UCLouvain
Institute for Information
and Communication Technologies,
Electronics and Applied Mathematics

Low-Rank Techniques for Exoplanet Detection in High Contrast Images

Hazan Daglayan Sevim

Thesis submitted in partial fulfillment
of the requirements for the degree of
Ph.D. in Engineering Sciences

Dissertation committee:

Prof. P.-A. Absil (UCLouvain, advisor)
Prof. Nicolas Gillis (University of Mons, advisor)
Prof. Laurent Jacques (UCLouvain, Belgium)
Prof. Mariya Ishteva (KULeuven, Belgium)
Prof. Jean-Charles Delvenne (UCLouvain, Belgium)
Dr. Faustine Cantalloube (Univ. Grenoble Alpes, France)

Version of December 30, 2024.

Abstract

"Is there life beyond Earth?" This question has captivated humanity for centuries, driving efforts to discover planets beyond our solar system. Among the various techniques developed for exoplanet detection, direct imaging stands out as one of the most promising, as it allows for the direct capture of signals from planets themselves.

Direct imaging offers crucial insights into planetary systems but encounters significant challenges, primarily due to the close angular separation between faint planetary signals and the bright host stars. Even with state-of-the-art instruments, advanced post-processing techniques are required to remove quasi-static speckles that obscure planetary signals. This thesis tackles these challenges by proposing novel algorithms and models using low-rank approximations aimed at enhancing exoplanet detection through the use of angular differential imaging datasets. The datasets, obtained from ground-based telescopes, track the host star over a few hours of a night, capturing the planet's motion induced by Earth's rotation. These three-dimensional datasets in the time and spatial domains contain both planetary signals and quasi-static speckles and star signals, which require separation for accurate detection.

The core contributions of this work involve developing methods to isolate and subtract the quasi-static components, enhancing the visibility of planetary signals. Techniques based on using Laplacian distribution rather than Gaussian, which we propose in this thesis, play a pivotal role. Additionally, we proposed alternative solutions such as sparse modeling and matrix completion—where planetary signals are removed and later reconstructed—to improve detection accuracy. All proposed methods were thoroughly tested on both real and synthetic datasets, and their performance was evaluated against well-known algorithms in the literature, demonstrating significant improvements.

To those who believe in and struggle for a better planet.

Acknowledgments

This thesis represents not just the culmination of years of work here in Belgium but a dream I have carried and nurtured for much longer. The journey began back in Türkiye, where I first envisioned this moment and started working towards it. Even during my time in Germany, this dream was always in the back of my mind. Along this long journey, I've had so many people to thank.

First and foremost, I must express my deepest gratitude to my supervisor, Prof. Pierre-Antoine Absil. This is not merely the traditional thanks given to an academic advisor; it is far more than that. Words fall short in describing my immense gratitude to him. His passion for academia, unbiased and realistic approach to problems, fairness, diligence, and willingness to support his students have been nothing short of inspirational. I was often amazed by how much he pondered my challenges, sometimes even more than I did myself, which motivated me to work even harder. Beyond helping me solve academic puzzles, he always found solutions to my personal challenges during this five-year journey, often before I even voiced them. I consider him a father figure in academia. His guidance has strengthened my faith in science and humanity, and I am convinced that as long as there are people like him in this planet, it will remain a better place.

I also owe deep gratitude to my co-advisor, Prof. Nicolas Gillis. His combination of intelligence, hard work, kindness, and warmth, along with the ease of reaching him during work, receiving detailed explanations to my questions, and having the privilege of being a student of a professor who approaches his work with such enthusiasm, made me incredibly fortunate.

Although not officially my advisor, I also owe thanks to Prof. Lau-

rent Jacques for his invaluable insights and thought-provoking questions, which often opened new doors in my thinking. His ideas and the way he approached problems were a constant source of admiration and inspiration. I am also grateful for the enjoyable conversations we occasionally shared, which added a touch of lightness to this challenging journey.

To my jury members — Prof. Mariya Ishteva, Prof. Jean-Charles Delvenne, and Dr. Faustine Cantalloube, — thank you for your valuable feedback and sincere demeanor that helped me navigate this stressful process. I am grateful for your precious time and ideas. A special thanks to Faustine for her friendly and insightful contributions to both my thesis and papers in the field of astrophysics.

I am also immensely grateful to Prof. Olivier Absil for his guidance over the past four years. His answers, often simple for him but groundbreaking for me, have been life-saving more times than I can count.

One of my greatest fortunes was working with Simon Vary. I'm so glad we were Pierre-Antoine's students/postdocs around the same time. I'm certain that if Simon hadn't been involved in this work, the process would have been much more difficult for me. He was always helpful and guiding. I owe him so much in this thesis.

My colleagues at Euler deserve heartfelt thanks for their warmth, which made me feel at home in Belgium. A special thanks to Philémon, with whom I had the pleasure of sharing an office. His constant kindness, warmth, and unwavering support have meant so much to me. He was always there to listen, offer solutions to my problems, and provide encouragement whenever I needed it, for which I am deeply grateful. Cécile and Loïc, for their kindness, friendly and welcoming attitudes, and for sharing an office with me, made every day more enjoyable. Ousmane, for his invaluable help that began even before I officially started my PhD, his warm and engaging conversations, and his constant support throughout this journey—I am truly thankful. To Yousef and Moslem, thank you for the enjoyable conversations during the final stages of my PhD. Brieuc, I deeply appreciate your help during our assignment preparation and evaluation process, and Rémi, thank you for the pleasant chats and transportation support—it meant a lot to me. Renato, your cheerful and warm demeanor consistently brought positivity to my days, and I am so grateful. My heartfelt thanks also go to the incredibly helpful secretaries, Marie-Christine and Pascale, who made navigating administrative challenges so much easier. To Etienne and John, thank you for always assisting me with technical issues, never failing to offer help with a smile. Lastly, my grati-

tude extends to everyone at Euler and the ICTEAM institute for contributing to such a welcoming and supportive environment.

I am deeply grateful to the STAR Team at the University of Liège for enriching my knowledge in astrophysics and for giving me numerous opportunities to discuss my work, provide suggestions, and extend their hospitality. A special thanks to Mariam, who from the very first moment we met, made me feel as if we had known each other for years—sharing a hotel room during our conference felt effortless and natural—, and Valentin, I will always cherish their warmth and kindness.

It is time to thank my dear friends, who never left me alone during this long and challenging journey. To Senem, who has been by my side since the earliest days of university and with whom I have long surpassed the boundaries of friendship to become like sister. Together, we even dreamed and planned for pursuing a PhD abroad. To Ozan, who consistently motivated both of us in our abilities, and to Atlas—thank you all from the bottom of my heart. I am so grateful to have you in my life.

Arka Bahçemiz (Cansu, Ferhat, Hasan, Gizem, and George), who have been my greatest source of support throughout this journey. You made even the hardest days of the pandemic bearable, and I know that no matter the obstacles, borders, or challenges, nothing could ever stop us from coming together if I needed help or faced tough times. I'm so incredibly lucky to have such a wonderful group by my side. Thank you for everything.

Ender and Çiğdem, my dear fellow doctoral travelers, thank you for sharing this PhD journey with me. From navigating similar challenges to making endless future plans—some as ambitious as daily new business ideas that never quite materialized—your presence made this experience so much more meaningful. Those coffee break conversations were like a breath of fresh air, and the strength I drew from our time together will always hold a special place in my heart. Thank you for being there.

To my dear friends and cousins in Belgium; Güncem, Tayfur, and Nora, thank you for moving in Belgium, which was my dream, and so having a family in Belgium. I'm so grateful for your presence in my life and the support you've given me; Yağmur, Eyüp, and Defne, thank you for your unwavering support during both my PhD and my job search. Knowing that I had two extra doors in Belgium I could knock on without hesitation was such a comforting feeling. And to Zeynep, who entered my life here in Belgium—thank you for being someone I could always call on, whether to ask for help, share a laugh, or just escape the routine. Even though your

house is guarded by your little “monsters”, which makes me hesitant to knock on your door, your warmth and friendship during these four years have been invaluable to me.

Eda, who is more than a friend - a member of my family, thank you for standing by me in all my experiences, celebrating my successes, comforting me in my struggles, and sharing countless moments that have shaped who I am today.

When we set out for Europe, pursuing a PhD here was always part of the plan. That’s why I must express my heartfelt thanks to Gülcan abla, Celal dayı, İhtilal dayı, and Arin, who provided incredible support during our time in Germany. They opened their doors and hearts to us, and their kindness and generosity mean the world to me. You all hold a very special place in my heart.

Continuing my acknowledgments with the Türkiye chapter of this journey, I want to thank my dear friend İmir, who was always there to support me during my long journay. Despite the distance, your unwavering friendship means a lot to me. To Çikom Özlem, thank you for being a constant source of encouragement and support, always making me feel that you are right there beside me.

I am deeply grateful to Prof. Murat Manguoglu, whose brilliance and dedication I greatly admire. He introduced me to topics like matrix approximations and sparse matrices, fostering a deep appreciation for these subjects and paving the way for me to work on something as exciting as exoplanet detection. During my years in Turkiye, it was a privilege to collaborate with him and learn from his expertise. I am also immensely thankful for his encouragement to pursue my studies abroad, supporting me in taking this significant step in my academic journey.

Finally, I want to express my deepest gratitude to my family—my mother Mercan, my father Ali, and my siblings Önder, Öner, Caner, and Sezer—who have stood by me every step of the way, supporting me in every decision I’ve made. You have always been my unwavering pillars of strength. To my family-in-law—Fatma, Baki, and Ekin—thank you for being there for me with the same love and support. I feel incredibly lucky to have such big-hearted families who care for me more than I could ever imagine. A special thanks to my parents for constantly motivating me to pursue my education, reminding me of the importance of girls’ education, and raising me with this mindset that has shaped who I am today.

And finally, to my husband, Emre — my best friend, sometimes my parent, sometimes my child, my chef, my listener, my supervisor, my greatest

helper, and, above all, my other half — thank you. I am absolutely certain that this thesis would not exist without you. I would not be here, and I could not have faced and overcome all these challenges on my own. You fed me, nurtured me, and brought order to my life, always thinking of me more than I thought of myself. There were even times when I couldn't express myself to others, and you stepped in to explain me to them. You gave me your share of everything and made sacrifices so that I could keep moving forward. Beyond your incredible emotional support, you enabled my journey abroad, enhanced my visuals, shifted my perspective, and always found ways to lift me higher. Words fail to truly express my gratitude. If I have achieved anything today, it is because you were steadfastly by my side. For all of this and more, I am eternally grateful.

Contents

List of Figures

List of Tables

Acronyms

ADI Angular differential imaging. 10

AMAT Alternating Minimization Algorithm with Trajectory. 76

AO Adaptive Optics. 9

EIDC Exoplanet Imaging Data Challenge. 96

ESA European Space Agency. xv, 3

FP False Positive. 34

FPR False Positive Rate. 34

FWHM Full Width at Half Maximum. 32

HCI High-contrast imaging. 8

IAU International Astronomical Union. 1

L1-LRA L1 low-rank approximation. 61

LAD least Absolute deviation. 45

LLSG Local Low-rank plus Sparse plus Gaussian-noise decomposition. 26

LRM likelihood ratio map. 45

LRPT low-rank plus sparse trajectory. 54

MAYO morphological component analysis. 53

MC Matrix Completion. 111

MCMC Markov Chain Monte Carlo. 100

NAHT Normalized Alternating Hard Threshold. 57

NEGFC Negative Fake Companion. 32

NMF Nonnegative Matrix Factorization. 24

PCA Principal Component Analysis. 20

PSF Point Spread Function. 11

ROC Receiver Operating Characteristic. 29

SDI Spectral Differential Imaging. 12

SNR or S/N Signal-to-Noise Ratio. 29

STIM Standardized Trajectory Intensity Mean. 29

SVD Singular Value Decomposition. 20

TP True Positive. 34

TPR True Positive Rate. 34

Nomenclature

$\hat{L}$ _____ low-rank solution of L2-norm (p. 62)

$\mathrm{H}_r^{\mathrm{SVD}}(\cdot)$ _____ randomized singular value decomposition (p. 57)

Λ _____ set of all feasible planet signatures (p. 76)

$\Lambda_g^{\mathrm{Gauss}}(R)$ _____ likelihood ratio using the Gaussian assumption (p. 46)

$\Lambda_g(R)$ _____ likelihood ratio (p. 46)

$\langle \cdot, \cdot \rangle$ _____ trace inner product (p. 56)

$\log \mathcal{L}_g^{\mathrm{Gauss}}$ _____ log-likelihood (p. 43)

$\mathcal{F}$ _____ Fourier transform (p. 140)

$\mathcal{F}^{-1}$ _____ inverse of the Fourier transform (p. 142)

$\mathcal{P}_\Omega$ _____ sampling operator (p. 112)

$\mathbb{1}_{i,j}$ _____ canonical basis matrix with a single one at index i, j (p. 56)

$\mathrm{rank}(A)$ _____ rank of the matrix A (p. 62)

Ω_g _____ the set of indices of pixels (p. 43)

Ω_{ρ_x,ρ_y} _____ mask for the Fourier transform with radi ρ_x and ρ_y (p. 146)

Ω_ρ _____ mask for the Fourier transform with radius ρ (p. 142)

$\Psi(P)$ _____ linear trajectorlet transform of P (p. 55)

Ψ^* _____ adjoint of Ψ (p. 56)

ρ disk size (p. 43)

ρ^2 coefficient of determination (p. 69)

σ_R^2 empirical variance of the residual frames (p. 43)

a_g^* flux which maximizes log-likelihood of Laplacian (p. 45)

a_g^{Gauss} flux which maximizes log-likelihood (p. 44)

a_g flux (p. 42)

E noise term (p. 42)

g unknown position of the planet (p. 42)

L low-rank background (p. 62)

M matrix of observations (p. 62)

$N \times N$ size of the images (p. 42)

P non-negative s-sparse matrix (p. 55)

P_g planet signature cube (p. 42)

R residual cube (p. 42)

T number of images (p. 42)

1

Introduction

The word 'planet' comes from the Ancient Greek word 'planetes', which means wanderers, lights moving across the sky. This definition also counted Earth's Moon and Sun as planets. It then evolved over time to mean objects orbiting around a star. Pluto was identified as the 9th planet when it was discovered in 1930. But because it was so much smaller than the others, even smaller than the moons of some planets, as discoveries increased over time, so did the debate over the definition of a planet. In 2006, the International Astronomical Union (IAU) announced a formal definition of a planet, reclassifying Pluto as a dwarf planet. According to this definition, a planet must:

(a) be in orbit around the Sun,

(b) has sufficient mass for its self-gravity to overcome rigid body forces so that it assumes a hydrostatic equilibrium (nearly round) shape, and

(c) have cleared the neighborhood around its orbit.

The IAU's definition of a planet is primarily tailored to our Solar System and does not fully encompass the complexities of defining exoplanets. This gap has led to alternative proposals and discussions. Initially, definitions focused on characteristics like mass, with distinctions drawn at the deuterium-burning threshold of approximately 13 Jupiter masses

(M_J) [SBM11]. However, subsequent observations, such as the discovery of a "dry region" in the mass distribution between 25–40 M_J, have prompted debates about whether this threshold adequately separates exoplanets from brown dwarfs [SDLS+11]. Some recent findings challenge the notion of a clear boundary, suggesting the need for more nuanced criteria that consider chemical composition and formation mechanisms [DK13, MV17], even if the recent IAU working definition categorized them according to limiting mass with no matter how they are formed [DEL22].

One key challenge in defining exoplanets is incorporating free-floating planetary-mass objects, which do not orbit a central star but are found in isolation. Such objects have been identified through methods like gravitational microlensing [SKU+11] and direct imaging [DGM+12]. According to the recent IAU definition, free-floating objects in young star clusters are not classified as planets [DEL22].

Future studies will likely integrate multiple lines of evidence, including mass, metallicity, formation scenarios, and observational properties, to refine exoplanet definitions. As our understanding and insight into the fundamental nature of these objects steadily progress, the definition will continue to adapt and evolve.

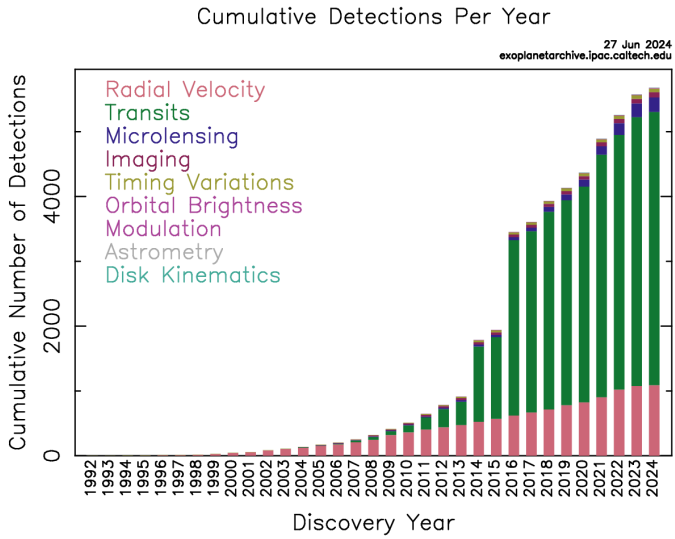

Fig. 1.1 Discoveries of exoplanets

Since the first detection of an exoplanet in 1992, more than 5000 plan-

ets have been discovered beyond our solar system. The majority of these discoveries have been made using indirect methods, as shown in Fig. 1.1. Only 1.4 percent of these planets have been observed by direct imaging, which captures light from the planet itself.

Before we delve into the direct imaging method, it is important to understand the strengths and weaknesses of these indirect methods to appreciate the challenges and advancements in the field of exoplanet research.

1.1 Indirect Methods

1.1.1 Transit method

When a planet passes in front of its host star, the star's brightness temporarily decreases a bit as visualized in Fig. 1.2. This drop in brightness depends on the size of the planet. The idea of detecting planets through this method was first proposed by O. Struve in 1952 [Str52], but the first actual transit observation of a known exoplanet was made by Charbonneau et al. in 2000 [CBLM99]. The first exoplanet discovery using this method was by Konacki et al. in 2003 [KTJS03]. Today, about a quarter of all detected exoplanets have been found using the transit method.

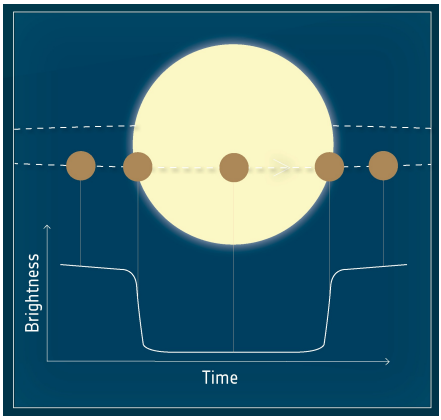

Fig. 1.2 Illustration of the transit method. Credits: European Space Agency (ESA) - Exoplanet detection methods

This method allows astronomers to determine the size of the planet based on the extent of the dip in the light of the host star. For the transit method to work, the orbital plane of the planet must align closely with our

line of sight from Earth. This alignment allows the planet to cross in front of the star relative to our viewpoint. The method is more effective for finding larger planets relative to their stars since they create a more noticeable dip in brightness. It is also more effective for planets located closer to their stars, since multiple transits are required to confirm a detection.

However, the transit method has several challenges and biases. A transit-like signal can be produced by other phenomena, such as eclipsing binary stars, brown dwarfs, background eclipses, or stellar spots, leading to false positives. Therefore, additional confirmation using other techniques, like radial velocity (RV), as suggested by Bouchy et al. (2009), or through in-depth analysis with transit timing variations (TTV), which involve detecting changes in the timing of a planet's transits due to gravitational interactions with other planets, as proposed by Holman et al. (2010), is often required to validate the presence of a planet.

1.1.2 Radial velocity

Typically, we think of planets orbiting their host stars due to the star's gravitational pull. However, gravity is a mutual force, meaning that both the planet and the star influence each other. In reality, both the planet and the star orbit a common center of mass, not just the planet around the star. As a result, if a planet is present, the star doesn't remain completely still but exhibits a slight wobble in its motion, which can be detected through the Doppler effect.

When the planet and star are orbiting, they are always on opposite sides of the center of mass, which means they both complete their orbits in the same amount of time. Knowing how long it takes the star to complete its orbit allows us to determine the orbital period of the planet as well. Additionally, we can estimate the mass of the planet from this motion. If the planet is heavier, it causes the star to move in a larger orbit. Since the planet and star must complete their orbits in the same time, a larger star orbit means the star has to move faster. This is why measuring the speed of the star gives us clues about the mass of the planet.

The radial velocity method is based on measuring the velocity of the star. As the star moves toward us during its orbit, its light becomes blue-shifted, and as it moves away, it becomes red-shifted, as shown in Fig. 1.3. By tracking these shifts, we can determine the planet's orbital period. If the planet is large enough, its gravitational influence will cause a noticeable change in the velocity of the star. However, for smaller planets, the

star's wobble is much smaller, making it more difficult to detect. Another limitation of this method is that if the star's orbit doesn't involve movement toward or away from Earth, there will be no blue or red shift, and the method won't work in such cases.

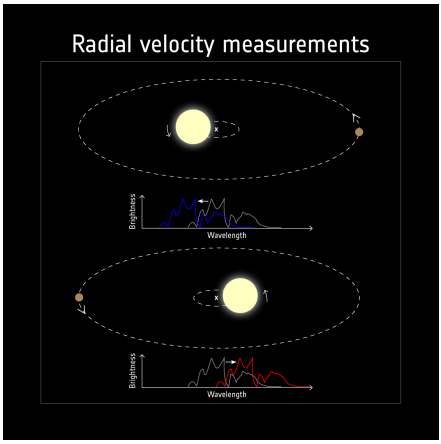

Fig. 1.3 Illustration of the radial velocity detection method. Credits: ESA - Exoplanet detection methods

The first discovery of a planetary companion with a few Jupiter masses using the radial velocity method was made in 1995 by Mayor and Queloz [MQ95]. This groundbreaking discovery earned them the Nobel Prize in Physics in 2019. Since then, around 19% of the exoplanets detected so far have been discovered through this method [NAS24]. The radial velocity method was one of the first successful techniques for identifying exoplanets, and it is also frequently used to confirm the existence of planets discovered through other methods.

1.1.3 Gravitational microlensing

Gravitational microlensing occurs when a star passes between Earth and a more distant star, causing the light from the distant star to bend due to the gravitational pull of the closer star. This effect acts like a lens from Earth's perspective, making the distant star appear as though it has two versions. The brightness of the distant star increases, peaking when the two stars and Earth are perfectly aligned. If the closer star has a planet orbiting it, the planet's presence creates a small blip in the brightness curve, indicating its existence, as shown in Fig. 1.4.

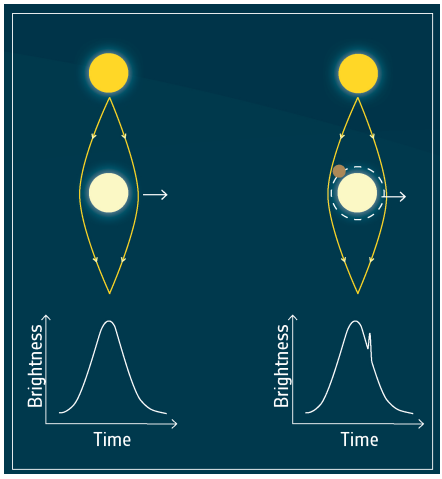

Fig. 1.4 Illustration of the microlensing detection method. Credits: ESA - Exoplanet detection methods

This method is particularly effective for detecting planets that are far from their host stars and can also be used to find smaller planets that other techniques might miss. However, a key limitation of microlensing is that it is non-reproducible, since it relies on a rare alignment of a background star, a foreground star, and Earth. The first exoplanet discovered using this technique was published by Bond et al. in 2001 [BAD+01], and to date, around 3.9% of all detected exoplanets have been found using microlensing [NAS24].

1.1.4 Pulsar timinig

Pulsar timing was the very first method used to discover an exoplanet in 1992 [WF92]. Pulsars are a special type of star called a neutron star, which is the leftover core of a intermediate mass star that exploded in a supernova. Neutron stars are incredibly dense and have intense gravitational and magnetic fields. Pulsars, however, are a subset of neutron stars that emit intense electromagnetic beams, often shining in the radio wavelength domain with very rapid and precise periods (ranging from seconds to milliseconds). These beams are detectable only when they align with our line of sight, which explains why we observe relatively few pulsars—the probability of being within their beam's path is rather low. These stars spin very fast, sending out regular pulses of radio waves.

In some cases, the timing of these pulses may be irregular, meaning the pulses arrive a little early or a little late. This annomaly in timing is caused by the gravitational pull of one or more planets orbiting the pulsar. By studying these small changes, astronomers can figure out details about the planets, like their mass and orbit.

While this method has been crucial for discovery, pulsars are rare, and so are planets around them. So far, only a couple of dozen exoplanets have been found this way. The planets discovered are often close to the pulsar and can be Earth-sized or even larger.

1.2 Direct Imaging

Direct imaging is a technique used to directly observe exoplanets by capturing the photons that are either emitted by the planets. For this method to be effective, the planet must have a wide enough angular distance from its host star to separate its light from the star's brightness, and the planet itself must be bright enough to stand out.

The first exoplanet discovered through direct imaging, which is given in Fig. 1.5, was a giant planet orbiting the brown dwarf discovered by Chauvin et al. in 2004 [CLD+05]. This planet was easier to detect due to the relatively low temperature of the brown dwarf and the large separation between the two. The contrast between the brown dwarf and the planet was only 10^{-2}, which made the planet's light detectable.

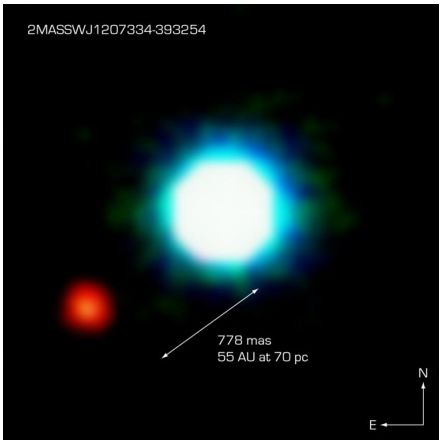

Fig. 1.5 The first exoplanet detected by direct imaging. Credits: ESA

1.2.1 Advantages of direct imaging

High-contrast imaging (HCI) is an essential observing technique for exoplanet discovery, especially because of its ability to directly observe young, massive planets orbiting their host stars at large distances [GM23] that indirect methods have difficulty detecting, as can be seen in Fig. 1.6. This direct observation method is complementary to indirect techniques and also provides crucial information such as bolometric luminosity, effective temperature, surface gravity, composition of the planets and, assuming a formation and evolutionary model, an estimate of their mass [Bow16].

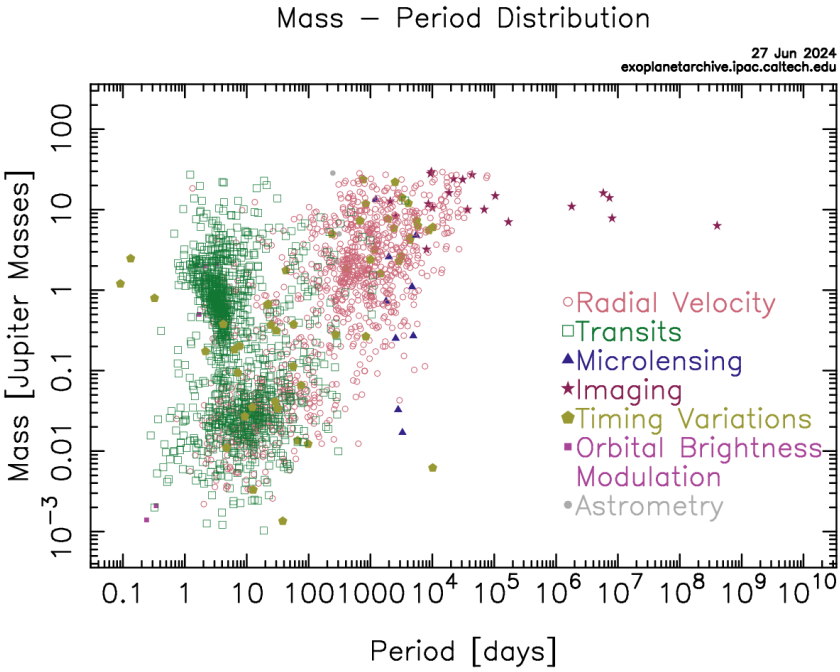

Fig. 1.6 The mass vs the distance from the host star for each planet categorized according to the detection methods.

1.2.2 Challenges of direct imaging

High-contrast direct imaging from ground-based telescopes faces three significant challenges:

1. Image distortion due to the Earth's turbulent atmosphere and the optical flaws of the telescope and instruments,
2. The small angular distance between stars and their planets, and
3. The vast difference in brightness between stars and planets, with contrasts usually ranging from 10^{-3} to 10^{-10}.

These challenges explain why direct imaging is only successful in detecting giant planets at large distances. To address these difficulties, progress in this field is supported by remarkable technological advances in extreme adaptive optics (AO) systems [Guy18] and coronagraphy, for better atmospheric turbulence correction and raw stellar light suppression, respectively. These technologies equip new generation HCI instruments, such as VLT/SPHERE [BVM+19], Gemini/GPI [MGI+14] or Subaru/SCExAO [MGL+09] allowing them to achieve unprecedented raw contrasts.

- Adaptive Optics

  Adaptive Optics (AO) enhances image quality by correcting environmental and optical distortions, leading to clearer and sharper images. It addresses image distortions caused by factors like atmospheric turbulence or optical system errors. The concept was introduced by Babcock in 1953 as a way to mitigate the effects of turbulence using deformable mirrors (DM) [Bab53]. In 1990, AO was first used in astronomy at the La Silla Observatory [RFK+90]. Figure 1.7 demonstrates the image-enhancing capabilities of AO in SPHERE, the Spectro-Polarimetric High-contrast Exoplanet REsearch instrument [BFD+08], which is a high-contrast imaging device installed on the VLT ESO telescope.

- Coronagraphy

  A coronagraph is a tool used to observe exoplanets that are typically obscured by the intense brightness of the stars they orbit. It works by minimizing the light from the star, making it easier to distinguish the light emitted by the planet. Essentially, a coronagraph blocks light from the central (on-axis) source while allowing light from off-axis sources to pass through with minimal interference. The original concept, introduced by Lyot in 1939, was created to observe the Sun's

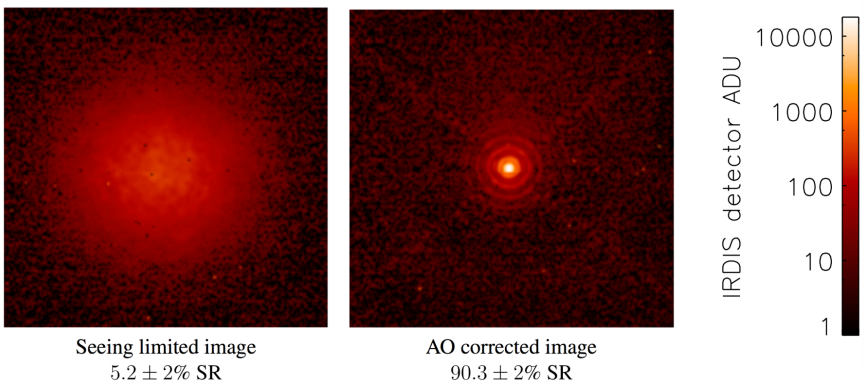

Fig. 1.7 Adaptive optics correction. Credits: [SFP+16]

corona [Lyo39]. The classical Lyot coronagraph typically consists of an opaque mask at the focal plane that blocks on-axis light and a smaller pupil, known as the "Lyot mask," which reduces diffracted starlight.

Despite these important steps, direct imaging of exoplanets remains a challenging task. Only 1.4% of known exoplanets have been discovered using this method [NAS24] and only a few of the known exoplanets have been directly imaged. It is because currently we can only directly image hot planets with a rather large separation to the star. This limitation underscores the inherent difficulty in distinguishing these faint planets from the overwhelming brightness of their host stars. Star light residuals under the form of speckles, can be seen in Fig. 1.8, caused by atmospheric turbulence and imperfections in telescopes and instruments, are very similar in shape and brightness to planets, posing a major challenge to direct imaging.

1.3 Observing Strategies

1.3.1 Angular differential imaging

Angular differential imaging (ADI), illustrated in Fig. 1.9, has emerged as a common strategy [MLD+06]. ADI involves capturing a sequence of frames in pupil-stabilized mode, where the field of view rotates by a given angle (the parallactic angle). This setup keeps the target star centered in

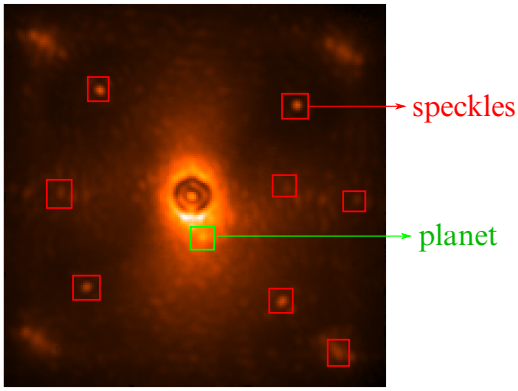

Fig. 1.8 A frame obtained with Gemini-S/GPI instrument. Green square represents the planet, while red squares represent the speckles.Although everthing except for planet is steller light in the image, this visualization emphasizes the similarities between the planet and the speckles.

the image. As a result, the star and speckles appearing static or quasi-static, while planets exhibit movement due to changes in the parallactic angle, a function of Earth's rotation and the star's position in the sky. By exploiting this differential motion, ADI enables the isolation and detection of exoplanetary signals from the surrounding noise, increasing the chances of successful direct imaging.

In the ADI sequence, a model called the reference Point Spread Function (PSF) is constructed to include both static and quasi-static signals from the star and surrounding speckles. Once the reference PSF is accurately modeled, it is subtracted from each frame in the ADI sequence, effectively removing the star's signals and leaving behind the residuals, which include the rotating planet. Usually, to enhance the planet's signal, the residual frames are then aligned using known parallectic angles to a common center based on the expected motion of the planet. These aligned frames are subsequently combined, typically using the mean or median to reduce noise, and it is called processed frame. The final step involves generating a detection map, which highlights the presence of the planet. This detection map is created using various algorithms and techniques described in Section 3.1, ensuring robust identification of the planetary signals.

The basic approach to obtain the reference PSF is to use the temporal median of frames in the ADI sequence [MLD+06]. The median frame,

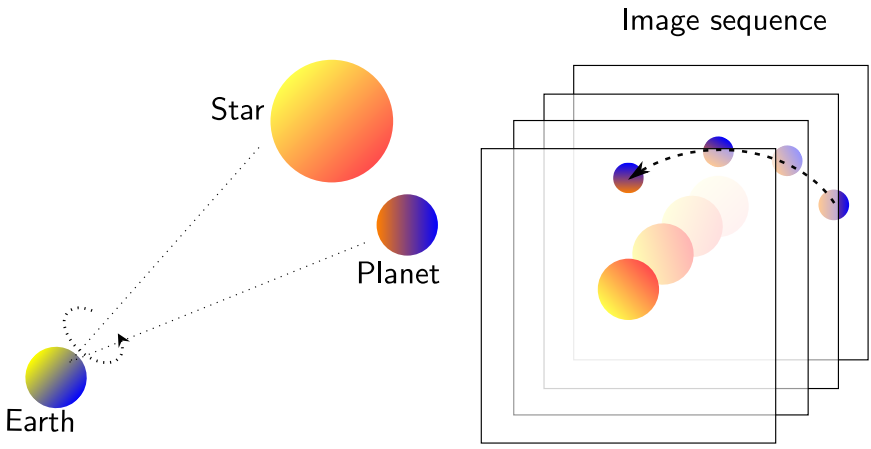

Fig. 1.9 The illustration of ADI

which captures the quasi-static stellar light, is then subtracted from each frame in the sequence. This approach has several variations, but the general idea is to construct a reference datacube, representing the stellar light, and subtract it from the original cube to enhance planet detection. This approach, while effective in isolating planet signals, introduces the concept of self-subtraction, where some of the planet's signal may also be removed during the subtraction process. So the algorithms aim to minimize self-subtraction by optimizing the separation of the planetary signal from the stellar light, ensuring more accurate detection.

1.3.2 Spectral differential imaging

Spectral Differential Imaging (SDI) is a technique designed to detect planetary signals by using differences in images captured at various wavelengths, rather than relying on azimuthal rotation as in ADI. The key idea behind SDI is that while the position of a planet remains fixed across different wavelengths, the star's speckles stretch in a wavelength-dependent manner. This property allows the planet's signal to stand out from the speckle noise. The concept was first introduced by Racine et al. in 1999 and has since become a widely used strategy in exoplanet detection [RWN+99].

In practice, SDI works by capturing images of the star and its surroundings at two or more wavelengths, so it is utilised the behaviour of the planet

signal and speckles being different. By scaling one image relative to the other, a reference PSF model is created [SF02]. This PSF model represents the star's signal and associated speckles, which can then be subtracted from the original images to isolate the planetary signal. The remaining signal contains the planetary signal, which can be further processed for detection. By rescaling the images, the planetary signals are aligned to the same position across different wavelengths, making it easier to confirm the presence of an exoplanet.

In recent developments, researchers have combined SDI with ADI to take advantage of both techniques [CCA+19, KBQ+21]. ADI uses the Earth's rotation to create differential motion between the star and planet over time, while SDI exploits the spectral differences. The combination of these methods results in a four-dimensional data cube, where the dimensions are wavelength, time, and spatial position.

1.3.3 Other differential imaging techniques

There are several other differential imaging techniques that have been developed to enhance the detection of exoplanets. Polarimetric Differential Imaging (PDI) works on the principle that light reflected from a planet or scattered by a circumstellar disk becomes polarized [KPP01]. By capturing both linear polarizations at the same time, the images can be subtracted from one another, revealing the circumstellar objects hidden by the star's light.

Another approach is Reference Differential Imaging (RDI), where the PSF model is created by observing a nearby reference star with nearly the same brightness as the target star [RCL+12, WMR+21]. By using this reference star, which has similar characteristics but lacks planetary signals, the PSF of the observed star can be more accurately subtracted, allowing the exoplanet to be detected. RDI relies on using reference stars to model and subtract the speckle pattern of the target star. This method requires high instrument stability because even small optical misalignments can affect the accuracy.

On the other hand, Binary Differential Imaging (BDI) utilizes the secondary star in a binary star system to estimate the PSF model for the primary star [RWM+15].

1.4 Objectives of the Thesis and Outline

The main objectives tackled in this thesis are the following:

- Analyze the ADI cubes and the noise present in the images.
- Develop new models to extract static and quasi-static signals, thereby obtaining the planet signals.
- Implement and test the developed algorithms in practical scenarios to confirm their effectiveness. Use common metrics to compare these algorithms with state-of-the-art methods.

Let us now give a short description of all chapters.

Chapter 2: Preliminaries and background

Before proceeding further, it is beneficial to review the basics. Additionally, a quick look at the literature will help contextualize our work.

Chapter 3: Metrics for exoplanet detection algorithms

To compare the algorithms, we need to use the same metrics. First, we explain the common metrics used in direct imaging. Additionally, we propose an alternative method to obtain the ROC curve, a well-known metric, specifically for direct imaging to ensure a fairer assessment.

Chapter 4: Likelihood ratio map for direct exoplanet detection

We propose a new detection map, which allows us to detect the planet's location after model identification and all other processes. We compare this map with methods used in the literature and evaluate its performance. Additional experiments are presented in all subsequent chapters.

Associated publication:

> Daglayan, H., Vary, S., Cantalloube, F., Absil, P.-A., & Absil O. (2022). Likelihood ratio map for direct exoplanet detection. In 2022 IEEE 5th International Conference on Image Processing Applications and Systems (IPAS), volume 5, pages1-5, doi:10.1109/IPAS55744.2022.10052997.

Chapter 5: Low-rank plus sparse trajectory decomposition for direct exoplanet imaging

In this chapter, we propose a direct imaging method for the detection of exoplanets, based on a combined low-rank plus structured sparse model.

For this purpose, we develop a dictionary of possible effective circular trajectories that a planet can take during the observation period. The elements of this dictionary can be efficiently computed using rotation and convolution operations. To address the non-convex optimization problem, we design a straightforward alternating iterative hard-thresholding algorithm that simultaneously promotes a low-rank background and a sparse exoplanet foreground.

Associated publication:

> Vary, S., Daglayan, H., Jacques, L. & P.-A. Absil. (2023). Low-rank plus sparse trajectory decomposition for direct exoplanet imaging. 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1-5. doi:10.1109/ICASSP49357.2023.10096197

Chapter 6: Direct exoplanet detection using L1 norm low-rank approximation

We analyze the noise present in the data and propose a new model based on the noise distribution, utilizing the L1 norm. This model is compared against the more commonly used L2 norm to evaluate their respective performances. By examining the differences between the L1 norm and the L2 norm, we aim to demonstrate the advantages and potential improvements offered by our approach. Extensive experiments are discussed in Chapter 7, where the algorithm utilizes L1 norm low-rank approximation.

Associated publication:

> Daglayan, H., Vary, S., Leplat, V., Gillis, N., & Absil, P.-A. (2023). Direct exoplanet detection using L1 norm low-rank approximation. Proceedings of BNAIC/BeNeLearn 2023, pp. 1-13. doi:10.48550/arXiv.2304.03619.

Chapter 7: An alternating minimization algorithm with trajectory for direct exoplanet detection

We propose a new method called an alternating minimization algorithm with trajectory (AMAT) that advances the use of the low-rank-plus-trajectory model by alternating between computing the low-rank approximation of the speckle field and the planet flux. Our approach formulates

the model as both L1 and L2 norm minimization problems, allowing us to leverage the strengths of each norm in optimizing the detection process.

Associated publications:

> Daglayan, H., Vary, S., Absil, & Absil, P.-A. (2023). An alternating minimization algorithm with trajectory for direct exoplanet detection. Proceedings of the 31st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN 2023), pages 417–422. doi:10.14428/ esann/2023.ES2023-137.

> Daglayan, H., Vary, S., Absil, O., Cantalloube, F., Christiaens, V., Gillis, N., Jacques, L., Leplat, V. & Absil, P. -A. (2024). An alternating minimization algorithm with trajectory for direct exoplanet detection-The AMAT algorithm. Astronomy & Astrophysics, 692, A126. doi:10.1051/0004-6361/202451242

Additionally, we submitted the results of the datesets provided by Exoplanet Imaging Data Challenge Phase 2. We succeed in retrieving the injection position with a very good accuracy and the results of flux estimation are relatively good in general. The results were published in the following publication:

> Cantalloube, F., et al. (2024). Exoplanet imaging data challenge, phase II: comparison of algorithms in terms of characterization capabilities. Proceedings Volume 13097, Adaptive Optics Systems IX; 1309713 doi:10.1117/12.3020278.

Chapter 8: Matrix completion for direct exoplanet detection

In this chapter, we propose a new model to detect exoplanets. In order to not leave some planet signal in the PSF model, we propose to remove planet signal from the dataset and find a PSF model without planet signal. By removing the pixels containing the planet and applying matrix completion method to fill the data, we aim to achieve a better fit for static/ quasi-static signals. We compare the results of our approach with other methods in the literature to evaluate its effectiveness.

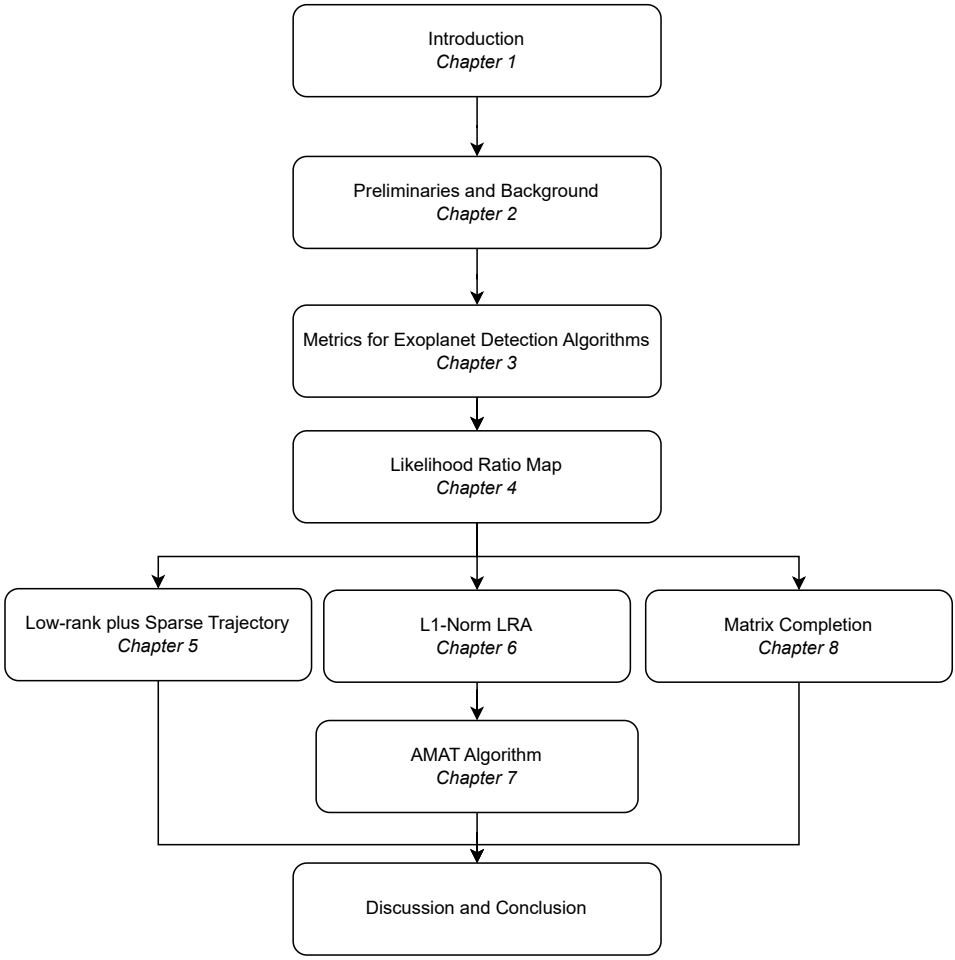

Fig. 1.10 The dependency structure of this thesis

2

Preliminaries and Background

2.1 An Overview of Post-processing Techniques for Direct Exoplanet Imaging

Using ADI sequences, several post-processing methods have been proposed. Among the most common ones are those that build a model for the stellar PSF (including the static and quasi-static speckle field), and subtract it to detect planets. To build a PSF model, methods such as principal component analysis (PCA) [AQ12, SPL12] or nonnegative matrix factorization (NMF) [GG+17, RPZ+18] aim to obtain a low-rank approximation of the time-by-pixel matrix containing the ADI observing sequence. After subtracting the PSF model from each frame, the residual matrix consists of both planetary signals and noise, which includes some non-dominant residual stellar signal that is dynamic or quasi-static signals in the images. Some studies suggest modeling the residual matrix as a sparse matrix and noise (LLSG) [GG+16]. Additionally, the Locally Optimized Combination of Images (LOCI) [LMD+07] employs a least-squares approach to construct a model PSF, which minimises the residuals locally, with variants. Further elaboration can be found in Sec. 2.2.

Besides PSF model subtraction, other post-processing methods are

based on inverse problem approaches to estimate the speckle field and planetary signal in a maximum-likelihood framework. One such method is The ANgular DiffeRential OptiMal Exoplanet Detection Algorithm (ANDROMEDA) [CMM+15], which uses maximum likelihood estimation to jointly estimate the planetary signal and the speckle noise. The Forward Model Matched Filter (FMMF), originally proposed by Pueyo [Pue16] and later refined by Ruffio et al. (2017) [RMW+17], applies the Karhunen-Loève transformation to analytically compute a forward-modeled planetary template. In order to reduce self and over-subtraction of planetary signals, PAtch COvariances (PACO) [FDTL18] combines a maximum likelihood estimator for flux estimation and a nonstationary multivariate Gaussian model to infer local spatial correlations and speckle noise statistics.

Furthermore, machine learning techniques have been increasingly applied to exoplanet detection. Supervised Detection and Reconstruction of Interferometric fringes using Recurrent Neural Networks (SODIRF) and Supervised Detection and Reconstruction of Interferometric fringes using Neural Networks (SODINN)[GAVD18] employ neural networks to distinguish planetary signals from speckle noise after PCA. Their refined variant, NA-SODINN (Noise-Aware SODINN)[CADVD23], further improves detection capabilities by incorporating noise characteristics into the training process of the neural networks, making it more robust to variations in speckle noise patterns.

2.2 PSF-Subtraction Techniques

2.2.1 PCA

Principal Component Analysis (PCA) is a statistical method that transforms a set of possibly correlated variables into a set of linearly uncorrelated variables known as principal components. When applying PCA to exoplanet detection, the main goal is to differentiate the light of an exoplanet from the overwhelming light of its host star. This process starts with a data matrix X, typically constructed from a series of flattened images from telescopic observations as shown in Fig. 2.1.

The PCA begins with the computation of the covariance matrix of the data set, which is the basis for determining the principal components since it captures the relationships between variables. The Singular Value Decomposition (SVD) is then applied to the covariance matrix, or directly to the centered data matrix, decomposing it into the product of three matrices:

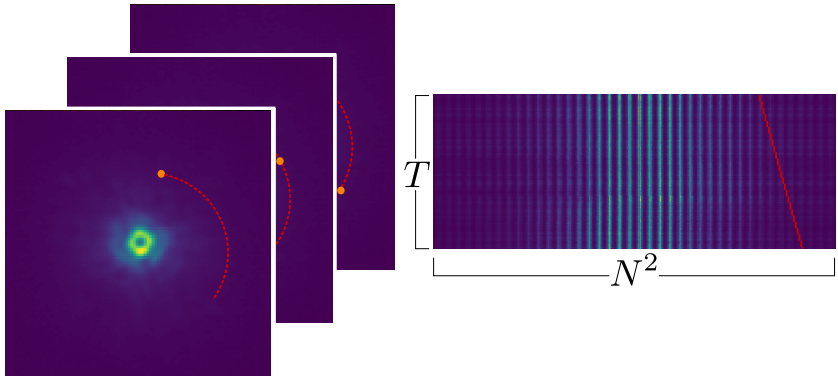

Fig. 2.1 Left: Data cube. Right: 2D representation of data cube. The red dashed path on left and the red line on right represent the trajectory of the planet.

$X = U\Sigma V^T$. In this decomposition, V (or its transpose depending on the orientation of X) contains the principal components, which are the eigenvectors of the covariance matrix, and represent the directions of maximum variance in the data space.

The columns u_i of U are called left singular vectors of X, and they form an orthonormal basis for the range of X. Each left singular vector u_i is associated with a singular value σ_i and a right singular vector v_i from the matrix V in the SVD. By applying a truncated SVD, which retains only the first r largest singular values and their corresponding singular vectors, we construct the low-rank matrix:

$$L = \sum_{i=1}^{r} \sigma_i u_i v_i^T, \tag{2.1}$$

which includes the static/quasi-static signals. By subtracting this low-rank matrix from X, the remaining residual matrix is expected to contain the planet signals.

In Fig. 2.2, the top sequence illustrates selected frames from the PSF model cube, with each frame representing the impact of incrementally increasing the rank from 1 to 4. The first principal components are adept at isolating the primary static signals within the datacube, resulting in an image that closely mirrors the image from the main datacube. Subsequent

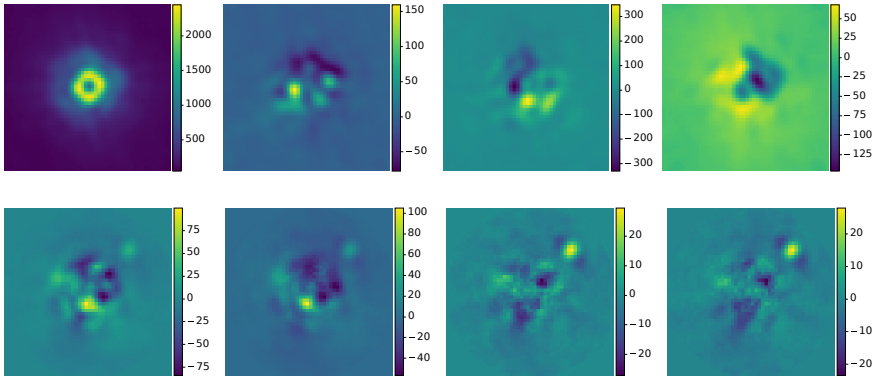

Fig. 2.2 Top: The effects of the principal components (PC). From left to right: 1st, 2nd, 3rd, and 4th PC. Bottom: Processing frames after we apply PCA using rank 1, 2, 3, and 4 from left to right. These images are generated using the β-Pictoris dataset, as detailed in Chapter 4.

ranks reveal additional nuances and complexities. Selecting the right rank for PCA is important since a rank that is too low will not capture the static or quasi-static signals, leaving a residual cube that is too noisy to detect planets. On the other hand, a high rank might incorporate planetary signals into the PSF model, thus obscuring these signatures in the residual datacube. The processed frames at different ranks are shown in the bottom row of the figure, which illustrates how increasing the rank gradually reveals the planetary signatures.

Annular PCA

In annular PCA [AMM+13], the data cube is divided into k annuli, denoted as X_{ann_i} for $i = \{1, \ldots, k\}$, as illustrated in Fig. 2.3. Each annulus is reshaped into a matrix whose columns correspond to the pixel values within the annulus over time. Consequently, the number of columns in this matrix equals the number of pixels in the annulus, while the number of rows corresponds to the total number of frames. A low-rank approximation L_{ann_i} is then computed for each annulus. Finally, all the low-rank matrices are combined to reconstruct the low-rank data cube L.

Randomized SVD

Although it is a perfect tool to crunch the data down into a low dimensional representation for extracting key features and pattern, com-

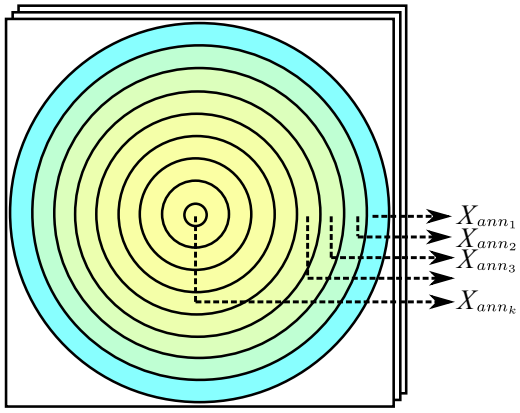

Fig. 2.3 Annuli splitting in PCA

puting the SVD of large datasets can be quite expensive. To decrese this computation time there is a technique that randomly samples the column space of the matrix X and with high probability, it identifies a subspace that captures the most important components of the left singular space of X [HMT11].

It is computed in 2 steps:

1) Compute a random projection matrix $G \in \mathcal{R}^{n \times r}$ and multiply with X

$$Y = XG. \tag{2.2}$$

It compresses the dimensionality from n to r. Then, using QR factorization, we obtain an orthonormal basis Q that approximates the column space of Y, and thus also the column space of X.

2) Project X onto the subspace spanned by Q

$$Z = Q^T X \tag{2.3}$$

The resulting matrix Z has dimensions $r \times n$, which is much smaller than the original matrix X. Then, compute the SVD of Z

$$Z \approx \hat{U}_r \Sigma_r V_r^T \tag{2.4}$$

where $\hat{U}_r$, Σ_r, and V_r^T the first r components of the SVD of Z. Using (2.3)

and (2.4), we obtain the low-rank approximation of X as

$$X \approx Q\hat{U}_r\Sigma_r V_r^T. \tag{2.5}$$

This shows that by computing the SVD of the smaller matrix Z and a matrix multiplication $U_r = Q\hat{U}_r$, we can compute the truncated SVD of the large matrix X

$$X \approx U_r\Sigma_r V_r^T. \tag{2.6}$$

In Fig. 2.4, we compare truncated SVD with randomized SVD in terms of relative error and computation time, using one of our datasets, 51 Eri, which is detailed in the following chapters. Both methods yield similar low-rank approximations; however, randomized SVD demonstrates significantly faster computation times compared to truncated SVD. The performance advantage of randomized SVD becomes more pronounced as the size of the matrix increases.

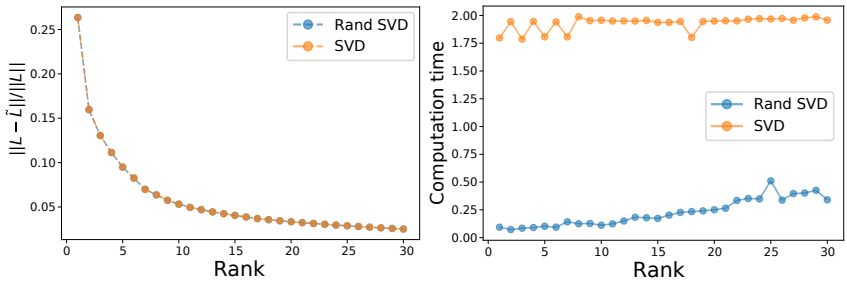

Fig. 2.4 Left: Relative error of randomized SVD and SVD. Right: Computation time of randomized SVD and SVD.

2.2.2 NMF

Nonnegative Matrix Factorization (NMF) is a useful technique that builds a low-rank approximation, similar to PCA, but has nonnegative components. NMF decomposes a matrix X into two nonnegative matrices $W \in \mathbb{R}^{m\times k}$ and $H \in \mathbb{R}^{k\times n}$, such that $X = WH$, instead of $X = U\Sigma V^T$ in PCA.

The columns of H represent nonzero components onto which the images are projected, while W contains the basis components. Unlike PCA, where the components are orthogonal, NMF components are not necessarily orthogonal but are constrained to be nonnegative.

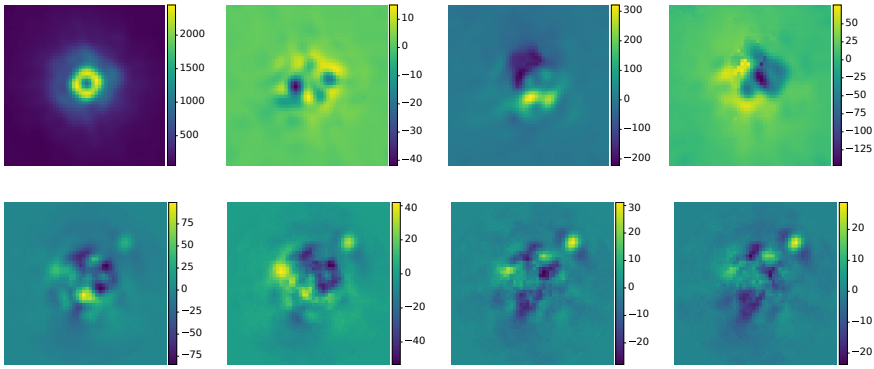

Fig. 2.5 Top: The effects of the rank in NMF. From left to right: 1st, 2nd, 3rd, and 4th rank. Bottom: Processing frames after we apply NMF using rank 1, 2, 3, and 4 from left to right. These images are generated using the β-Pictoris dataset, as detailed in Chapter 4.

The usefulness of nonnegative restrictions in such decompositions was clarified by Pueyo et al. [PCV+12]'s study of positive coefficients for the LOCI approach. Ren et al. [RPZ+18] promoted the usage of NMF in the context of PSF subtraction, drawing inspiration from this idea.

Unlike PCA, NMF does not have a unique solution and it depends on the initialization. Ren et. al adopted Zhu's approach [Zhu16] for solving NMF

$$
\begin{aligned}
W^{(i+1)} &= W^{(i)} \circ \frac{X {H^{(i)}}^T}{W^{(i)} H^{(i)} {H^{(i)}}^T} \\
H^{(i+1)} &= H^{(i)} \circ \frac{{W^{(i)}}^T X}{{W^{(i)}}^T W^{(i)} H^{(i)}}
\end{aligned}
\tag{2.7}
$$

where $\circ$ and $\frac{(\dots)}{(\dots)}$ denote element-wise multiplication and division of matrices, respectively. These update rules facilitate the convergence of W and H to a solution that minimizes the Euclidean distances between X and WH. In VIP-HCI [CGF+23, GG+17], the components are initialized by SVD, so the results of PCA and NMF may be similar.

Figure 2.5 illustrates how different ranks impact the PSF model using NMF. As with PCA, the initial rank primarily captures the core attributes of the images, revealing most of the stellar signals in the top left image. Subsequent ranks delve into finer details within the images displayed in

the top row of Fig. 2.5, with higher ranks typically aiding in the detection of planetary signals. While individual frames might exhibit variations, the processed frames shown in the bottom row resemblance to those produced using PCA. Since overall, the results of NMF are not significantly different from those of PCA, and since the interpretation of the matrix factors is not critical to our analysis, NMF was not further considered in this thesis.

2.2.3 LLSG

In the realm of exoplanetary astronomy, Principal Component Analysis (PCA) is leveraged in Angular Differential Imaging (ADI) sequences to enhance the visibility of rotating planets by separating them from static and quasi-static signals. This methodology is akin to the techniques used in video processing to segregate background and foreground pixels. Nevertheless, the literature on computer vision in recent years has seen a surge in the adoption of robust PCA (RPCA). This variant utilizes a low-rank matrix to model the stationary background while employing a sparse matrix to capture dynamic foreground elements.

The Local Low-rank plus Sparse plus Gaussian-noise decomposition (LLSG) framework, inspired by the principles of RPCA, advocates for the GoDec approach [ZT11] when applied to ADI image sequences, suggesting an innovative path to RPCA. Let $M \in \mathbb{R}^{T \times N^2}$ be an ADI sequence that consists of T unfolded images with size $N \times N$. It addresses the problem formulated as

$$M = L + S + G, \tag{2.8}$$

where L denotes a low-rank matrix, S represents a sparse matrix including planet signal, and G embodies Gaussian noise, with constraints $\text{rank}(L) \leq k$ and $\text{card}(S) \leq c$.

From this model arises a minimization problem:

$$\min \|M - L - S\|_F^2 \quad \text{subject to} \quad \text{rank}(L) \leq k \quad \text{and} \quad \text{card}(S) \leq c. \tag{2.9}$$

where $\|\cdot\|_F$ is the Frobenius norm. The resolution is pursued through the alternate optimization of two subproblems, iteratively refining estimates of L and S as follows:

- For the low-rank component, $L_t = \arg\min_L \|M - L - S_{t-1}\|_F^2$, which can be efficiently solved using truncated Singular Value Decomposition (SVD).

- For the sparse component, $S_t = \arg\min_S \|M - L_t - S\|_F^2$, which is approached through entry-wise hard thresholding.

These methods align with advanced computational strategies that aim to distill the essential features from high-dimensional data, offering robustness against noise and reducing the influence of irrelevant information. By integrating the GoDec algorithm into the ADI processing pipeline, astronomers can more effectively separate planetary signals from noise, a critical step in directly imaging exoplanets.

2.2.4 LOCI

Locally Optimized Combination of Images (LOCI) is an algorithm used to create a PSF model by combining images through a linear combination method. LOCI algorithm was initially proposed by Lafrenière et al. in 2007 as a technique to enhance direct imaging of exoplanets [LMD+07]. The algorithm divides each image into small subregions. Each subregion will be processed independently to optimize the starlight subtraction locally, which is crucial for reducing the influence of starlight while preserving any planetary signals. For each subregion, a set of reference images is created from the other images. These reference images should ideally represent the pattern without the planetary signal. The reference images are selected and combined in a way that optimizes the subtraction of starlight in terms of variance of the residual noise. This involves solving a least-squares problem to minimize the residuals after subtraction. A linear combination of the selected reference images is calculated for each subregion. The coefficients of this combination are chosen to minimize the residual starlight in the target subregion. The optimized linear combination is subtracted from the original image removing most of the starlight and leaving any potential planetary signals more visible. After processing all subregion, the resulting subtracted images are recombined to form the final processed image.

Several variations of LOCI exist, each with specific modifications to the original method:

Damped LOCI: This version explores different configurations of subtraction and optimization zones, concluding that spectral information of a synthetic companion could only be preserved when the cost function simultaneously minimized residuals in the optimization zone and maximized the planetary signal in the subtraction zone. This approach adjusts the least-squares fit by introducing a damping factor to control the linear coefficients more effectively [PCV+12].

Adapted LOCI: In this method, the images are filtered based on the degree of speckle correlation, improving the accuracy of PSF subtraction by reducing speckle noise [CFT+12].

Matched LOCI: Here, synthetic sources are injected into the data to generate the linear combination, which is then used to model the PSF more accurately [WCM+15].

The major limitation of these algorithms is that they assume the noise is Gaussian, so it can be challenging to find an effective PSF model when dealing with non-Gaussian noise.

3

Metrics for Exoplanet Detection Algorithms

There are various methods to detect exoplanets by direct imaging, which are described in Sec.2.1. The metrics used to evaluate the performance of exoplanet detection algorithms play a pivotal role. These metrics not only help in assessing the effectiveness of different methods but also guide the improvement and optimization of these techniques. A robust set of metrics allows researchers to benchmark algorithms, compare results, and ultimately enhance the reliability of exoplanet detection.

One of the fundamental tools in this field is the detection map, which provides a visual representation of potential exoplanetary signals in observational data. Among the various types of detection maps, the Signal-to-Noise Ratio (SNR or S/N) map which is a traditional and widely used method and a recent approach called the Standardized Trajectory Intensity Mean (STIM) map are introduced in this chapter.

Another critical aspect of evaluating exoplanet detection algorithms is the use of contrast curves. These curves measure the sensitivity of imaging instruments by plotting the detectable contrast between an exoplanet and its host star at varying distances.

Furthermore, the Receiver Operating Characteristic (ROC) curve is a powerful tool for assessing the performance of binary classification systems, including exoplanet detection algorithms. By plotting the true pos-

itive rate against the false positive rate, ROC curves provide insights into the trade-offs between sensitivity and specificity in detection methods. This helps in fine-tuning algorithms to achieve optimal performance. We not only introduce ROC curves, but also suggest a method to compute it for exoplanet detection.

As we delve deeper into these metrics, this chapter aims to provide an overview for common metrics used to evaluate exoplanet detection algorithms.

3.1 Detection Maps

3.1.1 SNR maps

The Signal-to-Noise Ratio (SNR) is a measure that compares the strength of a signal to the level of background noise. A higher SNR indicates a stronger, more distinguishable signal relative to the noise, while a lower SNR suggests that the signal is less discernible. In the context of a Z-test, the SNR is a measure that compares the observed mean (μ) to the variability in the data, represented by the standard deviation (σ)

$$\text{SNR} = \frac{\mu}{\sigma}. \tag{3.1}$$

Because the noise characteristics of the processed image vary with distance from the central star, it leads to higher false detection rates. Moreover, since the distribution of data is not Gaussian, Z-test does not standardize the noise.

To address these challenges, an alternative approach of SNR map (or S/N map) is outlined by [MMW+14] using Student's t-statistic to create a more reliable detection map. The t-test is particularly effective in scenarios with limited sample sizes because it adjusts for the uncertainty introduced by small datasets. Furthermore, it is better suited to account for annular regions that gather pixels where the noise has a similar distribution.

By applying the t-test to the fluxes within an annulus, the method evaluates whether the flux at a specific location significantly differs from the surrounding resolution elements. This provides a more statistically robust framework for detecting planetary signals.

The process involves comparing the flux of the speckles or planets against the flux in other resolution elements at the same radial distance

from the center of frame. Mathematically, this is expressed as:

$$\mathrm{S/N}(x) = \frac{\bar{x}_1 - \bar{x}_2}{s_2\sqrt{1 + \frac{1}{n_2}}} \tag{3.2}$$

In this formula, $\bar{x}_1$ represents the intensity of the test resolution element centered on x (blue in Fig. 3.1), while $\bar{x}_2$ and s_2 are the mean and standard deviation of the fluxes of the other apertures (red in Fig. 3.1) in the same annulus. The term n_2 denotes the number of these elements. The SNR map is developed by calculating the S/N ratio for each pixel in the processed frame.

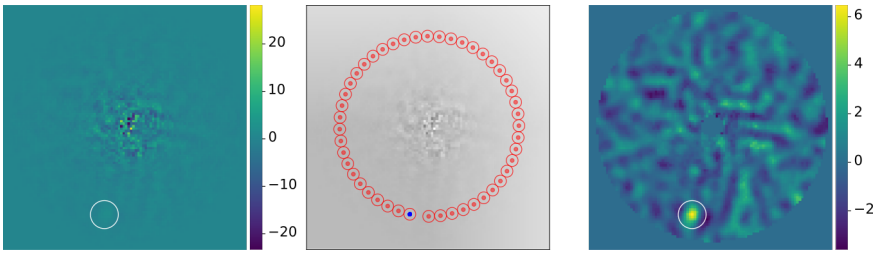

Fig. 3.1 Left: Processed frame. Middle: Resolution elements on processed frame. Right: SNR map. White circle highlights the planet location.

SNR has its limitations. For example, it does not account for scenarios where where the noise distribution deviates from a Gaussian assumption. In addition, SNR tends to have problems in situations where multiple planets or stellar disk are present within the same annular section, as this could lead to overestimate the noise in that region, thus reducing the effectiveness of SNR.

3.1.2 STIM maps

To overcome the challenges which SNR map has, a novel approach was proposed by [PCGG+19]. They introduced the standardized trajectory intensity mean (STIM) map, which operates directly on the residual data cube R. This method computes a detection value d_g for each trajectory g, defined as

$$d_g = \frac{\mu_g}{\sigma_g} \tag{3.3}$$

where μ_g is the empirical mean and σ_g the standard deviation of R along that trajectory. Since most of the correlated speckles are absorbed by the PSF model, the variance of the residual matrix along the trajectory is expected to be high, and this leads to decrease d_g. However, if the trajectory has a planet, the mean of the intensities along the trajectory is expected to be higher, resulting in a higher d_g.

One of the main advantages of the STIM map is its independence from small sample statistics. Since the mean and variance are computed over a consistent number of frames for each trajectory, the results are more reliable.

Additionally, the STIM map handles multiple companion scenarios in the same annulus. Since d_g is computed independently for each trajectory, the presence of multiple planets does not lead to an underestimation of the detection value. This method is also computationally efficient, since it involves simple calculations of the mean and variance on the derotated residual cube.

3.2 Contrast Curve

A key metric for evaluating the effectiveness of imaging instruments is the sensitivity limit, commonly represented by contrast curves. These curves plot the detectable level of contrast between a planet and its star against their separation distance.

For constructing these contrast curves, the process begins with an ADI dataset and its corresponding PSF. It is advisable to first exclude any significant real companions from the dataset, possibly using methods like Negative Fake Companion (NEGFC) [WAG+17]. The initial step involves quantifying noise σ_r at various angular distances in a processed ADI frame by calculating the standard deviation of fluxes within Full Width at Half Maximum (FWHM).

The next phase involves a practical procedure where we first inject fake planets into the data cube, which does not contain any planets. Following this, the algorithm is applied to generate a model of the PSF and to create a residual cube. During the planet injection process, which involves adding flux along a trajectory —rotating by parallactic angles— we assign a specific flux value, denoted as F_{in} for each fake planet. The flux ultimately recovered from the residual cube, denoted as F_r is then measured. Then, the algorithm's throughput T_r is defined as the ratio of the recovered flux F_r to the initially injected flux F_{in}. With these values, we can calculate the

contrast C_r as follows:

$$C_r = \frac{k\sigma_r}{T_r F_*} \tag{3.4}$$

where F_* is the aperture photometry of the non-coronagraphic PSF, used to scale the contrast. Here, k represents a chosen factor (often 5 for a five-sigma detection threshold). An example of such a contrast curve can be seen in Fig.3.2. In this figure, a lower curve indicates better performance, and any point above the curve represents a detectable signal.

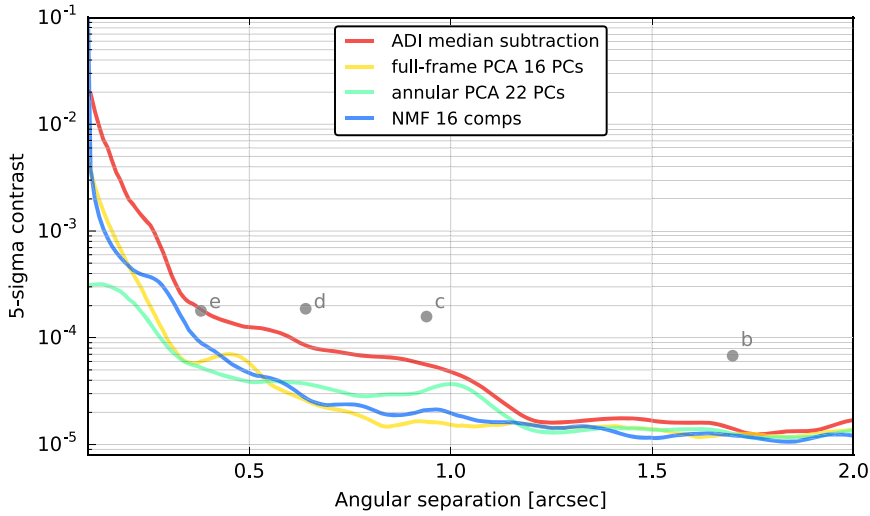

Fig. 3.2 Contrast curve of some classical algorithms. The visual is taken from [GG+17]. The points b, c, d, and e correspond to known planets, indicating their contrast and separation values on the curve.

3.3 ROC Curve Algorithms

The ROC (Receiver Operating Characteristic) curve is a commonly used method for evaluating the performance of a binary classifier model. It is a graphical plot that illustrates how well a model can distinguish between two classes (such as positive and negative) at various threshold settings. The threshold values can vary from minimum value to maximum value of

the detection maps such as SNR map, STIM map or likelihood ratio map which is described in the Chapter 4.

The ROC curve plots the True Positive Rate (TPR) against the False Positive Rate (FPR) at each threshold. The TPR, also known as sensitivity, measures the proportion of actual positives correctly identified as such (TPR = TP/(TP+FN)), while the FPR, or the probability of false alarm, measures the proportion of actual negatives that are incorrectly classified as positives (FPR = FP/(FP+TN)).

Here's a table representing these concepts:

	Actual Positive	**Actual Negative**
Predicted Positive	True Positive (TP)	False Positive (FP)
Predicted Negative	False Negative (FN)	True Negative (TN)

Table 3.1 Confusion Matrix for Binary Classification

In the context of the ROC curve, a perfect classifier would reach the top left corner of the graph, indicating a TPR of 1 and an FPR of 0. Conversely, a diagonal line from the bottom left to the top right corners of the graph represents a random classifier, with no discriminatory power between positive and negative classes. Typically, the closer the ROC curve is to the upper left corner, the more effective the classifier.

In the field of direct imaging, the ROC curve is widely used to compare different algorithms [GG$^+$16, CADVD23, DCA20]. It is one of the most popular methods for evaluating and contrasting the effectiveness of these algorithms in distinguishing between real and false detections.

In direct imaging algorithms, the calculation of False Positive Rates (FPR) is commonly replaced by the computation of mean false positives (FPs) [GG$^+$17]. This is done by forming segments of pixels exceeding a predetermined threshold. A segment is considered a True Positive (TP) if it predominantly covers pixels where a planet is located and most of the planetary aperture lies within that segment. Conversely, if the planetary aperture occupies only a small portion of a segment (even if the segment fully covers the aperture), it does not contribute to the TP. In such cases, the segment size is divided by the planetary aperture area to calculate False Positives (FPs). Additionally, any pixel that does not belong to the planetary aperture is similarly divided by the aperture area and counted as an FP.

There are two significant drawbacks to this method. First, the resulting ROC curves may never reach a TPR level of 1. In Fig. 3.3, we can see an ex-

ample to this situation from [DCA20]. This is because, as the threshold decreases, the segment identifying the planet might expand beyond its actual area, leading to an underestimation of the TP count. Consequently, this method does not produce the typical ROC curve shape. Second, and more importantly, such ROC curves do not provide a fair comparison. To illustrate this, consider a scenario where we simulate the injection of a planet with zero flux (effectively not altering the data cube at all). Although no flux is injected, the analysis proceeds as if a planet with zero flux were present. The resulting ROC curves, shown in Fig. 3.3, might misleadingly indicate that one algorithm outperforms another. However, since no actual planet was injected, fair ROC curves should ideally produce similar results across algorithms. This discrepancy highlights a critical limitation in the current methodology used for ROC curve analysis in direct imaging algorithms.

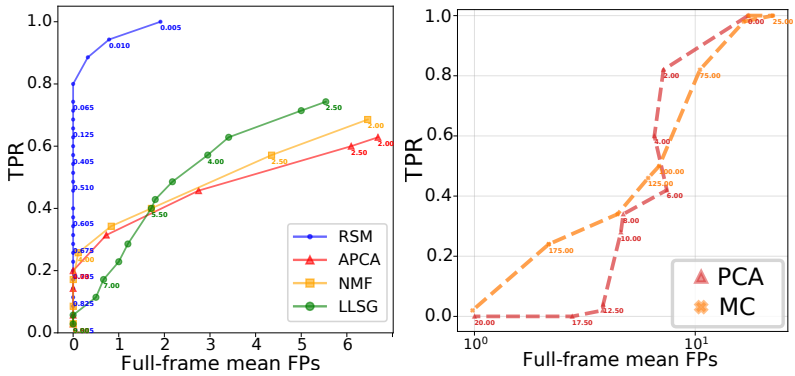

Fig. 3.3 Examples of ROC curves from the literature. Left: Adapted from [DCA20]. Right: Generated using the VIP-HCI package [CGF+23, GG+17], applying the same method as the left image, with zero-flux planets injected.

We propose a novel approach to address these issues in ROC curve analysis for direct imaging algorithms. This method focuses on the aperture where the planet is presumed to be. If a value exceeding the threshold is found within this aperture, it is classified as a True Positive (TP). Next, similar to the procedure used in SNR analysis, we define resolution elements. We then examine each resolution element individually. If any pixel within a resolution element exceeds the threshold, that element is counted as a False Positive (FP). Figure 3.4 illustrates a binary map obtained by ap-

plying a threshold to a detection map, such as an SNR map, STIM map, or likelihood ratio map. Based on the presence of white pixels within the resolution elements, we identify TPs, FPs, False Negatives (FNs), and True Negatives (TNs).

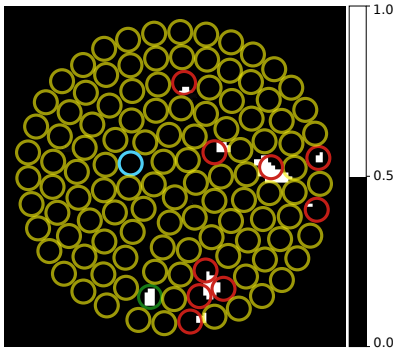

Fig. 3.4 Binary map used to determine TPs, FPs, FNs, and TNs after applying a threshold to a detection map. The green circle represents a TP, red circles represent FPs, the blue circle represents an FN, and yellow circles represent TNs.

This approach allows us to generate ROC curves that closely resemble the classical shape of ROC curves. Moreover, in the zero flux test scenario we created, as shown in Fig. 3.5, we obtain diagonal curves. This indicates that our method provides a fair and effective means of comparing different algorithms, as it yields the expected results even in a controlled test environment with no actual planetary signal.

In our analysis of ROC curves, our primary focus is on identifying the highest achievable TPR without incurring any FPs. Essentially, this means that the highest value of TPR achievable at zero FPR is a critical metric for comparing performances. This leads us to propose a new approach where we utilize the square root transformations of TPR and FPR for constructing ROC curves. By using $\sqrt{\text{TPR}}$ and $\sqrt{\text{FPR}}$, we can more effectively compare values on the axis as given in Fig. 3.6. Although it would be possible to use only $\sqrt{\text{FPR}}$, we opt for transforming both metrics to retain the characteristic diagonal result in zero flux scenarios. Thus, this square root approach offers a more nuanced and accurate metric for performance comparison in our studies.

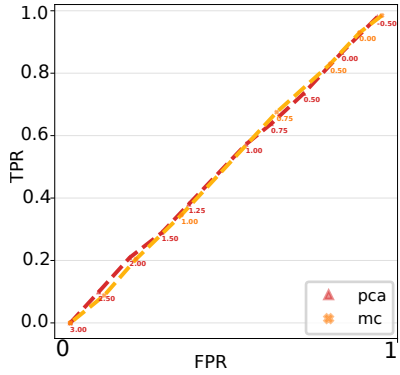

Fig. 3.5 ROC curve with zero flux.

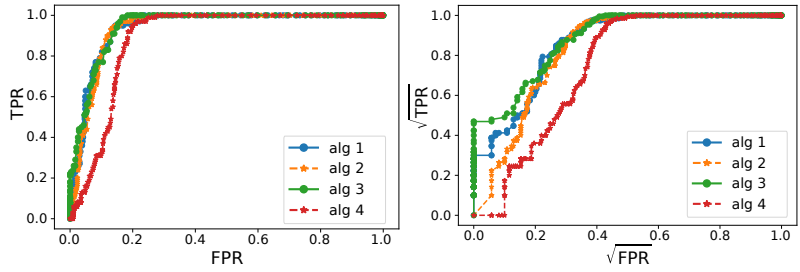

Fig. 3.6 Left: ROC curve with TPR and FPR on the axis. Right: ROC curve with $\sqrt{\text{TPR}}$ and $\sqrt{\text{FPR}}$ on the axis.

4

Likelihood Ratio Map for Direct Exoplanet Detection

DIRECT imaging of exoplanets is a challenging task due to the small angular distance and high contrast relative to their host star, and the presence of quasi-static noise. In this chapter, we propose a new statistical method for direct imaging of exoplanets based on a likelihood ratio detection map, which assumes that the noise after the background subtraction step obeys a Laplacian distribution.

The chapter is structured as follows. Our proposed algorithm for direct exoplanet detection is described in Section 4.1. We start with giving a general information about how a detection map is obtained using the classical methods and our proposed method. Then we continue with the mathematical expressions and reasons to propose our suggestion. The experimental results are presented in Section 4.2 to validate the performance of the proposed detection map. We compare the method with two detection approaches based on signal-to-noise ratio (SNR) map and STIM map after performing the background subtraction by the widely used Annular Principal Component Analysis (AnnPCA). The experimental results on the Beta Pictoris data set show the method outperforms SNR and STIM maps in terms of achieving the highest true positive rate (TPR) at zero false positive rate (FPR). Final remarks are given in Section 4.3.

This chapter builds upon the foundational work presented in [DVC+22] and [DVL+23], and extends these studies by incorporating additional data and methods, as detailed in next chapters, thereby enriching our understanding and application of the discussed concepts.

4.1 Detection Based on Likelihood Ratio Map

4.1.1 General information about obtaining a detection map

Angular differential imaging (ADI) partially overcomes the problem of quasi-static speckles by taking a sequence of images of a target star over a single night of observation without compensating for the Earth's rotation [MLD+06]. As a result, potential planet companions end up following circular trajectories across the image sequence, while the star and its quasi-static speckle field, being almost fixed with respect to the pupil of the telescope, remain roughly in the same location, exposing the detection problem to the use of dynamic-foreground/static-background separation methods [BSJ+17].

The complete pipeline of ADI detection consists of three steps. The first is the background subtraction, which estimates the *model PSF* containing the speckles and the *residual cube* that is meant to contain only the planets and some residual noise. The standard and most widely used methods are based on low-rank matrix models, such as principal component analysis (PCA) [AQ12, SPL12], its annular version (AnnPCA) [GG+17], and the low-rank plus sparse method (LLSG) [GG+16]. The rationale behind the low-rank models is that the bright, quasi-static speckles are captured by the first few principal components, while the higher-rank moving planets are excluded from the model.

The second step is the flux estimation of the residual cube, in which we, under some probabilistic model, compute the *flux*, i.e. the estimated light intensity of the planet, for each postulated planet trajectory. Classical approaches are median-based [MLD+06] and likelihood based [MCS+09, RMW+17].

The final third step is to compute the detection map for the image, which assigns a value to every pixel indicating how much we believe a planet is located in the pixel. From the detection map, pixels can then be predicted positive or negative according to whether the value is above or below a detection threshold. A common detection map based on the flux map is the signal-to-noise ratio (denoted by SNR or S/N), which uses a

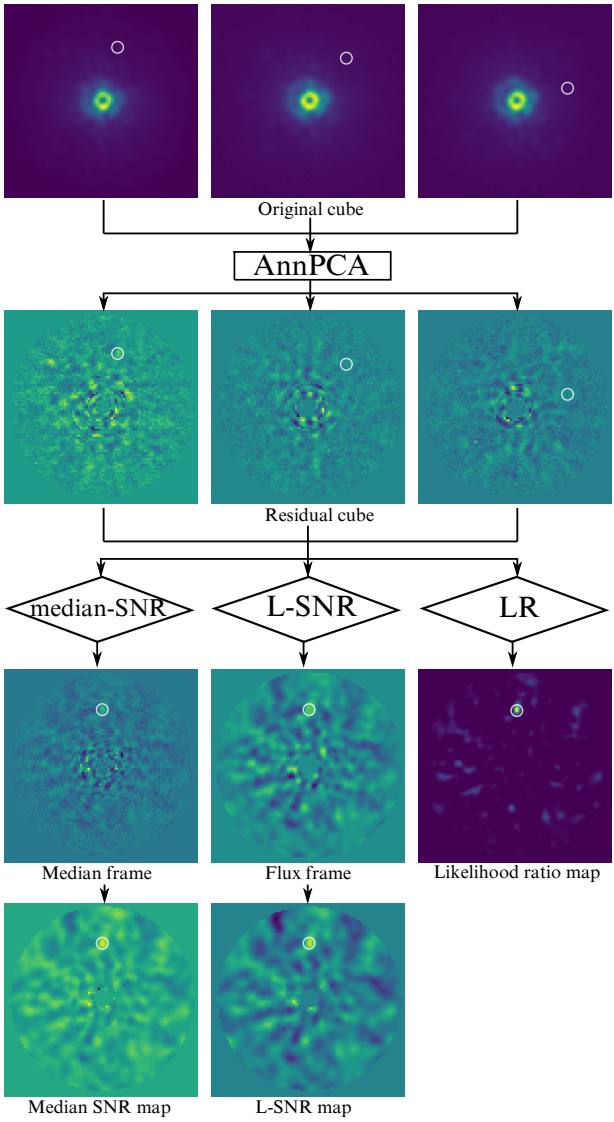

Fig. 4.1 Pipeline of the algorithms shown on the β-pictoris datacube with injected planet located at the white circle. For AnnPCA we choose rank 20.

two-sample t-test to compare the flux of the considered resolution element to the noise estimated over the remaining resolution elements over an an-

nulus [MMW+14]. Alternatively, STIM detection map [PCGG+19] utilizes from the standard deviation of each pixel across the frames in residual cube to be effective even when the faint planet is very close to the star. Other methods, such as the regime-switching model (RSM) [DCA20], employ more complex models that take into account the possible planet trajectories or the behaviour of quasi-static speckles.

Considering AnnPCA for the first step, median or maximum likelihood for the second step, and SNR for the third step yields two baseline methods that we term *median-SNR* and *L-SNR*, respectively. Our L-SNR is similar to the method proposed in [MCS+09, RMW+17] in that they are based on likelihood and SNR. However, it differs from the work of [MCS+09, RMW+17] by assuming a Laplacian instead of Gaussian distributed noise and by performing the background subtraction by AnnPCA instead of the difference of frames method [MCS+09] and the Karhunen-Loève image processing method [SPL12].

4.1.2 Maximum likelihood approach

Let $R \in \mathbb{R}^{T\times N\times N}$ be a residual cube that is the result of background subtraction, e.g. using AnnPCA from a sequence of T images of size $N \times N$, and should contain only the planet, starting at an unknown position $g \in [N]\times[N]$ and moving between frames with a known angular velocity with some small residual noise term E.

We develop a statistical model that postulates that a planet could be located at any pixel g of the first frame, which also defines its trajectory, and computes the optimal flux a_g such that the log-likelihood of observing the noise term E is maximised.

The method can be summarised as follows. Assuming there is a planet along the trajectory, the residual cube R is modelled as:

$$R = a_g P_g + E, \tag{4.1}$$

where a_g is the flux, P_g is a cube constructed by rotating the planet signature along the trajectory, based on the location g, illustrated in Fig. 4.2, and E is the residual noise. To construct P_g we place a copy of the normalized reference point spread function (PSF) defined by the optical instrument, depicted in Fig. 4.2, at coordinates g_t on each frame t of a cube with zero entries.

Even when multiple planets are present in the data cube, we analyze each potential pixel individually, with our model comprising a single a_g

(and P_g). In the algorithm's final results, all possibilities are combined, enabling a comprehensive evaluation of all cases together. The similar approaches are applied in the following chapters as well.

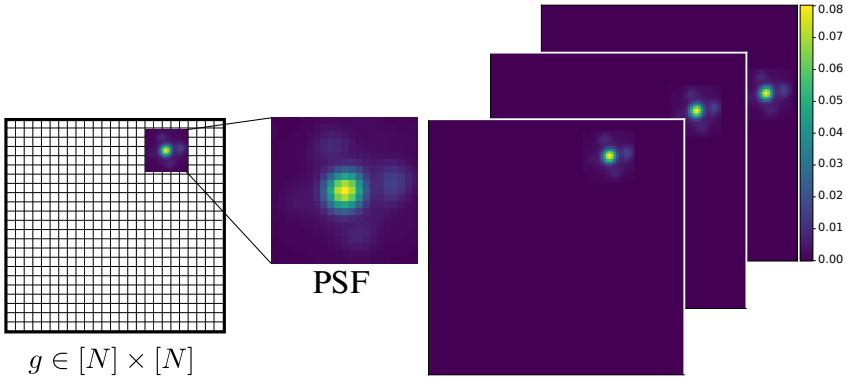

Fig. 4.2 The cube of the planet signature constructed by rotating position g of the PSF function along the trajectory.

Mugnier et al. [MCS+09] propose to model the noise term E as white and Gaussian. It is then possible to estimate the value of a_g by maximizing the following log-likelihood

$$\log \mathcal{L}_g^{\text{Gauss}}(a_g|R) \propto -\frac{1}{2} \sum_{(t,r)\in\Omega_g} \frac{(R(t,r) - a_g P_g(t,r))^2}{\sigma^2_{R(r)}}, \tag{4.2}$$

where σ_R^2 is the empirical variance of the residual frames computed along the time dimension and Ω_g is the set of indices (t,r) of pixels whose distance from the trajectory is smaller than the disk size ρ times half the diffraction limit

$$\Omega_g = \left\{ (t,r) \in [T]\times[N]^2 \,\middle|\, \|r - g_t\|_2 < \frac{1}{2}\frac{\lambda}{D}\rho \right\}. \tag{4.3}$$

In Fig. 4.3, we see the residual values $R(t,r)$ where $(t,r) \in \Omega_g$ for a trajectory with a planet and a trajectory without a planet. When we look at the residual values of the planetary and non-planetary trajectories in the figure, we cannot see a clear difference, but we can see that the planetary ones are slightly higher. To prove this scientifically, we can find the intensity a_g that maximizes the loglikelihood. Maximizing (4.2) translates to a

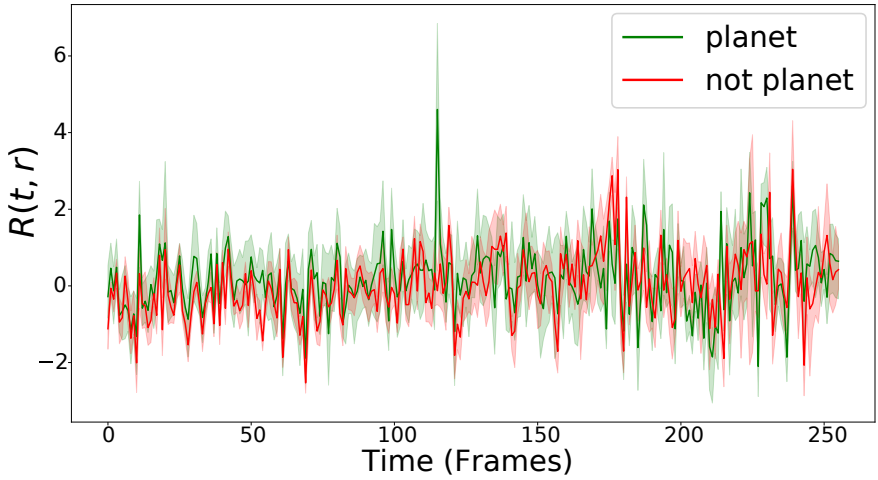

Fig. 4.3 Residual values in the sets Ω_g which has planet and does not have planet. The dark lines are mean values in the set Ω_g and light interval is limited with minimum and maximum intensities in the set Ω_g.

simple linear least-squares problem [MCS+09]

$$\frac{\partial \log \mathcal{L}_g^{\text{Gauss}}(a_g|R)}{\partial a_g} = \sum_{(t,r)\in\Omega_g} P_g(t,r)\frac{R(t,r) - a_g P_g(t,r)}{\sigma^2_{R(r)}} \tag{4.4}$$

which can be solved by setting this derivative to zero

$$a_g^{\text{Gauss}} = \frac{\sum_{(t,r)\in\Omega_g} \frac{R(t,r)P_g(t,r)}{\sigma^2_{R(r)}}}{\sum_{(t,r)\in\Omega_g} \frac{(P_g(t,r))^2}{\sigma^2_{R(r)}}}. \tag{4.5}$$

However, it has been observed that the tail of the distribution of the noise term E decays exponentially and is more accurately described by the Laplacian distribution [PCGG+19]. We replace the Gaussian assumption in (4.2) with a Laplacian

$$\log \mathcal{L}_g(a_g|R) \propto -\sum_{(t,r)\in\Omega_g} \frac{|R(t,r) - a_g P_g(t,r)|}{\sigma_{R(r)}}, \tag{4.6}$$

which leads to the following optimisation problem that estimates the planet flux

$$a_g^* = \text{argmax}_{a_g} \log \mathcal{L}_g(a_g|R) \tag{4.7}$$

$$= \text{argmin}_{a_g} \sum_{(t,r)\in\Omega_g} \frac{|R(t,r) - a_g P_g(t,r)|}{\sigma_{R(r)}}. \tag{4.8}$$

Solving (4.7) is an instance of the weighted least Absolute deviation (LAD) problem, which, unlike least squares, does not have a closed form solution. In general, ℓ_1 minimization can be solved by a number of efficient iterative methods, however, in our specific case, it is possible to compute the solution even more efficiently. Since the objective function is a convex piecewise linear function $\mathbb{R} \to \mathbb{R}$ with intervals between $R(t,r)/P_g(t,r)$, $(t,r) \in \Omega_g$, which is illustrated in Fig. 4.4 its minimum is attained at one of the (TN^2) points that can be easily searched exhaustively.

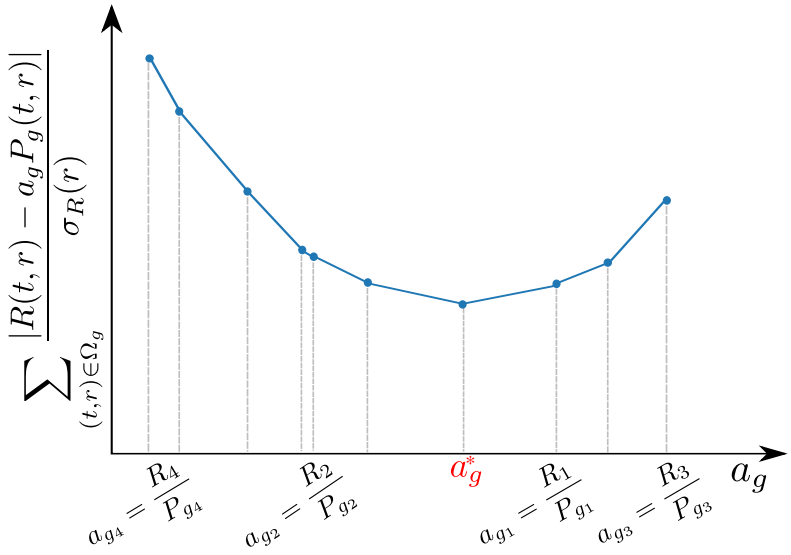

Fig. 4.4 A convex piecewise linear function

We propose a new detection map, termed likelihood ratio map (LRM), which compares how well two hypotheses, the null (H_0) and research (H_1) hypotheses, fit the data [CB21]. In our case, for each trajectory, H_0 corre-

sponds to the absence of a companion, and H_1 to the presence of one. We define the likelihood ratio $\Lambda_g(R)$ as the ratio of the maximum likelihood to the likelihood of the null hypothesis and compute its logarithm

$$\begin{aligned}\log \Lambda_g(R) &= \log\left(\frac{\mathcal{L}_g(a_g^*|R)}{\mathcal{L}_g(0|R)}\right) \\ &= -\sum_{(t,r)\in\Omega_g} \frac{|R(t,r) - a_g^* P_g(t,r)| - |R(t,r)|}{\sigma_{R(r)}}.\end{aligned} \tag{4.9}$$

Similarly, we can compute the likelihood ratio using the Gaussian assumption

$$\begin{aligned}\log \Lambda_g^{\text{Gauss}}(R) &= \log\left(\frac{\mathcal{L}_g^{\text{Gauss}}(a_g^{\text{Gauss}}|R)}{\mathcal{L}_g^{\text{Gauss}}(0|R)}\right) \\ &= -\frac{1}{2}\sum_{(t,r)\in\Omega_g} \frac{\left(R(t,r) - a_g^{\text{Gauss}} P_g(t,r)\right)^2 - \left(R(t,r)\right)^2}{\sigma_{R(r)}^2}.\end{aligned} \tag{4.10}$$

The LRM is then the map which assigns the value $\log \Lambda_g(R)$ to the starting pixel g of every trajectory. As shown in the shaded green area in Fig. 4.5, we compute the likelihood ratio values of all pixels except the edges and star pixels, and the combination of these values forms our LRM.

In the context of our approach, we investigate the performance of the algorithm using Gaussian and Laplacian assumptions, as visually compared in Fig. 4.6. Under Gaussian assumption, the algorithm yields high values in pixels external to the planetary region. This outcome can be attributed to the inherent sensitivity of the ℓ_2 norm towards outliers, which is a characteristic response to Gaussian distributions. Conversely, when the algorithm is exposed to Laplacian assumption, and thus under the influence of the ℓ_1 norm, there is a marked decrease in the values for non-planet pixels as compared to those within the planet.

Furthermore, in our study, we analyze the impact of the disk size ρ, which encompasses the planetary signals, on the performance of the algorithm. The radius of the disk is proportional to the ratio value λ/D described in detail in Chapter 2. Experiments with varying diameters reveal that when the disk size is set to a small value, the mask encompasses very few pixels, limiting its effectiveness. Conversely, if ρ is large, the mask

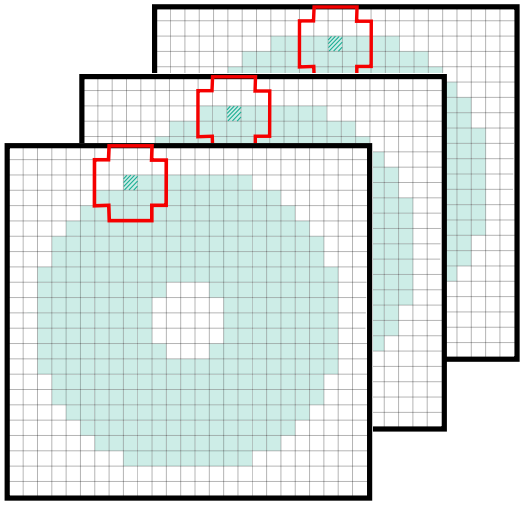

Fig. 4.5 The hatched pixel represents the coordinates g_t, the area inside of the red lines represents the set of indices Ω_g. Green shaded area shows the possible pixels to apply likelihood ratio given in (4.9).

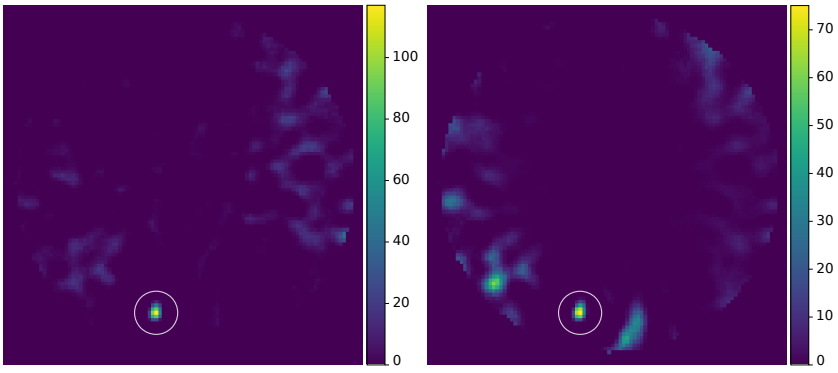

(a) ℓ_1 norm **(b)** ℓ_2 norm

Fig. 4.6 Likelihood ratio maps obtained by different norms. White circle represents the planet location. The results are obtained using 51 Eri dataset with the real planet.

includes an excessive number of non-planetary pixels. To optimize this parameter, we experiment different radius sizes, as illustrated in Fig. 4.7,

that the most effective results are achieved when the disk size is set to 1. This finding highlights the critical balance required in mask sizing to ensure both the inclusion of relevant planetary pixels and the exclusion of the background data, as evidenced by the very different SNR values obtained.

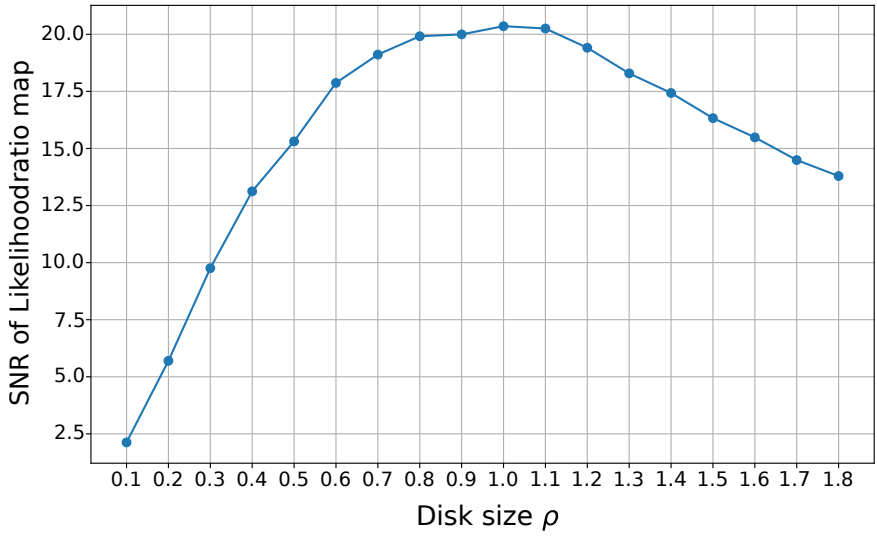

Fig. 4.7 Disk size vs SNR of likelihood ratio map

To make this comparison, we use the SNR of the LRM in the planetery pixels instead of the likelihood ratio of the planetary pixels. The reason for using this for comparison is that not only can the likelihood ratio be high in planetary pixels, but also in non-planetary pixels if we change ρ. To find a good ρ, we use the SNR of the LRM, which will allow us to show that the performance increases for planetary pixels but does not change or decreases for non-planetary pixels, or does not increase as much as for planetary pixels. For a more detailed explanation of how we obtain the SNR of the LRM, see the Chapter 7-Section 7.1.3.

4.2 Numerical Experiments

We compare the proposed LRM with the baseline median-SNR and L-SNR detection maps in terms of visual quality and Receiver Operating Characteristic (ROC) curves. We provide Python code based on the VIP

HCI package [GG+17] of our experiments publicly[1].

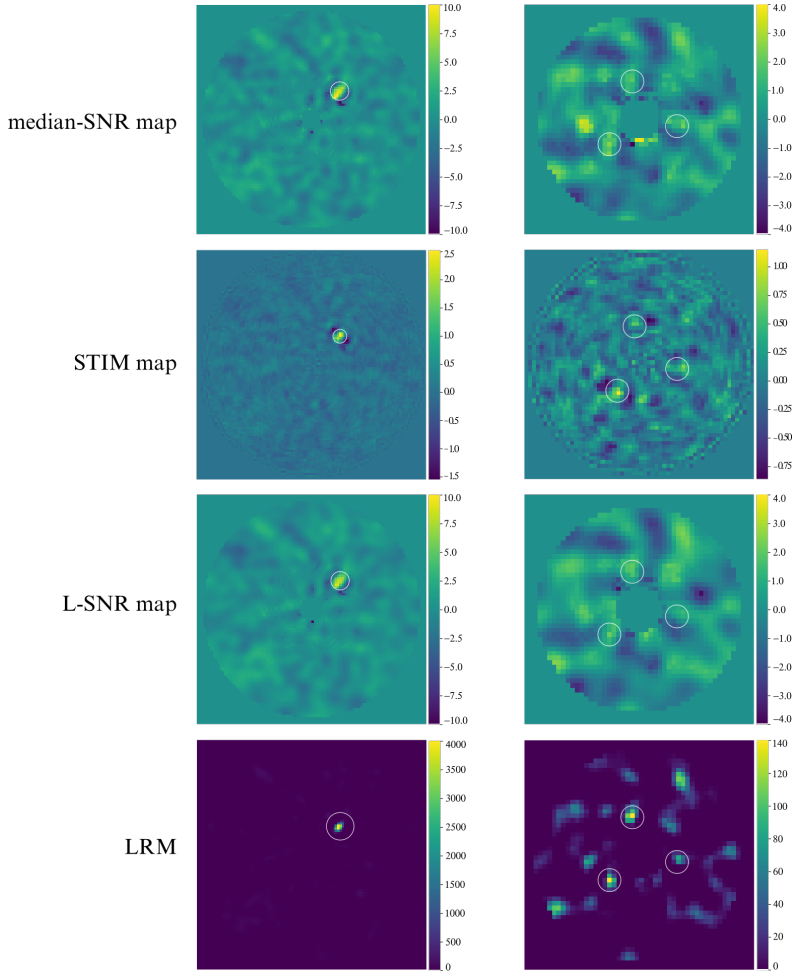

Fig. 4.8 Firstly, they are tested using the real planet (left column), then using three fake planets injected to the separation 1–2 λ/D with $1\sigma_{ann}$ (right column). They are obtained with AnnPCA using 20 PCs on β Pictoris data set. Note that the simultaneous injection of three fake planets is only for illustration; the ROC curves of Fig. 4.9 are obtained with single planet injections. We relied on the VIP package [GG+17] for step 1 and step 3 of the pipeline and for displaying the SNR and STIM maps.

[1]https://github.com/hazandaglayan/likelihoodratiomap

We test the efficacy of our method on the commonly used ADI cube VLT/NACO β-Pictoris in the infra-red L′ band (3.8μm), which has 612 frames covering 83° of parallactic angles and $\lambda/D \approx 4.6$ pixel [AMM+13]. In order to reduce the computation time, we consider only every third frame of the data set and crop it to 60 by 60 pixels, resulting into a cube of size $204 \times 60 \times 60$. In AnnPCA, we choose 20 as the number of components, which has been observed to perform well [DCA20]. To generate synthetic groundtruth examples, we inject the planet-free data cube (i.e., the data cube where the known planets have been removed) with fake planets using the VIP HCI package [GG+17].

Fig. 4.8 shows the visual quality of detection maps for two cases, the top row is the real data set with an actual planet, and the bottom row is the synthetically produced data set which we acquire by injecting three planet companions into an empty β-Pic cube.

We propose a novel deterministic approach to compute the ROC curves from synthetic planet injections. Our approach has the advantage of producing a diagonal ROC curve as expected when the flux of the fake planets is reduced to zero. We briefly recall it here, and refer to Section 3.3 for details.

Our procedure involves data cubes in which we choose one "actual positive" location where we inject a fake planet with prescribed flux, and several "actual negative" locations where we inject nothing. The detection algorithm receives as input all those locations, and it applies the exact same detection procedure to all the locations. Specifically, when at least one actual positive resolution element (of size λ/D) contains an above-threshold pixel of the detection map (median-SNR, STIM, L-SNR or LRM), we count one true positive (TP); when at least one actual negative resolution element contains an above-threshold pixel, we count one false positive (FP). This finally yields the true positive rate (TPR) and the false positive rate (FPR) that constitute the vertical and horizontal axes of the ROC curve.

Fig. 4.9 depicts the performance of the methods in terms of deterministic ROC curves computed for values of flux relative to the standard deviation of the annulus $c \cdot \sigma_{ann}$, where $c = 0, 0.5, 1, 1.5, 2$. We observe the performance of the LRM to be above the median-SNR, the STIM, and the L-SNR curves in the low FPR region. Specifically, as the threshold is lowered, more actual positives are detected with LRM than SNR and STIM before the first actual negatives are detected.

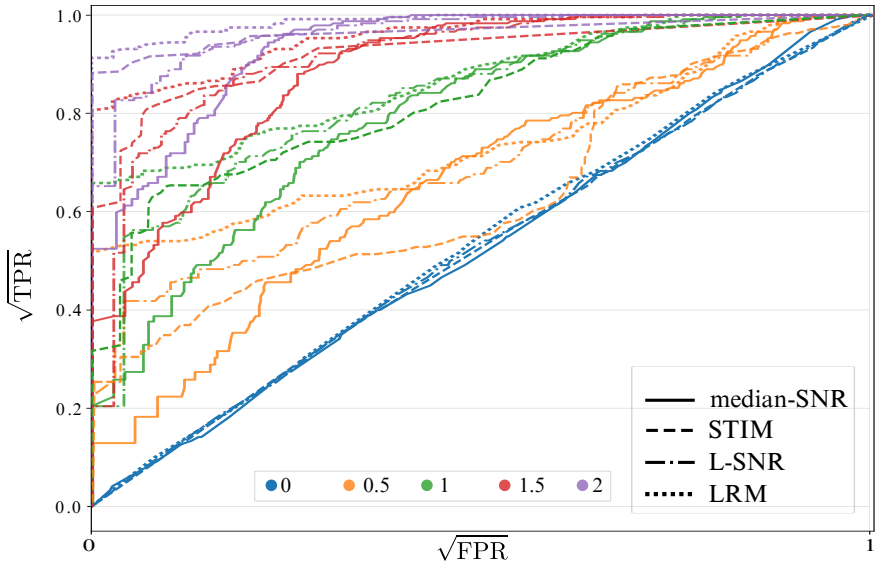

Fig. 4.9 ROC curves for β Pictoris data set. The planets are injected in 1–2 λ/D separation with the fluxes indicated in the legend multiplied by the standard deviation of the annulus. In order to better observe the low FPR regime, we scale the axes using the square root.

Table 4.1 Highest TPR at zero FPR.

	Detection Maps			
c	***median-SNR***	***STIM***	***L-SNR***	***LR***
0	0	0	0	0
0.5	0.017	0.083	0.050	**0.267**
1	0.042	0.121	0.042	**0.433**
1.5	0.142	0.394	0.042	**0.633**
2	0.275	0.780	0.425	**0.833**

Table 4.1 shows the highest TPR at a zero FPR for injected planets with different flux values, which reflects the fraction of detected actual planets before the first detection of an actual negative location. False positives are troublesome in high-contrast image processing because they require follow-up observations to be disproved, which are very costly in terms of telescope time. Therefore, compared to the more standard area under

curve (AUC) metric for ROC curves, we find the highest TPR at zero FPR to be more aligned with the goal of detecting exoplanets. The values in the table show that LRM outperforms the other methods.

4.3 Conclusions

In this chapter, we have presented a novel detection map for direct imaging of exoplanets that assumes a Laplacian distribution for noise, especially after background subtraction. This approach departs from traditional Gaussian noise assumptions and opens new perspectives on the detection and analysis of exoplanet data.

The essence of our method lies in the use of the LRM that excels in discriminating the presence of exoplanets, especially under low false positive rate (FPR) conditions. This advance is crucial in the field of astrophysics, where the accurate detection of exoplanets is a major challenge.

Our experimental evaluations on the β-Pictoris dataset show that our method is outstanding in detecting exoplanet signals, especially when compared to SNR-based and STIM map approaches. The innovative ROC curve methodology we use further emphasises the effectiveness of our approach. Unlike traditional methods, our ROC curves display a critical property: they form a diagonal line under zero flux injection scenarios, indicating a reliable fair balance.

Furthermore, our analysis of masking the pixels along the trajectory instead of applying a maximum likelihood approach to the entire frame, and even our decision of the disc size parameter ρ, highlights the delicate balance required in mask sizing for effective exoplanet detection. This finding is crucial for improving the sensitivity of exoplanet detection methodologies.

Our research, using a LRM based on the Laplacian noise model, not only supports previous studies on this noise, but also demonstrates performance that surpasses them, providing a promising path for future research and applications.

In future studies, the application of our method to a wider range of datasets and integration with other astrophysical data processing techniques can be explored to further validate and improve its applicability and effectiveness.

5

Low-rank plus sparse trajectory decomposition for direct exoplanet imaging

Many of the most successful methods for removal of the quasi-static speckled noise are based on low-rank matrix models, such as principal component analysis (PCA) [AQ12, SPL12], its annular version (AnnPCA) [GG+17], the low-rank plus sparse method (LLSG) [GG+16], and the morphological component analysis (MAYO) [PCJ21]. The rationale behind the low-rank models is that the bright, quasi-static speckles that appear as static elements with slowly changing light intensity, are captured by the first few principal components, while the higher-rank moving planets are excluded from the model.

Following the subtraction of the quasi-static speckled background, the typical next step in the ADI pipeline is the estimation of the light intensity of possible planet companions, referred to as the *flux estimation*. Classical approaches take into account the point spread function (PSF) of the optical device with the known angular velocity and are median-based [MLD+06] or likelihood based [MCS+09, DVC+22]. More recently, MAYO proposed to employ an image sparsifying transform to model general rotating objects which allows it to detect also circumstellar discs [PCJ21].

A key image processing challenge in the ADI pipeline is to be able to distinguish the moving planets from the quasi-static speckles as they are of similar intensities and resolution. A limitation of the ADI pipeline is that it works in steps: it assumes a low-rank model for the quasi-static speckles, subtracts the result, and then it proceeds to estimate the flux of possible planets.

Methods such as LLSG [GG+16] and the MAYO pipeline [PCJ21] have been developed with these issues in mind. In LLSG, the authors propose to perform low-rank plus sparse decomposition, which prevents accidentally subtracting the planet from the datacube by accounting for sparse outliers that can capture the planet. However, the issue here is that the sparse component could also capture a quasi-static speckle since it does not assume that the planet needs to move smoothly in a circular trajectory. The most similar to our approach is the MAYO method [PCJ21], which assumes both a low-rank model for the quasi-speckled background and a circular trajectory of the planet. However, whereas MAYO performs these tasks as part of a multi-step pipeline, in this work, we employ an alternating hard thresholding algorithm [TV23] that solves them at once as part of a single optimization problem.

In this chapter, we propose to employ a low-rank plus sparse trajectory (LRPT) method that estimates the low-rank component while trying to find a small number of viable planet trajectories based on a low-rank plus sparse decomposition [TV23]. We develop a linear transform, referred to as the *trajectorlet transform*, which takes into account the knowledge of possible planet trajectories and can be computed efficiently using convolution operations and rotation.

The chapter is structured as follows. The following Section 5.1 introduces the low-rank plus sparse model and the trajectorlet transform. In Section 5.2, we present experimental results and comparisons of the method with other standard methods. Finally, Section 5.3 concludes the chapter and suggests further research directions using this model.

This chapter builds upon the foundational work presented in [VDJA23].

5.1 Direct Imaging with Low-rank Plus Sparse Model

Let $M \in \mathbb{R}^{T \times N^2}$ be the matrix of observations whose rows are the $N \times N$ image frames of the video sequence with T frames that contains

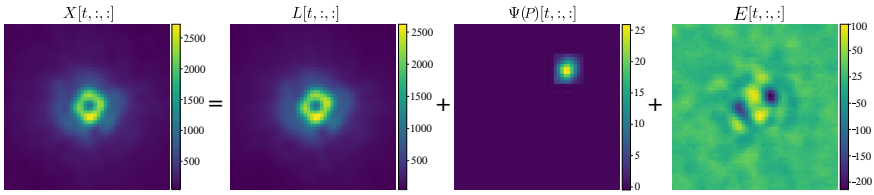

Fig. 5.1 The result of the low-rank plus sparse trajectory (LRPT) method applied to the β-Pictoris dataset with the real planet. We show a single frame t of the low-rank component, which captures the quasi-static speckled noise, and the $\Psi(P)$ component which detects the rotating planet.

an unknown small number of planets as well as the quasi-static speckle noise. Our model for M, depicted in Fig. 5.1, is expressed as

$$\begin{aligned} M &= L + \Psi(P) + E \\ \text{s.t.} \qquad & \operatorname{rank}(L) \leq r, \quad \|P\|_0 \leq s, \quad P_{i,j} \geq 0, \end{aligned} \tag{5.1}$$

where $L \in \mathbb{R}^{T \times N^2}$ is the rank-r component which models the quasi-static speckles, $\Psi(P)$ is the component containing the rotating planets with $P \in \mathbb{R}^{N \times N}$ being a non-negative s-sparse matrix constrained with the ℓ_0-norm counting the number of non-zero entries and $\Psi : \mathbb{R}^{N \times N} \to \mathbb{R}^{T \times N^2}$ is a linear *trajectorlet* transform specified below. The final component E is the residual noise which we assume to be normally distributed.

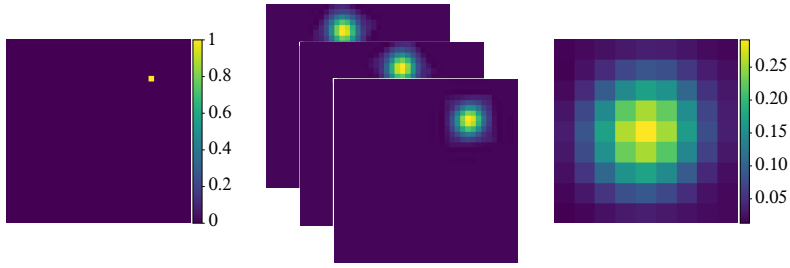

Fig. 5.2 Illustration of the trajectorlet transform: (left) depicts a matrix $P \in \mathbb{R}^{35 \times 35}$ with $P_{25,25} = 1$ and 0 otherwise, (center) shows $\Psi(P)$, i.e., a single element of the dictionary denoted $\Psi_{25,25}$, (right) shows the PSF of the optical device.

5.1.1 Trajectorlet transform

The trajectorlet transform $\Psi : \mathbb{R}^{N\times N} \to \mathbb{R}^{T\times N^2}$ is a linear mapping that expresses possible trajectories a planet can take in the $N \times N$ frame according to the known angular velocity and the point spread function (PSF) of the optical device.

The linear transform $\Psi(\cdot)$ is represented by a dictionary of N^2 elements in $\mathbb{R}^{T\times N^2}$ which we denote $\Psi_{i,j}$ for $\forall i, j \in [N] \times [N]$. A single item $\Psi_{i,j}$ of the dictionary is generated by placing the PSF of the optical device at the position i, j on the first frame and rotating it according to the known angular velocity of the telescope on the remaining frames; see Fig. 5.2 for an illustration. We normalize the elements of $\Psi_{i,j}$ to produce a unit norm frame. The linear transform Ψ and its adjoint Ψ^* are given by

$$\Psi(P) = \sum_{i,j=1}^{N,N} \Psi_{i,j}\, P_{i,j} \quad \Psi^*(M) = \sum_{i,j=1}^{N,N} \langle \Psi_{i,j},\, M \rangle\, \mathbb{1}_{i,j}, \tag{5.2}$$

where $\langle \cdot, \cdot \rangle$ is the trace inner product in $\mathbb{R}^{T\times N^2}$ and $\mathbb{1}_{i,j} \in \mathbb{R}^{N\times N}$ denotes the canonical basis matrix with a single one at index i, j and zero otherwise.

The underlying structure of the transform allows us to compute the operations in (5.2) in an computationally efficient way through the use of a rotation and a convolution operation. To compute the forward operation $\Psi(P)$, we apply the convolution of $P \in \mathbb{R}^{N\times N}$ with the PSF of the planet as the kernel and we place the rotated copy of the result along the T frames. The adjoint operation $\Psi^*(M)$ can be computed by reversing the operations: first, we derotate the rows of $M \in \mathbb{R}^{T\times N^2}$ summing them into an N^2 vector, we unfold it into a $N \times N$ matrix and perform the adjoint of the convolution with the PSF of the planet as the kernel. Detailed implementation using the `scipy.sparse` API [VGO+20] can be found in the code provided with our paper.[1] Note that the rotation and convolution procedure has also been used in MAYO [PCJ21].

[1] Link to the code `https://github.com/hazandaglayan/trajectorlets`.

5.1.2 Low-rank plus structured sparse model

We compute the best estimator for the model in (5.1) by solving the following non-convex optimisation problem

$$\begin{aligned} &\min_{L\in\mathbb{R}^{T\times N^2},\, P\in\mathbb{R}^{N\times N}} \frac{1}{2}\left\|M-(L+\Psi(P))\right\|_F^2, \\ &\text{s.t.} \qquad \operatorname{rank}(L)\leq r, \quad \|P\|_0\leq s, \quad P_{i,j}\geq 0, \end{aligned} \tag{5.3}$$

by applying a modification of the Normalized Alternating Hard Threshold (NAHT)algorithm [TV23], which is shown in Algorithm 1. The method solves (5.3) by taking alternating projected gradient steps. In line 5, NAHT computes an estimate of the quasi-static background by projecting $M-\Psi(P^j)$ on the set of rank-r matrices with $\mathrm{H}_r^{\mathrm{SVD}}(\cdot)$ that performs the randomized singular value decomposition and keeps only r largest singular values of the input matrix [HMT11]. Next, in line 6, it computes a gradient direction of the objective (5.3) with respect to P based on the current residual map. The algorithm estimates the adaptive step-size in line 7, where $R^j_{\Omega^{j-1}}$ denotes the matrix R^j with entries kept only on the support set Ω^{j-1} and zeroes otherwise. The step-size corresponds to the optimal gradient step-size when the current iterate P^j has the same support as the optimal sparse solution [BD10, TV23]. Finally, in line 8, it estimates the current foreground element $\mathrm{H}_s^{\ell_0}\left(P^j-\alpha_j R^j\right)$ by projecting the gradient updated P^j on the sparse positive constraint using $\mathrm{H}_s^{\ell_0}(\cdot)$ which keeps only s largest positive entries while setting others to zero.

To provide the algorithm with an initial guess, we first project the given data matrix M on the low-rank constraint as $L^0=\mathrm{H}_r^{\mathrm{SVD}}(M)$. We then fit the least-squares problem with $\hat{P}^0=\arg\min_P\left\|L^0+\Psi(P)-M\right\|_F^2$ with a least-squares solver for sparse systems and project the result on the sparse constraint as $P^0=\mathrm{H}_s^{\ell_0}\left(\hat{P}^0\right)$.

5.2 Numerical Experiments

We provide an experimental comparison of the proposed LRPT method with two variants of the widely used AnnPCA method, one based on the SNR detection map (AnnPCA-SNR) [GG+17] and the other one on likelihood ratio map (AnnPCA-LRM) [DVC+22]. We compare the methods in terms of the visual quality of the detection maps and the receiver operating characteristic (ROC) curves. The reproducible Python implementation

Algorithm 1: NAHT algorithm for (5.3)

Input: $M \in \mathbb{R}^{T\times N^2}$ and parameters $r, s \in \mathbb{N}$

$L^0 = \mathrm{H}_r^{\mathrm{SVD}}(M)$

$\hat{P}^0 = \arg\min_P \left\|L^0 + \Psi(P) - M\right\|_F^2$

$P^0 = \mathrm{H}_s^{\ell_0}\left(\hat{P}^0\right)$, $\Omega^0 = \mathrm{supp}(P^0)$, $j = 0$

while *not converged* **do**

  $L^{j+1} = \mathrm{H}_r^{\mathrm{SVD}}\left(M - \Psi(P^j)\right)$

  $R^j = -\Psi^*\left(M - L^{j+1} - \Psi(P^j)\right)$

  $\alpha_j = \left\|R^j_{\Omega^j}\right\|_F^2 \Big/ \left\|\Psi\left(R^j_{\Omega^j}\right)\right\|_F^2$

  $P^{j+1} = \mathrm{H}_s^{\ell_0}\left(P^j - \alpha_j R^j\right)$, $\Omega^{j+1} = \mathrm{supp}(P^{j+1})$

  $j = j + 1$

end

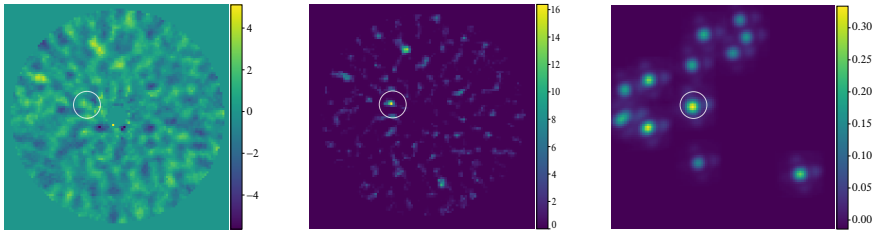

(a) AnnPCA-SNR **(b)** AnnPCA-LRM **(c)** LRPT

Fig. 5.3 Plot of the detection map for each of the methods from a single experiment on the β-Pictoris dataset with an injected planet at $3\lambda/D$ seperation for flux $1.5\sigma_{\mathrm{ann}}$, with the correct location circled in white. The different color scales do not necessarily affect the ROC curves, since the detection threshold is chosen relative to the highest observed value.

is publicly available[1].

We perform the tests on the common benchmark of the ADI cube VLT/NACO β-Pictoris in the infra-red L' band (3.8μm), with 612 frames covering 83° of parallactic angles and $\lambda/D \approx 4.6$ pixel [AMM+13]. In order to reduce the computation time, we subsample the cube to include only every third frame, resulting in a cube of size $204 \times 100 \times 100$.

To compare the methods, we set the rank in the methods to be the best performing rank in case of the flux of $1.5\sigma_{\mathrm{ann}}$, relative to the standard deviation of the pixel intensities in the annulus σ_{ann}, which is equal to $r = 10$

for LRPT and to $r = 35$ for AnnPCA. The sparsity parameter for LRPT is chosen to be $s = 10$. In order to provide a fair comparison, we implement an annular version of the LRPT method by solving the optimisation problem (5.3) for different annular regions as is also done in AnnPCA. Moreover, in AnnPCA we disable the additional heuristic for parallactic angle threshold in the VIP package [GG+17], which skips some frames if they are too close to each other. Turning off the angle heuristic ensures that the low-rank component is based on the same principle in both of the methods which results in the same model for the quasi-static field[2].

Fig. 5.3 shows the visual quality of detection maps for the methods on the β-Pictoris dataset for a synthetically injected planet with the value of flux $1.5\sigma_{\text{ann}}$ relative to the standard deviation of the pixel intensities in the annulus σ_{ann}. We see that whereas the AnnPCA-SNR fails to produce a detection map clearly identifying the exoplanet, AnnPCA-LRM and LRPT manage to produce meaningful guesses. Moreover, LRPT produces a detection map based on a positive sparsity criterion in (5.1) that could be further fine-tuned to detect a specific number of planets.

We compute the deterministic receiver operating characteristic (ROC) curves using the methodology for counting the true positive (TPR) and false positive (FPR) detections as in [DVC+22], whose main advantage is that it is easily reproducible and that for datasets without a planet any detection algorithm would generate a nearly diagonal line. To compute ROC curves, we generate synthetic groundtruth examples with varying flux, i.e., intensity of the planet, by injecting the planet-free data cube (the data cube where the known planets have been removed) with fake planets using the VIP HCI package [GG+17]. Although the injection of different levels of quasi-static noise is not possible using the VIP HCI package, we change the signal to noise ratio by controlling the intensity of the injected planets. Fig. 5.4 depicts the performance of the methods in terms of deterministic ROC curves computed for values of flux relative to the standard deviation of the annulus $c \cdot \sigma_{\text{ann}}$, where $c = 0, 1, 1.5, 2$. In order to lower the computational time we crop the β-Pictoris dataset around the star resulting in a cube of size $204 \times 60 \times 60$. We observe that the LRPT method consistently outperforms the AnnPCA-SNR in terms of the ROC for all tested values of flux and AnnPCA-LRM for flux $c = 1.5$.

[2]This is achieved by setting `delta_rot` parameter to zero and `max_frames_lib` parameter to the number of frames in the `vip_hci.psfsub.pca_local` module.

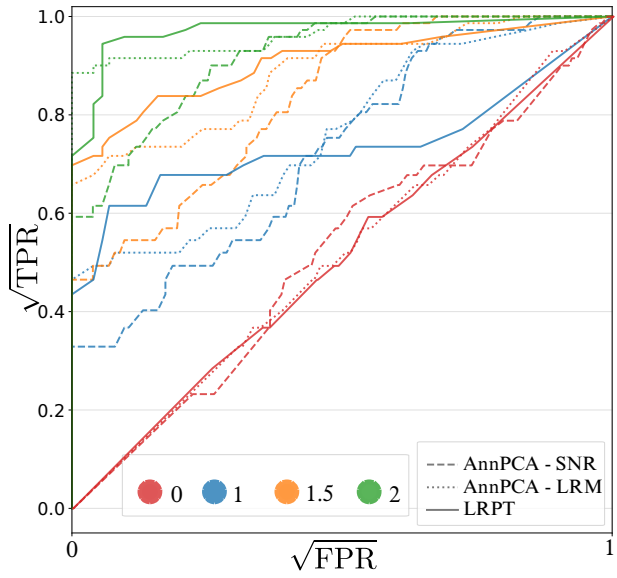

Fig. 5.4 ROC curves for β Pictoris data set. The planets are injected in 3 λ/D separation with the fluxes indicated in the legend multiplied by the standard deviation of the annulus. In order to better observe the low FPR regime, we scale the axes using the square root.

5.3 Conclusion & Future work

We proposed a model-based method for exoplanet detection that attempts to capture both the structure of the quasi-static speckled field and the shape of the circular planet trajectories can take in concurrence. This opens several avenues for further research, for example, by imposing sparsity in different domains to detect other celestial objects, such as circumstellar discs [PCJ21, SPL12]. Finally, one could adapt this method to work on images with multitude of spectral bands for which other geometric transformations, such as radial dilation, are expected.

6

Direct Exoplanet Detection Using L1 Norm Low-Rank Approximation

In the previous chapters, when we fit the low-rank approximation by minimizing the Frobenius norm, i.e., by PCA (or equivalently, truncated SVD), this corresponds to the maximum likelihood estimator under the i.i.d. white Gaussian noise assumption. However, several recent lines of work observed that the residual datacube, which is obtained by subtracting the low-rank part from the original data, is more compatible with the Laplacian distribution that has heavier tails instead of Gaussian [PCGG+19, DVC+22] as we also discussed in Chapter 4.

In this chapter, we propose to perform the low-rank background approximation using the component-wise L1 norm and to apply a data-dependent approximation. This approach leverages two distinct annuli in the frames, one close and the other far from the star. Firstly, it involves an analysis of the data to identify whether the data is better fitted by a Gaussian or Laplacian distribution. Subsequently, we evaluate the performance of the proposed L1 low-rank approximation (L1-LRA) algorithm based on visual comparisons and receiver operating characteristic (ROC) curve analysis. We compare the results of the L1-LRA with the widely used truncated

singular value decomposition (SVD) based on the L2 norm in two different annuli, one close to the star and one far away.

The rest of this chapter is structured as follows. In Section 6.1, we propose an alternative to PCA in the context of exoplanet detection, namely a component-wise L1 norm low-rank matrix approximation. We investigate different statistical assumptions on the data and analyze the performances, then apply the appropriate low-rank approximation to the data and present the experimental results in Section 6.2. Finally, we conclude in Section 6.3 and discuss potential future works.

This chapter is based on the paper [DVL+23].

6.1 Models and Methods

6.1.1 Planet flux estimation

Let $M \in \mathbb{R}^{T \times N^2}$ be a matrix of observations that consists of T unfolded frames with size $N \times N$, i.e., each row of the matrix represents a single vectorized frame. The model for M proposed in [DVC+22], assuming a single planet located at position $g \in [N] \times [N]$ in the first frame, is expressed as

$$M = L + a_g P_g + E, \quad \text{rank}(L) \leq k, \tag{6.1}$$

where L is the low-rank background, E is the noise, a_g is the intensity of the planet referred to as the *flux*, P_g is the planet signature along the trajectory.

For such models based on low-rank approximations, the choice for the rank value is crucial. Indeed, if the rank is too small, the signal of the speckles will remain in the residual $M - L$, making it difficult to separate the signal of the planet from the speckles. Conversely, if it is too large, the signal of the planet will be captured by the low-rank matrix, and it will be challenging to find the signal of the planets in the residual.

When the error E is Gaussian distributed, the maximum likelihood estimator for L is obtained by the minimization of the Frobenius norm. Classical methods, such as PCA and LLSG, fit the low-rank component as follows:

$$\hat{L} = \arg\min_L \|M - L\|_F \quad \text{subject to} \quad \text{rank}(L) \leq k, \tag{6.2}$$

where $\|A\|_F$ denotes the entry-wise L2-norm of A (the Frobenius norm), which can be solved using the truncated SVD. This step is followed by

subtracting the low-rank component and identifying the planet $a_g P_g$ in the residual matrix by solving the minimization problem

$$\hat{a}_g = \arg\min_{a_g > 0} \left\| M - \hat{L} - a_g P_g \right\|_2 , \tag{6.3}$$

as we did in (7.5), for all possible P_g.

In order to detect the presence of a planet, (6.3) has to be solved for all planet positions g of potential interest, which are all the pixels except for the ones that are too close to the star and those located at the corners since they cannot adequately capture the rotational motion of the planet; see Section 6.1.2 for the detection step.

Recently, it has been observed that the error term E has heavy tails and more closely follows the Laplacian distribution [PCGG+19]. Consequently, it was proposed in [DVC+22] to identify the planet using L1 minimization

$$\hat{a}_g = \arg\min_{a_g > 0} \left\| M - \hat{L} - a_g P_g \right\|_1 , \tag{6.4}$$

where $\| \cdot \|_1$ is the entry-wise L1 norm, that is, $\|M\|_1 = \sum_{i,j} |M[i,j]|$.

However, this approach introduces an inconsistency in our assumption between the speckle subtraction (6.2) and planet identification (6.4). While L1 minimization proves effective for planet identification, it raises the question of why not apply the L1 norm to stellar light subtraction as well.

Consequently, we propose to fit the low-rank component using the component-wise L1 norm

$$\hat{L} = \arg\min_{L} \|M - L\|_1 \quad \text{subject to} \quad \operatorname{rank}(L) \leq k, \tag{6.5}$$

followed by the planet estimation in L1 norm as stated in (6.4). This approach allows us to maintain a consistent noise assumption across both the low-rank speckle subtraction and the planet estimation phases. By doing so, we aim to enhance the overall integrity and reliability of the analysis.

We observe that the distribution of the residual cube $R = M - L$ and the error term $E = M - L - a_g P_g$ are quite similar on both norms which can be seen in Fig. 6.1.

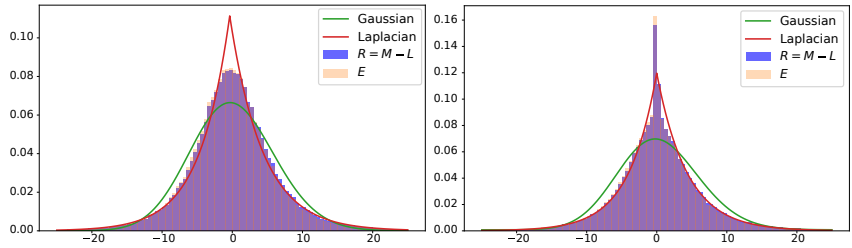

(a) PCA **(b)** L1-LRA

Fig. 6.1 Histograms of residual cubes and error after low-rank approximation applied. The distributions of residual cube and error are pretty similar.

Moreover, in PCA, a common issue is the sensitivity to outliers when using the L2 (Frobenius) norm. An example problem is given in Fig. 6.2, where fitting the data points with the L2 norm results in a poor fit due to an outlier. In contrast, the L1 norm demonstrates a robust approach to outliers, leading to a better fit as shown in the same figure. This makes the L1 norm a more suitable choice for data with potential outliers.

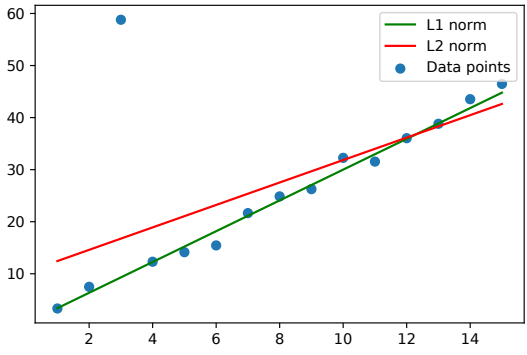

Fig. 6.2 An example of data fitting using L2 and L1 norms, with randomly selected data for visualization.

The L1 low-rank approximation in (6.5) is an NP-hard problem, even in the rank-one case [GV18]. Hence, most algorithms to tackle (6.5), such as alternating convex optimization [KK05], the Wiberg algorithm [EVDH10], and augmented Lagrangian approaches [ZLS+12], do not guarantee to find a global optimal solution, unlike in the case of PCA. Moreover, the computed solutions are sensitive to the initialization of the algorithms.

We use Algorithm 2 (L1-LRA) suggested by [GV18] to solve (6.5). It solves the problem using an exact block cyclic coordinate descent method, where the blocks of variables are the columns of $\hat{U}$ and the rows of $\hat{V}$ of the low-rank approximation $\hat{L} = \hat{U}\hat{V}$. We initialized the algorithm with the truncated SVD solution, denoted by $\mathrm{H}_k^{\mathrm{SVD}}(\cdot)$ in Algorithm 2. In our experiments, we apply an annular version, similar to annular PCA (AnnPCA) [GG+17, CGF+23], that selects only the pixels of M in a certain annulus. Indeed, as the values of pixels decrease away from the star, it is usually better to calculate the low-rank approximation of each annulus separately.

Algorithm 2: L1-LRA [GV18]

Input: Image sequence $M \in \mathbb{R}^{t \times n}$, rank k, maximum number of iteration maxiter

1 $\dot{U}, \dot{S}, \dot{V}^T = \mathrm{H}_k^{\mathrm{SVD}}(M)$
2 $\hat{U} = \dot{U}\dot{S}; \quad \hat{V} = \dot{V}^T$
3 **for** i =1: maxiter **do**
4 $\quad R = M - \hat{U}\hat{V}$
5 $\quad$ **for** $j = 1 : k$ **do**
6 $\quad\quad R = R + \hat{U}[:, j]\hat{V}[j, :]$
7 $\quad\quad \hat{U}[:, j] = \min_u \|R - u\hat{V}[j, :]\|_1$
8 $\quad\quad \hat{V}[j, :] = \min_v \|R^T - v\hat{U}[:, j]\|_1^T$
9 $\quad\quad R = R - \hat{U}[:, j]\hat{V}[j, :]$
10 $\quad$ **end**
11 **end**

To solve the minimization problem in steps 7-8 of Algorithm 2, we use the exact method from [GP11]; these subproblems are weighted median problems which can be solved in closed form.

In Fig. 6.3, we observe an outlier within a data cube frame, highlighted in a white square. When we apply PCA to the data cube containing this frame, PCA is able to capture this outlier. While initially, it seems beneficial to exclude the outlier from the residual matrix, this approach has unintended consequences. As shown in the red square, we can see the effects of this outlier on the same pixel across other frames. This effect is undesirable because it is a distortion of the overall data integrity. On the contrary, when we apply L1-LRA, this approximation may not capture the outlier, but importantly, it avoids influencing other frames. This char-

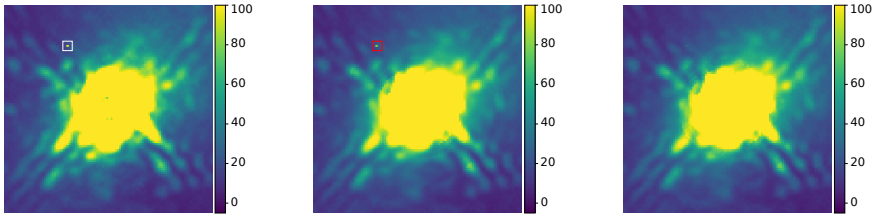

(a) Frame with outlier — original from datacube, low-rank by PCA, and low-rank by L1-LRA (from left to right).

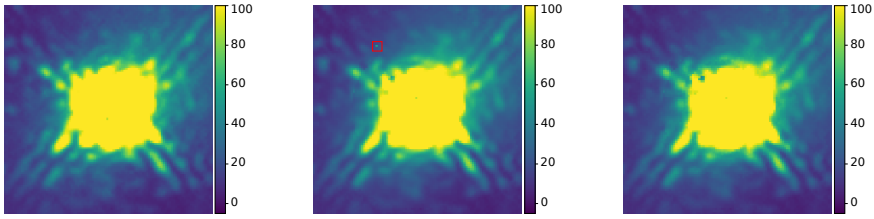

(b) Frame without outlier — original from datacube, low-rank by PCA, and low-rank by L1-LRA (from left to right)

Fig. 6.3 In each row, different frames of a datacube are shown. The top row highlights an outlier marked by a white square. The left column displays the original frames from the datacube M. PCA and L1-LRA are applied to M, with the middle column showing frames from the low-rank matrix after PCA and the right column showing frames after L1-LRA. PCA captures the outlier (red square) but introduces artifacts in frames without outliers, whereas L1-LRA avoids such artifacts.

acteristic makes L1-LRA a more favorable choice, as it affects the overall negative impact much less.

6.1.2 Planet detection

The detection procedure consists of declaring positive the positions g where a detection metric—e.g., an SNR or a likelihood ratio—exceeds a given threshold. A likelihood ratio map (LRM) was proposed in [DVC+22], which consists of L1 norm likelihood ratios $\Lambda_g(R)$ based on maximizing log-likelihood of Laplace distribution using the solution of (6.3) or (6.4) because it has been shown to provide better results in prac-

tice [DVC+22]

$$\log \Lambda_g(R) = -\sum_{(t,r)\in\Omega_g} \frac{|R(t,r) - \hat{a}_g P_g(t,r)| - |R(t,r)|}{\sigma_{R(r)}}, \tag{6.6}$$

where $R = M - \hat{L}$, σ_R is the standard deviation of R computed along the time dimension, and Ω_g is the set of indices (t,r) of pixels whose distance from the trajectory is smaller than half the diffraction limit, for more details see [DVC+22]. We will also use the L2 norm LRM which consists of L2 norm likelihood ratios $\Lambda_g(R)$ using the solution of (6.3) or (6.4)

$$\log \Lambda_g(R) = -\frac{1}{2}\sum_{(t,r)\in\Omega_g} \frac{|R(t,r) - \hat{a}_g P_g(t,r)|^2 - |R(t,r)|^2}{\sigma^2_{R(r)}}. \tag{6.7}$$

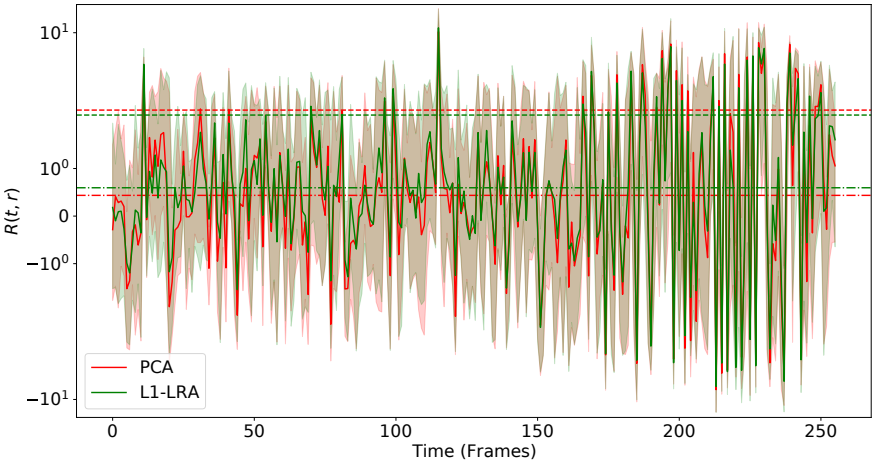

Fig. 6.4 The signals in the trajectory centered in the planet pixel in residual cube, obtained by PCA and L1-LRA. The dark lines are mean values in the aperture and the light interval is limited with minimum and maximum intensities in the aperture.The dashed lines and the dashed-dot lines represent the standard deviation and mean values of all pixels in the trajectory.

In Fig. 6.4, we examine the signals in the aperture centered on the planet along the trajectory when we apply L1-LRA and PCA for background subtraction. Although we seem to get similar results, small differences can have a significant impact on the planetary signal. In particular, when we apply the LRM, the higher values in the residual but lower standard devi-

ation in the trajectory containing the signals of the planet will lead to more successful results. When we use the L1-LRA approximation, the standard deviation is lower and the mean is higher, which is exactly the result we want.

6.2 Numerical Experiments

In order to analyze the performance of the L1-LRA algorithm for background subtraction in exoplanet detection, we compare the algorithms visually by plotting their LRMs and also using receiver operating characteristic (ROC) curves. Additionally, we empirically investigate the methods in terms of fitting the data to Gaussian and Laplacian distributions. The Python codes of the implementations are publicly available from `https://github.com/hazandaglayan/l1lra_for_exoplanets`.

We tested the algorithms using the publicly available dataset *sph3* for the exoplanet data challenge [C+20]. The ADI cube obtained with the VLT/SPHERE-IRDIS instrument has 228 frames with size 160×160, and it has a total field rotation of 80.5 degrees. It is a real dataset, but it has 5 synthetically injected planets at different distances from the star.

6.2.1 Empirical estimation of the noise distributions

In order to analyze the suitability of different noise assumptions, we fit the Gaussian and the Laplacian distribution to the residual data, i.e., the data after subtracting the low-rank component using PCA or L1-LRA. We look at two different annuli separately, one that is close to the star at $4\lambda/D$ separation and one more distant from the star at $10\lambda/D$, and measure the goodness of fit visually and by the coefficient of determination ρ^2.

In Figure 6.5, we observe that the residual data follows a Gaussian distribution after applying PCA and a Laplace distribution after applying L1-LRA in general. However, the Laplacian distribution provides a better fit for the residual cube distribution in the tails, regardless of whether PCA or L1-LRA is used, for both small and large separations.

To assess the quality of fit of a distribution to the residual data, we use the coefficient of determination as a metric. The coefficient of determination ρ^2 quantifies the relationship between two sample sets, denoted as

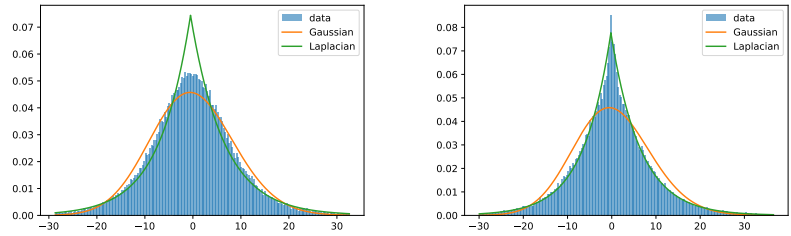

(a) PCA at $4\lambda/D$

(b) L1-LRA at $4\lambda/D$

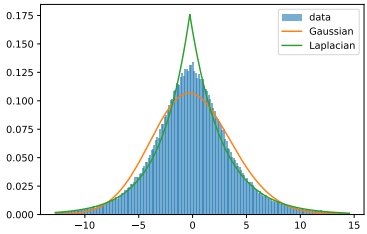

(c) PCA at $10\lambda/D$

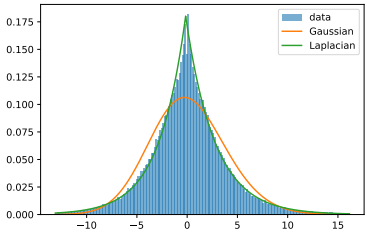

(d) L1-LRA at $10\lambda/D$

Fig. 6.5 Histograms of residual cubes after low-rank approximation applied for small and large separations

$X = (x_1, \ldots, x_n)$ and $Y = (y_1, \ldots, y_n)$, and is computed as follows:

$$\rho^2 = \frac{\left(\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})\right)^2}{\left(\sum_{i=1}^{n}(x_i - \bar{x})^2\right)\left(\sum_{i=1}^{n}(y_i - \bar{y})^2\right)}, \tag{6.8}$$

where $\bar{x}$ and $\bar{y}$ represent the means of the variables in samples X and Y. This metric aids in comparing the goodness of fit of different distributions. A ρ^2 value close to 1 indicates that the data is more consistent with the distribution. In our experiments, X represents the height of the bins of the data histogram, while Y corresponds to the values of the probability density function within these bins. In Table 6.1 and 6.2, we compute ρ^2 value between the tails of the data distribution and the Gaussian or Laplacian distributions. The highest ρ^2 is obtained with L1-LRA for both small and large separations.

Table 6.1 The coefficient of determination ρ^2 for small separation $4\lambda/D$

Rank	PCA		L1-LRA	
	Gaussian	Laplacian	Gaussian	Laplacian
5	0.9876	0.9918	0.9878	**0.9920**
10	0.9866	0.9890	0.9878	**0.9938**
15	0.9841	0.9948	0.9916	**0.9972**

Table 6.2 The coefficient of determination ρ^2 for large separation $10\lambda/D$

Rank	PCA		L1-LRA	
	Gaussian	Laplacian	Gaussian	Laplacian
5	0.9866	0.9920	0.9894	**0.9940**
10	0.9859	0.9936	0.9872	**0.9948**
15	0.9912	0.9954	0.9922	**0.9960**

6.2.2 Performance comparison of L1-LRA and PCA

We test the performance advantage of using the L1-LRA instead of PCA approximation in terms of visual quality and ROC curves tested on the sph3 test data [C+20]. We also provide ablation studies using combinations of mixed models, where the low-rank L is first subtracted from the original data M using L1-LRA or PCA, then used in the planet detection. As such, we have four possible algorithms to investigate; see Table 6.3.

Table 6.3 Descriptions of algorithms

Background subtraction	Planet detection	
	L1 (6.4) and (6.6)	L2 (6.3) and (6.7)
L1-LRA (6.5)	L1L1	L1L2
PCA (6.2)	L2L1	L2L2

We first compared the algorithms, visually in the LRM using the dataset with five synthetic planets. We selected the best performing rank according to the average likelihood ratio over the locations of the injected planets. Based on this, we chose ranks 5, 6, 5, and 10 for L1L1, L1L2, L2L1, and L2L2, respectively.

Figure 6.6 shows the LRM and the intensity of pixels in the LRM obtained with the four algorithms. In each right-hand plot, the black line shows the lowest threshold for which there is no false positive. At this threshold, we observe that L1L1 and L2L1 detect the five injected planets, whereas L1L2 and L2L2 miss one of them. Moreover, the likelihood ratios of the planet pixels (the red dots in the figures) obtained by L1L1 are much larger than the threshold.

To evaluate the performance of the four methods, we put them to the test using synthetically generated data and examine their results using ROC curves in Fig. 6.7. The detection map used as input for the ROC curve procedure of [DVC+22] is the LRM. We deleted the five injected planets from the dataset using VIP-HCI package [GG+17, CGF+23]. Then, we created 50 different datasets by injecting two planets in each, 180 degrees apart, and placed at the separation $4\lambda/D$. We set the intensity of each injected planet as 1σ where σ is the standard deviation of the annulus. We applied the same procedure to a larger separation of $10\lambda/D$ with the intensity 0.6σ. The four methods were tested for three different rank values $k = \{5, 10, 15\}$.

In the ROC curve results, we focus on the number of true positives before the first false positives are found, as done in [DVC+22]. In the ROC curves obtained by injecting the planets at the small separation (left-hand column of Figure 6.7), L1L1 always gives the best results. This is also the case for the large separation with $k = 5$ (top right plot of Figure 6.7). In the remaining two plots (i.e., large separation and $k = \{10, 15\}$), the four algorithms perform rather similarly. Moreover, we can see a relationship between the ROC curve performances of the algorithms and the ρ^2 values. L1L1, which corresponds to blue curves, has the highest ρ^2 value, while L2L2, which corresponds to red curves, has the lowest value. In addition, L2L1 generally has the second-highest ρ^2 value while generally performing the second-best on ROC curves.

We also compared the L1-LRA and PCA algorithms using SNR map as a detection map, which is a classical approach in direct imaging [MMW+14]. We use the implementation of SNR available in the VIP package [GG+17, CGF+23]. Figure 6.8 shows the ROC curves of SNR maps applied after PCA and L1-LRA. We use the same inputs and procedure as described above. Except for the results of $k = 5$, L1-LRA outperforms PCA in all ROC curves for both separations.

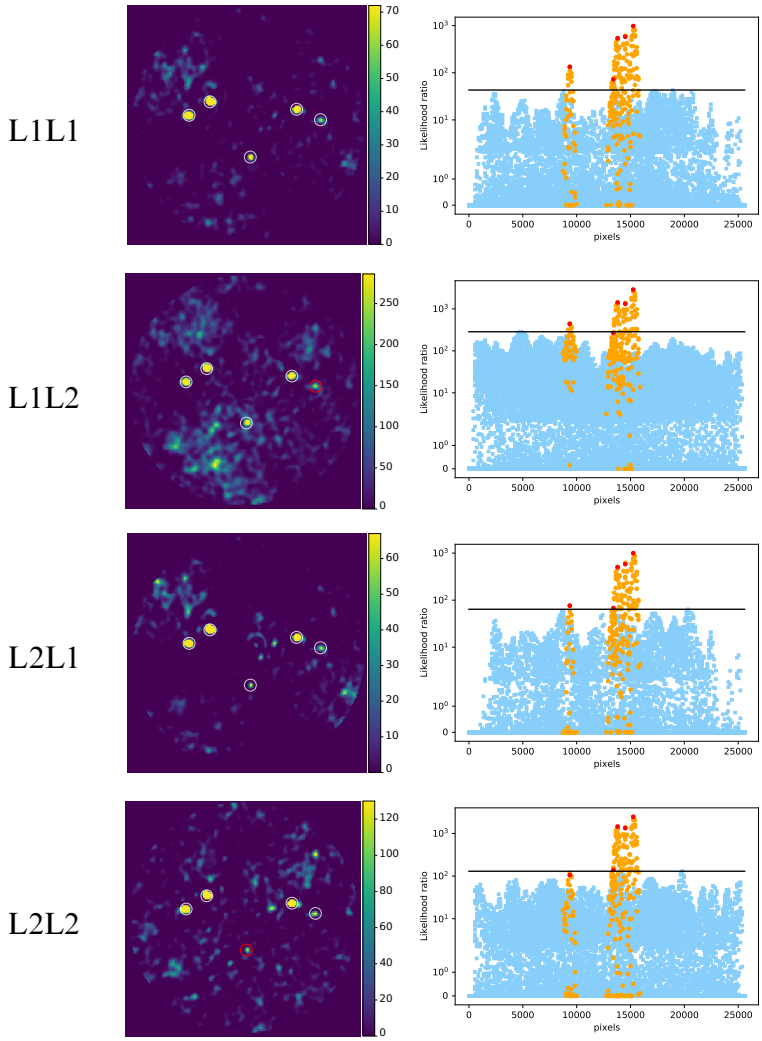

Fig. 6.6 Left column: LRM. The detected planets are circled in white and the non-detected planets are circled in red. Right column: Intensity of the pixels in the LRM. Those maps have been flattened such that the horizontal axis corresponds to the pixels of the LRM, and the vertical axis indicates the intensity of the pixels in the LRM. The center pixels of the planets and the pixels near the planet are shown in thick red and orange dots, respectively. Other pixels are displayed as blue squares. The black line corresponds to the highest blue square.

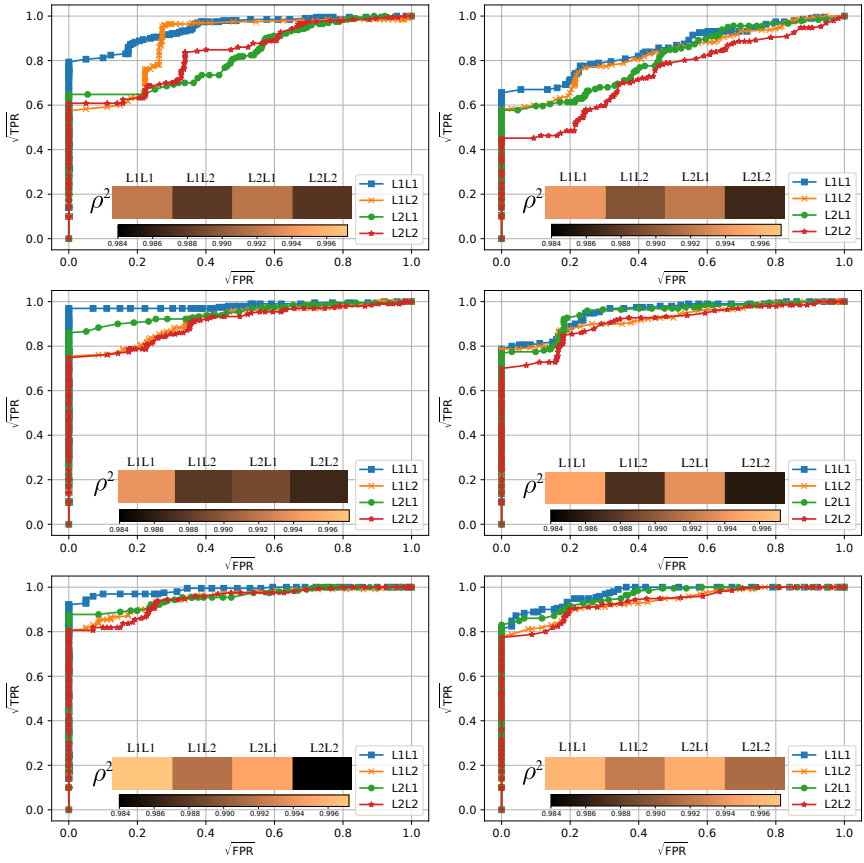

Fig. 6.7 ROC curves. Left column plots belong to data cubes where the planets are injected in small separation ($4\lambda/D$). Right column plots belong to data cubes where the planets are injected in large separation ($10\lambda/D$).The rank is 5, 10, and 15 from top to bottom. Brown bars show the coefficient of determination ρ^2 for each algorithm.

6.3 Conclusion

In this chapter, we proposed a low-rank approximation algorithm in the sense of the entry-wise L1-norm, termed L1-LRA, for stellar light subtraction to detect exoplanets. A comparison with the commonly used PCA based on the L2 norm was conducted within two distinct annuli: one near the host star and the other farther away. Through empirical analysis, we examined the statistical noise assumptions of our models and evaluated

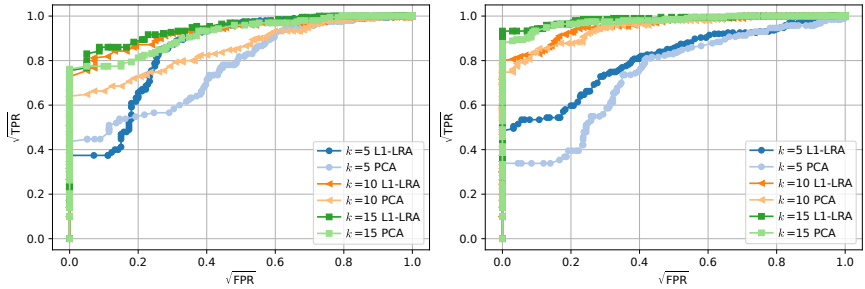

Fig. 6.8 ROC curves of SNR maps. The left column plot belongs to data cubes where the planets are injected in small separation ($4\lambda/D$), while the right column plot belongs to data cubes where the planets are injected in large separation ($10\lambda/D$).

the performance of the L1-LRA algorithm using visual comparisons and ROC curves. In addition to the success observed in these comparisons, our evaluation of performance metrics, including the coefficient of determination ρ^2, indicates that the L1L1 approach consistently outperforms alternatives. The proposed L1-LRA algorithm presents a promising solution for the challenging task of exoplanet detection.

There are several directions for future studies, such as using different initializations for the L1 norm approximation or designing more efficient algorithms (e.g., based on smoothing techniques).

7

An Alternating Minimization Algorithm with Trajectory for Direct Exoplanet Detection

Effective image post-processing algorithms are vital for the successful direct imaging of exoplanets. Existing algorithms use techniques based on a low-rank approximation to separate the rotating planet signal from the quasi-static speckles. Then, they try to obtain planet flux or signal-to-noise ratio to detect the planet. These steps do not interact or feed each other, leading to potential limitations in the accuracy and efficiency of exoplanet detection. Unlike traditional methods, we propose a novel approach that iteratively finds the planet's flux and the low-rank approximation of quasi-static signals, strengthening the existing model based on low-rank approximations. In this section, we also integrate the knowledge gained in the previous sections and the new iterative method. Moreover, thanks to the capability of our model to determine the flux, we also gain the opportunity to characterize the exoplanet, which means determining the contrast of the planet relative to its host star.

In Section 7.1, we propose a new method, dubbed Alternating Minimization Algorithm with Trajectory (AMAT), that makes a more advanced use of the low-rank-plus-trajectory model by alternating between computing the low-rank approximation of the speckle field and the planet flux. We pose our model as ℓ_1 and ℓ_2 norm minimization problems. We also give the details of Annular AMAT, highlighting the differences from Annular PCA. The construct of the residual cube in our method also differs from the existing methods in the literature. We explain how we obtain the residual cube and two alternatives of the processed frame. Moreover, we discuss the selection of parameters through experimentation.

AMAT also helps to estimate the flux of the exoplanet while trying to solve the proposed problem. In Section 7.2, we discuss how the algorithm estimates it and compared it with the alternative algorithms in the literature in Section 7.3.

Numerical experiments performed on dataset VLT/SPHERE-IRDIS 51 Eri show the potential of AMAT to improve upon the existing approaches in terms of more clear detections, receiver operating characteristic (ROC) curves, detection maps applying classical SNR map and likelihood ratio map suggested in Chapter 4, and contrast curve in Section 7.3. Moreover, for a data-independent comparison, we used datasets from the exoplanet data challange [C+20] and compared them to other algorithms in the challange. Finally, conclusions are drawn in Section 7.4.

This chapter builds upon the foundational concepts introduced in [DVA23], expanding them through the application of likelihood ratio map, experiments with many datasets, and estimation of fluxes of planets.

7.1 Alternating Minimization Algorithm with Trajectory

Let $M \in \mathbb{R}^{T\times N^2}$ be an ADI sequence that consists of T unfolded images with size $N \times N$. The standard low-rank model [DVC+22] can be expressed as

$$M = L + a_g P_g + E, \quad \text{rank}(L) \leq r, \quad P_g \in \Lambda, \tag{7.1}$$

where L is a low-rank matrix representing the background with quasi-static speckles, a_g, referred to as *flux*, corresponds to the intensity coefficient of the planet signature $P_g \in \Lambda \subset \mathbb{R}^{T\times N^2}$, where $g \in [N] \times [N]$ stands for the position of the planet on the initial frame, Λ is the set of all feasible planet signatures, and E is the noise term.

In existing approaches, the low-rank matrix L is estimated first. In particular, the PCA approach computes L in (7.1) by truncated SVD. As an alternative, L can be obtained by L1 low-rank approximation (L1-LRA), as proposed in [DVL+23]. Once L is estimated, a residual cube is then obtained as $M - L$, from which the planet flux a_g is estimated for each planet trajectory g by various methods; see, e.g., [AQ12, SPL12, DVC+22].

Instead of sequentially estimating first the component L followed by estimating the flux a_g, we propose to estimate them simultaneously in the following optimization problem

$$\min_{L\in\mathbb{R}^{T\times N^2}, a_g\in\mathbb{R}} \left\|M - L - a_g P_g\right\| \quad \text{s.t.} \quad \operatorname{rank}(L) \leq r, \quad P_g \in \Lambda \tag{7.2}$$

where $\|\cdot\|$ denotes either ℓ_1 or ℓ_2 norm depending on the assumed distribution of the error E in (7.1). The optimization problem (7.2) is addressed by alternatingly solving the following two sub-problems until a stopping criterion is met:

$$L^{(i)} = \arg\min_{L\in\mathbb{R}^{T\times N^2}} \|M - L - a_g^{(i-1)} P_g\| \tag{7.3}$$

$$a_g^{(i)} = \arg\min_{a_g\in\mathbb{R}} \|M - L^{(i)} - a_g P_g\|. \tag{7.4}$$

The stopping criterion is defined as either reaching a maximum number of iterations or ensuring that the relative changes in the intensity $a_g^{(i)}$ are less than a specified threshold.

When we select the norm L2 in (7.2), computing $L^{(i)}$ in Eq. (7.3) amounts to a k-truncated SVD, denoted by $\mathrm{H}_k^{\mathrm{SVD}}(\cdot)$. In practice, in order to speed up the computations, we compute $L^{(i)}$ by a randomized SVD of $M - a_g^{(i-1)} P_g$ [HMT11]. We have checked that the resulting approximation does not affect the planet detection performance. The optimal value of $a_g^{(i)}$ can be computed by

$$a_g^{(i)} = \frac{\sum_{(\theta,r)\in\Omega_g} R_g(\theta,r) P_g(\theta,r) \sigma_{R_g(r)}^{-2}}{\sum_{(\theta,r)\in\Omega_g} \left(P_g(\theta,r)\right)^2 \sigma_{R_g(r)}^{-2}}, \tag{7.5}$$

where $R_g = M - L^{(i)}$, σ_R^2 is the empirical variance of the residual frames computed along the time dimension and Ω_g is the set of indices (θ, r) of

pixels whose distance from the trajectory is smaller than half the chosen aperture diameter $\rho\lambda/D$

$$\Omega_g = \left\{ (\theta, r) \in [t] \times [n]^2 \mid \|r - g_t\|_2 < \tfrac{1}{2} \tfrac{\lambda}{D} \rho \right\}. \tag{7.6}$$

When the norm is ℓ_1, we solve the problem (7.3) with an algorithm suggested by [DVL+23] using an exact block-cyclic coordinate descent method. Moreover, (7.4) becomes a least absolute deviation problem, which we solve by an efficient exhaustive search over the kink points of the piecewise linear objective function as in [DVC+22].

Algorithm 3: AMAT

Input: Image sequence $M \in \mathbb{R}^{t \times N^2}$, possible trajectories Λ, maximum number of iteration maxiter

for *all trajectories* $P_g(r, \theta) \in \Lambda$ **do**
  $a_0 = 0$
  for $i = 1$: maxiter **do**
    $L_i = \arg\min_L \|M - L - a_{i-1} P_g\|$
    if norm $= 1$ **then**
      $a_g^{(i)} = \arg\min\{\frac{M - L^{(i)}}{P_g}\}$
    end
    else if norm $= 2$ **then**
      $a_g^{(i)} = \langle P_g, M - L_i \rangle / \|P_g\|^2$
    end
    $R(r, \theta) = M - L_i$
  end
end

We produce a flux map, consisting of all possible trajectories, $g \mapsto a_g$ by simply returning (7.4) in the final iterate of the alternating iteration, then compute the SNR map with VIP [GG+17, CGF+23] and use this as the detection map. We define the set of positions of the planet on the initial frame

$$G = \{ g \in \mathbb{Z}^+ \mid \lceil \frac{\lambda}{D} \rceil \leq g \leq \lfloor \frac{N}{2} - \frac{\lambda}{2D} \rfloor \}, \tag{7.7}$$

so that we evaluate the probability of having a planet in all positions except the pixels of the star and the pixels in the corners of the images.

7.1.1 Annular AMAT

Since we know that annular PCA is more successful than full PCA, we can adapt the AMAT algorithm by splitting it into annuli. But our approach differs from that used in PCA:

- In PCA, the data cube M is divided into annuli $M_{ann_1}, M_{ann_2}, M_{ann_3}, \ldots, M_{ann_k}$, where k is the number of annuli. For each annulus, low-rank matrices $L_{ann_1}, L_{ann_2}, L_{ann_3}, \cdots, L_{ann_k}$ are obtained and then combine to form the data cube L.
- In our method, it is not necessary to process all annuli. We focus only on the annulus that contains the trajectory. We take this annulus, denoted as M_{ann}, as the basis for our algorithm. We then find its low-rank matrix L_{ann} and the coefficient a_g corresponding to the annulus that contains the planetary signature $P_{g_{ann}}$.

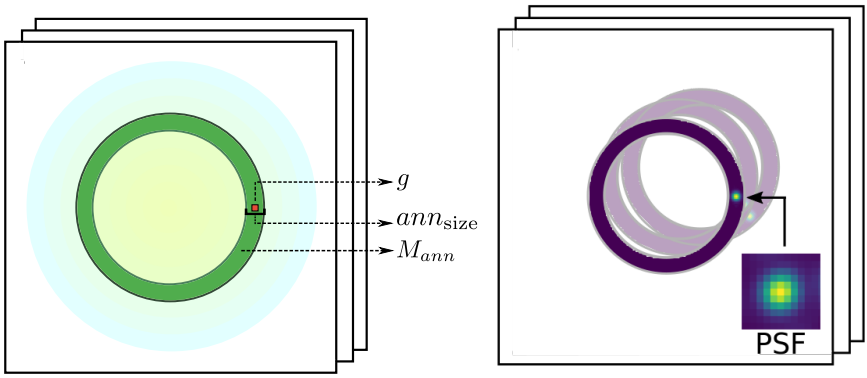

Fig. 7.1 Annuli splitting in AMAT. AMAT is applied to the one annulus represented with green color for the specified trajectory. Similarly, P_g is bounded by the same annulus properties to obtain $P_{g_{ann}}$.

- Unlike in annular PCA, the annular size ann_{size} should not be any chosen value. It must be at least λ/D to cover the planetary signature contained in $P_{g_{ann}}$.

7.1.2 Construct the residual cube & processed frame

In our algorithm, we initiate the process by selecting a pixel, denoted as (x, y), as our starting point. Using the coordinates of this pixel, we con-

struct a trajectory for the potential exoplanet. We obtain this trajectory with the help of the information on how the position of planet evolves with respect to known rotation angles and the PSF.

The next step in our procedure is the application of the AMAT algorithm. Regardless of whether we opt for its full or annular version, the objective remains the same: to estimate a low-rank matrix that encapsulates the background, including quasi-static speckles, and the intensity or flux associated with the exoplanet.

With the low-rank matrix in hand, we proceed to subtract it from the original data cube. This subtraction operation yields a residual matrix that has the potential to contain planetary signals.

To analyze this further, we convert the residual matrix back into a cube format. This helps us align it with our reference pixel (x, y). We then rotate the cube to align with the exoplanet position. The pixel values from each frame corresponding to (x, y) are placed into the same positions in a new empty cube.

We repeat this process for all pixels except those at the frame's corners or overlapping with the star. This approach combines the data obtained from each trajectory, allowing us to create a composite cube that encapsulates the exoplanet signal in all possible scenarios.

Finally, to represent the exoplanet signal in a single processed frame, we can either take the median of the composite cube or use a combination of intensity values a_g derived from the solution of the minimization problem for all possible trajectories. The entire process is illustrated in Fig. 7.2.

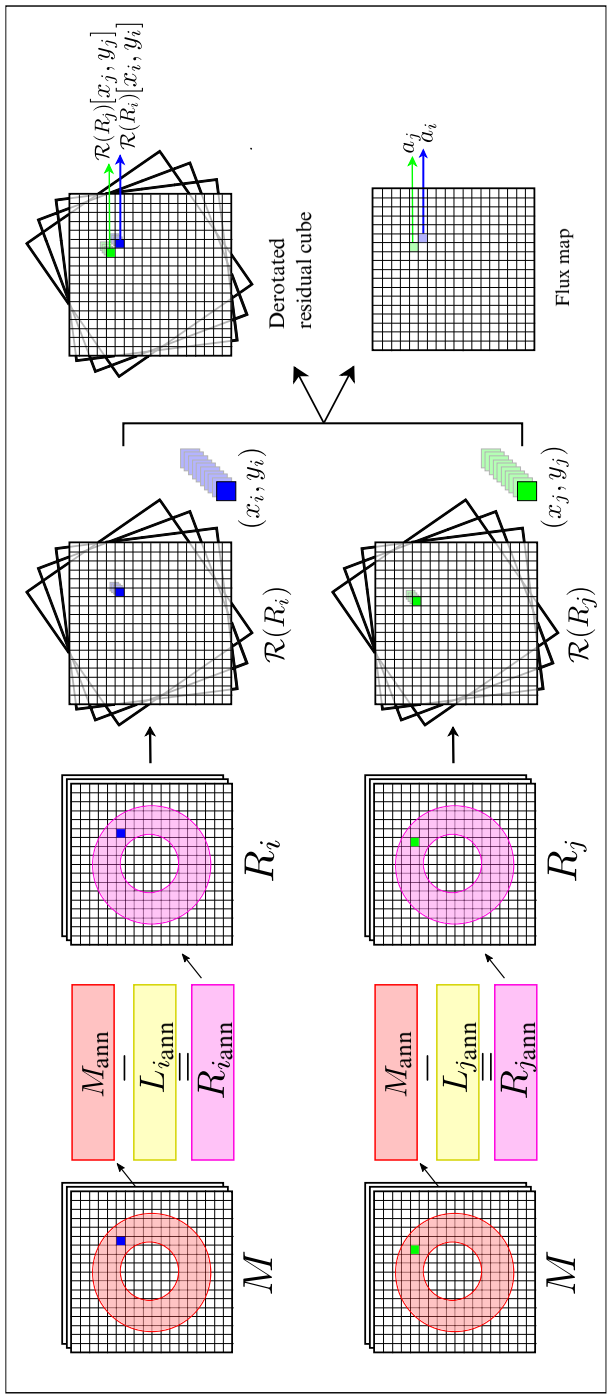

Fig. 7.2 Pipeline of AMAT Algorithm

7.1.3 Parameter selection

Rank: The rank r in the context of ADI sequence processing for exoplanet detection is a critical parameter that directly affects the algorithm's ability to discriminate between the planetary signal and the speckle noise. This parameter directly influences the complexity of the low-rank model used to approximate the star and speckle noise, enabling the capture of quasi-static patterns across the image sequence.

Selecting the most suitable rank involves finding a delicate balance. If the rank is set too low, the model may not capture the speckle patterns. This could lead to mistakenly identifying noise as planets or even losing actual planetary signals. On the other hand, if the rank is set too high, the model may capture the planet signal as part of the speckle noise. This would result in incorrectly subtracting those signals from the image, leading to the inability to detect them.

Ideally, the chosen rank r should encapsulate the essential structure of the starlight and speckles, enabling the algorithm to effectively separate this from the planetary signals. In practice, we see that the selection of rank r may be informed by the location of the planet. When the planet is very close to the star, we should choose higher ranks, and for the far planet, we should choose smaller ranks. We came to this conclusion when we examined the best ranks for planets whose locations we knew.

In Fig. 7.4, a planet close to the star is selected, while in Fig. 7.5 a planet far away from the star is selected. When deciding on the best rank, we choose the rank that gives the highest SNR in the likelihoodratio map. The reason we chose this method is that when we used LR, it was not a correct recommendation to choose based on the LR values because the LR was high where there were planets, but the values were also high in nonplanet pixels, which can be seen in Fig. 7.3. So, we used the likelihood ratio map to calculate the SNR using the pixels in the aperture where the planet was. We obtain the SNR of the likelihood ratio map as follows

$$x_{\text{SNR}} = (x_p - x_{np})/\sigma_{np} \tag{7.8}$$

where x_p is the mean value of the pixels in the planet aperture, while x_{np} and σ_{np} are the mean value and standard deviation of the pixels outside of the aperture.

Figures 7.4 and 7.5 suggest that lower ranks may perform better for small separations, while higher ranks are more suitable for larger separations. This is because noise increases radially and varies significantly

between large and small separations.

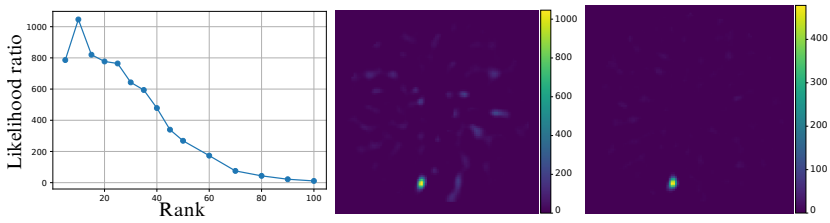

Fig. 7.3 Left: Likelihood ratio vs. rank at the planet's location. Middle: Likelihood ratio map using rank 10. Right: Likelihood ratio map using rank 40. This analysis examines how the likelihood ratio varies with different ranks at the planet's location, alongside a visual comparison of likelihood ratio maps for different ranks. The real planet data from the 51 Eri dataset is used for this comparison.

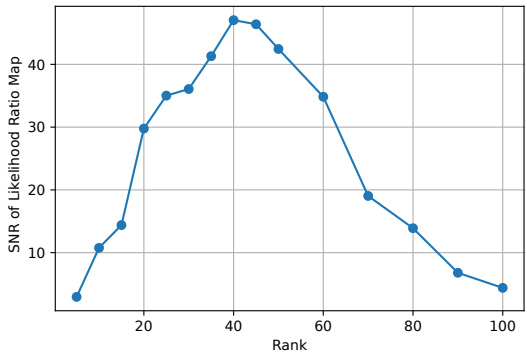

Fig. 7.4 SNR of likelihood ratio map vs. rank at the planet's location using the 51 Eri dataset with an injected planet at a small separation.

Annular size: The annular size parameter defines the size of the annuli into which the data is divided, directly impacting the quality and the sensitivity of the planet detection process. When the annular size is too large, the heterogeneity among pixels increases, making it difficult to compare planetary signals with other pixels around it. This scenario may yield results similar to a full-frame analysis of the ADI sequence, which could make it difficult to detect planets that have faint signals against the bright background of the host star. On the other hand, selecting an annular size

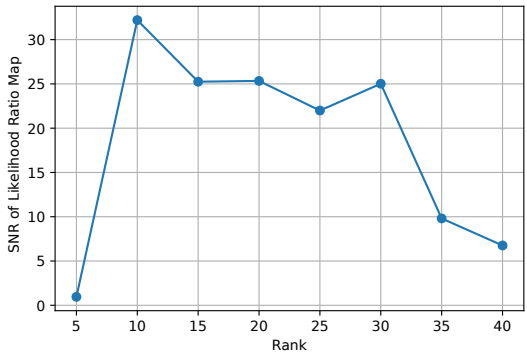

Fig. 7.5 SNR of likelihood ratio map vs. rank at the planet's location using the 51 Eri dataset with an injected planet at a large separation.

that is too small may result in annuli that lack a sufficient number of pixels to construct a reliable statistical model, undermining the ability of the method to characterize the speckle noise accurately.

Therefore, the ring size must be chosen to balance these extremes, optimizing the trade-off between having a sufficient pixel sample and maintaining the high-resolution characteristics critical for detecting exoplanets. We chose this parameter through experimental studies. In Fig. 7.6, we applied our AMAT algorithm by selecting different numbers of pixels. As we did when selecting the rank, we applied SNR to the likelihood ratio maps and decided to select the most appropriate annular size. According to the results of our experiments, we concluded that it is usually appropriate to decide the annular size based on $2\lambda/D$, since the $2\lambda/D$ range is close to the annular size that gives the highest SNR.

7.2 Flux Estimation

The detection of exoplanets is followed by the characterization of the planets, which encompasses estimating their positions and the intensity relative to the host star. Different algorithms, such as the negative fake companion (NEGFC) method [LBC+10, MMV10, WAG+17], ANDROMEDA via maximum likelihood estimation [CMM+15], PACO estimation [FDTL18] are employed for this purpose.

In our work, when using the AMAT algorithm, we produce a flux map

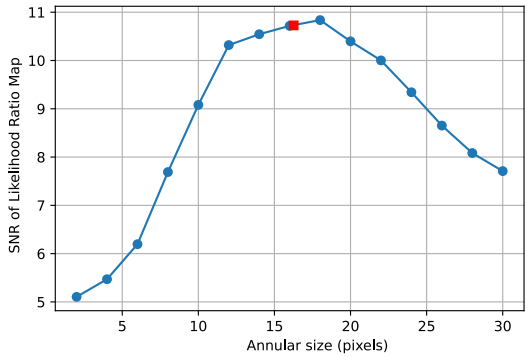

Fig. 7.6 SNR vs annular size in the location of the planet. The red square represents the SNR value when we choose $2\lambda/D$ for annular size.

consisting of the intensities a_g in (7.2) corresponding to each trajectory. Following the planet detection, we use this flux map to estimate the intensity of the planet at its detected location.

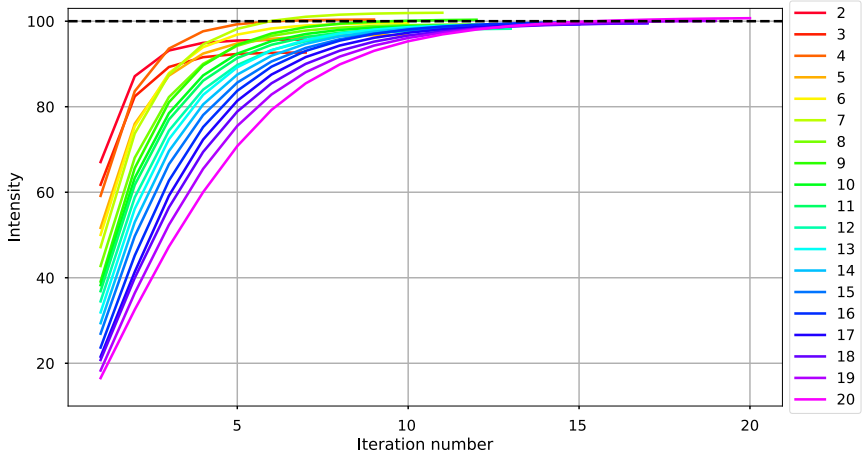

Fig. 7.7 AMAT_{L2} for the flux estimation of an exoplanet. Black dashed line represents the intensity of the injected planet. The algorithm is applied with different ranks ranging from 2 to 20. The iteration number for each rank varies (terminated before reaching the maximum iteration) as the changes in intensity become smaller than the given threshold.

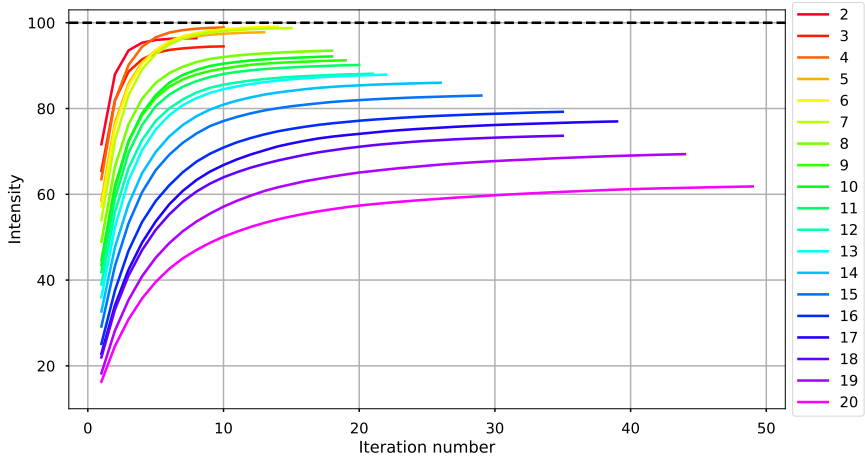

Fig. 7.8 AMAT_{L1} for the flux estimation of an exoplanet. Same inputs as Fig. 7.7.

To test this method, we inject a fake planet into our existing dataset. In this scenario, we already know the intensity of the injected planet. In Fig. 7.7, we apply the L2 norm version of the AMAT algorithm (AMAT_{L2}) for different ranks from 2 to 20. Although we observe that the algorithm requires more iterations at large ranks, the algorithm approximates the injected intensity for many rank possibilities. Similarly, we apply L1 norm version of the AMAT algorithm (AMAT_{L1}) in Fig. 7.8. However, unlike AMAT_{L2}, AMAT_{L1} only gets close to the injected intensity at small ranks, but not at all when using large ranks. The underlying reason for this is traced back to AMAT_{L1}'s initialization process; it begins with a randomized SVD for its initial approximation but subsequently relies on the outcomes of preceding steps for further initialization within its iterative process. This dependency introduces a sensitivity to the initial conditions, which coupled with the inability of the algorithm to guarantee recovery of the global optimum, restricts the effectiveness of AMAT_{L1} at accurately capturing the injected intensity.

We also aimed to assess the effect of injecting a planet at different positions within the pixel. Initially, we injected a planet with an intensity of 100 at the center of the pixel, which resulted in an intensity of 100.07 after the AMAT_{L2} algorithm ran. Next, we injected planets with the same intensity at each of the four corners of the pixel in four separate instances. Upon analysis of the trajectories originating from this pixel in the algorithm, the

resulting intensities were observed to be 97.23, 96.31, 97.19, and 96.36 respectively. We attribute this flux loss to the intra-pixel variation of flux in the injected PSF, considering the fact that the flux map is only estimated at the center of each pixel.

7.3 Numerical Experiments

We begin to investigate the performance of the AMAT algorithm in the field of direct exoplanet detection and perform a series of comparisons and evaluations that shed light on the different characteristics of the algorithm. With each comparative analysis, we try to uncover the algorithm's strengths and distinct advantages. We navigate through diverse scenarios, comparing AMAT against traditional methodologies, examining its full and annular versions, comparing the fluxmaps and processed median frames, assessing its likelihood-based approach versus SNR, and inspecting its performance using the datasets of the Exoplanet Data Challenge. Apart from making detection map comparisons, we try to prove the performance of the algorithm with contrast curve and ROC curve comparisons. Through this comprehensive evaluation, we reveal the multifaceted potential of the AMAT algorithm and demonstrate its important contributions to the field of exoplanet exploration and direct imaging.

We apply our algorithm to the publicly available ADI data set VLT/SPHERE-IRDIS 51 Eri in the K1 (2.11 μm) band with a known planet. The cube has 256 frames covering 42° of parallactic angles and the frames are cropped to 100 by 100 pixels for detection map and ROC curve comparisons and 150 by 150 pixels for contrast curves.

In our preliminary experiments, we simply stop the algorithm after the maximum of 50 iterations or when the relative change of the intensity $a_g^{(i)}$ is small, i.e. $|a_g^{(i)} - a_g^{(i-1)}|/|a_g^{(i)}| < 0.001$.

7.3.1 One iteration vs many iterations

As part of our comprehensive evaluation of the AMAT algorithm's performance in direct exoplanet detection, we delve into the impact of iteration counts on its results. Iterative processes are integral to the algorithm, making it essential to investigate how its outcomes evolve over successive iterations. To this end, we consider two distinct trajectory scenarios—one situated within the pixels of a planet and the other within the same annulus but positioned where no planet is present. Our analysis explores

the changes in the flux estimates (a_g) during iterations, shedding light on the algorithm's behavior as it refines its planet identification. In the subsequent sections, we examine and contrast the performance of AMAT under different conditions, including locations with and without planets, and analyze the influence of iteration counts. Additionally, we compare the results obtained using ℓ_1 and ℓ_2 norms, offering insights into the trade-offs between computational complexity and signal-to-noise ratio (SNR) achieved by these estimators.

In order to see the performance, we choose two different trajectories: one is located in the pixels of the planet, and the other is in the same annulus and located in the pixels without a planet. Figure 7.9 shows how the flux a_g changes in each iteration. The results for both norms show that there is a considerable amount of change in the flux a_g for the P_g in the planet pixels, whereas the change in the flux a_g for the P_g in the pixels without a planet is very small.

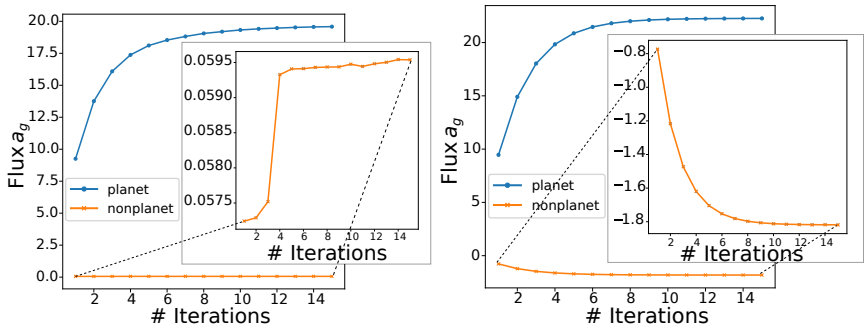

Fig. 7.9 Intensity a_g of the planet against the number of iterations. Left: The intensity $a_g^{(i)}$ is obtained using ℓ_1 norm. Right: The intensity $a_g^{(i)}$ is obtained using ℓ_2 (Frobenius) norm. The blue plots show how the intensity changes in each iteration when we choose P_g in the location of the planet. The orange plots show how the intensity changes in each iteration when we choose P_g in a location without a planet.

In Fig. (7.10a, 7.11a) and Fig. (7.10b, 7.11b), we examine the performance of AMAT when testing a location without a planet. We see that the estimated flux is of very small magnitude, and with additional iterations, this estimation does not change by much.

In contrast, when AMAT is tested on a location with a planet in Fig (7.10c, 7.11c) and Fig (7.10d, 7.11d), we see that it detects the flux much

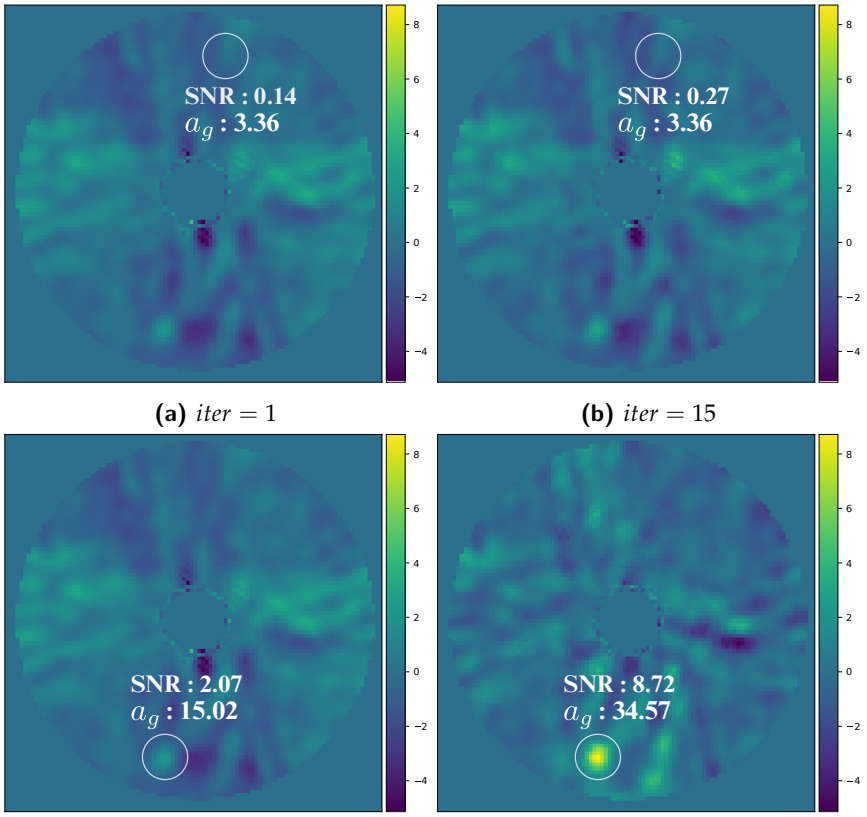

(a) $iter = 1$ **(b)** $iter = 15$

(c) $iter = 1$ **(d)** $iter = 15$

Fig. 7.10 The SNR maps using ℓ_1 norm. Here, we apply SNR maps to the interim frame obtained by using only one P_g to show the performance of applying the algorithm with multiple iterations as a visual, but normally, we apply SNR maps to the flux map obtained by combining all fluxes a_g. The white circle represents the location of P_g. Top: P_g is located in the pixels without a planet. Bottom: P_g is located in the planet pixels.

more strongly, which gets more pronounced after the additional iterations.

When comparing the results for the ℓ_1 and ℓ_2 norms in Fig. 7.10 and 7.11, we see that, in general, the ℓ_1 estimator is able to produce a much better SNR at the price of solving a more challenging non-smooth optimization problem.

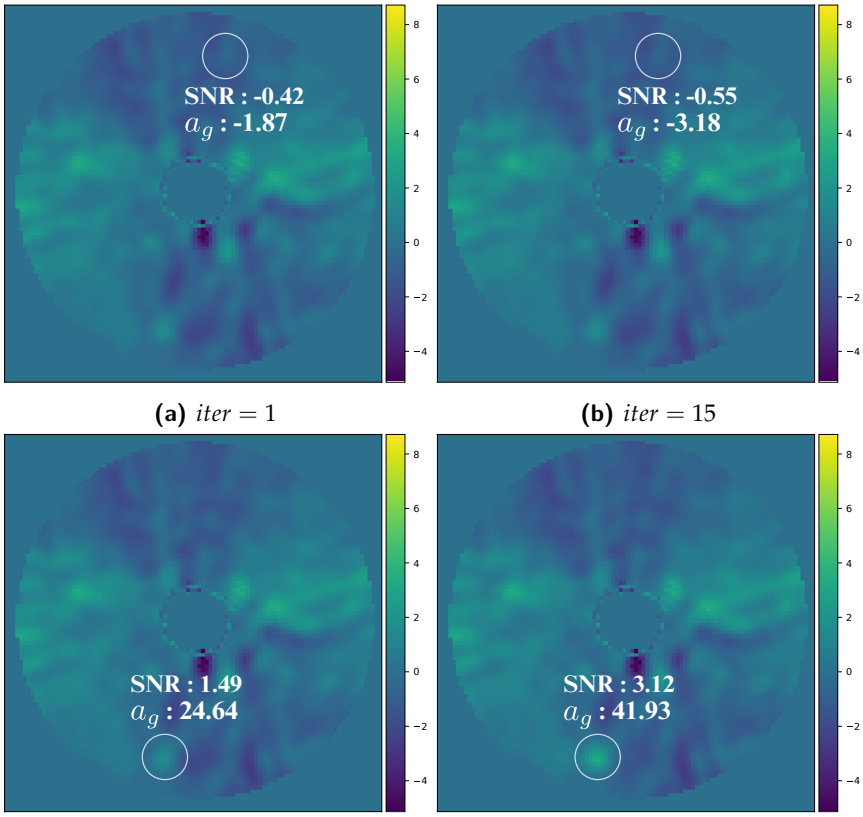

(a) $iter = 1$ **(b)** $iter = 15$

(c) $iter = 1$ **(d)** $iter = 15$

Fig. 7.11 The SNR maps using Frobenius norm. Here, we apply SNR maps to the interim frame obtained by using only one P_g to show the performance of applying the algorithm with multiple iterations as a visual, but normally, we apply SNR maps to the flux map obtained by combining all fluxes a_g. The white circle represents the location of P_g. Top: P_g is located in the pixels without a planet. Bottom: P_g is located in the planet pixels.

7.3.2 Full vs annular AMAT

The recommendation of annular AMAT algorithms for direct exoplanet detection was motivated by the idea of annular methods, which have been validated by significant research in exoplanet imaging. The advantage of annular PCA and the successful results of LLSG obtained by dividing the

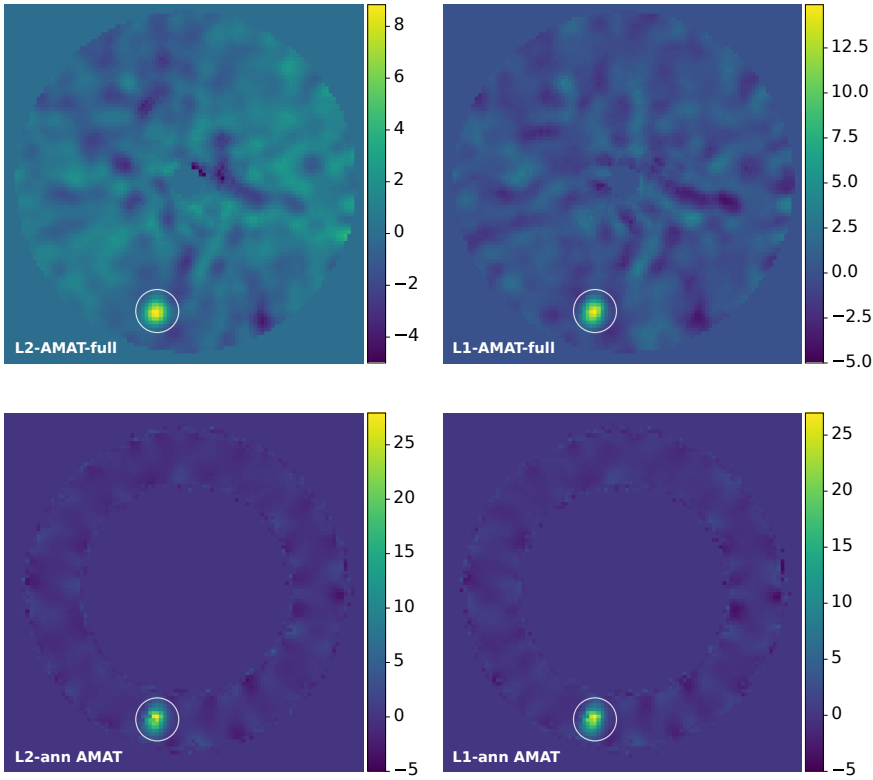

Fig. 7.12 SNR map after applying full and annular AMAT using both norms. In these maps, P_g is located in the planet pixels.

image into annulus allows for a more localized and subtle noise reduction. This technique can potentially increase the detection rates of exoplanets closer to their stars by obtaining a low-rank matrix that adapts to the changing characteristics of the noise, which increases closer to the star.

In Fig. 7.12, we compare the SNR maps using annular and full versions of the AMAT algorithm with two different norms. In these experiments, instead of applying the AMAT algorithm to all pixels, we select the location of the planet and applied the algorithm only to this pixel. There is a significant difference between the annular and full versions of both norms. So, the annular version not only helps us to decrease computation time because it obtains a low-rank matrix only for selected annulus but also has an important effect on SNR visuals.

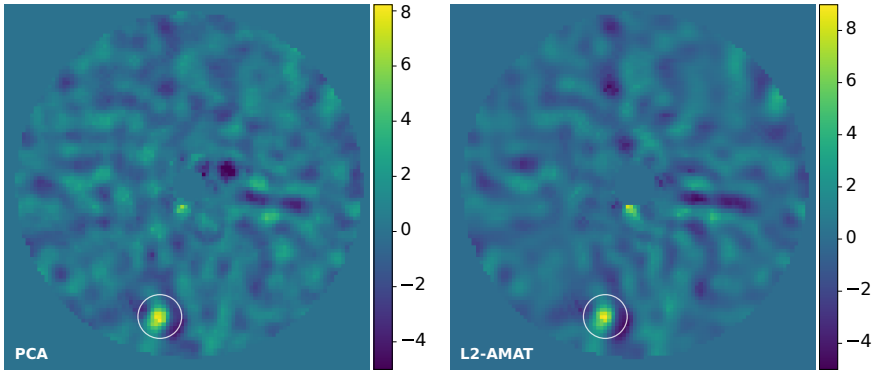

Fig. 7.13 SNR map after applying AMAT with L2 norm and PCA.

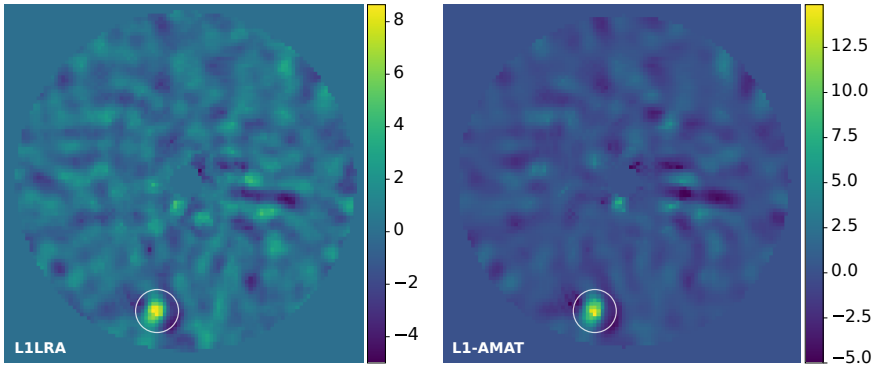

Fig. 7.14 SNR map after applying AMAT with L1 norm and L1LRA.

7.3.3 AMAT vs PCA/L1LRA

In our comparative analysis, we delve into a comprehensive examination of two techniques for direct exoplanet detection: classical annular PCA (AnnPCA) and annular AMAT. To evaluate their performance, we employ L2 and L1 norms, enabling us to distinguish their effectiveness in noise reduction and signal enhancement.

When applying AnnPCA and L2-norm Annular AMAT, we observe similar results given in Fig. 7.13. Both methods demonstrate remarkable similarities in their ability to elevate the SNR in the detection of exoplanets. Yet, L2-norm Annular AMAT exhibits a slight advantage over AnnPCA.

Under the L1 norm, our assessment involves a comparison of Annu-

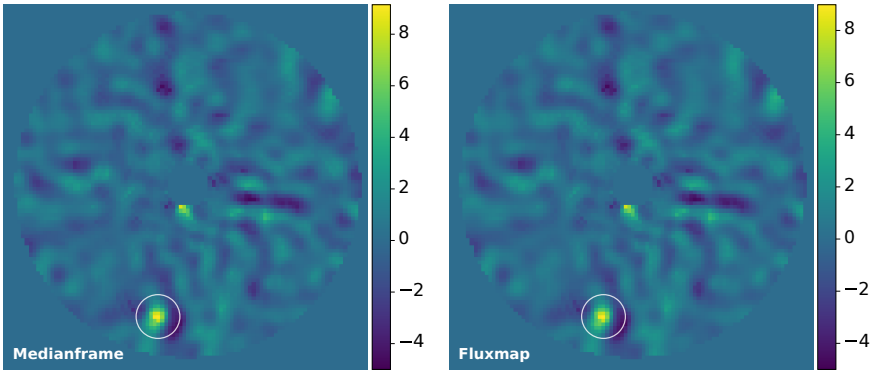

Fig. 7.15 SNR maps of the median frame and fluxmap after applying AMAT with L2 norm.

lar L1LRA, which has a similar idea to AnnPCA except for applying the L1 norm to obtain low-rank approximation as described in Chapter 6, and Annular AMAT using SNR maps represented in Fig. 7.14. While the L1LRA shares significant similarities with L2-norm Annular AMAT, the standout performer in this comparison is L1-norm Annular AMAT. This distinction underscores the unique strengths of L1-norm AMAT in optimizing the SNR in the detection of exoplanets.

7.3.4 Fluxmap vs processed median frame

In our comprehensive evaluation of SNR maps of the AMAT algorithm, we aim to distinguish the effectiveness of two different approaches. The first approach relies on the median frame from the composite cube, while the second approach leverages intensity values a_g derived from the results of possible trajectories. We visually compare these two strategies and evaluate their performance using both L2 and L1 norms.

When examining the results of our SNR maps under the L2 norm in Fig. 7.15, we observed a striking similarity between the median frame and fluxmap. Both approaches delivered SNR maps that closely resembled each other. While these results underscore the capability of both methods, the L2 norm does not significantly differentiate between them.

In contrast, our investigation under the L1 norm revealed interesting differences between the median frame and fluxmap approaches in Fig. 7.16. When employing the L1 norm, the fluxmap, combining intensity values, exhibited a notably higher SNR compared to the median frame.

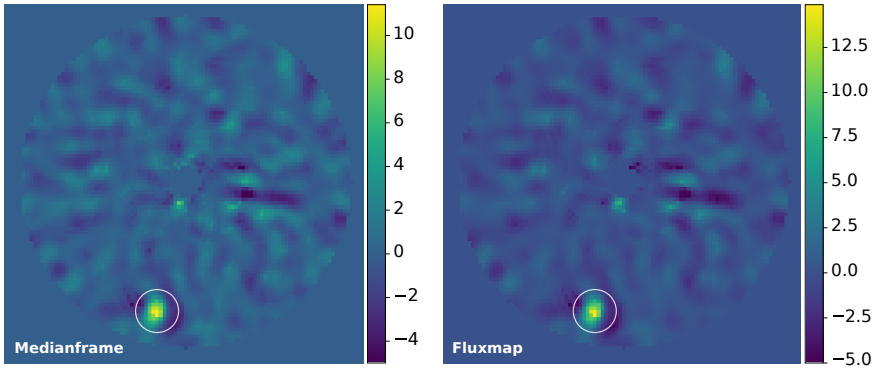

Fig. 7.16 SNR maps of the median frame and fluxmap after applying AMAT with L1 norm.

Here, while we would normally expect the L1 norm to be close to the result obtained with the median due to the ability of L1 norm to reduce outliers with the median, on the contrary, we observed a significant difference.

7.3.5 Contrast curves

A more comprehensive way to assess the sensitivity of various algorithms is to build contrast curves [MMW+14]. Contrast curves are a common tool in the comparison of direct imaging algorithms for exoplanet detection. They are used to determine the level of detectability of planets in relation to their host stars. Essentially, a contrast curve illustrates the minimum detectable contrast ratio between a star and its potential orbiting planets at various angular separations. To do so, we relied on the VIP package. In Fig. 7.17, we compare PCA, AnnPCA, and AnnL1-LRA with our AMAT algorithm using both norms. To generate the AnnPCA contrast curves, we use VIP with its default values except for the rank. We use the same intensities, which are used to obtain the contrast curve of AnnPCA, for the injected planets. Because the different ranks might yield varying results, we applied each algorithm at various ranks, increasing from 5 to 30 in steps of 5. For each algorithm, we selected the deepest contrast curve to represent the best outcome. The optimal ranks were found to be 25 for PCA, 15 for AnnPCA and AMAT_{L2}, 20 for AnnL1-LRA and AMAT_{L1}. The intensities of the injected planets used to build the contrast curve for each algorithm are based on the noise values in each annulus obtained when

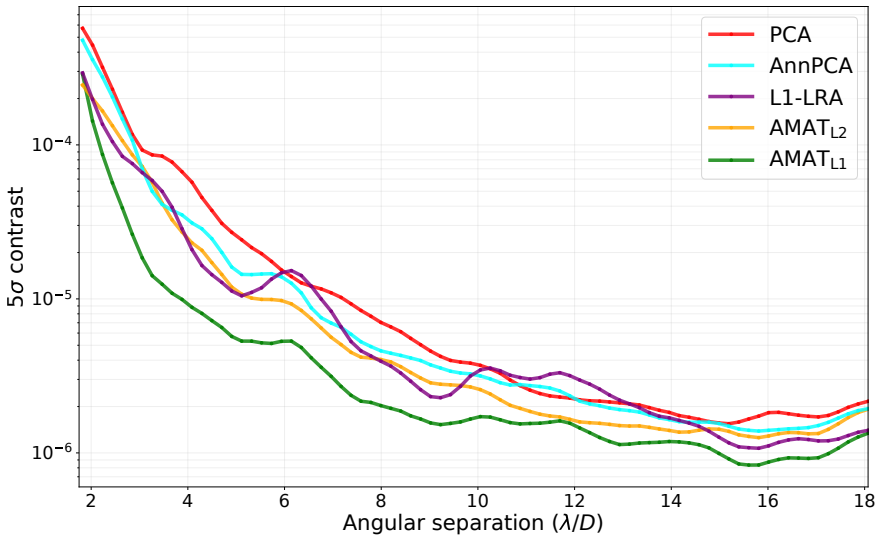

Fig. 7.17 Contrast curves for PCA, AnnPCA and AnnL1-LRA, AMAT$_{L2}$, and AMAT$_{L1}$.

applying AnnPCA with default values other than rank. This ensured that the contrast curves were obtained using the dataset with the same injected planets. As can be observed in Fig. 7.17, the performance using PCA tends to be the worst among the compared methods. It is followed by AnnPCA, which shows slightly better performance. AMAT$_{L2}$ outperforms AnnPCA by a small margin. The performance of AnnL1-LRA fluctuates based on the separation, indicating a performance where it sometimes outperforms or underperforms compared to the other methods. Finally, the AMAT$_{L1}$ algorithm demonstrates increased efficacy, outperforming other algorithms across all separations.

7.3.6 ROC curves

While S/N and contrast curves are useful to illustrate the gain provided by AMAT, they do not explore the behaviour of the algorithms as a function of the detection threshold, which can be done through ROC curves. Building ROC curves relies on injecting a large quantity of synthetic planets into the chosen data set. Our process begins with the removal of the real planet present in the dataset using the VIP package [GG$^+$17, CGF$^+$23]. Subsequently, we inject synthetic planets with an intensity of 1.5 times the

standard deviation of the values in the cube at a distance of $2\lambda/D$ from the star. The injections are placed methodically, starting from 0 to 360 degrees in increments of 3.6 degrees, effectively placing a synthetic planet per scenario. This approach results in a total of 100 different cases for evaluation, effectively covering the entire 360-degree span around the star. For each scenario, we apply the algorithms and then examine the location where the planet was injected. A detection within a specified aperture exceeding a predefined threshold is counted as a true positive (TP); then, we check the other apertures for the presence of the signal. If a signal above the threshold is found within these apertures, it is classified as a false positive (FP); the absence of such a signal results in a true negative (TN) classification. This examination across all apertures facilitates the construction of a ROC curve. Given the critical importance of maintaining a low number of FPs in exoplanet detection, our analysis primarily focuses on achieving a high true positive rate (TPR) without incurring false positives. To better visualize and compare ROC curves, especially to highlight the performance at minimal false positive rate (FPR), we employ a transformed plot of the square root of TPR versus FPR, allowing for an enhanced representation of algorithmic efficiency.

We compare our AMAT algorithms with the results of the PCA, AnnPCA, and AnnL1-LRA algorithms. We use the ranks where we get the best contrast curves for each algorithm. The results displayed in Fig. 7.18 show that our method consistently outperforms the results of PCA, AnnPCA, and AnnL1-LRA in terms of ROC curves. Moreover, similar to the findings for the S/N map, the results obtained using the L1 norm are better than those obtained with the L2 norm. Furthermore, AnnL1-LRA shows slightly better performance compared to both AnnPCA and PCA.

7.3.7 Exoplanet data challenge dataset

As part of our evaluation process, we utilize the datasets from the Exoplanet Imaging Data Challenge (EIDC) [C+20], as specified in Table 7.1 to provide varying datasets for comparing and assessing our algorithm alongside state-of-the-art HCI algorithms. This challenge encompasses a diverse array of ADI sequences, including nine ADI datasets, with a total of 20 injected planet signals. These signals exhibit varying contrasts and positional coordinates to create an evaluation from three different instruments: VLT/SPHERE-IRDIS, Keck/NIRC2, and LBT/LMIRCam each

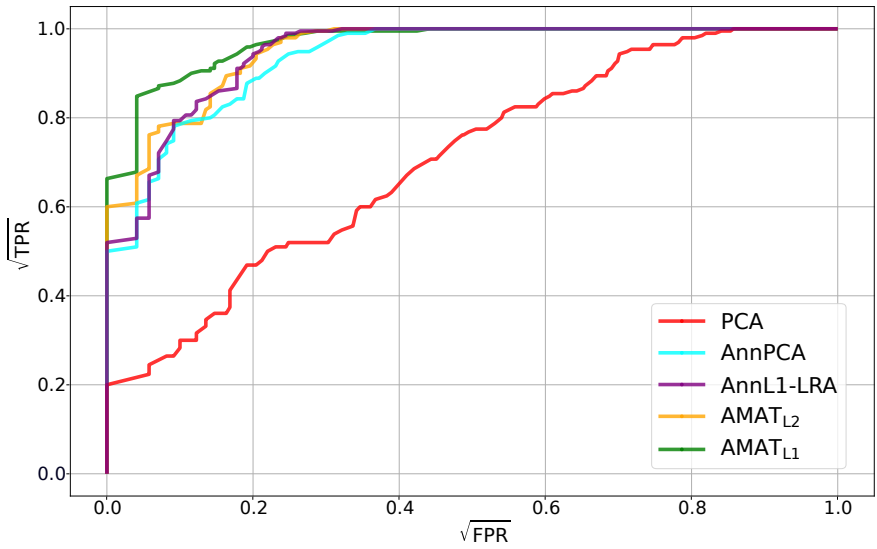

Fig. 7.18 ROC curve of S/N maps. We compare PCA, AnnPCA and AnnL1-LRA, AMAT_{L2}, and AMAT_{L1}. We use the square root to scale the axes in order to better see the low FPR regime.

providing three datasets.

As with the EIDC, our assessment involves the generation of detection maps by each algorithm for every ADI sequence. In the result report of EIDC, a set of standard metrics was employed to compare the performance of these detection maps. Therefore, to apply a common metric for comparing the EIDC algorithms, we calculate F1-scores using the same definition

$$\text{F1-score} = \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}}. \tag{7.9}$$

To determine the values of TP, FP, and FN, we need to decide on a threshold value, as this can significantly impact performance. To determine the appropriate threshold, we inject synthetic planets into the 51 Eri dataset. We create three distinct datasets illustrated in Fig. 7.19, each injected with three, two, and four planets respectively, placed at different locations, and subsequently run our algorithm. We experiment with varying

Table 7.1 The properties of EIDC datasets. N_{img} is the size of the images, N_t is the number of frames, Δ_{field} is the total field rotation of the planets, and Inj is the number of the injected planet.

ID	Telescope/Instr	N_{img} (px × px)	N_t	Δ_{field} (°)	Inj
sph1	VLT/SPHERE-IRDIS	160×160	252	40.3	1
sph2	VLT/SPHERE-IRDIS	160×160	80	31.5	0
sph3	VLT/SPHERE-IRDIS	160×160	228	80.5	5
nrc1	Keck/NIRC2	321×321	29	53.0	4
nrc2	Keck/NIRC2	321×321	40	37.3	3
nrc3	Keck/NIRC2	321×321	50	166.9	0
lmr1	LBT/LMIRCam	200×200	4838	153.4	2
lmr2	LBT/LMIRCam	200×200	3219	60.6	2
lmr3	LBT/LMIRCam	200×200	4620	91.0	3

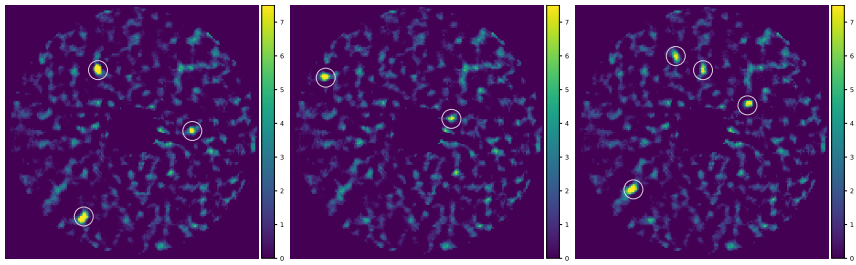

Fig. 7.19 S/N maps for 51 Eri datasets with injected planets: In these S/N maps, white circles represent TP.

thresholds to identify the threshold yielding the highest F1-score. Fig. 7.20 illustrates the plot of F1-score against threshold for the 51 Eri dataset with synthetic planets. Based on this analysis, if we select a threshold between 6.6-7.7 for S/N maps after $AMAT_{L1}$, we can detect all planets without false positives and this allows us to achieve the highest F1-score. Additionally, we tested the use of a 5σ threshold but observed better performance with our proposed approach.

In the published results report, we have knowledge of the locations of the planets and evaluated various threshold values as evidenced by the F1 scores presented in Fig. 7.20. Since we also established our threshold

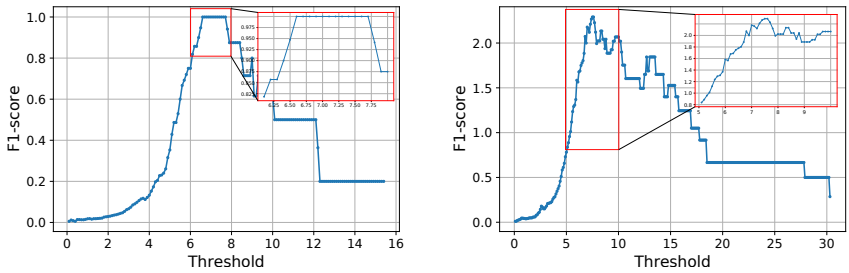

Fig. 7.20 F1-score of $AMAT_{L1}$ algorithm with S/N map using various thresholds for 51 Eri datasets (left) and EIDC datasets (right).

range by testing injections into the 51 Eri dataset, we conducted a posteriori verification to ensure that the thresholds also yield a high F1 score when applied to the EIDC datasets. Participating in the EIDC, where we had multiple opportunities to submit and observe the F1-score, allows us to undergo a similar process. Selecting the value that yields the highest F1-score represents a fair and comparable approach. In our comparison, we used the threshold of 7.5 within the 6.6-7.7 range which we obtain the highest F1-score for the 51 Eri dataset.

The effectiveness of the $AMAT_{L1}$ algorithm is illustrated in Fig. 7.23 through S/N maps. The analysis of the VLT/SPHERE-IRDIS datasets showcases the proficiency of the algorithm, where it successfully identified all six planets, with only one false positive. In the Keck/NIRC2 dataset, the algorithm accurately detected four out of seven planets with three false positives, and in the LBT/LMIRCam dataset five out of seven planets without any false positives. The results of our experiments using AnnPCA on the EIDC data set are shown in Fig. 7.21 for reference, to illustrate the significant gain provided by AMAT in terms of TPR. Additionally, the $AMAT_{L2}$ algorithm exhibits a performance between $AMAT_{L1}$ and AnnPCA in Fig. 7.22. It detects five out of six planets with one false positive in the VLT/SPHERE-IRDIS datasets, three out of seven planets in the Keck/NIRC2 dataset with one false positive, and five out of seven planets with two false positives in the LBT/LMIRCam dataset. We compute F1-scores based on these maps and compare them with the best results of the published results report in Fig. 7.20.

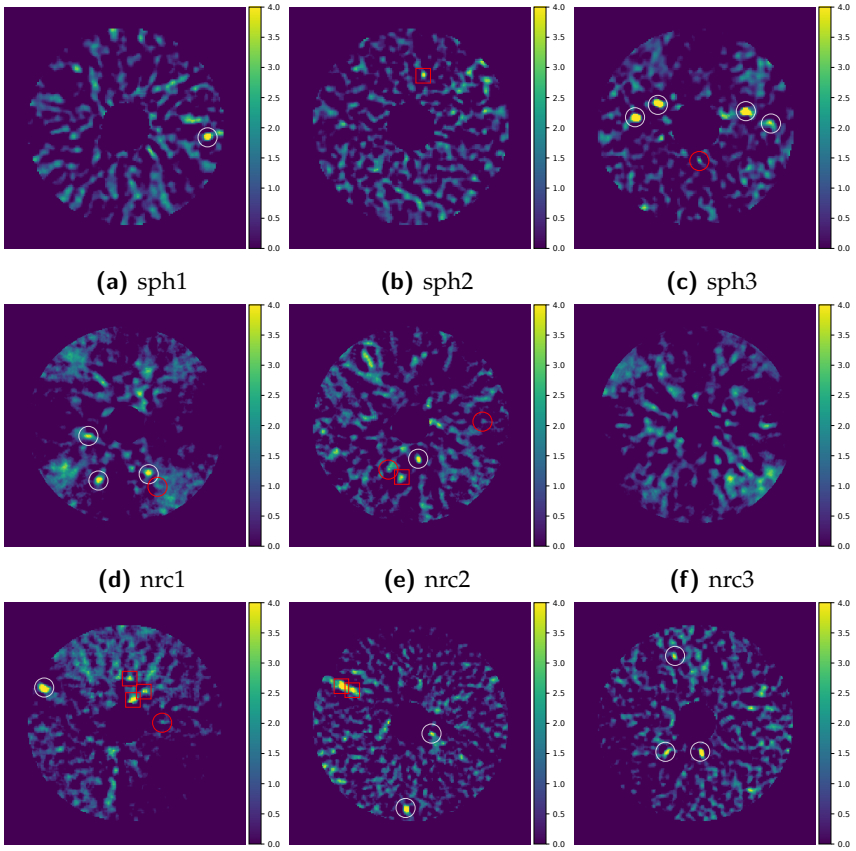

(a) sph1 **(b)** sph2 **(c)** sph3

(d) nrc1 **(e)** nrc2 **(f)** nrc3

(g) lmr1 **(h)** lmr2 **(i)** lmr3

Fig. 7.21 S/N maps after AnnPCA using VIP package: In these S/N maps, white circles represent TP, red squares denote FP, and red circles signify FN.

7.3.8 Flux estimation performance

For the purpose of flux estimation, we conducted a comparative analysis of our AMAT_{L2} algorithm against the NEGFC method implemented in VIP, where the companion parameters are estimated either through a Nelder Mead minimization method or through a Markov Chain Monte Carlo (MCMC) method [WAG+17]. Some other forward modeling or inverse problem approaches [CMM+15, RMW+17] also have the potential to provide accurate flux measurements and may outperform both NegFC and

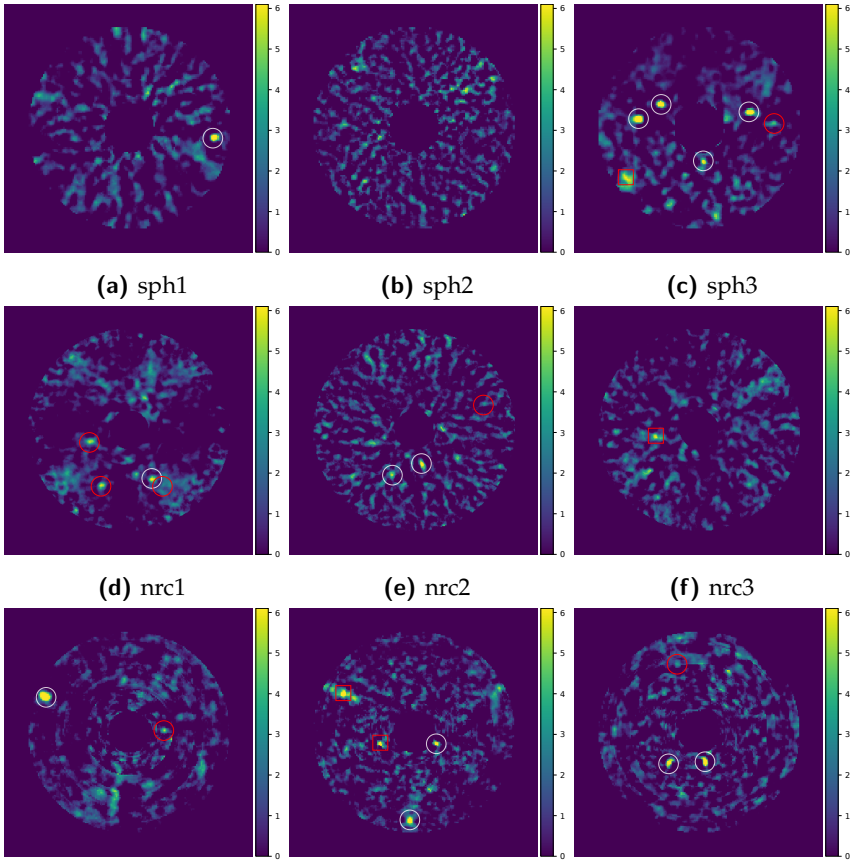

(a) sph1 **(b)** sph2 **(c)** sph3

(d) nrc1 **(e)** nrc2 **(f)** nrc3

(g) lmr1 **(h)** lmr2 **(i)** lmr3

Fig. 7.22 S/N maps after AMAT_{L2}: In these S/N maps, white circles represent TP, red squares denote FP, and red circles signify FN.

AMAT. However, a full comparison of these methods is beyond the scope of this study. We designed two distinct sets of comparisons, one with the synthetic planets injected at a small separation ($4\lambda/D$) and the other at a large separation ($10\lambda/D$). In both scenarios, we created various cases by rotating the position of the injected planet from 0 to 350 degrees in increments of 10 degrees. This resulted in 36 different cases for each scenario. This systematic approach was applied to both the $4\lambda/D$ and the $10\lambda/D$ separations, thereby ensuring a comprehensive evaluation. Planets were injected at specified radii and angles without consideration for whether

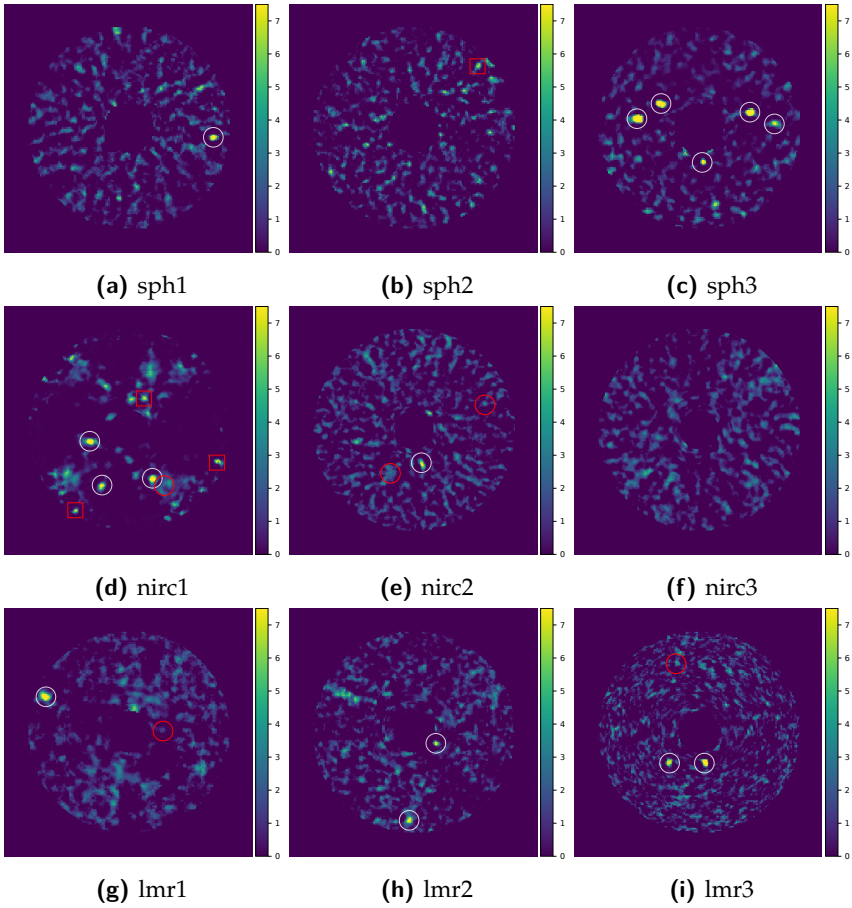

(a) sph1 **(b)** sph2 **(c)** sph3

(d) nirc1 **(e)** nirc2 **(f)** nirc3

(g) lmr1 **(h)** lmr2 **(i)** lmr3

Fig. 7.23 S/N maps after AMAT_{L1}: In these S/N maps, white circles represent TP, red squares denote FP, and red circles signify FN.

the location falls precisely on the center or the edge of the pixel, ensuring consistency across all cases. To decide on the rank for each algorithm, we use the approach suggested in VIP [CGF+23, GG+17], which finds the optimal number of principal components in terms of S/N using PCA. In the AMAT algorithm, after a maximum of 50 iterations, or when the relative change of the intensity $a_g^{(i)}$ is minor, *i.e.*, when $|a_g^{(i)} - a_g^{(i-1)}|/|a_g^{(i)}| < 10^{-3}$, we simply stop the algorithm.

In Fig. 7.24, the AMAT_{L2} algorithm demonstrates excellent accuracy in flux estimations for both small and large separations, as evidenced by the

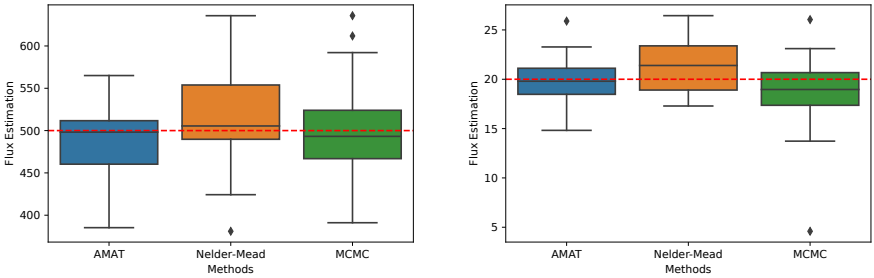

Fig. 7.24 Box plot of flux estimation for planets injected at $4\lambda/D$ (left) and $10\lambda/D$ (right) separations. The dashed red line indicates the injected flux value, while the black lines inside the boxes represent the median values. The boxes encompass the range between the first quartile (Q1) and the third quartile (Q3) of the results. The horizontal lines extending from the boxes correspond to Q1−1.5IQR and Q3+1.5IQR, where IQR = Q3 - Q1. Diamonds indicate outliers beyond these ranges.

median line of the boxplots precisely aligning with the injected intensities. Contrastingly, the boxplots for the Nelder-Mead and MCMC methods exhibit the presence of outliers, indicating lower precision in their flux estimations. Specifically, the Nelder-Mead method tends to produce higher flux values, often deviating significantly from the injected values, while the MCMC method frequently results in lower flux estimations.

7.3.9 Improving detection performance with likelihood ratio maps

In [DVC+22], we presented an enhanced approach for exoplanet detection utilizing likelihood ratio maps (LRM), which offers improvements over traditional S/N maps. The LRM is derived through a maximum likelihood estimation, employing Laplacian distributions. It consists of the ratio of the maximum likelihood $\mathcal{L}_g(\hat{a}_g|R)$ to the likelihood of the null hypothesis $\mathcal{L}_g(0|R)$, which corresponds to the absence of a planet

$$\begin{aligned} \log \Lambda_g(R) &= \log\left(\frac{\mathcal{L}_g(\hat{a}_g|R)}{\mathcal{L}_g(0|R)}\right) \\ &= -\sum_{(\theta,r)\in\Omega_g} \frac{|R(\theta,r) - \hat{a}_g P_g(\theta,r)| - |R(\theta,r)|}{\sigma_R(r)} \end{aligned} \tag{7.10}$$

where R is the residual cube and $\hat{a}_g$ is the estimated flux by solving the

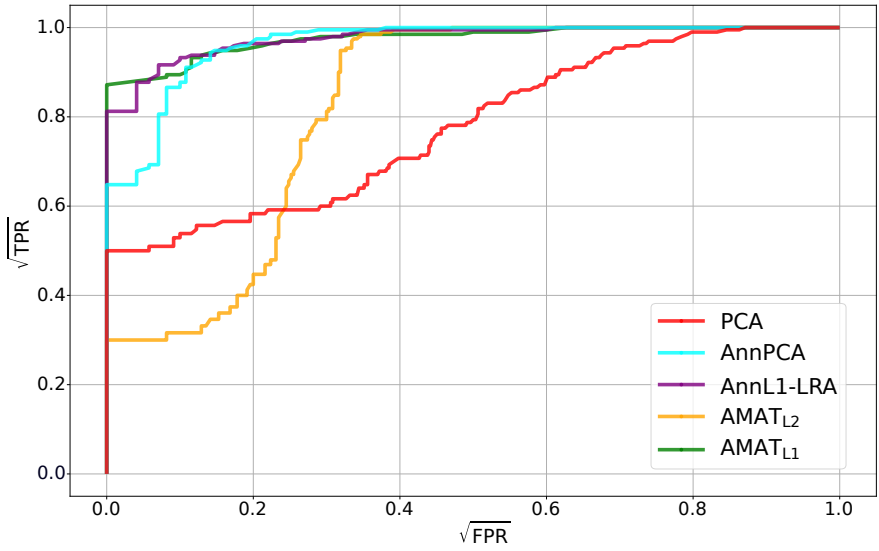

Fig. 7.25 ROC curve of LRMs. We compare PCA, AnnPCA, AnnL1-LRA, AMAT$_{L2}$, and AMAT$_{L1}$.

following optimization problem

$$\begin{aligned}\hat{a}_g &= \text{argmax}_a \log \mathcal{L}_g(a|R) \\ &= \text{argmin}_a \sum_{(\theta,r)\in\Omega_g} \frac{|R(\theta,r) - aP_g(\theta,r)|}{\sigma_R(r)}.\end{aligned} \tag{7.11}$$

In light of this information, we have combined the AMAT algorithm with the LRM. After obtaining the residual cube, we apply the LRM algorithm to our results. The steps for generating the residual cube with the AMAT algorithm and applying LRM are incorporated in the AMAT python package, which we developed based on the model outlined in this study [1]

We obtained the ROC curve using the residual cubes presented in Section 7.3.6 and applied the LRM instead of the S/N map in Fig. 7.25. Each algorithm, except AMAT$_{L2}$, demonstrated higher performance compared to the S/N map results in Fig. 7.18. We believe this discrepancy arises

[1]The AMAT algorithm python package is available on GitHub: https://github.com/hazandaglayan/AMAT.

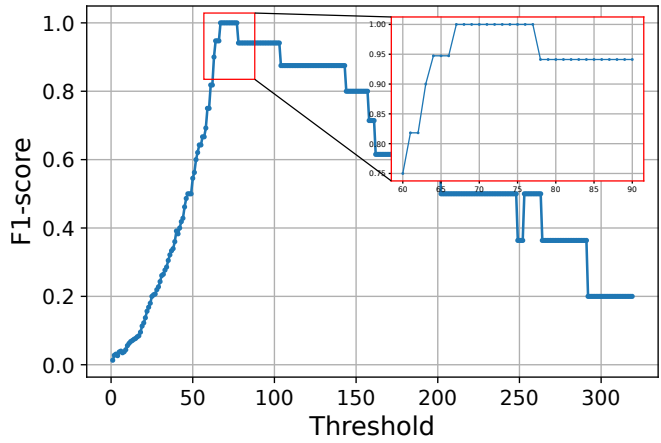

Fig. 7.26 F1-score of AMAT_{L1} algorithm with LRM using various thresholds for 51 Eri datasets.

from AMAT_{L2} being based on a Gaussian distribution, while LRM relies on a Laplacian distribution. Among the algorithms, AMAT_{L1} showed the best performance, while AnnL1-LRA also exhibited strong results. Both algorithms are Laplacian-based, which likely contributes to their superior performance.

To decide on the threshold, we used the same method in Fig. 7.26-7.27 as described in Section 7.3.6 for finding the threshold of S/N maps. Based on this analysis, if we choose a value between 67-77 for the LRMs after AMAT_{L1}, we get the highest F1 score for the Eri 51 dataset, and as for the case of AMAT-S/N, we check that the threshold for LRMs after AMAT_{L1} is indeed in this range to maximize the F1 score on the EIDC.

Figure 7.28 shows that the version AMAT_{L1} using LRM proficiently identifies all exoplanets within the VLT/SPHERE-IRDIS datasets and accurately detects four out of seven planets, and six out of seven planets in the LBT/LMIRCam datasets without any false positive in any dataset. These findings showcase the algorithm's high success rate, especially when compared with other algorithms reported in the EIDC results. Moreover, AMATL1 is the only algorithm capable of detecting both planets in the lmr1 dataset among all algorithms in the EIDC.

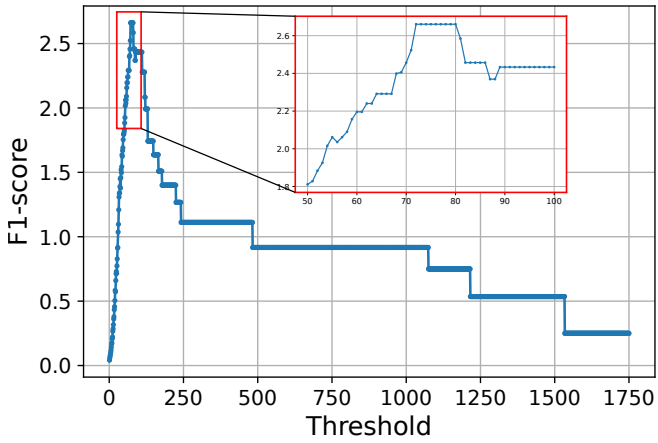

Fig. 7.27 F1-score of AMAT_{L1} algorithm with LRM using various thresholds for EIDC datasets.

In the paper reporting on the EIDC results [C+20], the algorithms are classified into four categories: classical speckle subtraction providing residual maps, advanced speckle subtraction building detection maps, inverse problems, and supervised machine learning. A comparison of the most successful of these methods in their categories, focusing on the $F1_{\text{score}}$, is presented in Fig. 7.29. Our AMAT_{L1} method with both S/N map and LRM comes out as the most effective within the category of classical speckle subtraction methods. Furthermore, the AMAT_{L1} method with LRM outperforms advanced algorithms like RSM and FMMF, demonstrating the highest level of success. This underlines the significant potential and robustness of AMAT_{L1} in the realm of exoplanet detection.

7.4 Conclusion

In this chapter, we investigated AMAT, a new exoplanet detection method designed to improve the separation of planetary flux from static and quasi-static signals using iterative techniques. The method enhances models based on low-rank approximations like PCA. Current approaches typically assume Gaussian noise, but recent studies suggest that residuals often follow a Laplacian distribution. Most low-rank approximation techniques still rely on PCA, which assumes Gaussian noise. To address this inconsistency, we proposed L1-LRA, which is based on assuming Laplace-

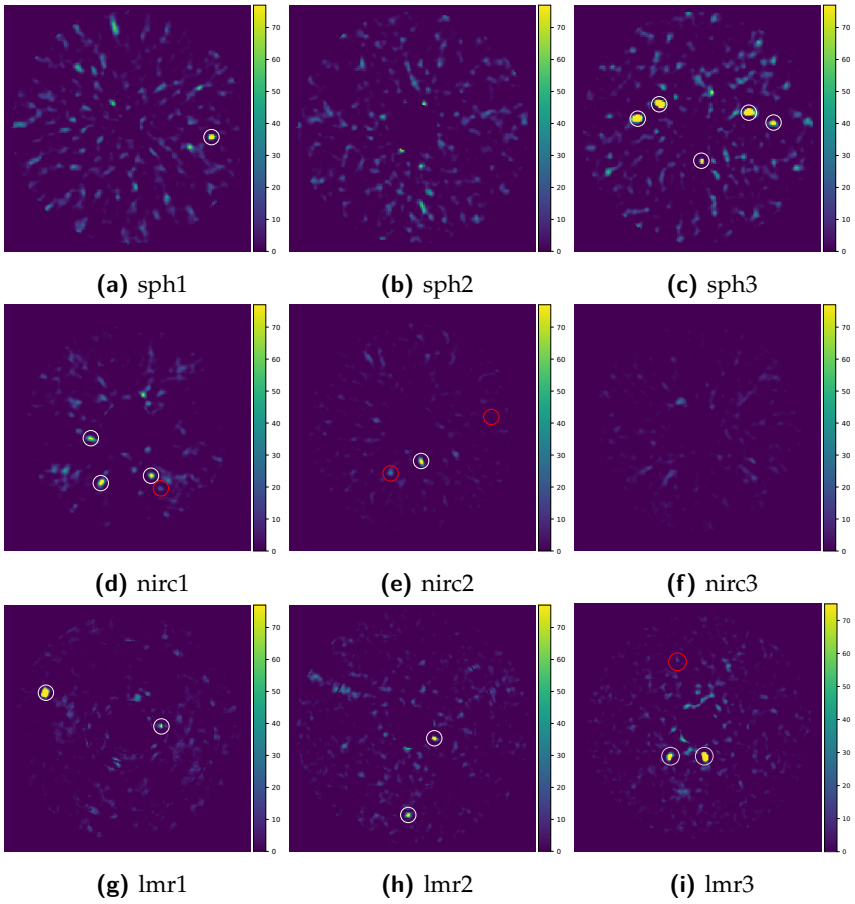

(a) sph1 **(b)** sph2 **(c)** sph3

(d) nirc1 **(e)** nirc2 **(f)** nirc3

(g) lmr1 **(h)** lmr2 **(i)** lmr3

Fig. 7.28 LRMs after AMAT_{L1} for EIDC Datasets: In these maps, white circles represent TP, red square denotes FP, and red circles signify FN.

distributed noise, and integrated it within the AMAT algorithm.

The AMAT algorithm was thoroughly tested with various approaches. First, we compared sensitivity limits using S/N comparisons and contrast curves on the 51 Eri dataset acquired with the VLT/SPHERE-IRDIS instrument. We evaluated the performance of AMAT alongside AnnPCA and AnnL1-LRA, and our results demonstrated significant performance improvements. This was supported by ROC curve comparisons. We then benchmarked AMAT against state-of-the-art algorithms using EIDC datasets. When comparing similar speckle subtraction algorithms, AMAT

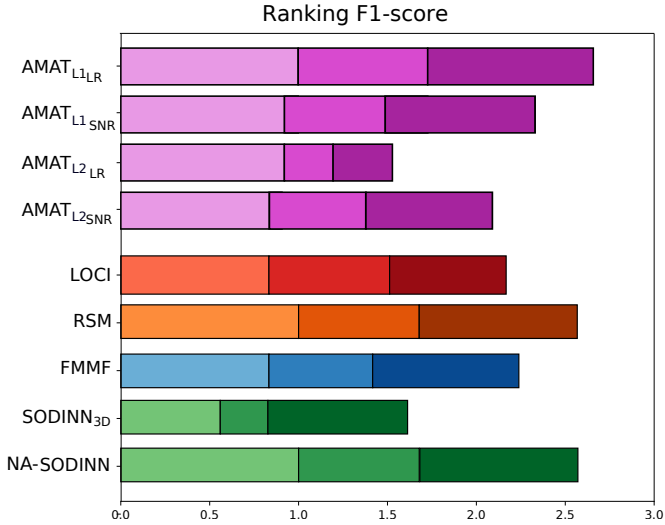

Fig. 7.29 Ranking based on F1-score of the different algorithms in the EIDC results report. The algorithms are classified as AMAT algorithms with S/N maps and LRMs (purple), classical speckle subtraction providing residual maps (red), advanced speckle subtraction building detection maps (orange), inverse problems (blue), and supervised machine learning (green). The light, medium, and dark colors correspond to the three VLT/SPHERE-IRDIS, Keck/NIRC2, and LBT/LMIRCam data sets, respectively.

delivered competitive performance, achieving results comparable to the best in its category. By utilizing the LRM based on Laplacian noise, we further enhanced the algorithm's performance, allowing AMAT$_{\text{L1}}$ with LRM to achieve the highest F1 score among all categories. Additionally, the flux map generated by the AMAT algorithm provided accurate planetary flux measurements, even for faint planets.

One limitation of this study is the high computational cost due to the iterative nature of the algorithm and the need to apply it to each pixel, given the possibility of the planet appearing in any pixel — despite this being less relevant in the specific context of exoplanets. The AMAT$_{\text{L1}}$ algorithm, while producing better results, relies on L1-LRA, which is computationally more demanding than PCA. Future studies should explore faster implementations of L1-LRA to improve the algorithm's speed and applicability. Another limitation involves determining the threshold for the LRM, which

should be standardized rather than observation-based, ensuring a more automated and reliable method.

8

Matrix Completion for Direct Exoplanet Detection

Low-rank plus trajectory models such as AMAT in Chapter 7 face the difficulty that the low-rank term, meant to collect the noise, may also capture some of the planet light. This may result in a less accurate planet flux estimation and a lower planet detection performance. A rather radical remedy, that we pursue in this chapter, is to hold back from the low-rank estimation method all the pixels that a postulated planet significantly affects. This approach leads us into the realm of matrix completion (MC).

We begin by introducing the matrix completion algorithm and then describe its application in exoplanet detection. To assess the effectiveness of this algorithm, we compare its performance against classical detection methods.

8.1 General Definition of Matrix Completion

Matrix Completion (MC) is a problem of estimating the missing entries of a matrix $A \in \mathbb{R}^{m \times n}$ by observing a subset $\Omega \subset [m] \times [n]$ of the known entries with $[x] := \{1, \ldots, x\}$. The problem can be expressed using sampling

operator $\mathcal{P}_\Omega$. The sampling operation $\mathcal{P}_\Omega(A)$ of a matrix A is defined as

$$[\mathcal{P}_\Omega(A)]_{ij} = \begin{cases} A_{ij} & (i,j) \in \Omega \\ 0 & \text{otherwise.} \end{cases} \tag{8.1}$$

This problem arises in many different applications including recommender systems [KBV09], phase retrieval [CESV15], localization in IoT networks [NKKS19, KNKS21], traffic monitoring [DCY+14], image inpainting [XZC17], weather data filling [LAJ22].

In matrix completion, a common prior is to suppose that the underlying matrix has low-rank rank $r \ll \min(m,n)$. So the *rank minimization* (RM) problem tries to find a minimum rank matrix X

$$\min_{X \in \mathbb{R}^{m \times n}} rank(X) \tag{8.2}$$

$$\text{subject to } \mathcal{P}_\Omega(X) = \mathcal{P}_\Omega(M). \tag{8.3}$$

It is an NP-hard problem [CR12], however, it is possible to relax the problem into a convex nuclear norm minimization problem introduced by [CR12]

$$\min_{X \in \mathbb{R}^{m \times n}} \|X\|_*$$
$$\text{subject to } \mathcal{P}_\Omega(X) = \mathcal{P}_\Omega(M)$$

where $\|X\|_*$ is the sum of the singular values of X. This method provides a practical way to approach the rank minimization problem in real-world scenarios.

[CR12] uses semidefinite programming to solve this problem. Later, [CCS08] improved the performance of this approach by providing a specialized *singular values soft thresholding* algorithm. Returning to the RM problem, [MJD09] developed an approximate solution based on singular value *hard* thresholding, giving results similar to its soft thresholding counterpart while being more robust to noise.

Both soft and hard singular value thresholding rely on computing the SVD of large matrices, which is typically quite slow. If the actual rank r is known (or if a good estimate can be found), it is possible to avoid the SVD by replacing the matrix $X \in \mathbb{R}^{m \times n}$ with a product of two smaller matrices $W \in \mathbb{R}^{m \times r}$ and $H \in \mathbb{R}^{r \times n}$, such that $X = WH$. Then, rather than the rank of X, one minimizes the ℓ_2-norm discrepancy between A and WH over the

observed entries:

$$\min_{W \in \mathbb{R}^{m \times r}, H \in \mathbb{R}^{r \times n}} \|\mathcal{P}(A) - \mathcal{P}(WH)\|_2^2. \tag{8.4}$$

To solve this problem, many different approaches have been proposed for years. These suggestions are categorized in 4 approaches [NKS19]: heuristic greedy algorithms [LB10, WLL+14], alternating minimization techniques [TW16, WYZ12], optimization over smooth Riemannian manifold [Van13, BA15], and truncated nuclear norm minimization [HZY+12].

8.2 Matrix Completion for Exoplanet Detection

Data loss or missing observations are common issues, but we use matrix completion in this case not to deal with data loss, but to construct a low-rank matrix that represents the background (without the exoplanet) by removing the exoplanet's signal from the data. The reason for this is that the background data generated through methods like PCA still contains traces of the exoplanet signal. By removing the planetary signals from the original data before calculating the background, we ensure that the background data is free of planetary signals, allowing the planet's signal to stand out more clearly in the residual data.

In matrix completion, we assume the presence of a planet at a given location. Based on this assumption, we remove the pixels along the planet's trajectory and then use matrix completion to fill in the removed pixels. The difference between the original data cube and the filled data cube gives us the planet's flux in the residual cube. In Fig. 8.1, you can see an illustration of the algorithm.

In real cases, since the planet's location is unknown, we apply this algorithm across nearly all pixels in the frame. We skip the pixels where the star is located and those in the corners (since we rotate the frames in the final step). The algorithm is applied to the annulus region surrounding the star.

Let M be an ADI cube consisting of T images, each with $N \times N$ pixels. Let $\mathcal{M}$ be the chosen matrix completion method, where $\mathcal{M} : X \in \mathbb{R}^{T \times N^2}, \Omega \subset [T] \times [N^2] \to \mathcal{M}(X, \Omega; \theta) \in \mathbb{R}^{T \times N^2}$, with X as the input matrix, Ω as the set of matrix indices that are not removed, and θ representing parameters for $\mathcal{M}$. The planet's trajectory in the t-th frame is denoted by $g_t \in \mathbb{N}^2$.

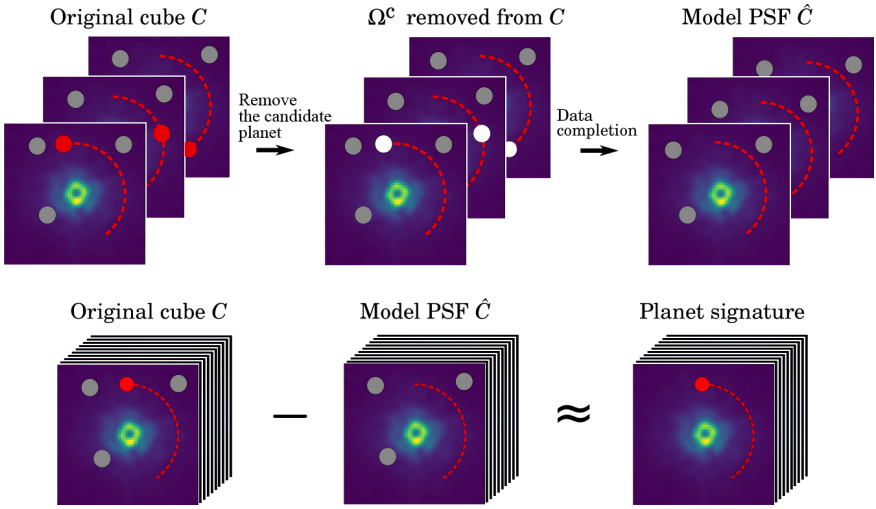

Fig. 8.1 Illustration of the matrix completion-based PSF estimation technique. The dashed lines represent the trajectory g. We assume the presence of a planet at a specific location Ω^c represented by the white disk. We remove Ω^c from the original cube M. Then we fill the data using a matrix completion method. We subtract the model PSF $\hat{M}$ from the original cube M in order to get the planet signature.

In the first step, we create the set of pixel indices Ω_g that are not affected by the planet, centered on the trajectory g_t. This set, Ω_g, is defined as:

$$\Omega_g = \left\{ (t,r) \in [T] \times [N]^2 \;\middle|\; \|r - g_t\|_2 > \frac{1}{2}\rho_{ds}\frac{\lambda}{D} \right\}. \tag{8.5}$$

where D is the telescope diameter, λ is the observed wavelength, and ρ_{ds} is the disk size of the region impacted by the planet. The diameter of the disk removed along g is proportional to λ/D.

To apply matrix completion, we first transform the 3D data cube into a $T \times N^2$ matrix. The indices in Ω_g^c (which correspond to the removed pixels) are filled using low-rank matrix completion, and the matrix is then converted back into a 3D data cube.

In our experiments, we use the LMaFit algorithm, which solves a low-rank matrix fitting problem through nonlinear successive over-relaxation [WYZ12]. It is solved iteratively using linear least squares, ini-

tialized with singular value decomposition (SVD) as follows:

$$W_0 = US^{1/2}$$
$$H_0 = S^{1/2}V^T$$

Then, we solve the following at each iteration:

$$W_{i+1} = \arg\min_W \|WH_i - X_i\|_F^2 = X_i H_i^\dagger$$
$$H_{i+1} = \arg\min_H \|W_{i+1}H - X_i\|_F^2 = W_{i+1}^\dagger X_i$$
$$X_{i+1} = W_{i+1}H_{i+1} + \mathcal{P}(X^0 - W_{i+1}H_{i+1})$$

where $X^\dagger$ is the Moore-Penrose pseudoinverse of X.

Figure 8.2 shows the results of applying matrix completion and PCA to produce a model PSF. The matrix completion model is computed along the trajectory g of a known companion, and the residual R_g clearly reveals the companion. In comparison, the residuals from PCA are less clear. Figure 8.3 shows a similar example for a different companion, where matrix completion provides clearer results than PCA.

8.2.1 Annular matrix completion

We observe that the pixels within the same annulus tend to have similar values, while the values differ significantly between different annuli. This variation occurs because pixels closer to the star are brighter, and their intensity decreases as we move away from the star. These differences are illustrated in Fig. 8.4. Therefore, rather than using all the data indiscriminately, we propose filling the missing pixels based on the annulus in which they are located. This way, the removed pixels can be filled using values from pixels at indices Ω'_g within the same annulus

$$\Omega'_g = \Big\{(t,r) \in [T]\times[N]^2 \;\Big|\; \|r - g_t\|_2 > \frac{1}{2}\rho_{ds}\frac{\lambda}{D}, \quad \|g_t\|_2 - \frac{r_{\text{ann}}}{2} \le \|r\| \le \|g_t\|_2 + \frac{r_{\text{ann}}}{2}\Big\} \tag{8.6}$$

where r_{ann} is the annular width. This means Ω'_g is the list of the cube indices of the annulus with width r_{ann} that are not inside of a circle with a distance of a trajectory g. Figure 8.5 illustrates Ω'_g as defined in (8.6).

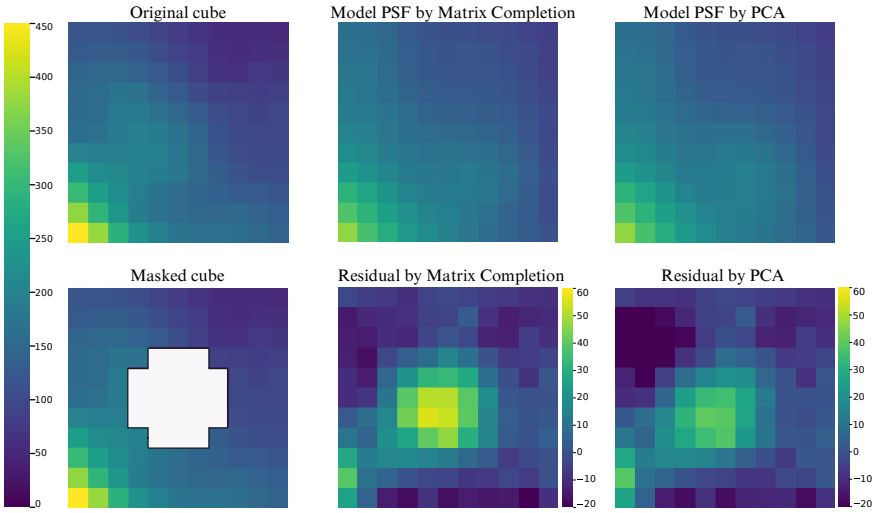

Fig. 8.2 Example of using matrix completion and PCA to produce a model PSF. Each figure shows a zoom over the same area of a specific frame, in different cubes. That area of the frame is known to contain a planet. Top-left: the original cube C. Bottom-left: a representation of the mask as a white area superimposed on the original cube. Top-middle: the model PSF produced with the matrix completion technique. Bottom-middle: the residual produced by matrix completion technique. Top-right: the model PSF produced with PCA. Bottom-right: the residual produced by PCA. The mask has size $\rho_{\text{ds}} = 1$ and is applied over the location of the companion.

So the complementary indices Ω_g are filled with using the annulus with width r_{ann} instead of all the indices in the data cube.

8.2.2 Pipeline of matrix completion algorithm

In classical methods like PCA, a low-rank approximation algorithm is applied to the dataset, and the resulting low-rank part is referred to as the background cube. The residual cube is obtained by subtracting the background cube from the original dataset. This residual cube is then derotated, and the median frame of the derotated cube is selected as the processed frame. Finally, an SNR map algorithm is applied to the median frame.

In the matrix completion method, a trajectory centered at a pixel in the annulus is removed. Then, a low-rank matrix completion algorithm is ap-

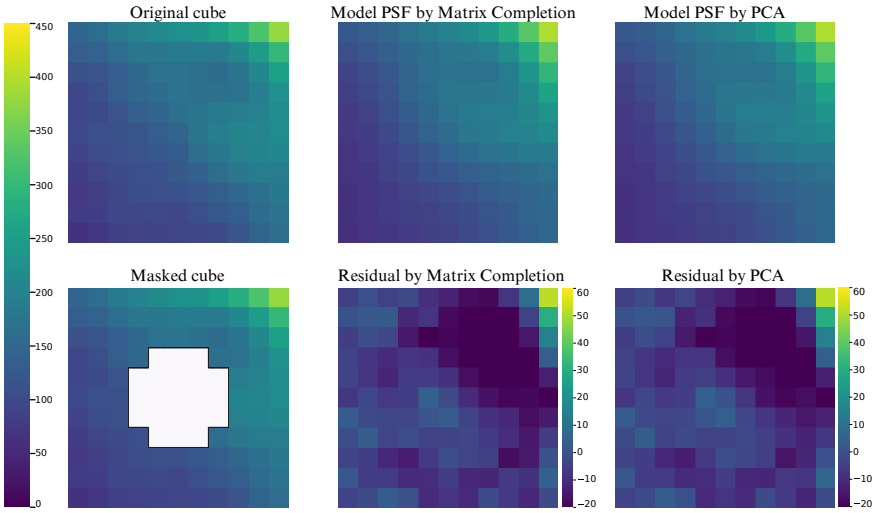

Fig. 8.3 Another example of matrix completion and PCA. The dataset and the techniques are the same as in Fig. 8.2, but another trajectory, which does not include planetary signals, is used.

plied, and the residual cube is obtained by subtracting the completed data cube from the original data cube. This procedure is repeated for each pixel in the annulus, resulting in n residual cubes, where n is the number of pixels in the annulus.

To generate a single final cube from these n residual cubes, all residual cubes are first derotated. Then, the vector from the residual cube representing the time series of values for the pixel at the center of the removed trajectory is copied into the corresponding pixel of the final cube. This process is repeated for all pixels in the annulus, merging the residual cubes into a complete annulus in the final residual cube. This process is depicted in Fig. 8.6. Once the final residual cube is constructed, the median frame and SNR map algorithms are applied.

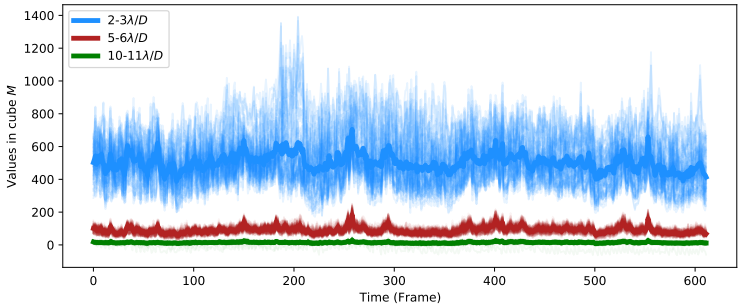

Fig. 8.4 Illustration of the indices of annular for the annular matrix completion. The indices in annulus except for the disk constructs Ω'_g.

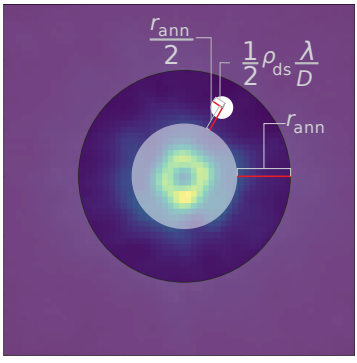

Fig. 8.5 Illustration of the indices of annular for the annular matrix completion. The indices in annulus except for the disk constructs Ω'_g.

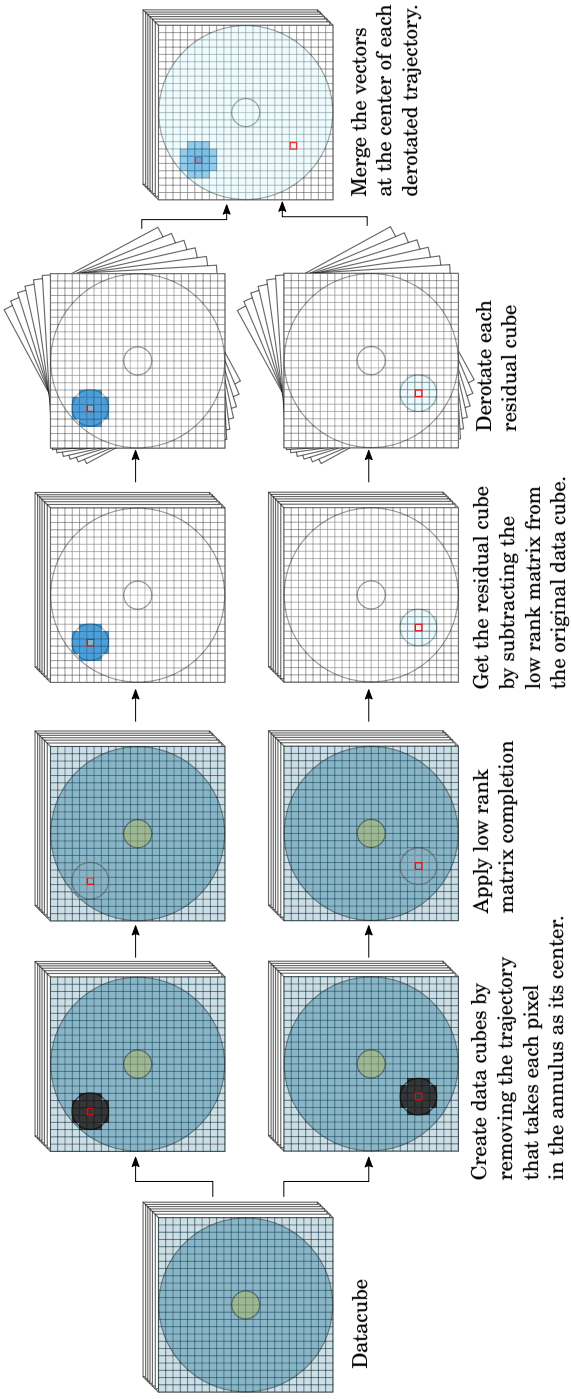

Fig. 8.6 Pipeline of the Matrix Completion Algorithm

8.2.3 Likelihood ratio map of matrix completion algorithm

The likelihood ratio map is derived using the residual cube obtained through the Likelihood Ratio Map algorithm, as described in Chapter 4. Unlike previous approaches, which considered only the final residual cube, we take a different approach here by generating a separate residual cube for each trajectory. This is done by calculating the difference between the original matrix M and the completed matrix $X_g = W_g H_g$, expressed as:

$$R_g = M - X_g. \tag{8.7}$$

The next step is to compute the likelihood ratio for each residual cube. The likelihood is defined as in (4.9):

$$\log \Lambda_g(R_g) = -\sum_{(t,r)\in\Omega_g} \frac{|R_g(t,r) - a_g^* P_g(t,r)| - |R_g(t,r)|}{\sigma_{R_g(r)}}. \tag{8.8}$$

Here, the intensity a_g^* is obtained by minimizing the following equation:

$$a_g^* = \text{argmin}_{a_g} \sum_{(t,r)\in\Omega_g} \frac{|R_g(t,r) - a_g P_g(t,r)|}{\sigma_{R_g(r)}}. \tag{8.9}$$

The key difference in this approach is that, instead of using R (as in the previous definition of likelihood ratio in (4.9)), we use R_g, which corresponds to each specific trajectory. Finally, the likelihood ratios $\log \Lambda_g(R_g)$ from all trajectories are combined to create the overall likelihood ratio map.

8.3 Experimental Results

We tested and compared our algorithms using the β-Pictoris and 51 Eri datasets, which are described in detail in other sections. The testing process began with the detection of real planets in these datasets. Next, we added synthetic (fake) planets to the datasets to further evaluate the performance of our algorithms. Finally, we conducted ROC curve comparisons to assess the accuracy and effectiveness of the detection methods.

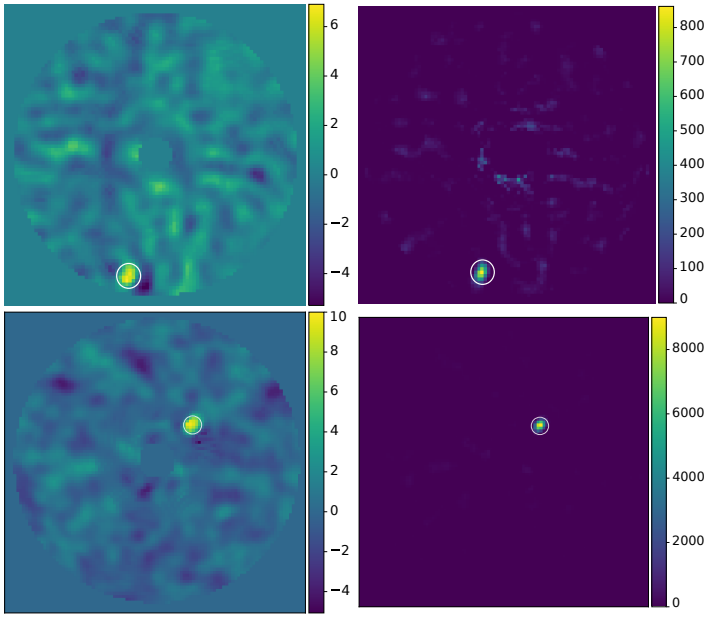

Fig. 8.7 SNR maps (left) and likelihood maps (right) after MC using 51 Eridani dataset (top) and β-Pictoris dataset (bottom).

8.3.1 Detection map results

Real planets

In Fig. 8.7, we present the results from the real dataset tests. The SNR maps on the left and the likelihood ratio maps on the right were used to detect the planets. In all cases, the planets were successfully detected without any false positives.

Afterward, we examined the effect of the ρ_{ds} value on the performance of the likelihood ratio map algorithm. To do this, we generated likelihood ratio maps using different ρ_{ds} values, ranging from 0.2 to 2 in increments of 0.2, applying the 51 Eri dataset. For all non-zero values, we ensured that the ρ value used to obtain the likelihood ratio map and the ρ_{ds} value used in matrix completion remained consistent. When ρ_{ds} was set to 0, there was no data left for the likelihood ratio calculation, so we began the comparison with ρ_{ds}=0.2.

To evaluate the performances, we used the S/N ratio of the likelihood ratio map, as described in Section 7.1.3. This approach allowed us to assess

how different values of ρ_{ds} affect the algorithm's efficiency across various data scenarios.

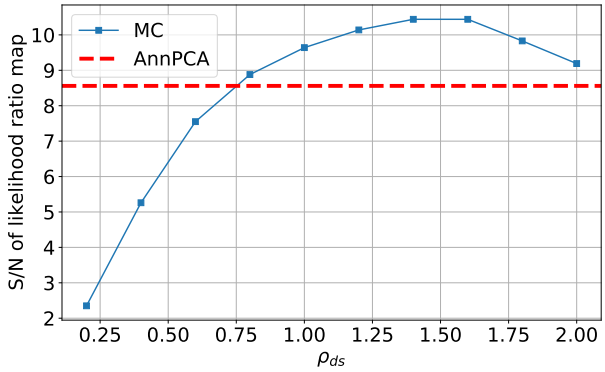

Fig. 8.8 S/N of likelihood ratio map in planet location using different ρ_{ds}.

Fig. 8.8 illustrates that increasing the ρ_{ds} value initially improves performance, peaking between 1.4 and 1.6, before beginning to decline. This trend indicates that removing only a small portion of the data does not yield particularly successful results. However, when we expand the size of the disk, we also remove points unrelated to the exoplanet, leading to suboptimal outcomes. This observation highlights a balance needed in setting the ρ_{ds} value to effectively differentiate between relevant and irrelevant data points without compromising the overall accuracy of our results. In this figure, we also see that when we use a good ρ_{ds} we can have better values than annular PCA.

Fake planets

To effectively assess the performance of our algorithm, only examining results for real planets is insufficient. Hence, we utilize a comparative approach involving fake planets.

In our experiments, we employed the 51 Eri dataset. Within this dataset, we strategically injected four synthetic planets at separations of $2\lambda/D$, $5\lambda/D$, $8\lambda/D$, and $15\lambda/D$. The detection results of this exercise are presented in Fig. 8.9.

Upon analyzing the detection maps, it becomes evident that the SNR map after applying Annular PCA could only identify one of the four plan-

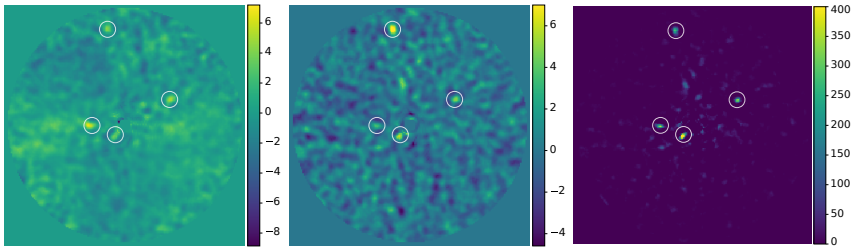

Fig. 8.9 SNR after annular PCA, SNR after MC, and likelihood after MC from left to right.

ets. In contrast, SNR map following Matrix Completion (MC) successfully detected three out of four planets. Notably, the likelihood ratio method applied after MC demonstrated the most effective performance, successfully identifying all four injected planets.

8.3.2 ROC curve results

For our ROC (Receiver Operating Characteristic) curve comparisons, we start with ablation studies, where we compare the performance of algorithms with and without MC. Specifically, when MC is applied as described in Section 8.2.1, the trajectory is removed and replaced using the MC process. In contrast, without MC, the same algorithm is used, but there is no removal or filling, which means $\rho_{\text{ds}} = 0$ as in Equation (8.6).

Next, we aim to investigate the impact of the disk size, denoted as ρ_{ds}, on the performance of algorithm. This will help us understand the effectiveness of disk size adjustments in different scenarios.

Finally, we compare our algorithm with the traditional Annular PCA method on various separations.

MC vs without MC

To validate the effectiveness of the Matrix Completion (MC) process, we performed several sanity checks. These involved two distinct procedures: first, we removed the trajectory and applied MC, and second, we ran the same algorithm without removing the trajectory, which we refer to as without MC (w/o MC). Without MC means we applied MC algorithm with $\rho_{\text{ds}} = 0$. Our tests focused on assessing the performance of algorithm for detecting planets both close to (at a separation of $3\lambda/D$) and far from (at separations of 8-9 λ/D) the star.

The outcomes of these tests are illustrated in the form of ROC curves. The results for planets at smaller separations are presented in Fig. 8.10, while the results for planets at larger separations can be found in Fig. 8.11. These figures provide a visual representation of the algorithm's effectiveness under different scenarios, comparing the performance of the algorithm with and without the application of MC.

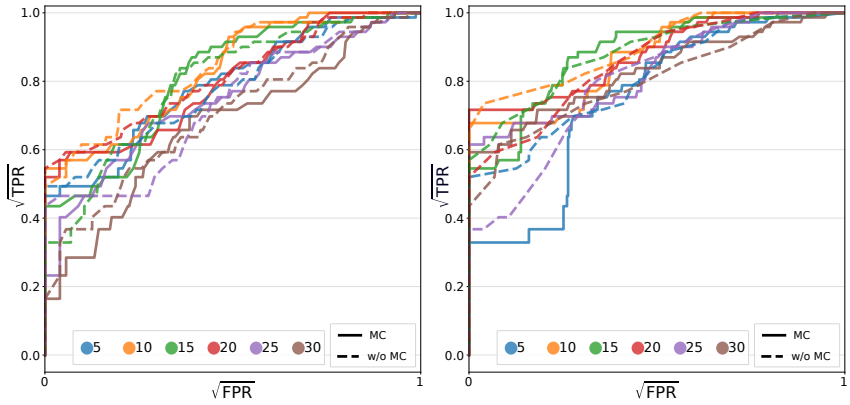

Fig. 8.10 Comparison of MC and without MC for various ranks from 5 to 30. The planets are injected into the separation $3\lambda/D$ with σ flux. We use β-Pictoris in these curves. The left column is SNR comparison and the right column is likelihood comparison. For likelihood ratio map we use $\rho = 1$ even if we use use $\rho_{\text{ds}} = 0$ for without MC.

In the figures mentioned, our observations generally indicate that whether we apply MC or not, the resulting ROC (Receiver Operating Characteristic) curves are quite similar. Specifically, when we utilize the same rank for comparison, we observe notably comparable values. This similarity suggests that the MC process, in this context, does not significantly alter the ability of algorithm to distinguish between true and false positives, as evidenced by the overlapping or closely aligned ROC curves in both scenarios.

MC vs AnnPCA

We performed a comparison between a classical method AnnPCA and MC. Both are a technique based on low-rank approximation. For this comparison, we injected planets at a separation of $3\lambda/D$. The β-Pictoris dataset was chosen for this study, and the flux of the injected planets was set to 1 λ. We tested each algorithm with varying rank values, incrementally increas-

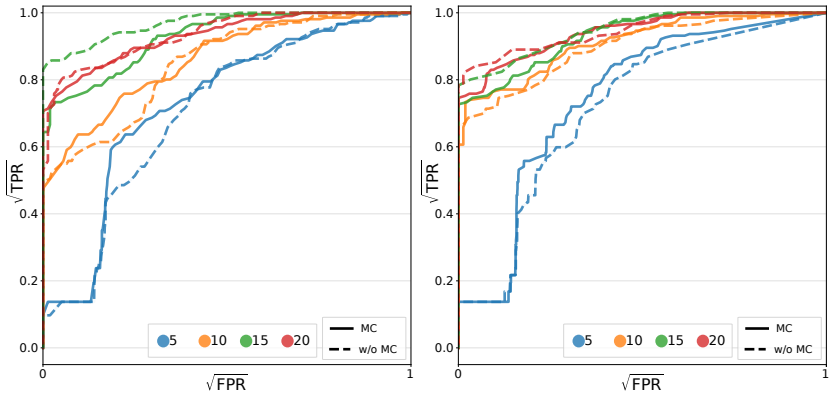

Fig. 8.11 Comparison of MC and without MC for various ranks from 5 to 20. The planets are injected into the separation $8 - 9\lambda/D$ with 1.0σ flux. We use 51 Eri dataset in these curves. The left column is SNR comparison and the right column is likelihood comparison.And we applied the same procedure in Fig. 8.10 for mask sizes.

ing from 10 to 35 in steps of 5 (you can see the results for each algorithm with different rank values in Figures 8.12 and 8.13). The curve that yielded the best results was then used for a comparative analysis in Fig 8.14.

When we examine the SNR comparisons, we observe that both Annular PCA and MC exhibit similar performance levels. However, in the likelihood ratio comparisons, MC shows a slight edge, delivering slightly better results.

8.4 Discussion

An observed trend is that, the brighter the planet, the more MC outperforms standard PCA. An interpretation is that if the planet is sufficiently bright, then PCA favors capturing it to the detriment of its purpose of capturing the noise. Overall, we found that MC is not competitive with the AMAT approach that we described in Chapter 7. We nevertheless present it in detail as we find that it mobilizes an appealing concept.

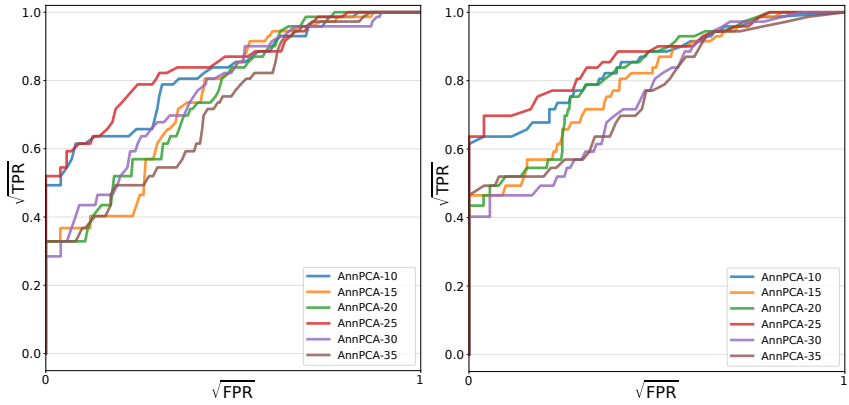

Fig. 8.12 Annular PCA using different ranks. The planets are injected into the separation $3\lambda/D$ with 1.0σ flux. The dataset is β-pictoris. The left column is SNR comparison and the right column is likelihood comparison.

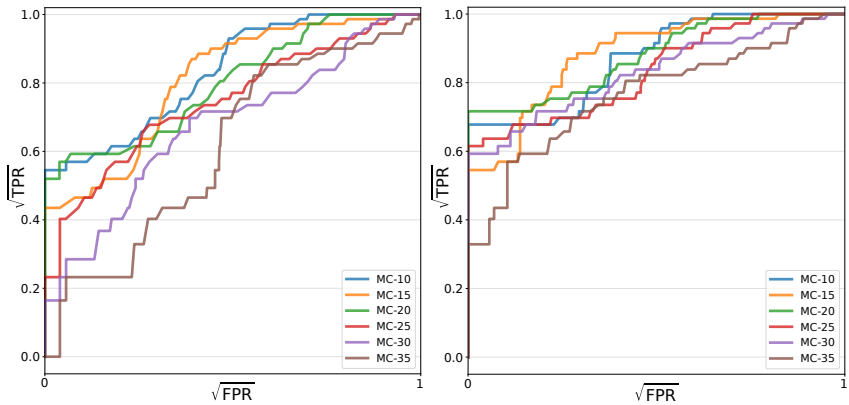

Fig. 8.13 MC using different ranks. The planets are injected into the separation $3\lambda/D$ with 1.0σ flux. The dataset is β-pictoris. The left column is SNR comparison and the right column is likelihood comparison.

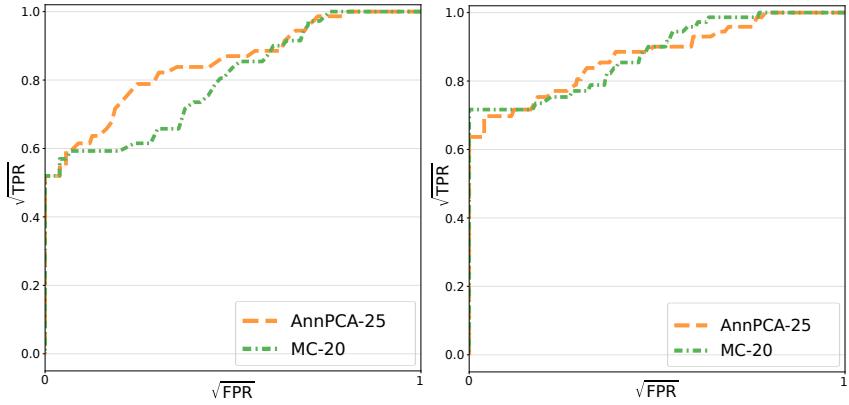

Fig. 8.14 MC vs. Annular PCA. The planets are injected into the separation $3\lambda/D$ with 1.0σ flux. The dataset is β-pictoris.The left column is SNR comparison and the right column is likelihood comparison.

9

Discussion and Conclusion

9.1 Summary

In this thesis, we proposed novel approaches to direct imaging, a technique with numerous advantages for exoplanet detection due to the ability to capture signals directly from the planet itself. We introduced various methods that could enhance direct imaging throughout different stages of the process and evaluated them by comparing them with existing techniques.

In the first two chapters, we provided general information about exoplanet detection and direct imaging, discussing the state-of-the-art methods. In Chapter 3, we explained how evaluations related to direct imaging can be conducted, drawing on current techniques. We emphasized the importance of performing thorough and fair comparisons by using ROC curves and proposed our own approach to obtaining these curves.

In Chapter 4, we introduced a method called the likelihood ratio map. This technique involves applying our model to the residual data cube obtained after PSF model subtraction, aiming to highlight and reveal the planet's signal. We modeled it as $R = a_g P_g + E$, where R is the residual cube, P_g is the PSF copied according to the planet's location, a_g is the flux, and E is the noise. Using this model, we first determined a_g using the Maximum Likelihood approach and then calculated the likelihood ratio

for all possible positions. By combining these likelihood ratios, we constructed the likelihood ratio map. A significant point in our approach was recognizing that the noise might follow a Laplacian rather than a Gaussian distribution, and we adapted our method accordingly. We compared this with the commonly used SNR map and found that our approach provided the most effective results. The likelihood ratio map was used throughout the other chapters.

In Chapter 5, we proposed a model that combines PSF modeling with sparse matrix estimation. Instead of applying an SNR or STIM map to the residual matrix, we used rotation and convolution operations to obtain a sparse matrix. We designed simple alternating iterative hard-thresholding algorithm to solve this problem. We compared this with the standard PCA + SNR and PCA + likelihood ratio map (LRM) approaches, achieving better results than the PCA + SNR approach and obtaining mixed results compared to PCA + LRM, depending on the input.

In Chapter 6, we developed a solution to incorporate the Laplacian noise model into the PSF subtraction process. Whereas Gaussian noise models use the L2 norm for solving the problem, we adapted it to the L1 norm for the Laplacian noise model. We use L1-LRA to solve this problem. We explored whether the residual data fits better with Gaussian or Laplacian distributions and demonstrated the success of using the L1 norm through various comparison methods.

In Chapter 7, we further improved the PSF model introduced in Chapter 5. Using the model $M = L + a_g P_g + E$, we solved for L and a_g simultaneously. L was obtained using the method proposed in Chapter 5, and then a_g was determined by solving the least absolute deviation problem. We alternated between solving these two subproblems, achieving improved results. We compared the performance using various metrics and further enhanced it with the LRM approach. Although L2 norm results were not as effective as L1 norm to detect planets, L2 was still useful for estimating the planet's intensity. When compared to flux estimation methods in the literature, our method produced promising results.

Since the low-rank component, intended to capture noise, may also absorb some of the planet's signals in the algorithms which we proposed in the other chapters, we explored matrix completion techniques in Chapter 8. We assumed knowledge of the planet's location and removed its signals along the planet's orbit. We then applied matrix completion algorithms to reconstruct the missing data. When we subtract the reconstructed data from the original version, we expect to see the planet signal in the residual.

Although the algorithm provided successful results, it did not outperform PCA, yielding similar outcomes.

9.2 Discussion and Future Works

During my PhD studies, I observed that algorithms assuming Laplacian noise, such as our AMAT_{L1} algorithm, delivered successful results compared to approaches relying on Gaussian noise models. This improvement was particularly notable when these algorithms were paired with likelihood ratio maps, where the Laplacian noise assumption further enhanced performance.

Given the demonstrated success of the Laplacian noise models in this thesis, it is clear that future research should prioritize the development and application of L1 norm-based techniques. Promising research directions include the adaptation of L1 norm matrix completion algorithms to handle datasets with varying noise characteristics. Additionally, future work could explore the integration of noise models that go beyond the Gaussian and Laplacian assumptions. Mixed noise models, which combine elements of different noise distributions, could provide a more nuanced representation of the variability inherent. These models could be evaluated using metrics such as Kullback-Leibler divergence, which would offer insights into their alignment with the underlying data distributions and their suitability for specific datasets.

The L1-LRA and AMAT_{L1} algorithms, which are grounded in L1-LRA principles, have demonstrated successful results in this work, outperforming many conventional approaches. However, one significant drawback of these methods is their computational expense, which can make them less practical for large-scale datasets. To address this limitation, future research could focus on accelerating low-rank matrix computations within the L1 norm framework. Moreover, one potential solution involves leveraging GPU-accelerated computations, which could significantly reduce processing times by exploiting parallel processing capabilities.

One of the main challenges of the AMAT_{L1} algorithm is its high computational cost, primarily due to the exhaustive evaluation of all possible trajectories in the dataset. To mitigate this issue, several strategies can be explored. One approach, detailed in Section B.2.3, involves splitting the annulus into smaller sections for evaluation rather than processing the entire annulus simultaneously. Additionally, initializing the algorithm with a better starting value for a_g, such as the average a_g of neighboring pix-

els in the above row instead of 0, could potentially decrease the number of iterations required for convergence. While our preliminary attempts introduced vertical shifts in processed images—such as the rolling shutter effect in the processed frame—refining a better approach could yield substantial improvements in computational efficiency and accuracy. Another promising optimization could involve selectively applying the algorithm to pre-identified pixels within the dataset, focusing only on specific trajectories, rather than the entire set of positions. A faster pre-processing step could identify these pixels, ensuring computational resources are concentrated on regions of interest without sacrificing accuracy.

Extending our algorithms to accommodate Spectral Differential Imaging (SDI) and combined SDI-ADI data represents another promising direction for future research. This extension would involve applying Angular Differential Imaging (ADI) techniques to each wavelength individually, effectively transforming the current 3D data model into a 4D framework. However, given the complexity of 4D data, alternative approaches might be more efficient or yield better results after we modify our 3D data model to 4D data. The adoption of tensor methods could significantly enhance the capabilities of direct imaging techniques, enabling more detailed and accurate analyses.

Rank selection in low-rank approximation methods remains a critical factor influencing algorithm performance. In our experiments, we observed that determining a suitable rank is both challenging and essential for achieving the best results. Although suitable ranks depend on the dataset and the algorithm, it is possible to estimate an approximate rank range. This variability is influenced by the radial changes in noise intensity and distribution: for closely separated sources, a higher rank may perform better, while a lower rank could be more effective for larger separations. To ensure fairness in our comparisons, we tested multiple ranks for each algorithm and used their best results for performance evaluation. While this approach provided a baseline for comparison, it highlighted the need for more systematic methods to choose the rank dynamically, based on the characteristics of the dataset. Future research should focus on developing adaptive techniques for rank selection, which could significantly improve the efficiency and performance of these algorithms.

Preprocessing steps also play a pivotal role in improving algorithmic performance and reducing noise in exoplanet data. Our preliminary work using Fourier transforms to reduce noise in the datacube, as described in Chapter B, has shown promising initial results. Although this work is still

in its early stages, it underscores the potential for further refinement and application of such techniques. By systematically addressing noise reduction through preprocessing, researchers could improve the overall performance and reliability of algorithms like AMAT_{L1}, particularly in datasets with high noise levels.

Finally, distinguishing circumstellar disks from exoplanets remains a significant challenge, especially when the planet and disk signals are closely intertwined. One promising approach involves estimating the intensity of the disk and separating its signal from the planetary signal through matrix completion methods. By creating a mask to remove the disk signals, matrix completion can reconstruct the missing data and isolate the disk's intensity. This process enables more accurate separation of the disk and planetary signals, paving the way for improved detection and characterization of exoplanets. Although this approach has shown potential, its success depends heavily on the initial mask and the reliability of the reconstructed matrix. Addressing these challenges through further research and refinement will be essential for advancing the field of direct imaging.

A

Appendix

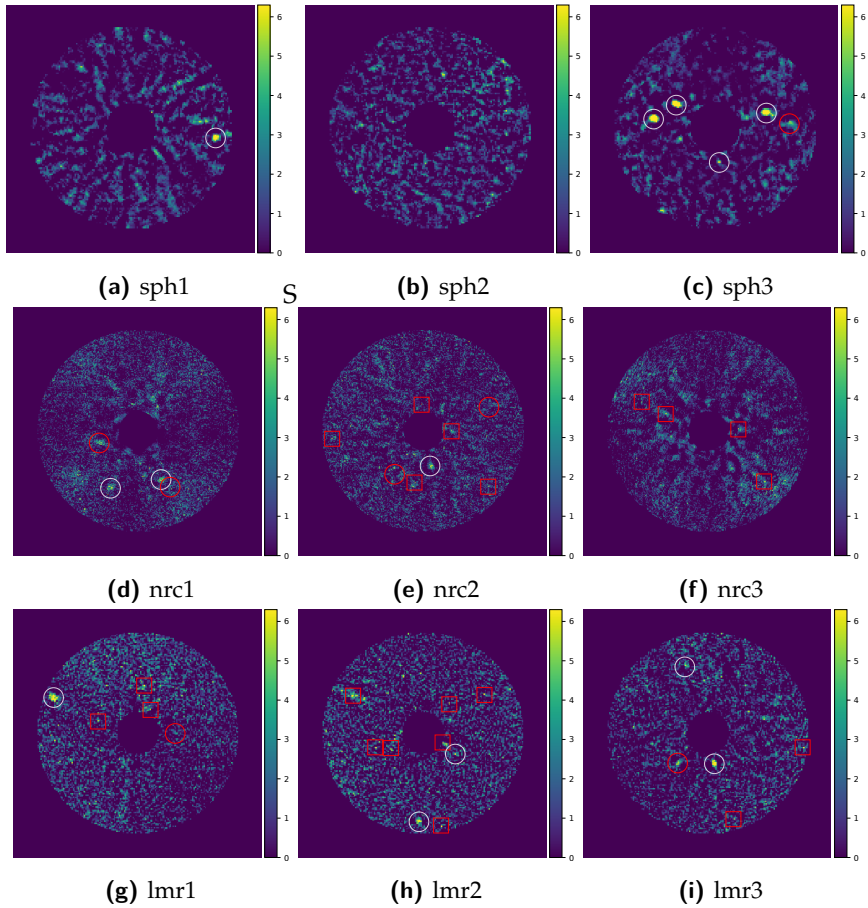

Fig. A.1 S/N maps, proposed by [BGG+23], after AnnPCA: In these S/N maps, white circles represent TP, red squares denote FP, and red circles signify FN.

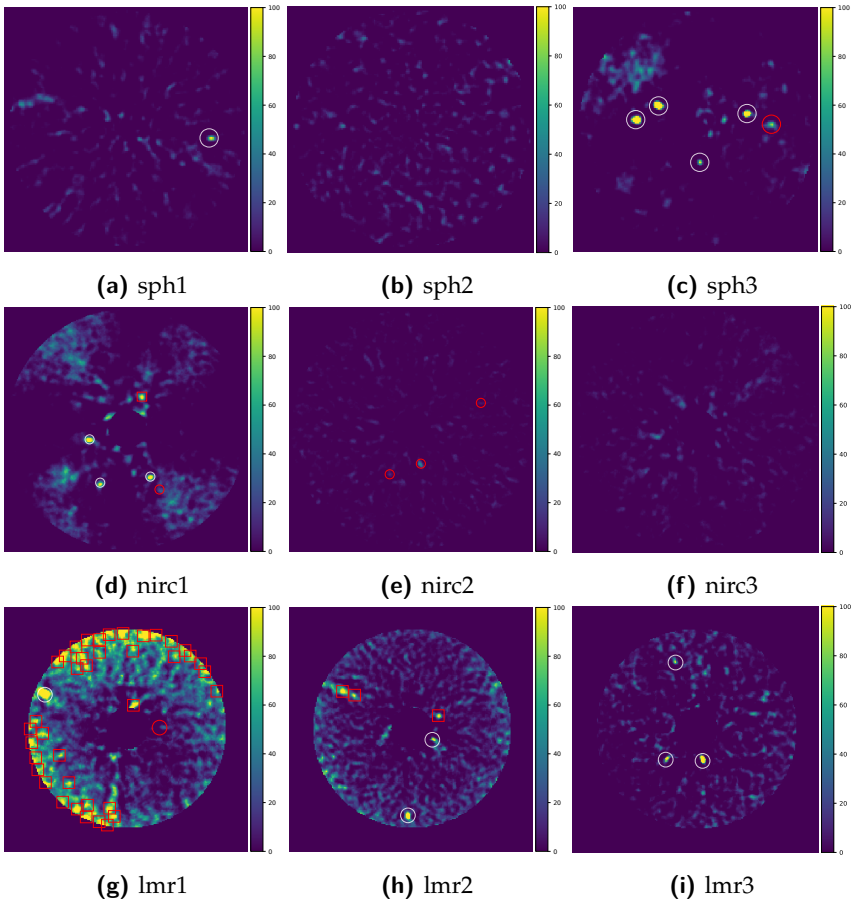

(a) sph1 **(b)** sph2 **(c)** sph3

(d) nirc1 **(e)** nirc2 **(f)** nirc3

(g) lmr1 **(h)** lmr2 **(i)** lmr3

Fig. A.2 LRM after AnnPCA for EIDC Datasets: In these maps, white circles represent TP, red squares denote FP, and red circles signify FN.

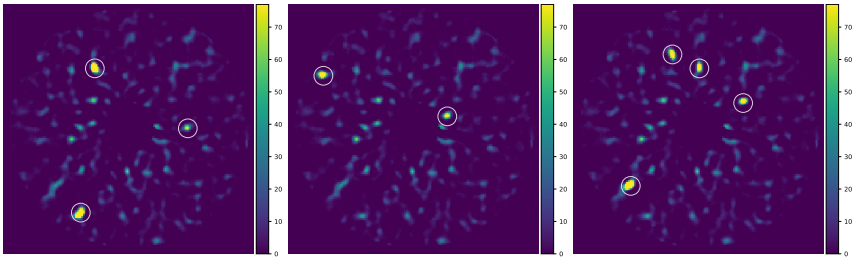

Fig. A.3 LRMs after AMAT_{L1} for 51 Eri datasets with injected planets: In these maps, white circles represent TP.

B

Fourier Transform for Exoplanet Detection

In this appendix, we report the development of a preprocessing step to prepare the dataset for the procedures outlined in previous chapters, particularly for the AMAT algorithm introduced in Chapter 7, and the likelihood map method detailed in Chapter 4, which we used for detecting exoplanets within the obtained data cube. Given that datasets used in exoplanet detection often contain significant levels of noise, we employ a Fourier transform to better prepare the data for subsequent analysis.

The Fourier transform helps reduce noise and improves the quality of the data, making it suitable for input into the AMAT algorithm. We explain the specific steps involved in adapting the data using the Fourier transform in accordance with the requirements of the AMAT algorithm.

In Section B.1, we provide a detailed overview of the Fourier transform, discussing its purpose and its general applications. Section B.2 explores how this method can be tailored for exoplanet detection, particularly in adapting it to an annular form. Given the significant computational cost associated with applying the algorithm to all possible trajectories, and the additional slowdown caused by the Fourier transform process, we propose strategies to improve the algorithm's efficiency.

Finally, in Section B.3, we evaluate the effectiveness of the proposed method through visual comparisons using SNR maps and likelihood ra-

tio maps. These comparisons are further strengthened by presenting ROC curves to provide a robust analysis of the algorithm's performance.

B.1 Fourier Transform

The Fourier Transform is a fundamental mathematical technique that plays a pivotal role in signal processing, image analysis, physics, and engineering. It serves as a bridge between the time or spatial domain and the frequency domain, allowing us to analyze complex waveforms by decomposing them into simpler sinusoidal components at various frequencies [Bra00, Kai94].

In essence, it takes a complex waveform, such as a signal or image, and represents it as the sum of simpler sinusoidal waves, specifically sine and cosine functions, each with its unique frequency and amplitude. This decomposition process enables us to gain insights into the underlying frequency components of the original signal.

The Continuous Fourier Transform is the primary tool for continuous-time signals or functions. It is defined by an integral equation, expressed as follows:

$$\mathcal{F}(w) = \int_{-\infty}^{\infty} f(t) \cdot e^{-jwt}, dt \tag{B.1}$$

Here, $\mathcal{F}(w)$ represents the complex-valued frequency domain, t is the time variable, and w is the angular frequency. This mathematical operation is essential for understanding the spectral content of signals and is widely applied in various fields.

B.1.1 Discrete fourier transform

In many practical scenarios, we deal with digital data sampled at discrete intervals. The Discrete Fourier Transform (DFT) is a discretized version of the Fourier Transform, specially designed for digital data. It takes a finite sequence of data points and transforms it into a corresponding sequence of complex numbers that represent the amplitudes and phases of sinusoidal components at discrete frequencies [OWH97].

The DFT is a critical tool in applications such as audio processing, image processing, and data analysis. For a discrete signal $x[n]$ of length N, the DFT is defined as:

$$\mathcal{X}[k] = \sum_{n=0}^{N-1} x[n] \cdot e^{-j2\pi kn/N} \tag{B.2}$$

Here, $\mathcal{X}[k]$ is the complex-valued frequency domain representation of $x[n]$, and k takes discrete values from 0 to $N-1$. This transformation enables us to analyze digital signals in terms of their frequency components.

B.1.2 Fast fourier transform

While the DFT is a powerful tool, it can be computationally intensive for large datasets. The Fast Fourier Transform (FFT) is an algorithmic technique that accelerates the computation of the DFT [Bra00]. It achieves this by exploiting symmetries and reducing the number of arithmetic operations required. As a result, it significantly reduces the computation complexity from $O(N^2)$ to $O(N \log N)$, where N is the number of data points.

The FFT has applications in diverse fields, including image processing, audio analysis, and scientific simulations. It speeds up the transformation process, making it feasible to analyze large datasets in real-time or near real-time. Moreover, the FFT is widely used in multidimensional data, such as images. In the case of a two-dimensional data version, denoted as FFT2, it is applied to matrices, like images, and expressed as:

$$\mathcal{F}(k_x, k_y) = \sum_{n=0}^{N_x-1} \sum_{m=0}^{N_y-1} f(n,m) e^{-j2\pi(k_x n/N_x + k_y m/N_y)} \tag{B.3}$$

Here, $f(x,y) \in \mathbb{R}^{m \times n}$ represents the input matrix, and (k_x, k_y) correspond to the discrete frequencies in the horizontal and vertical directions. The FFT2 is instrumental in applications involving two-dimensional data, such as image processing and spatial signal analysis.

B.2 Fourier Transform for Exoplanet Detection

B.2.1 FFT as a preprocessing

Let $M_{3D} \in \mathbb{R}^{T \times N \times N}$ be datacube which includes T images with size $N \times N$. We propose a new minimization called Fourier Transform of AMAT (FoTAMAT) that removes high frequencies

$$\hat{L_{3D}}, \hat{a}_g = \underset{L_{3D}, a_g}{\arg\min} \|\mathcal{F}^{-1}(\mathcal{F}_\Omega(M_{3D})) - L_{3D} - a_g \mathcal{F}^{-1}(\mathcal{F}_\Omega(P_{g_{3D}}))\| \quad \text{(B.4)}$$

$$\text{subject to} \quad \text{rank}(L_{3D}) \leq r \quad \text{(B.5)}$$

$$P_{g_{3D}} \in \Lambda \quad \text{(B.6)}$$

$$\|P_{g_{3D}}\| \approx \|\mathcal{F}^{-1}(\mathcal{F}_\Omega(P_{g_{3D}}))\|. \quad \text{(B.7)}$$

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ define FFT2 and inverse of FFT2 along spatial dimension of datacube M_{3D}, respectively, and $\mathcal{F}_\Omega$ is the masked version of $\mathcal{F}$. In Fig. B.1, we present an illustrative frame of both M_{3D} and $P_{g_{3D}}$ shown in both spatial and frequency domains. Ω indicates the mask, an example of which is shown in Fig. B.3, applied to the frames in the frequency domain to remove high frequencies.

In FFT processing, the handling of high frequency outliers is critical. By applying a mask to these outliers, we can either eliminate or mitigate their effects. This is illustrated in Fig. B.2, where two different mask sizes are applied to an outlier in a data frame. First, a small-radius mask is applied. This effectively removes the outlier, but also alters other pixels in the frame. Next, a larger-radius mask is used. This approach better preserves the integrity of the pixels, but only reduces, rather than eliminates, the effect of the outlier. The choice of mask size therefore represents a trade-off between removing the outlier and preserving the other pixels in the frame.

We decide the radius of the mask by setting a threshold below which the application of the Fourier transform has a very small effect on P_{g_i} where P_{g_i} is any frame of $P_{g_{3D}}$, like choosing the point where the break occurs in the example in Fig. B.3. Here we obtain the max error with

$$E(\rho) = \max_{j=1,\ldots,m} \max_{k=1,\ldots,n} |\mathcal{F}(P_{g_i}) - \mathcal{F}_{\Omega_\rho}(P_{g_i})| \quad \text{(B.8)}$$

where ρ represents the radius of the mask and (x_0, y_0) is the center of the frame

$$\Omega_\rho(x, y) = \begin{cases} 1, & \text{if } (x - x_0)^2 + (y - y_0)^2 \leq \rho^2 \\ 0, & \text{otherwise} \end{cases} \quad \text{(B.9)}$$

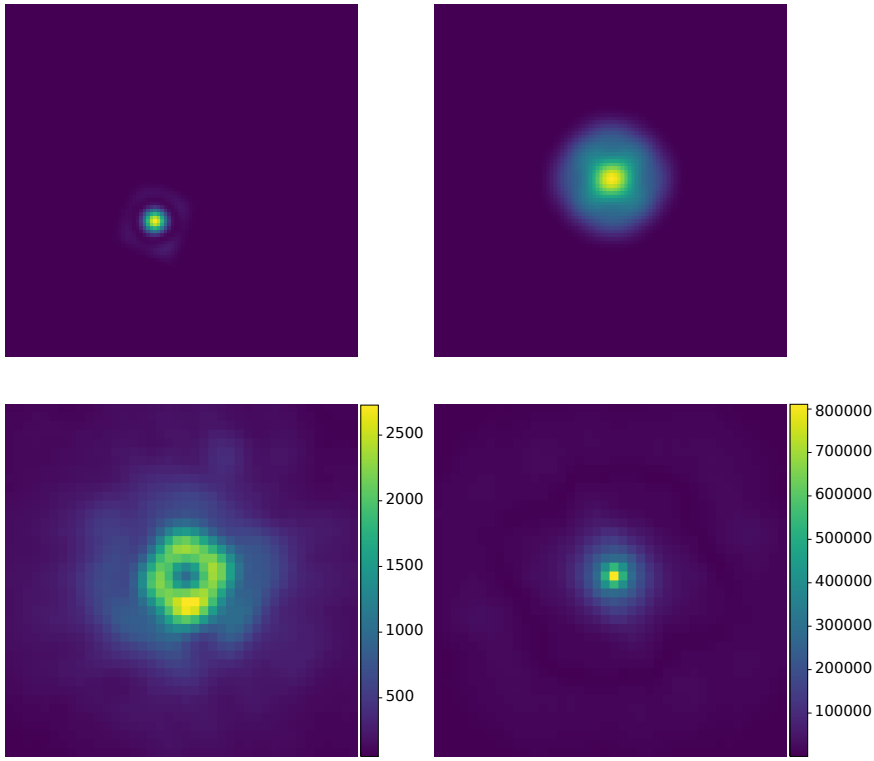

Fig. B.1 A frame from $P_{g_{3D}}$ (top) and the cube M_{3D} (bottom) in the spatial (left) and frequency (right) domain.

B.2.2 Annular fourier transform - polar fourier transform

In direct imaging techniques, we hypothesize that the high-frequency components of pixels located near the star differ from those farther away, and analyzing these components separately could enhance the overall performance. Similar to the concept in Annular PCA, evaluating the correlation and discrepancies between pixels in an annular pattern can provide a clearer understanding of their compatibility. By constructing matrices based on annular regions, rather than uniformly across the field, we aim to achieve more accurate results. In our annular system analysis, we propose to initially convert points from a cartesian coordinate system to polar coordinate system using standard transformations.

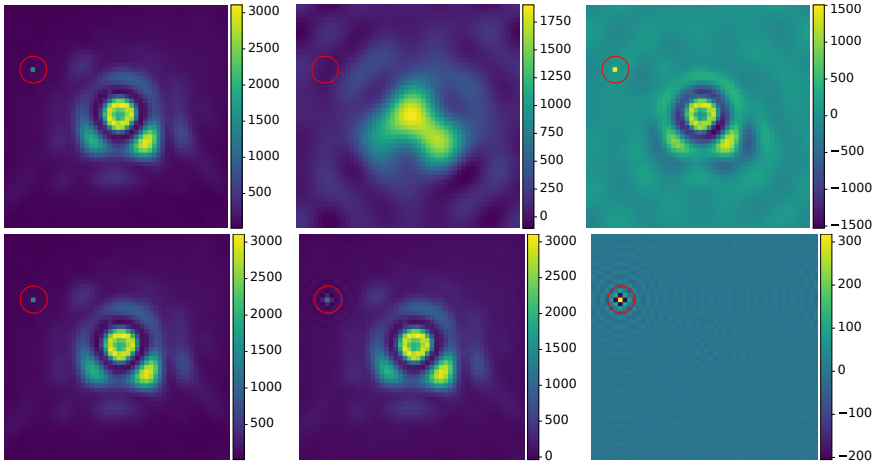

Fig. B.2 Left: A frame from a datacube. Middle: The frame after applying FFT. Right: The difference between the frames. We compare the frames when we use a small (top) and large (bottom) radius ρ for the mask. Red circle highlights the a noise which has high frequency.

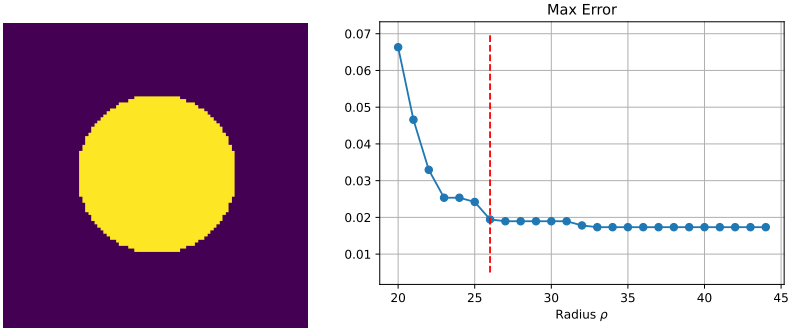

Fig. B.3 Mask to remove high frequencies and error vs radius ρ of mask.

$$r = \sqrt{x^2 + y^2}$$
$$\theta = \arctan\left(\frac{y}{x}\right) \tag{B.10}$$

This is essential for our subsequent processes. However, not all pixel values are available in this transformed space, leading to missing data

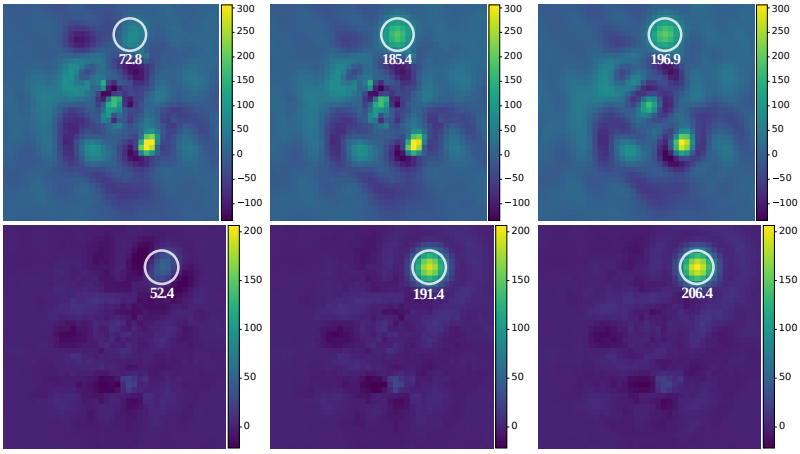

Fig. B.4 First frame after applying PCA, AMAT, and FoTAMAT, recursively (top), and median frame after applying PCA, AMAT, and FoTAMAT, recursively (bottom).

points. To address this, we use cubic interpolation, specifically using the `map_coordinates` function from SciPy, to estimate these missing values accurately.

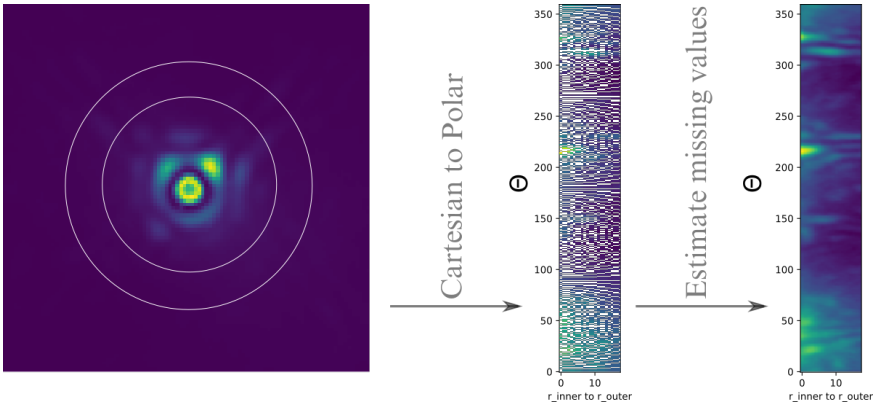

Fig. B.5 Cartesian to polar coordinates transformation

After the conversion from cartesian to polar coordinates, each frame of the dataset is transformed into a rectangular matrix. In order to apply a Fourier transform and then a mask to it, our circular system requires an

Algorithm 4: Cartesian to Polar

Input: Frame $I \in \mathbb{R}^{n \times n}$, inner radius r_{in}, outer radius r_{out}, shape of output (h, w)

1 $\Delta_h = (r_{out} - rin)/h$
2 $\Delta_w = 2\pi/w$
3 $x(i,j) = i\Delta_h \cos(j\Delta_w) + n/2$ for $i = 1, ..., h$ and $j = 1, ..., w$
4 $y(i,j) = i\Delta_h \sin(j\Delta_w) + n/2$ for $i = 1, ..., h$ and $j = 1, ..., w$
5 Cubic spline $I_{new} = p_3(x, y)$

elliptical mask due to the rectangular nature of our matrices, unlike the typical square or circular masks used in Fourier analysis. We define this mask as:

$$\Omega_{\rho_x,\rho_y}(x,y) = \begin{cases} 1, & \text{if } \left(\frac{x-x_0}{\rho_x}\right)^2 + \left(\frac{y-y_0}{\rho_y}\right)^2 \leq 1 \\ 0, & \text{otherwise} \end{cases} \tag{B.11}$$

where ρ_x and ρ_y are the radii of the ellipse along x and y axes, respectively.

After applying the new elliptical mask to our frames, we can proceed by applying the inverse of the Fourier transform to revert the data from the frequency domain back to the spatial domain. We then apply the AMAT algorithm to these transformed images. Finally, to complete our processing sequence, we convert the data back into cartesian coordinates. This conversion is accomplished by the following equations:

$$\begin{aligned} x &= r\cos(\theta) \\ y &= r\sin(\theta). \end{aligned} \tag{B.12}$$

Algorithm 5: Polar to Cartesian

Input: Frame $I \in \mathbb{R}^{h \times w}$, inner radius r_{in}, outer radius r_{out}, shape of output (n, n)

1 $x = [1, ..., n] - n/2$ and $y = [1, ..., n] - n/2$
2 $r = \sqrt{x^2 + y^2}$
3 $r = n(r - r_{in})/(r_{out} - r_{in})$
4 $t = \arctan(y/x)$
5 $t = n * t/(2\pi)$
6 Cubic spline $I_{new} = p_3(r, t)$

Algorithm 6: FoTAMAT - annular

Input: Image sequence $M \in \mathbb{R}^{t \times n \times n}$, possible trajectories Λ, maximum number of iteration maxiter, stopping criteria ϵ, radius of mask (ρ_x, ρ_y), annular size r_a

for *all trajectories* $P_g(r,\theta) \in \Lambda$ **do**
  for *all frames* M_j **do**
    Obtain M'_j (similarly P'_j) using cartesian to polar algorithm with the inputs: $I = M_j$ $(I = P_j)$, $r_{in} = r - r_a$, $r_{out} = r + r_a$, and $(h, w) = (2r_a, 360)$
    $\hat{M}_{\Omega_j} = \mathcal{F}^{-1}(\mathcal{F}(M'_j) \cdot \Omega_{\rho_x,\rho_y})$ and $\hat{P}_{\Omega_j} = \mathcal{F}^{-1}(\mathcal{F}(P'_j) \cdot \Omega_{\rho_x,\rho_y})$
  end
  $a_g^{(0)} = 0$
  for $i = 1$: maxiter **do**
    $L_i = \arg\min_L \|\hat{M}_\Omega - L - a_g^{(i-1)} P_{g_\Omega}\|$
    if norm $= 1$ **then**
      $a_g^{(i)} = \min\{\frac{\hat{M}_\Omega - L^{(i)}}{P_{g_\Omega}}\}$
    end
    else if norm $= 2$ **then**
      $a_g^{(i)} = \langle P_{g_\Omega}, \hat{M}_\Omega - L_i \rangle / \|P_{g_\Omega}\|^2$
    end
    if $\frac{|a_g^{(i)} - a_g^{(i-1)}|}{|a_g^{(i)}|} < \epsilon$ **then**
      break
    end
  end
  $\hat{R}(r,\theta) = \hat{M}_\Omega - L_i$
  for *all frames* $\hat{R}_{\Omega_j}$ **do**
    Obtain R'_j (similarly P'_j) using polar to cartesian algorithm with the inputs: $I = \hat{R}_j$, $r_{in} = r - r_a$, and $r_{out} = r + r_a$
  end
  $R(r,\theta) = R'$
end

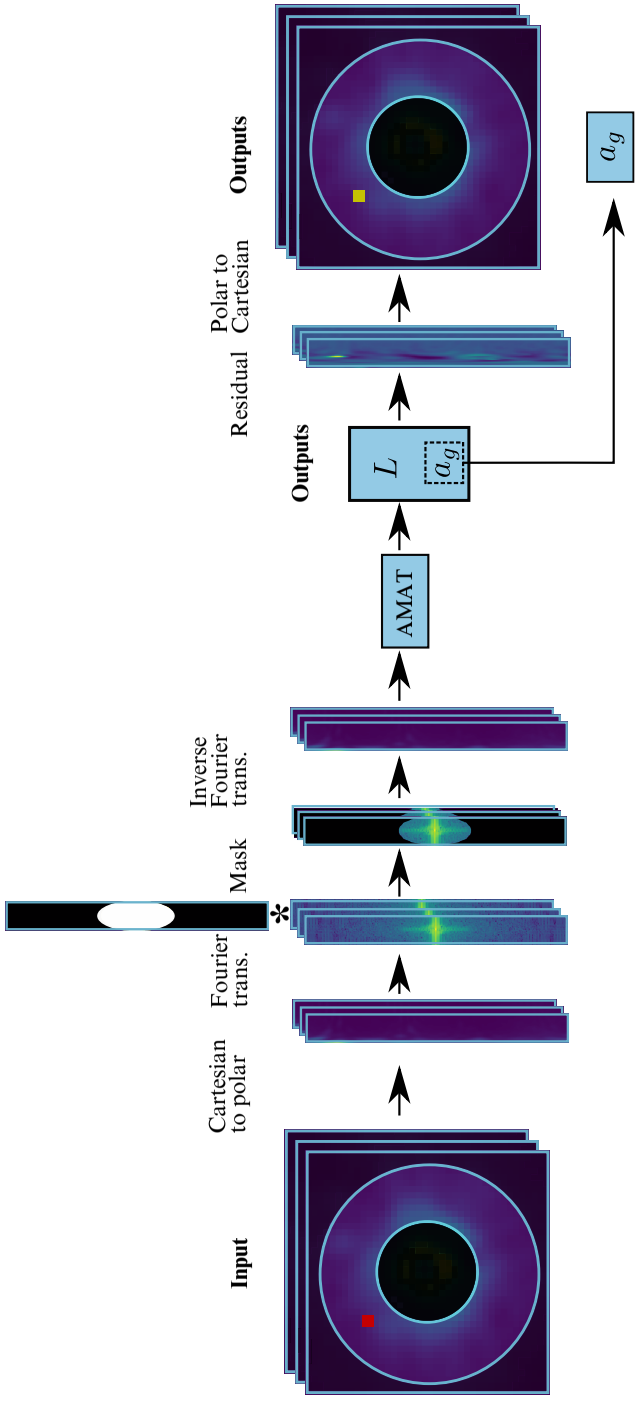

Fig. B.6 Pipeline of FoTAMAT Algorithm

B.2.3 Fast version of FoTAMAT

When we apply the AMAT algorithm with these pre- and post-processing steps, our algorithm slows down significantly. This is because, like the AMAT algorithm, we apply the entire process to all possible pixels in the frame (excluding edge pixels and star pixels). This results in significant time consumption even for a single datacube. Therefore, we are trying to speed up this algorithm.

Instead of taking all the pixels in the annulus, we opt to use only a subset of them, as they are likely to exhibit similar values. Recognizing that small separations include fewer pixels, we developed three different scenarios to ensure that a comparable number of pixels are analyzed in both small and large separations. Instead of dividing the annuli into a fixed number of segments, we propose to divide them into segments with a similar number of pixels. In practice, these scenarios allow for a more balanced and representative analysis across different separation sizes, improving the accuracy and reliability of our results. We suggest three options for these segments illustrated in Fig. B.7.

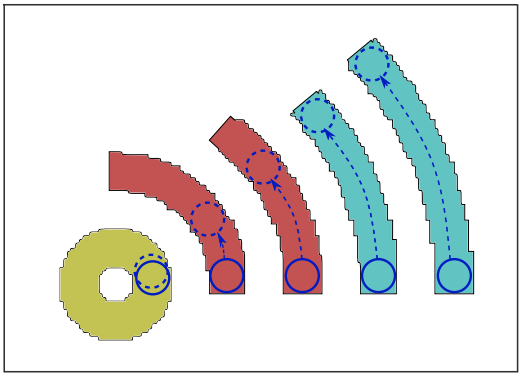

Fig. B.7 The visualisation of different segments according to the separations. The yellow, red and blue segments are the examples of the options (i), (ii), and (iii), respectively.

(i) If the distance r of selected pixel to the center of the frame is smaller than $2\lambda/D$, we take all the values in the annulus.

(ii) If it's larger than $2\lambda/D$, we define a segment within the annulus corresponding to the area between $2\lambda/D \pm r_{\text{ann}}$. To determine the

specifics of this segment, we calculate α_{FoTAMAT} using the following equations:

$$\frac{\alpha_{\text{FoTAMAT}}}{2\pi}\underbrace{\pi\left(\underbrace{\left(r+\frac{r_{\text{ann}}}{2}\right)^2}_{\text{outer radius}}-\underbrace{\left(r-\frac{r_{\text{ann}}}{2}\right)^2}_{\text{inner radius}}\right)}_{\text{Area of annulus}}=\underbrace{\pi\left((2\frac{\lambda}{D}+\frac{r_{\text{ann}}}{2})^2-(2\frac{\lambda}{D}-\frac{r_{\text{ann}}}{2})^2\right)}_{\text{Area between } 2\lambda/D\pm r_{\text{ann}}}$$
$$\iff \alpha_{\text{FoTAMAT}}=\frac{4}{r}\frac{\lambda}{D}\pi \tag{B.13}$$

In this process, we select an arc of α_{FoTAMAT} to define the segment that will cover the selected pixel.

(iii) If the distance is very large, α_{FoTAMAT} may be smaller than the covered angles α_{PA} of parallactic angles. This means that the segment does not cover the entire rotation of the planet, in this case we use α_{PA} instead of α_{FoTAMAT}.

In Fig. B.8, a frame from a dataset is displayed, highlighting an outlier encircled in red. We apply the AMAT and FoTAMAT algorithms to this dataset. We also apply FFT, masking and the inverse of FFT to the entire dataset as a preprocessing, then apply AMAT to evaluate and compare the results of a standard FFT preprocessing and FoTAMAT. In the dataset, the intensity of the outlier initially is about 400. When processed with AMAT, this value is reduced to approximately 200 in the residual cube. Incorporating FFT as a preprocessing step further improves AMAT's performance, lowering the intensity of outlier to about 100. However, after FoTAMAT, the effect of the outlier is reduced to an almost negligible level, indicating the outstanding effectiveness of FoTAMAT in softening such anomalies.

B.3 Experimental Results

B.3.1 Detection maps comparisons

We utilized a publicly accessible dataset from the VLT/SPHERE-IRDIS instrument, specifically observing the star 51 Eri in the K1 band at 2.11μm, which is known to have a planet. This dataset comprises 256 images, covering a parallactic angle of 42 degrees, and we cropped each image to a 100×100 pixel grid to create a detection map. For a more challenging task

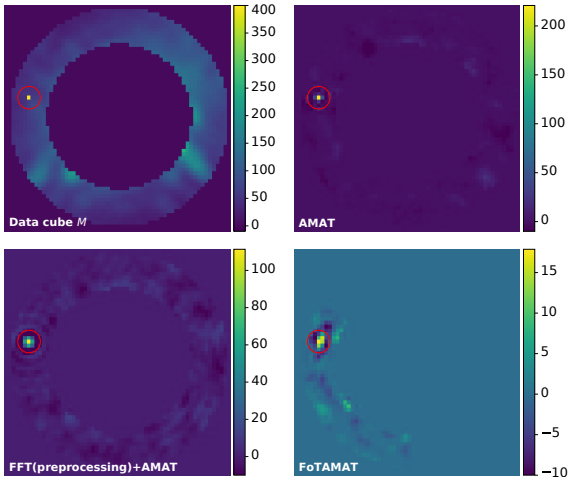

Fig. B.8 Top left: An annular of a frame from the data cube M which has an outlier. Top right: The frame from the residual cube obtained by AMAT algorithm using a trajectory including the outlier. Bottom left: The same procedure given in top right. The only difference is before applying AMAT, the datacube is smoothed by FFT. Bottom right: The frame from the residual cube obtained by FoTAMAT algorithm using a trajectory including the outlier.

and to decrease the computation load, we analyzed only every third image in the series, which reduced our dataset to a total of 86 images. For our analysis, we applied AnnPCA, AMAT and FoTAMAT with a rank of 20 for each of the algorithms, which is suggested previously for AnnPCA using the same dataset [DCA20].

First, we compared the results obtained with AnnPCA, AMAT and FoTAMAT using an SNR map. The SNR map was generated using VIP-HCI [GG+17, CGF+23]. In the maps produced by each of the three algorithms, the highest SNR values belongs to the planets. While the AnnPCA algorithm resulted in the lowest SNR values for the planet, both AMAT and FoTAMAT yielded similar values. However, FoTAMAT achieved slightly higher values.

In addition to this, we conducted a comparison using likelihood ratio maps. When evaluating each map individually, the planets consistently showed the highest values. However, to determine which map is better, we calculate the SNR at the planet's location in the likelihood map

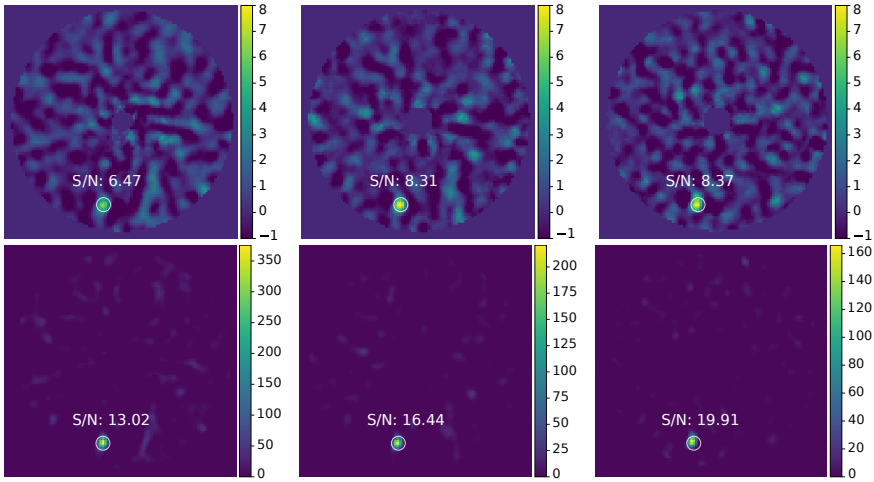

Fig. B.9 SNR (top) and likelihood ratio (bottom) maps of AnnPCA, AMAT and FoTAMAT from left to right.

as suggested in Section 7.1.3. We recommend this method to compare the difference between the targeted planet and the surrounding pixels. Upon comparing the SNR values of these three maps, we once again found that FoTAMAT provided the highest values.

B.3.2 ROC curve comparisons

We compared our FoTAMAT algorithm with AMAT, using deterministic ROC curves as we suggested in the previous chapters. Our focus was primarily on the true positive rate (TPR) achieved without any false positives, a critical measure in ROC analysis. Firstly, we used two datasets; one is 51 Eri explained in Section B.3.1 and second is β-Pictoris from VLT/NACO instrument in the infra-red L' band (3.8μm), which has 612 frames selected one of three of them, covering 83° of parallactic angles. We removed the real planets from the datasets using Negative fake companion (NEGFC) technique [WAG+17], and using opposite values of the parallactic angles [MLMD08]. Then, we injected a total of 100 planets into the datasets, each at a distance of $3\lambda/D$, with a flux of 1.5 times the standard deviation for 51 Eri and 1 times the standard deviation for β-Pictoris.

In Fig. B.10, we can see the results for both datasets. In the ROC curves obtained with the 51 Eri dataset on the left, we observe that the FoTAMAT algorithm gives better results both when using likelihood ratio map and

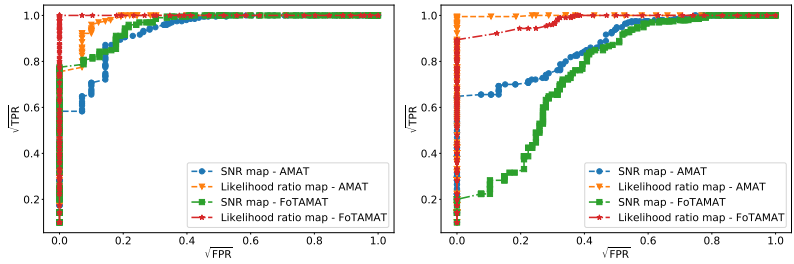

Fig. B.10 Compare of algorithm on 51 Eri (left) and β-Pictoris (right) datasets.

when using SNR map. But in the figure on the right, obtained with the β-Pictoris dataset, we have the opposite graph. Here we get better results in both metrics using the AMAT algorithm.

We think that the reason why we get such different results may be due to the different characteristics of the datasets. According to our observations, the 51 Eri dataset contains much more noise than the β-Pictoris. We add extra noise to the β-Pictoris dataset to see if this happens because it is cleaner.

We divide the frames to annuli. We add uniform distribution $U(-\sigma/4, \sigma/4)$ where σ is the standard deviation of data in the annulus. Moreover, we additinonally added the maximum value of the annulus to a pixel which is chosen with a probability 10^{-4}

$$M_{\text{noisy}} = M + U\left(-\frac{\sigma_i}{4}, \frac{\sigma_i}{4}\right) + \begin{cases} \max(M_i) & \text{if the pixel is chosen} \\ & \text{with probability } p = 10^{-4}, \\ 0 & \text{otherwise.} \end{cases} \tag{B.14}$$

In Fig. B.11, we can we how the dataset changes as a visual.

We tested the noisy β-Pictoris dataset using the ROC curves. First, we injected fake exoplanets with 1σ as we did with the clean dataset, but as shown in Fig. B.12 we got very low values in the ROC curves, although FoTAMAT gives slightly better results. Then we changed the flux of the injected planets to 1.5σ and we observed that FOTAMAT gives better results than AMAT, which supports our explanation of why we got different results in the β-Pictoris and 51 Eri datasets.

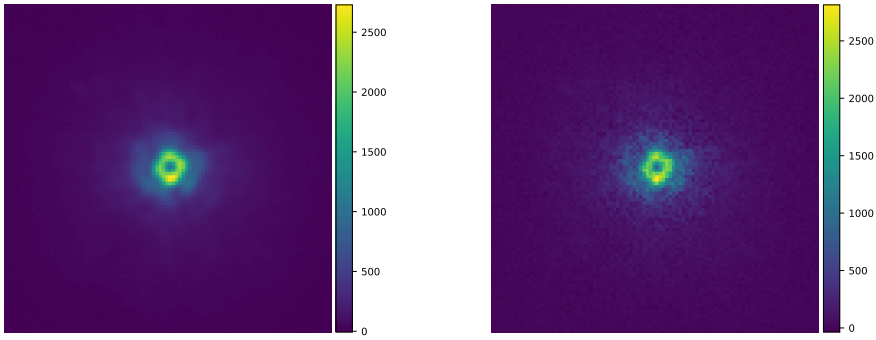

Fig. B.11 An example frame from clean (left) and noisy (right) data.

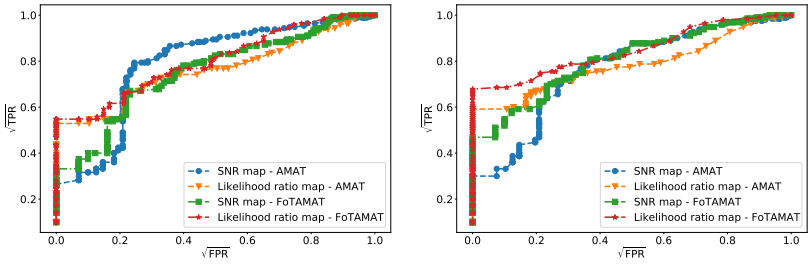

Fig. B.12 Comparison of the algorithms on noisy data by injecting fake exoplanets with flux levels of 1σ (left) and 1.5σ (right).

B.4 Conclusion

In this chapter, we examined the benefits of applying a FFT as a preprocessing step in the AMAT algorithm for exoplanet detection. Building on the success of the Fourier transform in mitigating noise, we demonstrated how applying FFT to noisy datasets improves the performance of the algorithm and prevents potential misinterpretations that could arise from unfiltered noise. We investigated how this preprocessing step enhances the overall success of the exoplanet detection process.

By comparing the results obtained through PCA and AMAT, we showed that FFT positively impacts visual performance in noise reduction. Furthermore, when conducting a broader comparison using ROC curves, we found that in datasets containing significant noise, FFT yielded more accurate results than both PCA and AMAT. This reinforces the value of

FFT as a preprocessing step in enhancing detection accuracy within noisy datasets.

Bibliography

[AMM+13] Olivier Absil, Julien Milli, Dimitri Mawet, A-M Lagrange, Julien Girard, Gaël Chauvin, Anthony Boccaletti, Christian Delacroix, and Jean Surdej. Searching for companions down to 2 AU from β pictoris using the L'-band AGPM coronagraph on VLT/NACO. *Astronomy & Astrophysics*, 559:L12, 2013.

[AQ12] A. Amara and S. P. Quanz. PYNPOINT : An image processing package for finding exoplanets: PYNPOINT. *Monthly Notices of the Royal Astronomical Society*, 427(2):948–955, December 2012. doi:10.1111/j.1365-2966.2012.21918.x.

[BA15] Nicolas Boumal and P.-A. Absil. Low-rank matrix completion via preconditioned optimization on the grassmann manifold. *Linear Algebra and its Applications*, 475:200–239, 2015. URL: https://www.sciencedirect.com/science/article/pii/S0024379515001342, doi:https://doi.org/10.1016/j.laa.2015.02.027.

[Bab53] Horace W Babcock. The possibility of compensating astronomical seeing. *Publications of the Astronomical Society of the Pacific*, 65(386):229–236, 1953.

[BAD+01] IA Bond, F Abe, RJ Dodd, JB Hearnshaw, M Honda, J Jugaku, PM Kilmartin, A Marles, K Masuda, Y Matsubara, et al. Real-time difference imaging analysis of moa galactic bulge observations during 2000. *Monthly Notices of the Royal Astronomical Society*, 327(3):868–880, 2001.

[BD10] T. Blumensath and M.E. Davies. Normalized iterative hard thresholding: Guaranteed stability and performance. *IEEE*

Journal of Selected Topics in Signal Processing, 4(2):298–309, April 2010. doi:10.1109/JSTSP.2010.2042411.

[BFD+08] Jean-Luc Beuzit, Markus Feldt, Kjetil Dohlen, David Mouillet, Pascal Puget, Francois Wildi, Lyu Abe, Jacopo Antichi, Andrea Baruffolo, Pierre Baudoz, et al. Sphere: a planet finder instrument for the vlt. In *Ground-based and airborne instrumentation for astronomy II*, volume 7014, pages 476–487. SPIE, 2008.

[BGG+23] Markus J Bonse, Emily O Garvin, Timothy D Gebhard, Felix A Dannert, Faustine Cantalloube, Gabriele Cugno, Olivier Absil, Jean Hayoz, Julien Milli, Markus Kasper, et al. Comparing apples with apples: Robust detection limits for exoplanet high-contrast imaging in the presence of non-gaussian noise. *The Astronomical Journal*, 166(2):71, 2023.

[Bow16] Brendan P Bowler. Imaging extrasolar giant planets. *Publications of the Astronomical Society of the Pacific*, 128(968):102001, 2016.

[Bra00] Ronald N. Bracewell. *The Fourier Transform and Its Applications*. McGraw-Hill Education, 3rd edition, 2000.

[BSJ+17] Thierry Bouwmans, Andrews Sobral, Sajid Javed, Soon Ki Jung, and El-Hadi Zahzah. Decomposition into low-rank plus additive matrices for background/foreground separation: A review for a comparative evaluation with a large-scale dataset. *Computer Science Review*, 23:1–71, 2017.

[BVM+19] J-L Beuzit, Arthur Vigan, David Mouillet, Kjetil Dohlen, Raffaele Gratton, Anthony Boccaletti, J-F Sauvage, Hans Martin Schmid, Maud Langlois, Cyril Petit, et al. Sphere: the exoplanet imager for the very large telescope. *Astronomy & Astrophysics*, 631:A155, 2019.

[C+20] Faustine Cantalloube et al. Exoplanet imaging data challenge: benchmarking the various image processing methods for exoplanet detection. In *Adaptive Optics Systems VII*, volume 11448, pages 1027–1062. SPIE, 2020.

[CADVD23] Carles Cantero, Olivier Absil, C-H Dahlqvist, and Marc Van Droogenbroeck. Na-sodinn: A deep learning algo-

rithm for exoplanet image detection based on residual noise regimes. *Astronomy & Astrophysics*, 680:A86, 2023.

[CB21] George Casella and Roger L Berger. *Statistical inference*. Cengage Learning, 2021.

[CBLM99] David Charbonneau, Timothy M Brown, David W Latham, and Michel Mayor. Detection of planetary transits across a sun-like star. *The Astrophysical Journal*, 529(1):L45, 1999.

[CCA+19] Valentin Christiaens, Simon Casassus, Olivier Absil, Faustine Cantalloube, C Gomez Gonzalez, J Girard, R Ramírez, B Pairet, V Salinas, Daniel J Price, et al. Separating extended disc features from the protoplanet in pds 70 using vlt/sinfoni. *Monthly Notices of the Royal Astronomical Society*, 486(4):5819–5837, 2019.

[CCS08] Jian-Feng Cai, Emmanuel J. Candes, and Zuowei Shen. A singular value thresholding algorithm for matrix completion. 10 2008. URL: `http://arxiv.org/abs/0810.3286`, `arXiv:0810.3286`.

[CESV15] Emmanuel J Candes, Yonina C Eldar, Thomas Strohmer, and Vladislav Voroninski. Phase retrieval via matrix completion. *SIAM review*, 57(2):225–251, 2015.

[CFT+12] Thayne Currie, Misato Fukagawa, Christian Thalmann, Soko Matsumura, and Peter Plavchan. Direct detection and orbital analysis of the exoplanets hr 8799 bcd from archival 2005 keck/nirc2 data. *The Astrophysical Journal Letters*, 755(2):L34, 2012.

[CGF+23] Valentin Christiaens, Carlos Gonzalez, Ralf Farkas, Carl-Henrik Dahlqvist, Evert Nasedkin, et al. VIP: A Python package for high-contrast imaging. *The Journal of Open Source Software*, 8(81):4774, January 2023. `doi:10.21105/joss.04774`.

[CLD+05] Gaël Chauvin, A-M Lagrange, Ch Dumas, B Zuckerman, D Mouillet, I Song, J-L Beuzit, and P Lowrance. Giant planet companion to 2massw j1207334-393254. *Astronomy & Astrophysics*, 438(2):L25–L28, 2005.

[CMM+15] F Cantalloube, D Mouillet, LM Mugnier, J Milli, Olivier Absil, CA Gomez Gonzalez, G Chauvin, J-L Beuzit, and A Cornia. Direct exoplanet detection and characterization using the andromeda method: Performance on vlt/naco data. *Astronomy & Astrophysics*, 582:A89, 2015.

[CR12] Emmanuel Candes and Benjamin Recht. Exact matrix completion via convex optimization. *Communications of the ACM*, 55(6):111–119, 2012.

[DCA20] C-H Dahlqvist, Faustine Cantalloube, and Olivier Absil. Regime-switching model detection map for direct exoplanet detection in adi sequences. *Astronomy & Astrophysics*, 633:A95, 2020.

[DCY+14] Rong Du, Cailian Chen, Bo Yang, Ning Lu, Xinping Guan, and Xuemin Shen. Effective urban traffic monitoring by vehicular sensor networks. *IEEE transactions on Vehicular Technology*, 64(1):273–286, 2014.

[DEL22] A Lecavelier Des Etangs and Jack J Lissauer. The iau working definition of an exoplanet. *New Astronomy Reviews*, 94:101641, 2022.

[DGM+12] P Delorme, J Gagné, L Malo, C Reylé, E Artigau, L Albert, T Forveille, X Delfosse, F Allard, and D Homeier. Cfbdsir2149-0403: a 4–7 jupiter-mass free-floating planet in the young moving group ab doradus? *Astronomy & Astrophysics*, 548:A26, 2012.

[DK13] Gaspard Duchêne and Adam Kraus. Stellar multiplicity. *Annual Review of Astronomy and Astrophysics*, 51(1):269–310, 2013.

[DVA23] Hazan Daglayan, Simon Vary, and P-A Absil. An alternating minimization algorithm with trajectory for direct exoplanet detection. *Proceedings of the 31st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN 2023)*, pages 417–422, 2023. doi:10.14428/esann/2023.ES2023-137.

[DVC+22] Hazan Daglayan, Simon Vary, Faustine Cantalloube, P.-A. Absil, and Olivier Absil. Likelihood ratio map for direct exoplanet detection. In *2022 IEEE 5th International Conference on*

Image Processing Applications and Systems (IPAS), volume Five, pages 1–5, 2022. doi:10.1109/IPAS55744.2022.10052997.

[DVL+23] Hazan Daglayan, Simon Vary, Valentin Leplat, Nicolas Gillis, and Pierre-Antoine Absil. Direct exoplanet detection using l1 norm low-rank approximation. *Proceedings of BNAIC/BeNeLearn 2023*, page 1, 2023.

[EVDH10] Anders Eriksson and Anton Van Den Hengel. Efficient computation of robust low-rank matrix approximations in the presence of missing data using the l 1 norm. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 771–778. IEEE, 2010.

[FDTL18] Olivier Flasseur, Loïc Denis, Éric Thiébaut, and Maud Langlois. Exoplanet detection in angular differential imaging by statistical learning of the nonstationary patch covariances-the paco algorithm. *Astronomy & Astrophysics*, 618:A138, 2018.

[GAVD18] CA Gomez Gonzalez, Olivier Absil, and Marc Van Droogenbroeck. Supervised detection of exoplanets in high-contrast imaging sequences. *Astronomy & Astrophysics*, 613:A71, 2018.

[GG+16] C. A. Gomez Gonzalez et al. Low-rank plus sparse decomposition for exoplanet detection in direct-imaging ADI sequences: The LLSG algorithm. *Astronomy & Astrophysics*, 589:A54, May 2016. doi:10.1051/0004-6361/201527387.

[GG+17] C. A. Gomez Gonzalez et al. VIP : Vortex Image Processing Package for High-contrast Direct Imaging. *The Astronomical Journal*, 154(1):7, June 2017. doi:10.3847/1538-3881/aa73d7.

[GM23] Raphaël Galicher and Johan Mazoyer. Imaging exoplanets with coronagraphic instruments. *Comptes Rendus. Physique*, 24(S2):1–45, 2023.

[GP11] Nicolas Gillis and Robert J Plemmons. Dimensionality reduction, classification, and spectral mixture analysis using non-negative underapproximation. *Optical Engineering*, 50(2):027001–027001, 2011.

[Guy18] Olivier Guyon. Extreme adaptive optics. *Annual Review of Astronomy and Astrophysics*, 56:315–355, 2018.

[GV18] Nicolas Gillis and Stephen A Vavasis. On the complexity of robust pca and ℓ_1-norm low-rank matrix approximation. *Mathematics of Operations Research*, 43(4):1072–1084, 2018.

[HMT11] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, January 2011. doi:10.1137/090771806.

[HZY+12] Yao Hu, Debing Zhang, Jieping Ye, Xuelong Li, and Xiaofei He. Fast and accurate matrix completion via truncated nuclear norm regularization. *IEEE transactions on pattern analysis and machine intelligence*, 35(9):2117–2130, 2012.

[Kai94] Gerald Kaiser. *A Friendly Guide to Wavelets*. Birkhäuser, 1994.

[KBQ+21] Sven Kiefer, Alexander J Bohn, Sascha Patrick Quanz, M Kenworthy, and Tomas Stolker. Spectral and angular differential imaging with sphere/ifs-assessing the performance of various pca-based approaches to psf subtraction. *Astronomy & Astrophysics*, 652:A33, 2021.

[KBV09] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.

[KK05] Qifa Ke and T. Kanade. Robust l_1 norm factorization in the presence of outliers and missing data by alternative convex programming. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 739–746 vol. 1, 2005. doi:10.1109/CVPR.2005.309.

[KNKS21] Sunwoo Kim, Luong T Nguyen, Junhan Kim, and Byonghyo Shim. Deep learning based low-rank matrix completion for iot network localization. *IEEE Wireless Communications Letters*, 10(10):2115–2119, 2021.

[KPP01] Jeffrey R Kuhn, D Potter, and Berengere Parise. Imaging polarimetric observations of a new circumstellar disk system. *The Astrophysical Journal*, 553(2):L189, 2001.

[KTJS03] Maciej Konacki, Guillermo Torres, Saurabh Jha, and Dimitar D Sasselov. An extrasolar planet that transits the disk of its parent star. *Nature*, 421(6922):507–509, 2003.

[LAJ22] Benoît Loucheur, Pierre-Antoine Absil, and Michel Journée. Gap filling in air temperature series by matrix completion methods. In *Proceedings of the 30th European Symposium on Artifical Neural Networks, Computational Intelligence and Machine Learning (ESANN2022)*, pages 369–374, 2022.

[LB10] Kiryung Lee and Yoram Bresler. Admira: Atomic decomposition for minimum rank approximation. *IEEE Transactions on Information Theory*, 56(9):4402–4416, 2010.

[LBC+10] A-M Lagrange, M Bonnefoy, G Chauvin, D Apai, D Ehrenreich, A Boccaletti, D Gratadour, D Rouan, D Mouillet, S Lacour, et al. A giant planet imaged in the disk of the young star β pictoris. *Science*, 329(5987):57–59, 2010.

[LMD+07] David Lafreniere, Christian Marois, René Doyon, Daniel Nadeau, and Étienne Artigau. A new algorithm for point-spread function subtraction in high-contrast imaging: a demonstration with angular differential imaging. *The Astrophysical Journal*, 660(1):770, 2007.

[Lyo39] Bernard Lyot. The study of the solar corona and prominences without eclipses (george darwin lecture, 1939). *Monthly Notices of the Royal Astronomical Society, Vol. 99, p. 580*, 99:580, 1939.

[MCS+09] Laurent M Mugnier, Alberto Cornia, Jean-François Sauvage, Gérard Rousset, Thierry Fusco, and Nicolas Védrenne. Optimal method for exoplanet detection by angular differential imaging. *JOSA A*, 26(6):1326–1334, 2009.

[MGI+14] Bruce Macintosh, James R Graham, Patrick Ingraham, Quinn Konopacky, Christian Marois, Marshall Perrin, Lisa Poyneer, Brian Bauman, Travis Barman, Adam S Burrows, et al. First light of the gemini planet imager. *proceedings of the National Academy of Sciences*, 111(35):12661–12666, 2014.

[MGL+09] Frantz Martinache, Olivier Guyon, Julien Lozi, Vincent Garrel, Célia Blain, and Gaetano Sivo. The subaru coronagraphic extreme ao project, 2009.

[MJD09] Raghu Meka, Prateek Jain, and Inderjit S. Dhillon. Guaranteed rank minimization via singular value projection. 09 2009. URL: http://arxiv.org/abs/0909.5457, arXiv:0909.5457.

[MLD+06] Christian Marois, David Lafreniere, René Doyon, Bruce Macintosh, and Daniel Nadeau. Angular differential imaging: a powerful high-contrast imaging technique. *The Astrophysical Journal*, 641(1):556, 2006.

[MLMD08] Christian Marois, David Lafreniere, Bruce Macintosh, and René Doyon. Confidence level and sensitivity limits in high-contrast imaging. *The Astrophysical Journal*, 673(1):647, 2008.

[MMV10] Christian Marois, Bruce Macintosh, and Jean-Pierre Véran. Exoplanet imaging with loci processing: photometry and astrometry with the new sosie pipeline. In *Adaptive Optics Systems II*, volume 7736, pages 595–606. SPIE, 2010.

[MMW+14] Dimitri Mawet, Julien Milli, Zahed Wahhaj, Didier Pelat, Olivier Absil, Christian Delacroix, Anthony Boccaletti, Markus Kasper, Matthew Kenworthy, Christian Marois, et al. Fundamental limitations of high contrast imaging set by small sample statistics. *The Astrophysical Journal*, 792(2):97, 2014.

[MQ95] Michel Mayor and Didier Queloz. A jupiter-mass companion to a solar-type star. *nature*, 378(6555):355–359, 1995.

[MV17] J Maldonado and E Villaver. Searching for chemical signatures of brown dwarf formation. *Astronomy & Astrophysics*, 602:A38, 2017.

[NAS24] NASA Exoplanet Archive. Planetary systems, 2024. URL: https://catcopy.ipac.caltech.edu/dois/doi.php?id=10.26133/NEA12, doi:10.26133/NEA12.

[NKKS19] Luong Trung Nguyen, Junhan Kim, Sangtae Kim, and Byonghyo Shim. Localization of iot networks via low-rank matrix completion. *IEEE Transactions on Communications*, 67(8):5833–5847, 2019.

[NKS19] Luong Trung Nguyen, Junhan Kim, and Byonghyo Shim. Low-rank matrix completion: A contemporary survey. *IEEE Access*, 7:94215–94237, 2019.

[OWH97] Alan V. Oppenheim, Alan S. Willsky, and S. Nawab Hamid. *Signals and Systems*. Prentice Hall, 2nd edition, 1997.

[PCGG+19] Benoît Pairet, Faustine Cantalloube, Carlos A Gomez Gonzalez, Olivier Absil, and Laurent Jacques. Stim map: detection map for exoplanets imaging beyond asymptotic gaussian residual speckle noise. *Monthly Notices of the Royal Astronomical Society*, 487(2):2262–2277, 2019.

[PCJ21] B. Pairet, F. Cantalloube, and L. Jacques. MAYONNAISE: A morphological components analysis pipeline for circumstellar discs and exoplanets imaging in the near-infrared. *Monthly Notices of the Royal Astronomical Society*, 503(3):3724–3742, April 2021. doi:10.1093/mnras/stab607.

[PCV+12] Laurent Pueyo, Justin R Crepp, Gautam Vasisht, Douglas Brenner, Ben R Oppenheimer, Neil Zimmerman, Sasha Hinkley, Ian Parry, Charles Beichman, Lynne Hillenbrand, et al. Application of a damped locally optimized combination of images method to the spectral characterization of faint companions using an integral field spectrograph. *The Astrophysical Journal Supplement Series*, 199(1):6, 2012.

[Pue16] Laurent Pueyo. Detection and characterization of exoplanets using projections on karhunen–loeve eigenimages: Forward modeling. *The Astrophysical Journal*, 824(2):117, 2016.

[RCL+12] J Rameau, G Chauvin, A-M Lagrange, P Thébault, J Milli, JH Girard, and M Bonnefoy. High-contrast imaging of the close environment of hd 142527-vlt/naco adaptive optics thermal and angular differential imaging. *Astronomy & Astrophysics*, 546:A24, 2012.

[RFK+90] Gérard Rousset, Jean-Claude Fontanella, Pierre Kern, Pierre Gigan, and Francois Rigaut. First diffraction-limited astronomical images with adaptive optics. *Astronomy and Astrophysics (ISSN 0004-6361), vol. 230, no. 2, April 1990, p. L29-L32.*

Research supported by the European Southern Observatory, Ministere de la Recherche et de la Technologie, Ministere de l'Education Nationale, INSU, DRET, and Ministere de la Defense., 230:L29–L32, 1990.

[RMW+17] Jean-Baptiste Ruffio, Bruce Macintosh, Jason J Wang, Laurent Pueyo, Eric L Nielsen, Robert J De Rosa, Ian Czekala, Mark S Marley, Pauline Arriaga, Vanessa P Bailey, et al. Improving and assessing planet sensitivity of the gpi exoplanet survey with a forward model matched filter. *The Astrophysical Journal*, 842(1):14, 2017.

[RPZ+18] Bin Ren, Laurent Pueyo, Guangtun Ben Zhu, John Debes, and Gaspard Duchêne. Non-negative matrix factorization: robust extraction of extended structures. *The Astrophysical Journal*, 852(2):104, 2018.

[RWM+15] Timothy J Rodigas, Alycia Weinberger, Eric E Mamajek, Jared R Males, Laird M Close, Katie Morzinski, Philip M Hinz, and Nathan Kaib. Direct exoplanet detection with binary differential imaging. *The Astrophysical Journal*, 811(2):157, 2015.

[RWN+99] René Racine, Gordon AH Walker, Daniel Nadeau, René Doyon, and Christian Marois. Speckle noise and the detection of faint companions. *Publications of the Astronomical Society of the Pacific*, 111(759):587, 1999.

[SBM11] David S Spiegel, Adam Burrows, and John A Milsom. The deuterium-burning mass limit for brown dwarfs and giant planets. *The Astrophysical Journal*, 727(1):57, 2011.

[SDLS+11] Jean Schneider, Cyrill Dedieu, Pierre Le Sidaner, Renaud Savalle, and Ivan Zolotukhin. Defining and cataloging exoplanets: the exoplanet. eu database. *Astronomy & Astrophysics*, 532:A79, 2011.

[SF02] William B Sparks and Holland C Ford. Imaging spectroscopy for extrasolar planet detection. *The Astrophysical Journal*, 578(1):543, 2002.

[SFP+16] Jean-Francois Sauvage, Thierry Fusco, Cyril Petit, Anne Costille, David Mouillet, Jean-Luc Beuzit, Kjetil Dohlen, Markus

Kasper, Marcos Suarez, Christian Soenke, et al. Saxo: the extreme adaptive optics system of sphere (i) system overview and global laboratory performance. *Journal of Astronomical Telescopes, Instruments, and Systems*, 2(2):025003–025003, 2016.

[SKU+11] T. Sumi, K. Kamiya, A. Udalski, D.P. Bennett, I.A. Bond, F. Abe, C.S. Botzler, A. Fukui, K. Furusawa, J.B. Hearnshaw, Y. Itow, P.M. Kilmartin, A. Korpela, W. Lin, C.H. Ling, K. Masuda, Y. Matsubara, N. Miyake, M. Motomura, Y. Muraki, M. Nagaya, S. Nakamura, K. Ohnishi, T. Okumura, Y.C. Perrott, N. Rattenbury, To. Saito, T. Sako, D.J. Sullivan, W.L. Sweatman, P.J. Tristram, P.C.M. Yock, M.K. Szymanski, M. Kubiak, G. Pietrzynski, R. Poleski, I. Soszynski, L. Wyrzykowski, and K. Ulaczyk. Unbound or distant planetary mass population detected by gravitational microlensing. *Nature*, 473(7347):349–352, May 2011. URL: `http://dx.doi.org/10.1038/nature10092`, `doi:10.1038/nature10092`.

[SPL12] R. Soummer, L. Pueyo, and J. Larkin. Detection and characterization of exoplanets and disks using projections on Karhunen-Loève eigenimages. *The Astrophysical Journal*, 755(2):L28, August 2012. `doi:10.1088/2041-8205/755/2/L28`.

[Str52] Otto Struve. Proposal for a project of high-precision stellar radial velocity work. *The Observatory, Vol. 72, p. 199-200 (1952)*, 72:199–200, 1952.

[TV23] J. Tanner and S Vary. Compressed sensing of low-rank plus sparse matrices. *Applied and Computational Harmonic Analysis*, 64:254–293, 2023. `doi:https://doi.org/10.1016/j.acha.2023.01.008`.

[TW16] Jared Tanner and Ke Wei. Low rank matrix completion by alternating steepest descent methods. *Applied and Computational Harmonic Analysis*, 40(2):417–429, 2016.

[Van13] Bart Vandereycken. Low-rank matrix completion by riemannian optimization. *SIAM Journal on Optimization*, 23(2):1214–1236, 2013.

[VDJA23] Simon Vary, Hazan Daglayan, Laurent Jacques, and Pierre-Antoine Absil. Low-rank plus sparse trajectory decomposition for direct exoplanet imaging. *arXiv preprint arXiv:2301.07018*, 2023. in press.

[VGO+20] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ Polat, Y Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi:10.1038/s41592-019-0686-2.

[WAG+17] Olivier Wertz, Olivier Absil, CA Gómez González, Julien Milli, Julen H Girard, Dimitri Mawet, and Laurent Pueyo. Vlt/sphere robust astrometry of the hr8799 planets at milliarcsecond-level accuracy-orbital architecture analysis with pyastrofit. *Astronomy & Astrophysics*, 598:A83, 2017.

[WCM+15] Zahed Wahhaj, Lucas A Cieza, Dimitri Mawet, Bin Yang, Hector Canovas, Jozua de Boer, Simon Casassus, François Ménard, Matthias R Schreiber, Michael C Liu, et al. Improving signal-to-noise in the direct imaging of exoplanets and circumstellar disks with mloci. *Astronomy & Astrophysics*, 581:A24, 2015.

[WF92] Aleksander Wolszczan and Dail A Frail. A planetary system around the millisecond pulsar psr1257+ 12. *Nature*, 355(6356):145–147, 1992.

[WLL+14] Zheng Wang, Ming-Jun Lai, Zhaosong Lu, Wei Fan, Hasan Davulcu, and Jieping Ye. Rank-one matrix pursuit for matrix completion. In *International Conference on Machine Learning*, pages 91–99. PMLR, 2014.

[WMR+21] Z Wahhaj, J Milli, C Romero, L Cieza, A Zurlo, Arthur Vigan, E Peña, G Valdes, F Cantalloube, J Girard, et al. A search

for a fifth planet around hr 8799 using the star-hopping rdi technique at vlt/sphere. *Astronomy & Astrophysics*, 648:A26, 2021.

[WYZ12] Zaiwen Wen, Wotao Yin, and Yin Zhang. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. *Mathematical Programming Computation*, 4(4):333–361, 2012.

[XZC17] Hongyang Xue, Shengming Zhang, and Deng Cai. Depth image inpainting: Improving low rank matrix completion with low gradient regularization. *IEEE Transactions on Image Processing*, 26(9):4311–4320, 2017.

[Zhu16] Guangtun Zhu. Nonnegative matrix factorization (nmf) with heteroscedastic uncertainties and missing data. *arXiv preprint arXiv:1612.06037*, 2016.

[ZLS+12] Yinqiang Zheng, Guangcan Liu, Shigeki Sugimoto, Shuicheng Yan, and Masatoshi Okutomi. Practical low-rank matrix approximation under robust l 1-norm. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1410–1417. IEEE, 2012.

[ZT11] Tianyi Zhou and Dacheng Tao. Godec: Randomized low-rank & sparse matrix decomposition in noisy case. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*, 2011.