torres et al. 2008), ¹² (Metchev & Hillenbrand 2009), ¹³ (Zuckerman et al. 2011), ¹⁴ (da Silva, L. et al. 2009), ¹⁵ (Vican 2012), ¹⁶ (Moor et al. 2006), ¹⁷ (King et al. 2003), ¹⁸ (Taberner et al. 2012), ¹⁹ (Zorec & Royer 2012), ²⁰ (Mamajek et al. 2002), ²¹ (Mamajek & Hillenbrand 2008), ²² (Delgado Mena, E. et al. 2014), ²³ (van Leeuwen 2007), ²⁴ (Kim et al. 2005), ²⁵ (Matthews et al. 2018), ²⁶ (Mamajek & Bell 2014), ²⁷ (West et al. 2008), ²⁸ (Smith et al. 2008), ²⁹ (Meshkat et al. 2017), ³⁰ (Gullikson et al. 2016), ³¹ (Delorme et al. 2017).

Name	RA	DEC	V mag	H mag	Sp. type	Age (My)	Distance (pc)	# Epochs
HD 105	00:05:53	-41:45:11	7.53	6.19	G0V	30 ¹	38.85	1
HD 203	00:06:50	-23:06:27	6.17	5.33	F3V	23 ²	39.97	1
HD 377	00:08:26	+06:37:00	7.59	6.15	G2V	170 ³	38.52	1
HD 3003	00:32:44	-63:01:53	5.09	5.16	A0V	30 ¹	45.89	1
HD 3670	00:38:57	-52:32:03	8.21	7.15	F5V	30 ⁴	77.58	1
HD 9672	01:34:38	-15:40:34	5.61	5.53	A1V	40 ⁶	57.08	1
HD 10472	01:40:24	-60:59:56	7.61	6.69	F2IV/V	30 ⁷	71.17	2
HD 10638	01:44:23	+32:30:57	6.73	6.19	A3	100 ⁸	68.68	1
HD 13246	02:07:26	-59:40:45	7.50	6.30	F7V	40 ⁹	45.60	1
HD 14082B	02:17:25	+28:44:30	7.74	6.36	G2V	21 ⁹	39.75	1
AG-Tri	02:27:29	+30:58:24	10.12	7.24	K8	23 ¹	41.05	4
HD 15257	02:28:10	+29:40:09	5.29	4.82	F0III	1000 ⁸	49.93	1
HD 16743	02:39:08	-52:56:05	6.77	5.97	F1III/IV	200 ⁸	57.94	1
HD 17390	02:46:45	-21:38:22	6.47	5.63	F3IV/V	610 ¹⁰	48.19	1
HD 21997	03:31:54	-25:36:50	6.37	6.12	A3IV/V	30 ¹¹	69.64	1
HD 22179	03:35:30	+31:13:37	8.93	7.49	G5IV	63 ¹²	70.37	1
HD 24636	03:48:11	-74:41:38	7.13	6.22	F3IV/V	30 ¹³	57.05	1
HD 25457	04:02:37	-00:16:08	5.38	4.34	F6V	70 ¹³	18.77	1
HD 31392	04:54:04	-35:24:16	7.61	5.89	G9V	3690 ¹⁰	25.77	1
HD 35650	05:24:30	-38:58:10	9.05	6.11	K6V	70 ¹	17.48	1
HD 274255	05:30:14	-42:41:50	9.71	6.47	M0V	1000 ²⁹	19.15	1

Table 5.2: Age, distance, spectral-type and magnitude distributions along with the number of ADI sequences for each SHARDDS target. For the definition of the star age multiple papers have been used: ¹ (Zuckerman & Song 2004), ² (Zuckerman et al. 2001), ³ (Chen et al. 2014), ⁴ (Moór et al. 2011), ⁶ (Rodriguez & Zuckerman 2012), ⁷ (Fernández et al. 2008), ⁸ (Rhee et al. 2007), ⁹ (Malo et al. 2013), ¹⁰ (Casagrande, L. et al. 2011), ¹¹ (Torres et al. 2008), ¹² (Metchev & Hillenbrand 2009), ¹³ (Zuckerman et al. 2011), ¹⁴ (da Silva, L. et al. 2009), ¹⁵ (Vican 2012), ¹⁶ (Moor et al. 2006), ¹⁷ (King et al. 2003), ¹⁸ (Taberner et al. 2012), ¹⁹ (Zorec & Royer 2012), ²⁰ (Mamajek et al. 2002), ²¹ (Mamajek & Hillenbrand 2008), ²² (Delgado Mena, E. et al. 2014), ²³ (van Leeuwen 2007), ²⁴ (Kim et al. 2005), ²⁵ (Matthews et al. 2018), ²⁶ (Mamajek & Bell 2014), ²⁷ (West et al. 2008), ²⁸ (Smith et al. 2008), ²⁹ (Meshkat et al. 2017), ³⁰ (Gullikson et al. 2016), ³¹ (Delorme et al. 2017).

Name	RA	DEC	V mag	H mag	Sp. type	Age (My)	Distance (pc)	# Epochs
HD 37484	05:37:40	-28:37:34	7.25	6.29	F3V	30 ¹⁴	59.10	2
HD 38207	05:43:21	-20:11:21	8.47	7.55	F2V	534 ¹⁵	110.99	1
HD 38206	05:43:22	-18:33:26	5.73	5.84	A0V	30 ¹⁴	71.43	2
HD 40540	05:57:53	-34:28:34	7.54	6.93	A8IV	170 ³	88.26	1
HD 53842	06:46:14	-83:59:29	8.62	6.40	F5V	30 ¹⁶	57.87	1
HD 60491	07:34:26	-06:53:48	8.14	6.14	K2V	500 ¹⁷	23.51	1
HD 69830	08:18:24	-12:37:55	5.95	4.36	G8V	5670 ¹⁵	12.56	1
HD 71722	08:26:25	-52:48:26	6.04	5.91	A0V	324 ³	69.35	1
HD 73350	08:37:50	-06:48:24	6.73	5.32	G5V	600 ¹⁸	24.34	1
HD 76582	08:57:35	+15:34:52	5.68	5.21	F0IV	538 ¹⁹	48.80	1
HD 80950	09:17:28	-74:44:04	5.86	5.92	A0V	138 ¹⁹	77.34	1
HD 82943	09:34:51	-12:07:46	6.53	5.25	F9V	430 ³	27.61	4
HD 84075	09:36:18	-78:20:41	8.59	7.24	G2V	40 ¹³	64.10	1
HD 107649	12:22:25	-51:01:34	8.78	7.76	F5V	17 ²⁰	108.34	1
HIP63942	13:06:15	+20:43:45	9.40	6.21	K5	4500 ²⁷	18.80	1
HD 114082	13:09:16	-60:18:30	8.21	7.23	F3V	17 ²⁰	95.69	1
HD 120534	13:50:40	-31:12:23	7.02	6.33	A5V	320 ¹⁶	86.81	3
HD 122652	14:02:32	+31:39:39	7.15	5.94	F8	500 ⁸	39.54	2
HD 133803	15:07:15	-29:30:16	8.12	7.36	A9V	16 ²⁰	110.74	2
HD 135599	15:15:59	+00:47:46	6.91	5.12	K0V	1300 ²¹	15.82	2
HD 138965	15:40:11	-70:13:40	6.42	6.34	A1V	348 ²⁵	78.08	1

Table 5.3: Age, distance, spectral-type and magnitude distributions along with the number of ADI sequences for each SHARDDS target. For the definition of the star age multiple papers have been used: ¹ (Zuckerman & Song 2004), ² (Zuckerman et al. 2001), ³ (Chen et al. 2014), ⁴ (Moór et al. 2011), ⁶ (Rodriguez & Zuckerman 2012), ⁷ (Fernández et al. 2008), ⁸ (Rhee et al. 2007), ⁹ (Malo et al. 2013), ¹⁰ (Casagrande, L. et al. 2011), ¹¹ (Torres et al. 2008), ¹² (Metchev & Hillenbrand 2009), ¹³ (Zuckerman et al. 2011), ¹⁴ (da Silva, L. et al. 2009), ¹⁵ (Vican 2012), ¹⁶ (Moor et al. 2006), ¹⁷ (King et al. 2003), ¹⁸ (Taberner et al. 2012), ¹⁹ (Zorec & Royer 2012), ²⁰ (Mamajek et al. 2002), ²¹ (Mamajek & Hillenbrand 2008), ²² (Delgado Mena, E. et al. 2014), ²³ (van Leeuwen 2007), ²⁴ (Kim et al. 2005), ²⁵ (Matthews et al. 2018), ²⁶ (Mamajek & Bell 2014), ²⁷ (West et al. 2008), ²⁸ (Smith et al. 2008), ²⁹ (Meshkat et al. 2017), ³⁰ (Gullikson et al. 2016), ³¹ (Delorme et al. 2017).

Name	RA	DEC	V mag	H mag	Sp. type	Age (My)	Distance (pc)	# Epochs
HD 145229	16:09:26	+11:34:28	7.44	6.06	G0	650 ²⁴	33.74	1
HD 157728	17:24:06	+22:57:37	5.72	5.22	A7V	100 ¹⁶	42.74	1
HD 164249A	18:03:03	-51:38:56	7.01	6.02	F6V	1800 ²³	49.60	1
HD 172555	18:45:26	-64:52:16	4.77	4.25	A7V	20 ²⁶	28.79	1
HD 181296	19:22:51	-54:25:26	5.02	5.15	A0V	12 ²⁸	47.37	1
HD 182681	19:26:56	-29:44:35	5.64	5.66	B8.5V	107 ³⁰	71.42	1
HD 192758	08:18:16	-42:51:36	7.03	6.30	A5V	45 ¹⁶	66.53	2
HD 201219	21:07:56	+07:25:58	0	46.5	G5	5370 ¹⁰	37.89	1
HD 205674	21:37:21	-18:26:28	7.17	6.25	F4IV	850 ¹⁰	56.40	2
HD 206893	21:45:22	-12:47:00	6.67	5.69	F5V	250 ³¹	40.80	1
HD 218340	23:08:12	-63:37:41	8.44	7.07	G3V	2050 ²²	56.18	1
HD221853	23:35:36	+08:22:57	7.34	6.44	F0	20 ¹⁶	65.40	1

5.3 Data reduction

5.3.1 Pre-processing and extraction of environmental data

The first reduction steps consist in applying standard calibrations to the raw IRDIS images (sky subtraction, flat-field correction, and bad-pixel correction), and registering the frames. This was done using a dedicated pipeline in python ². The frame registration was done using the four satellite spots imprinted on the IRDIS images by a specific waffle pattern applied on the deformable mirror of SPHERE. The output of the pre-processing consists of a temporal cube of frames (individual detector integrations), cosmetically cleaned and recentered, called hereafter an ADI sequence. This cube is accompanied by the corresponding list of parallactic angles for the dedicated high-contrast image processing steps (see section 5.3.2).

For the clustering of data and to guide the interpretation, we also extracted environmental data from either the adaptive optics telemetry³ or the Astronomical Site Monitor (ASM) of the Cerro Paranal Observatory⁴. We collected, among other, data on the seeing, coherence time, relative humidity, temperature, wind speed and direction at various heights above the platform, Strehl ratio, precipitable water vapour.

5.3.2 Image processing

The resulting corrected sets of ADI sequences have been cropped to a 199×199 pixels size, corresponding to the innermost region of the field of view (FOV). We consider angular separations below 1.25 arcsec to take advantage of the higher sensitivity of the RSM map algorithm in the region near the host star, while limiting the computation time. Indeed the increased performance of the RSM map algorithm compared to other PSF-subtraction techniques reduces above 1 arcsec, which makes it less suitable for larger angular distances when considering its high computational cost.

The computation time is also reduced by limiting the size of the ADI sequences to a maximum of 300 frames. We relied on image binning to reduce the size of the ADI sequence. The images were binned by pair (resp. by three) for sequences containing between 300 and 600 images (resp. between 600 and 900 images). The binning procedure started by the derotation of the images based on the parallactic angles. Then the pixel-wise average was computed along the temporal axis with a common step size and window size of two for sequences containing between 300 and 600 images and three for sequences containing between 600 and 900 images. The same procedure was applied

²available at https://github.com/jmilou/sphere_pipeline.git

³The SPHERE real time controller called SPARTA stores a summary of the adaptive optics telemetry during each observation. Those files are available on the ESO archive as described in Milli et al. (2017). We developed an automatic script to query and analyse the SPARTA and ASM data available at <https://github.com/jmilou/sparta.git>

⁴<http://archive.eso.org/cms/eso-data/ambient-conditions.html>

on the parallactic angles. The resulting set of images were then rotated back to their original orientation. The noise content of these ADI sequences should be reduced by the binning procedure via partial time-averaging.

5.3.3 Clustering

In order to take full advantage of the RSM algorithm, we rely on the auto-RSM optimisation framework to define the optimal sets of parameters for the PSF-subtraction techniques and the RSM algorithm itself. This optimisation pipeline being computationally expensive, we propose to apply it on a subset of targets representative of the whole data set. The obtained optimal parametrisations can then be used to compute the RSM detection maps for all targets. The analysis of Section 4.4.4 showed a relatively high degree of similarity in the optimal parametrisations of both the PSF-subtraction techniques and the RSM algorithm, when using ADI sequences generated with the VLT SPHERE instrument. Dividing the SHARDDS data set into multiple subsets should nevertheless allow to account for small variations in the optimal parametrisations depending on the ADI sequence characteristics.

The subdivision of the SHARDDS data set in multiple subsets is based on a set of observables characterising the ADI sequences. The subdivision itself is done via the K-means algorithm, a centroid-based clustering procedure aiming to find the centroids that minimise the within-cluster sum-of-squares, also called inertia:

$$\sum_{i=0}^n \min_{\mathbf{u}_i \in C} (\|\mathbf{x}_i - \mathbf{u}_j\|^2), \quad (5.1)$$

where \mathbf{u}_i is the centroid of the cluster j and \mathbf{x}_i a member of cluster j . The inertia can be seen as a measure of the internal coherence of the clusters. The K-means algorithm consists of three main steps. It starts with an initialisation step during which initial centroids are randomly drawn from the data set. Each sample is then assigned to the nearest centroid and the centroids positions are updated. The algorithm loops between these last two steps until convergence.

The K-means algorithm was selected as it provides a good estimate of the centroids positions. This is a key element to define properly which ADI sequence within a cluster is the most representative. These centroids being often not associated to a sample, we define the most representative ADI sequences as the ones closest to the cluster centroids. Once defined, the auto-RSM optimisation framework is applied on the selected set of ADI sequences. The optimal parametrisations are then used to compute the RSM detection maps for the remaining ADI sequences of each cluster, following the standard RSM map procedure.

Table 5.4: Set of performance indicators selected for the subdivision of the SHARDDS Survey data set via clustering.

Mean seeing	Strehl ratio	Coherence time
Number of images	Field rotation	Raw contrast
Autocorrelation	Mean wind speed	WDH

Clustering parameters

The K-means algorithm needs to be applied on a set of parameters that characterise the properties of the ADI sequences. For our cluster analysis, we chose metrics providing information about the sequence, the observing conditions, and the noise distribution within the set of frames. This set of observables consists in the mean seeing, the Strehl ratio, the mean coherence time, the number of images, the total field rotation in term of parallactic angle, the raw contrast at 200, 500 and 700 mas, the autocorrelation timescale between images, the mean wind speed at 30 meters, and the wind driven halo strength and asymmetry (Cantalloube et al. 2020a).

The seeing, Strehl ratio, and coherence time are commonly used performance indicators to assess the observing conditions. Considering the 40 minutes integration time used for the SHARDDS survey, the number of images contained in the ADI sequence affects the sampling frequency, and therefore both the performance and the parametrisation of the PSF-subtraction algorithm (e.g. optimal number of principal components). The field rotation also impacts the performance because of the higher self-subtraction of the signal in the case of small field rotation. When mitigating self-subtraction, it translates into a reduced set of available images to compute the reference PSF.

The raw contrasts were estimated by placing apertures of 1 FWHM diameter in the selected annuli and computing the ratio between the mean encircled flux and the star flux. The autocorrelation timescale between the ADI sequence images was estimated by considering the region between 300-600 mas, where the adaptive optics is affecting the most the performance. The flux within a one FWHM aperture was computed for each pixel in the selected region and for each image. An exponential function was then fitted on the temporal autocorrelation of these fluxes and its exponential factor was kept as a measure of the autocorrelation decay rate. We expect that a slower autocorrelation decay will result in lower performance.

The wind-driven halo (WDH) strength and asymmetry were computed using the method presented in Cantalloube et al. (2020a). The WDH is a feature observable when atmospheric turbulence, mainly in the jet stream layer, vary faster than the adaptive optics can compensate. The WDH appears as a bright elongated structure centred on the coronagraph in post-processed high contrast images. The WDH is characterised by a direction, a strength and an asymmetry of the lobes. The WDH cannot

be easily treated with standard PSF-subtraction techniques and affects therefore the achievable contrast at small angular separations (below 1000 mas). Along with the WDH, the low wind effect (LWE, Milli et al. 2018) is also a wind-driven phenomenon degrading the performance of high contrast imaging. LWE arises from uncorrected wavefront aberrations due to air temperature inhomogeneities in large telescope pupil, caused by the radiative cooling of the spiders, which dominates in the absence of wind. We included the wind speed at 30 m to account for this potential effect.

The number of images included in the ADI sequences was identified as a key metrics for the definition of the optimal parametrisation during the development of the auto-RSM framework. We have therefore decided to divide the SHARDDS data set into two subsets before applying the clustering algorithm. We defined a threshold of 151 frames to separate the two subsets, as this value ensures that the standard deviation of the number of images within each subset is equivalent. This ensures a similar distribution in terms of sequence size within the two subsets.

Application and results of the K-means clustering

The K-means algorithm being based on Euclidean distance, the selected set of metrics must be standardised before applying the clustering algorithm, to avoid that metrics with larger values dominate the calculation. Before applying the K-means algorithm, we looked for possible multicollinearity between the selected set of observables. Relying on the variance inflation factor and Pearson correlations, we removed the contrast at 200 and 700 mas, which led to multicollinearity, affecting potentially the definition of the clusters. The initialisation of the K-means algorithm consisting in the random selection of initial centroids, the results may lack consistency and differ from one estimation to another. The algorithm can also be affected by the order of the observables. In order to tackle these two issues, we initialised our estimation by running the K-means algorithm 100 times, selecting at each iteration a different permutation of the parameters. We then took the mean of these centroids positions to initialise the final cluster definition.

We finally defined the number of clusters. This definition was based on the analysis of the evolution of the total squared distance between cluster members and their centroid when changing the number of clusters. Looking at Figure 5.1, we see that the largest fraction of the total squared distance reduction occurs between 1 and 4 clusters. We therefore selected for both subsets a number of clusters equal to 4, implying a total of 8 ADI sequences on which auto-RSM will be applied. The 8 cluster centroids, as well as the composition of their respective clusters are presented in Table 5.5.

After the subdivision of the data set into 8 clusters, we made several consistency checks by relying on principal component analysis to reduce the dimensionality of our set of observables and eliminate residual correlations between the variables. We tested

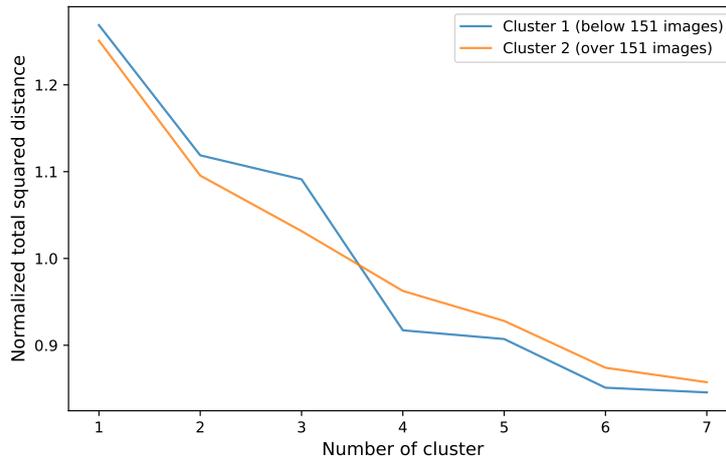


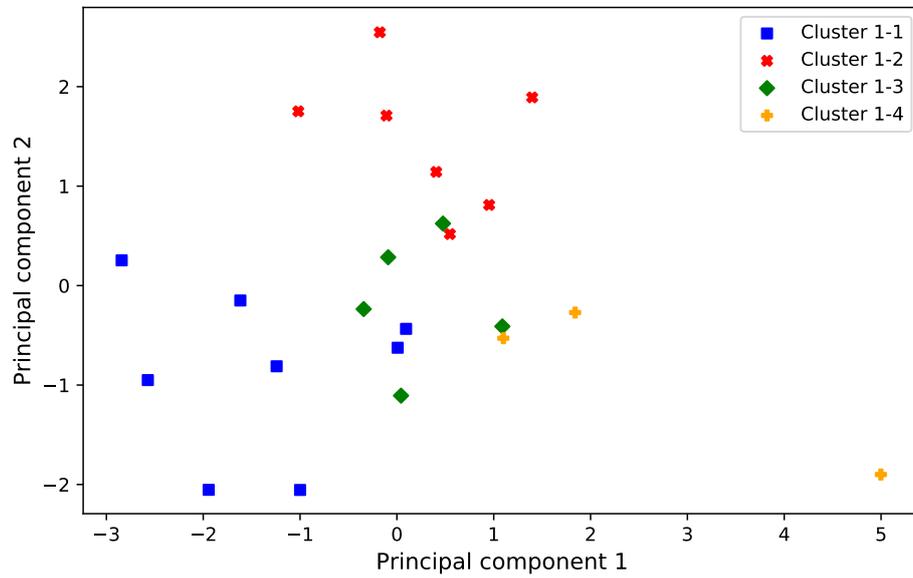
Figure 5.1: Evolution of the standardised total square distance between every cluster member and their centroid, depending on the number of cluster for the two subsets (i.e. ADI sequences with a number of images higher or lower than 151 frames)

the K-means algorithm with different numbers of principal components and retrieved almost every time the same set of clusters. Figure 5.2 illustrates the repartition between the different clusters in the space formed by the first two principal components. As can be seen, the different clusters are relatively well defined except for cluster 2-2 and 2-4, for which a larger set of principal components are necessary to make a clear distinction. We finally applied a Gaussian mixture model instead of the K-means algorithm as a last consistency check. The Gaussian mixture model considers on top of the number of clusters and the centroid position, the standard deviation of the distance between cluster members to characterise clusters. The obtained cluster repartitions were very close although not exactly the same.

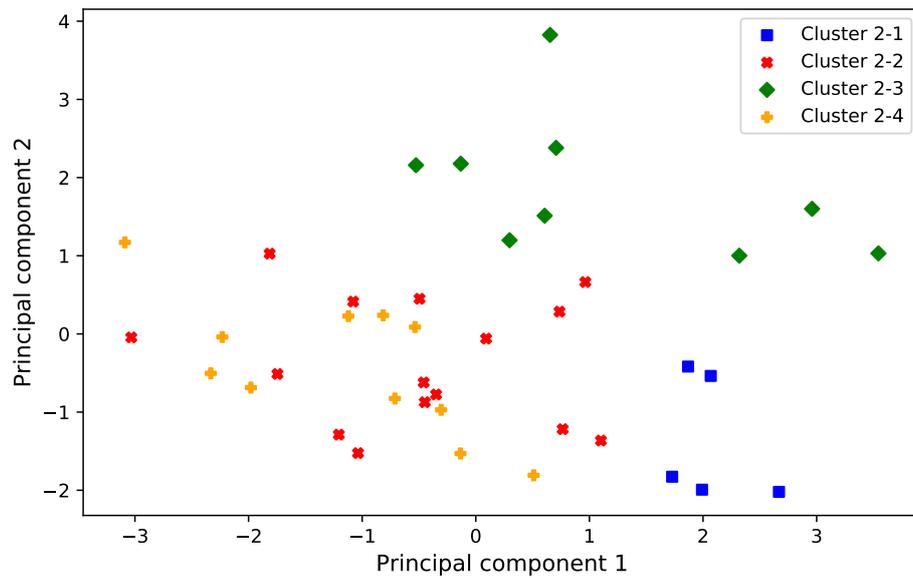
Two targets were excluded from these clusters, HD 133803 and HD 205674. They were treated separately as they were imaged at two epochs separated by only a couple of days. We therefore took advantage of the ability of the RSM algorithm to deal with multiple ADI sequences at once to generate a single detection map per target. This was not possible for the other multi-epoch targets due to the longer time span separating the image sequences, implying a potential movement of planetary candidates.

5.3.4 High contrast image processing

This section is devoted to the computation of RSM detection maps for all the targets included in the SHARDDS survey, as well as the computation of the contrast curves. This computation starts with the optimisation of the model parameters via the auto-RSM framework for the 8 selected targets (see cluster center in Table 5.5). The auto-RSM framework requires the selection of the PSF-subtraction techniques as well as the definition of the parameter ranges to be considered during the optimisa-



(a)



(b)

Figure 5.2: Projection of the SHARDDS survey targets on the first two principal components computed based on their observational characteristics. The top (respectively bottom) graph provides the targets with a number of frames in their ADI sequence below 151 (respectively above 151). The colours indicates to which cluster the target has been assigned.

Table 5.5: Subdivision of the SHARDDS data set into 8 clusters.

Cluster center	Cluster members
Cluster 1-1	
HD 192758	HD 38207, HD 37484, HD 10472, AG Tri, HD 84075, HD 192758 2 nd epoch, HD 274255
Cluster 1-2	
HD 3670	HD 37484 2 nd epoch, HD 22179, AG Tri 2 nd epoch, AG Tri 3 rd epoch, HD 82943 3 rd epoch, HD 114082
Cluster 1-3	
HD 201219	HD 53842, AG Tri 4 th epoch, HD 218340, HD 221853
Cluster 1-4	
HD 14082B	HD 82943, HD 107649
Cluster 2-1	
HD 21997	HD 24636, HD 15257, HD 10472 2 nd epoch, HD 145229 2 nd epoch
Cluster 2-2	
HD 206893	HD 40540, HD 35650, HD 31392, HD 25457, HD 17390, HD 16743, HD 9672, HD 105, HD 69830, HD 71722, HD 120534, HD 182681, HD 120534 2 nd epoch, HD 164249A
Cluster 2-3	
HD 181296	HD 14082B 2 nd epoch, HD 13246, HD 203, HD 60491, HD 122652, HD 135599 3 rd epoch, HD 145229, HD 172555, HD 181296
Cluster 2-4	
HD 3003	HD 377, HD 73350, HD 76582, HD 80950, HD 82943 2 nd epoch, HD 82943 4 th epoch, HD 138965, HD 157728, HIP63942, HD 122652 2 nd epoch

tion. We considered in this chapter all the PSF-subtraction techniques that the RSM algorithm can accommodate: annular PCA (APCA, Gomez Gonzalez et al. 2017), non-negative matrix factorisation (NMF, Ren et al. 2018), the local low rank plus sparse plus Gaussian decomposition (LLSG, Gonzalez et al. 2016), locally optimised combination of images (LOCI, Lafreniere et al. 2007b), and forward-model versions of KLIP and LOCI (see chapter 3). In order to further improve the RSM map algorithm performance, we considered for APCA, NMF, LLSG and KLIP two different ranges for the number of principal components/ranks, which are considered as separate models by the algorithm. Planetary signals and residual speckle noise evolve differently with the number of principal components used to generate the reference PSF. We take advantage of this differentiated evolution to further boost the ability of the RSM map

algorithm to disentangle faint planetary signals from residuals speckle noise, following to some extent the approach proposed in Gomez Gonzalez et al. (2018).

The considered ranges of principal components for APCA, NMF and KLIP, the ranks for LLSG, and the tolerance for LOCI are selected by a new function of the PyRSM python package which regroups the different functions of the auto-RSM framework. This function studies the evolution of the contrast at different angular separations when modifying the number of principal components/ranks/tolerance. The upper boundary of the considered ranges is defined as the value for which the contrast, averaged over the different angular separations, reaches a peak. We performed this analysis for all cluster centroids, with for APCA, NMF, LLSG and KLIP the final ranges divided in two equal size ranges, to form two separate models. For HD 201219, the range of considered ranks being limited to $[1, 5]$, a single LLSG model was considered with range $[1, 5]$. Regarding the other parameters of the PSF-subtraction techniques, a single range was defined for all cluster centroids. The range for the number of segments was fixed to $[1, 4]$, the FOV rotation threshold to $[0.25, 1]$ and the crop size to $[3, 5]$ for standard PSF-subtraction techniques and $[7, 9]$ for the forward model versions to account for the side lobes due to self-subtraction (see chapter 2-4 for more information about these parameters).

The computation of the PSF forward model being computationally very intensive and side lobes due to self-subtraction becoming fainter for increasing angular separation, we considered the forward model versions for only the first 400 mas. Indeed the main interest of PSF forward modelling is to account for the signature of over- and self-subtraction of the PSF due to the reference PSF subtraction. This effect becoming weaker when the angular separation increases, due to larger field rotation, it is not necessary to use such advanced model at large distance from the host star. Having defined all the parameters, the auto-RSM optimisation framework was applied on each centroid, using the full frame mode to optimise the PSF-subtraction techniques and RSM algorithm parameters, the forward model to compute the RSM detection maps, and the bottom-up approach to select the optimal set of likelihoods. Following the original auto-RSM framework, the parameters optimisation was performed using the reversed parallactic angles. Considering the low probability of detecting a planet, we also tried to use the original parallactic angle to optimise the parameters, but it did not lead to a performance increase in terms of contrast. We therefore relied on reversed parallactic angles to avoid any potential planetary signal suppression during the optimisation process.

We investigate in Appendix D.1 the similarities existing between the optimal parametrisations obtained for the 8 cluster centroids, as well as the relationships between these optimal parameters and the set of metrics characterising the ADI sequences. The comparison of the optimal parametrisations is done via the computation of dissimilarity measures between cluster centroids, for both the PSF-subtraction tech-

niques and the RSM algorithm. The results demonstrate a relatively high degree of similarity between the different parametrisations, confirming the conclusions drawn in chapter 4 about the high stability of the ADI sequence imaged by the VLT/SPHERE instrument. The Pearson correlations between the 10 observables characterising our ADI sequences, and the PSF-subtraction techniques parameters show a sensible correlation for some observables, with the contrast at 500 mas showing the highest average correlation rate, and the exponent of the autocorrelation function the lowest one.

Detection maps

Following the definition of the optimal set of parameters for the cluster centroids, we computed the RSM detection maps for every target of the SHARDDS survey. Two sets of detection maps were computed using the original and the reverse parallactic angles. The detection maps with the reversed parallactic angles allowed the computation of a radially dependent residual noise measure⁵. A polynomial fit was applied on the obtained values to limit the influence of potential outliers (see below) and smooth the curve. This residual noise measure was finally subtracted from the original detection maps and any negative value was set to zero. This subtraction reduces the background residual noise and therefore eases the detection of potential planetary candidates.

The resulting detection maps were then analysed to uncover potential planetary signals or other bright structures. From this analysis, we rejected HD 107649 due to the presence of extended speckle-like bright structures. For other targets, some redundant epochs presenting a high degree of residual noise were also removed⁶. From the remaining ADI sequences, we identified 16 targets containing a point-like source or an extended bright structure above a probability threshold of 0.05. To insure that these detections were not the result of a sub-optimal parametrisation of the RSM algorithm, we applied the auto-RSM algorithm to 15 of these targets⁷.

We performed a correlation analysis similar to the one made in Appendix D.1 on these 15 targets, in order to assess the influence of a stronger speckle field on the optimal parametrisations. We found much lower correlation rates between these optimal parameters and the set of metrics characterising the ADI sequences. We also observed a higher degree of dissimilarity between the parametrisations of these 15 targets, especially for the PSF-subtraction techniques parameters. These results highlight the limit of a clustering approach based solely on the parameters characterising the ADI sequence, when facing noisier samples. They also demonstrate the necessity to adopt

⁵The residual noise measure is estimated by taking, for each annulus, the largest value observed in the detection map generated with reversed parallactic angles.

⁶These ADI sequences include AG Tri, AG Tri 2nd epoch, AG Tri 3rd epoch, HD 82943 and HD 82943 3rd epoch

⁷From the set of 16 targets including detections above a 0.05 probability threshold, one was a cluster centroid (HD 206893) for which we kept the original RSM detection maps.

an empirical approach, such as the auto-RSM optimisation framework, to optimise the parametrisation when the samples noise structure cannot be well captured by the set of ADI sequence characteristics. However, the low residual noise level in the detection maps shown in Figures D.2-D.4, as well as the large fraction of the survey data set (70%) that did not require the use of auto-RSM, still favour the use of a limited number of optimal parameter sets computed for well chosen targets.

Following this individual optimisation, the analysis of the resulting 16 detection maps allowed the detection of three already known point-like sources that will be further analysed in the next section (see Figure 5.5). The detection maps containing no plausible planetary candidates are shown in Appendix B. As can be seen from Figures D.2-D.4, the residual noise level is most of the time very low, except for bright structures observed in HD 53842 and HD 80950. These structures are diffraction patterns due to the presence of a bright companion just outside the 199×199 pixels window considered in this analysis⁸.

As mentioned earlier, a threshold was computed based on the detection map generated with the reversed parallactic angles. This threshold should however not be considered as a sufficient condition to classify any signal above it as a planetary candidate. As can be seen from Figure 5.3, bright structures may appear in the detection map generated with the reversed parallactic angles (right), which explains the use of a polynomial fit when estimating the threshold. Most of the time, the residual noise distributions are similar in the two detection maps, as illustrated with HD 122652 (2nd epoch). But in some cases, very bright artefacts appear in the detection map with reversed parallactic angles although only a weak level of noise is visible in the original detection map (see HD 157728). Considering all ADI sequences, around 20% of the detection maps computed with the parallactic angles reversed show point-like sources or bright structures above a 0.05 threshold while this percentage falls to 9% for the original detection maps. It is therefore preferable to avoid using reversed parallactic angles to define a detection threshold. Detection maps generated with reversed parallactic angles may however be used to reduce the level of residual noise in the original detection maps, as described in the beginning of this section.

Contrast curves

Following the computation of the detection maps, we relied on an optimised version of the approach proposed in chapter 3 to compute contrast curves for every target. When relying on probability detection maps, standard S/N-based approaches involving the estimation of the throughput and the noise standard deviation (Mawet et al. 2014)

⁸For HD 80950, the companion is situated at a projected separation of 130 au with an apparent magnitude in H band of 9.97. HD 53842 is a very young binary system, with a primary spectral type F5 star and a secondary M-dwarf situated at a projected separation of 82 au, with an estimated orbital period of 300 years (C. del Burgo, in prep).

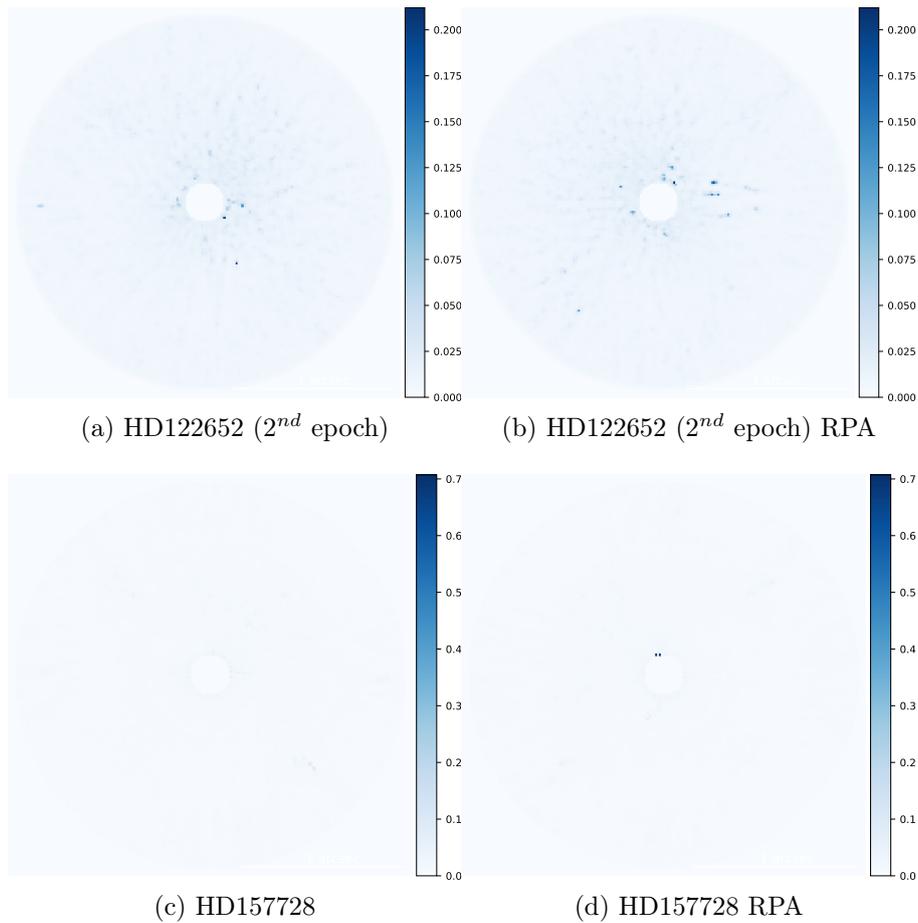


Figure 5.3: RSM detection maps with and without the parallactic angles reversed for 2 data sets (RPA stands for reversed parallactic angles). A square root scale has been selected to highlight residual speckle noise.

cannot be used. We replace this definition by an empirical estimation of the contrast corresponding to a predefined detection rate (also called true positive rate) computed at a specific threshold. As it is not possible to reach a 5σ confidence level empirically, this threshold corresponds simply to the first detection of a false positive within the entire detection map. The detection rate is computed, for a given angular separation, via the injection of fake companions at different azimuths. The computation of the contrast follows the iterative procedure, where the contrast is increased or decreased depending on the obtained detection rate and the previously tested contrasts (see Section 3.6 for more details).

We selected a detection rate of 95%, which is the traditional completeness level for the computation of planet detection probability or occurrence rate (see Section 5.6). This detection rate requires the successive injection of 20 fake companions per considered annulus. We considered nine angular separations ranging from 60 to 1150 mas. When multiple epochs were available, the lowest contrast was kept for each

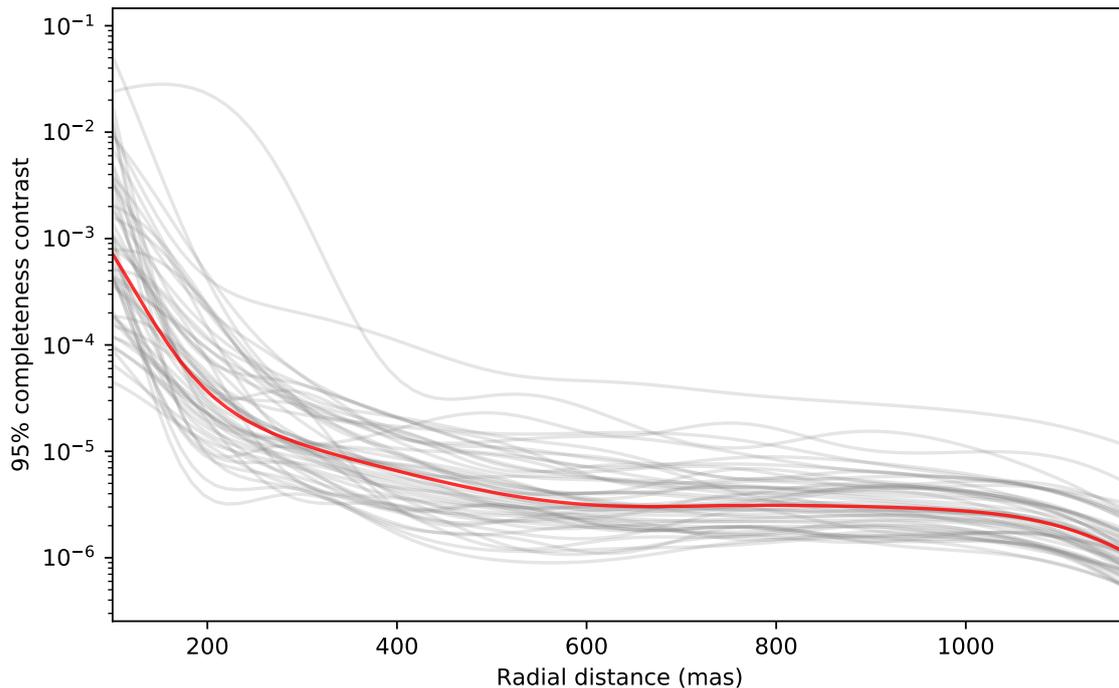


Figure 5.4: Contrast curves computed for the set of considered targets (gray) and median contrast curve (thick red) computed at a 95% completeness level.

considered angular separation, to generate a single contrast curve per target. A radial basis multiquadric function (RBF) was then used to perform the interpolation (Hardy 1971) between the nine angular separations for which a contrast was estimated.

Figure 5.4 provides a consolidated view of the contrast curves, with gray curves showing the individual contrast curves corresponding to each target and the red line providing the median. As can be seen from these curves, the contrast decreases quickly with the angular separation, with a median contrast below 10^{-5} at already 300 mas. Considering the high completeness level we have selected, it demonstrates the good performance of the RSM map algorithm at angular separations below 1 arcsec. However, we observe a relatively high dispersion of the contrasts at close separations, with the contrast ranging from 3×10^{-1} to 3×10^{-4} at 100 mas. This high dispersion can be directly linked to the observing conditions. This relationship between the performance in terms of contrasts and the observing conditions will be further investigated in Section 5.5.

5.4 Identification of planetary candidates

Figure 5.5 presents the two ADI sequences containing a signal above the previously defined threshold of 0.05, after having applied auto-RSM on the 16 sequences for which a signal was previously detected. The two ADI sequences contain already known

targets, with HD 206893 B identified in Milli et al. (2016b) , and the debris disk from HD 114082 in Wahhaj et al. (2016) which includes also two background stars. In the rest of the section, we propose a new way to extract the photometry and astrometry of point-like sources based on the RSM framework, and apply it to these two data sets.

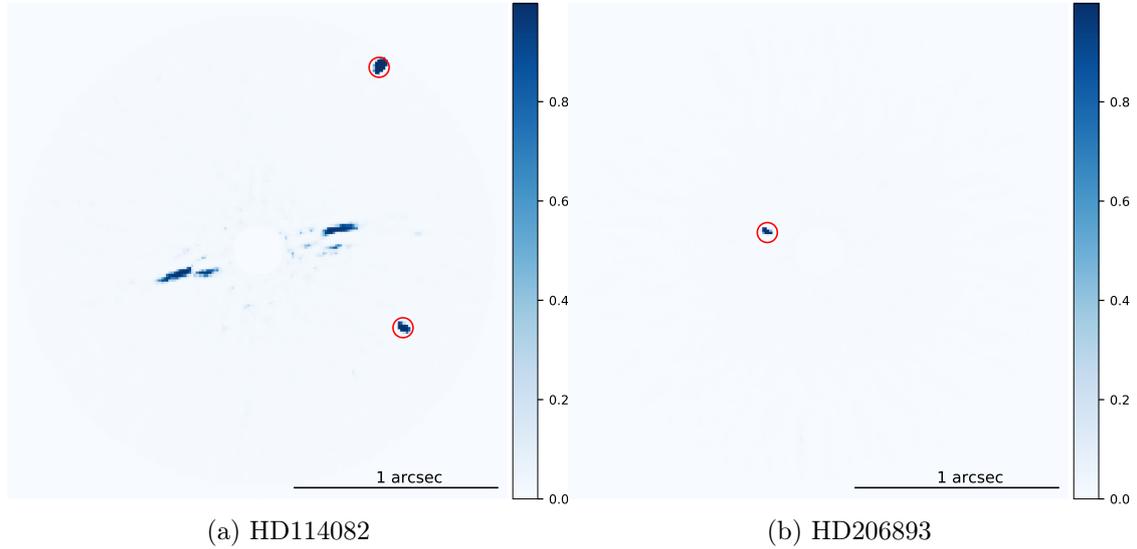


Figure 5.5: RSM detection maps generated using auto-RSM. These detection maps led to the detection of one or multiple planetary candidates. The color scale is expressed in terms of inferred detection probabilities. A square root scale has been selected to highlight potential residual speckle noise.

5.4.1 RSM NEGFC algorithm

Like in the negative fake companion (NEGFC) approach (Lagrange et al. 2010; Marois et al. 2010; Wertz, O. et al. 2017), the astrometry and photometry are determined by injecting a fake companion at the expected position of the planet, with a negative flux providing the photometry. Multiple positions and fluxes are tested and their optimum is defined as the values minimising a loss function defined as the average probability inside an aperture of two FWHM centred on the expected location of the planet. The minimisation relies on a particle swarm optimisation (PSO) framework (Kennedy & Eberhart 1995). A series of particles, each defining a set of parameters, travel within the parameter space following an iterative procedure. At each step the velocity of these particles in the parameter space is updated based on the knowledge of the particle's own optimum and the global optimum of the entire swarm. The velocity equation is given by

$$\mathbf{v}_{i,t} = \alpha \mathbf{v}_{i,t-1} + \beta_p r_p (\mathbf{l}_{i,t-1} - \mathbf{x}_{i,t-1}) + \beta_g r_g (\mathbf{g}_{i,t-1} - \mathbf{x}_{i,t-1}) \quad (5.2)$$

where $\mathbf{v}_{i,t}$ is the velocity at time t , α is the inertia factor, β_p and β_g are respectively the cognitive and social coefficients, $\mathbf{l}_{i,t-1}$ and $\mathbf{g}_{i,t-1}$ the local and global optimum, $\mathbf{x}_{i,t-1}$ the position at previous step $t-1$, and r_p and r_g two random numbers drawn from a uniform probability with support $[0, 1]$. The PSO framework was chosen as it showed, during tests, a higher convergence rate than Bayesian optimisation and allowed multi-core estimation⁹, reducing the computation time. More standard minimisation frameworks (Nelder-Mead, Newton, or Broyden-Fletcher-Goldfarb-Shanno) were tested without success because of the non linear behaviour of the selected loss function near the optimum and the presence of multiple local optima. The inertia, the cognitive and social coefficient of the PSO help defining the right balance between exploitation of known minima and exploration of the parameter space. Several sets of parameters were tested and the set $[\alpha = 0.5, \beta_p = 1, \beta_g = 1]$ was selected, as it led to a high convergence rate while avoiding local minima.

The algorithm is initialised by relying on a detection map generated with the RSM map algorithm using the forward-backward mode, which considers both past and future observations to infer the detection probability. This mode has demonstrated a higher precision in terms of astrometry in chapter 3. A Gaussian fit similar to the one used in Section 3.5 allows the determination of the initial target astrometry. Once the initial astrometry has been defined, a range of fluxes is tested to get an initial estimation of the photometry. The PSO framework is then used to minimise the average probability in the two-FWHM aperture centred on the expected position. We relied on 10 particles with a maximum number of iterations equal to 20. At the end of the PSO minimisation, the global minimum is kept and a confidence interval is computed based on the computation of the inverted Hessian matrix¹⁰.

We have tested several versions of this planetary signal characterisation algorithm. We tried to replace the average probability within a FWHM aperture by a more advanced metrics based on the noise probability distribution. We relied on the detection probabilities within the annulus containing the signal to compute this probability distribution of non detection. We then estimate the probability that the detection probabilities within the aperture is part of this Gaussian distribution to generate our loss function. However, this approach was unsatisfactory, leading to much higher astrometric and photometric errors during the performance comparison. We also tried to subtract from the average probability within the aperture, a local measure of the noise. This local noise was computed as the detection probabilities averaged over two sections of the annulus with a width of one FWHM containing the signal, situated at a distance of 1.5 FWHM on either sides of the expected target position. We did not consider the entire annulus, as local features may be observed in the detection map,

⁹Multi-core optimisation is not possible with usual minimisation algorithms such as Nelder-Mead, Newton or Broyden-Fletcher-Goldfarb-Shanno.

¹⁰The Hessian matrix is calculated with finite difference derivative approximation.

leading to a potential bias. We finally considered replacing the PSO minimisation by a Bayesian optimisation. We tested these different versions along with the NEGFC function provided by the VIP package (Gomez Gonzalez et al. 2017) which relies on Nelder-Mead minimisation.

We based our performance comparison on the ADI sequence obtained on HD 3003, considering an intermediate angular separation of $4\lambda/D$. We injected fake companions at eight different azimuths and considered eight different contrasts ranging from 1×10^{-5} to 8×10^{-5} . This range goes from a non detection in a traditional S/N map (a detection just above the background with the RSM map) to a very bright planetary signal. This should allow to investigate the behaviour of the planetary signal characterisation algorithms in two very different regimes. The astrometric error is computed as the root mean squared (rms) position error between the obtained position and the injected fake companion true position, averaged over the eight considered azimuths. The photometric error follows the same approach but comparing in terms of rms the estimated photometry and the true underlying one.

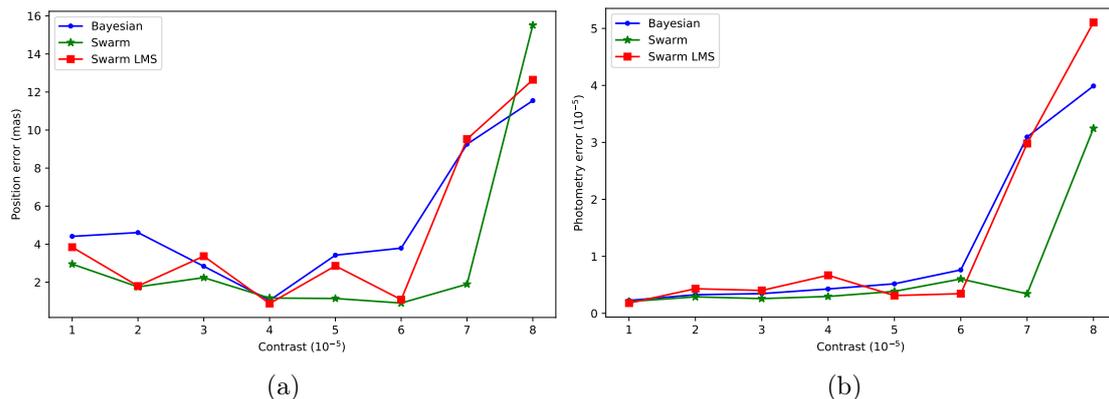


Figure 5.6: Astrometric and photometric errors for the RSM based planetary signal characterisation algorithm using the PSO approach with and without subtraction of the local mean noise (resp. red and green) and using Bayesian optimisation (blue). The left graph shows the dependence of the averaged rms position error on the contrast, while the right one shows the dependence of the photometric rms error (computed at an angular separation of $4\lambda/D$).

Figure 5.6 shows the evolution of the astrometric and photometric mean error with the contrast. The left graph shows a higher performance of the PSO approach without local mean subtraction, except for the highest contrast value. When comparing with the NEGFC algorithm in Figures 5.7, we see that our approach provides a more accurate estimation of the astrometry and photometry for low contrast values, while breaking at high contrast values. This lower performance for very bright companions comes from the fact that a slight shift of the negative injected fake companion compared to the true underlying position, leads to the appearance of bright artefacts near

the companion position, and therefore to a high loss function value which prevents its effective minimisation. This is explained by the very high sensitivity of the RSM map algorithm, which is a drawback in this particular case. A way to prevent this behaviour is to apply as an initialisation step the NEGFC algorithm and then use the RSM-based PSO approach. We see from Figures 5.7, that this approach reduces drastically the error for very bright companions, while unfortunately decreasing the astrometric accuracy when facing faint signals (but increasing the overall photometric accuracy). The optimal solution would be one combining both approaches, relying on the NEGFC approach to initialise the PSO algorithm as from a given brightness threshold.

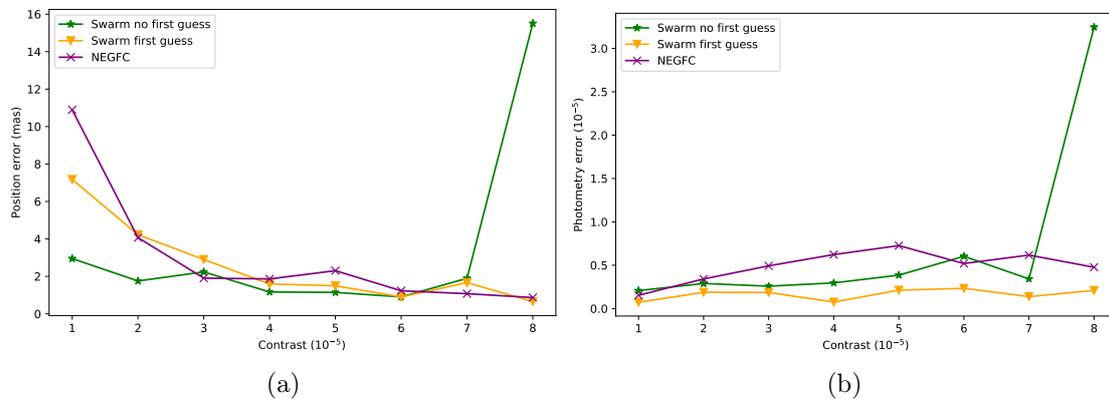


Figure 5.7: Astrometric and photometric errors for the NEGFC approach (purple) and for RSM based planetary signal characterisation algorithm using the PSO approach with and without the initialisation step relying on the NEGFC approach (resp. yellow and green). The upper graph shows the dependence of the averaged rms position error on the contrast, while the bottom one shows the dependence of the photometric rms error (computed at a radial distance of $4\lambda/D$).

5.4.2 Point-source characterisation

We applied the RSM-based planetary signal characterisation algorithm on the two targets for which signals were detected. The results are presented in Table 5.6. Besides the astrometry and photometry, we estimated contrast curves for HD 206893 at two additional completeness levels, 50% and 5%. This could further help us classifying the detected signal between planetary candidates and bright speckle, by considering its relative distance to these contrast curves. In contrast with S/N-based analysis, which relies on Gaussian assumption, there is no linear relationship between companion brightness and completeness level in RSM detection map. The distance between a companion and contrast curves estimated at different completeness levels should therefore give information about the uncertainty associated with the detection. The

Table 5.6: Detected targets photometry and astrometry.

Target	Radial distance (mas)	Position Angle ($^{\circ}$)	Contrast
Confirmed detections			
HD 206893 b	266.58 ± 3.25	159.76 ± 0.65	$4.59 \pm 0.37 \times 10^{-5}$
HD 114082 BKG (1)	803.93 ± 1.06	332.10 ± 0.08	$7.49 \pm 0.11 \times 10^{-6}$
HD 114082 BKG (2)	1082.67 ± 0.93	56.75 ± 0.05	$1.69 \pm 0.01 \times 10^{-5}$

contrast curves were computed after removing the detected signal via the negative fake companion subtraction technique, using the parameters from Table 4.1. Figure 5.8 presents the contrast curves along with the detected signal positioned at its estimated contrast and angular separation. No contrast curves were computed for HD 114082 pertaining to the difficulty of removing the disk via fake companion injections.

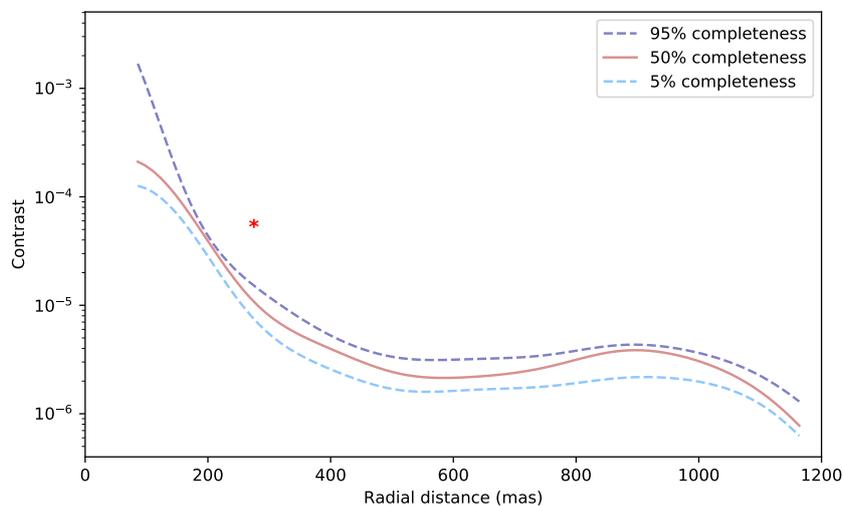


Figure 5.8: Contrast curve of the data set for which one or multiple planetary candidates have been detected. The contrast curves have been computed at a 95%, 50% and 5% completeness level. The planetary candidate is indicated by a red star at its respective contrast and angular separation from the host star.

We finally computed additional detection maps. We ran the Auto-RSM framework replacing the bottom-up approach by a top-down selection method to define the set of likelihoods cubes used to generate the final RSM detection maps. We also relied on the Auto-SNR framework (Dahlqvist et al. 2021) to generate optimised SNR maps. This framework uses the optimised parameters of the Auto-RSM framework for the PSF-subtraction techniques, but relies on dedicated function to select and combine the optimal set of SNR maps. We eventually computed SNR maps with APCA, NMF, LLSG and LOCI and simply mean combined them to generate an averaged SNR map. All these detection maps are presented in Figures D.5, with a yellow circle indicating the position of the detected signals.

HD 206893

The first detection of HD 206893 B dates back to 2015 (Milli et al. 2016b), with numerous papers devoted to its characterisation published since (e.g. Grandjean, A. et al. 2019; Kammerer et al. 2021). We see from Figures 5.8 and D.5 that HD 206893 B is a very bright companion, located well above the 95 % completeness contrast curve, and visible in all detection maps. We estimate a contrast of $4.59 \pm 0.37 \times 10^{-5}$, which translates into a mass of $24.76_{-0.62}^{+0.67} M_{Jup}$ and $33.22_{-0.34}^{+0.37} M_{Jup}$ for respectively the AMES-COND and AMES-DUSTY evolutionary models, using an estimated stellar age of 0.25 Gy. These estimated masses lie inside or close to the $[5 - 30] M_{Jup}$ range defined in (Kammerer et al. 2021), while the estimated angular separation of 266.58 ± 3.25 mas (10.88 au) is very close to the one determined for the same epoch in Milli et al. (2016b).

HD 114082

Although the RSM approach is not designed to unveil large structures, the debris disk around HD 114082, first detected by Wahhaj et al. (2016), is clearly visible. Two point-like sources are also visible. They are situated at an estimated distance of 803.93 ± 1.06 mas and 1082.67 ± 0.93 mas from HD 114082. These signals are visible in all detection maps from Figure D.5. HD 114082 being in a dense field, we rely on TRILEGAL stellar population model (Girardi et al. 2012) to infer the density of background stars around HD 114082. This density is then used to estimate the probability of observing two or more background stars at a distance below 1082.67 mas from HD 114082, using a spatial Poisson point process. This probability is equal to 63.5 %, and increase to 88.5 % when considering the probability of observing one or more background stars. Considering these high probabilities and the high inclination of these objects compared to the debris disk, these detections are most likely background stars. A second-epoch follow-up and an astrometric analysis is presented in Engler et al. (submitted) and confirmed that those two sources are background sources without proper motion.

5.5 Contrast curves analysis

The contrast curves computed in Section 5.3.4 are used throughout this section as a measure of the ADI sequences quality, as well as a metric for the RSM map algorithm performance.

5.5.1 Influence of clustering

We start by comparing the contrast curves obtained for the cluster centroids and the ones obtained by applying the centroids optimal parameters on the remaining targets of the cluster. The comparison aims to determine if the cluster centroids, for which the optimal parametrisations were computed, do perform better than the other members of the cluster in terms of achievable contrast. This should provide an idea of how far from the optimum we are, the optimum being the case where auto-RSM is applied on every target. We have estimated the difference between each of the members and their cluster centroid in terms of $\Delta \text{ mag}$ ¹¹, and report in Figure 5.9 the radial evolution of this measure averaged, for each cluster, over their set of members.

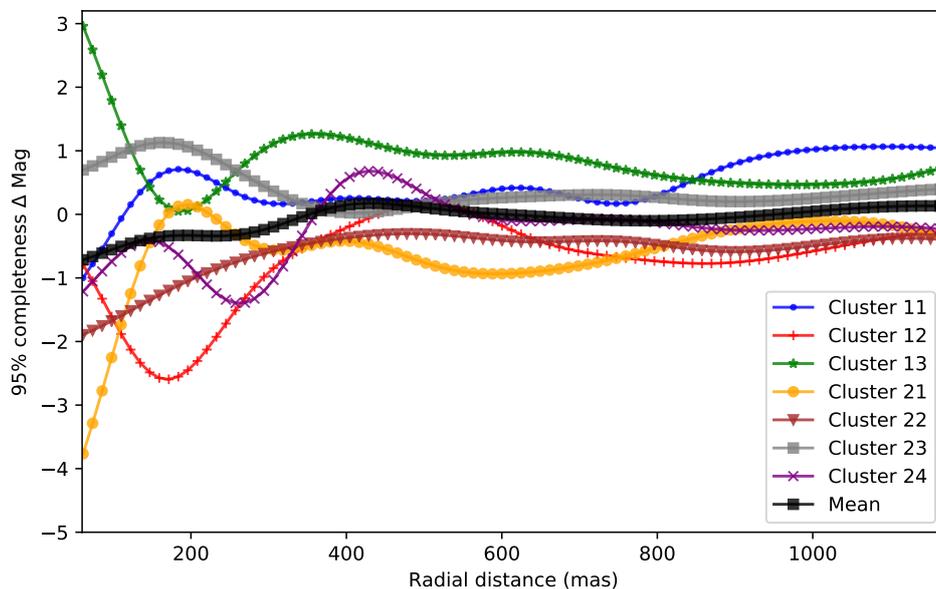


Figure 5.9: Angular evolution of the average contrast difference between cluster members and their cluster center expressed in terms of $\Delta \text{ mag}$. A positive $\Delta \text{ mag}$ indicates a poorer contrast for the cluster members.

Looking at the seven curves¹², the center seems to perform better for some clusters (see clusters 11, 13 and 23), while for others the cluster members show a higher performance (see clusters 12, 21, 22, and 24). Surprisingly, we observe on average a small increase of the performance in terms of contrast for the cluster members at close angular separations. The average performance gain is close to zero at larger separations. This seems to support the use of a reduced number of optimal parameters, as it does not seem to negatively impact the performance within the different clusters.

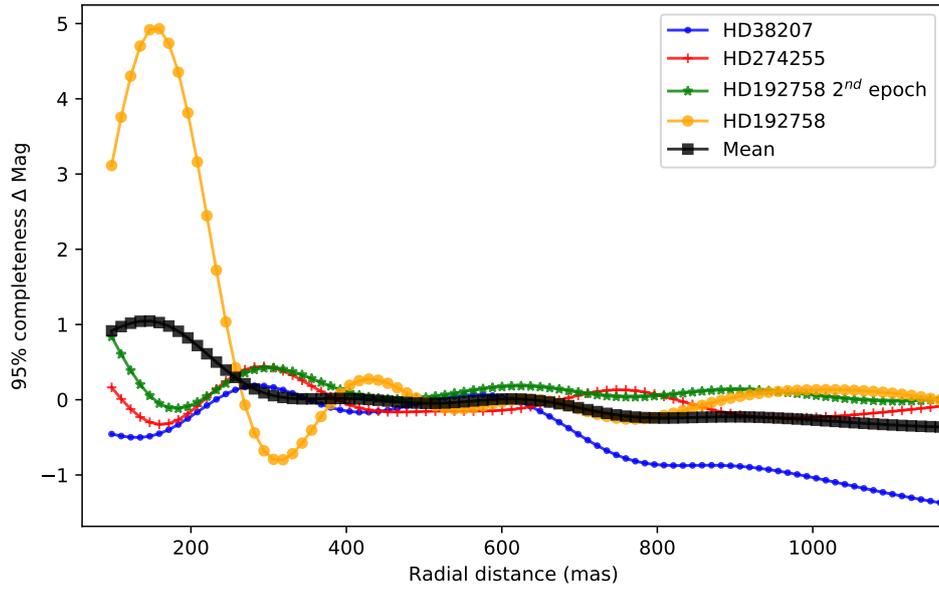
¹¹We expressed both contrast curves in terms of magnitude and then subtracted the magnitudes of the members from the one of the cluster center.

¹²The cluster composed of HD 14082B, HD 82943, HD 107649 was not included in the analysis as two of the three cluster members were rejected, due to the presence of multiple extended speckle-like bright structures in the HD 107649 detection map and the existence of better epochs for HD 82943.

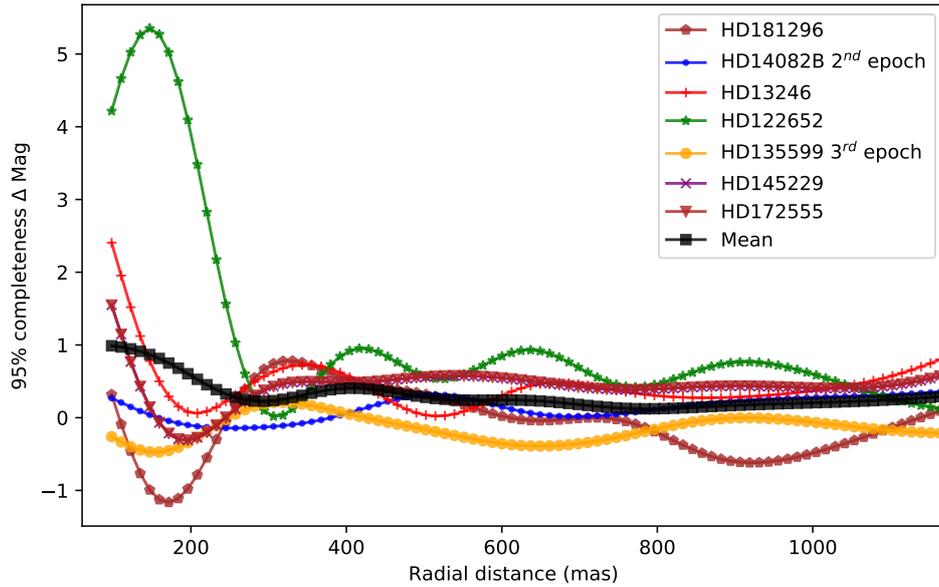
We used the same approach to assess the necessity to rely on multiple optimal parametrisations instead of a single one for the entire survey. This allows us to investigate also the impact of the degree of dissimilarity between optimal parametrisations on the performance, measured in terms of contrast. We considered two sets of clusters, one set of clusters close in terms of parametrisation, cluster 11 and 12 and one set of clusters presenting a larger level of dissimilarity, cluster 23 and 24 (see Figure D.1). We computed for cluster 11 and 23, a new set of contrast curves using respectively the optimal parameters of cluster 12 and 24 (obtained for respectively HD 3670 and HD 3003). We then estimated the difference between these new contrast curves and the contrast curves obtained with the optimal parametrisation of their own cluster centroid (respectively HD 192758 and HD 181296). These contrast differences, expressed in terms of Δ mag, are shown in Figures 5.10. As can be seen from the mean curves, using the optimal parameters estimated for their own cluster centroid leads on average to a better performance, especially at small angular separation. We see also that the mean distance is larger for the cluster 23, which showed a higher degree of dissimilarity in Figure D.1. These results highlight the added value, at close separation, of the definition of local optimal parametrisation via auto-RSM. The reasons for this higher performance are twofold. First, regions with a high level of background residual noise are more difficult to treat and are therefore more sensitive to parametrisation. Secondly, auto-RSM focuses mainly on close separations to optimise the model parameters, which explain its better performance at these distances compared to other approaches. This confirms the interest of computing several sets of optimal parameters for a large survey to account for dissimilarities in the ADI sequences' characteristics.

5.5.2 Influence of environmental parameters

We perform a similar correlation analysis as the one made in Appendix D.1, but focusing here on the relationships existing between the parameters characterising the ADI sequences and the performance in terms of achievable contrast. We start by re-expressing every contrast curve in terms of magnitude and average these magnitudes over the set of considered angular distances. We then compute the Pearson correlations between the parameters characterising the ADI sequences and the median contrast, considering the entire SHARDDS data set. As can be seen from Figure 5.11, the raw contrast at 500 mas, the Strehl and the WDH asymmetry show relatively high correlations and have the expected sign. A higher asymmetry of the WDH is indeed more difficult to treat by the PSF-subtraction techniques, which do not cope well with anisotropy in the speckle field. Despite their lower correlation, the other parameters show also the expected sign. As in Table D.1, the lowest correlation is associated to the autocorrelation measure, indicating that the decay rate of the autocorrelation function is not the best measure of the temporal relationships between the frames.



(a)



(b)

Figure 5.10: Angular evolution of the contrast difference, for the cluster 11 (top) and 23 (bottom), between the contrast obtained with their optimal parametrisation (corresponding to the optimal parametrisation of their respective cluster center, HD 192758 and HD 181296) and the contrast obtained with the optimal parametrisation of another cluster center (resp. HD 3670 and HD 3003, i.e. the center of the cluster 12 and 24), expressed in terms of Δ mag. A positive Δ mag indicates a poorer contrast achieved with the optimal parametrisation of the other cluster centres compared to their own cluster center. The black curve provide the Δ mag averaged over the set of considered targets.

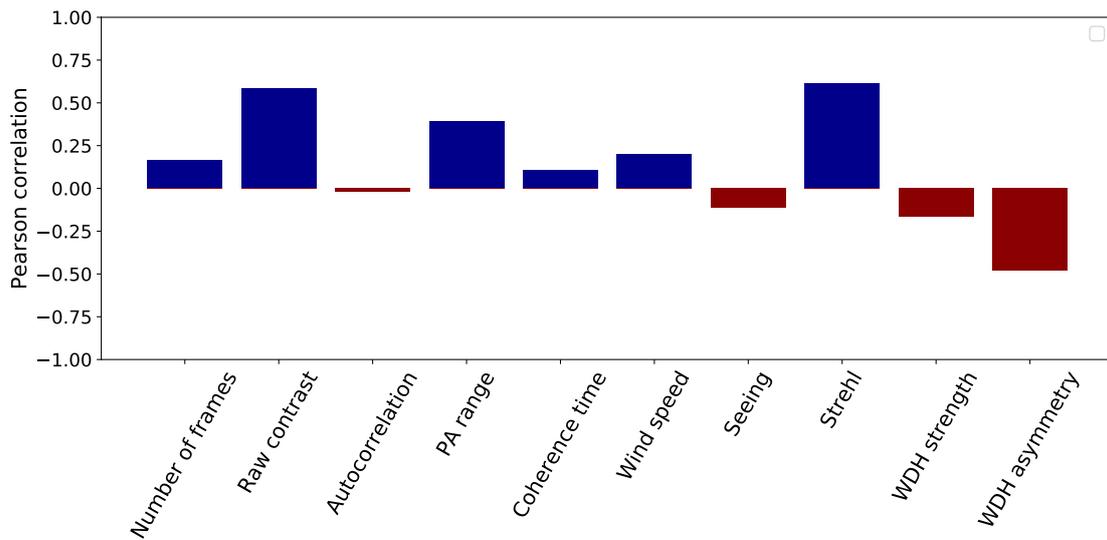


Figure 5.11: Pearson correlations between the contrast curve median values expressed in Δ mag and the parameters characterising the ADI sequences.

In order to further investigate the relationship between the achievable contrast and parameters characterising the ADI sequences, we propose to rely on linear regression to highlight the parameters contributing the most to the quality of the ADI sequences. Considering the relatively low number of data points with only 60 fully treated targets, and the potential co-linearities existing between the parameters, we rely on a bottom-up approach based on the Akaike information criterion (AIC) to select one by one the parameters to be included in our model. The AIC provides a measure of the amount of information lost by a model. This measure includes a penalty term increasing with the number of parameters, providing a good trade-off between the model complexity and its goodness of fit. We start by computing the AIC for every parameter and select the parameter having the lowest AIC. We include this parameter to the model and compute again the AIC of this model after adding one at a time each of the remaining parameters. The parameter leading to the highest reduction of the AIC is then included in the model. This procedure is repeated until no more reduction of the AIC is observed.

Table 5.7 gives the set of parameters that were selected using this method, along with the parameter values in the linear regression, their standard error, and p-value. We retrieve all three parameters that were already identified as highly correlated to the contrast in Figure 5.11, with in addition the wind speed showing a positive coefficient most probably attributable to the low wind effect. All the selected parameters show a high significance, especially the raw contrast at 500 mas and the WDH asymmetry. This highlights the importance of finding mitigation strategies to tackle the WDH to increase the quality of the ADI sequences (see Cantalloube et al. 2020a). With a R^2 adjusted for the number of parameters equal to 0.699, this simple model provides

already a good indication of the expected contrast, relying on only four parameters that can be quickly computed or are already available in the metadata.

Table 5.7: Linear regression coefficients, standard error, and p-value for the five parameters selected via the minimisation of the AIC with as dependent variable the contrast curve median values expressed in Δ mag. The minimum AIC and the adjusted R^2 are respectively equal to 60.04 and 0.699.

Parameters	Coefficient	Standard error	p-value
Contrast at 500 mas	0.5863	0.082	0.000
WDH asymmetry	-0.0468	0.013	0.001
Strehl	2.2234	1.020	0.034
Wind speed	0.0192	0.010	0.063

Following this analysis of the parameters driving the most the quality of the ADI sequences in terms of achievable contrast, we propose to look at existing observation quality ratings. In Figure 5.12, we report the different ADI sequences of the SHARDDS survey classified in terms of ESO observation quality grading and their respective mean contrast. As can be seen from this graph, apart from a single ADI sequence graded C showing a very low mean contrast, there are no major differences between the contrast distribution among the three grades. The ESO grading system used for this survey was mainly based on the seeing. A more robust multi-factor grading system was introduced in April 2018 (Milli et al. 2019). However, a more HCI-oriented grading system based on a multi-factor linear regression, such as the one presented in Table 5.7, could be an interesting tool to grade HCI observations at the telescope, and/or inform the post-processing of large surveys.

5.6 Survey Sensitivity

5.6.1 Target detection probability

The median contrast curves provide a good metric for the quality of the ADI sequences of the SHARDDS survey, and its relationship with the observing conditions. However, this analysis did not provide information about the global sensitivity of the SHARDDS survey to planets. In this section, we translate these contrast curves into upper limits on the detectability of planets depending on their semi-major axis and their mass, using respectively an astrodynamical and an evolutionary model. The astrodynamical model relies on Keplerian motion to determine the range of angular separations covered by a planet depending on its orbital elements. The evolutionary model describes how planets cool down over time depending on their mass.

Different evolutionary models were developed and refined in the past decades. For the sake of continuity with previous studies, we choose two well-known models, namely

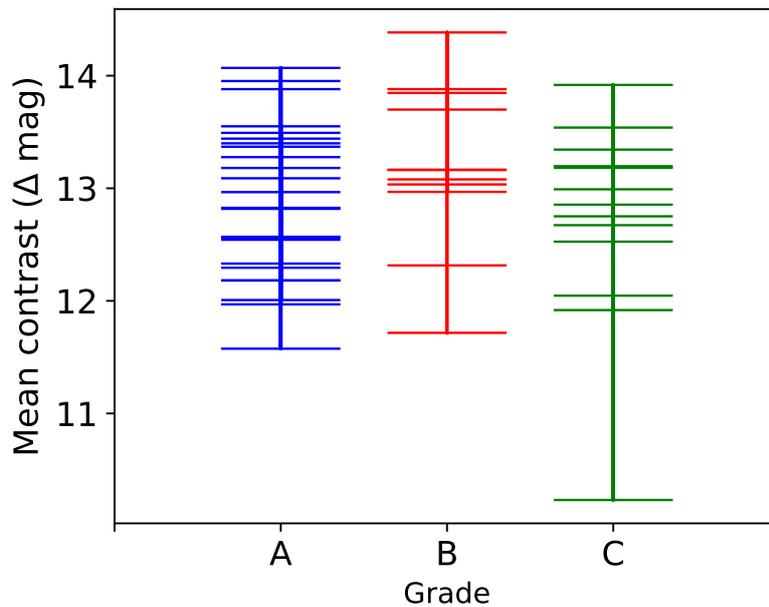


Figure 5.12: Distribution of the targets mean contrast (expressed in Δ mag) in the different grading categories.

the AMES-DUSTY (Chabrier et al. 2000) and AMES-COND (Baraffe et al. 2003) models. Both models assume planet formation via direct collapse of part of the disk due to gravitational instabilities. Direct collapse is much faster than planet formation via core-accretion. The gas forming the protoplanet retains more of its initial entropy, explaining the 'hot start' appellation often associated with the gravitational instability scenario, although this formation process allows a range of initial entropies (Spiegel & Burrows 2012). Disk instabilities are assumed to be the main scenario for the formation of giant planets and brown dwarfs at large distance from their host star (>10 au). The tables of cloud-free atmosphere AMES-COND, and dusty atmosphere AMES-DUSTY models for SPHERE were used to convert the contrast curves (Δmag) into planetary mass curves, knowing the age and the magnitude in H-band of the host star.

Having computed the planetary mass sensitivity curves for all targets, we have now to determine the accessible range of angular separations corresponding to a given semi-major axis. This range of angular separations is used alongside the planetary mass curves to compute the detection probabilities for the set of masses and semi-major axis which form the grid points of the planetary detection probability map. We define the range of angular separations for a given semi-major axis, by computing the projected distance between the planet and the host star, as seen from the Earth, for multiple sets of randomly generated orbital elements (eccentricity, inclination, argument of the periapsis, longitude of the ascending node and mean anomaly). The true anomaly is obtained through the definition of the eccentric anomaly, which is computed via an expansion in terms of Bessel functions of the mean anomaly. The position vector in

the perifocal frame is then estimated using the elliptic orbit equation based on the true anomaly, the eccentricity and the semi major axis. The angular separation is finally obtained by estimating the norm of the position vector, expressed in equatorial coordinates via three Euler rotations. The detailed computation of the angular separation is provided in Appendix D.4.

For each target of the survey, 150 semi-major axes, ranging from 0.1 to 1 000 au and 100 planetary masses, ranging from 0.1 to 100 M_{Jup} are uniformly distributed in log space to form our grid. For each point in the grid, 5000 sets of orbital elements are defined, using a uniform distribution for the inclination, the argument of the periapsis, the longitude of the ascending node, and the mean anomaly. For the inclination, we rely on a uniform distribution in sine to take into account the higher number of configurations for near edge-on orientations compared to face-on orientations, and ensure isotropy. The eccentricity follows a Beta distribution with parameters $\alpha = 0.95$ and $\beta = 1.30$, corresponding to the best fit to the full sample of wide substellar companions obtained by Bowler et al. (2020). The planetary detection rate is then computed for each target and each grid point, as the fraction of the 5000 drawn angular distances for which the considered mass lies above the planetary mass sensitivity curves. The obtained values are then averaged over the entire set of targets and multiplied by 0.95 to account for the selected completeness of the contrast curves.

Figure 5.13 shows the resulting planet detection probability maps as a function of companion mass and semi-major axis. We see that higher detection rates are obtained for a semi-major axis range of [10, 100] au with masses above 10 M_{Jup} . We have superimposed on this plot, the predicted planets derived from the dynamical constraints presented in Pearce et al. (2022). This study inferred the planet properties (mass, semi-major axis and eccentricity) if the inner edge of the disk is sculpted by one or several planets, and modelled the disk morphology based on ALMA, Herschel or the star spectral energy distribution (SED). We have plotted in Figure 5.13 the minimum masses and maximum semi-major axes of the planets predicted to be sculpting the inner edges of the disks if one planet is responsible in each of the 21 systems that are common between the SHARDDS sample and that of Pearce et al. (2022). These 21 targets are presented in Appendix D.5. These are the minimum masses and maximum semi-major axes that a single planet would need to sculpt the inner edge of the disk. Alternatively, a more massive planet located further inwards could also have the same effect. The planet masses could also theoretically be lower if multiple planets sculpt each disc, rather than just one planet, **or if the inner edge of the disk is smaller than estimated**. The disk inner edge was estimated from either a blackbody fit the Spectral Energy Distribution (SED) or if available from resolved observation with Herschel or ALMA (see Fig. 9 left in Pearce et al. 2022, the data being reproduced here in Appendix D.5). Considering the conservative limits we computed for the detection probabilities (95 % completeness), these planets are relatively close to the detection

limit when considering the AMES-COND evolutionary model.

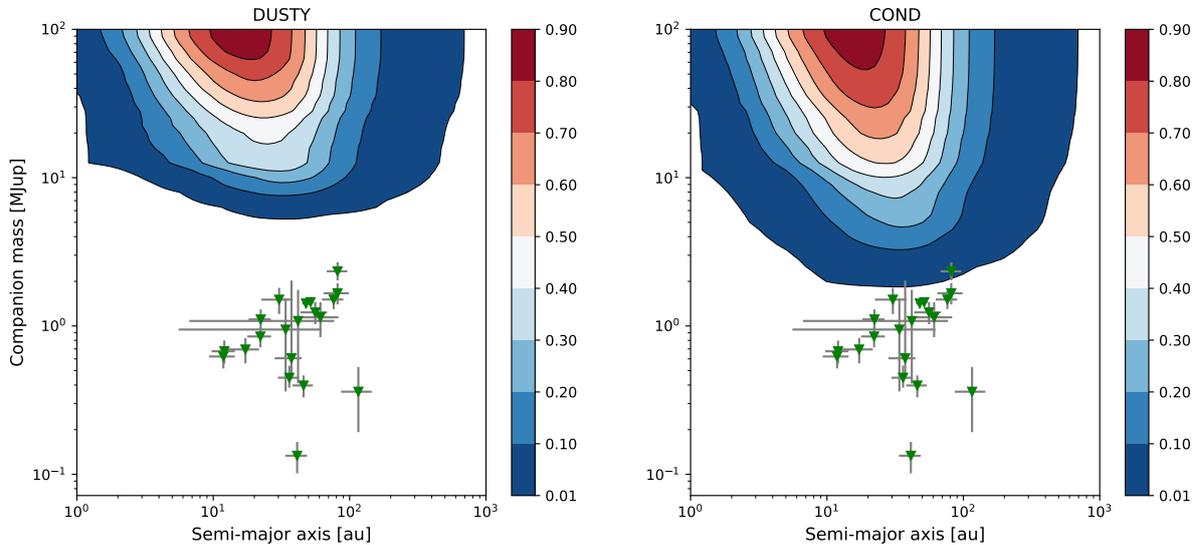


Figure 5.13: The figures give the detection probability as function of companion mass and semi-major axis. The contour plots have been calculated using the AMES-COND and AMES-DUSTY evolutionary models, relying on the contrast curves generated for the 53 targets of the SHARDDS survey (Fomalhaut C and HD 107649 have been rejected due to respectively adverse observing conditions and the presence of extended bright structures). The estimated mass and semi-major axis estimated for 21 targets of the SHARDDS survey by Pearce et al. (2022) are injected in the probability map along with the associated uncertainties.

Figure 5.14 shows the contrast curve of HD38206, the most favourable target in terms of mass and semi-major axis, translated into mass curves using the AMES-COND and AMES-DUSTY evolutionary models. We computed the probability distribution of the companion’s expected projected separation, using the orbital elements provided in (Pearce et al. 2022) and assuming a Gaussian distribution for these different orbital elements. As can be seen, the mass curve obtained with AMES-COND is very close to the expected mass of the companion for the region with the highest probability for the projected separation.

5.6.2 Occurrence rate

The definition of planetary detection probabilities allows us to derive statistical constraints on the planet occurrence rate. We consider the statistical approach proposed by Lafreniere et al. (2007a) who build confidence intervals for the planet occurrence rate relying on a Bayesian approach. We start by defining the likelihood of observing a planet characterised by a mass $m \in [m_{min}, m_{max}]$ and a semi-major axis

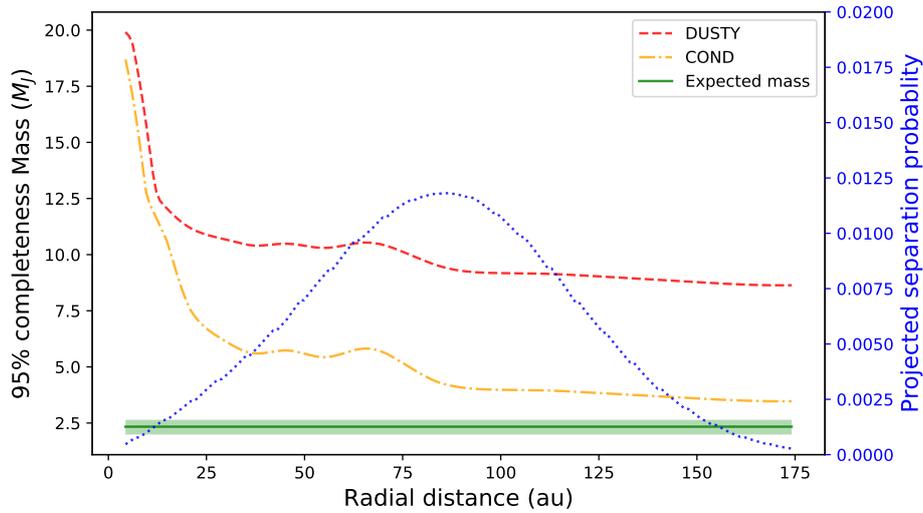


Figure 5.14: Translation of the HD38206 95 % completeness contrast curve into mass curves, using the AMES-COND (yellow dashed-dotted line) and AMES-DUSTY (red dashed line) evolutionary models. The green line provides the expected mass of the companion (Pearce et al. 2022) and the blue line the probabilities associated with the considered range of projected separations.

$a \in [a_{min}, a_{max}]$ around star $i \in [1, N]$ ¹³ as follows:

$$\mathcal{L}([d_j]|f) = \prod_{i=0}^N (1 - fp_i)^{(1-d_i)} (fp_i)^{d_i} \quad (5.3)$$

where f is the planet occurrence rate we are looking for, p_i the previously derived planet detection probability, and d_i the detections, with $d_i = 1$ for the detection of a planet with $m \in [m_{min}, m_{max}]$ and $a \in [a_{min}, a_{max}]$ around target i . The occurrence rates are computed for specific points in the mass-semi-major axis space defined for the estimation of the planet detection probabilities. We replace therefore each of the ranges $m \in [m_{min}, m_{max}]$ and $a \in [a_{min}, a_{max}]$ by a single mass and semi-major axis point.

Following Bayes' theorem, we estimate the posterior probability distribution from the likelihood and the prior probability distribution, which we set to $p(f) = 1$, assuming no prior knowledge about the distribution of the occurrence rate. The posterior probability reads:

$$p(f|[d_j]) = \frac{\mathcal{L}([d_j]|f)p(f)}{\int_0^1 \mathcal{L}([d_j]|f)p(f)df}, \quad (5.4)$$

¹³For the SHARDDS survey $N = 53$ as we removed two targets from the initial set of 55 stars because of adverse observing conditions, i.e. Fomalhaut C and HD 107649

from which we derive the minimum and maximum occurrence rate at a given level of confidence α by solving:

$$\frac{1 - \alpha}{2} = \int_0^{f_{min}} p(f|[d_j])df, \quad \frac{1 - \alpha}{2} = \int_{f_{max}}^1 p(f|[d_j])df \quad (5.5)$$

These last expressions simplify for grid points where no detection has been made within the considered set of targets. This is the case for all grid points except the one associated with HD 206893 B. The simplified expression provides only the maximum occurrence rate, f_{max} :

$$\alpha = \int_0^{f_{max}} p(f|[d_j])df \quad (5.6)$$

For each considered grid point, the occurrence rates are obtained via simplex minimisation using the Nelder-Mead approach imposing a confidence level $\alpha = 0.95$.

Figure 5.15 presents the upper limit of the companion occurrence rate obtained for the two considered evolutionary models, as a function of semi-major axis and mass. We see that the occurrence rate is especially low (below 10%) for companion with masses above $20 M_{Jup}$ with a semi major axis ranging between 10 and 60 au, because of the high sensitivity of our survey to this region. The lower sensitivity towards the larger semi-major axis, and the sensitivity peak at 30 au are explained by the stellar distances limited to 100 pc in the SHARDDS survey, as well as the field of view of 1.25 arcsec used in this study. Having considered a completeness level of 95%, we discarded a large fraction of the cumulative probability distribution of the contrast versus the detection probability. This approach is therefore conservative as it considers the lower bound of the planet detection probability, providing an upper limit of the planet occurrence rates.

5.7 Conclusion

In this chapter, we present an in-depth analysis of the SHARDDS survey in terms of point-source detection, based on the auto-RSM framework. Although the SHARDDS survey was mainly designed to image bright debris disks in near-infrared scattered light, the detection of point sources may provide a better understanding of the interaction between planets and debris disks, and give information about the formation and evolution of circumstellar systems.

Considering the computational cost of the auto-RSM framework, as well as the high degree of similarity observed between the optimal parametrisations of different ADI sequences in chapter 4, we decided to rely on clustering to reduce the number of required optimisations. We divided our data set into eight clusters using K-means

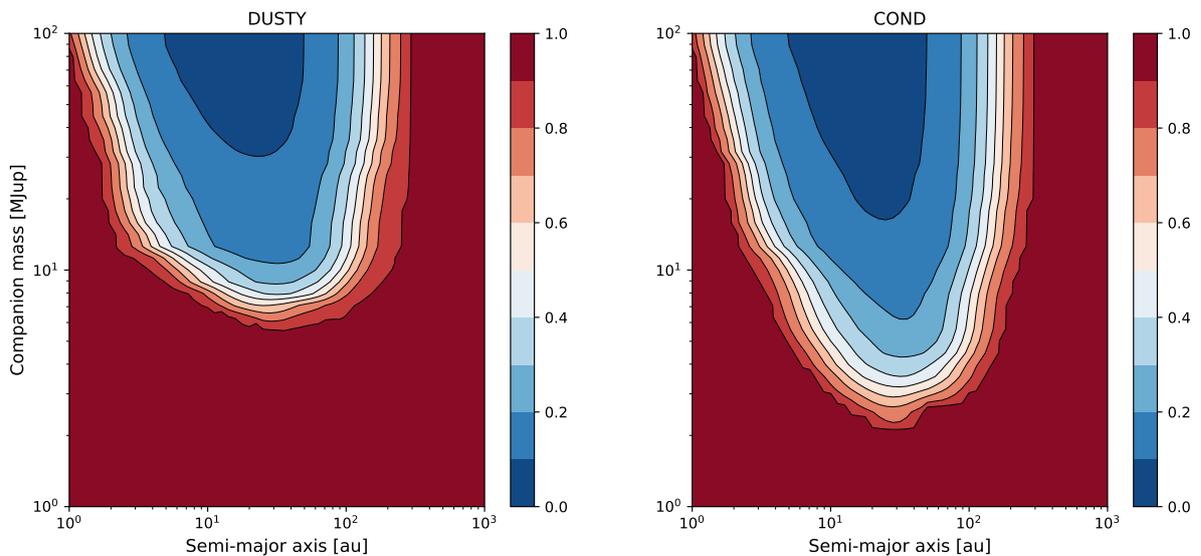


Figure 5.15: The two graphs provide the 95% confidence upper limits to the planet occurrence rate as function of companion mass and semi-major axis. The contour plots have been calculated using the AMES-COND and AMES-DUSTY evolutionary models, relying on the contrast curves generated for the 53 targets of the SHARDDS survey.

clustering algorithm, based on parameters characterising the ADI sequence itself and the related observing conditions. For each cluster, the most representative ADI sequence was selected and the auto-RSM framework was applied on it. The generated set of optimal parameters for both the PSF-subtraction techniques and the RSM algorithm was then used to generate detection maps for all the ADI sequences contained in the cluster. The analysis of the obtained detection maps showed the presence of a higher number of bright speckles when reversing the parallactic angles, providing an important reminder that care should be taken when computing detection thresholds based on reversed parallactic angles.

Based on the detection maps, we identified high-probability signals in only two ADI sequences: HD 206893 B which had already been previously detected, and the bright debris disk around HD 114082. Although these astrophysical objects had already been identified, we proposed a multi-factor detection and characterisation pipeline to confirm the detections and characterise the signals in terms of astrometry and photometry.

Following the analysis of the detection maps, we computed for each target a contrast curve at a 95% completeness level, subtracting the detected signal via the negative fake companion approach when necessary. The median contrast curve demonstrated the high performance of the auto-RSM framework, reaching a contrast of 10^{-5} at 300 mas and 3×10^{-6} at 600 mas. These contrast curves were then used to assess the performance of the proposed clustering approach. Using the contrast as a performance

metric, we found that on average the optimal parametrisation led to slightly higher performance for cluster members compared to cluster centroids. Shifting the optimal parametrisation between clusters led to lower performance in term of contrast, especially at close separation, highlighting the interest of a clustering approach to account for dissimilarities in the ADI sequences characteristics. The quality of an ADI sequence is also shown to be driven by some key observing condition metrics such as the WDH, the Strehl, the wind speed, or the raw contrast, which could allow to develop a simple and efficient HCI-oriented grading measure.

A planet detection probability map was then generated based on these contrast curves and on two different evolutionary models, AMES-COND and AMES-DUSTY. The planet detection probability map showed a high detection probability for a semi-major axis range of $[10, 100]$ au with mass above $10 M_{Jup}$. We finally computed two planet occurrence rate maps based on the estimated detection probabilities, which showed a very small occurrence rate for companions with masses above $20 M_{Jup}$ having a semi-major axis between 10 and 60 au.

The analysis of the SHARDDS survey allowed the development of new tools as well as the improvement of the auto-RSM framework, allowing it to gain in maturity and become a robust HCI post-processing pipeline, achieving good performance in terms of contrasts.

Chapter 6

Conclusion

6.1 Summary

Each chapter of this thesis has contributed to the literature either methodologically or empirically. While the first three chapters focused on the development of a new post-processing framework for ADI sequences, leading to increased performance in terms of detectability and reachable contrast, the last chapter applied this new framework to a survey in order to infer planetary occurrence rates.

The goal of the second chapter was to introduce the main concepts behind the proposed RSM map algorithm, which aims to improve exoplanet detection using a Regime Switching Model based on a Markov chain. This approach considers two regimes to describe the evolution of the pixels intensity within ADI sequences, one regime representing the planetary signal in addition to the speckle noise, and the other one only the speckle noise. The Regime Switching Model being based on a Markov chain, the pixels intensity is modelled via a probabilistic approach. This allows to compute for each pixel a probability associated to each regime, and therefore the creation of a probability map for the planetary regime, also called detection map. Beyond the use of a regime switching framework, one of the main advantages of the RSM algorithm is its ability to accommodate multiple cubes of residuals obtained with different ADI-based post-processing techniques. The speckle field being treated differently by different PSF-subtraction techniques, considering several techniques helps reducing the residual speckle noise level, and therefore better discriminating planetary signals from residual speckles. This reduction of the residual speckle noise level is attributable to the short memory process at the heart of our RSM detection map, which allows quasi-static speckles to be treated more effectively.

The performance of the vanilla RSM map algorithm was tested by injecting fake companions at different angular separations with different contrasts into two data sets provided by the VLT/NACO and VLT/SPHERE instruments. The performance of the RSM detection map was compared to standard S/N maps obtained with three PSF-

subtraction techniques (annular PCA, NMF and LLSG), using ROC curves. The ROC curves were computed at three angular separations by considering multiple contrasts for the injecting fake companions. The ROC curves demonstrate clearly the interest of our model as it outperforms all the other methods, especially at close separation from the host star. The performance assessment via ROC curve allowed also to highlight the importance of selecting the right probability distribution, within the RSM algorithm, to describe the speckle noise. Indeed, the Laplacian distribution performs clearly better for close separations while the Gaussian one provides better results for larger angular distances. The radially dependent optimal selection of the probability distribution was therefore included in the RSM algorithm, relying on a best fit approach.

Building on the good results obtained with the RSM algorithm, the third chapter was dedicated to improving the algorithm at different levels. We first considered a forward-model version of the RSM map for two different PSF-subtraction techniques, KLIP and LOCI. The computation of forward-modelled PSFs allowed to take into account PSF distortions due to the speckle field subtraction. We relied again on ROC curves to assess the added value of forward-modelled PSFs. The results showed a higher performance of the two forward-model versions of the RSM map algorithm at small angular separations, while the gap between forward-model and standard approach reduces at larger angular separations, as the self-subtraction patterns fade away, the relative movement of astrophysical signals increasing with the distance to the host star. These ROC curves highlight also the radial dependence of the optimal forward-modelled PSF crop size, with larger crop sizes leading to better results for small angular separations while the reverse is true for smaller crop sizes. This could be explained by the fact that larger PSF crop sizes better account for the self-subtraction patterns, whose intensity decreases with radial distance.

The third chapter also investigated the question of the optimal selection of the PSF subtraction techniques to be included in the estimation of the RSM map. Relying again on ROC curves, we demonstrated the importance of this selection, with sometimes large differences between the performance of the considered combinations of PSF-subtraction techniques. The question of the optimal combination, which seems to depend on the instrument generating the ADI sequence, is further explored in the fourth chapter. The next improvement considered in the third chapter relates to the way probabilities are estimated within the RSM algorithm, by replacing the original forward approach by a forward-backward approach. This approach relies on both past and future observations to compute a final probability. This allows us to deal more efficiently with background noise as speckles are not treated in the same way when relying on a forward and backward approach. Another advantage of the forward-backward version of the RSM map is its ability to extract more precisely the planetary astrometry.

Besides these improvements of the RSM map algorithm, the third chapter allowed

the development of a new framework to compute contrast curves in the case of probability maps, which cannot be treated as S/N maps. The 50% true positive rate (TPR) is kept, but the Gaussian-based noise threshold definition is replaced by a threshold based on the detection of the first false positive within the entire frame. The contrast curves are computed using a procedure based on successive linear interpolations to define, at successive angular separations, the contrast associated with a TPR of 50 % at the predefined threshold. Contrast curves computed with both forward and forward-backward version of the RSM map algorithm highlighted the ability of the forward version to detect fainter companions at small angular separations (below 400 mas), and the overall similar results for larger separations. However, the forward-backward approach reduces drastically the speckle background noise and gives a higher precision in terms of astrometry. It should therefore be considered when characterising a detected signal, while the forward version of the RSM map algorithm is better suited for detections.

Following the improvements made in the third chapter, the RSM map algorithm can accommodate up to seven PSF-subtraction techniques, increasing the complexity of the parameter selection and the associated arbitrariness. The fourth chapter addresses this shortcoming via the development of an automated optimisation framework for the RSM approach, called auto-RSM, which should provide users with a simple framework to compute reliable detection maps. The proposed multi-step parameter-optimisation framework considers the optimisation of the PSF-subtraction techniques, as well as the optimisation of the RSM approach parametrisation. Building on the findings made in chapter three, it also considers the selection of the optimal set of PSF-subtraction techniques and ADI sequences to be considered when generating the final detection map. The selection of the optimal set of parameters for the PSF-subtraction techniques is based on the minimisation of the mean contrast within the selected set of annuli, while the optimisation of the RSM approach and selection of the optimal set of cubes of likelihoods are based on the probability ratio between injected fake companion peak probability and background residual probabilities. As some PSF-subtraction techniques have a continuous parameter space, a Bayesian optimisation framework is proposed to explore the parameter space and select the optimal set of parameters.

Two versions of the auto-RSM algorithm were considered, an annular version where the set of optimal parameters evolves with radial distance, and a full-frame version where a single set of parameters is selected for all angular separations. Different flavours of the full-frame and annular auto-RSM algorithm were tested to investigate the added value of different methods to select the optimal set of cubes of likelihoods (bottom-up or top-down approach), or to compute the final probabilities (forward or forward-backward approach). The auto-RSM framework being not limited to the RSM algorithm, we also tested a version of our optimisation framework replacing the RSM algorithm by the computation of standard S/N maps. The auto-S/N algorithm

relies on the same first step, i.e., the optimisation of the PSF-subtraction techniques parametrisation, but replaces the last two steps by the computation of a single S/N map via an optimal weighted sum of the S/N maps.

The performance of these different versions were assessed by relying on the data sets of the Exoplanet Imaging Data Challenge and on the performance assessment framework proposed in Cantalloube et al. (2020b). Several performance metrics were computed based on the detection maps generated for the nine ADI sequences included in the EIDC data sets: a data set-dependent F1 score at a predefined threshold, as well as the estimation of the area under the curve (AUC) for the TPR, FPR, and FDR. In most cases, auto-RSM provided better performance than the original RSM-detection map submitted to the EIDC, which was already at or close to the top of the ranking for all performance metrics in the EIDC. The full-frame auto-RSM using the bottom-up approach to select the optimal set of cubes of likelihood and the forward approach to compute the RSM probabilities provided the best overall performance in terms of detection. Considering the longer computation time and lower performance of the annular version, the full-frame auto-RSM was selected for the survey considered in chapter 5. Regarding the auto-S/N algorithm, it showed a degraded performance compared to auto-RSM, although its reduced computation time could make it sometimes a good complement to auto-RSM.

Considering the high computational requirement of auto-RSM, we studied the commonalities existing between the parametrisations of the nine data sets of the EIDC in order to investigate the possibility of using a common set of parameters for each instrument. The computed dissimilarity metrics showed that the distance between the parametrisations for a common instrument is smaller than the distance between the parametrisations of different instruments. Potential differences between the noise characteristics of different data sets generated with a common instrument should nevertheless be taken into account, as illustrated by the NIRC2 data sets. These findings suggest the possibility of using a reduced number of optimal parametrisations for large surveys, allowing to reduce sensibly the required computation time.

The last chapter of this thesis was dedicated to an empirical application of the RSM framework developed in the three previous chapters. The aim of the chapter was twofold: the first goal was to process an entire survey with the RSM framework in order to obtain detection maps and contrast curves, allowing the computation of exoplanet occurrence rate maps, and identify potential planetary candidates. Although the SHARDDS survey was mainly dedicated to near-infrared imaging of bright debris disks in scattered light, the detection of point sources should help to better understand the interaction between planets and debris disks, and give information about the formation and evolution of protoplanetary systems. The second objective was to identify potential shortcomings in the RSM framework, solve them, and improve the overall performance and reliability of the PyRSM python package.

Following the findings made in the fourth chapter regarding the high degree of similarity existing between optimal parametrisations for a given instrument and considering the computational cost of the auto-RSM framework, we decided to rely on clustering approach to reduce the number of required optimisations. Relying on the K-means clustering algorithm, we divided our data set into eight clusters, based on parameters characterising the ADI sequence and on the observing condition. For each cluster, we defined the ADI sequence closest to the cluster center as the most representative sequence and apply the auto-RSM framework. The obtained set of optimal parameters for the eight centroids were then used to generate detection maps for all the ADI sequences contained in their respective cluster. As already suggested in chapter 4, the optimal parametrisations of the different cluster centroids showed a relatively high degree of similarity. However, relying on contrast curves as performance metrics, we showed in the last chapter that using the optimal parametrisation of another cluster led to degraded performance, especially at close separations. These results demonstrate the interest of relying on a few well-selected ADI sequences to compute local optimal parametrisations, instead of relying on a single one for the entire survey.

Looking at the detection maps, only two targets seemed to include a potential planetary candidate. Although these astrophysical objects had already been identified, we proposed a multi-factor detection and characterisation pipeline to confirm the detections and characterise the signals in terms of astrometry and photometry. We generated for these targets multiple detection maps using several algorithms to confirm the detection. We also computed contrast curves at several levels of completeness in order to compare them with the companion angular separation and contrast that were estimated using a new planetary signal characterisation algorithm based on the RSM framework. We identified HD206893 B, which had already been previously detected, as well as the bright debris disk around HD114082, which also contained two background stars.

Following the identification of potential planetary candidates and bright speckles, we computed for each ADI sequence a contrast curve at a 95% completeness level, subtracting the detected signals via the negative fake companion approach when necessary. The median contrast curves demonstrated the high performance of the auto-RSM framework, reaching a mean contrast of 10^{-5} at 300 mas and 3×10^{-6} at 600 mas. A planet detection probability map was then generated based on these contrast curves and on two different evolutionary models: AMES-COND and AMES-DUSTY. The planet detection probability map showed a high detection probability for a semi-major axis range of [10, 100] AU with mass above $10 M_{Jup}$. We finally computed two planet occurrence frequency maps based on the estimated detection probabilities, which showed a very small occurrence frequency for companions with masses above $20 M_{Jup}$ having a semi-major axis between 10 and 60 AU. These occurrence rate maps were generated at a completeness level of 95%, providing an upper limit as only a

small fraction of the contrast distribution in terms of completeness levels was taken into account.

6.2 Perspectives

Throughout these four chapters, we went from a simple concept to a complex but highly efficient post-processing framework for ADI (and SDI) sequences, adding in each chapter many new features and simplifying its use. The PyRSM package includes all the functions described along these lines, and offers a parameter-free detection map computation algorithm with a very low level of residual speckles, allowing a simple detection threshold selection. This package has largely increased in maturity thanks to the SHARDDS survey and has become a robust HCI post-processing pipeline, achieving good performance in terms of contrasts.

Several improvements could however be considered in the near future to further increase the sensitivity of the RSM framework to faint planetary signals. The probability distribution describing the residual speckle noise and its parametrisation has a critical impact on the ability of the RSM algorithm to effectively remove background noise. An annular-based approach was proposed in the second chapter to account for the radial evolution of the noise statistics, by switching between different probability distributions. Different regions were also considered to compute the probability distribution parameters. Although these different approaches increased significantly the performance of the RSM framework, more advanced distributions could be considered such as mixtures of multiple probability distributions, or distributions with evolving skewness and kurtosis (e.g. CanteroMitjans et al. in prep). The development of a faster approach to define the most suitable probability distribution and its optimal parametrisation should also help reducing the computation time, as this step represents a large fraction of the RSM likelihood estimation.

The identification of planetary signals in probability maps requires the definition of a detection threshold. As shown in chapter 5, probability maps with reversed parallactic angles cannot directly be used to define such a detection threshold. However, a probabilistic analysis of a large sample of detection maps with and without reversed parallactic angles could allow the definition of a threshold probability distribution giving an uncertainty associated to each threshold value. The main challenge will be to relate the probability distribution of the maximum probabilities in the detection maps with and without reversed parallactic angles. We indeed saw in Section 5.3.4 that the probability of having high maximum probabilities increased when using reversed parallactic angles, requiring some truncation of the distribution. It could also be interesting to determine the source of these bright artefacts.

An SDI version of the auto-RSM framework has been included in PyRSM and

tested successfully on a single SDI data set. A performance assessment of this version could be considered, relying for example on the EIDC SDI data sets. This performance assessment could also help defining the right balance between the computation time and the number of spectral channels to be considered to reach high performance in terms of achievable contrast. Besides the use of multiple spectral channels, the RSM framework can also accommodate multiple ADI cubes of a single target, taken at different epochs. This multi-epoch approach has been tested only on a few targets, but a more thorough analysis is necessary to determine its potential added value, by comparing its performance with single epoch estimations.

We estimated in the last chapter exoplanet detection probability maps and occurrence rate maps by considering a 95 % completeness level, which provide very conservative estimate, as we consider only a small fraction of the contrast distribution in terms of completeness levels. The analysis of future surveys should rather rely on the full distribution in terms of completeness levels to get a better estimate of the occurrence rate. In contrast with S/N-based contrast curve estimation, there is no linear relationship between the considered completeness level and achievable contrast with the RSM model (see Figure 5.8). The computation of the contrast distribution in terms of completeness levels should be done empirically by computing the contrasts for a range of completeness levels and interpolate for intermediate values. A new function is now included in PyRSM and VIP packages to compute such completeness maps with the RSM model but also with any PSF-subtraction techniques. This new function should provide more accurate estimates for the exoplanet occurrence rate computation, and also give new insights about the radial evolution of the residual noise homogeneity, broader contrast distribution in terms of completeness levels indicating a larger heterogeneity.

Following the good results obtained for the SHARDDS survey, the clustering approach proposed in the last chapter could be further refined and apply to other surveys such as the SHINE survey. The list of parameters providing information about the ADI sequence quality and the observing conditions could be adapted by considering the correlation analysis performed in the last chapter. The approach used to determine the optimal number of clusters could also be improved. Building on the findings made in the last chapter, a new grading system for HCI observations could be developed and proposed to ESO. Such a grading system should however be tested with other HCI post processing techniques to confirm its interest.

Considering the maturity of the current RSM framework and the many potential developments presented along these lines, PyRSM will hopefully be used for many more surveys and provide unprecedented detection limits, allowing the detection of many exoplanets with the next generation of telescopes and instruments.

Bibliography

- Absil, O., Milli, J., Mawet, D., et al. 2013, *A&A*, 559, L12
- Ackerman, A. S. & Marley, M. S. 2001, *ApJ*, 556, 872
- Allard, F., Hauschildt, P. H., Alexander, D. R., Tamanai, A., & Schweitzer, A. 2001, *ApJ*, 556, 357
- Allard, N. F., Spiegelman, F., & Kielkopf, J. F. 2016, *A&A*, 589, A21
- Allard, N. F., Spiegelman, F., Leininger, T., & Molliere, P. 2019, *A&A*, 628, A120
- Amara, A. & Quanz, S. P. 2012, *Monthly Notices of the Royal Astronomical Society*, 427, 948
- Babcock, H. W. 1953, *The Astronomical Society of the Pacific*, 65, 229
- Baraffe, I., Chabrier, G., Barman, T. S., Allard, F., & Hauschildt, P. H. 2003, *A&A*, 402, 701
- Baudino, J.-L., Bezard, B., Boccaletti, A., et al. 2015, *A&A*, 582, A83
- Baudoz, P., Dorn, R., Lizon, J.-L., et al. 2010, McLean, Ian S.; Ramsay, Suzanne K.; Takami, Hideki: *Ground-based and Airborne Instrumentation for Astronomy III*, SPIE, 77355B-77355B-7 (2010), 7735
- Berardo, D., Cumming, A., & Marleau, G.-D. 2017, *The Astrophysical Journal*, 834, 149
- Beuzit, J.-L., Vigan, A., Mouillet, D., et al. 2019, *Astronomy & Astrophysics*, 631, A155
- Bonavita, M., Gratton, R., Desidera, S., et al. 2021, *arXiv e-prints*, arXiv:2103.13706
- Bond, I. A., Abe, F., Dodd, R. J., et al. 2001, *MNRAS*, 327, 868
- Booth, M., del Burgo, C., & Hambaryan, V. V. 2021, *MNRAS*, 500, 5552
- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Science*, 327, 977

- Borucki, W. J., Koch, D. G., Basri, G., et al. 2011, *ApJ*, 736, 19
- Boss, A. P. 2000, *The Astrophysical Journal Letters*, 536, L101
- Boss, A. P. 2002, *The Astrophysical Journal*, 576, 462
- Boss, A. P., Butler, R. P., Hubbard, W. B., et al. 2005, *Proceedings of the International Astronomical Union*, 1, 183
- Bottom, M., Ruane, G., & Mawet, D. 2017, *Research Notes of the AAS*, 1, 30
- Bouwmans, T. & Zahzah, E. H. 2014, *Computer Vision and Image Understanding*, 122, 22
- Bowler, B. P., Blunt, S. C., & Nielsen, E. L. 2020, *AJ*, 159, 63
- Bryan, M. L., Chiang, E., Bowler, B. P., et al. 2020, *The Astronomical Journal*, 159, 181
- Cantalloube, F., Farley, O. J. D., Milli, J., et al. 2020a, *A&A*, 638, A98
- Cantalloube, F., Gomez-Gonzalez, C., Absil, O., et al. 2020b, in *Adaptive Optics Systems VII*, ed. L. Schreiber, D. Schmidt, & E. Vernet, Vol. 11448, *International Society for Optics and Photonics (SPIE)*, 1027 – 1062
- Cantalloube, F., Mouillet, D., Mugnier, L. M., et al. 2015, *A&A*, 582, A89
- Casagrande, L., Schönrich, R., Asplund, M., et al. 2011, *A&A*, 530, A138
- Chabrier, G. & Baraffe, I. 1997, *A&A*, 327, 1039
- Chabrier, G., Baraffe, I., Allard, F., & Hauschildt, P. 2000, *ApJ*, 542, 464
- Charbonneau, D., Brown, T. M., Latham, D. W., & Mayor, M. 2000, *The Astrophysical Journal*, 529, L45
- Chauvin, G., Lagrange, A. M., Bonavita, M., et al. 2010, *A&A*, 509, A52
- Chauvin, G., Lagrange, A. M., Dumas, C., et al. 2004, *A&A*, 425, L29
- Chauvin, G., Lagrange, A.-M., Dumas, C., et al. 2005, *A&A*, 438, L25
- Cheetham, A. C., Samland, M., Brems, S. S., et al. 2019, *A&A*, 622, A80
- Chen, C. H., Mittal, T., Kuchner, M., et al. 2014, *The Astrophysical Journal Supplement Series*, 211, 25
- Choquet, É., Milli, J., Wahhaj, Z., et al. 2017, *ApJL*, 834, L12

- Claudi, R., Maire, A.-L., Mesa, D., et al. 2019, *A&A*, 622, A96
- Cosslett, S. R. & Lee, L.-F. 1985, *Journal of Econometrics*, 27, 79
- Cronin-Coltsmann, P. F., Kennedy, G. M., Kalas, P., et al. 2021, *MNRAS*, 504, 4497
- Cumming, A., Butler, R. P., Marcy, G. W., et al. 2008, *Publ.Astron.Soc.Pac.*, 120, 531
- Currie, T., Biller, B., Lagrange, A.-M., et al. 2022, In prep
- Currie, T., Burrows, A., Madhusudhan, N., et al. 2013, *ApJ*, 776, 15
- Currie, T., Fukagawa, M., Thalmann, C., Matsumura, S., & Plavchan, P. 2012, *The Astrophysical Journal*, 755, L34
- da Silva, L., Torres, C. A. O., de la Reza, R., et al. 2009, *A&A*, 508, 833
- Dahlqvist, C.-H., Cantalloube, F., & Absil, O. 2020, *A&A*, 633, A95
- Dahlqvist, C.-H., Cantalloube, F., & Absil, O. 2021, *A&A*, 656, A54
- Delgado Mena, E., Israelian, G., González Hernández, J. I., et al. 2014, *A&A*, 562, A92
- Delorme, P., Gagne, J., Malo, L., et al. 2012, *A&A*, 548, A26
- Delorme, P., Meunier, N., Albert, D., et al. 2017, in SF2A proceedings
- Delorme, P., Schmidt, T., Bonnefoy, M., et al. 2017, *A&A*, 608, A79
- Desidera, S., Chauvin, G., Bonavita, M., et al. 2021, *A&A*, 651, A70
- Dierickx, P. 1992, *Journal of Modern Optics*, 39, 569
- Dong, R., Zhu, Z., Rafikov, R. R., & Stone, J. M. 2015, *The Astrophysical Journal*, 809, L5
- D’Orazi, V., Desidera, S., Gratton, R. G., et al. 2017, *A&A*, 598, A19
- Duchene, G. & Kraus, A. 2013, *Annual Review of Astronomy and Astrophysics*, 51, 269
- Eiroa, C., Marshall, J. P., Mora, A., et al. 2013, *A&A*, 555, A11
- Engler, N., Milli, J., Vigan, A., et al. submitted, *A&A*, 582, A89
- Fernandes, R. B., Mulders, G. D., Pascucci, I., Mordasini, C., & Emsenhuber, A. 2019, *ApJ*, 874, 81

- Fernández, D., Figueras, F., & Torra, J. 2008, *A&A*, 480, 735
- Flasseur, O., Denis, L., Thiébaud, E., & Langlois, M. 2018, *A&A*, 618, A138
- Flasseur, O., Denis, L., Thiébaud, É., & Langlois, M. 2020, *A&A*, 637, A9
- Fulton, B. J., Rosenthal, L. J., Hirsch, L. A., et al. 2021, *The Astrophysical Journal Supplement Series*, 255, 14
- Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, *A&A*, 595, A1
- Geballe, T. R., Saumon, D., Golimowski, D. A., et al. 2009, *ApJ*, 695, 844
- Gebhard, T. D., Bonse, M. J., Quanz, S. P., & Schölkopf, B., eds. 2022, Half-sibling regression meets exoplanet imaging: PSF modeling and subtraction using a flexible, domain knowledge-driven, causal framework, in press
- Girardi, L., Barbieri, M., Groenewegen, M. A. T., et al. 2012, in *Astrophysics and Space Science Proceedings*, Vol. 26, Red Giants as Probes of the Structure and Evolution of the Milky Way, 165
- Goebel, S. B., Guyon, O., Hall, D. N. B., et al. 2018, *Publications of the Astronomical Society of the Pacific*, 130, 104502
- Goldfeld, S. M. & Quandt, R. E. 1973, *Journal of Econometrics*, 1, 3
- Golimowski, D. A., Ardila, D. R., Krist, J. E., et al. 2006, *The Astronomical Journal*, 131, 3109
- Gomez Gonzalez, C. A., Absil, O., & Van Droogenbroeck, M. 2018, *A&A*, 613, A71
- Gomez Gonzalez, C. A., Wertz, O., Absil, O., et al. 2017, *Aj*, 154, 7
- Gonzalez, C. G., Absil, O., Absil, P.-A., et al. 2016, *A&A*, 589, A54
- Goździewski, K. & Migaszewski, C. 2020, *The Astrophysical Journal*, 902, L40
- Grandjean, A., Lagrange, A.-M., Beust, H., et al. 2019, *A&A*, 627, L9
- Gullikson, K., Kraus, A., & Dodson-Robinson, S. 2016, *The Astronomical Journal*, 152, 40
- Hamilton, J. D. 1988, *Journal of Economic Dynamics and Control*, 12, 385
- Hamilton, J. D. 1994, *Time Series Analysis* (Princeton University Press.)
- Hardy, R. L. 1971, *Journal of Geophysical Research (1896-1977)*, 76, 1905

- Hubickyj, O., Bodenheimer, P., & Lissauer, J. J. 2005, *Icarus*, 179, 415
- Isaac, N., I. Bernard, C., & Anne, W. 1999, *The Principia: A New Translation and Guide*, ed. U. of California Press
- Jensen-Clem, R., Mawet, D., Gonzalez, C. A. G., et al. 2017, *The Astronomical Journal*, 155, 19
- Jones, D., Schonlau, M., & Welch, W. 1998, *Journal of Global Optimization*, 13, 455
- Kalas, P., Graham, J. R., Chiang, E., et al. 2008, *Science*, 322, 1345
- Kalas, P., Graham, J. R., & Clampin, M. 2005, *Nature*, 435, 1067
- Kammerer, J., Lacour, S., Stolker, T., et al. 2021, *A&A*, 652, A57
- Kennedy, J. & Eberhart, R. 1995, in *Proceedings of ICNN'95 - International Conference on Neural Networks*, Vol. 4, 1942–1948 vol.4
- Keppler, M., Benisty, M., Müller, A., et al. 2018, *A&A*, 617, A44
- Kiefer, S., Bohn, A. J., Quanz, S. P., Kenworthy, M., & Stolker, T. 2021, *A&A*
- Kim, J. S., Hines, D. C., Backman, D. E., et al. 2005, *ApJ*, 632, 659
- King, J. R., Villarreal, A. R., Soderblom, D. R., Gulliver, A. F., & Adelman, S. J. 2003, *AJ*, 125, 1980
- Kuzuhara, M., Tamura, M., Kudo, T., et al. 2013, *ApJ*, 774, 11
- Lacour, Wang, J. J., Rodet, L., et al. 2021, *A&A*, 654, L2
- Lafreniere, D., Doyon, R., Marois, C., et al. 2007a, *The Astrophysical Journal*, 670, 1367
- Lafrenière, D., Marois, C., Doyon, R., & Barman, T. 2009, *ApJL*, 694, L148
- Lafreniere, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, E. 2007b, *ApJ*, 660
- Lagrange, A. M., Boccaletti, A., Milli, J., et al. 2012, *A&A*, 542, A40
- Lagrange, A.-M., Bonnefoy, M., Chauvin, G., et al. 2010, *Science (New York, N.Y.)*, 329, 57
- Lagrange, A. M., Gratadour, D., Chauvin, G., et al. 2009a, *A&A*, 493, L21
- Lagrange, A. M., Gratadour, D., Chauvin, G., et al. 2009b, *A&A*, 493, L21
- Lagrange, A. M., Rubini, P., Nowak, M., et al. 2020, *A&A*, 642, A18

- Lee, D. D. & Seung, H. S. 1999, *Nature*, 401, 788
- Leggett, S. K., Tremblin, P., Esplin, T. L., Luhman, K. L., & Morley, C. V. 2017, *ApJ*, 842, 118
- Lenzen, R., Close, L., Brandner, W., Biller, B., & Hartung, M. 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, 970–977
- Lestrade, J.-F. & Thilliez, E. 2015, *A&A*, 576, A72
- Liu, M. C. 2004, *Science*, 305, 1442
- Lyot, B. 1939, *MNRAS*, 99, 580
- Ma, B. & Ge, J. 2014, *MNRAS*, 439, 2781
- Macintosh, B., Graham, J. R., Barman, T., et al. 2015, *Science*, 350, 64
- Maire, A.-L., Rodet, L., Cantalloube, F., et al. 2019, *A&A*, 624, A118
- Maldonado, J. & Villaver, E. 2017, *A&A*, 602, A38
- Malo, L., Doyon, R., Lafrenière, D., et al. 2013, *ApJ*, 762, 88
- Mamajek, E. E. & Bell, C. P. M. 2014, *MNRAS*, 445, 2169
- Mamajek, E. E. & Hillenbrand, L. A. 2008, *ApJ*, 687, 1264
- Mamajek, E. E., Meyer, M. R., & Liebert, J. 2002, *AJ*, 124, 1670
- Margot, J.-L. 2015, *The Astronomical Journal*, 150, 185
- Marley, M. S., Fortney, J. J., Hubickyj, O., Bodenheimer, P., & Lissauer, J. J. 2007, *The Astrophysical Journal*, 655, 541
- Marois, C., Correia, C., Vigan, J.-P., & Currie, T. 2013, *Proceedings of the International Astronomical Union*, 8, 48
- Marois, C., Doyon, R., Nadeau, D., et al. 2005, *The Astronomical Society of the Pacific*, 117, 745
- Marois, C., Lafreniere, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, 641
- Marois, C., Lafreniere, D., Macintosh, B., & Doyon, R. 2008a, *ApJ*, 673
- Marois, C., Macintosh, B., Barman, T., et al. 2008b, *Science*, 322, 1348

- Marois, C., Macintosh, B., Barman, T., et al. 2008, *Science*, 322, 1348
- Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. 2010, *Nature*, 468, 1080
- Marshall, J. P., Milli, J., Choquet, É., et al. 2018, *ApJ*, 869, 10
- Matsuo, T., Shibai, H., Ootsubo, T., & Tamura, M. 2007, *ApJ*, 662, 1282
- Matthews, E., Hinkley, S., Vigan, A., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 480, 2757
- Mawet, D., Absil, O., Girard, J. H., et al. 2013, *The Messenger*, 152, 8
- Mawet, D., Hirsch, L., Lee, E. J., et al. 2019, *The Astronomical Journal*, 157, 33
- Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, *ApJ*, 792, 97
- Mawet, D., Pueyo, L., Lawson, P., et al. 2012, *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*
- Mawet, D., Riaud, P., Absil, O., & Surdej, J. 2005, *The Astrophysical Journal*, 633, 1191
- Mayor, M. & Queloz, D. 1995, *Nature*, 378, 355
- Meshkat, T., Mawet, D., Bryan, M. L., et al. 2017, *The Astronomical Journal*, 154, 245
- Metchev, S. A. & Hillenbrand, L. A. 2009, *The Astrophysical Journal Supplement Series*, 181, 62
- Milli, J., Banas, T., Mouillet, D., et al. 2016a, *SPIE proceedings*, 9909
- Milli, J., Gonzalez, R., Fluxa, P. R., et al. 2019, *arXiv e-prints*, arXiv:1910.13767
- Milli, J., Hiben, P., Christiaens, V., et al. 2016b, *A&A*, 597, L2
- 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
- Milli, J., Mouillet, D., Fusco, T., et al. 2017, *arXiv e-prints*, arXiv:1710.05417
- Mockus, J. 1994, *Journal of Global Optimization*, 4, 347
- Mockus, J., Tiesis, V., & Zilinskas, A. 1978, *The application of Bayesian methods for seeking the extremum*, Vol. 2 (Elsevier), 117–128

- Moor, A., Abraham, P., Derekas, A., et al. 2006, *The Astrophysical Journal*, 644, 525
- Moór, A., Pascucci, I., Kóspál, Á., et al. 2011, *ApJs*, 193, 4
- Morton, T. D., Bryson, S. T., Coughlin, J. L., et al. 2016, *The Astrophysical Journal*, 822, 86
- Mouillet, D., Larwood, J. D., Papaloizou, J. C. B., & Lagrange, A. M. 1997, *MNRAS*, 292, 896
- Mugnier, L. M., Cornia, A., Sauvage, J.-F., et al. 2009, *Journal of the Optical Society of America A*, 26, 1326
- Nair, V. & Hinton, G. E. 2010, *ICML*, eds. J. Furnkranz & T. Joachims (Madison, WI: Omnipress), 807
- Nielsen, E. L., De Rosa, R. J., Macintosh, B., et al. 2019a, *AJ*, 158, 13
- Nielsen, E. L., De Rosa, R. J., Macintosh, B., et al. 2019b, *AJ*, 158, 13
- Pairet, B., Cantalloube, F., Gomez Gonzalez, C., Absil, O., & Jacques, L. 2019, *Monthly Notices of the Royal Astronomical Society*, 487, 2262
- Pawellek, N., Wyatt, M., Matrà, L., Kennedy, G., & Yelverton, B. 2021, *MNRAS*, 502, 5390
- Pearce, T. D., Launhardt, R., Ostermann, R., et al. 2022, *A&A*, 659, A135
- Perryman, M. 2018, *The Exoplanet Handbook* (Cambridge University Press)
- Phillips, M. W., Tremblin, P., Baraffe, I., et al. 2020, *A&A*, 637, A38
- Pollack, J. B., Hubickyj, O., Bodenheimer, P., et al. 1996, *Icarus*, 124, 62
- Pueyo, L. 2016, *ApJ*, 824, 117
- Pueyo, L. 2018, *Direct Imaging as a Detection Technique for Exoplanets*, ed. H. J. Deeg & J. A. Belmonte, 10
- Pueyo, L., Crepp, J. R., Vasisht, G., et al. 2012, *The Astrophysical Journal Supplement Series*, 199, 6
- Quanz, S. P., Crossfield, I., Meyer, M. R., Schmalzl, E., & Held, J. 2015, *International Journal of Astrobiology*, 14, 279
- Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., & Marois, C. 1999, *Pub. Astron. Soc. Pacific.*, 111, 587

- Ragazzoni, R. & Farinato, J. 1999, *A&A*, 350, L23
- Rajan, A., Rameau, J., Rosa, R. J. D., et al. 2017, *The Astronomical Journal*, 154, 10
- Rameau, J., Nielsen, E. L., De Rosa, R. J., et al. 2016, *ApJL*, 822, L29
- Rauer, H., Catala, C., Aerts, C., et al. 04 September 2014, *Experimental Astronomy*
- Ren, B., Pueyo, L., Zhu, G. B., Debes, J., & Duchêne, G. 2018, *ApJ*, 852
- Ren, D., Dou, J., Zhang, X., & Zhu, Y. 2012, *ApJ*, 753, 99
- Rhee, J. H., Song, I., Zuckerman, B., & McElwain, M. 2007, *The Astrophysical Journal*, 660, 1556
- Roddier, F. & Roddier, C. 1988, in *European Southern Observatory Conference and Workshop Proceedings, Vol. 30, Very Large Telescopes and their Instrumentation, Vol. 2*, 667
- Rodriguez, D. R. & Zuckerman, B. 2012, *The Astrophysical Journal*, 745, 147
- Romero, C., Milli, J., Lagrange, A. M., et al. 2021, *A&A*, 651, A34
- Rousset, G., Fontanella, J. C., Kern, P., Gigan, P., & Rigaut, F. 1990, *A&A*, 230, L29
- Ruane, G., Ngo, H., Mawet, D., et al. 2019, *The Astronomical Journal*, 157, 118
- Ruffio, J.-B., Macintosh, B., Wang, J. J., & Pueyo, L. 2017, *ApJ*, 842
- Samland, M., Mollière, P., Bonnefoy, M., et al. 2017, *A & A*, 603, A57
- Santerne, A., Díaz, R. F., Moutou, C., et al. 2012, *A& A*, 545, A76
- Saumon, D. & Marley, M. S. 2008, *ApJ*, 689, 1327
- Saumon, D., Marley, M. S., Lodders, K., & Freedman, R. S. 2003, in *Brown Dwarfs*, ed. E. Martín, Vol. 211, 345
- Sauvage, J.-F., Costille, A., Fusco, T., et al. 2016, *Journal of Astronomical Telescopes, Instruments, and Systems*, 2, 1
- Schneider, G., Weinberger, A. J., Becklin, E. E., Debes, J. H., & Smith, B. A. 2008, *The Astronomical Journal*, 137, 53
- Schneider, J., Dedieu, C., Le Sidaner, P., Savalle, R., & Zolotukhin, I. 2011, *A&A*, 532, A79
- Serabyn, E., Huby, E., Matthews, K., et al. 2017, *The Astronomical Journal*, 153, 43

- Shack, R. B. & Platt, B. C. 1971, *J. Opt.Soc. Am.*, 61
- Shi, X., Chen, Z., Wang, H., & et al. 2015, NIPS (Cambridge, MA: MIT Press)
- Showman, A. P. & Kaspi, Y. 2013, *The Astrophysical Journal*, 776, 85
- Skrutskie, M., Jones, T., Hinz, P., et al. 2010, in *Ground-Based and Airborne Instrumentation for Astronomy III*, part 1 edn., *Proceedings of SPIE - The International Society for Optical Engineering* No. PART 1, ground-Based and Airborne Instrumentation for Astronomy III ; Conference date: 27-06-2010 Through 02-07-2010
- Smith, R., Churcher, L. J., Wyatt, M. C., Moerchen, M. M., & Telesco, C. M. 2008, *Astronomy & Astrophysics*, 493, 299
- Soummer, R. 2004, *The Astrophysical Journal*, 618, L161
- Soummer, R., Hagan, J. B., Pueyo, L., et al. 2011, *The Astrophysical Journal*, 741, 55
- Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJ*, 755
- Sparks, W. B. & Ford, H. C. 2002, *The Astrophysical Journal*, 578, 543
- Spiegel, D. S. & Burrows, A. 2012, *The Astrophysical Journal*, 745, 174
- Struve, O. 1952, *The Observatory*, 72, 199
- Sumi, T., Kamiya, K., Bennett, D. P., et al. 2011, *Nature*, 473, 349
- Tabernerero, H. M., Montes, D., & González Hernández, J. I. 2012, *A&A*, 547, A13
- Tan, X. & Showman, A. 2021, *Monthly Notices of the Royal Astronomical Society*, 502
- Tan, X. & Showman, A. P. 2017, *The Astrophysical Journal*, 835, 186
- Thompson, W. & Marois, C. 2021, *The Astrophysical Journal*, 161
- Torres, C. A. O., Quast, G. R., Melo, C. H. F., & Sterzik, M. F. 2008, in *Handbook of Star Forming Regions, Volume II*, ed. B. Reipurth, Vol. 5, 757
- Udry, S. 2010
- van Leeuwen, F. 2007, *A&A*, 474, 653
- Vican, L. 2012, *The Astronomical Journal*, 143, 135
- Vigan, A., Bonavita, M., Biller, B., et al. 2017, *A&A*, 603, A3
- Vigan, A., Fontanive, C., Meyer, M., et al. 2021, *A&A*, 651, A72

- Wahhaj, Z., Milli, J., Kennedy, G., et al. 2016, *A&A*, 596, L4
- Walker, G. A. H., Bohlender, D. A., Walker, A. R., et al. 1992, *ApJL*, 396, L91
- Wertz, O., Absil, O., Gómez González, C. A., et al. 2017, *A&A*, 598, A83
- West, A. A., Hawley, S. L., Bochanski, J. J., et al. 2008, *AJ*, 135, 785
- Wolszczan, A. & Frail, D. A. 1992, *Nature*, 355, 145
- Yip, K. H., Nikolaou, N., Coronica, P., et al. 2020, Pushing the Limits of Exoplanet Discovery via Direct Imaging with Deep Learning, 322–338
- Zhou, T. & Tao, D. 2011, in *International Conference on Machine Learning*
- Zorec, J. & Royer, F. 2012, *A&A*, 537, A120
- Zuckerman, B., Rhee, J. H., Song, I., & Bessell, M. S. 2011, *ApJ*, 732, 61
- Zuckerman, B. & Song, I. 2004, *Annual Review of Astronomy and Astrophysics*, 42, 685
- Zuckerman, B., Song, I., Bessell, M. S., & Webb, R. A. 2001, *ApJL*, 562, L87
- Zurlo, A., Vigan, A., Galicher, R., et al. 2016, *A&A*, 587, A57

Appendices

Appendix A

RSM detection map for direct
exoplanet detection in ADI sequences

A.1 Mathematical notations for the RSM map model

Table A.1: Description of the mathematical notations for the variables used in the RSM detection map computation.

Symbol	Dimension	Comments
RSM map		
X_{i_a}	TL_a	Vector of residuals for the annulus a
\mathbf{x}_{i_a}	$\theta \times \theta \times TL_a$	Patches of residuals centred on pixel i_a
F_{i_a}	TL_a	Realisation of a two-state Markov chain representing the state in which the system is for pixel i_a
\mathbf{m}	$\theta \times \theta$	Cropped planetary signal (off-axis PSF)
$\boldsymbol{\varepsilon}_{s,i_a}$	$2\theta \times \theta \times TL_a$	Error terms associated with the two regimes
S_{i_a}	TL_a	State in which the system is for every pixel i_a
ξ_{s,i_a}	$2 \times TL_a$	Probability associated with state s for every pixel i_a
η_{s,i_a}	$2 \times TL_a$	Likelihood of being in each state for every pixel i_a
$p_{q,s}$	2×2	Transition probabilities between the regimes
μ	1	Mean of the residuals contained in an annulus a , with width equal to θ
σ	1	Standard deviation of the residuals contained in an annulus a , with width equal to θ
β	1	Parameter representing the intensity of the planetary signal in the cube of residual
a	1	Annulus index
L_a	1	Number of pixels included in the annulus a
T	1	Number of frames in the cube of residuals
i_a	1	Index associated with every pixel from every frame in the annulus a (ranges from 1 to TL_a)
θ	1	Angular size of the considered planetary signal (set to $1\lambda/D$)

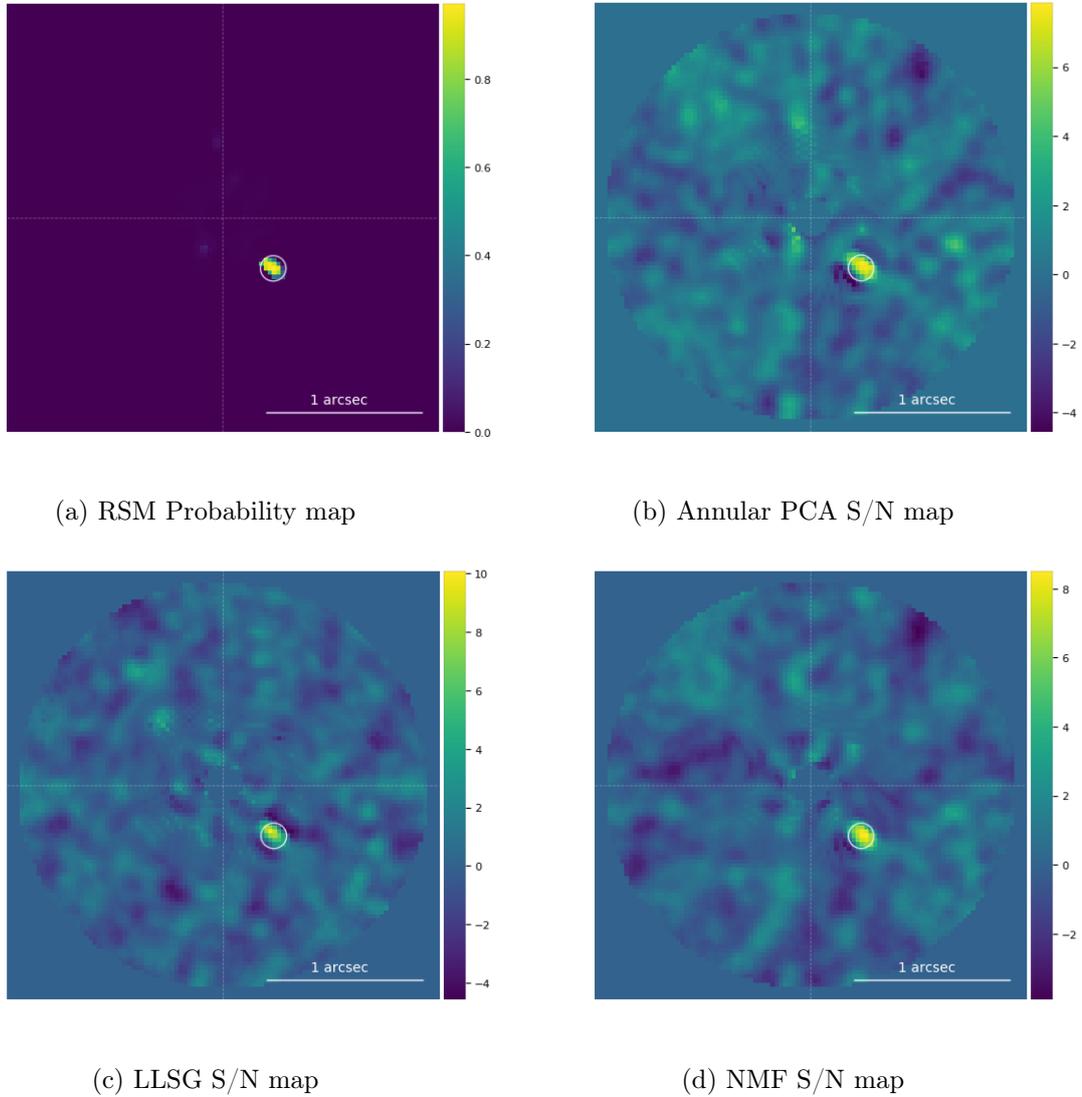
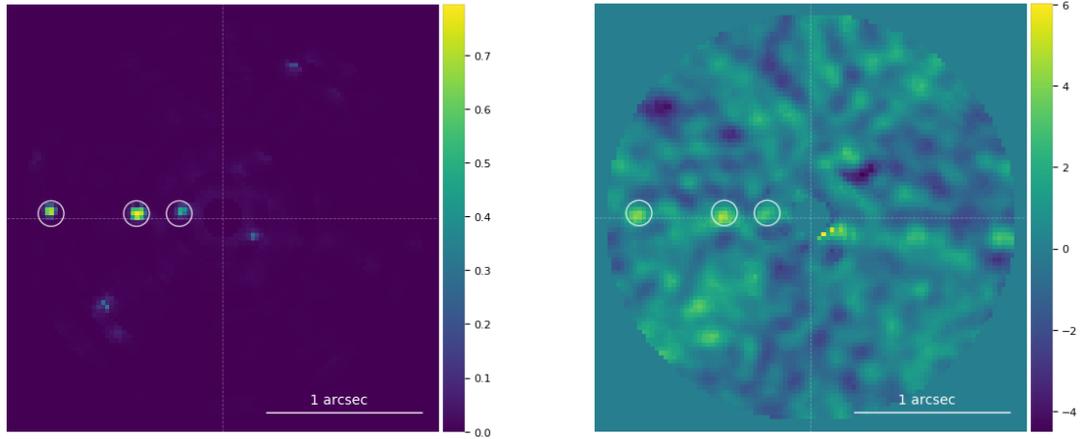
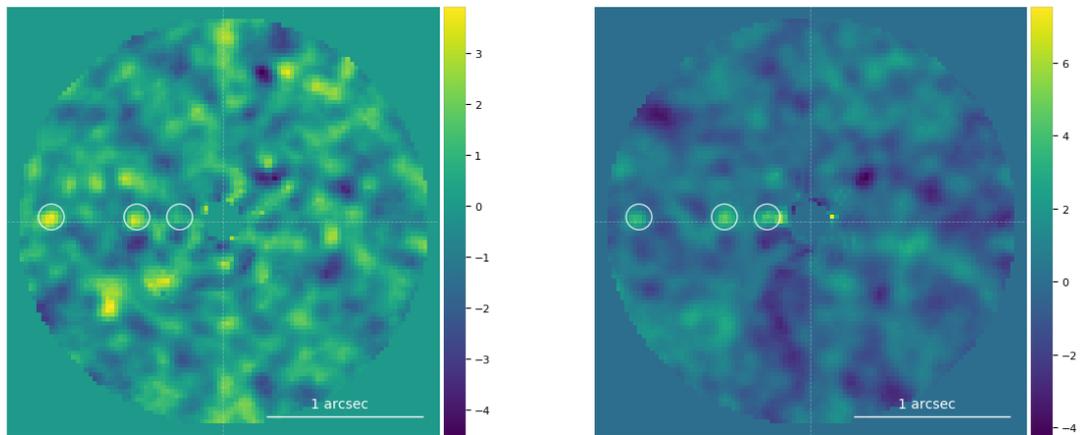
A.2 NACO β Pictoris

Figure A.1: Probability map obtained for the NACO β Pictoris data set, with the RSM using a Gaussian distribution along the S/N map generated with the cube of residuals obtained with Annular PCA, LLSG and NMF with respectively 20 components for the Annular PCA and the NMF and a rank of 5 for the LLSG. The colour scale indicates the probability for the RSM map and the signal-to-noise ratio for the three S/N maps. The maps are centred on the star β Pictoris while β Pictoris b is identified by the white circle in the lower left quadrant.



(a) RSM Probability map

(b) Annular PCA S/N map



(c) LLSG S/N map

(d) NMF S/N map

Figure A.2: Detection map obtained after injecting three fake companions in the NACO β Pictoris reference cube used for the ROC estimation, at a distance of 2, 4 and $8 \lambda/D$ with respectively a contrast of 3.3×10^{-4} , 0.4×10^{-4} and 1.7×10^{-5} . The colour scale indicates the probability for the RSM map and the signal-to-noise ratio for the three S/N maps. The maps are centred of the star β Pictoris while the fake companions are identified by the white circles.

Appendix B

Improving the RSM map algorithm:
PSF forward modelling and optimal
selection of PSF subtraction
techniques

B.1 Mathematical notations for the RSM KLIP FM and LOCI FM

Table B.1: Description of the mathematical notations for the variables used in the RSM KLIP FM, and LOCI FM computation.

Symbol	Dimension	Comments
KLIP FM		
\mathbf{i}	N^{pix}	Vectorised science image/annulus before speckle subtraction
\mathbf{m}	N^{pix}	Vectorised planetary signal (off-axis PSF) inside the selected annulus
\mathbf{p}	N^{pix}	Vectorised forward-modelled planetary signal inside the selected annulus
\mathbf{x}	N^{pix}	Vectorised processed annulus after speckle subtraction
$\mathbf{p}_{i_a}/\mathbf{p}_j$	$\theta \times \theta$	Derotated and cropped forward-modelled planetary signal for pixel i_a /frame j
\mathbf{R}	$N^R \times N^{pix}$	Reference library matrix
\mathbf{V}_K	$K \times N^R$	Eigenvector matrix of the covariance matrix $\mathbf{R}\mathbf{R}^\top$ with $\mathbf{V}_K = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k]$
$\mathbf{\Lambda}$	$N^R \times N^{pix}$	Diagonal matrix with the eigenvalues of the covariance matrix $\mathbf{R}\mathbf{R}^\top$ with $\mathbf{\Lambda} = \text{diag}(\mu_1, \mu_2, \dots, \mu_k)^\top$
\mathbf{Z}_K	$K \times N^R$	Karhunen-Loève image matrix
$\Delta\mathbf{Z}_K$	$K \times N^R$	Perturbation of the Karhunen-Loève image matrix
\mathbf{M}	$N^R \times N^{pix}$	Planet signal component in the reference library \mathbf{R}
\mathbf{C}_{MR}	$N^R \times N^R$	Covariance Matrix between \mathbf{M} and \mathbf{R}
N^{pix}	1	Number of pixels in the selected annulus of width equal to one FWHM
N^R	1	Number of reference frames used for the speckle field computation
K	1	Number of principal components used for the speckle field computation
LOCI FM		
c_k	N^R	Factors of the linear combination used to model the speckle field
\mathbf{o}_i	N^{pix}	Selected annulus in the science image used for the estimation of c_k
\mathbf{o}_k	N^{pix}	Selected annulus in frame k of the reference library used for the estimation of c_k
\mathbf{r}_k	N^{pix}	Selected annulus in frame k of the reference library used for the computation of the speckle field
\mathbf{m}_i	N^{pix}	Selected annulus in frame corresponding to the science image in the planetary signal matrix
\mathbf{m}_k	N^{pix}	Selected annulus in frame k of the planetary signal
N^{pix}	1	Number of pixels in the annulus with width of three FWHM used for the factors estimation
N^{pix}	1	Number of pixels in the annulus with width of one FWHM used for speckle field computation
N^R	1	Number of reference frames used for the speckle field computation

B.2 Forward-backward model and β parameter estimation

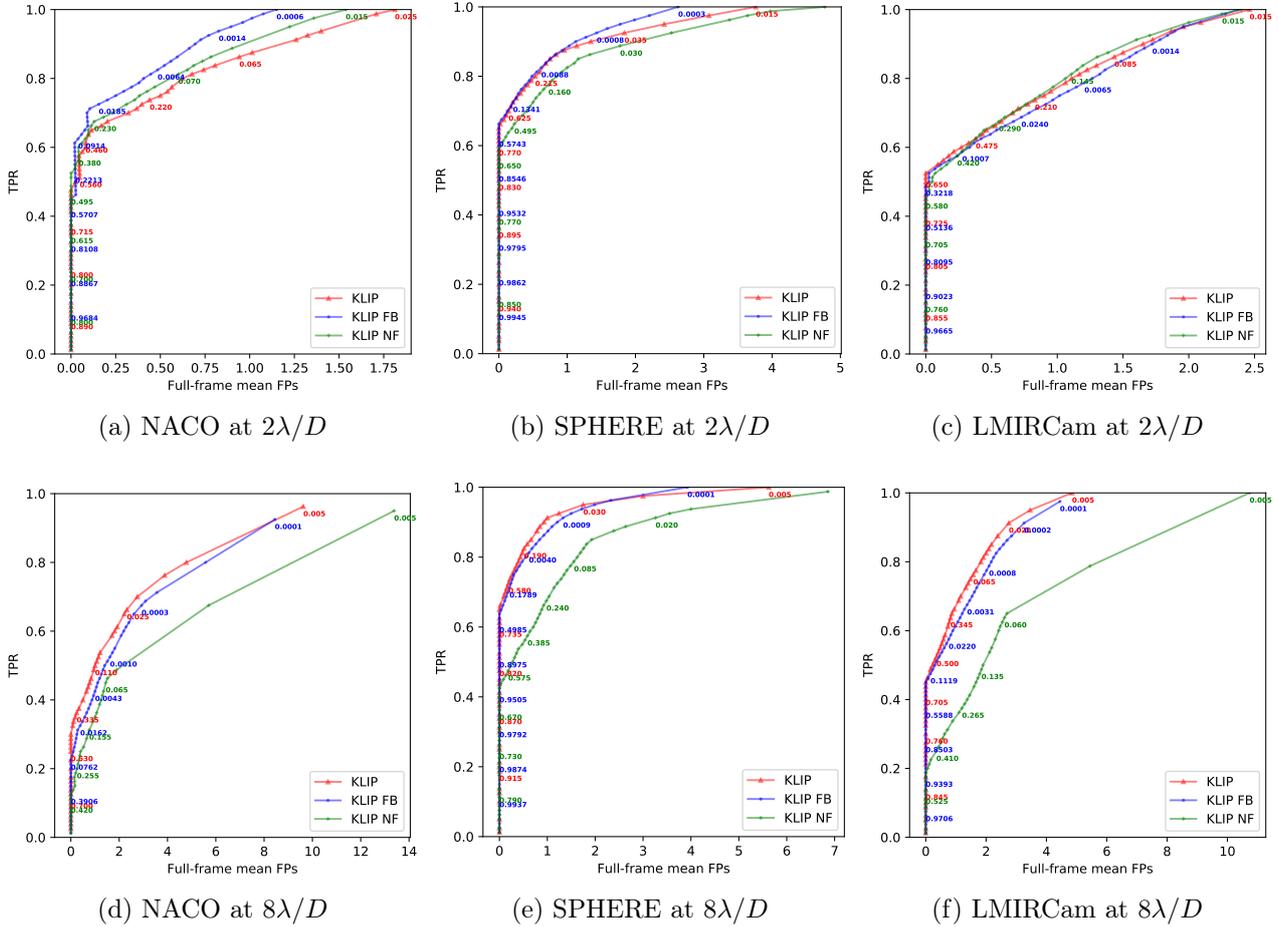


Figure B.1: ROC curves for the NACO, SPHERE, and LMIRCam data sets, with respectively the KLIP RSM map using the Gaussian maximum likelihood for the pre-optimisation of the flux parameter β (red), the KLIP RSM map with no flux pre-optimisation (NF), which relies on the standard maximum likelihood used in the original RSM map for the estimation of flux parameter β (blue), and the KLIP RSM map using the forward-backward approach for the probability estimation instead of the original forward approach (green).

B.3 Crop size comparison for FM LOCI RSM

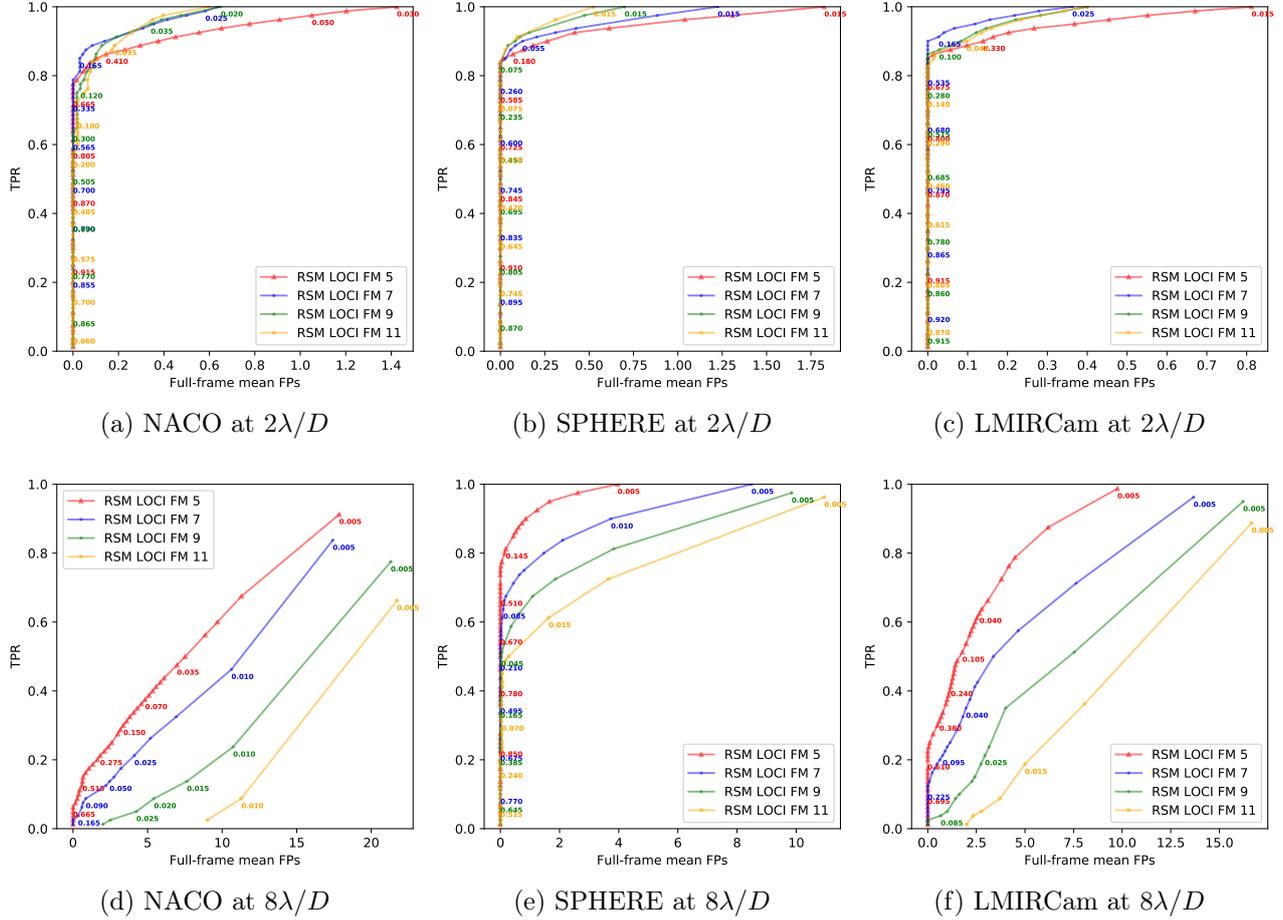


Figure B.2: ROC curves for the NACO, SPHERE, and LMIRCam data sets, with the LOCI-FM RSM map using respectively a crop size for the forward modelled PSF of 5 (red), 7 (blue), 9 (green), 11 (orange) pixels ($\text{FWHM} \approx 5$ pixels for all three data sets).

B.4 PSF-subtraction techniques combinations

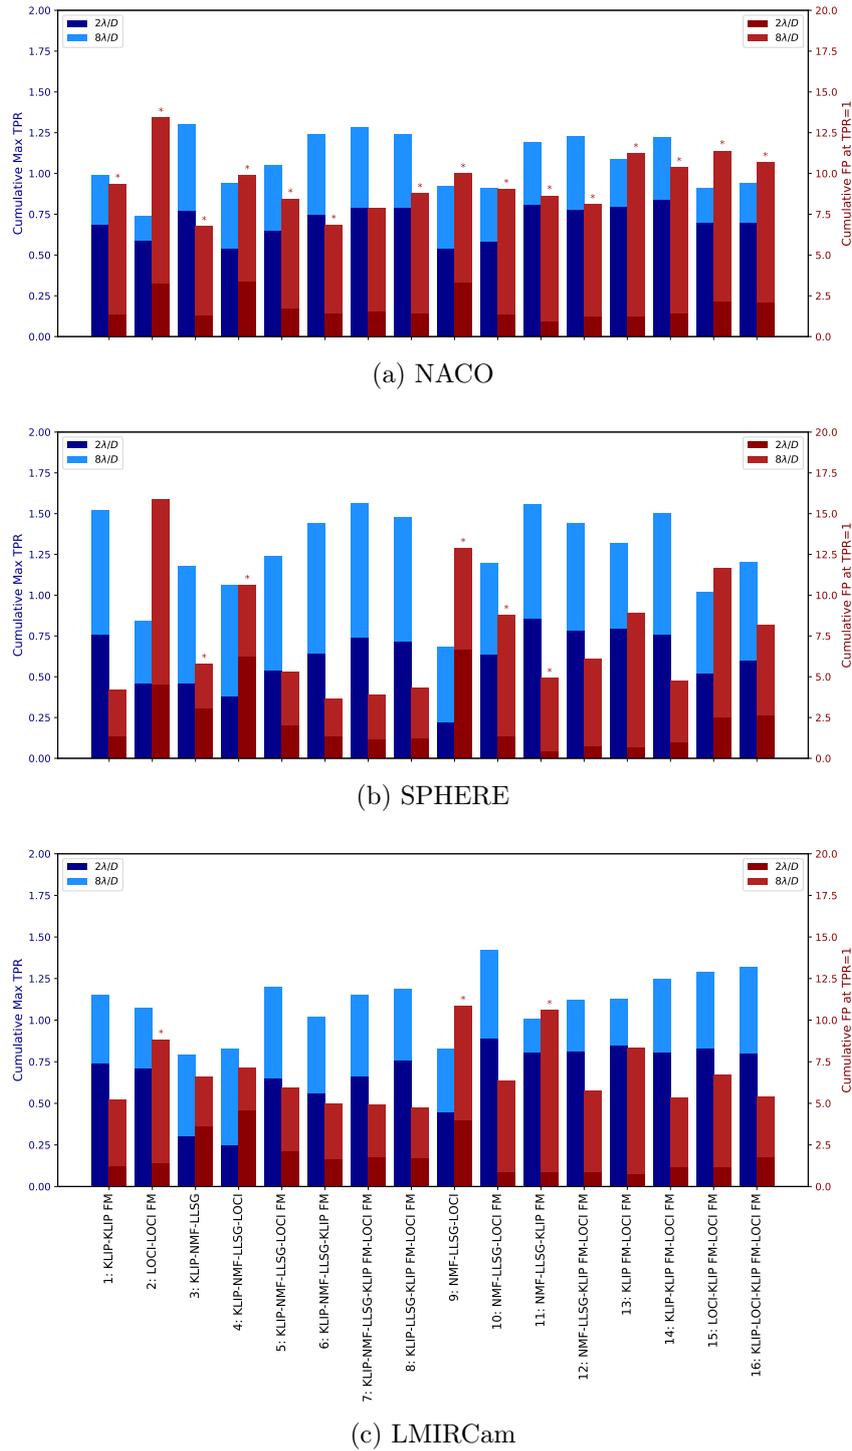


Figure B.3: Cumulative maximum TPR reached without any FP (blue), and cumulative average number of FPs inside the entire frame at TPR=1 (red) for 16 different combinations of PSF subtraction techniques used to generate the RSM map and two radial distances. The dark-coloured bars are obtained at a radial distance of $2\lambda/D$ while the light-coloured bars correspond to $8\lambda/D$ for, respectively, the NACO (top), SPHERE (middle), and LMIRCam (bottom) data sets. The asterisks for some values of FP at TPR=1 indicate that a TPR of 1 has not been reached at a distance of $8\lambda/D$ and that the smallest probability threshold was chosen instead (highest TPR). A high performance for a combination of PSF subtraction techniques corresponds to a tall blue histogram alongside a short red histogram.

Table B.2: Maximum TPR reached without any FP and average number of FPs inside the entire frame at TPR=1 for 16 different combinations of PSF subtraction techniques used to generate the RSM map. The left values are obtained at a radial distance of $2\lambda/D$ while the right values correspond to $8\lambda/D$ for respectively the NACO, SPHERE and LMIRCam data sets. The asterisks for some values of FP at TPR=1 indicate that a TPR of 1 has not been reached and that the smallest probability threshold has been taken (highest TPR).

Selected PSF subtraction techniques	NACO		SPHERE		LMIRCam	
	Max TPR	Max FP	Max TPR	Max FP	Max TPR	Max FP
1: KLIP-KLIP FM	0.69/0.30	1.36/8.00*	0.76/0.76	1.34/2.86	0.74/0.41	1.20/4.01
2: LOCI-LOCI FM	0.59/0.15	3.26/10.15*	0.46/0.38	4.56/11.34	0.71/0.36	1.40/7.43*
3: KLIP-NMF-LLSG	0.66/0.53	1.28/5.47*	0.46/0.72	3.06/2.72*	0.30/0.49	3.63/2.99
4: KLIP-NMF-LLSG-LOCI	0.54/0.40	3.36/6.53*	0.38/0.68	6.26/4.36*	0.25/0.58	4.59/2.56
5: KLIP-NMF-LLSG-LOCI FM	0.65/0.40	1.74/6.71*	0.54/0.70	2.02/3.28	0.65/0.55	2.15/3.80
6: KLIP-NMF-LLSG-KLIP FM	0.75/0.49	1.42/5.43*	0.64/0.80	1.34/2.28	0.56/0.46	1.64/3.31
7: KLIP-NMF-LLSG-KLIP FM-LOCI FM	0.79/0.49*	1.54/6.35	0.74/0.82	1.20/2.70	0.66/0.49	1.77/3.15
8: KLIP-LLSG-KLIP FM-LOCI FM	0.79/0.45	1.42/7.35*	0.72/0.76	1.24/3.08	0.76/0.43	1.71/3.01
9: NMF-LLSG-LOCI	0.54/0.38	3.31/6.68*	0.22/0.46	6.68/6.20*	0.45/0.38	3.96/6.89*
10: NMF-LLSG-LOCI FM	0.58/0.33	1.39/7.65*	0.64/0.56	1.38/7.40*	0.89/0.53	0.91/5.46
11: NMF-LLSG-KLIP FM	0.81/0.38	0.95/7.70*	0.86/0.70	0.44/4.47*	0.81/0.20	0.84/9.76*
12: NMF-LLSG-KLIP FM-LOCI FM	0.78/0.45	1.24/6.86*	0.78/0.66	0.76/5.34	0.81/0.31	0.89/4.91
13: KLIP FM-LOCI FM	0.80/0.29	1.22/10.01*	0.80/0.52	0.70/8.24	0.85/0.28	0.75/7.59
14: KLIP-KLIP FM-LOCI FM	0.84/0.38	1.41/8.96*	0.76/0.74	1.00/3.78	0.81/0.44	1.19/4.18
15: LOCI-KLIP FM-LOCI FM	0.70/0.21	2.18/9.23*	0.52/0.50	2.52/9.14	0.83/0.46	1.18/5.53
16: KLIP-LOCI-KLIP FM-LOCI FM	0.70/0.24	2.08/8.60*	0.60/0.60	2.66/5.54	0.80/0.52	1.74/3.65

B.5 Contrast curve for the SPHERE data set using KLIP RSM

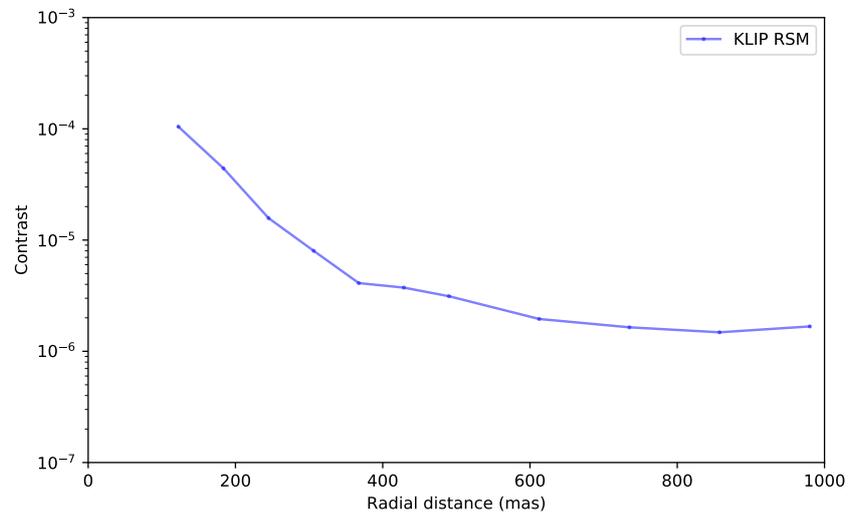


Figure B.4: Contrast curves for the SPHERE data set using KLIP RSM with 20 principal components and a FOV rotation expressed in terms of FWHM of 0.3. The region $[2\lambda/D, 16\lambda/D]$ has been considered to get the contrast in the first arcsecond.

Appendix C

Auto-RSM: an automated parameter selection algorithm for the RSM map exoplanet detection algorithm

C.1 Mathematical notations for auto-RSM

Table C.1: Description of the mathematical notations for the variables used in the Bayesian optimisation algorithm, the RSM detection map and auto-RSM optimisation framework.

Symbol	Dimension	Comments
Bayesian optimisation		
$\mathcal{O}_{1:t}$	t	Set of observations of the loss function
t	1	Number of tested sets of parameters
$\mathbf{p}_{1:t}$	$n_p \times t$	Tested sets of parameters
$f(\mathbf{p})$	1	Loss function evaluated with the set of parameters \mathbf{p}
$\mathcal{GP}(m(\mathbf{p}_{1:t}), \mathbf{K})$	1	Gaussian process returning the mean and variance of a Gaussian distribution over the possible values of f at \mathbf{p}
$m(\mathbf{p})$	1	Mean of the Gaussian distribution of the loss function at \mathbf{p}
\mathbf{K}	1	Covariance function of the tested set of parameters $\mathbf{p}_{1:t}$
$\mu(\mathbf{p}_{t+1})$	1	Mean of the Gaussian posterior distribution at \mathbf{p}_{t+1}
$\sigma(\mathbf{p}_{t+1})$	1	Variance of the Gaussian posterior distribution at \mathbf{p}_{t+1}
$\hat{\mathbf{p}}$	1	Set of parameters providing the current known maximum value of the loss function $f(\hat{\mathbf{p}})$

Table C.2: Description of the mathematical notations for the variables used in the Bayesian optimisation algorithm, the RSM detection map and auto-RSM optimisation framework.

Symbol	Dimension	Comments
Auto-RSM		
\mathbf{i}	N^{pix}	Vectorised science image/annulus before speckle subtraction
\mathbf{m}	N^{pix}	Vectorised planetary signal (off-axis PSF) inside the selected annulus
a_{max}	1	Largest angular separation considered for the detection map computation
$C_{a,m,c}$	$L_a \times N_{technique} \times N_{sequence}$	Average contrast obtained with the optimal set of parameters
$M_{a,m,c}$	$L_a \times N_{technique} \times N_{sequence}$	Position of the median contrast aperture within the selected annulus
$P_{a,m,c}^{PSF}$	$L_a \times N_{technique} \times N_{sequence}$	Optimal set of parameters for the reference PSF computation
$P_{a,m,c}^{PSF,*}$	$L_a \times N_{technique} \times N_{sequence}$	Optimal set of parameters for the reference PSF smoothed via moving average
$P_{a,m,c}^{RSM}$	$L_a \times N_{technique} \times N_{sequence}$	Optimal set of parameters for the computation of the RSM map
$P_{a,m,c}^{RSM,*}$	$L_a \times N_{technique} \times N_{sequence}$	Optimal set of parameters for the computation of the RSM map interpolated via RBF
\mathbf{Y}^a	L_a	Set of likelihood time series available (one per couple of ADI sequence and PSF-subtraction technique) for the computation of the RSM map
$Y_{c,m}^a$	$L_a \times N_{technique} \times N_{sequence}$	Likelihood time series
Y_{h*}^a	$L_a \times T$	Likelihood time series maximising the RSM performance index at a given iteration
\mathbf{Z}^a	L_a	Selected likelihood time series for the computation of the final detection map
T_a	L_a	Annulus-wise thresholds computed with flipped parallactic angles
T_a^*	L_a	Smoothed annulus-wise thresholds computed with flipped parallactic angles

C.2 Computation of the expected improvement and update of posterior probability moments

The objective of the expected improvement approach is to estimate the magnitude of the improvement that a set of parameters can potentially yield in terms of loss function value. As the true maximum value of the loss function $f(\mathbf{p}^*)$ is not known, Mockus et al. (1978) propose maximising the expected improvement with respect to a known maximum $f(\hat{\mathbf{p}})$. They define the improvement function as

$$I(\mathbf{p}_{t+1}) = \max \{0, f(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}})\}, \quad (\text{C.1})$$

where $I(\mathbf{p}_{t+1})$ is positive when the prediction is higher than the current maximum loss function value and zero otherwise. The new set of parameters \mathbf{p}_{t+1} is found by maximising the expected improvement, as follows:

$$\mathbf{p}_{t+1} = \operatorname{argmax}_{\mathbf{p}_{t+1}} \mathbb{E}(\{0, f(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}})\} \mid \mathcal{O}_{1:t}). \quad (\text{C.2})$$

The likelihood of improvement I when considering the Gaussian process giving the posterior distribution is then

$$\frac{1}{\sqrt{2\pi}\sigma(\mathbf{p}_{t+1})} \exp\left(-\frac{(\mu(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}}) - I)^2}{2\sigma^2(\mathbf{p}_{t+1})}\right), \quad (\text{C.3})$$

with $\mu(\mathbf{p}_{t+1})$ and $\sigma(\mathbf{p}_{t+1})$ being the mean and standard deviation, respectively, of the posterior probability $f(\mathbf{p}_{1:t}) \sim \mathcal{N}(0, \mathbf{K})$ for the new set of parameters \mathbf{p}_{t+1} . The expected improvement is then simply the integral over this likelihood function:

$$\mathbb{E}(I(\mathbf{p}_{t+1})) = \int_0^\infty \frac{I}{\sqrt{2\pi}\sigma(\mathbf{p}_{t+1})} \exp\left(-\frac{(\mu(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}}) - I)^2}{2\sigma^2(\mathbf{p}_{t+1})}\right) dI, \quad (\text{C.4})$$

which gives after integration by part

$$\mathbb{E}(I(\mathbf{p}_{t+1})) = \sigma(\mathbf{p}_{t+1}) \left[\frac{\mu(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}})}{\sigma(\mathbf{p}_{t+1})} \Phi\left(\frac{\mu(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}})}{\sigma(\mathbf{p}_{t+1})}\right) + \phi\left(\frac{\mu(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}})}{\sigma(\mathbf{p}_{t+1})}\right) \right]. \quad (\text{C.5})$$

Considering the improvement function definition, we obtain the expression of Eq. 4.7,

$$\text{EI}(\mathbf{p}_{t+1}) = \begin{cases} (\mu(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}}))\Phi(Z) + \sigma(\mathbf{p}_{t+1})\phi(Z) & \text{if } \sigma(\mathbf{p}_{t+1}) > 0 \\ 0 & \text{if } \sigma(\mathbf{p}_{t+1}) = 0 \end{cases}, \quad (\text{C.6})$$

with $Z = \frac{\mu(\mathbf{p}_{t+1}) - f(\hat{\mathbf{p}})}{\sigma(\mathbf{p}_{t+1})}$.

We see that the computation of EI requires an estimation of the mean $\mu(\mathbf{p}_{t+1})$ and standard deviation $\sigma(\mathbf{p}_{t+1})$ of the posterior probability of $f(\mathbf{p}_{1:t})$. Starting from the posterior probability $f(\mathbf{p}_{1:t}) \sim \mathcal{N}(0, \mathbf{K})$ and taking into account the new set of parameters \mathbf{p}_{t+1} we get

$$\begin{pmatrix} f(\mathbf{p}_{1:t}) \\ f(\mathbf{p}_{t+1}) \end{pmatrix} = \mathcal{N}\left(0, \begin{pmatrix} \mathbf{K} & \mathbf{k} \\ \mathbf{k}^T & k(\mathbf{p}_{t+1}, \mathbf{p}_{t+1}) \end{pmatrix}\right), \quad (\text{C.7})$$

where $\mathbf{k} = \{k(\mathbf{p}_1, \mathbf{p}_{t+1}), k(\mathbf{p}_2, \mathbf{p}_{t+1}), \dots, k(\mathbf{p}_t, \mathbf{p}_{t+1})\}$.

We then get the following expression for the posterior distribution using the Sherman-Morrison-Woodbury formula:

$$P(f(\mathbf{p}_{t+1}) \mid \mathcal{O}_{1:t}, \mathbf{p}_{t+1}) = \mathcal{N}(\mu(\mathbf{p}_{t+1}), \sigma^2(\mathbf{p}_{t+1})), \quad (\text{C.8})$$

with the mean and variance given by

$$\mu(\mathbf{p}_{t+1}) = \mathbf{k}^T \mathbf{K}^{-1} f(\mathbf{p}_{1:t}), \quad (\text{C.9})$$

$$\sigma^2(\mathbf{p}_{t+1}) = k(\mathbf{p}_{t+1}, \mathbf{p}_{t+1}) - \mathbf{k}^T \mathbf{K}^{-1} \mathbf{k}. \quad (\text{C.10})$$

C.3 Average contrast computation via multiple fake companion injections

In this section we aim to assess the validity of our approximation when computing the average contrast by considering the agreement between the average contrasts generated using multiple injections and sequential injections. When relying on multiple injection, the self- and over-subtraction associated with an injected fake companion may affect neighbouring apertures, especially at small angular separations. We impose, for multiple injections, a minimal separation of two FWHMs between the positions of two injected fake companions in order to reduce the impact of these interferences on the es-

timization of the average contrast. The number of injected fake companions is therefore limited to half the number of apertures contained in a given annulus with a maximum of eight fake companions, which should provide a reliable estimate of the speckle field within the annulus while limiting the interference between fake companions. As can be seen from Figure C.1, the intensity patterns for multiple injections are similar to the ones observed for sequential injections, with the companions having the smallest or largest flux positioned at the same azimuth.

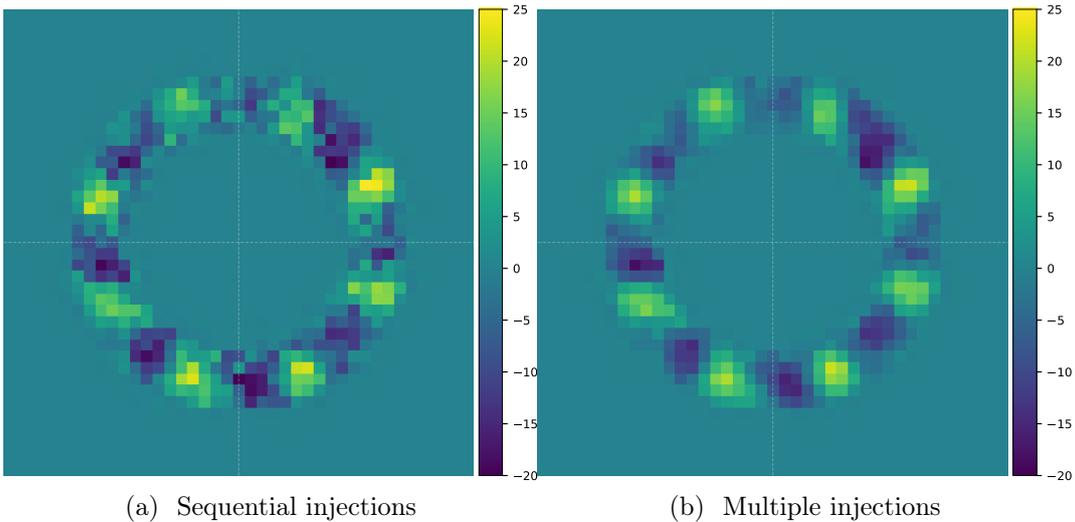


Figure C.1: Comparison of the recovered intensities of eight fake companions injected sequentially or at once, at a radial distance of $2 \lambda/D$, using the SPHERE 1 data set of the EIDC, and relying on annular PCA to generate the reference PSF with a number of principal components equal to 20.

Figure C.2 provides the evolution of the average contrast computed with the sequential and multiple injections for an increasing number of fake companions. As expected the average contrast does not vary significantly for the sequential injections. However, for multiple injections, the average contrast starts to strongly diverge for distances between the injected positions of neighbouring companions below two FWHMs. A distance of two FWHMs corresponds to 9, 18, and 33 companions for an angular distance of 2, 4, and $8 \lambda/D$, respectively. Looking at Figure C.2, we see that for eight injected fake companions, the average contrasts generated with the multiple and sequential injections are very similar.

Besides the distance between the average contrasts generated by the two approaches, the behaviour of these average contrasts when modifying the parameters of the PSF-subtraction techniques is the most important element, as it drives the optimal parameter selection. We computed the average contrasts for several different numbers of principal components in the case of annular PCA to determine if the behaviour of the contrast curves generated with multiple and sequential injections was

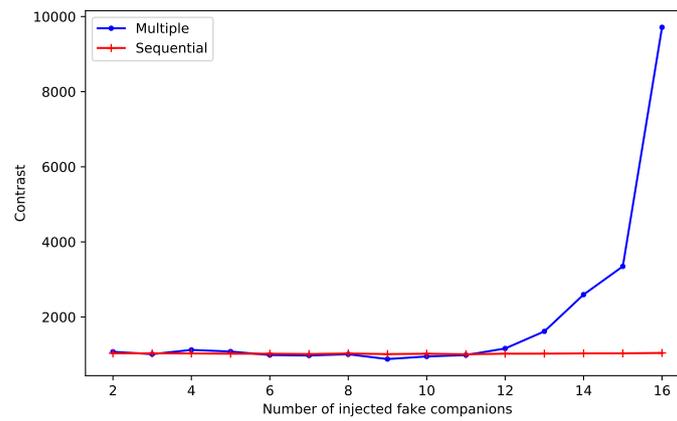
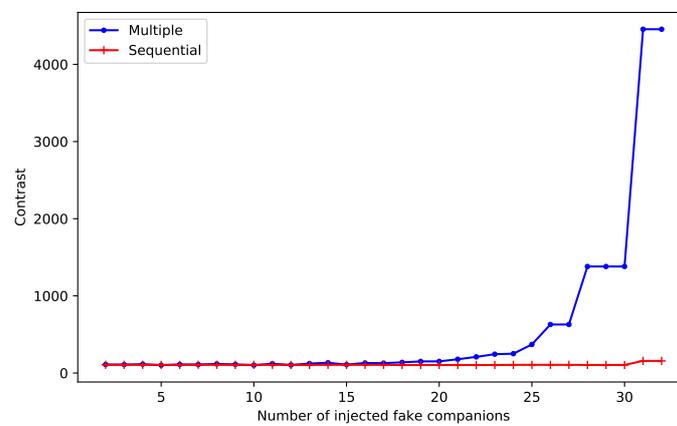
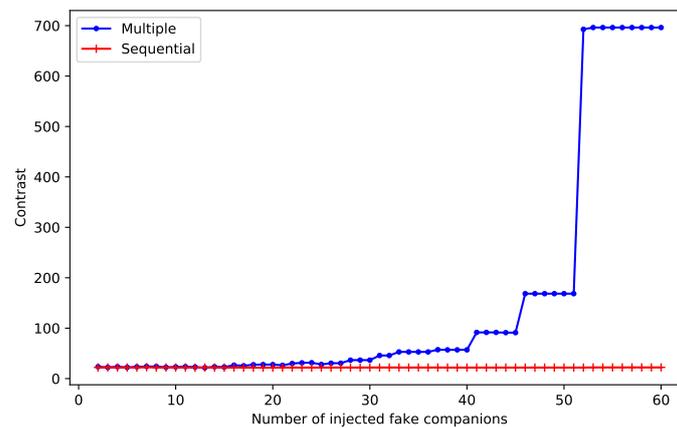
(a) $2 \lambda/D$ (b) $4 \lambda/D$ (c) $8 \lambda/D$

Figure C.2: Comparison of the average contrasts obtained with sequential and multiple injections for an increasing number of injected fake companions. The curves have been computed at three different angular separations (2 , 4 , and $8 \lambda/D$), using the SPHERE 1 data set of the EIDC, and relying on annular PCA to generate the reference PSF with a number of principal components equal to 20.

similar. Figure C.3 shows the evolution of the average contrast with the number of principal components used for the reference PSF computation for different angular separations. Although there exists a gap between both curves, the two curves seem to evolve in parallel, especially for smaller angular distances. We observe high Pearson correlations between the curves generated with multiple and sequential injections, with a correlation of 0.996, 0.992, and 0.704 for an angular distance of respectively 2, 4, and $8 \lambda/D$. This seems to indicate that the same set of parameters will minimise the average contrast and confirm the validity of our approximation when relying on multiple injections to compute the average contrast.

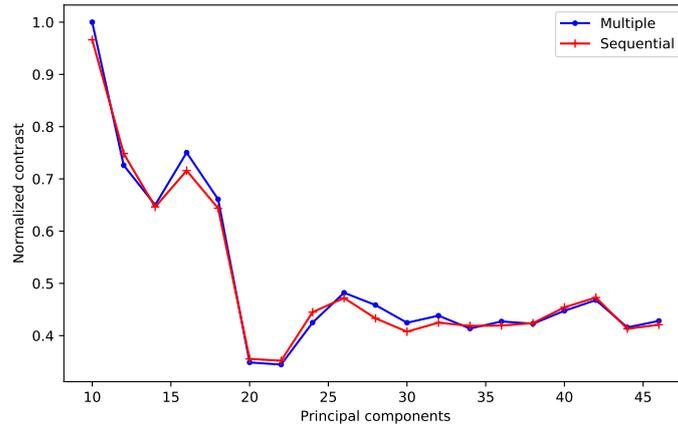
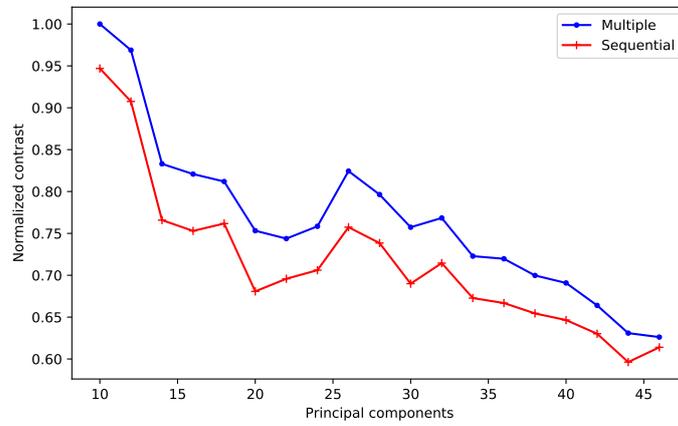
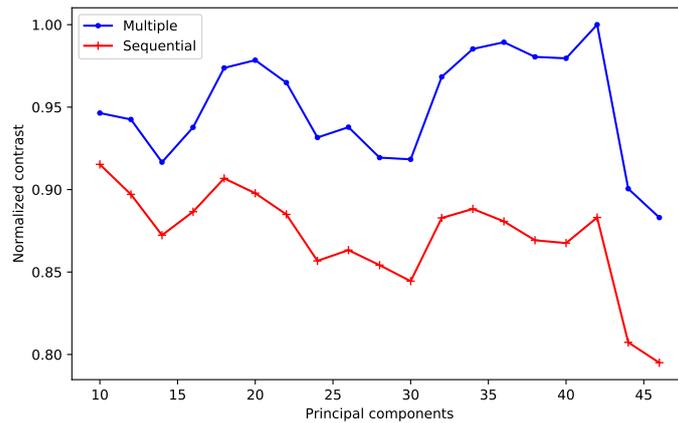
(a) $2 \lambda/D$ (b) $4 \lambda/D$ (c) $8 \lambda/D$

Figure C.3: Comparison of the average contrasts obtained with sequential and multiple injections for a range of principal components used by the annular PCA to generate the reference PSF (between 10 and 45 principal components). The curves were computed at three different angular separations (2 , 4 , and $8 \lambda/D$) using the SPHERE 1 data set of the EIDC.

C.4 Auto-S/N

C.4.1 Algorithm definition

We define in this section, the auto-S/N algorithm which is derived from the auto-RSM framework. The first step of the auto-RSM algorithm, i.e. the parameter optimisation of the PSF-subtraction techniques, is used to generate an optimised cube of residuals for every considered PSF-subtraction technique. As in the case of auto-RSM, the auto-S/N aims to combine the obtained cubes of residuals to generate a final detection map. As the cubes of residuals generated by the different PSF-subtraction techniques have their own noise distribution, a simple mean-combination is not possible. A simple way to overcome this limitation is to mean-combine the S/N maps instead of the cubes of residuals. As the dissimilarities in the noise structure of the different cubes of residuals are reflected in their respective S/N maps, part of the residual speckle noise should average out. The main difficulty pertains to the proper definition of the throughput to estimate the contrast used for the optimal selection, as we are combining S/N maps.

Considering the detection map obtained by averaging N different S/N maps, each pixels S/N is defined as:

$$\text{S/N}_{i_a} = \frac{1}{N} \sum_{j=1}^N \frac{F_{i_a,j}}{N_{a,j}} \quad (\text{C.11})$$

$$= \frac{1}{N} \frac{\sum_{j=1}^N F_{i_a,j} \prod_{k \neq j}^N N_{a,k}}{\prod_{k=1}^N N_{a,k}}, \quad (\text{C.12})$$

with $F_{i_a,j}$ the flux associated with the aperture centred on pixel i_a , where a is the considered annulus in the mean-combined de-rotated cube of residuals generated with the PSF-subtraction technique j , and $N_{a,j}$ is the associated noise computed in annulus a . Following this expression, the throughput obtained from the injection of a fake companion at pixel i_a is given in the case of a combination of N S/N maps:

$$\text{throughput}_{i_a} = \frac{\sum_{j=1}^N \text{IF}_{i_a,j} \prod_{k \neq j}^N N_{a,k}}{\sum_{j=1}^N \text{RF}_{i_a,j} \prod_{k \neq j}^N N_{a,k}}, \quad (\text{C.13})$$

where IF stands for injected flux and RF for retrieved flux. The throughput becomes simply a sum of fluxes weighted by the noises of the other considered PSF-subtraction techniques. This implies that the throughput associated with a less noisy mean-combined de-rotated cube of residuals has a higher weight as both the injected flux and retrieved flux are multiplied by larger noise values than the others. The noise appearing in the expression of the contrast (see eq. 4.1) is then computed as

the noise averaged over the different S/N maps for the considered angular separation a . Similarly to the parameter optimisation for the PSF-subtraction techniques, fake companions are injected at different azimuths to obtain an average contrast. The obtained average contrast is then used to select the optimal set of S/N maps either via the bottom-up or the top-down approach described respectively in Tables C.3 and C.4. As for the auto-RSM, the auto-S/N can also use either the full-frame or the annular optimisation mode.

C.4.2 Performance assessment

We follow the same procedure as the one proposed in Section 4.4 to assess the performance of different parametrisations of the auto-S/N. We consider the full-frame case as well as the annular and annular full-frame optimisation mode along the bottom-up and top-down approach for the PSF-subtraction techniques selection, similarly to the auto-RSM performance assessment in Section 4.4. The S/N maps combinations generated by the four parametrisations of the auto-S/N may be found in Figure C.4 and C.5. As can be seen from these graphs, the auto-S/N clearly performs better than the baseline proposed in (Cantalloube et al. 2020b), although the results are degraded compared to the ones obtained with the auto-RSM (see Figure 4.4). This degraded performance was expected, considering the higher performance of RSM probability maps compared to standard S/N maps as demonstrated in Dahlqvist et al. (2020).

These results are confirmed by the performance metrics shown in Figures C.6 and C.7, with both a lower TPR and a higher FPR for the auto-S/N versions compared to the full-frame-bottom-up forward auto-RSM. As in the case of the auto-RSM, the full-frame auto-RSM versions perform better than the annular and hybrid annular full-frame versions.

Considering the shorter computation time compared to auto-RSM and the better performance compared to standard S/N based PSF-subtraction techniques, the auto-S/N can be considered as an interesting alternative to the auto-RSM for large surveys. The auto-S/N may also represent a good complement to the auto-RSM¹, as it may lead to the identification of planetary signals missed by the auto-RSM as illustrated by the LMIRCam-3 data set (to be compared with Figure 4.4).

¹ The computation time is further reduced when the auto-RSM has already been applied to a data set.

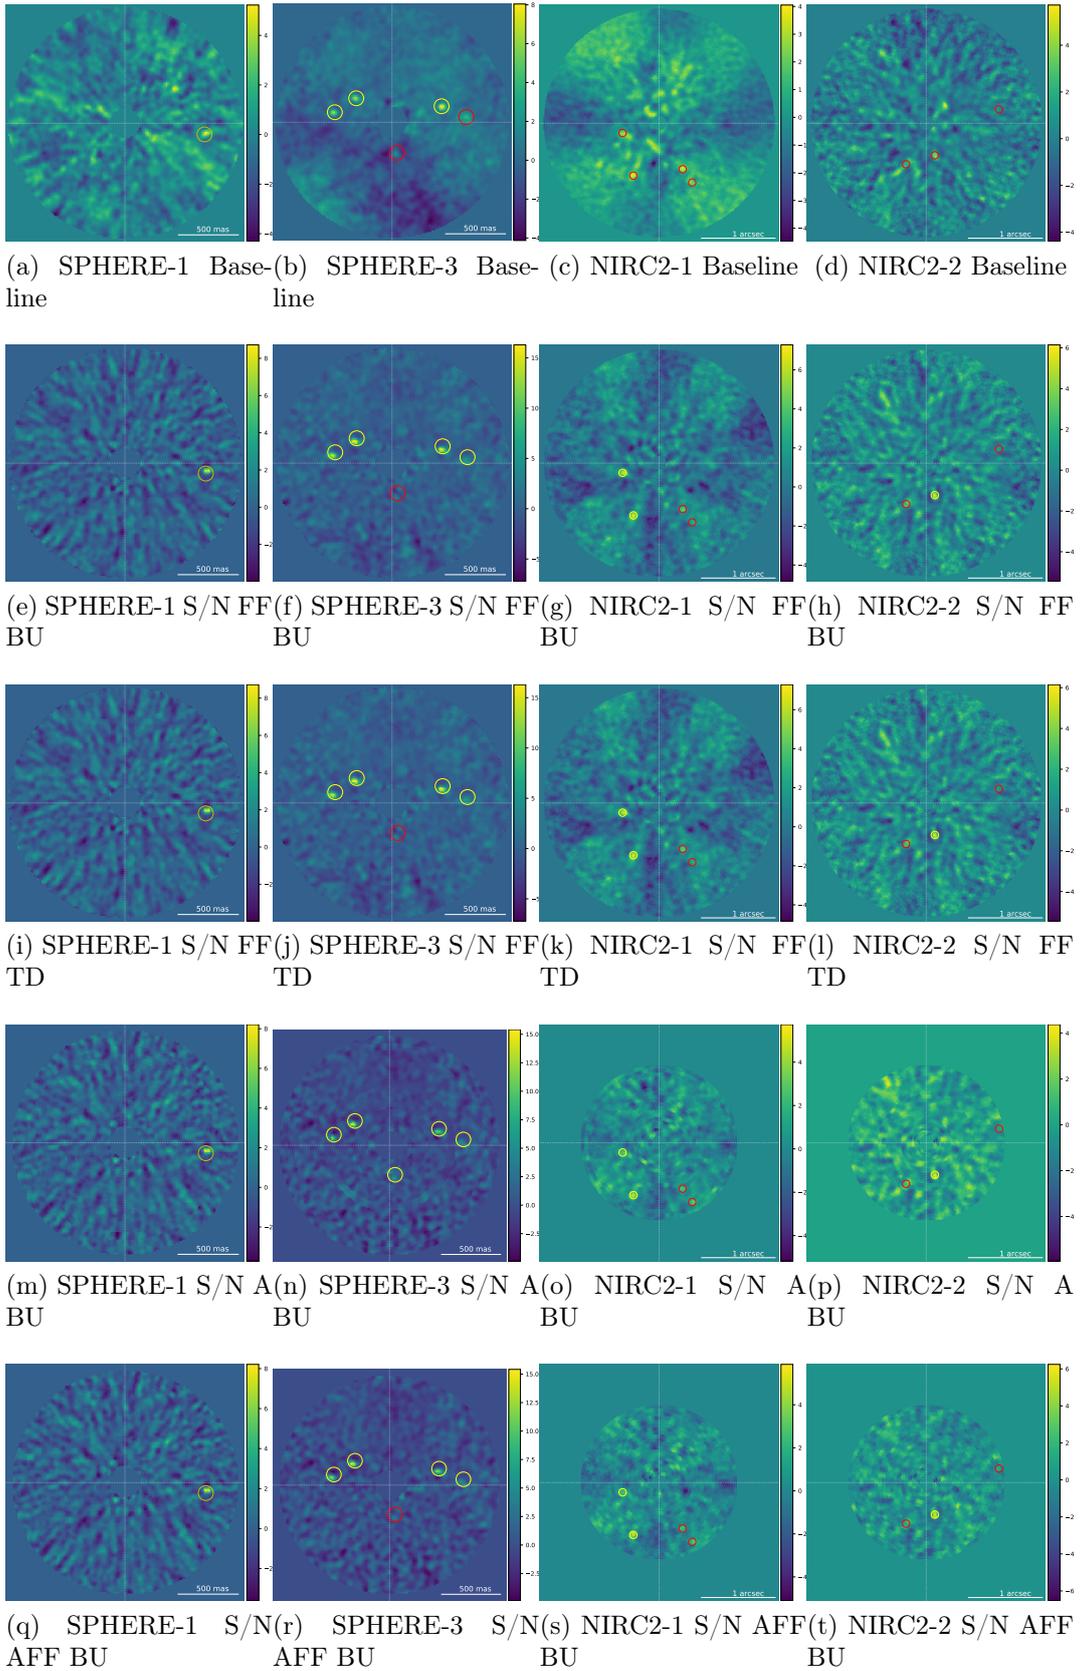


Figure C.4: Detection maps corresponding to the SPHERE and NIRC2 data sets generated with different parametrisations of the full-frame and annular auto-S/N along the baseline model presented in (Cantalloube et al. 2020b). The SPHERE-2 and NIRC2-3 detection maps are not shown, as no fake companions were injected in these two data sets. See Section 4.3.1. for the definition of the acronyms used to characterise the auto-RSM versions. The yellow circles are centred on the true position of the detected targets (TP) and the red circles give the true positions of FNs.

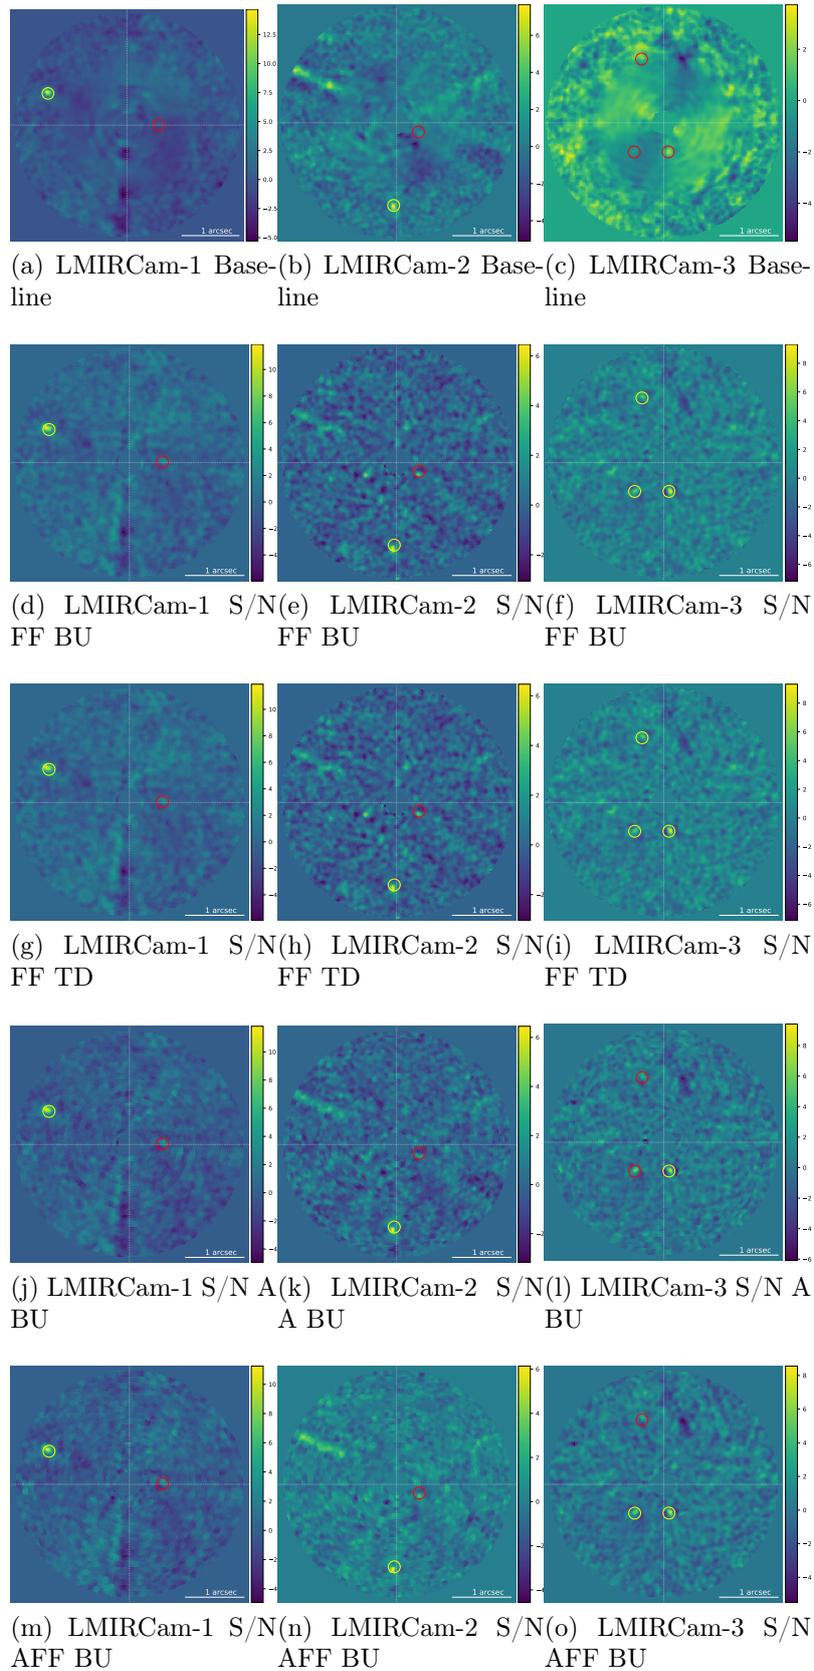


Figure C.5: S/N maps corresponding to the LMIRCam data sets generated with different parametrisations of the full-frame and annular auto-S/N along the baseline model presented in (Cantalloube et al. 2020b). See Section 4.3.1. for the definition of the acronyms used to characterise the auto-RSM versions. The yellow circles are centred on the true position of the detected targets (TP) and the red circles give the true positions of FNs.

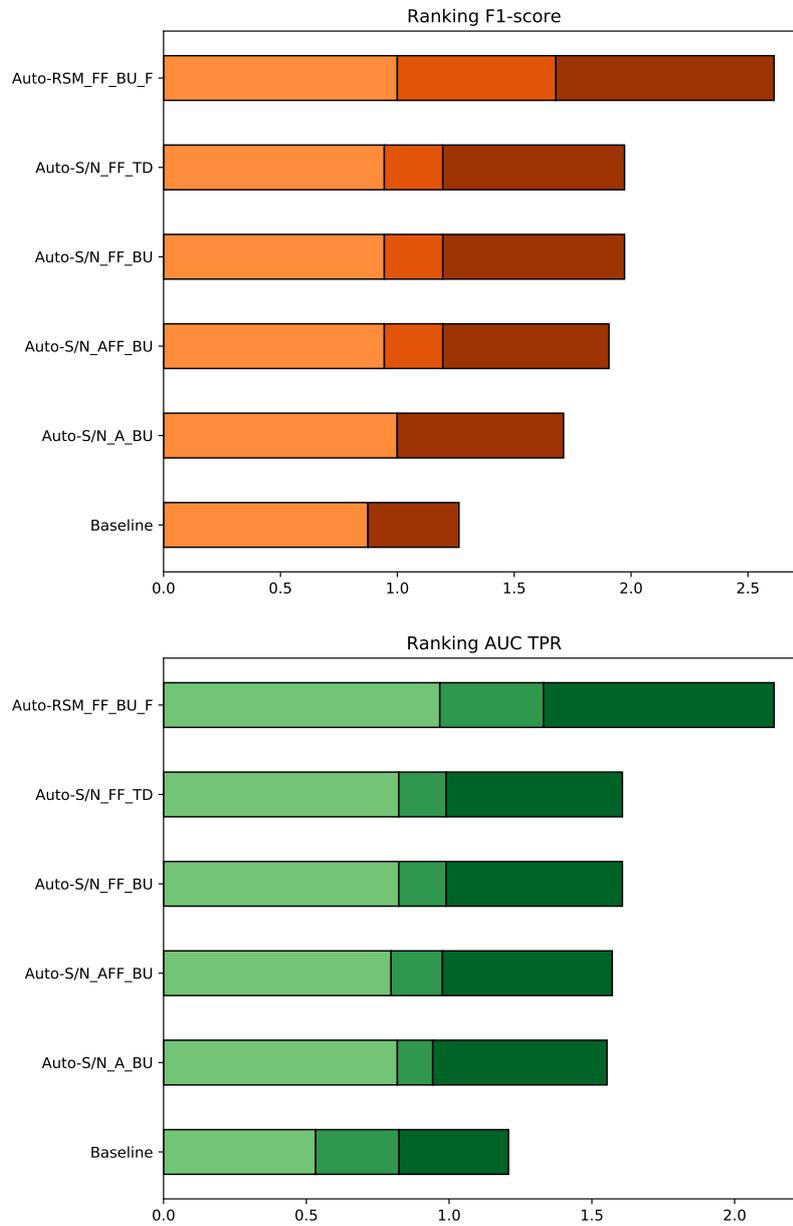


Figure C.6: Ranking of the different parametrizations of the full-frame and annular versions of the auto-S/N along the full-frame bottom-up forward auto-RSM and the baseline presented in (Cantalloube et al. 2020b). Figure (a) provides the ranking based on the F1 score obtained at the selected threshold. Figure (b) gives the ranking based on the AUC of the TPR. See Section 4.4.3 for the definition of the acronyms used to characterise the auto-RSM versions. The light, medium, and dark colours correspond to the VLT/SPHERE-IRDIS, Keck/NIRC2, and LBT/LMIRCam data sets, respectively.

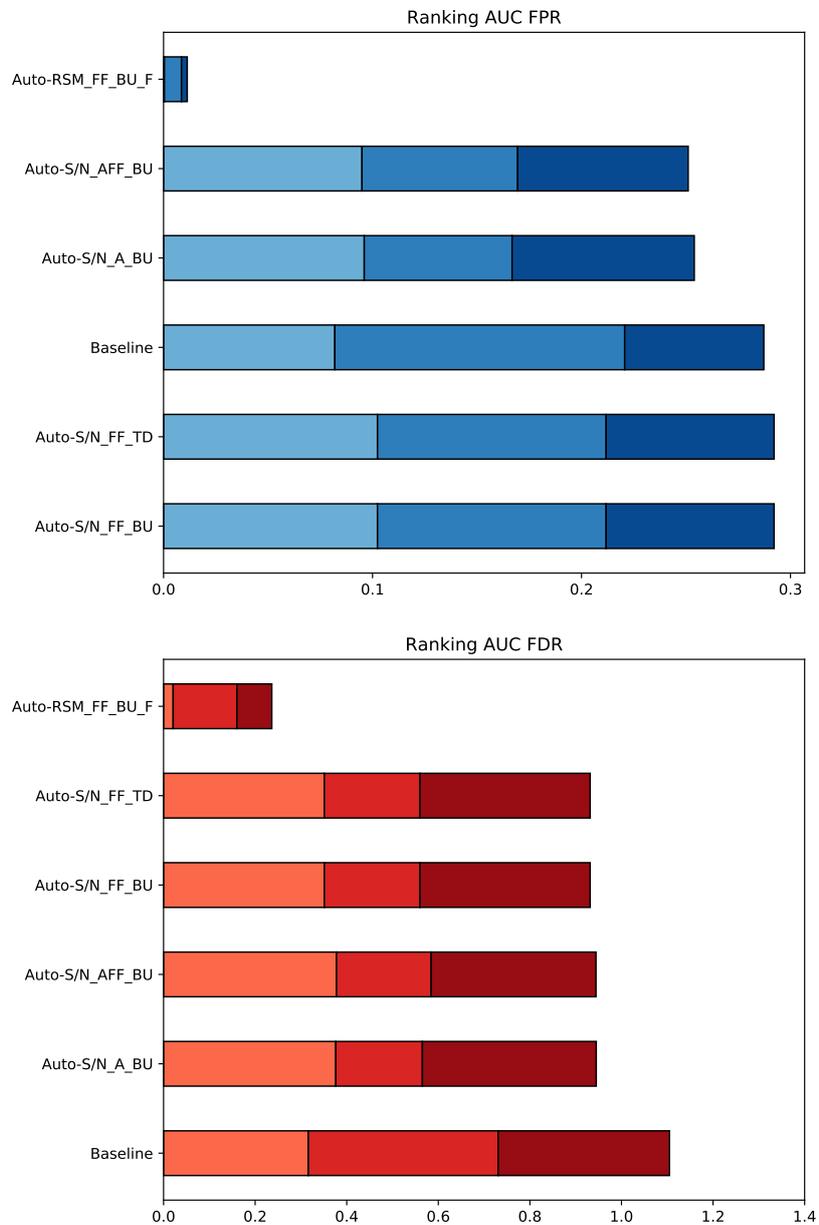


Figure C.7: Ranking of the different parametrizations of the full-frame and annular versions of the auto-S/N along the full-frame bottom-up forward auto-RSM and the baseline presented in (Cantalloube et al. 2020b). Figure (a) gives the ranking based on the AUC of the FPR, while Figure (b) provides the ranking based on the AUC of the FDR. See Section 4.4.3 for the definition of the acronyms used to characterise the auto-RSM versions. The light, medium, and dark colours correspond to the VLT/SPHERE-IRDIS, Keck/NIRC2, and LBT/LMIRCam data sets, respectively.

C.5 Auto-RSM pseudo-code

This Appendix presents first the pseudo-codes of the greedy algorithms used to select the optimal set of likelihood cubes generating the final RSM detection map (see Table C.3 and C.4 for respectively the bottom-up and top-down approaches). Tables C.5 and C.6 then provide the pseudo-codes for the auto-RSM optimisation algorithm in the full-frame mode and annular mode.

Table C.3: Pseudo-code for the bottom-up greedy selection algorithm. The PI symbol represents the RSM performance metrics.

```

1:  $\mathbf{Y}^a = \{Y_{c,m}^a, \forall c \in [0, N_{sequence}], m \in [0, N_{technique}]\}$ 
2:  $\mathbf{Z}^a = \emptyset$ 
3:  $PI_p^a = 0$ 
4: While  $\mathbf{Y}^a \neq \emptyset$  do
5:   For  $h = 1$  to  $length(\mathbf{Y}^a)$  do
6:      $PI_h^a = PI([\mathbf{Z}^a, \mathbf{Y}_h^a])$ 
7:   End for
8:   If  $max(PI_h^a) > PI_p^a$ 
9:      $Y_{c^*,m^*}^a = Y_{h^*}^a = argmax(PI_h^a)$ 
10:     $\mathbf{Y}^a = \mathbf{Y}^a \setminus \{Y_{h^*}^a, Y_h^a \mid (PI_h^a - PI_p^a) < 0\}$ 
11:     $\mathbf{Z}^a = \mathbf{Z}^a \cup Y_{h^*}^a$ 
12:     $PI_p^a = max(PI_h^a)$ 
13:   Else
14:      $\mathbf{Y}^a = \emptyset$ 
15:   End If
16: End While

```

Table C.4: Pseudo-code for the top-down greedy selection algorithm. The PI symbol represents the RSM performance metrics.

```

1:  $\mathbf{Z}^a = \{Y_{c,m}^a, \forall c \in [0, N_{sequence}], m \in [0, N_{technique}]\}$ 
2:  $PI_c^a = PI(\mathbf{Z}^a)$ 
3:  $PI_p^a = 0$ 
4: While  $PI_c^a > PI_p^a$  do
5:    $PI_p^a = PI_c^a$ 
6:   For  $h = 1$  to  $length(\mathbf{Z}^a)$  do
7:      $PI_h^a = PI(\mathbf{Z}^a \setminus Y_h^a)$ 
8:   End for
9:   If  $max(PI_h^a - PI_p^a) > 0$ 
10:     $Y_{c^*,m^*}^a = Y_{h^*}^a = argmax(PI_h^a - PI_p^a)$ 
11:     $\mathbf{Z}^a = \mathbf{Z}^a \setminus Y_{h^*}^a$ 
12:     $PI_c^a = max(PI_h^a - PI_p^a)$ 
13:   End If
14: End While

```

Table C.5: Pseudo-code for the auto-RSM algorithm in full-frame mode. The centre of the selected set of annuli $Range_{sel}$ is computed based on the rule provided in eq.4.9, starting at 1.5 FWHM and ending at a_{max} .

```

1: Flipping parallactic angle sign:  $PA = -PA$ 
2: For  $c = 1$  to  $N_{sequence}$  do
3:   For  $m = 1$  to  $N_{technique}$  do
4:     PSF-subtraction technique parameters optimisation
5:     For  $a$  in  $Range_{sel}$  do
6:       Contrast computation via fake companions injection
7:       Contrast normalisation
8:     End For
9:     Optimal parameters and contrast selection based on
       summed normalised contrast:  $[C_{a,m,c}, P_{m,c}^{PSF}]$ 
10:    Optimisation of RSM algorithm parameters
11:    For  $a$  in  $Range_{sel}$  do
12:      Median position computation:  $M_{a,m,c}$ 
13:      Performance metric computation via fake companion
        injection using  $[M_{a,m,c}, C_{a,m,c}, P_{m,c}^{PSF}]$ 
14:    End For
15:    Optimal parameters selection:  $P_{m,c}^{RSM}$ 
16:  End For
17: End For
18: Optimal combination  $Z$  selection via bottom-up or
    top-down approach using  $[P_{m,c}^{RSM}, M_{a,m,c}, C_{a,m,c}, P_{m,c}^{PSF}]$ 
19: For  $a = FWHM$  to  $a_{max}$  do
20:   Threshold  $T_a$  computation using  $[P_{m,c}^{RSM}, P_{m,c}^{PSF}, Z]$ 
21: End For
22: Threshold smoothing via polynomial fit:  $T_a^*$ 
23: Flipping back parallactic angle sign:  $PA = -PA$ 
24: Final RSM map computation using  $[P_{m,c}^{RSM}, P_{m,c}^{PSF}, Z, T_a^*]$ 
25: Threshold subtraction from final RSM map

```

Table C.6: Pseudo-code for the auto-RSM algorithm in annular mode. The centre of the selected set of annuli $Range_{sel}$ starts at 1,5 FWHM and end at a_{max} with the centre of every selected annulus separated by one FWHM.

```

1: Flipping parallactic angle sign:  $PA = -PA$ 
2: For  $c = 1$  to  $N_{sequence}$  do
3:   For  $m = 1$  to  $N_{technique}$  do
4:     For  $a$  in  $Range_{sel}$  do
5:       PSF-subtraction technique parameters optimisation
6:         Contrast computation via fake companions injection
7:         Optimal parameters and contrast  $[C_{a,m,c}, P_{a,m,c}^{PSF}]$ 
8:       End For
9:       Outliers suppression in  $P_{a,m,c}^{PSF}$  via Hampel Filter
10:      Parameters smoothing via moving average:  $P_{a,m,c}^{PSF,*}$ 
11:      For  $a$  in  $Range_{sel}$  do
12:        Optimisation of RSM algorithm parameters
13:          Median position computation  $M_{a,m,c}$ 
14:          Performance metric computation via fake companion
            injection using  $[M_{a,m,c}, C_{a,m,c}, P_{a,m,c}^{PSF,*}]$ 
15:          Optimal parameters selection:  $P_{a,m,c}^{RSM}$ 
16:        End For
17:        Outliers suppression in  $P_{a,m,c}^{RSM}$  via Hampel Filter
18:        Interpolation of optimal parameters via RBF:  $P_{a,m,c}^{RSM,*}$ 
19:      End For
20:    End For
21:    Optimal combination  $Z^a$  selection via bottom-up or
      top-down approach using  $[P_{a,m,c}^{RSM,*}, M_{a,m,c}, C_{a,m,c}, P_{a,m,c}^{PSF,*}]$ 
22:    For  $a = FWHM$  to  $a_{max}$  do
23:      Threshold  $T_a$  computation using  $[P_{a,m,c}^{RSM,*}, P_{a,m,c}^{PSF,*}, Z^a]$ 
24:    End For
25:    Flipping back parallactic angle sign:  $PA = -PA$ 
26:    Final RSM map computation using  $[P_{a,m,c}^{RSM}, P_{a,m,c}^{PSF}, Z^a, T_a]$ 
27:    Threshold subtraction from final RSM map

```

C.6 Description of the ADI sequences

This Appendix provides a description of the data sets of the EIDC ADI subchallenge used in the performance assessment of the different modes of the auto-RSM optimisation algorithm (Table C.7).

Table C.7: Characteristics of the nine ADI sequences included in the EIDC ADI subchallenge. The original number of frames for the LMIRCam ADI sequences was reduced to limit the computation time, relying on a moving average applied along the time axis on the de-rotated ADI cubes. The step and window sizes have been set to 20 frames for LMIRCam 1 and 3, and 15 frames for LMIRCam 2.

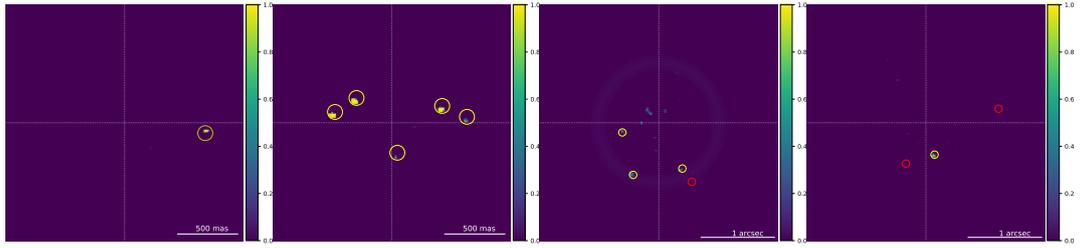
Instruments/ID	Number of frames	Frame size	Plate-scale (mas/pixel)	FOV rotation
SPHERE 1	252	160 × 160	12.255	40.3°
SPHERE 2	80	160 × 160	12.255	31.5°
SPHERE 3	228	160 × 160	12.255	80.5°
NIRC2 1	29	321 × 321	20	53.0°
NIRC2 2	40	321 × 321	20	37.3°
NIRC2 3	50	321 × 321	20.2	166.9°
LMIRCam 1	241	200 × 200	10.7	153.4°
LMIRCam 2	214	200 × 200	10.7	60.6°
LMIRCam 3	231	200 × 200	10.7	91.0°

C.7 Definition of the parameter ranges

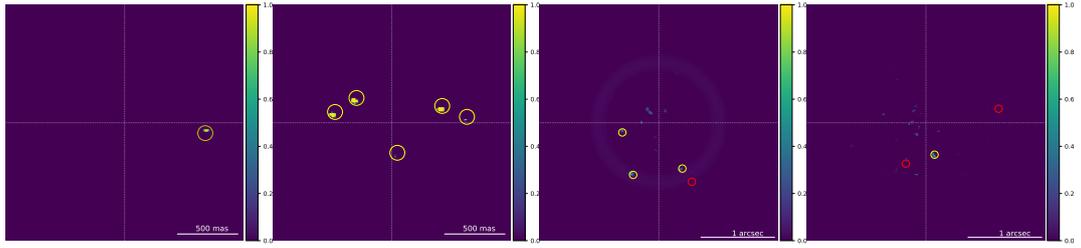
This Appendix provides the boundaries of the parameter space for the different data sets of the EIDC ADI subchallenge used for the performance assessment of the auto-RSM optimisation algorithm (Table C.8).

C.8 Detection maps for auto-RSM parametrisations

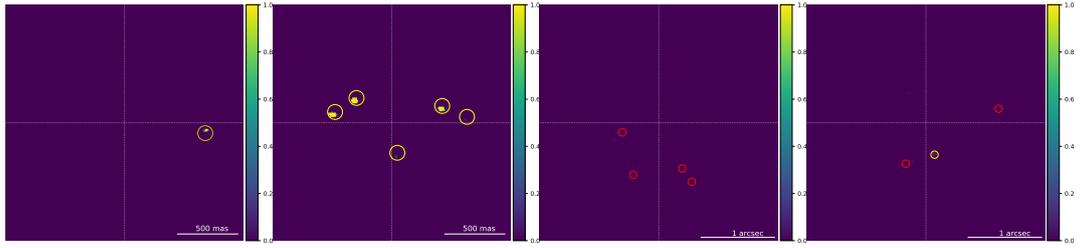
Here, Figure C.8 and Figure C.9 show the detection maps obtained with four parametrisations of the auto-RSM algorithm for the data sets of the EIDC ADI subchallenge.



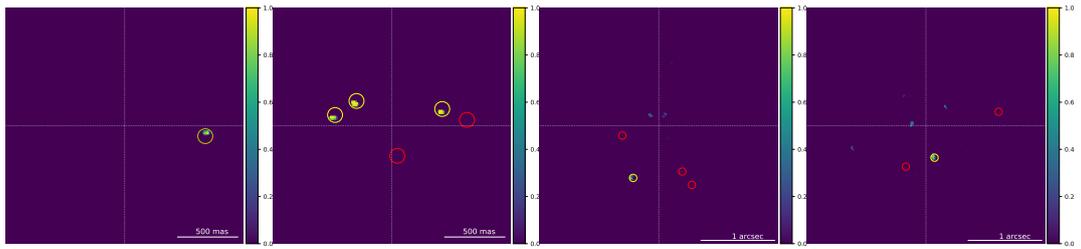
(a) SPHERE-1 FF BU (b) SPHERE-3 FF BU (c) NIRC2-1 FF BU (d) NIRC2-2 FF BU
F F F F



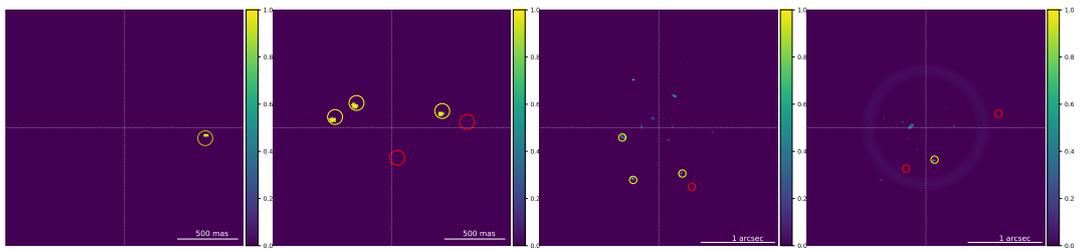
(e) SPHERE-1 FF TD (f) SPHERE-3 FF TD (g) NIRC2-1 FF TD (h) NIRC2-2 FF TD
F F F F



(i) SPHERE-1 FF BU (j) SPHERE-3 FF BU (k) NIRC2-1 FF BU (l) NIRC2-2 FF BU
FB FB FB FB



(m) SPHERE-1 A BU (n) SPHERE-3 A BU (o) NIRC2-1 A BU (p) NIRC2-2 A BU
F F F F



(q) SPHERE-1 AFF (r) SPHERE-3 AFF (s) NIRC2-1 AFF BU (t) NIRC2-2 AFF BU
BU F BU F F F

Figure C.8: Detection maps corresponding to the SPHERE and NIRC2 data sets generated with different parametrisations of the full-frame and annular auto-RSM. The SPHERE-2 and NIRC2-3 detection maps are not shown, as no fake companions were injected in these two data sets. See Section 4.4.3 for the definition of the acronyms used to characterise the auto-RSM versions. The yellow circles are centred on the true position of the detected targets (TP) and the red circles give the true positions of FNs.

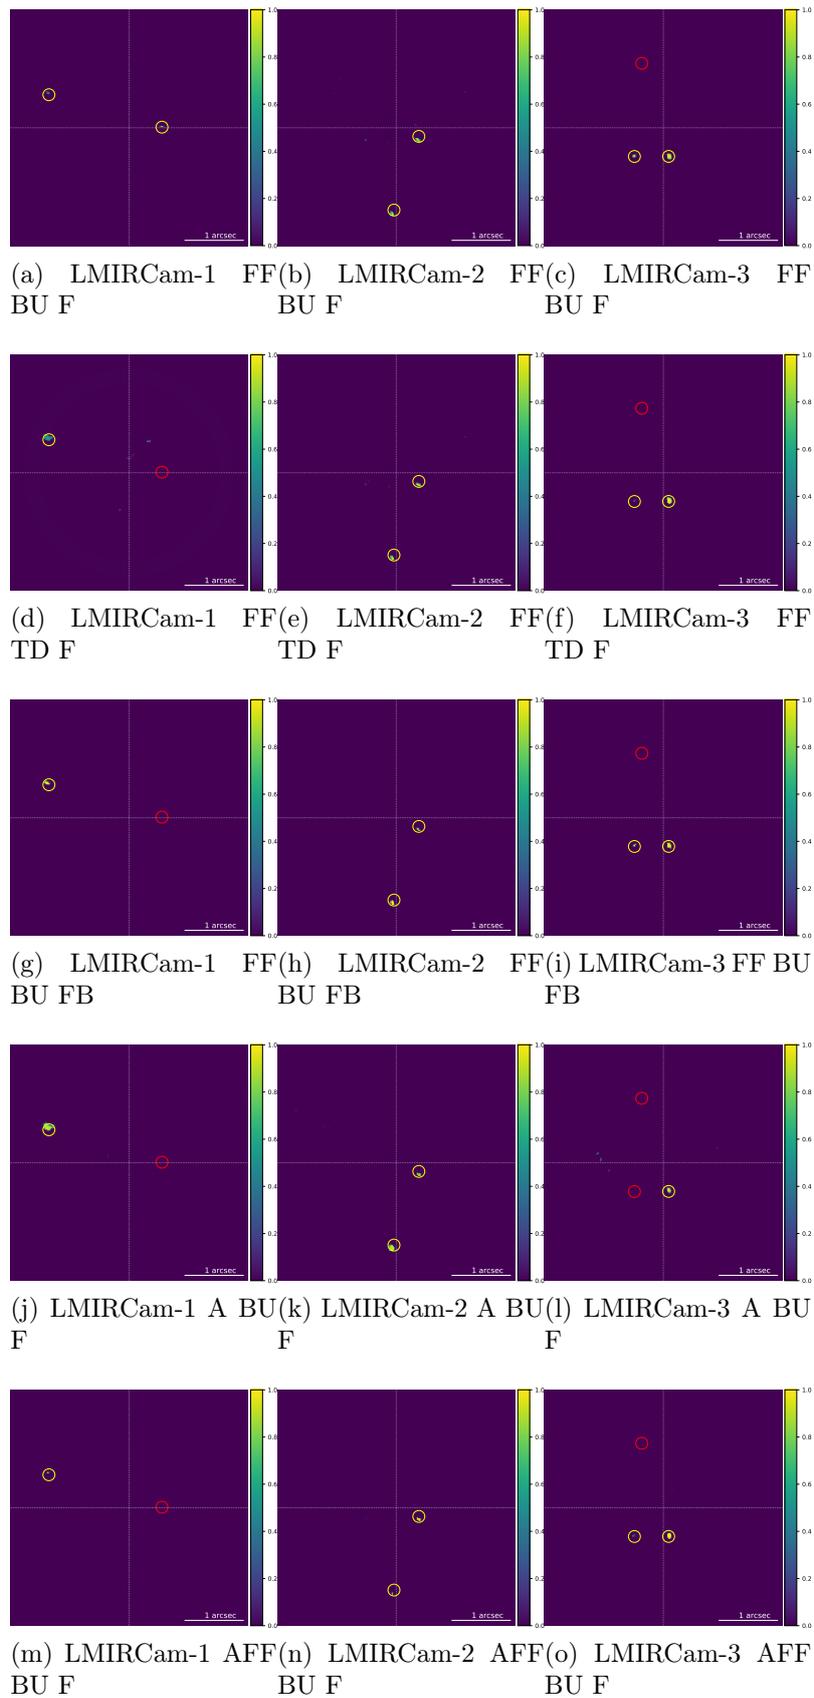


Figure C.9: Detection maps corresponding to the LMIRCam data sets, generated with different parametrisations of the full-frame and annular auto-RSM. See Section 4.4.3 for the definition of the acronyms used to characterise the auto-RSM versions. The yellow circles are centred on the true position of the detected targets (TP) and the red circles give the true positions of FNs.

C.9 Parametrisation for the full-frame auto-RSM

Here, Table C.9 regroups the optimal parameters selected with the auto-RSM FF_BU_F for the nine ADI sequences of the EIDC ADI subchallenge.

Table C.9: Optimal set of parameters for the PSF-subtraction techniques and RSM algorithm for the nine ADI sequences obtained with the auto-RSM FF_BU_F. The ‘fit’ row indicates if the noise properties have been estimated using a best-fit approach while the β row indicates if a Gaussian maximum likelihood has been used to compute the intensity parameter. The variance row provides information about the region used for the noise properties computation and translates as follows: ST-Spatio-Temporal estimation, F-Frame based estimation, FM-Frame with mask estimation, SM -Segment with mask estimation, and T-Temporal estimation.

Parameters/ID	NIRC2-1	NIRC2-2	NIRC2-3	SPHERE-1	SPHERE-2	SPHERE-3	LMIRCam-1	LMIRCam2	LMIRCam-3
APCA components	11	17	25	25	18	23	40	42	21
APCA segments	4	4	4	2	1	3	4	2	2
APCA FOV rotation	0.679	0.296	0.984	0.311	0.261	0.389	0.298	0.300	0.256
NMF components	12	2	14	18	18	12	14	20	18
LLSG rank	4	5	3	8	8	8	8	6	8
LLSG segments	4	4	1	4	2	2	4	2	3
LOCI tolerance	0.00752	0.00138	0.00425	0.00242	0.00128	0.00887	0.00112	0.00104	0.00218
LOCI FOV rotation	0.355	0.268	0.447	0.250	0.267	0.255	0.252	0.261	0.326
APCA δ	5	5	5	5	5	5	5	5	5
NMF δ	5	5	5	5	5	5	5	5	5
LLSG δ	5	5	3	4	1	2	4	2	5
LOCI δ	5	5	2	5	5	5	4	5	5
APCA crop size	3	3	3	3	3	3	3	3	3
NMF crop size	3	3	3	3	3	3	3	3	3
LLSG crop size	3	3	3	3	3	3	3	3	3
LOCI crop size	3	3	3	3	3	3	3	3	3
APCA Fit	True	True	True	True	True	True	True	True	True
NMF Fit	True	True	True	True	True	True	True	True	True
LLSG Fit	True	True	True	True	True	True	True	True	True
LOCI Fit	True	True	True	True	True	True	True	True	True
APCA β	True	True	False	False	False	False	False	False	False
NMF β	False	False	False	False	False	False	False	False	False
LLSG β	True	True	False	False	False	False	False	False	False
LOCI β	False	True	False	False	False	False	False	False	False
APCA variance	T	FM	FM	SM	T	FM	ST	SM	FM
NMF variance	FM	FM	ST	SM	FM	ST	ST	SM	ST
LLSG variance	FM	FM	SM	SM	SM	SM	ST	FM	ST
LOCI variance	FM	T	FM	T	ST	SM	SM	SM	SM

Appendix D

SHARDDS Survey: Limits on Planet Occurrence Rates Based on Point Sources Analysis via the Auto-RSM Framework

D.1 Parametrisation commonalities and relationship with ADI sequence characteristics

Following the computation of the optimal set of parameters for the 8 cluster centroids, we propose to investigate the similarities existing between these 8 optimal parametrisations. We also consider the relationships existing between the centroids optimal parameters and the set of metrics characterising their ADI sequence. We start by comparing in Figure D.1, the obtained optimal set of parameters via a normalised distance for the PSF-subtraction techniques and a dissimilarity index for the RSM algorithm. These measures were computed for each pair of cluster centroids and then averaged over the three possible pairs within each size subset (e.g. for HD 192758, we have HD 192758-HD 3670, HD 192758-HD 201219, and HD 192758-HD 14082B). The normalised distance was computed considering the 19 parameters required by the 10 selected PSF subtraction techniques. For each pair of cluster centroids, we computed the absolute value of the distance between their parameters and normalised them with the mean values of these pairs of parameters¹. We then averaged the resulting distances over the 19 parameters. The normalisation ensures a proper comparison between the different parameters when consolidating the distances. For the RSM algorithm, a dissimilarity metric replaces the normalised distance as most parameters are non-numerical. This dissimilarity index is simply computed as one minus the percentage of common RSM parameters between a pair of centroids, averaged over the

¹For centroid A with 20 principal components for APCA and centroid B with 24 principal components, the normalised distance is equal to $4/22 = 0.18$

five parameters of the RSM algorithm.

Looking at the degree of similarity of the parametrisations within the two size subsets, Figure D.1 shows an overall higher degree of similarity. We observe a lower degree of dissimilarity for the RSM parametrisation and a lower normalised distance for the PSF subtraction-techniques for the centroids of the subset containing less than 151 observations. For the subset containing more than 151 observations, the slightly higher normalised distance pertain to the high degree of dissimilarity of HD 181296, which affects strongly the averaged normalised distance. The main drivers of the dissimilarity is the number of segments used for APCA and LLSG², the tolerance parameter of LOCI and the method used to compute the residual speckle noise statistics within the RSM algorithm. These results tend to demonstrate the relatively high stability of the ADI sequence imaged by the VLT/SPHERE instrument and confirm the conclusions drawn in Dahlqvist et al. (2021). The impact of the dissimilarities in the optimal parametrisations on the performance in terms of achievable contrast is further investigated in Section 5.5.

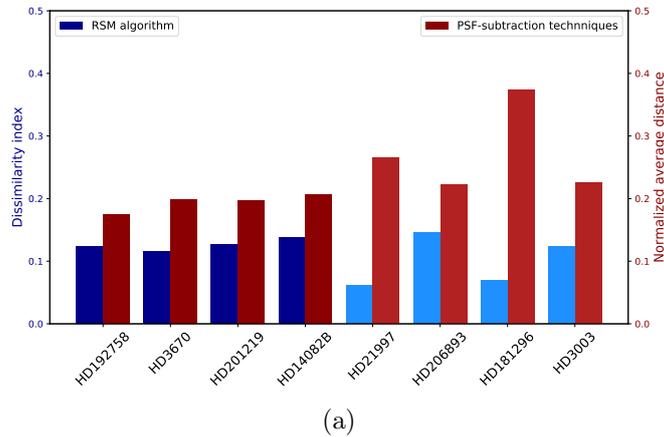


Figure D.1: Comparison of the optimal parametrisation of the cluster centroids in terms of dissimilarity index of the RSM algorithm parametrisation and normalised average distance for the PSF-subtraction techniques, for the subset containing ADI sequences with less than 151 frames (dark colors) and the one with a number of frames above 151 images (light colors).

We now turn to the analysis of the relationship existing between the parameters that we selected in Section 5.3.3 to describe our data set and the parametrisations of the PSF-subtraction techniques³. We computed the Pearson correlation between the 10 parameters characterising our sample and the PSF-subtraction techniques parameters, considering the 8 cluster centroids as data-points. The raw correlations show a significant correlation between these sets of parameters, with overall, around 25% of the obtained values over 0.5. Table D.1 gives the absolute values of the obtained correlations averaged over five classes of parameters, the number of principal compo-

²The number of segments correspond to the number of subdivisions of every annulus during the estimation of the reference PSF when relying on APCA and LLSG.

³Such an analysis is not possible with the parametrisation of the RSM map algorithm as most parameters are non numerical.

nents, the FOV rotation threshold, the number of segments, the rank of LLSG and the tolerance of LOCI. Looking at these consolidated results, the contrast at 500 mas shows the highest average correlation rate, while the exponent of the autocorrelation function has the lowest one. Once averaged over the 5 considered classes, the percentage of consolidated correlations above 0.5 reach only 16%, indicating the existence of some discrepancies between the different PSF-subtraction techniques relying on the same parameter.

Table D.1: Average absolute Pearson correlations between the PSF-subtraction techniques parameters and the parameters selected to characterised the SHARDDS survey data set. WDH S and WDH A stand respectively for wind driven halo strength and asymmetry.

Parameters	# frames	Contrast	Auto-corr exp	PA	Coherence	Wind speed	Seeing	Strehl	WDH S	WDH A
Principal components	0.44	0.41	0.32	0.44	0.36	0.28	0.37	0.39	0.33	0.54
FOV rotation threshold	0.32	0.54	0.36	0.16	0.35	0.7	0.37	0.13	0.55	0.16
Number of segment	0.41	0.34	0.22	0.42	0.33	0.42	0.29	0.45	0.15	0.24
Rank	0.36	0.29	0.14	0.33	0.19	0.14	0.19	0.41	0.17	0.51
Tolerance	0.49	0.6	0.17	0.71	0.24	0.34	0.64	0.31	0.21	0.13

D.2 RSM detection maps

This section contains the RSM detection maps containing no plausible planetary signals.

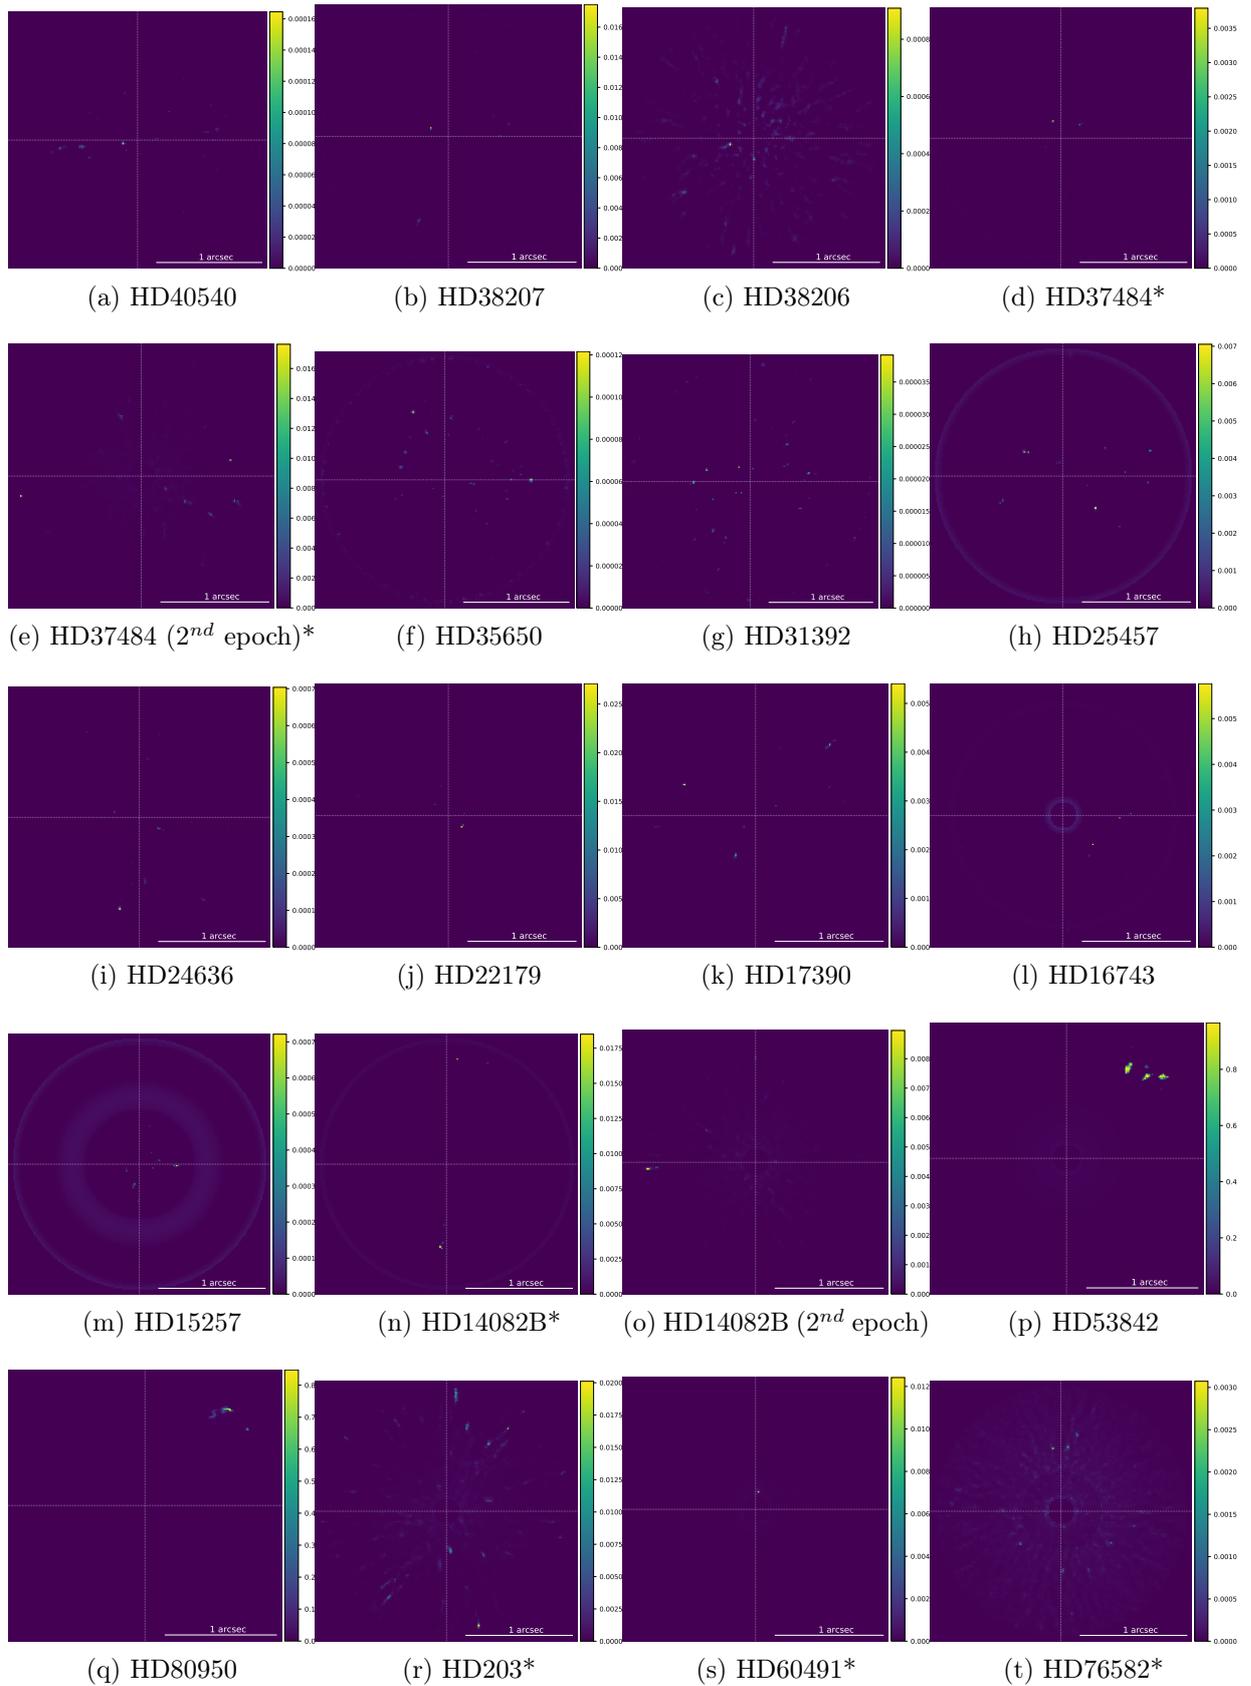


Figure D.2: RSM detection maps generated using auto-RSM or the optimal parameters obtained with auto-RSM for the data set at the center of the clusters (see Table 5.5). These detection maps did not lead to the detection of a target. The asterisks indicate the targets on which the full auto-RSM framework was applied.

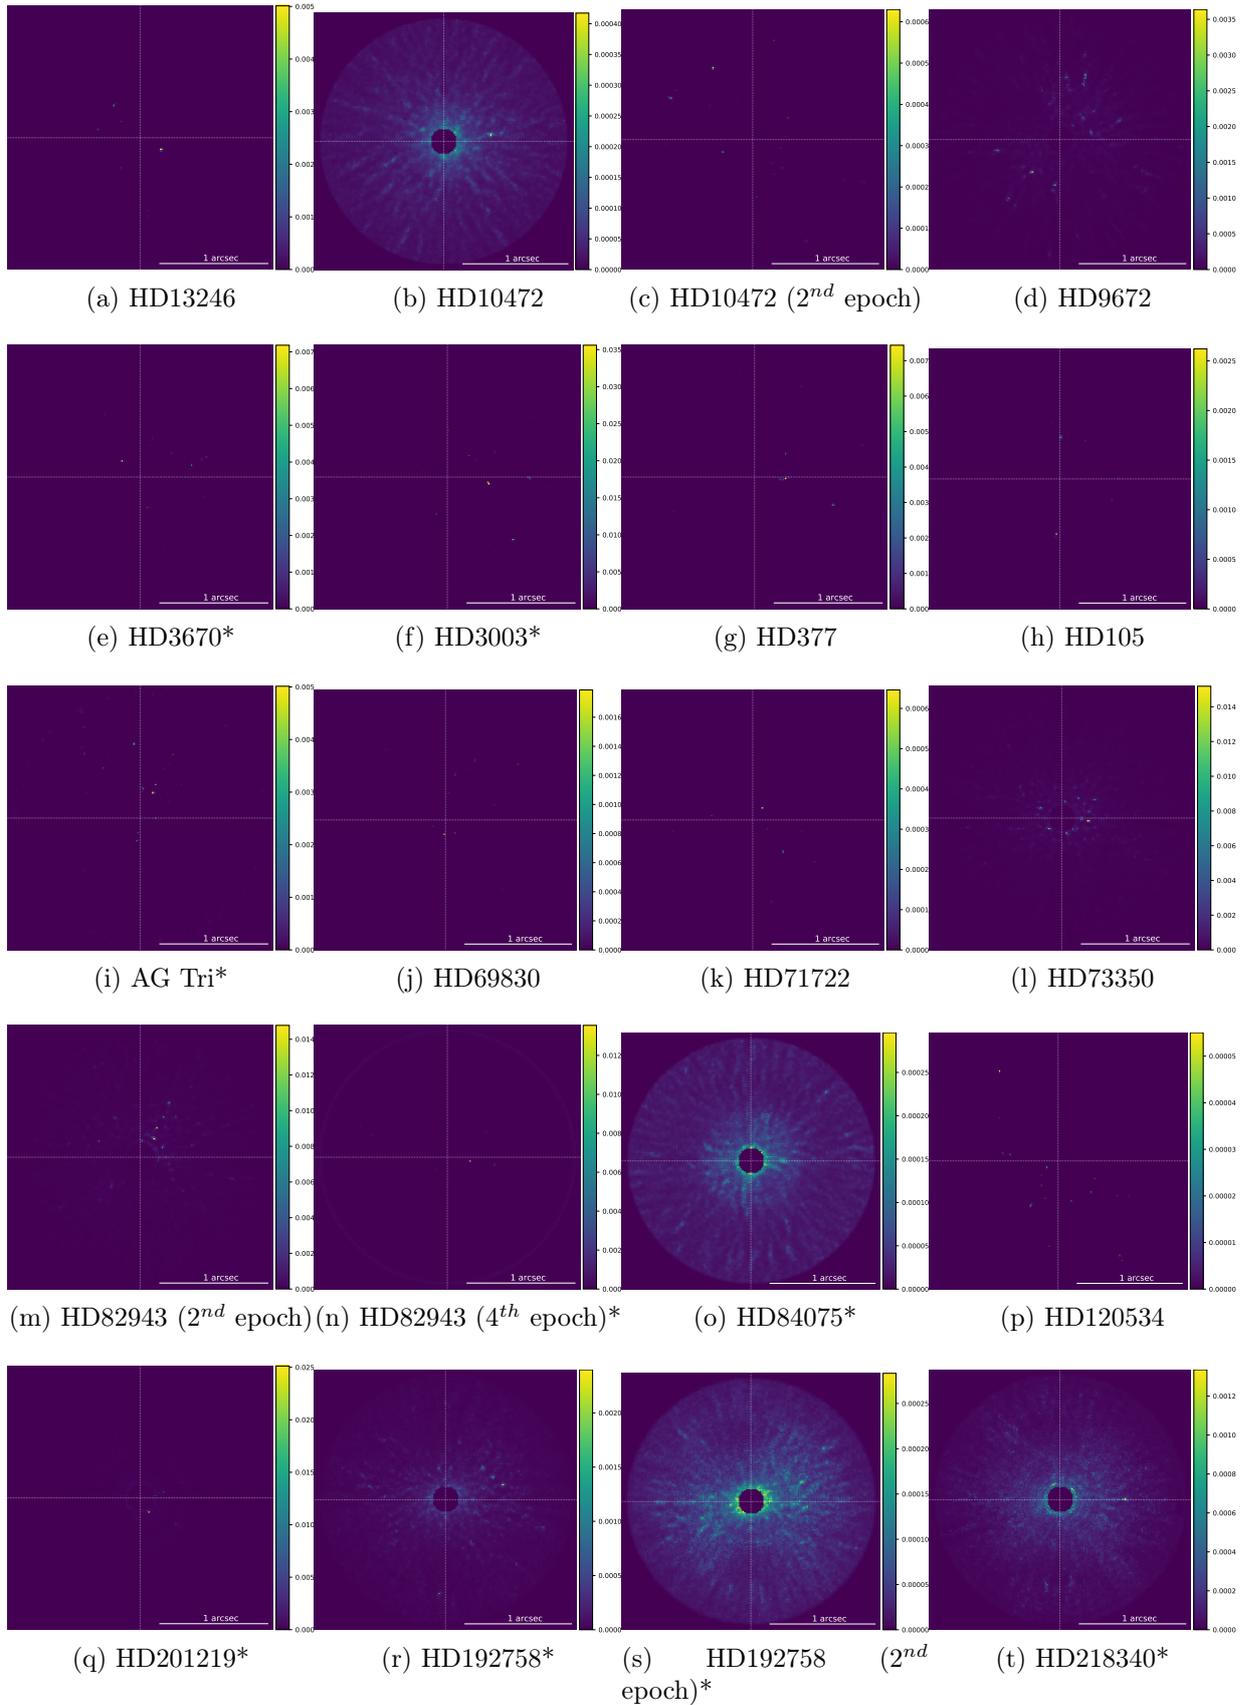


Figure D.3: RSM detection maps generated using auto-RSM or the optimal parameters obtained with auto-RSM for the data set at the center of the clusters (see Table 5.5)). These detection maps did not lead to the detection of a planetary candidate. The asterisks indicate the targets on which the full auto-RSM framework was applied.

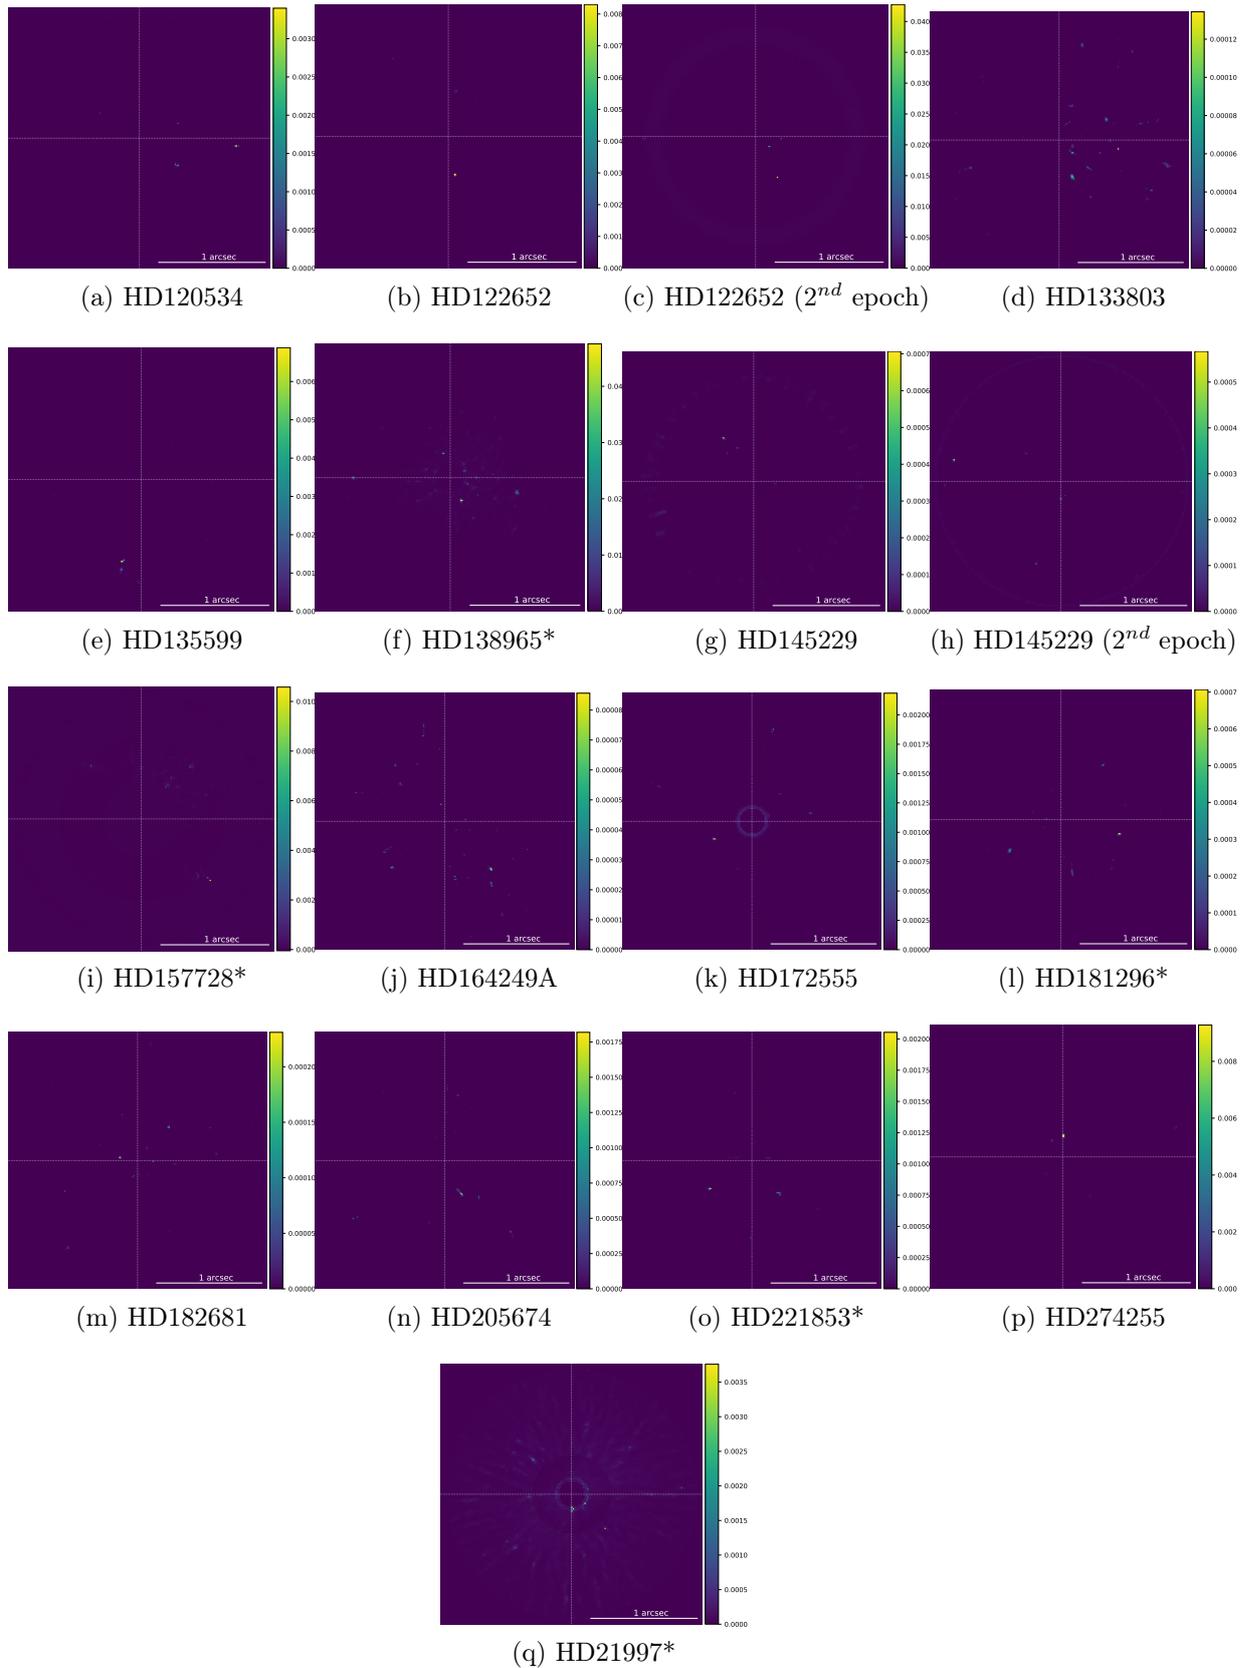


Figure D.4: RSM detection maps generated using auto-RSM or the optimal parameters obtained with auto-RSM for the data set at the center of the clusters (see Table 5.5). These detection maps did not lead to the detection of a planetary candidate. The asterisks indicate the targets on which the full auto-RSM framework was applied.

D.3 Detection maps for planetary candidates

This appendix regroups the RSM detection maps obtained with auto-RSM using either the bottom-up or top-down approaches to select the optimal set of likelihoods cubes (each likelihoods cube corresponding to a PSF-subtraction technique), as well as S/N maps generated via the Auto-S/N approach (Dahlqvist et al. 2021) or obtained by averaging the S/N maps generated with APCA, NMF, LLSG and LOCI, for the two samples containing a potential planetary signal.

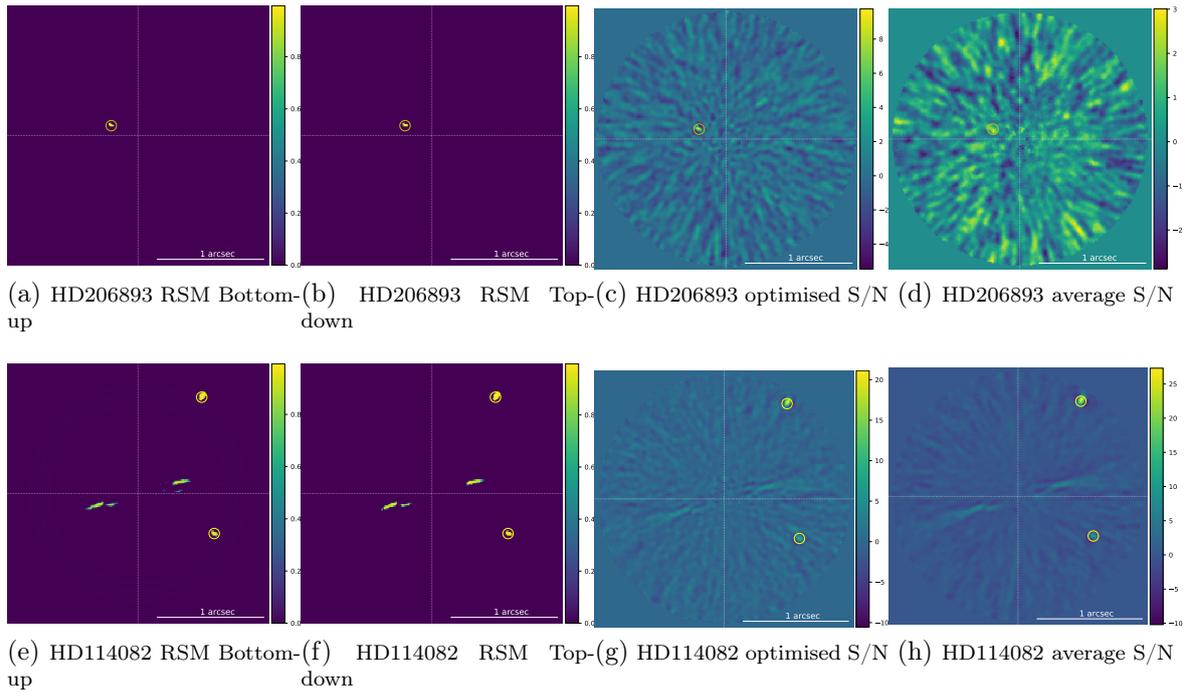


Figure D.5: RSM detection maps generated using the auto-RSM with the bottom-up (left) and top-down (middle left) approaches to select the optimal set of likelihoods cubes (each likelihood cube corresponds to a PSF-subtraction technique), S/N maps generated using the Auto-S/N to select the optimal set of S/N maps (middle right) and S/N maps obtained by averaging the S/N map generated with APCA, NMF, LLSG and LOCI (right). The expected position of the planetary signal as estimated via the RSM based planetary signal characterisation algorithm is indicated by a yellow circle.

D.4 Determination of projected angular separation

The true anomaly is defined as the angle between the direction of the periapsis and the current position vector of the body in the perifocal plane. Its estimation starts by the definition of the mean anomaly, which provides the fraction of the elliptical orbit that was covered since the periapsis expressed in radian $[0, 2\pi]$. The mean anomaly is linked to the eccentric anomaly by the following relationship:

$$M = \frac{2\pi}{T}t = E - e \sin(E), \quad (\text{D.1})$$

with T the orbital period and e the eccentricity. This transcendental equation relating time and eccentric anomaly cannot be directly solved. However, there exists a unique solution for every value of the mean anomaly M . We rely on the expansion of E in terms of Bessel functions to relate eccentric anomaly and mean anomaly.

$$E = M + \sum_{n=1}^{\infty} \frac{2}{n} J_n(ne) \sin(nM), \quad (\text{D.2})$$

with $J_n(x)$ the Bessel function of the first kind and $n = \frac{2\pi}{T}$. The sum over n is truncated to $N = 100$. The true anomaly θ is then computed via the following relationship:

$$\theta = 2 \tan^{-1} \left(\sqrt{\frac{1+e}{1-e}} \tan\left(\frac{E}{2}\right) \right), \quad (\text{D.3})$$

Once the true anomaly has been estimated, the position vector in the perifocal frame is computed using the elliptic orbit equation:

$$\mathbf{r}_p = \frac{h^2}{\mu} \frac{1}{1 + e \cos(\theta)} (\cos(\theta)\hat{\mathbf{p}} + \sin(\theta)\hat{\mathbf{q}}), \quad (\text{D.4})$$

where the coordinates are normalised, such as $\hat{\mathbf{p}} = [1, 0, 0]$ and $\hat{\mathbf{q}} = [0, 1, 0]$. Using $h = \sqrt{\mu a(1 - e^2)}$, we get:

$$\mathbf{r}_p = \frac{a(1 - e^2)}{1 + e \cos(\theta)} (\cos(\theta)\hat{\mathbf{p}} + \sin(\theta)\hat{\mathbf{q}}). \quad (\text{D.5})$$

We project this position vector in the equatorial frame via three Euler rotations:

$$\mathbf{r}_e = [\mathbf{Q}] \mathbf{r}_p, \quad (\text{D.6})$$

with the Euler rotations given by:

$$[\mathbf{Q}] = [\mathbf{R}_3(w)] [\mathbf{R}_1(i)] [\mathbf{R}_3(\Omega)], \quad (\text{D.7})$$

where i is the inclination, w the argument of the periapsis, and Ω the longitude of the ascending node. The normalised distance to the star is then obtained by computing the norm of the position vector in the equatorial frame.

$$r = \|\mathbf{r}_e\| \quad (\text{D.8})$$

The angular separation expressed in mas is finally defined as the normalised distance to the star multiplied by the semi-major expressed in mas:

$$a_{sep} = ra \frac{1000 \times 3600 \times 180}{(206265\pi d)} \quad (\text{D.9})$$

with a the semi-major axis expressed in au and d the distance from the star expressed in pc.

D.5 Disks analysed in Pearce et al. (2022)

There are 21 targets in common between the SHARDDS sample analysed in this paper and the sample of Pearce et al. (2022). We do not consider here Fomalhaut C, part of SHARDDS and in Pearce et al. (2022) because of the very poor quality of the data. We present these targets in Table D.2, with the location of the disk inner radius⁴ used in the analysis by Pearce et al. (2022) to estimate the planet minimum masses.

Table D.2: The 21 common targets between SHARDDS and the sample analysed in Pearce et al. (2022). This table is an extract from Table A.1 in (Pearce et al. 2022, see references therein)

Target	Disk data	Disk location and extent
HD203	SED	29 ± 6
HD377	SED	60 ± 10
HD3003	SED	21 ± 6
HD3670	SED	100 ± 20
HD9672	ALMA	$62 \pm 4 \rightarrow 210 \pm 4$
HD10472	SED	110 ± 20
HD13246	SED	80 ± 30
HD16743	Herschel 100 μm	$50 \pm 50 \rightarrow 260 \pm 70$
HD21997	ALMA	$68 \pm 4 \rightarrow 120 \pm 4$
HD25457	SED	45 ± 8
HD37484	SED	70 ± 20
HD38206	ALMA	$0 \pm 20, 140^{+30} \rightarrow 190 \pm 30, 320^{+50}$
HD69830	SED	0.8 ± 2
HD107649	SED	15 ± 3
HD114082	SED	29 ± 6
HD135599	SED	49 ± 9
HD172555	SED	15 ± 3
HD181296	SED	81 ± 10
HD192758	Herschel 100 μm	$40 \pm 40 \rightarrow 180 \pm 50$
HD218340	SED	140 ± 40
HD221853	SED	47 ± 9

⁴The 'Disc location and extent' column describes the location and shape of the debris disc inner and outer edges: if the disc is resolved and fitted with an asymmetric model (case of HD38206), then the column shows the inner edge pericentre, q_i , inner edge apocentre, Q_i , outer edge pericentre, q_o , and outer edge apocentre, Q_o , as ' $q_i, Q_i \rightarrow q_o, Q_o$ '. Alternatively, if the disc is resolved and fitted with an axisymmetric model, then the column shows the disc inner edge, a_i , and outer edge, a_o , as ' $a_i \rightarrow a_o$ '. Finally, if the disc location is estimated from SED data, then only the corrected blackbody radius is shown.

Appendix E

Summary of main contributions and publication status

E.1 Publications as first author

Chapter 2: RSM detection map for direct exoplanet detection in ADI sequences

Beyond the choice of wavefront control systems or coronagraphs, advanced data processing methods play a crucial role in disentangling potential planetary signals from bright quasi-static speckles. Among these methods, angular differential imaging (ADI) for data sets obtained in pupil tracking mode (ADI sequences) is one of the foremost research avenue, considering the many observing programs performed with ADI-based techniques and the associated discoveries. Inspired by the field of econometrics, we propose in this chapter a new detection algorithm for ADI sequences, deriving from the regime-switching model first proposed by Hamilton (1988). The proposed model is very versatile as it allows the use of PSF-subtracted data sets (residual cubes) provided by various ADI-based techniques, separately or together, to provide a single detection map. The temporal structure of the residual cubes is used for the detection as the model is fed with concatenated series of pixel-wise time sequences. The algorithm provides a detection probability map by considering for concentric annuli two possible regimes, the first one accounting for the residual noise and the second one for the planetary signal in addition to the residual noise. The algorithm performance is tested on data sets from two instruments, VLT/NACO and VLT/SPHERE. The results show an overall better performance in the receiver operating characteristic space when compared with standard signal-to-noise ratio maps for several state-of-the-art ADI-based post-processing algorithms.

"RSM detection map for direct exoplanet detection in ADI sequences" published in A&A, 2020, 633, A95.

Chapter 3: Improving the RSM map exoplanet detection algorithm

High-contrast imaging (HCI) is one of the most challenging techniques for exoplanet detection. It relies on sophisticated data processing to reach high contrasts at small angular separations. Most data processing techniques of this type are based on the angular differential imaging (ADI) observing strategy to perform the subtraction of a reference point spread function (PSF). In addition, such techniques generally make use of signal-to-noise (S/N) maps to infer the existence of planetary signals via thresholding. An alternative method for generating the final detection map was recently proposed with the regime-switching model (RSM) map, which uses a regime-switching framework to generate a probability map based on cubes of residuals generated by different PSF subtraction techniques. In this chapter, we present several improvements to the original RSM map, focusing on novel PSF subtraction techniques and their optimal combinations, as well as a new procedure for estimating the probabilities involved. We started by implementing two forward-model versions of the RSM map algorithm based on the LOCI and KLIP PSF subtraction techniques. We then addressed the question of optimally selecting the PSF subtraction techniques to optimise the overall performance of the RSM map. A new forward-backward approach was also implemented to take into account both past and future observations to compute the RSM map probabilities, leading to improved precision in terms of astrometry and lowering the background speckle noise. We tested the ability of these various improvements to increase the performance of the RSM map based on data sets obtained with three different instruments: VLT/NACO, VLT/SPHERE, and LBT/LMIRCam via a computation of receiver operating characteristic (ROC) curves. These results demonstrate the benefits of these proposed improvements. Finally, we present a new framework to generate contrast curves based on probability maps. The contrast curves highlight the higher performance of the RSM map compared to a standard S/N map at small angular separations.

"Improving the RSM map exoplanet detection algorithm: PSF forward modelling and optimal selection of PSF subtraction techniques" published in A&A, 2021, 646, A49

Forward modelled matched filter algorithm for KLIP and LOCI introduced in VIP:
https://github.com/vortex-exoplanet/VIP/tree/master/vip_hci/invprob

Chapter 4: Auto-RSM: an automated parameter selection algorithm for the RSM map exoplanet detection algorithm

Most of the high-contrast imaging (HCI) data processing techniques used over the last 15 years have relied on angular differential imaging (ADI) observing strategy, along with reference point spread function (PSF) subtraction to generate exoplanet detection maps. Recently a new algorithm, called regime switching model (RSM) map, has been proposed to take advantage of these numerous PSF-subtraction techniques, by using several of them to generate a single probability map. The selection of the optimal parameters for these PSF-subtraction techniques as well as for the RSM map is not straightforward, time consuming, and can be biased by assumptions made on the underlying data set. We propose in this chapter a novel optimisation procedure that can be applied to each of the PSF-subtraction techniques alone, or to the entire RSM framework. The optimisation procedure consists of three main steps, (i) the definition of the optimal set of parameters for the PSF-subtraction techniques using the contrast as performance metrics, (ii) the optimisation of the RSM algorithm, and (iii) the selection of the optimal set of PSF-subtraction techniques and ADI sequences used to generate the final RSM probability map. The optimisation procedure is applied to the data sets of the exoplanet imaging data challenge (EIDC), which provides tools to compare the performance of HCI data processing techniques. The data sets consist of ADI sequences obtained with three state-of-the-art HCI instruments: SPHERE, NIRC2 and LMIRCam. The results of our analysis demonstrate the interest of the proposed optimisation procedure, with better performance metrics compared to the earlier version of RSM, as well as to other HCI data processing techniques.

"Auto-RSM: an automated parameter selection algorithm for the RSM map exoplanet detection algorithm" published in A&A 656, A54 (2021)

Python package for the auto-RSM:

<https://github.com/chdahlqvist/RSMmap>

Chapter 5: The SHARDDS Survey: Limits on Planet Occurrence Rates Based on Point Sources Analysis via the auto-RSM Framework

In the past decade, HCI surveys provided new insights about the frequency and properties of substellar companions at separation larger than 5 au. In this context, our study aims to detect and characterise potential exoplanets and brown dwarfs within debris disks, by considering the SHARDDS survey, which gathers 55 Main Sequence stars with known bright debris disk. We rely on the auto-RSM framework to perform an in-depth analysis of the targets, via the computation of detection maps and contrast curves. A clustering approach is used to divide the set of targets in multiple subsets, in order to reduce the computation time by estimating a single optimal parametrisation for each considered subset. The use of auto-RSM allows to reach high contrast at short separations, with a median contrast of 10^{-5} at 300 mas, for a completeness level of 95%. Detection maps generated with different approaches are used along with contrast curves, to identify potential planetary companions. A new planetary characterisation algorithm, based on the RSM framework, is developed and tested successfully, showing a higher astrometric and photometric precision for faint sources compared to standard approaches. A correlation study between achievable contrasts and parameters characterising HCI sequences highlights the importance of the Strehl, wind speed and wind driven halo to define the quality of high contrast images. Finally, planet detection and occurrence frequency maps are generated and show, for the SHARDDS survey, a high detection rate between 10 and 100 au for substellar companions with mass $>10 M_J$.

"The SHARDDS Survey: Limits on Planet Occurrence Rates Based on Point Sources Analysis via the auto-RSM Framework" accepted for publication in A&A

Completeness curve and completeness matrix functions introduced in PyRSM:

<https://github.com/chdahlqvist/RSMmap>

Completeness curve and completeness matrix functions introduced in VIP:

https://github.com/vortex-exoplanet/VIP/tree/master/vip_hci/metrics

E.2 Publications as co-author

Exoplanet imaging data challenge: benchmarking the various image processing methods for exoplanet detection, SPIE Proceedings Volume 11448, Adaptive Optics Systems VII; 114485A (2020)

Constraints on the nearby exoplanet ϵ Indi Ab from deep near- and mid-infrared imaging limits, A&A, 651 (2021) A89

E.3 Conferences and seminars

RSM detection map for direct exoplanet detection in ADI sequences, AO4ASTRO workshop, Laboratoire d'Astrophysique de Marseille, France (28/03/2019)

RSM detection map for direct exoplanet detection in ADI sequences (Poster), In the Spirit of Lyot, Astrobiology Center, National Institutes of Natural Sciences, Tokyo, Japan (21/10/2019)

Improving the RSM map exoplanet detection algorithm, HCI post-processing Workshop, Harnack Haus, Max Planck Institut für Physik, Berlin (28/01/2020)

The SHARDDS Survey: Limits on Planet Occurrence Rates Based on Point Sources Analysis via the auto-RSM Framework, In the Spirit of Lyot, University of Leiden, Leiden, The Netherlands (27/06/2022)



Advanced Data Processing Techniques for Exoplanet Detection in High Contrast Images

Carl-Henrik Dahlqvist

Carl-Henrik Dahlqvist holds a Master's degree in Business Engineering from the Université catholique de Louvain (LSM). After having worked 2 years for SWIFT as a financial analyst, he decided to go back to university to pursue a Ph.D. in Economics in co-supervision between the Université catholique de Louvain (LSM) and the Université de Namur (CeReFiM), where he was also a teaching assistant in Finance. Along his Ph.D. in Econophysics, he started in 2014, a 2 years preparatory program for the M.Sc. in Particle Physics and Cosmology at the Université Catholique de Louvain. He completed this M.Sc. in Particle Physics and Cosmology in June 2018, being awarded the Prix Pierre Macq and the Belgian Physical Society prize for his master's thesis.

After having received his Ph.D. in Econophysics from the Université de Namur and the Université catholique de Louvain, he started in October 2018, a second Ph.D. fellowship in exoplanet direct imaging at the Space sciences, Technologies and Astrophysics Research (STAR) Institute of the Université de Liège. His current interests lie in Exoplanets Detection, High Contrast Imaging, Optics, Deep Learning, Spaces and Planetary Sciences as well as Satellite Engineering. Following his second Ph.D. in Space Sciences, he joined Spacebel in September 2022 as Space Physics Engineering to contribute to the development of the Belgian space sector.