

Coronagraphic data post-processing using projections on instrumental modes

by

Yeyuan (Yinzi) Xin

B.S., California Institute of Technology (2018)

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2020

© Massachusetts Institute of Technology 2020. All rights reserved.

Author
Department of Aeronautics and Astronautics
May 18, 2020

Certified by
Kerri Cahoy
Associate Professor
Thesis Supervisor

Accepted by
Sertac Karaman
Chair, Graduate Program Committee

Coronagraphic data post-processing using projections on instrumental modes

by

Yeyuan (Yinzi) Xin

Submitted to the Department of Aeronautics and Astronautics
on May 18, 2020, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

High contrast astronomy has yielded the direct observations of over a dozen exoplanets and a multitude of brown dwarfs and circumstellar disks. Despite advances in coronagraphy and wavefront control, high contrast observations are still plagued by residual wavefront aberrations. Post-processing techniques can provide an additional boost in separating residual aberrations from an astrophysical signal. This work explores using a coronagraph instrument model to guide post-processing. We consider the propagation of signals and wavefront error through a coronagraphic instrument, and approach the post-processing problem using "robust observables." We model and approximate the instrument response function of a classical Lyot coronagraph (CLC) and find from it a projection that removes the dominant error modes. We use this projection to post-process synthetically generated data, and assess the performance of the new model-based post-processing approach compared to using the raw intensity data by calculating their respective flux ratio detection limits. We extend our analysis to include the presence of a dark hole using a simulation of the CLC on the High-contrast imager for complex aperture telescopes (HiCAT) testbed. We find that for non-time-correlated wavefront errors, using the robust observables modestly increases our sensitivity to the signal of a binary companion for most of the range of separations over which our treatment is valid, for example, by up to 50% at $7.5\lambda/D$. For time-correlated wavefront errors, the results vary depending on the test statistic used and degree of correlation. The modest improvement using robust observables with non-time-correlated errors is shown to extend to a CLC with a dark hole created by the stroke minimization algorithm. Future work exploring the inclusion of statistical whitening processes will allow for a more complete characterization of the robust observables with time-correlated noise. We discuss the dimensionality of coronagraph self-calibration problem and motivate future directions in the joint study of coronagraphy and post-processing.

Thesis Supervisor: Kerri Cahoy
Title: Associate Professor

Acknowledgments

I want to thank Kerri Cahoy, my advisor, for supporting me and my various research interests throughout these couple of years, and for the genuine care and support she has for me and all of her students. I want to thank Laurent Pueyo for somehow believing in this project from the moment I (basically a stranger at the time) mentioned it, for sponsoring my work on it at STScI, and for all of the great mentorship and guidance he has provided since. I want to thank Ewan Douglas for teaching me much of what I know about optics and what real astronomy is like, for maintaining the perfect balance of interest and skepticism for my half-baked ideas, and for being a supportive mentor both for this project and in general. I want to thank Romain Laugier for taking the time to help me with many of the topics involved in this investigation, for his valuable suggestions for directions to explore, and for enlightening me about the term "coronagram." Thank you as well to Ben Pope, Frantz Martinache, Mamadou N'Diaye, and Alban Ceau for our helpful conversations on this topic.

I want to acknowledge Rachel Morgan and Greg Allan, my labmates and friends at MIT, as well as the rest of the MIT DeMi team, with whom I've shared many fun experiences and many stressful experiences, and who have made those experiences memorable. Thank you also to the Stuffed Animal Party: Jules Fowler, Iva Leginja, Scott Will, Theo Jolivet, Kelsey Glazer, Evelyn McChesney, and Remi Soummer, who have made me feel incredibly welcome at STScI.

Software: This research made use of PROPER, an optical propagation library [1]; hicat-package, the software infrastructure that models and runs the HiCAT testbed; POPPY, an open-source optical propagation Python package originally developed for the James Webb Space Telescope project [2]; HCIPy, an open-source object-oriented framework written in Python for performing end-to-end simulations of high-contrast imaging instruments [3]; Astropy [4, 5]; NumPy [6]; SciPy [7]; and Matplotlib [8].

Contents

1	Introduction and Background	12
1.1	Introduction	12
1.2	Instrumentation	13
1.2.1	Coronagraphy	14
1.2.2	Wavefront Sensing and Control (WFSC)	15
1.3	Post-processing	17
1.3.1	Self-calibration, Robust Observables, and Kernel Methods	17
1.3.2	Coronagraphic Post-processing	19
1.3.3	Compatibility of Methods	20
1.4	Thesis Contributions	20
2	Signals and Wavefront Error in Coronagraphs	22
3	Classical Lyot Coronagraph (Simulation)	26
3.1	Optical Model	26
3.2	Coronagraph Properties	27
3.3	Building Response Matrices	31
3.4	Singular Value Decomposition, Singular Modes, and the Projection K	33
3.5	Synthetic Data Analysis	35
3.5.1	Synthetic Data Generation	35
3.5.2	Processing Synthetic Data	36
3.5.3	Detection Tests	38
3.5.4	Receiver Operating Characteristic Curves	39

3.5.5	Detection Limits	40
3.5.6	Time-correlated Wavefront Errors	42
3.6	Summary of Results	42
4	Classical Lyot Coronagraph with Dark Hole (HiCAT Simulation)	46
4.1	Background	46
4.2	Optical Model	47
4.3	Dark Holes with Stroke-Minimization	49
4.4	Building a Response Matrix	51
4.5	Singular Value Decomposition, Singular Modes, and the Projection K	52
4.6	Synthetic Data Analysis	53
4.6.1	Synthetic Data Generation	53
4.6.2	Synthetic Data Analysis	55
4.6.3	ROCs and Detection Limits	55
4.7	Summary of Results	58
5	Discussion	59
5.1	Lack of Statistical Whitening	59
5.2	Kernel-Nulling	60
5.3	Connection with PASTIS Formalism	60
5.4	Dimensions for Continuous Coronagraphy	61
6	Future Work and Conclusions	65
A	Heuristic Details of Robust Observables with the CLC	67

List of Figures

1-1	a) The nominal coronagraph (in \log_{10} scale) of a central star with an off-axis point source at $7.5\lambda/D$ with a contrast of 4×10^{-4} . b) An example of phase errors (nm) in the wavefront at the primary telescope pupil. c) A coronagraph if the same system in a) in the presence of the wavefront errors shown in (b).	15
1-2	A typical AO system with a deformable mirror and wavefront sensor in closed loop control, and a downstream science instrument such as a high-resolution camera. [9]	16
1-3	Pictorial representation of kernel phases. The operator \mathbf{A}_ϕ injects phase aberrations ϕ into a subspace of baseline phases Φ . There is another subspace of baseline phases, $\mathbf{K}_\phi\Phi_0$, which are not affected by these phase aberrations, and we call these kernel phases. [10]	19
3-1	Optical model for the Lyot coronagraph, included as an example in the PROPER optical propagation library [1]. The coronagraph operates at 585 nm. The occulter is a circular obscuration with radius 0.561 μm , and the Lyot stop is a circular aperture with radius 2.02 m. The Fraunhofer regime is assumed, so the model consists of only 2D Fourier transforms and masks, and resizing and propagation distances are not considered.	27
3-2	The nominal coronagraph of just a central star in \log_{10} scale of arbitrary intensity units.	28
3-3	A radial slice of the nominal coronagraph residuals on a central star.	29

3-4	a) The radial slices of a mostly linear, a mostly quadratic, and an in-between sinusoidal speckle in the detector plane plotted against the coronagraph residuals. b) The intensity with a predominantly linear speckle, which mostly interferes with the nominal star PSF. c) The intensity with a predominantly quadratic speckle that is mostly the square-amplitude of the wavefront error itself. d) The intensity with a speckle that is neither dominantly linear nor quadratic, but somewhere in-between.	30
3-5	The throughput of an off-axis point source as a function of its spatial separation, indicated by the solid line. The inner working angle is the separation at which the throughput is 50% (indicated by the dashed line), which in this case is $5\lambda/D$	31
3-6	The singular values of the transfer matrix A. We choose the 102nd through 10,000th singular modes to make up our projection matrix K, which will project out the first 101 modes from our data and result in $N_k = 9899$	34
3-7	The first 10 singular modes of A_Z as represented in the the detector plane intensity basis.	35
3-8	Top left: A single aberrated snapshot (in linear scale) with a companion with a flux ratio of 4×10^{-4} at $7.5\lambda/D$. Top right: The data frame with the model coronagraph subtracted. Bottom left: The mean of 20 frames of data, with the companion "location" (visible by eye) marked with the red box. Bottom right: The standard deviation of 20 frames of data.	37
3-9	Left: The histograms over 1000 datasets of the test statistic $\Delta\chi_r^2$ for raw intensity and observables, both with a companion with flux ratio of 4×10^{-4} at $7.5\lambda/D$ and without it. Right: The histograms over 1000 datasets of the test statistic $\Delta\epsilon^2$ for raw intensity and observables, with and without the same companion.	40
3-10	Left: The receiver operating characteristic curve, or true positive rate as a function of false positive rate, for the $\Delta\chi_r^2$ statistic for a companion with a flux ratio of 4×10^{-4} at $7.5\lambda/D$. Right: The receiver operating characteristic curve for the $\Delta\epsilon^2$ statistic for the same companion	41

3-11	Left: The $\text{FPR} = 0.01$, $\text{TPR} = 0.90$ detection limits (at the 10^{-4} resolution) as a function of separation for raw intensities and the observables, using the $\Delta\chi_r^2$ test statistic. Right: The $\text{FPR} = 0.01$, $\text{TPR} = 0.90$ detection limit (at the 10^{-4} resolution) as a function of separation for raw intensities and the observables, using the $\Delta\epsilon^2$ test statistic.	41
3-12	Left: Mean of 20 frames with energy correlation parameter $\alpha = 0.1$. Right: Mean of 20 frames with $\alpha = 0.5$	43
3-13	$\text{FPR} = 0.01$ and $\text{TPR} = 0.90$ detection curves. Top left: $\alpha = 0.1$ using the $\Delta\chi_r^2$ statistic. Top right: $\alpha = 0.1$ using the $\Delta\epsilon^2$ statistic. Bottom left: $\alpha = 0.5$ using the $\Delta\chi_r^2$ statistic. Bottom right: $\alpha = 0.5$ using the $\Delta\epsilon^2$ statistic.	44
3-14	Histograms of the remaining errors after processing a dataset of 20 frames, uncorrelated in time, for intensities and observables.	45
4-1	Full optomechanical layout of the HiCAT testbed. The testbed has the capability to transition between an APLC mode (with the Iris-AO DM simulating a hexagonally segmented pupil) and a CLC mode (with a simple circular pupil). [11].	47
4-2	Left: view of the entire HiCAT testbed in its (opened) enclosure. Right: detail of the central part of the testbed. [11].	48
4-3	Top left: Contrast in intensity after 80 iterations of stroke minimization. Top middle: Estimated complex electric field in dark hole, where brightness indicates arbitrary amplitude units and color indicates phase as per visualization by HCIPy [3]. Top right: The radial profile in contrast before and after the 80th iteration, with the probe magnitude shown as a diagnostic. Bottom left: The surface profile of DM 1. Bottom right: The surface profile of DM 2. All two-dimensional maps are square, where the Y axis has the same units and spacing as the X axis.	50
4-4	The singular values of the transfer matrix A. We choose the 101st through 31,684th singular modes to make up our projection matrix K, which will project out the first 100 modes from our data and result in $N_k = 31584$	53

4-5	The first 10 singular modes of A represented in the the detector plane intensity basis. For visualization, the intensity map is zoomed in and pixels not in the dark hole are masked out.	54
4-6	Top left: A single aberrated snapshot (in linear scale) with a companion with a flux ratio of 4×10^{-8} at $10.0\lambda/D$. The location of the companion is unmarked as it is not actually identifiable by eye. Top right: The data frame with the model coronagram subtracted. Bottom left: The mean of 20 frames of data. Bottom right: The standard deviation of 20 frames of data.	56
4-7	Left: The histograms over 40 datasets of the test statistic $\Delta\chi_r^2$ for raw intensity and observables, both with a companion with a flux ratio of $c = 4 \times 10^{-8}$ at $10\lambda/D$ and without it. Right: The histograms over 40 datasets of the test statistic $\Delta\epsilon^2$ for raw intensity and observables, with and without the same companion.	57
4-8	Left: The receiver operating characteristic curve, or true positive rate as a function of false positive rate, for the $\Delta\chi_r^2$ statistic for a companion with a flux ratio of $c = 4 \times 10^{-8}$ at $10\lambda/D$. Right: The receiver operating characteristic curve for the $\Delta\epsilon^2$ test statistic for the same companion. With both of these test statistics, using the robust observables results in better performance for detections.	57
4-9	Left: The $FPR = 0.1$, $TPR = 0.9$ detection limit (at the 10^{-8} resolution) as a function of separation for raw intensities and the observables, using the $\Delta\chi_r^2$ statistic. Right: The $FPR = 0.1$, $TPR = 0.9$ detection limit (at the 10^{-8} resolution) as a function of separation for raw intensities and the observables, using the $\Delta\epsilon^2$ statistic.	58

5-1	The singular value spectrum of a full-dimensional 10000×38226 A matrix. The drop at around 2000 seems to be a numerical effect independent of the size of the input space. Applying the left singular matrix U to synthetic data reveals that almost all of the astrophysical signal projects onto modes before the drop. The singular values are smooth before the drop, indicating that there is no kernel that also overlaps well with astrophysical signal.	63
5-2	The first 10 singular modes of the full dimensional A matrix as represented in the detector plane intensity basis. These modes are qualitatively similar to the approximated singular modes in Figure 3-7, and using a projection determined from this matrix to analyze synthetic data does not result in noticeable differences in flux ratio detection limits compared to using the approximation neglecting cross-terms.	64
A-1	Top left) The mean intensity of 20 non-time-correlated ($\alpha = 1$) frames with the model star coronagraph subtracted of a companion with a flux ratio of 4×10^{-4} at $7.5\lambda/D$. Top right) The mean intensity data plotted against the model intensity data for the system, where the black line indicates a 1-to-1 correspondence. Bottom left) The mean observables data plotted against the model observables data for the system, where the black line indicates a 1-to-1 correspondence. Bottom right) The mean $U^T I$ data plotted against the model $U^T I$ data for the system, where the black line indicates a 1-to-1 correspondence.	68
A-2	Top) Data corresponding to a companion with a flux ratio of 4×10^{-4} at $7.5\lambda/D$, and no time-correlation $\alpha = 1$. Middle) Data corresponding to a companion with a flux ratio of 5×10^{-4} at $7.5\lambda/D$, and time-correlation parameter $\alpha = 0.5$. Bottom) Data corresponding to a companion with a flux ratio of 9×10^{-4} at $7.5\lambda/D$, and time-correlation parameter $\alpha = 0.1$.	70

List of Tables

3.1	Sampling quantities for the classical Lyot coronagraph model	27
3.2	Quantities and dimensions for analysis of the classical Lyot coronagraph . . .	32
4.1	Physical dimensions for the HiCAT CLC	47
4.2	Sampling quantities for the HiCAT CLC	49
4.3	Quantities and dimensions for analysis of the HiCAT CLC	51

Chapter 1

Introduction and Background

1.1 Introduction

High contrast astronomy is the direct observation of faint objects near bright objects, such as exoplanets, brown dwarfs, or circumstellar disks around much brighter central stars. High contrast techniques are extremely useful for answering many scientific questions, including those in binary and planetary system population statistics, planet and disk formation and evolution, planetary atmospheres, and planet habitability and the search for biosignatures [12]. However, many technical challenges must be overcome in order to discern and measure the signal of the faint and often close-in objects in the presence of multiple sources of noise.

In general, high contrast astronomy faces two major obstacles: 1. photon noise from the light of the central star, and 2. wavefront error from phase and amplitude aberrations of the optical wavefront, such as those caused by atmospheric turbulence, imperfections of the optics themselves, or thermo-mechanical changes in the telescope or instrument.

In the visible spectrum, photon noise is dominated by shot noise associated with the detection of discrete photons, for which the standard deviation scales as the square root of the number of photons detected. Practical matters such as detector saturation aside, if the star is many orders of magnitude brighter than its companion, the shot noise associated with the fainter outer lobes of star's point-spread-function (PSF) can obfuscate any signal from the companion, even if the on-axis star's signal is perfectly known. Although the ratio of signal to shot noise can be improved by increasing the integration time, detecting

and characterizing faint exoplanets with typical (non-coronagraphic) imagers may require impractical exposure times [13].

The other primary source of noise is wavefront noise, which distorts the signal of the on-axis point source. At an instantaneous point in time, the effect of a perturbation to the wavefront is to scatter energy from the core of the PSF into speckles throughout the image that can resemble off-axis sources. If the magnitude of the electric field scattered outwards by the wavefront error is smaller than that of the underlying electric field from the point-spread-function, then the speckles are symmetric about zero in the detector plane intensity and will average out over time. However, when the wavefront error is larger, as in the case of atmospheric wavefront error in the absence of correction, then the speckles are predominantly positive and will average out to a halo that can obscure off-axis signals. The increase in intensity at larger spatial separations also contributes to an increase in the shot noise at those locations in the detector plane intensity.

In this chapter, we briefly overview the relevant instrumentation and post-processing techniques that have been developed to mitigate the effects of both shot and wavefront noise, in order to provide context for the work in this thesis on developing an instrument model-based approach to post-processing.

1.2 Instrumentation

The primary goal of high contrast instrumentation is to be able to separate the signal of the on-axis star from any off-axis sources. One approach is to physically remove as much of the on-axis starlight as possible from the measurement, while maintaining the throughput of off-axis sources. Coronagraphs are passive optical elements that produce a coronagram¹ from which the signal of an unaberrated on-axis source has been filtered out. Adaptive optics and other wavefront control techniques such as electric field conjugation are active methods of reducing the starlight scattered by wavefront aberrations by measuring those aberrations

¹The term "coronagram" is not typically used in the literature for the data resulting from a coronagraphic instrument, which is usually referred to as a coronagraphic image. However, "image" is a bit of a misnomer as it implies a translationally invariant imaging system, of which a coronagraph is not one. We thus use "coronagram" instead to refer to the intensity measurements from coronagraphs.

and correcting for them. One example of an active element used for wavefront correction is a deformable mirror (DM). Together, coronagraphy and wavefront control provide significant on-axis starlight suppression, and play essential roles in current and proposed high contrast instruments.

1.2.1 Coronagraphy

The first successful design and use of a coronagraph is attributed to Bernard Lyot, who used it to block the light of the sun in order to observe the solar corona. This classical Lyot coronagraph design involves using a solid occulter at a focal plane to scatter the light to the outer regions of the optical beam, where it is then blocked by an iris dubbed the Lyot stop [14]. A schematic of a classical Lyot coronagraph instrument can be found in Figure 3-1. Coronagraph designs have since proliferated, and the ability of coronagraphs to remove on-axis light, both suppressing its associated photon noise and allowing for deeper exposures without saturating the detector, has become indispensable in the hunt for exoplanets and circumstellar disks.

Many space telescopes include coronagraphic instruments. The Hubble Space Telescope (HST) is equipped with several simple coronagraphs, in both its Advanced Camera for Surveys (ACS) [15] and Space Telescope Imaging Spectrograph (STIS) [16] instruments. Future space telescopes with coronagraphic instruments include the James Webb Space Telescope (JWST) [17], Wide Field Infrared Survey Telescope (WFIRST) [18], the Habitable Exoplanet Observatory (HabEx) [19], and the Large Ultraviolet Optical Infrared Surveyor (LUVOIR) [20].

An overview of some promising coronagraph designs for extrasolar high contrast astronomy can be found in Guyon *et al.* (2006) [21]. These coronagraphs can be grouped into four main categories: interferometric coronagraphs based on interfering multiple copies of the optical beam, apodization coronagraphs based on masking the pupil, Lyot coronagraph designs with modified phase masks, and Lyot coronagraph designs with modified amplitude masks. This work will focus on the latter three types of coronagraphs, where there is only one copy of the entrance pupil that propagates through a series of masks and other optical elements.

Although there are "perfect" coronagraph design for circular apertures at a single wavelength, *i.e.* coronagraphs designed to remove all of the on-axis starlight when it is a perfectly flat wavefront [21], the presence of wavefront aberrations will still cause light to leak through (see Figure 1-1). Thus, coronagraphy used in conjunction with wavefront control techniques is needed in order to reach the highest sensitivities.




Figure 1-1: a) The nominal coronagraph (in log10 scale) of a central star with an off-axis point source at $7.5\lambda/D$ with a contrast of 4×10^{-4} . b) An example of phase errors (nm) in the wavefront at the primary telescope pupil. c) A coronagraph if the same system in a) in the presence of the wavefront errors shown in (b).

1.2.2 Wavefront Sensing and Control (WFSC)

The wavefront of the on-axis light can be actively measured, and a deformable mirror (DM) can be used to control it. There are various methods used to measure and control the wavefront, but two main approaches (which can also be used in conjunction with one another) are discussed here.

Adaptive Optics

Adaptive optics (AO) is a technique that uses a designated wavefront sensor (WFS) to measure phase perturbations in the input pupil, and uses the DM to flatten the wavefront. AO is typically used on ground-based telescopes to counteract the effect of atmospheric wavefront error for a wide variety of astronomical applications. Figure 1-2 shows a typical AO setup, in which a wavefront with perturbations (such as those caused by the atmosphere) enters the telescope and is reflected off the DM. The wavefront is then split by a beamsplitter,

and a portion of the light goes to the WFS which provides measurements of the wavefront error residuals that are fed back to the DM in closed loop control. The other portion of the light goes to focal plane science camera, or alternatively, to another downstream science instrument such as a coronagraph.




Figure 1-2: A typical AO system with a deformable mirror and wavefront sensor in closed loop control, and a downstream science instrument such as a high-resolution camera. [9]

A common metric for the performance of an AO system is the Strehl ratio, which is defined as the ratio of the peak of the aberrated PSF to the peak of the ideal PSF without aberrations [22]. Assuming that phase errors dominate the wavefront error, the Strehl ratio S is also approximately related to their standard deviation σ_{ph} by $S \approx e^{-\sigma_{ph}^2}$.

In general, AO, motivated by achieving good imaging resolution, aims to mitigate phase errors to achieve a flat wavefront as measured by the wavefront sensor regardless of the actual downstream science instrument. AO does not correct for any optical aberrations resulting from optics within a downstream instrument such as a coronagraph, as such non-common-path aberrations are not sensed by the wavefront sensor. Thus, other wavefront control methods are required to correct for non-common-path errors.

Focal Plane Sensing and Dark Holes

Another WFSC approach is to measure the electric field directly at the final focal (detector) plane of the science instrument, and control that electric field with DMs. In high contrast astronomy applications, the goal is typically to minimize the electric field from the central star in a portion of the detector plane and create a dark hole in which faint astrophysical signals can be observed. The two algorithms commonly used for control are known as electric field conjugation (EFC) and stroke-minimization, and a detailed review of both, including how they are mathematically equivalent despite algorithmic differences, can be found in Groff *et al.* (2015) [23].

Focal plane wavefront sensing provides two major additional capabilities when compared to an adaptive optics approach: the correction of non-common-path aberrations and the mitigation of leftover light from telescope geometry diffraction or imperfect coronagraphs. However, focal plane WFSC is limited in speed by photon noise (as speckles within the dark hole would be very faint) as well as the computational demands of the algorithms. Focal plane WFSC therefore cannot combat quickly varying sources of wavefront error (such as atmospheric turbulence), and are thus usually used to correct for static and quasistatic effects. Consequently, AO features prominently in ground-based telescopes, while focal plane sensing and control are typically used in space-based designs, although elements of both are used in conjunction both on the ground and in space instrument designs [13].

1.3 Post-processing

Although coronagraphy and wavefront control removes much of the unwanted on-axis light, the observations will still contain residual errors. Many post-processing strategies exist to further mitigate the effects of residuals and increase sensitivity to real astrophysical signals.

1.3.1 Self-calibration, Robust Observables, and Kernel Methods

Before discussing post-processing approaches for coronagraphic systems with dark-hole enabling wavefront control, we first review the notion of self-calibration. Self-calibration in

astronomy originates from closure phases in radio interferometry, in which phases measured from three different baselines can be added together such that the phase errors of each subaperture cancel out [24]. The resulting measurements, the closure phases, are thus self-calibrating with respect to phase errors at the subapertures. Closure phases can be applied to other nonredundant arrays (where no baselines are identical) as well, including optical interferometers or full aperture telescopes with a nonredundant mask of subapertures applied over the pupil.

The concept of self-calibration can be extended to full-aperture telescope imagers if the wavefront error is small, such as with the use of an AO system. In the limit of low wavefront aberration ($\ll 1$ rad), the relationship between subaperture phases and visibility phases (the phases of the 2D Fourier transform of the image) is linear to first order even with redundant baselines, allowing for the treatment of unobscured, full-aperture telescope images with kernel phases [25]. Kernel phases are the full-aperture generalization of closure phases, and leverage this linearized relationship between subaperture phases and visibility phases. The image can be projected into the null space, or kernel, of the response matrix describing that relationship [25]. The resulting data are consequently immune to wavefront aberrations within that linearized regime, to first order. The same approach can also be applied to pupil plane amplitudes and visibility amplitudes, leading to kernel amplitudes analogous to kernel phases [10]. In general, self-calibration takes advantage of having more data than ways that errors can manifest, allowing for the extraction of data that are free from aberrations (see Figure 1-3 for such a pictorial representation of kernel phases).

The relationship between the pupil phases and amplitudes and the visibility phases and amplitudes is only approximately linear in the low wavefront error regime, as they will still contain higher order error terms. Nevertheless, they are an example of finding in the data robust observables, which are less sensitive to aberrations and can thus exhibit a higher signal-to-noise ratio (SNR).

However, these kernel methods cannot be applied to coronagraphic data, as the coronagraphic instrument is not translation invariant and breaks the linear relationship between the pupil plane and the visibility. Thus, when using this technique, we cannot also take advantage of the photon noise suppression that coronagraphs (along with coronagraphic




Figure 1-3: Pictorial representation of kernel phases. The operator \mathbf{A}_ϕ injects phase aberrations ϕ into a subspace of baseline phases Φ . There is another subspace of baseline phases, $\mathbf{K}_\phi \Phi_0$, which are not affected by these phase aberrations, and we call these kernel phases. [10]

focal-plane WFSC systems) provide.

1.3.2 Coronagraphic Post-processing

Many approaches exist to reduce the impact of wavefront error residuals on coronagraphic data. One of the most successful families of methods involves using reference observations to determining the structure of the on-axis PSF in the presence of aberrations in order to subtract it from the data. This family of methods is collectively known as reference star differential imaging (RSDI), and relies on data for a star similar to the target star but with a known absence of any astrophysical companions, either through observational strategy or a built-up library (see Ruane *et al.* for an overview [26]). Within this family, there are several approaches to constructing the PSF to be subtracted, including just scaling the reference PSF [27], principle component analysis (PCA) based approaches [28] [29] [30], and non-negative matrix factorization (NMF) [31].

Another family of post-processing methods uses instrumental data from the target star

observation itself. For example, knowledge of the wavefront error residuals that are left over from the AO system but still measured by a wavefront sensor can also be used in post-processing [32]. Additionally, focal plane electric field sensing results in a series of estimations of the coherent electric field from the central star, from which the incoherent intensity (such as that of an astrophysical companion) not captured by the electric field measurements also be estimated in post-processing [33].

1.3.3 Compatibility of Methods

The methods described in Section 1.3.2 are not limited to only coronagraphs and are generally applicable to different types of data. For example, additional post-processing of kernel-phase data can be done using reference star calibrators, similarly to the RSDI family of methods for coronagraphic post-processing [34] [35]. The kernel methods are also in theory compatible with approaches using wavefront sensing data propagated through the appropriate model. Robust observables thus complement these other post-processing approaches by providing them with an instrumentally determined robust subspace of data to work in. Any coronagraphic robust observables should also be compatible with post-processing steps such as RSDI or the incorporation of wavefront sensing data.

1.4 Thesis Contributions

This thesis explores one approach to finding robust observables for coronagraphic instruments. Chapter 2 reviews the propagation of signals and wavefront error through a coronagraphic instrument and motivates the problem of finding a projection onto robust observables. In Chapter 3, a CLC is explored in simulation, the algorithm to find a projection onto robust observables is applied, the projection is used to post-process synthetic data, and its effects on the detection of hypothetical astrophysical companions are characterized. Chapter 4 repeats the same analysis with a simulation of the CLC on the HiCAT testbed, with a key difference in the introduction of a dark hole created using the stroke-minimization algorithm. Chapter 5 discusses some additional details about our work, including connections to other related works in the modeling and analysis of high contrast instruments and signal processing

as well as a discussion of the full-dimensional representation of the quadratic coronagraph response. Chapter 6 concludes with an outlook on potential future work on coronagraphy and robust observables.

Chapter 2

Signals and Wavefront Error in Coronagraphs

In order to motivate our approach to finding robust observables for coronagraphic data, we first overview how signals and wavefront error propagate through a coronagraphic instrument into the final detected intensities. In the following equations, electric fields at the input pupil plane are assumed to have an amplitude of 1 across the entire pupil, and the brightness differences between sources are captured through scaling by the flux ratio c in the final image.

A coronagraph can be modeled as a linear transformation of the electric field, described by the matrix \mathcal{C} . \mathcal{C} is a constant for the coronagraph considered in this paper, and for most coronagraph designs in general, transforming the electric field at the pupil plane into the electric field at the image plane. If E_{0*} is the electric field of the central source (star) at the pupil plane in the absence of aberrations, and ΔE_* a small perturbation to that electric field representing wavefront aberration, then the electric field at the image plane is

$$E_{im} = \mathcal{C}E_* = \mathcal{C}E_{0*} + \mathcal{C}\Delta E_* \quad (2.1)$$

The intensity measured is the element-wise norm squared of the image plane electric field (here, \bar{x} indicates the element-wise complex conjugate of x).

$$I_* = |E_{im}|^2 \quad (2.2)$$

$$I_* = |\mathcal{C}E_{0*}|^2 + \overline{(\mathcal{C}E_{0*})}\mathcal{C}\Delta E_* + \overline{(\mathcal{C}\Delta E_*)}\mathcal{C}E_{0*} + |\mathcal{C}\Delta E_*|^2 \quad (2.3)$$

$$I_* = |\mathcal{C}E_{0*}|^2 + 2\operatorname{Re}\left\{\overline{(\mathcal{C}E_{0*})}\mathcal{C}\Delta E_*\right\} + |\mathcal{C}\Delta E_*|^2 \quad (2.4)$$

Simiarly, for an off-axis source with flux ratio c and pupil plane wavefront E_p , its image plane intensity can be expressed as

$$I_p = c(|\mathcal{C}E_{0p}|^2 + 2\operatorname{Re}\left\{\overline{(\mathcal{C}E_{0p})}\mathcal{C}\Delta E_p\right\} + |\mathcal{C}\Delta E_p|^2) \quad (2.5)$$

The total intensity of the star and off-axis source is:

$$I_{tot} = |\mathcal{C}E_{0*}|^2 + 2\operatorname{Re}\left\{\overline{(\mathcal{C}E_{0*})}\mathcal{C}\Delta E_*\right\} + |\mathcal{C}\Delta E_*|^2 + c(|\mathcal{C}E_{0p}|^2 + 2\operatorname{Re}\left\{\overline{(\mathcal{C}E_{0p})}\mathcal{C}\Delta E_p\right\} + |\mathcal{C}\Delta E_p|^2) \quad (2.6)$$

Since the effect of aberrations on the off-axis source is smaller by a factor of c , the dominant errors are expected to come from the effect of aberrations on the central source. Thus, the total image is approximately

$$I_{tot} \approx |\mathcal{C}E_{0*}|^2 + 2\operatorname{Re}\left\{\overline{(\mathcal{C}E_{0*})}\mathcal{C}\Delta E_*\right\} + |\mathcal{C}\Delta E_*|^2 + c|\mathcal{C}E_{0p}|^2 \quad (2.7)$$

The first $|\mathcal{C}E_{0*}|^2$ term is the residual starlight not blocked by the coronagraph in the case of no aberrations. The second $2\operatorname{Re}\left\{\overline{(\mathcal{C}E_{0*})}\mathcal{C}\Delta E_*\right\}$ term is linear in the wavefront aberration and corresponds to the interference of the aberration with the underlying residual starlight from the coronagraph, analogous to speckle pinning [36]. The third term $|\mathcal{C}\Delta E_*|^2$ is the quadratic term, corresponding to the norm squared of the wavefront error at the focal plane. The last term $c|\mathcal{C}E_{0p}|^2$ is the nominal off-axis signal, which is the signal of interest.

Whether the effect of wavefront errors at some location in the image plane are dominated by the linear term or the quadratic term depends the attenuation of starlight by the

coronagraph and the level of wavefront error at that location. In the case that wavefront aberrations have much larger complex amplitudes than the residual starlight after the coronagraph with no aberrations, the quadratic term will be dominant. For this work, we will work in the regime where the quadratic term $|\mathcal{C}\Delta E_*|^2$ is dominant, which is the more likely and realistic scenario [37]. This is particularly the case with modern perfect coronagraph designs, where the coronagraph is designed so that the residual starlight in the case of no aberrations is zero. In these quadratically-dominated situations, the image plane intensity is approximately

$$I_{tot} \approx |\mathcal{C}E_{0*}|^2 + |\mathcal{C}\Delta E_*|^2 + c|\mathcal{C}E_{0p}|^2 \quad (2.8)$$

We approximate the quadratic error term with the response matrix A_q .

$$I_{tot} \approx |\mathcal{C}E_{0*}|^2 + A_q|\Delta E_*|^2 + c|\mathcal{C}E_{0p}|^2 \quad (2.9)$$

Since $|\mathcal{C}\Delta E_*|^2$ is not actually linear in $|\Delta E_*|^2$, A_q is simply an approximation as determined in one basis. Thus, calculating A_q in a different basis will capture the effects of different combinations of cross-terms (see Section 5.4 for discussion about the full dimensional representation). A key part of this investigation is testing how well a matrix A_q can capture the quadratic instrumental response despite neglecting the crossterms of the basis vectors over which it is calculated.

Given the approximate instrument response, A_q , we can explore if there are linear combinations of the measurements that will yield observables that are robust to wavefront error. If a left-hand null space of A_q exists, we can choose K_q to be that left-hand null space, such that $K_q A_q = 0$. If A_q perfectly captured the quadratic response of the instrument, left-multiplying the data by K_q would give observables O_q independent of the quadratic error.

$$O_q = K_q I_{tot} \approx K_q (|\mathcal{C}E_{0*}|^2 + c|\mathcal{C}E_{0p}|^2) \quad (2.10)$$

In reality, A_q is only an approximation, so even if it has a null space K_q , applying K_q to the quadratic error as $K_q|\mathcal{C}\Delta E_*|^2$ will not actually yield zero because of all the cross-terms

that are unaccounted for by A_q . However, we find that our linearized A_q does not even have a clear null space, as discussed in Section 3.4.

Consequently, the rest of our work thus mainly explores how well an approximation of A_q can help identify the most dominant error modes as predicted by the instrument response and project them out of the data. The remaining observables, although not identically self-calibrating, can nevertheless exhibit a certain amount of robustness to input wavefront error. We study whether extracting these robust observables increases the the signal to noise ratio of the data, which, for example, could allow for a binary signal detection with a deeper flux ratio than using the raw intensity data.

Chapter 3

Classical Lyot Coronagraph (Simulation)

3.1 Optical Model

The Lyot coronagraph model under consideration is an example included in the PROPER optical propagation library [1]. The optical system is as follows (see Figure 3-1): a single wavelength of $\lambda = 585$ nm is used, and the pupil is a circular with a 2.4 m diameter, resulting in the natural optical unit of distance in the detector plane of λ/D to be 2.4375×10^{-7} radians. The wavefront after the pupil is focused to a point, where a circular central occulter with a $4\lambda/D$ radius is applied. The wavefront is then collimated into another 2.4 m beam before a circular Lyot stop with 2.02 m diameter is applied. The resulting beam is then focused onto the detector plane where the intensity measurements are made.

The corresponding coronagraph operator \mathcal{C} is $\mathcal{C} = F\Theta F^{-1}\Omega F$, where F is the 2D Fourier transform (and F^{-1} its inverse), Ω the occulter mask, and Θ the Lyot stop mask. Although PROPER is consulted for the coronagraph design and for mask generation in order to build the matrix \mathcal{C} , for the rest of this study, the matrix \mathcal{C} is used for signal propagation instead of PROPER. The relevant quantities related to the representation and sampling of the system are given in Table 3.1.

Table 3.1: Sampling quantities for the classical Lyot coronagraph model

Quantity	Description	Calculation	Value
N_{grid}	Number of samples across the grid	—	100
$N_{samples}$	Number of total samples	N_{grid}^2	10,000
f_{beam}	Fraction of the grid the beam occupies	—	0.3
dx_{pup}	Pupil sampling (m/sample)	$D/(f_{beam}N_{grid})$	0.08
dx_{det}	Detector plane sampling $((\lambda/D)/\text{sample})$	f_{beam}	0.3




Figure 3-1: Optical model for the Lyot coronagraph, included as an example in the PROPER optical propagation library [1]. The coronagraph operates at 585 nm. The occulter is a circular obscuration with radius $0.561 \mu\text{m}$, and the Lyot stop is a circular aperture with radius 2.02 m. The Fraunhofer regime is assumed, so the model consists of only 2D Fourier transforms and masks, and resizing and propagation distances are not considered.

3.2 Coronagraph Properties

Several properties of the coronagraph are useful to quantify in order to guide some of the choices and interpretation of results in this analysis.

First, Figure 3-2 shows the coronagram of a central star in the absence of any wavefront error, corresponding to the $|\mathcal{C}\Delta E_{0*}|^2$ term. This is the intensity contribution that is known to come from the central star, and is thus subtracted from the data before performing any

post-processing or analysis to search for off-axis sources. Figure 3-3 shows the radial profile of the magnitude of the electric field distribution at the detector plane. The classical Lyot coronagraph is decidedly not a perfect coronagraph, and the magnitude of the starlight's electric field sets the limit for which wavefront errors transition from being predominantly linear to predominantly quadratic.




Figure 3-2: The nominal coronagraph of just a central star in \log_{10} scale of arbitrary intensity units.

We can propagate wavefront errors from the pupil plane to the detector plane to quantify their electric field magnitude at the detector plane. For example, Figure 3-4 shows how a sinusoidal phase pattern at the pupil plane, depending on its amplitude and spatial frequency, can manifest as a predominantly linear speckle (Figure 3-4b), a predominantly quadratic speckle (Figure 3-4c), or in between where both terms have comparable contributions (Figure 3-4d). This analysis supports the choice to focus on the quadratic term, even for this




Figure 3-3: A radial slice of the nominal coronagraph residuals on a central star.

imperfect, un-optimized classical Lyot coronagraph. The predominantly linear speckle and even the in-between speckle shown are extremely contrived situations where unrealistically low aberration amplitudes were chosen specifically based on the residual light from the coronagraph, and most high-order wavefront errors of realistic amplitudes will map to the predominantly quadratic regime. For example, the in-between speckle shown is from a phase sinusoid that is 3 nm in amplitude, and it would already be well in the quadratic regime at further spatial separations. The wavefront error at those spatial frequencies is expected to be of higher amplitude than this example, and thus, predominantly quadratic.

Another important property to quantify is the inner working angle (IWA) of the coronagraph, which is defined as the separation for which the throughput is 50%, or where 50% of the light of a point source at that separation is transmitted. To quantify this property,

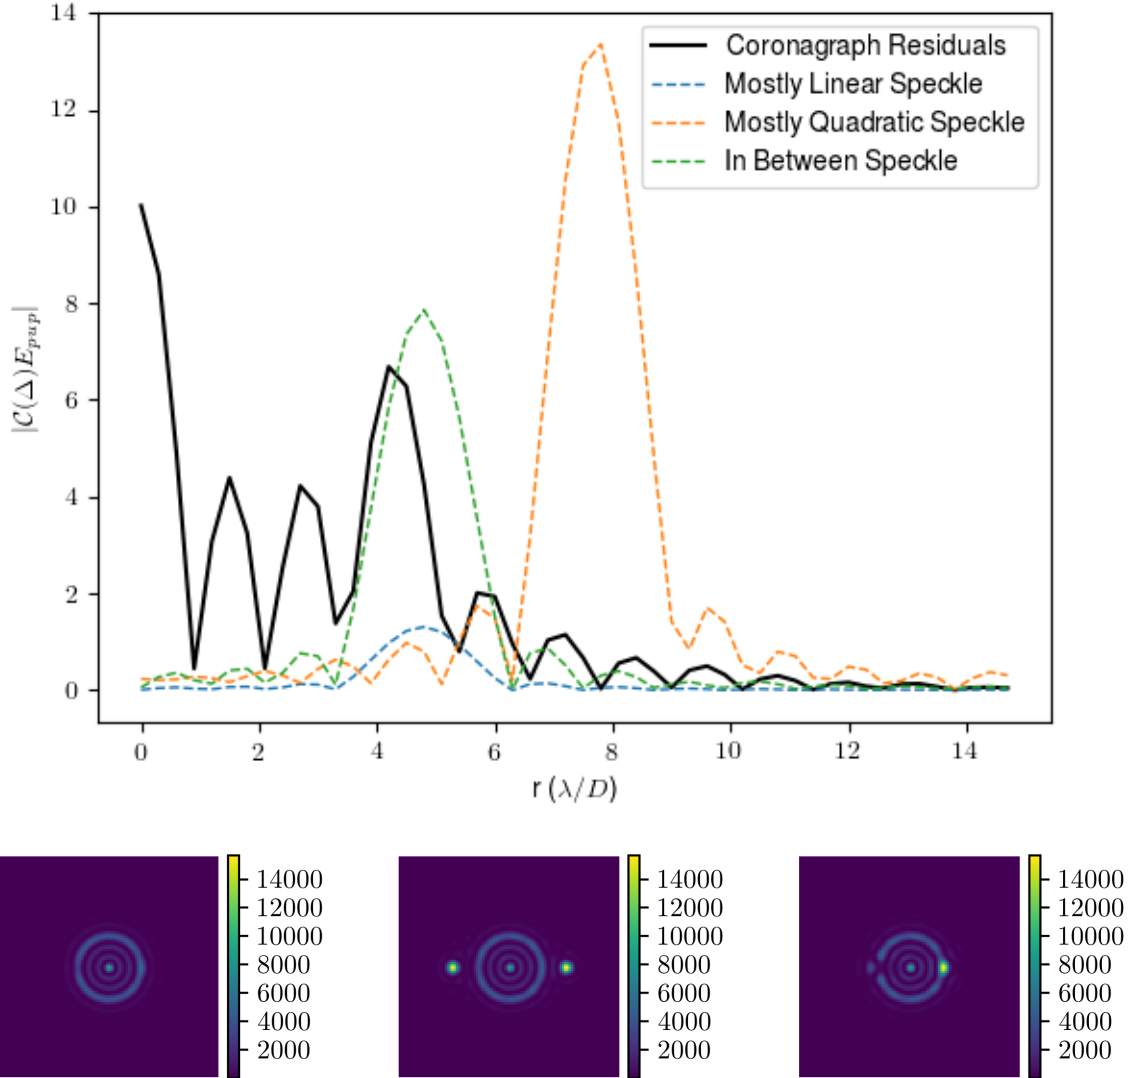


Figure 3-4: a) The radial slices of a mostly linear, a mostly quadratic, and an in-between sinusoidal speckle in the detector plane plotted against the coronagraph residuals. b) The intensity with a predominantly linear speckle, which mostly interferes with the nominal star PSF. c) The intensity with a predominantly quadratic speckle that is mostly the square-amplitude of the wavefront error itself. d) The intensity with a speckle that is neither dominantly linear nor quadratic, but somewhere in-between.

the light of point sources at various separations is propagated, and the total intensity in the detector plane is summed up. This total intensity is then divided by the summed up intensity of a point source very far away from central star (the outer edge of the range of

separations propagated, $10\lambda/D$ in this case), and the resulting throughput curve is shown in Figure 3-5. From this throughput curve, the IWA of the coronagraph is identified to be $5\lambda/D$. Since the signal of any off-axis source will be strongly attenuated by the coronagraph around and within this IWA, our analysis will focus on point sources that are sufficiently outside of it.




Figure 3-5: The throughput of an off-axis point source as a function of its spatial separation, indicated by the solid line. The inner working angle is the separation at which the throughput is 50% (indicated by the dashed line), which in this case is $5\lambda/D$.

3.3 Building Response Matrices

This section details the numerical linearization of the quadratic instrument response. Table 3.2 summarizes the dimensions of the relevant quantities in this analysis. Note that some

of the quantities in the table are generic. E_{pup} encompasses any vector of electric fields as represented in the pupil plane, such as E_{0*} or ΔE_* . However, its dimension will depend on N_{basis} , the number of basis elements used to represent it. For full representation of electric fields for propagation through the optical model, N_{basis} is the total number of samples, $N_{samples}$. For the linearization of the response matrix, however, N_{basis} will be the number of basis vectors used in the calculation of the response matrix. E_{det} also encompasses any electric field propagated to the detector plane, and in this analysis, always has a dimension of $N_{samples}$. N_k is either the dimension of the kernel, if one exists, or, it is a choice that is made during the determination of the projection matrix as discussed in Section 3.4.

Table 3.2: Quantities and dimensions for analysis of the classical Lyot coronagraph

Quantity	Description	Dimension
\mathcal{C}	Coronagraph propagation matrix	$N_{samples} \times N_{samples}$
E_{pup}	Vector of electric field in pupil plane	$N_{basis} \times 1$
E_{det}	Vector of detector plane E-field	$N_{samples} \times 1$
I	Vector of intensity measurements	$N_{samples} \times 1$
A	Instrument response matrix	$N_{samples} \times N_{basis}$
U	Left singular matrix of A	$N_{samples} \times N_{samples}$
S	Singular values of A	$N_{basis} \times 1$
V	Right singular matrix of A	$N_{basis} \times N_{basis}$
K	Projection matrix	$N_k \times N_{samples}$
O	Vector of observables	$N_k \times 1$

The general method of this numerical linearization is as follows. Aberrations corresponding to unit vectors of a chosen basis are induced in a wavefront that is propagated through the optical system, and their effect on the corresponding error term is registered into the response matrix. Since we are operating in the regime where the quadratic term is dominant, we register the quantity $|\mathcal{C}\Delta E_*|^2$ into the response matrix A (the subscript q is dropped, since only the quadratic term is relevant from now on). Because of the quadratic response, this linearization approach will fail to capture the effects of cross-terms of whatever basis is chosen to represent the wavefront errors. Different bases will capture the effects of different combinations of wavefront error, so the choice of basis will have an impact on the results. We choose to represent the pupil wavefront in the Zernike basis for the construction of our response matrix.

The response matrix is linearized with the wavefront aberrations represented in the Zernike basis, where ΔE_{*Z_n} is the aberration induced by the n^{th} Noll ordered Zernike polynomial [38], and N is the total number of polynomials chosen to construct the response matrix:

$$\Delta E_{*Z} = \begin{pmatrix} \Delta E_{Z_1} \\ \dots \\ \Delta E_{Z_N} \end{pmatrix} \quad (3.1)$$

This response matrix is linearized through applying 1 nm RMS (an arbitrary number small enough such that the phase is largely the imaginary component of the electric field) of each Zernike mode up to the 276^{th} mode, for an N_{basis} of 276. The effect of each wavefront error mode on the quantity $|\mathcal{C}\Delta E_*|^2$ is registered into the $10,000 \times 276$ matrix A_Z . The only reason that the number of Zernikes used is 276 is that the PROPER package that is used to generate the wavefronts with Zernike polynomials stops working beyond the 276th Zernike for a currently unknown reason. In future work, this limitation can be overcome by simply using a different method to generate the wavefronts. For now, this analysis focuses on the range of spatial separations that are well-spanned by the first 276 Zernikes. This range can be roughly determined from a dimensional argument: the dimension of the Zernikes basis corresponds to the dimension of the resolution elements they span. The circular area spanned by the 276 Zernikes is approximately $276 (\lambda/D)^2$ resolution elements, resulting in a spanned radius of $\sqrt{276/\pi} \approx 9\lambda/D$. Consequently, this analysis focuses on signals within $9\lambda/D$.

3.4 Singular Value Decomposition, Singular Modes, and the Projection K

A singular value decomposition (SVD) of $A_Z = USV^T$ (see Table 3.2 for matrix definitions) is performed, revealing its singular modes and corresponding singular values. The singular

values are plotted in Figure 3-6, while the first 10 singular modes as represented in the detector plane intensity basis are shown in Figure 3-7.




Figure 3-6: The singular values of the transfer matrix A . We choose the 102nd through 10,000th singular modes to make up our projection matrix K , which will project out the first 101 modes from our data and result in $N_k = 9899$.

The left-hand null space, or kernel, of each response matrix is comprised of the vectors in the respective U that correspond to singular values of 0. Figure 3-6 shows that there is no well defined cutoff in the singular values of A_Z , and thus no kernel. Our simulations show that response matrices that capture the instrument response well in general do not have kernels.

Although a kernel does not exist for A_Z , its most dominant error modes can still be projected out. There is a choice of N_k , or how many modes to keep in the projection, but we heuristically find that the choice of how many modes to project out does not really affect the results as long as it is larger than a few modes, but not the majority of N_{basis} . In this analysis,




Figure 3-7: The first 10 singular modes of A_Z as represented in the the detector plane intensity basis.

we choose to project out the first 101 modes, resulting in $N_k = N_{samples} - 101 = 9899$. This projection is onto the modes that correspond to the green shaded region and beyond of singular values in Figure 3-6. The following section discusses applying this projection to the data and its effect on the detection of companion point sources.

3.5 Synthetic Data Analysis

3.5.1 Synthetic Data Generation

In order to characterize the performance of our projection matrices, synthetic datasets first need to be generated. The three parameters that characterize the astrophysical signal of a

binary companion are its flux ratio, separation (in λ/D units), and position angle. Since the classical Lyot coronagraph is circularly symmetric, the position angle is set to zero. Ultimately, data over a range of flux ratio values and separations are needed, but for this first example dataset, the flux ratio is chosen to be 4×10^{-4} with a separation of $7.5\lambda/D$.

The dataset is chosen to have 20 frames, corresponding to 20 snapshots at different points in time. These snapshots are assumed to be instantaneous, which is valid if the exposure times are much shorter than the timescale over which the phase errors change. For each frame, the PSD given in Equation 3.2 is used to generate a phase error map [1].

$$PSD_{2D}(k) = \frac{a}{[1 + (k/b)^2]^{\frac{c+1}{2}}} \quad (3.2)$$

In Equation 3.2, k is the spatial frequency. We choose a (the low spatial frequency power) to be $a = 4.935 \times 10^{-18} \text{ m}^4$, b (the correlation length parameter) to be $b = 212.26 \text{ cycles/m}$, and c (the high frequency falloff power law exponent) to be $c = 7.8$.

The resulting phase error map is smoothed by a two-dimensional Gaussian with a one pixel standard deviation so that the errors are Nyquist sampled, and then renormalized to the original peak error. Both the central star wavefront and the off-axis companion wavefront are propagated through the model with this error map applied. Their intensities (which are mutually incoherent) are then summed together, with the companion intensity map multiplied by its flux ratio. These 20 snapshots form the dataset to be analyzed. As a control case, which will be important for detection tests, the same phase error maps are used to generate a dataset in which there is no companion.

3.5.2 Processing Synthetic Data

At this point, the model coronagraph of the central star in the absence of wavefront errors can be subtracted off from each frame, as it is a known contribution to the intensity that can be removed. For the baseline analysis using the raw intensity, the resulting twenty model-subtracted frames are simply averaged to obtain the final measurement ΔI , while the standard deviation of the twenty frames is used as the measurement uncertainty. Figure 3-8 shows this process and the resulting data frames for the dataset with the injected companion.




Figure 3-8: Top left: A single aberrated snapshot (in linear scale) with a companion with a flux ratio of 4×10^{-4} at $7.5\lambda/D$. Top right: The data frame with the model coronagraph subtracted. Bottom left: The mean of 20 frames of data, with the companion "location" (visible by eye) marked with the red box. Bottom right: The standard deviation of 20 frames of data.

Then, each frame of the model-subtracted intensity data (in vector shape as ΔI) is first left-multiplied by the projection matrix in order to obtain the observables $O = K\Delta I$ for that frame. Finally, the mean of the observables across the 20 frames is taken to be the final measurement, and the standard deviation the measurement uncertainty.

The characterization of the performance of this post-processing approach ultimately depends on the goal of the data analysis. This work analyzes the effect of the post-processing on the detection of the binary point source. A detection test is applied to the resulting measurements in order to characterize the detectability of a companion from those measurements.

3.5.3 Detection Tests

Detections are typically made using hypothesis testing. Discussion of hypothesis testing and test statistics as applied to binary companion detection can be found in the literature [39] [40] [41], and the key elements are summarized here.

The null hypothesis, \mathcal{H}_0 , is the case that the data is entirely noise, while \mathcal{H}_1 is the hypothesis that the data contains both noise and the signal of the binary companion. A test statistic T is calculated from the data and compared to a threshold ξ . \mathcal{H}_1 (a detection) is claimed if $T \geq \xi$, and \mathcal{H}_0 (lack of a detection) is claimed otherwise.

The optimal test statistic for this situation is unknown, but we motivate and examine the results of using two different test statistics for detection tests.

Delta Reduced Chi-Squared

The first test statistic we use is $\Delta\chi_r^2$, or the difference in the reduced χ^2 of the data assuming it contains only noise, and the reduced χ^2 of the data assuming it contains noise and data. The formula for calculating this test statistic from the data is given by Equation 3.3.

$$\Delta\chi_r^2 = \frac{y^2}{\nu\sigma^2} - \frac{(y-x)^2}{\nu\sigma^2} \quad (3.3)$$

In this formula, y is the data, x is the model of the binary signal, σ is the estimated uncertainty, and ν is the degrees of freedom, or the number of data elements minus the number of free parameters, for which there are three for a binary (flux ratio, separation, and position angle).

If the noise is uncorrelated and normally distributed, then the exponential of this test statistic is also the likelihood ratio of the existence of the companion to the non-existence of the companion. In this situation, this test statistic would be equivalent to the Neyman-Pearson test, which is the most powerful statistical test [42]. However, the noise in our data is not uncorrelated nor is it normally distributed, so using $\Delta\chi_r^2$ as the test statistic is not equivalent to the Neyman-Pearson test, and thus not necessarily the most powerful statistical test.

Delta Reduced Error Squared

A second test statistic that is explored is similar to the $\Delta\chi_r^2$ statistic, except without the inclusion of the uncertainties σ . The formula for the $\Delta\epsilon^2$, or the difference in the squared error (per degree of freedom) of assuming noise-only data and noise-and-signal data is given by Equation 3.4.

$$\Delta\epsilon^2 = \frac{y^2}{\nu} - \frac{(y-x)^2}{\nu} \quad (3.4)$$

The inclusion of this test statistic is motivated by the nongaussianity and correlation of the errors in the data, as well as the consequent unrealistic estimates of the uncertainties obtained by calculating the standard deviation (see Section 3.6 for a discussion). These factors become even more prominent when analyzing data that is also correlated in time, as in Section 3.5.6. Thus, using $\Delta\epsilon^2$ reflects the extreme of complete agnosticism about the uncertainties and correlations of the measurements.

3.5.4 Receiver Operating Characteristic Curves

We perform detection tests on our synthetic data using the two test statistics $\Delta\chi_r^2$ and $\Delta\epsilon^2$. Continuing with the case of a companion at a flux ratio of 4×10^{-4} at $7.5\lambda/D$, we generate 1000 synthetic datasets of 20 frames each, along with the accompanying datasets in which no companion is injected. Each dataset is processed as described in Section 3.5.2, and the test statistics are calculated for each dataset. Histograms of both of the resulting test statistics are shown in Figure 3-9.

There are four possible outcomes when comparing the test statistic calculated from a dataset to a value of the test statistic that is set as the detection threshold. The first is a true positive, in that there is a companion in the data that is detected. The fraction of real companions that are detected is the true positive rate (TPR). A second possible outcome is that there is a real companion that is not detected, which occurs at a rate of $1 - \text{TPR}$. A third outcome is that there is no companion in the data, but the detection test incorrectly claims a detection. The rate at which this occurs is the false positive rate (FPR). The fourth and last outcome is that there is no companion, and a detection is also correctly not claimed,




Figure 3-9: Left: The histograms over 1000 datasets of the test statistic $\Delta\chi_r^2$ for raw intensity and observables, both with a companion with flux ratio of 4×10^{-4} at $7.5\lambda/D$ and without it. Right: The histograms over 1000 datasets of the test statistic $\Delta\epsilon^2$ for raw intensity and observables, with and without the same companion.

occurring at a rate of $1 - \text{FPR}$.

Choosing a threshold for the test statistic is a balancing act between the TPR and FPR. As the threshold is increased, detecting real companions becomes more likely, but false detections also become more likely. This dependency can be characterized by going through every test statistic value in the histograms shown above and calculating the TPR and FPR if that value were chosen to be the detection threshold. Plotting the TPR as a function of the FPR results in a receiver operating characteristic (ROC) curve (whose name derives from radar receiver applications), as shown in Figure 3-10 for the $\Delta\chi_r^2$ statistic and the $\Delta\epsilon^2$ statistic.

This process is repeated for companions at a range of flux ratios (to a resolution of 10^{-4}) and spatial separations (to a resolution of $0.5\lambda/D$), resulting in ROC curves at those flux ratios and separations. These ROC curves are then used to determine detection limits.

3.5.5 Detection Limits

The flux ratio detection limit is dependent on the desired TPR and FPR. The choice of TPR and FPR is flexible and depends on the observational context, but for illustrative purposes, we choose to have $\text{FPR} \leq 0.01$ and $\text{TPR} \geq 0.90$. From the ROC curves, we identify the




Figure 3-10: Left: The receiver operating characteristic curve, or true positive rate as a function of false positive rate, for the $\Delta\chi_r^2$ statistic for a companion with a flux ratio of 4×10^{-4} at $7.5\lambda/D$. Right: The receiver operating characteristic curve for the $\Delta\epsilon^2$ statistic for the same companion

deepest flux ratio for which a companion at a certain spatial separation can be detected the specified TPR and FPR. This results in the detection curves for the $\Delta\chi_r^2$ statistic in Figure 3-11a and for the $\Delta\epsilon^2$ statistic in Figure 3-11b. With both of the test statistics used, detections using the robust observables outperform detections using the raw intensities.




Figure 3-11: Left: The $\text{FPR} = 0.01$, $\text{TPR} = 0.90$ detection limits (at the 10^{-4} resolution) as a function of separation for raw intensities and the observables, using the $\Delta\chi_r^2$ test statistic. Right: The $\text{FPR} = 0.01$, $\text{TPR} = 0.90$ detection limit (at the 10^{-4} resolution) as a function of separation for raw intensities and the observables, using the $\Delta\epsilon^2$ test statistic.

3.5.6 Time-correlated Wavefront Errors

In Section 3.5.1, the phase errors introduced were uncorrelated in time, as for each frame, an entirely new phase error map was drawn from the PSD. We now examine cases where the phase errors are also correlated in time. This corresponds to the situation of more slowly varying perturbations that occur over the timescale of many frames. In order to model and generate time-correlated phase errors, we introduce a parameter $0 \leq \alpha \leq 1$, corresponding to the fraction of the error energy that is new in each snapshot. Each new portion of the phase map, ϕ , is generated using the same method as described in Section 3.5.1. The phase error of a subsequent frame Φ_{k+1} given the current frame's phase error Φ_k and the α parameter is given by Equation 3.5.

$$\Phi_{k+1} = \sqrt{\alpha}\phi + \sqrt{1-\alpha}\Phi_k \quad (3.5)$$

This method of generating phase errors maintains the expected value of the error energy across all frames, as well as consistency with the expected value of error energy of all the frames in the datasets without time-correlation. As limiting cases, $\alpha = 0$ corresponds to an entirely static wavefront error while $\alpha = 1$ corresponds to wavefront error that is uncorrelated in time.

Using α values of 0.1 and 0.5 results in 20-frame mean intensity measurements and standard deviations as shown in Figure 3-12 (no companion was injected into these datasets).

Repeating the detection test analysis on these time-correlated datasets results in the $\text{FPR} = 0.01$ and $\text{TPR} = 0.90$ detection curves given in Figure 3-13.

3.6 Summary of Results

We discuss key results from the detection limit curves. First, in the case of non-time-correlated wavefront error, using the observables increases the detectability of a binary companion using both of the test statistics explored by up to 50%. That the $\Delta\epsilon^2$ statistic seems to work better for both the observables data and the raw intensity data indicates that neither falls under the Gaussian hypothesis, and that $\Delta\chi_r^2$ is not the likelihood ratio for the existence




Figure 3-12: Left: Mean of 20 frames with energy correlation parameter $\alpha = 0.1$. Right: Mean of 20 frames with $\alpha = 0.5$.

of the companion. The non-Gaussianity can be confirmed by examining the histograms of the residual errors for the intensities and observables (see Figure 3-14).

For the time-correlated wavefront errors, the results are different for each of the test statistics and also vary with the degree of correlation. For both correlation parameters of $\alpha = 0.1$ and $\alpha = 0.5$, raw intensities perform better using $\Delta\chi_r^2$ while observables perform better using $\Delta\epsilon^2$. Although these comparisons between different test statistics may not be particularly meaningful, in the less correlated case of $\alpha = 0.5$, observables using $\Delta\epsilon^2$ still provide a very marginal improvement over raw intensities using $\Delta\chi_r^2$. In the more correlated case of $\alpha = 0.1$, the opposite case is true. It is likely that neither of the test statistics explored here are the optimal statistic for companion detection under these noise distributions. Future work may include analyzing the behavior of test statistics under different error distributions

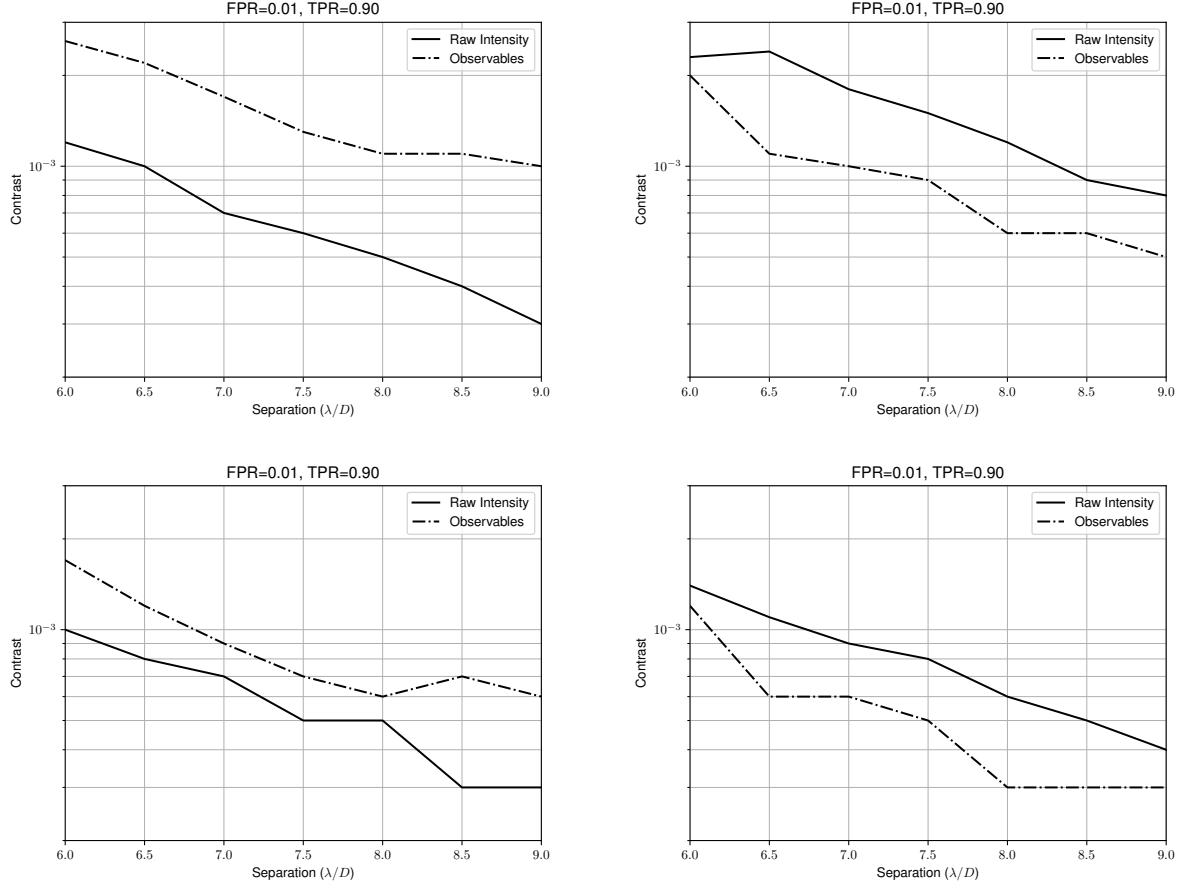


Figure 3-13: FPR = 0.01 and TPR = 0.90 detection curves. Top left: $\alpha = 0.1$ using the $\Delta\chi_r^2$ statistic. Top right: $\alpha = 0.1$ using the $\Delta\epsilon^2$ statistic. Bottom left: $\alpha = 0.5$ using the $\Delta\chi_r^2$ statistic. Bottom right: $\alpha = 0.5$ using the $\Delta\epsilon^2$ statistic.

in order to identify the most optimal test statistic. In practice, it would be more worthwhile to analyze the combination of this projection with an additional whitening post-processing step such as with projections on Karhunen-Loève eigenmodes, which would remove many of these issues arising from time-correlated data [28].

See Appendix A for additional details and discussion on a heuristic view of the model-based projection.




Figure 3-14: Histograms of the remaining errors after processing a dataset of 20 frames, uncorrelated in time, for intensities and observables.

Chapter 4

Classical Lyot Coronagraph with Dark Hole (HiCAT Simulation)

4.1 Background

The High-contrast imager for complex aperture telescopes (HiCAT) testbed is an optical testbed at the Space Telescope Science Institute (STScI), developed for testing high-contrast instrumentation techniques for telescopes with complex geometry arising from, for example, segmentation, obstruction, or support structures [43]. HiCAT can be configured with a circular pupil or an Iris-AO segmented deformable mirror to simulate a segmented telescope pupil. The coronagraph used on HiCAT can also be a CLC or an APLC that is designed for the segmented pupil. These coronagraphs are dark-hole enabled and can operate at deeper contrasts closer to the coronagraphs in modern instruments.

A Python package developed by the Makidon Optics Lab at STScI serves as the infrastructure for modeling HiCAT and for conducting experiments both in simulation and on the physical testbed. This infrastructure includes scripts to set up and run the stroke-minimization algorithm to dig a dark hole, as well as functionality for commanding the DMs, for taking both direct and coronagraphic exposures, and for retrieving the modeled final focal plane electric field in any given testbed state. We use this functionality to characterize this post-processing algorithm in the presence of a dark hole in simulation, and with additional possibility of smoothly transitioning to collecting and analyzing real testbed data for future

work.

4.2 Optical Model

For this experiment, we simulate the HiCAT testbed in the classical Lyot coronagraph configuration. The opto-mechanical layout of the HiCAT testbed is shown in Figure 4-1, and pictures of the testbed in Figure 4-2. The physical dimensions that are relevant to the optical model are summarized in Table 4.1.

Table 4.1: Physical dimensions for the HiCAT CLC

Quantity	Description	Value
λ	Wavelength (nm)	638
D	Pupil diameter (mm)	18
D_{DM}	Diameter of DMs (mm)	9.9
D_{DM1lp}	DM 1 diameter mapped to Lyot plane (mm)	22.22
z_{DM}	Distance between DMs (m)	0.3
D_{fpm}	Focal plane mask diameter (λ/D)	4.15
D_{ls}	Lyot stop diameter (mm)	15
f_{ls}	Lyot stop fraction	0.83




Figure 4-1: Full optomechanical layout of the HiCAT testbed. The testbed has the capability to transition between an APLC mode (with the Iris-AO DM simulating a hexagonally segmented pupil) and a CLC mode (with a simple circular pupil). [11].




Figure 4-2: Left: view of the entire HiCAT testbed in its (opened) enclosure. Right: detail of the central part of the testbed. [11].

This CLC is architecturally similar to the CLC from the previous chapter, but with different model parameters and the inclusion of two deformable mirrors in the collimated space before the focal plane occulter. On the physical testbed, these two DMs are both microelectromechanical systems (MEMS) devices from Boston Micromachines, each of which has 952 actuators arranged in a circle that has a diameter of 9.9 mm and is 34 actuators across. The HiCAT testbed is simulated using the Physical Optics Propagation in PYthon (POPPY) software package [2]. In the model, Fourier propagation is used for most of the instrument, but Fresnel propagation is used for the propagation around the DMs. The sizing and propagation distances for the DM section of the testbed is important and are shown in the table of physical dimensions, Table 4.1.

The testbed is simulated using a single wavelength of $\lambda = 638$ nm. The circular pupil has a diameter of $D = 18$ mm, and is first remapped to the plane of the first deformable mirror, DM 1, with a beam diameter of 8 mm. (Note that the DM is oversized by about 23% with respect to the beam, and its diameter when mapped back to the pupil/lyot stop plane is 22.22 mm.) The wavefront propagates a distance of $z_{DM} = 0.3$ m to the second deformable mirror, DM 2.

After the DMs, the wavefront is propagated to a plane conjugate to the pupil before it is focused onto the focal plane mask (FPM), a circular mask with diameter $D_{fpm} = 4.15\lambda/D$. The beam is then expanded to its original diameter of 18 mm, where a circular Lyot stop of

$D_{ls} = 15$ mm (or beam fraction of $f_{ls} = 0.83$) is applied. The beam is finally focused onto the detector where the intensity measurement is made.

The corresponding coronagraph operator \mathcal{C} is $\mathcal{C} = F\Theta F^{-1}\Omega\Lambda F$, where F is the 2D Fourier transform (and F^{-1} its inverse), Ω the occulter mask, Θ the Lyot stop mask, and Λ the operator that describes the effect of the DMs in a given state.

The relevant sampling quantities for the system are given in Table 4.2. The HiCAT simulator generally significantly oversamples for its propagation, so the sampling given only refers to the final (binned) data products that are used in our analysis. Additionally, instead of modeling wavefront error by directly modifying the electric field at the input pupil plane, we use the first DM (conjugate to the input pupil plane) to create electric field phase perturbations. This choice allows us to seamlessly transition to repeating experiments on the real testbed, which would use DM 1 to inject perturbations to the electric field. Thus, the sampling for the representation of the wavefront error is not the sampling of the wavefront that is used for propagation, but rather based on the spacing of the DM actuators.

Table 4.2: Sampling quantities for the HiCAT CLC

Quantity	Description	Value
N_{grid}	Number of pixels across detector	178
$N_{samples}$	Total number of detector pixels	31,684
N_{DH}	Number of pixels within the dark hole	2,912
dx_{det}	Detector plane sampling $((\lambda/D)/\text{pixel})$	0.363
N_{act_x}	Number of actuators across DM	34
N_{act}	Total number of DM actuators	952
n_{act_x}	Number of actuators sampling across pupil	≈ 27
dx_{DM}	DM actuator pitch ($\mu\text{m}/\text{actuator}$)	300

4.3 Dark Holes with Stroke-Minimization

Since our goal is to test our post-processing algorithm in the presence of a dark hole resulting from focal plane sensing and nulling, we first set up and run the stroke-minimization algorithm to generate the dark hole. The dark hole ranges from $6\lambda/D_{ls}$ to $11\lambda/D_{ls}$, corresponding to $7.2\lambda/D$ to $13.2\lambda/D$. Key aspects of the dark hole, including the coronagraph intensity and

electric field, as well as the DM profiles at the end of 80 iterations of stroke-minimization, are shown in Figure 4-3.

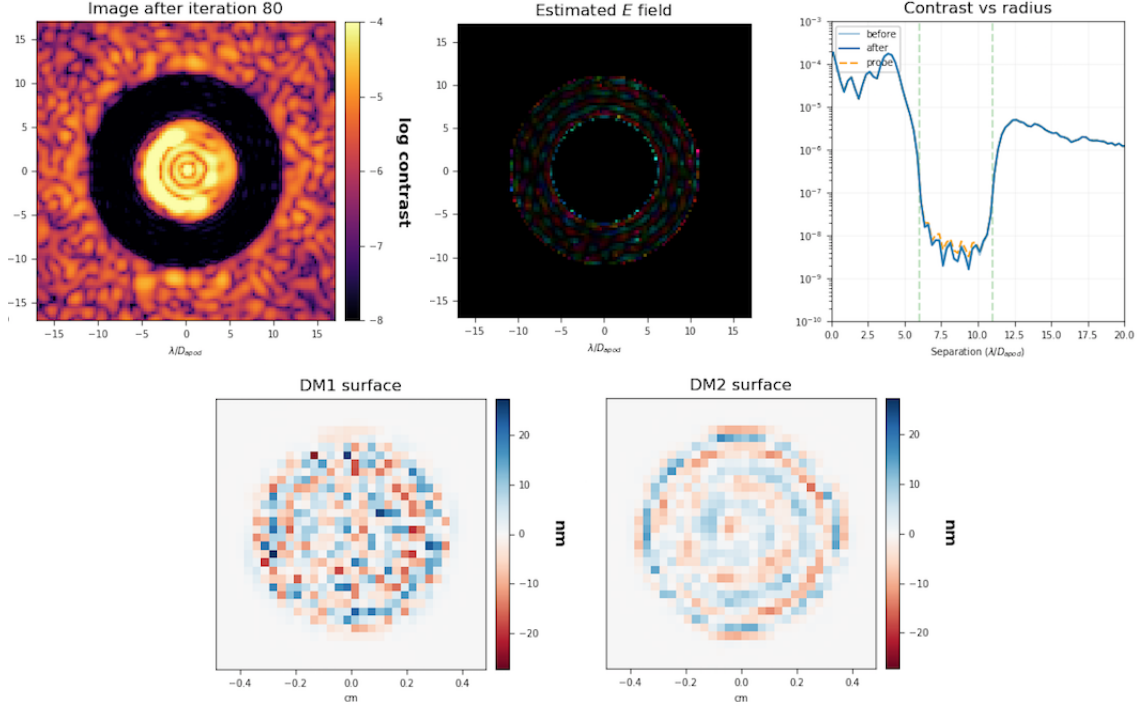


Figure 4-3: Top left: Contrast in intensity after 80 iterations of stroke minimization. Top middle: Estimated complex electric field in dark hole, where brightness indicates arbitrary amplitude units and color indicates phase as per visualization by HCIPy [3]. Top right: The radial profile in contrast before and after the 80th iteration, with the probe magnitude shown as a diagnostic. Bottom left: The surface profile of DM 1. Bottom right: The surface profile of DM 2. All two-dimensional maps are square, where the Y axis has the same units and spacing as the X axis.

The DM commands that correspond to this dark hole are saved and can be loaded in again to return the simulation to the dark hole state. This dark hole state serves as the linearization point about which our analysis takes place, and the coronagraph model thus includes within it the effects of the two DMs.

Since the simulator included in the HiCAT software package is intended to reflect the testbed state and functionality, the ability to simulate substantially off-axis sources is not included. We thus assume that the throughput of off-axis sources throughout the dark hole is 1 and that the shape of the off-axis PSF is the same as the noncoronagraphic PSF of the

on-axis source. This assumption is likely not true in reality, but its primary consequence is that there should be a scaling factor on the flux ratios of all of the injected off-axis sources. While this has an effect on the absolute flux ratio values that are shown in this work, our analysis focuses on the relative flux ratios that are detectable using intensity alone and with the post-processing projection. Thus, this assumption does not impact the core result of this analysis, the fractional improvement provided using the projection.

4.4 Building a Response Matrix

The relevant dimensions of the objects used in this analysis are listed in Table 4.3.

Table 4.3: Quantities and dimensions for analysis of the HiCAT CLC

Quantity	Description	Dimension
E_{pup}	Vector of electric in pupil plane	$N_{basis} \times 1$
E_{det}	Vector of detector plane E-field	$N_{samples} \times 1$
I_{det}	Vector of detector plane intensity	$N_{samples} \times 1$
E_{DH}	Vector of detector plane E-field in dark hole	$N_{DH} \times 1$
I_{DH}	Vector of detector plane intensity in dark hole	$N_{DH} \times 1$
A	Instrument response matrix	$N_{samples} \times N_{basis}$
U	Left singular matrix of A	$N_{samples} \times N_{samples}$
S	Singular values of A	$N_{basis} \times 1$
V	Right singular matrix of A	$N_{basis} \times N_{basis}$
K	Projection matrix	$N_k \times N_{DH}$
O	Vector of observables	$N_k \times 1$

The quadratic term $|\mathcal{C}\Delta E_*|^2$ is again numerically linearized. Instead of propagation by applying a pre-built \mathcal{C} operator, these propagations simply use the existing simulation infrastructure provided by the HiCAT software package, using DM 1 to inject wavefront errors accordingly. However, applying Zernikes to the commands of DM 1 results in the input perturbations not being true Zernikes in phase. While they are Zernikes with respect to DM actuator commands, they result in a different map in phase due to convolution with the DM actuator influence function as well as the oversizing of the DM with respect to the beam. This does not detract from the construction of the response matrix, as the response matrix is simply linearized in a basis that is not *quite* the Zernike phase basis, but one that

is qualitatively similar to it.

The method of numerical linearization is as follows. Zernike maps are applied to DM 1, and their effects are propagated through the optical system. The corresponding error term of $|\mathcal{C}\Delta E_*|^2$ is registered into the response matrix A . The number of Zernikes used is heuristically determined by examining their propagation to the final focal plane and ensuring that the entire dark hole is well sampled by the Zernikes used. This results in using 1,500 Zernikes to sample the entire dark hole, or a N_{basis} of 1,500. The response matrix A is consequently a $31,684 \times 1,500$ matrix.

4.5 Singular Value Decomposition, Singular Modes, and the Projection K

A singular value decomposition of $A = USV^T$ is performed, revealing its singular modes and corresponding singular values. The singular values are plotted in Figure 4-4, while the first 10 singular modes as represented in the detector plane intensity basis are shown in Figure 4-5 (zoomed in and with pixels that are not in the dark hole masked out).

Unsurprisingly, there is no left-hand null space of this response matrix, but we choose to project out the first 100 modes (again, the results are not very sensitive to this exact number), resulting in $N_k = N_{samples} - 100 = 31,584$. This corresponds to the green shaded region and beyond of the singular values in Figure 4-4. We only keep in the projection the rows acting on pixels that are within the defined dark hole by extracting the rows with indices corresponding to the indices denoting the dark hole. The resulting projection K has a dimension of $N_k \times N_{DH}$, or $31,584 \times 2,912$. The following section discusses applying this projection to the data and its effect on the detection of companion point sources.




Figure 4-4: The singular values of the transfer matrix A . We choose the 101st through 31,684th singular modes to make up our projection matrix K , which will project out the first 100 modes from our data and result in $N_k = 31584$.

4.6 Synthetic Data Analysis

4.6.1 Synthetic Data Generation

Due to the testbed not having a separate off-axis source, we do not perform a full simulation of off-axis sources through the optical system, but rather approximate off-axis signals by using the direct (non-coronagraphic) exposures taken with the same DM states that generate the dark hole and shifting the (2D) intensity map by the desired distance. This process is something that can be easily repeated with the real testbed. As discussed in Section 4.3, this approximation impacts the overall scale of the flux ratios used in this study, but should have a negligible impact on the fractional change in flux ratios detectable with and without




Figure 4-5: The first 10 singular modes of A represented in the the detector plane intensity basis. For visualization, the intensity map is zoomed in and pixels not in the dark hole are masked out.

the use of the projection.

We illustrate an example dataset of a companion with a flux ratio of 4×10^{-8} at $10\lambda/D$. The dataset is chosen to have 20 frames, corresponding to 20 snapshots at different points in time. These snapshots are again assumed to be instantaneous. The optical system is first initialized in the dark hole state from the end of the stroke-minimization algorithm. For each frame, the PSD given in Equation 4.1 is used to generate an error map in DM actuator commands.

$$PSD_{2D}(k) = ak^b \quad (4.1)$$

In this equation, k is the spatial frequency. We choose a (the normalization parameter) to be 7 nm, and b (the power law exponent) to be -2.5. The resulting DM error map is added to DM 1, and both coronagraphic and non-coronagraphic exposures are taken in this DM state. In order to create data with an injected companion, the direct exposure is scaled by the desired flux ratio and shifted along the x-axis to the desired separation (to the nearest pixel). Its intensity is summed with the mutually incoherent dark hole intensity.

Twenty such snapshots are generated and form the dataset to be analyzed. The corresponding dark hole intensities without the injected companion are used for the control case.

4.6.2 Synthetic Data Analysis

The model coronagraph signal used in this case is the initial dark hole state without any extra WFE applied, and this signal is first subtracted off from each frame as it is a known contribution to the intensity. For the baseline analysis using the raw intensity, the resulting twenty model-subtracted frames are simply averaged to obtain the final measurement, while the standard deviation of the twenty frames is used as the measurement uncertainty. Figure 4-6 shows this process and the resulting data frames for the dataset with the injected companion.

Then, only the pixels that are within the defined dark hole, or the pixels targeted for nulling by the stroke minimization minimization, are gathered into the vector ΔI_{DH} , which is left-multiplied by the K matrix to obtain the observables $O = K\Delta I_{DH}$. Finally, the mean of the observables across the 20 frames is taken to be the final measurement, and the standard deviation the measurement uncertainty. We then apply detection tests to these measured intensities and observables in order to characterize the detectability of a companion with these measurements.

4.6.3 ROCs and Detection Limits

For the detection tests, we use the same test statistics as discussed in Section 3.5.3, the $\Delta\chi_r^2$ and $\Delta\epsilon^2$ test statistics. Running the simulation with the HiCAT software package is both




Figure 4-6: Top left: A single aberrated snapshot (in linear scale) with a companion with a flux ratio of 4×10^{-8} at $10.0\lambda/D$. The location of the companion is unmarked as it is not actually identifiable by eye. Top right: The data frame with the model coronagraph subtracted. Bottom left: The mean of 20 frames of data. Bottom right: The standard deviation of 20 frames of data.

more computationally intensive and harder to integrate with scientific computing clusters than the previous matrix-based propagation, thus only 40 datasets of 20 frames each are generated. Since the off-axis signal is simply the scaled and shifted non-coronagraphic exposure, we are able to use the same non-coronagraphic datasets for companions of all separations (in addition to all flux ratios), since the higher order effects are already approximated away in this treatment.

Each dataset is processed as described in Section 4.6.2, and the two test statistics are calculated for each dataset. Histograms of the resulting test statistics for the $c = 4 \times 10^{-8}$ at $10\lambda/D$ companion are shown in Figure 4-7, and the corresponding ROC curves are shown

in Figure 4-8.




Figure 4-7: Left: The histograms over 40 datasets of the test statistic $\Delta\chi_r^2$ for raw intensity and observables, both with a companion with a flux ratio of $c = 4 \times 10^{-8}$ at $10\lambda/D$ and without it. Right: The histograms over 40 datasets of the test statistic $\Delta\epsilon^2$ for raw intensity and observables, with and without the same companion.




Figure 4-8: Left: The receiver operating characteristic curve, or true positive rate as a function of false positive rate, for the $\Delta\chi_r^2$ statistic for a companion with a flux ratio of $c = 4 \times 10^{-8}$ at $10\lambda/D$. Right: The receiver operating characteristic curve for the $\Delta\epsilon^2$ test statistic for the same companion. With both of these test statistics, using the robust observables results in better performance for detections.

This process is repeated for a range of flux ratios (to a resolution of 10^{-8}) and separations (to a resolution of $0.5\lambda/D$). The range of separations was chosen to be from $8 - 12\lambda/D$, the range over which the companion PSF (or at least the majority of it) would fall within

the dark hole. Since we only have 40 datasets, it is not possible to access FPRs as low as in Chapter 3, so we choose a $FPR = 0.1$ and $TPR = 0.9$ detection threshold. The resulting detection limits are shown in Figure 4-9.




Figure 4-9: Left: The $FPR = 0.1$, $TPR = 0.9$ detection limit (at the 10^{-8} resolution) as a function of separation for raw intensities and the observables, using the $\Delta\chi_r^2$ statistic. Right: The $FPR = 0.1$, $TPR = 0.9$ detection limit (at the 10^{-8} resolution) as a function of separation for raw intensities and the observables, using the $\Delta\epsilon^2$ statistic.

4.7 Summary of Results

The results from simulating the HiCAT CLC with a dark hole are not very different from the CLC simulation around a flat wavefront, despite operating at contrasts deeper by several orders of magnitude. Using the projection modestly increases sensitivity, aside from separations near the inner edge of the dark hole with the $\Delta\chi_r^2$ test statistic. For raw intensities, $\Delta\epsilon^2$ seems to decrease sensitivity relative to $\Delta\chi_r^2$ at smaller separations, while increasing the relative sensitivity at larger separations, although this trend may not be very significant.

This work shows that results of the previously simulated CLC around a flat wavefront can also be extended to dark-hole enabled coronagraphs with deeper contrasts. Future work includes analyzing realistic time-correlated aberrations with the HiCAT CLC in the presence of a dark hole, once we understand the combination of our projection approach with whitening post-processing steps.

Chapter 5

Discussion

In this section we consider additional aspects of the approaches used in this work, leading into the summary and future work directions in Chapter 6.

5.1 Lack of Statistical Whitening

We emphasize that we explicitly chose to not combine our analysis of projections on instrument modes with other post-processing steps such as those aimed at statistically whitening the measurement error distribution (in space or in time), as our goal is to characterize the approximated quadratic instrument response and the projection determined from it separately from the effects of other post-processing operations. As discussed in Section 1.3.3, this approach can be combined with statistical whitening steps, and such a whitening step would be recommended in any practical implementation of this algorithm. The best approach to statistical whitening can depend on many complicated factors, including the spatial and temporal PSDs of the wavefront errors, how well those PSDs are understood, and the availability of reference data that matches the target observation conditions well. Thus, for our analysis, we omit statistical whitening and focus on analyzing the instrument model projection alone. However, accurate statistical whitening over time and space may alleviate some of the confusion arising from the behavior of different test statistics in the presence of time-correlated errors. In particular, we suspect that our varying results with time-correlated wavefront error is related to how well a calculation of the standard deviation of the frames actually

estimates the measurement error, but this has not been tested. Related potential future work is discussed in Chapter 6.

5.2 Kernel-Nulling

Coronagraphs are not currently designed with robust observables in mind, but a nulling interferometer that enables the extraction of kernel observables has been proposed [44]. Such a kernel-nuller is specifically designed to produce more output channels than input beams such that kernel observables robust to pupil phase differences to second order can be extracted, while also encoding visibilities in such a way that adequate information on the source can be obtained.

In summary, the proposed system in [44] has 4 input beams, with 3 nulled dark channels phase-mixed into 6 measurement channels. The second order contribution to Δx of 3 pupil phase errors ϕ (with the fourth pupil as the reference) on the 6 intensity measurements is given by $\Delta x = A \cdot [\frac{\partial^2 x}{\partial \phi_1^2}, \frac{\partial^2 x}{\partial \phi_2^2}, \frac{\partial^2 x}{\partial \phi_3^2}, \frac{\partial^2 x}{\partial \phi_1 \partial \phi_2}, \frac{\partial^2 x}{\partial \phi_1 \partial \phi_3}, \frac{\partial^2 x}{\partial \phi_2 \partial \phi_3}]^T$.

Since there are 6 measurements and 6 contributing degrees of freedom, the existence of the kernel is dependent on the null space of the square matrix A . In [44], an A matrix that gives 3 kernel observables is analyzed, and a corresponding instrument design is proposed.

5.3 Connection with PASTIS Formalism

Studies on the quadratic response of a coronagraph can be useful for different applications, including modeling the quadratic response of a coronagraph to the segment disturbances of a segmented telescope, such as with the Pair-based Analytical model for Segmented Telescopes Imaging from Space (PASTIS) [45] [46].

PASTIS considers the effects of the quadratic response on the *average* intensity contrast over the entire dark hole (or region of interest). The contribution of the quadratic term is given by the matrix product $a_l M_l a_l^T$, where a_l is a vector of Zernike coefficients (the subscript l denotes the Zernike) describing aberrations on each segment, and M_l a matrix describing the effects of each pair-wise combination of segments with the l th Zernike applied on the

average contrast in the dark hole. The goal of PASTIS is to understand which modes the coronagraph is most sensitive to in order to calculate stability tolerances for the segments.

The post-processing approach in this work is somewhat akin to doing PASTIS backwards - understanding which modes the coronagraph is most sensitive to in order to take those modes out of the data (assuming those modes have not been tolerated and engineered out). A few key differences in the approaches are as follows. We consider only "pistons" on pupil elements instead of various Zernikes on segments. Instead of reducing dimension by averaging over the dark hole, we consider the response on each individual pixel. This would result in a huge dimensional space, so we reduce the dimension by taking the shortcut of only considering the norm-squared of the propagation of each basis vector, rather than all possible pair-wise combinations. However, this connection reminds us that a full dimensional representation for the quadratic response does exist.

5.4 Dimensions for Continuous Coronagraphy

Insights from PASTIS and kernel-nulling help build a full dimensional view of measurements in continuous coronagraphy.

Extending the notation of PASTIS, the quadratic response of each detector plane pixel intensity ΔI_q^k to wavefront error (represented as complex pistons of pupil elements) can be written as $\Delta I_q^k = a_i M_j^{ik} a^j$. The quantity M_j^{ik} relates each pairwise combination of pupil elements a^i and a^j to its effect on each detector plane pixel k , and is a 3-tensor containing the Hessians of each pixel intensity with respect to the wavefront error. The dimension of M_j^{ik} can be fairly large: with N_{basis} pupil basis vectors and $N_{samples}$ pixels, M_j^{ik} would have $(N_{basis})^2 \times N_{samples}$ entries. With sufficient computational resources, it is possible to build this 3-tensor using methods similar to how M_l is built for PASTIS, and to investigate this quantity directly.

Since Hessians are symmetric as a result of partial derivatives commuting, many of the entries of M_j^{ik} are redundant. The kernel-nulling formalism represents only the nonredundant entries of M_j^{ik} with the A matrix, which acts on a vector of all nonredundant pairwise combinations of pupil elements. This formalism can also be extended to the continuous coro-

nagraphy case. When removing these straightforward redundancies, the input dimensions scales as $N_{basis} + \binom{N_{basis}}{2}$ instead of $(N_{basis})^2$.

We review the dimensions of our analysis of continuous coronagraphy, and show that the input dimension quickly outgrows the output dimension. We chose an N_{basis} of 276 Zernikes, and claimed that this approximately corresponds to sampling 276 resolution elements in the detector plane. If we assume that sampling beyond Nyquist in the detector plane does not provide much additional meaningful information, this results in $276 \times 4 = 1104$ meaningful output measurements. Meanwhile, the input dimension is $276 + \binom{276}{2} = 38226$. A transfer matrix A describing this relationship would thus be 1104×38226 , which makes the existence of a null space pretty unlikely.

We build a full dimensional A matrix with our existing CLC model by including all possible pairwise combinations of Zernikes. Without changing the existing sampling, we obtain a 10000×38226 A matrix. A singular value decomposition results in the singular value spectrum shown in Figure 5-1, and the singular modes in Figure 5-2. The apparent drop in singular values around 2000 seems to be a numerical effect independent of the actual size of the input space, and further inspection shows that virtually no astrophysical signal gets propagated to the modes after the drop-off. Although this has not been rigorously explored, we suspect that the drop-off indicates the actual size of the output space of our Fourier optics *model*, as fixed by the pupil sampling and the rank of the coronagraph operator \mathcal{C} , regardless of how much we oversample the detector plane or how we choose to represent the pupil.

The singular value spectrum is completely smooth up to the drop-off, indicating that there is no kernel in the space that astrophysical signal can propagate to. However, similar to the dimensionally-reduced version explored for most of this work, the most dominant error modes can still be projected out. In fact, the dominant modes using the approximation that neglects cross-terms seem to resemble the dominant modes of the full-dimensional matrix fairly well, and using a projection determined from the A matrix that includes the cross-terms to analyze synthetic data does not seem to result in noticeable differences in the flux ratio detection limits compared to using the approximation neglecting cross-terms.

Although a preliminary look at the full dimensional representation is shown here, a




Figure 5-1: The singular value spectrum of a full-dimensional 10000×38226 A matrix. The drop at around 2000 seems to be a numerical effect independent of the size of the input space. Applying the left singular matrix U to synthetic data reveals that almost all of the astrophysical signal projects onto modes before the drop. The singular values are smooth before the drop, indicating that there is no kernel that also overlaps well with astrophysical signal.

more careful analysis and treatment of input and output sampling is needed to analyze the properties of A . The full dimensional treatment can help inform the joint analysis of instrumentation and robust observables, and just as kernel-nulling modifies the output of a traditional nulling-interferometer to enable robust observable extraction, it may be possible to modify coronagraphs (a kind of continuous nulling-interferometer) to do so as well.




Figure 5-2: The first 10 singular modes of the full dimensional A matrix as represented in the detector plane intensity basis. These modes are qualitatively similar to the approximated singular modes in Figure 3-7, and using a projection determined from this matrix to analyze synthetic data does not result in noticeable differences in flux ratio detection limits compared to using the approximation neglecting cross-terms.

Chapter 6

Future Work and Conclusions

In this chapter, we summarize our work on coronagraphic robust observables, and suggest several related topics to further explore.

In this work, we explore a post-processing algorithm, which projects the data onto approximated instrumental error modes in order to remove the dominant modes from the data. This algorithm provides modest gains for binary companion detection with errors that are not time-correlated. More work is needed to understand effects of the projection on time-correlated errors, especially how it combines with (possibly imperfect) statistical whitening processes. Given that most observational data is limited by highly time-correlated, quasi-static wavefront error, we do not yet expect this algorithm to provide improvements on most observational data with our current approach to error estimation without additional whitening post-processing steps.

We suggest a few directions for future work.

1. Investigate the effects of non-Gaussian and particularly time-correlated error distributions on error estimation, test statistics, and hypothesis testing. This is not necessarily related to robust observables, and is a general problem on which there is plenty of existing literature and active research [47]. However, our work on attempting hypothesis testing with robust observables highlights the importance of understanding error distributions and correlations, and future work can include combining our instrumental projection with additional statistical whitening. In theory, full statistical whitening

over time and space would remove many of these issues, but such whitening can be difficult to achieve in practice for non-Gaussian error from PSDs that are unknown a priori, especially at small angular separations [47]. Understanding the effects of imperfect statistical whitening thus may also play an important part in the analysis of post-processing algorithms.

2. Treat the continuous coronagraphy problem carefully with its full dimensions. Although it is unlikely that a full dimensional A (or equivalently M_j^{ik}) for a continuous coronagraph would have a null space, building and analyzing it may still provide some valuable insights. For example, A removes the obvious redundancies from M_j^{ik} , but there are likely additional degeneracies amongst the remaining degrees of freedom, especially with highly symmetric coronagraphs. Analyzing A would reveal how much larger the input space actually is than the output space when considering these possible degeneracies. From this understanding, attempts can be made to extend coronagraph design in a way that enables robust observable extraction with self-calibrating properties.
3. From a very broad perspective, it may be worthwhile to examine instrumentation and post-processing together from an information point of view. This process would begin with high-level science objectives that define a few quantities to be measured. Then, instrumentation (possibly including both optical design and PSD engineering) and post-processing can be jointly optimized to maximize the Fisher information of those scientific quantities. This would result in instruments that are perhaps not flexible for general purpose use, but rather super-optimized for a specific science case.

The main outcome of this work is an exploration of one approach to dimension reduction of the coronagraphic robust observables problem, understanding the limitations of the approach, and motivating further exploration of the joint analysis of coronagraphic instrumentation and post-processing with the full dimensional representation.

Appendix A

Heuristic Details of Robust Observables with the CLC

We provide a few details on a more heuristic view of the effects of using the projection on the data, using a few example datasets. Although these qualitative pictures are not directly useful for rigorously quantifying the effects of the projection, they are somewhat helpful in aiding human interpretation of the results.

Figure A-1 simply plots the aberrated signal (as determined from the mean of 20 frames) against the model (unaberrated) signal, and includes a new quantity that is calculated by applying the transpose of the U matrix (as determined by the SVD of A) to the intensities. The quantity $U^T I$ is simply the projection of the intensities onto the instrumental modes, but *without* the removal of the most dominant error modes. Only the $\alpha = 1$ case is shown here, as the plots for the time-correlated data are very similar. These correlation plots show an illustrative picture of what the post-processing projection entails: aberrations cause the intensity measurements to no longer follow the 1-to-1 correspondence with the model signal. The data can be projected onto the instrument modes with $U^T I$, while the projection KI excludes data that projects onto the most sensitive instrumental modes, which does apparently also remove the measurements with the largest scatter away from the model signal. However, this outlook does not take into account the spatial distribution of the modes (for example, most of the scatter actually occurs in very different spatial locations than the peak of the companion intensity), and is thus an incomplete picture.

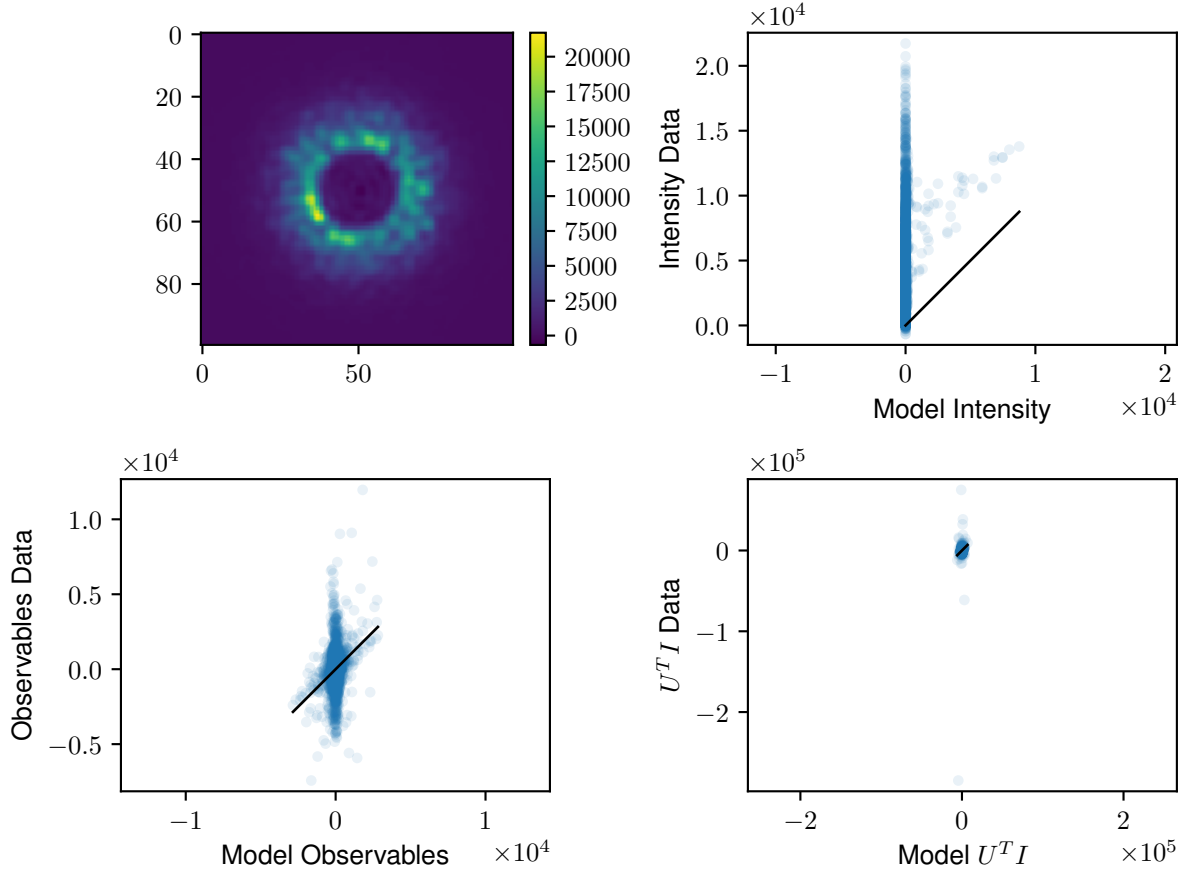


Figure A-1: Top left) The mean intensity of 20 non-time-correlated ($\alpha = 1$) frames with the model star coronagraph subtracted of a companion with a flux ratio of 4×10^{-4} at $7.5\lambda/D$. Top right) The mean intensity data plotted against the model intensity data for the system, where the black line indicates a 1-to-1 correspondence. Bottom left) The mean observables data plotted against the model observables data for the system, where the black line indicates a 1-to-1 correspondence. Bottom right) The mean $U^T I$ data plotted against the model $U^T I$ data for the system, where the black line indicates a 1-to-1 correspondence.

The closer the measurements in blue fall along the black line, the better the measurements are at capturing the true astrophysical signal. Farther separations from the black line correspond to larger deviations due to wavefront error.

Another qualitative effect is shown in Figure A-2, which shows the effects of applying the Moore-Penrose pseudoinverse of K , denoted as K^{-1} , on the mean observables over the 20 frames to reconstruct the post-projection signal in the detector plane Y (given by Equation A.1), alongside the original mean intensities.

$$Y = K^{-1}O \tag{A.1}$$

Perhaps it is incorrect for the reconstruction to be human-interpreted, but very qualitatively, it resembles (but is not exactly) an anti-symmetrization of the signal. This is related to the symmetry of the speckles arising from the linear and quadratic terms of the coronagraph response to wavefront error, previously examined in the context of "speckle pinning" [36] [48]. To summarize, if the optics are described as purely a Fourier transform, speckles that arise from phase errors are anti-symmetric for the linear term and symmetric for the quadratic term, and an increase in SNR is achievable by either symmetrizing or anti-symmetrizing the image depending on which term is dominant.

The coronagraph operators discussed in this work are not Fourier transforms, but they are somewhat close to Fourier transforms, particularly for high spatial frequency speckles. Empirically, most of the dominant modes for the quadratic term are symmetric, but anti-symmetric modes (or at least modes with anti-symmetric components) do occur with finite singular values. We suspect that this approach of projecting out the dominant quadratic modes thus combines an approximate anti-symmetrization (taking into account the anti-symmetric contributions) with a search for signal in the pixels that are less affected by phase errors.

While it is not clear by eye if the positive peak of the reconstructed signal is particularly brighter (relative to the speckles) than in the raw intensities before the projection, there is also a negative peak that also contributes to the total signal. More quantitative approaches such as the hypothesis tests performed in this work are needed to actually determine in which cases this "spreading" of the signal information actually contributes to a increase in the total SNR of the entire signal.

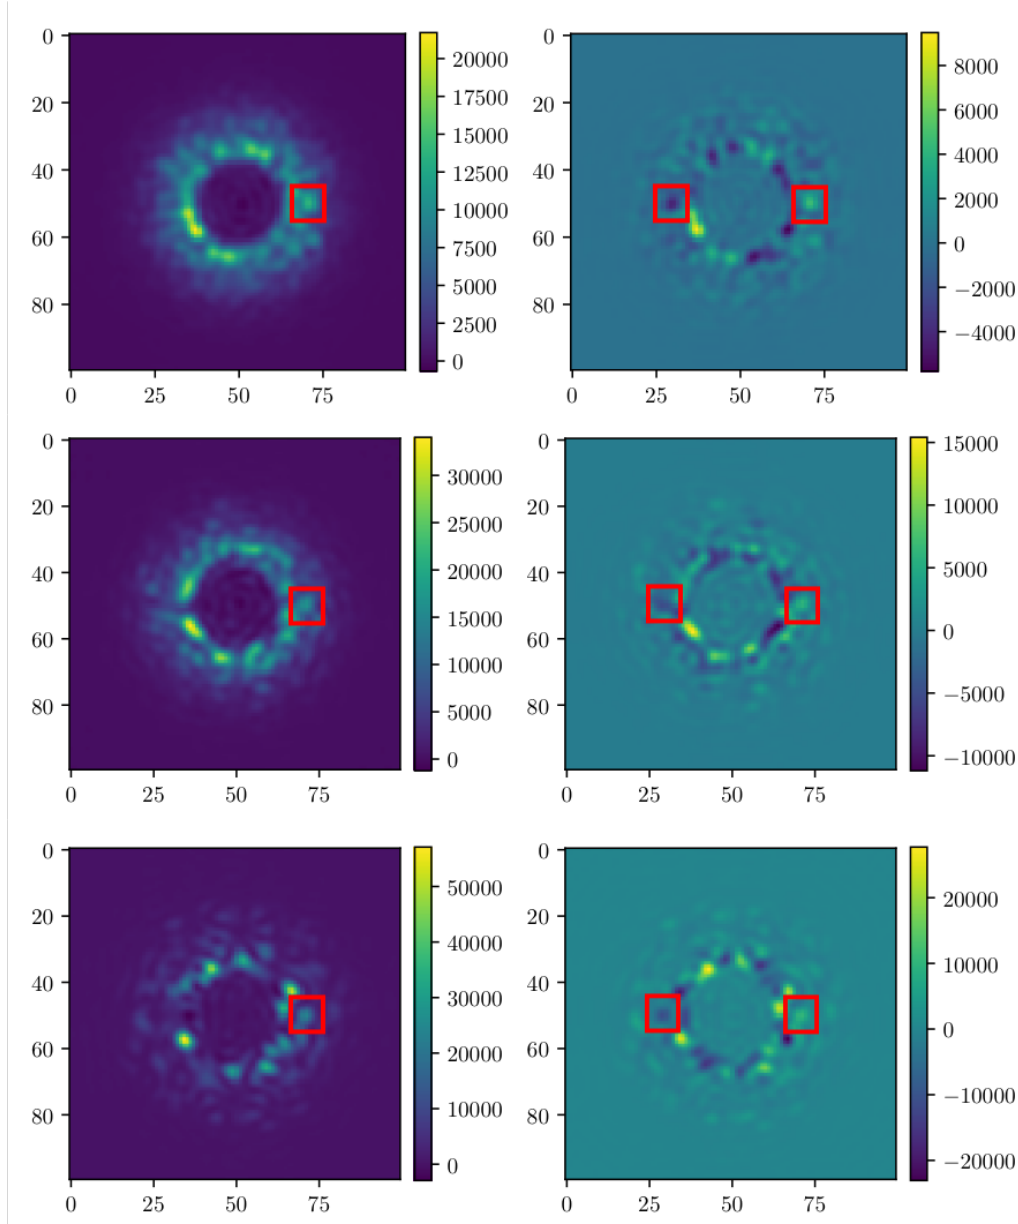


Figure A-2: Top) Data corresponding to a companion with a flux ratio of 4×10^{-4} at $7.5\lambda/D$, and no time-correlation $\alpha = 1$. Middle) Data corresponding to a companion with a flux ratio of 5×10^{-4} at $7.5\lambda/D$, and time-correlation parameter $\alpha = 0.5$. Bottom) Data corresponding to a companion with a flux ratio of 9×10^{-4} at $7.5\lambda/D$, and time-correlation parameter $\alpha = 0.1$.

Left) The mean intensities of 20 frames. Right) A reconstruction Y obtained by applying the Moore-Penrose pseudoinverse of the projection matrix as described by Equation A.1.

"Locations" of peak companion signal are identified with red boxes.

Bibliography

- [1] John E. Krist. PROPER: an optical propagation library for IDL, 2007.
- [2] Marshall D. Perrin, Rémi Soummer, Erin M. Elliott, Matthew D. Lallo, and Anand Sivaramakrishnan. Simulating point spread functions for the James Webb Space Telescope with WebbPSF. In *Proceedings of SPIE*, volume 8442 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 84423D, September 2012.
- [3] E. H. Por, S. Y. Haffert, V. M. Radhakrishnan, D. S. Doelman, M. Van Kooten, and S. P. Bos. High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator. In *Adaptive Optics Systems VI*, volume 10703 of *Proc. SPIE*, 2018.
- [4] Astropy Collaboration, T. P. Robitaille, E. J. Tollerud, P. Greenfield, M. Droettboom, E. Bray, T. Aldcroft, M. Davis, A. Ginsburg, A. M. Price-Whelan, W. E. Kerzendorf, A. Conley, N. Crighton, K. Barbary, D. Muna, H. Ferguson, F. Grollier, M. M. Parikh, P. H. Nair, H. M. Unther, C. Deil, J. Woillez, S. Conseil, R. Kramer, J. E. H. Turner, L. Singer, R. Fox, B. A. Weaver, V. Zabalza, Z. I. Edwards, K. Azalee Bostroem, D. J. Burke, A. R. Casey, S. M. Crawford, N. Dencheva, J. Ely, T. Jenness, K. Labrie, P. L. Lim, F. Pierfederici, A. Pontzen, A. Ptak, B. Refsdal, M. Servillat, and O. Streicher. Astropy: A community Python package for astronomy. *Astronomy and Astrophysics*, 558:A33, October 2013.
- [5] A. M. Price-Whelan, B. M. Sipőcz, H. M. Günther, P. L. Lim, S. M. Crawford, S. Conseil, D. L. Shupe, M. W. Craig, N. Dencheva, A. Ginsburg, J. T. VanderPlas, L. D. Bradley, D. Pérez-Suárez, M. de Val-Borro, (Primary Paper Contributors, T. L. Aldcroft, K. L. Cruz, T. P. Robitaille, E. J. Tollerud, (Astropy Coordination Committee, C. Ardelean, T. Babej, Y. P. Bach, M. Bachetti, A. V. Bakanov, S. P. Bamford, G. Bar-entsen, P. Barmby, A. Baumbach, K. L. Berry, F. Biscani, M. Boquien, K. A. Bostroem, L. G. Bouma, G. B. Brammer, E. M. Bray, H. Breytenbach, H. Buddelmeijer, D. J. Burke, G. Calderone, J. L. Cano Rodríguez, M. Cara, J. V. M. Cardoso, S. Cheedella, Y. Copin, L. Corrales, D. Crichton, D. D’Avella, C. Deil, É. Depagne, J. P. Dietrich, A. Donath, M. Droettboom, N. Earl, T. Erben, S. Fabbro, L. A. Ferreira, T. Finethy, R. T. Fox, L. H. Garrison, S. L. J. Gibbons, D. A. Goldstein, R. Gommers, J. P. Greco, P. Greenfield, A. M. Groener, F. Grollier, A. Hagen, P. Hirst, D. Homeier, A. J. Horton, G. Hosseinzadeh, L. Hu, J. S. Hunkeler, Ž. Ivezić, A. Jain, T. Jenness, G. Kanarek, S. Kendrew, N. S. Kern, W. E. Kerzendorf, A. Khvalko, J. King, D. Kirkby, A. M. Kulkarni, A. Kumar, A. Lee, D. Lenz, S. P. Littlefair, Z. Ma, D. M. Macleod,

- M. Mastropietro, C. McCully, S. Montagnac, B. M. Morris, M. Mueller, S. J. Mumford, D. Muna, N. A. Murphy, S. Nelson, G. H. Nguyen, J. P. Ninan, M. Nöthe, S. Ogaz, S. Oh, J. K. Parejko, N. Parley, S. Pascual, R. Patil, A. A. Patil, A. L. Plunkett, J. X. Prochaska, T. Rastogi, V. Reddy Janga, J. Sabater, P. Sakurikar, M. Seifert, L. E. Sherbert, H. Sherwood-Taylor, A. Y. Shih, J. Sick, M. T. Silbiger, S. Singanamalla, L. P. Singer, P. H. Sladen, K. A. Sooley, S. Sornarajah, O. Streicher, P. Teuben, S. W. Thomas, G. R. Tremblay, J. E. H. Turner, V. Terrón, M. H. van Kerkwijk, A. de la Vega, L. L. Watkins, B. A. Weaver, J. B. Whitmore, J. Woillez, V. Zabalza, and (Astropy Contributors). The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package. *Astronomical Journal*, 156:123, September 2018.
- [6] Travis E Oliphant. *A guide to NumPy*, volume 1. Trelgol Publishing USA, 2006.
- [7] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, CJ Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [8] J. D. Hunter. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.
- [9] Lick Observatory. Introduction to adaptive optics and its history. http://cfao.ucolick.org/EO/resources/History_AO_Max.pdf.
- [10] Benjamin J. S. Pope. Kernel phase and kernel amplitude in Fizeau imaging. *Monthly Notices of the Royal Astronomical Society*, 463(4):3573–3581, Dec 2016.
- [11] Rémi Soummer, Gregory R. Brady, Keira Brooks, Thomas Comeau, Élodie Choquet, Tom Dillon, Sylvain Egron, Rob Gontrum, John Hagopian, Iva Luginja, Lucie Leboulleux, Marshall D. Perrin, Peter Petrone, Laurent Pueyo, Johan Mazoyer, Mamadou N’Diaye, A. J. Eldorado Riggs, Ron Shiri, Anand Sivaramakrishnan, Kathryn St. Laurent, Ana-Maria Valenzuela, and Neil T. Zimmerman. High-contrast imager for complex aperture telescopes (HiCAT): 5. first results with segmented-aperture coronagraph and wavefront control. In , volume 10698 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1069810, August 2018.
- [12] W. A. Traub and B. R. Oppenheimer. *Direct Imaging of Exoplanets*, pages 111–156. 2010.
- [13] Laurent Pueyo. *Direct Imaging as a Detection Technique for Exoplanets*, pages 1–61. Springer International Publishing, Cham, 2018.
- [14] The Lyot Project. Coronagraphy. <http://lyot.org/background/coronagraphy.html>.

- [15] D. R. Ardila, J. Krist, M. Clampin, D. A. Golimowski, G. Hartig, H. C. Ford, G. D. Illingworth, C. Burrows, F. Bartko, N. Benítez, J. P. Blakeslee, R. Bouwens, T. J. Broadhurst, R. A. Brown, E. Cheng, N. Cross, P. D. Feldman, M. Franx, C. Gronwall, L. Infante, R. A. Kimble, M. Lesser, A. R. Martel, F. Menanteau, G. R. Meurer, G. Miley, M. Postman, P. Rosati, M. Sirianni, W. B. Sparks, JHU Tran, Z. I. Tsvetanov, R. L. White, and W. Zheng. The ACS Coronagraph. In *American Astronomical Society Meeting Abstracts*, volume 201 of *American Astronomical Society Meeting Abstracts*, page 41.03, December 2002.
- [16] Space Telescope Science Institute. STIS Instrument Handbook. <https://hst-docs.stsci.edu/stisihb>.
- [17] Space Telescope Science Institute. MIRI Coronagraphic Imaging. <https://jwst-docs.stsci.edu/mid-infrared-instrument/miri-observing-modes/miri-coronagraphic-imaging>.
- [18] Goddard Space Flight Center. WFIRST Observatory. <https://wfirst.gsfc.nasa.gov/observatory.html>.
- [19] Jet Propulsion Laboratory. HabEx Coronagraph. <https://www.jpl.nasa.gov/habex/mission/instruments/coronagraph/>.
- [20] Stuart Shaklan. LUVOIR Tech Note Series: Starlight Suppression with Coronagraphs. https://asd.gsfc.nasa.gov/luvoir/technology/notes/Coronagraphs-LUVOIR-Tech-Note_Final.pdf.
- [21] O. Guyon, E. A. Pluzhnik, M. J. Kuchner, B. Collins, and S. T. Ridgway. Theoretical limits on extrasolar terrestrial planet detection with coronagraphs. *The Astrophysical Journal Supplement Series*, 167(1):81–99, Nov 2006.
- [22] Robert K. Tyson. *Introduction to Adaptive Optics*. The Society of Photo-optical Instrumentation Engineers, Bellingham, 2000.
- [23] Tyler D. Groff, A. J. Eldorado Riggs, Brian Kern, and N. Jeremy Kasdin. Methods and limitations of focal plane sensing, estimation, and control in high-contrast imaging. *Journal of Astronomical Telescopes, Instruments, and Systems*, 2(1):1 – 15, 2015.
- [24] R. C. Jennison. A Phase Sensitive Interferometer Technique for the Measurement of the Fourier Transforms of Spatial Brightness Distributions of Small Angular Extent. *Monthly Notices of the Royal Astronomical Society*, 118(3):276–284, 06 1958.
- [25] Frantz Martinache. Kernel-phases for high-contrast detection beyond the resolution limit. In *Techniques and Instrumentation for Detection of Exoplanets V*, volume 8151 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 815111, Oct 2011.
- [26] Garreth Ruane, Henry Ngo, Dimitri Mawet, Olivier Absil, Elodie Choquet, Therese Cook, Carlos Gomez Gonzalez, Elsa Huby, Keith Matthews, Tiffany Meshkat, Madalena Reggiani, Eugene Serabyn, Nicole Wallack, and W. Jerry Xuan. Reference star

differential imaging of close-in companions and circumstellar disks with the NIRC2 vortex coronagraph at W.M. Keck Observatory, 2019.

- [27] G. Schneider, A. J. Weinberger, E. E. Becklin, J. H. Debes, and B. A. Smith. STIS Imaging of the HR 4796A Circumstellar Debris Ring. *The Astronomical Journal*, 137(1):53–61, January 2009.
- [28] Rémi Soummer, Laurent Pueyo, and James Larkin. Detection and Characterization of Exoplanets and Disks Using Projections on Karhunen-Loève Eigenimages. *Astrophysical Journal, Letters*, 755(2):L28, August 2012.
- [29] Laurent Pueyo. Detection and Characterization of Exoplanets using Projections on Karhunen Loeve Eigenimages: Forward Modeling. *Astrophysical Journal*, 824(2):117, June 2016.
- [30] David Lafrenière, Christian Marois, René Doyon, Daniel Nadeau, and Étienne Artigau. A New Algorithm for Point-Spread Function Subtraction in High-Contrast Imaging: A Demonstration with Angular Differential Imaging. *Astrophysical Journal*, 660(1):770–780, May 2007.
- [31] Bin Ren, Laurent Pueyo, Guangtun Ben Zhu, John Debes, and Gaspard Duchêne. Non-negative Matrix Factorization: Robust Extraction of Extended Structures. *Astrophysical Journal*, 852(2):104, January 2018.
- [32] Frédéric P. A. Vogt, Frantz Martinache, Olivier Guyon, Takashi Yoshikawa, Kaito Yokochi, Vincent Garrel, and Taro Matsuo. Coronagraphic low-order wavefront sensor: Postprocessing sensitivity enhancer for high-performance coronagraphs. *Publications of the Astronomical Society of the Pacific*, 123(910):1434–1441, Dec 2011.
- [33] Leonid Pogorelyuk and N. Jeremy Kasdin. Dark hole maintenance and a posteriori intensity estimation in the presence of speckle drift in a high-contrast space coronagraph. *The Astrophysical Journal*, 873(1):95, Mar 2019.
- [34] Benjamin Pope, Peter Tuthill, Sasha Hinkley, Michael J. Ireland, Alexandra Greenbaum, Alexey Latyshev, John D. Monnier, and Frantz Martinache. The Palomar kernel-phase experiment: testing kernel phase interferometry for ground-based astronomical observations. *Monthly Notices of the Royal Astronomical Society*, 455(2):1647–1653, Jan 2016.
- [35] Jens Kammerer, Michael J. Ireland, Frantz Martinache, and Julien H. Girard. Kernel phase imaging with VLT/NACO: high-contrast detection of new candidate low-mass stellar companions at the diffraction limit. *Monthly Notices of the Royal Astronomical Society*, 486(1):639–654, Jun 2019.
- [36] E. E. Bloemhof. Suppression of speckle noise by speckle pinning in adaptive optics. *The Astrophysical Journal*, 582(1):L59–L62, Dec 2002.

- [37] Marshall Perrin, Anand Sivaramakrishnan, Russell Makidon, Ben Oppenheimer, and James Graham. The structure of high strehl ratio point-spread functions. *The Astrophysical Journal*, 596, 06 2003.
- [38] Robert J. Noll. Zernike polynomials and atmospheric turbulence*. *J. Opt. Soc. Am.*, 66(3):207–211, Mar ts , url = <http://www.osapublishing.org/abstract.cfm?URI=josa-66-3-207>, doi = 10.1364/JOSA.66.000207, abstract = This paper discusses some general properties of Zernike polynomials, such as their Fourier transforms, integral representations, and derivatives. A Zernike representation of the Kolmogoroff spectrum of turbulence is given that provides a complete analytical description of the number of independent corrections required in a wave-front compensation system.,.
- [39] N. Jeremy Kasdin and Isabelle Braems. Linear and bayesian planet detection algorithms for the Terrestrial planet finder. *The Astrophysical Journal*, 646(2):1260–1274, Aug 2006.
- [40] Rebecca Jensen-Clem, Dimitri Mawet, Carlos A. Gomez Gonzalez, Olivier Absil, Ruslan Belikov, Thayne Currie, Matthew A. Kenworthy, Christian Marois, Johan Mazoyer, Garreth Ruane, Angelle Tanner, and Faustine Cantalloube. A new standard for assessing the performance of high contrast imaging systems. *The Astronomical Journal*, 155(1):19, Dec 2017.
- [41] Alban Ceau, David Mary, Alexandra Greenbaum, Frantz Martinache, Anand Sivaramakrishnan, Romain Laugier, and Mamadou N’Diaye. Kernel-phase Detection Limits : Hypothesis Testing and the Example of JWST NIRISS Full Pupil Images. *arXiv e-prints*, page arXiv:1908.03130, Aug 2019.
- [42] J. Neyman and E. S. Pearson. On the Problem of the Most Efficient Tests of Statistical Hypotheses. *Philosophical Transactions of the Royal Society of London Series A*, 231:289–337, January 1933.
- [43] Mamadou N’Diaye, Elodie Choquet, Laurent Pueyo, Erin Elliot, Marshall D. Perrin, J. Kent Wallace, Tyler Groff, Alexis Carlotti, Dimitri Mawet, Matt Sheckells, Stuart Shaklan, Bruce Macintosh, N. Jeremy Kasdin, and Rémi Soummer. High-contrast imager for complex aperture telescopes (HiCAT): 1. testbed design. In Stuart Shaklan, editor, *Techniques and Instrumentation for Detection of Exoplanets VI*, volume 8864, pages 558 – 567. International Society for Optics and Photonics, SPIE, 2013.
- [44] Frantz Martinache and Michael J. Ireland. Kernel-nulling for a robust direct interferometric detection of extrasolar planets. *Astronomy Astrophysics*, 619:A87, Nov 2018.
- [45] Lucie Leboulleux, Jean-François Sauvage, Laurent Pueyo, Thierry Fusco, Rémi Soummer, Johan Mazoyer, Anand Sivaramakrishnan, Mamadou N’Diaye, and Olivier Fauvarque. Pair-based analytical model for segmented telescopes imaging from space (PASTIS) for sensitivity analysis, 2018.
- [46] Iva Laginja, Lucie Leboulleux, Laurent Pueyo, Rémi Soummer, Jean-François Sauvage, Laurent Mugnier, Laura E. Coyle, J. Scott Knight, Kathryn St. Laurent, Emiel H. Por,

- and James Noss. Wavefront error tolerancing for direct imaging of exo-Earths with a large segmented telescope in space. In *Techniques and Instrumentation for Detection of Exoplanets IX*, volume 11117, pages 382 – 396. International Society for Optics and Photonics, SPIE, 2019.
- [47] Dimitri Mawet, Julien Milli, Zahed Wahhaj, Didier Pelat, Olivier Absil, Christian Delacroix, Anthony Boccaletti, Markus Kasper, Matthew Kenworthy, Christian Marois, Bertrand Mennesson, and Laurent Pueyo. Fundamental limitations of high contrast imaging set by small sample statistics. *The Astrophysical Journal*, 792, 07 2014.
- [48] Remi Soummer, Andre Ferrari, Claude Aime, and Laurent Jolissaint. Speckle noise and dynamic range in coronagraphic images. *The Astrophysical Journal*, 669(1):642–656, Nov 2007.