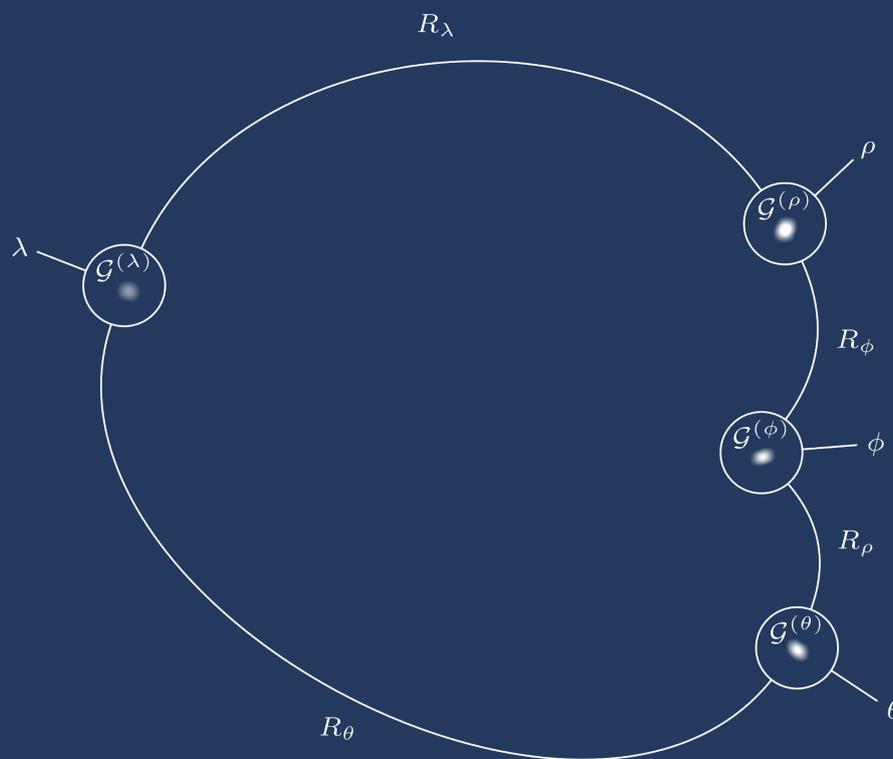
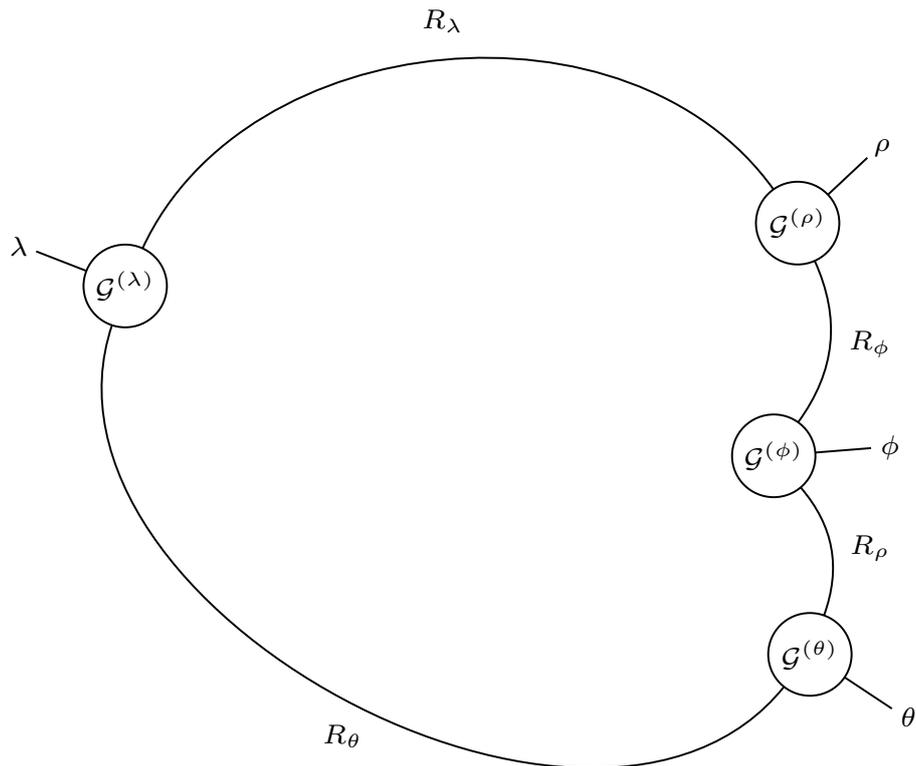


Direct Exoplanet Imaging with Tensor Decompositions

Lukas Welzel





The cover page shows the exoplanets HR 8799 b, c, d and e behind a tensor network diagram of the annular tensor ring decomposition used to reduce the VLT/IRDIS observations of the system. Each exoplanet is placed inside one of the order-3 factors of the tensor ring. The angular-spectral differential imaging technique was used for the observations of HR 8799 from October 2017. These observations were part of the SHINE survey (PI Gael Chauvin).

Direct Exoplanet Imaging with Tensor Decompositions

Thesis report

by

Lukas Welzel

to obtain the degree of Master of Science
at Leiden University
&
the degree of Master of Science
at the Delft University of Technology
to be defended publicly on October 16, 2024 at 09:00

Thesis committee:

Chair: Dr. Stéphanie Cazaux

Supervisors: Dr. Tomas Stolker
Dr. Jérôme Loicq

External examiner: Dr. Sijme-Jan Paardekooper

Place: Lecture Hall H, Faculty of Aerospace Engineering, Delft

Project Duration: November, 2023 - October, 2024

Student number: 4551486

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Preface

Completing this thesis has been both a challenging and rewarding journey, marking the culmination of my double master's degree in Aerospace Engineering at the Faculty of Aerospace Engineering, Delft University of Technology, and Astronomy at Leiden Observatory, Faculty of Science, Leiden University. The interdisciplinary nature of this project - bridging space engineering and astronomical sciences - presented unique organizational and academic challenges that tested my resilience and adaptability.

I am profoundly grateful to my supervisors, Tomas and Jérôme, for your unwavering support and enthusiasm. Your willingness to embrace a project that I defined, and in Jérôme's case, one that extended beyond your primary field of research, was instrumental in bringing this work to fruition. Your insightful guidance and encouragement at every stage have been invaluable; I could not have asked for better mentors.

In the early stages, when the viability of this idea was still uncertain, Rico's support was pivotal. Your belief in the potential of this work gave me the confidence to pursue it wholeheartedly. Similarly, Matt was a constant source of wisdom and encouragement, helping me find the right approach to this project and supporting me over the last year, just as you did during my first thesis.

I would also like to express my sincere appreciation to those who provided technical feedback and insights that enriched this work. Valentin Christiaens offered an open ear from the direct imaging community, providing valuable perspectives. Kim Batselier shared his expertise in tensor decompositions for signal processing, which was crucial to the development of the methods presented here. Rafael Ballester-Ripoll was always ready to assist with questions about tensor computation, helping me navigate this complex topic.

My friends have been a source of strength and positivity throughout this journey. Mark and Nick, your patience in listening to my thoughts and your unwavering support have been indispensable; I could not have done this without you. To all my friends who lent an open ear and helped me stay optimistic during the toughest times, thank you for being there.

I am deeply thankful to my parents, my brothers, for being a solid rock in the storm. Your love and support provided me with the foundation to persevere and achieve my goals.

"But it is one thing to read about dragons and another to meet them."
- Ursula K. Le Guin, *A Wizard of Earthsea*

Abstract

Direct imaging of exoplanets relies on advanced post-processing techniques to differentiate planetary signals from the bright stellar point spread function (PSF) and quasi-static speckle noise. While angular-spectral differential imaging (ASDI) with integral field spectrographs has significantly enhanced the capability to induce diversity in high-contrast imaging (HCI) observations, current post-processing methods like matrix Principal Component Analysis (PCA) are limited in their ability to fully exploit the multi-modal structure of ASDI data.

This thesis introduces tensor decomposition methods as generalizations of matrix-PCA for modeling the coronagraphic PSF in ASDI observations. These methods preserve the higher-order structure of ASDI data, enabling the modeling of complex cross-modal interactions while maintaining the strengths of PCA-based approaches. By extending PCA to higher-order tensors, these methods offer a more natural representation of the multi-dimensional nature of ASDI data.

The proposed tensor methods are evaluated using both synthetic and real observations from the SPHERE instrument on the VLT. They successfully recover known exoplanets in benchmark systems like HR 8799 and β Pictoris, and detect a new collision-induced feature in the AU Microscopii debris disk. Notably, a new sub-stellar companion candidate is identified in the HD 108767 B system. Their performance is assessed through both quantitative metrics and qualitative analysis of residual images. Results demonstrate that tensor methods are competitive with matrix-PCA, outperforming it at small angular separations where speckle noise is most problematic.

This work establishes tensor decompositions as a powerful tool for HCI post-processing, offering increased flexibility and interpretability of factorizations while building upon the established framework of PCA. Furthermore, this approach lays the groundwork for tractable deep learning techniques on HCI datasets, and more generally for higher-order tensor methods in direct exoplanet imaging.

Lay Summary

Directly photographing planets around other stars is extremely challenging due to the overwhelming brightness of the stars themselves. Astronomers use advanced image processing techniques to separate the faint light of planets from the intense glare of their host stars and other unwanted "noise" in the images.

This thesis introduces new mathematical techniques called "tensor decompositions" to process these high-contrast images. These methods improve upon current approaches like Principal Component Analysis (PCA) by better preserving the complex structure of astronomical data, which contains information across multiple dimensions such as color, time, and space.

The new techniques were tested using both simulated data and real observations from the SPHERE instrument on the Very Large Telescope. They successfully detected known planets in well-studied star systems like HR 8799 and β Pictoris, demonstrating their effectiveness. Importantly, they also enabled two new discoveries: A previously unseen structure in the dust disk surrounding the star AU Microscopii, likely caused by collisions of objects within the disk and a potential new planet-like object orbiting the star HD 108767 B.

When compared to current methods, these tensor-based techniques performed competitively overall and showed improved performance in detecting planets very close to their stars, where separating planet light from starlight is most difficult.

This research establishes tensor decompositions as a powerful new tool for processing images of distant planets. It offers astronomers more flexibility in analyzing their data and provides a clearer understanding of the results. The methods are particularly valuable because they build upon the established framework of PCA while extending its capabilities to handle more complex data structures. By preserving the multi-dimensional nature of the data, these methods create a foundation for more sophisticated analysis techniques, potentially improving other techniques which are used to search for and study worlds beyond our solar system.

In summary, this thesis introduces a new class of image processing methods that better handle the complexity of exoplanet imaging data, leading to improved planet detection capabilities and new astronomical discoveries.

Contents

Cover	i
Abstract	iii
Lay Summary	iv
Contents	vi
Nomenclature	viii
List of Figures	ix
List of Tables	xii
1 Introduction	1
1.1 Direct Imaging of Exoplanets	2
1.2 Research Objectives & Methodology	4
1.3 Contributions	5
2 Background	9
2.1 Exoplanets	9
2.1.1 Detection Methods	10
2.1.2 Formation	12
2.1.3 Characteristics	13
2.1.4 Directly Imaged Exoplanet Population	15
2.2 Direct Imaging of Exoplanets	16
2.2.1 Image Formation	17
2.2.2 Instrumentation	18
2.2.3 Speckle Pattern	20
2.2.4 Differential Imaging	22
2.2.5 Post-Processing	27
2.3 Multi-Linear Algebra	30
2.3.1 Tensors	31
2.3.2 Operations on Tensors	36
2.3.3 Tensor Network Notation	38
2.4 Tensor Decompositions	40
2.4.1 Canonical-Polyadic	41
2.4.2 Motivating Example	42
2.4.3 Tucker	44
2.4.4 Tensor-Train	47
2.4.5 Tensor-Ring	47
2.4.6 Other Decompositions	48
2.4.7 Computation	48
2.4.8 Limitations	50
2.5 Failure of Matrix-PCA on Higher-Order Tensors	50
2.6 Chapter Summary & Contributions	53
3 Multi-Modal PSF Modeling	54
3.1 CODI with matrix-PCA	55
3.2 CODI with Tensor Decompositions	55
3.2.1 Full-Frame PSF Model	56
3.2.2 Unfolding	59
3.2.3 Tensorization	59
3.2.4 Polar Full-Frame PSF Model	60

3.2.5	Annular PSF Model	63
3.2.6	Index Permutation	64
3.3	Implementation	66
3.4	Chapter Summary & Contributions	67
4	Instrumentation & Data	68
4.1	VLT/SPHERE	68
4.1.1	Coronagraph	68
4.1.2	Adaptive Optics	70
4.1.3	Speckle Pattern & Combined Adaptive Optics-Coronagraphic Imaging	71
4.2	Observations	72
4.2.1	SHINE	72
4.2.2	BEAST	73
4.2.3	DIVA	73
4.2.4	EIDC	73
4.3	Chapter Summary & Contributions	75
5	Results	76
5.1	Detection Metrics	76
5.2	Post-Processed Residuals	77
5.2.1	HR 8799	77
5.2.2	Beta Pictoris	79
5.2.3	Features in the AU Microscopii Debris Disk	81
5.2.4	HD 108767 B	82
5.3	Assessment	84
5.3.1	Verification	84
5.3.2	Receiver Operating Characteristics	84
5.3.3	Runtime Analysis	87
5.4	Chapter Summary & Contributions	89
6	Discussion	91
6.1	Detection Performance	91
6.2	Noise Characteristics	92
6.3	Theoretical Considerations	94
6.4	Implications of the HD 108767 B Candidate	95
6.5	Broader Impact on Computational Methods in Astronomy	95
6.6	Limitations & Challenges	96
6.7	Chapter Summary & Contributions	97
7	Conclusion	98
8	Recommendations	100
	Bibliography	120
A	Algorithms	121
B	Observation Data	123
B.1	HD 108767 B Observations	123
B.2	Assessment Systems & Observations	124
B.3	Full Observation Sample	126
C	Reductions	144
D	Receiver Operating Characteristics	147
E	Tools	151
F	TIDE Framework Implementation	152
G	High-Level Tensor Software	154

Nomenclature

List of Abbreviations

ADI	Angular Differential Imaging	LR	Likelihood Ratio
ANOVA	Analysis of Variance	LRA	Low-Rank Approximation
ASDI	Angular-Spectral Differential Imaging	LRTA	Low-Rank Tensor Approximation
AUC	Area-Under-Curve	NMF	Non-negative matrix Factorization
BTD	Block-Term Decomposition	NTD	Non-Negative Tucker Decomposition
CODI	Combined Differential Imaging	NTF	Non-Negative Tensor Factorization
CPD	Canonical-Polyadic Decomposition	OWA	Outer Working Angle
EIDC	Exoplanet Imaging Data Challenge	PC	Principal Component
ESO	European Southern Observatory	PCA	Principal Component Analysis
FN	False Negatives	PPV	Positive Prediction Value
FNR	False Negative Rate	PRC	Precision-Recall
FOV	Field of View	PSF	Point Spread Function
FP	False Positive Rate	RDI	Reference Star Differential Imaging
FPR	False Positive Rate	ROC	Receiver Operating Characteristics
HC-DC	High Contrast Data Center	ROI	Region of Interest
HCI	High-Contrast Imaging	RPCA	Robust Principal Component Analysis
HOI	Higher Order Orthogonal Iteration	SDI	Spectral Differential Imaging
HOSVD	Higher-Order Singular Value Decomposition	SNR	Signal to Noise Ratio
Hot-SVD	Higher-Order Tensor Singular Value Decomposition	SPHERE	Spectro-Polarimetric High-contrast Exoplanet REsearch (instrument)
HSI	Hyper-Spectral Imaging	SVD	Singular Value Decomposition
HTD	Hierarchical Tucker Decomposition	t-SVD	truncated Singular Value Decomposition
IFS	Integral Field Spectrograph	TD	Tucker Decomposition
IRDIS	Infra-Red Dual Imaging Spectrograph	TN	True Negatives
IWA	Inner Working Angle	TND	Tensor Network Diagrams
KLIP	Karhunen-Loève Image Projection	TNR	True Negative Rate
LLSG	Local Low-rank Plus Sparse Plus Gaussian-Noise Decomposition	TP	True Positives
		TPR	True Positive Rate
		TR	Tensor Ring

TRD Tensor Ring Decomposition
 TT Tensor Train
 TTD Tensor Train Decomposition
 VLT Very Large Telescope

List of Symbols

λ Wavelength

$\langle \mathcal{A}, \mathcal{B} \rangle$ Tensor inner product

A Matrix

a Vector

A * B Matrix Hadamard product

A \odot B Matrix Khatri–Rao product

A \otimes B Matrix Kronecker product

\mathcal{A} Tensor (with order ≥ 3)

$\mathcal{A} \bullet \mathcal{B}$ tensor-tensor product

Tr Trace operator

ϕ Azimuthal angle

ψ Tensor model parameters

ρ Radius

θ Parallax angle

$\varrho \in \rho$ Radius of annulus

a Scalar

D Aperture diameter

\star Convolution operator

\times_n Mode- n (contracted) tensor product

$R_{\lambda_0}^T$ Telescope response at wavelength λ_0

List of Figures

1	Introduction	
1.1	Post-processing of direct imaging observations.	1
1.2	Post-processing spectral differential imaging observations with matrix principal component analysis (PCA) using the truncated matrix singular value decomposition (SVD).	2
1.3	Post-processing spectral differential imaging observations with tensor methods using the Canonical-Polyadic Decomposition (CPD).	3
1.4	Thesis topics, structure and contributions.	8
2	Background	
2.1	High-multiplicity planetary systems and cumulative discoveries of exoplanets by discovery method.	10
2.2	Exoplanet mass versus plotted against their semi-major axis.	11
2.3	Low-resolution emission spectra of directly imaged exoplanets in the near-mid infrared by ground-based telescopes and their contrast ratio.	15
2.4	Directly imaged exoplanets with known semi-major axis, mass and equilibrium temperature.	16
2.5	Spreading out of a PSF with increasing wavelength.	18
2.6	The four pillars of ground-based direct imaging of exoplanets.	19
2.7	Normalized PSF of a circular pupil telescope.	20
2.8	Detection limits for current and planned observatories.	21
2.9	Lenslet array-type IFS.	22
2.10	Unprocessed lenslet array-type IFS detector data and the definition of the IWA and OWA.	22
2.11	The primary components in high-contrast imaging observations of exoplanets	23
2.12	Modeling and subtracting stationary components in observations obtained by differential imaging to obtain a model of the non-stationary components in observations.	24
2.13	Aberrations in a simulated VLT/SPHERE analogue using an apodized pupil Lyot coronagraph.	25
2.14	Observations obtained by angular-spectral differential imaging and their angular and spectral diversity.	26
2.15	Post-processing angular spectral differential imaging observations.	28
2.16	Fibers and slices in a third-order tensor generalize columns (or rows) of matrices by partially fixing indices.	32
2.17	Example tensor.	34
2.18	Diagonal tensors.	34
2.19	Graphic example of a mode- n unfolding.	38
2.20	Vector, Matrix, and order-3 tensor.	39
2.21	Order- N tensor and superdiagonal tensor.	39
2.22	Graphic representation of the canonical-polyadic decomposition.	41
2.23	Mean combined residuals from a rank-1 CPD of the residual spectral cube of β Pictoris.	42
2.24	A stationary point source with spectrum $f(\lambda) \approx \mathbf{u}^{(\lambda)}$ in a spectral-spatial data cube approximated by a rank-1 CPD.	43
2.25	Factorizing a residual spectral cube of β Pictoris b using a rank-1 CPD.	44
2.26	Vector space representation of approximating a tensor with the CPD.	45
2.27	Vector space representation of approximating a tensor with the Tucker decomposition.	46
2.28	Approximating a third-order tensor with a Tucker decomposition.	46
2.29	Computing the factor matrices $\mathbf{U}^{(n)}$ for the Tucker decomposition of an order-3 tensor using the HOSVD.	49
2.30	Basis vectors $\mathbf{v}^{(I)}$ (red), $\mathbf{v}^{(J)}$ (green) and $\mathbf{v}^{(K)}$ (blue) as well as the corresponding factors of the CPD (orange, light-green, cyan).	51
2.31	6^{th} lateral slice ($I \times K$) of the ground truth \mathcal{X} , observed data \mathcal{Y} and models $\hat{\mathcal{Y}}$ by PCA with 1 PC, 3 PC and CPD.	52

3 Multi-Modal PSF Modeling	
3.1 Computational graph of CODI.	54
3.2 Normalized factors of a rank-5 CPD PSF model.	57
3.3 Frontal slices (λ_0, θ_0) of CPD PSF models of GPI observations.	58
3.4 Tensorizing a periodic signal vector into a matrix.	60
3.5 Factors of a rank-1 CPD an Airy pattern.	61
3.6 Monochromatic diffraction pattern at the focal plane of a telescope with a circular pupil and no secondary obstruction.	62
3.7 Reconstructed images rank-1 CPD Airy pattern models in Cartesian and polar coordinates.	62
3.8 Rank-1 CPD Airy pattern model and the theoretic Airy function discretized to the image grid.	65
4 Instrumentation & Data	
4.1 Ideal residual noise levels for SPHERE's IRDIS and IFS as well as the SHINE, BEAST, and GPIES samples.	69
4.2 Apodized Pupil Lyot Coronagraph architecture of VLT/SPHERE.	70
4.3 Correlation between coronagraphic raw contrast and coherence time.	71
4.4 SHINE survey depth and SPHERE detection probabilities.	72
5 Results	
5.1 HR 8799 IFS residual images from full-frame algorithms.	78
5.2 HR 8799 IFS residual images from annular algorithms.	79
5.3 Annular CPD using robust power iteration on HR 8799 IFS observations.	80
5.4 TRD on IFS observations of HR 8799	80
5.5 HR 8799 IRDIS residual images from algorithms on observations in polar coordinates.	81
5.6 β Pictoris (full frame).	82
5.7 New feature in AU Microscopii and HD 108767 B candidate contrast curve.	83
5.8 Residual images from IFS observations of HD 108767 B reduced with annular TTD.	83
5.9 Joint fit of the position and spectrum of the source in the HD 108767 B system for the 2016 observations.	84
5.10 Evaluation metrics for the ROC AUC and F_1 comparisons across various tensor decomposition methods and PCA.	85
5.11 The ROC curves of the top-performing algorithms for each class.	86
5.12 Comparative STIM maps for HR 8799 using the Tucker decomposition.	87
5.13 Comparative analysis of runtime efficiency and performance across various tensor decomposition methods and PCA for HCI.	88
6 Discussion	
6.1 True positive rate of best performing algorithms from each group for sources with separations smaller than a specific radius.	92
6.2 Distribution of residuals for β Pictoris reduced with the TTD (full frame), showing the difference between pixels at the location of β Pictoris b and residual coronagraphic PSF flux.	93
6.3 Weights of the CPD PSF models showing inflection point.	94
8 Recommendations	
8.1 Tucker decomposed molecule map. Since the factors are physically meaningful, structure can be imposed on the decomposition: the spectral factors are template spectra and the spatial components model the PSF intensity distribution.	101
8.2 Rotation and projection of spatial factor matrices of the CPD to approximate the (affine) rotation of a matrix.	102
8.3 Directional derivative of residual components.	104
C Reductions	
C.1 HR 8799 IFS residual images from annular TRD.	144

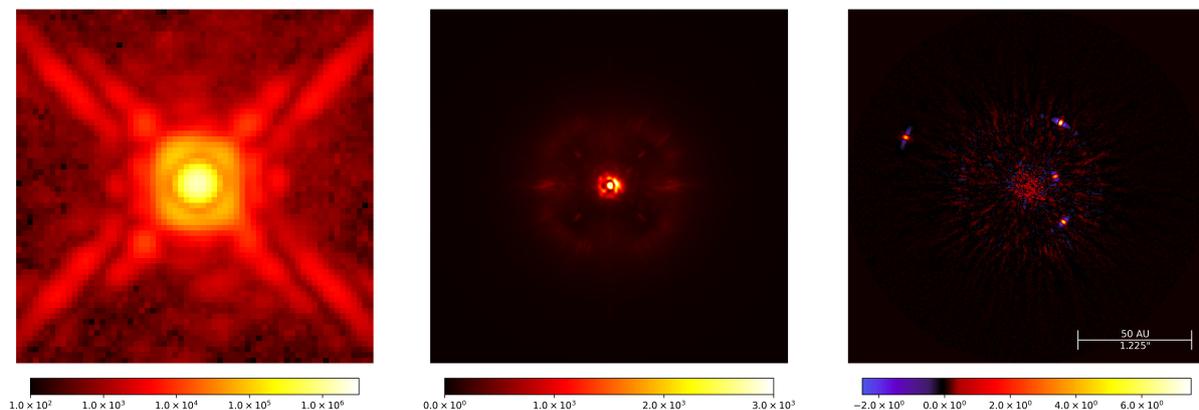
C.2	β Pictoris residual images from algorithms using a single annulus centered on β Pictoris b. The residuals are very similar for all methods using spatial unfolding.	145
C.3	Fast moving features in the AU Mic. disk in IFS observations. These features were originally detected by Boccaletti et al., 2015 and Boccaletti et al., 2018. Reductions with tensor methods very clearly detect the features.	146
C.4	STIM maps of IFS observations of HD 108767 B reduced with annular TTD.	146
D Receiver Operating Characteristics		
D.1	PCA: Extended receiver operating characteristic curves.	147
D.2	CPD: Extended receiver operating characteristic curves.	148
D.3	Tucker decomposition: Extended receiver operating characteristic curves.	149
D.4	Annular TTD: Extended receiver operating characteristic curves.	150
E Tools		
E.1	td-viz drawing of a rank-[36, 150, 150] tensor train decomposition of an order-3 tensor. . . .	151

List of Tables

B.1	HD 108767 B: Publicly available ESO observations.	123
B.2	Objects used for synthetic injections for assessing the performance of algorithms.	124
B.3	Full sample of observations processed in this thesis.	126
G.1	Overview of relevant tensor processing and multi-linear algebra software.	154

Introduction

Exoplanets, planets orbiting stars other than the Sun, have become a central focus in astrophysics. Direct imaging of exoplanets, a technique that isolates planetary light from its host star, provides crucial insights into planetary system formation and evolution. This method uniquely allows for direct spectroscopic characterization of exoplanetary atmospheres, their orbital dynamics, and interactions with circumstellar disks. It yields data that is complementary to and independent from other detection techniques. Direct imaging excels at detecting and characterizing wide-separation companions. This capability is essential for comprehensively mapping the architecture and evolution of exoplanetary systems and for probing the physical properties of young, self-luminous planets. Direct imaging, together with other exoplanet detection methods, is key to advancing our understanding of planetary systems and their formation. (Currie et al., 2023)



(a) Observation without direct imaging instrumentation. The image is dominated by the bright star in the center and the ring-like diffraction pattern due to the circular telescope pupil. The elongated streaks at 45° angles are caused by star light diffracting around the support structure of the telescope's secondary mirror (spider vanes). Note that the color is in log-scale.

(b) Observation with direct imaging instrumentation. The star (center) is partially masked but the image is still dominated by its diffracted light and atmospheric effects. The larger ring-shaped structure is due to the imperfect correction of atmospheric turbulence by the adaptive optics system. Four artificial guidance stars are visible as dim streaks at 45° angles.

(c) Post-processed image after subtracting the stellar light and speckle pattern, revealing the presence of four exoplanets (bright dots). Residual noise and a remaining speckle pattern is also visible, particularly prominent at small separations near the center of the image.

Figure 1.1: Direct imaging and post-processing of observations of the HR 8799 system obtained by the IRDIS instrument on VLT/SPHERE. The contrast between the host star and the exoplanets is roughly 10^6 and the distance between the inner-most exoplanet (HR 8799 e) and the star is ~ 16.25 AU or ~ 397.6 mas. This work addresses the step between Figure 1.1b and Figure 1.1c. The residual image (right) shown here was obtained using a method developed in this work. This method uses the Tucker decomposition to model and remove the stellar flux from observations like Figure 1.1b. The color-scale is in detector units of counts per wavelength band pixel per second.

Direct imaging of exoplanets faces significant challenges, primarily due to the extreme contrast between the bright star and the faint planet. To address these challenges, various post-processing methods have been developed. This work develops novel tensor-based methods which extend the traditional matrix-based techniques. (Currie et al., 2023)

1.1. Direct Imaging of Exoplanets

Extracting the signal of potential companions from the combined image of the star and companions presents significant challenges in direct imaging. The extreme contrast ratios between stellar and companion flux, often reaching 10^6 to 10^9 , make it difficult to distinguish faint companion signals from the bright stellar background. Furthermore, the quasi-static speckle pattern arising from instrumental and atmospheric effects further complicates the isolation of companion signals. [Figure 1.1](#) illustrates how direct imaging instrumentation and post-processing are used together to find exoplanets.

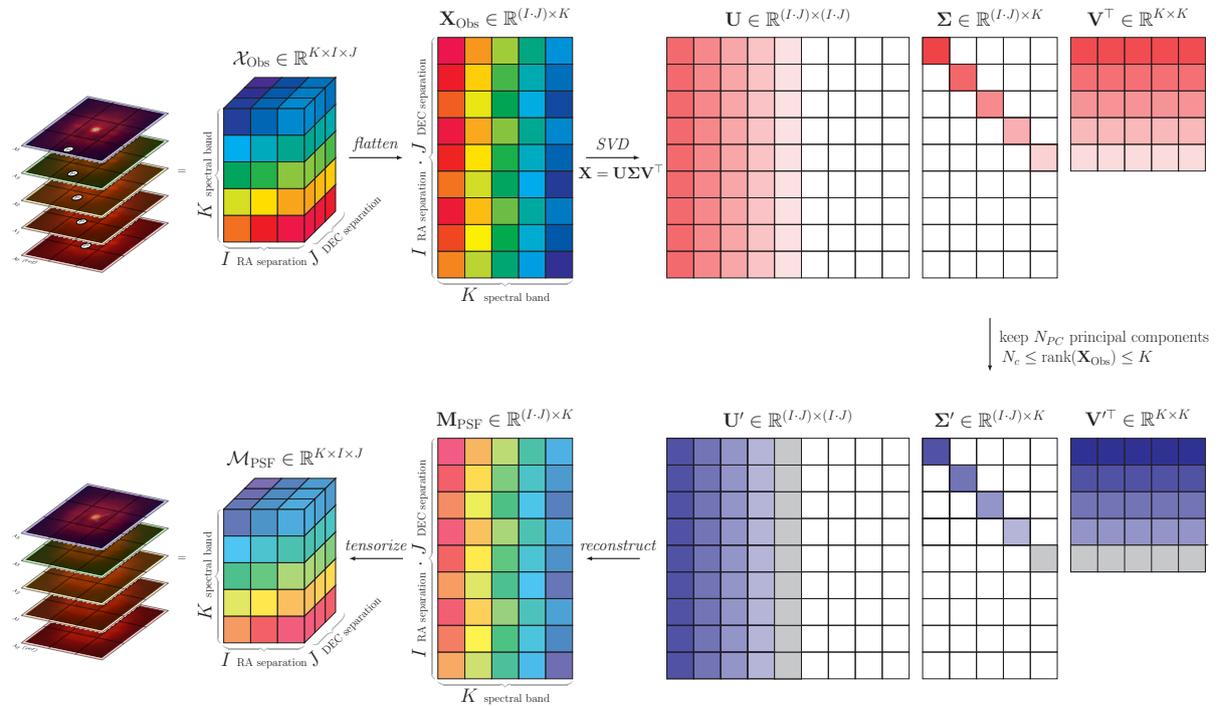


Figure 1.2: Post-processing spectral differential imaging observations with matrix principal component analysis (PCA) using the truncated matrix singular value decomposition (SVD). First the observations are prepared by aligning the coronagraphic point spread function (PSF) in the observations \mathcal{X}_{Obs} . This misaligns the planet (P). In this schematic, the observations of the star and planet are frames with 3×3 pixels and cover 5 spectral bands from the optical (blue) to the infrared (red). The observations are thus a $5 \times 3 \times 3$ data cube. The observations are then flattened into a 9×5 matrix. This is necessary to use matrix-PCA to find a PSF model. Flattening the observations disrupts the relationships between the data points. The example in [Section 2.5](#) will illustrate why this is problematic. The flattened observation matrix \mathbf{X} is factorized using the SVD. This factorization is exact so that $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ holds. The columns of \mathbf{U} are the principal components (PC) and the rows of \mathbf{V}^T are the principal directions, ordered from most "important" (dark) to least "important" (light). The "importance" of each PC is given by its singular values on the diagonal of $\mathbf{\Sigma}$. The unimportant PC contain typically only noise and the planet. The white cells are zero. To model the PSF without including the planet, only the important PC are considered, thus truncating the matrices. In this example 4 of the 5 PC are included in the model. This is equivalent to finding a rank-4 model \mathbf{M}_{PSF} that optimally approximates the rank-5 observations \mathbf{X}_{Obs} in a least-squares sense. The matrix PSF model \mathbf{M}_{PSF} can then be tensorized into its original data-cube shape. By including only a few PC in the model, the planet has been successfully removed from the data. Subtracting the PSF model from the original observations leaves only the flux from the planet and noise.

To address these challenges, post-processing techniques typically employ a PSF subtraction approach. This method involves modeling the coronagraphic Point Spread Function (PSF) and quasi-static speckles in [Figure 1.1b](#), then subtracting this model from the observations. The resulting residual image should ideally contain only the planetary signal and as little remaining noise as possible, as illustrated in [Figure 1.1c](#).

Over the past two decades, various image processing techniques have been developed to refine this PSF subtraction approach. The current standard method in high-contrast imaging (HCI) is principal component analysis (PCA). [Figure 1.2](#) illustrates how PCA is applied to spectral differential imaging observations. ([Pueyo, 2018](#))

PCA uses the singular value decomposition (SVD) to find a low-rank approximation of observations

obtained by angular-spectral differential imaging (ASDI). This low-rank approximation serves as the model for the stellar PSF and quasi-static speckles. (Amara et al., 2012; Soummer et al., 2012) However, as shown in Figure 1.2, to apply SVD, the higher-dimensional observation cube must be reshaped into a 2-dimensional matrix. This flattening process disrupts inter-modal relations (Favier, 2021), potentially causing PCA to miss important information in the observations and limiting its effectiveness in modeling the coronagraphic PSF and quasi-static speckles.

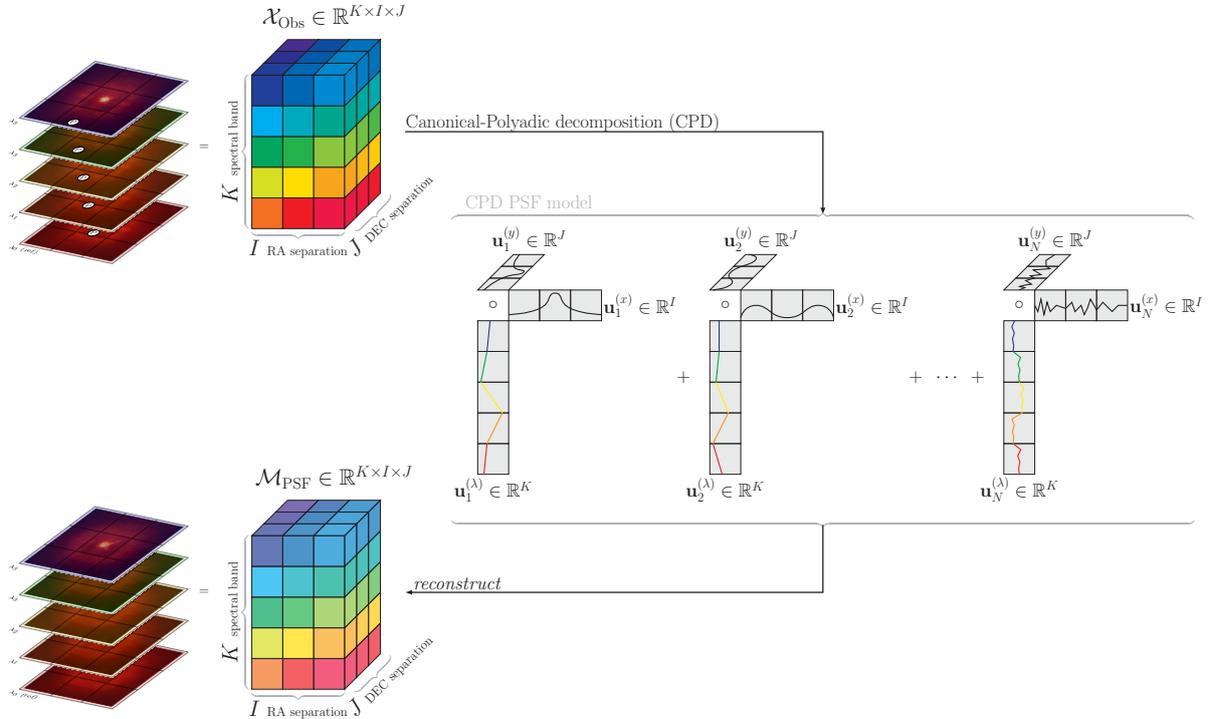


Figure 1.3: Post-processing spectral differential imaging observations with tensor methods using the Canonical-Polyadic Decomposition (CPD). The observations are pre-processed as explained in Figure 1.2. However, instead of modeling the PSF as a matrix, the tensor methods find a model that has the same shape as the original data. In the case of the canonical-polyadic decomposition, the PSF model is the sum of vector products. Each vector product has a factor for the spectrum $\mathbf{u}_n^{(\lambda)}$, right ascension $\mathbf{u}_n^{(x)}$ and declination separation $\mathbf{u}_n^{(y)}$. These vector products are conceptually similar to the principal components of the singular value decomposition. However, because the canonical-polyadic decomposition preserves of the original data shape, it is able to exploit covariance between features in the spectral and spatial modes. This means that its factors can correspond to “physical” components in the observations like the coronagraphic PSF, individual speckles, but also off-axis point sources like exoplanets and extended sources like disks. The example in Section 2.5 will illustrate why this is advantageous.

This work proposes tensor decompositions as an alternative to matrix-PCA for modeling the coronagraphic PSF and quasi-static speckles. Like matrix-PCA, tensor decompositions find factors that approximate the original data (Vasilescu et al., 2002). However, they preserve the original higher-dimensional structure, preventing the information loss associated with flattening observations. Figure 1.3 illustrates one of the proposed tensor PSF models which uses the canonical-polyadic decomposition (Hitchcock, 1927).

Note that the canonical-polyadic PSF model in Figure 1.3 maintains the same higher-dimensional structure as the observation cube, whereas the matrix-PCA model in Figure 1.2 requires reshaping the data into a matrix. This is because the canonical-polyadic model approximates the coronagraphic PSF and quasi-static speckles using a sum of outer vector products. By avoiding reshaping the observations, tensor decompositions can use the multi-modal correlations of PSF components when computing the PSF model. The correlations are informative for separating the stellar PSF and speckle pattern from exoplanets. The flexibility and expressiveness of tensor decompositions together with their ability to capture mode cross-coupling thus make them well-suited for extracting low-rank PSF models from HCI observations.

1.2. Research Objectives & Methodology

The main research objective of this thesis is to develop tensor decomposition methods that can capture multi-modal structure when modeling the stellar flux in observations obtained by angular-spectral differential imaging. These methods need to model the coronagraphic point spread function (PSF) and quasi-static speckle pattern, both of which are known to be well approximated by low-rank matrix factorizations (Gonzalez et al., 2016). Since the locations and spectra of coronagraphic PSF components and speckles in differential imaging observations are quasi-static, they have inter-model relationships that are disrupted by flattening. The new methods use these relationships to improve the low-rank model without needing to flatten the observation data.

Objectives 1.1: Primary

Matrix factorization-based Principal Component Analysis (PCA) is commonly employed for modeling the coronagraphic Point Spread Function (PSF) for post-processing exoplanet observations obtained by differential imaging. However, this approach suffers from the **loss of information** due to the flattening of the multi-modal observations which disrupts inter-modal relationships of PSF components. Unlike matrix factorization, tensor decompositions preserve these relationships.

The main objective of this thesis is to develop data post-processing methods that use tensor decompositions to model the coronagraphic PSF and quasi-static speckles in exoplanet observations obtained by angular-spectral differential imaging.

The challenges in achieving this objective are closely related to the general challenges in HCI post-processing, as illustrated in Figure 1.1 through Figure 1.3. All post-processing methods must contend with the high contrast between the stellar and companion flux, the complex, varying structure of the PSF together with a quasi-static speckle pattern, and the non-Gaussian nature of the noise. Additionally, PSF modeling methods must be suitable for spectrophotometry and astrometry of exoplanets, as these are crucial for confirming and characterizing detected companions. (Pueyo, 2018)

Currently, no method in HCI post-processing preserves the higher-order relations between the modes of HCI observations, demonstrated by the difference between Figure 1.2 and Figure 1.3. Additionally, the impact of the Rician-distributed noise in HCI observations remains underexplored for tensor decompositions. Moreover, the effect of exoplanet flux on tensor decomposition-based PSF models is unknown since PSF subtraction methods fundamentally differ from other applications of low-rank tensor models. To be used by the HCI community, new methods must be easily integrable into existing computational frameworks and openly accessible. Lastly, these methods need to demonstrate competitive or superior performance compared to matrix-PCA in terms of exoplanet detection capabilities and computational demands to be adopted.

Addressing the research objective and challenges allows answering the following research question:

Research Questions 1.1: Primary

Can tensor decomposition methods provide a competitive alternative to matrix-PCA in the modeling of the coronagraphic PSF and quasi-static speckle pattern in exoplanet observations obtained by angular-spectral differential imaging?

To address the main research question, the following sub-questions need to be answered:

1. *Modeling Capability*: Can tensor decomposition methods model the coronagraphic PSF and speckle pattern?
2. *Detection Performance*: What are the receiver operating characteristic curves of tensor decomposition methods?
3. *Detection Limits*: What are the detection limits of tensor decomposition methods in terms of stellar-to-sub-stellar companion contrast and angular separation?
4. *Performance Comparison*: How do tensor decomposition methods compare to matrix-PCA?

The proposed methodology consists of the following key points: Instead of flattening the high-dimensional observation into a 2-dimensional matrix, the data is treated as a *higher-order* tensor. This preserves

the inherent structure and inter-modal relationships present in the data. To separate exoplanet flux from background, this work uses the PSF subtraction technique on real and synthetic observations obtained by the SPHERE instrument of the VLT obtained by ASDI. The main contribution of this work is the development and evaluation of several tensor decomposition techniques for modeling the coronagraphic PSF and quasi-static speckles. Their suitability for spectrophotometry and astrometry is demonstrated. The new methods are assessed as binary classifiers to determine their performance and compared against the current standard, matrix-PCA. Lastly, their efficacy for HCI is demonstrated by using them to detect a new sub-stellar companion candidate and discover a new feature in a circumstellar disk.

This approach aims to develop tensor decomposition-based methods that effectively model the coronagraphic PSF and quasi-static speckles in ASDI observations while preserving multi-modal structure and inter-modal relationships. A comprehensive assessment of direct exoplanet imaging methods is challenging due to varied observing conditions, instruments, and campaign goals, however, this work provides a valuable indication of tensor decomposition methods' efficacy relative to matrix-PCA in expected use cases. As the first extension of existing HCI post-processing methods to higher-order data, this approach offers a competitive alternative to matrix-PCA for direct imaging.

1.3. Contributions

This section outlines the structure of this thesis, explaining the contents and contributions of each chapter to the overall work. Detailed contributions of each chapter are provided in a dedicated section at the end of the respective chapter. The main contributions of this thesis are listed below.

Contributions 1.1: Main Contributions

This thesis makes several contributions to the field of exoplanet direct imaging:

1. **Development of Multi-Modal Point-Spread Function Models using Tensor Decompositions:** The thesis introduces a new class of PSF modeling methods that use tensor decompositions to model the coronagraphic PSF and speckle pattern in angular-spectral differential imaging observations. These methods preserve the higher-order structure of ASDI data, improving modeling of spectral-temporal-spatial features compared to matrix-based approaches.
2. **Quantitative Assessment of Tensor-based Models:** The proposed tensor methods are evaluated using synthetically injected exoplanets into a large sample of ASDI observations. This assessment provides a comprehensive comparison with existing matrix-based techniques.
3. **New Discoveries:** The tensor-based methods developed in this thesis have led to two new discoveries:
 - A debris disk feature in the AU Microscopii system
 - A sub-stellar companion candidate in the HD 108767 B system
 These detections underscore the practical utility of the new methods.
4. **Open Source Software Implementation:** An open-source implementation of the multi-modal point-spread function models using tensor decompositions is provided. The software package is available at [this repository](#)^a. The implementation is compatible with the [Vortex Image Processing](#)^b package.

These contributions collectively establish a new framework for PSF modeling and subtraction in high-contrast imaging, offering improved exoplanet detection sensitivity and characterization capabilities.

^aDirect Exoplanet Imaging with Tensor Decompositions: <https://github.com/lwelzel/tide>.

^bVortex Image Processing: <https://github.com/vortex-exoplanet/VIP>. (Christiaens et al., 2023; Gonzalez et al., 2017)

The main limitation of this study is described below.

Main Limitation

Comprehensively assessing methods for direct exoplanet imaging is challenging due to the wide variety of observing conditions, instruments, and goals for observation campaigns. This work gives an indication of the efficacy of tensor decomposition methods in relative to matrix-PCA in the *expected use case*, however, any such work is necessarily *not comprehensive*.

Chapter 2: Background

This chapter provides the necessary foundation for understanding the methods developed in this thesis, bridging the previously unconnected fields of exoplanet imaging and multi-linear algebra.

Contents:

- Introduction to exoplanet science, with focus on direct imaging
- Overview of current post-processing methods in direct imaging
- Introduction to multi-linear algebra and tensor decompositions
- Example of the problem addressed by this thesis

Contributions:

- Establishes the interdisciplinary context of the research
- Provides essential background in direct imaging of exoplanets and multi-linear algebra
- Motivates the need for tensor-based post-processing methods in direct imaging

Chapter 3: Method

This chapter develops the new tensor-based methods for post-processing direct imaging observations.

Contents:

- Transposition of principal component analysis into multi-linear framework
- Description of new tensor-based methods
- Discussion of modifications and their impacts on the methods

Contributions:

- Development of tensor-based methods
- Illustrates how new methods generalize and improve upon current techniques
- Provides a comprehensive exploration of tensor decomposition applications in direct imaging

Chapter 4: Observations & Data

This chapter describes the observational data used to validate and assess the new methods.

Contents:

- Description of instruments used to obtain observations
- Overview of surveys and observational campaigns

Contributions:

- Provides context for interpreting results
- Ensures reproducibility of the research

Chapter 5: Results

This chapter demonstrates the efficacy of the new methods and compares them with current standard techniques.

Contents:

- Qualitative assessment using well-known directly imaged systems

- Demonstration of method applicability to circumstellar disk features
- Quantitative assessment using synthetic exoplanet injection
- Discovery of new exoplanet candidates, astrometry and spectrum model

Contributions:

- Validates the new methods through detection of known exoplanets and disk features
- Provides quantitative performance metrics for new methods
- Demonstrates practical utility through new astronomical discoveries

Chapter 6: Discussion

This chapter contextualizes the new methods within the broader field of direct imaging.

Contents:

- Positioning of new methods among recent developments in direct imaging
- Assessment of residual noise properties
- Discussion of practical concerns in applying multi-linear algebra to post-processing

Contributions:

- Identifies connections between new methods and existing techniques
- Provides insights into the theoretical underpinnings of tensor-based approaches

Through this structure, the thesis comprehensively introduces, develops, validates, and contextualizes tensor-based methods for post-processing direct imaging observations of exoplanets. The topics, structure and contributions of this thesis are illustrated in [Figure 1.4](#).

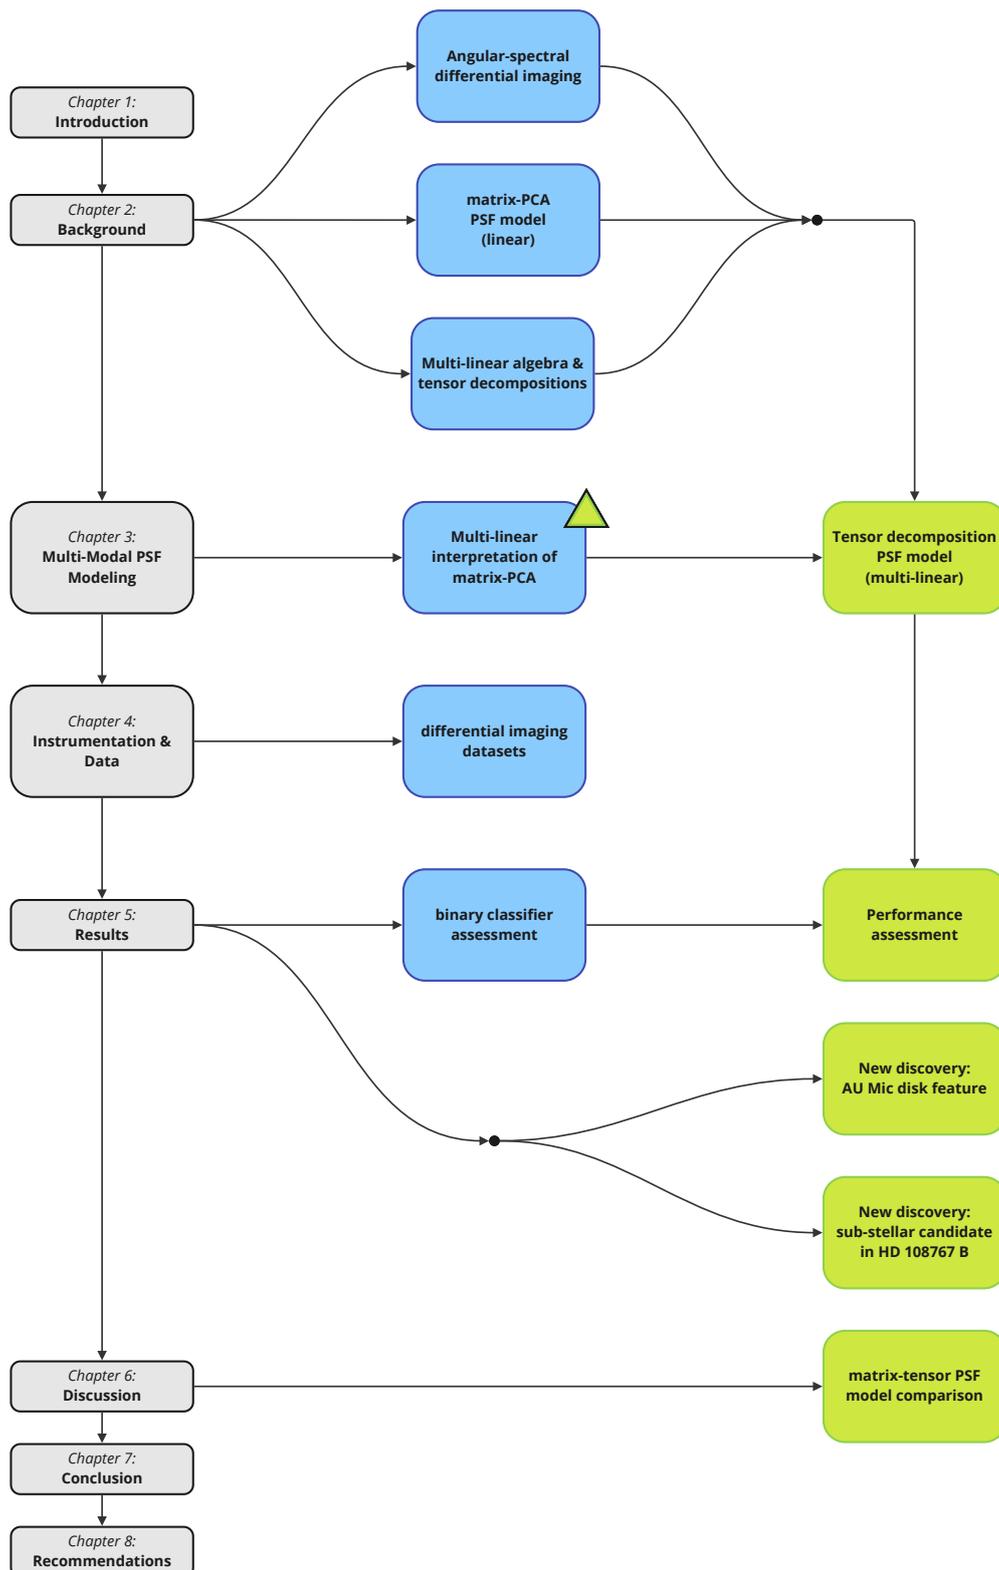


Figure 1.4: The main topics in this thesis, together with its structure and contributions. This flow chart visually represents how existing concepts and methods are the foundation for the contributions of the thesis. Chapters are in gray, existing concepts or methods are in blue, and this thesis' contributions are in green.

Background

This chapter provides the essential background and context for understanding the application of tensor decomposition methods to the post-processing of direct exoplanet observations. It is structured into five sections, each building upon the previous to establish a comprehensive foundation for the research presented in this thesis.

The chapter begins with an overview of exoplanets, delineating their characteristics, detection methodologies, and the current state of exoplanetary science. Particular attention is given to the subset of the exoplanet population which is accessible through direct imaging techniques. This section underscores the significance of exoplanet research and the inherent challenges in their detection and characterization.

The second section focuses specifically on the direct imaging of exoplanets. It explains the principles behind this technique, the general instrumentation involved, and the unique challenges it presents. This section also introduces the concept of angular-spectral differential imaging, which is crucial for understanding the data structure in direct imaging observations.

The second section focuses specifically on the direct imaging of exoplanets. It explains the principles behind this technique, the instrumentation involved, and the unique challenges it presents. This section also introduces the concept of angular-spectral differential imaging, which is crucial for understanding the data structure in direct imaging observations.

The third section introduces multi-linear algebra, providing the basic mathematical framework necessary for working with higher-order tensors. This foundation is essential for understanding the tensor decomposition methods that are central to this thesis.

Building on this mathematical groundwork, the fourth section expands on the concept of factorizing tensors through tensor decompositions. It explores various decomposition techniques, their properties, and their potential applications in data analysis, with a particular focus on their relevance to exoplanet imaging data.

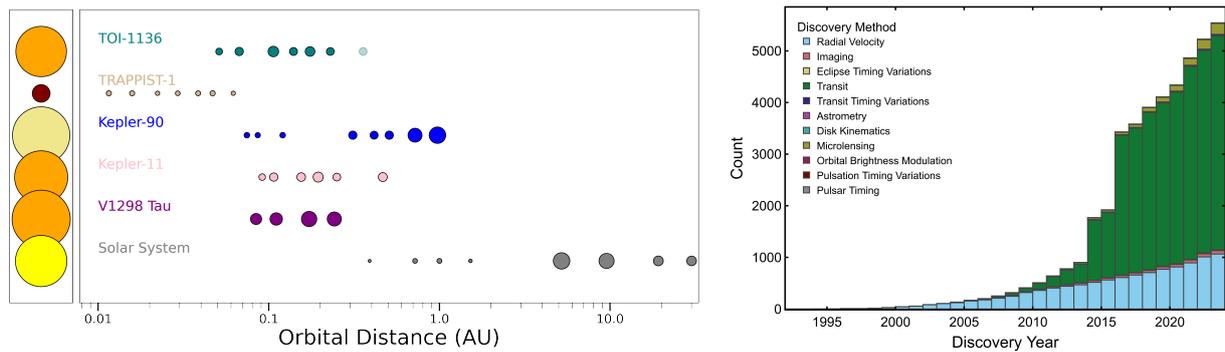
Finally, the chapter concludes with a section demonstrating the limitations of traditional matrix-based principal component analysis when applied to higher-order tensor data. This section presents a simplified example that illustrates why new approaches are needed for effectively processing and analyzing multi-dimensional exoplanet imaging data.

Together, these sections provide the reader with the necessary background to appreciate the motivation behind this research and to understand the novel tensor-based methods proposed in the subsequent chapters of this thesis.

2.1. Exoplanets

Exoplanet research has evolved into a cornerstone of modern astrophysics, concerned with understanding the broader context of the Solar System and the potential for life beyond Earth. The discovery of exoplanets - planets orbiting stars outside the Solar System - has revolutionized our understanding of planetary systems and the vast diversity of existing worlds.

The surprising detection of planets around the millisecond pulsar "Lich" (PSR B1257+12, [Wolszczan et al., 1992](#)) and the discovery of the first confirmed exoplanet orbiting a Sun-like star, 51 Pegasi b ([Mayor et al., 1995](#)), started the search for exoplanets in earnest in the 1990s. Since then, advancements in space-based and large ground-based observatories have quickly increased the number of known exoplanets. At this point, more than 5600 exoplanets have been detected ([NASA Exoplanet Archive, 2024](#)). The cumulative exoplanet discoveries over the last three decades are shown in [Figure 2.1b](#). These exoplanets are astonishingly diverse, ranging from hot Jupiters - gas giants in close orbits around their stars- to Earth-like



(a) High-multiplicity planetary systems in comparison with the solar system. While the systems shown here might look similar, each system is unique in its own way and differs in key aspects from the others. Stellar and planet radii are scaled for comparison, however, planet-star size is not to scale. (Beard et al., 2024)

(b) Cumulative discoveries of exoplanets by first discovery method. The vast majority of discoveries are from the transit and radial velocity method, which are both indirect methods. The large year-on-year increases are due to new telescopes becoming operational, for example the Kepler and TESS spacecraft. (NASA Exoplanet Archive, 2024)

Figure 2.1

rocky planets in the habitable zone, where conditions might allow liquid water to exist. This diversity is compounded by the diversity of the environments they exist in, from protoplanets, still embedded in their planet forming disk (Keppler et al., 2018), to 8-planet systems around stars similar to the Sun (Shallue et al., 2018). Figure 2.1a shows the architecture of the solar systems and Kepler-90 in comparison with other high-multiplicity planetary systems.

This stunning diversity and ubiquity of exoplanets encourages perhaps the most profound aspect of exoplanet research: the search for life beyond Earth. This involves identifying potentially habitable exoplanets and probing their atmospheres for biosignatures. These are molecular species which are indicators of life, such as oxygen together with methane and other organic molecules. The discovery of a habitable exoplanet with a stable climate and an atmosphere conducive to life would be a groundbreaking milestone, fundamentally altering our understanding of life's potential pervasiveness across the universe.

This is also why studying exoplanet atmospheres is particularly intriguing as it allows scientists to understand their composition, weather patterns, climate, and potential habitability. Atmospheric studies can reveal the presence of key molecules, such as water vapor, methane, and oxygen, which are crucial for assessing the potential for life. Additionally, atmospheres play a significant role in regulating planetary temperatures and protecting surfaces from harmful radiation, making them central to understanding the overall environmental conditions on exoplanets. (Hardegree-Ullman et al., 2023; Madhusudhan, 2019) Nevertheless, to study an exoplanet, and its potential atmosphere, it needs to be discovered first.

2.1.1. Detection Methods

Detecting exoplanets is a complex endeavor that employs a variety of techniques, broadly categorized into indirect and direct methods. Each method has its unique advantages, limitations, and requirements for its use. The choice of detection method often depends on the type of exoplanet being sought, the distance to the exoplanetary system, and the observational resources available. Figure 2.2 illustrates which detection methods were used to discover exoplanets, depending on their mass and the semi-major axis of their orbit.

Indirect Detection Methods

Indirect detection methods infer the presence of exoplanets by observing the effects these planets have on their host stars. These methods have been highly successful in identifying and characterizing a large number of exoplanets.

The **radial velocity** method, also known as the Doppler method, detects exoplanets by measuring variations in the velocity of a star along the line of sight from Earth. These variations are caused by the gravitational pull of an orbiting planet, which induces a wobble in the star's motion. As the star moves towards or away from the observer, its light is shifted to shorter or longer wavelengths, respectively, a phenomenon known as the Doppler effect. This method is particularly effective for detecting massive planets in close orbits

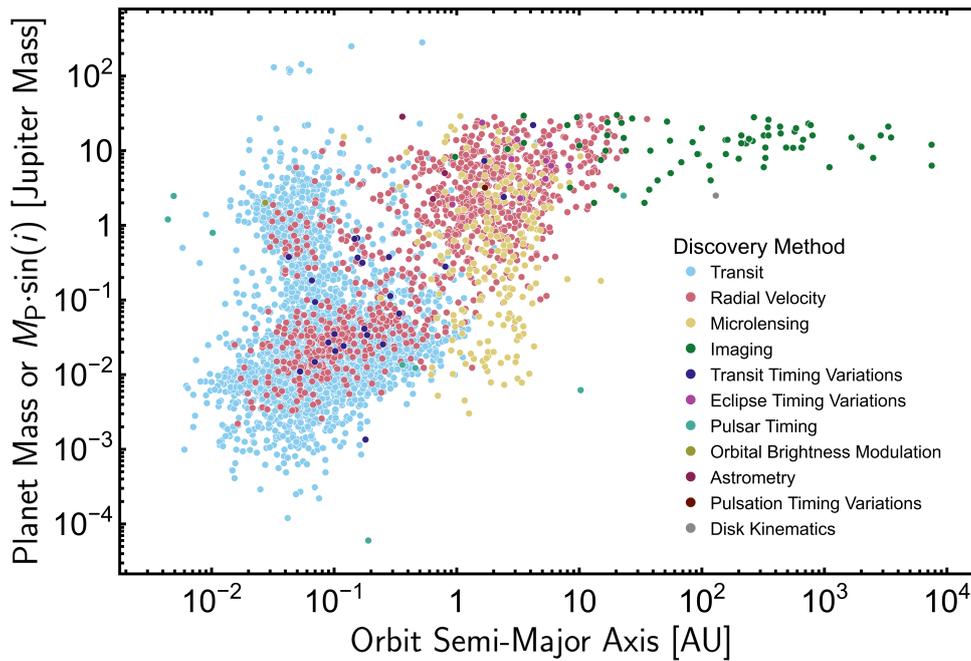


Figure 2.2: Exoplanet mass versus plotted against their semi-major axis. The first detection method is indicated by color. The population of directly imaged exoplanets is visibly separated from the other methods as it performs best for massive exoplanets on wide orbits. ([NASA Exoplanet Archive, 2024](#)).

around their stars, known as hot Jupiters, as can be seen in [Figure 2.2](#). ([Hara et al., 2023](#); [Mayor et al., 1995](#))

Transit photometry detects exoplanets by measuring the periodic dimming of a star's light caused by a planet passing in front of it, an event known as a transit. When a planet transits its host star, it blocks a small fraction of the star's light, leading to a temporary decrease in brightness. These light curves, can be used to infer the planet's size, orbital period, and atmospheric properties if multiple wavelengths are observed. This method is responsible for discovering the majority of exoplanets as well as characterizing their atmospheres. However, in order to be applied the exoplanet must transit its host star while being observed. As is visible from [Figure 2.2](#), the probability of catching a transit is higher for planets on short orbits so that the methods efficiency quickly drops with increasing semi-major axis and orbital period. ([Charbonneau et al., 1999](#); [Henry et al., 1999](#); [Perryman, 2018](#))

Astrometry involves precise measurements of a star's position in the sky. A star with an orbiting planet will exhibit a small, periodic shift in its position due to the gravitational influence of the planet. This method can provide accurate measurements of a planet's mass and orbit, particularly for planets that are farther from their host stars. However, the technique requires extremely high precision and is technically challenging. ([Sozzetti, 2005](#))

Gravitational microlensing takes advantage of the gravitational field of a star to act as a lens, magnifying the light from a more distant background star. If a planet orbits the foreground star, it can cause a noticeable additional increase in brightness. This method is especially powerful for detecting planets that are far from their host stars and even those in other galaxies. However, microlensing events are rare and unpredictable, requiring extensive monitoring of dense star fields.

Several other indirect techniques have been used in specific scenarios, however, the above methods account for the vast majority of currently known exoplanets. ([Dai et al., 2021](#); [Perryman, 2018](#))

Direct Detection Methods

Direct detection methods aim to observe the planets themselves, rather than inferring their presence through their effects on host stars. These methods are very challenging due to the faintness of exoplanets compared to their bright host stars. This ratio, called contrast, is typically lower than 10^{-3} . Observing Jupiter and Earth from a different solar system, it would be around 10^{-9} and 10^{-10} .

Direct imaging involves capturing pictures of exoplanets by blocking out most of the light of their host stars. Advanced instruments such as coronagraphs and starshades are used to suppress starlight, allowing the faint light from exoplanets to be detected. As visible from [Figure 2.2](#), this method is most effective for young, massive planets that are far from their stars, as they are brighter and easier to distinguish from the starlight. There are a variety of techniques that are used in direct imaging, such as reference-star differential imaging ([Lafrenière et al., 2009](#); [Smith et al., 1984](#)), polarimetric differential imaging ([Baba et al., 2003](#)), and coherent differential imaging ([Bottom et al., 2017](#); [Codona et al., 2004](#)). The most successful methods, however, are angular-, and spectral differential imaging ([Marois et al., 2006](#); [Smith, 1987](#)). This is the group of methods to which this work contributes. A more in-depth introduction will be given in [Section 2.2](#). ([Currie et al., 2023](#))

Polarimetry is closely related to direct imaging and measures the polarization of light reflected from a planet's atmosphere. Since starlight is originally unpolarized but becomes polarized when it scatters off a planet's atmosphere, it can be a powerful tool in characterizing exoplanets. The polarized light can provide information about the planet's atmosphere and surface properties, such as the presence of liquid water. Still, the method has seen more success when applied to extended objects such as circumstellar disks. ([Currie et al., 2023](#); [Vaughan et al., 2023](#))

Comparison of Methods

The choice between indirect and direct detection methods depends on the specific goals of the observation and the characteristics of the exoplanetary system. Indirect methods, such as radial velocity and transit photometry, have been highly successful in discovering and characterizing a large number of exoplanets, particularly those close to their host stars. These methods are well-suited for large-scale surveys and statistical studies of exoplanet populations. Direct methods, such as direct imaging and polarimetry, provide unique insights into the physical and atmospheric properties of exoplanets, allowing for direct studies of individual planets. However, these methods are technically challenging and often require advanced instrumentation and observatories. Together, these detection methods complement each other, providing a comprehensive toolkit for detecting and characterizing worlds beyond the Solar System. ([Currie et al., 2023](#))

2.1.2. Formation

The formation of exoplanets, involves complex processes that begin within protoplanetary disks. These disks, composed mostly of hydrogen and helium gas together with dust, orbit young stars and provide the material from which planets form. Several mechanism likely contribute to the formation of planets; the dust particles in the disks collide and, if the collision is sufficiently gentle, stick together forming larger particles. However, when the particles become larger, this mechanism becomes inefficient as the collision of larger particles is more likely to lead to their breakup. Streaming instabilities in the disks potentially overcome this barrier. These instabilities are caused by the differential orbital velocities of gas and particles and lead to local over-densities in the disks around isolated particles. These over-densities can become self-gravitating which leads to positive feedback, growing the particles into planetesimals. These planetesimals are large enough to start accumulating mass through further accretion without breaking up, eventually forming protoplanets. ([Raymond et al., 2022](#))

An alternative mechanism is spontaneous gravitational collapse due to instabilities in the disk. Gravitational collapse and subsequent disk fragmentation can lead to the rapid formation of gas giants around young stars, which are otherwise difficult to explain. As these objects become larger and more massive they start to significantly interact with the wider disk. This process leads to spiral arms in the disk and often the migration of the protoplanet through the disk. Furthermore, the protoplanet can begin to accrete significant amounts of the surrounding gas, opening a gap at its orbit in the disk, again, potentially accompanied by the protoplanets migration. As the gas in the disk further depletes the final formed planets begin to settle. Their final orbits are often determined by the resonant chains in the (multi-planet) system. ([Kley et al., 2012](#); [Paardekooper et al., 2023](#); [Raymond et al., 2022](#))

While the gas in the circumstellar disks is almost entirely accreted onto the star or expelled by stellar winds, some dust and smaller sized bodies can remain in the form of a debris disk. These debris disks can be long lived and continuously replenish dust that is blown away by its host star from particle collisions and fragmentation. ([Hughes et al., 2018](#); [Najita et al., 2022](#); [Wyatt, 2008](#)) The formation history significantly impacts not only an exoplanets orbit but also its composition. ([Bergin et al., 2024](#))

2.1.3. Characteristics

The characteristics of exoplanets involve primarily their bulk, interior and atmospheric composition, inferred from measured densities and theoretical models, as well as analyzing their spectra across various wavelengths.

Composition

The bulk composition of exoplanets can be inferred from their densities, which are determined by measuring their masses and radii. These measurements, combined with theoretical models, allows classification of exoplanets into two broad categories based on their composition: rocky planets and gaseous planets.

Gaseous planets, broadly known as gas giants, are composed predominantly of hydrogen and helium, likely with diffuse cores of heavier elements. The low density of gaseous exoplanets indicates that they have thick atmospheres of light gases. Hot Jupiters, such as 51 Pegasi b, are an important category of gaseous exoplanets. These gas giants orbit very close to their host stars, resulting in extremely high surface temperatures and significant atmospheric evaporation. This makes them the easiest to study out of the entire exoplanet population. The hot Jupiter HD 209458 b, for example, is the first exoplanet around which an atmosphere was detected. (Madhusudhan, 2019; Miguel et al., 2023)

Rocky planets, also known as terrestrial planets, are composed primarily of silicate rocks and metals. The density of rocky exoplanets suggests a composition similar to that of the terrestrial planets in the Solar System, with a core of iron and nickel surrounded by a mantle of silicate minerals. These planets can have atmospheres and complex climates. Most currently known rocky exoplanets are larger, more massive, and hotter than the rocky planets in the solar system because these kinds of planets are easier to detect with the current technologies and methods. Even then, directly detecting rocky exoplanets is exceedingly difficult. (Chao et al., 2021; Wordsworth et al., 2022)

Intermediate types also exist, such as mini-Neptunes and super-Earths, which exhibit a range of compositions and densities. Super-Earths, with masses up to ten times that of Earth, may have thick atmospheres or substantial water ice content, complicating the simple rocky versus gaseous classification. (Madhusudhan, 2019) In the following, this work will be mostly concerned with young gas giants.

Atmospheres

The atmospheres of giant exoplanets are predominantly composed of light gases, primarily hydrogen (H_2) and helium (He), much like the gas giants in the own solar system. However, their compositions can vary significantly due to factors such as formation history, migration patterns, and proximity to their host stars.

Composition The high abundance of hydrogen and helium stems from the protoplanetary disk from which these planets formed, which was rich in these primordial gases. Additionally, giant planet atmospheres can contain various molecular compounds that give rise to diverse chemical processes. Key molecules that have been observed in many gas giants are water (H_2O , typically as vapour), methane (CH_4), carbon bearing species such as carbon monoxide (CO) and carbon dioxide (CO_2), nitrogen (N_2), and nitrogen bearing species such as ammonia (NH_3) and hydrogen cyanide (HCN). The impact of these species in near infrared exoplanet emission spectra is shown in Figure 2.3a.

Water vapor is one of the most abundant and crucial molecules in the atmospheres of giant exoplanets. It significantly influences the atmospheric opacity and is a key indicator of planetary formation and migration history. It exhibits strong absorption features in the infrared spectrum, particularly around 1.1, 1.4, and 1.9 μm . This means it also contributes to the greenhouse effect, trapping heat and affecting the thermal profile of the atmosphere.

Methane is commonly found in cooler giant planets, where it is more stable. Methane has prominent absorption bands in the infrared, especially around 3.3 μm , as well as 1.6 and 2.3 μm . Methane is one of the primary carbon carriers in gas giants and can contribute to the greenhouse effect, though its impact is temperature-dependent. This is because in strongly irradiated or hot atmospheres, methane (photo-) dissociates, reducing its abundance.

Carbon monoxide and carbon dioxide are prevalent in giant planets. Together with CH_4 , these molecules are the main drivers for carbon chemistry and have a major impact on the thermal dynamics of exoplanets. They have relatively weak spectral features in the near infra-red, but have a strong impact in the mid infrared

beyond $3 \mu\text{m}$. Both CO and CO₂ are potent greenhouse gases, strongly influencing the atmospheric temperature structure. Their presence can lead to higher atmospheric temperatures and affect the overall climate dynamics.

Nitrogen is often abundant in giant planet atmospheres, serving as a major component alongside hydrogen and helium. N₂ itself does not have strong infrared absorption features due to its lack of a dipole moment. However, its presence is inferred through the detection of nitrogen-bearing compounds. These compounds, especially NH₃ (1.3, 1.6, 2.1, 2.3 μm), as well as HCN, have features in the near, mid, and far-infrared, and are the main indicators for nitrogen chemistry. Ammonia is prevalent in cooler (regions of) gas giants where it is more likely to condense. Hydrogen cyanide forms under specific conditions, often linked to high levels of nitrogen and carbon. It is particularly interesting from a chemical standpoint due to its reactivity and potential as a tracer for atmospheric processes. (Fortney, 2024; Fortney et al., 2021; Madhusudhan, 2018, 2019)

Structure The atmospheres of gas giant exoplanets exhibit complex and dynamic structures, driven by a combination of radiative, chemical, and dynamical processes. Typically, these atmospheres consist of multiple layers, each characterized by distinct temperature and pressure conditions. The deepest layer, known as the troposphere, is where temperature decreases with altitude and where most of the atmospheric mass resides. Above the troposphere lies the stratosphere, which may feature temperature inversions due to the absorption of stellar radiation by molecules such as titanium oxide (TiO) and vanadium oxide (VO) in hotter exoplanets. (Fortney, 2024; Gandhi et al., 2023) The outermost layers include the thermosphere and exosphere, where temperatures can rise again due to the absorption of high-energy ultraviolet (UV) and X-ray radiation from the host star. This complex vertical structure is further influenced by the presence of clouds and hazes composed of various condensates like silicates, sulfides, and metallic elements, which form at different altitudes depending on the local temperature and pressure. (Dyrek et al., 2024; Lothringer et al., 2022; Vos et al., 2023) Additionally, atmospheric circulation patterns, driven by rapid rotation and significant temperature differences between the day and night sides of close-in or tidally-locked gas giants, lead to high-speed winds and jet streams, redistributing heat and shaping the overall atmospheric dynamics. (Pierrehumbert et al., 2019) While the complex climates of gas giants in the Solar System have been long known, it has only recently become clear that exoplanets also host complex climates, due to the difficulty in modeling and observing these dynamic and localized phenomena. (Bisikalo et al., 2021; Showman et al., 2020)

Chemical Processes The atmospheric composition of gas giant exoplanets is significantly influenced by various chemical processes and interactions, which shape their climates. The balance between these reactions determines the steady-state composition of the atmosphere, influencing the planet's overall chemical environment. One key process is photochemistry, driven by stellar radiation. High-energy photons from the host star break down molecules in the atmosphere, leading to the formation of new species. These reactions not only alter the chemical makeup of the atmosphere but also affect its thermal and spectral properties. At higher temperatures, thermochemical reactions become dominant, playing a crucial role in the formation and destruction of various molecular species. (Madhusudhan, 2019)

Condensation and cloud formation are also vital processes that impact the atmospheres of gas giants. Depending on the local temperature and pressure, certain molecular compounds condense to form clouds, which significantly influence the planet's albedo and thermal emission. In the upper atmospheres of hot gas giants, clouds composed of silicates, iron, and other refractory materials are common. In contrast, cooler gas giants tend to have ammonia and water clouds. These clouds not only modify the planet's appearance but also play a critical role in its climate dynamics, affecting heat distribution and atmospheric circulation. (Helling, 2019; Showman et al., 2020)

Characterizing Exoplanet Atmospheres

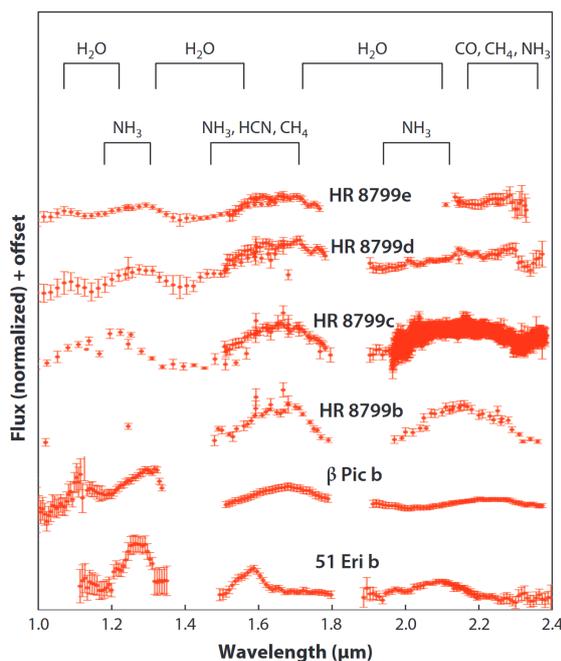
For studies of exoplanet atmospheres, transit spectroscopy and direct imaging are particularly effective. The spectrum that is emitted by, or passes through, an exoplanet atmosphere is determined by its atmospheric composition. Different molecules in the atmosphere emit and absorb specific wavelengths of light, creating emission and absorption lines that can be used to infer the atmospheric composition, temperature, and pressure. (Currie et al., 2023; Madhusudhan, 2018)

High-resolution spectroscopy resolves exoplanet atmosphere spectra at a spectral resolution of $R \gtrsim 5000$. At this resolution individual molecular features become identifiable, so that spectra can provide detailed insights into the atmosphere environment and its composition. However, the ground-based instruments that archive this high spectral resolution are generally not well suited for high contrast imaging because of their low field-of-view (FOV). (Currie et al., 2023)

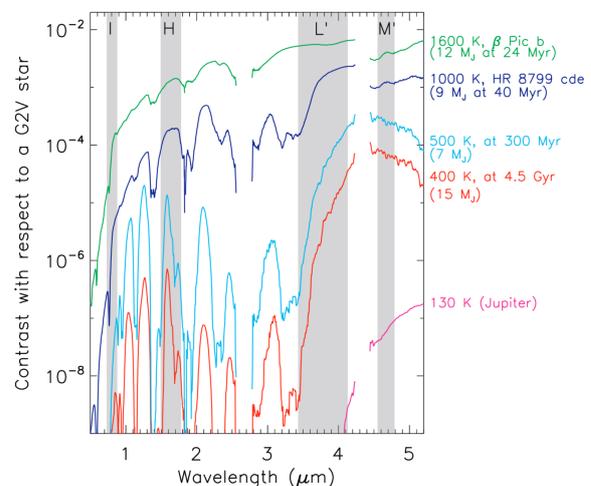
In contrast, low-resolution spectroscopy ($R \lesssim 100$) only provides broad spectral features and can be used to simultaneously find exoplanets over a large FOV and infer the dominant atmospheric constituents such as water vapor, methane, and carbon dioxide (Hoeijmakers et al., 2018; de la Roche et al., 2018). An example of low-resolution spectra obtained by ground-based direct imaging is shown in Figure 2.3a. Most systems that directly image exoplanets and capture their emission spectra operate at low-spectral resolution. (Landman et al., 2023)

2.1.4. Directly Imaged Exoplanet Population

Direct imaging has primarily succeeded in detecting young, self-luminous gas giants, typically with masses greater than 2 Jupiter masses (M_J) and orbiting at wide separations from their host stars (10 – 300 AU). This means they are not detectable with other methods, compare Figure 2.2 with Figure 2.4. Notably, these separations are much larger than the typical habitable zone distances and the planets are too massive to be habitable themselves. They are generally young, often less than 100 million years old. Their youth means they retain significant residual heat from their formation and ongoing contraction, which manifests as strong thermal emission detectable in the infrared spectrum. The most common hosts of these planets are stars in young stellar associations, particularly B, A, and F type stars which are more massive than the Sun and often surrounded by debris disks. (Currie et al., 2023)



(a) Low-resolution emission spectra of directly imaged exoplanets in the near-mid infrared by ground-based telescopes. The spectra are offset for clarity and show the impact of the different atmosphere species and climate. (Madhusudhan, 2019).



(b) Contrast ratio in the optical and infrared of directly imaged exoplanets and models. The contrast is given with respect to a G2V type star for easier comparison. The actual host star of the HR 8799 and β Pictoris planets are hotter, so that their actual contrast is $\sim 5 - 10$ times higher. (Currie et al., 2023).

Figure 2.3

The thermal emission of an exoplanet can be approximated by a black-body spectrum. The peak wavelength of this radiation shifts according to the planet's temperature, following Wien's displacement law. For instance, a young gas giant may peak in the infrared, while its much hotter host star peaks in the visible spectrum. This is why directly imaged exoplanets are typically observed in the infrared spectrum between 0.95 and 2.4 μm . Moving to longer wavelengths is challenging because of Earth's atmosphere and problems

with detecting low-energy photons. Figure 2.3b shows the brightness ratio between exoplanet and G2V type star emission as a function of wavelength, illustrating it becoming more favourable in the infrared. Real exoplanets often deviate from this idealized spectrum due to molecular absorption and emission in their atmospheres. In the near-infrared, exoplanetary spectra are significantly marked by these molecular features. (Currie et al., 2023; Pueyo, 2018)

Notable systems that have been imaged are the HR 8799 system which hosts 4 giant planets with almost face-on orbits (Marois et al., 2008b, 2010), and the β Pictoris system which hosts two giant planets on almost edge-on orbits (Lagrange et al., 2009; Lagrange et al., 2010; Nowak et al., 2020) as well as a dust belt (Wahhaj et al., 2003). Both systems also host debris disks (Smith et al., 1984; Su et al., 2009).

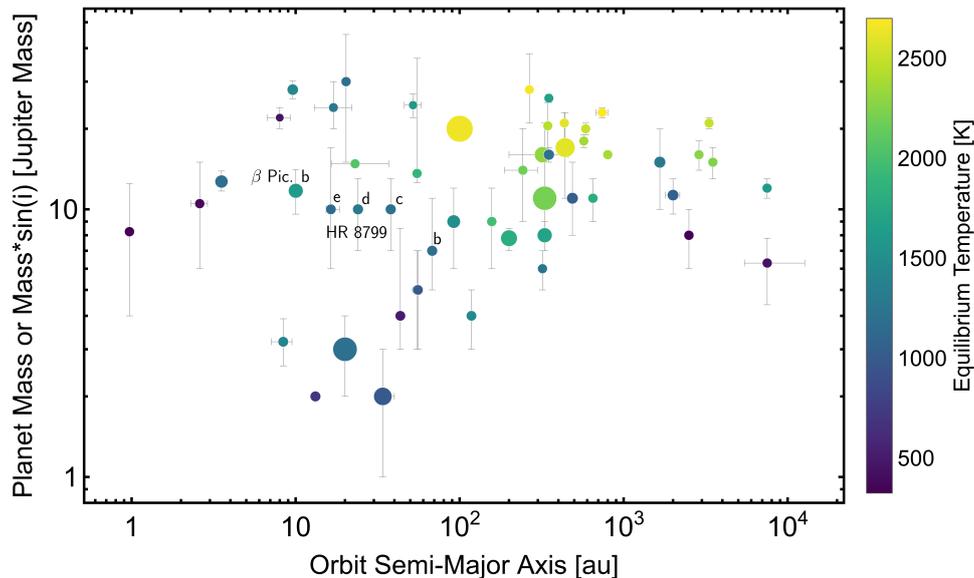


Figure 2.4: Directly imaged exoplanets with known semi-major axis, mass and equilibrium temperature. The color scale indicates the inferred equilibrium temperature of the planets and the radii of the markers are scaled according to the exoplanets inferred radius. Directly imaged exoplanets are much hotter than expected from pure stellar irradiation, indicating that they are young and still contracting. The exoplanets of the HR 8799 system as well as β Pictoris b are marked. (NASA Exoplanet Archive, 2024)

Recently, protoplanets have been directly imaged for the first time in PDS 70, where both a protoplanetary disk and forming planets embedded in it have been observed (Haffert et al., 2019; Keppler et al., 2018). These observations demonstrate that direct imaging can bridge the gap between the initial stages of planet formation and fully formed planets.

Having provided the essential background on exoplanet science, the next section will expand on the methods, technologies and problems in directly imaging exoplanets.

2.2. Direct Imaging of Exoplanets

Direct imaging of exoplanets aims at spatially resolving the emission of a host star and its accompanying exoplanets. Separating the emissions of star and planet facilitates the spectroscopic analysis of planetary atmospheres and enables directly determining the orbital motion of exoplanets. In turn, this provides an independent probe into the physical and chemical processes in exoplanets.

The main challenge that direct imaging has to overcome is twofold: firstly, the star is much brighter than the planet, see e.g. Figure 2.3b. The ratio between the brightness of a planet and its host star is called the *contrast ratio*, or contrast in short. Secondly, the distance between the host star and telescopes on Earth is so large that the angular distance between host star and planet is extremely small. This angular distance is called *angular separation*, or separation in short.

Info box 2.1: Contrast ratio (contrast)

The ratio between the exoplanet F_{planet} and its host star brightness or flux F_{star} at wavelength λ : $c(\lambda) = \frac{F_{planet}(\lambda)}{F_{star}(\lambda)}$. In ground direct imaging the contrast is typically between 10^{-2} and 10^{-10} . Currently, the best achievable contrast with ground-based telescopes is between 10^{-6} and 10^{-7} .

To image exoplanets their angular separation needs to be large enough for the telescope to distinguish the photons originating from the star and planet. This is because telescopes have a limited *resolving power*. The resolving power of a telescope is the smallest angular separation for which it can distinguish two objects. The resolving power of a circular aperture telescope scales according to the Rayleigh criterion with the wavelength λ the telescope is observing at and its aperture diameter $D_{aperture}$ as $\propto \lambda/D_{aperture}$. The origin of this scaling will be explained in [Section 2.2.1](#). The ideal targets to observe with direct imaging are thus luminous, to improve the contrast, and orbit their host at wide angular separations, to reduce the required resolving power of telescopes. The planetary companions for which these conditions are typically fulfilled are the young, self-luminous gas giants on wide orbits which were discussed in [Section 2.1](#). While the study of these exoplanets is very fruitful, they are unlikely to host life themselves. If to find Earth-like, habitable planets, the contrast at which exoplanets can be detected and the effective resolving power of telescopes need to be increased. Beyond imaging Earth-like exoplanets, improvements in both contrast and minimum separation would also increase the fraction of the total exoplanet population that is observable using direct imaging.

Info box 2.2: Angular separation (separation)

The separation is the apparent angle between an exoplanet and its host star as seen from Earth, typically measured in arcseconds (") or milliarcseconds (mas). It determines how far apart the planet appears from the star in telescope images. The angular separation can be converted into a projected distance, typically measured in astronomical units (AU), if the distance to the star is known. Telescopes are currently limited to separations larger than $\sim 100''$.

To increase the contrast at which telescopes can detect exoplanets, they need to collect more photons from the exoplanet without also receiving more from the host star. This means that, without the technical solutions in [Section 2.2.2](#), just building bigger telescopes or observing for longer periods will not enable finding Earth-like exoplanets as the contrast between star and planet stays constant.

To reduce the angular separation at which exoplanets can be detected, astronomers would either need to observe at smaller wavelengths or increase the diameter of the telescopes. However, as was discussed in [Section 2.1](#), the contrast between star and planet typically increases when observing at smaller wavelengths. This is because stars have much higher surface temperatures than planets, so that their black-body radiation at smaller wavelengths is much more than the planets. This is why the next generation telescopes will have aperture diameters of up to 30 m compared to today's 8 m apertures.

Nevertheless, both problems can also be addressed using the observation strategies discussed in [Section 2.2.4](#) together with using the post-processing methods in [Section 2.2.5](#), without having to improve the existing technology and telescopes itself. To demonstrate why this is, understanding the image formation process is crucial. ([Currie et al., 2011](#); [Galicher et al., 2023](#))

2.2.1. Image Formation

Fundamentally, when observing exoplanets, telescopes collect the photons they emit. Because the planet and star are so far from the telescopes, both can be treated as point sources. Imaging a point source produces a finite sized point spread function (PSF) due to the diffraction of light. A simple example is shown in [Figure 2.7](#).

The planet and star are close together compared to the distance between them and the telescopes. This means that photons emitted by the star will be collected together with the planet's photons. There might also be background objects like other stars, or dust which emit photons that are collected. Moreover, the atmosphere of the earth and the telescope itself also emit photons according to their black-body temperature. These photons are all collected by a photon detector, or detector for short. Describing the

highly advanced detectors that are used in high contrast imaging is beyond the scope of this work, however suffice it to say that they consist of a grid of pixels. If a photon hits this pixel the detector "counts" it. The more photons hit a pixel the higher its photon count will be and the brighter the image at its grid position will be. Each pixel of the detector corresponds to an angular field on the sky.

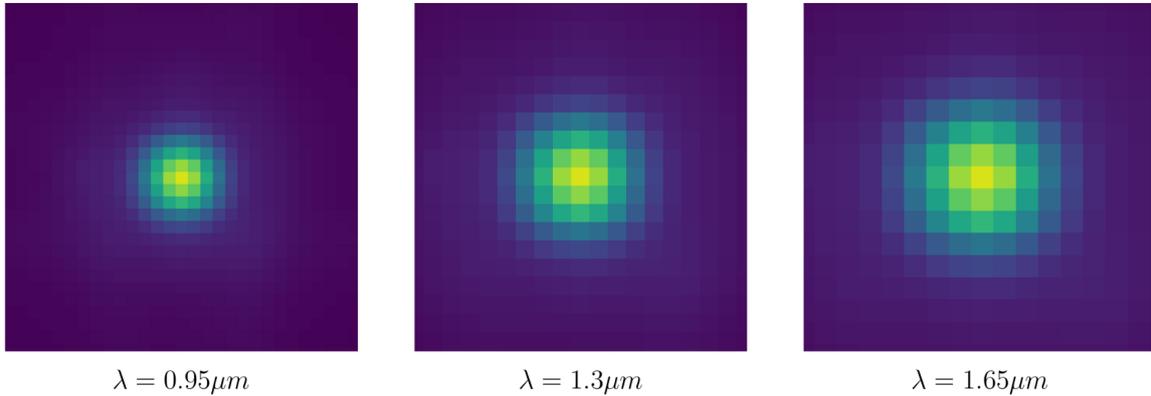


Figure 2.5: Spreading out of a PSF with increasing wavelength.

Simplifying greatly, the flux of a star and exoplanet as seen by the telescope can be written as

$$F_{\lambda}(\boldsymbol{\beta}) = F_{star, \lambda}(\boldsymbol{\beta}) + F_{planet, \lambda}(\boldsymbol{\beta} - \boldsymbol{\alpha}_P), \quad (2.1)$$

where $\boldsymbol{\beta}$ are the angular coordinates on the sky, and $\boldsymbol{\alpha}_P$ denotes the angular separation between the star and the exoplanet. This flux then propagates through the telescope optics to the detector. The telescope optics modify the flux according to its response $R_{\lambda_0}^T(\boldsymbol{\beta})$. The response function determines the PSF of the telescope. The observed image at the detector at wavelength λ_0 is then the convolution (\star) of the flux with the response:

$$R_{\lambda_0}^T(\boldsymbol{\beta}) \star F_{\lambda_0}(\boldsymbol{\beta})|_{\alpha} \quad (2.2)$$

Importantly, the response R_{λ}^T scales approximately linearly with the wavelength of incoming light λ , respectively to the reference wavelength λ_0 :

$$R_{\lambda}^T(\boldsymbol{\alpha}) = R_{\lambda_0}^T\left(\frac{\lambda_0}{\lambda}\boldsymbol{\alpha}\right) \quad (2.3)$$

Because the star and exoplanet are point sources, the convolution in Equation 2.2 produces two PSF copies with differing brightness at the star and planet position. The earth's atmosphere and the specific optical elements that are used to image exoplanets each leave a distinctive impact on the shape of the PSF. In the next section each main component of this process and its impact on the final science images will be discussed. (Galicher et al., 2023; Pueyo, 2018)

Info box 2.3: Point spread function (PSF)

Imaging point sources like stars and exoplanets results in a finite sized PSF due to diffraction. The shape of the PSF is determined by the imaging system. In post-processing high contrast imaging observations, "PSF" generally refers to the stellar PSF, degraded by the earth's atmosphere, corrected by the adaptive optics system, and suppressed by a coronagraph.

2.2.2. Instrumentation

Imaging exoplanets is possible with both space and ground-based telescopes. Figure 2.8 shows the limits to the direct imaging of exoplanets in terms of minimum contrast and separation from the host star for a number of existing and near-future systems. None of the existing systems are able to image solar-system analogues, however the next generation of telescopes might reasonably achieve high enough performance to image Jupiter and Saturn analogues. Space-based observatories have successfully been used for direct imaging. Figure 2.8 further shows that ground-based facilities currently provide the highest resolution and

contrast for exoplanet imaging. For ground-based direct imaging, all, advanced post-processing, specific observational techniques, coronagraphy, adaptive optics, large telescopes and advanced imaging systems, are typically employed in concert to recover the faint signal of exoplanets.

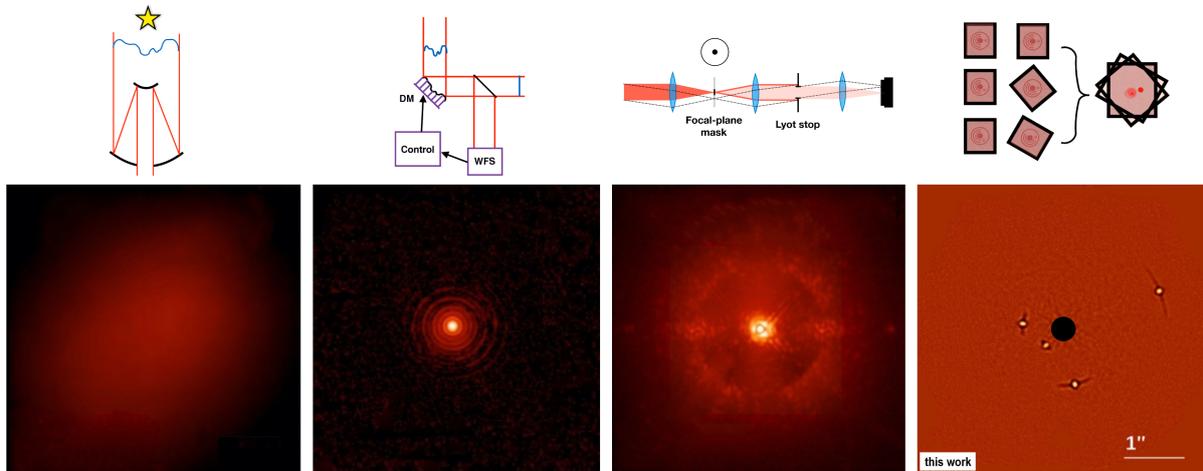


Figure 2.6: The four pillars of ground-based direct imaging of exoplanets. The images show observations of the HR 8799 system observed with the SPHERE instrument at the VLT. From left to right: To capture enough photons, large ground-based telescopes are used. These telescopes however are seeing limited. To overcome the earth's atmospheric distortions of the light from the star and planets, adaptive optics correct for the atmosphere turbulence in real-time. Shown here is the SPHERE extreme adaptive optics system which allows imaging exoplanets at contrasts down to 10^{-3} . To further improve the contrast, the stellar light needs to be suppressed as it is still much brighter than planets, especially at small separations. For this a coronagraph is used, for SPHERE this is an apodized Lyot coronagraph. Current coronagraphs allow imaging planets down to contrasts of 10^{-4} . Lastly, by post-processing the data the achievable contrast can be enhanced by another order of magnitude down to 10^{-5} . Together with the imaging system these components enable the direct imaging of exoplanets when used together. Adapted from [Doelman, 2021](#) (top) and [Cantalloube et al., 2022](#) (bottom).

To capture enough photons from faint exoplanets, large ground-based telescopes are essential. Today's $\sim 8\text{m}$ -class telescopes, such as the Very Large Telescope (VLT), Keck, Gemini, and Subaru, provide the light-gathering power and resolution needed to detect planets orbiting nearby stars. These telescopes have circular pupils for which the incoming light from a centered point source is diffracted into an Airy pattern, shown in [Figure 2.7](#).

However, Earth's atmosphere poses a significant challenge for these ground-based telescopes. Atmospheric turbulence, caused by mixing of air layers with different temperatures and wind speeds, distorts the wavefront of light from astronomical sources ([Figure 2.6](#), leftmost image). This results in a blurred PSF, with a characteristic size determined by the coherence length of the turbulence, known as the Fried parameter (r_0). For typical seeing conditions at a wavelength of 500 nm, r_0 is approximately 10 cm, meaning that telescope diameters larger than this do not provide improved resolution without adaptive optics correction. ([Doelman, 2021](#))

To move beyond seeing-limited imaging, adaptive optics systems are used to measure and correct the wavefront distortions in near real-time. Extreme adaptive optics systems, like the SPHERE extreme adaptive optics system, use high-order deformable mirrors and efficient wavefront sensors to achieve Strehl ratios approaching the diffraction limit. The impact of atmospheric turbulence is reduced at longer wavelengths, with r_0 scaling as $\lambda^{(6/5)}$. However, even in the near-infrared (NIR, e.g. $r_0(1.6\mu\text{m}) \approx 40$ cm) where most ground-based exoplanet imaging is conducted, adaptive optics systems are essential for correcting the wavefront distortions and achieving diffraction-limited performance ([Figure 2.6](#), left-center image). Often, the photon efficiency of the wavefront sensor, together with the temporal lag in the correction due to the finite servo loop speed, limits an adaptive optics systems performance. These systems face further challenges relating to sensing and correcting non-common path aberrations between the wavefront sensor and science camera. Achieving high Strehl ratios ($> 90\%$) and stable PSFs through extreme adaptive optics is crucial for reaching the contrast levels needed to detect faint exoplanets amidst the residual stellar speckles. ([Tyson et al., 2022](#))

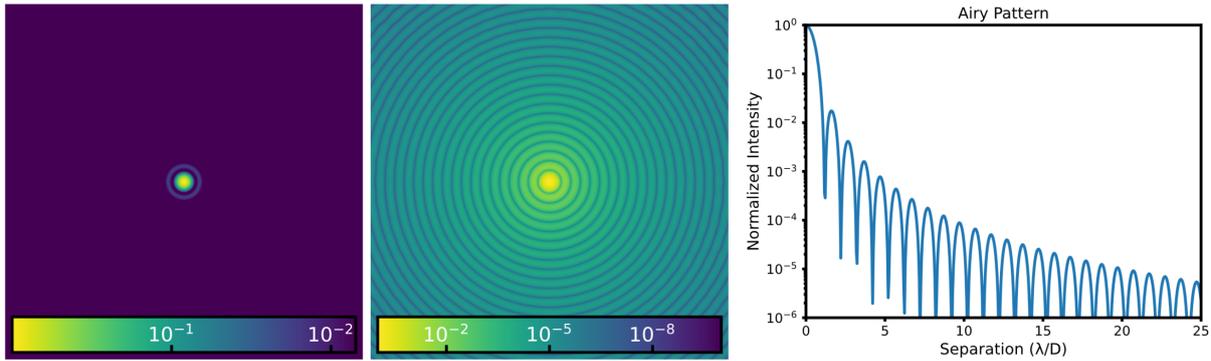


Figure 2.7: Normalized PSF of a circular pupil telescope. The incoming light of a centered point source is diffracted into an Airy pattern. Most of the flux is concentrated within $\lambda/D_{\text{aperture}}$ in the core of the PSF. With increasing separation, the flux slowly decreases. (Galicher et al., 2023)

Even with adaptive optics correction, the stellar light is typically still several orders of magnitude brighter than the faint planet signal at small angular separations. To remove the stellar flux two main methods have been used, coronagraphy and interferometry. Interferometry uses multi-aperture telescopes to archive destructive interference for the starlight, thereby canceling it out, while constructive interference enhances the light from the planet (Lacour et al., 2019). Coronagraphy, on the other hand, employs a physical mask or an optical technique to block or dim the starlight. Diverse coronagraph designs exist, including focal-plane coronagraphs like the classical Lyot coronagraph and phase-mask coronagraphs, as well as pupil-plane coronagraphs like the apodizing phase plate. The Apodized Pupil Lyot Coronagraph (APLC) is commonly used in high-contrast imagers like SPHERE. The speckle field when using the SPHERE APLC is shown in the center-right figure of Figure 2.6. Focal-plane coronagraphs place a mask or a phase-modifying element at the focal plane to block or redirect the starlight while allowing the off-axis planet light to pass through. The classical Lyot coronagraph, is the simplest example of a focal-plane coronagraph. It consists of an opaque mask at the focal plane to block the central star, followed by a slightly undersized pupil stop (the Lyot stop) to suppress the diffracted starlight. The inner working angle (IWA, Figure 2.10b) of a coronagraph is the minimum angular separation where half the flux of a source is still transmitted. For Lyot coronagraphs the IWA is mainly determined by the size of the focal-plane mask, typically a few $\lambda/D_{\text{aperture}}$. The APLC together with the impact of the adaptive optics on HCI observations is discussed in more detail in Section 4.1. (Galicher et al., 2023; Guerri et al., 2011)

The last component of the instrumentation is the imaging system which disperses the light onto the detector, integrates the incoming photons for each pixel, and finally amplifies them into an analog or digital electronic signal. Integral Field Spectrograph (IFS) are especially noteworthy in exoplanet imaging as they enable simultaneous spatial and spectral data acquisition. The lenslet array type stands out for its proven design and high throughput. In this configuration, an array of tiny lenses, or lenslets, is positioned at the focal plane of a telescope. Each lenslet captures light from a specific region of the FOV, effectively segmenting the image into a grid of micro-pupils. The light from these micro-pupils is then dispersed by a spectrograph, resulting in a spectrum for each spatial element. Figure 2.10a shows the result of this dispersion on the image at the detector. The image then is pre-processed to convert it from an data array to a data cube. The overall process is shown in Figure 2.9, which depicts the sequence from focal plane to the final data cube. The final data cube has axes in spatial coordinates (x, y) and wavelength λ .

2.2.3. Speckle Pattern

Even with high-order adaptive optics systems, the correction of atmospheric turbulence is never perfect, resulting in residual wavefront errors that manifest as speckles in the focal plane. Speckles are coherent, spatially localized intensity fluctuations in the PSF that resemble the diffraction-limited core. They arise from the interference of the residual wavefront errors with the underlying diffraction pattern, creating a speckle pattern that can mask or mimic the signals of faint exoplanets.

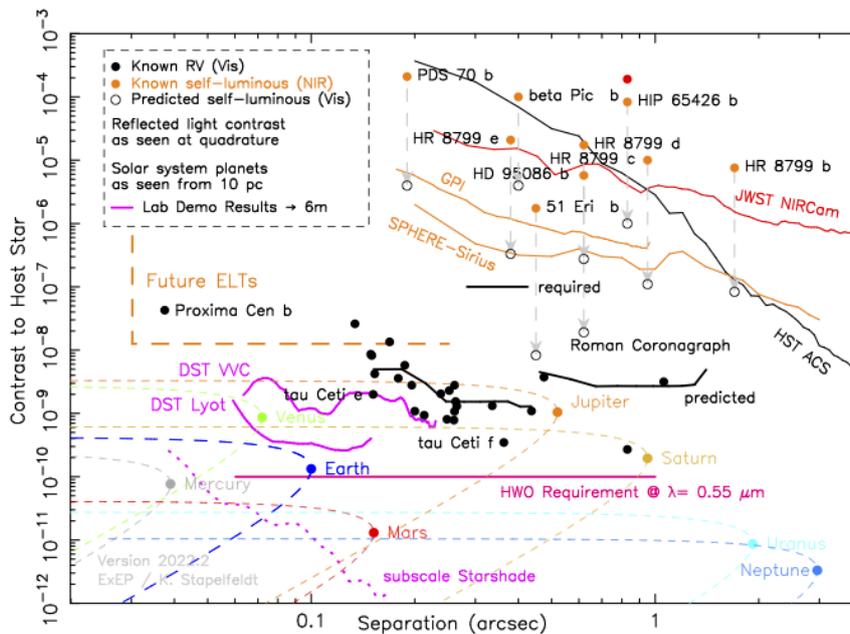


Figure 2.8: Detection limits for current and planned observatories. Increasing the performance of image processing techniques enables future extremely large telescopes or space-based observatories to image terrestrial exoplanets with minimal impact on project cost or schedule. Improved methods can also be used to re-reduce previously collected data and increase its utility. (Boss et al., 2023)

Info box 2.4: Speckle pattern

A speckle pattern consists of speckles which are spatially localized intensity fluctuations in the PSF. These speckles arise from residual wavefront errors due to imperfect atmospheric turbulence correction. They can mask or mimic the signals of faint exoplanets.

Speckles have two main origins in ground-based high-contrast imaging: non-common path aberrations and temporal lag in the adaptive optics correction. Non-common path aberrations result from the wavefront differences between the light path to the wavefront sensor and the science camera. These aberrations, caused by imperfect optics, misalignment, and chromatic effects, are not sensed by the main adaptive optics loop and therefore remain uncorrected. Non-common path aberrations induced speckles vary slowly on timescales of minutes to hours, due to changes in the gravity vector, temperature, and humidity affecting the instrument. These quasi-static speckles are particularly problematic as they can closely resemble the signal of a companion. This is illustrated by comparing the individual speckles in the simulated, degraded PSF in Figure 2.13 (left) with the exoplanet PSF model in Figure 2.24a (left). (Galicher et al., 2023; Pueyo, 2018)

The second major source of speckles is the temporal lag between the wavefront measurement and the correction applied by the deformable mirror. The finite servo loop speed, typically on the order of a few hundred Hz to a few kHz, cannot keep up with the rapidly evolving high-frequency components of the atmospheric turbulence. This results in a servo-lag error, which appears as a butterfly-shaped speckle pattern (called the wind-driven halo) elongated in the direction of the dominant wind velocity at the telescope pupil. The center right image of Figure 2.6 includes a wind-driven halo visible as a broad, horizontal line of increased speckle intensity. (Tyson et al., 2022)

The statistical properties of speckles depend on the wavefront error power spectrum, which is influenced by the turbulence profile, adaptive optics system performance, and instrumental aberrations. In the case of fully developed speckles, where the wavefront errors are much larger than the wavelength, the speckle intensity follows an exponential distribution. The spatial scale of the speckles is governed by the diffraction limit ($\lambda/D_{\text{aperture}}$), while their lifetime depends on the decorrelation timescale of the atmospheric turbulence and the instrument's thermal and mechanical stability. The presence of speckle noise significantly deteriorates the achievable contrast in high-contrast imaging, especially at small angular separations where the speckles

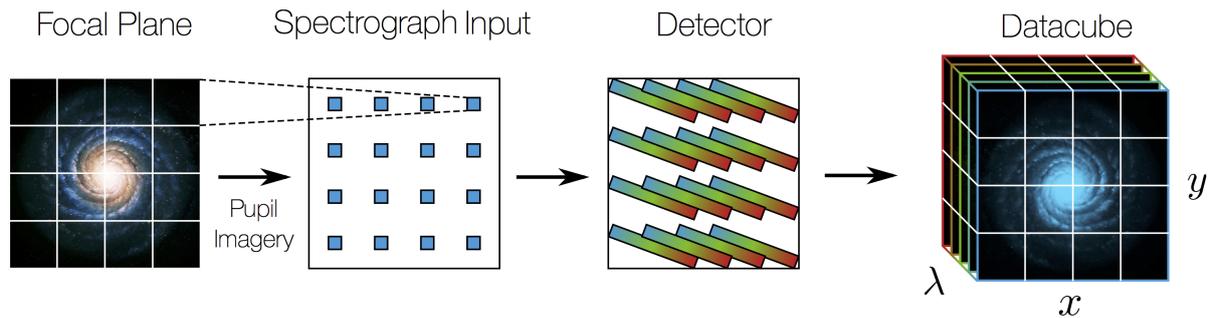
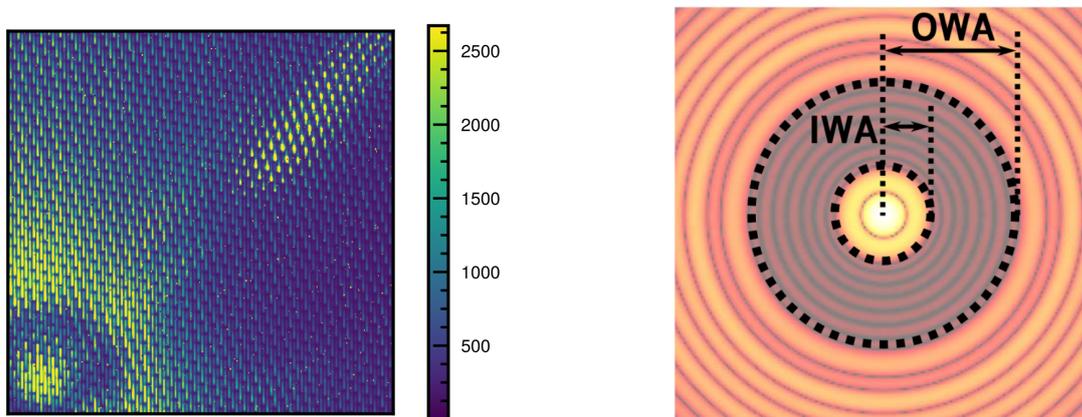


Figure 2.9: Lenslet array-type IFS. The incoming light is split up by lenslets where spectral elements separate the white light into their spectrum. For each lenslet (spatial dimensions) the spectrum is recorded by a column of detector pixels. The final data cube has axes in spatial coordinates (x, y) and wavelength λ . Adapted from [National Aeronautics and Space Administration, 2017](#).



(a) Cutout from an unprocessed observation of 51 Eridani b by a lenslet array-type IFS detector. The dispersed spectra from each lenslet are visible as interleaving vertical columns on the detector. The host star is in the lower left corner. Brighter color means higher flux. ([Samland et al., 2022](#))

(b) Dark zone geometry. The inner working angle (IWA) of coronagraph is the minimum angular separation where half a the flux of a source is still transmitted. The outer working angle (OWA) is defined equivalently but typically less critical. ([Por, 2020](#))

Figure 2.10

are most intense. The typical speckle contrast, defined as the ratio of the speckle intensity to the peak of the PSF, can range from 10^{-3} to 10^{-5} depending on the adaptive optics performance and instrumental design. Post-processing aims to subtract the speckle pattern and enhance the planet signal-to-noise ratio. However, the temporal evolution and chromatic dependence of speckles limit the effectiveness of current techniques. ([Bonse et al., 2018](#); [Galicher et al., 2023](#)) Methods that improve speckle suppression by capturing these cross-modal effects have recently been at the forefront of advances in direct exoplanet imaging. ([Cantero et al., 2023](#); [Dahlqvist et al., 2020](#); [Flasseur et al., 2020](#); [Samland et al., 2021](#))

2.2.4. Differential Imaging

Observations of exoplanets in HCI using ASDI feature three primary components, shown in [Figure 2.11](#). The most significant component is typically the coronagraphic point spread function (PSF) of the host star. The PSF is the intensity distribution on the detector after the light from a point-source passes through the telescope. While much of the light of the host star is filtered out by the coronagraph, it typically is still much brighter than its companions. The host star is in the center of the observations, on the optical axis of the telescope. In contrast, exoplanets need to be at a certain (angular) distance from the star to be observable, and are thus off the optical axis, or "off-axis".

In contrast to many other astronomical targets, increasing only the integration time for a direct exoplanet observation does not improve the image beyond a certain threshold. This is because after the fast changing effects discussed above have been averaged out a longer integration collects more photons of both the star and exoplanet, leaving the final contrast unchanged. While advanced instrumentation is crucial for direct

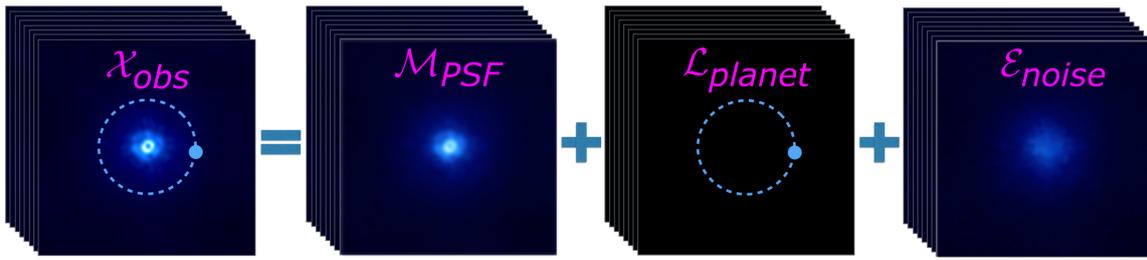


Figure 2.11: The primary components in high-contrast imaging observations of exoplanets: the coronagraphic point spread function \mathcal{M}_{PSF} of the host star at the center, the fainter off-axis exoplanet \mathcal{L}_{planet} , and noise together with other nuisance components \mathcal{E}_{noise} . (adapted from [Gonzalez et al., 2016](#))

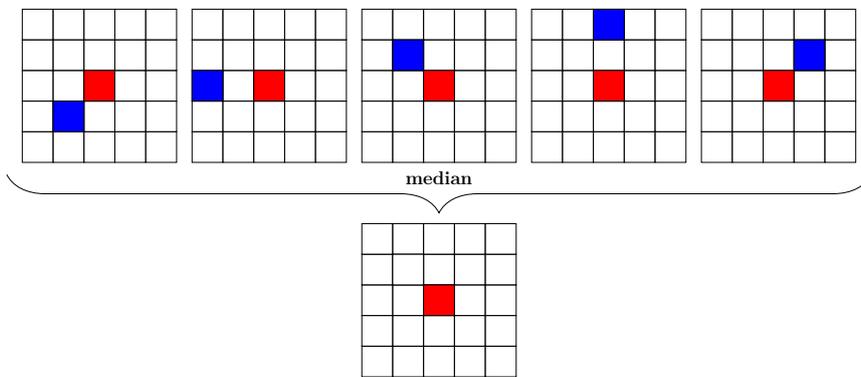
imaging, observational techniques are equally important for the detection of exoplanets. Two break-through observation techniques have to be used to separate the stellar and planet flux. These techniques are Angular Differential Imaging (ADI, [Marois et al., 2006](#)) and Spectral Differential Imaging (SDI, [Smith, 1987](#)). Both techniques aim at keeping the stellar PSF and speckle pattern as static as possible while changing the position of the off-axis companion as much as possible. These techniques are particularly effective in dealing with the quasi-static speckle patterns discussed above, as they help distinguish between the static stellar PSF and the moving planetary signal. If components are static (or constant) across multiple observations, a model can be easily found that only includes these static parts. Subtracting this model from the original observations leaves only the non-static components, i.e. the off-axis companion.

To illustrate how differential imaging techniques work, consider the example in [Figure 2.12](#), where there are 5 observations of a stationary components (red) and a moving component (blue). These observations are 5×5 -matrices, where the two colors would correspond to specific flux levels. Based on the 5 observations a model of the stationary parts can be constructed. In this example the median operation is applied to the observations to obtain the stationary component model. Subtracting the model from the observations leaves only the non-stationary components, called residuals. Had a different choice be made for the PSF model function, e.g. the mean instead of the median, the PSF model would "self-subtract" some of the non-stationary component flux. In direct exoplanet imaging, the red stationary component is the stellar PSF and the blue non-stationary component is the substellar companion.

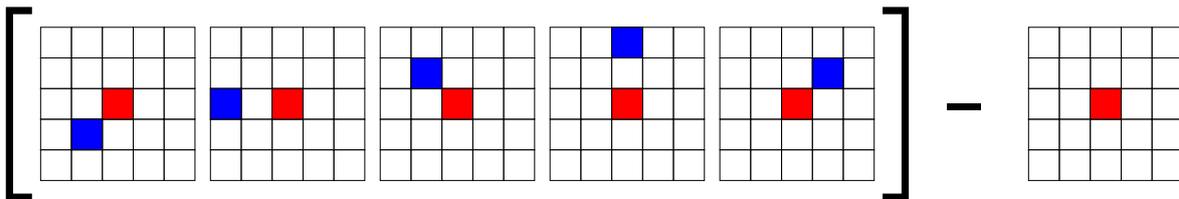
This example, while simplified, illustrates the fundamental principle behind overcoming the contrast and angular separation challenges discussed earlier in direct imaging of exoplanets. The example in [Figure 2.12](#) shows the ideal case where the PSF is full stationary across all observations, there is no overlap between the exoplanet position in other frames and no overlap between the stellar PSF and the exoplanet. Real observations are much more complex than this idealized example, particularly due to the evolving nature of speckle patterns discussed earlier. Non-stationary components need to be included in the PSF model while avoiding the inclusion of exoplanet flux. However, differentiating between non-stationary components belonging to the stellar PSF and the exoplanet is challenging. This is why PSF modeling methods that are more advanced than median and mean combination are necessary.

Currently, the majority of image post-processing methods aim to model and subtract the coronagraphic PSF of the host star and the speckle pattern from the observations, while retaining as much of the companion signal as possible. This is enabled by both angular and spectral differential imaging and is most effective when both are used together which is known as Angular-Spectral Differential Imaging (ASDI). ([Kiefer et al., 2021](#))

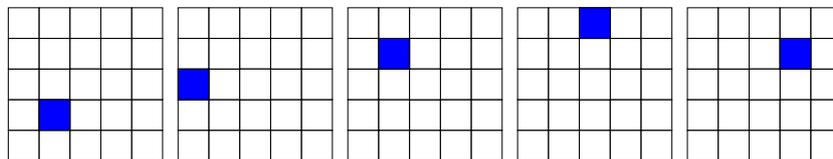
The angular diversity is obtained by centering the target star on the optical axis in pupil tracking mode while the telescope rotates under the sky. Pupil-tracking mode, in contrast to sky-tracking mode, keeps the optics of the instrument in the same alignment during the entire observation. In pupil-tracking mode, off-axis companions traverse a circular arc around the optical axis, as shown in the rows of [Figure 2.14](#) first column in [Figure 2.15](#), because of the earth's rotation under the sky. Furthermore, pupil-tracking mode is also necessary to keep the PSF mostly stationary because changes in the optical configuration lead to large changes of the PSF and speckle pattern. As explained in [Section 2.2.3](#), it is not possible to keep the optical system completely unchanged during an observation. Hence, the resulting PSF is called quasi-static as the stellar PSF and speckle pattern change over time.



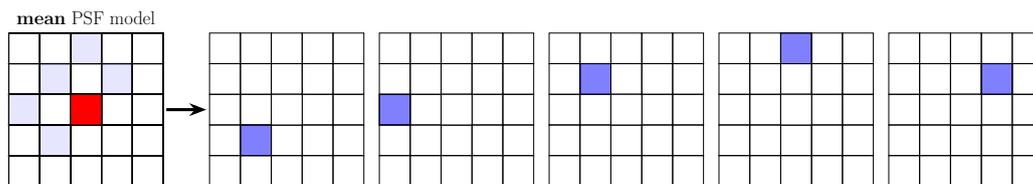
(a) The stationary components across the 5 observation frames (matrices) can be modeled by applying the median operator to the elements at the same position in the matrices.



(b) Subtracting the stationary component model from each observation frame.



(c) After subtracting the stationary components only the non-stationary components remain. These are also called residuals.



(d) Choosing a different PSF model, e.g. using the mean, alters the residuals. Since the mean model includes some "blue" signal, a part of the non-stationary component flux is subtracted. This is also called self-subtraction.

Figure 2.12: Modeling and subtracting stationary components in observations obtained by differential imaging to obtain a model of the non-stationary components in observations.

The spectral diversity is due to the PSF of point sources spreading out with increasing wavelength. This is due to the wavelength-dependent diffraction of light, see Equation 2.3, and is illustrated in Figure 2.5. The position of the point sources remains static over one observation, shown in the columns of Figure 2.14 and the first column of Figure 2.15. The spectral diversity is inherent to observations with integral field spectrographs. (Pueyo, 2018)

Info box 2.5: Angular Differential Imaging (ADI)

An observational technique that uses the rotation of the sky relative to the telescope pupil to distinguish between the quasi-static stellar PSF and orbiting companions. It allows for improved PSF subtraction and enhanced detection of faint, off-axis sources like exoplanets.

Info box 2.6: Spectral Differential Imaging (SDI)

An observational technique that exploits the wavelength dependence of the quasi-static stellar PSF to differentiate it from planetary companions. It uses simultaneous multi-wavelength observations to enhance the contrast between the star and potential exoplanets.

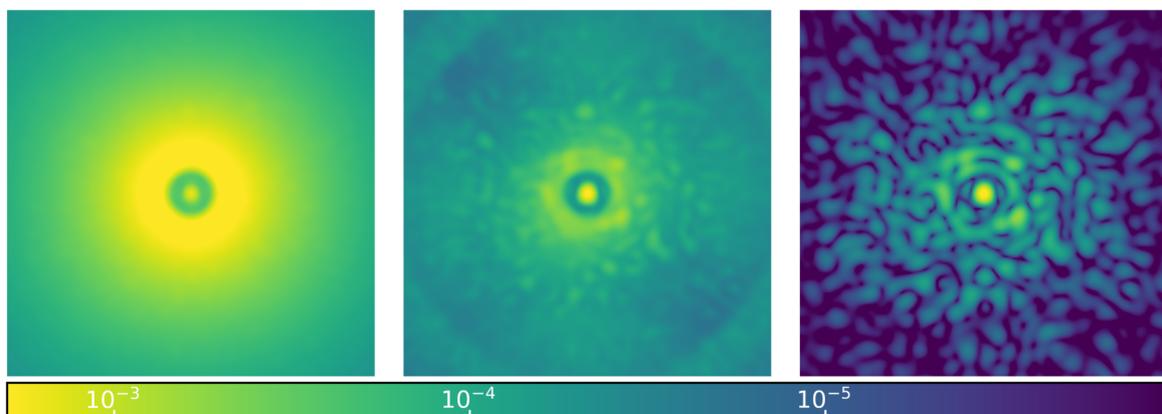
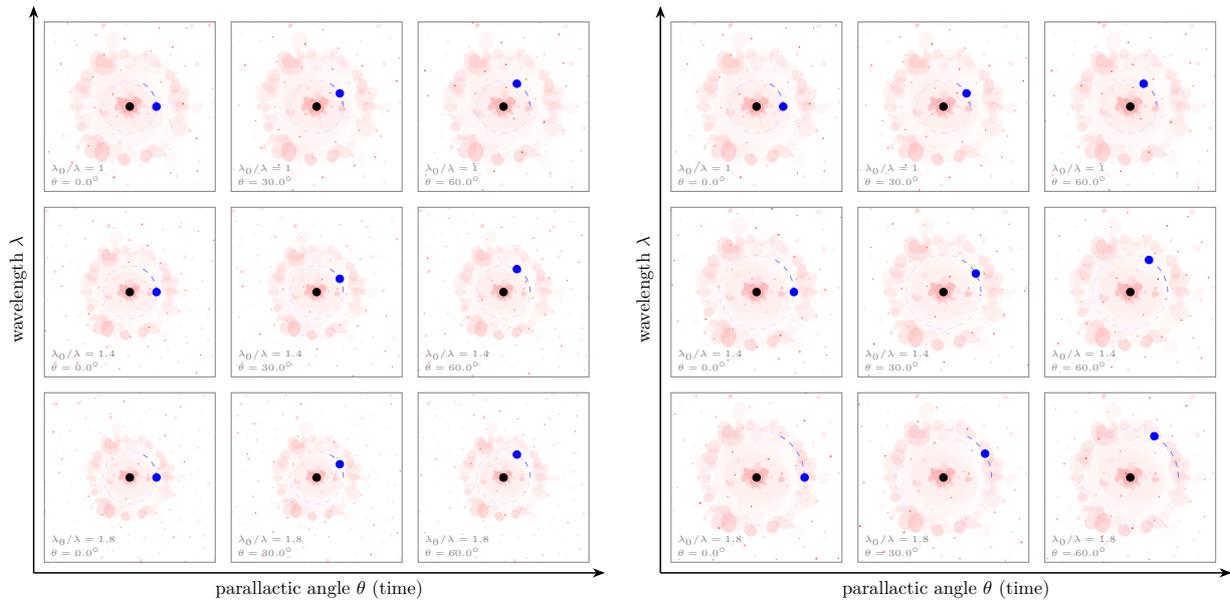


Figure 2.13: Focal plane normalized intensity of aberrations in a simulated VLT/SPHERE analogue using an apodized pupil Lyot coronagraph. Without adaptive optics the coronagraph is ineffective and the telescope is seeing limited like in the leftmost figure of Figure 2.6 (left). A realistic coronagraph working in concert with an adaptive optics system suppresses rapidly varying speckles due to atmospheric turbulence inside corrected region (center). An ideal adaptive optics system corrects all except for quasi-static speckles (right). (Galicher et al., 2023)

By exploiting both angular and spectral diversity the PSF of the host star can be modeled after scaling the observations (second column of Figure 2.15) so that the stellar PSF is aligned along the angular and spectral dimension. This misaligns the companion along both dimensions. The stellar PSF is modeled as the quasi-static part of the observations. The PSF model is then subtracted from the observations, ideally leaving only the companion signal, as shown in Figure 2.15. To improve the performance of PSF modeling techniques, the mean or median of the observations is typically subtracted beforehand.

ASDI exploits both angular and spectral diversity to enhance the detection of exoplanets. When both angular and spectral diversity are available Kiefer et al., 2021 found that modeling the combined differential imaging observations simultaneously consistently outperforms modeling them sequentially. The simultaneous modeling is known as Combined Differential Imaging (CODI) in contrast to Angular-Spectral Differential Imaging (ASDI), and Spectral-Angular Differential Imaging (SADI), which model and reduce the observations sequentially. In this work, ASDI refers to the observation technique, whereas CODI refers to reducing observations obtained by ASDI with a simultaneous modeling approach.

The algorithms presented in this work simultaneously exploit the dual diversity induced by angular and spectral differential imaging and are thus CODI methods. Nevertheless, they are also applicable to other differential imaging methods. However, the proposed methods are expected to particularly excel when modeling features in higher-dimensional data for which they are known to surpass matrix methods (Liu



(a) Pre-processed observations. The PSF and speckle pattern spread out with increasing wavelength due to the diffraction of light. Due to the rotation of the earth under the sky, the off-axis source moves on an arc through the observations.

(b) Scaled observations. By scaling the frames in Figure 2.14a by the ratio of a reference wavelength λ_0 (typically the largest wavelength in an observation) over the wavelength of a frame, also called the scale factor $\lambda_0/\lambda = s$, the PSF and speckle pattern is aligned throughout the entire cube. This also misaligns off-axis sources, both radially and azimuthally. Ideally, the PSF and speckle pattern is now the same in every frame and can be easily modeled.

Figure 2.14: Schematic observations obtained by angular-spectral differential imaging. Each frame (gray border) in the 3-by-3 grids is an image of the observed solar system at wavelength λ and parallax angle θ . The coordinate system in each frame is not wavelength and angle, but instead (projected) distances e.g. right ascension and declination. The PSF and speckle pattern are shown in red, and an off-axis source (like an exoplanet) is shown in blue, and its trajectory through the observation cube is shown as a dashed blue arc. The center of each frame, coinciding with the position of the star, is indicated by a black circle. While only 9 frames are shown here, full observations typically consist of thousands of frames. The grids shown here correspond to the columns in Figure 2.15.

et al., 2022; Wang et al., 2023). Hence, the expected performance difference is largest for higher-order tensors.

Info box 2.7: Angular-Spectral Differential Imaging (ASDI)

An observational technique that induces both angular and spectral diversity into direct imaging observations. This maximizes the differentiation between the stellar PSF and exoplanet signals. It leverages the benefits of both ADI and SDI to achieve improved contrast and detection capability in high-contrast imaging. The data obtained by ASDI is a 4D data-cuboid, and thus has a spectral, temporal (parallax angles) and two spatial modes. ASDI can also refer to post-processing methods which reduce the data first along the angular and then the spectral mode. If the post-processing is first along the spectral and then angular mode it is called SADI.

Info box 2.8: Combined Differential Imaging (CODI)

A post-processing technique for observations obtained by ASDI that simultaneously uses the angular and spectral diversity to build the PSF model. In the literature it is sometimes used as a synonym for ASDI, however in this work it specifically refers to the combined post-processing method.

Info box 2.9: ASDI vs CODI

In this work ASDI will refer to the observational technique of the data-products that result from it. CODI will refer to post-processing methods which are used on data obtained by ASDI (abbreviated as ASDI data or ASDI observations) and build a spectral and angular model simultaneously. This is used to distinguish CODI from techniques that build the spectral and angular models sequentially. In the literature ASDI and CODI are sometimes used interchangeably.

2.2.5. Post-Processing

The general goal for post-processing is to remove all remaining stellar signal (coronagraphic PSF, speckle pattern) from the observations or, equivalently, extract only the exoplanet signal. As explained above, this requires two elements; Firstly, information that helps in distinguishing the exoplanet signal from the stellar signal. This information is induced by differential imaging. The second requirement is a method that uses this information to build a model for the stellar or exoplanet component. This method is the post-processing algorithm, sometimes also called the reduction algorithm. The differential imaging techniques and the post-processing are closely linked; without post-processing differential imaging does not help with detecting exoplanets, and without differential imaging the post-processing methods do not work.

If the ASDI technique has been used for obtaining the observations, the angular and spectral diversity induced by it can be exploited when post-processing the observations. The development of advanced post-processing methods and their use in concert with the instrumentation and observatories described above has enabled the discovery of the majority of directly imaged exoplanets. Post-processing is the final step in direct imaging, and results in fully reduced residual images as shown in the rightmost figure of [Figure 1.1c](#) and [Figure 2.6](#). In contrast with the other aspects of direct imaging, improved post-processing methods enable better utilization of the previously collected data. ([Chomez et al., 2023a](#))

Before the data is post-processed it is typically pre-processed to transpose it from the matrix-like detector-format into a data-cube where each axis corresponds to a single index. The individual images are also sub-pixel centered on the star, artifacts such as hot pixels are removed and the detector flat-field is subtracted. Pre-processed observations produced by contemporary high-contrast imaging systems like VLT/SPHERE (see [Section 4.1](#)) are large data products, with the size of individual observations ranging from 500 megabyte to well over 20 gigabyte. The unprocessed, raw observation data can exceed this by a factor of a few. Auxiliary data which is required to further post-process the observations adds another few megabyte. The data products are generally stored in the `.fits` format ([Pence et al., 2010](#)). Compressing the data is counterproductive as it would lead to unrecoverable loss of signal. ([Delorme et al., 2017](#); [Vigan et al., 2017](#))

This work mainly considers the joint angular and spectral post-processing approach where a combined angular and spectral PSF model is found, known as CODI. For CODI, the general post-processing method is illustrated in [Figure 2.15](#) and proceeds as follows; first the median observation cube along the temporal mode is subtracted from the observations (not shown in [Figure 2.15](#) for clarity). Then the observations are scaled according to their spectral channel to align the Point Spread Function and misalign any off-axis sources, both radially and azimuthally. This misalignment induced by scaling observations is shown in [Figure 2.14](#) as well as [Figure 2.15](#) (first to second column). A PSF model is then constructed (column three in [Figure 2.15](#)) and subtracted from the scaled observations. Because off-axis sources are at different positions in each frame of the ASDI observation cube, they are not included in the PSF model. The residual cube (column four in [Figure 2.15](#)) is then re-scaled and derotated to align off-axis sources and misalign PSF residuals from the PSF subtraction (column five in [Figure 2.15](#)). Finally, the residual is median combined along the temporal mode to create a residual spectral cube (column six in [Figure 2.15](#)), or further mean combined to create a residual image. Detailed steps are provided in [algorithm 1](#).

The current standard for modeling the stellar PSF in high contrast imaging of exoplanets is principal component analysis (PCA, [Amara et al., 2012](#); [Soummer et al., 2012](#)). PCA is a fundamental technique in data science and statistics aimed at finding the most significant patterns and features in matrix-shaped data.

In PCA the dataset is linearly transformed to project data onto a subspace that explains the maximum variance of the data for each axis and typically features fewer dimensions than the original dataset.

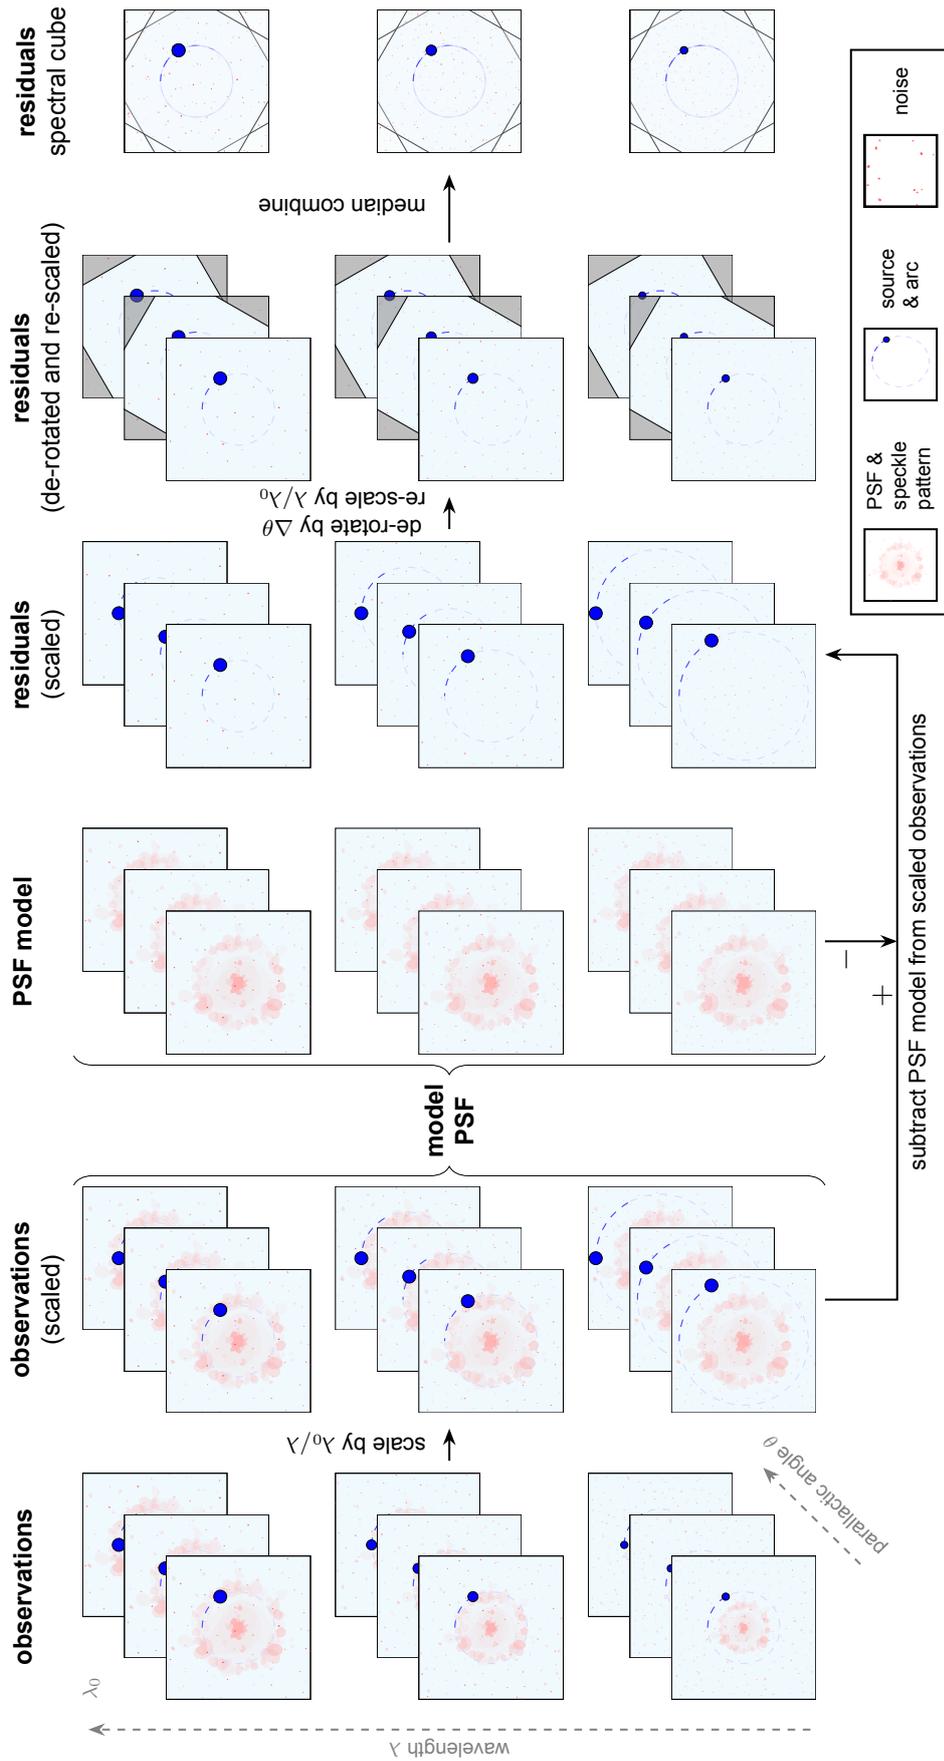


Figure 2.15: Post-processing angular spectral differential imaging (ASDI) observations. An observation obtained by ASDI is a 4-dimensional data cube, or an order-4 tensor. The data can be seen as frames (tensor slices, in light blue) in the spatial dimensions (modes) which vary with the angular and spectral dimensions (modes). These slices are centered on the host star point spread function (PSF) in light red. This figure uses depth to indicate the diversity along the angular mode and offset heights to indicate the diversity along the spectral mode. Post-processing the data to detect exoplanets proceeds as follows: First, the ASDI cubes are scaled to align the PSF across all modes. The coronagraphic PSF and the quasi-static speckle pattern are modeled using the scaled observation data. This PSF model can be found using a variety of methods, most commonly using matrix-PCA. The contribution of this work is the development of algorithms capturing higher-order, low-rank properties using tensor decompositions. The PSF model is then subtracted from the scaled observations, leaving only residuals which contain the exoplanet signal and noise. The residuals are then rescaled and derotated to align the planet in each frame. The signal to noise ratio can be further increased by mean-combining the residual spectral cube in the last column.

Algorithm 1: Combined differential imaging by PSF subtraction.

```

Input: observation cube  $\mathcal{X} \in \mathbb{R}^{I_\lambda \times I_\theta \times I_x \times I_y}$ ,
          spectral channels  $\lambda \in \mathbb{R}^{I_\lambda}$ ,
          parallactic angles  $\theta \in \mathbb{R}^{I_\theta}$ 
Output: residual cube  $\mathcal{R} \in \mathbb{R}^{I_\lambda \times I_x \times I_y}$ 
1 Initialize:
2  $\lambda_{ref} \leftarrow \max(\lambda)$ 
3  $\mathbf{s} \leftarrow \frac{\lambda_{ref}}{\lambda}$  // define linear scale factors
4 map :  $\mathbb{R}^{I_x \times I_y} \rightarrow \mathbb{R}^{I_x \times I_y}$  // scale & rotate (affine) transform w.r.t.  $x, y$  center
5 model :  $\mathbb{R}^{I_\lambda \times I_\theta \times I_x \times I_y} \rightarrow \mathbb{R}^{I_\lambda \times I_\theta \times I_x \times I_y}$  // PSF model function, e.g. PCA

6  $\mathcal{X} \leftarrow \mathcal{X} - \text{median}_\theta(\mathcal{X})$  // subtract median observation cube
7  $\mathcal{X}' \leftarrow \text{map}(\mathcal{X}; \mathbf{s}, \mathbf{0})$  // scale observation cube to align PSF
8  $\mathcal{M}_{\text{PSF}} \leftarrow \text{model}(\mathcal{X}')$  // construct scaled PSF model
9  $\mathcal{R}' \leftarrow \mathcal{X}' - \mathcal{M}_{\text{PSF}}$  // subtract scaled model PSF from scaled observation cube
10  $\mathcal{R} \leftarrow \text{map}(\mathcal{R}'; \mathbf{s}^{-1}, -\lambda)$  // rescale and derotate the scaled residual cube
11  $\mathcal{R} \leftarrow \text{median}_\theta(\mathcal{R})$  // collapse the residual cube along the parallactic angles
12 return  $\mathcal{R}$ 

```

The orthonormal basis vectors of this subspace are known as Principal Components (PC). By finding the components that account for the most variance, PCA reduces noise and redundancy, highlighting the underlying trends and relationships within the dataset. Notably, the PC are a set of uncorrelated variables. By reducing the dimensionality of the observations, PCA in HCI is used to capture the quasi-static components which describe the PSF and speckle pattern. PCA thus results in a low-dimensional, low-rank PSF model, which is then subtracted from the observations, as shown in [Figure 2.15](#). ([Bouwman et al., 2016](#)) What "low-rank" means and why a low-rank PSF model is desirable will be discussed in [Section 2.3.1](#) and [Section 3.1](#) respectively.

The implementation of PCA in HCI typically relies on the truncated matrix singular value decomposition (SVD). The SVD is a fundamental method in linear algebra and operates on matrices. Hence, PCA requires reshaping the original observations into a matrix form. For example, in CODI, the observation \mathcal{X} is a four-way (or four dimensional) array in $\mathcal{X} \in \mathbb{R}^{I_\lambda \times I_\theta \times I_y \times I_x}$, which needs to be reshaped into a matrix $\mathbf{X}_{[\lambda, \theta]} \in \mathbb{R}^{I_{(\lambda, \theta)} \times I_{(y, x)}}$. Here, $\lambda, \theta, y,$ and x represent the "dimensions" (or modes) of the observation, with $I_{(\cdot)}$ indicating the size of each mode. [Section 2.3](#) will introduce a mathematically more rigorous notion of this data format and [Section 2.3.2](#) will explain the reshaping operation.

Finding uncorrelated variables which maximize variance is the eigenvector problem in [Equation 2.4](#) and [Equation 2.5](#), where the estimate of \mathbf{M} is optimal in the least-squares sense. [Equation 2.4](#) which can be solved using the matrix truncated SVD in [Equation 2.5](#). ([Amara et al., 2012](#); [Gonzalez et al., 2016](#); [Soummer et al., 2012](#))

$$\min_M \|\mathbf{X} - \mathbf{M}\|, \text{ s.t. } \text{rank}(\mathbf{M}) \leq k, \quad (2.4)$$

where \mathbf{M} is the low-rank model of the observation matrix \mathbf{X} with rank k , also called N_{comp} in the HCI literature.

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^T, \quad (2.5)$$

where σ are the singular values of \mathbf{X} , and \mathbf{u} and \mathbf{v} are the left and right singular vectors.

Info box 2.10: Matrix-based Principal Component Analysis (PCA)

The standard post-processing technique in high-contrast imaging. It uses the singular value decomposition to model and subtract the stellar PSF from observations. It works by flattening multi-modal data into a matrix, identifying principal components that capture the most variance, and constructing a low-rank approximation of the quasi-static PSF. While effective, it does not preserve the structure of multi-modal data like ASDI observations.

An extension of PCA that addresses some of its limitations when applied to high-contrast imaging data is annular PCA. This technique applies principal component analysis separately to concentric annuli in the image, rather than to the full frame at once. The image is divided into thin annular regions, typically a few λ/D in width, and PCA is performed independently on each annulus to model and subtract the stellar PSF. This approach allows the PCA parameters, such as the number of components, to be optimized separately for different angular separations from the star. Annular PCA can provide better performance than full-frame PCA, especially at small separations, by accounting for the radially-varying PSF structure and the different degrees of companion motion due to field rotation at different separations. It helps mitigate the impact of the strong radial dependence of speckle behavior in A(S)DI sequences. By applying PCA in this localized manner, annular PCA can achieve improved detection sensitivity, particularly at small angular separations near the coronagraph's inner working angle. (Absil et al., 2013)

Deep learning methods have recently emerged as a promising approach for post-processing direct imaging observations. These techniques leverage neural networks to learn complex patterns and relationships in high-contrast imaging data. Supervised machine learning methods, like statistical models with convolutional neural networks (Cantero et al., 2023; Flasseur et al., 2024) and convolutional autoencoders (Wolf et al., 2023), have demonstrated improved detection and characterization capabilities compared to more traditional post-processing algorithms like PCA. By training on large datasets of simulated planetary signals and residual noise, these models capture features that discriminate true companions from speckle noise. However, the application of deep learning to direct imaging faces significant challenges due to the enormous data volumes involved, particularly for ASDI observations which can exceed tens of gigabytes per sequence. This places substantial computational demands on both training and inference. Methods and computational capabilities need to advance for deep learning to be fully adapted by the HCI community.

While PCA and other low-rank matrix decompositions have been extremely successful in HCI, the required flattening of the observations highlights a problem with its use for multi-way datasets such as observations with angular and spectral diversity. This is because the flattening of multi-modal observations into matrices can fail to preserve valuable structure in the observations. (Cichocki et al., 2015; Favier, 2021; Lathauwer et al., 1996) Hence, conventional PCA, and matrix models in general cannot discover hidden components in ASDI as they are inadequate for capturing the cross-couplings across the modes of the observations. While this work has stressed this fundamental limitation of matrix-PCA, it is not obvious. To illustrate this problem an example will be given in Section 2.5 after introducing some more necessary background and a solution in Section 2.3.

Furthermore, using matrix PCA can become inconvenient as the amount of data that needs to be processed for HCI increases further due to its computational complexity and storage requirements. This problem is not unique to direct imaging. Other fields such as hyper-spectral imaging, facial recognition and outlier detection, have faced with similar challenges. These fields have often adopted tensor methods to overcome these problems. (Kolda et al., 2009; Liu et al., 2022)

2.3. Multi-Linear Algebra

Tensors, which can be viewed as multi-dimensional arrays, play a crucial role in high-contrast imaging. Despite their relevance, the rich body of work on tensor, or multi-linear, algebra has not been recognized in the field or leveraged to post-process observations. Multi-linear algebra is the extension of linear algebra to higher-order tensors and provides a framework for manipulating multi-dimensional data arrays and understanding their properties.

Since the concept of multi-linear algebra has not yet been applied to the post-processing of exoplanet observations, this section will introduce the essential background and notations for both multi-linear algebra and fundamental tensor decompositions within the context of high-contrast imaging. Given the

extensive scope of multi-linear algebra, this overview will focus on the basics. Relevant further reading is recommended where appropriate. For an excellent introduction, the work by [Kolda et al., 2009](#), "*Tensor Decompositions and Applications*", is highly recommended. For more advanced introductory reading, chapters 1, 2, and 4 of [Liu et al., 2022](#), "*Tensor Computation for Data Analysis*", are also suggested. Lastly, "*Matrix and Tensor Decompositions in Signal Processing*" by [Favier, 2021](#) is recommended for its comprehensive, integrated, and comparative review of matrix and tensor factorization methods as well as their applications.

2.3.1. Tensors

Tensors are mathematical objects that generalize the concepts of scalars, vectors, and matrices to higher dimensions. They provide a natural framework for representing and analyzing datasets with multiple modes or attributes, such as time series of images, multi-sensor measurements, or user-item-context interactions in recommendation systems. Hence, tensors are used to capture and represent complex, multi-modal data in a way that preserves its inherent multi-dimensional relationships. ([Liu et al., 2022](#))

Motivating Tensor Models

Tensor approaches offer several significant advantages over traditional matrix-based methods in handling multi-modal data structures. These benefits stem from the inherent ability of tensors to preserve and exploit the natural multi-modal structure of data, leading to more nuanced and powerful analytical capabilities.

One key advantage is the essential uniqueness property exhibited by certain tensor decompositions. This property allows for unique solutions under relatively mild conditions, in contrast to matrix decompositions which often require more restrictive constraints like orthogonality or non-negativity to achieve uniqueness. This characteristic enables more robust and interpretable factorizations of multi-way data.

Tensor methods also excel in solving certain problems directly from measured signals, without the need for computing higher-order statistics or using extensive pilot sequences. This capability is particularly valuable in areas like communication channel identification, where tensor-based approaches can operate effectively with shorter signal recordings compared to traditional statistical methods that rely on estimating high-order moments or cumulants.

In the realm of blind source separation, tensor techniques offer a distinct advantage in addressing underdetermined systems - scenarios where the number of sources exceeds the number of sensors. This capability extends the applicability of source separation methods to a broader range of practical situations.

For handling large-scale datasets, tensor approaches provide efficient compression techniques through data tensorization and low multi-linear rank approximations. This not only reduces storage requirements but also preserves the inherent structure of the data, allowing for more faithful representations of complex, multi-modal information.

The flexibility of tensor representations in processing multi-modal data is another significant benefit. By treating each modality separately rather than vectorizing or matricizing the data, tensor methods preserve the multilinear structure and capture intricate interactions between modes. This approach allows for a more natural and comprehensive analysis of complex, multi-faceted datasets. Furthermore, tensor models can incorporate a greater number of modalities compared to matrix-based approaches. The increased modality enables the integration of more complementary information, potentially enhancing the performance of various systems in fields such as wireless communication, recommendation systems, diagnostics, and monitoring. The additional modalities can lead to improved detection, interpretation, recognition, and classification tasks.

In summary, tensor approaches provide a powerful framework for analyzing and manipulating multi-dimensional data, offering unique advantages in terms of decomposition uniqueness, direct signal processing, underdetermined system handling, data compression, multi-modal flexibility, and the incorporation of multiple information sources. These capabilities make tensor methods increasingly valuable across a wide range of scientific and engineering applications dealing with complex, multi-faceted datasets. ([Favier, 2020](#))

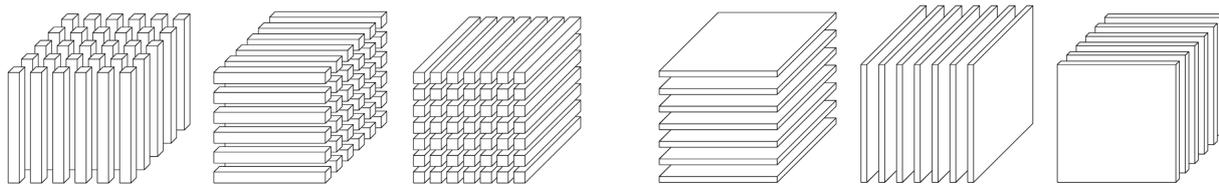
Order

A tensor is defined as a multi-dimensional array spanning the product of N vector spaces, called a tensor space. In that way tensors generalize several familiar concepts: a scalar (a single number) is a tensor of

order-0, a vector (a list of numbers) is a tensor of order-1, and a matrix (an array of numbers) is a tensor of order-2. The 'order' of a tensor refers to the number of indices required to uniquely identify each element within the tensor. For example, each element $x_{i,j}$ in a matrix $\mathbf{X} \in \mathbb{R}^{I \times J}$ with can be indexed by using two indices $i \in I$ and $j \in J$:

$$\mathbf{X} = \begin{pmatrix} x_{1,1} & x_{1,2} & x_{1,3} \\ x_{2,1} & x_{2,2} & x_{2,3} \\ x_{3,1} & x_{3,2} & x_{3,3} \end{pmatrix},$$

where $I = J = 3$. Higher-order tensors can be viewed to be constructed from these lower-order tensors, known as fibers (vectors, Figure 2.16a) and slices (matrices, Figure 2.16b).



(a) Mode-1 to 3 fibers: column ($\mathbf{x}_{:,jk}$), row, ($\mathbf{x}_{i:,k}$), and tube ($\mathbf{x}_{i,j,:}$), respectively.

(b) Horizontal ($\mathbf{X}_{i::}$), lateral ($\mathbf{X}_{:,j}$), and frontal ($\mathbf{X}_{::k}$) slices.

Figure 2.16: Fibers and slices in a third-order tensor generalize columns (or rows) of matrices by partially fixing indices. (Kolda et al., 2009)

For example an order-3 tensor $\mathcal{X} \in \mathbb{R}^{I \times J \times K}$ can be seen as being constructed from K slices, with each slice a matrix $\mathbf{X}_k \in \mathbb{R}^{I \times J}$:

$$\mathcal{X} = \begin{pmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \\ \mathbf{X}_3 \end{pmatrix},$$

where each slice \mathbf{X}_k is a matrix:

$$\mathbf{X}_1 = \begin{pmatrix} x_{1,1,1} & x_{1,2,1} & x_{1,3,1} \\ x_{2,1,1} & x_{2,2,1} & x_{2,3,1} \\ x_{3,1,1} & x_{3,2,1} & x_{3,3,1} \end{pmatrix}, \quad \mathbf{X}_2 = \begin{pmatrix} x_{1,1,2} & x_{1,2,2} & x_{1,3,2} \\ x_{2,1,2} & x_{2,2,2} & x_{2,3,2} \\ x_{3,1,2} & x_{3,2,2} & x_{3,3,2} \end{pmatrix}, \quad \mathbf{X}_3 = \begin{pmatrix} x_{1,1,3} & x_{1,2,3} & x_{1,3,3} \\ x_{2,1,3} & x_{2,2,3} & x_{2,3,3} \\ x_{3,1,3} & x_{3,2,3} & x_{3,3,3} \end{pmatrix},$$

where $I = J = K = 3$.

Here, and in the following, tensors of order-3 and higher will be denoted using calligraphic capital letters \mathcal{X} , whereas matrices will be in bold capital letters \mathbf{X} , vectors in bold lowercase letters \mathbf{x} , and scalars in lowercase letters x . Note that while 'dimension' is a familiar term in the field of post-processing differential imaging observations, this work will instead use the more rigorous term 'order' to describe the number of indices of a tensor. "Dimension" will instead refer to the number of possible values that an index can take. In the examples above the indices I , J , and K can each take up to three values so that their dimension is three.

Info box 2.11: Order

The order of a tensor, is the number of indices, or modes required to uniquely identify each element within the tensor. For example, a scalar is a tensor of order 0, a vector is a tensor of order 1, a matrix is a tensor of order 2, and so forth. A tensor $\mathcal{X} \in \mathbb{R}^{I \times J \times K}$ is order 3. In multi-linear algebra the word "dimension" denotes the range of values a single index (or mode) can take, so it will not be used to describe the order of a tensor in this work.

Mode

Up to now, the word "index" has been used to refer to the axes of a tensor. This was convenient when describing the notation that is used, however, in data science the axes are more commonly referred to as "modes" of a tensor.

In the context of tensor analysis, various interpretations of modes have been proposed by [Favier, 2021](#):

- a source of information (e.g., user, patient, client, trial)
- a type of entity attached to the data (e.g., items/products, types of music, types of film)
- a tag that characterizes an item, piece of music, film, etc.
- a recording modality that captures diversity in various domains (e.g., space, time, frequency, wavelength, polarization, color)

In the following this work will be referring to a tensors indices as modes as this is more natural in the context of processing observations. For angular spectral differential imaging, this will mainly be the following four modes: spectral mode (channels), temporal mode (parallactic angles), and two spacial modes (right ascension and declination). Hence, we interpretation of modes as a recording modality is most natural in the context of this work.

Info box 2.12: Mode (index)

The modes, or indices, of a tensor are the ways or axes along which the tensor can be indexed. An observation obtained by angular spectral differential imaging has four modes: a spectral mode, temporal mode (along which the parallactic angle varies), and two spacial modes along which the right ascension (RA) and declination (DEC) varies. The following indices refer to these modes in this work:

spectral mode (channels): λ	
temporal mode (parallactic angles): θ	
spatial mode 1 (right ascension): x ,	or radial mode (separation): ρ (or ϱ)
spatial mode 2 (declination): y ,	or azimuthal mode (azimuthal angle): ϕ

So that observations obtained by ASDI are written as $\mathcal{X} \in \mathbb{R}^{I_\lambda \times I_\theta \times I_x \times I_y}$ or $\mathcal{X} \in \mathbb{R}^{I_\lambda \times I_\theta \times I_\rho \times I_\phi}$.

Example 2.1: Tensor notation

Consider a video of a colorful bicyclists riding along a dike.^a Each frame of this video is a matrix of pixel values. These pixels are ordered into rows and columns. To show the colors of the riders pixel is expressed with one value for red, green, and blue each. Since this is a video recording, it has also a temporal axis.

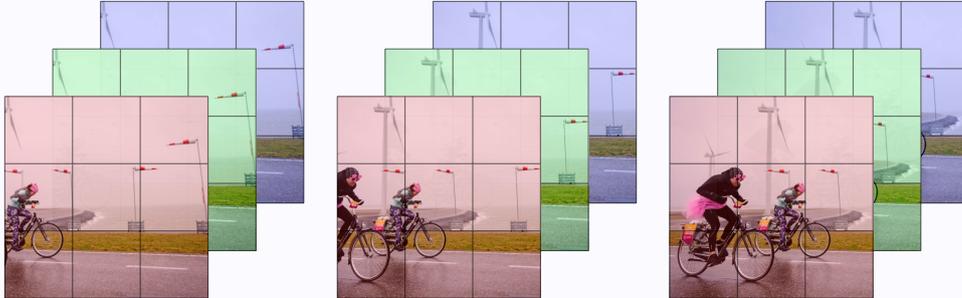


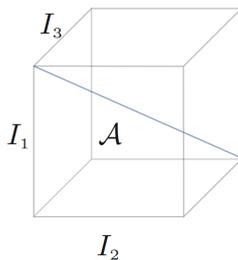
Figure 2.17: Example tensor.

This video is a tensor with two spacial modes (pixel rows and columns), a spectral mode, and a temporal mode. Since it has four modes in total it is an order-4 tensor.

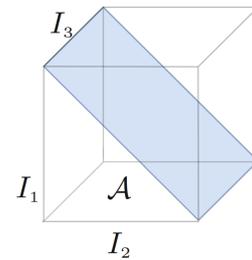
^aImage credit: AFP, 2023.

Diagonal Tensors

Two special types of tensors are the diagonal and f -diagonal tensors, shown in Figure 2.18. Like diagonal matrices, diagonal tensors are useful when factorizing tensors. A diagonal tensor $\mathcal{D} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ has non-zero elements d_{i_1, i_2, \dots, i_N} only on the superdiagonal where all indices are equal: $i_1 = i_2 = \dots = i_N$. An f -diagonal tensor $\mathcal{F} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is more general, with each frontal slice \mathcal{F}_i is a diagonal matrix. The diagonal and f -diagonal tensors occur in various applications and have useful algebraic properties. (Liu et al., 2022)



(a) Superdiagonal order-3 tensor, for which the elements on the diagonal (in blue) are non-zero, and all other elements are zero. (adapted from Liu et al., 2022)



(b) f -diagonal order-3 tensor. (adapted from Liu et al., 2022)

Figure 2.18: Diagonal tensors.

Rank

Like matrices, tensors have a rank. The rank of a matrix is defined as the maximum number of linearly independent columns in the matrix. In other words, the rank of a matrix is the minimum number of linearly independent vectors needed to span the column space of the matrix.

For the general tensor case the notion of a rank-1 tensor must be introduced first before the general case can be defined. A order- N , rank-1 tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is a tensor for which the following holds exactly

$$\mathcal{X} = \mathbf{u}^{(1)} \circ \mathbf{u}^{(2)} \circ \dots \circ \mathbf{u}^{(N)}, \quad (2.6)$$

where $\mathbf{u}^{(n)} \in \mathbb{R}^{I_n}$ are vectors in the mode- n and \circ is the vector outer product. That is, an rank-1 tensor can be written as a single outer vector product. For the previous order-3 case, this is depicted graphically in Equation 2.7:

$$\mathcal{X} \in \mathbb{R}^{I \times J \times K} = \mathbf{u}^{(1)} \circ \mathbf{u}^{(2)} \circ \mathbf{u}^{(3)}, \quad (2.7)$$

where the equality holds exactly. Its elements x_{i_1, i_2, \dots, i_N} are computed using the product of the corresponding vector elements u_{i_n}

$$x_{i_1, i_2, \dots, i_N} = u_{i_1}^{(1)} u_{i_2}^{(2)} \dots u_{i_N}^{(N)} \quad \forall 1 \leq i_n \leq I_n. \quad (2.8)$$

The rank of a tensor, denoted as $R = \text{rank}(\mathcal{X})$, is minimum number of rank-1 tensors needed to express the tensor exactly as their sum.

Info box 2.13: Rank

A rank-1 tensor is an order- N tensor that can be written as the outer vector product of N vectors $\mathbf{u}^{(n)}$:

$$\mathcal{X} = \mathbf{u}^{(1)} \circ \mathbf{u}^{(2)} \circ \dots \circ \mathbf{u}^{(N)}$$

The rank R of a tensor \mathcal{Y} is the minimum number of rank-1 tensors needed to express the tensor $\mathcal{Y} = \sum_{r=1}^R \mathbf{u}_r^{(1)} \circ \mathbf{u}_r^{(2)} \circ \dots \circ \mathbf{u}_r^{(N)}$.

While the concepts of matrix and tensor rank are analogous, one major difference between them is that there is no algorithm to determine the rank of a tensor in general. This means that, in practice the concept of a border rank is often used when finding rank- R approximations of tensors. The border rank of a tensor is the rank of a tensor that approximates another tensor up to arbitrary precision. (Kolda et al., 2009)

Info box 2.14: Border rank

The border rank of a tensor is relevant when considering approximations of tensors. The border rank of a tensor is the minimum number of rank-1 tensors needed to express the tensor up to an arbitrary error threshold. Formally, a tensor \mathcal{X} has border rank s if for every $\epsilon > 0$, there exists a tensor \mathcal{X}_ϵ of rank s such that the euclidean distance between \mathcal{X} and \mathcal{X}_ϵ is less than ϵ . The border rank of a tensor is always less than or equal to the rank of the tensor. When referring to the rank of a tensor of observations this work means the border rank of this tensor.

Low-Rank Tensor Models

Low-rank tensors are particularly useful in various fields, including signal processing, data compression, and machine learning, because they provide a way to represent high-dimensional data in a compact manner while preserving the important data features or removing noise. A low-rank tensor model approximates a tensor with a reduced number of ranks. Since the rank of a tensor in practical applications is generally not known, the notion of low-rank in practice is driven by the application and depends on the desired properties of the low-rank model.

Low-rank tensor models can be viewed from several compatible angles;

- Dimensionality reduction: Low-rank models reduce the number of parameters needed to represent the data. This helps in simplifying models and finding useful sub-spaces. This view of low-rank tensors will be used in Section 2.4.1 and Section 2.4.3 to compare them with matrix-PCA.

- **Noise reduction:** In many applications, data is corrupted by noise. Low-rank approximations can filter out noise by capturing patterns in the data and neglecting minor variations that may be due to noise. The importance of removing noise in a smart manner is highlighted in [Section 2.5](#).
- **Computational efficiency:** Operations such as multiplications, inversions, and factorizations are computationally cheaper for low-rank tensors. This efficiency is crucial for large-scale data analysis and real-time applications. This feature will be investigated in [Section 5.3.3](#).
- **Interpretability:** Low-rank models often provide more interpretable results. For instance, in the example of [Figure 2.17](#), foreground and background can be separated and treated as two independent objects. This will be discussed both in [Section 3.2.4](#) and [Chapter 8](#).
- **Compression:** Representing data using low-rank approximations allows for efficient storage. For example, in image compression, a low-rank approximation can significantly reduce the file size while maintaining visual quality. In [Section 3.2.4](#) the impact of different ways to compress tensors will be highlighted.

Info box 2.15: Low-rank tensor

Low-rank (reduced-rank) tensors approximate high-dimensional tensors with lower-dimensional structures, capturing the most significant patterns in the data.

2.3.2. Operations on Tensors

In the subsequent chapters two specific operations on tensors will become important, namely, the mode- n unfolding of a tensor, and the tensor-tensor product.

Matrization & Tensorization

Matrization (also known as unfolding or flattening) reshapes a tensor by reducing its order. This reshaping is typically done along a specific mode, see [Section 2.3.2](#) where this is explained in detail. At maximum, tensors can be flattened into a vector. This vector can be seen as the least informative description of the original tensor, where each element of the tensor is its own dimension. In this way, a fully flattened tensor can be viewed as a point in a very-high dimensional space, such as illustrated in [Figure 2.26a](#). Tensorization (or folding) is the inverse operation, where the tensor is reshaped and the resulting tensor is of higher-order. Tensorization creates new modes, whereas unfolding combines or mixes modes. Both unfolding and tensorization operations must respect the original tensor shape so that the total number of elements is unchanged. Both matrization and tensorization are very common operations in image post-processing and multi-linear algebra as demonstrated by [Section 2.2.5](#) and [Section 2.4.7](#).

This work indicates general unfolding by writing the indices of the modes that have been combined in round brackets (\dots) . For example, unfolding the spatial modes x and y of an observation obtained by ASDI $\mathcal{X} \in \mathbb{R}^{I_\lambda \times I_\theta \times I_x \times I_y}$ would be indicated like

$$\mathcal{X} \xrightarrow{\text{unfold } x,y} \mathcal{X}', \text{ where } \mathcal{X}' \in \mathbb{R}^{I_\lambda \times I_\theta \times (I_x \cdot I_y)}, \quad (2.9)$$

thus reducing the tensor order from four to three. This work puts further emphasis on tracking the modes by writing combined modes, also known as multi-indices, as $(m_1 \cdot m_2)$, where the brackets indicate that this is a single mode consisting of the unfolded modes m_1 and m_2 . ([Liu et al., 2022](#)) The representations introduced in [Section 2.3.3](#) will greatly simplify this notation.

Mode- n Unfolding

The mode- n unfolding operator on a tensor is a specific matrization which reshapes an order- N tensor into a matrix where the n^{th} mode becomes the row(s) and all other modes become the column(s). The mode- n unfolding of a tensor $\mathcal{A} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_n \times \dots \times I_N}$ is written as $\mathbf{A}_{[n]}$:

$$\mathbf{A}_{[n]} \in \mathbb{R}^{I_n \times (I_1 \dots I_{n-1} I_{n+1} \dots I_N)} \quad (2.10)$$

The elements of $\mathbf{A}_{[n]}$ are mapped from \mathcal{A} as follows:

$$\mathbf{A}_{[n]}(i_n, j) = \mathcal{A}(i_1, i_2, \dots, i_N) \quad (2.11)$$

where j is the column index in the unfolded matrix, computed as:

$$j = 1 + \sum_{k=1, k \neq n}^N (i_k - 1)J_k, \quad \text{with } J_k = \prod_{m=1, m \neq n}^{k-1} I_m. \quad (2.12)$$

Example 2.2: Mode- n unfolding - worked tensor

Consider the example given by [Kolda et al., 2009](#). The frontal slices of a tensor $\mathcal{X} \in \mathbb{R}^{3 \times 4 \times 2}$ are:

$$\mathbf{X}^{(1)} = \begin{bmatrix} 1 & 4 & 7 & 10 \\ 2 & 5 & 8 & 11 \\ 3 & 6 & 9 & 12 \end{bmatrix}, \quad \mathbf{X}^{(2)} = \begin{bmatrix} 13 & 16 & 19 & 22 \\ 14 & 17 & 20 & 23 \\ 15 & 18 & 21 & 24 \end{bmatrix}$$

The three mode- n unfoldings are then:

$$\mathbf{X}_{[1]} = \begin{bmatrix} 1 & 4 & 7 & 10 & 13 & 16 & 19 & 22 \\ 2 & 5 & 8 & 11 & 14 & 17 & 20 & 23 \\ 3 & 6 & 9 & 12 & 15 & 18 & 21 & 24 \end{bmatrix}$$

$$\mathbf{X}_{[2]} = \begin{bmatrix} 1 & 2 & 3 & 13 & 14 & 15 \\ 4 & 5 & 6 & 16 & 17 & 18 \\ 7 & 8 & 9 & 19 & 20 & 21 \\ 10 & 11 & 12 & 22 & 23 & 24 \end{bmatrix}$$

$$\mathbf{X}_{[3]} = \begin{bmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 \\ 13 & 14 & 15 & 16 & 17 & 18 & 19 & 20 & 21 & 22 & 23 & 24 \end{bmatrix}$$

Example 2.3: Mode- n unfolding - graphical

Consider the unfolding operation shown in Figure 2.19 that is used to illustrate the post-processing of spectral differential imaging observations with PCA in Figure 1.2. In this example, the observation tensor is $\mathcal{X}_{\text{Obs}} \in \mathbb{R}^{5 \times 3 \times 3}$, so that the transposed observation matrix $\mathbf{X}_{\text{Obs}}^{\top}$ is the mode-3 unfolding $\mathbf{X}_{[3]}$ of the observation tensor:

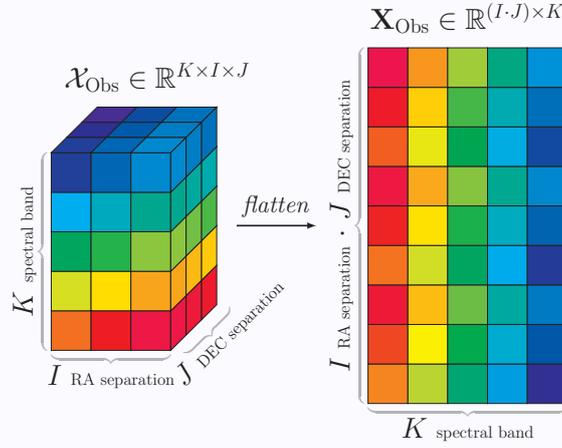


Figure 2.19: Graphic example of a mode- n unfolding.

In Figure 2.19 the mode-3 the mode-3 unfolding is transposed to make the data easier to process with PCA.

Unfolding mode-1 of an observation obtained by ASDI $\mathcal{X} \in \mathbb{R}^{I_\lambda \times I_\theta \times I_x \times I_y}$ reshapes it to be $\mathbf{X}_{[1]} = \mathbf{X}_{[\lambda]} \in \mathbb{R}^{I_\lambda \times (I_\theta \cdot I_x \cdot I_y)}$.

The mode- n unfolding is a crucial operation in many tensor decomposition algorithms, as it allows for the application of matrix-based techniques to higher-order tensors. (Liu et al., 2022)

Contracted Tensor Product

The contracted tensor product generalizes matrix multiplication to higher-order tensors with a set of shared indices. For tensors $\mathcal{A} \in \mathbb{R}^{I_1 \times \dots \times I_N \times J_1 \times \dots \times J_L}$ and $\mathcal{B} \in \mathbb{R}^{J_1 \times \dots \times J_L \times K_1 \times \dots \times K_M}$, their contracted product $\mathcal{C} = \mathcal{A} \bullet \mathcal{B} \in \mathbb{R}^{I_1 \times \dots \times I_N \times K_1 \times \dots \times K_M}$ is defined as:

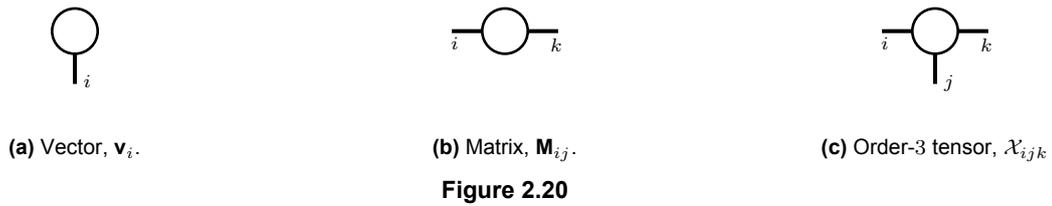
$$c_{i_1, \dots, i_N, k_1, \dots, k_M} = \sum_{j_1, \dots, j_L} a_{i_1, \dots, i_N, j_1, \dots, j_L} b_{j_1, \dots, j_L, k_1, \dots, k_M}. \quad (2.13)$$

This operation contracts (sums over) the L shared indices $\{J_1, \dots, J_L\}$ between \mathcal{A} and \mathcal{B} . If the contraction is over a single index n it is also known as a mode- n contraction (or product). (Liu et al., 2022)

2.3.3. Tensor Network Notation

In order to simplify the notation of tensor algebra and the representation of both tensors, their decomposition and multi-linear functions, tensor network notation and tensor network diagrams (TND) can be used. For a complete introduction to tensor network notation the excellent lecture notes “*Hand-Waving and Interpretive Dance*” by Bridgeman et al., 2017 are recommended. In TND, tensors are represented by shapes, typically circles, with lines (also called legs) representing the tensor indices.

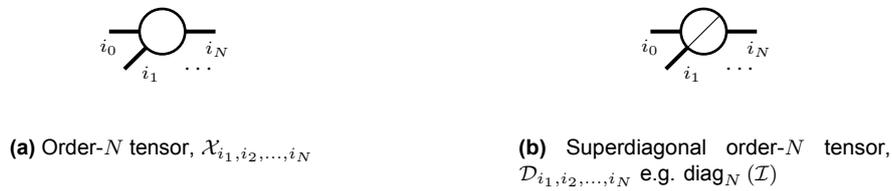
A vector \mathbf{v}_i with index i , is represented as a node with a single leg (Figure 2.20a), a matrix \mathbf{M}_{ij} with indices i, j is a node with two legs (Figure 2.20b), and a order-3 tensor \mathcal{X}_{ijk} as a node with three legs (Figure 2.20c). An observation data cube obtained by angular-spectral differential imaging is written in



TND as Equation 2.14:

$$\mathcal{X}_{\lambda,\theta,x,y} = \text{Diagram of a circle with four legs labeled } \lambda, \theta, x, y, \quad (2.14)$$

where the indices $\lambda, \theta, x,$ and y respectively correspond to the spectral, temporal, and two spatial modes. More generally, an order- N tensor with indices I_1 through I_N has N legs (Figure 2.21a), and the previously mentioned superdiagonal tensor is indicated with a diagonal line through the node (Figure 2.21b).



This notation is very helpful for communicating operations involving higher-order tensors, as the more familiar multi-linear algebra quickly becomes difficult to read. Consider the matrix product

$$\mathbf{AB} = \mathbf{A}_{ij} \mathbf{B}_{jk}, \quad (2.15)$$

where $\mathbf{A} \in \mathbb{R}^{I \times J}$ and $\mathbf{B} \in \mathbb{R}^{J \times K}$. This product is graphically written as Equation 2.16:

$$\text{Diagram of two circles labeled A and B connected by a horizontal line, with legs i, j, k.} \quad (2.16)$$

where each node has two legs, and they share index I . The connected leg indicates one or more shared indices over which tensors are contracted. The contraction shown in Equation 2.16 is known as a contraction along mode J , and a mode- n contraction (in Equation 2.13) in general:

$$\text{Diagram of a circle with n legs labeled i_0, i_1, ..., i_{n-1}, i_n and another circle with leg j connected to i_n.} \quad (2.17)$$

The mode- n product is an important operation in multi-linear algebra as it appears frequently when working with tensors. A leg connecting a node to itself denotes its (partial) trace:

$$\text{Diagram of a circle with a loop on one leg labeled x.} \quad (2.18)$$

which is simply the sum of the elements on the diagonal. (Bridgeman et al., 2017; Google/TensorNetwork 2023; Liu et al., 2022)

Often, the indexing in TND is dropped in favour of more concise notation. This work will keep indices when they are physically meaningful or necessary for describing the modes or a transform of the data.

After this quick tour of multi-linear algebra the next chapter will continue with one of its primary applications: factorizing higher-order tensors.

2.4. Tensor Decompositions

Tensor decompositions are a fundamental unsupervised machine learning method that extend linear algebra tools like the Singular Value Decomposition (SVD) and Principal Component Analysis (PCA) to multi-linear algebra. These methods decompose tensors into smaller components, revealing underlying structure, provide a compressed representation that approximates the original tensor, or make it easier to interpret.

Their ability to preserve the higher-order structure in multi-modal data makes them valuable across various fields, including psychometrics, signal processing, and deep learning (Liu et al., 2022). Despite their limited application in astronomy, recent advancements have led to tentative uses in radio interferometry (Sørensen et al., 2023), adaptive optics control (Joseph, 2024; Siquin et al., 2018).

In this work, tensor decompositions are used to model the coronagraphic PSF in ASDI observations. The proposed methods model the low-rank structure in the observations and do not rely on flattening the observations. This section introduces the decompositions, presenting them in four complementary ways to provide a comprehensive understanding:

- **Multi-linear equations:** The traditional notation of each decomposition, showing how they generalize matrix factorizations to higher-order tensors. This allows for a rigorous definition and comparison of the different methods.
- **Tensor Network Diagrams (TND):** Building on the TND notation introduced in the previous section, this section uses their graphical representation to visualize the structure of each decomposition. TNDs provide an intuitive way to understand the connection between different factors and factorizations.
- **3D graphics:** This section utilizes three-dimensional visualizations to illustrate how decompositions work on order-3 tensors. These graphics help in understanding how the decompositions separate and represent different modes of variation in the data. However, these graphics are limited to the order-3 case and only used for the decompositions that benefit from them.
- **Vector spaces:** Some decompositions are represented in terms of vector spaces, showing how they project the high-dimensional tensor data onto lower-dimensional subspaces. This perspective helps in understanding how these methods achieve dimensionality reduction and capture the most significant features of the data. Not all decompositions benefit from such a representation.

By presenting the tensor decompositions through these four lenses, this section aims to provide a comprehensive understanding of their mathematical foundations, practical implementation, and geometric interpretation. This diverse approach will be particularly useful for readers from the high-contrast imaging field, as it bridges the gap between the familiar matrix-based methods and the new tensor-based approaches. The application of tensor decompositions in the context of post-processing direct imaging data will be discussed in detail in [Section 3.2](#).

The following subsections will expand on specific tensor decomposition methods, starting with the Canonical Polyadic Decomposition (CPD), followed by the Tucker Decomposition, Tensor Train Decomposition (TTD), and Tensor Ring Decomposition (TRD). For each method, the relevant subsection lays out its formulation, properties, and potential applications in the context of exoplanet direct imaging.

Info box 2.16: Tensor Decomposition

Data analysis techniques that extend matrix factorization methods to higher-order tensors. These methods break down complex, high-dimensional data into simpler components, revealing underlying structures and patterns. Tensor decompositions preserve the multi-way nature of data, allowing for the simultaneous analysis of interactions between different modes. They are particularly useful in fields dealing with large, multi-modal datasets, offering improved data compression, feature extraction, and pattern recognition capabilities compared to traditional matrix-based approaches.

2.4.1. Canonical-Polyadic

The Polyadic Decomposition (PD) is a method to decompose a tensor as a sum of rank-one tensors, shown in [Figure 2.22](#). Consider the order- N tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$. Most fundamentally, this tensor can be approximated to arbitrary precision by a rank- R decomposition, also called canonical-polyadic decomposition (CPD, [Kolda et al., 2009](#); [Tucker, 1966](#)). Using multi-linear notation,

$$\mathcal{X} = \sum_{r=1}^R \lambda_r \mathbf{u}_r^{(1)} \circ \mathbf{u}_r^{(2)} \circ \dots \circ \mathbf{u}_r^{(N)} \quad (2.19a)$$

$$= \sum_{r=1}^R \lambda_r \mathcal{U}_r \quad (2.19b)$$

$$= \mathbf{\Lambda} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \dots \times_N \mathbf{U}^{(N)} \quad (2.19c)$$

$$= \llbracket \boldsymbol{\lambda}; \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \dots, \mathbf{U}^{(N)} \rrbracket, \quad (2.19d)$$

where $\boldsymbol{\lambda} = [\lambda_1 \ \lambda_2 \ \dots \ \lambda_R]$ is a vector of the scaling factors $\lambda_r \in \mathbb{R}$ so that $\mathbf{\Lambda} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$, $\mathbf{\Lambda} = \text{diag}_N(\boldsymbol{\lambda})$ is a superdiagonal tensor, and $\mathbf{u}_r^{(n)} \in \mathbb{R}^{I_n}$ are the r th column of a factor matrix $\mathbf{U}^{(n)} \in \mathbb{R}^{I_n \times R}$. Each vector product in [Equation 2.19a](#) is a rank-1 tensor \mathcal{U} in [Equation 2.19b](#). ([Favier, 2020](#))

The fundamental concept expressed by [Equation 2.19a](#) through [Equation 2.19d](#) is that tensors can be approximated as a sum of vector products. This is analogous to approximating matrices using matrix factorizations. However, the tensor decomposition approach enables the modeling of the higher-order structure in tensor-shaped data. Specifically, [Equation 2.19a](#) illustrates this concept by explicitly highlighting the vectors that constitute each rank-1 tensor.

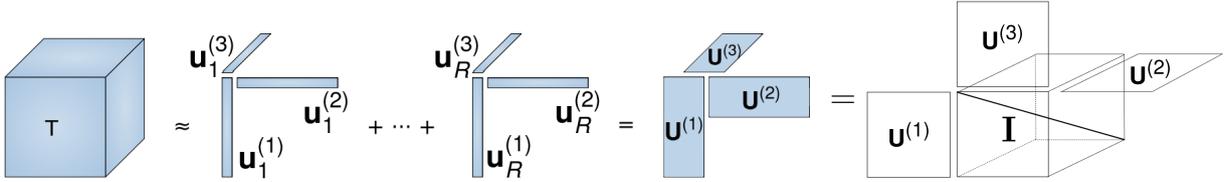


Figure 2.22: Approximating a third-order tensor (left) with a rank- R Canonical-Polyadic decomposition as a sum of vector products (middle-left), factor matrices (middle-right, [Vervliet et al., 2016](#)), and b) as product of an identity core tensor $\mathcal{I} = \text{diag}_3(1)$ and factor matrices ([Equation 2.19](#), right, [Sidiropoulos et al., 2017](#)). In b) the core is an identity tensor because the weights $\boldsymbol{\lambda}$ have been absorbed into the factors. Note that the CPD is a special case of a Tucker decomposition with non-orthogonal factor matrices [Sidiropoulos et al., 2017](#), see [Figure 2.28](#).

It should be stressed that in this work a data-driven approach is adopted, as this is common when post-processing ASDI observations ([Flasseur et al., 2018](#); [Gonzalez et al., 2016](#)). This means that the vectors in [Equation 2.19a](#) are learned from the data itself. [Section 2.5](#) demonstrates what is meant by a "data-driven" approach. It will also demonstrate why the tensor decomposition framework is intrinsically more constrained than current matrix-based methods by showing how data-driven CPD factor estimation recovers a set of basis functions. Naturally, choosing basis functions for the vectors or matrices in [Equation 2.19a](#) and [Equation 2.19c](#) would also be possible. Nevertheless, constrained tensor decompositions are possible and a promising direction for future work, see [Chapter 8](#). ([Favier et al., 2014](#))

[Equation 2.19b](#) emphasizes that due to the data-driven nature of the model, the individual factors are insufficient on their own to model the tensor effectively. Instead, the tensor model must be viewed as a whole because the rank-1 tensors are computed to jointly approximate the tensor. This means the models are not sequential; the vector product in a rank-1 CPD is not the same as the first rank-1 vector product in a rank- N decomposition. This is demonstrated in [Figure 3.2](#), where individual rank-1 components would yield a poor PSF model (consider the factors in the azimuthal mode), but together they form an excellent model, as shown in [Figure 3.3](#). Finally, [Equation 2.19d](#) provides a more concise notation for writing [Equation 2.19c](#), common in the literature.

Info box 2.17: Canonical-Polyadic Decomposition (CPD)

A tensor decomposition method that represents a higher-order tensor as a sum of rank-one tensors. It expresses the data as a combination of simple, interpretable components, each formed by the outer product of vectors. CPD is particularly useful for extracting meaningful patterns from high-dimensional data.

Equation 2.19 can be written as the TND in Equation 2.20. In TND the contraction in Equation 2.19c need to be shown in their full tensor form, so that the superdiagonal core tensor from Figure 2.22 connects the CPD factor matrices $\mathbf{U}^{(n)}$.

$$\begin{array}{c} i_1 \\ \diagdown \\ \mathcal{X} \\ \diagup \\ i_2 \quad \dots \quad i_N \end{array} \approx \begin{array}{c} \text{diag}_N(\boldsymbol{\lambda}) \\ \diagdown \\ \mathbf{U}^{(1)} \quad \mathbf{U}^{(2)} \quad \dots \quad \mathbf{U}^{(N)} \\ \diagup \\ i_1 \quad i_2 \quad \dots \quad i_N \end{array} \quad (2.20)$$

Fundamentally, a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ can be represented as a point $\vec{\mathcal{X}}$ in a $(I_1 \times I_2 \times \dots \times I_N)$ -dimensional space spanned by basis vectors $\hat{\mathbf{x}}_i$. This point $\vec{\mathcal{X}}$ and space is shown in Figure 2.26a. Observations obtained by ASDI with instruments like SPHERE are in a roughly $10^7 - 10^9$ -dimensional space. Figure 2.26 illustrates how the combination of individual factors of the CPD approximates the observations. Each rank-1 component is in the same vector space as the original tensor, and each column of the CPD factor matrices is in a sub-space corresponding to its mode. This figure shows that only the sum of factors gives the best rank- R approximation of the tensor.

2.4.2. Motivating Example

To motivate the tensor decompositions in the context of HCI, consider the CPD model in Figure 2.24 which shows the rank-1 decomposition of an off-axis source in the residual spectral cube illustrated in Figure 2.23. This residual cube was obtained by post-processing observations of β Pictoris obtained by ASDI. This results in a spectral residual data cube from which most of the PSF and speckle pattern has been removed. Only the stars companion, β Pictoris b, and residual speckle noise remains in the data. The mean combined residual frame is shown in the left frame of Figure 2.23. The residual speckle noise is better visible in the right frame of Figure 2.23.

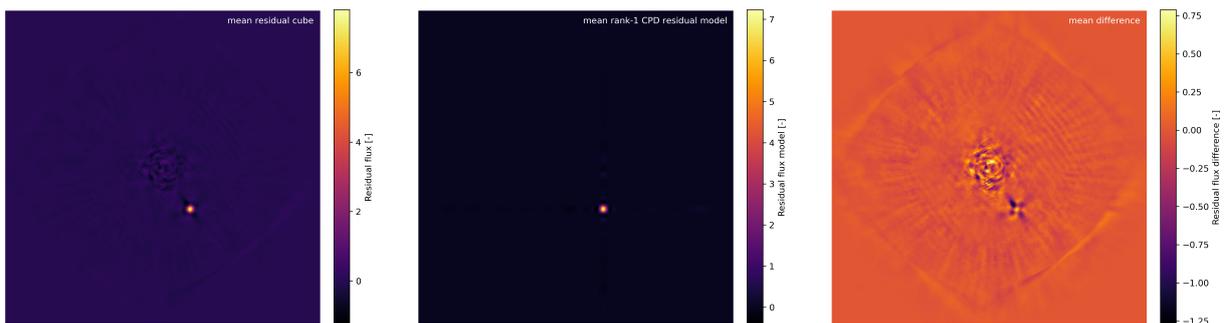
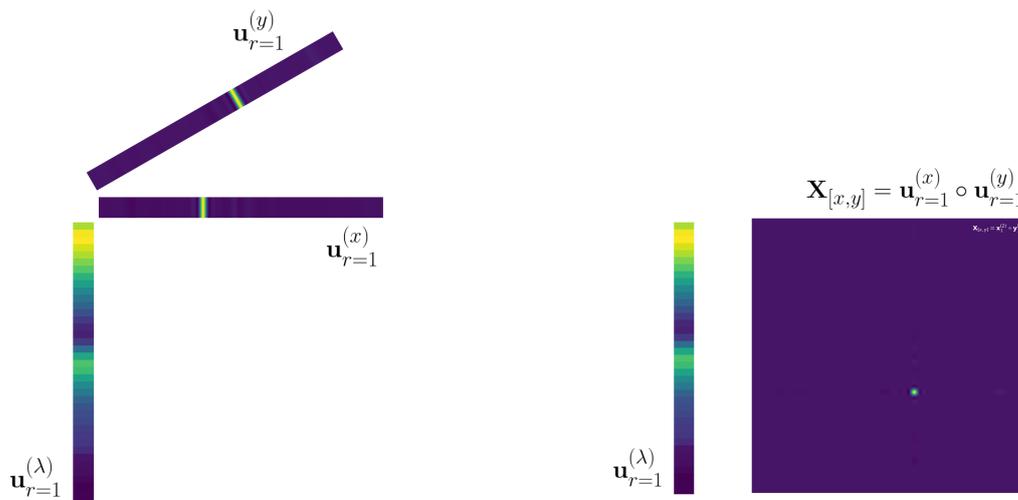


Figure 2.23: Mean combined residuals from a rank-1 CPD of the residual spectral cube of β Pictoris. The observations were reduced using a full-frame Tensor-Train decomposition. The residual image after post-processing the observations with the full-frame Tensor-Train decomposition shows the bright companion β Pictoris b (left). The rank-1 Canonical-Polyadic decomposition includes the bright companion, but not the residual noise at small separations (center). The difference between the previous images shows that the PSF model approximates β Pictoris b well, and does not include most of the self-subtraction lobes.

The residual spectral cube, an order-3 tensor, can be approximated as the vector product of rank-1 CPD factors; vectors in the spectral and spatial modes shown in Figure 2.24a. Taking the product of the spatial

factors shows the off-axis source as a rank-1 matrix in [Figure 2.24b](#), with its spectrum modulated by the spectral factor in the same figure.



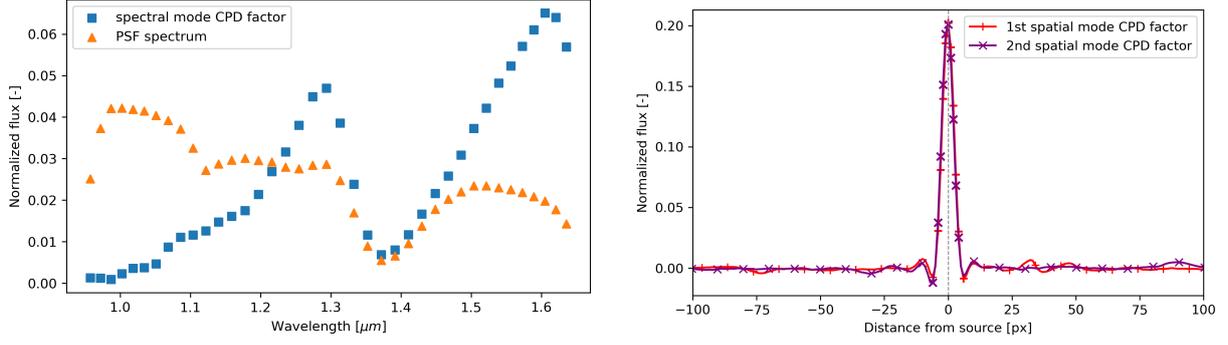
(a) Rank-1 CPD factors are vectors capturing most of the signal in the spectral (λ) and two spatial (x, y) modes. The product of these factors is an approximation of the residual spectral cube of β Pictoris containing only β Pictoris b. In this example, the spatial factors $\mathbf{u}^{(x)}$ and $\mathbf{u}^{(y)}$ have a much higher dimension than the spectral factor $\mathbf{u}^{(\lambda)}$.

(b) Multiplying the spatial factors shows the spatial component as a matrix $\mathbf{X}_{[x,y]}$. The full residual tensor (order-3) is reconstructed by multiplying the spatial matrix with the spectral factor along the spectral mode.

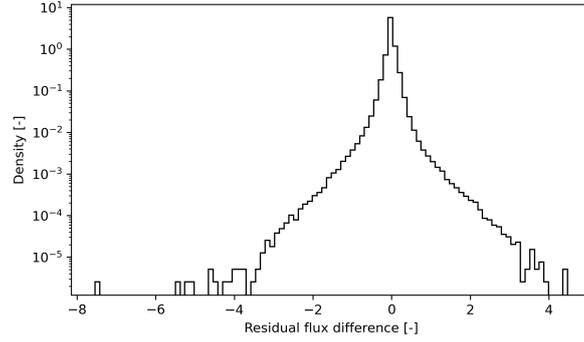
Figure 2.24: A stationary point source with spectrum $f(\lambda) \approx \mathbf{u}^{(\lambda)}$ in a spectral-spatial data cube can be approximated by a rank-1 CPD. In the figures above, this CPD is graphically represented by its factors. The color-scale is normalized for each factor, where a brighter element indicates a higher value. The CPD factors shown here correspond to the factorization of the residual spectral cube of β Pictoris in [Figure 2.25](#) and [Figure 2.23](#).

The same rank-1 CPD factors shown graphically in [Figure 2.24a](#) are shown in [Figure 2.25a](#) to highlight the difference between the spectrum of β Pictoris b (blue) and the spectrum of the host star β Pictoris (orange), both attenuated by the Earth's atmosphere. The common attenuation is most visible at the $\sim 1.4\mu\text{m}$ dip due to water vapour in the Earth's atmosphere.

[Figure 2.25b](#) shows the spatial factors translated so that the maximum of a 2D Gaussian fitted to the source is at zero. In both spatial modes the source has been partially self-subtracted which is indicated by the factors being negative. The distribution of the difference between model and residual ([Figure 2.23](#)) has both a central peak due to the small residuals at large separations and a bulge at larger residual values due to the residual speckle noise at small separations. This is due to the difference of the noise statistics at small and larger separations. ([Bonse et al., 2023](#); [Mawet et al., 2014](#))



(a) The residual spectral factor of the CPD model is clearly differentiated from the con-coronagraphic stellar spectrum (see also Figure 2.24a). (b) Spatial factors peak at the source location and include self subtraction lobes (see also Figure 2.24a).



(c) Distribution of the difference between the rank-1 CPD model and residual spectral cube.

Figure 2.25: Factorizing a residual spectral cube of β Pictoris b using a rank-1 CPD. The observations were reduced using a full-frame Tensor-Train decomposition.

2.4.3. Tucker

The CPD is limited in its representation power by the superdiagonality of the core tensor (\mathcal{I} in Figure 2.22b, or Λ in Equation 2.19). This means that the CPD components only consider interactions of the same columns of its factors. The Tucker decomposition (Tucker, 1966) generalizes the CPD to non-diagonal core tensors and PCA to higher-order analysis.

In Figure 2.22, the CPD is shown with a superdiagonal core tensor \mathcal{I} , which restricts the interaction of the factor matrices. The intuition behind the Tucker decomposition is to extend the core to include off-diagonal elements. This means each component considers the interactions between distinct columns of its factors. This achieves several objectives. Firstly, the sum in Equation 2.19a is relaxed to the matrix-tensor product in Equation 2.21a, allowing the factor matrices to project the core along each mode. In contrast, the rank-1 component tensors (vector products in Equation 2.19a) complement each other only through summation. The projection of the core tensor quantifies the interactions between the modes in a manner similar to matrix-PCA. Finally, viewing the Tucker decomposition as a compression algorithm, it has greater expressive power than a CPD. (Cohen et al., 2016) "Expressive power" refers to the number of parameters (elements in factors and/or cores) needed to approximate a tensor to an arbitrary threshold. (Favier, 2021; Liu et al., 2022)

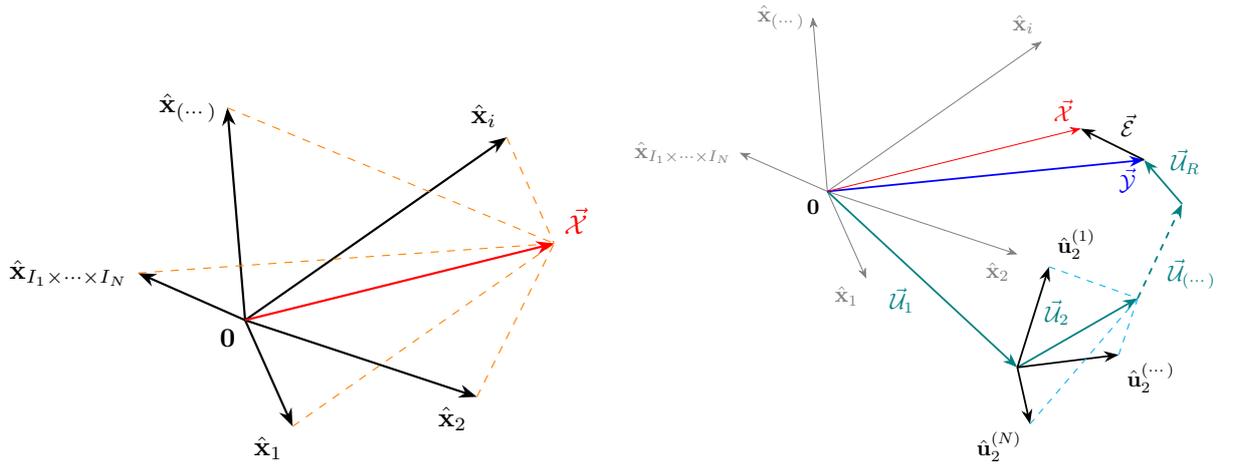
The Tucker decomposition of a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is defined in Equation 2.21 and Equation 2.22,

$$\mathcal{X} = \sum_{r_1=1}^{R_1} \sum_{r_2=1}^{R_2} \dots \sum_{r_N=1}^{R_N} g_{r_1 r_2 \dots r_N} \mathbf{u}_{r_1}^{(1)} \circ \mathbf{u}_{r_2}^{(2)} \circ \dots \circ \mathbf{u}_{r_N}^{(N)} \quad (2.21a)$$

$$= \mathcal{G} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \dots \times_N \mathbf{U}^{(N)} \quad (2.21b)$$

$$= \left[\mathcal{G}; \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \dots, \mathbf{U}^{(N)} \right], \quad (2.21c)$$

where $\mathbf{u}_{r_n}^{(n)} \in \mathbb{R}^{I_n}$ are the r th column of factor matrix $\mathbf{U}^{(n)} \in \mathbb{R}^{I_n \times R_n}$, and $\mathcal{G} \in \mathbb{R}^{R_1 \times R_2 \times \dots \times R_N}$ is the core



(a) Observation space. An observation $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is a point $\vec{\mathcal{X}}$ in a high-dimensional space spanned by basis vectors $\hat{\mathbf{x}}_i$.

(b) Approximation of $\vec{\mathcal{X}}$ with a sum of R outer vector products \vec{u}_r . Each vector product is in a subspace of \mathcal{X} , and not necessarily orthogonal.

Figure 2.26: Approximating a tensor with the CPD. In this representation a tensor \mathcal{A} is represented by a vector $\vec{\mathcal{A}}$ which is the original tensor reshaped into a vector.

(or interaction) tensor which is of order- N corresponding with the N -ranks of the Tucker decomposition. \times_n is the mode- n product. The similarity between Equation 2.21b and Equation 2.19c highlights the close connection between the CP and Tucker decomposition. Again, Equation 2.21c is a concise way of writing Equation 2.21b which is common in the literature.

$$\begin{array}{c} i_1 \\ \circlearrowleft \mathcal{X} \\ \circlearrowright i_2 \dots i_N \end{array} \approx \begin{array}{c} i_1 \\ \circlearrowleft \mathbf{U}^{(1)} \\ \circlearrowright R_1 \end{array} \begin{array}{c} \circlearrowleft R_2 \\ \circlearrowright \mathcal{G} \\ \circlearrowleft R_N \end{array} \begin{array}{c} \circlearrowleft R_N \\ \circlearrowright \mathbf{U}^{(N)} \\ \circlearrowright i_N \end{array} \quad (2.22)$$

Info box 2.18: Tucker Decomposition

A tensor decomposition that generalizes PCA to higher-order tensors. It represents a tensor as a core tensor multiplied by factor matrices along each mode. This method allows for the flexible modeling of multi-way data interactions, capturing complex relationships between different dimensions. It generalizes principal component analysis to higher-order data.

Equivalence for Matrix Data

To demonstrate that the Tucker decomposition generalized PCA to higher-order tensors and reduces to PCA in the order-2 case, consider the principal components (PC) \mathbf{T} of matrix $\mathbf{X} \in \mathbb{R}^{N \times M}$,

$$\mathbf{T} = \mathbf{X}\mathbf{W}, \quad (2.23)$$

where \mathbf{W}^T is the whitening transformation so that the columns of $\mathbf{W} \in \mathbb{R}^{M \times M}$ are the eigenvectors of $\mathbf{X}^T \mathbf{X}$. In HCI the PC are often truncated to find a rank- R approximation of \mathbf{X} , keeping only the first R eigenvectors in \mathbf{W} ;

$$\mathbf{T}_R = \mathbf{X}\mathbf{W}_R, \quad (2.24)$$

where $\mathbf{W}_R \in \mathbb{R}^{M \times R}$ and $\mathbf{T}_R \in \mathbb{R}^{N \times R}$.

Comparing the Tucker decomposition Equation 2.21b with PCA Equation 2.23 or SVD (Equation 2.5), one

finds that by taking identity matrices \mathbf{I} as the trailing factor matrices in the Tucker decomposition,

$$\mathcal{X} = \llbracket \mathcal{G}; \mathbf{U}, \mathbf{I}, \dots, \mathbf{I} \rrbracket = \mathcal{G} \times_1 \mathbf{U} \tag{2.25a}$$

$$\Leftrightarrow \mathbf{X}_{[1]} = \mathbf{U} \mathbf{G}_{[1]}, \tag{2.25b}$$

it is equivalent to PCA. This is why the Tucker decomposition is said to generalize PCA and SVD to higher-order tensors. (Liu et al., 2022)

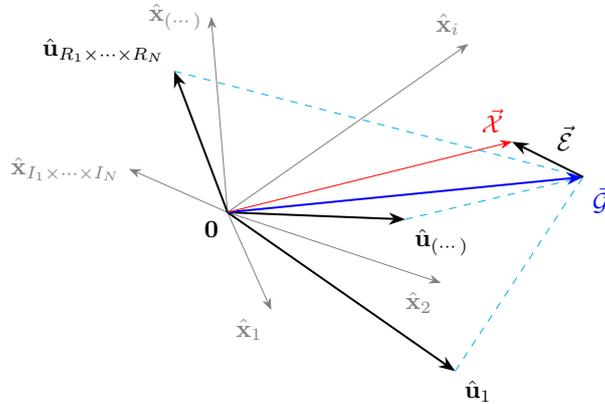
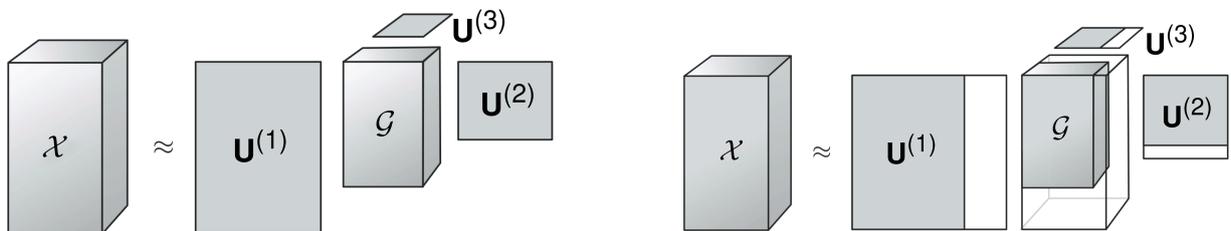


Figure 2.27: Tucker decomposition approximation space. The Tucker core \mathcal{G} is rotated and scaled by the Tucker factor matrices $\mathbf{U}^{(n)}$ so that the point $\vec{\mathcal{G}}$ is the closest point to $\vec{\mathcal{X}}$ in the lower-dimensional subspace spanned by the columns $\mathbf{u}_{r_n}^{(n)}$ of $\mathbf{U}^{(n)}$. The columns are shown here as basis vectors $\hat{\mathbf{u}}_i$. In this representation, a tensor \mathcal{X} is represented by a vector $\vec{\mathcal{X}}$ which is the original tensor reshaped into a vector.

Generalizing Matrix-PCA to Higher-Order Data: Subspace Projection View

The Tucker decomposition finds a core \mathcal{G} and factor matrices $\mathbf{U}^{(n)}$. The core is in a lower-dimensional subspace defined by the columns $\mathbf{u}_{r_n}^{(n)}$ of $\mathbf{U}^{(n)}$, called the approximation space. The objective of the Tucker decomposition is to find a core and factors so that the distance $\vec{\epsilon}$ between the point $\vec{\mathcal{G}}$ defined by the core and the original tensor is minimal in a least-squares sense. This is shown in Figure 2.27, where the columns of the factors are indicated as basis vectors $\hat{\mathbf{u}}_i = \mathbf{u}_{r_n}^{(n)} / \|\mathbf{u}_{r_n}^{(n)}\|$. The dimension of the approximation space is the product of the ranks of the Tucker decomposition, $R_1 \times R_2 \times \dots \times R_N$, which can be much lower than the original observation so that only important features are captured. In the case of matrix-PCA, the subspace in Figure 2.27 would be spanned by the principal components and instead of the Tucker core the point $\vec{\mathcal{G}}$ would be the projection of the observation onto this lower dimensional subspace. Note, that while Figure 2.27 graphically illustrates how the Tucker decomposition generalizes PCA, it is a simplified representation. In contrast with Figure 2.27, the core and factors are actually computed by repeatedly projecting mode- n unfoldings of the data onto a subspace of the specific mode until convergence or sufficiently small error. This will be elaborated on in Section 2.4.7. (Kolda, 2006; Lu et al., 2008; Pajarola et al., 2016)



(a) Tucker decomposition of an order-3 tensor \mathcal{X} into a core tensor \mathcal{G} and factor matrices \mathbf{A} , \mathbf{B} , and \mathbf{C} (Equation 2.21, Kolda et al., 2009).

(b) Truncated Tucker decomposition of an order-3 tensor \mathcal{X} into a truncated core tensor \mathcal{G} and truncated factor matrices \mathbf{A} , \mathbf{B} , and \mathbf{C} (Equation 2.21, Kolda et al., 2009).

Figure 2.28: Approximating a third-order tensor with a Tucker decomposition. Comparing Figure 2.22 with Figure 2.28 shows that the Tucker decomposition generalizes the CPD.

Block-Term Decomposition

The CPD and Tucker decomposition are further generalized by the Block-Term decomposition (BTD) which is the sum of Tucker decompositions (De Lathauwer, 2008a,b; De Lathauwer et al., 2008). Combining the CPD and Tucker decomposition means the BTD shares advantages of both decompositions. Nevertheless, it has not been implemented in the libraries used in this work and thus is not considered further. The CPD, TD, BTD and related decompositions (e.g. tensor-SVD, De Lathauwer et al., 2000) and hot-SVD (Wang et al., 2022b) all generalize the SVD to higher-order tensors. Each is a valid, but different interpretation of the matrix SVD for higher-order tensors. While the BTD is not used in this work, Chapter 8 highlights a potential avenue for its application in HCI.

2.4.4. Tensor-Train

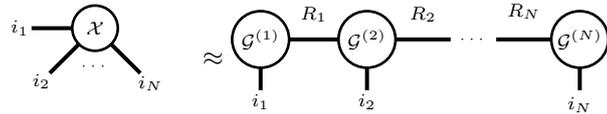
The Tensor Train Decomposition (TTD, Oseledets, 2011) has become an essential algorithm in low-rank decompositions in deep learning and physics. TTD is known as the matrix product state in the (quantum) physics community where it is used to model the state of N-body quantum systems (Biamonte, 2020). It factorizes a tensor into a sequence of order-3 cores, see Equation 2.27. Due to the use of repeated factorization into order-3 cores, the TTD captures low-rank properties while remaining parameter efficient in higher orders as there is a linear increase in parameters with the tensor order. In contrast with the Tucker decomposition of an order- N tensor which has an order- N core, no TTD core has an order larger than 3. This means that to approximate a tensor with I^N elements, I being the dimension of the N modes, the Tucker decomposition needs $\mathcal{O}(R^N)$ parameters, whereas the TTD only requires $\mathcal{O}(NIR^2)$ elements.

The TTD of an order- N tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is defined as follows:

$$\mathcal{X} = \sum_{r_1=1}^{R_1} \dots \sum_{r_{N+1}=1}^{R_{N+1}} \mathcal{G}^{(1)}(r_1, :, r_2) \circ \mathcal{G}^{(2)}(r_2, :, r_3) \circ \dots \circ \mathcal{G}^{(N)}(r_N, :, r_{N+1}) \quad (2.26a)$$

$$= \mathcal{G}^{(1)} \bullet \mathcal{G}^{(2)} \bullet \dots \bullet \mathcal{G}^{(N-1)} \bullet \mathcal{G}^{(N)}, \quad (2.26b)$$

where each $\mathcal{G}_n \in \mathbb{R}^{R_n \times I_n \times R_{n+1}}$ for $n = 1, \dots, N$ are the order-3 cores, with the added constraints ($R_1 = R_{N+1} = 1$) for the ends of the TT and \bullet is the tensor-tensor product using the folding operation. To ensure the proper dimensionality of the factorization the ranks of neighbouring cores need to be consistent. For neighbouring cores the trailing and next leading mode have dimension r_n so that $\mathcal{G}^{(n-1)}(r_{n-1}, :, r_n) \circ \mathcal{G}^{(n)}(r_n, :, r_{n+1})$. This also means that the first and last core of the TTD are matrices.



The diagram illustrates the Tensor Train Decomposition (TTD). On the left, a tensor \mathcal{X} is shown as a circle with N legs representing modes i_1, i_2, \dots, i_N . This is shown to be approximately equal to a sequence of cores $\mathcal{G}^{(1)}, \mathcal{G}^{(2)}, \dots, \mathcal{G}^{(N)}$ connected in a chain. Each core $\mathcal{G}^{(n)}$ is a circle with three legs: one vertical leg representing mode i_n , and two horizontal legs representing ranks R_{n-1} and R_n . The first core $\mathcal{G}^{(1)}$ has a rank R_1 connecting to the second core $\mathcal{G}^{(2)}$, and so on, up to the last core $\mathcal{G}^{(N)}$ which has a rank R_N connecting to the next core. The equation is labeled (2.27).

Due to the train-like structure of the decomposition the computation of the TTD is efficient, which is an especially important property for very high-dimensional tensors. However, this also leads to the main disadvantage of the TTD; the cores are computed according to the ordering of the indices of the original tensor. This means that the decomposition is not permutation invariant so that pairs of cores are generally not separable (Zhao et al., 2016).

Info box 2.19: Tensor Train Decomposition (TTD)

A tensor decomposition method that represents a high-order tensor as a sequence of interconnected low-order tensors. This "train" of core tensors allows for efficient representation of higher-order data while maintaining a relatively low number of parameters. TTD is particularly useful for very high order problems (order ≥ 4).

2.4.5. Tensor-Ring

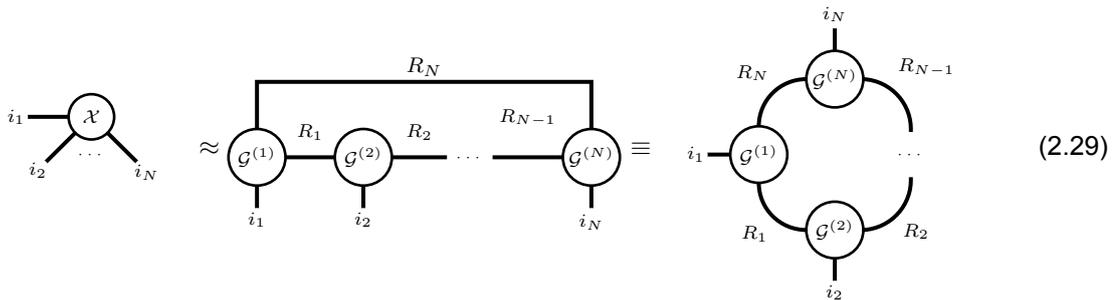
The Tensor Ring Decomposition (TRD, Zhao et al., 2016) builds on the TTD, addressing some of its limitations and introducing a more flexible structure. TRD represents an order- N tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ as a sequence of 3rd-order tensors, with the leading and trailing factors connected. This connection is

archived with the tensor-trace operation $\text{Tr}(\mathcal{X})$. This essentially means that the TRD extends TTD by enforcing periodic boundary conditions on Equation 2.26. The TRD of a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is

$$\mathcal{X} = \sum_{r_1=1}^{R_1} \dots \sum_{r_N=1}^{R_N} \mathcal{G}^{(1)}(r_1, :, r_2) \circ \mathcal{G}^{(2)}(r_2, :, r_3) \circ \dots \circ \mathcal{G}^{(N)}(r_N, :, r_1) \quad (2.28a)$$

$$= \text{Tr}[\mathcal{G}^{(1)}, \mathcal{G}^{(2)}, \dots, \mathcal{G}^{(N)}], \quad (2.28b)$$

where $\mathcal{G}^{(n)} \in \mathbb{R}^{R_n \times I_n \times R_{n+1}}$ again are the core factors. The relation between the TTD and TRD, as well as the effect of the trace operation is shown in the TND of the decompositions, Equation 2.27 and Equation 2.29 respectively: the TTDs leading and trailing cores are matrices, the corresponding cores of the TRD are order-3.



Info box 2.20: Tensor Ring Decomposition (TRD)

An extension of the Tensor Train Decomposition where the first and last cores are connected, forming a circular structure. This "ring" format provides more flexibility in modeling correlations between different modes and can capture more complex patterns in the data compared to TTD.

2.4.6. Other Decompositions

While this work covers many of the fundamental decompositions, many extensions have been developed. Most notable in relation to HCI are the non-negative tensor factorization (NTF) methods like the non-negative CPD and Tucker decomposition (NTD) as well as robust tensor factorization (RTD/RTF) methods. Additionally, tensor decompositions can be easily combined into hybrid and hierarchical formats, e.g. by decomposing the cores of a TTD using the CPD. Recently, a tensor analysis of variances (ANOVA) decomposition based global variance partitioning using Sobol-TT has been proposed (Ballester Ripoll et al., 2018, 2019) which might be of interest in HCI. This work verified that the NTD, RTD and ANOVA decomposition can be used for PSF modeling, however, the full evaluation of these methods is beyond the scope of this work.

2.4.7. Computation

The computation of tensor decompositions is a crucial aspect of their application in high-contrast imaging. This section provides an overview of the computational methods for finding factorizations, focusing on direct computation techniques and iterative optimization approaches. The importance of error terms and optimization methods in handling non-Gaussian data distributions which are often encountered in real-world applications are also discussed.

Direct Computation

Tensor decompositions are computed by systematically transforming and factorizing higher-order tensors into simpler, constituent components. The process typically involves unfolding the tensor into matrices along each mode, followed by applying the truncated matrix-SVD to these matrices to extract core components and factor matrices. This sequential approach allows for efficient processing and updating of the core and factor matrices iteratively. For instance, in Tucker decomposition via Higher-Order Singular Value

Decomposition (HOSVD), the tensor is projected onto subspaces spanned by singular vectors to form a core tensor and factor matrices. A single mode- n unfolding-then-SVD step, finding the factor $\mathbf{U}^{(n)}$ (line 2 of [algorithm 6](#)), is shown graphically in [Figure 2.29](#) for an order-3 tensor. ([Kolda et al., 2009](#); [Pajarola et al., 2016](#))

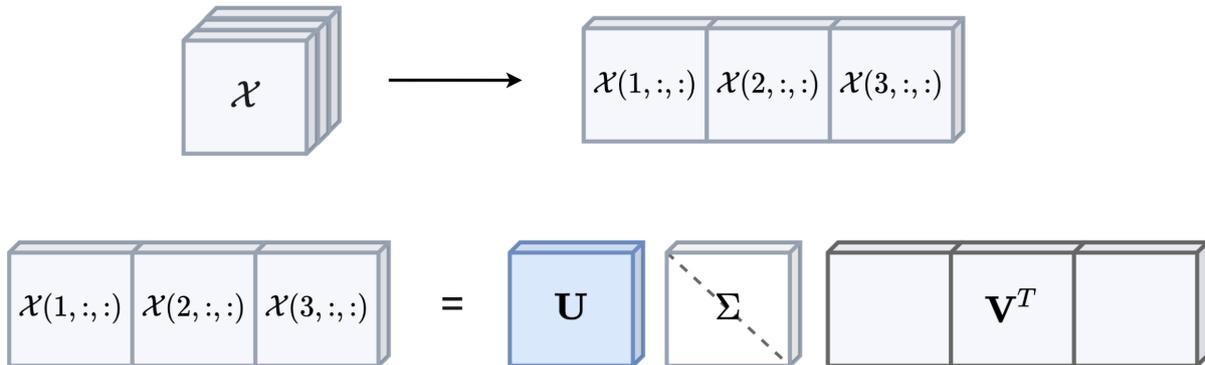


Figure 2.29: Computing the factor matrices $\mathbf{U}^{(n)}$ for the Tucker decomposition of an order-3 tensor using the HOSVD. This involves a mode- n unfolding operation, here mode-1, and the SVD of the matricized tensor. Because the left singular vectors matrix of the mode- n unfolding, $\mathbf{U}^{(n)}$, is small for higher-order tensors the computation of the factors $\mathbf{U}^{(n)}$ is efficient. ([Denarié, 2022](#))

Similarly for the TTD, the tensor is unfolded and decomposed into a series of low-rank cores connected in a chain, each found iteratively through SVD. In contrast to the Tucker decomposition, the matrix-SVD is used on the residuals of the mode- n unfoldings sequentially. This means that instead of considering the full mode- n unfolding to compute factor $\mathbf{U}^{(n+1)}$, only the residual matrix $\Sigma_{(n)} \mathbf{V}_{(n)}^T$ from the previous factor is used. ([Oseledets, 2011](#))

Iterative Computation

A number of more advanced algorithms exist to compute decompositions. Most notably alternating least-squares for the CPD and algorithms building on the HOSVD for the Tucker and TT decomposition ([algorithm 7](#)). Similar to the algorithms outlined above, these algorithms iteratively update the decomposition factors. ([Liu et al., 2022](#))

Gaussian & Non-Gaussian Residuals

In this work, Gaussian residuals are assumed for simplicity in tensor computations. Methods hence minimize the error $(\mathcal{X} - \mathcal{M})^2$ between the observation \mathcal{X} and model \mathcal{M} . However, it is crucial to recognize that HCI data often deviates from Gaussian distributions, necessitating different approaches. The choice of error term and optimization method can substantially influence the decompositions, particularly when factoring tensors with non-Gaussian error distributions. ([Chi et al., 2012](#))

Error Terms & Optimization Methods

Moving beyond traditional algorithms can be advantageous when the data violates the Gaussian error assumption. Various solvers and error terms can then be considered for tensor computations, including:

- Generalized Alternating Least-Squares (ALS) methods ([Hong et al., 2020](#))
- Gradient-based optimization techniques ([Kolda et al., 2020](#))
- Probabilistic approaches using maximum likelihood estimation ([Giampouras et al., 2022](#); [Wei et al., 2019](#))
- Robust estimation methods for handling outliers ([Gu et al., 2014](#); [Wang et al., 2020](#))

Stochastic gradient descent methods ([Kingma et al., 2014](#)) can directly minimize a loss function that incorporates specific assumptions about the noise distribution. This approach allows for more flexible and robust tensor factorizations, adapting to the underlying statistical properties of the data ([Goldfarb et al., 2014](#); [Hong et al., 2020](#); [Kolda et al., 2020](#)).

Optimization of non-Gaussian tensor model can also be accommodated by selecting appropriate tensor norms. For instance, using the L_1 norm instead of the Frobenius norm (L_2) can provide robustness against

outliers. These methods can be particularly effective for large-scale tensor problems, offering improved computational efficiency and scalability compared to traditional approaches. (Daglayan et al., 2023; Wang et al., 2022a)

The choice of solver, error term, and norm should be based on the specific characteristics of the data, the computational requirements of the problem at hand, and the desired properties of the solution. Careful consideration of these factors can lead to more accurate and meaningful tensor decompositions, especially when dealing with complex, real-world datasets that do not conform to simple Gaussian assumptions.

2.4.8. Limitations

Tensor decompositions, such as CPD, Tucker, TTD, and TRD, present an alternative approach to modeling the PSF and interpreting both the PSF and residual in HCI. While these methods build on the capabilities of classical matrix decompositions like SVD, they also introduce specific limitations that impact their application in HCI.

Tensor decompositions inherently deal with higher-order data, which can increase the computational complexity compared to PCA and SVD. For instance, while the SVD uses linear transforms to capture the most significant variance, tensor methods use multiple SVDs across modes. Depending on both the implementation and data this can either reduce or increase the use of resources, as this work demonstrates in Section 5.2.1. Selecting the appropriate decomposition together with suitable parameters, is essential for their use in HCI.

One of the most significant differences between matrix SVD and tensor decompositions arises from their interpretability. PCA and SVD provide a clear and intuitive understanding of data variance and principal components, and are straightforward to interpret as major variance directions in the data, however the unfolded dimensions are not inherently meaningful. In contrast, decompositions like the CPD can significantly ease the the interpretation of factors compared to PCA, see e.g. Figure 2.24 and Figure 8.1 which shows that they are physically meaningful. Nevertheless, advanced tensor decompositions like the Tucker decomposition, TTD and TRD, while powerful in capturing multi-way relationships, yield factors that are not as intuitively understandable. Their core tensors typically represent interactions that are difficult to visualize and interpret. Trading off the expressiveness of the decomposition with ease of interpretability is a key challenge and opportunity for their use in HCI.

Tensor decompositions are generally more sensitive to initial conditions compared to PCA/SVD (Ge et al., 2017; Zhang et al., 2018b). This is because the optimization landscape of tensor decompositions is non-convex and most methods are not guaranteed to converge to global minima. (Chen et al., 2019; Ge et al., 2017) Hence, tensor methods typically require a good initialization for convergence to a meaningful local or global minimum. This is because the full, non-convex problem is solved by separating it into convex sub-problems (Kolda et al., 2009). This can partially be overcome by performing several decompositions using the randomized SVD with differing initial random states. This is less of an issue with PCA, where SVD provides deterministic and globally optimal results. (Tomioka et al., 2011)

From a practical point of view it should be noted that the mathematical foundations of tensor decompositions are less mature than those of matrix decompositions. The rank determination problem in tensors is NP-hard, and there are often no efficient algorithms for decompositions comparable to those available for PCA/SVD. This can lead to practical challenges in determining the optimal number of components to retain, which in PCA/SVD is guided by more straightforward criteria like the explained variance. (Kolda et al., 2009)

In conclusion, while tensor decompositions are powerful tools for modeling and data interpretation in HCI, their complexity and the challenges inherent in their application require careful consideration to fully leverage their advantages for HCI.

2.5. Example: Failure of Matrix-PCA on Higher-Order Tensors

To illustrate the possibility of matrix-PCA failing catastrophically when used on higher-order data, consider an order-3 tensor $\mathcal{X} \in \mathbb{R}^{I \times J \times K}$. The tensor is constructed using sine functions:

$$\mathcal{X} = \mathbf{v}^{(I)} \circ \mathbf{v}^{(J)} \circ \mathbf{v}^{(K)}, \quad (2.30)$$

where the vectors $\mathbf{v}^{(I)}$, $\mathbf{v}^{(J)}$, and $\mathbf{v}^{(K)}$ are constructed using $\mathbf{n} = \{r : r = \frac{1}{N} \times n, n \in \{0, 1, \dots, N\}\}$ with $N = I = J = K = 25$:

$$\mathbf{v}^{(I)} = 0.25 \sin(2\pi\mathbf{n}) \quad (2.31a)$$

$$\mathbf{v}^{(J)} = 0.25 \sin\left(2\pi\mathbf{n} - \frac{\pi}{2}\right) \quad (2.31b)$$

$$\mathbf{v}^{(K)} = -0.25 \sin\left(2\pi\mathbf{n} + \frac{\pi}{2}\right). \quad (2.31c)$$

The elements of $\mathbf{v}^{(I)}$, $\mathbf{v}^{(J)}$, and $\mathbf{v}^{(K)}$ are shown in Figure 2.30 as circles. While this is an example tensor specifically chosen to highlight the problems with matrix-PCA on higher-order tensors, it is easily analogous to an observation where one of the modes is spectral or temporal, and the two other modes are spatial and contain large speckles evolving slowly over time or the spectrum.

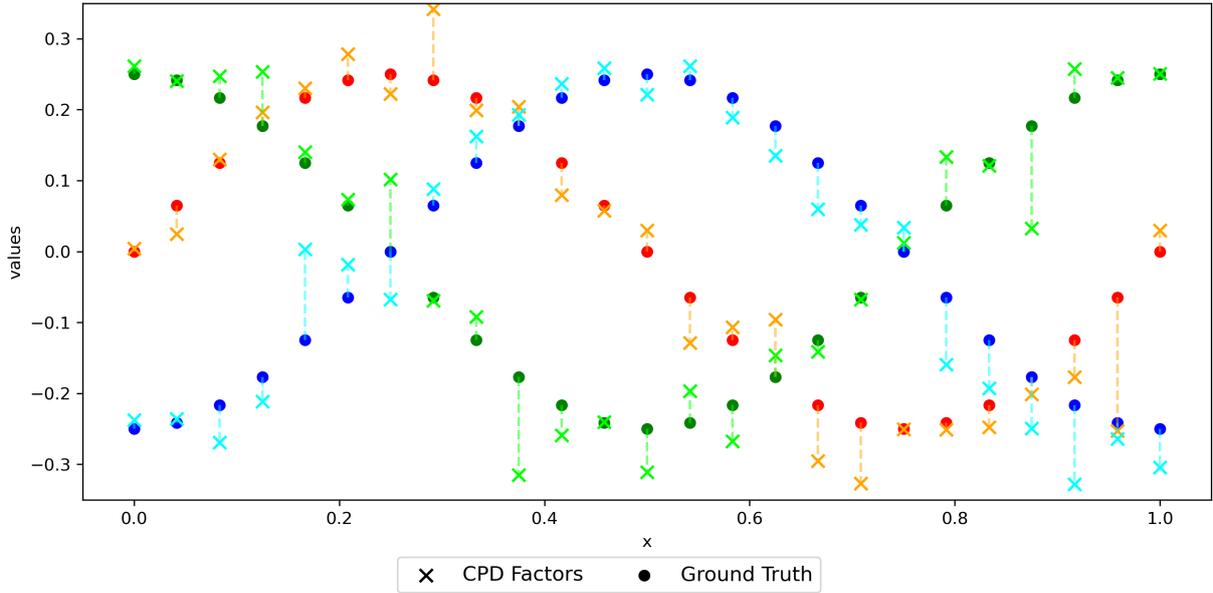


Figure 2.30: Basis vectors $\mathbf{v}^{(I)}$ (red), $\mathbf{v}^{(J)}$ (green) and $\mathbf{v}^{(K)}$ (blue) as well as the corresponding factors of the CPD (orange, light-green, cyan).

\mathcal{X} is the underlying data generated by some process; however, the actual observation of this process is heavily corrupted by normally distributed noise. To highlight the difference between matrix-PCA and tensor methods, the noise is taken as $\mathcal{N}(\mu = 0, \sigma = 5.0 \times 10^{-2}) \in \mathbb{R}^{I \times J \times K}$. This noise level is larger than for pre-processed ASDI observations; however, when subtracting the median spectral PSF cube, as is common for post-processing (line 6 of algorithm 1), the chosen noise level is more appropriate. Nevertheless, the noise in ASDI observations is typically not Gaussian (Bonse et al., 2023; Mawet et al., 2014). The observed data is then

$$\mathcal{Y} = \mathcal{X} + \mathcal{N}. \quad (2.32)$$

The aim is to build a model $\hat{\mathcal{Y}}$ of the observation \mathcal{Y} to recover the underlying data so that $\|\mathcal{X} - \hat{\mathcal{Y}}\|$ is minimal. Discovering the vectors $\mathbf{v}^{(I)}$, $\mathbf{v}^{(J)}$, and $\mathbf{v}^{(K)}$ that generated this tensor might also be desirable to compare the data-driven model $\hat{\mathcal{Y}}$ with theoretic models of the data generation process.

Figure 2.31 shows the 6th lateral slice ($I \times K$) of the ground truth \mathcal{X} (left), observed data \mathcal{Y} (center-left), and models $\hat{\mathcal{Y}}$. Clearly, the noise level is very high when considering just this slice, however, the slices in the observation are highly correlated with this slice so that recovery should be possible.

The center and center-right frame of Figure 2.31 shows the matrix-PCA model using 1 and 3 principal components and demonstrates that matrix-PCA cannot recover the underlying data. While the 1 component model discovers some structure in the data, the unfolding operation mixes the modes to such a degree that the model in one mode is modulated by the other mode (center-frame). The 3 component model

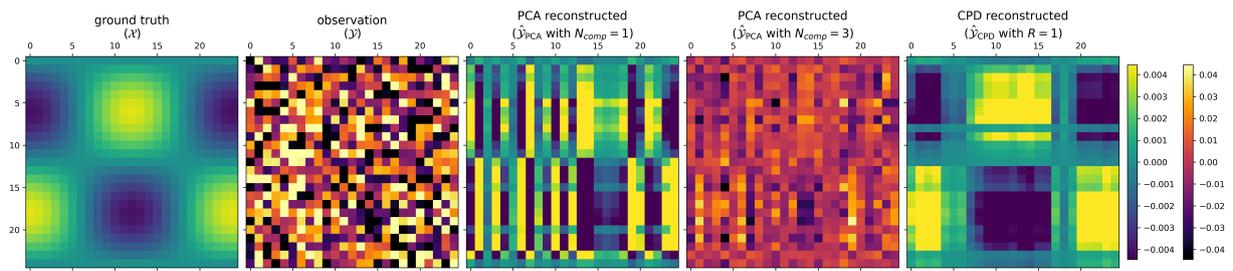


Figure 2.31: 6^{th} lateral slice ($I \times K$) of the ground truth \mathcal{X} , observed data \mathcal{Y} and models $\hat{\mathcal{Y}}$ by PCA with 1 PC, 3 PC and CPD.

(2^{nd} frame from the right, 12% completeness) already starts increasingly modeling noise, indicated by the absolute value of the pixels. While Figure 2.31 shows only one specific way to perform matrix-PCA on this data, it can never recover the underlying higher-order data structure, no matter how \mathcal{Y} is unfolded. Again, in this example, \mathcal{Y} is corrupted by Gaussian noise and standardized before applying matrix-PCA, so that this result is a best-case scenario which does not reflect actual observations obtained by ASDI.

In contrast, the rank-1 tensor model recovers a reasonable approximation of the low-rank data (rightmost frame). Furthermore, the factors of this model, shown in Figure 2.30 (crosses), are a close match to the ground truth, so that the data generation processes can plausibly be guessed at.

While this example is not a direct analogue to the observations obtained by ASDI, partly because the data generation process in direct imaging is non-linear, it demonstrates the fundamental problem that is encountered when applying matrix methods to higher-order tensors.

2.6. Chapter Summary & Contributions

This chapter provides essential background for understanding the application of tensor decomposition methods to post-processing direct exoplanet observations.

The first section introduces exoplanets science. It emphasizes that direct imaging primarily detects young, self-luminous gas giants at wide separations, typically ≥ 2 Jupiter masses and 10 – 300 AU from their host stars. This population is distinct from those found by other detection methods, highlighting the unique contribution of direct imaging to exoplanet science.

The following section explains the fundamental problem of direct imaging: separating the faint planetary emission from the much brighter stellar emission. It details the key components of an integrated direct imaging pipeline: large ground-based telescopes, adaptive optics, coronagraphs, observational strategies, and advanced post-processing techniques. The necessity and impact of each component on the physical signal from the exo-solar system is detailed. The section specifically introduces the ASDI observational technique, which produces four-dimensional data cubes (wavelength, time, and two spatial dimensions) that are operated on by the tensor methods developed in this thesis.

The next section establishes the mathematical framework of multi-linear algebra for working with higher-order tensors. It introduces key concepts such as tensor order, rank, and various operations on tensors. Furthermore, the section describes the two primary ways of describing multi-linear algebra; multi-linear notation and tensor network notation. This section highlights how higher-order tensors like ASDI data fundamentally differ from matrices.

The subsequent section describes tensor decomposition methods (CPD, Tucker, TTD, TRD) for factorizing higher-order data. It explains how these methods generalize matrix factorizations to higher-order data, preserving multi-modal relationships. This section provides the theoretical basis for the novel post-processing methods developed in the thesis.

The chapter concludes with an example demonstrating how matrix-based PCA fails when applied to higher-order tensor data, while a simple tensor method succeeds. This practical illustration motivates the need for tensor-based approaches in processing ASDI observations.

Contributions 2.1: Background

This chapter makes several key contributions to the broader goal of the thesis:

- It provides essential context on exoplanets and direct imaging techniques, explaining the importance and challenges of this field, and motivating the need for advanced post-processing methods.
- It explains angular-spectral differential imaging and the resulting data structure, crucial for understanding why tensor methods are particularly well-suited for this application.
- It introduces multi-linear algebra and tensor decompositions, bridging the gap between traditional matrix-based approaches and higher-order tensor methods. The background on different tensor decompositions (CPD, Tucker, TTD, TRD) lays the groundwork for understanding the specific methods proposed and evaluated in the thesis.
- It describes the limitations of matrix-based PCA when applied to higher-order data, providing motivation for exploring tensor-based alternatives. A practical example demonstrates the failure of matrix PCA on a synthetic higher-order tensor, illustrating in concrete terms why a tensor approach is needed.

Overall, this chapter establishes the interdisciplinary nature of the research, connecting astrophysics, astronomical imaging and multi-linear techniques, and provides the necessary context for readers to engage with the novel contributions presented in the rest of the thesis.

Having established this foundation, the next chapter will describe tensor-based methods for post-processing direct imaging observations.

Multi-Modal Point Spread Function Modeling with Tensor Decompositions

This chapter introduces a novel class of PSF subtraction methods, termed *Tensor Image DEcompositions* (TIDE), which leverage tensor decompositions to model the coronagraphic PSF in observations obtained by differential imaging of exoplanets. Angular-spectral differential imaging observations are structured as an order-4 tensor with a spectral, temporal and two spatial modes $\mathcal{X}_{\lambda,\theta,x,y}$. Unlike traditional matrix-PCA-based methods, tensor decompositions preserve a tensors higher-order and multi-modal structure. A tensor decomposition approach thus prevents the loss of information associated with flattening higher-order data into a matrix. Moreover, the components of the decompositions are physically meaningful as they describe and correlate spectral, temporal and spatial features of the PSF and speckle pattern. Furthermore, tensor methods can overcome some of the computational challenges of post-processing large observations.

TIDE methods build on the general procedural framework established for CODI PSF subtraction techniques, which are detailed in this chapter (algorithm 1, Figure 2.15). PSF subtraction can also be represented using their computational graph. This graph, shown in Figure 3.1, shows the sequence of post-processing operations and the parameter dependence of each step. The effectiveness and reliability of TIDE hinge on the use of mature, robust implementations of underlying multi-linear algebra algorithms, such as in TensorNetwork (Roberts et al., 2019), TensorLy (Kossaifi et al., 2019), TNTorch (Usvyatsov et al., 2022) and PyTorch (Paszke et al., 2019). These tools provide the computational foundation necessary for implementing the tensor operations required by TIDE, ensuring that the methods can be broadly applied within the direct imaging community. Details on these tools and their integration into the TIDE framework can be found in Appendix G and Table G.1.

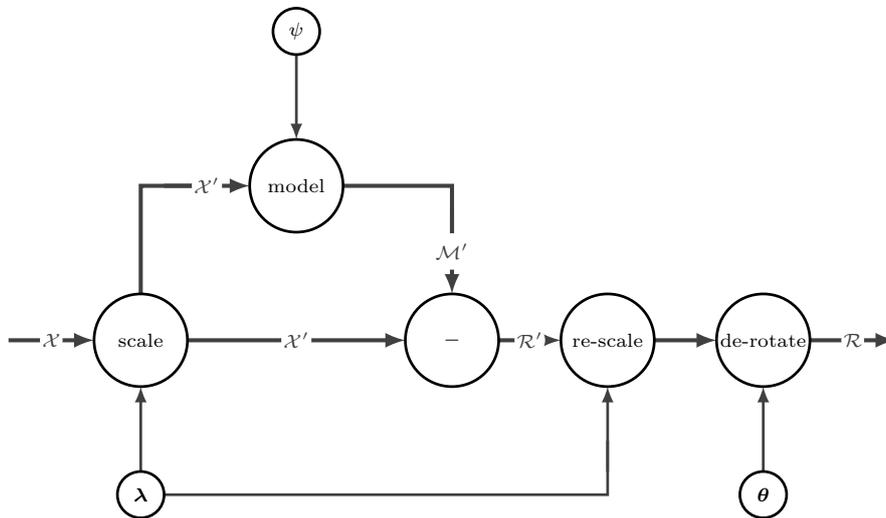
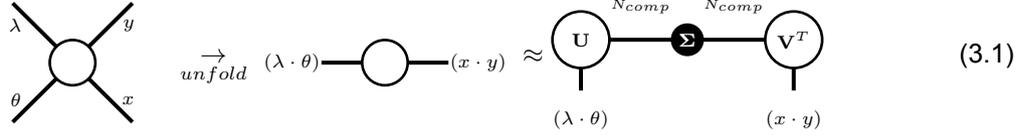


Figure 3.1: Computational graph of CODI. Each large vertex is an operation (function), the small vertices are parameter sets and each labeled edge is a tensor. In this graph the prime symbols indicate that the tensor is in the scaled and rotated frame, whereas tensors without prime symbols are in the de-rotated frame. \mathcal{X} is the observation tensor, \mathcal{M} the PSF model, \mathcal{R} the residual, and ψ are PSF model parameters.

By advancing beyond matrix-PCA, this work establishes a framework for modeling the PSF and speckle pattern as higher-order tensors, addressing both theoretical and practical challenges in post-processing of direct imaging observations.

3.1. CODI with matrix-PCA

Today, the prevailing reduction technique for ASDI observations involves matrix-PCA, using the matrix truncated-SVD to model the PSF, as demonstrated in Equation 3.1. This approach is generally optimal in a least-squares sense for modelling the coronagraphic PSF from matrix-shaped direct imaging data. (Amara et al., 2012; Lu et al., 2008; Soummer et al., 2012) Since the observations obtained by ASDI are not in matrix shape, they need to be matricized. Specifically, the unfolding in Equation 3.1 is generally employed. This is known in multi-linear algebra as the 2-unfolding matrix of the observation tensor.



$$\begin{array}{c} \lambda \\ \theta \end{array} \begin{array}{c} y \\ x \end{array} \xrightarrow{\text{unfold}} (\lambda \cdot \theta) \text{---} (x \cdot y) \approx \begin{array}{c} \text{U} \\ (\lambda \cdot \theta) \end{array} \text{---} \begin{array}{c} \Sigma \\ N_{comp} \end{array} \text{---} \begin{array}{c} \text{V}^T \\ (x \cdot y) \end{array} \quad (3.1)$$

In this work PCA is implemented using the truncated matrix-SVD which was found to be more robust at the cost of speed when compared with randomized and low-rank matrix factorization methods (e.g. algorithm 5.1 in Halko et al., 2011, implemented by Paszke et al., 2019), especially for low-SNR observations. To facilitate comparison, this work also uses the t-SVD over the QR decomposition or randomized SVD when computing tensor decompositions unless noted otherwise, e.g. for the power iteration method or randomised partial CPD methods.

3.2. CODI with Tensor Decompositions

This work proposes to model the quasi static components using a tensor decomposition which solves the problem in Equation 3.3, where it is assumed that the observations \mathcal{X} can be expressed using Equation 3.2.

$$\mathcal{X} = \mathcal{M} + \mathcal{E}, \quad (3.2)$$

where \mathcal{M} is a low-rank PSF model analogously to Equation 2.4, and the tensor \mathcal{E} captures both nuisance components but also off-axis sources.

$$\min_{\mathcal{M}} \|\mathcal{X} - \mathcal{M}\| \text{ s.t. } \widetilde{\text{rank}}_i(\mathcal{M}) = R_i, \quad (3.3)$$

where R is the mode- i border rank of \mathcal{M} , with the tensor L2 norm defined as:

$$\|\mathcal{A}\| = \sqrt{\sum_{i_1=1}^{I_1} \sum_{i_2=1}^{I_2} \cdots \sum_{i_N=1}^{I_N} a_{i_1 i_2 \dots i_N}^2} \text{ for } \mathcal{A} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N} \quad (3.4)$$

Tensor decompositions can serve as a direct substitute for PCA-based methods in modelling the PSF. Moreover, they offer several advantages over traditional PCA-based methods. First, tensor-based methods can capture the multi-modal structure of the data, allowing for a more direct representation of spatial, temporal and spectral features in the PSF. Secondly, the flexibility in choosing the multi-linear rank of decomposition allows tuning the completeness of the model for each mode. Moreover, having physically meaningful factors for the PSF model enables a closer connection between PSF subtraction methods and physics informed methods or causal noise models in HCI. Lastly, the ongoing development in tensor decomposition algorithms and their increasing integration into machine learning frameworks suggests a growing potential for these methods to become standard tools in areas requiring high-dimensional data analysis (Cichocki, 2014; Panagakis et al., 2021; Sidiropoulos et al., 2017). Thus, tensor decompositions can be instrumental in overcoming the significant challenges that deep learning on HCI data is currently facing (Cantero et al., 2023; Flasseur et al., 2024; Wolf et al., 2023).

The data-driven nature of the models explored in this work needs to be stressed again. Modeling the PSF and speckle pattern to the required accuracy for detecting off-axis sources is greatly simplified by this approach. To demonstrate this, consider a perfectly stable optical system where the PSF and speckles

only vary in brightness due to the airmass over the telescope, changing non-trivially over time. To model this temporal variance a data-driven model would only require a single factor to exactly describe the relative brightness at each time. A model constrained by a specific set of basis functions, however, might require many factors. Experimentally, this work finds that speckle pattern modeling with tensor decompositions is not tractable using the discrete cosine transform (Yeo et al., 1995), nor Legendre, Chebyshev, or Hermite polynomials as basis functions (Usvyatsov et al., 2022). Fourier (Chiueh et al., 1997) and wavelet transform (Rodler, 1999) bases are also possible, and have been previously applied in direct imaging (Bonse et al., 2018), but were not explored in this work. Nevertheless, physically informed constraints are poised to be a promising direction for future work, see Chapter 8.

In this work, three groups of algorithms are explored; tensor decomposition PSF models in the original Cartesian reference frame (Section 3.2.1, full-frame), tensor decomposition in the polar reference frame (Section 3.2.4, single annulus), as well as in concentric annuli transformed to the polar reference frame (Section 3.2.5, annular). Furthermore, there is a trade-off to consider when choosing whether to keep the modes physically meaningful, or to partially unfold or to further tensorize them, which is discussed in Section 3.2.2, Section 3.2.3, and Section 3.2.6.

3.2.1. Full-Frame PSF Model

This section describes full-frame tensor decomposition methods for modeling the coronagraphic PSF and speckle pattern in ASDI observations. The full-frame approach applies tensor decompositions to the entire observational data cube, treating it as a single, coherent data structure. This allows for global modeling of the PSF and speckle pattern across all modes simultaneously. The four main tensor decomposition techniques are discussed in this context: CPD, Tucker Decomposition, TTD, and TRD. Each of these methods offers distinct advantages and trade-offs in terms of expressiveness, computational efficiency, and interpretability.

These full-frame methods aim to provide more accurate and physically meaningful representations of the PSF compared to matrix-based techniques, with the potential of improving contrast and sensitivity in exoplanet detection. The following subsections detail each decomposition method, discussing their formulation, strengths, and limitations in the context of high-contrast imaging data analysis.

Canonical-Polyadic PSF Model

The CPD is somewhat similar to PCA, in that instead of selecting the number of components N_{comp} which are kept to model the quasi-static PSF, the number of rank-1 tensors, R , is selected. Under the constraint of the rank the observation is approximated with the CPD up to an error threshold, see Equation 3.5. This preserves only the low-rank components of the observation, which are the quasi-static components of the coronagraphic PSF.

Figure 3.2 shows the factors of a rank-5 CPD PSF model. Each row are the factors in a specific mode and each row forms a rank-1 tensor of the total CPD PSF model. The last column shows the weight λ of each factor, and the frontal slice of the full tensor PSF model is shown in Figure 3.3a. The factors of this PSF model can be directly interpreted. For example, the factors in the first column are factors along the spectral mode of the observation and, together, describe components of the PSF spectrum. In this case the factors are very similar to each other due to the strong influence of earth's atmosphere. The factors in the temporal mode differ much more strongly, indicating that the PSF shape and brightness of the target varied over the total observation time. All temporal factors share a common dip at the same observation indices (roughly half-way), which shows that the PSF was less bright at these times. Note however, that each factor also has finer structure.

By increasing the rank, components which account for less static components are included in the tensor model. Since the CPD has only a single parameter which needs to be tuned, its use in HCI corresponds strongly with the current use of PCA. Thus, the CPD provides a simple yet effective way to model quasi-static components of the PSF in HCI.

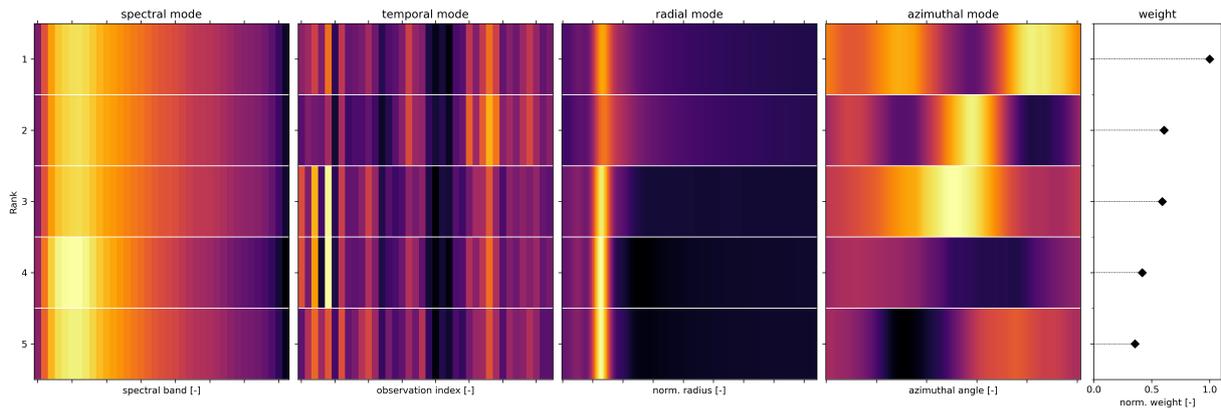


Figure 3.2: Normalized factors of a rank-5 CPD PSF model. The weights λ are the tensor L2 norm (Equation 3.4) of each rank-1 tensor, scaled by the largest weight. Each column shares a common color-scale.

$$(3.5)$$

However, the major limitation of CPD is its restricted expressive power, as a single parameter adjustment across all modes may not optimally represent the data. This limitation is apparent in ASDI observations where the spectral and angular modes exhibit distinct variances. For instance, SPHERE-IFS data typically shows less variance along the spectral mode compared to the temporal mode, suggesting that a lower rank could suffice for the spectral mode while a higher rank might be needed for the temporal mode.

Tucker PSF Model

The Tucker decomposition, Equation 3.6, addresses this issue by introducing a full interaction core alongside a multi-linear rank, significantly enhancing the model’s expressive power. The multi-linear rank allows for fine-tuning of the model’s complexity in each mode, capturing essential features while selectively filtering out noise and flux from off-axis sources. Nevertheless, this has the notable drawback that the Tucker decomposition includes a core that exponentially increases in parameter count with the order of the tensor, which can lead to substantial storage requirements for PSF modeling.

$$(3.6)$$

When using PCA or the CDP, including more variance in the model will increase the fidelity of the model with respect both the spectral and temporal modes simultaneously. PCA finds directions of maximum

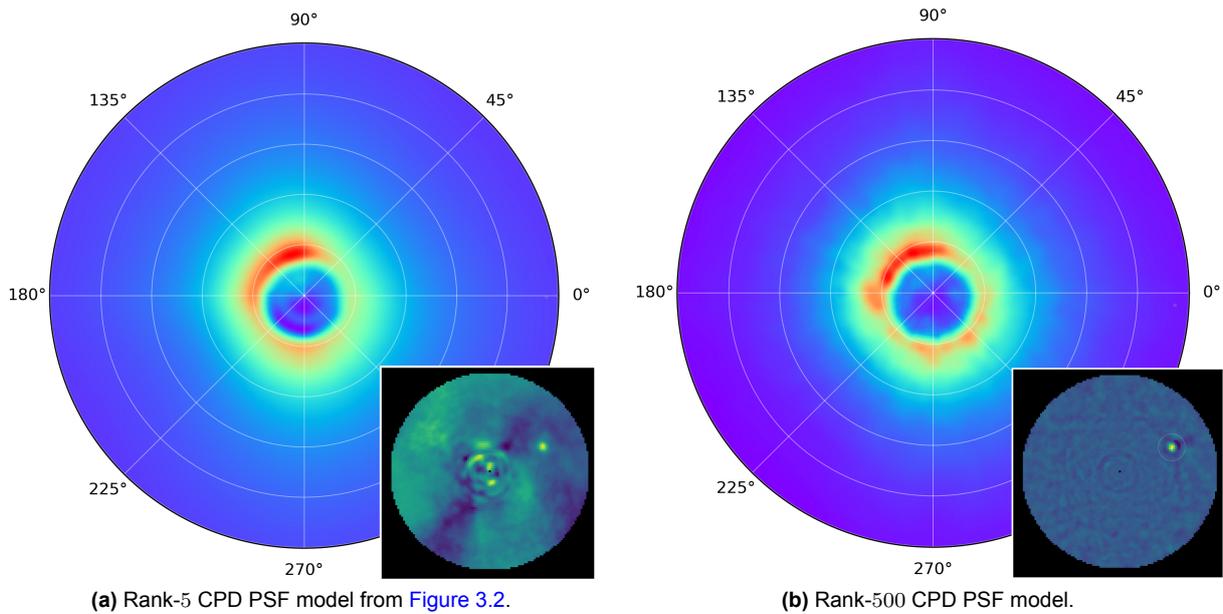


Figure 3.3: Frontal slices (λ_0, θ_0) of CPD PSF models of GPI observations. The rank-5 model captures most of the reference PSF, however it does not capture the speckle pattern. In contrast, the rank-500 model captures both the reference PSF and the speckle field. The residual images are inset in the lower right corners.

variance space spanned by the flattened observations without the ability to differentiate which mode the variation is coming from. In contrast, the CPD essentially assumes that each component shares the same level of significance across all modes. This is particularly effective when the elements of different modes contribute equally to the data structure or when the interaction between modes is relatively uniform. However, in cases where different modes exhibit different levels of complexity or information content, the CPD may not capture these nuances as effectively as the Tucker decomposition. Finally, the Tucker decomposition both preserves the structure of the data and the multi-linear rank enables increasing fidelity in one specific mode without necessarily increasing it in another mode. This is particularly useful where the variability or information content differs significantly across modes as can be the case in ASDI. Consider e.g. that observations with a quickly and slowly varying parallactic angle can have the same spectral resolution and number of channels.

The performance of tensor decompositions on observations in Cartesian coordinates is predominantly limited by their ability to cope with strong variance over the mode domains (Liu et al., 2022). The core in the Tucker decomposition needs to account for interactions in all dimensions simultaneously, which can be challenging when the flux varies strongly with separation. This is where the TTD (Equation 3.7) and TRD (Equation 3.8) offer distinct advantages by restricting the cores to order- ≤ 3 .

Tensor-Train PSF Model

The TTD is particularly useful for high-order tensors, such as order-4 ASDI tensors, since it decomposes the tensor into a sequence of smaller cores, each linking only a pair of modes. This decomposition structure makes TTD less demanding in terms of parameter count compared to the Tucker decomposition, thereby reducing storage requirements. However, a potential drawback of TTD is that the leading and trailing factors do not share an interaction core, which can limit the model's ability to capture global interactions between the first and last modes in the data sequence.

$$\begin{array}{c} \lambda \\ \theta \end{array} \begin{array}{c} \diagup \\ \circ \\ \diagdown \end{array} \begin{array}{c} y \\ x \end{array} \approx \begin{array}{c} \mathcal{G}^{(\lambda)} \\ \lambda \end{array} - \begin{array}{c} \mathcal{G}^{(\theta)} \\ \theta \end{array} - \begin{array}{c} \mathcal{G}^{(x)} \\ x \end{array} - \begin{array}{c} \mathcal{G}^{(y)} \\ y \end{array} \quad (3.7)$$

Tensor Ring PSF Model

The TRD addresses this limitation by connecting the ends of the tensor train, thereby forming a circular tensor trace that enables each mode to interact more flexibly with its neighbors, including the first and last modes. This circular configuration enhances the ability to model cyclic cross-couplings, which can be particularly beneficial in HCI cross-couplings. For example, they could be significant between the spectral and spatial modes. Despite these advantages, in practice, the performance differences between TTD and TRD are typically minimal, suggesting that the choice between these two methods should be guided by specific requirements of the application and computational resources available.

While TRD offers an improvement over TTD in terms of modeling cyclic interactions, both decompositions provide significant benefits over CPD and the Tucker decomposition when handling complex multi-modal data in HCI. They allow for more granular control over the interactions between different data modes without becoming untenable to compute due to memory requirements. In practice, this is a critical factor in accurately modeling the PSF, improving the residual noise characteristics, and increasing the achievable contrast in the residual cubes.

3.2.2. Unfolding

Tensor decomposition can be applied directly to the observation tensor in its original form (order-4 tensor). However, the tensor can also be partially unfolded the tensor along specific modes before computing the decompositions (e.g. TTD in Equation 3.9). The choice between these approaches depends on the characteristics of the data and the specific goals of the analysis. By partially unfolding the tensor along the spatial modes, the observation is no longer spatially coherent and each pixel is modeled separately. This method is particularly useful when the spatial structure of the data is not inherently low-rank, such as in observations where a central region near the inner working angle (IWA) is neither low-rank nor informative for the rest of the frame. By unfolding, the tensor decomposition can then model each pixel's angular-spectral slice separately, allowing it to more easily capture non-low-rank spatial components. The impact of this is especially strong for datasets where the cross-couplings vary significantly across the domains of the spatial modes. Consider here e.g. a fiber across one of the spatial modes of an observation tensor which contains parts of the region inside the IWA. The fiber would vary strongly between the outer regions (very low flux), the close-in regions (very high flux), and the region masked by the coronagraph (both high and low flux). However, unfolding increases computational and storage demands, as it increases the size of the spatial modes from $I_x + I_y$ to $I_x \cdot I_y$ during the decomposition process.

3.2.3. Tensorization

Conversely, tensorization increases the dimensionality of the original observations to divide the domains into smaller, subdomains. This process implicitly reduces variance across each dimension and allows

aligns radially streaking speckles and structure like the wind-driven halo, shown in [Figure 2.6](#). There is however, also a case to be made that polar transformed model more easily incorporates exoplanet signal as its position varies along the radial and azimuthal mode. The impact of the trade-off between lower-rank models and self-subtraction is discussed in [Section 5.2](#).

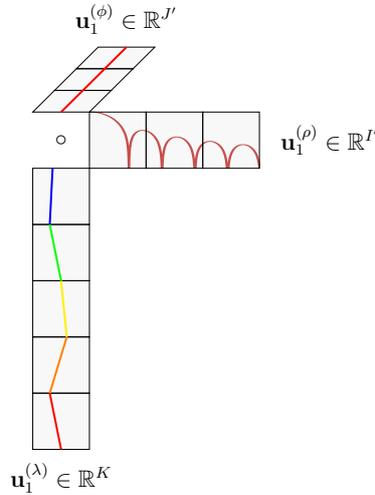


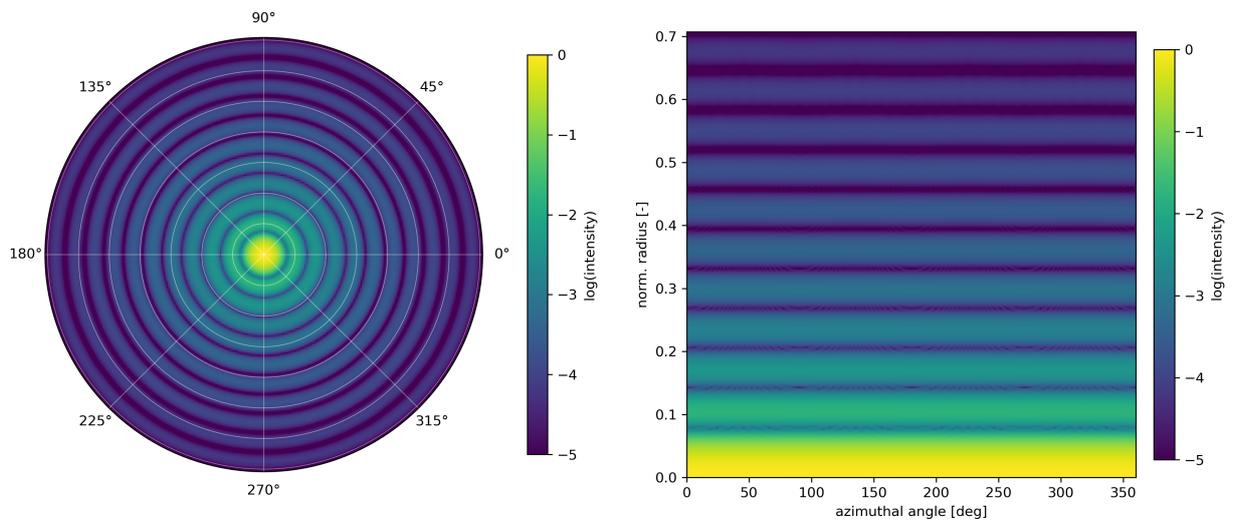
Figure 3.5: Factors of a rank-1 CPD an Airy pattern (like the one in [Figure 2.7](#)). The factors are drawn for a spectral observation cube with varying brightness in each spectral band. To better illustrate the point the factor in the radial mode ($\mathbf{u}_1^{(\rho)}$) is plotted in log-scale. The factor in the azimuthal mode ($\mathbf{u}_1^{(\phi)}$) is constant.

Overall, a transform to polar coordinates simplifies the capturing the low-rank structure of the PSF (or equivalently, it reduces the rank of the model), while not negatively impacting the capture of the speckle pattern. Further, it offers the advantage of masking the central and outer regions implicitly by only considering radial coordinates between $0 \leq \rho_0$ and $\rho_1 \leq R$. Note, that this converts the PSF modeling problem into an annular PSF modelling problem as proposed by [Absil et al., 2013](#). To archive this, the scaled observations are transformed to polar coordinates, the PSF is modeled and subtracted from the transformed observations, see [Equation 3.10](#), and the residuals are transformed back to Cartesian coordinates, rescaled and derotated.

The equation shows a 2D coordinate system on the left with axes λ and ϕ . This is followed by an approximation symbol \approx and a sequence of four operations: $\mathcal{G}^{(\lambda)}$, $\mathcal{G}^{(\theta)}$, $\mathcal{G}^{(\rho)}$, and $\mathcal{G}^{(\phi)}$. Each operation is represented by a circle with its corresponding coordinate below it.

where ρ is the radial and ϕ the azimuthal mode. The polar transform and its inverse are implemented as row- and column-wise discrete coordinate flows, see [algorithm 2](#) and [algorithm 3](#). Since the in- and output of the transforms are in discrete domains, bilinear interpolation is used to transform the observations. Typically, it is desirable to increase the resolution of the observations in polar coordinates to not lose information in the transform, however experimentally, the up-scaling factors for HCI observations did not impact the images significantly, as long as the radial dimension is sampled 1 : 1 and the azimuthal dimension is up-sampled at 1 : 2π so that frame with shape (100 px, 100 px) in Cartesian coordinates will have shape (71 px, 315 px) in polar coordinates.

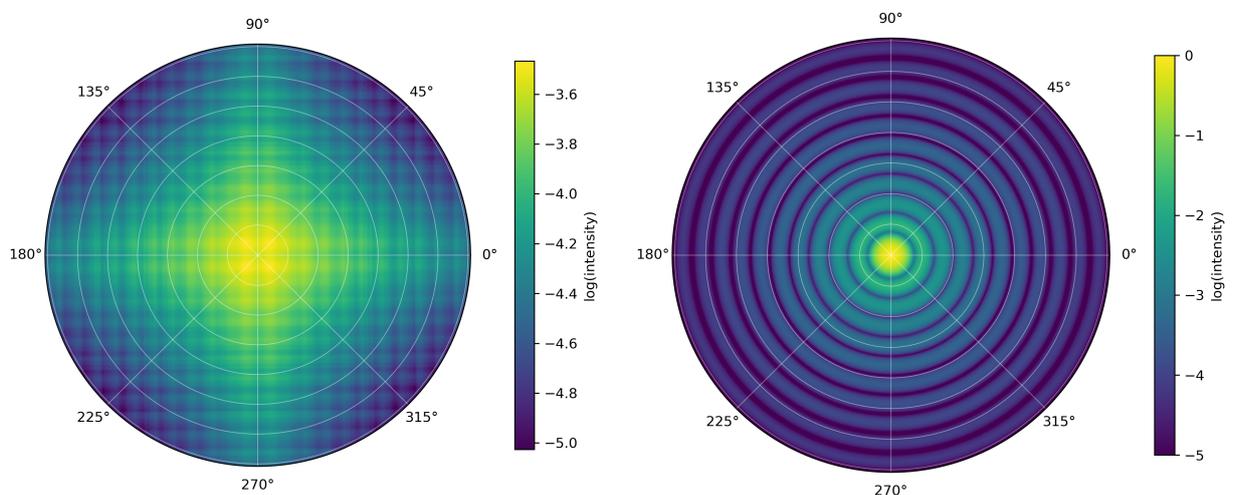
Still, both tensor decompositions in Cartesian and polar coordinates still suffer from the issue that a large fraction of the voxels is concentrated at larger radial separations. This is somewhat counteracted by much of the PSF flux being at smaller separations. Log-polar transforms could thus offer an advantage in terms of approximation performance and storage requirements.



(a) Airy pattern in Cartesian coordinates with overlaid polar coordinate frame.

(b) The same Airy pattern as in Figure 3.6a aligned with the polar coordinate frame. The image is obtained using the polar transform in algorithm 2 without scaling. Only flux within $1/\sqrt{2}$ of the maximum radius in the Cartesian frame is considered here to better illustrate the pattern.

Figure 3.6: Monochromatic diffraction pattern at the focal plane of a telescope with a circular pupil and no secondary obstruction. This is the Airy pattern shown in Figure 2.7.



(a) Reconstructed rank-1 CPD model from the image in Cartesian coordinates (Figure 3.6a).

(b) Reconstructed rank-1 CPD model from the image in polar coordinates (Figure 3.6b), transformed back to Cartesian coordinates using algorithm 3.

Figure 3.7: Reconstructed images rank-1 CPD Airy pattern models in Cartesian and polar coordinates. Note the differing color scales. The polar rank-1 CPD model is almost exact because the Airy model can be separated into two linear functions in the radial and azimuthal mode in polar coordinates. In Cartesian coordinates the Airy pattern can be asymptotically approximated into two multi-linear functions.

Algorithm 2: Cartesian to polar transform for tensors. Since each pixel transforms independently lines 9 to 18 can be fully vectorized. The flows F and G are the same for each slice. The flows can be reused for transforms of tensors in the same $x - y$ domain as \mathcal{X} , so that repeated transforms are very fast.

```

Input: tensor in cartesian coordinates (trailing dimensions)  $\mathcal{X} \in \mathbb{R}^{I_1 \times \dots \times I_{N-2} \times I_x \times I_y}$ ,
         radius  $R$ , s.t.  $R \geq I_x/2$ 
Output: tensor in polar coordinates (trailing dimensions)  $\mathcal{X}' \in \mathbb{R}^{I_1 \times \dots \times I_{N-2} \times I_\rho \times I_\phi}$ 
1 Initialize:
2  $\rho \leftarrow [0 \ 1 \ \dots \ R]^T$  // define radial resolution elements
3  $\phi \leftarrow [0 \ 1 \ \dots \ 2\pi R]^T$  // define azimuthal resolution elements
4  $\mathbf{F}, \mathbf{G} : (\rho, \phi) \mapsto (x, y)$  // discrete (pixel) flow from cartesian to polar coordinates
5  $\mathbf{P} \leftarrow \rho \circ \phi$  // create radial map
6  $\Phi \leftarrow \phi \circ \rho$  // create azimuthal map
7  $\mathbf{F}^{I_{N-1}} \leftarrow \mathbf{P} * \cos(\Phi R^{-1})$  // flow from cartesian to polar coordinates, column-wise
8  $\mathbf{G}^{I_N} \leftarrow \mathbf{P} * \sin(\Phi R^{-1})$  // flow from cartesian to polar coordinates, row-wise

/* remap each element in the cartesian tensor to its new position in the polar
   tensor using bilinear interpolation */
9 for  $i_1 = 1$  to  $I_1$  do
10 | ...
11 | for  $i_{N-2} = 1$  to  $I_{N-2}$  do
12 | | for  $j = 1$  to  $I_x$  do
13 | | | for  $k = 1$  to  $I_y$  do
14 | | | | /* remap each slice, in normalized coordinates:
15 | | | | |  $\text{dst}(x, y) = \text{src}(\mathbf{F}(x, y), \mathbf{G}(x, y))$ ,  $x = j/I_x$  and  $y = k/I_y$  */
16 | | | |  $\mathcal{X}'_{i_1, \dots, i_{N-2}, j, k} \leftarrow \mathcal{X}_{i_1, \dots, i_{N-2}, \mathbf{F}_{j,k}, \mathbf{G}_{j,k}}$ 
17 | | | end
18 | | end
19 end
20 return  $\mathcal{X}'$ 

```

3.2.5. Annular PSF Model

While the transform to polar coordinates partially alleviates the problem of the strong variance over mode domains, it also requires more memory to preserve all information in the observation. Hence, separating the radial domain into smaller subdomains aids with reducing memory requirements and further reduces the variance as much of the variance in the observation cubes is along the radial mode. Equation 3.11 gives the example of the annular TTD for CODI.

$$\begin{array}{c}
 \lambda \\
 \diagdown \\
 \circ \\
 \diagup \\
 \phi \\
 \theta \\
 \diagdown \\
 \circ \\
 \diagup \\
 \rho \supseteq \rho
 \end{array}
 \approx
 \begin{array}{c}
 \textcircled{G^{(\lambda)}} \\
 | \\
 \lambda
 \end{array}
 \text{---}
 \begin{array}{c}
 \textcircled{G^{(\theta)}} \\
 | \\
 \theta
 \end{array}
 \text{---}
 \begin{array}{c}
 \textcircled{G^{(\rho)}} \\
 | \\
 \rho
 \end{array}
 \text{---}
 \begin{array}{c}
 \textcircled{G^{(\phi)}} \\
 | \\
 \phi
 \end{array},
 \quad (3.11)$$

where $\rho \in \rho$ is the radial mode in the specific radial sub-domain of the annulus.

Each annulus can be independently transformed using algorithm 2 and algorithm 3, applied to sub-frames of observations defined by the annulus's outer radius. This method reduces the memory required for both the transformation and the PSF modeling.

Algorithm 3: Polar to cartesian transform for tensors. Since each pixel transforms independently lines 8 to 17 can be fully vectorized. The flows F and G are the same for each slice. The flows can be reused for transforms of tensors in the same $\rho - \phi$ domain as \mathcal{X} , so that repeated transforms are very fast.

```

Input: tensor in polar coordinates (trailing dimensions)  $\mathcal{X} \in \mathbb{R}^{I_1 \times \dots \times I_{N-2} \times I_\rho \times I_\phi}$ ,
        radius  $R$ , s.t.  $R \geq I_x/2$ 
Output: tensor in cartesian coordinates (trailing dimensions)  $\mathcal{X}' \in \mathbb{R}^{I_1 \times \dots \times I_{N-2} \times I_x \times I_y}$ 
1 Initialize:
2  $\mathbf{x}, \mathbf{y} \leftarrow [-R \quad -R+1 \quad \dots \quad R]^T$  // define cartesian resolution elements
3  $\mathbf{X} \leftarrow \mathbf{x} \circ \mathbf{y}$  //
4  $\mathbf{Y} \leftarrow \mathbf{y} \circ \mathbf{x}$  //
5  $\mathbf{F}, \mathbf{G} : (\rho, \phi) \mapsto (x, y)$  // discrete (pixel) flow from polar to cartesian coordinates
6  $\mathbf{F}^{I_{N-1}} \leftarrow \sqrt{\mathbf{X}^2 + \mathbf{Y}^2} + R$  // flow from polar to cartesian coordinates, column-wise
7  $\mathbf{G}^{I_N} \leftarrow \arctan2(\mathbf{X}, \mathbf{Y})$  // flow from polar to cartesian coordinates, row-wise

/* remap each pixel in the polar tensor to its new position in the cartesian
   tensor using bilinear interpolation */
8 for  $i_1 = 1$  to  $I_1$  do
9   ...
10  for  $i_{N-2} = 1$  to  $I_{N-2}$  do
11    for  $j = 1$  to  $I_\rho$  do
12      for  $k = 1$  to  $I_\phi$  do
13        /* remap each slice, in normalized coordinates:
14           $\text{dst}(x, y) = \text{src}(\mathbf{F}(x, y), \mathbf{G}(x, y))$ ,  $x = j/I_\rho$  and  $y = k/I_\phi$  */
15         $\mathcal{X}'_{i_1, \dots, i_{N-2}, j, k} \leftarrow \mathcal{X}_{i_1, \dots, i_{N-2}, \mathbf{F}_{j,k}, \mathbf{G}_{j,k}}$ 
16      end
17    end
18  end
19 return  $\mathcal{X}'$ 

```

$$(3.12)$$

Naturally, the strategies for improving the performance of tensor decomposition methods can be combined, so that e.g. the the annular tensors are unfolded along the spatial modes, as shown in [Equation 3.12](#).

3.2.6. Index Permutation

Many tensor decompositions, e.g. the TTD, are sensitive to the ordering (permutation) of the modes, which might impact their performance on ASDI observations. This means that there can be a choice of index permutation under which the decomposition is optimal in terms of [Equation 3.3](#). This is because of the computation of the TTD (see [Section 2.4.7](#)) which is performed sequentially on the reduced mode- n unfoldings so that specific permutations can capture the inherent structure of the tensor better. Conversely, the objective might instead be to find a tensor PSF model which is smooth and optimal under an arbitrary index permutation. This is still an open problem, and it is not known if satisfactory algorithms exist in the moderate-to-low SNR case of HCI. ([Lee et al., 2021](#); [Zhang et al., 2018b](#))

Experimentally, this work finds that for the TTD and TRD the index permutation has a negligible impact for the order-4 case. In the following the naive permutation of ASDI as λ, θ, x, y , and $\lambda, \theta, \rho, \phi$ for the polar

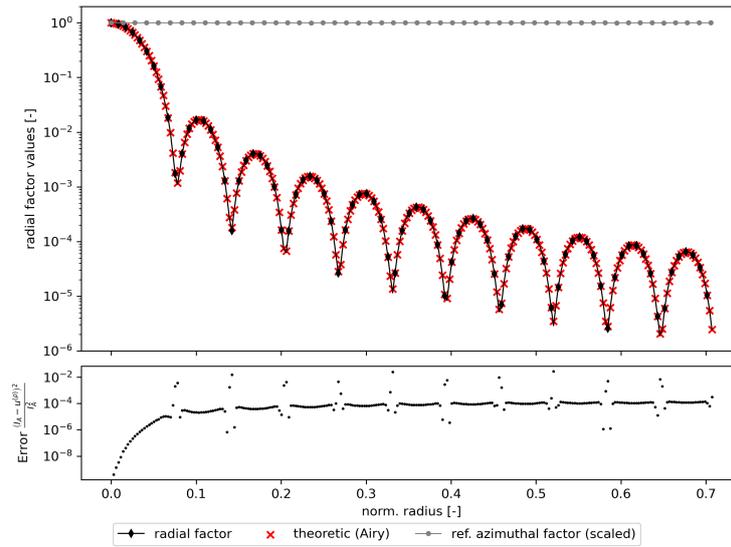


Figure 3.8: Rank-1 CPD Airy pattern model and the theoretic Airy function discretized to the image grid. The radial factor matches the theoretic pattern almost exactly, except at the radii with perfect destructive interference. The lower plot shows the signal-power normalized square error between CPD model and theoretic function. Note that this is the practically achievable error since the data is transformed to polar, and the CPD model back to Cartesian coordinates. The azimuthal factor is shown for reference projected onto the radial domain and is constant up to machine precision.

case, is exclusively considered.

3.3. Implementation

The implementation of the tensor-based PSF modeling and subtraction methods is structured as a modular and extensible Python framework that generates specific post-processing methods through composition. Built around a hierarchy of abstract base classes, the framework supports various coordinate systems and processing strategies, including polar and annular techniques. A key feature is the dynamic class composition achieved through a factory module, allowing for automatic and flexible creation of specific CODI algorithms. The implementation emphasizes compatibility with existing tools, particularly the Vortex Image Processing package (VIP, [Christiaens et al., 2023](#); [Gonzalez et al., 2017](#)), facilitating easy adoption and comparison with traditional methods. By leveraging GPU acceleration and providing a user-friendly interface, the framework enables efficient processing of high-contrast imaging data while maintaining flexibility for future extensions and improvements. A detailed description of the implementation can be found in [Appendix F](#).

3.4. Chapter Summary & Contributions

This chapter describes a novel class of PSF subtraction methods which use tensor decompositions to model the coronagraphic PSF and speckle pattern in observations obtained by angular-spectral differential imaging of exoplanets. These methods are a significant and original contribution to the field of exoplanet direct imaging.

The chapter begins by revisiting the current standard in post-processing, angular-spectral differential imaging with matrix-PCA, establishing a baseline for comparison. It then introduces the general framework for the new methods, explaining how they are implemented and how they preserve the higher-order structure of ASDI data.

The core of the chapter details various approaches to modeling the coronagraphic PSF and speckle pattern, including full-frame methods, methods using unfolding of the spatial modes, methods in the polar domain, and annular methods. Each approach is explored in terms of its motivation, advantages, and limitations. Lastly, the implementation of the methods is explained.

Throughout the chapter, the benefits of the new methods for improving PSF modeling, enhancing contrast, and increasing sensitivity in exoplanet detection are highlighted.

Contributions 3.1: Methodology

This chapter makes several key contributions to the field of high-contrast imaging:

- It describes a new class of PSF modeling and subtraction methods based on tensor decompositions which preserve the multi-modal structure of ASDI data.
- It provides a comprehensive exploration of various tensor decomposition techniques applied to exoplanet imaging data.
- It develops the methods beyond their naive implementation by utilizing existing ideas such as annular models and novel approaches enabled by tensor decompositions such as tensorization and index-permutation.
- It discusses the theoretical advantages of tensor methods over traditional matrix-based approaches, including improved modeling of spectral-temporal-spatial features and potential computational benefits.
- It lays the groundwork for future research in tensor-based methods for high-contrast imaging, including potential applications of more advanced tensor techniques.

Overall, this chapter establishes a new framework for PSF modeling in high-contrast imaging, offering potential improvements in exoplanet detection sensitivity.

The methods presented in this chapter form the main theoretical contribution of this work. Real observational data is required to assess the methods. To this end, the next chapter describes the instrumentation which produced these observations and the resulting datasets used in this study. This context is crucial for interpreting the post-processing results and understanding the limitations of the systematic assessment of the proposed methods.

Instrumentation & Data

This chapter provides an overview of the observational data and instruments central to this work.

It begins by introducing the Spectro-Polarimetric High-contrast Exoplanet REsearch instrument of the Very Large Telescope, which has revolutionized the field of direct imaging of exoplanets and circumstellar disks. The chapter expands on the components used in exoplanet imaging, including its adaptive optics system, coronagraph, and imaging spectrographs, to highlight the technical foundations underpinning the data.

Following this instrument overview, the chapter introduces the primary surveys and datasets that form the basis of this work's analysis. These include the SPHERE infrared survey for exoplanets, the B-star Exoplanet Abundance Study, and data from the Direct Imaging Virtual Archive. Each of these surveys contributes unique insights into exoplanet demographics and detection methods. The chapter also discusses the Exoplanet Imaging Data Challenge, which provides independently generated synthetic datasets for validating the post-processing methods.

By exploring both the instrumentation and the observational datasets, this chapter lays the groundwork for understanding the origins, quality, and limitations of the data used in subsequent analyses.

4.1. VLT/SPHERE

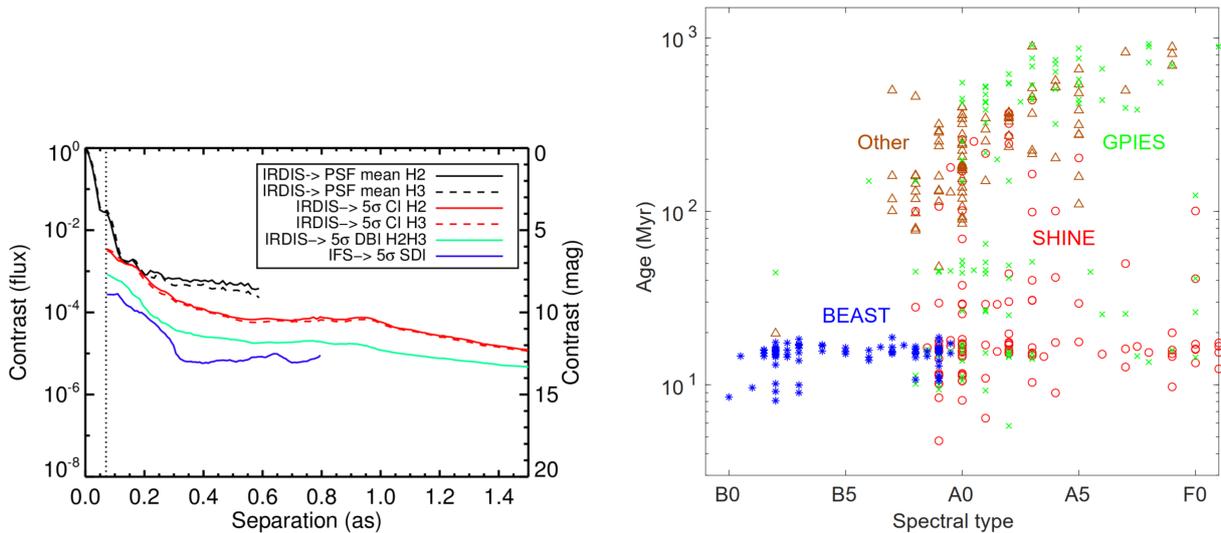
The Very Large Telescope (VLT), operated by the European Southern Observatory (ESO), is one of the world's most advanced ground-based optical and infrared telescopes, located at the Paranal Observatory in Chile. The VLT consists of four Unit Telescopes (UTs), each with an $\sim 8\text{m}$ diameter primary mirror. These telescopes can work together to achieve high-resolution imaging or operate individually with specialized instruments. SPHERE (Spectro-Polarimetric High-contrast Exoplanet REsearch, [Beuzit et al., 2008](#)) is mounted on UT3 (Melipal) and stands out among the VLT's suite of instruments for its focus on directly imaging exoplanets and circumstellar disks. Its main subsystems are the Integral Field Spectrograph (IFS) and Infra-Red Dual Imaging Spectrograph (IRDIS) together with their common path and infrastructure. The common path and infrastructure contains the elements that both the IFS and IRDIS need to image at high contrasts, namely the derotator, adaptive optics, and coronagraph. Depending on the instrument operating mode, both the IFS and IRDIS can operate simultaneously in different spectral bands. SPHERE also includes the Zurich Imaging Polarimeter (ZIMPOL) for polarimetry.

The IFS within SPHERE is designed to capture both spatial and spectral information simultaneously. It operates in the near infrared from ~ 0.95 to $1.66\mu\text{m}$ over 39 spectral bands and uses a hexagonal lenslet array to segment and disperse the light, combined with an infrared detector. Complementing the IFS, IRDIS specializes in dual-band imaging in the near-infrared spectrum. The dual-band technique captures images in two spectral bands at a much larger FOV. ([Mesa et al., 2015](#); [Samland et al., 2022](#); [Zurlo et al., 2014](#))

4.1.1. Coronagraph

SPHERE uses an improved version of the classical Lyot coronagraph, the Apodized Pupil Lyot Coronagraph (APLC, [Guerra et al., 2011](#); [Martinez, 2010](#)). The APLC is a type of hybrid coronagraph that combines elements of both pupil-plane and focal-plane coronagraphy to achieve high contrast performance and robustness to aberrations. It is widely used in current ground-based high-contrast imaging instruments, such as SPHERE, GPI, and SCExAO. The APLC consists of three main components:

Apodizer The first stage of the APLC is a pupil-plane apodizer (B in [Figure 4.2](#)), which is a gray-scale amplitude mask that reshapes the telescope pupil to optimize the brightness distribution of the stellar PSF.



(a) Ideal 5σ residual noise levels for SPHERE's IRDIS and IFS in a lab environment. The curves show data with synthetic planets injected (see Section 5.3). The vertical dotted line marks the coronagraph's IWA. Displayed are the mean radial PSF profile (black) and coronagraphic profile (red) for IRDIS channels H2 (solid) and H3 (dashed). Contrast levels after SDI reduction for IRDIS (green, dual-band) and SDI reduction for IFS (blue, multi-band) are also shown. (Zurlo et al., 2014)

(b) SHINE (red circles), BEAST (blue asterisks), as well as GPIES stellar system samples (green crosses). While SHINE focuses on a much broader population of stars, BEAST samples exclusively from the Scorpius-Centaurus association which have a common formation history over several star-forming molecular clouds. (Janson et al., 2021b)

Figure 4.1

The apodizer is designed to minimize the diffraction effects caused by the telescope's aperture, including the central obscuration, spiders, and segmentation. By carefully tailoring the transmission profile of the apodizer, the APLC can create a PSF with reduced diffraction rings and a more concentrated core.

Focal-plane mask Following the apodizer, the light is focused onto a focal-plane mask (C in Figure 4.2), similar to that used in a classical Lyot coronagraph. This mask, typically a circular occulter, blocks the central part of the stellar PSF, effectively suppressing the starlight. The size of the focal-plane mask strongly determines the IWA of the coronagraph, see Figure 2.10b.

Lyot stop After the focal-plane mask, the light is re-collimated, and a Lyot stop (D in Figure 4.2) is placed in the subsequent pupil plane. The Lyot stop is an amplitude mask that blocks the residual starlight that is diffracted around the edges of the focal-plane mask. The size and shape of the Lyot stop are optimized to maximize the suppression of the diffracted starlight while preserving the off-axis planet signal.

The combination of these three elements in the APLC results in a high-contrast region, or "dark hole", in the final focal plane, where the starlight is effectively suppressed, and faint companions or circumstellar material can be detected. This dark hole is shown in the left most image of Figure 4.2 at (E).

One of the key advantages of the APLC is its improved tolerance to low-order aberrations, such as tip-tilt errors and defocus, compared to other coronagraph designs. The apodizer helps to reduce the sensitivity of the coronagraph to these aberrations, making it more robust in the presence of atmospheric turbulence and instrumental vibrations. Another benefit of the APLC is its ability to provide high contrast over a wide spectral bandwidth, making it well-suited for spectroscopic characterization of exoplanets. (Laurent et al., 2018; Por, 2020; Zimmerman et al., 2016)

In practice, the APLC has demonstrated impressive performance on current ground-based instruments. For example, SPHERE routinely achieves contrasts of 10^{-5} to 10^{-6} at angular separations of a few hundred milliarcseconds, enabling the detection and characterization of young, self-luminous exoplanets and circumstellar disks. (Galicher et al., 2023)

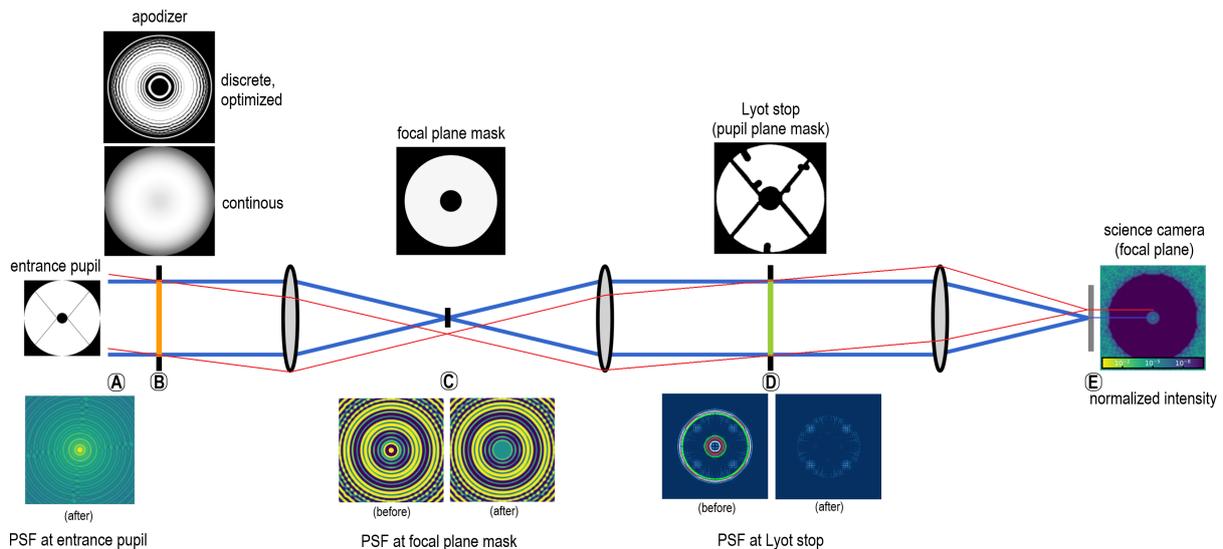


Figure 4.2: Apodized Pupil Lyot Coronagraph (APLC) architecture of VLT/SPHERE. In the center the APLC primary elements are shown schematically as a cut-through. Light from an on-axis source (host star) is shown in blue and light from an off-axis source (companion) is shown in red. On top the physical elements in the optical path are shown as seen from the left. On the bottom the log-intensity or phase pattern of the light at the respective locations are shown, again viewed from the left. The intensity images show the intensity distribution assuming it is imaged at that plane. The log-intensity and phase images do not include light from the off-axis source. **(A)** Light from the primary aperture of the VLT enters the coronagraph via the entrance pupil. The light is already diffracted by the spiders and around the secondary reflector. The light is then apodized at **(B)** to archive a better distribution of the incoming flux. SPHERE uses a discrete apodizer as shown on top. At **(C)** the focal plane mask removes most of the light from the on-axis source. The images below (C) are the real part of the electric field. The Lyot stop at **(D)** removes the light that is diffracted around the focal plane mask. In the intensity distribution images below (D) the Lyot stop removes blocks all light outside the green circle. This step however, also removes light from the off-axis source, decreasing the amount of flux received from potential companions. Finally, at **(E)**, the light is refocused and can be imaged by a detector. The final normalized intensity image shows the dark hole in which exoplanets can be imaged. Adapted from [Guerra et al., 2011](#); [Leboulloux et al., 2022](#); [Por, 2020](#).

4.1.2. Adaptive Optics

The adaptive optics system on SPHERE, known as the SPHERE Adaptive Optics for exoplanet Observation (SAXO, [Fusco et al., 2016](#); [Petit et al., 2008](#)), corrects turbulence induced distortions in near real-time. SAXO consists of a 40x40 segment deformable mirror, a Shack-Hartmann wavefront sensor, a tip-tilt mirror, and control algorithms. SAXO sits in front of and operates in conjunction with the SPHERE's APLC, providing a wavefront error equivalent to a Strehl ratio of > 0.9 for IRDIS and IFS. ([Petit et al., 2008](#))

Wavefront Sensor, Deformable and Tip-Tilt Mirror SAXO uses a 40x40 Shack-Hartmann wavefront sensor to measure the incoming wavefront distortions. The wavefront sensor operates within a spectral range of 0.45 to $0.95\mu\text{m}$ and has a temporal sampling frequency of ~ 1.2 kHz. The high-frequency operation and low read-out noise of the wavefront sensor's electron-multiplying CCD detector are critical for obtaining useful wavefront measurements. The deformable mirror, located in a pupil plane, corrects high-order wavefront errors caused by atmospheric turbulence and instrumental vibrations. The tip-tilt mirror compensates for low-order wavefront errors at 800 Hz. ([Jones et al., 2022](#); [Sauvage et al., 2010](#))

Control SAXO's control system employs a mixed numerical control law. The tip-tilt control is handled by a Kalman filter, which optimally deals with turbulence and reduces residual tip-tilt errors caused by vibrations. For the deformable mirror, an optimal modal gain integrator is used to control higher-order modes. The global adaptive optics loop delay, including wavefront sensor read-out and real-time computer delays, is below 1 ms to ensure rapid correction of wavefront errors. ([Petit et al., 2008](#); [Sauvage et al., 2010](#))

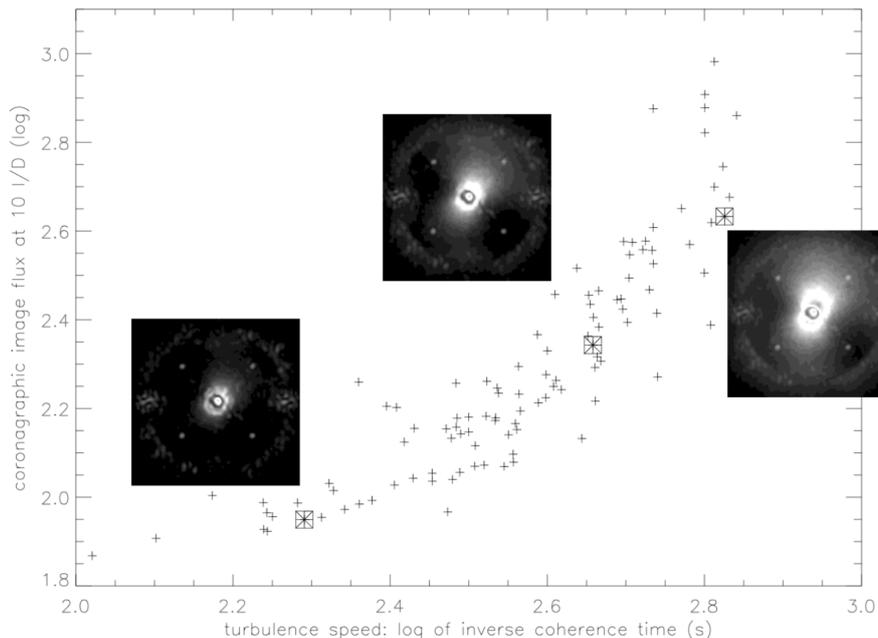


Figure 4.3: Correlation between coronagraphic raw contrast at $10\lambda/D$ (red lines in Figure 4.1a) and turbulence speed (inverse coherence time in log scale: slow turbulence on the left, faster towards the right). Adaptive optics servo-lag error and turbulence significantly impact image quality, with raw contrast scaling inversely with the turbulence speed, or roughly proportional to $r_0^{-3/2}$. (Boccaletti et al., 2020).

Non-Common Path Aberrations Non-common path aberrations refer to the disparities in optical distortions between the wavefront sensing pathway and the scientific observation channel. These distortions compromise the quality of scientific images and pose a unique challenge as they remain undetectable by the adaptive optics system’s primary correction loop. Non-common path aberrations result in systematic discrepancies between the actual wavefront reaching the science detectors and the wavefront measured by the sensor. Consequently, even when the wavefront sensor suggests optimal correction, the science path may still suffer from residual aberrations that can obscure or imitate real astronomical signals. In the context of high-contrast imaging, particularly for exoplanet detection, these systematic errors can produce enduring speckle patterns in the image plane. These speckles are challenging to differentiate from potential planetary signals, thus significantly reducing the system’s ability to achieve high contrast ratios. To mitigate this issue, SAXO implements a two-pronged approach: it conducts offline calibrations and performs real-time corrections using phase diversity techniques. This method involves comparing both in-focus and slightly defocused images to infer the aberrations, which are subsequently compensated for, thereby enhancing the overall image quality in the science channel. (Henault, 2019; Petit et al., 2008; Sauvage et al., 2010)

4.1.3. Speckle Pattern & Combined Adaptive Optics-Coronagraphic Imaging

The speckle pattern produced in the adaptive optics-corrected images is a combination of residual atmospheric turbulence and instrumental aberrations. To mitigate this, SPHERE employs the APLC described above. However, the performance of the coronagraph is heavily influenced by the quality of the adaptive optics correction. Imperfect correction results in speckles within the dark hole created by the coronagraph, which can be mistaken for or obscure faint companions. This speckle pattern is removed by the post-processing methods described in this work. (Jones et al., 2022)

In Figure 2.13 a simulated speckle pattern is shown. However, other effects further contribute to image degradation in complex, co-dependent way. Next to the non-common path aberrations discussed previously, this is wind and adaptive optics servo-lag. Figure 4.3 shows their combined impact on the science images. In the rightmost image, the wind-driven halo spans a large fraction of the dark hole with a wavelength dependent speckle pattern. The final images are thus degraded in several ways which are wavelength dependent and often change throughout an observation. (Boccaletti et al., 2020)

4.2. Observations

This section details the datasets and observational campaigns utilized in this thesis. For each survey and associated dataset, the objectives, methodologies, sample, and key findings are described. Each dataset represents a significant contribution to the field of exoplanet research. The observations span a range of targets, instruments and methods, providing a robust foundation for the analysis presented in this work.

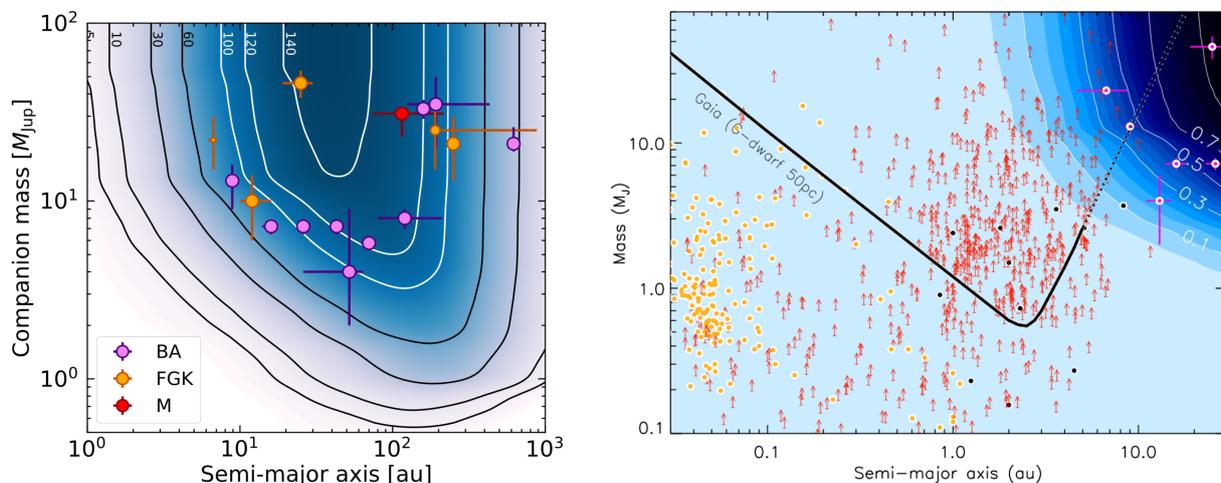
The data sources include the SPHERE infrared survey for exoplanets, the B-star Exoplanet Abundance Study, the Direct Imaging Virtual Archive and the Exoplanet Imaging Data Challenge. The pre-processed and post-processed observations from these sources are essential for the subsequent analysis and findings discussed in this thesis.

The full sample of observations is presented in [Section B.3](#) and [Table B.3](#).

4.2.1. SHINE

This work makes use of the SPHERE infrared survey for exoplanets (SHINE, [Chauvin et al., 2017](#)) which observed young stars in the solar neighbourhood between 2015 and 2017 with IFS and IRDIS on SPHERE. This survey spans roughly 150 stars observed in the YJ and YJH bands over 200 telescope nights and uses a coronagraph in addition to an angular and spectral differential imaging strategy. While the survey was designed to target substellar companions, several stars with circumstellar disks are also part of the sample. This survey has been selected as the de-facto dataset in this work as it is one of the largest, uniformly processed, and publicly accessible direct imaging surveys to-date ([Desidera et al., 2021](#); [Langlois et al., 2021](#); [Vigan et al., 2021](#)).

The SHINE survey has archived several key objectives: identifying and characterizing new planetary and brown dwarf companions, studying the architecture of planetary systems, investigating the link between planets and disks, determining the occurrence rate of giant planets beyond 10 AU, and exploring the impact of stellar mass on planetary characteristics ([Vigan et al., 2021](#)). To accomplish these goals, the survey originally targeted 500 stars from a sample of 800 nearby young stars (see [Figure 4.1b](#)), ensuring a broad range of stellar masses and optimal scheduling to maximize detection capabilities ([Langlois et al., 2021](#)).



(a) Depth of search of the SHINE survey. The black and white contour lines show the number of stars around which the survey can detect substellar companions based on mass and semi-major axis. The mass is approximated using a stellar evolution model. Colored circles denote the detected substellar companions, with the color indicating the spectral type of the host star. ([Vigan et al., 2021](#))

(b) SPHERE detection probabilities from the SHINE sample compared to the current exoplanet population detected by various methods: transit (yellow dot), radial velocity (red arrow), microlensing (black dot), and direct imaging (pink dot). Projected Gaia (DR4) detection limits for a solar-type star at 50 pc are shown for reference by the black line. The detection probability is related to the fractional value of detected substellar companions in [Figure 4.4a](#). ([Boccaletti et al., 2020](#))

Figure 4.4

In addition, it has achieved significant discoveries, initially detecting sixteen substellar companions, including eight brown dwarfs and eight planetary-mass companions. Notable discoveries from the survey include HIP 65426 b and HIP 64892 B ([Langlois et al., 2021](#)). The breadth of the survey and its detection performance

is illustrated in [Figure 4.4a](#), showing the depth of search based on mass and semi-major axis, with the detected substellar companions highlighted.

Statistical analysis of the SHINE data provides constraints on the occurrence and characteristics of exoplanets at wide orbits, offering valuable insights into planetary formation and evolution mechanisms. The survey's detection probabilities compared to other exoplanet detection methods are shown in [Figure 4.4b](#), underscoring its contributions to understanding exoplanet demographics ([Boccaletti et al., 2020](#)).

Additionally, the SHINE survey has imaged twelve circumstellar disks, including new detections around HIP 73145, HIP 86598, and HD 106906, despite not being initially designed for disk searches. These findings highlight the survey's comprehensive contributions to the study of young, nearby stars and their surrounding environments ([Langlois et al., 2021](#)).

Part of the observations from SHINE have been re-processed for the 2018 B-star Exoplanet Abundance Study (see below). When available, this work uses the most recently processed version of observations to make use of the improved pipelines ([Delorme, 2024](#)). The pre-processed observations were obtained from the High Contrast Data Center (HC-DC, [Delorme et al., 2017](#)).

4.2.2. BEAST

This work further draws on some observations from the aforementioned B-star Exoplanet Abundance Study (BEAST, [Janson et al., 2019, 2021b](#)), which is a VLT-SPHERE direct imaging survey of 85 B-type stars aimed at detecting wide-separation companions to constrain stellar mass limits for planet formation. Like SHINE, BEAST uses IFS and IRDIS and makes use of the ASDI observation technique. The BEAST survey is designed to explore the occurrence and properties of giant planets at wide separations around B-type stars. It seeks to determine whether an upper stellar mass limit exists for planet formation. By comparing the BEAST sample with surveys of lower-mass stars in the same region, the survey controls for factors such as age and metallicity, isolating stellar mass as a key variable influencing planet occurrence rates.

To ensure uniformity, the survey focuses on Scorpius-Centaurus members, maintaining a sample with consistent ages and metallicities ([Figure 4.1b](#)), which enhances the reliability of statistical analyses. Targets already observed by SPHERE and those with intermediate separation binaries were excluded.

The initial observations of 67 out of the 85 targets have identified 708 candidate companions, which require follow-up observation to confirm common proper motion and physical association with the primary stars. From these candidates, [Chomez et al., 2023b](#); [Gratton et al., 2023](#); [Janson et al., 2021a,b](#); [Squicciarini et al., 2022](#); [Viswanath et al., 2023](#) and [Herz et al., 2024](#) report the discovery of seven new binary companions. This survey is expected to lead to more new companion discoveries as follow-up observations target promising candidates.

This work relies on [Janson et al., 2021b](#) and the previously mentioned papers for pre-processed observations, obtained from the HC-DC, and for identifying background objects and known sources in observations. The BEAST dataset complements the SHINE survey, providing valuable insights into planetary formation around massive stars and contributing to the broader analysis of exoplanet detection and characterization techniques presented in this thesis.

4.2.3. DIVA

This work implements existing post-processing methods like matrix-PCA to ensure a homogeneous data reduction. Nevertheless, the methods are verified against existing frameworks and datasets. The SHINE survey produced the most comprehensive dataset of SPHERE direct imaging observations. This data is provided by the Direct Imaging Virtual Archive (DIVA, [Vigan et al., 2017](#)). Specifically, this work uses the post-processed residual images, cubes, detection maps and detections of the SHINE survey hosted by DIVA. The post-processed images and cubes are used to verify the correct implementation of matrix-PCA-based reduction. The detection maps and table of detected point sources parameters provided by DIVA are used to cross-reference detections and background objects in residuals obtained by using tensor methods.

4.2.4. EIDC

To validate the method on synthetically generated observations, this work uses the diverse set of pre-processed observations obtained from the Exoplanet Imaging Data Challenge (EIDC, [Cantalloube et al.,](#)

[2022, 2020](#)) Phase II, which include synthetically injected companions. Synthetic injection of companions is discussed in [Section 5.3](#). The observations in the EIDC include data obtained from the SPHERE-IFS and the Gemini Planet Imager (GPI [Chilcote et al., 2020](#); [Macintosh et al., 2006, 2014](#)).

4.3. Chapter Summary & Contributions

This chapter provides an overview of the observational data and instrumentation used to verify and assess the methods proposed in this work. It revisits the concept of exoplanet imaging instrumentation, focusing specifically on the SPHERE instrument on the VLT, which was used to obtain the observations used in this study. The chapter details SPHERE's key components for exoplanet imaging: the adaptive optics system, coronagraph, and imaging spectrographs, expanding on their individual impact on the observed data and the image formation process.

The chapter then introduces the primary datasets used in this study, all obtained by VLT/SPHERE: the SPHERE infrared survey for exoplanets (SHINE), the B-star Exoplanet Abundance Study (BEAST), and data from the Direct Imaging Virtual Archive (DIVA). It also discusses the Exoplanet Imaging Data Challenge (EIDC), which provides independently generated synthetic datasets for method validation.

Contributions 4.1: Observations

This chapter contributes to the thesis by:

- Providing an extended description of the specific instrumentation used to obtain this study's datasets, crucial for understanding the origin and characteristics of the data.
- Detailing the impact of the instrument's components on the observed data, including a description of the image formation process in the coronagraph used in this study.
- Introducing the major observational surveys and datasets used, including their objectives, methodologies, key findings and limitations.
- Introducing the independently generated synthetic datasets used for method validation.

This detailed overview of instrumentation and datasets provides the necessary foundation for interpreting the results of the tensor-based methods developed in this thesis and understanding their practical limitations.

The next chapter presents the results of using tensor-based methods to post-process the observational data described here. It evaluates their performance through statistical analysis, providing practical insights into the benefits and limitations of these approaches in high-contrast imaging.

Results

This chapter presents the results of applying tensor decomposition methods to direct imaging observations, focusing on the analysis of residual images from observations of several well-known systems. The chapter is structured to systematically evaluate the performance of the proposed tensor-based algorithms in comparison to traditional matrix-based PCA techniques.

The chapter begins by examining post-processed residual images from observations of the HR 8799 and β Pictoris systems, demonstrating the effectiveness of tensor decompositions in detecting known exoplanets and circumstellar features. This analysis serves to validate the practical utility of the proposed methods in real astronomical datasets. The application of tensor decompositions is also explored in more challenging scenarios, including the detection of fast moving disk features in the AU Microscopii system. This case study highlights the versatility of tensor-based methods in addressing complex imaging scenarios beyond point source detection. Additionally, the chapter reports the detection of a new companion candidate in the HD 108767 B system, as well as its spectrophotometric and astrometric characterization.

The chapter then progresses to a quantitative assessment of algorithm performance using synthetic companion injection. Detection capabilities are evaluated across a range of contrasts and separations, with a focus on the receiver operating characteristic curves and F_1 score to rigorously compare the various methods.

Finally, a runtime analysis is presented to assess the practical computational efficiency of tensor decomposition approaches relative to matrix-based approaches. This evaluation is crucial for understanding the practical implications of implementing these methods in exoplanet imaging pipelines.

Through this comprehensive examination, this chapter aims to demonstrate the capabilities, limitations, and potential advantages of tensor decomposition methods in the context of high-contrast imaging for exoplanet detection and characterization.

5.1. Detection Metrics

The main detection criterion used here is based on the signal-to-noise ratio (SNR) calculated using the method described by [Mawet et al., 2014](#). This approach accounts for small sample statistics, applying a t-student correction to the significance levels. This correction is particularly important at small separations where fewer resolution elements are available for noise estimation. The method employs a two-sample t-test to rigorously assess the statistical significance of potential detections. It defines a new SNR metric as

$$(\bar{x}_1 - \bar{x}_2) > \tau \times \sigma_2 \sqrt{1 + \frac{1}{n_2}}, \quad (5.1)$$

where \bar{x}_1 is the intensity of the test resolution element at separation r , \bar{x}_2 is the mean intensity of the remaining elements, σ_2 is their standard deviation, and $n_2 = 2\pi r$ is the number of remaining elements. $\tau = 5$ is the significance threshold above which a detection is considered a true positive, set as 5σ in this work. This SNR follows a Student's t-distribution with $n_2 - 1$ degrees of freedom, rather than a normal distribution. Consequently, it yields more conservative but statistically robust detection thresholds at small separations, with the threshold needing to be up to 10 times higher at $1\lambda/D$ and 2 times higher at $2\lambda/D$ compared to larger separations to maintain the same confidence level. This method is called SNR-map. ([Marois et al., 2008a](#); [Mawet et al., 2014](#))

While the SNR-map is generally a good baseline estimate for detections, other methods perform better in terms of true positive, false positive and false negative fractions ([Cantalloube et al., 2022, 2020](#)). One such

method is the Standardized Trajectory Intensity Mean map (STIM map, [Pairet et al., 2019](#)). The STIM map improves upon SNR-based detection maps by using a temporal model for reducing the temporal residual cube rather than just the final processed frame. It computes the ratio of the mean to standard deviation along each pixel trajectory through the data cube after PSF subtraction, leveraging the fact that residual speckle noise should have a low mean and high variance, while real planets will have a higher mean and lower variance. Unlike the SNR-map, the STIM map does not rely on comparing pixels at the same radius, avoiding issues with multiple planets or extended structures like disks at the same separation. It maintains consistent sample statistics across all separations from the star, improving performance at small inner working angles. The method provides a theoretical basis for setting detection thresholds based on the number of frames, accounting for temporal correlations. Empirically, the STIM map demonstrates improved true positive rates and lower false positive rates compared to SNR maps, especially for faint planets and at small separations. ([Bonse et al., 2023](#); [Pairet et al., 2019](#))

5.2. Post-Processed Residuals

To demonstrate the viability of tensor decomposition based methods for detecting point sources, the residual images of observations from two well known systems, HR 8799 and β Pictoris, are shown below.

5.2.1. HR 8799

HR 8799 is a young (approximately 30 million years old) A-type main-sequence star located about 129 light-years away in the Pegasus constellation. This system is notable for being one of the first directly imaged multi-planet systems, with four known giant planets orbiting the star. These exoplanets, designated HR 8799 b, c, d, and e, were discovered between 2008 and 2010 by [Marois et al., 2008b](#) and [Marois et al., 2010](#). The system has since become a key target for studying exoplanet atmospheres, orbital dynamics, and planetary system formation. ([Lafrenière et al., 2009](#))

The data for the HR 8799 system presented here are ASDI observations obtained using the SPHERE/IFS instrument on the 12. October 2017 and SPHERE/IRDIS on the 19. June 2018 as part of the SHINE observation campaign. The data has been pre-processed by the HC-DC to remove instrumental signatures and bad pixels, and to align the frames spatially.

The full-frame residual images (RI) of HR 8799, using the proposed methods, are shown in [Figure 5.1](#) along with a PCA reduction for comparison. Sources within the residual frames are identified using local peaks, and the SNR and significance are refined using the SNR-map ([Christiaens et al., 2023](#)), taking into account the noise in an annulus at the source position. In the IFS observation from October 2017 HR 8799 c appears only in a small fraction of the frames, making it much fainter than HR 8799 c and d in the residual images. In both the IFS and IRDIS observations, the known HR 8799 companions within the respective fields of view are detected above 5σ by all proposed algorithms, considering both full-frame and annular reductions.

In [Figure 5.1](#), the residual images from the Tucker, TT, and TR decompositions, as seen in [Figure 5.4](#), utilize spatial unfolding (approximating $\mathcal{X} \in \mathbb{R}^{I_\lambda \times I_\theta \times (I_x \cdot I_y)}$). This approach reduces the compressive performance of the decompositions, as they no longer approximate low-rank components in the spatial dimensions of the observations. Instead, each pixel is treated independently. For full-frame reductions in Cartesian coordinates, this significantly enhances the performance of the tensor decompositions by eliminating the issues with dimension alignment discussed in [Section 3.2.4](#). However, for the CPD in [Figure 5.1b](#), spatial unfolding is not feasible as it substantially increases the memory required to compute the factors of the CPD, necessitating the calculation of $\mathbf{U}^{(x,y)} \in \mathbb{R}^{(I_x \cdot I_y) \times R}$ instead of $\mathbf{U}^{(x)} \in \mathbb{R}^{I_x \times R}$ and $\mathbf{U}^{(y)} \in \mathbb{R}^{I_y \times R}$. Consequently, computing the factors using ALS (`tntorch`) and Robust Tensor Power Iteration (`tensorly`) is not feasible, even with low ranks.

Note the concentric circular artifacts visible in the CPD annular residual image in [Figure 5.2b](#). These structures are not natural features of the system but rather artifacts likely arising from the radial scaling in the annular approach together with the limited expressiveness of the CPD compared with the other decompositions. These artifacts degrade the noise statistics for any source in the image, as they increase the variance of the background.

In contrast to the poor reduction using the CPD on the full-frame in Cartesian coordinates ([Figure 5.1b](#)), the annular CPD reduction shown in [Figure 5.2b](#) is more comparable to the other methods. However,

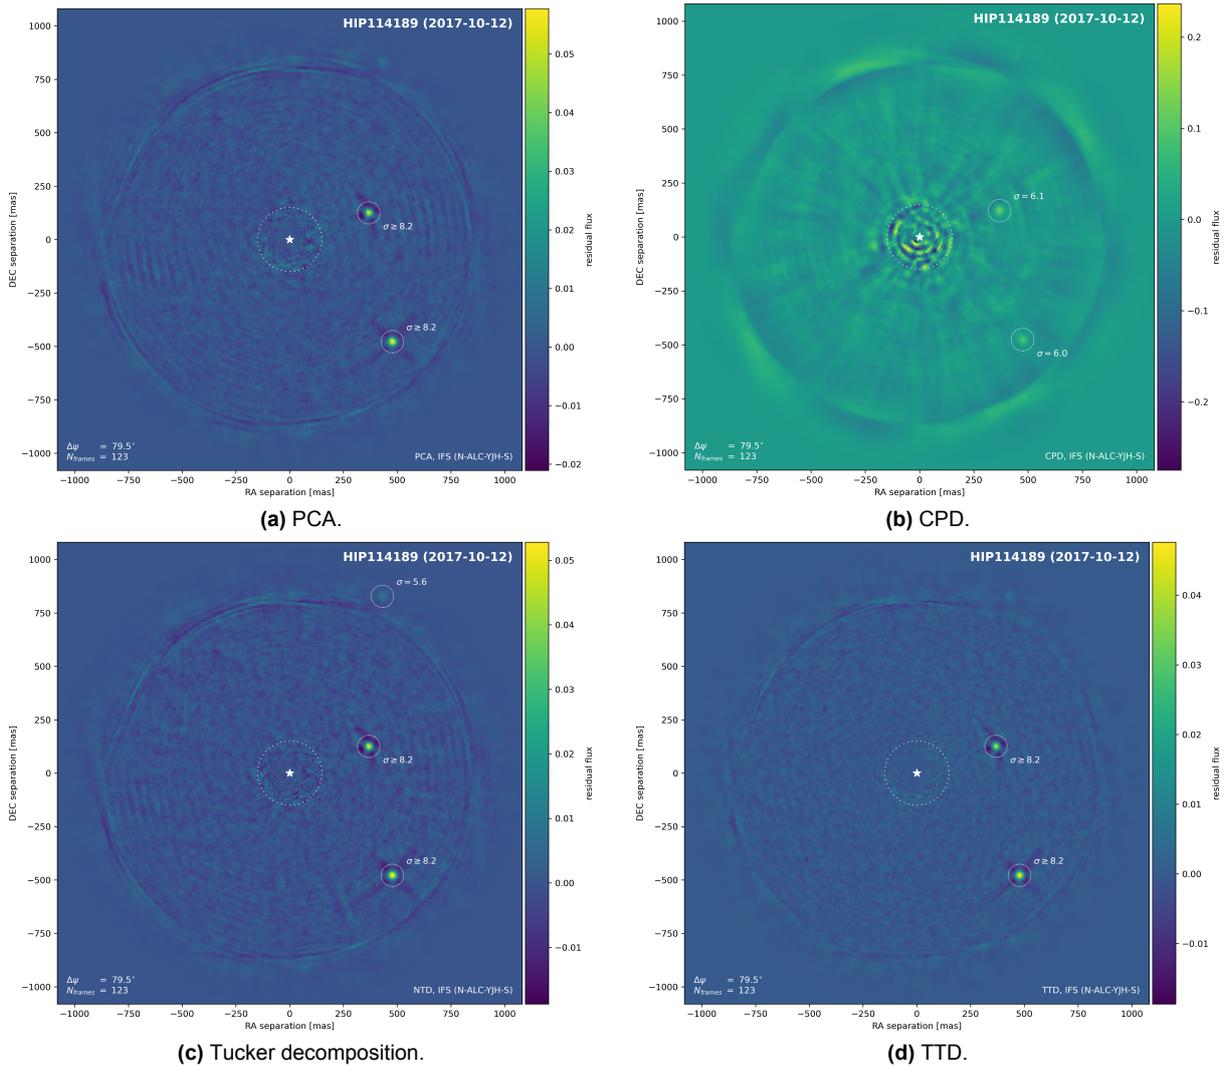


Figure 5.1: HR 8799 IFS residual images from full-frame algorithms. **(b) CPD:** The large, structured residuals at small separations show that the full-frame CPD on order-4 observations struggles to model the speckles near the IWA. Additionally, the much larger mean absolute residual flux indicates that the coronagraphic PSF is not fully subtracted. **(c) Tucker decomposition:** Spatial unfolding together with considering factor interactions leads to a significant performance increase over the CPD.

the requirement for low-rank spatial components results in inferior performance compared to PCA and other tensor decompositions. This impact is particularly pronounced at small separations where the PSF varies more strongly overall and especially at smaller spatial frequencies across the temporal and spectral dimensions. The annular versions of the Tucker and TTD methods generally enhance performance in this noise regime at small separations, as shown in [Figure 5.2](#).

The TRD performs well (compare [Figure 5.4](#) with [Figure 5.1](#)); however, this work considers it in less detail due to its similarity to the TTD. Additionally, while not a practical concern in the high SNR scenario of HR 8799, the TRD cannot find the best low-rank approximation, which presents a problem in cases where the SNR is below the threshold at which the TRD still recovers a good PSF approximation ([Batselier, 2018](#)).

To demonstrate the ability of tensor decompositions to post-process IRDIS observations, [Figure 5.5](#) shows the RI of the full HR 8799 system. The reduced dimension of the spectral mode in IRDIS observations compared to IFS (2 versus 39) means that these observations more closely resemble the matrix-case for which tensor decompositions are reduced to PCA. [Figure 5.5](#) demonstrates that in this scenario, the RI from PCA and tensor decompositions on spatially unfolded observations become very similar, as expected. In contrast to the IFS RI in [Figure 5.1](#) and [Figure 5.2](#), the IRDIS observations were transformed into polar

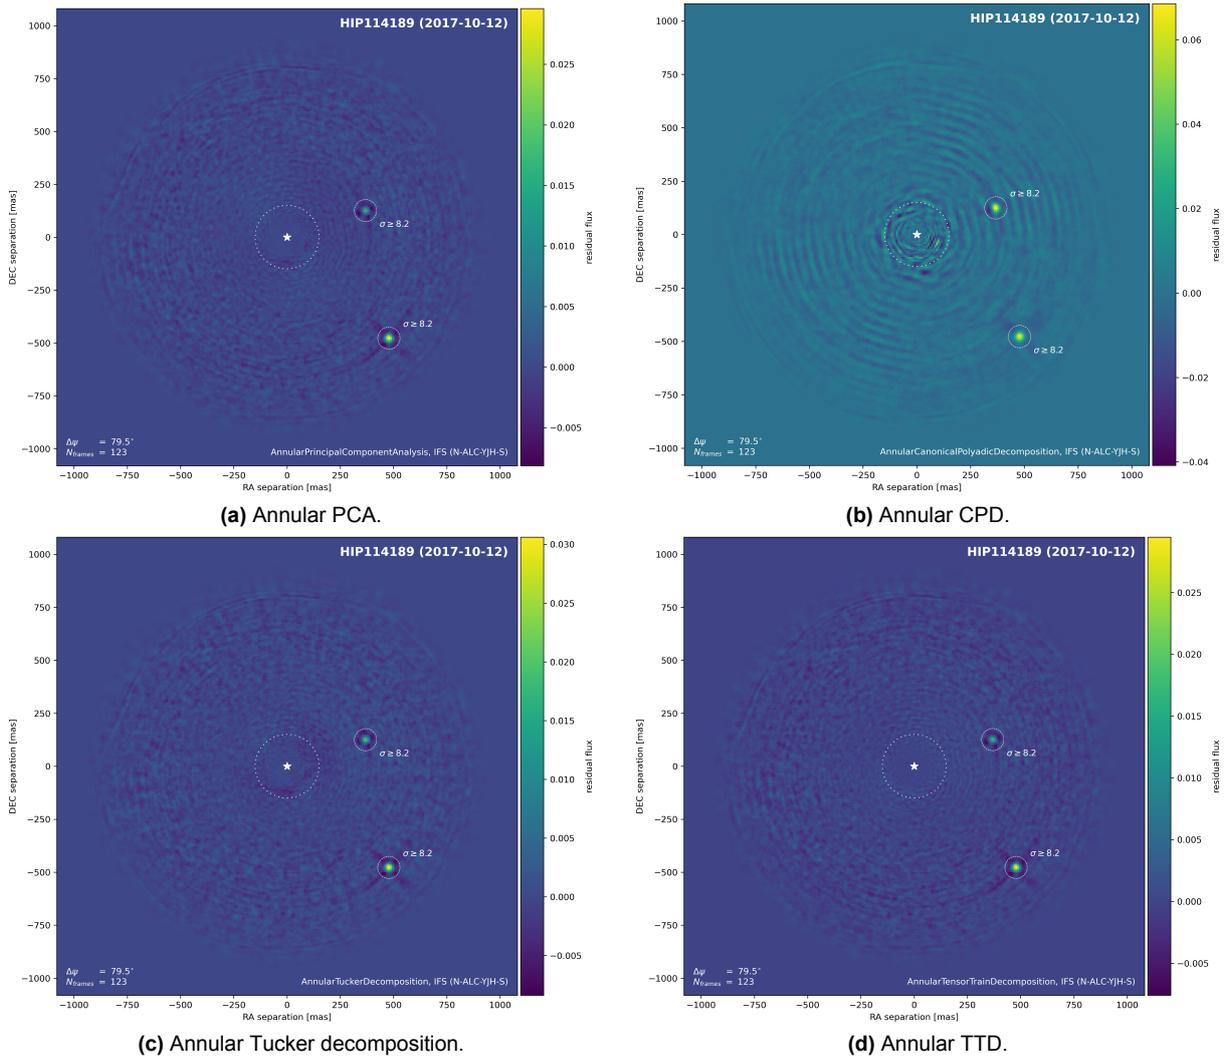


Figure 5.2: HR 8799 IFS residual images from annular algorithms. Note that for annular PCA, Tucker, and TTD the completeness fraction increases with decreasing separation so that HR 8799 e (and more generally sources at small separations) are more self-subtracted compared with Figure 5.1. However, the amount of self subtraction depends on the parameterization of the algorithms. **(b)** Annular CPD: Performing the CPD on annuli increases its performance significantly, bringing it more in line with the other algorithms. Nevertheless, note that the remaining background has structure that is not properly removed. **(c)** Annular Tucker decomposition: The annular version slightly improves the performance, especially at small and large separations. **(d)** Annular TTD: The performance improvement is similar to PCA and the Tucker decomposition.

coordinates before the entire polar frame was reduced, essentially using a single large annulus compared to many smaller ones. This indicates that aligning the coordinate system of the decomposition with the image generation process improves the RI, even without separately modeling regions with different noise regimes.

Having demonstrated the effectiveness of tensor decomposition methods on the HR 8799 system, the next section will consider another well-known exoplanetary system, β Pictoris, to further validate our approach.

5.2.2. β Pictoris

The ~ 23 Myr old β Pictoris system features a bright giant embedded in a nearly edge-on debris disk, as seen in Figure 5.6a and Figure 5.6b respectively (Lagrange et al., 2009; Mamajek et al., 2014). β Pictoris b was first detected in 2003 before passing behind its host star and being observed again in 2010 by Lagrange et al., 2010. This makes it an excellent case study for demonstrating the effectiveness of tensor decomposition in detecting exoplanets at small separations using a single annulus. More recently,

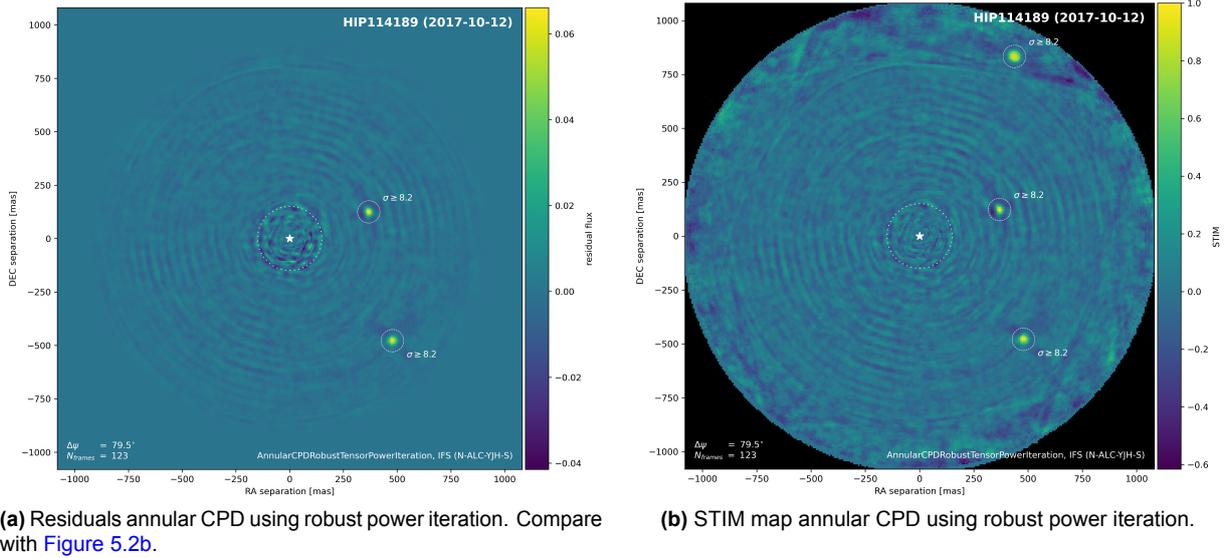


Figure 5.3: Annular CPD using robust power iteration on HR 8799 IFS observations.

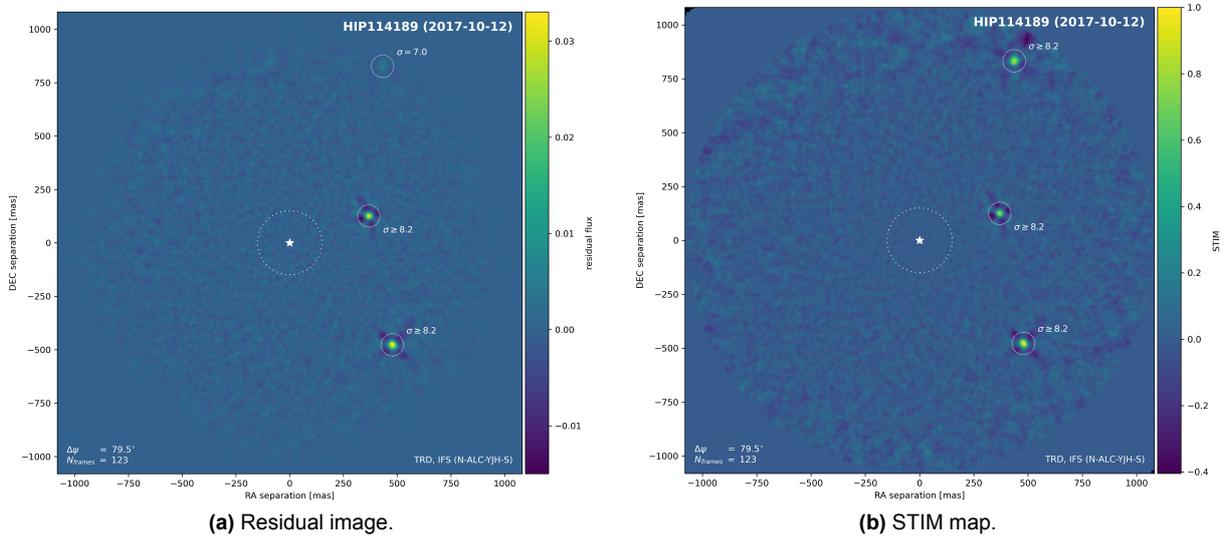


Figure 5.4: TRD on IFS observations of HR 8799. Note the improved structure of the residual noise compared to the algorithms in Figure 5.1 due to the periodic boundary of the decomposition.

a second close-in exoplanet was discovered using a combination of direct imaging and other methods (Lagrange et al., 2020).

The data for the β Pictoris system presented here are ASDI observations obtained using the SPHERE/IFS instrument on the 15. December 2018 and SPHERE/IRDIS on the 18. October 2018 as part of the SHINE observation campaign. The data has been pre-processed by the HC-DC to remove instrumental signatures and bad pixels, and to align the frames spatially.

Figure C.2 shows the RI of β Pictoris b. When considering spatially unfolded decompositions, the residual speckle pattern in the annulus is almost equivalent to the RI with PCA; however, the factors of the decompositions have significantly fewer parameters than PCA. This indicates that the PSF in this annulus has a low-rank structure in the spectral-temporal slices of the observations, which is not exploited when the observations are matrixized for PCA. However, when the full tensor shape is preserved, only low-rank components are captured in the spatial slices, which precludes the inclusion of high spatial frequency features or spatially sparse components in the PSF model.

While HR 8799 and β Pictoris showcase the ability of tensor decompositions to detect exoplanets, the

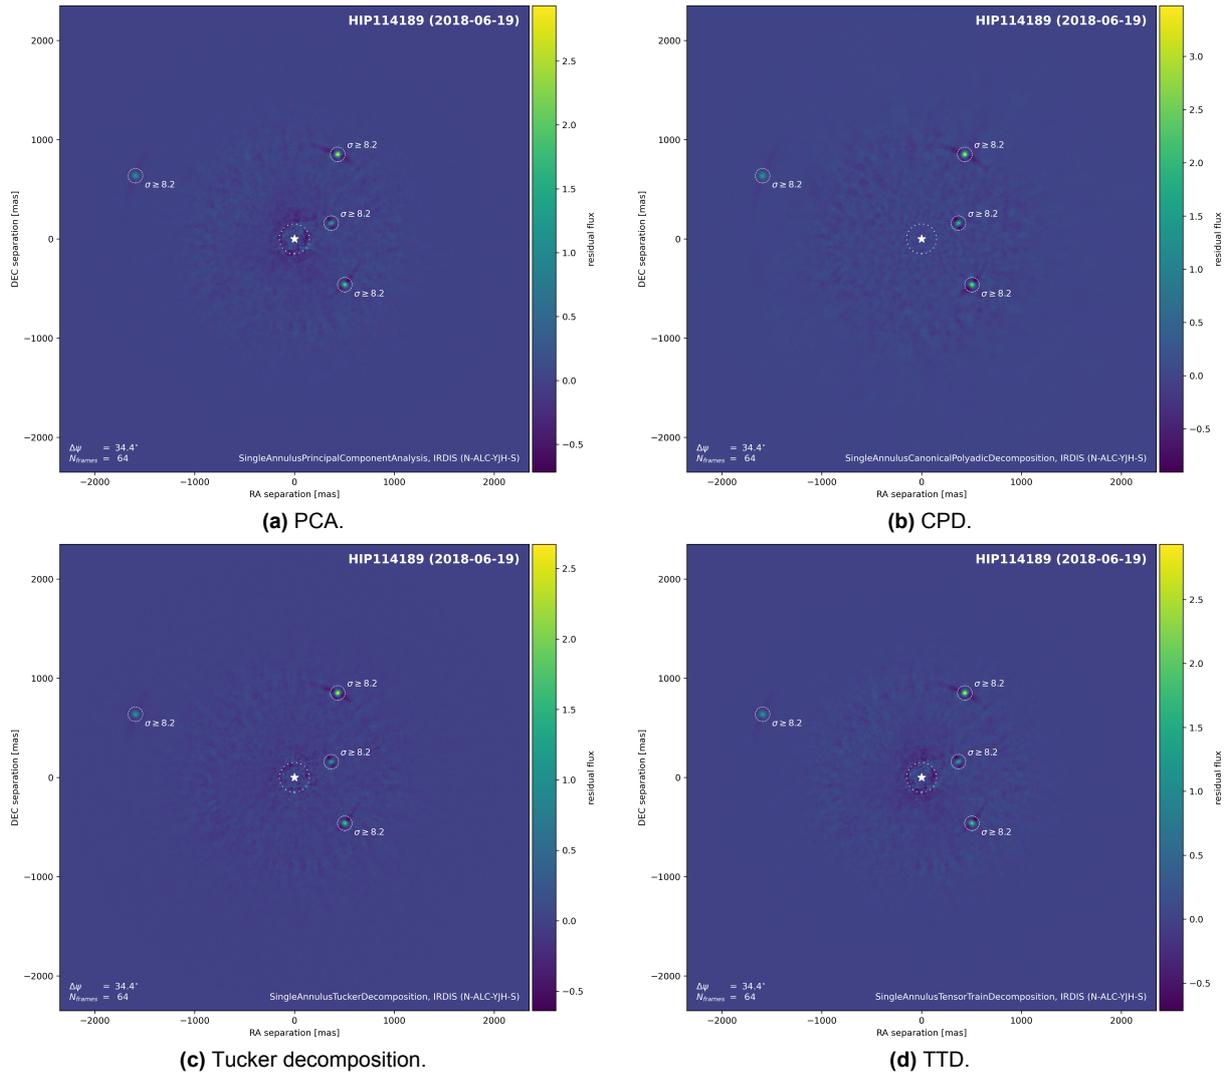


Figure 5.5: HR 8799 IRDIS residual images from algorithms on observations in polar coordinates. All algorithms perform similarly well on IRDIS data since the size of the spectral mode is very small.

AU Microscopii system presents an opportunity to demonstrate the method’s capability in identifying non-planetary features in debris disks.

5.2.3. Features in the AU Microscopii Debris Disk

The young star AU Microscopii, also a member of the β Pictoris moving group, hosts a low-inclination debris disk and several close-in exoplanets. Since 2014, the debris disk has garnered increased attention due to the detection of several fast-moving features with direct imaging (Boccaletti et al., 2018; Boccaletti et al., 2015). To demonstrate the ability of tensor decompositions to identify non-PSF-shaped features in disks, this work processed several previously published IFS observations of the AU Microscopii system, from Figure C.3a to Figure C.3d. The RI using TTD clearly confirm the presence of the previously detected NW1 and NW2 features below the disk’s mid-plane in the IFS observations. By tuning the rank of the decomposition, extended spatial features like the disk can be removed while preserving components at higher spatial frequencies to enhance the SNR of small-scale features. Importantly, the features are visually inconsistent with unresolved objects. The reductions obtained in this work, shown in Figure C.3a, show that tensor decomposition methods can improve residuals compared to PCA-processed RI by Boccaletti et al., 2018 (their Fig. 10).

In addition to confirming the NW features, this work detects a new SE feature at approximately $0.25''$ separation in 2019 IFS observations of AU Microscopii, Figure 5.7a. This feature was not present in

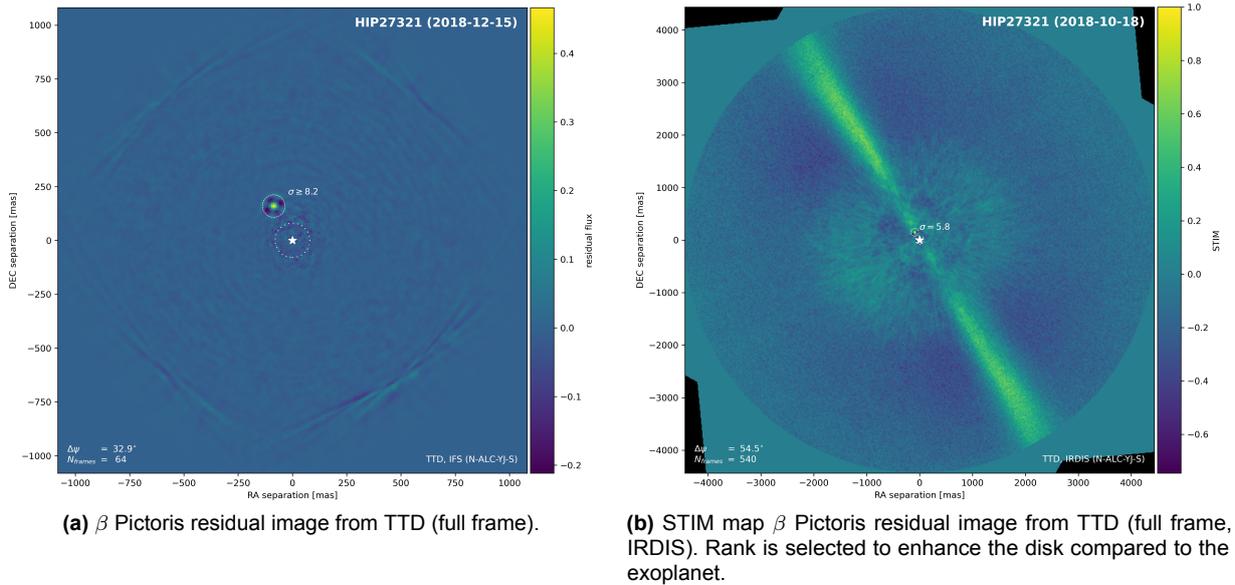


Figure 5.6: β Pictoris (full frame).

observations prior to 2016 and was not detected by [Boccaletti et al., 2018](#) in 2017 IFS observations or reported elsewhere since 2017. However, [Boccaletti et al., 2018](#) did report a candidate feature at the SE (SE6) which they observed over several epochs between 2016 and 2017 NE of the disk mid-plane, moving toward larger separations. This work does not find evidence of the SE6 feature in observations prior to 2016, and its previous apparent motion makes it unlikely to be the SE feature in [Figure 5.7a](#). Compared to the NW features, the structure of the new SE feature is much less spread out; however, its sudden appearance suggests it is produced by a process similar to the other features in the disk. Since the feature is on the opposite side of the disk to the confirmed features, its presence and motion can impose strong constraints on the formation process of the features in the AU Microscopii disk. Models that suggest collisional cascades and require either out-of-plane belts intersecting the main disk or the destruction of large bodies within the disk ([Chiang et al., 2017](#); [Grady et al., 2020](#)) are especially impacted by this feature.

Having examined systems with known sources, the next section will present the discovery of a new sub-stellar companion candidate, detected by the application of tensor decomposition methods.

5.2.4. HD 108767 B

HD 108767 B (δ Corvi B) is a K1 D star with an estimated age of 180^{+170}_{-80} Myr, located in the Corvus constellation at a distance of 26.7 pc. Its primary partner, HD 108767 A, is at a separation of 24.2" ([Chomez et al., 2023a](#); [Desidera et al., 2021](#); [Wenger et al., 2000](#)). Annular TTD has identified a source in IFS observations from 2016 and 2018, with significances of 5.3σ and 4.4σ respectively, at a separation of approximately 558 mas. This source is also detected at a slightly lower significance with annular PCA, as shown in [Figure 5.7b](#). The RI of both observations are shown in [Figure 5.8](#), and the corresponding STIM maps in [Figure C.4](#) indicate that the source has some structure that is obscured in the RI.

Between the 2016 and 2018 observations, the source shows an apparent motion of approximately 8 mas. Jointly fitting the position ([Figure 5.9a](#)) and spectrum of the source using the negative fake companion technique using nested sampling ([Buchner, 2016, 2019, 2021](#)) in the 2016 observation reveals an almost constant contrast with the host star ([Figure 5.9b](#)). In IRDIS observations from 2016 and 2018, no source above 3σ is detected at the same position.

Considering the PSF shape of the source in both the RI and STIM maps, particularly in the residual of the 2018 observation, along with the flat contrast and non-detection in the IRDIS bands, it is challenging to determine what process produces the signature. Neither HD 108767 B nor A exhibit an infrared excess or a known disk ([Ertel et al., 2014](#); [Su et al., 2006](#)). Additionally, this work does not find a disk around HD 108767 B in the IFS and IRDIS bands. Previously, no source was detected around HD 108767 B using VLT/NACO observations ([Kammerer et al., 2019](#)), and while results from the SHINE survey do not mention

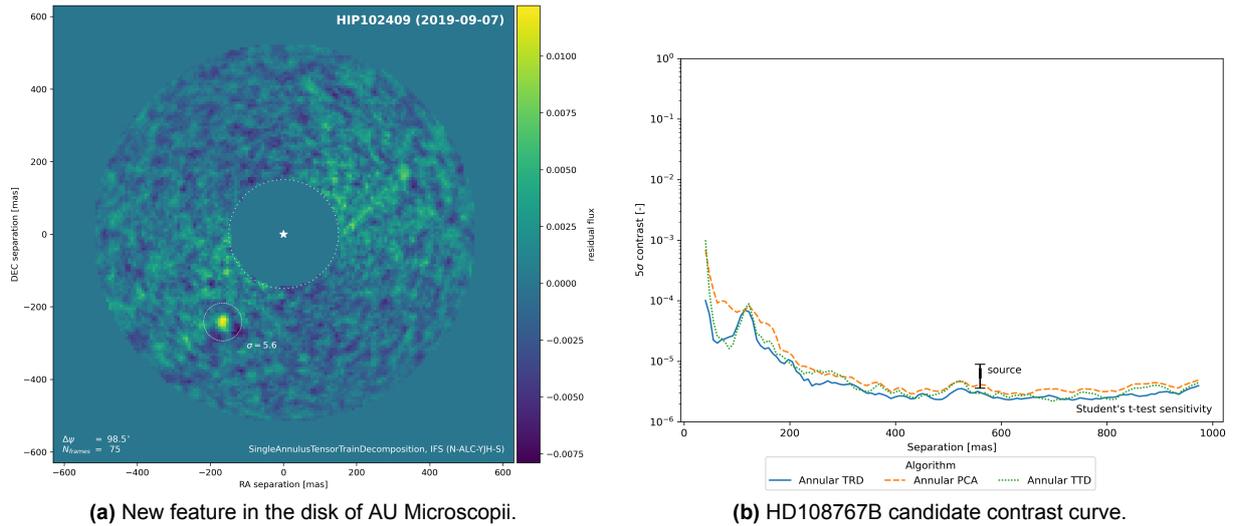


Figure 5.7

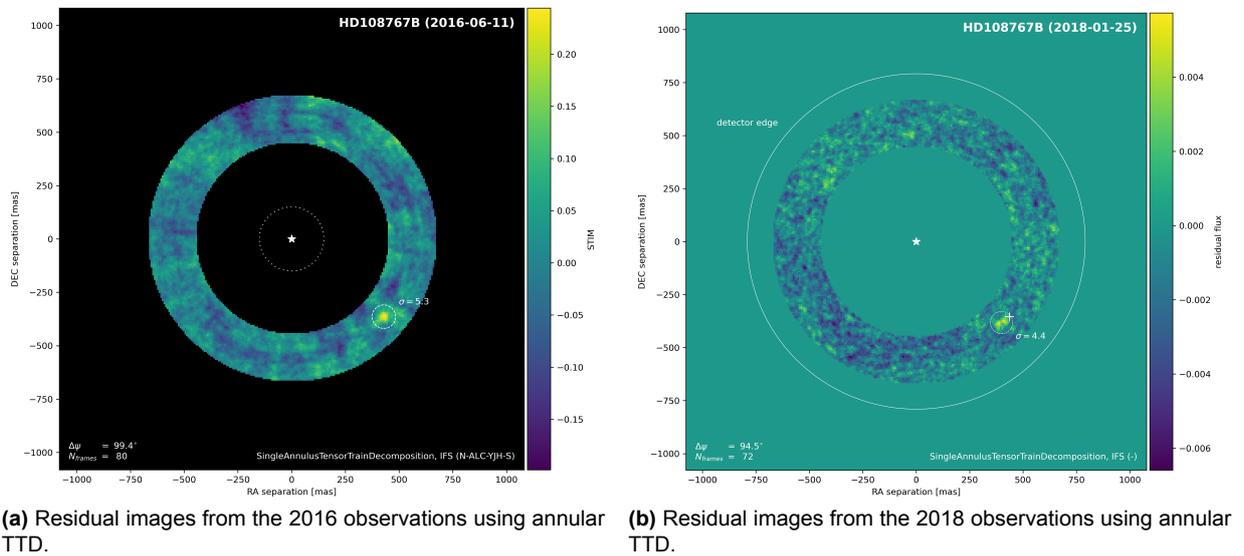


Figure 5.8: Residual images from IFS observations of HD 108767 B reduced with annular TTD. (b) The white cross indicates the best fit location of the source in the 2016 observations. The detector edge line indicates the approximate position of pixels that are effected by the edge of the IFS detector. When reducing observations, blops outside of this line might be due to overlapping regions from derotated frames instead of true sources, see e.g. the outer regions of Figure 5.6a.

HD 108767 B, this might be due to the follow-up observations being delayed until 2018 (Vigan et al., 2017).

Table B.1 lists all publicly accessible observations of HD 108767 B; however, none of the telescopes listed would be capable of detecting a source around HD 108767 B at the estimated contrast and position.

The detection of this potential companion to HD 108767 B, if confirmed, would have significant implications for the developed methods as it would serve as an independent confirmation of the methods utility. The unusual spectral characteristics and the lack of detection in IRDIS bands make this a particularly intriguing case. Further follow-up observations, together with a more in-depth forward modeling approach accounting for the potential presence of a circumstellar disk will be crucial to confirm the nature of this source and to constrain its physical properties.

To rigorously evaluate the performance of the tensor decomposition methods demonstrated in the previous case studies, the following sections turn to a systematic assessment using synthetic companion injection and statistical analysis.

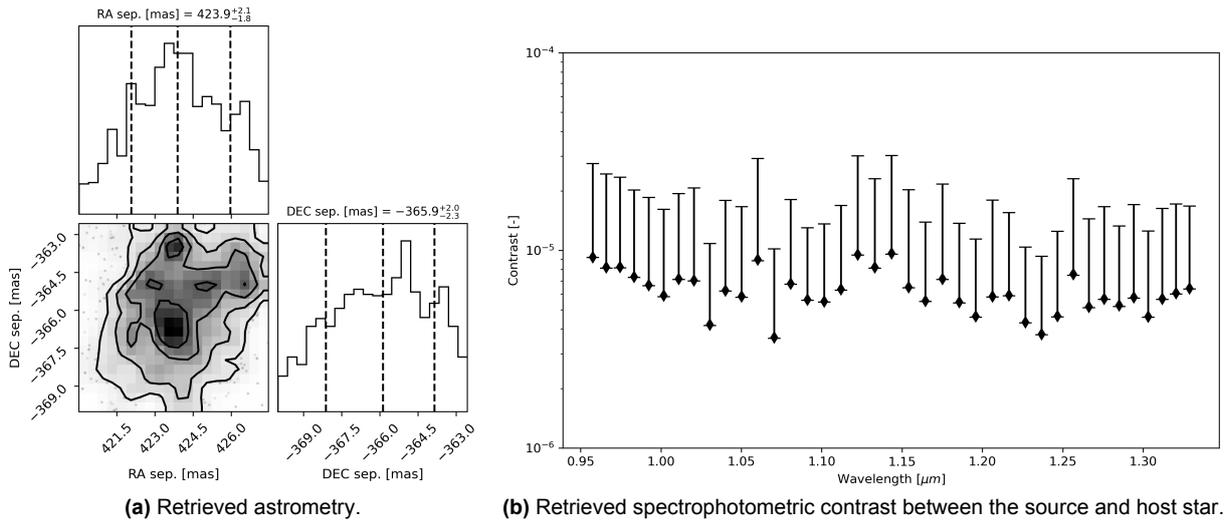


Figure 5.9: Joint fit of the position and spectrum of the source in the HD 108767 B system for the 2016 observations.

5.3. Assessment

To rigorously evaluate tensor decomposition methods, this work employs a multi-faceted approach centered on injecting synthetic companions into real data cubes. This technique allows for controlled testing of each method's detection capabilities under realistic conditions, while avoiding the challenges of fully simulating complex observational processes.

5.3.1. Verification

To rigorously assess the performance of tensor decomposition methods, a combination of approaches is employed. Primarily, synthetic companions are injected into real data cubes which are then post-processed to test whether a method recovers the known sources. This method is preferred over producing full synthetic data cubes, which is extremely challenging due to the complexity of accurately simulating all aspects of the observing process and instrumental effects. (Cantalloube et al., 2022, 2020)

The injection of synthetic companions allows finding the detection limits of the algorithms under realistic conditions. The contrast and position of these injected sources is carefully controlled, so that the resulting contrast curves (see the example in Section 5.2.4) and operating characteristics curves (see Section 5.3.2) accurately estimate the sensitivity of each method.

It's important to note that the development of the algorithms was carried out 'blind' to the presence of these injected sources, ensuring an unbiased assessment of their performance. Additionally, the successful detection of known companions using the new methods serves as a form of verification, demonstrating that the algorithms can recover real astronomical signals. Detecting confirmed exoplanets such as HR 8799 b, c, d and e as well as β Pictoris b is generally used as a test for new methods and the most robust way to verify their basic operation.

However, it's crucial to acknowledge the limitations of the verification process. The lack of absolute ground truth in real astronomical data makes complete validation challenging. The relative performance between methods can be assessed only for a specific data set. Comparing this assessment to theoretical expectations allows the limited generalization of the performance estimates, however, potential unknown systematics or biases in this approach cannot be fully eliminated. (Cantalloube et al., 2022)

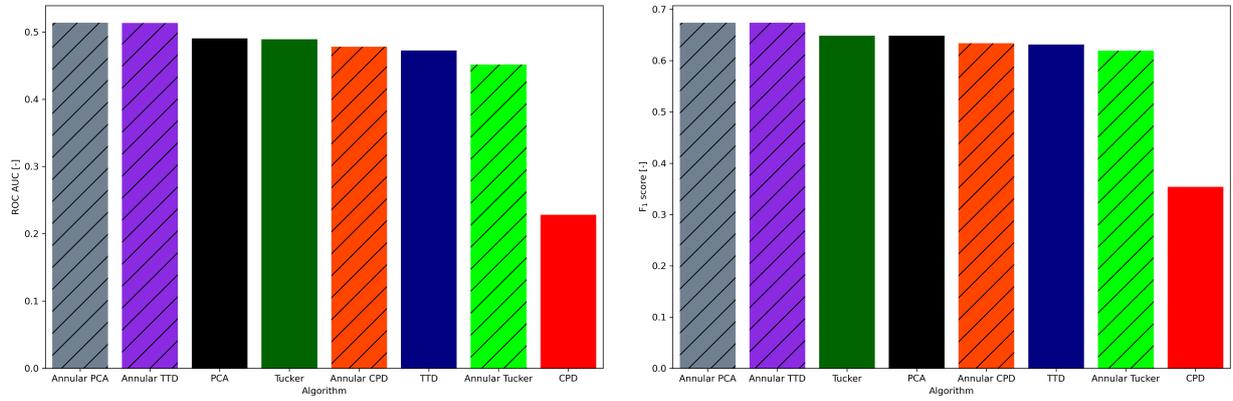
5.3.2. Receiver Operating Characteristics

The algorithm assessment employs methods developed for evaluating binary classifiers (Tharwat, 2020) and incorporates experiences from the Exoplanet Imaging Data Challenge (EIDC, Cantalloube et al., 2022, 2020). Accordingly, detection limits will be determined using the true positive and false positive fractions, the F_1 score, the receiver operating characteristics (ROC) curve, and the ROC area-under-curve (AUC). (Tharwat, 2020) These metrics are chosen to ensure a comprehensive analysis. Nevertheless, considering the complexity and imbalanced nature of the datasets typical in exoplanet detection, the

following assessment cannot be comprehensive.

The evaluation of the most promising tensor decomposition methods utilizes several IFS observations selected from the SHINE survey to ensure consistent preprocessing. Each observation has been further processed to remove bad frames and pixels. It is assumed that the star was accurately centered by the HC-DC pipeline in the remaining frames. The systems within the sample neither host candidates nor background sources, and the observations were manually inspected to verify this condition. The sample encompasses observations with field rotations ranging from very small (4°) to large (145°) under varying observational conditions. IRDIS observations are excluded from this evaluation because they cover only two spectral bands, leading to a smaller expected performance differential between PCA and tensor decomposition methods compared to IFS data. Furthermore, to address the same concern in the temporal dimension, the minimum number of temporal slices in the sample has been set to 40. Observations with more than 200 temporal slices were excluded to reduce computational demands. The selected sample is detailed in Table B.2 and includes B, A, F, and K type stars located at distances ranging from 30 to 170 pc.

Using opposite parallactic angles, synthetic point sources with both constant and chromatically varying spectra are injected into the observations at contrasts ranging from 10^{-3} to 10^{-7} . These sources are modeled using the non-coronagraphic PSF of the corresponding observations. The flux of the injected sources depends on their separation from the star and the noise in an annulus at that separation. This procedure leverages the forward modeling module of VIP to accurately inject synthetic sources, with a total of 264 synthetic companions injected.



(a) Comparison of the ROC AUC for the best-performing algorithms, demonstrating their overall effectiveness in distinguishing between true positive and false positive detections.

(b) Best F_1 scores achieved by the algorithms.

Figure 5.10: Evaluation metrics for the ROC AUC and F_1 comparisons across various tensor decomposition methods and PCA.

To assess the performance of the algorithms as binary classifiers, each algorithm processes the observations and the STIM map is utilized to construct a detection map. Originally developed for ADI sequences, the STIM map is applied sequentially to create a detection map for each channel; these are then mean combined and normalized to produce the final detection map. Contrary to the implementation in VIP, which normalizes the STIM map using the inverse STIM map in each channel, each pixel $d^{x,y}$ is normalized by the maximum STIM value in its spectrum $\mathbf{d}_\lambda^{x,y}$, as illustrated in Equation 5.2. This method results in a more robust detection map by penalizing pixels with outliers in the inverse STIM map more severely, without affecting true sources.

$$d^{x,y} = \frac{1}{N_\lambda \max_\lambda(\mathbf{d}_\lambda^{x,y})} \sum_{\lambda=1}^{N_\lambda} d_\lambda^{x,y} \quad (5.2)$$

Based on this detection map, sources detected at a specific threshold are classified as either true positives, false positives, or false negatives. For each algorithm-parameter combination, the ROC AUC and F_1 score (Equation 5.3) are then calculated. While PCA has a single parameter, N_{comp} , tensor decompositions allow for a broader range of parameters to be adjusted. Notably, these include the ranks, as well as the

matricization or tensorization modes and parameters specific to the computation of the decomposition, such as tolerance and initialization for higher-order orthogonal iteration. However, some areas of this parameter space are restricted or inaccessible due to their high computational resource demands.

$$\frac{\text{TPR}}{\text{TPR} + \frac{1}{2}(\text{FPR} + \text{FNR})} \quad (5.3)$$

All annular versions of the algorithms divide the observations into concentric annuli, each with a width of 10 pixels. Each annulus is individually transformed from Cartesian to polar coordinates. While the number of pixels along the radial axis remains constant for each annulus, the number of pixels along the azimuthal axis increases with separation. This transformation significantly reduces the memory requirements for the processed annuli.

Figure 5.11 presents the ROC of the best-performing algorithms, namely those with the largest ROC AUC and the highest true positive rate (TPR) at a false positive rate (FPR) of 0, for each algorithm class. PCA and all tensor decomposition methods, except for Tucker, greatly benefit from an annular method for modeling the PSF.

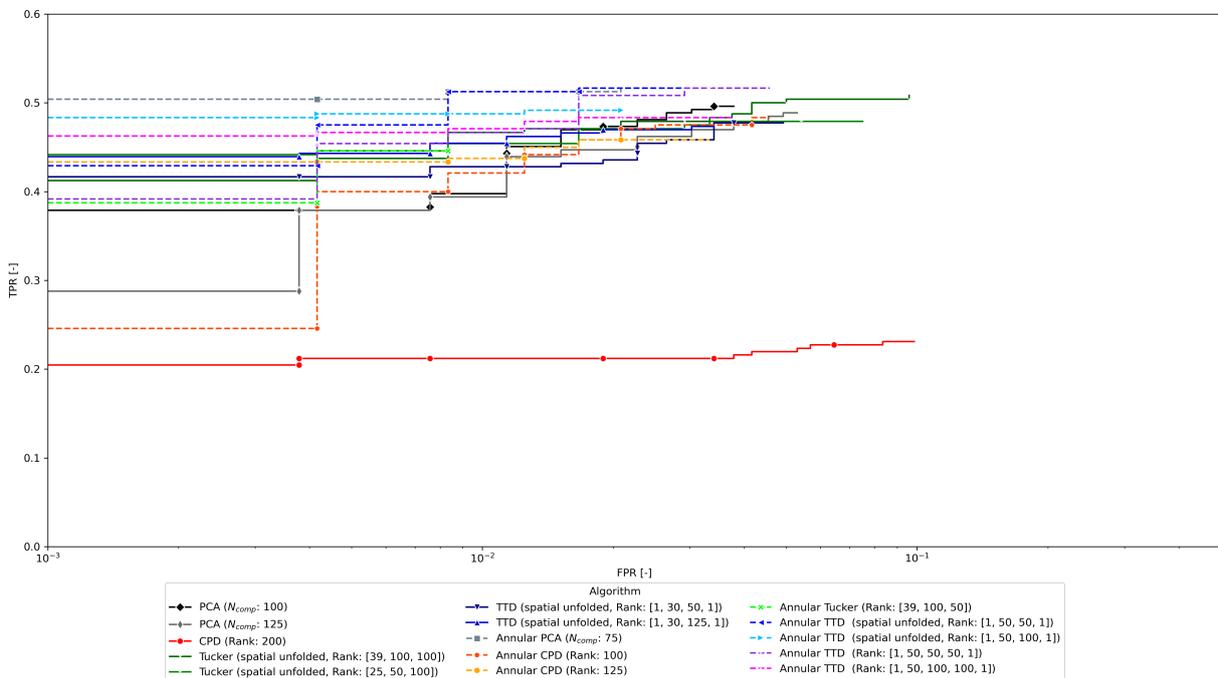
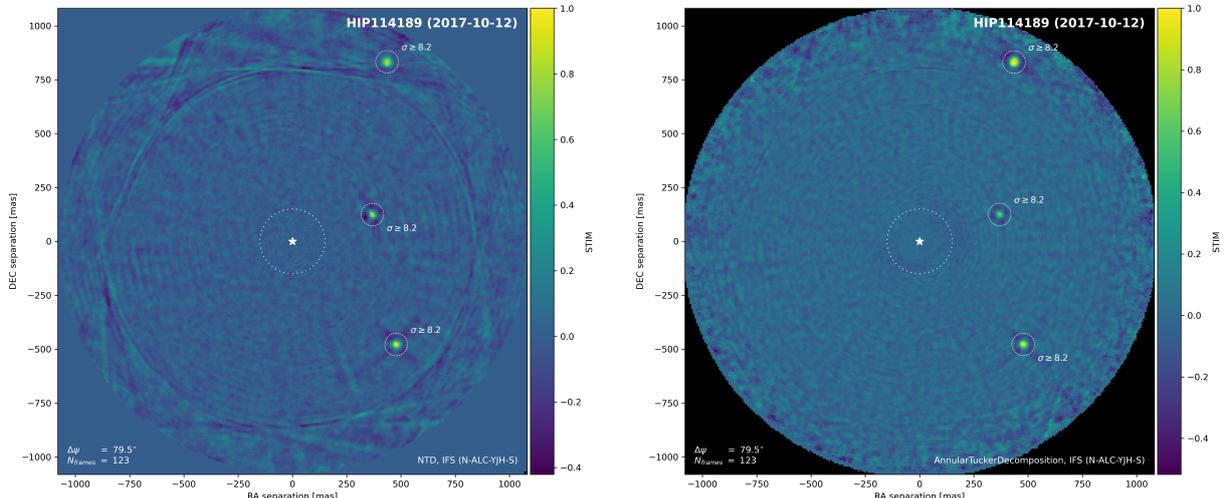


Figure 5.11: The ROC curves of the top-performing algorithms for each class.

The full-frame CPD performs significantly worse compared to other tensor decomposition methods in terms of ROC, as it is the only method that does not utilize spatial unfolding, as explained in Section 5.2.1. The annular version of the CPD shows improved performance, achieving a ROC AUC and F_1 score more similar to those of the other methods. This improvement underscores the importance of aligning the physically meaningful dimensions with the modes of the decompositions to more effectively model the PSF. This aspect is particularly crucial for the CPD since its components do not interact, which limits its expressive power. Comparisons of the ROC curves from the full CPD parameter sweep, Figure D.2a and Figure D.2b, further demonstrate that while the full-frame CPD does not accurately model the PSF even with $R = 200$, the annular CPD already achieves good performance with $R = 20$.

The Tucker decomposition is the only algorithm for which the annular version does not outperform the full-frame version in terms of ROC AUC and F_1 score in this experiment. This might be due to the fact that the performance of the Tucker decomposition exhibits much larger variance compared to the CPD and sequential decompositions like TTD and TRD. As shown in Figure D.3, the decomposition is especially sensitive to the rank of the spatial dimension. Since the number of pixels in the azimuthal axis changes

for each annulus, the completeness fraction also changes, assuming a constant spatial rank. Although a low-rank PSF would not typically be problematic, the spatial frequency of speckles changes with radius, so that a constant rank does not necessarily model only low-rank components at small and large separations. Nevertheless, as can be seen from [Figure 5.12](#) as well as [Figure 5.1c](#) and [Figure 5.2c](#), it is possible to select a spatial rank $R_{(\rho,\phi)}$ that is appropriate for each annulus. Conversely, being able to easily select a suitable number of ranks directly for each dimension in the Cartesian case makes it easier to find a good parameterization compared to TTD and TRD, with the added benefit of potentially finding an over-complete PSF model in one of the dimensions, as is the case with CPD.



(a) STIM map of HR 8799 using the full-frame Tucker decomposition.

(b) STIM map of HR 8799 using the annular Tucker decomposition, showing the change in detection performance when using annuli to model the PSF.

Figure 5.12: Comparative STIM maps for HR 8799 using the Tucker decomposition.

The annular TTD is the best performing tensor decomposition in this survey and performs very closely to the annular PCA. However, the Cartesian version of the TTD slightly underperforms compared to PCA and Tucker. As previously mentioned, this may be due to the parameter sweep for the TTD being too sparse compared to that for the Tucker decomposition. When comparing the spatially unfolded and order-4 performance of the annular TTD—for instance, $R_2^{(I_\rho \times I_\phi)} = R_2^{I_\rho \times I_\phi} = R_3^{I_\rho \times I_\phi} = 50$ or 100 as shown in [Figure D.4a](#) and [Figure D.4b](#), respectively. This again demonstrates that unfolding consistently enhances the performance for annular decompositions.

While these assessment methods provide a robust framework for comparing the performance of different algorithms, it's important to acknowledge their limitations in the context of exoplanet imaging. The complexity and variability of real astronomical data mean that performance on injected sources may not always perfectly reflect performance on real companions. Furthermore, the assessment metrics used here may not capture all aspects of algorithm performance that are relevant to astronomers, such as the ability to preserve extended structures.

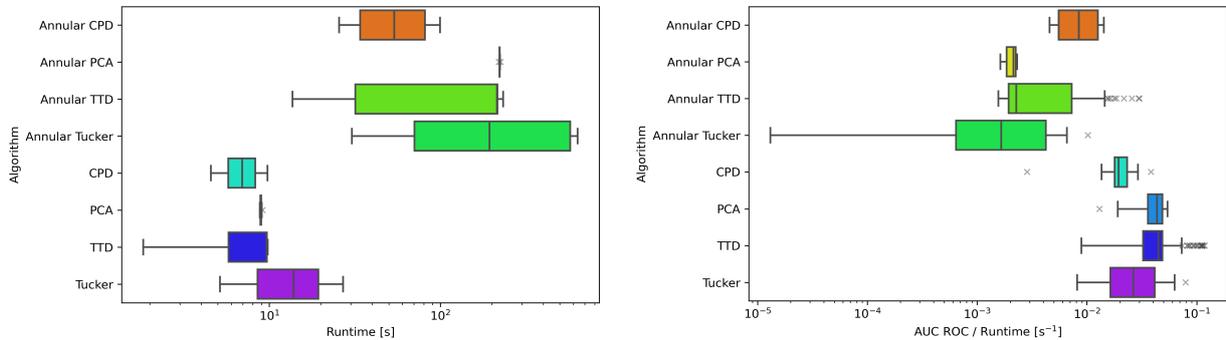
5.3.3. Runtime Analysis

One of the primary motivations for employing tensor decompositions in recent years has been their potential speed advantage over traditional matrix-SVD methods applied to flattened tensors. To quantify this advantage in the context of exoplanet imaging, a comprehensive runtime analysis of the algorithms is presented in [Chapter 3](#).

To compare the speed of different algorithms this work measures their runtime. [Figure 5.13a](#) shows the distribution of runtime for the algorithms evaluated in [Section 5.3](#). The runtime is measured as the wall-time required to complete lines 6 to 12 of [algorithm 1](#); this measurement excludes setup, as well as data reading and transfer prior to the actual algorithm execution. The annular versions process each annulus sequentially, and the conversions detailed in [algorithm 2](#) and [algorithm 3](#) are not included in the runtime calculation. This sequential processing of the residuals in the annuli means that the runtime of

annular methods is significantly longer than that of full-frame methods.

PCA shows a very tight runtime distribution because the SVD and projection are performed over the entire unfolded tensor, varying only with the number of components used as the basis. In contrast to PCA, tensor decomposition methods showed greater variability in runtime. This is largely due to the influence of rank choice on the required number of operations, leading to significant variations in computation time.



(a) Distribution of runtime for different tensor decomposition and PCA methods, excluding setup and data transfer times. This highlights the computational efficiency of each group of methods when processing the same direct imaging data.

(b) Ratio of ROC AUC to runtime for each algorithm group, demonstrating the balance between computational speed and performance effectiveness.

Figure 5.13: Comparative analysis of runtime efficiency and performance across various tensor decomposition methods and PCA for HCI.

Naturally, fast computation is only advantageous if the algorithms perform comparably well in the context of direct imaging. Figure 5.13b shows the ratio of ROC AUC to runtime for each algorithm group. Although the ROC AUC varies significantly among the full-frame methods, they cluster tightly in terms of ROC AUC per second, with the TTD displaying notable outliers.

When considering the performance of the algorithms, the distinction between full-frame and annular methods diminishes. However, the rapid computation of the TTD, even for annuli, results in outliers that perform similarly to full-frame methods. Given that the computation of the CPD (for order-4 tensors) is relatively quick the transform to polar coordinates brings its performance close to that of PCA and TTD, and it compares favorably with other annular methods.

These runtime analyses highlight the trade-offs between computational speed and algorithm performance in exoplanet imaging applications. While tensor decomposition methods offer potential speed advantages, the choice of method and parameterization can significantly impact both runtime and detection performance. The benefit of tensor methods likely grows quickly with the size of observations. Optimization of tensor decomposition algorithms, particularly for annular processing, narrows the gap with full-frame methods in terms of computational efficiency. This is because the scaling and rotation operation are cheaper in an annular frame since they involve just scaling in one axis and translation, respectively. Additionally, implementing parallel processing techniques for annular methods will improve performance without sacrificing the benefits of localized PSF modeling.

5.4. Chapter Summary & Contributions

This chapter applies tensor decomposition methods to post-process observations of exoplanets and circumstellar disks obtained by angular-spectral differential imaging. It proves that tensor decompositions are a competitive alternative to matrix-based PCA for high-contrast imaging. This is achieved by demonstrating their efficacy for reducing the data from several well-known systems, detecting known exoplanets, disks and features, as well as performing a systematic assessment of the new methods using synthetic data. Both approaches also serve to verify the correct implementation of the new methods.

The chapter demonstrates that the new methods are robust to different instruments like the SPHERE/IFS and SPHERE/IRDIS, a wide range of observing conditions, and targets. Additionally, it indicates that the proposed methods might be comparatively well suited to difficult target systems which host disks as well as fast-moving features or embedded exoplanets. This is achieved by a qualitative analysis of the residual images obtained by applying selected tensor methods and matrix-PCA to IFS and IRDIS observations of the HR 8799, β Pictoris and AU Microscopii system.

The chapter then progresses to a quantitative assessment of algorithm performance using synthetic companion injection into real observations. A representative sample of systems without known companions or disks is selected for this task. Detection capabilities are evaluated across a range of contrasts and separations, with a focus on the receiver operating characteristic curves and F_1 score to statistically compare the various methods. The results show that tensor decomposition methods, particularly annular versions are competitive with matrix-PCA, with some indication of superior performance at small separations.

A runtime analysis is included to assess the practical computational efficiency of tensor decomposition approaches relative to matrix-based approaches. While tensor methods show potential speed advantages, the analysis reveals that performance can vary significantly based on method choice and parameterization.

Lastly, this chapter details two new discoveries made using the new tensor decomposition methods. Firstly, a new sub-stellar companion candidate is detected in the HD 108767 B system using IFS observations from 2016 and 2018 at a significance of 5.3σ and 4.4σ respectively. The candidate spectrum and astrometry is jointly determined using a tensor decomposition based reduction together with a nested sampling approach. This demonstrates that the new tensor methods are suitable for exoplanet atmosphere characterization and other down-stream science tasks. Secondly, a new south-east debris disk feature is detected in 2019 IFS observations of AU Microscopii at a significance of 5.6σ . This is likely a newly formed feature indicating a collisional cascade and not one of the previously known fast-moving features. The previously known disk features were also detected with tensor decomposition methods with a substantially improved significance relative to matrix-PCA post-processing. Together, this clearly demonstrates that tensor decomposition methods are not only able to confirm known exoplanets, disks and disk features but also make new discoveries.

These results collectively establish tensor decompositions as a powerful and versatile method for direct exoplanet imaging.

Contributions 5.1: Results

This chapter makes several key contributions to the field of high-contrast imaging:

- It provides a comprehensive demonstration of tensor decomposition methods applied to real astronomical datasets:
 - Validates the practical utility of these methods for exoplanet and circumstellar disk detection in well-known systems like HR 8799 and β Pictoris.
 - Demonstrates the effectiveness of tensor methods in identifying and characterizing circumstellar disk features, particularly in the AU Microscopii system.
 - Shows the robustness of these methods across different instruments like SPHERE/IFS and SPHERE/IRDIS.
 - Shows the robustness of these methods to several observation conditions.
- It offers a rigorous, quantitative comparison of tensor decomposition methods with traditional PCA:
 - Presents a systematic assessment using synthetic companion injection into real observational data.
 - Evaluates detection capabilities across a range of contrasts and separations using ROC curves and f_1 score.
 - Highlights the competitive performance of tensor methods, particularly annular versions, compared to matrix-PCA.
 - Indicates potential superior performance of tensor methods at small separations.
- It presents two new discoveries:
 - A companion candidate in the HD 108767 B system, detected in multiple observation epochs with high significance.
 - A new disk feature in the AU Microscopii system, indicating ongoing dynamical processes in the disk.
- It demonstrates the applicability of tensor methods to downstream science tasks:
 - Shows how tensor decomposition can be used for joint determination of companion spectra and astrometry.
 - Illustrates the potential of these methods for exoplanet atmosphere characterization.
- It discusses the practical computational aspects of implementing tensor methods:
 - Provides a runtime analysis comparing the efficiency of tensor methods to matrix-based approaches.
 - Offers insights into the trade-offs between different tensor decomposition techniques in terms of computational demands and performance.
 - Highlights considerations when integrating the new methods into existing exoplanet imaging pipelines.

Through this comprehensive analysis, this chapter demonstrates the capabilities, limitations, and potential advantages of tensor decomposition methods in the context of high-contrast imaging for exoplanet detection and characterization. It lays the groundwork for further development and application of these techniques in future exoplanet research.

The findings presented in this chapter raise important questions about the theoretical underpinnings of tensor decomposition methods in the context of high-contrast imaging, their practical implications, and their potential impact on future exoplanet research. The next chapter provides a detailed discussion of above results, exploring their broader implications and identifying areas for future investigation.

Discussion

This chapter presents a critical analysis of tensor decomposition methods applied to HCI for exoplanet detection and characterization. The performance of tensor-based approaches relative to matrix-PCA techniques is examined. Fundamental noise characteristics of HCI data and their impact on tensor decompositions are explored, considering the implications of various signal-to-noise regimes.

The theoretical foundations for applying tensor decompositions to HCI are investigated, including how different decomposition methods reflect varying assumptions about data structure. Implications of a potential new discovery in the HD 108767 B system are discussed, demonstrating the capability of tensor methods to detect challenging objects.

Broader impacts on computational approaches in astronomy are considered, with potential applications in other fields. Limitations and challenges in applying tensor decomposition methods to HCI are critically assessed, including issues of mode selection, computational efficiency, and extended source handling. Through this analysis, the results are contextualized within exoplanet detection and characterization, and directions for future research in tensor-based HCI methods are suggested.

6.1. Detection Performance

The performance of tensor decompositions relative to PCA varies across different observational scenarios, with particularly interesting results at small separations. This variability can be attributed to the inherent structure of the data and the nature of the noise present in ASDI observations.

In terms of the collapsed residuals, ROC AUC and F_1 -score, tensor decomposition methods are close to PCA, especially when evaluated in annuli. Especially annular TTD performs very similar to annular PCA. This also extends to its limits in terms of contrast and separation where there is no consistent significant difference between the best tensor decompositions and PCA. However, this work finds that, for some observations, tensor decompositions can outperform PCA in their contrast limit at small separations (see e.g. [Figure 5.7b](#) and [Figure 6.1](#)). This is likely related to the quality of the observations where a more slowly varying speckle field is sufficiently low-rank and cross-coupling along modes is strong so that tensor decompositions are able to capture structured components that PCA is unable to.

Considering that a) neither PCA nor the tensor decompositions in this work have an explicit noise model (see [Figure 6.2](#)), b) both are implicitly assuming normally distributed tensor elements in their likelihood model ([Kolda et al., 2009, 2020](#)), and c) the moderate-to-low SNR of the off-axis sources ([Zhang et al., 2018b](#)), the residual speckles in the PSF might not be sufficiently low-rank for tensor decompositions to offer a direct advantage over matrix-PCA.

In the context of HCI, several methods that directly estimate the noise and sparse components, like the Local Low-rank plus Sparse plus Gaussian-noise decomposition (LLSG, [Gonzalez et al., 2016](#)) and the low-rank plus sparse trajectory method (LRPT, [Vary et al., 2023](#)) have been developed. Most of these make use of the matrix Robust PCA (RPCA) methods or the (SS) Go Decomposition ([Guo et al., 2018; Zhou et al., 2011](#)) which have been generalized to higher-order tensors. Nevertheless, while RTD methods have been widely used in related fields like hyper-spectral imaging ([Wang et al., 2023](#)), they are not implemented in the well-maintained tensor decomposition libraries (see [Table G.1](#)). The implementations that are public typically consider only the order-3 case.

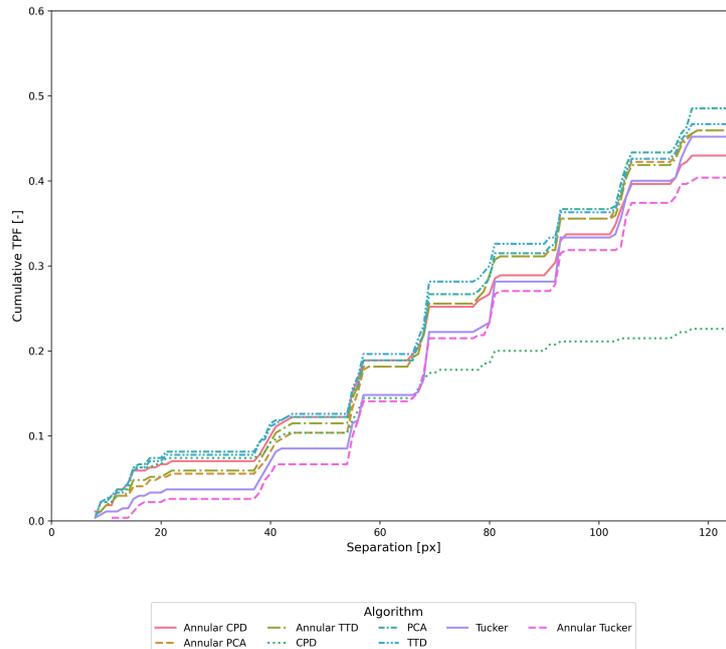


Figure 6.1: True positive rate of best performing algorithms from each group for sources with separations smaller than a specific radius.

6.2. Noise Characteristics

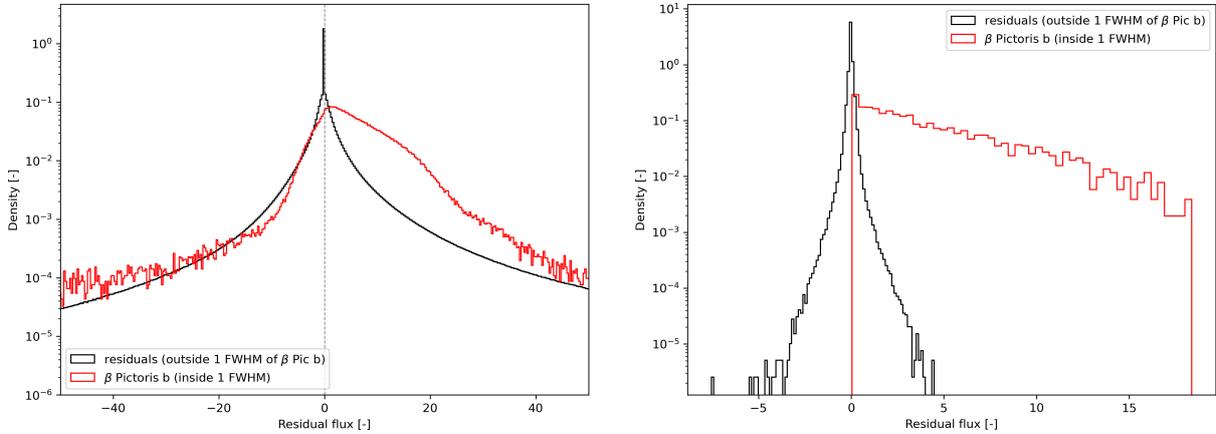
Similarly, it is well known that the noise in HCI is Rician distributed (Mawet et al., 2014) and tensor models for data corrupted by Laplacian or Rician noise have been explored. In the context of blind source separation, and found to be preferable over normal noise models, especially in the case of low SNR (Chatzichristos et al., 2019). However, the existing implementations (see e.g. Psarras et al., 2022) are currently unsuitable for use in HCI.

Tensor decompositions behave differently to PCA as the SNR of the signal decreases. This is particularly relevant to HCI observations, where the SNR can vary significantly across different parts of the tensor. The higher-order singular value decomposition is specifically well understood, and the extraction of hidden low-rank structure undergo significant phase transitions between the strong, moderate and low SNR cases. (Zhang et al., 2018b) The impact of noise in HCI observations on low-rank tensor approximations of the coronagraphic PSF and speckle pattern can be further understood through the lens of these phase transitions:

- **Strong SNR regime:** In this regime, efficient algorithms like higher-order singular value decomposition and higher-order orthogonal iteration can successfully recover the low-rank structure. For HCI, this suggests that tensor decomposition methods likely outperform PCA in regions of the image where the signal from the speckle pattern is sufficiently strong relative to the noise.
- **Moderate SNR regime:** In this phase, recovery of the low-rank structure is possible but may require computationally intensive methods such as Maximum Likelihood Estimation. This regime might explain the computational challenges observed when applying tensor decomposition methods to HCI data, particularly in regions where the residual speckle pattern varies strongly along the spectral or temporal modes.
- **Weak SNR regime:** In this phase, reliable recovery of the low-rank structure is not possible. This could correspond to regions in HCI data where the speckle pattern is too weak relative to the noise for any method, including both tensor decomposition and PCA, to effectively separate the signal from the noise.

The observation that tensor methods sometimes outperform PCA at small separations could be explained if these cases fall into the strong SNR regime for tensors but not for matrices as the low-rank structure is disrupted by the flattening operation. Conversely, the similarity in performance between tensor methods and PCA in many cases might be due to both approaches falling into the same SNR regime.

The phase transition framework also sheds light on the computational challenges associated with tensor decomposition in HCI (Section 5.3.3). The existence of a moderate SNR regime where polynomial-time algorithms may not be able to reliably recover the low-rank structure suggests that the computational difficulties encountered in applying tensor methods to HCI are not merely implementation issues, but may be fundamental to the problem.



(a) Distribution of the full order-4 residual tensor after performing TTD (full frame).

(b) Distribution of the median reduced order-3 residual spectral cube after performing TTD (full frame).

Figure 6.2: Distribution of residuals for β Pictoris reduced with the TTD (full frame), showing the difference between pixels at the location of β Pictoris b and residual coronagraphic PSF flux.

Figure 6.3 presents the weight distributions of two CPD models with different ranks, extracted from observations using the Gemini Planet Imager. These graphs plot the magnitude of the weights associated with each rank-1 tensor component, which contribute to the PSF and speckle pattern model. Such plots are commonly referred to as scree plots in the context of matrix-PCA, which visualize the contribution of each component to the model, helping in determining the number of components to include for optimal data representation.

Figure 6.3a shows an exponential decrease in the weight magnitudes over all components. The absence of any inflection points indicates that the model is likely not complex enough to explain the full low-rank structure. In contrast, Figure 6.3b exhibits noticeable inflection points around the 20th and $\sim 400^{\text{th}}$ component. This inflection suggests a change in the significance of the components, where those above 20 contribute progressively less to the model and those above ~ 400 capture mostly noise. This behavior aligns with observations from similar analyses using matrix-based PCA, where after a certain point, additional components contribute minimally to improving model accuracy and increasingly remove exoplanet signal. (Soummer et al., 2012)

The differences in the factor weights of Figure 6.3 highlight the trade-offs between PSF model completeness and subtracting exoplanet signal which exist in matrix-PCA methods are also present in tensor methods. In HCI, where data often suffer from low signal-to-noise ratios, the ability to identify and restrict model complexity to only significant components can drastically affect the quality of the post-processed residuals.

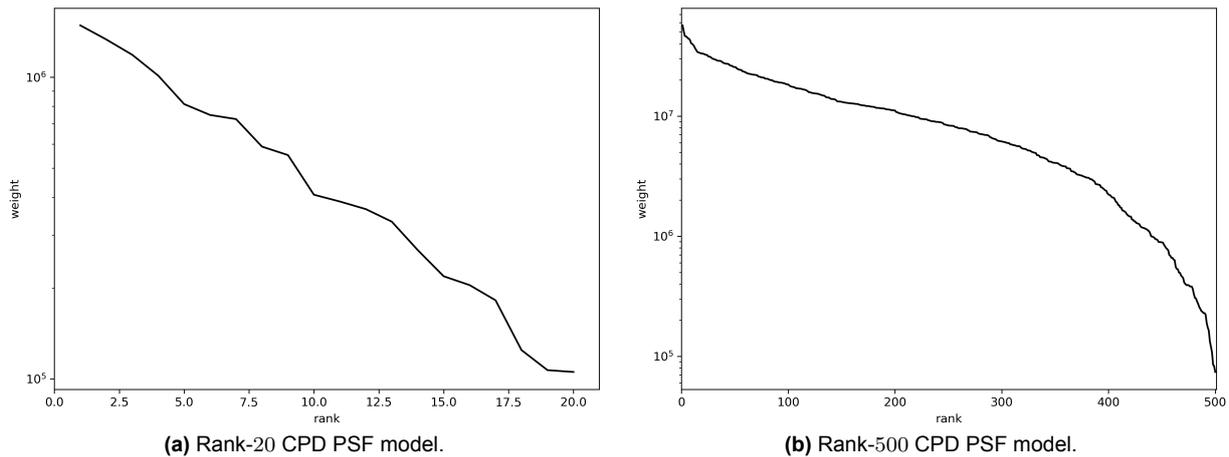


Figure 6.3: Weights λ of the CPD PSF models from GPI observation in Figure 3.3. Here, the x-axis enumerates the component rank-1 tensors (r_i), not the rank of the decomposition (R). The distribution of the weights resembles the total variance as a function of N_{comp} in residual images reduced with matrix-PCA (Soummer et al., 2012). Further, the distribution of the weights for the rank-500 CPD model has an inflection point around the 20th component.

6.3. Theoretical Considerations

The application of tensor decompositions to exoplanet imaging is underpinned by several key theoretical considerations that link these mathematical constructs to the physical processes involved in HCI.

At its core, the use of tensor decompositions in HCI is based on the assumption that the observed data can be well-approximated by a low-rank tensor structure. This assumption is justified by the physical nature of the imaging process: the observed speckle pattern is primarily a result of the interaction between the incoming wavefront and the telescope optics, which can be described by a relatively small number of modes. In the context of ASDI observations, these modes correspond to the spatial, spectral, and temporal dimensions of the data.

The choice of tensor decomposition method (e.g., CPD, Tucker, TTD) reflects different assumptions about the underlying data structure. For instance:

- CPD assumes that the data can be represented as a sum of rank-one tensors. In HCI, these rank-one tensors might correspond to individual speckles or distinct instrumental effects. For instance, a rank-one tensor could represent a speckle with a specific spatial pattern, spectral signature, and temporal evolution. The limitation of CPD in modeling complex speckle patterns, as observed in this work (e.g. Figure 5.1b), suggests that the speckle field in HCI data often requires more complex interactions between modes than CPD can provide.
- Tucker decomposition allows for more complex interactions between modes, potentially capturing more intricate speckle patterns or instrumental effects. Its similarity to PCA in several key ways makes it the decomposition for which the current understanding of ASDI post-processing is most transferable. This is evident from the systematic assessment (Section 5.3) where Tucker decomposition often performed similarly to PCA (e.g. Figure 5.11), especially when spatial unfolding was applied.
- TTD provides a way to model hierarchical correlations in the data, which could be particularly relevant for capturing the multi-scale nature of atmospheric turbulence and its effects on the speckle pattern. This work showed that TTD often outperformed other methods, especially in annular processing (e.g. Figure 5.11), suggesting that this hierarchical structure effectively captures the complex spatio-temporal-spectral relationships in HCI data.

The theoretical relationship between tensor ranks and the complexity of the observed phenomena warrants further investigation. This work indicates that the optimal rank varies depending on the separation from the star (e.g. Figure 6.3), suggesting a complex relationship between the rank of the decomposition and the physical properties of the speckle field. Understanding this relationship could provide insights into the fundamental limits of exoplanet detection and characterization using HCI. Furthermore, it was expected that the tensor methods can significantly exploit local structure in the spatial modes (i.e. a speckle extends

over several pixels) of observations pre-processed by subtracting the median spectral cube. Instead, the assessment shows that methods which do not account for a low-rank spatial structure by using spatial unfolding. The presence, or absence, of residual low-rank structure should be investigated by repeated rank-1 PSF approximation and subtraction.

The impact of noise on tensor decompositions is a crucial theoretical consideration. While most current implementations of tensor decompositions assume Gaussian noise, the actual noise in HCI data follows a Rician distribution, as noted by [Mawet et al., 2014](#). This work, in [Figure 6.2](#), shows that the residuals after tensor decomposition are not perfectly Gaussian, indicating that developing tensor decomposition methods that explicitly account for Rician noise could significantly improve their performance in HCI applications.

The connection between tensor decompositions and other mathematical frameworks used in HCI, such as the KLIP algorithm or the LOCI method, is an area ripe for theoretical exploration. For instance, the Tucker decomposition can be seen as a generalization of PCA to higher-order tensors, suggesting potential ways to extend existing PCA-based methods like KLIP to fully exploit the tensor structure of ASDI data.

Lastly, the theoretical foundations of tensor decompositions provide a framework for understanding the information content and fundamental limits of HCI data. By analyzing the tensor structure of the data, it may be possible to derive new bounds on the achievable contrast or detection limits. For example, the distribution of weights in CPD models ([Figure 6.3](#)) suggests an inflection point that could be related to the theoretical limit of extractable information from the data cube.

6.4. Implications of the HD 108767 B Candidate

The potential detection of a new source in the HD 108767 B system using tensor decomposition methods has significant implications for both the specific system and the broader field of exoplanet detection. This detection, if confirmed, would not only add to the understanding of the HD 108767 system but also serve as a powerful validation of tensor decomposition techniques in HCI.

The characteristics of the detected source - its apparent motion, flat spectrum, and visibility in IFS but not IRDIS observations - present an intriguing puzzle. These properties do not clearly align with expectations for either a typical exoplanet, a background star or disk feature.

To confirm and characterize this potential new source, follow-up studies would be necessary:

- Additional high-contrast imaging observations to confirm the source's presence and track its motion over time. It's apparent motion needs to be better constrained to differentiate it from a background object and confirm that it is co-moving with HD 108767 B.
- In-depth spectrophotometric forward model combining IFS and IRDIS observations from 2016 and potentially from 2018. Comparison of the apparent spectrum and the stellar spectrum is necessary to differentiate it from a disk feature in reflected star light. The presence of a disk around HD 108767 B would be a significant discovery in itself. Furthermore, the spectrum of the source can also be used to distinguish it from a background object. Lastly, the low-resolution spectrum can be used to characterize the candidates atmosphere. ([Pueyo, 2016](#))
- Detailed modeling of the system to understand how this object, if real, fits into the broader context of the HD 108767 (B) system.
- If the source is neither a sub-stellar companion, disk feature nor a background object, alternative explanations for the apparent source, like instrument systematic, need to be examined.

The confirmation of this source would demonstrate the capability of tensor decomposition methods to detect objects that might be missed by other techniques. This could potentially open up a new parameter space for exoplanet and brown dwarf detection, particularly for objects with unusual properties or in challenging observational regimes.

6.5. Broader Impact on Computational Methods in Astronomy

The development and application of tensor decomposition methods in exoplanet imaging have the potential to significantly influence computational approaches in other areas of astronomy. As astronomical datasets continue to grow in size and complexity, the need for efficient, multi-dimensional data analysis techniques becomes increasingly critical.

The success of tensor decompositions in HCI demonstrates their potential for handling high-dimensional, noisy data with complex correlations. This capability could be particularly valuable in fields such as:

- Radio astronomy: Analysis of large-scale interferometric data, where tensor methods could help in source separation and radio frequency interference mitigation. [Sørensen et al., 2023](#) already have successfully applied the CPD to the source separation problem in radio astronomy using delay-and-sum beam-former data.
- Cosmology: Exploration or visualization of cosmic web structures or analysis of CMB data. Tensor decompositions could be especially useful for efficiently compressing the large high-dimensional data encountered in cosmology. The existing work in tensor models for processing point-cloud data might be especially applicable in this field. ([Panagakis et al., 2021](#))
- Transit-method spectroscopy: Processing of repeat-visit transit light-curves which have with spatial, temporal, and spectral dimensions.

Moreover, the computational challenges addressed in applying tensor decompositions to HCI data - such as dealing with large datasets and developing efficient algorithms for specific tensor structures - are likely to be relevant in these other fields as well. Solutions developed for HCI could potentially be adapted or serve as inspiration for tackling similar challenges in different astronomical contexts.

The interpretability aspect of tensor decompositions could also drive new approaches to data analysis in astronomy. The ability to extract physically meaningful components from complex datasets could lead to new insights across various subfields, potentially revealing underlying patterns or phenomena that are not easily discernible with traditional methods.

6.6. Limitations & Challenges

This work identifies the selection of modes for the decompositions as a primary concern. Especially, in the low-rank case it is essential either to transform the observations to polar coordinates or unfold the spatial dimensions. This aligns the low-rank structure in the observations with the modes of the decompositions, greatly increasing their performance. Furthermore, computing the PSF model in annuli reduces the variance over the domain of the observation (annulus). However, care should be taken to consider the context in which the reductions are used. Unlike point sources, extended sources like disks are much more likely to be removed if observations are transformed to polar coordinates (compared e.g. [Figure 5.6b](#) and [Figure 5.7a](#)). Unfolding the spatial dimension similarly does not suffer from the same problem as the quasi-low-rank structure of extended sources in the spatial modes is not preserved. Still, all methods are susceptible to self-subtraction comparable with PCA when a decomposition rank is chosen that starts to capture non-static components. Enforcing sparsity in the decompositions, especially in cores, can partially counteract this as it encourages the exclusion of non-static components [Figure 5.3](#). ([Shah et al., 2015](#); [Zhang et al., 2018a](#))

$$\text{Var}_\theta(\mathcal{X}) = \left\| \frac{\partial^2 \mathcal{X}_{\lambda, \theta, \varrho, \phi}}{\partial \theta^2} \right\| \quad (6.1)$$

Importantly, this work identifies several limitations when using tensor decomposition methods, while low-rank PSF models can be computed very efficiently, [Section 5.2](#) and [Section 5.3](#) demonstrates that the best performing algorithms are never truly low-rank. This is likely due to speckles, which are especially severe at small separations where the speckle field varies strongly across the temporal mode. For the observations in the synthetic sample, even the low-rank speckle field captured by annular tensor decomposition PSF models is strongly varying (non-smooth, [Equation 6.1](#)) over the temporal mode, with its second partial derivative a factor of $\sim 100 - 1000$ larger than along the other modes. This indicates that, fundamentally, low-rank models which have no other method of accounting for moderately structured noise such as speckles can struggle to fully model the coronagraphic PSF. Requiring the rank of the decomposition to be relatively high, leads to high memory usage, as was observed especially for the CPD.

The developed tensor methods have been implemented in such a way that they are compatible with VIP and can be used as a drop-in when using the framework. The methods are publicly accessible at the [TIDE repository](#).

6.7. Chapter Summary & Contributions

This chapter provides a discussion of the results presented in the previous chapter, exploring the implications, limitations, and broader impacts of applying tensor decomposition methods to high-contrast imaging for exoplanet detection. It begins by summarizing the detection performance of tensor methods compared to matrix-PCA, highlighting their competitive performance and potential advantages, particularly at small separations.

The chapter then explores the noise characteristics of high-contrast imaging data and how they affect tensor decompositions. It discusses the impact of different SNR regimes on the performance of tensor methods and discusses the theoretical considerations underlying the application of these techniques to exoplanet imaging. This includes an examination of how different tensor decomposition methods (CPD, Tucker, TTD) reflect various assumptions about the underlying data structure.

The implications of the potential new discovery in the HD 108767 B system are discussed, emphasizing the need for follow-up studies and the significance of this detection for validating tensor decomposition techniques.

The chapter also explores the broader impact of tensor methods on computational approaches in astronomy, suggesting potential applications in fields such as radio astronomy, cosmology, and transit spectroscopy.

Finally, the chapter addresses the limitations and outstanding challenges of applying tensor decomposition methods to high-contrast imaging, including issues related to mode selection, computational efficiency, and the handling of extended sources.

Contributions 6.1: Discussion

This chapter contributes to the understanding and development of tensor decomposition methods in high-contrast imaging by:

- Analyzing performance of tensor methods vs. matrix-PCA, revealing potential advantages at small separations and for identifying multi-modal low-rank data structures.
- Exploring the impact of different SNR regimes on tensor method performance, linking to phase transitions in low-rank structure recovery.
- Examining how CPD, Tucker, and TTD reflect different assumptions about HCI data structure, relating to speckle patterns and instrumental effects.
- Discussing implications of the HD 108767 B candidate, demonstrating tensor methods' potential to detect challenging objects.
- Proposing applications in radio astronomy, cosmology, and transit spectroscopy, suggesting broader impact on computational methods.
- Identifying key challenges, including mode selection, computational efficiency, and extended sources, guiding future research.

This discussion contextualizes the results within exoplanet detection and characterization, pointing towards future directions in tensor-based HCI methods.

The following chapter will conclude this work and recommend areas for further research.

Conclusion

This work has developed tensor decompositions as robust alternatives to traditional PCA for the modeling of the coronagraphic PSF in the direct imaging of exoplanets. Tensor decompositions generalize matrix factorization methods to higher-order tensors which they break down into their factors. In contrast to matrix factorizations, they preserve the multi-modal structure of observations with angular and spectral diversity which allows them to capture cross-couplings between modes. Moreover, tensor methods offer increased flexibility for modeling the coronagraphic PSF and companion sources by treating each mode individually. Lastly, factors in tensor decompositions are physically meaningful which allows a closer coupling between PSF subtraction, forward modeling and statistical techniques in direct imaging. This work demonstrates that reducing the observations from angular-spectral differential imaging as higher-order tensors is a valid and promising alternative to the current approach.

Several tensor decomposition methods suitable to HCI, including CPD, Tucker, TTD, and TRD, have been developed and assessed, both in full-frame and annular configurations. These methods model the coronagraphic PSF by decomposing the observations into several factors which are found by selectively factoring modes of the tensor using the SVD. A key trade-off when using tensor methods is their interpretability for which the CPD and Tucker decomposition excel and their expressiveness which favours the TTD and TRD. The methods are validated using observations from the SPHERE instrument of the VLT. Assessing the methods as binary classifiers on synthetically injected companions shows that their performance is comparable to PCA, however, they can exhibit superior performance in certain scenarios. The assessment is limited by the large parameter space which needs to be explored and the computational effort to test them on the full diversity of HCI observations. The methods developed in this work are publicly available and compatible with the HCI data processing package VIP.

Research Questions 7.1: Main Result

The main research question in this study is: *Can tensor decomposition methods provide a competitive alternative to matrix-PCA in the modeling of the coronagraphic PSF and quasi-static speckles in observations obtained by ASDI?* This study concludes that tensor decomposition methods can provide a competitive alternative to matrix-PCA in modeling the coronagraphic PSF and quasi-static speckles in ASDI observations. The results show that tensor methods, particularly in their annular configurations, perform comparably to matrix-PCA and in some cases offer potential advantages, especially at small separations.

To answer the main research question, several sub-questions have been answered.

Research Questions 7.2: Secondary Result 1

Modeling Capability: Can tensor decomposition methods model the coronagraphic PSF? The successful application of tensor decomposition methods in recovering synthetically injected exoplanets and known exoplanets confirms that they can model the coronagraphic PSF. Their successful application to systems with luminous debris disks such as β Pictoris and debris disk features such as AU Microscopii show that tensor methods can differentiate between the coronagraphic PSF and extended sources. This indicates that tensor models robustly model the coronagraphic PSF, separating it from both point-like and extended sources.

Research Questions 7.3: Secondary Result 2

Detection Performance: What are the receiver operating characteristic curves of tensor decomposition methods? This study determines the receiver operating characteristic curves for a number of tensor decomposition methods and parametrizations thereof.

Research Questions 7.4: Secondary Result 3

Detection Limits: What are the detection limits of tensor decomposition methods in terms of stellar-to-sub-stellar companion contrast and angular separation? Tensor methods demonstrate competitive detection limits in terms of stellar-to-sub-stellar companion contrast and angular separation, with successful detections ranging down to contrasts of 10^{-6} at a separation of 500 mas in this studies' sample. In some cases, particularly at small separations, tensor methods show potential for improved contrast limits compared to PCA.

Research Questions 7.5: Secondary Result 4

Performance Comparison: How do tensor decomposition methods compare to matrix-PCA? Overall, tensor decomposition methods perform similarly well compared to matrix-PCA, measured by ROC AUC and F_1 scores. They offer increased flexibility in modeling spatio-temporal-spectral components and PSF features. Their factors can be physically meaningful, so that a direct inspection and interpretation of decomposed PSF model is possible. They scale well with increasing dataset size compared to matrix-PCA in terms of parameter count and operations.

Looking forward, the tensor decompositions are promising a new class of PSF subtraction methods that naturally extend existing techniques in HCI. In addition to their suitability for forward modeling methods and causal noise models, they are a key enabler for deep learning on HCI datasets, especially in the angular-spectral differential imaging context.

In summary, this thesis establishes tensor decompositions as a powerful tool for advancing the direct imaging of exoplanets. By offering a method that aligns closely with the physical processes and variability of the data they model, tensor decompositions provide a promising avenue for future work in astronomical imaging and analysis.

Recommendations

While this text lays the ground work for the use of tensor decompositions in HCI, there are several promising areas for future work.

Constrained Decompositions & Interpreting Factors

One of the major advantages of low-rank tensor decompositions over PCA is the interpretability of their factors. This work does not address the potential of using the factors either directly to model the source or indirectly when finding a PSF model. As is demonstrated in this work, the rank of the decompositions needs to be moderately high to detect the faint companions at high contrast. Nevertheless, in derotated observations with off-axis sources aligned, the off-axis sources should be well approximated with a low-rank model as it is local, smooth and highly sparse. Furthermore, for decompositions like CPD and Tucker the spectral factor can likely be rank-1 as the source is unresolved. Hence, a rank-1 spectral factor can be directly interpreted as the best-fit spectrum of the source, as [Equation 8.1](#). Alternatively, for decompositions with spectral rank- > 1 , the factors can be interpreted as spectral components. Adapting an approach similar to molecule mapping ([Hoeijmakers et al., 2018](#)) the observation tensor can be partially decomposed via HOI ([algorithm 7](#)) on the non-spectral modes which are either fixed at known species spectra or cross-approximated from the observation data and template spectra according to the targets (consistent) radial velocity ([Oseledets et al., 2010](#)).

$$\mathcal{X}_{Obs} \approx \mathcal{M}_{PSF} + \mathcal{L} + \mathcal{E}, \quad (8.1)$$

where the observations are modelled as the sum of PSF model \mathcal{M}_{PSF} , the sparse off-axis companion model \mathcal{L} , and error \mathcal{E} so that $\|\mathcal{M}_{PSF} + \mathcal{L} - \mathcal{X}_{Obs}\|$ is minimal. If the PSF and companion model are tensor decompositions then \mathcal{E} is the sum of reconstruction errors of the decompositions with rank- R . Note, that [Equation 8.1](#) is itself a BTM.

$$\mathcal{L} = \mathcal{G}_L \times_1 \mathbf{U}^{(\lambda)} \times_2 \mathbf{U}^{(x)} \times_3 \mathbf{U}^{(y)}, \quad (8.2a)$$

$$= \mathcal{G}_L \times_1 \begin{bmatrix} \mathbf{s}_1(\lambda') & \mathbf{s}_2(\lambda') & \dots & \mathbf{s}_N(\lambda') \end{bmatrix} \times_2 \mathbf{U}^{(x)} \times_3 \mathbf{U}^{(y)} \quad (8.2b)$$

where the spectral factor matrix $\mathbf{U}^{(\lambda)}$ is a matrix of the template molecule spectra $\mathbf{s}_i(\lambda')$ at the Doppler shifted wavelengths λ' with maximum correlation coefficient. Considering the rank- $[5, 1, 1]$ Tucker decomposition of an reduced spectral cube for simplicity, shown in [Figure 8.1](#). The companion $\llbracket \mathcal{G}_L; \mathbf{U}^{(\lambda)}, \mathbf{x}, \mathbf{y} \rrbracket$ is modelled using 5 template spectra of molecule species and two free spatial components in x and y. Since the PSF might be known, the problem can be further reduced by setting the spatial components to a decomposition of the PSF translated in x and y. The final companion model then depends only on the radial velocity v and PSF translation vector $\begin{bmatrix} \Delta x & \Delta y \end{bmatrix}^T$. In addition to retrieving a best-fit spectrum as with the negative fake companion technique in this way, the translation vector is similarly useful for astrometry.

The decompositions above interpret the spectral factor and translation vector as the best-fit models, however, the factors can also be computed in a Bayesian way which yields the uncertainties on the estimates ([Guhaniyogi et al., 2017](#)). Specific models for rank-3 BTM like in [Equation 8.1](#) have been developed by [Giampouras et al., 2022](#).

Polar Models

In the setup of the previous section, the polar (full-frame) methods ([Section 3.2.4](#)) can offer a further advantage. In polar coordinates off-axis sources in scaled ASDI observations move along either azimuthal or radial trajectories when fixing the index in the spectral or temporal mode respectively. This can be a

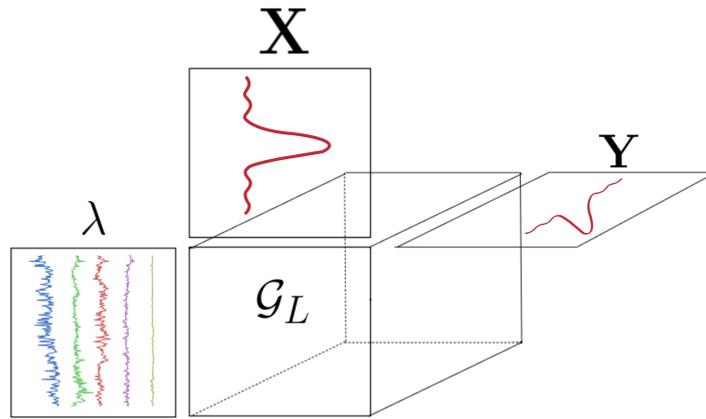


Figure 8.1: Tucker decomposed molecule map. Since the factors are physically meaningful, structure can be imposed on the decomposition: the spectral factors are template spectra and the spatial components model the PSF intensity distribution.

useful discriminator between speckles and exoplanets for a model which finds a PSF and exoplanet model simultaneously. A similar approach has been taken by [Vary et al., 2023](#), however, as for all matrix based methods, their method cannot fully exploit the higher-order structure in the observations. A tensor model using e.g. the CPD or Tucker decomposition could exploit the known trajectory of off-axis sources much more easily.

Affine Transforms of Decomposed Tensors for HCI

While approximating SDI observations as a block-term decomposition is straightforward, [Equation 8.1](#), it also highlights an issue when using tensor decomposition in the HCI context, especially for ASDI observations. The PSF model \mathcal{M}_{PSF} is modeled using the scaled and rotated observations, however, the companion model \mathcal{L} is found using re-scaled derotated observations which align the companion. The repeated scaling-rescaling-derotation is computationally expensive, especially for large data-sets. The repeated interpolation involved in image operations such as scaling and rotation is a more general computational problem in HCI ([Christiaens, 2024](#))¹.

However, with tensor decompositions the PSF model can be physically meaningful as highlighted above. Being able to perform the scaling-rescaling-derotation directly on the decomposition factors could be computationally beneficial and allow for the joint optimization of a PSF and exoplanet model. ([Cichocki, 2014](#); [Vervliet et al., 2014](#)) As highlighted above, currently each optimization step for finding [Equation 8.2](#) needs to perform at least two affine transforms (each scaling and rotating) to model \mathcal{M}_{PSF} and \mathcal{L} jointly if \mathcal{L} has rank- $[R_\lambda, \dots, 1, 1]$. To perform these transformations the full tensor needs to be computed from the decompositions which is inefficient. However, affine transformations are not defined for decomposed tensors and it is not known if finding general transformations for decomposed higher-order tensors is possible ([Ballester Ripoll, 2023](#)).

Affine transforms of CPD and Tucker factor matrices can conceivably be approximated using tensor cross-approximation ([Aguado et al., 2019](#); [Oseledets et al., 2010](#)), via a Fourier transform ([Rakhuba et al., 2015](#)), or simply via rotation and projection (see [Figure 8.2](#)). However, the corresponding transform of the Tucker core or for TTD factors is not known. The alternative to affine transform is to model each temporal tensor separately in the scaled and rotated observations by modeling the trajectory of the companion in the tensor ([Vary et al., 2023](#)).

Other Differential Imaging Methods

The ability to build a low-rank PSF model with tensor decompositions is naturally not limited to ASDI. Instead, the methods can similarly extend reference star differential imaging ([Xie et al., 2022](#)) by cross-approximating the PSF model jointly from reference and target and differential imaging over several epochs ([Dallant et al., 2023](#)) by interpreting the (aligned) epochs as an additional mode. Coherent Differential

¹See this [tutorial](#) for a practical demonstration and comparison. ([Christiaens, 2024](#))

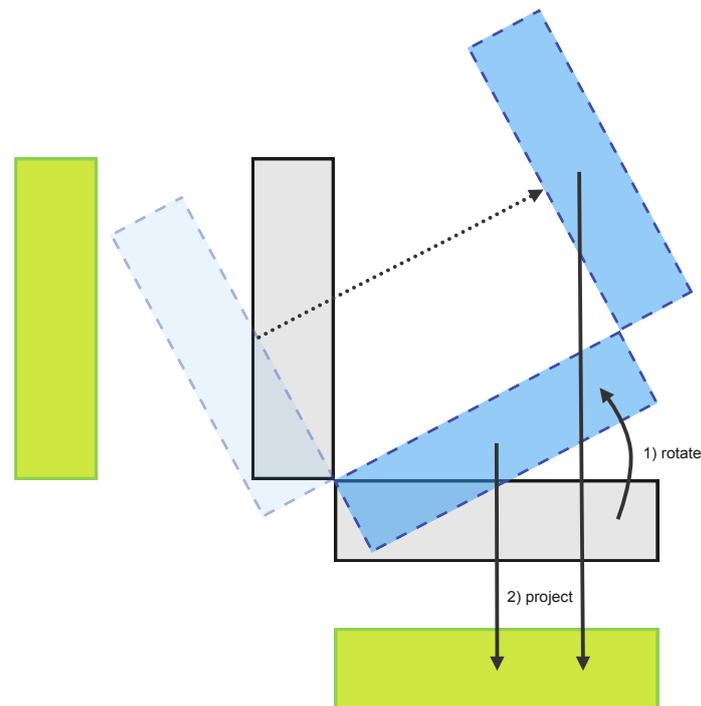


Figure 8.2: Rotation and projection of spatial factor matrices ($U^{(x)}$ and $U^{(y)}$) of the CPD to approximate the (affine) rotation of a matrix (observation frame). The original factors (in gray) are first rotated (blue) using an affine transform and then jointly projected onto the new basis (green). Note that the resulting factors might no longer be the best rank- R approximation of the rotated frame.

Imaging in HCI might also benefit from tensor approaches to speed up imaging (Afrakhteh et al., 2021), improve SNR (Fang et al., 2023), and integrated adaptive optics with phase aberration correction (Joseph, 2024).

Deep Learning in HCI

Deep learning on HCI is extremely challenging both from the perspective of the dimension as well as size of data sets and high dynamic contrast. Recently, significant advances have been made by Cantero et al., 2023, Wolf et al., 2023 and Flasseur et al., 2024, however deep learning on real HCI observations, especially in the ASDI context, remains limited and computationally expensive. Further, many of the methods that perform well for common images (CNN, ViT, Dosovitskiy et al., 2021; Al-Saffar et al., 2017) are not suited or intractable for large A(S)DI cubes. Deep learning based on tensor decompositions, tensorized neural networks, or decomposed higher-order convolutions has revolutionized image sequence processing (Jahromi et al., 2023; Ma et al., 2022; Novikov et al., 2015). These methods should be investigated as a potential key enabler of large scale deep learning in HCI.

Real-Time PSF Modeling

This work repeatedly highlights the computational advantages of tensor decompositions with respect to matrix factorization methods. The computational effort of computing matrix factorizations has repeatedly been tried to address in order to develop methods that can model or subtract the stellar PSF in real- or near-real-time and dig deeper dark holes in the regions of interest or more generally improve adaptive optics performance (Guyon et al., 2022; Long et al., 2021). Tensor decompositions, especially streaming tensor decompositions should be explored for this area both due to their computational efficiency but also because of their advantages in regards to selective and iterative updating. (Fanaee-T et al., 2016; Mardani et al., 2015; Thanh et al., 2023; Yu et al., 2015)

Tensor Completion & Data Imputation

Recently, [Ren et al., 2020](#) have considered HCI from the data imputation perspective. Due to their extensive use in recommender systems and denoising, tensor decompositions are ideally suited for handling missing or masked data in high-contrast imaging (HCI) datasets. By applying tensor completion techniques, it is possible to reconstruct observations from their low-rank factors ([Liu et al., 2013](#); [Zhou et al., 2018](#)). This can prove particularly beneficial in HCI to robustly denoise observations or approximate the unbiased PSF model at the location of a known source.

Tensorization, Permutation & Spatial Unfolding

As discussed in [Section 3.2.3](#) and [Section 3.2.6](#), tensorization and index permutation might be beneficial for tensor methods. This work does not systematically assess the impact either has on the performance of the proposed methods. As such further tensorization might be able to break the curse of dimensionality that impacts the memory usage for methods like CPD, allowing PSF modelling for very large ASDI observations all at once. This work finds that spatial unfolding can have a positive impact on the performance of decompositions due to the decompositions reaching the required contrast for very faint companions by not modeling spatial low-rank structure. Nevertheless, more expressive decompositions like TTD and TRD performs well without spatial unfolding. Determining whether this is due to the ability of factors to interact and if there are specific cross-couplings that are preferential for detecting companions can guide the rank selection.

Gradient-based Discriminator

This work mostly addresses the PSF modeling and subtraction part of HCI post-processing. However, as discussed in [Chapter 5](#), the detection of exoplanets also requires a discriminator to differentiate between exoplanet and noise. This discriminator is typically a statistical model like the SNR- or STIM map. Because the tensor models described in this work implicitly construct a computational graph, automatic differentiation methods can be used on the models. This potentially allows the construction of a gradient-based discriminator for differentiating between exoplanets and speckles. This idea is described in the following.

Assume that one has access to a residual image (this could as well be a tensor) which is obtained from ASDI observations \mathcal{X} . Neglecting the parameters of the reduction model ξ itself, the final residual \mathcal{R} is a function of the ASDI parameters, specifically, the parallactic angle variation $\theta = [\theta_1 \ \theta_2 \ \dots \ \theta_N]$ and wavelength scaling $\lambda = [\lambda_1 \ \lambda_2 \ \dots \ \lambda_M]$:

$$\mathcal{R} = \xi(\mathcal{X}, \theta, \lambda). \quad (8.3)$$

Presume that the reduction operation optimally recovers the exoplanet signal \mathcal{L} when the actual values for the parallactic angle variation and wavelength scaling are used, but still contains residual speckle noise \mathcal{E}_R :

$$\mathcal{R} = \mathcal{L} + \mathcal{E}_R \quad (8.4)$$

In other words, for each element in the exoplanet signal tensor it holds that:

$$\mathcal{L}_{i_1, i_2, \dots, i_K}(\theta, \lambda) > \mathcal{L}_{i_1, i_2, \dots, i_K}(\theta + \Delta_\theta, \lambda + \Delta_\lambda) \quad \forall \Delta_\theta, \Delta_\lambda \neq \mathbf{0}, \forall i_k \quad (8.5)$$

where Δ_θ and Δ_λ are errors on the true values of the true ASDI parameters and i_1, i_2, \dots, i_K are the indices of the residual, e.g. two spatial modes for a fully collapsed residual. This requires that the recovered exoplanet signal is positive, which is desirable ([Gonzalez et al., 2017](#)). Note that the same is not the case for the residual speckle noise elements since it is typically distributed around zero:

$$\mathcal{E}_{R_{i_1, i_2, \dots, i_K}}(\theta + \Delta_\theta, \lambda + \Delta_\lambda) \stackrel{\leq}{\geq} \mathcal{E}_{R_{i_1, i_2, \dots, i_K}}(\theta, \lambda) \quad \forall \Delta_\theta, \Delta_\lambda \neq \mathbf{0}, \forall i_k \quad (8.6)$$

Note that in practice, [Equation 8.5](#) and [Equation 8.6](#) might only hold for sufficiently large Δ_θ and Δ_λ due to the imperfect PSF modeling and reduction operation.

Now, consider the gradient of the exoplanet model element $\mathcal{L}_{i_1, i_2, \dots, i_K}$ with respect to the ASDI parameters $\Delta = \{\Delta_\theta, \Delta_\lambda\}$:

$$\nabla_\Delta \mathcal{L}_{i_1, i_2, \dots, i_K} = \begin{bmatrix} \frac{\partial \mathcal{L}_{i_1, i_2, \dots, i_K}}{\partial \Delta_{\theta_1}} & \frac{\partial \mathcal{L}_{i_1, i_2, \dots, i_K}}{\partial \Delta_{\theta_2}} & \dots & \frac{\partial \mathcal{L}_{i_1, i_2, \dots, i_K}}{\partial \Delta_{\theta_N}} \\ \frac{\partial \mathcal{L}_{i_1, i_2, \dots, i_K}}{\partial \Delta_{\lambda_1}} & \frac{\partial \mathcal{L}_{i_1, i_2, \dots, i_K}}{\partial \Delta_{\lambda_2}} & \dots & \frac{\partial \mathcal{L}_{i_1, i_2, \dots, i_K}}{\partial \Delta_{\lambda_M}} \end{bmatrix}, \quad (8.7)$$

so that $\nabla_{\Delta} \mathcal{L}_{i_1, i_2, \dots, i_K}$ is the Jacobian. For illustrative purposes, consider the directional derivative of the exoplanet model element with respect to the ASDI parameters:

$$\Delta \mathcal{L}_{i_1, i_2, \dots, i_K} \approx \nabla_{\Delta} \mathcal{L}_{i_1, i_2, \dots, i_K} \cdot \Delta. \quad (8.8)$$

Assuming local linearity a first-order Taylor expansion can be used to approximate the change in \mathcal{L} :

$$\Delta \mathcal{L}_{i_1, i_2, \dots, i_K} \approx \sum_{n=1}^N \frac{\partial \mathcal{L}_{i_1, i_2, \dots, i_K}}{\partial \Delta \theta_n} \Delta \theta_n + \sum_{m=1}^M \frac{\partial \mathcal{L}_{i_1, i_2, \dots, i_K}}{\partial \Delta \lambda_m} \Delta \lambda_m. \quad (8.9)$$

Since it was assumed that the exoplanet model is optimal for $\Delta = \mathbf{0}$ in Equation 8.5, the directional derivative is strictly positive:

$$\Delta \mathcal{L}_{i_1, i_2, \dots, i_K} = \nabla_{\Delta} \mathcal{L}_{i_1, i_2, \dots, i_K} \cdot \Delta > 0 \quad \forall \Delta \neq \mathbf{0}, \forall i_k. \quad (8.10)$$

This means that $-\Delta \mathcal{L}_{i_1, i_2, \dots, i_K}$ is locally convex for non-zero error $\Delta \alpha \in \{\Delta \mid \Delta \neq \mathbf{0}\}$, shown in Figure 8.3a. Conversely, the directional derivative of the residual noise $\Delta \mathcal{E}_{R_{i_1, i_2, \dots, i_K}}$ can be positive or negative or even constant for the same error, shown in Figure 8.3b.

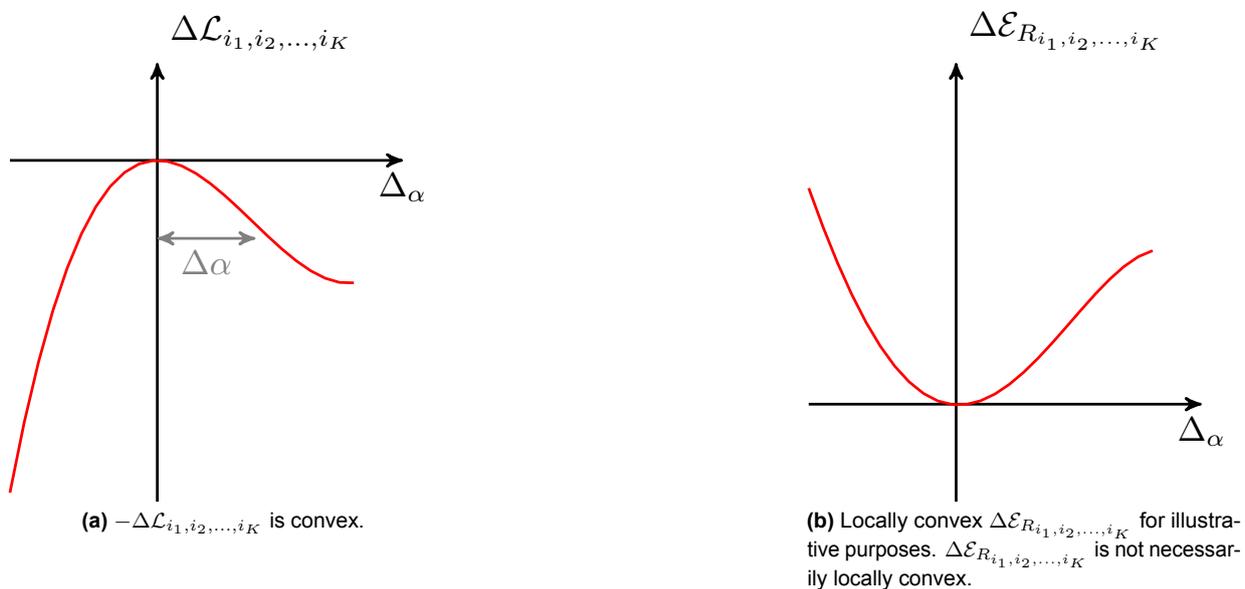


Figure 8.3: Directional derivative of residual components.

This difference in the directional derivative of exoplanet signal and speckle noise might allow discriminating between them in \mathcal{R} . This can be done directly for a residual obtained by independent component analysis (ICA) (Denarié, 2022), since the residual \mathcal{R} is obtained directly, however ICA methods lack behind PCA in maturity. Hence, the directional derivatives need to be found by tracing back the computational graph of the methods developed in Chapter 3. The goal is to discriminate between \mathcal{L} and \mathcal{E}_R based on the directional derivative of elements in \mathcal{R} with respect to error terms Δ added to θ and λ . Hence, $\Delta \mathcal{R}_{i_1, i_2, \dots, i_K}$ needs to be evaluated. This can be archived by applying the chain rule to the computational graph (see Figure 3.1) of the tensor based PSF subtraction methods. For the derivative with respect to θ this is simply

$$\frac{\partial \mathcal{R}}{\partial \theta}, \quad (8.11)$$

since θ is only used in the de-rotation operation. To find the derivative with respect to λ the chain rule can be used on the three paths originating from the λ -node

$$\frac{\partial \mathcal{R}}{\partial \lambda} = \frac{\partial \mathcal{R}}{\partial \lambda} + \frac{\partial \mathcal{R}}{\partial \mathcal{X}'} \cdot \left(\frac{\partial \mathcal{X}'}{\partial \lambda} + \frac{\partial \mathcal{X}'}{\partial \mathcal{M}'} \cdot \frac{\partial \mathcal{M}'}{\partial \mathcal{X}'} \cdot \frac{\partial \mathcal{X}'}{\partial \lambda} \right). \quad (8.12)$$

Using Equation 8.11 and Equation 8.12 the effect of perturbations to θ and λ can be found using

$$\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} \approx \sum_{n=1}^N \frac{\partial\mathcal{R}_{i_1, i_2, \dots, i_K}}{\partial\Delta\theta_n} \Delta\theta_n + \sum_{m=1}^M \frac{\partial\mathcal{R}_{i_1, i_2, \dots, i_K}}{\partial\Delta\lambda_m} \Delta\lambda_m. \quad (8.13)$$

However, back propagation e.g. using PyTorch, is more suitable to approximating the gradients. Based on the previous treatment of the idealized gradients for the two classes L (planet) and E (noise), the following observations can be made:

- $P(L|\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} \leq 0) > 0$
- $P(L|\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} > 0) = 0$
- $P(E|\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} \leq 0) > 0$
- $P(E|\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} > 0) = 1$

This leaves the case where $\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} \leq 0$. Since both $P(E)$ and $P(L)$ are unknown, Bayes theorem cannot be used. Instead, the likelihood ratio (LR) test could be used. The likelihood ratio considered here will be

$$LR = \frac{P(\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} \leq 0|L)}{P(\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} \leq 0|E)}, \quad (8.14)$$

where $P(\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} \leq 0|L) = 1$, so that the LR becomes:

$$LR = 1/P(\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} \leq 0|E). \quad (8.15)$$

Like for STIM, this requires the definition of a decision threshold. Finding a decision threshold could be done empirically by observing the proportion of $\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} \leq 0$ in regions that are certainly from class E and trading off sensitivity and specificity. Finding a provably optimal decision threshold is beyond the scope of this work.

Assuming a threshold τ has been found, instances in \mathcal{R} can be classified using:

- If $\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} > 0$, classify as E: noise
- If $\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} \leq 0$, calculate LR and apply decision rule:
 - If $LR > \tau$, classify as L: exoplanet
 - If $LR \leq \tau$, classify as E: noise

This can be interpreted as $LR > 1$ suggesting the instance is more likely to be from class L, $LR < 1$ suggesting the instance is more likely to be from class E, and the magnitude of LR indicating the strength of this suggestion. The limitation of an empirical threshold is the ability to estimate $P(\Delta\mathcal{R}_{i_1, i_2, \dots, i_K} \leq 0|E)$, in a comparable region of the residual.

Interpreting this approach in the framework of classifiers for exoplanet detection, the approach outlined here asks the question: "How does input Δ affect the predicted probability for the exoplanet and noise class?"

The computational methods for tensor decomposition used in this work directly allow the computation of sets of partial derivatives and gradients on the decomposed tensors. Hence, gradient-based discriminators can be directly applied to residuals resulting from the methods proposed in this work.

Acknowledgements

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

The Direct Imaging Virtual Archive (DIVA), operated at CeSAM/LAM, Marseille, France. The Exoplanet Imaging Data Challenge, collaboratively managed and maintained by members of the high-contrast imaging community, mainly led at IPAG/CNRS/UGA (Grenoble, France) and PSILab/STAR institute/ULiège (Liège, Belgium).

The international Gemini Observatory, a program of NSF NOIRLab, which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the U.S. National Science Foundation on behalf of the Gemini Observatory partnership: the U.S. National Science Foundation (United States), National Research Council (Canada), Agencia Nacional de Investigación y Desarrollo (Chile), Ministerio de Ciencia, Tecnología e Innovación (Argentina), Ministério da Ciência, Tecnologia, Inovações e Comunicações (Brazil), and Korea Astronomy and Space Science Institute (Republic of Korea).

The SIMBAD database, operated at CDS, Strasbourg, France. The NASA Exoplanet Archive, which is operated by the California Institute of Technology, under contract with the National Aeronautics and Space Administration under the Exoplanet Exploration Program.

Data obtained from and tools provided by the portal `exoplanet.eu` of The Extrasolar Planets Encyclopaedia.

Furthermore, the following software packages have been used in this thesis: numpy, matplotlib, seaborn, corner, Astropy, UltraNest, CuPy, PyTorch, TensorNetwork, Tensorly, and TNTorch.

I extend my sincere gratitude to the dedicated teams responsible for the operation and ongoing development of these instruments and telescopes, reduction pipelines and data centers. Your efforts in making high-quality astronomical data accessible have been indispensable.

Special thanks are also due to the contributors to the software used in this thesis, with a particular acknowledgment to those involved in open scientific software development. Open source software has been crucial to the accomplishments of this interdisciplinary research. Your commitment to open science practices not only enriched this work but also fosters a culture of collaboration and transparency that is vital for the future of high-contrast imaging.

Bibliography

- Absil, O., J. Milli, D. Mawet, A.-M. Lagrange, J. Girard, et al. (Nov. 2013). “Searching for Companions down to 2 AU from Beta Pictoris Using the L'-Band AGPM Coronagraph on VLT/NACO”. In: *A&A* 559, p. L12. DOI: [10.1051/0004-6361/201322748](https://doi.org/10.1051/0004-6361/201322748).
- AFP (Feb. 2023). *Storm Ciaran Forces Dutch to Scrap 'Headwind Cycle Championships'*. URL: <https://www.flashscore.co.uk/news/storm-ciaran-forces-dutch-to-scrap-headwind-cycle-championships/QcyDk7JU/> (visited on 06/03/2024).
- Afrakhteh, S. and H. Behnam (Oct. 2021). “Coherent Plane Wave Compounding Combined With Tensor Completion Applied for Ultrafast Imaging”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 68.10, pp. 3094–3103. DOI: [10/gs4zcd](https://doi.org/10/gs4zcd).
- Aguado, J. V., D. Borzacchiello, K. S. Kollepara, F. Chinesta, and A. Huerta (Oct. 2019). “Tensor Representation of Non-linear Models Using Cross Approximations”. In: *J Sci Comput* 81.1, pp. 22–47. DOI: [10.1007/s10915-019-00917-2](https://doi.org/10.1007/s10915-019-00917-2).
- Amara, A. and S. Quanz (Dec. 2012). “PynPoint: An Image Processing Package for Finding Exoplanets”. In: *Monthly Notices of the Royal Astronomical Society* 427.2, pp. 948–955. DOI: [10.1111/j.1365-2966.2012.21918.x](https://doi.org/10.1111/j.1365-2966.2012.21918.x).
- Baba, N. and N. Murakami (Dec. 2003). “A Method to Image Extrasolar Planets with Polarized Light”. In: *PASP* 115.814, p. 1363. DOI: [10.1086/380422](https://doi.org/10.1086/380422).
- Ballester Ripoll, R. (Jan. 2023). *Affine Transformations of Decomposed Tensors. Email Exchange from 01-10-2023*.
- Ballester Ripoll, R., E. G. Paredes, and R. Pajarola (Jan. 2018). “Tensor Algorithms for Advanced Sensitivity Metrics”. In: *SIAM/ASA J. Uncertainty Quantification* 6.3, pp. 1172–1197. DOI: [10.1137/17M1160252](https://doi.org/10.1137/17M1160252).
- (Mar. 2019). “Sobol Tensor Trains for Global Sensitivity Analysis”. In: *Reliability Engineering & System Safety* 183, pp. 311–322. DOI: [10.1016/j.ress.2018.11.007](https://doi.org/10.1016/j.ress.2018.11.007).
- Batselier, K. (Nov. 2018). *The Trouble with Tensor Ring Decompositions*. DOI: [10.48550/arXiv.1811.03813](https://doi.org/10.48550/arXiv.1811.03813). (Visited on 11/13/2023).
- Beard, C., P. Robertson, F. Dai, R. Holcomb, J. Lubin, et al. (Jan. 2024). “The TESS-Keck Survey. XVII. Precise Mass Measurements in a Young, High-multiplicity Transiting Planet System Using Radial Velocities and Transit Timing Variations”. In: *AJ* 167.2, p. 70. DOI: [10.3847/1538-3881/ad1330](https://doi.org/10.3847/1538-3881/ad1330).
- Bergin, E. A., R. A. Booth, M. J. Colmenares, and J. D. Ilee (June 2024). *C/O Ratios and the Formation of Wide Separation Exoplanets*. DOI: [10.48550/arXiv.2406.12037](https://doi.org/10.48550/arXiv.2406.12037). (Visited on 07/05/2024).
- Beuzit, J.-L., A. Vigan, D. Mouillet, K. Dohlen, R. Gratton, et al. (Nov. 2019). “SPHERE: The Exoplanet Imager for the Very Large Telescope”. In: *A&A* 631, A155. DOI: [10.1051/0004-6361/201935251](https://doi.org/10.1051/0004-6361/201935251).
- Beuzit, J.-L., M. Feldt, K. Dohlen, D. Mouillet, P. Puget, et al. (July 2008). “SPHERE: A Planet Finder Instrument for the VLT”. In: *Ground-Based and Airborne Instrumentation for Astronomy II*. Vol. 7014. SPIE, pp. 476–487. DOI: [10.1117/12.790120](https://doi.org/10.1117/12.790120).
- Biamonte, J. (Jan. 2020). *Lectures on Quantum Tensor Networks*. DOI: [10.48550/arXiv.1912.10049](https://doi.org/10.48550/arXiv.1912.10049). (Visited on 11/20/2023).
- Bisikalo, D. V., V. I. Shematovich, P. V. Kaygorodov, and A. G. Zhilkin (Aug. 2021). “Gas Envelopes of Exoplanets — Hot Jupiters”. In: *Phys.-Usp.* 64.8, p. 747. DOI: [10.3367/UFNe.2020.11.038879](https://doi.org/10.3367/UFNe.2020.11.038879).
- Boccaletti, A., G. Chauvin, D. Mouillet, O. Absil, F. Allard, et al. (Mar. 2020). *SPHERE+: Imaging Young Jupiters down to the Snowline*. DOI: [10.48550/arXiv.2003.05714](https://doi.org/10.48550/arXiv.2003.05714). (Visited on 06/24/2024).

- Boccaletti, A., E. Sezestre, A.-M. Lagrange, P. Thébault, R. Gratton, et al. (June 2018). “Observations of Fast-Moving Features in the Debris Disk of AU Mic on a Three-Year Timescale: Confirmation and New Discoveries”. In: *A&A* 614, A52. DOI: [10.1051/0004-6361/201732462](https://doi.org/10.1051/0004-6361/201732462).
- Boccaletti, A., C. Thalmann, A.-M. Lagrange, M. Janson, J.-C. Augereau, et al. (Oct. 2015). “Fast-Moving Features in the Debris Disk around AU Microscopii”. In: *Nature* 526.7572, pp. 230–232. DOI: [10.1038/nature15705](https://doi.org/10.1038/nature15705).
- Bonse, M. J., E. O. Garvin, T. D. Gebhard, F. A. Dannert, F. Cantalloube, et al. (Aug. 2023). “Comparing Apples with Apples: Robust Detection Limits for Exoplanet High-Contrast Imaging in the Presence of Non-Gaussian Noise”. In: *AJ* 166.2, p. 71. DOI: [10/gs4zb9](https://doi.org/10/gs4zb9).
- Bonse, M. J., S. P. Quanz, and A. Amara (Apr. 2018). *Wavelet Based Speckle Suppression for Exoplanet Imaging - Application of a de-Noiseing Technique in the Time Domain*. DOI: [10.48550/arXiv.1804.05063](https://doi.org/10.48550/arXiv.1804.05063).
- Boss, A., R. Oppenheimer, J. Pitman, L. Poyneer, and S. Ridgway (Aug. 2023). *Progress in Technology for Exoplanet Missions (Technology Plan Appendix 2023)*. Tech. rep.
- Bottega, F. and Diniz, Felipe B. (Sept. 2023). *Tensor Fox*. URL: <https://github.com/felipebottega/Tensor-Fox> (visited on 12/01/2023).
- Bottom, M., J. K. Wallace, R. D. Bartos, J. C. Shelton, and E. Serabyn (Jan. 2017). “Speckle Suppression and Companion Detection Using Coherent Differential Imaging”. In: *Monthly Notices of the Royal Astronomical Society* 464.3, pp. 2937–2951. DOI: [10.1093/mnras/stw2544](https://doi.org/10.1093/mnras/stw2544).
- Bousse, M., O. Debals, and L. De Lathauwer (Jan. 2017). “A Tensor-Based Method for Large-Scale Blind Source Separation Using Segmentation”. In: *IEEE Trans. Signal Process.* 65.2, pp. 346–358. DOI: [10.1109/TSP.2016.2617858](https://doi.org/10.1109/TSP.2016.2617858).
- Bouwmans, T., N. S. Aybat, and E.-h. Zahzah (Sept. 2016). *Handbook of Robust Low-Rank and Sparse Matrix Decomposition: Applications in Image and Video Processing*. CRC Press.
- Bridgeman, J. C. and C. T. Chubb (May 2017). “Hand-Waving and Interpretive Dance: An Introductory Course on Tensor Networks”. In: *J. Phys. A: Math. Theor.* 50.22, p. 223001. DOI: [10.1088/1751-8121/aa6dc3](https://doi.org/10.1088/1751-8121/aa6dc3).
- Buchner, J. (Jan. 2016). “A Statistical Test for Nested Sampling Algorithms”. In: *Stat Comput* 26.1-2, pp. 383–392. DOI: [10.1007/s11222-014-9512-y](https://doi.org/10.1007/s11222-014-9512-y).
- (Oct. 2019). “Collaborative Nested Sampling: Big Data vs. Complex Physical Models”. In: *PASP* 131.1004, p. 108005. DOI: [10.1088/1538-3873/aae7fc](https://doi.org/10.1088/1538-3873/aae7fc).
- (Apr. 2021). *UltraNest – a Robust, General Purpose Bayesian Inference Engine*. DOI: [10.48550/arXiv.2101.09604](https://doi.org/10.48550/arXiv.2101.09604). (Visited on 05/12/2024).
- Cantalloube, F., V. Christiaens, C. Cantero, E. Nasedkin, A. Cioppa, et al. (Aug. 2022). “Exoplanet Imaging Data Challenge, Phase II: Characterization of Exoplanet Signals in High-Contrast Images”. In: *Adaptive Optics Systems VIII*, p. 4. DOI: [10.1117/12.2627968](https://doi.org/10.1117/12.2627968).
- Cantalloube, F., C. Gomez-Gonzalez, O. Absil, C. Cantero, R. Bacher, et al. (Dec. 2020). “Exoplanet Imaging Data Challenge: Benchmarking the Various Image Processing Methods for Exoplanet Detection”. In: *Adaptive Optics Systems VII*, p. 321. DOI: [10.1117/12.2574803](https://doi.org/10.1117/12.2574803).
- Cantero, C., O. Absil, C.-H. Dahlqvist, and M. Van Droogenbroeck (Oct. 2023). “NA-SODINN: A Deep Learning Algorithm for Exoplanet Image Detection Based on Residual Noise Regimes”. In: *A&A*. DOI: [10.1051/0004-6361/202346085](https://doi.org/10.1051/0004-6361/202346085).
- Chao, K.-H., R. deGraffenried, M. Lach, W. Nelson, K. Truax, et al. (May 2021). “Lava Worlds: From Early Earth to Exoplanets”. In: *Geochemistry* 81.2, p. 125735. DOI: [10.1016/j.chemer.2020.125735](https://doi.org/10.1016/j.chemer.2020.125735).
- Charbonneau, D., T. M. Brown, D. W. Latham, and M. Mayor (Dec. 1999). “Detection of Planetary Transits Across a Sun-like Star”. In: *ApJ* 529.1, p. L45. DOI: [10.1086/312457](https://doi.org/10.1086/312457).

- Chatzichristos, C., M. Vandecastelle, E. Kofidis, S. Theodoridis, L. D. Lathauwer, et al. (Nov. 2019). “Tensor-Based Blind fMRI Source Separation Without the Gaussian Noise Assumption — A β -Divergence Approach”. In: *2019 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pp. 1–5. DOI: [10.1109/GlobalSIP45357.2019.8969150](https://doi.org/10.1109/GlobalSIP45357.2019.8969150).
- Chauvin, G., S. Desidera, A. -. Lagrange, A. Vigan, M. Feldt, et al. (Dec. 2017). *SHINE, The SpHere Infrared Survey for Exoplanets*. Proceedings of the Annual Meeting of the French Society of Astronomy and Astrophysics SF2A-2017. Paris: French Society of Astronomy and Astrophysics.
- Chen, Z., Y. Li, and J. Lu (May 2019). *Tensor Ring Decomposition: Optimization Landscape and One-loop Convergence of Alternating Least Squares*. DOI: [10.48550/arXiv.1905.07101](https://doi.org/10.48550/arXiv.1905.07101). (Visited on 06/07/2024).
- Chi, E. C. and T. G. Kolda (Jan. 2012). “On Tensors, Sparsity, and Nonnegative Factorizations”. In: *SIAM J. Matrix Anal. & Appl.* 33.4, pp. 1272–1299. DOI: [10.1137/110859063](https://doi.org/10.1137/110859063).
- Chiang, E. and J. Fung (Oct. 2017). “Stellar Winds and Dust Avalanches in the AU Mic Debris Disk”. In: *ApJ* 848.1, p. 4. DOI: [10.3847/1538-4357/aa89e6](https://doi.org/10.3847/1538-4357/aa89e6).
- Chilcote, J., Q. Konopacky, R. J. D. Rosa, R. Hamper, B. Macintosh, et al. (Dec. 2020). “GPI 2.0: Upgrading the Gemini Planet Imager”. In: *Ground-Based and Airborne Instrumentation for Astronomy VIII*. Vol. 11447. SPIE, pp. 394–407. DOI: [10.1117/12.2562578](https://doi.org/10.1117/12.2562578).
- Chiueh, T.-c., C.-k. Yang, T. He, H. Pfister, and A. Kaufman (1997). “Integrated volume compression and visualization”. In: *Proceedings. Visualization '97 (Cat. No. 97CB36155)*. IEEE, pp. 329–336.
- Chomez, A., A.-M. Lagrange, P. Delorme, M. Langlois, G. Chauvin, et al. (July 2023a). “Preparation for an Unsupervised Massive Analysis of SPHERE High-Contrast Data with PACO - Optimization and Benchmarking on 24 Solar-Type Stars”. In: *A&A* 675, A205. DOI: [10.1051/0004-6361/202245723](https://doi.org/10.1051/0004-6361/202245723).
- Chomez, A., V. Squicciarini, A.-M. Lagrange, P. Delorme, G. Viswanath, et al. (Aug. 2023b). “An Imaged 15 MJup Companion within a Hierarchical Quadruple System”. In: *A&A* 676, p. L10. DOI: [10.1051/0004-6361/202347044](https://doi.org/10.1051/0004-6361/202347044).
- Christiaens, V. (Mar. 2024). *Image Operations: FFT- vs Interpolation-Based Methods - Vortex Image Processing*. URL: https://vip.readthedocs.io/en/latest/tutorials/08_imlib_and_interpolation.html (visited on 07/07/2024).
- Christiaens, V., C. A. G. Gonzalez, R. Farkas, C.-H. Dahlqvist, E. Nasedkin, et al. (Jan. 2023). “VIP: A Python Package for High-Contrast Imaging”. In: *Journal of Open Source Software* 8.81, p. 4774. DOI: [10.21105/joss.04774](https://doi.org/10.21105/joss.04774).
- Cichocki, A. (Aug. 2014). *Era of Big Data Processing: A New Approach via Tensor Networks and Tensor Decompositions*. DOI: [10.48550/arXiv.1403.2048](https://doi.org/10.48550/arXiv.1403.2048). (Visited on 07/07/2024).
- Cichocki, A., D. Mandic, L. De Lathauwer, G. Zhou, Q. Zhao, et al. (Mar. 2015). “Tensor Decompositions for Signal Processing Applications: From Two-Way to Multiway Component Analysis”. In: *IEEE Signal Processing Magazine* 32.2, pp. 145–163. DOI: [10.1109/MSP.2013.2297439](https://doi.org/10.1109/MSP.2013.2297439).
- Codona, J. L. and R. Angel (Mar. 2004). “Imaging Extrasolar Planets by Stellar Halo Suppression in Separately Corrected Color Bands”. In: *ApJ* 604.2, p. L117. DOI: [10.1086/383569](https://doi.org/10.1086/383569).
- Cohen, N., O. Sharir, and A. Shashua (2016). “On the expressive power of deep learning: A tensor analysis”. In: *Conference on learning theory*. PMLR, pp. 698–728.
- Currie, T., B. Biller, A. Lagrange, C. Marois, O. Guyon, et al. (July 2023). “Direct Imaging and Spectroscopy of Extrasolar Planets”. In: 534, p. 799. DOI: [10.48550/arXiv.2205.05696](https://doi.org/10.48550/arXiv.2205.05696).
- Currie, T., A. Burrows, Y. Itoh, S. Matsumura, M. Fukagawa, et al. (Mar. 2011). “A Combined Subaru/VLT/MMT 1-5 M m Study of Planets Orbiting HR 8799: Implications for Atmospheric Properties, Masses, and Formation”. In: *The Astrophysical Journal* 729, p. 128. DOI: [10.1088/0004-637X/729/2/128](https://doi.org/10.1088/0004-637X/729/2/128).
- Daglayan, H., S. Vary, V. Leplat, N. Gillis, and P.-A. Absil (Oct. 2023). *Direct Exoplanet Detection Using $L1$ Norm Low-Rank Approximation*. DOI: [10.48550/arXiv.2304.03619](https://doi.org/10.48550/arXiv.2304.03619). (Visited on 11/13/2023).

- Dahlqvist, C.-H., J. Milli, O. Absil, F. Cantalloube, L. Matra, et al. (Oct. 2022). “The SHARDDS Survey: Limits on Planet Occurrence Rates Based on Point Sources Analysis via the Auto-RSM Framework”. In: *A&A* 666, A33. DOI: [10.1051/0004-6361/202244145](https://doi.org/10.1051/0004-6361/202244145).
- Dahlqvist, C.-H., F. Cantalloube, F. Cantalloube, F. Cantalloube, and O. Absil (2020). “Regime-Switching Model Detection Map for Direct Exoplanet Detection in ADI Sequences”. In: *Astronomy and Astrophysics* 633. DOI: [10/gs4zc2](https://doi.org/10/gs4zc2).
- Dai, Z., D. Ni, L. Pan, and Y. Zhu (Sept. 2021). “Five Methods of Exoplanet Detection”. In: *J. Phys.: Conf. Ser.* 2012.1, p. 012135. DOI: [10.1088/1742-6596/2012/1/012135](https://doi.org/10.1088/1742-6596/2012/1/012135).
- Dallant, J., M. Langlois, O. Flasseur, and É. Thiébaud (Nov. 2023). “PACOME: Optimal Multi-Epoch Combination of Direct Imaging Observations for Joint Exoplanet Detection and Orbit Estimation”. In: *Astronomy and Astrophysics* 679, A38. DOI: [10.1051/0004-6361/202346427](https://doi.org/10.1051/0004-6361/202346427).
- De Lathauwer, L. (Jan. 2008a). “Decompositions of a Higher-Order Tensor in Block Terms—Part I: Lemmas for Partitioned Matrices”. In: *SIAM J. Matrix Anal. Appl.* 30.3, pp. 1022–1032. DOI: [10.1137/060661685](https://doi.org/10.1137/060661685).
- (Jan. 2008b). “Decompositions of a Higher-Order Tensor in Block Terms—Part II: Definitions and Uniqueness”. In: *SIAM J. Matrix Anal. Appl.* 30.3, pp. 1033–1066. DOI: [10.1137/070690729](https://doi.org/10.1137/070690729).
- (June 2017). *Dealing with Curse and Blessing of Dimensionality through Tensor Decompositions*. KU Leuven.
- De Lathauwer, L., B. De Moor, and J. Vandewalle (Jan. 2000). “A Multilinear Singular Value Decomposition”. In: *SIAM J. Matrix Anal. Appl.* 21.4, pp. 1253–1278. DOI: [10.1137/S0895479896305696](https://doi.org/10.1137/S0895479896305696).
- De Lathauwer, L. and D. Nion (Jan. 2008). “Decompositions of a Higher-Order Tensor in Block Terms—Part III: Alternating Least Squares Algorithms”. In: *SIAM J. Matrix Anal. Appl.* 30.3, pp. 1067–1083. DOI: [10.1137/070690730](https://doi.org/10.1137/070690730).
- Delorme, P., N. Meunier, D. Albert, E. Lagadec, H. L. Coroller, et al. (Dec. 2017). *The SPHERE Data Center: A Reference for High Contrast Imaging Processing*. DOI: [10.48550/arXiv.1712.06948](https://doi.org/10.48550/arXiv.1712.06948). (Visited on 12/08/2023).
- Delorme, P. (Mar. 2024). *SHINE and BEAST HC-DC Observation Reduction and Pre-Processing. Email Exchange on 14-03-2024*.
- Denarié, P.-A. (2022). “Tensor Decomposition for Independent Component Analysis: Through Implicit Cumulant Tensor Manipulation”. In.
- Desidera, S., G. Chauvin, M. Bonavita, S. Messina, H. LeCoroller, et al. (July 2021). “The SPHERE Infrared Survey for Exoplanets (SHINE) - I. Sample Definition and Target Characterization”. In: *A&A* 651, A70. DOI: [10.1051/0004-6361/202038806](https://doi.org/10.1051/0004-6361/202038806).
- Doelman, D. S. (June 2021). “It’s Just a Phase: High-Contrast Imaging with Patterned Liquid-Crystal Phase Plates to Facilitate Characterization of Exoplanets”. PhD thesis. Leiden University. URL: <https://hdl.handle.net/1887/3191978> (visited on 05/20/2024).
- Dosovitskiy, A., L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, et al. (June 2021). *An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale*. DOI: [10.48550/arXiv.2010.11929](https://doi.org/10.48550/arXiv.2010.11929). (Visited on 12/10/2023).
- Dyrek, A., M. Min, L. Decin, J. Bouwman, N. Crouzet, et al. (Jan. 2024). “SO₂, Silicate Clouds, but No CH₄ Detected in a Warm Neptune”. In: *Nature* 625, pp. 51–54. DOI: [10.1038/s41586-023-06849-0](https://doi.org/10.1038/s41586-023-06849-0).
- Ertel, S., O. Absil, D. Defrère, J.-B. L. Bouquin, J.-C. Augereau, et al. (Oct. 2014). “A Near-Infrared Interferometric Survey of Debris-Disk Stars - IV. An Unbiased Sample of 92 Southern Stars Observed in H Band with VLT/PIONIER”. In: *A&A* 570, A128. DOI: [10.1051/0004-6361/201424438](https://doi.org/10.1051/0004-6361/201424438).
- Fanaee-T, H. and J. Gama (Apr. 2016). “Tensor-Based Anomaly Detection: An Interdisciplinary Survey”. In: *Knowledge-Based Systems* 98, pp. 130–147. DOI: [10.1016/j.knosys.2016.01.027](https://doi.org/10.1016/j.knosys.2016.01.027).

- Fang, Y., X. Shao, B. Liu, and H. Lv (July 2023). “Optical Coherence Tomography Image Despeckling Based on Tensor Singular Value Decomposition and Fractional Edge Detection”. In: *Heliyon* 9.7, e17735. DOI: [10/g4zcg](https://doi.org/10/g4zcg).
- Favier, G. (Jan. 2020). *From Algebraic Structures to Tensors*. 1st ed. Vol. 1. Digital Signal and Image Processing: Matrices and Tensors in Signal Processing. John Wiley & Sons.
- (Aug. 2021). *Matrix and Tensor Decompositions in Signal Processing*. 1st ed. Vol. 2. Digital Signal and Image Processing: Matrices and Tensors in Signal Processing. John Wiley & Sons.
- Favier, G. and A. L. de Almeida (Sept. 2014). “Overview of Constrained PARAFAC Models”. In: *EURASIP J. Adv. Signal Process.* 2014.1, p. 142. DOI: [10.1186/1687-6180-2014-142](https://doi.org/10.1186/1687-6180-2014-142).
- Flasseur, O., T. Bodrito, J. Mairal, J. Ponce, M. Langlois, et al. (Jan. 2024). “Deep PACO: Combining Statistical Models with Deep Learning for Exoplanet Detection and Characterization in Direct Imaging at High Contrast”. In: *Monthly Notices of the Royal Astronomical Society* 527, pp. 1534–1562. DOI: [10.1093/mnras/stad3143](https://doi.org/10.1093/mnras/stad3143).
- Flasseur, O., L. Denis, É. Thiébaud, and M. Langlois (Apr. 2018). “Exoplanet Detection in Angular Differential Imaging by Statistical Learning of the Nonstationary Patch Covariances. The PACO Algorithm”. In: *Astronomy and Astrophysics* 618. DOI: [10.1051/0004-6361/201832745](https://doi.org/10.1051/0004-6361/201832745).
- (May 2020). “PACO ASDI: An Algorithm for Exoplanet Detection and Characterization in Direct Imaging with Integral Field Spectrographs”. In: *A&A* 637, A9. DOI: [10.1051/0004-6361/201937239](https://doi.org/10.1051/0004-6361/201937239).
- Fortney, J. J. (May 2024). *Characterizing Exoplanetary Atmospheres: New Views on Exoplanet Composition and Chemistry*. DOI: [10.48550/arXiv.2405.04587](https://doi.org/10.48550/arXiv.2405.04587). (Visited on 06/03/2024).
- Fortney, J. J., R. I. Dawson, and T. D. Komacek (2021). “Hot Jupiters: Origins, Structure, Atmospheres”. In: *Journal of Geophysical Research: Planets* 126.3. DOI: [10.1029/2020JE006629](https://doi.org/10.1029/2020JE006629).
- Fusco, T., J. -S. Sauvage, D. Mouillet, A. Costille, C. Petit, et al. (July 2016). “SAXO, the SPHERE Extreme AO System: On-Sky Final Performance and Future Improvements”. In: *Adaptive Optics Systems*. DOI: [10.1117/12.2233319](https://doi.org/10.1117/12.2233319).
- Galicher, R. and J. Mazoyer (2023). “Imaging Exoplanets with Coronagraphic Instruments”. In: *Comptes Rendus. Physique* 24.S2, pp. 1–45. DOI: [10.5802/crphys.133](https://doi.org/10.5802/crphys.133).
- Gandhi, S., A. Kesseli, Y. Zhang, A. Louca, I. Snellen, et al. (May 2023). “Retrieval Survey of Metals in Six Ultrahot Jupiters: Trends in Chemistry, Rain-out, Ionization, and Atmospheric Dynamics”. In: *AJ* 165.6, p. 242. DOI: [10.3847/1538-3881/accd65](https://doi.org/10.3847/1538-3881/accd65).
- Ge, R. and T. Ma (2017). “On the Optimization Landscape of Tensor Decompositions”. In: *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc.
- Giampouras, P. V., A. A. Rontogiannis, and E. Kofidis (2022). “Block-Term Tensor Decomposition Model Selection and Computation: The Bayesian Way”. In: *IEEE Transactions on Signal Processing* 70, pp. 1704–1717. DOI: [10.1109/TSP.2022.3159029](https://doi.org/10.1109/TSP.2022.3159029).
- Goldfarb, D. and Z. Qin (Jan. 2014). “Robust Low-rank Tensor Recovery: Models and Algorithms”. In: *SIAM J. Matrix Anal. & Appl.* 35.1, pp. 225–253. DOI: [10.1137/130905010](https://doi.org/10.1137/130905010).
- Gonzalez, C. A. G., O. Absil, P.-A. Absil, M. Van Droogenbroeck, D. Mawet, et al. (May 2016). “Low-Rank plus Sparse Decomposition for Exoplanet Detection in Direct-Imaging ADI Sequences. The LLSG Algorithm”. In: *A&A* 589, A54. DOI: [10.1051/0004-6361/201527387](https://doi.org/10.1051/0004-6361/201527387).
- Gonzalez, C. A. G., O. Wertz, O. Absil, V. Christiaens, D. Defrère, et al. (June 2017). “VIP: Vortex Image Processing Package for High-contrast Direct Imaging”. In: *AJ* 154.1, p. 7. DOI: [10.3847/1538-3881/aa73d7](https://doi.org/10.3847/1538-3881/aa73d7).
- Google/TensorNetwork (Nov. 2023). Google. URL: <https://github.com/google/TensorNetwork> (visited on 12/02/2023).
- Grady, C. A., J. P. Wisniewski, G. Schneider, A. Boccaletti, A. Gaspar, et al. (Jan. 2020). “The Eroding Disk of AU Mic”. In: *ApJL* 889.1, p. L21. DOI: [10.3847/2041-8213/ab65bb](https://doi.org/10.3847/2041-8213/ab65bb).

- Gratton, R., M. Bonavita, D. Mesa, S. Desidera, A. Zurlo, et al. (May 2024). “Stellar Companions and Jupiter-like Planets in Young Associations”. In: *A&A* 685, A119. DOI: [10.1051/0004-6361/202348393](https://doi.org/10.1051/0004-6361/202348393).
- Gratton, R., V. Squicciarini, V. Nascimbeni, M. Janson, S. Reffert, et al. (Oct. 2023). “Multiples among B Stars in the Scorpius-Centaurus Association”. In: *A&A* 678, A93. DOI: [10.1051/0004-6361/202346806](https://doi.org/10.1051/0004-6361/202346806).
- Gray, J. (Sept. 2018). “Quimb: A Python Package for Quantum Information and Many-Body Calculations”. In: *JOSS* 3.29, p. 819. DOI: [10.21105/joss.00819](https://doi.org/10.21105/joss.00819).
- (Oct. 2023). *Jcmgray/Cotengra*. URL: <https://github.com/jcmgray/cotengra> (visited on 12/02/2023).
- Gray, J. and S. Kourtis (Mar. 2021). “Hyper-Optimized Tensor Network Contraction”. In: *Quantum* 5, p. 410. DOI: [10.22331/q-2021-03-15-410](https://doi.org/10.22331/q-2021-03-15-410).
- Gu, Q., H. Gui, and J. Han (2014). “Robust Tensor Decomposition with Gross Corruption”. In: *Advances in Neural Information Processing Systems*. Vol. 27. Curran Associates, Inc.
- Guerra, G., J.-B. Daban, S. Robbe-Dubois, R. Douet, L. Abe, et al. (May 2011). “Apodized Lyot Coronagraph for SPHERE/VLT: II. Laboratory Tests and Performance”. In: *Experimental Astronomy* 30, pp. 59–81. DOI: [10.1007/s10686-011-9220-y](https://doi.org/10.1007/s10686-011-9220-y).
- Guhaniyogi, R., S. Qamar, and D. B. Dunson (2017). “Bayesian Tensor Regression”. In: *Journal of Machine Learning Research* 18.79, pp. 1–31.
- Guo, K., L. Liu, X. Xu, D. Xu, and D. Tao (June 2018). “GoDec+: Fast and Robust Low-Rank Matrix Decomposition Based on Maximum Correntropy”. In: *IEEE Transactions on Neural Networks and Learning Systems* 29.6, pp. 2323–2336. DOI: [10.1109/TNNLS.2016.2643286](https://doi.org/10.1109/TNNLS.2016.2643286).
- Guyon, O., B. Norris, M.-A. Martinod, K. Ahn, V. Deo, et al. (Aug. 2022). “High Contrast Imaging at the Photon Noise Limit with WFS-based PSF Calibration”. In: *Adaptive Optics Systems VIII*. Vol. 12185. SPIE, pp. 101–112. DOI: [10.1117/12.2630503](https://doi.org/10.1117/12.2630503).
- Haffert, S. Y., A. J. Bohn, J. de Boer, I. a. G. Snellen, J. Brinchmann, et al. (Aug. 2019). “Two Accreting Protoplanets around the Young Star PDS 70”. In: *Nat Astron* 3.8, pp. 749–754. DOI: [10.1038/s41550-019-0780-5](https://doi.org/10.1038/s41550-019-0780-5).
- Halko, N., P. G. Martinsson, and J. A. Tropp (Jan. 2011). “Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions”. In: *SIAM Rev.* 53.2, pp. 217–288. DOI: [10.1137/090771806](https://doi.org/10.1137/090771806).
- Hara, N. C. and E. B. Ford (Mar. 2023). “Statistical Methods for Exoplanet Detection with Radial Velocities”. In: *Annual Review of Statistics and Its Application* 10. Volume 10, 2023, pp. 623–649. DOI: [10.1146/annurev-statistics-033021-012225](https://doi.org/10.1146/annurev-statistics-033021-012225).
- Hardegree-Ullman, K. K., D. Apai, G. J. Bergsten, I. Pascucci, and M. López-Morales (June 2023). “Bioverse: A Comprehensive Assessment of the Capabilities of Extremely Large Telescopes to Probe Earth-like O₂ Levels in Nearby Transiting Habitable-zone Exoplanets”. In: *AJ* 165.6, p. 267. DOI: [10.3847/1538-3881/acd1ec](https://doi.org/10.3847/1538-3881/acd1ec).
- Harris, C. R., K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, et al. (Sept. 2020). “Array Programming with NumPy”. In: *Nature* 585.7825, pp. 357–362. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2).
- Helling, C. (May 2019). “Exoplanet Clouds”. In: *Annual Review of Earth and Planetary Sciences* 47. Volume 47, 2019, pp. 583–606. DOI: [10.1146/annurev-earth-053018-060401](https://doi.org/10.1146/annurev-earth-053018-060401).
- Henault, F. (Sept. 2019). *New Concepts for Calibrating Non-Common Path Aberrations in Adaptive Optics and Coronagraph Systems*. DOI: [10.48550/arXiv.1909.05224](https://doi.org/10.48550/arXiv.1909.05224). (Visited on 08/05/2024).
- Henry, G. W., G. W. Marcy, R. P. Butler, and S. S. Vogt (Dec. 1999). “A Transiting “51 Peg-like” Planet*”. In: *ApJ* 529.1, p. L41. DOI: [10.1086/312458](https://doi.org/10.1086/312458).
- Herz, P., M. Samland, and C. A. L. Bailer-Jones (Feb. 2024). “Distinguishing Exoplanet Companions from Field Stars in Direct Imaging Using Gaia Astrometry”. In: *A&A* 682, A92. DOI: [10.1051/0004-6361/202348496](https://doi.org/10.1051/0004-6361/202348496).

- Hitchcock, F. L. (1927). “The Expression of a Tensor or a Polyadic as a Sum of Products”. In: *Journal of Mathematics and Physics* 6.1-4, pp. 164–189. DOI: [10.1002/sapm192761164](https://doi.org/10.1002/sapm192761164).
- Hoeijmakers, H. J., H. Schwarz, I. a. G. Snellen, R. J. de Kok, M. Bonnefoy, et al. (Sept. 2018). “Medium-Resolution Integral-Field Spectroscopy for High-Contrast Exoplanet Imaging - Molecule Maps of the β Pictoris System with SINFONI”. In: *A&A* 617, A144. DOI: [10.1051/0004-6361/201832902](https://doi.org/10.1051/0004-6361/201832902).
- Hong, D., T. G. Kolda, and J. A. Duersch (Jan. 2020). “Generalized Canonical Polyadic Tensor Decomposition”. In: *SIAM Rev.* 62.1, pp. 133–163. DOI: [10.1137/18M1203626](https://doi.org/10.1137/18M1203626).
- Hughes, A. M., G. Duchêne, and B. C. Matthews (Sept. 2018). “Debris Disks: Structure, Composition, and Variability”. In: *Annual Review of Astronomy and Astrophysics* 56. Volume 56, 2018, pp. 541–591. DOI: [10.1146/annurev-astro-081817-052035](https://doi.org/10.1146/annurev-astro-081817-052035).
- Jahromi, S. S. and R. Orus (Jan. 2023). *Variational Tensor Neural Networks for Deep Learning*. DOI: [10.48550/arXiv.2211.14657](https://doi.org/10.48550/arXiv.2211.14657). (Visited on 12/10/2023).
- Janson, M., R. Asensio-Torres, D. André, M. Bonnefoy, P. Delorme, et al. (June 2019). “The B-Star Exoplanet Abundance Study: A Co-Moving 16–25 MJup Companion to the Young Binary System HIP 79098”. In: *A&A* 626, A99. DOI: [10.1051/0004-6361/201935687](https://doi.org/10.1051/0004-6361/201935687).
- Janson, M., R. Gratton, L. Rodet, A. Vigan, M. Bonnefoy, et al. (Dec. 2021a). “A Wide-Orbit Giant Planet in the High-Mass β Centauri Binary System”. In: *Nature* 600.7888, pp. 231–234. DOI: [10.1038/s41586-021-04124-8](https://doi.org/10.1038/s41586-021-04124-8).
- Janson, M., V. Squicciarini, P. Delorme, R. Gratton, M. Bonnefoy, et al. (Feb. 2021b). “BEAST Begins: Sample Characteristics and Survey Performance of the B-star Exoplanet Abundance Study”. In: *A&A* 646, A164. DOI: [10.1051/0004-6361/202039683](https://doi.org/10.1051/0004-6361/202039683).
- Jones, M. I., J. Milli, I. Blanchard, Z. Wahhaj, R. J. D. Rosa, et al. (Nov. 2022). “SPHERE Adaptive Optics Performance for Faint Targets”. In: *A&A* 667, A114. DOI: [10.1051/0004-6361/202243820](https://doi.org/10.1051/0004-6361/202243820).
- Joseph, D. (Jan. 2024). *Ricci-Notation Tensor Framework for Model-Based Approaches to Imaging*. DOI: [10.48550/arXiv.2312.04018](https://doi.org/10.48550/arXiv.2312.04018). (Visited on 04/08/2024).
- Kammerer, J., M. J. Ireland, F. Martinache, and J. H. Girard (June 2019). “Kernel Phase Imaging with VLT/NACO: High-Contrast Detection of New Candidate Low-Mass Stellar Companions at the Diffraction Limit”. In: *Monthly Notices of the Royal Astronomical Society* 486.1, pp. 639–654. DOI: [10.1093/mnras/stz882](https://doi.org/10.1093/mnras/stz882).
- Keppler, M., M. Benisty, A. Müller, T. Henning, R. van Boekel, et al. (Sept. 2018). “Discovery of a Planetary-Mass Companion within the Gap of the Transition Disk around PDS 70”. In: *A&A* 617, A44. DOI: [10.1051/0004-6361/201832957](https://doi.org/10.1051/0004-6361/201832957).
- Kiefer, S., A. J. Bohn, S. P. Quanz, M. A. Kenworthy, and T. Stolker (Aug. 2021). “Spectral and Angular Differential Imaging with SPHERE/IFS: Assessing the Performance of Various PCA-based Approaches to PSF Subtraction”. In: *Astronomy & Astrophysics* 652, A33. DOI: [10.1051/0004-6361/202140285](https://doi.org/10.1051/0004-6361/202140285).
- Kingma, D. P. and J. Ba (Dec. 2014). “Adam: A Method for Stochastic Optimization”. In: *CoRR* 1412.6980.
- Kley, W. and R. P. Nelson (Sept. 2012). “Planet-Disk Interaction and Orbital Evolution”. In: *Annual Review of Astronomy and Astrophysics* 50. Volume 50, 2012, pp. 211–249. DOI: [10.1146/annurev-astro-081811-125523](https://doi.org/10.1146/annurev-astro-081811-125523).
- Kolda, T. (Apr. 2006). “Multilinear Operators for Higher-Order Decompositions.” In: vol. SAND2006-2081, 923081. Sandia National Laboratories. DOI: [10/cz74dp](https://doi.org/10/cz74dp).
- Kolda, T. G. and B. W. Bader (Aug. 2009). “Tensor Decompositions and Applications”. In: *SIAM Rev.* 51.3, pp. 455–500. DOI: [10/dzcrv6](https://doi.org/10/dzcrv6).
- Kolda, T. G. and D. Hong (Jan. 2020). “Stochastic Gradients for Large-Scale Tensor Decomposition”. In: *SIAM Journal on Mathematics of Data Science* 2.4, pp. 1066–1095. DOI: [10.1137/19m1266265](https://doi.org/10.1137/19m1266265).
- Kossaifi, J., Y. Panagakis, A. Anandkumar, and M. Pantic (2019). “TensorLy: Tensor Learning in Python”. In: *Journal of Machine Learning Research* 20.26, pp. 1–6.

- Kral, Q., J. E. Pringle, L. Matrà, and P. Thébault (Jan. 2023). “Potential Effects of Stellar Winds on Gas Dynamics in Debris Disks Leading to Observable Belt Winds”. In: *A&A* 669, A116. DOI: [10.1051/0004-6361/202243729](https://doi.org/10.1051/0004-6361/202243729).
- Lacour, S., M. Nowak, J. Wang, O. Pfuhl, F. Eisenhauer, et al. (Mar. 2019). “First Direct Detection of an Exoplanet by Optical Interferometry - Astrometry and K-band Spectroscopy of HR 8799 e”. In: *A&A* 623, p. L11. DOI: [10.1051/0004-6361/201935253](https://doi.org/10.1051/0004-6361/201935253).
- Lafrenière, D., C. Marois, R. Doyon, and T. Barman (Mar. 2009). “HST/NICMOS Detection Of HR 8799 B In 1998”. In: *ApJ* 694.2, p. L148. DOI: [10.1088/0004-637X/694/2/L148](https://doi.org/10.1088/0004-637X/694/2/L148).
- Lagrange, A. .-, D. Gratadour, G. Chauvin, T. Fusco, D. Ehrenreich, et al. (Jan. 2009). “A Probable Giant Planet Imaged in the β Pictoris Disk. VLT/NaCo Deep L'-Band Imaging”. In: *Astronomy and Astrophysics* 493, pp. L21–L25. DOI: [10.1051/0004-6361:200811325](https://doi.org/10.1051/0004-6361:200811325).
- Lagrange, A. M., P. Rubini, M. Nowak, S. Lacour, A. Grandjean, et al. (Oct. 2020). “Unveiling the β Pictoris System, Coupling High Contrast Imaging, Interferometric, and Radial Velocity Data”. In: *A&A* 642, A18. DOI: [10.1051/0004-6361/202038823](https://doi.org/10.1051/0004-6361/202038823).
- Lagrange, A., M. Bonnefoy, G. Chauvin, D. Apai, D. Ehrenreich, et al. (July 2010). “A Giant Planet Imaged in the Disk of the Young Star β Pictoris”. In: *Science* 329, p. 57. DOI: [10.1126/science.1187187](https://doi.org/10.1126/science.1187187).
- Landman, R., I. a. G. Snellen, C. U. Keller, M. N'Diaye, F. Fagginger-Auer, et al. (July 2023). “Trade-Offs in High-Contrast Integral Field Spectroscopy for Exoplanet Detection and Characterisation - Young Gas Giants in Emission”. In: *A&A* 675, A157. DOI: [10.1051/0004-6361/202245169](https://doi.org/10.1051/0004-6361/202245169).
- Langlois, M., R. Gratton, A.-M. Lagrange, P. Delorme, A. Boccaletti, et al. (July 2021). “The SPHERE Infrared Survey for Exoplanets (SHINE) - II. Observations, Data Reduction and Analysis, Detection Performances, and Initial Results”. In: *A&A* 651, A71. DOI: [10.1051/0004-6361/202039753](https://doi.org/10.1051/0004-6361/202039753).
- Lathauwer, L. D. and B. Moor (1996). “From Matrix to Tensor: Multilinear Algebra and Signal Processing”. In: *Conf. Digest of the 4th IMA International Conference on Mathematics in Signal Processing*. Warwick, U.K.: IMA, pp. 1–11.
- Laurent, K. S., K. Fogarty, N. T. Zimmerman, M. N'Diaye, C. C. Stark, et al. (July 2018). “Apodized Pupil Lyot Coronagraphs Designs for Future Segmented Space Telescopes”. In: *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*, p. 102. DOI: [10.1117/12.2313902](https://doi.org/10.1117/12.2313902).
- Leboulleux, L., A. Carlotti, M. N'Diaye, F. Cantalloube, J. Milli, et al. (July 2022). “Redundant Apodized Pupils (RAP) for High-Contrast Imagers Robust to Segmentation-Due Aberrations and Island Effect”. In: *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation V*. Vol. Proceedings of the SPIE, Volume 12188. Montréal, Canada: SPIE, p. 63. DOI: [10.1117/12.2629276](https://doi.org/10.1117/12.2629276).
- Lee, C. and M. Wang (Nov. 2021). *Smooth Tensor Estimation with Unknown Permutations*. DOI: [10.48550/arXiv.2111.04681](https://doi.org/10.48550/arXiv.2111.04681). (Visited on 05/06/2024).
- Liu, J., P. Musialski, P. Wonka, and J. Ye (Jan. 2013). “Tensor Completion for Estimating Missing Values in Visual Data”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35.1, pp. 208–220. DOI: [10.1109/TPAMI.2012.39](https://doi.org/10.1109/TPAMI.2012.39).
- Liu, Y., J. Liu, Z. Long, and C. Zhu (2022). *Tensor Computation for Data Analysis*. Cham: Springer International Publishing. DOI: [10.1007/978-3-030-74386-4](https://doi.org/10.1007/978-3-030-74386-4).
- Long, J. D. and J. R. Males (Mar. 2021). “Unlocking Starlight Subtraction in Full-data-rate Exoplanet Imaging by Efficiently Updating Karhunen–Loève Eigenimages”. In: *AJ* 161.4, p. 166. DOI: [10.3847/1538-3881/abe12b](https://doi.org/10.3847/1538-3881/abe12b).
- Lothringer, J. D., D. K. Sing, Z. Rustamkulov, H. R. Wakeford, K. B. Stevenson, et al. (Apr. 2022). “UV Absorption by Silicate Cloud Precursors in Ultra-Hot Jupiter WASP-178b”. In: *Nature* 604, pp. 49–52. DOI: [10.1038/s41586-022-04453-2](https://doi.org/10.1038/s41586-022-04453-2).
- Lu, H., K. N. Plataniotis, and A. N. Venetsanopoulos (2008). “MPCA: Multilinear principal component analysis of tensor objects”. In: *IEEE transactions on Neural Networks* 19.1, pp. 18–39.

- Ma, T.-Y., H.-C. Li, R. Wang, Q. Du, X. Jia, et al. (2022). “Lightweight Tensorized Neural Networks for Hyperspectral Image Classification”. In: *IEEE Transactions on Geoscience and Remote Sensing* 60, pp. 1–16. DOI: [10.1109/TGRS.2022.3225438](https://doi.org/10.1109/TGRS.2022.3225438).
- Macintosh, B., J. Graham, D. Palmer, R. Doyon, D. Gavel, et al. (June 2006). “The Gemini Planet Imager”. In: 6272, p. 62720L. DOI: [10.1117/12.672430](https://doi.org/10.1117/12.672430).
- Macintosh, B., J. R. Graham, P. Ingraham, Q. Konopacky, C. Marois, et al. (Sept. 2014). “First Light of the Gemini Planet Imager”. In: *Proceedings of the National Academy of Sciences* 111.35, pp. 12661–12666. DOI: [10.1073/pnas.1304215111](https://doi.org/10.1073/pnas.1304215111).
- Madhusudhan, N. (Jan. 2018). “Atmospheric Retrieval of Exoplanets”. In: *Handbook of Exoplanets*, p. 104. DOI: [10.1007/978-3-319-55333-7_104](https://doi.org/10.1007/978-3-319-55333-7_104).
- (Aug. 2019). “Exoplanetary Atmospheres: Key Insights, Challenges, and Prospects”. In: *Annual Review of Astronomy and Astrophysics* 57. Volume 57, 2019, pp. 617–663. DOI: [10.1146/annurev-astro-081817-051846](https://doi.org/10.1146/annurev-astro-081817-051846).
- Mamajek, E. E. and C. P. M. Bell (Dec. 2014). “On the Age of the β Pictoris Moving Group”. In: *Monthly Notices of the Royal Astronomical Society* 445.3, pp. 2169–2180. DOI: [10.1093/mnras/stu1894](https://doi.org/10.1093/mnras/stu1894).
- Mardani, M., G. Mateos, and G. B. Giannakis (May 2015). “Subspace Learning and Imputation for Streaming Big Data Matrices and Tensors”. In: *IEEE Transactions on Signal Processing* 63.10, pp. 2663–2677. DOI: [10.1109/TSP.2015.2417491](https://doi.org/10.1109/TSP.2015.2417491).
- Marois, C., D. Lafreniere, R. Doyon, B. Macintosh, and D. Nadeau (Apr. 2006). “Angular Differential Imaging: A Powerful High-Contrast Imaging Technique”. In: *ApJ* 641.1, pp. 556–564. DOI: [10.1086/500401](https://doi.org/10.1086/500401).
- Marois, C., D. Lafrenière, B. Macintosh, and R. Doyon (Jan. 2008a). “Confidence Level and Sensitivity Limits in High-Contrast Imaging”. In: *ApJ* 673.1, p. 647. DOI: [10.1086/523839](https://doi.org/10.1086/523839).
- Marois, C., B. Macintosh, T. Barman, B. Zuckerman, I. Song, et al. (Nov. 2008b). “Direct Imaging of Multiple Planets Orbiting the Star HR 8799”. In: *Science* 322.5906, pp. 1348–1352. DOI: [10.1126/science.1166585](https://doi.org/10.1126/science.1166585).
- Marois, C., B. Zuckerman, Q. M. Konopacky, B. Macintosh, and T. Barman (Dec. 2010). “Images of a Fourth Planet Orbiting HR 8799”. In: *Nature* 468.7327, pp. 1080–1083. DOI: [10.1038/nature09684](https://doi.org/10.1038/nature09684).
- Martinez, P. (Sept. 2010). “Apodized-Pupil Lyot Coronagraphs: Multistage Designs for Extremely Large Telescopes”. In: *Astronomy and Astrophysics* 519, A61. DOI: [10.1051/0004-6361/201014818](https://doi.org/10.1051/0004-6361/201014818).
- Mawet, D., J. Milli, Z. Wahhaj, D. Pelat, O. Absil, et al. (Aug. 2014). “Fundamental Limitations of High Contrast Imaging Set by Small Sample Statistics”. In: *ApJ* 792.2, p. 97. DOI: [10/fb32](https://doi.org/10/fb32).
- Mayor, M. and D. Queloz (Nov. 1995). “A Jupiter-mass Companion to a Solar-Type Star”. In: *Nature* 378.6555, pp. 355–359. DOI: [10.1038/378355a0](https://doi.org/10.1038/378355a0).
- Mesa, D., M. Bonavita, S. Benatti, R. Gratton, S. Marino, et al. (Sept. 2022). “Constraining Masses and Separations of Unseen Companions to Five Accelerating Nearby Stars”. In: *A&A* 665, A73. DOI: [10.1051/0004-6361/202244033](https://doi.org/10.1051/0004-6361/202244033).
- Mesa, D., R. Gratton, A. Zurlo, A. Vigan, R. U. Claudi, et al. (Apr. 2015). “Performance of the VLT Planet Finder SPHERE - II. Data Analysis and Results for IFS in Laboratory”. In: *A&A* 576, A121. DOI: [10.1051/0004-6361/201423910](https://doi.org/10.1051/0004-6361/201423910).
- Miguel, Y. and A. Vazan (Jan. 2023). “Interior and Evolution of the Giant Planets”. In: *Remote Sensing* 15, p. 681. DOI: [10.3390/rs15030681](https://doi.org/10.3390/rs15030681).
- Najita, J. R., S. J. Kenyon, and B. C. Bromley (Jan. 2022). “From Pebbles and Planetesimals to Planets and Dust: The Protoplanetary Disk-Debris Disk Connection”. In: *The Astrophysical Journal* 925, p. 45. DOI: [10.3847/1538-4357/ac37b6](https://doi.org/10.3847/1538-4357/ac37b6).
- NASA Exoplanet Archive (June 2024). *Planetary Systems Composite Data*. URL: <https://catcopy.ipac.caltech.edu/doi/doi.php?id=10.26133/NEA13> (visited on 12/08/2023).

- National Aeronautics and Space Administration (Aug. 2017). *Types of Integral Field Units.Png - Wikipedia*. URL: https://commons.wikimedia.org/wiki/File:Types_of_Integral_Field_Units.png (visited on 06/24/2024).
- Novikov, A., P. Izmailov, V. Khrukov, M. Figurnov, and I. Oseledets (2020). “Tensor Train Decomposition on TensorFlow (T3F)”. In: *Journal of Machine Learning Research* 21.30, pp. 1–7.
- Novikov, A., D. Podoprikin, A. Osokin, and D. P. Vetrov (2015). “Tensorizing Neural Networks”. In: *Advances in Neural Information Processing Systems*. Vol. 28. Curran Associates, Inc.
- Nowak, M., S. Lacour, A.-M. Lagrange, P. Rubini, J. Wang, et al. (Oct. 2020). “Direct Confirmation of the Radial-Velocity Planet β Pictoris c”. In: *A&A* 642, p. L2. DOI: [10.1051/0004-6361/202039039](https://doi.org/10.1051/0004-6361/202039039).
- Oseledets, I. V. (Jan. 2011). “Tensor-Train Decomposition”. In: *SIAM J. Sci. Comput.* 33.5, pp. 2295–2317. DOI: [10/bc9qwf](https://doi.org/10/bc9qwf).
- Oseledets, I. and E. Tyrtyshnikov (Jan. 2010). “TT-cross Approximation for Multidimensional Arrays”. In: *Linear Algebra and its Applications* 432.1, pp. 70–88. DOI: [10.1016/j.laa.2009.07.024](https://doi.org/10.1016/j.laa.2009.07.024).
- Paardekooper, S., R. Dong, P. Duffell, J. Fung, F. Masset, et al. (2023). “Planet-Disk Interactions and Orbital Evolution”. In: *Protostars and Planets*. Vol. 9. Protostars and Planets VII, pp. 685–716.
- Pairet, B., F. Cantalloube, C. A. Gomez Gonzalez, O. Absil, and L. Jacques (Aug. 2019). “STIM Map: Detection Map for Exoplanets Imaging beyond Asymptotic Gaussian Residual Speckle Noise”. In: *Monthly Notices of the Royal Astronomical Society* 487.2, pp. 2262–2277. DOI: [10/gs4zcc](https://doi.org/10/gs4zcc).
- Pajarola, R. and R. Ballester-Ripoll (2016). “Tensor Decomposition Methods in Visual Computing: Tensor Decomposition Models”. IEEE Visualization Conference 2016. URL: <https://www.ifi.uzh.ch/en/vmml/links/VISTensorTutorial.html> (visited on 11/17/2023).
- Panagakis, Y., J. Kossaifi, G. G. Chrysos, J. Oldfield, M. A. Nicolaou, et al. (May 2021). “Tensor Methods in Computer Vision and Deep Learning”. In: *Proceedings of the IEEE* 109.5, pp. 863–890. DOI: [10.1109/JPROC.2021.3074329](https://doi.org/10.1109/JPROC.2021.3074329).
- Paszke, A., S. Gross, F. Massa, A. Lerer, J. Bradbury, et al. (Dec. 2019). *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. DOI: [10.48550/arXiv.1912.01703](https://doi.org/10.48550/arXiv.1912.01703). (Visited on 11/13/2023).
- Pence, W. D., L. Chiappetti, C. G. Page, R. A. Shaw, and E. Stobie (Dec. 2010). “Definition of the Flexible Image Transport System (FITS), Version 3.0”. In: *A&A* 524, A42. DOI: [10.1051/0004-6361/201015362](https://doi.org/10.1051/0004-6361/201015362).
- Perryman, M. (Aug. 2018). *The Exoplanet Handbook*. 2nd. Cambridge University Press.
- Petit, C., T. Fusco, J. Charton, D. Mouillet, P. Rabou, et al. (July 2008). “The SPHERE XAO System: Design and Performance”. In: *Adaptive Optics Systems*. Vol. 7015. SPIE, pp. 504–515. DOI: [10.1117/12.790308](https://doi.org/10.1117/12.790308).
- Pierrehumbert, R. T. and M. Hammond (Jan. 2019). “Atmospheric Circulation of Tide-Locked Exoplanets”. In: *Annual Review of Fluid Mechanics* 51. Volume 51, 2019, pp. 275–303. DOI: [10.1146/annurev-fluid-010518-040516](https://doi.org/10.1146/annurev-fluid-010518-040516).
- Por, E. H. (Jan. 2020). “Phase-Apodized-Pupil Lyot Coronagraphs for Arbitrary Telescope Pupils”. In: *ApJ* 888.2, p. 127. DOI: [10.3847/1538-4357/ab3857](https://doi.org/10.3847/1538-4357/ab3857).
- Psarras, C., L. Karlsson, J. Li, and P. Bientinesi (June 2022). *The Landscape of Software for Tensor Computations*. DOI: [10.48550/arXiv.2103.13756](https://doi.org/10.48550/arXiv.2103.13756). (Visited on 12/01/2023).
- Pueyo, L. (June 2016). “Detection and Characterization of Exoplanets Using Projections on Karhunen-Loeve Eigenimages: Forward Modeling”. In: *ApJ* 824.2, p. 117. DOI: [10.3847/0004-637X/824/2/117](https://doi.org/10.3847/0004-637X/824/2/117).
- (2018). “Direct Imaging as a Detection Technique for Exoplanets”. In: *Handbook of Exoplanets*. Ed. by H. J. Deeg and J. A. Belmonte. Cham: Springer International Publishing, pp. 1–61. DOI: [10.1007/978-3-319-30648-3_10-1](https://doi.org/10.1007/978-3-319-30648-3_10-1). (Visited on 11/17/2023).
- PyTensor (2023). *PyTensor: A Python Library for Defining, Optimizing, and Efficiently Evaluating Mathematical Expressions Involving Multi-Dimensional Arrays*. URL: <https://github.com/pymc-devs/pytensor/tree/main> (visited on 12/02/2023).

- Rakhuba, M. V. and I. V. Oseledets (Jan. 2015). “Fast Multidimensional Convolution in Low-Rank Tensor Formats via Cross Approximation”. In: *SIAM J. Sci. Comput.* 37.2, A565–A582. DOI: [10.1137/140958529](https://doi.org/10.1137/140958529).
- Raymond, S. N. and A. Morbidelli (2022). “Planet Formation: Key Mechanisms and Global Models”. In: *Demographics of Exoplanetary Systems: Lecture Notes of the 3rd Advanced School on Exoplanetary Science*. Ed. by K. Biazzo, V. Bozza, L. Mancini, and A. Sozzetti. Cham: Springer International Publishing, pp. 3–82. DOI: [10.1007/978-3-030-88124-5_1](https://doi.org/10.1007/978-3-030-88124-5_1). (Visited on 06/03/2024).
- Ren, B., L. Pueyo, L. Pueyo, L. Pueyo, Christine Chen, et al. (Jan. 2020). “Using Data Imputation for Signal Separation in High Contrast Imaging”. In: *The Astrophysical Journal*. DOI: [10/gnkz8](https://doi.org/10/gnkz8).
- Roberts, C., A. Milsted, M. Ganahl, A. Zalcman, B. Fontaine, et al. (May 2019). *TensorNetwork: A Library for Physics and Machine Learning*. DOI: [10.48550/arXiv.1905.01330](https://doi.org/10.48550/arXiv.1905.01330). (Visited on 12/02/2023).
- de la Roche, D., H. J. Hoeijmakers, and I. A. G. Snellen (Aug. 2018). “Molecule Mapping of HR8799b Using OSIRIS on Keck: Strong Detection of Water and Carbon Monoxide, but No Methane”. In: *A&A* 616, A146. DOI: [10.1051/0004-6361/201833384](https://doi.org/10.1051/0004-6361/201833384).
- Rodler, F. F. (1999). “Wavelet based 3D compression with fast random access for very large volume data”. In: *Proceedings. Seventh Pacific Conference on Computer Graphics and Applications (Cat. No. PR00293)*. IEEE, pp. 108–117.
- Al-Saffar, A. A. M., H. Tao, and M. A. Talab (Oct. 2017). “Review of Deep Convolution Neural Network in Image Classification”. In: *2017 International Conference on Radar, Antenna, Microwave, Electronics, and Telecommunications (ICRAMET)*, pp. 26–31. DOI: [10.1109/ICRAMET.2017.8253139](https://doi.org/10.1109/ICRAMET.2017.8253139).
- Samland, M., J. Bouwman, D. W. Hogg, W. Brandner, T. Henning, et al. (Feb. 2021). “TRAP: A Temporal Systematics Model for Improved Direct Detection of Exoplanets at Small Angular Separations”. In: *A&A* 646, A24. DOI: [10.1051/0004-6361/201937308](https://doi.org/10.1051/0004-6361/201937308).
- Samland, M., T. D. Brandt, J. Milli, P. Delorme, and A. Vigan (Dec. 2022). “Spectral Cube Extraction for the VLT/SPHERE IFS - Open-source Pipeline with Full Forward Modeling and Improved Sensitivity”. In: *A&A* 668, A84. DOI: [10.1051/0004-6361/202244587](https://doi.org/10.1051/0004-6361/202244587).
- Sauvage, J. .-, T. Fusco, C. Petit, S. Meimon, E. Fedrigo, et al. (July 2010). “SAXO, the eXtreme Adaptive Optics System of SPHERE: Overview and Calibration Procedure”. In: 7736, 77360F. DOI: [10.1117/12.856942](https://doi.org/10.1117/12.856942).
- Shah, P., N. Rao, and G. Tang (2015). “Sparse and Low-Rank Tensor Decomposition”. In: *Advances in Neural Information Processing Systems*. Vol. 28. Curran Associates, Inc.
- Shallue, C. J. and A. Vanderburg (Jan. 2018). “Identifying Exoplanets with Deep Learning: A Five-planet Resonant Chain around Kepler-80 and an Eighth Planet around Kepler-90”. In: *AJ* 155.2, p. 94. DOI: [10.3847/1538-3881/aa9e09](https://doi.org/10.3847/1538-3881/aa9e09).
- Showman, A. P., X. Tan, and V. Parmentier (Dec. 2020). “Atmospheric Dynamics of Hot Giant Planets and Brown Dwarfs”. In: *Space Sci Rev* 216.8, p. 139. DOI: [10.1007/s11214-020-00758-8](https://doi.org/10.1007/s11214-020-00758-8).
- Sidiropoulos, N. D., L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, et al. (July 2017). “Tensor Decomposition for Signal Processing and Machine Learning”. In: *IEEE Transactions on Signal Processing* 65.13, pp. 3551–3582. DOI: [10.1109/TSP.2017.2690524](https://doi.org/10.1109/TSP.2017.2690524).
- Sinquin, B. and M. Verhaegen (Sept. 2018). “Tensor-Based Predictive Control for Extremely Large-Scale Single Conjugate Adaptive Optics”. In: *J. Opt. Soc. Am. A, JOSAA* 35.9, pp. 1612–1626. DOI: [10.1364/JOSAA.35.001612](https://doi.org/10.1364/JOSAA.35.001612).
- Smith, B. A. and R. J. Terrell (Dec. 1984). “A Circumstellar Disk Around β Pictoris”. In: *Science* 226.4681, pp. 1421–1424. DOI: [10.1126/science.226.4681.1421](https://doi.org/10.1126/science.226.4681.1421).
- Smith, W. H. (Dec. 1987). “Spectral Differential Imaging Detection Of Planets About Nearby Stars.” In: *PASP* 99.622, p. 1344. DOI: [10.1086/132124](https://doi.org/10.1086/132124).

- Sobral, A., T. Bouwmans, and E. H. Zahzah (Nov. 2023). *LRSLibrary: Low-Rank and Sparse Tools for Background Modeling and Subtraction in Videos*. CRC Press, Taylor and Francis Group. URL: <https://github.com/andrewssobral/lrslibrary> (visited on 11/29/2023).
- Sørensen, M. and N. D. Sidiropoulos (June 2023). “Radio-Astronomy Imaging and Interference Excision Using Tensor Decomposition and Canonical Correlation Analysis”. In: *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. DOI: [10.1109/ICASSP49357.2023.10096197](https://doi.org/10.1109/ICASSP49357.2023.10096197).
- Soummer, R., L. Pueyo, and J. Larkin (Aug. 2012). “Detection and Characterization of Exoplanets and Disks Using Projections on Karhunen-Loeve Eigenimages”. In: *ApJ* 755.2, p. L28. DOI: [10.1088/2041-8205/755/2/L28](https://doi.org/10.1088/2041-8205/755/2/L28).
- Sozzetti, A. (Sept. 2005). “Astrometric Methods and Instrumentation to Identify and Characterize Extrasolar Planets: A Review”. In: *PASP* 117.836, p. 1021. DOI: [10.1086/444487](https://doi.org/10.1086/444487).
- Squicciarini, V., R. Gratton, M. Janson, E. E. Mamajek, G. Chauvin, et al. (Aug. 2022). “A Scaled-up Planetary System around a Supernova Progenitor”. In: *A&A* 664, A9. DOI: [10.1051/0004-6361/202243675](https://doi.org/10.1051/0004-6361/202243675).
- Su, K. Y. L., G. H. Rieke, J. A. Stansberry, G. Bryden, K. R. Stapelfeldt, et al. (Dec. 2006). “Debris Disk Evolution around A Stars”. In: *ApJ* 653.1, p. 675. DOI: [10.1086/508649](https://doi.org/10.1086/508649).
- Su, K. Y. L., G. H. Rieke, K. R. Stapelfeldt, R. Malhotra, G. Bryden, et al. (Oct. 2009). “The Debris Disk Around HR 8799”. In: *ApJ* 705.1, p. 314. DOI: [10.1088/0004-637X/705/1/314](https://doi.org/10.1088/0004-637X/705/1/314).
- Tang, Y.-H. and E. Rebrova (Nov. 2023). *FunFact: Build Your Own Tensor Decomposition Model in a Breeze*. Lawrence Berkeley National Laboratory (LBNL). URL: <https://github.com/yhtang/FunFact> (visited on 12/01/2023).
- Thanh, L. T., K. Abed-Meraim, N. L. Trung, and A. Hafiane (Nov. 2023). “A Contemporary and Comprehensive Survey on Streaming Tensor Decomposition”. In: *IEEE Transactions on Knowledge and Data Engineering* 35.11, pp. 10897–10921. DOI: [10.1109/TKDE.2022.3230874](https://doi.org/10.1109/TKDE.2022.3230874).
- Tharwat, A. (Jan. 2020). “Classification Assessment Methods”. In: *Applied Computing and Informatics* 17.1, pp. 168–192. DOI: [10.1016/j.aci.2018.08.003](https://doi.org/10.1016/j.aci.2018.08.003).
- The Tensor Network* (2023). URL: <https://www.tensornetwork.org> (visited on 11/15/2023).
- Tomioka, R., T. Suzuki, K. Hayashi, and H. Kashima (2011). “Statistical Performance of Convex Tensor Decomposition”. In: *Advances in Neural Information Processing Systems*. Vol. 24. Curran Associates, Inc.
- Tucker, L. R. (Sept. 1966). “Some Mathematical Notes on Three-Mode Factor Analysis”. In: *Psychometrika* 31.3, pp. 279–311. DOI: [10.1007/BF02289464](https://doi.org/10.1007/BF02289464).
- Tyson, R. K. and B. W. Frazier (Feb. 2022). *Principles of Adaptive Optics*. 5th ed. Boca Raton: CRC Press. DOI: [10.1201/9781003140191](https://doi.org/10.1201/9781003140191).
- Usvyatsov, M., R. Ballester-Ripoll, and K. Schindler (2022). “Tntorch: Tensor Network Learning with PyTorch”. In: *Journal of Machine Learning Research* 23.208, pp. 1–6.
- Vary, S., H. Daglayan, L. Jacques, and P.-A. Absil (June 2023). “Low-Rank Plus Sparse Trajectory Decomposition for Direct Exoplanet Imaging”. In: *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. DOI: [10.1109/ICASSP49357.2023.10096197](https://doi.org/10.1109/ICASSP49357.2023.10096197).
- Vasilescu, M. A. O. and D. Terzopoulos (2002). “Multilinear Analysis of Image Ensembles: TensorFaces”. In: *Computer Vision — ECCV 2002*. Ed. by A. Heyden, G. Sparr, M. Nielsen, and P. Johansen. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 447–460. DOI: [10.1007/3-540-47969-4_30](https://doi.org/10.1007/3-540-47969-4_30).
- Vaughan, S. R., T. D. Gebhard, K. Bott, S. L. Casewell, N. B. Cowan, et al. (Oct. 2023). “Chasing Rainbows and Ocean Glints: Inner Working Angle Constraints for the Habitable Worlds Observatory”. In: *Monthly Notices of the Royal Astronomical Society* 524.4, pp. 5477–5485. DOI: [10.1093/mnras/stad2127](https://doi.org/10.1093/mnras/stad2127).

- Vervliet, N., O. Debals, and L. De Lathauwer (Nov. 2016). “Tensorlab 3.0 — Numerical Optimization Strategies for Large-Scale Constrained and Coupled Matrix/Tensor Factorization”. In: *2016 50th Asilomar Conference on Signals, Systems and Computers*, pp. 1733–1738. DOI: [10.1109/ACSSC.2016.7869679](https://doi.org/10.1109/ACSSC.2016.7869679).
- Vervliet, N., O. Debals, L. Sorber, and L. De Lathauwer (Sept. 2014). “Breaking the Curse of Dimensionality Using Decompositions of Incomplete Tensors: Tensor-based Scientific Computing in Big Data Analysis”. In: *IEEE Signal Processing Magazine* 31.5, pp. 71–79. DOI: [10.1109/MSP.2014.2329429](https://doi.org/10.1109/MSP.2014.2329429).
- Vigan, A., M. Bonavita, B. Biller, D. Forgan, K. Rice, et al. (July 2017). “The VLT/NaCo Large Program to Probe the Occurrence of Exoplanets and Brown Dwarfs at Wide Orbits - IV. Gravitational Instability Rarely Forms Wide, Giant Planets”. In: *A&A* 603, A3. DOI: [10.1051/0004-6361/201630133](https://doi.org/10.1051/0004-6361/201630133).
- Vigan, A., C. Fontanive, M. Meyer, B. Biller, M. Bonavita, et al. (July 2021). “The SPHERE Infrared Survey for Exoplanets (SHINE) - III. The Demographics of Young Giant Exoplanets below 300 Au with SPHERE”. In: *A&A* 651, A72. DOI: [10.1051/0004-6361/202038107](https://doi.org/10.1051/0004-6361/202038107).
- Viswanath, G., M. Janson, R. Gratton, V. Squicciarini, L. Rodet, et al. (July 2023). “BEAST Detection of a Brown Dwarf and a Low-Mass Stellar Companion around the Young Bright B Star HIP 81208”. In: *A&A* 675, A54. DOI: [10.1051/0004-6361/202346154](https://doi.org/10.1051/0004-6361/202346154).
- Vos, J. M., B. Burningham, J. K. Faherty, S. Alejandro, E. Gonzales, et al. (Feb. 2023). “Patchy Forsterite Clouds in the Atmospheres of Two Highly Variable Exoplanet Analogs”. In: *The Astrophysical Journal* 944, p. 138. DOI: [10.3847/1538-4357/acab58](https://doi.org/10.3847/1538-4357/acab58).
- Wahhaj, Z., D. W. Koerner, M. E. Ressler, M. W. Werner, D. E. Backman, et al. (Jan. 2003). “The Inner Rings of β Pictoris”. In: *ApJ* 584.1, p. L27. DOI: [10.1086/346123](https://doi.org/10.1086/346123).
- Wang, A., Z. Jin, and G. Tang (Feb. 2020). “Robust Tensor Decomposition via T-SVD: Near-optimal Statistical Guarantee and Scalable Algorithms”. In: *Signal Processing* 167, p. 107319. DOI: [10.1016/j.sigpro.2019.107319](https://doi.org/10.1016/j.sigpro.2019.107319).
- Wang, A., Q. Zhao, Z. Jin, C. Li, and G. Zhou (June 2022a). “Robust Tensor Decomposition via Orientation Invariant Tubal Nuclear Norms”. In: *Sci. China Technol. Sci.* 65.6, pp. 1300–1317. DOI: [10.1007/s11431-021-1976-2](https://doi.org/10.1007/s11431-021-1976-2).
- Wang, M., D. Hong, Z. Han, J. Li, J. Yao, et al. (Mar. 2023). “Tensor Decompositions for Hyperspectral Data Processing in Remote Sensing: A Comprehensive Review”. In: *IEEE Geoscience and Remote Sensing Magazine* 11.1, pp. 26–72. DOI: [10.1109/MGRS.2022.3227063](https://doi.org/10.1109/MGRS.2022.3227063).
- Wang, Y. and Y. Yang (Apr. 2022b). *Hot-SVD: Higher-Order t-Singular Value Decomposition for Tensors Based on Tensor-Tensor Product*. DOI: [10.48550/arXiv.2204.10229](https://doi.org/10.48550/arXiv.2204.10229). (Visited on 11/17/2023).
- Wei, K. and Y. Fu (Feb. 2019). “Low-Rank Bayesian Tensor Factorization for Hyperspectral Image Denoising”. In: *Neurocomputing* 331, pp. 412–423. DOI: [10.1016/j.neucom.2018.10.023](https://doi.org/10.1016/j.neucom.2018.10.023).
- Wenger, M., F. Ochsenbein, D. Egret, P. Dubois, F. Bonnarel, et al. (Apr. 2000). “The SIMBAD Astronomical Database. The CDS Reference Database for Astronomical Objects”. In: *Astronomy and Astrophysics Supplement Series* 143, pp. 9–22. DOI: [10.1051/aas:2000332](https://doi.org/10.1051/aas:2000332).
- Wolf, T. N., B. A. Jones, and B. P. Bowler (Dec. 2023). *Direct Exoplanet Detection Using Deep Convolutional Image Reconstruction (ConStruct): A New Algorithm for Post-Processing High-Contrast Images*. DOI: [10.48550/arXiv.2312.03671](https://doi.org/10.48550/arXiv.2312.03671). (Visited on 12/14/2023).
- Wolszczan, A. and D. A. Frail (Jan. 1992). “A Planetary System around the Millisecond Pulsar PSR1257 + 12”. In: *Nature* 355.6356, pp. 145–147. DOI: [10.1038/355145a0](https://doi.org/10.1038/355145a0).
- Wordsworth, R. and L. Kreidberg (Sept. 2022). “Atmospheres of Rocky Exoplanets”. In: *Annual Review of Astronomy and Astrophysics* 60. Volume 60, 2022, pp. 159–201. DOI: [10.1146/annurev-astro-052920-125632](https://doi.org/10.1146/annurev-astro-052920-125632).
- Wyatt, M. C. (Sept. 2008). “Evolution of Debris Disks”. In: *Annual Review of Astronomy and Astrophysics* 46. Volume 46, 2008, pp. 339–383. DOI: [10.1146/annurev.astro.45.051806.110525](https://doi.org/10.1146/annurev.astro.45.051806.110525).

- Xie, C., E. Choquet, A. Vigan, F. Cantalloube, M. Benisty, et al. (Oct. 2022). “Reference-Star Differential Imaging on SPHERE/IRDIS”. In: *Astronomy and Astrophysics* 666, A32. DOI: [10.1051/0004-6361/202243379](https://doi.org/10.1051/0004-6361/202243379).
- Yeo, B.-L. and B. Liu (1995). “Volume rendering of DCT-based compressed 3D scalar data”. In: *IEEE Transactions on Visualization and Computer Graphics* 1.1, pp. 29–43.
- Yu, R., D. Cheng, and Y. Liu (June 2015). “Accelerated Online Low Rank Tensor Learning for Multivariate Spatiotemporal Streams”. In: *Proceedings of the 32nd International Conference on Machine Learning*. PMLR, pp. 238–247.
- Zhang, A. and R. Han (Sept. 2018a). *Optimal Sparse Singular Value Decomposition for High-dimensional High-order Data*. DOI: [10.48550/arXiv.1809.01796](https://doi.org/10.48550/arXiv.1809.01796). (Visited on 11/28/2023).
- Zhang, A. and D. Xia (Nov. 2018b). “Tensor SVD: Statistical and Computational Limits”. In: *IEEE Transactions on Information Theory* 64.11, pp. 7311–7338. DOI: [10.1109/TIT.2018.2841377](https://doi.org/10.1109/TIT.2018.2841377).
- Zhao, Q., G. Zhou, S. Xie, L. Zhang, and A. Cichocki (June 2016). *Tensor Ring Decomposition*. DOI: [10.48550/arXiv.1606.05535](https://doi.org/10.48550/arXiv.1606.05535). (Visited on 11/13/2023).
- Zhou, P., C. Lu, Z. Lin, and C. Zhang (Mar. 2018). “Tensor Factorization for Low-Rank Tensor Completion”. In: *IEEE Transactions on Image Processing* 27.3, pp. 1152–1163. DOI: [10.1109/TIP.2017.2762595](https://doi.org/10.1109/TIP.2017.2762595).
- Zhou, T. and D. Tao (June 2011). “GoDec: Randomized Lowrank & Sparse Matrix Decomposition in Noisy Case”. In: *International Conference on Machine Learning*.
- Zimmerman, N. T., A. J. E. Riggs, N. J. Kasdin, A. Carlotti, and R. J. Vanderbei (Jan. 2016). “Shaped Pupil Lyot Coronagraphs: High-Contrast Solutions for Restricted Focal Planes”. In: *J. Astron. Telesc. Instrum. Syst* 2.1, p. 011012. DOI: [10.1117/1.JATIS.2.1.011012](https://doi.org/10.1117/1.JATIS.2.1.011012).
- Zurlo, A., A. Vigan, D. Mesa, R. Gratton, C. Moutou, et al. (Dec. 2014). “Performance of the VLT Planet Finder SPHERE - I. Photometry and Astrometry Precision with IRDIS and IFS in Laboratory”. In: *A&A* 572, A85. DOI: [10.1051/0004-6361/201424204](https://doi.org/10.1051/0004-6361/201424204).

Algorithms

This appendix contains additional algorithms to further detail the specific methods developed or used in this project.

Algorithm 4: PSF modeling using principal component analysis by singular value decomposition. (Amara et al., 2012; Soummer et al., 2012)

Input: scaled observation cube $\mathcal{X}' \in \mathbb{R}^{I_\lambda \times I_\theta \times I_x \times I_y}$,
number of PCA components N_{comp}
Output: scaled PSF model $\mathcal{M}_{PSF} \in \mathbb{R}^{I_\lambda \times I_\theta \times I_x \times I_y}$

- 1 **Initialize:**
- 2 **tSVD** : $\mathbb{R}^{(I_\lambda \cdot I_\theta) \times (I_x \cdot I_y)} \rightarrow \mathbb{R}^{(I_\lambda \cdot I_\theta) \times N_{comp}}, \mathbb{R}^{N_{comp} \times N_{comp}}, \mathbb{R}^{N_{comp} \times (I_x \cdot I_y)}$ // truncated SVD
- 3 $\mathbf{X}' \leftarrow \text{unfold}_{(1,2),(3,4)}(\mathcal{X}')$ // unfold the scaled observation cube: $\mathbf{X}' \in \mathbb{R}^{(I_\lambda \cdot I_\theta) \times (I_x \cdot I_y)}$
- 4 $\mathbf{U}, \Sigma, \mathbf{V} \leftarrow \text{tSVD}(\mathbf{X}')$ // compute the truncated SVD of the scaled observation matrix
- 5 $\tilde{\mathbf{X}}' \leftarrow \mathbf{X}' \mathbf{V}_{N_{comp}}$ // project \mathbf{X}' onto the space spanned by the first N_{comp} of the SVD
- 6 $\mathbf{X}' \leftarrow \tilde{\mathbf{X}}' \mathbf{V}_{N_{comp}}^T$ // reconstruct the reduced \mathbf{X}' to get the PSF model
- 7 $\mathcal{M}_{PSF} \leftarrow \text{fold}_{(1,2),(3,4)}(\mathbf{X}')$ // fold the PSF model back to its original shape
- 8 **return** \mathcal{M}_{PSF}

Algorithm 5: PSF modeling using the rank- R canonical polyadic decomposition via alternating least squares (Kolda et al., 2009).

Input: scaled observation cube $\mathcal{X}' \in \mathbb{R}^{I_\lambda \times I_\theta \times I_x \times I_y}$,
rank R ,
error threshold ϵ
Output: scaled PSF model $\mathcal{M}_{PSF} \in \mathbb{R}^{I_\lambda \times I_\theta \times I_x \times I_y}$

- 1 **Initialize:**
- 2 $\mathbf{U}^{(n)} \in \mathbb{R}^{I_n \times R}$ for $n = 1, 2, 3, 4$ // initialize factors, e.g. random
- 3 $\lambda \leftarrow \mathbf{1} \in \mathbb{R}^R$ // initialize normalization factor
- 4 **repeat**
- 5 **for** $n = 1, 2, \dots, N$ **do**
- 6 $\mathbf{V} \leftarrow \prod_{\forall i \in \{1,2,3,4\} \setminus \{n\}}^* [\mathbf{U}^{(i)T} \mathbf{U}^{(i)}]$ // Hadamard product
- 7 $\mathbf{U}^{(n)} \leftarrow \mathbf{X}^{(n)} \prod_{\forall i \in \{4,3,2,1\} \setminus \{n\}}^\circ [\mathbf{U}^{(i)}] \mathbf{V}^\dagger$ // find factor using Khatri-Rao product
- 8 $\lambda \leftarrow \left[\|\mathbf{u}_1^{(n)}\| \quad \|\mathbf{u}_2^{(n)}\| \quad \dots \quad \|\mathbf{u}_R^{(n)}\| \right]^T$ // find normalization from columns of $\mathbf{U}^{(n)}$
- 9 $\mathbf{U}^{(n)} \leftarrow \mathbf{U}^{(n)} \lambda^{-1}$ // normalize $\mathbf{U}^{(n)}$
- 10 **end**
- 11 $e \leftarrow \|\mathcal{X}' - \llbracket \lambda; \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \mathbf{U}^{(3)}, \mathbf{U}^{(4)} \rrbracket\|$ // calculate error
- 12 **until** $\epsilon > e$
- 13 $\mathcal{M}_{PSF} \leftarrow \llbracket \lambda; \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \mathbf{U}^{(3)}, \mathbf{U}^{(4)} \rrbracket$
- 14 **return** \mathcal{M}_{PSF}

Algorithm 6: Higher-order singular value decomposition. (Kolda et al., 2009)

Input: tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$,
rank $\{R_1, R_2, \dots, R_N\}$,
error threshold ϵ

Output: Tucker core \mathcal{G} and factors $\mathbf{U}^{(n)}$

1 **Initialize:**

2 $\mathbf{U}^{(n)} \leftarrow R_n$ left leading singular vectors of $\mathbf{X}_{(n)}$ for $n = 1, \dots, N$ // initialize factors

3 $\mathcal{G} \leftarrow \mathcal{X} \times_1 \prod_{i \in \mathcal{N}} \left[\mathbf{U}^{(i)T} \right]$ // find Tucker core using mode-n product

4 **return** $\mathcal{G}, \mathbf{U}^{(n)}$ for $n = 1, \dots, N$

Algorithm 7: Higher-order orthogonal iteration computes a Tucker decomposition using alternating least squares. (Kolda et al., 2009)

Input: tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$,
rank $\{R_1, R_2, \dots, R_N\}$,
error threshold ϵ

Output: Tucker core \mathcal{G} and factors $\mathbf{U}^{(n)}$

1 **Initialize:**

2 $\mathcal{N} \leftarrow \{1, 2, \dots, N\}$

3 $\mathbf{U}^{(n)} \in \mathbb{R}^{I_n \times R_n}$ for $n = 1, \dots, N \leftarrow \text{HOSVD}(\mathcal{X})$ using [algorithm 6](#) // initialize factors

4 **repeat**

5 **for** $n = 1, 2, \dots, N$ **do**

6 $\mathbf{Y} \leftarrow \mathcal{X} \times_1 \prod_{i \in \mathcal{N} \setminus \{n\}} \left[\mathbf{U}^{(i)T} \right]$ // mode-n product

7 $\mathbf{U}^{(n)} \leftarrow R_n$ left leading singular vectors of $\mathbf{Y}_{(n)}$ for $n = 1, \dots, N$ // update factor

8 **end**

9 $e \leftarrow \left\| \mathcal{X}' - \llbracket \boldsymbol{\lambda}; \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)} \rrbracket \right\|$ // calculate error

10 **until** $\epsilon > e$

11 $\mathcal{G} \leftarrow \mathcal{X} \times_1 \prod_{i \in \mathcal{N}} \left[\mathbf{U}^{(i)T} \right]$ // find Tucker core using mode-n product

12 **return** $\mathcal{G}, \mathbf{U}^{(n)}$ for $n = 1, \dots, N$

Observation Data

This appendix lists the objects used for evaluating the algorithms in [Section 5.3](#) and the available observations of HD 108767 B.

B.1. HD 108767 B Observations

Table B.1: HD 108767 B: Publicly available ESO observations.

Instrument	Filter	Exposure Time (h:min:s)	No. of Observations
CRIRES	total	-	90
	HX1E-3	0:2:40	-
	HX5E-2	0:27:20	-
EFOSC	OIII#687	0:0:0	4
FEROS	-	0:5:10	3
GRAVITY	total	-	7
	IN,MED,IN,5MAG	0:21:4	-
	OUT,HIGH,OUT,5MAG	0:8:54	-
HARPS	-	7:0:3	28
ISAAC	SL,OPEN	0:0:29	36
NAOS+CONICA	L_PRIME	0:20:44	62
OMEGACAM	total	-	20
	G_SDSS	0:3:20	-
	I_SDSS	0:3:0	-
	R_SDSS	0:4:30	-
	U_SDSS	0:4:0	-
	Z_SDSS	0:1:30	-
PIONIER	H,FREE	0:7:41	18
SOFI	total	-	6
	GBF,OPEN	0:0:47	6
	KS,OPEN	0:4:0	-
SPHERE	total	-	307
	SPHERE	0:53:36	-
	B_H,D_H23	0:51:16	-
	B_J	0:1:50	-
Total		10:41:54	581

B.2. Assessment Systems & Observations

This section gives an overview of the sample of stellar systems and observations for the assessment of the post-processing methods. The sample includes a diverse range of spectral types (B, A, F, and K) and observational conditions, providing a comprehensive sample for algorithm evaluation. The selected observations exhibit vary over characteristics such as parallactic angle rotation, observation time, seeing conditions, and coherence time, allowing for a thorough examination of method performance under different scenarios. To ensure consistency, all observations were processed using the updated pre-processing pipelines of the HC-DC from the BEAST survey (see [Section 4.2.2](#)). The observations are summarized in [Table B.2](#). Notable inclusions are systems with debris disks (HIP 88399, HIP 95261) where previous in-depth analysis found no sub-stellar companions in the IFS field of view and those previously studied in other surveys. This increases confidence about the absence of sub-stellar companions. The sample balances typical cases (CD-54 8168, HIP 82430) with more challenging ones (HIP 12394, HIP 95261), ensuring a robust assessment of the new post-processing techniques.

Table B.2: Objects used for synthetic injections for assessing the performance of algorithms. Additional data collected from SIMBAD ([Wenger et al., 2000](#)). PID: process ID, PA: parallactic angle, Cor.: coronagraph.

Main ID	Type	Spec. Type	Obs. date	PID	Band	IR filter	Cor.	Total Obs. Time [s]	Average Air-mass [-]	Average Seeing ["]	Average Coherence Time [s]	Pre-Proc. Date	Total PA var. [deg]	ESO program	Motivation
CD-54 8168	HighPM*	K2V	2018-04-18	480356	YJ	DB_H23	N_ALC_YJH_S	978	1.16	0.64	0.0134	2023-01-21	36.95	1100.C-0481(F)	K star, average conditions
HD 159911	SB*	K7	2016-05-21	468187	YJ	DB_H23	N_ALC_YJH_S	329	1.03	0.5	0.0058	2023-01-20	292.193	097.C-0865(B)	K7 star, extreme field rotation
HIP 12394	HighPM*	B9V	2017-09-30	476533	H	DB_K12	N_ALC_YJH_S	51	1.44	0.9	0.0022	2020-12-01	17.332	099.C-0841(A)	Challenging case
HIP 59505	Star	A9V	2018-03-29	480019	YJ	DB_H23	N_ALC_YJH_S	978	1.17	0.45	0.0135	2022-02-18	48.011	1100.C-0481(E)	Optimal observing conditions
HIP 61468	HighPM*	A7V	2018-06-18	480848	YJ	DB_H23	N_ALC_YJH_S	395	1.04	1.91	0.0016	2022-11-22	41.465	1100.C-0481(H)	Poor seeing, good coherence
HIP 82430	Star	B9V	2017-05-17	475430	YJ	DB_H23	N_ALC_YJH_S	395	1.04	0.92	0.0025	2022-02-18	68.881	198.C-0209(G)	Typical case, excellent coherence
HIP 88399	HighPM*	F6V	2018-04-10	480143	YJ	DB_H23	N_ALC_YJH_S	978	1.13	0.52	0.0056	2022-02-18	32.778	1100.C-0481(F)	Debris disk, known candidates (outside IFS FOV)
HIP 95261	HighPM*	A0V	2017-06-14	475733	YJ	DB_H23	N_ALC_YJH_S	406	1.23	1.42	0.0023	2022-11-21	2.355	198.C-0209(H)	Very challenging, minimal field rotation, debris disk

CD-54 8168

This K2 star observation represents an average case with typical parallactic angle variation, making it valuable for assessing the general performance of post-processing methods. Its inclusion in the Bioverse catalogue ([Hardegree-Ulman et al., 2023](#)) adds significance. The Bioverse study, which assessed the likelihood of detecting Earth-like O2 levels on nearby exoplanets, provides a comprehensive framework for evaluating potential biosignatures.

HD 159911

As a K7 star observation with very large total parallactic angle variation and average observation time, HD 159911 presents an excellent opportunity to evaluate how post-processing methods perform under very good conditions. This case can help identify strengths or limitations of new techniques in handling substantial field rotation.

HIP 12394

This observation poses a challenging case due to its small total parallactic angle variation and very short observation time. It was previously used in a similar assessment by (Chomez et al., 2023a) as a test star, ruling out potential companions.

HIP 59505

With average parallactic angle variation, excellent seeing conditions, and long integration time, this observation of HIP 59505 provides an ideal scenario for testing the performance of post-processing methods under optimal observing conditions.

HIP 61468

This observation offers an interesting combination of average parallactic angle variation, poor seeing, but excellent average coherence time. It is included in the Bioverse catalogue. This case allows for evaluation of how the post-processing methods handle suboptimal seeing conditions while benefiting from good coherence time for the adaptive optics correction.

HIP 82430

This observation also has an excellent coherence time benefiting the adaptive optics correction. The otherwise typical conditions, provide a reference for comparison with more extreme cases.

HIP 88399

This system is particularly valuable due to its debris disk and known candidates in the IRDIS field of view (Gratton et al., 2024; Mesa et al., 2022). However, there are no known candidates in the IFS field of view. Its inclusion in multiple studies (Hardegree-Ullman et al., 2023, Chomez et al., 2023a, as well as disk-focused studies) makes it a well-characterized target. The presence of a debris disk allows for evaluation of how new methods handle extended structures, while the absence of confirmed planets in differential imaging and RV studies provides a good test star for comparison.

HIP 95261

With a debris disk (Kral et al., 2023), very small total parallactic angle variation, but excellent average coherence time, the observation of HIP 95261 presents a the most challenging case in this sample, especially at small separations. Its inclusion in the AutoRSM sample and SHARDDS survey add confidence that there are no detectable substellar-companions (Dahlqvist et al., 2022). The SHARDDS survey, utilizing the Auto-RSM framework, performs in-depth analysis of its sample. This system allows for evaluation of new post-processing methods' performance in detecting planets at small angular separations under minimal field rotation conditions.

This selection of targets provides a representative dataset for evaluating the strengths and limitations of the new tensor-based methods across a range of realistic observational scenarios.

B.3. Full Observation Sample

This thesis reduced a comprehensive sample comprising 744 observation datasets of 180 target systems, totaling over 800 GB of pre-processed data. The sample includes observations from both the SHINE (Chauvin et al., 2017) and BEAST (Janson et al., 2019, 2021b) surveys, pre-processed and made available through the High Contrast Data Center (HC-DC) (Beuzit et al., 2019). For each observation, the most up-to-date pre-processing result was selected (Delorme, 2024). The observations in the sample were post-processed using a tensor-train PSF model, and automatically detected candidates were cross-referenced with candidate lists from previous post-processing efforts to ensure consistency and completeness (Chomez et al., 2023b; Desidera et al., 2021; Gratton et al., 2023; Herz et al., 2024; Janson et al., 2021a,b; Squicciarini et al., 2022; Vigan et al., 2017; Viswanath et al., 2023). The final residual images *did not undergo in-depth manual inspection* and *sensitivity maps were not constructed* for each residual, however, each residual image was visually inspected. Table B.3 provides a comprehensive list of the observations processed in this work.

Table B.3: Full sample of observations processed in this thesis. The sample consists of SHINE (Chauvin et al., 2017) and BEAST (Janson et al., 2019, 2021b) observations. The observations were pre-processed and made available by the HC-DC (Beuzit et al., 2019). For each observation the most up-to-date pre-processing result was chosen (Delorme, 2024). The dataset can be obtained from the HC-DC upon request. PID: HC-DC pre-processing process ID.

Object	Obs. date	PID	Instr.
2MASSJ09053033-4918382	2015-10-24	531322	IRDIS
GaiaDR31750318749306157824	2018-05-12	480775	IFS
GaiaDR31750318749306157824	2018-05-12	480566	IRDIS
GaiaDR31750318749306157824	2018-05-12	388934	IRDIS
GaiaDR31750318749306157824	2016-09-16	476400	IRDIS
GaiaDR31750318749306157824	2015-06-05	466626	IRDIS
GaiaDR32885863400349980288	2019-11-27	485502	IFS
GaiaDR32885863400349980288	2019-03-05	484881	IFS
GaiaDR32885863400349980288	2019-03-05	482208	IRDIS
GaiaDR32885863400349980288	2018-01-24	478881	IFS
GaiaDR32885863400349980288	2018-01-24	479876	IRDIS
GaiaDR32885863400349980288	2017-02-07	477130	IRDIS
GaiaDR32885863400349980288	2016-02-18	466698	IFS
GaiaDR32885863400349980288	2016-02-18	474512	IRDIS
GaiaDR32885863400349980288	2016-01-18	466418	IFS
GaiaDR32885863400349980288	2016-01-18	474502	IRDIS
GaiaDR32885863400349980288	2015-11-28	467921	IRDIS
GaiaDR32962658549474035584	2017-02-06	476992	IRDIS
GaiaDR32980822034888930176	2016-01-17	416759	IRDIS
GaiaDR33363049071685181056	2017-02-09	477180	IRDIS
GaiaDR33363049071685181056	2015-02-06	469039	IRDIS
GaiaDR34069167258796371712	2016-07-20	481737	IRDIS
GaiaDR34069167258796371712	2016-07-20	481735	IRDIS
GaiaDR34069167258796371712	2016-07-20	481734	IRDIS
GaiaDR34069167258796371712	2016-07-20	476254	IRDIS
GaiaDR34082257533151797760	2016-07-20	481740	IRDIS
GaiaDR34082257533151797760	2016-07-20	481739	IRDIS
GaiaDR34082257533151797760	2016-07-20	476255	IRDIS
GaiaDR34216298411563165440	2015-09-24	531254	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
GaiaDR34222102218052861440	2015-05-14	468107	IRDIS
GaiaDR34289366113217575424	2018-06-18	480849	IFS
GaiaDR34289366113217575424	2018-06-18	480630	IRDIS
GaiaDR34289366113217575424	2016-05-22	468275	IFS
GaiaDR34289366113217575424	2016-05-22	475724	IRDIS
GaiaDR34289366113217575424	2016-05-22	299676	IRDIS
GaiaDR34289366113217575424	2015-05-30	468428	IRDIS
GaiaDR34632037879874910080	2015-09-27	482646	IRDIS
GaiaDR34707563810327288192	2017-11-05	478349	IFS
GaiaDR34707563810327288192	2017-11-05	478337	IRDIS
GaiaDR34707563810327288192	2016-09-16	476392	IRDIS
GaiaDR34707563810327288192	2015-09-30	467326	IRDIS
GaiaDR34707563810327288192	2015-09-30	467325	IRDIS
GaiaDR34846466733467661184	2015-12-27	474509	IRDIS
GaiaDR34846466733467661184	2015-09-23	469299	IRDIS
GaiaDR34913138232358905216	2019-10-29	485387	IFS
GaiaDR34913138232358905216	2019-10-29	369976	IFS
GaiaDR34913138232358905216	2019-10-27	485359	IFS
GaiaDR34913138232358905216	2019-10-27	369974	IFS
GaiaDR34913138232358905216	2018-12-16	484867	IFS
GaiaDR34913138232358905216	2018-12-16	482196	IRDIS
GaiaDR34913138232358905216	2018-09-17	484239	IFS
GaiaDR34913138232358905216	2018-09-17	484238	IFS
GaiaDR34913138232358905216	2018-09-17	300340	IRDIS
GaiaDR34913138232358905216	2018-09-14	482212	IFS
GaiaDR34913138232358905216	2018-09-14	482159	IFS
GaiaDR34913138232358905216	2018-09-14	262682	IFS
GaiaDR34913138232358905216	2016-01-16	475046	IRDIS
GaiaDR34913138232358905216	2015-09-24	482096	IRDIS
GaiaDR35047124994197379456	2016-01-01	474642	IRDIS
GaiaDR35109911781067864064	2018-12-16	484868	IFS
GaiaDR35109911781067864064	2018-12-16	484866	IFS
GaiaDR35109911781067864064	2018-12-16	482195	IRDIS
GaiaDR35109911781067864064	2017-11-05	478351	IFS
GaiaDR35109911781067864064	2017-11-05	478350	IFS
GaiaDR35109911781067864064	2017-11-05	478338	IRDIS
GaiaDR35109911781067864064	2016-09-17	416785	IRDIS
GaiaDR35109911781067864064	2016-09-17	416774	IRDIS
GaiaDR35209023133585929984	2016-01-16	475049	IRDIS
GaiaDR35444751795151480320	2017-02-06	476993	IRDIS
GaiaDR35444751795151480320	2015-05-09	446243	IRDIS
GaiaDR35460728077057780864	2016-03-29	467260	IFS
GaiaDR35460728077057780864	2016-03-29	475455	IRDIS
GaiaDR35460728077057780864	2015-05-05	466655	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
GaiaDR35481843025342799104	2015-02-06	469038	IRDIS
GaiaDR35481843025342799104	2015-02-06	299869	IRDIS
GaiaDR35491506843495850240	2017-02-10	477221	IRDIS
GaiaDR35506101790904438656	2015-11-29	417895	IRDIS
GaiaDR35507465876817019136	2018-01-25	479979	IRDIS
GaiaDR35507465876817019136	2015-11-28	467922	IRDIS
GaiaDR35570338906629079936	2016-12-12	476815	IRDIS
GaiaDR35924485966955008896	2015-06-02	465706	IRDIS
GaiaDR36110141563309613056	2018-05-17	480571	IRDIS
GaiaDR36110141563309613056	2018-05-17	320112	IRDIS
GaiaDR36110141563309613056	2018-05-17	305015	IRDIS
GaiaDR36110141563309613056	2015-05-03	304498	IRDIS
GaiaDR36123612264372461312	2015-05-12	467215	IRDIS
GaiaDR36354671987249126784	2016-06-27	469236	IFS
GaiaDR36354671987249126784	2016-06-27	476198	IRDIS
GaiaDR36354671987249126784	2015-10-26	467614	IRDIS
GaiaDR36382640367603744128	2016-06-26	469213	IFS
GaiaDR36382640367603744128	2016-06-26	476153	IRDIS
GaiaDR36382640367603744128	2015-10-24	467483	IRDIS
GaiaDR36571446820698811904	2016-06-10	468737	IFS
GaiaDR36571446820698811904	2016-06-10	468736	IFS
GaiaDR36571446820698811904	2016-06-10	476043	IRDIS
GaiaDR36631685008336771072	2015-05-15	468130	IRDIS
GaiaDR36640057480144348928	2017-06-14	475734	IFS
GaiaDR36644173192685587328	2018-04-18	480356	IFS
GaiaDR36644173192685587328	2018-04-18	480496	IRDIS
GaiaDR36644173192685587328	2015-06-26	469404	IRDIS
GaiaDR36666178680963755008	2015-05-29	468251	IRDIS
GaiaDR36666178680963755008	2015-05-08	466851	IRDIS
GaiaDR36736232346363422336	2018-04-16	480301	IFS
GaiaDR36736232346363422336	2018-04-16	480435	IRDIS
GaiaDR36736232346363422336	2015-05-14	468178	IRDIS
GaiaDR36747106443324127488	2016-05-23	468515	IFS
GaiaDR36747106443324127488	2016-05-23	475807	IRDIS
GaiaDR36747106443324127488	2015-09-29	466933	IRDIS
GaiaDR36747467224874108288	2015-06-08	468449	IRDIS
GaiaDR36747467224874108288	2015-05-17	468148	IRDIS
GaiaDR36747467224874108288	2015-05-03	466577	IRDIS
HD108767B	2016-06-11	468832	IFS
HD108767B	2016-06-11	476102	IRDIS
HD116664	2017-05-02	475090	IFS
HD116664	2017-05-02	475088	IFS
HD141943	2017-04-29	474842	IFS
HD141943	2017-04-29	474898	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HD141943	2016-05-30	468603	IFS
HD141943	2016-05-30	466825	IRDIS
HD141943	2016-05-24	466563	IRDIS
HD156097	2016-06-27	476197	IRDIS
HD159911	2019-04-14	485087	IFS
HD159911	2019-04-14	482229	IRDIS
HD159911	2018-04-11	480190	IFS
HD159911	2018-04-11	480274	IRDIS
HD159911	2016-05-21	468187	IFS
HD159911	2016-05-21	475714	IRDIS
HD159911	2016-04-15	467538	IFS
HD159911	2016-04-15	472177	IRDIS
HD159911	2015-05-11	467124	IRDIS
HD189285	2018-05-06	480688	IFS
HD189285	2018-05-06	480497	IRDIS
HD189285	2016-05-28	475855	IRDIS
HD189285	2015-06-02	465707	IRDIS
HD201919	2015-06-02	465708	IRDIS
HD25284	2015-09-29	466937	IRDIS
HD317617	2015-05-30	468426	IRDIS
HD326277	2016-03-29	467263	IFS
HD326277	2016-03-29	475457	IRDIS
HD51797	2016-01-02	474816	IRDIS
HD75505	2018-01-05	478583	IFS
HD75505	2018-01-05	479748	IRDIS
HD75505	2015-02-05	468458	IRDIS
HD75505	2015-02-05	445787	IRDIS
HD75505	2015-02-05	299840	IRDIS
HD75505	2015-02-05	299839	IRDIS
HIP102409	2019-09-06	485172	IFS
HIP102409	2019-09-06	485170	IFS
HIP102409	2019-09-06	417697	IFS
HIP102409	2019-09-06	482537	IRDIS
HIP102409	2016-05-20	468153	IFS
HIP102409	2016-05-20	468152	IFS
HIP102409	2016-05-20	417607	IFS
HIP102409	2016-05-20	417939	IRDIS
HIP102409	2016-05-20	417938	IRDIS
HIP102409	2016-05-20	417937	IRDIS
HIP102409	2015-10-06	466909	IRDIS
HIP102409	2015-10-06	466830	IRDIS
HIP102409	2015-09-30	417578	IFS
HIP102409	2015-09-30	417852	IRDIS
HIP102409	2015-09-30	417848	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP102409	2015-09-30	417847	IRDIS
HIP102409	2015-09-24	417572	IFS
HIP102409	2015-09-24	417841	IRDIS
HIP102409	2015-06-26	417563	IFS
HIP102409	2015-06-26	417833	IRDIS
HIP102409	2015-06-26	417830	IRDIS
HIP102409	2015-06-26	417829	IRDIS
HIP102409	2015-05-29	417550	IFS
HIP102409	2015-05-29	417549	IFS
HIP102409	2015-05-29	417820	IRDIS
HIP102409	2015-05-29	417819	IRDIS
HIP102409	2015-05-15	417545	IFS
HIP102409	2015-05-15	417813	IRDIS
HIP102409	2015-05-15	417812	IRDIS
HIP102626	2015-06-03	465883	IRDIS
HIP104365	2015-09-28	466644	IRDIS
HIP105388	2016-05-30	468606	IFS
HIP105388	2016-05-30	468605	IFS
HIP105388	2016-05-30	466826	IRDIS
HIP105388	2015-05-30	447164	IRDIS
HIP10602	2017-08-31	476174	IFS
HIP10602	2017-08-31	248409	IFS
HIP10602	2017-08-31	405900	IRDIS
HIP10602	2017-08-31	405899	IRDIS
HIP10602	2016-09-16	476405	IRDIS
HIP10602	2016-09-16	476404	IRDIS
HIP10602	2016-09-16	476393	IRDIS
HIP10602	2015-09-25	482170	IRDIS
HIP107345	2015-07-04	416749	IRDIS
HIP107345	2015-07-04	416747	IRDIS
HIP107345	2015-05-31	468583	IRDIS
HIP107350	2015-09-29	466934	IRDIS
HIP107412	2019-11-27	485500	IFS
HIP107412	2019-11-27	482550	IRDIS
HIP107412	2018-06-19	480924	IFS
HIP107412	2018-06-19	480676	IRDIS
HIP107412	2016-09-15	469269	IFS
HIP107412	2016-09-15	467169	IRDIS
HIP107412	2015-10-03	465820	IRDIS
HIP109285	2016-10-13	476517	IRDIS
HIP109285	2016-09-16	476401	IRDIS
HIP109285	2015-06-30	467251	IRDIS
HIP1113	2016-10-14	467500	IRDIS
HIP112312	2015-09-30	467322	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP112312	2015-09-30	467321	IRDIS
HIP113201	2017-06-01	477942	IRDIS
HIP113201	2017-06-01	477941	IRDIS
HIP113201	2016-10-14	476627	IRDIS
HIP113201	2016-09-17	416772	IRDIS
HIP113283	2015-06-26	483997	IRDIS
HIP113283	2015-06-26	468494	IRDIS
HIP113283	2015-06-26	468493	IRDIS
HIP113368	2015-09-23	481416	IRDIS
HIP113368	2015-06-29	466865	IRDIS
HIP113368	2015-06-29	466864	IRDIS
HIP113368	2015-06-29	466848	IRDIS
HIP11360	2015-10-25	467574	IRDIS
HIP114189	2018-06-18	480850	IFS
HIP114189	2018-06-18	480631	IRDIS
HIP114189	2017-10-12	478327	IFS
HIP114189	2017-10-12	478304	IRDIS
HIP114189	2017-10-12	394715	IRDIS
HIP114189	2017-10-11	477086	IFS
HIP114189	2017-10-11	478298	IRDIS
HIP114189	2017-10-11	478297	IRDIS
HIP114189	2017-10-11	478296	IRDIS
HIP114189	2017-10-11	393389	IRDIS
HIP114189	2017-10-11	393388	IRDIS
HIP114189	2017-10-11	393387	IRDIS
HIP114189	2017-10-07	476635	IFS
HIP114189	2017-10-07	260651	IFS
HIP114189	2017-10-07	478191	IRDIS
HIP114189	2017-10-07	478190	IRDIS
HIP114189	2017-10-07	399863	IRDIS
HIP114189	2017-10-07	399862	IRDIS
HIP114189	2017-06-14	475735	IFS
HIP114189	2017-06-14	475504	IRDIS
HIP114189	2016-11-17	467759	IRDIS
HIP114189	2015-09-27	466215	IRDIS
HIP114189	2015-07-30	490918	IRDIS
HIP114189	2015-07-29	467436	IRDIS
HIP114189	2015-07-03	460400	IRDIS
HIP114189	2014-12-08	466686	IRDIS
HIP114189	2014-12-07	466589	IRDIS
HIP114189	2014-12-06	466548	IRDIS
HIP114189	2014-12-05	466390	IRDIS
HIP114189	2014-12-04	466288	IRDIS
HIP114189	2014-08-11	460544	IFS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP114189	2014-08-11	460543	IFS
HIP114189	2014-08-11	460542	IFS
HIP114189	2014-07-15	460251	IFS
HIP114189	2014-07-12	460202	IRDIS
HIP114530	2015-09-29	466935	IRDIS
HIP114948	2015-09-30	467324	IRDIS
HIP114948	2015-09-30	467323	IRDIS
HIP114952	2015-10-25	467579	IRDIS
HIP115738	2015-09-24	482095	IRDIS
HIP118008	2015-10-24	467484	IRDIS
HIP118121	2016-06-27	476209	IRDIS
HIP118121	2016-06-27	476200	IRDIS
HIP118121	2016-06-27	476199	IRDIS
HIP118121	2015-10-26	467615	IRDIS
HIP12361	2015-10-02	465638	IRDIS
HIP12394	2017-09-30	476533	IFS
HIP12394	2017-09-30	258010	IFS
HIP12394	2017-09-30	250949	IFS
HIP12394	2017-09-30	478135	IRDIS
HIP12394	2016-09-15	476279	IRDIS
HIP12394	2015-09-30	467328	IRDIS
HIP12394	2015-09-30	467327	IRDIS
HIP13402	2016-10-14	310075	IFS
HIP13402	2016-10-14	467501	IRDIS
HIP14551	2016-10-12	443532	IRDIS
HIP1481	2016-09-17	416784	IRDIS
HIP1481	2016-09-17	416773	IRDIS
HIP1481	2015-10-25	467573	IRDIS
HIP15457	2015-12-26	485231	IRDIS
HIP17764	2016-01-02	474814	IRDIS
HIP17764	2015-10-23	467472	IRDIS
HIP17764	2015-10-02	484092	IRDIS
HIP17764	2015-10-02	468822	IRDIS
HIP1993	2015-11-28	467920	IRDIS
HIP21547	2019-11-27	485501	IFS
HIP21547	2019-11-27	482552	IRDIS
HIP21547	2018-09-17	484240	IFS
HIP21547	2018-09-17	481771	IRDIS
HIP21547	2017-09-27	476446	IFS
HIP21547	2017-09-27	475681	IRDIS
HIP21547	2016-12-12	476820	IRDIS
HIP21547	2016-12-11	476742	IRDIS
HIP21547	2016-12-09	476721	IRDIS
HIP21547	2016-01-15	331401	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP21547	2015-12-25	469175	IRDIS
HIP21547	2015-09-25	482171	IRDIS
HIP21547	2015-09-24	482097	IRDIS
HIP22226	2016-12-05	569251	IFS
HIP22226	2016-12-05	466312	IRDIS
HIP22226	2015-12-28	474720	IRDIS
HIP22226	2015-11-30	469101	IRDIS
HIP22295	2015-09-29	466938	IRDIS
HIP23200	2015-09-27	466217	IRDIS
HIP23309	2016-01-02	474815	IRDIS
HIP23451	2016-12-18	476951	IRDIS
HIP24947	2017-02-11	477247	IRDIS
HIP25283	2015-12-27	474510	IRDIS
HIP25283	2015-10-14	467002	IRDIS
HIP25544	2015-12-28	474721	IRDIS
HIP2578	2016-09-15	476278	IRDIS
HIP2578	2015-10-02	468821	IRDIS
HIP2578	2015-09-27	466216	IRDIS
HIP27288	2017-02-10	477220	IRDIS
HIP27321	2019-12-21	485536	IFS
HIP27321	2019-12-21	354895	IFS
HIP27321	2019-11-02	485400	IFS
HIP27321	2019-11-02	363453	IFS
HIP27321	2019-03-10	484900	IFS
HIP27321	2019-03-10	417686	IFS
HIP27321	2018-12-14	484849	IFS
HIP27321	2018-12-14	417675	IFS
HIP27321	2018-10-17	484389	IFS
HIP27321	2018-10-17	481943	IRDIS
HIP27321	2018-09-16	484140	IFS
HIP27321	2018-09-16	417652	IFS
HIP27321	2018-09-16	481510	IRDIS
HIP27321	2018-09-13	482128	IFS
HIP27321	2018-09-13	262458	IFS
HIP27321	2018-09-13	481501	IRDIS
HIP27321	2018-09-13	481500	IRDIS
HIP27321	2018-09-13	393752	IRDIS
HIP27321	2018-09-13	393751	IRDIS
HIP27321	2018-09-11	482086	IFS
HIP27321	2018-09-11	262454	IFS
HIP27321	2018-09-11	481495	IRDIS
HIP27321	2018-09-11	480818	IRDIS
HIP27321	2018-09-11	392871	IRDIS
HIP27321	2018-09-11	392870	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP27321	2017-02-12	417631	IFS
HIP27321	2016-11-17	417623	IFS
HIP27321	2016-11-17	476707	IRDIS
HIP27321	2016-11-17	476706	IRDIS
HIP27321	2016-11-17	418000	IRDIS
HIP27321	2016-10-13	417621	IFS
HIP27321	2016-10-13	476519	IRDIS
HIP27321	2016-10-13	417995	IRDIS
HIP27321	2016-09-15	469273	IFS
HIP27321	2016-09-15	417615	IFS
HIP27321	2016-09-15	476284	IRDIS
HIP27321	2016-09-15	417987	IRDIS
HIP27321	2016-04-15	467537	IFS
HIP27321	2016-04-15	417604	IFS
HIP27321	2016-04-15	472174	IRDIS
HIP27321	2016-04-15	472173	IRDIS
HIP27321	2016-04-15	417920	IRDIS
HIP27321	2016-04-15	417919	IRDIS
HIP27321	2016-03-26	466862	IFS
HIP27321	2016-03-26	417601	IFS
HIP27321	2016-03-26	475034	IRDIS
HIP27321	2016-03-26	475033	IRDIS
HIP27321	2016-03-26	417910	IRDIS
HIP27321	2016-01-19	466504	IFS
HIP27321	2016-01-19	417596	IFS
HIP27321	2016-01-19	475160	IRDIS
HIP27321	2015-12-25	465850	IFS
HIP27321	2015-12-25	417592	IFS
HIP27321	2015-12-25	484843	IRDIS
HIP27321	2015-11-29	467941	IRDIS
HIP27321	2015-11-27	417584	IFS
HIP27321	2015-11-27	467898	IRDIS
HIP27321	2015-11-27	467897	IRDIS
HIP27321	2015-11-27	417889	IRDIS
HIP27321	2015-11-27	417888	IRDIS
HIP27321	2015-09-30	417579	IFS
HIP27321	2015-09-30	467421	IRDIS
HIP27321	2015-02-04	417540	IFS
HIP27321	2015-02-04	467061	IRDIS
HIP27321	2015-02-04	417684	IRDIS
HIP27321	2014-12-07	460549	IFS
HIP27321	2014-12-07	417534	IFS
HIP27321	2014-12-07	466590	IRDIS
HIP27321	2014-12-07	417539	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP27321	2014-12-07	417538	IRDIS
HIP30030	2017-02-08	310079	IFS
HIP30030	2017-02-08	477174	IRDIS
HIP30030	2015-10-27	467745	IRDIS
HIP30034	2015-02-05	445786	IRDIS
HIP30034	2015-02-05	299837	IRDIS
HIP30314	2016-01-15	331402	IRDIS
HIP31878	2016-01-01	474643	IRDIS
HIP32235	2015-12-27	474511	IRDIS
HIP33737	2017-02-07	477131	IRDIS
HIP33737	2015-12-26	485232	IRDIS
HIP34899	2016-03-29	467261	IFS
HIP34899	2016-03-29	475461	IRDIS
HIP34899	2016-03-28	467079	IFS
HIP34899	2016-03-28	475343	IRDIS
HIP34899	2015-12-28	474722	IRDIS
HIP36349	2016-03-27	299443	IRDIS
HIP36948	2015-03-29	466223	IRDIS
HIP36948	2015-03-29	445893	IRDIS
HIP36948	2015-02-02	467003	IRDIS
HIP36948	2015-02-02	445340	IRDIS
HIP37288	2017-02-12	477305	IRDIS
HIP37288	2015-11-30	484779	IRDIS
HIP37766	2015-12-25	465851	IFS
HIP37766	2015-12-25	417903	IRDIS
HIP37766	2015-12-25	417901	IRDIS
HIP39826	2017-02-11	477248	IRDIS
HIP41307	2018-12-14	484850	IFS
HIP41307	2018-12-14	484848	IFS
HIP41307	2018-12-14	484847	IFS
HIP41307	2018-12-14	482099	IRDIS
HIP41307	2016-12-15	221566	IFS
HIP41307	2016-12-15	221565	IFS
HIP41307	2016-12-05	221320	IFS
HIP41307	2016-01-18	466421	IFS
HIP41307	2016-01-18	466420	IFS
HIP41307	2016-01-18	466419	IFS
HIP41307	2016-01-18	474554	IRDIS
HIP41307	2016-01-18	474553	IRDIS
HIP41307	2016-01-18	474504	IRDIS
HIP41307	2016-01-18	474503	IRDIS
HIP42637	2016-12-12	476822	IRDIS
HIP42637	2016-01-01	474644	IRDIS
HIP42808	2016-12-11	417302	IFS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP42808	2016-12-11	476743	IRDIS
HIP42808	2016-03-27	466951	IFS
HIP42808	2016-03-27	475280	IRDIS
HIP42808	2016-01-02	474826	IRDIS
HIP42808	2016-01-02	474817	IRDIS
HIP44526	2017-02-12	302452	IFS
HIP44526	2017-02-12	477306	IRDIS
HIP44526	2016-04-16	475553	IRDIS
HIP47135	2017-02-10	477222	IRDIS
HIP47135	2016-03-14	474750	IRDIS
HIP47135	2016-01-17	416760	IRDIS
HIP47135	2015-02-04	417683	IRDIS
HIP47135	2015-02-04	299524	IRDIS
HIP490	2015-09-25	482169	IRDIS
HIP490	2015-09-25	482168	IRDIS
HIP50191	2018-02-26	479136	IFS
HIP50191	2018-02-26	480031	IRDIS
HIP50191	2018-01-04	478544	IFS
HIP50191	2018-01-04	478492	IRDIS
HIP50191	2016-03-26	466863	IFS
HIP50191	2016-03-26	475037	IRDIS
HIP50191	2016-03-26	475036	IRDIS
HIP50191	2016-03-26	475035	IRDIS
HIP50191	2016-03-26	417912	IRDIS
HIP50191	2015-05-10	467008	IRDIS
HIP50191	2015-03-29	466261	IRDIS
HIP50191	2015-03-29	466224	IRDIS
HIP51228	2016-01-02	474818	IRDIS
HIP51317	2015-12-31	474485	IRDIS
HIP51317	2015-02-06	469041	IRDIS
HIP51317	2015-02-06	469040	IRDIS
HIP51317	2015-02-06	299873	IRDIS
HIP53524	2019-05-17	485125	IFS
HIP53524	2019-05-17	485119	IFS
HIP53524	2019-05-17	482253	IRDIS
HIP53524	2019-04-12	485026	IFS
HIP53524	2019-04-12	477129	IRDIS
HIP53524	2018-03-27	479769	IFS
HIP53524	2018-03-27	480056	IRDIS
HIP53524	2018-02-23	479052	IFS
HIP53524	2018-02-23	476656	IRDIS
HIP53524	2018-01-05	478584	IFS
HIP53524	2018-01-05	476385	IRDIS
HIP53524	2017-05-09	475263	IFS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP53524	2017-05-09	475251	IRDIS
HIP53524	2017-05-02	475087	IFS
HIP53524	2017-05-02	475047	IRDIS
HIP53524	2016-05-30	468602	IFS
HIP53524	2016-05-30	466827	IRDIS
HIP53524	2016-04-15	467536	IFS
HIP53524	2016-04-15	417921	IRDIS
HIP53524	2016-04-15	417918	IRDIS
HIP53524	2016-01-17	466196	IFS
HIP53524	2016-01-17	416761	IRDIS
HIP53524	2015-05-11	446292	IRDIS
HIP53524	2015-05-04	446062	IRDIS
HIP53524	2015-05-04	446061	IRDIS
HIP54155	2017-02-08	477175	IRDIS
HIP54155	2016-04-20	467827	IFS
HIP54155	2016-04-20	472150	IRDIS
HIP54231	2015-02-05	468743	IRDIS
HIP54231	2015-02-05	468461	IRDIS
HIP54231	2015-02-05	468460	IRDIS
HIP54231	2015-02-05	299846	IRDIS
HIP54231	2015-02-05	299845	IRDIS
HIP56354	2016-01-22	475349	IRDIS
HIP56354	2016-01-20	475243	IRDIS
HIP56354	2016-01-19	475162	IRDIS
HIP56354	2016-01-15	331403	IRDIS
HIP56354	2015-04-09	466500	IRDIS
HIP56379	2018-03-28	480091	IRDIS
HIP56379	2017-02-07	299391	IRDIS
HIP56379	2016-05-31	475985	IRDIS
HIP56379	2016-04-16	299589	IRDIS
HIP56379	2016-04-16	299588	IRDIS
HIP56379	2016-03-26	475012	IRDIS
HIP56379	2016-01-16	475051	IRDIS
HIP56379	2015-05-29	468249	IRDIS
HIP56379	2015-05-03	466575	IRDIS
HIP56379	2015-02-06	469042	IRDIS
HIP56379	2015-02-06	468496	IRDIS
HIP56379	2015-02-06	299875	IRDIS
HIP57013	2016-06-01	323900	IFS
HIP57013	2016-06-01	476005	IRDIS
HIP57632	2015-05-31	468579	IRDIS
HIP57632	2015-05-30	468425	IRDIS
HIP58167	2019-05-18	485146	IFS
HIP58167	2019-05-18	482268	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP58167	2015-06-02	465703	IRDIS
HIP58465	2016-03-29	467262	IFS
HIP58465	2016-03-29	475456	IRDIS
HIP58465	2015-06-03	465879	IRDIS
HIP59505	2018-03-29	480019	IFS
HIP59505	2018-03-29	480149	IRDIS
HIP59505	2017-05-06	475151	IFS
HIP59505	2017-05-06	477730	IRDIS
HIP59505	2016-06-13	425980	IRDIS
HIP59960	2015-05-11	467127	IRDIS
HIP60183	2016-01-18	466423	IFS
HIP60183	2016-01-18	466422	IFS
HIP60183	2016-01-18	474505	IRDIS
HIP60183	2015-06-05	450112	IRDIS
HIP60459	2019-03-09	484891	IFS
HIP60459	2019-03-09	482223	IRDIS
HIP60459	2015-02-03	460370	IRDIS
HIP60459	2015-02-03	299515	IRDIS
HIP60965	2018-01-24	478880	IFS
HIP60965	2018-01-24	479863	IRDIS
HIP61468	2018-06-18	480848	IFS
HIP61468	2018-06-18	480625	IRDIS
HIP61468	2016-04-20	467829	IFS
HIP61468	2016-04-20	467828	IFS
HIP61468	2016-04-20	472151	IRDIS
HIP61498	2018-05-05	480505	IFS
HIP61498	2015-02-02	445341	IRDIS
HIP61960	2016-06-10	468730	IFS
HIP61960	2016-06-10	476039	IRDIS
HIP62171	2019-03-05	484882	IFS
HIP6276	2016-11-16	467672	IRDIS
HIP6276	2015-12-01	305952	IFS
HIP6276	2015-12-01	305799	IFS
HIP6276	2015-12-01	469147	IRDIS
HIP63839	2016-04-16	467637	IFS
HIP63839	2016-04-16	467636	IFS
HIP63839	2016-04-16	466212	IRDIS
HIP64792	2018-04-12	480313	IRDIS
HIP64792	2017-02-09	477182	IRDIS
HIP64792	2016-03-28	467081	IFS
HIP64792	2016-03-28	475344	IRDIS
HIP64792	2015-06-06	310973	IRDIS
HIP64792	2015-06-05	466624	IRDIS
HIP64792	2015-06-04	466531	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP64792	2015-06-02	465704	IRDIS
HIP6485	2016-10-13	467458	IRDIS
HIP6485	2015-10-27	467744	IRDIS
HIP6485	2015-09-23	469298	IRDIS
HIP64892	2018-04-13	480283	IFS
HIP64892	2018-04-13	480403	IRDIS
HIP64892	2017-02-07	477133	IRDIS
HIP64892	2016-03-31	467426	IFS
HIP64892	2016-03-31	475522	IRDIS
HIP64995	2017-05-07	477743	IRDIS
HIP64995	2015-06-03	465880	IRDIS
HIP65426	2018-05-12	480774	IFS
HIP65426	2018-05-12	480564	IRDIS
HIP65426	2018-05-12	388931	IRDIS
HIP65426	2017-05-03	475144	IFS
HIP65426	2017-05-02	475091	IFS
HIP65426	2017-05-02	475089	IFS
HIP65426	2017-02-08	477176	IRDIS
HIP65426	2017-02-06	477001	IRDIS
HIP65426	2016-06-26	476152	IRDIS
HIP65426	2016-05-30	468601	IFS
HIP65426	2016-05-30	475960	IRDIS
HIP66252	2016-06-27	476195	IRDIS
HIP66566	2015-05-14	468105	IRDIS
HIP66651	2016-06-01	323902	IFS
HIP66651	2016-06-01	323901	IFS
HIP66651	2016-06-01	476006	IRDIS
HIP682	2018-08-03	481156	IFS
HIP682	2018-08-03	261621	IFS
HIP682	2016-10-13	310038	IFS
HIP682	2016-10-13	476512	IRDIS
HIP682	2015-10-03	465822	IRDIS
HIP682	2015-10-03	465821	IRDIS
HIP6856	2017-09-28	476494	IFS
HIP6856	2017-09-28	478108	IRDIS
HIP6856	2016-11-16	476634	IRDIS
HIP6856	2016-11-16	467671	IRDIS
HIP6856	2015-09-29	466936	IRDIS
HIP68781	2016-03-29	467264	IFS
HIP68781	2016-03-29	475464	IRDIS
HIP68781	2015-03-30	416738	IRDIS
HIP69989	2016-06-11	468834	IFS
HIP69989	2016-06-11	468833	IFS
HIP69989	2016-06-11	476103	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP71724	2017-03-18	474495	IFS
HIP71724	2017-03-18	477434	IRDIS
HIP71724	2016-04-16	466213	IRDIS
HIP71743	2016-05-31	468705	IFS
HIP71743	2016-05-31	475986	IRDIS
HIP72940	2016-05-06	475651	IRDIS
HIP73145	2015-05-13	468058	IRDIS
HIP73990	2018-04-12	480300	IRDIS
HIP73990	2016-06-27	469233	IFS
HIP73990	2016-06-27	476196	IRDIS
HIP73990	2016-04-20	467830	IFS
HIP73990	2016-04-20	472152	IRDIS
HIP73990	2016-03-17	474761	IRDIS
HIP73990	2016-03-06	474542	IRDIS
HIP74824	2015-03-29	466225	IRDIS
HIP76063	2016-03-28	467078	IFS
HIP76063	2016-03-28	475340	IRDIS
HIP77317	2016-04-20	467832	IFS
HIP77317	2016-04-20	467831	IFS
HIP77317	2016-04-20	472153	IRDIS
HIP77457	2015-06-08	468448	IRDIS
HIP77457	2015-03-30	416739	IRDIS
HIP77464	2015-06-04	466532	IRDIS
HIP78092	2019-05-17	482252	IRDIS
HIP78099	2015-05-08	466849	IRDIS
HIP78196	2015-06-02	465705	IRDIS
HIP78196	2015-06-02	300087	IRDIS
HIP78530	2015-05-03	466576	IRDIS
HIP78541	2016-06-10	468732	IFS
HIP78541	2016-06-10	468731	IFS
HIP78541	2016-06-10	476040	IRDIS
HIP78541	2015-05-18	468179	IRDIS
HIP78541	2015-05-08	466850	IRDIS
HIP78581	2017-03-19	477530	IRDIS
HIP78581	2015-06-03	465882	IRDIS
HIP78581	2015-06-03	465881	IRDIS
HIP78581	2015-05-31	468580	IRDIS
HIP80324	2018-06-17	480834	IFS
HIP80324	2018-06-17	480623	IRDIS
HIP80324	2016-04-16	475559	IRDIS
HIP80324	2016-03-27	466952	IFS
HIP80324	2016-03-27	475283	IRDIS
HIP80324	2015-05-29	468250	IRDIS
HIP80591	2016-03-28	467080	IFS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP80591	2016-03-28	475341	IRDIS
HIP80591	2015-04-04	465701	IRDIS
HIP81084	2015-08-03	469264	IRDIS
HIP81084	2015-05-10	467068	IRDIS
HIP82388	2016-06-11	468835	IFS
HIP82388	2016-06-11	476104	IRDIS
HIP82430	2019-05-19	485157	IFS
HIP82430	2019-05-19	482280	IRDIS
HIP82430	2017-05-17	475430	IFS
HIP82430	2017-05-17	477886	IRDIS
HIP82430	2017-05-17	398976	IRDIS
HIP82430	2016-03-31	467428	IFS
HIP82430	2016-03-31	467427	IFS
HIP82430	2016-03-31	475523	IRDIS
HIP82588	2015-06-04	466533	IRDIS
HIP85647	2016-05-30	475963	IRDIS
HIP85647	2015-05-05	466657	IRDIS
HIP86305	2019-05-19	485158	IFS
HIP86305	2016-06-10	468733	IFS
HIP86305	2016-06-10	476041	IRDIS
HIP86598	2016-05-22	468274	IFS
HIP86598	2016-05-22	299675	IRDIS
HIP86598	2015-05-12	467216	IRDIS
HIP87174	2017-05-02	477675	IRDIS
HIP88399	2019-09-07	485336	IFS
HIP88399	2019-09-07	482300	IRDIS
HIP88399	2018-04-10	480143	IFS
HIP88399	2018-04-10	480191	IRDIS
HIP88399	2016-04-16	475554	IRDIS
HIP88399	2015-05-31	468581	IRDIS
HIP88399	2015-05-09	467000	IRDIS
HIP89829	2019-04-13	485079	IFS
HIP89829	2017-06-02	475697	IFS
HIP89829	2017-06-02	478032	IRDIS
HIP89829	2016-06-13	469109	IFS
HIP89829	2016-06-13	469108	IFS
HIP89829	2016-06-13	425983	IRDIS
HIP92024	2015-05-10	467027	IRDIS
HIP92024	2015-05-10	466211	IRDIS
HIP92680	2019-09-06	485171	IFS
HIP92680	2018-05-12	310083	IFS
HIP92680	2018-05-12	480565	IRDIS
HIP92680	2018-05-12	388933	IRDIS
HIP92680	2017-05-17	475431	IFS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP92680	2017-05-17	477888	IRDIS
HIP92680	2017-05-17	477887	IRDIS
HIP92680	2017-05-17	398978	IRDIS
HIP92680	2017-05-17	398977	IRDIS
HIP92680	2016-09-16	476399	IRDIS
HIP92680	2016-06-22	469202	IFS
HIP92680	2016-06-22	476145	IRDIS
HIP92680	2016-04-17	467808	IFS
HIP92680	2016-04-17	475589	IRDIS
HIP92680	2015-05-30	468427	IRDIS
HIP92680	2015-05-30	300035	IRDIS
HIP92680	2015-05-05	466660	IRDIS
HIP92680	2014-08-07	460304	IFS
HIP92680	2014-08-07	465929	IRDIS
HIP92680	2014-07-14	460181	IFS
HIP92680	2014-07-14	460217	IRDIS
HIP92984	2016-09-15	469270	IFS
HIP92984	2016-09-15	476280	IRDIS
HIP92984	2016-06-11	468837	IFS
HIP92984	2016-06-11	468836	IFS
HIP92984	2016-06-11	476105	IRDIS
HIP93375	2015-06-06	468419	IRDIS
HIP93747	2016-05-30	468604	IFS
HIP93747	2016-05-30	475961	IRDIS
HIP93747	2015-06-27	466197	IRDIS
HIP94114	2015-06-05	466625	IRDIS
HIP95261	2018-05-07	480743	IFS
HIP95261	2018-05-07	480530	IRDIS
HIP95261	2017-06-14	475733	IFS
HIP95261	2017-06-14	475732	IFS
HIP95261	2017-06-14	446022	IRDIS
HIP95261	2016-06-14	469189	IFS
HIP95261	2015-05-04	445725	IRDIS
HIP95270	2015-05-09	467001	IRDIS
HIP95347	2019-05-17	485126	IFS
HIP95347	2018-05-07	480742	IFS
HIP95347	2018-05-07	480529	IRDIS
HIP95347	2018-05-07	480528	IRDIS
HIP95347	2016-05-20	475659	IRDIS
HIP95347	2015-05-12	467225	IRDIS
HIP95347	2015-05-12	467217	IRDIS
HIP97255	2015-07-04	465758	IRDIS
HIP98470	2015-05-13	468060	IRDIS
HIP98470	2015-05-13	468059	IRDIS

Continued on next page

Table B.3 continued from previous page

Object	Obs. date	PID	Instr.
HIP98495	2016-05-24	468535	IFS
HIP98495	2016-05-24	475824	IRDIS
HIP98495	2016-05-24	475812	IRDIS
HIP98495	2015-06-30	467250	IRDIS
HIP98495	2015-05-31	468582	IRDIS
HIP98495	2015-05-11	467125	IRDIS
HIP99742	2016-06-10	468735	IFS
HIP99742	2016-06-10	468734	IFS
HIP99742	2016-06-10	476042	IRDIS
HIP99742	2015-09-27	466214	IRDIS
HIP99742	2015-06-21	453756	IRDIS
End of Table B.3			

Reductions

This appendix contains further residual images and detection maps for the algorithms discussed in [Section 5.2](#).

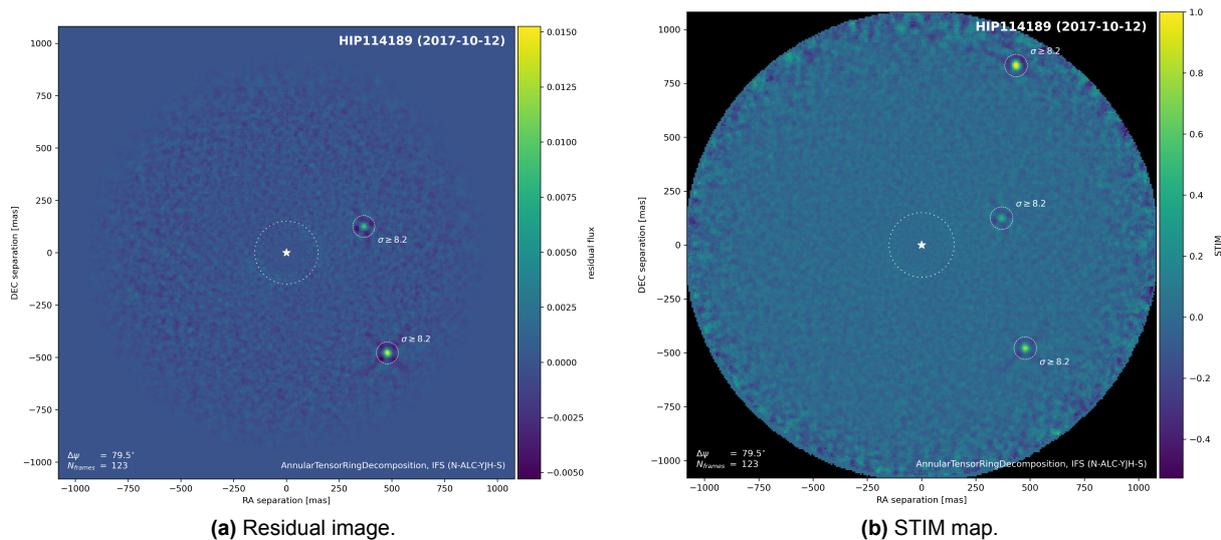


Figure C.1: HR 8799 IFS residual images from annular TRD.

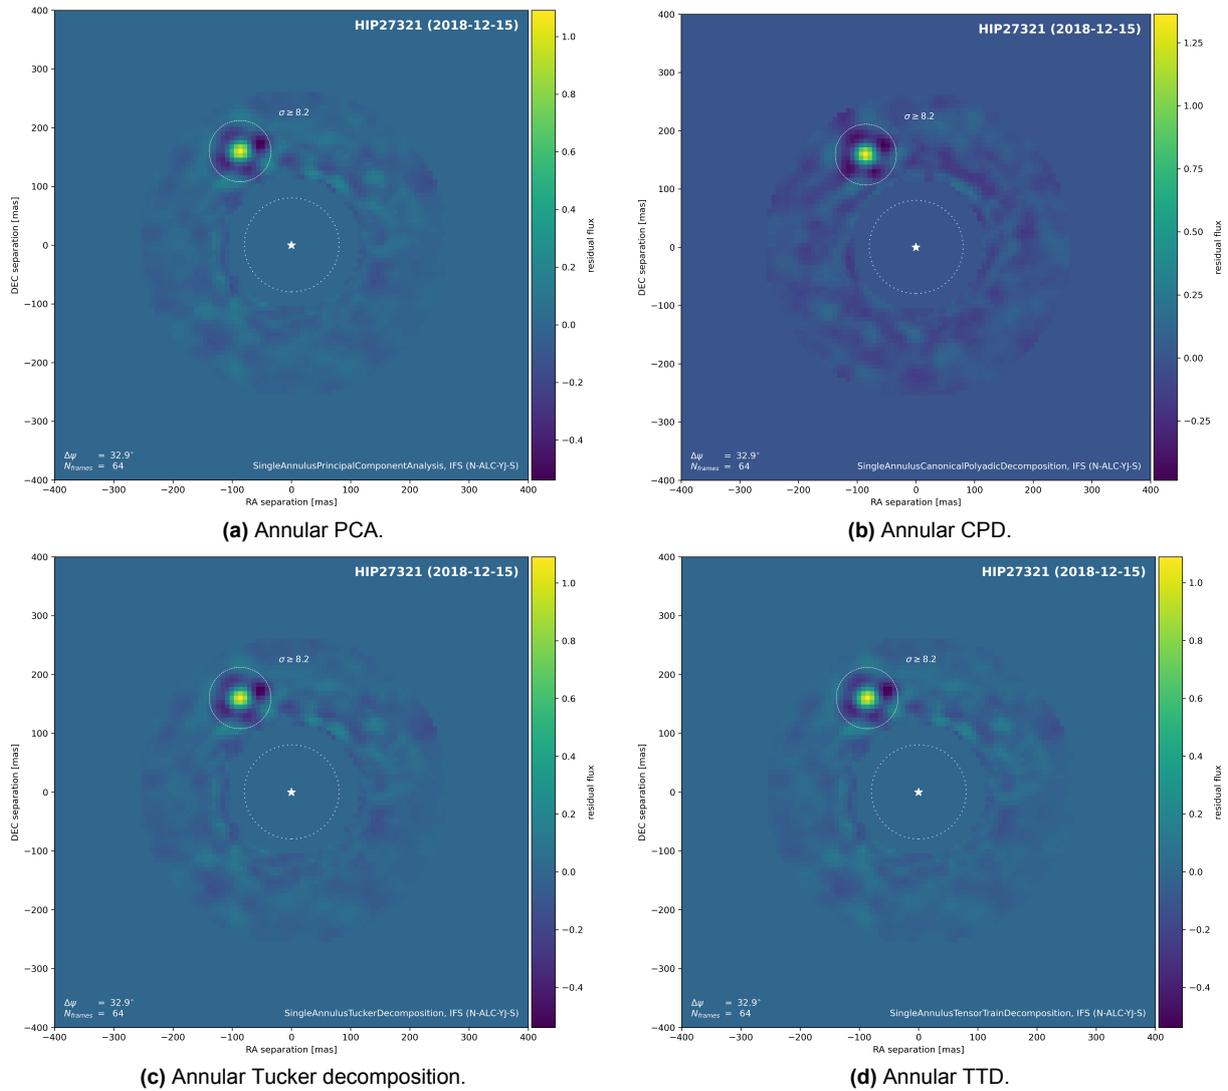


Figure C.2: β Pictoris residual images from algorithms using a single annulus centered on β Pictoris b. The residuals are very similar for all methods using spatial unfolding.

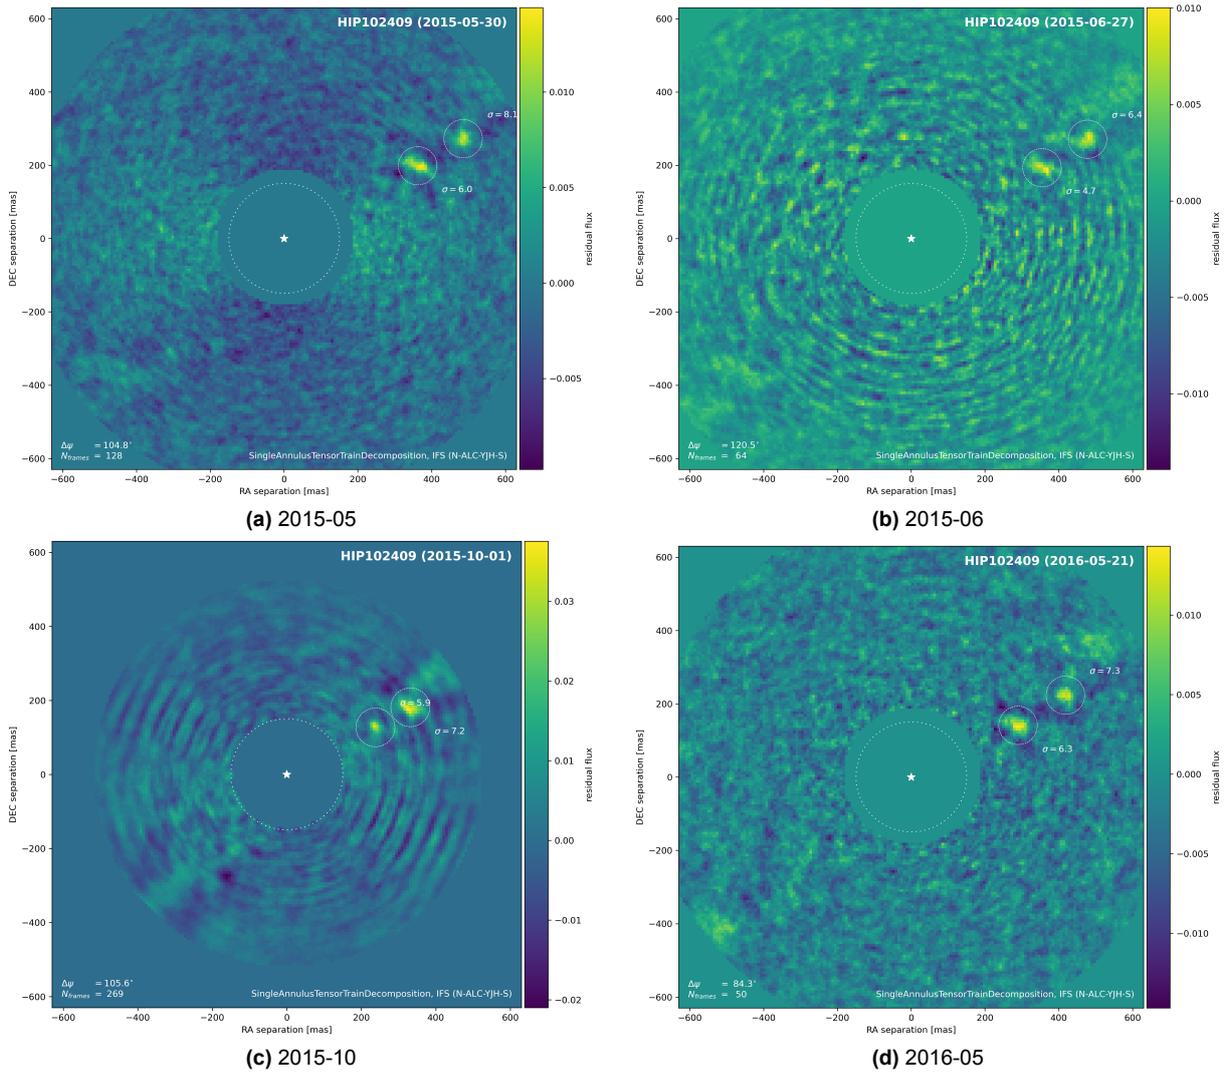


Figure C.3: Fast moving features in the AU Mic. disk in IFS observations. These features were originally detected by [Boccaletti et al., 2015](#) and [Boccaletti et al., 2018](#). Reductions with tensor methods very clearly detect the features.

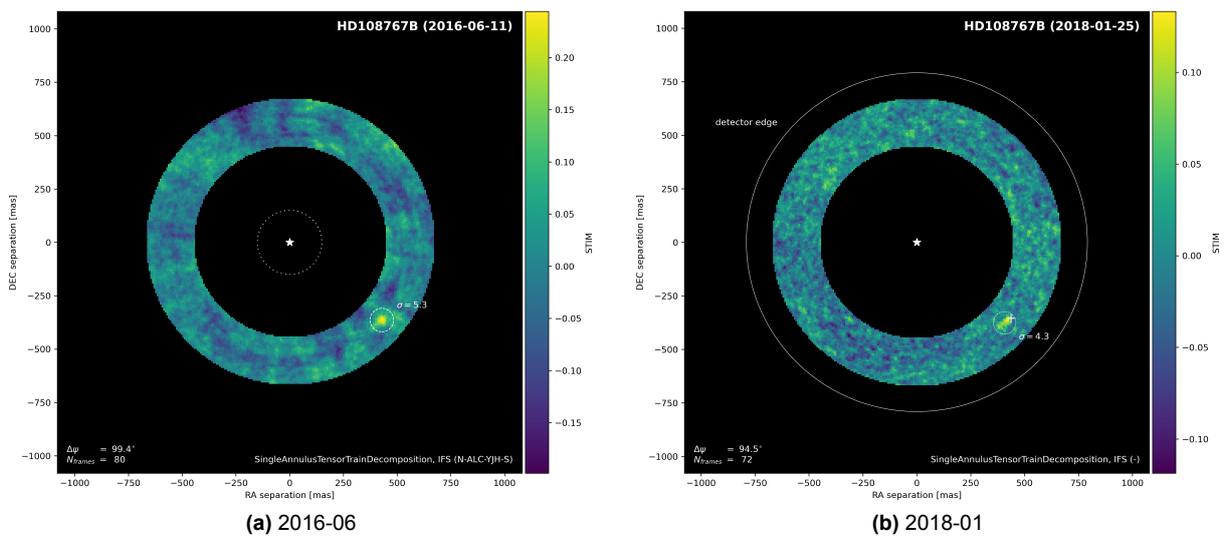
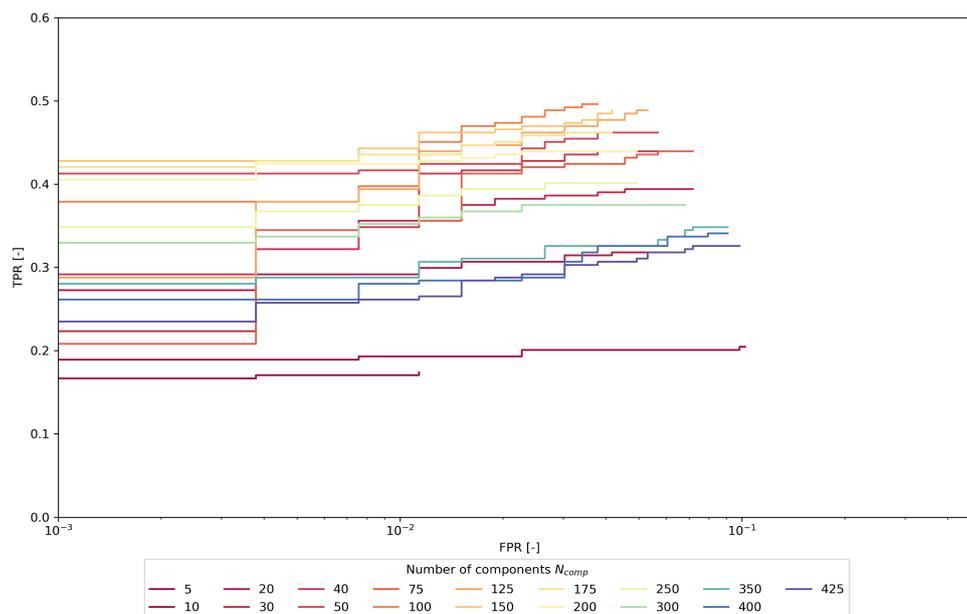


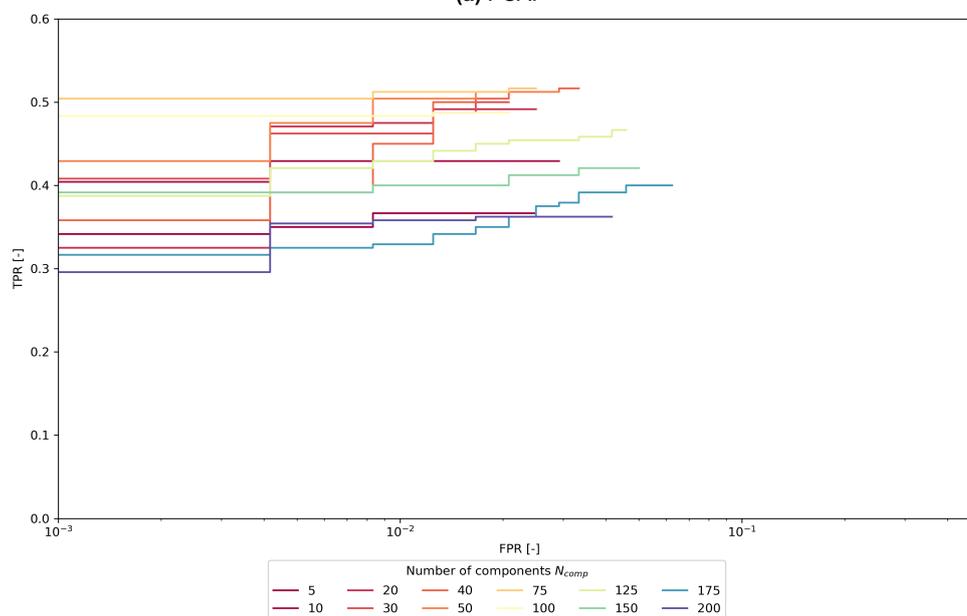
Figure C.4: STIM maps of IFS observations of HD 108767 B reduced with annular TTD.

Receiver Operating Characteristics

This appendix includes the full ROC results from the algorithm assessment in [Section 5.3](#).

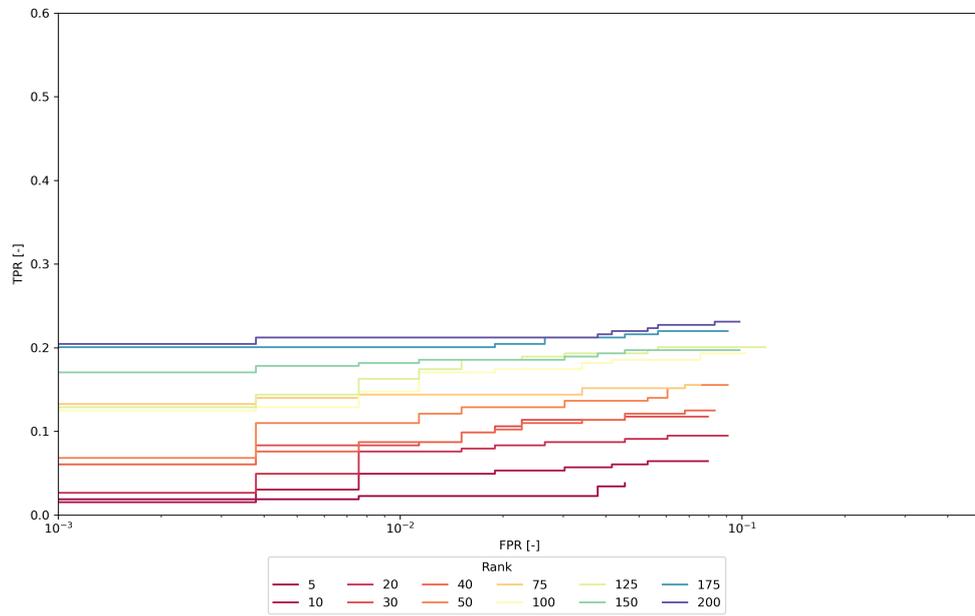


(a) PCA.

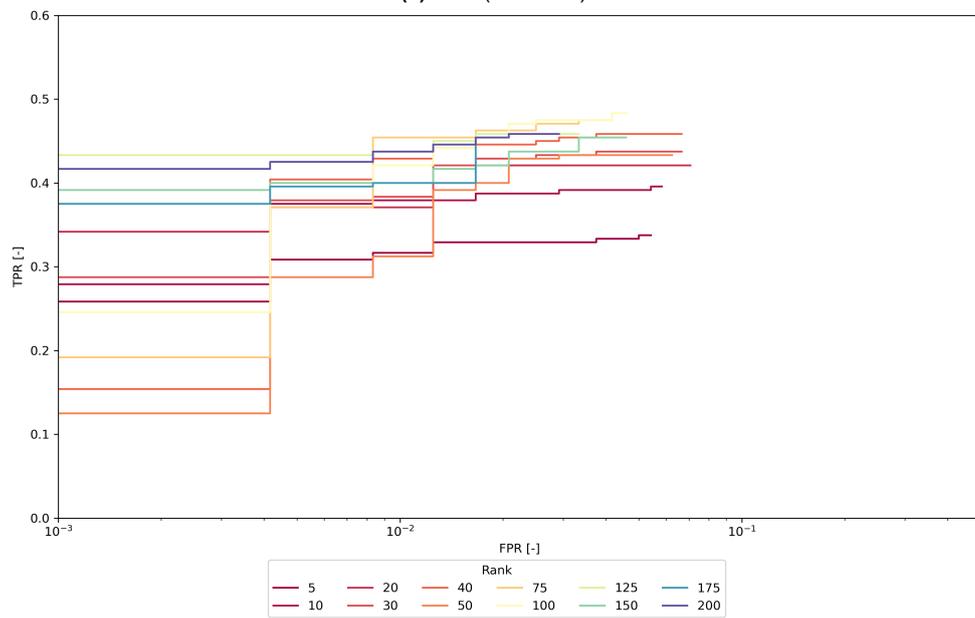


(b) Annular PCA.

Figure D.1: PCA: Extended receiver operating characteristic curves.

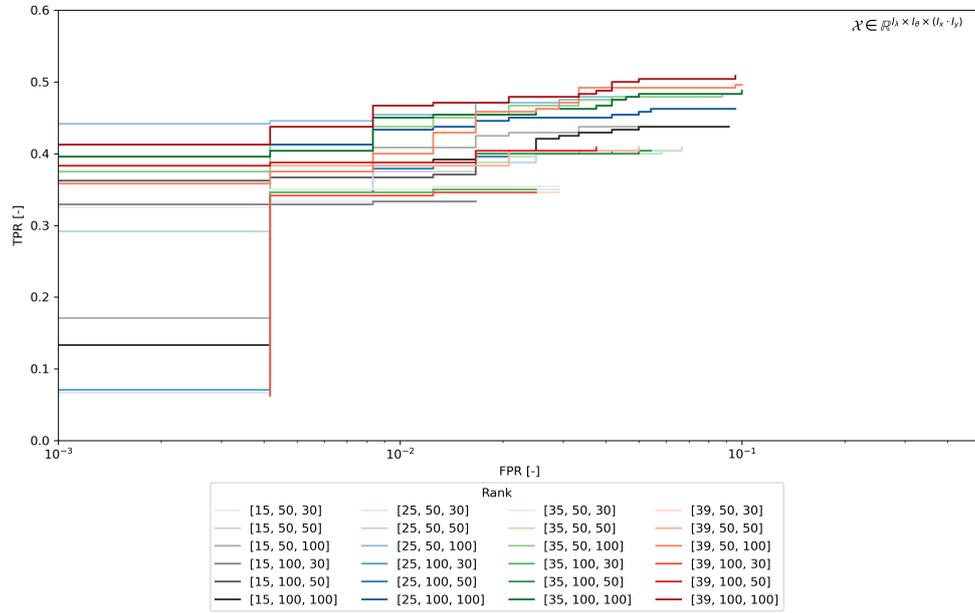


(a) CPD (full-frame).

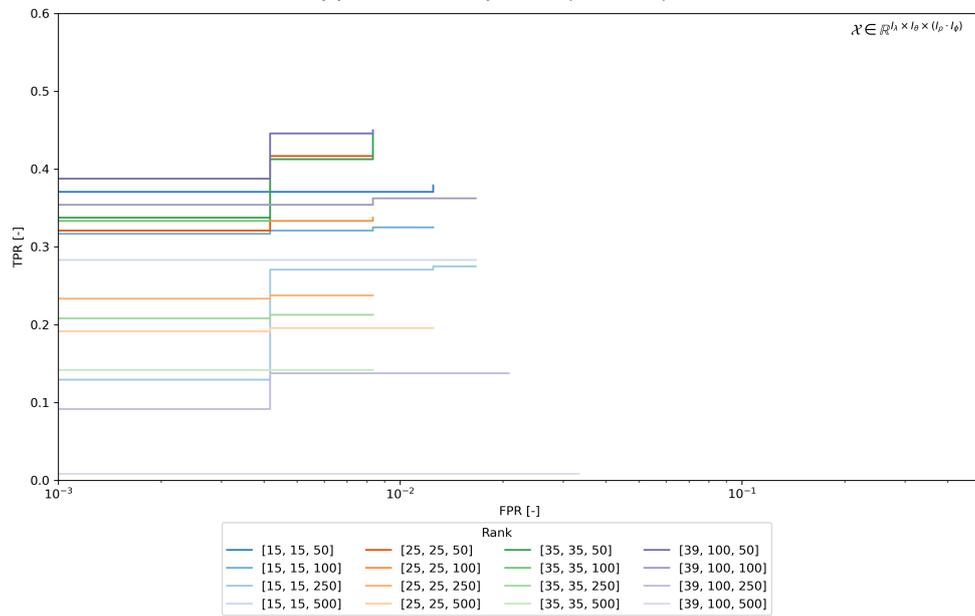


(b) Annular CPD.

Figure D.2: CPD: Extended receiver operating characteristic curves.

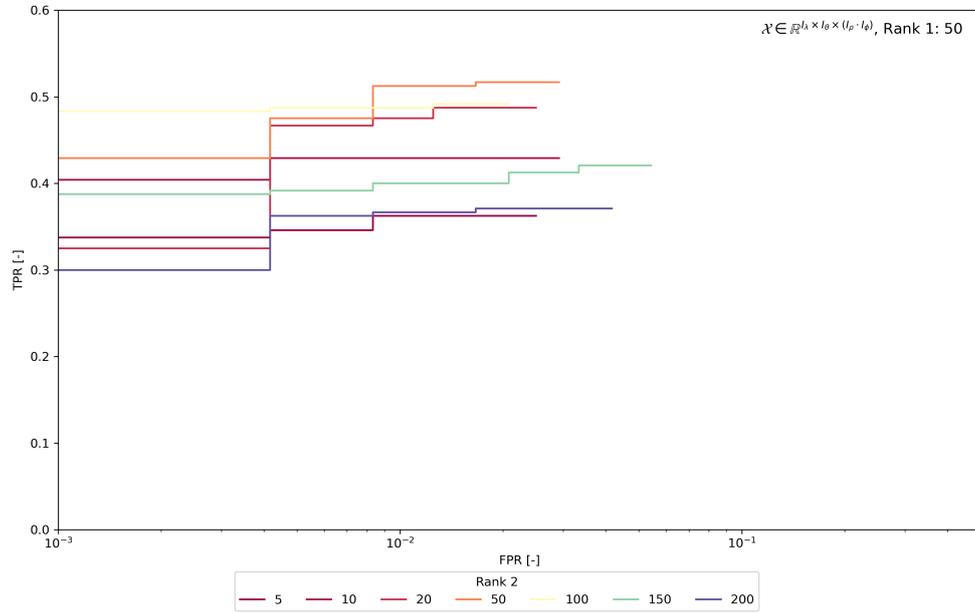


(a) Tucker decomposition (full-frame).

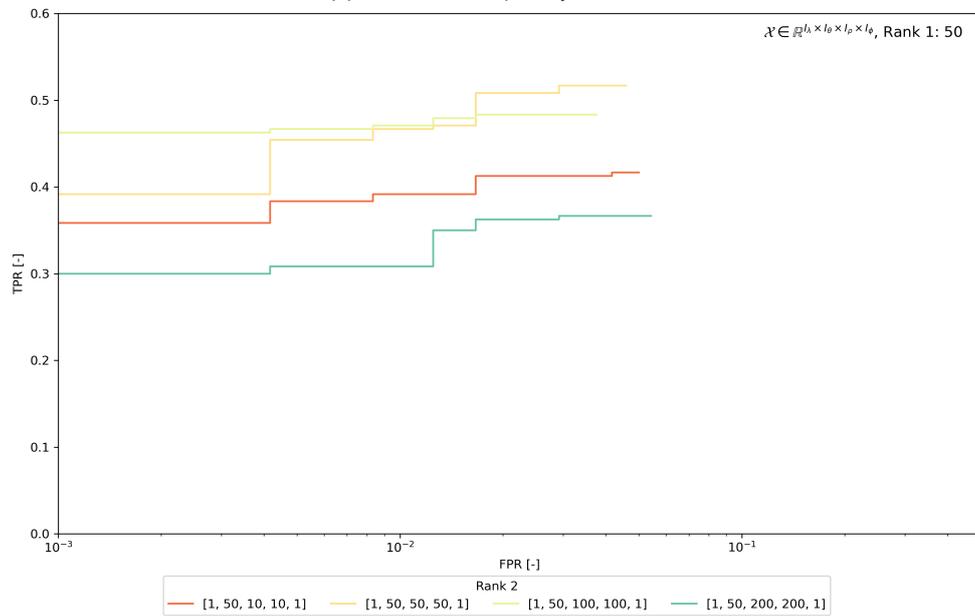


(b) Annular Tucker decomposition.

Figure D.3: Tucker decomposition: Extended receiver operating characteristic curves.



(a) Annular TTD, spatially unfolded.



(b) Annular TTD, order-4.

Figure D.4: Annular TTD: Extended receiver operating characteristic curves.

The algorithms developed for this work are publicly available at the [TIDE repository](#). The algorithms are compatible as with VIP as drop-in contributions, as long as the requirements for the tensor decompositions are installed. There are two principal methods of using the algorithms; a) they can be accessed using the factory by specifying the name of the decomposition and its properties. This dynamically creates the appropriate subclass and is the intended interface for their use. b) alternatively the user can manually subclass the designed algorithms. This interface is intended for users who want to build on the existing algorithms. TIDE includes the polar transforms that were developed for the annular methods, however, they are also published in a separate [polar-transform repository](#). The TIDE framework is documented in-depth in [Appendix F](#).

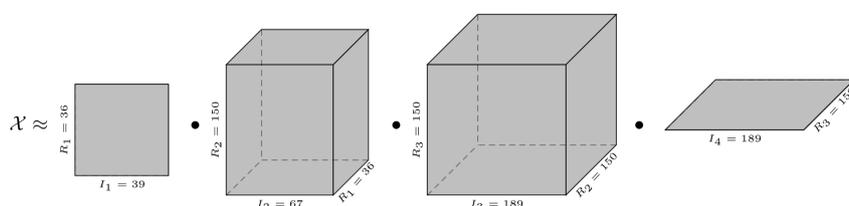


Figure E.1: `td-viz` drawing of a rank-[36, 150, 150] tensor train decomposition of an order-3 tensor.

In addition to the algorithms specific to HCI, this work also developed a library to produce publication quality `tikz` figures of tensor decompositions for use in Latex, which is publicly available at the [td-viz repository](#). The library can automatically produce both graphic 3D visualizations of tensor decompositions and tensor network diagrams of tensor decompositions. Both diagram generators support automatic labeling of indices and ranks. The authors hope that this tool can support research into tensor decompositions by significantly simplifying the TND drawing process and increasing its consistency and quality. The TND in this work were created using `td-viz`, and [Figure E.1](#) shows an example of the 3D graphics that can be produced with `td-viz`.

TIDE Framework Implementation

The tensor-based PSF modeling and subtraction methods have been developed as a modular and extensible Python framework which generates the specific post-processing methods via composition. This design choice allows for seamless integration of the newly developed methods into existing high-contrast imaging packages. Moreover, the decoupling of image transforms, CODI, and the PSF modeling methods provides the flexibility to independently evolve or expand and add to the tensor post-processing framework.

The implementation is organized around a hierarchy of abstract base classes, each representing a different aspect of the CODI methodology. At the root of this hierarchy is an abstract base CODI class, which encapsulates the core functionality common to all PSF subtraction methods.

Building on this abstract base class, the framework implements specialized classes for different coordinate systems and processing strategies. The abstract base polar class extends the base functionality to work in polar coordinates, utilizing a polar transform module for efficient coordinate transformations. Further specialization is achieved through the abstract base annular abstract base single annulus classes, which implement annular processing techniques.

The tensor decomposition PSF modeling methods, are implemented as separate classes. These classes define the specific model that perform the tensor decompositions which are used to find a PSF model. They inherit from an abstract base tensor decomposition class which defines the common interface between the CODI methods and the PSF models.

A key feature of the framework is its use of dynamic class composition to generate specific implementations. This is achieved through a factory module that constructs a requested post-processing class. The factory dynamically creates this class through composition of a CODI method and PSF model. This approach allows for great flexibility in creating tailored PSF modeling techniques without the need for extensive code duplication. It also facilitates easy experimentation with different combinations of processing strategies and decomposition methods.

Example F.1: Creation and Use of Post-Processing Algorithm

Assume that a user wishes to reduce an ASDI observation with annular CODI and a TT PSF model. Since the user is familiar with the Vortex Image Processing package they would like to use it to load and process the data. For this the user simply declares the specific algorithm to the constructor function and optimally passes the parameters of the PSF model:

```
algorithm = get_dynamic_subclass_by_name("VIPAnnularTensorTrainDecomposition")()
```

The specific subclass is then automatically constructed and fully usable. If the user has the ASDI observations in memory, e.g. as numpy arrays via the Vortex Image Processing package, they can simply process them like:

```
residuals = algorithm(
    cube=observation,          # ASDI data
    angle_list=angle_list,    # parallactic angles
    scale_list=scale_list     # wavelength bins
)
```

where the implementation follows the VIP syntax. Internally, the implementation processes the data on CPU or GPU depending on the available hardware.

The implementation places a strong emphasis on compatibility with existing high-contrast imaging tools, particularly the Vortex Image Processing package (VIP, [Christiaens et al., 2023](#); [Gonzalez et al., 2017](#)). This compatibility is achieved through a wrapper interface that allows the new methods to be seamlessly integrated into VIP workflows. The wrapper handles the necessary data conversions and ensures that tensor methods can be used as drop-in replacements for existing PSF subtraction techniques within the VIP framework.

This integration strategy serves multiple purposes. First, it allows for easy adoption of the new methods by researchers already familiar with VIP, lowering the barrier to entry for using these new techniques. Second, it facilitates direct comparisons between tensor methods and traditional matrix-based approaches, as both can be run within the same processing pipeline. Finally, it ensures that the new methods can immediately benefit from VIP's extensive suite of pre- and post-processing tools, data handling capabilities, and performance metrics.

By default the implementation leverages native GPU acceleration for fast and parallel computation. However, the methods are compatible with other Python computing frameworks like `numpy` in case users do not have access to accelerators. This connection with the `PyTorch` ecosystem also allows the methods to be accessible as modules (`torch.nn.Module`), which enables TIDE to accomplish two tasks: fast and efficient re-use of existing modules for repeated processing of similar data and automatic differentiation with `torch.autograd`. ([Paszke et al., 2019](#))

Processing similar data is a common task in HCI, e.g. the repeated reduction of observations with injected false companions to find an exoplanet model or the processing of a larger dataset obtained from a specific instrument. TIDE methods implemented as modules intrinsically allow for large-scale, and distributed processing via integration with the `PyTorch` ecosystem. The main motivation behind this choice is a large observation post-processing effort at the HC-DC by [Chomez et al., 2023a](#).

Automatic differentiation is relatively new to HCI, remaining mostly exclusive to ADI due to the data volumes associated with ASDI observation post-processing. Nevertheless, it is becoming clear that deep learning approaches are powerful for detecting exoplanet in direct imaging observations. ([Cantero et al., 2023](#); [Flasseur et al., 2024](#); [Wolf et al., 2023](#)) Additionally, tensor decomposition methods are promising for deep learning methods in HCI as noted in [Chapter 8](#). Hence, TIDE allows for automatic differentiation of its methods via `PyTorch`. Nevertheless, this could not be extensively verified beyond simplified ASDI test cases due to hardware limitations together with required additional computational resources for tracking gradients in the computational graph. ([Paszke et al., 2019](#))

High-Level Tensor Software

A number of computational tensor frameworks have been released, targeting different communities and languages. The two main communities developing tensor software are data science and quantum many-body systems, with the main language in the former being MatLab and both groups using C++ and Python. The fragmented nature of the tensor computation landscape when compared to matrix computations seems to be mostly due to the application driven development of tensor software (Psarras et al., 2022). This has led to a large amount of redundant work and lack of common frameworks and standardization. In Table G.1 this work collects a number of tools that are relevant to this work, focusing on the Python language due to its predominance in the HCI community. Special formatting indicates software of interest in this work. The purpose of this section is to develop an understanding of the available high-level tools and how they can be used for work on tensor methods for HCI. For a more complete overview see Psarras et al., 2022 as well as Novikov et al., 2020, The Tensor Network 2023, and Sobral et al., 2023 for complete, tensor network, Tensor Train, and RPCA focused surveys respectively.

Table G.1: Overview of relevant tensor processing and multi-linear algebra software. Software marked in ***bold italics*** was further explored and is recommended for use in this work. Software marked in **bold** might be considered at a later stage since it offers utility that is unique to it but is less flexible and general than other software. Software marked in *italics* was considered for this work but would require a radical shift in the HCI community either in the way that images are understood or in the software that is used. Lastly, software that is not specifically marked is less suitable to HCI than the other options. Note that for all software based on Python the list of backends does not mention NumPy, which is available as a backend for each of the listed packages.

Name	Language	Backend	Uses	Notes	References
PyTensor	Python	C, Numba, JAX	multi-linear algebra	Extends NumPy like functionality to higher-order tensors.	PyTensor, 2023
TensorNetwork	Python	PyTorch, TensorFlow, JAX	tensor networks	Deep learning and quantum circuit focus, with wide support for algorithms in both fields but at very low level. Interprets images as high-order (entangled) quantum systems.	Google/TensorNetwork 2023 ; Roberts et al., 2019
<i>TensorLy</i>	Python	PyTorch, TensorFlow, JAX	tensor learning	Wide support for multi-linear algebra, from low-level to high-level methods. Wide range of decompositions, RPCA, CPD, TTD, TD, TRD, but not very flexible. Focus on deep and tensor learning. Similar to tntorch.	Kossaifi et al., 2019
<i>tntorch</i>	Python	PyTorch	tensor learning	Very big library building on TensorNetwork. Supports several formats, CPD, TTD, TD, ANOVA, and hybrids. Flexible support for decomposition, algebra, auto-differentiation, cross-approximation and sensitivity analysis. Build for deep learning and tensorized neural networks. Similar to TensorLy.	Usvyatsov et al., 2022

Continued on next page

Table G.1 – continued from previous page

Name	Language	Backend	Uses	Notes	References
Tensor Fox	Python	Numba	CPD	Multilinear algebra routines mainly based around CPD, supports sparse tensors. Includes rank determination solvers. Has some support for other decompositions like TTD. Has IO for MatLab based tensor software.	Bottega et al., 2023
<i>FunFact</i>	Python	PyTorch, Tensor-Flow, JAX	tensor factorization, expressions	Provides support for tensor factorization and abstract tensor expressions. Main focus is on quantum circuits but the tensor expression utility is very flexible.	Tang et al., 2023
quimb	Python	autoray (PyTorch, Tensor-Flow, JAX)	tensor networks	Tensor network manipulations for quantum information and many-body calculations. Includes algorithms (DMRG, TEBD) and decompositions (MPS, PEPS, PEPO, MERA) from the quantum circuit community.	Gray, 2018
contenga	Python	quimb: autoray (PyTorch, Tensor-Flow, JAX)	tensor network contraction	Focus on high-order tensor decompositions with a hyper-optimizer for efficient contraction path sampling. Main application is quantum circuit simulations.	Gray, 2023 ; Gray et al., 2021
<i>TensorLab</i>	MatLab	NA	tensor methods	Flexible toolbox for tensor computations and decompositions for low-medium order. Supports incomplete, sparse, or structured tensors. Wide range of factor transformations. Support for tensorization, structured data fusion, ICA, and PCA. Manual is a great reference.	Vervliet et al., 2016

This work recommends two software packages, TensorLy ([Kossaifi et al., 2019](#)) and tntorch ([Usvyatsov et al., 2022](#)), both of which build on several deep learning frameworks for low-level functionality and NumPy ([Harris et al., 2020](#)) like high-level abstraction. These packages are mature and well tested and allow for an easy drop in for the HCI community which is familiar with Python and NumPy. Both TensorLy and tntorch offer tensor methods on a low- and high level of abstraction for tensor methods, especially for decompositions. While TensorLy offers a more diverse set of decompositions and more flexibility in handling the decompositions, tntorch has better support for cross-approximation and auto-differentiation as well as optimization. tntorch also includes some variance based decompositions that are unique to it. Both packages can make use of hardware acceleration, are massively parallel and thus very fast for tensor decompositions. Since tntorch is only available for PyTorch ([Paszke et al., 2019](#)), which is also the dominant deep learning framework for research, this work recommends focusing on PyTorch as the backend.

Tensor Fox ([Bottega et al., 2023](#)) and FunFact ([Tang et al., 2023](#)) are both libraries that offer some unique methods which are not included in TensorLy and tntorch. For Tensor Fox this is a strong focus on CPD with more flexibility than any other software. FunFact implements a flexible interface for abstract tensor expressions which could be of use for more fine-grained control of decompositions in HCI. The packages that have a strong focus on quantum circuits (quimb ([Gray, 2018](#)) and contenga ([Gray et al., 2021](#))) are interesting if understanding images less like spatial tensors and more like entangled quantum states (each voxel) is useful for source separation. This exploration is beyond the scope of this work.