

NON-LINEAR REDUCED-ORDER ESTIMATION
AND CONTROL OF SPACE-BASED
CORONAGRAPHS

LEONID POGORELYUK

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
MECHANICAL AND AEROSPACE ENGINEERING
ADVISERS: PROFESSOR CLARENCE W. ROWLEY AND
PROFESSOR N. JEREMY KASDIN

APRIL 2020

© Copyright by Leonid Pogorelyuk, 2020.

All rights reserved.

Abstract

To directly image an Earth-like planet orbiting another star, coronagraphs would have to suppress the light from that star to within one part per ten billion. Such observations would have to be carried out from space and even then, the images will be extremely sensitive to structural and thermal fluctuations in the optical elements of the telescope. With just tens of photons reaching each detector during a single exposure, sensing time variations in the residual starlight (speckles) and discerning them from planets proves to be a challenging task. This thesis describes algorithms for estimating the speckles both during the observations and in post-processing. It states these estimation problems in terms of the electric field rather than the intensity of the light, which then become non-linear and high-dimensional. On the other hand, this approach allows taking into account the influence of deformable mirrors, and seamlessly incorporating probabilistic and reduced-order formulations for the underlying optimization tasks.

This thesis also present numerical simulations of a realistic model of a space coronagraph in various observation scenarios. They suggest that it is possible to continuously maintain a high image contrast even when pointing at a dim star. Moreover, employing a reduced-order model of the electric field allows increasing the accuracies of both online and offline estimators without increasing their time complexity. Under favorable conditions, the corresponding planet detectability thresholds lie within less than an order of magnitude of the instrument's theoretical limits.

Acknowledgements

First, I want to thank my advisers, Dr. C. W. Rowley and Dr. N. J. Kasdin, for their guidance in choosing research problems, approaching them with mathematical rigor and communicating results to the scientific community.

I would also like to thank present and past members of my PhD committee, Dr. R. J. Vanderbei, Dr. A. Majumdar, Dr. N. E. Leonard and Dr. Y. G. Kevrekidis, for providing constructive criticism throughout my research.

Finally, I wish to thank my collaborators, Dr. L. Pueyo, Dr. R. Soummer, S. E. Otto, S. M. Redmond, I. Luginja, and S. D. Will.

Part of my research was sponsored by the Army Research Office and was accomplished under Grant Number W911NF-17-1-0512. The views and conclusions contained in this thesis are mine and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government.

This dissertation carries T#3392 in the records of the Department of Mechanical and Aerospace Engineering.

Contents

Abstract	iii
Acknowledgements	iv
1 Introduction	1
1.1 Detecting and Characterizing Exoplanets	1
1.2 High Contrast Imaging	3
1.3 Organization and Contributions	7
1.4 Achieving High Contrast	9
1.4.1 A Linear Coronagraph Model	10
1.4.2 Electric Field Conjugation (EFC)	12
2 Dark Hole Maintenance	15
2.1 Non-linear Recursive Estimation	16
2.1.1 State Equations	17
2.1.2 Estimation with an Extended Kalman Filter (EKF)	18
2.2 Closing the Loop and Dithering	20
2.2.1 Dark Hole Maintenance Control Law	21
2.2.2 Recalibration	21
2.2.3 Broadband Light	22
2.3 Numerical Results	24
2.3.1 Effectiveness of Closed-Loop Dark Hole Maintenance	27

3	Reduced-order Approximation of the Filtering Problem	29
3.1	A Block-diagonal Approximation	33
3.1.1	Efficient Computation of the Block-diagonal Approximation	37
3.2	Numerical Results	39
3.3	Properties (Fixed Matrices Case)	41
3.3.1	Convergence	41
3.3.2	Stability	43
3.3.3	Asymptotic Decoupling	44
4	Reduced-order Approach to Post Processing	48
4.1	An Electric Field Formulation of the Post-processing Problem	50
4.1.1	A Maximum Likelihood (ML) Formulation	50
4.1.2	Regularization via Drift Increments Prior	52
4.1.3	Regularization via Electric Field Order Reduction (EFOR)	53
4.1.4	Numerical Optimization Procedure	55
4.1.5	Realistic Telescope Effects with EFOR	56
4.1.6	Reference Observations	58
4.2	Numerical Results	59
4.2.1	Effects of Dimensionality	62
4.2.2	Effects of Drift Magnitude	63
5	Future Directions	65
5.1	EFOR for Ground-based Coronagraphy	65
5.2	Low-order Coordinates Extraction via Debiased Isomaps	66
5.2.1	Choice of the Distance Metric	69
5.2.2	Estimation of Metric Bias	72
5.2.3	Choice of Neighborhood Criterion	74
5.2.4	Results	75

6	Conclusions	77
	Bibliography	80

Chapter 1

Introduction

1.1 Detecting and Characterizing Exoplanets

Coronagraphs are optional elements in modern telescopes and were invented in the 1930's to study the solar corona by blocking most of the light from the Sun[1]. Since the 2000's, they have been used to detect planets orbiting around stars other than the Sun (exoplanets)[2, 3, 4, 5] – a field known as “high-contrast imaging”. The “contrast” above refers to the ratio of light intensities between the host star and the detected exoplanet.

However, out of more than four thousand confirmed planets discovered by the year 2020, less than fifty were detected via direct imaging (see fig. 1.1)[6]. Methods which yield the vast majority of detections rely on variations in the starlight: it either dims slightly when the star is partially blocked by a transiting planet[7, 8] or is Doppler shifted by small variations in the radial velocity of the star[9] due to the gravitational pull of its planets.

The difference between the nominal spectrum of the star and its spectrum when partially obstructed by a planet, gives some insight into the contents of the atmosphere of the transiting planet[10]. Such high precision transmission/emission spectra are




Figure 1.1: Confirmed exoplanets from NASA’s exoplanet archive[6]. Most are gas giants detected by indirect methods.

currently available only for a small fraction of the exoplanets, mostly gas giants with thick atmospheres orbiting close to their host stars[11]. On the other hand, rocky planets such as Earth or Mars, have a very thin to non-existent atmosphere, and are therefore extremely hard to characterize by indirect methods[12]. Direct imaging would circumvent this issue by measuring the *reflected* spectrum of such exoplanets although, so far, it has been proven to be even more challenging.

1.2 High Contrast Imaging

To illustrate the difficulties involved in direct imaging, consider the planet 51 Eridani b which was discovered in 2014 at a distance of 13 AU ($1.9 \cdot 10^{12}$ m) away from the its host star, 51 Eridani[4]. Since 51 Eridani is located 96 ly ($9.0 \cdot 10^{17}$ m) away from Earth, the angular separation between the two objects is just

$$\theta = \frac{1.9 \cdot 10^{12}}{9.0 \cdot 10^{17}} = 2 \cdot 10^{-6} \text{ rad} = 0.44 \text{ arcsec}.$$

51 Eridani b was detected by a $D = 8.1$ m wide telescope[13] in near infrared light with wavelengths on the order of $\lambda \sim 1 \mu\text{m}$. The ratio between the two,

$$\lambda/D = \frac{1 \mu\text{m}}{8.1 \text{ m}} = 1.2 \cdot 10^{-7} = 0.025 \text{ arcsec},$$

is a close approximation of the angular resolution of the telescope. In other words, the planet was detected at an angular separation of about $18\lambda/D$, which is considered relatively close.

At such small angular separations, the wave-like nature of the light plays an important role[14]. A star 96 light years away is, for all purposes, a point source, but when observed through a finite diameter circular aperture, it creates a diffraction pattern (similar to the one shown on figure 1.2, top left), also known as the point spread function (PSF). Although at $18\lambda/D$ this PSF is 10^5 dimmer than at its peak, 51 Eridani b would still not be detectable by a conventional telescope as it is about 10^6 times dimmer than its host. Put differently, a circular aperture $D = 8.1$ m wide telescope alone, doesn't provide the 10^6 contrast necessary to detect a planet at 0.44 arcsec separation (in infrared). Indeed, the first directly observed exoplanets were hot self-luminous giants located far away from the star with contrasts between 10^2 and 10^4 [15, 16].




Figure 1.2: Point Spread Functions (on a logarithmic scale) of a star (top) and a planet (bottom). The diffraction pattern of a star observed through a circular aperture is many orders of magnitude brighter than that of a planet at $3.5 \lambda/D$ (left). A passive shaped masked at the pupil plane reduces the brightness in some regions of the image significantly, but the remaining speckles are still many times brighter than the planet (middle). By modifying the shapes of the DM upstream of the pupil plane, the intensity of the speckles is reduced further to create regions of very high contrast – the dark hole (right). *The PSF were generated by a numerical model of the shaped pupil coronagraph at Princeton's High Contrast Imaging Laboratory*[17, 18].

Coronagraphs block some of the light from the star (an on-axis source), effectively reducing the intensity of its PSF without significantly affecting the off-axis PSF, i.e., without significantly reducing the intensity of the planet. They allowed the discoveries of slightly cooler giants located closer to the star with contrast of about 10^6 , e.g. 51 Eridani b. However, Earth-like planets orbiting close to their hosts are expected to have contrasts on the order of 10^{10} [19], far above what has been achieved to date outside of a lab.

The major limiting factor for Earth-based observatories is atmospheric turbulence: spatial fluctuations in the density of the air above the telescope which distort the incoming light (wavefront) and the delicate diffraction pattern behind the coronagraph. Modern telescopes employ adaptive optics (AO[20]) – fast deformable mirrors which correct the wavefront, and are able to reduce the residual wavefront errors to a ~ 100 nm level[21]. Still, the finite number of photons arriving from the star and the chaotic nature of turbulence, limit the theoretical effectiveness of AO systems to a contrast of about 10^6 in the visible light[22, 23].

It is important to note that the above limit applies to the contrast during the observations phase: the intensity ratio between the planets and the residual starlight, also called “speckles” (or, in the context of ground based coronagraphy, “atmospheric halo”). The threshold of the dimmest detectable planets and the signal-to-noise ratio (S/N) depend additionally on the length of the observation and the image post-processing strategy[24, 25]. In other words, planets dimmer than the speckles can, in theory, be detected and characterized, given enough exposures and some algorithmic means to differentiate their signal from the time-varying speckles. Assuming a “perfect knowledge” of the speckles and the probabilistic nature of photon arrival rates (known as shot-noise), the time to achieve a given S/N is inversely proportional to the contrast[26],

$$\text{observation time} \propto \frac{(\text{desired S/N})^2}{\text{contrast}}.$$

One way to avoid the atmospheric limitation and directly image exoplanets within a reasonable time frame, is to put a coronagraph in space. One such mission is the upcoming Wide-Field Infra-Red Survey Telescope (WFIRST) Coronagraph Instrument (CGI)[27] which is planned to be able to detect Jupiter-like planets. Future telescopes such as the Large UV/Optical/IR Surveyor (LUVOIR)[28] and the Habitable Exoplanet Imaging Mission (HabEx)[29] are projected to detect tens of Earth-like exoplanets[30].

To achieve a contrast of 10^8 at around $3 - 6 \lambda/D$ the WFIRST-CGI will contain a variety of passive and active optical components. Besides the 2.4 m primary mirror, it will employ several optimized optical masks[31], fast steering mirrors[32] (the FSM reject “low-order” wavefront errors such as line-of-sight jitter[33]) and slower deformable mirrors (DM). The FSM and the DM will be positioned upstream of the so called “pupil plane” after which most of the light is blocked by shaped masks.

As illustrated in Fig. 1.2, actively controlling the DM is crucial to achieve contrasts above the passive 10^4 level[34] in a high-contrast region of the image commonly referred to as the “dark hole”. Although yet to be tested in space[35], the algorithm for creating dark holes, the Electric Field Conjugation[36], was shown to reach contrasts above 10^9 in a lab setting[37].

This thesis, however, will mostly focus on the tens of hours long observation phase that follows the creation of the dark hole. The numerical simulations of WFIRST-CGI in Sec. 2.3 show that the dark hole is extremely sensitive to accumulation of wavefront errors: a non-evenly distributed displacement of the primary mirror by just 1 nm root-mean-square (RMS), can halve the contrast. In fact, low sensitivities of the wavefront errors to thermal and mechanical stresses are one of the major design requirements for WFIRST-CGI[38, 39]. For example, the error budget for the quadrafoil error (a higher error mode not attenuated by the FSM) is on the order of just 0.1 nm[40] - the diameter of a Hydrogen atom.

1.3 Organization and Contributions

The stability requirements for future space based telescopes are expected to increase even further[41]. It therefore becomes clear that the high contrast in the dark hole has to be actively maintained.

Chapter 2 presents an algorithm for dark hole maintenance. It belongs to the field of Focal Plane Wavefront Sensing and Control since it continuously estimates the speckle field in the dark hole by measuring its intensity in the focal (image) plane. This estimation task proves to be challenging due to the small amount of light available at high contrasts, although it avoids non-common path errors (i.e., discrepancies between the speckles and their estimates based on measurements outside of the dark hole). This chapter is based on

- L. Pogorelyuk and N. J. 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, 2019.

Chapter 3 discusses the potential benefits of reduced-order estimation for systems coupled via low-dimensional stochastic input. Specifically, it introduces the computationally efficient “block-diagonal filter” which exploits the large number of measurements and the “smooth” nature of the speckles in the dark hole, potentially relaxing its stability requirements even further. This chapter is based on

- L. Pogorelyuk, C. W. Rowley, and N. J. Kasdin. An efficient approximation of the Kalman filter for multiple systems coupled via low-dimensional stochastic input. *submitted*.

Chapter 4 presents a formulation of the “post-processing” problem that incorporates the history of the DM actuations throughout the observation. These actuations introduce the so called “phase diversity”[42] which can help detect planets significantly dimmer than the speckles. Moreover, the probabilistic framework provided

in this chapter allows incorporating various constraints on the electric field of the speckles, e.g. their spatial and temporal “smoothness”, to help distinguish them from other sources including planets. It is based on

- L. Pogorelyuk, N. J. Kasdin, and C. W. Rowley. Reduced order estimation of the speckle electric field history for space-based coronagraphs. *The Astrophysical Journal*, 881(2):126, 2019.

Chapter 5 discusses future potential application of the above “smoothness” notion. In particular, in the context of “learning” equations of motion from videos (an ongoing work with S. E. Otto and Dr. C. W. Rowley), it may allow extracting physically meaningful coordinates from sequential data.

Chapters 2 through 4 also incorporate results from the following conference proceeding and article:

- L. Pogorelyuk and N. J. Kasdin. Maintaining a dark hole in a high contrast coronagraph and the effects of speckles drift on contrast and post processing factor. In *Techniques and Instrumentation for Detection of Exoplanets IX*, volume 11117, pages 397 – 403. International Society for Optics and Photonics, SPIE, 2019.
- L. Pogorelyuk, L. Pueyo, and N. J. Kasdin. On the effects of pointing jitter, actuators drift, telescope rolls and broadband detectors in dark hole maintenance and electric field model reduction. *submitted*.

The next section introduces notations, physical quantities and assumptions that are used throughout this thesis.

1.4 Achieving High Contrast

Before a hybrid coronagraph can begin its observations, it has to create a dark hole – a region of the image where the speckles are very dim, i.e., the contrast is very high. One of the several[43] methods that achieve this is the Electric Field Conjugation (EFC[36]) which is described below. It formulates the problem in terms of intensities and electric fields, which will be denoted by I and E respectively. For the purposes of this thesis, intensities and electric fields can be thought of merely as mathematical constructs which describe the *probability* of detecting photons and the way in which this probability is affected by deformable mirrors and telescope instabilities.

Each individual photon has a probability of reaching a particular region of the image depending on its energy (or wavelength) and direction of arrival (DOA). In Fourier optics[44], this probability is given by the squared magnitude of the complex electric field associated with a particular wavelength and DOA. The more photons are likely to arrive from a particular direction (with a particular wavelength), the higher the squared magnitude of the corresponding electric field is and the brighter the source of the light in this direction (and wavelength) is. The sum over the squared field magnitudes of all the sources – the intensity at the image plane, is proportional to the mean rate of photons arrival at the detectors.

The necessity of representing the probabilities as squared magnitudes of a complex fields becomes evident when analyzing the effects of various optical elements. In particular, the first order impact of a given DM perturbation on some image intensity is not uniquely defined by just those perturbation and intensity. Instead, the linearization needs to be carried out in terms of the electric field, as shown below.

1.4.1 A Linear Coronagraph Model

In a simplified model of a coronagraph, the dark hole is a collection of n pixels enumerated by i . At any given frame (or image) k , the detector at pixel i makes a measurement $y_{i,k}$. This measurement is assumed to follow a known distribution parameterized by an unknown intensity, $I_{i,k}$. In the absence of detector “read noise”, $y_{i,k}$ is Poisson distributed[45],

$$y_{i,k} \sim \text{Poiss}(\alpha \cdot I_{i,k}), \quad (1.1)$$

where α is an instrument-dependent constant which describes the probability of detecting a certain number of photons in a given period of time. Without loss of generality, it is possible to scale the physical quantities of the system such that $\alpha = 1$.

As mentioned earlier, the intensity, $I_{i,k}$, is a sum (also known as an “incoherent sum”) of scaled photon arrival probabilities for various sources. For simplicity, these sources will also be assigned intensities and split into three categories,

$$I_{i,k} = I_{i,k}^S + I_{i,k}^P + I_{i,k}^D, \quad (1.2)$$

where $I_{i,k}^S$ is the intensity of the speckles (residual starlight), $I_{i,k}^P$ includes all distant sources other than the star (planets, debris disks, zodiacal dust, etc.), and $I_{i,k}^D$ represents the dark current noise – a source of Poisson distributed thermal electrons internal to the telescope.

The high contrast is achieved by modifying the diffraction pattern of the speckles, I^S , in the dark hole. Below, the star will be considered as a monochromatic point source, i.e., having a single DOA and a single wavelength. Its diffraction pattern therefore corresponds to a single electric field, E (see Sec. 2.2 for a broadband

formulation),

$$I_{i,k}^S = |E_{i,k}|^2 \in \mathbb{R}. \quad (1.3)$$

This diffraction pattern depends on the configuration of various optical elements: deformable mirrors whose a control inputs, $\mathbf{u}_k^{DM} \in \mathbb{R}^a$, are assumed to be known; passive elements that vary slowly between time frames, k ; and FSM whose effects will be ignored until Sec. 4.1 (where it will be shown that they can be attributed to the $I_{i,k}^D$ term). This dependency will be expressed as

$$E_{i,k} = E_i(k, \mathbf{u}_k^{DM}) \in \mathbb{C}.$$

The linear model of the coronagraph describes the effects of small DM perturbations, \mathbf{u} , about some nominal DM setting $\mathbf{u}^{DM} \approx \mathbf{u}_k^{DM}$. To first order, this can be expressed as

$$E_i(k, \mathbf{u}_k^{DM} + \mathbf{u}) \approx E_i(k, \mathbf{u}_k^{DM}) + \frac{\partial E_i}{\partial \mathbf{u}^{DM}} \mathbf{u}, \quad (1.4)$$

which neglects the non-linearities in E_i and variations in the Jacobian, $\mathbf{g}_i^U = \left(\frac{\partial E_i}{\partial \mathbf{u}^{DM}} \right)^T \in \mathbb{C}^a$, as a function of k and \mathbf{u}_k^{DM} .

It is convenient to denote the electric field as a vector over the reals,

$$\mathbf{E} = \begin{bmatrix} \text{Re}\{E_1\} \\ \text{Im}\{E_1\} \\ \text{Re}\{E_2\} \\ \text{Im}\{E_2\} \\ \vdots \end{bmatrix} \in \mathbb{R}^{2n},$$

and the Jacobian (across all pixels) as

$$G^U \approx \frac{\partial \mathbf{E}}{\partial \mathbf{u}^{DM}} = \begin{bmatrix} (\text{Re}\{\mathbf{g}_1^U\})^T \\ (\text{Im}\{\mathbf{g}_1^U\})^T \\ (\text{Re}\{\mathbf{g}_2^U\})^T \\ (\text{Im}\{\mathbf{g}_2^U\})^T \\ \vdots \end{bmatrix} \in \mathbb{R}^{2n \times a}, \quad (1.5)$$

where n is the number of pixels in the dark hole and a is the number of DM actuators (possibly across multiple mirrors). The linear approximation (1.4) then becomes

$$\mathbf{E}(k, \mathbf{u}_k^{DM} + \mathbf{u}) \approx \mathbf{E}(k, \mathbf{u}_k^{DM}) + G^U \mathbf{u}, \quad (1.6)$$

and is the basis for the EFC algorithm below.

1.4.2 Electric Field Conjugation (EFC)

The process of creating the dark hole is relatively short and the magnitudes of the electric field are large, hence the fluctuations of the passive optical elements in time (k) are typically neglected during this stage. Put differently, the goal of this process is to estimate and reduce the magnitude of electric field as a function of just the controls, $\|\mathbf{E}(0, \mathbf{u}_k^{DM})\|$.

The estimation is done via pair-probing[46]. A set of images is taken with some predetermined probes, $\{\pm \mathbf{u}_j\}_{j=1}^r$, added on top of the controls \mathbf{u}_l^{DM} (here l denotes the iteration of the EFC algorithm as opposed to the time k). Since the mean of the Poisson distribution in (1.1) is $\alpha I_{i,k} = I_{i,k}$, it follows from (1.2)-(1.4) that the measurements, $y_{i,\pm j,l}$, approximately satisfy

$$y_{i,\pm j,l} \approx |E_i(0, \mathbf{u}_l^{DM}) \pm \mathbf{g}_i^U \cdot \mathbf{u}_j|^2 + I_i^P + I_i^D, \quad (1.7)$$

where $y_{i,\pm j,l}$ is the number photons measured at pixel i when the probe $\pm \mathbf{u}_j$ was added to the control during the l th iteration of the EFC algorithm (here, I_i^P and I_i^D lost their time index, k , because they are assumed to be either very small or constant or both).

Although E_i are unknown, subtracting the measured intensities for each pair of probes gives r equations in terms of the electric field estimates, \hat{E}_i ,

$$y_{i,+j,l} - y_{i,-j,l} = 4\text{Re} \left\{ \hat{E}_i(0, \mathbf{u}_l^{DM}) \right\} \text{Re} \left\{ \mathbf{g}_i^U \cdot \mathbf{u}_j \right\} + 4\text{Im} \left\{ \hat{E}_i(0, \mathbf{u}_l^{DM}) \right\} \text{Im} \left\{ \mathbf{g}_i^U \cdot \mathbf{u}_j \right\}.$$

These equations, put together for all probes, can be rearranged as one overdetermined linear equation,

$$\begin{bmatrix} 4\text{Re} \left\{ \mathbf{g}_i^U \cdot \mathbf{u}_1 \right\} & 4\text{Im} \left\{ \mathbf{g}_i^U \cdot \mathbf{u}_1 \right\} \\ 4\text{Re} \left\{ \mathbf{g}_i^U \cdot \mathbf{u}_2 \right\} & 4\text{Im} \left\{ \mathbf{g}_i^U \cdot \mathbf{u}_2 \right\} \\ \vdots & \vdots \end{bmatrix} \begin{bmatrix} \text{Re} \left\{ \hat{E}_i(0, \mathbf{u}_l^{DM}) \right\} \\ \text{Im} \left\{ \hat{E}_i(0, \mathbf{u}_l^{DM}) \right\} \end{bmatrix} = \begin{bmatrix} y_{i,+1,l} - y_{i,-1,l} \\ y_{i,+2,l} - y_{i,-2,l} \\ \vdots \end{bmatrix}, \quad (1.8)$$

and solved (in least-squares sense) for $\hat{E}_i(0, \mathbf{u}_l^{DM})$.

The EFC control law aims at reducing $\mathbf{E}(0, \mathbf{u}_l^{DM})$ based on the pair-probing estimate, $\hat{\mathbf{E}}(0, \mathbf{u}_l^{DM})$, and is given by,

$$\mathbf{u}_{l+1}^{DM} - \mathbf{u}_l^{DM} = -K_{EFC,l} \hat{\mathbf{E}}(0, \mathbf{u}_l^{DM}) \quad (1.9)$$

where \mathbf{u}_{l+1}^{DM} is the nominal control chosen for the next iteration and $K_{EFC,l}$ is the EFC gain. The gain,

$$K_{EFC,l} = \left((G^U)^T (G^U) + \gamma_l I_{a \times a} \right)^{-1} (G^U)^T \quad (1.10)$$

takes the energy of the controls into account by carefully choosing the γ_l parameter (which multiplies the identity matrix $I_{a \times a}$) at each iteration[18] (although for simulation purposes, γ_l can be chosen by trial and error).

The crucial assumption that allows disregarding wavefront instabilities during the creation of the dark hole is that the discrete number of detected photons is large enough so that (1.7) holds and therefore the electric field can be accurately estimated. However, when observing dim stars, the time scales required to collect a large number of photons are long enough for wavefront instabilities to have a significant impact. A method which takes those instabilities and the probabilistic nature of photon arrival rates is presented in the next chapter.

Chapter 2

Dark Hole Maintenance

A typical WFIRST-CGI observation scenario consists of pointing at a relatively dim (~ 5 mag) star for tens of hours, during which it will detect a total of just $\sim 10^3$ photons per pixel[26]. At the same time, the wavefront will change by hundreds of picometers and the contrast will decrease significantly[48]. While the effects of time




Figure 2.1: The intensity of dark-hole speckles relative to the maximum intensity of the star (the inverse of the contrast) as it evolves during a simulated 30 hour observation (for details see Sec. 2.3). In the baseline WFIRST-CGI scenario[47] (dashed blue line), the DM are kept constant for 8 hours while wavefront errors accumulate and the contrast decreases. The telescope is then pointed at a reference star for 2 hours to re-create the dark hole. On the other hand, a closed-loop scenario (solid red line) allows continuously maintaining a fixed contrast, albeit at slightly lower level than achievable by pointing at bright reference star.

varying wavefront errors may be mitigated by periodically re-creating the dark hole, it can only be done by pointing at a bright reference star for several hours[47]. This is a consequence of the pair-probing procedure (Eq. (1.8)) which relies on a large ($\gg 1$) number photons per exposure. Otherwise, Poisson statistics, (1.1), dictate that the standard deviation of each measurement becomes as large as the measurement itself and the approximation in (1.7) breaks down.

Additionally, periodic changes in the orientation of the telescope might further accelerate wavefront variations by exciting structural and thermal modes, resulting in peaks of decreased contrast after each such maneuver[48]. Putting these effects aside, the contrast profile corresponding to a WFIRST-CGI observation strategy resembles a sawtooth pattern – slowly decreasing when pointed at a target star and collecting no scientific data during dark-hole recreation phases (dashed blue line in figure 2.1).

An alternative closed-loop approach[49] is presented below. While continuously observing the dim target, it maintains a fixed contrast at a slightly lower level than in a “perfect” dark hole but significantly higher than the average open-loop contrast (red line in figure 2.1). Unlike pair probing and EFC, this approach does not perform “batch estimation” since it cannot rely on just a few low-intensity measurements (ground-based approaches for maintaining high contrast based on a single intensity measurement[50, 51], can also be considered batch estimators and would similarly suffer from issues with low photon arrival rates in the context of space coronagraphs).

2.1 Non-linear Recursive Estimation

Recursive estimation allows merging all prior intensity measurements to get significantly more accurate estimates of the electric field that would otherwise be possible with a small number of low magnitude probes (such that they do not significantly increase the intensity of the speckles). Although each new measurement is very noisy

(S/N ~ 1), and past measurements become increasingly irrelevant, their information content can be appropriately scaled and combined in real time. In the contexts of linear dynamical systems this can be done optimally via the Kalman Filter (KF)[52, 53]. However, since the intensity measurements depend non-linearly on the electric field, the algorithm below uses the sub-optimal Extended Kalman Filter (EKF).

2.1.1 State Equations

To formulate the estimation problem, it is convenient to define the hypothetical “open-loop” electric field,

$$\mathbf{E}_k^{OL} = \mathbf{E}(k, \mathbf{u}_{DH}^{DM}),$$

where \mathbf{u}_{DH}^{DM} is the nominal dark hole DM setting (e.g. the last command computed by the dark hole creation procedure described in Sec. 1.4). The actual, “closed-loop”, electric field depends on the correction, \mathbf{u}_k , added to the nominal control at the k th step of the maintenance procedure. In light of (1.6) it is approximately given by,

$$\mathbf{E}(k, \mathbf{u}_{DH}^{DM} + \mathbf{u}_k) \approx \mathbf{E}_k^{OL} + G^U \mathbf{u}_k. \quad (2.1)$$

Note that the first order effects of the wavefront drift and the effects of control are independent, i.e., $\mathbf{E}_{k+1}^{OL} - \mathbf{E}_k^{OL}$ evolves on its own. A further simplifying assumption, to be somewhat relaxed in chapters 3 and 4, is that the electric field increments are uncorrelated in time, Normally distributed and independent between pixels,

$$\begin{bmatrix} \text{Re} \{ E_{i,k+1}^{OL} - E_{i,k}^{OL} \} \\ \text{Im} \{ E_{i,k+1}^{OL} - E_{i,k}^{OL} \} \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \Sigma_i \right), \quad (2.2)$$

with $\Sigma_i \in \mathbb{R}^{2 \times 2}$, $\forall i$ fixed and known. This completes the description of the evolution of the state of the system, \mathbf{E}^{OL} (which excludes the intensities incoherent with the speckles, I_i^P , I_i^D , that are assumed to be fixed).

The influence of the controls shows up in the measurement equation,

$$y_{i,k} \sim \text{Pois} (I_{i,k}) \cong \text{Pois} \left(|E_{i,k}^{OL} + \mathbf{g}_i^U \cdot \mathbf{u}_k|^2 + I_i^P + I_i^D \right), \quad (2.3)$$

which, for the purposes of formulating an EKF, is approximated by a Normal distribution with the same mean and variance,

$$y_{i,k} \sim \mathcal{N} (I_{i,k}, I_{i,k}). \quad (2.4)$$

2.1.2 Estimation with an Extended Kalman Filter (EKF)¹

In this section, the EKF will be formulated for each pixel separately and the estimates $\hat{E}_{i,k|k}^{OL}$, $\forall i$ will be computed in parallel. The Kalman Filter was also experimented with in the context of dark hole creation with pair probing[54], but turned out to be unstable when applied directly to the above state equations (or a variation thereof, see “zero probes” EKF in [55]).

Below is a standard formulation of the EKF (e.g. [53]), assuming noise sources with known fixed intensity, $I_{i,k}^D = I_i^D$, and negligible planet signal $I_{i,k}^P \ll I_{i,k}^S$ (the extension to an EKF which accounts for varying $I_i^D + I_{i,k}^P$ is straightforward but cumbersome[49]). The state equation (2.2) gives difference equations for the predic-

¹ A code example is available at
<https://github.com/leonidprinceton/DHMaintenanceExample>

tion step of the filter,

$$\begin{bmatrix} \text{Re} \left\{ \hat{E}_{i,k|k-1}^{OL} \right\} \\ \text{Im} \left\{ \hat{E}_{i,k|k-1}^{OL} \right\} \end{bmatrix} = \begin{bmatrix} \text{Re} \left\{ \hat{E}_{i,k-1|k-1}^{OL} \right\} \\ \text{Im} \left\{ \hat{E}_{i,k-1|k-1}^{OL} \right\} \end{bmatrix}, \quad (2.5)$$

$$P_{i,k|k-1} = P_{i,k-1|k-1} + \Sigma_i, \quad (2.6)$$

where $\hat{E}_{i,k-1|k-1}^{OL}$ is the estimate of the open-loop field, $E_{i,k-1}^{OL}$, at frame $k-1$ and $P_{i,k-1|k-1}$ is the (approximation of the) covariance of the error $\hat{E}_{i,k-1|k-1}^{OL} - E_{i,k-1}^{OL}$.

Incorporating the k th measurement is done in the update step,

$$\begin{bmatrix} \text{Re} \left\{ \hat{E}_{i,k|k}^{OL} \right\} \\ \text{Im} \left\{ \hat{E}_{i,k|k}^{OL} \right\} \end{bmatrix} = \begin{bmatrix} \text{Re} \left\{ \hat{E}_{i,k|k-1}^{OL} \right\} \\ \text{Im} \left\{ \hat{E}_{i,k|k-1}^{OL} \right\} \end{bmatrix} + K_{i,k}(y_{i,k} - \hat{y}_{i,k}), \quad (2.7)$$

which follows from the non-linear measurement equation (2.4) with,

$$\hat{y}_{i,k} = \left| \hat{E}_{i,k|k-1}^{OL} + \mathbf{g}_i^U \cdot \mathbf{u}_k \right|^2 + I_i^D, \quad (2.8)$$

$$K_{i,k} = \frac{1}{s_{i,k}} P_{i,k|k-1} H_{i,k}^T, \quad (2.9)$$

and

$$H_{i,k}^T = \frac{\partial \hat{y}_{i,k}}{\partial \begin{bmatrix} \text{Re} \left\{ \hat{E}_{i,k|k-1}^{OL} \right\} & \text{Im} \left\{ \hat{E}_{i,k|k-1}^{OL} \right\} \end{bmatrix}} = \begin{bmatrix} 2\text{Re} \left\{ \hat{E}_{i,k|k-1}^{OL} + \mathbf{g}_i^U \cdot \mathbf{u}_k \right\} \\ 2\text{Im} \left\{ \hat{E}_{i,k|k-1}^{OL} + \mathbf{g}_i^U \cdot \mathbf{u}_k \right\} \end{bmatrix}, \quad (2.10)$$

$$s_{i,k} = H_{i,k} P_{i,k|k-1} H_{i,k}^T + \hat{y}_{i,k}. \quad (2.11)$$

It is also necessary to specify the initial distribution of the electric field. For example,

$$\begin{bmatrix} \text{Re} \left\{ E_{i,0}^{OL} \right\} \\ \text{Im} \left\{ E_{i,0}^{OL} \right\} \end{bmatrix} \sim \mathcal{N}(\mu_{i,0}, \Sigma_{i,0}), \quad (2.12)$$

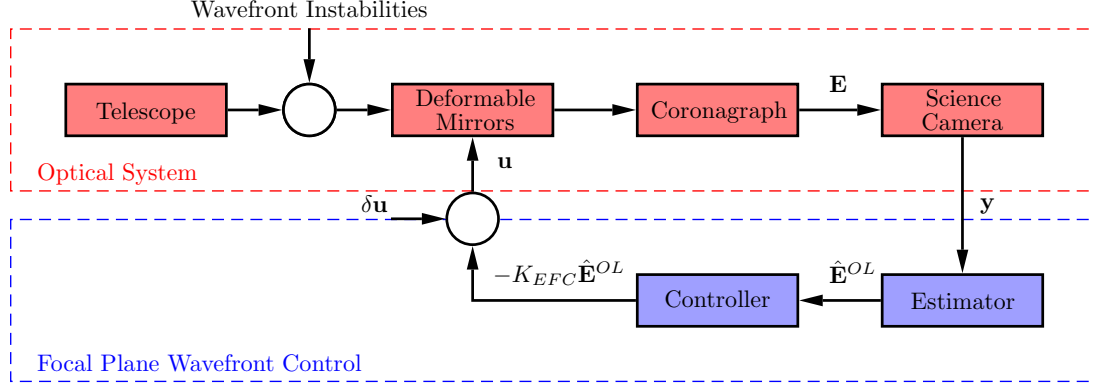


Figure 2.2: Dark hole maintenance loop. Based on the measurements, \mathbf{y} , the EKF computes an estimate of the hypothetical open-loop electric field, $\hat{\mathbf{E}}^{OL}$, used to choose a DM correction, \mathbf{u} , to reduce the magnitude of the actual speckle field, \mathbf{E} , which varies in time due to wavefront instabilities. The correction, \mathbf{u} , combines an EFC control law with small random actuations (dithering) to ensure the stability of the EKF. *adapted from [49]*

yields a natural initialization of the EKF with $\hat{E}_{i,0|0}^{OL} = \begin{bmatrix} 1 & \sqrt{-1} \end{bmatrix} \cdot \mu_{i,0}$ and $P_{i,0|0} = \Sigma_{i,0}$. For numerical purposes, it is convenient to choose $\Sigma_{i,0}$ to be a multiple of Σ_i and $\hat{E}_{i,0}^{OL}$ to be either 0 or the last estimate from the preceding observation (or dark hole creation).

2.2 Closing the Loop and Dithering

Equipped with an EKF estimate of the open-loop electric field, $\hat{\mathbf{E}}_{k-1|k-1}^{OL}$, at frame $k-1$, the goal is to choose a DM command correction for the next frame, \mathbf{u}_k , (the total DM command would then be $\mathbf{u}_k^{DM} = \mathbf{u}_{DH}^{DM} + \mathbf{u}_k$). As shown in Fig. 2.2, the correction is meant to reduce the magnitude of the closed-loop electric field of the speckles. However, it is also supposed to ensure that the EKF remains stable which is not guaranteed in this non-linear setting.

2.2.1 Dark Hole Maintenance Control Law

Closing the loop by simply applying an EFC correction, as during the dark hole creation in Eq. (1.9), causes the EKF to diverge due to lack of phase diversity. Indeed, if $H_{i,k}$ in Eq. (2.10) remains close to constant, the EKF is close to not being observable: the pair $(I_{2 \times 2}, H)$ is unobservable for any fixed $H \in \mathbb{R}^{2 \times 1}$ [56], where $I_{2 \times 2}$ stands for the identity transformation describing the evolution of the electric field in (2.2). Fortunately, this issue can be resolved by dithering the DM, i.e., adding small random actuations to \mathbf{u}_k which directly modify $H_{i,k}$ between iterations.

A closed loop control law that is capable of maintaining high contrast while introducing enough phase diversity is therefore given by,

$$\mathbf{u}_k = -K_{EFC} \hat{\mathbf{E}}_{k-1|k-1}^{OL} + \delta \mathbf{u}_k, \quad (2.13)$$

where $K_{EFC} = K_{EFC,l}$ is the EFC gain used in the last step (last l) of the dark hole creation, Eq. (1.10). The dither, $\delta \mathbf{u}_k$, is randomly sampled for example via

$$\delta \mathbf{u}_k \sim \mathcal{N}(\mathbf{0}, \sigma_u^2 I_{a \times a}), \quad (2.14)$$

where $\mathbf{0} \in \mathbb{R}^a$ is the zero vector and σ_u is the dither magnitude – a parameter which can be optimized (see, for example, Fig. 2.4 in the next section). The DM dither therefore plays a similar role to the DM probes but is slightly more general.

2.2.2 Recalibration

The main difference between creating the dark hole and maintaining it is that the latter keeps track of the open-loop electric field rather than the actual, closed-loop electric field in the dark hole. This approach seems to give a more stable EKF since

the linearization and model errors appear in the measurement equation (2.3) and only indirectly affect the estimate.

Eventually, though, the electric fields drifts enough so that the linearization in (1.4), begins to break down making the estimates increasingly biased and the control less effective. A possible solution is to perform one “recalibration” step every N maintenance steps. Formally, this is done by updating the dark hole setting, \mathbf{u}_{DH}^{DM} , and the open-loop estimate, $\hat{\mathbf{E}}_{k|k}^{OL}$, when $k \bmod N = N - 1$ via

$$\begin{aligned}\mathbf{u}_{DH}^{DM} &\leftarrow \mathbf{u}_{DH}^{DM} - K_{EFC} \hat{\mathbf{E}}_{k|k}^{OL}, \\ \hat{\mathbf{E}}_{k|k}^{OL} &\leftarrow \hat{\mathbf{E}}_{k|k}^{OL} - G^U K_{EFC} \hat{\mathbf{E}}_{k|k}^{OL},\end{aligned}$$

where “ \leftarrow ” denotes assignment. In a wide range of drift magnitudes simulated for the numerical results section below, any choice of $10 < N < 100$ prevented the nonlinearities from destabilizing the EKF or diminishing the performance of the EFC.

2.2.3 Broadband Light

The derivations of the EFC and EKF assumed that the star is a monochromatic source, hence the vector \mathbf{E} corresponded to values of just one electric field at various dark-hole pixels. This formulation can be readily extended[57] to broadband light assuming that it is well approximated by several monochromatic intensities generated by their corresponding electric fields. In other words, the intensity of the speckles at the i th pixel is approximately given by the incoherent sum

$$I_{i,k}^S \approx \sum_{\lambda \in \Lambda_i} |E_{\lambda,i,k}|^2, \quad (2.15)$$

where $\Lambda_i = \{\lambda_{i,1}, \lambda_{i,2}, \dots\}$ is some spectrum discretization. Note that this discretization is parameterized by the pixel index, since each pixel may “receive” a different part of the spectrum.

In vector form, (2.15) can be written as,

$$I_{i,k}^S \approx \mathbf{1}_i \cdot (\mathbf{E}_{i,k} \circ \mathbf{E}_{i,k}) = \mathbf{1}_i \cdot \mathbf{E}_{i,k}^{\circ 2} \quad (2.16)$$

where $\mathbf{1}_i$ is a vector of ones, \circ stands for the elementwise (Hadamard) product and

$$\mathbf{E}_{i,k} = \begin{bmatrix} \text{Re} \{E_{\lambda_{i,1},i,k}\} \\ \text{Im} \{E_{\lambda_{i,1},i,k}\} \\ \text{Re} \{E_{\lambda_{i,2},i,k}\} \\ \text{Im} \{E_{\lambda_{i,2},i,k}\} \\ \vdots \end{bmatrix}.$$

The remainder of the derivation of the broadband EKF follows the procedure in sections 1.4 and 2.1.

The open-loop electric field(s) defined in (2.1) becomes

$$\mathbf{E}_{i,k}^{OL} \equiv \mathbf{E}_{i,k} - G_i^U \mathbf{u}_k,$$

where the Jacobian $G_i^U = \partial \mathbf{E}_i / \partial \mathbf{u}$ is the broadband extension of \mathbf{g}_i^U . The drift in (2.2) retains the same form in a higher dimension,

$$\mathbf{E}_{i,k+1}^{OL} - \mathbf{E}_{i,k}^{OL} \sim \mathcal{N}(\mathbf{0}_i, \Sigma_i),$$

and so do most of the EKF equations, (2.5)-(2.11),

$$\begin{aligned}
\hat{\mathbf{E}}_{i,k|k-1}^{OL} &= \hat{\mathbf{E}}_{i,k-1|k-1}^{OL}, \\
P_{i,k|k-1} &= P_{i,k-1|k-1} + \Sigma_i, \\
\hat{\mathbf{E}}_{i,k|k}^{OL} &= \hat{\mathbf{E}}_{i,k|k-1}^{OL} + K_{i,k} (y_{i,k} - \hat{y}_{i,k}), \\
K_{i,k} &= \frac{1}{s_{i,k}} P_{i,k|k-1} H_{i,k}^T, \\
s_{i,k} &= H_{i,k} P_{i,k|k-1} H_{i,k}^T + \hat{y}_{i,k}.
\end{aligned}$$

A minor subtlety arises due to the incoherent addition of intensities in (2.16). The intensity estimate becomes,

$$\hat{y}_{i,k} = \mathbf{1}_i \cdot \left(\hat{\mathbf{E}}_{i,k}^{OL} + G_i^U \mathbf{u}_k \right)^{\circ 2} + I_i^D$$

and therefore,

$$H_{i,k}^T = \frac{\partial \hat{y}_{i,k}}{\partial \left(\hat{\mathbf{E}}_{i,k}^{OL} \right)^T} = 2 \left(\hat{\mathbf{E}}_{i,k}^{OL} + G_i^U \mathbf{u}_k \right).$$

In the numerical simulations below, each detector corresponded to a single wavelength, $|\Lambda_i| = 1$, but multiple wavelengths were simulated in parallel. As a consequence, the estimator used monochromatic EKF equations (2.5)-(2.11), while the controller used a broadband estimate since it technically contained monochromatic estimates at multiple wavelengths.

2.3 Numerical Results

The numerical results in chapters 2 to 4 were simulated using the Fast Linearized Coronagraph Optimizer (FALCO)[58], and in particular its model of the WFIRST hybrid Lyot coronagraph. All simulations began after a dark hole in a ring between 3 and $9 \lambda/D$ was created (by FALCO) with a contrast of $4.2 \cdot 10^{-10}$. Figure 2.3 shows the




Figure 2.3: WFIRST-CGI point spread functions of an on-axis star and an off-axis planet in various wavelengths (the intensities are in photons per frame). Top: lowest speckle intensity - immediately after creating the dark hole. Middle: speckle intensity after 10 hours of open-loop observation. Bottom: the planet PSF whose main lobe is visible at $6 \lambda/D$. The variations of its shape with wavelength and DM commands are unnoticeable. *simulated by FALCO*[58]

initial dark hole PSF of the star and the planet in five different wavelengths uniformly spaced between 546 nm and 602 nm. The measurements were derived from five sets of simulated monochromatic intensities, one for each of the above wavelengths.

In a hypothetical observation scenario the speckles initially had an average intensity of $1.2 \frac{\text{photons}}{\text{frame}}$ across $n = 2608 \times 5$ pixels, and the frames were chosen to be $\Delta t = 100$ sec long to make the total observation time consistent with WFIRST-CGI projections[26]. The drift of the speckles was introduced through a random walk of

the first 21 Zernike coefficients of the wavefront in the pupil plane,

$$z_{k+1}^{p,p-2j} - z_k^{p,p-2j} \sim \mathcal{N} \left(0, \left(\frac{\sigma_d}{(p+1)^2 \cdot \lambda} \right)^2 \Delta t \right), \quad 0 \leq j \leq p \quad (2.17)$$

where p is the order of the polynomial, j is its azimuthal degree and σ_d is the magnitude of the drift in units of $\frac{\text{nm}}{\sqrt{\text{hr}}}$.

The model incorporated two deformable mirror with 48×48 actuators each ($a = 4608$) and their Jacobian, G^U , was computed via straightforward numerical differentiation. Similarly, the effects of drift on the electric field, Σ_i in (2.2), were numerically estimated for the trajectories of $E_{i,k}$ given a sample path generated by (2.17). The DM dithering magnitude, σ_u in (2.14), was on the order of 10 mV where an input of 1 V to a particular actuator translated to a maximum of 1 nm change in DM surface shape at the location of that actuator.

In this section, the only source of noise other than the speckles, was dark current – a uniform flux of thermal electrons that is indistinguishable from a uniform intensity source in photon counting mode[59]. As such, it was also Poisson distributed as part of (2.3) with $I_i^D = 0.25 \frac{\text{photons}}{\text{frame}} \forall i$.

The contrast history of a sample scenario with $\sigma_d = 0.2 \frac{\text{nm}}{\sqrt{\text{hr}}}$ and $\sigma_u = 5 \text{ mV}$ is shown on figure 2.1. The open-loop contrast decreases by about a factor of 2 in the first eight hours and, in the baseline WFIRST-CGI scenario, the telescope would be pointed at a reference star for two hours to recreate the dark hole. The combination of the EKF and the closed-loop control law in Eq. (2.13), allows maintaining a constant contrast indefinitely. This closed-loop contrast, although higher than the average open-loop contrast, is lower than in the initial dark hole due to dithering. This loss of contrast is addressed next.

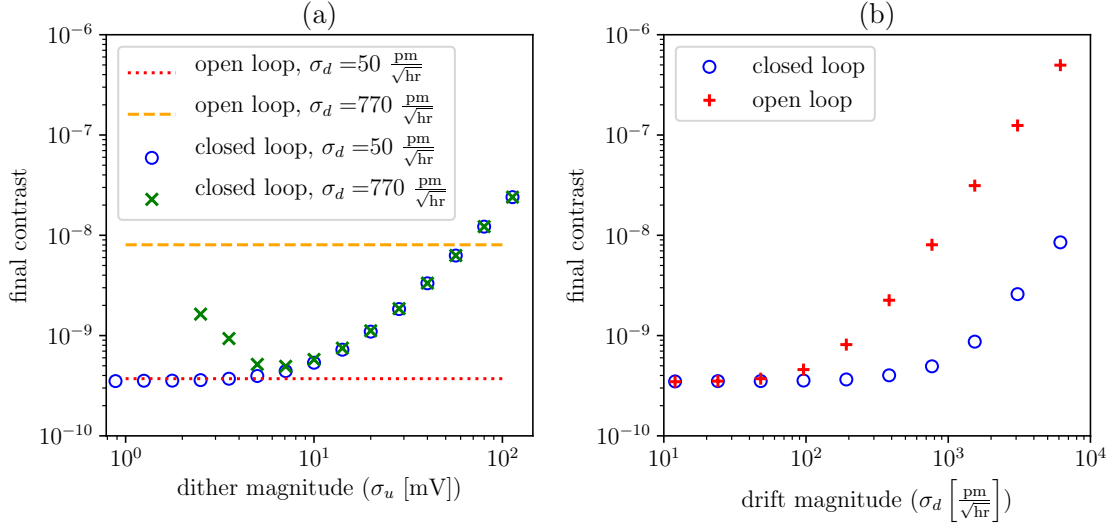


Figure 2.4: (a) The average dark hole (inverse) contrast during a 14.2 hour long observation as a function of DM dithering magnitude, σ_u , for two closed-loop scenarios: with low, $\sigma_d = 50 \frac{\text{pm}}{\sqrt{\text{hr}}}$, and high, $\sigma_d = 770 \frac{\text{pm}}{\sqrt{\text{hr}}}$, drift magnitudes. The open-loop contrast at the end of the observation is shown for comparison. The “optimal” σ_u correspond to the “best” contrast. (b) The average closed-loop (with optimal σ_u) and the final open-loop contrasts as a function of drift magnitude σ_d . *from [60]*

2.3.1 Effectiveness of Closed-Loop Dark Hole Maintenance

The phase diversity provided by the dithering of the DM (σ_u) is necessary for the EKF to keep up with the drift (σ_d). For σ_u/σ_d that are too small, the estimation error constantly increases, the loop opens and the contrast decreases as illustrated on the left of figure 2.4(a). On the other hand, the larger the dither, the more adversely it impacts the delicate speckles and the lower the steady-state contrast becomes (right of Fig. 2.4(a)). Intuitively, there exists an optimal value of σ_u for which the EKF remains accurate with minimal increase in speckle intensity.

Another source of process noise is the drift of the DM actuators[61] which is indistinguishable from electric field drift under the linearity assumptions, (1.4) and (2.2). However, the covariances Σ_i need to be increased accordingly to account for the adverse effects of the DM. Jitter[62], on the other hand, appears as a source of incoherent intensity (see discussion below (4.14) in Sec. 4.1) which varies slowly in

time. As such it can be assumed roughly known throughout the observation and attributed to I^D as an additional source of measurement noise.

Figure 2.4(b) presents the final open-loop contrast (at the end of the observation) and the closed-loop contrast (throughout the observation) after a hypothetical 14.2 hour long observation for various values of σ_d . The advantage of closing the loop can be interpreted as either delaying the onset of drift effects from $100 \frac{\text{pm}}{\sqrt{\text{hr}}}$ to $1 \frac{\text{nm}}{\sqrt{\text{hr}}}$ or maintaining a significantly higher contrast when the drift is large. Moreover, by considering a more elaborate model of the coronagraph, it is possible to further increase the accuracy of the EKF and potentially stabilize the dark hole with even less dither.

Chapter 3

Reduced-order Approximation of the Filtering Problem

In previous chapters there was a discrepancy in the treatment of the electric field between its estimation and its control. Both the pair probing, Eq. (1.8), and the EKF, Eqs. (2.5)-(2.11), were estimating the electric field at a single pixel at a time. On the other hand, the EFC law for the creation of the dark hole, Eq. (1.9), as well




Figure 3.1: The complex electric field in the WFIRST-CGI dark hole and its response to Zernike phase perturbations in the pupil plane. The uncertainties in the initial distribution of the electric field are high-dimensional while its increments may potentially be well approximated by a small number of drift modes. *simulated by FALCO*[58]

as during its maintenance, Eq. (2.13), relies on the estimate of the field at all pixels. Of course, the disturbances in the field are correlated between pixels both due to DM actuations and due to wavefront errors. The difference lies in the computation cost required to account for this correlation: the EFC gain is computed just once, while the full EKF gain would have to be computed at every iteration of the maintenance scheme (the flight computer on WFIRST[63], for example, will not be able to perform multiplications and inversions of matrices with millions of entries after every exposure).

In general, the time complexity of an n -dimensional Kalman Filter (KF) is $\mathcal{O}(n^3)$. However, by ignoring the correlation between subsystems (see (2.2) where Σ_i are defined for each pixel separately), the process noise covariance matrix for the whole system becomes block diagonal,

$$\Sigma = \begin{bmatrix} \Sigma_1 & 0 & \cdots & 0 \\ 0 & \Sigma_2 & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & \Sigma_n \end{bmatrix}. \quad (3.1)$$

Moreover, since the measurement noise is uncorrelated between pixels, its covariance is also block diagonal, and so is the initial error covariance (this is implicit in the choice of $P_{i,0|0}$). Therefore, the subsystems remain decoupled, the error covariances are always block diagonal, and the EKF is advanced in $\mathcal{O}(n)$ by effectively ignoring the $\mathcal{O}(n^2)$ zero terms on the off-diagonals of the “full” covariance matrices. This will be referred to as the “banded filter”[64]. While fast, the banded filter completely ignores the correlation between the fluctuations of the “smooth” speckles at separate pixels.

One approach to modeling the “smoothness” of the speckles stems from the assumption that most of the “energy” of wavefront drift is associated with a small

number of structural modes, r (e.g. the 21 Zernike modes in (2.17)). This can be expressed by splitting the open-loop electric field increments across *all* pixels into

$$\mathbf{E}_{k+1}^{OL} - \mathbf{E}_k^{OL} = G^V \Delta \mathbf{v}_k + \Delta \mathbf{w}_k, \quad (3.2)$$

$$\Delta \mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, V_k),$$

$$\Delta \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, W_k),$$

where $G^V \in \mathbb{R}^{2n \times r}$ and $V_k \in \mathbb{R}^{r \times r}$ are low rank, and $W_k \in \mathbb{R}^{2n \times 2n}$ is block diagonal. The purpose of W_k is to model the remainder of the process-noise (not captured by the low rank variations, V_k) as uncorrelated. The columns of G^V will be referred to as speckle modes, several of which are shown in Fig. 3.1, and their estimation will be discussed in chapter 4. The increments of the coefficients of the modes, $\Delta \mathbf{v}_k \in \mathbb{R}^r$, act as a “global” source of process noise whereas $\Delta \mathbf{w}_k \in \mathbb{R}^{2n}$ correspond to a “local” source with a much smaller impact than $G^V \Delta \mathbf{v}_k$.

Presumably, the coupling, $\Delta \mathbf{v}_k$, is low dimension compared to the number of measurements, $r \ll n$. Therefore, a KF formulated according to Eq. (3.2) will potentially estimate the coupling with high precision and eliminate this major source of uncertainty. Unfortunately, despite V_k being low rank and W_k being block diagonal, the covariance of the full KF is neither. Since advancing KF matrices in $\mathcal{O}(n^3)$ is prohibitively time consuming, a sub-optimal estimator would be necessary to utilize the convenient but incompatible numerical properties of the two noise sources.

In the literature, the low rank cross-correlation term, $\Delta \mathbf{v}_k$, is sometimes viewed as an “unknown input”[65] and can be “decoupled” and optimally estimated under certain assumptions[66]. Although these estimators are as computationally expensive as the full KF, they can be significantly “accelerated” by employing reduced order filters[67] and their extensions to systems with unknown inputs[68]. The faster performance of reduced order filters stems from the restriction of the estimates to a low

dimensional subspace of the original state space. This restriction, however, results in poor estimates of the speckle field since it is generated by an unrelated process of dark hole creation, has a high-dimensional drift component, $\Delta \mathbf{w}_k$, and therefore does not lie in any low dimensional subspace.

Another family of algorithms known as “sub-optimal schemes” comes from the area of atmosphere and ocean data assimilation[69, 70]. The multi-dimensional high-resolution models required to simulate large bodies of fluid, maintain an order of magnitudes more parameters than the number of pixels in a space telescope, hence the $\mathcal{O}(n^3)$ complexity of the KF remains prohibitive even for land-based supercomputers. In these schemes, the measurements are “assimilated” using sub-optimal gains based on increasingly sophisticated approximations of the state covariance matrix. The banded filter[64] simply discards all off-diagonal terms of the process-noise covariance matrix which, when used to estimate the speckles, results in mediocre accuracy (see Sec. 3.2). Asymptotic filters (e.g. [70]), use a precomputed covariance approximation, but still require $\mathcal{O}(n^3)$ operations to compute the gain for time-varying systems (i.e., when the matrices $H_{i,k}$ in Eq. 2.10 are not known ahead of time). Reduced order equivalents[71, 72] suffer from the above mentioned problem of projecting the state onto a low dimensional subspace – an approximation uncharacteristic of a realistic speckle distribution.

Several other methods employ a low-rank representation of the covariance matrix, but allow the estimates to remain in a high dimensional space. This can be achieved by propagating the state along the most dominant “error subspaces”[73, 74] or using low dimensional Monte Carlo approximations of the covariance (i.e., the Ensemble Kalman Filter[75]). In both cases the sub-optimal gains are based on low rank covariance matrices, hence the estimates are “updated” exclusively along low-dimensional error subspaces. The convergence of such methods would be too slow for speckle estimation purposes because the initial covariance matrix is high-dimensional, but the

number of measurements is relatively low due to the lengthy exposure times at high contrasts.

Other methods for large distributed systems where the EKF is prohibitively expensive include the Interlaced Kalman Filter[76] (used for robot localization), the Decoupled Kalman Filter[77] (training of neural networks) and the Fusion Kalman Filter[78] (for multiple sensors with dynamical states). These and other approaches, presume no cross-correlation between process noises of a large number of otherwise decoupled subsystems.

The next section describes a sub-optimal filter which allows for high-dimensional state updates in $\mathcal{O}(n)$ time complexity, while also accounting for the “nice” properties of the sources of the process noise. Sections 3.2 and 3.3 study the cost-efficiency of the filter and its analytical properties in some simple cases.

3.1 A Block-diagonal Approximation

To make the discussion more general, the states of the i th subsystem will be denoted as $\mathbf{x}_k^{(i)} \in \mathbb{R}^c$, which can stand for the electric field of just one pixel ($c = 2$ and $\mathbf{x}_{i,k} = \begin{bmatrix} \text{Re}\{E_{i,k}^{OL}\} & \text{Im}\{E_{i,k}^{OL}\} \end{bmatrix}^T$), the electric field and the incoherent intensity of just one pixel ($c = 3$), or of groups of multiple pixels ($c \geq 4$). The state of the full system composed of n subsystems will be denoted as,

$$\mathbf{x}_k = \begin{bmatrix} \mathbf{x}_{1,k} \\ \mathbf{x}_{2,k} \\ \vdots \end{bmatrix} \in \mathbb{R}^{cn},$$

and the measurement as,

$$\mathbf{y}_k = \begin{bmatrix} \mathbf{y}_{1,k} \\ \mathbf{y}_{2,k} \\ \vdots \end{bmatrix} \in \mathbb{R}^{dn},$$

for some d .

The most general linear state equation corresponding to Eq. (3.2) is

$$\mathbf{x}_{i,k} = F_{i,k} \mathbf{x}_{i,k-1} + G_i^V \Delta \mathbf{v}_k + \Delta \mathbf{w}_{i,k}, \quad \forall i,$$

with $G_i^V \in \mathbb{R}^{c \times r}$ and $\Delta \mathbf{w}_{i,k} \in \mathbb{R}^c$ being the i th row elements of G^V and $\Delta \mathbf{w}_k$ respectively. Alternatively,

$$\mathbf{x}_k = F_k \mathbf{x}_{k-1} + G^V \Delta \mathbf{v}_k + \Delta \mathbf{w}_k,$$

where $F_k \in \mathbb{R}^{cn \times cn}$ consists of $F_{i,k} \in \mathbb{R}^{c \times c}$ on its main diagonal,

$$F_k = \begin{bmatrix} F_{1,k} & & \\ & F_{2,k} & \\ & & \ddots \end{bmatrix}$$

(note that the random walk assumption in Eq. (2.2) translates into $F_{i,k} = I_{2 \times 2}$).

The error covariance, $P_{k|k-1}$, of the of the state estimate, $\hat{\mathbf{x}}_{k|k-1}$, of the full (Extended) Kalman Filter is based on the error covariance of the previous estimate, $P_{k-1|k-1}$, and of the process noise, Q_k ,

$$\begin{aligned} P_{k|k-1} &= F_k P_{k-1|k-1} F_k^T + Q_k, \\ Q_k &= G^V V_k (G^V)^T + W_k. \end{aligned}$$

The (E)KF then incorporates the measurements, \mathbf{y}_k , based on their noise covariance, $R_k = \text{cov}\{\mathbf{y}_k, \mathbf{y}_k\} \in \mathbb{R}^{dn \times dn}$, and $H_k = \partial \hat{\mathbf{y}}_k / \partial \hat{\mathbf{x}}_{k|k-1} \in \mathbb{R}^{dn \times cn}$ – the dependence of

the predicted measurement, $\hat{\mathbf{y}}_k$, on the state estimate. Since, the measurements are assumed to be uncorrelated between pixels, Eqs.(2.8) and (2.10), R_k and H_k are also block diagonal,

$$R_k = \begin{bmatrix} R_{1,k} & & \\ & R_{2,k} & \\ & & \ddots \end{bmatrix}, \quad H_k = \begin{bmatrix} H_{1,k} & & \\ & H_{2,k} & \\ & & \ddots \end{bmatrix}.$$

The dense and high-dimensional Kalman gain is given by,

$$K_k = P_{k|k-1} H_k^T (H_k P_{k|k-1} H_k^T + R_k)^{-1},$$

and the state of the filter is updated via,

$$\begin{aligned} \hat{\mathbf{x}}_{k|k} &= \hat{\mathbf{x}}_{k|k-1} + K_k (\mathbf{y}_k - \hat{\mathbf{y}}_k), \\ P_{k|k} &= (I - K_k H_k) P_{k|k-1}. \end{aligned}$$

Presumably, as the number of subsystems, n , grows, the information available for estimating the coupling input, $\Delta \mathbf{v}_k$, increases as well. In the limit $n \rightarrow \infty$, one would expect to be able to perfectly estimate $\Delta \mathbf{v}_k$ so that it effectively becomes deterministic, and the only remaining source of uncertainty is $\Delta \mathbf{w}_k$. While the numerical results in Sec. 3.2 suggest that this is indeed what happens in a model of a coronagraph, it is not necessarily the case for general linear systems as illustrated by an example in Sec. 3.3.

The filter in Sec. 2.1 used the block-diagonal process-noise covariance, Σ , from Eq. (3.1) instead of Q_k , i.e., ignored the correlation of the electric fields increments between pixels. It therefore allowed formulating a separate filter for each pixel which is equivalent to keeping the covariance matrices, $P_{k|k-1}$ and $P_{k|k}$, and gains, K_k , also

block diagonal. By propagating n filters with a $c = 2$ dimensional state each, it avoided the $\mathcal{O}(n^3)$ time complexity of the full EKF. Also, by discarding the off-diagonal elements of Q_k , it ignored almost all the information about the low dimensional coupling, $\Delta \mathbf{v}_k$, resulting in a relatively poor accuracy.

Alternatively, one can choose to discard the off-diagonal elements of $P_{k|k}$, rather than Q_k . This would lead to the following “block-diagonal” approximation, denoted by $\tilde{\cdot}$,

$$\tilde{P}_{k|k-1} = F_k \tilde{P}_{k-1|k-1} F_k^T + Q_k, \quad (3.3)$$

$$\begin{aligned} \tilde{K}_k &= \tilde{P}_{k|k-1} H_k^T \left(H_k \tilde{P}_{k|k-1} H_k^T + R_k \right)^{-1}, \\ \tilde{P}_{k|k} &= \mathcal{D} \left\{ \left(I - \tilde{K}_k H_k \right) \tilde{P}_{k|k-1} \right\}, \end{aligned} \quad (3.4)$$

where the operator for “discarding” the off-diagonal elements is formally given by,

$$\mathcal{D} \left\{ \begin{bmatrix} M_{1,1} & M_{1,2} & \cdots & M_{1,n} \\ M_{2,1} & M_{2,2} & & M_{2,n} \\ \vdots & & \ddots & \vdots \\ M_{n,1} & M_{n,2} & \cdots & M_{n,n} \end{bmatrix} \right\} = \begin{bmatrix} M_{1,1} & 0 & \cdots & 0 \\ 0 & M_{2,2} & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & M_{n,n} \end{bmatrix},$$

for all $M_{1 \leq i, j \leq n}$. In this case, the sub-optimal gain, \tilde{K}_k , does take the cross-correlation into account and the update-step covariance *approximation*, $\tilde{P}_{k|k} = \mathcal{D} \left\{ \tilde{P}_{k|k} \right\}$, has just $\mathcal{O}(n)$ elements. An efficient way to advance this filter that avoids explicitly computing \tilde{K}_k is given below[79].

3.1.1 Efficient Computation of the Block-diagonal Approximation¹

Note that the prediction-step covariance approximation is not block diagonal, $\mathcal{D}\{\tilde{P}_{k|k-1}\} \neq \tilde{P}_{k|k-1}$, and it remains to be shown that advancing $\tilde{P}_{k-1|k-1}$ to $\tilde{P}_{k|k}$ can be done in $\mathcal{O}(n)$ operations, i.e., without computing most of the elements of $\tilde{P}_{k|k-1}$. The matrix inversion lemma[80] allows expressing $\tilde{P}_{k|k}$ in terms of a block diagonal matrix, A_k , and products of rank r matrices,

$$\tilde{P}_{k|k} = A_k + \mathcal{D}\left\{B_k C_{k,1} B_k^T + G^V C_{k,2} (G^V)^T + G^V C_{k,3} B_k^T + B_k C_{k,3} (G^V)^T\right\}, \quad (3.5)$$

which are defined via

$$\begin{aligned} A_k &= L_k - L_k H_k^T M_k^{-1} H_k L_k, & \in \mathbb{R}^{cn \times cn} \\ B_k &= L_k H_k^T M_k^{-1} H_k G^V, & \in \mathbb{R}^{cn \times r} \\ C_{k,1} &= (V_k^{-1} + N_k)^{-1}, & \in \mathbb{R}^{r \times r} \\ C_{k,2} &= V_k (N_k C_{k,1} N_k - N_k) V_k + V_k, & \in \mathbb{R}^{r \times r} \\ C_{k,3} &= V_k N_k C_{k,1} - V_k, & \in \mathbb{R}^{r \times r} \end{aligned}$$

and

$$\begin{aligned} L_k &= F_k \tilde{P}_{k-1|k-1} F_k^T + W_k, & \in \mathbb{R}^{cn \times cn} \\ M_k &= R_k + H_k L_k H_k^T, & \in \mathbb{R}^{dn \times dn} \\ N_k &= (G^V)^T H_k^T M_k^{-1} H_k G^V. & \in \mathbb{R}^{r \times r} \end{aligned}$$

¹ A code example is available at <https://github.com/leonidprinceton/block-diagonal-filter>

To get $\tilde{P}_{k|k}$ from $\tilde{P}_{k-1|k-1}$, one first computes, in $\mathcal{O}(n)$, A_k, L_k, M_k since they are block-diagonal. Then, one proceeds by obtaining N_k and B_k in $\mathcal{O}(nr^2)$ by computing all the sparse products before multiplying them by G^V . The $C_{k,1}, C_{k,2}, C_{k,3}$ are then computed in $\mathcal{O}(r^3)$ since they are $r \times r$ matrices. Finally, the $\mathcal{O}(n)$ non-zero elements of $\tilde{P}_{k|k}$ are computed in $\mathcal{O}(r^2)$ each, as products of r dimensional vectors inside the brackets in Eq. (3.5).

Moreover, the corresponding approximation of the Kalman gain can be expressed as

$$\tilde{K}_k = \left(L_k - B_k C_{k,1} (G^V)^T - G C_{k,3} (G^V)^T \right) H_k^T M_k^{-1}, \quad (3.6)$$

and the estimate update, $\tilde{K}_k (\mathbf{y}_k - \hat{\mathbf{y}}_k)$, can be computed in $\mathcal{O}(n)$ by carefully choosing the order of multiplication.

Although not used for control purposes, the estimate of the coupling, $\Delta \mathbf{v}_k$, and an approximation of its error covariance are given by,

$$\begin{aligned} \Delta \hat{\mathbf{v}}_k &= -C_{k,3} (G^V)^T H_k^T M_k^{-1} (\mathbf{y}_k - \hat{\mathbf{y}}_k), \\ \text{cov} \{ \Delta \hat{\mathbf{v}}_k - \Delta \mathbf{v}_k, \Delta \hat{\mathbf{v}}_k - \Delta \mathbf{v}_k \} &\approx (V_k^{-1} + N_k)^{-1} = C_{k,1}. \end{aligned} \quad (3.7)$$

If $\tilde{P}_{k-1|k-1}$ happens to be the true error covariance of $\hat{\mathbf{x}}_{k-1}$, then the above $\Delta \hat{\mathbf{v}}_k$ and $C_{k,1}$ are exactly the conditional mean and variance of $\Delta \mathbf{v}_k$ given the measurement \mathbf{y}_k . Eq. (3.7) also suggests that as the number of subsystems increases, so does the norm of N_k , and therefore the uncertainty in the coupling estimate decreases. This could possibly lead to increasingly better estimates of the block-diagonal filter as discussed in Sec. 3.3.

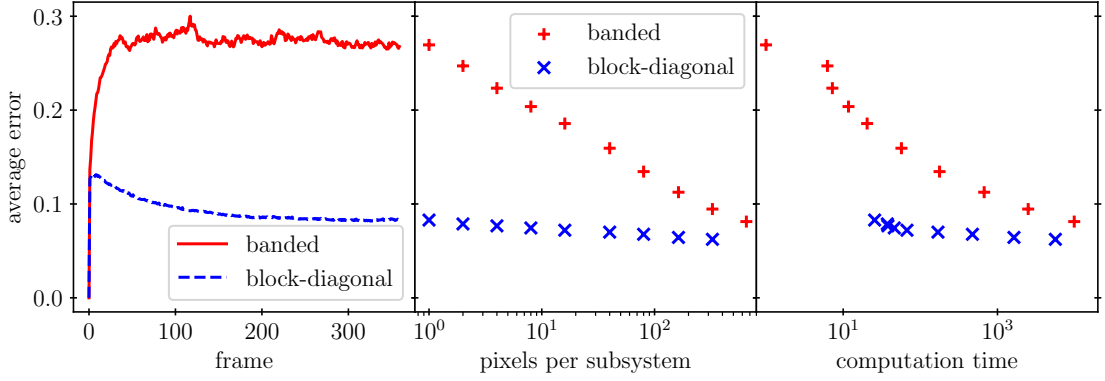


Figure 3.2: Average errors of the banded (solid red line and +) and block-diagonal (dashed blue line and ×) filters across all pixels in the dark hole (in $\sqrt{\frac{\text{photon}}{\text{frame}}}$) as a function of (from left to right): time, number of pixels assigned to each subsystem, computation time (arbitrarily scaled). The banded filter ignores the cross-correlation between subsystems and is therefore faster and less accurate than the block-diagonal filter. As more pixels are grouped together, its accuracy increases at the expense of computation time (in the limit of just one subsystem, it becomes the Kalman filter). The block-diagonal filter has a faster performance for the same accuracy for a wide range subsystem sizes.

3.2 Numerical Results

The cost-efficiency of the block-diagonal filter is analyzed below for the data simulated in the scenario described in detail in Sec. 2.3. The empirical covariance matrix of open-loop electric field increments,

$$Q = \text{cov} \{ \mathbf{E}_k^{OL} - \mathbf{E}_{k-1}^{OL}, \mathbf{E}_k^{OL} - \mathbf{E}_{k-1}^{OL} \},$$

was computed in full and then split into $G^V V (G^V)^T$ and W via Singular Value Decomposition (SVD). Specifically, if

$$Q = T^V \begin{bmatrix} \sigma_1 & & \\ & \sigma_2 & \\ & & \ddots \end{bmatrix} (T^V)^T,$$

then G^V was chosen as the first 21 (per (2.17)) columns of T^V , and hence

$$V = \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_{21} \end{bmatrix}, \quad W = \mathcal{D} \left\{ Q - G^V V (G^V)^T \right\}.$$

Figure 3.2 shows the average error of the estimate of the banded filter employed in Sec. 2.3 and the block-diagonal filter given by Eqs. (3.5) and (3.6). The left panel depicts the evolution of the errors in time, with the block-diagonal filter clearly outperforming the banded filter. However, the former was also about 30 times slower than the latter.

The accuracy of both filters can be increased at the expense of their performance. By bundling adjacent pixels together into the same block, the dimension of each subsystem, c , can be increased while keeping the dimension of the whole system fixed, i.e., $n \propto c^{-1}$. As shown in the middle panel of Fig. 3.2, the accuracy of the banded filter greatly benefits from increasing the size of subsystem and in the limit $c = 2n$ (not computed) it becomes the full KF. Yet, since the computational complexity of the filters grows as c^2 , the slight improvements in accuracy slow down the estimation process by orders of magnitude.

The cost-efficiency of the block-diagonal filter becomes evident when the errors are plotted against computation time (right panel of Fig. 3.2). The banded filter with 16 pixels per block is less than half as accurate and almost as slow as the block-diagonal filter with one pixel per block. To reach the same accuracy, the banded filter would have to be about 400 times slower with 650 pixels per block. Therefore, if a good low-dimensional model of the drift is available and the computational power allows, the block-diagonal filter would be preferable to the pixel-based EKF in Sec.2.1 since it would allow reducing the dither magnitude and help reject faster wavefront drifts.

3.3 Properties (Fixed Matrices Case)

The block-diagonal filter is clearly sub-optimal, but one could hope to show that it is stable and gets “close” to the Kalman Filter as the “relative importance” of the coupling input decreases. Unfortunately, neither property is necessarily true although they can be proven to hold under some specific conditions. The discussion will, for simplicity, be constrained to fixed matrices even though a precomputed gain would be preferable to the block-diagonal filter in that case.

First, it will be shown that the gain of the block-diagonal filter “converges” in time which will then allow formulating conditions for the asymptotic stability of the filter. Finally, the question of asymptotic decoupling will be addressed, i.e., whether as the number of subsystem increases, the coupling term ($G^V \Delta \mathbf{v}_k$ in Eq. (3.2)) becomes effectively “deterministic”.

3.3.1 Convergence

For fixed system matrices ($F_k = F$, etc.) such that the pair $[F, H]$ is stabilizable and $[F, \sqrt{Q}]$ is detectable (with $\sqrt{Q}(\sqrt{Q})^T = Q$), the prediction-step covariance of the full KF, $P_{k|k-1}$, converges to a unique limit, $P_- = \lim_{k \rightarrow \infty} P_{k|k-1}$. This limit satisfies the discrete Algebraic Ricatti Equation (ARE)[81],

$$P_- = F \left(P_- - P_- H^T (H P_- H^T + R)^{-1} H P_- \right) F^T + Q. \quad (3.8)$$

The update-step covariance and the asymptotic Kalman gain are also unique,

$$K = \lim_{k \rightarrow \infty} K_k = P_- H^T (H P_- H^T + R)^{-1},$$

$$P = \lim_{k \rightarrow \infty} P_{k|k} = (I - KH) P_-,$$

and the closed loop transition matrix of the KF, $F_C = (I - KH)F$, is a stability matrix.

Similarly, for $Q > 0$, the prediction-step covariance *approximation* of block-diagonal filter, $\tilde{P}_{k|k-1}$, has a limit, $\tilde{P} = \lim_{k \rightarrow \infty} \tilde{P}_{k|k-1}$. It is a solution of an analogous equation,

$$\tilde{P} = \mathcal{D} \left\{ F \left(P - P H^T (H P H^T + R)^{-1} H P \right) F^T \right\} + Q, \quad (3.9)$$

that also happens to be unique[79]. The corresponding asymptotic gain and update-step covariance approximation are given by

$$\begin{aligned} \tilde{K} &= \lim_{k \rightarrow \infty} \tilde{K}_k = \tilde{P} H^T \left(H \tilde{P} H^T + R \right)^{-1} \\ \tilde{P} &= \lim_{k \rightarrow \infty} \tilde{P}_{k|k} = \mathcal{D} \left\{ (I - KH) \tilde{P} \right\}. \end{aligned}$$

Note that, \tilde{P} and \tilde{P} are not the actual covariance matrices of the error estimates, $\hat{\mathbf{x}}_k - \mathbf{x}_k$ in the limit $k \rightarrow \infty$ and, in some cases, might be very poor approximations thereof. Some counter-examples and conditions for when \tilde{P} does get close to P are given below.

3.3.2 Stability

The first case in which the block-diagonal approximation, \tilde{P} , is poor, is when it results in an unstable block-diagonal filter. For example, consider

$$F = \begin{bmatrix} 0.998 & -0.052 & & \\ 0.052 & 0.998 & & \\ & & 0.993 & -0.110 \\ & & 0.110 & 0.993 \end{bmatrix}, \quad Q = \begin{bmatrix} 1.01 & 1 & 1 & 1 \\ 1 & 1.01 & 1 & 1 \\ 1 & 1 & 1.01 & 1 \\ 1 & 1 & 1 & 1.01 \end{bmatrix}, \quad (3.10)$$

$$H = \begin{bmatrix} 1 & 0 & & \\ & & 1 & 0 \end{bmatrix}, \quad R = \begin{bmatrix} 1 & \\ & 1 \end{bmatrix},$$

with $n = c = 2$, and note that the system itself is stable (the eigenvalues of F are all inside the unit circle). In this case, the matrix \tilde{P} can be computed as the limit of equations (3.3)-(3.4) to be

$$\tilde{P} = \begin{bmatrix} 0.57 & -0.46 & & \\ -0.46 & 24.87 & & \\ & & 0.58 & -0.39 \\ & & -0.39 & 11.65 \end{bmatrix},$$

and it results in an unstable filter. That is, it can be directly verified that one of the eigenvalues of the closed loop state transition matrix, $\tilde{F}_C = (I - \tilde{K}H)F$, is just outside the unit circle.

The numbers in (3.10) were chosen carefully so that the system is “barely” observable and Q is far from being block diagonal. However, when the coupling, $\Delta\mathbf{v}$, is “weak”, it can be seen as a perturbation of Eq. (3.9) around a solution of the ARE, (3.8) (since the equations coincide when Q is block diagonal).

Formally, let

$$\eta = \left\| G^V V (G^V)^T \right\|_F \quad (3.11)$$

be a measure of the effect of the coupling. For small η , the difference between the covariance of the KF and its block diagonal approximation, $\left\| P - \tilde{P} \right\|$, is linearly bounded by η ,

$$\left\| P - \tilde{P} \right\| = O(\eta),$$

when all other parameters are held constant[79]. The difference between the gains and the closed loop transition matrices can similarly be bound by η ,

$$\begin{aligned} \left\| K - \tilde{K} \right\| &= O(\eta), \\ \left\| F_C - \tilde{F}_C \right\| &= O(\eta). \end{aligned}$$

It follows from eigenvalue perturbation theory[82] that the eigenvalues of \tilde{F}_C also remain within $O(\eta)$ of the eigenvalues of F_C . Finally, since F_C is a stability matrix, its eigenvalues are inside the unit circle, and therefore there exists η small enough such that \tilde{F}_C is a stability matrix too. Still, this criterion for the stability of the block-diagonal filter is not useful when the number of subsystems is large since η increases with the number of rows in G^V .

3.3.3 Asymptotic Decoupling

As the number of subsystems, n , increases, their measurements provide more information about the coupling, $\Delta \mathbf{v}$, whose dimension remains fixed. If $\Delta \mathbf{v}$ was perfectly known, it would have no effect on the KF covariance (i.e., $V = 0$) and the steady-state update-step covariance of both the KF and the block-diagonal filter would coincide and satisfy,

$$P_0 = \mathcal{D} \{P_0\} = \left((F P_0 F^T + W)^{-1} + H^T R^{-1} H \right)^{-1}.$$

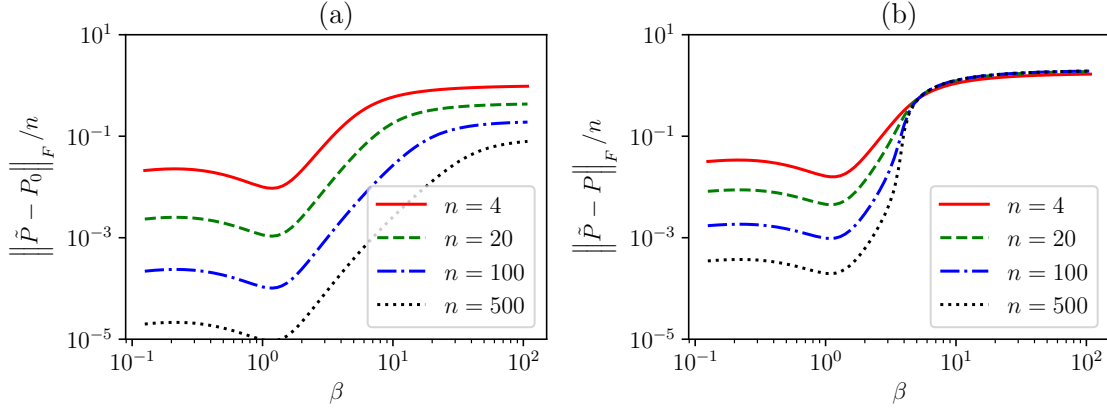


Figure 3.3: (a) The scaled difference between the steady-state covariance of the decoupled filter, P_0 , and of the block-diagonal approximation of the coupled system, \tilde{P} , corresponding to Eqs. (3.12) and (3.13). The blocks on the main diagonal of these two matrices get close to one another as the number of subsystems, n , increases regardless of β . (b) The scaled difference between \tilde{P} and the covariance of the full KF of the coupled system, P . For β beyond some critical value, \tilde{P} becomes a poor approximation of P . It follows that P also remains far from P_0 , i.e., the KF remains coupled despite the increasingly available information about the coupling, $\Delta \mathbf{v}$. *from* [79]

Hence, one might expect that as $n \rightarrow \infty$, the coupling will become better estimated, $\Delta \hat{\mathbf{v}}_k \rightarrow \Delta \mathbf{v}_k$, and the filters will become decoupled, $P, \tilde{P} \rightarrow P_0$.

The next example shows otherwise. Consider a set of identical systems parameterized by β

$$F_i = \begin{bmatrix} 0.9 & \beta \\ 0 & 0.9 \end{bmatrix}, G_i^V = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, H_i = \begin{bmatrix} 1 & 1 \end{bmatrix}, \quad (3.12)$$

$$V = \begin{bmatrix} 1 \end{bmatrix}, W_i = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, R_i = \begin{bmatrix} 1 \end{bmatrix}, \quad (3.13)$$

which is another case in which the block-diagonal approximation, \tilde{P} , might diverge from the true covariance. Figure 3.3(a) shows the norm of the difference between each block on the main diagonals of P_0 and \tilde{P} (or equivalently $\|P_0 - \tilde{P}\|_F / n$). It decreases

linearly with the number of subsystems regardless of the value of the parameter β , i.e., the block-diagonal filter approximates the system as becoming increasingly decoupled.

Yet, the actual system and the corresponding full KF covariance, P , do not necessarily become decoupled. Indeed, for β above some critical value, $\|\tilde{P} - P\|_F/n$ does not decrease with n thus rejecting the possibility that P becomes close to \tilde{P} (or P_0). Since P is the true covariance of the (optimal) KF, the true covariance of the block-diagonal filter must be larger and even further away from P_0 (and from its approximation, \tilde{P}). For other β , one has $\|P - \tilde{P}\|_F/n \rightarrow 0$ and $\|P - P_0\|_F/n \rightarrow 0$ and, as a consequence of the triangle inequality, also $\|P_0 - \tilde{P}\|_F/n \rightarrow 0$, i.e., the full KF becomes decoupled.

To quantify the convergence of \tilde{P} to P_0 , it is helpful to define the relative influence of the coupling,

$$C = \left(V^{-1} + (G^V)^T H^T \left(H \left(F \dot{P} F^T + W \right) H^T + R \right)^{-1} H G^V \right)^{-1} \in \mathbb{R}^{r \times r} \quad (3.14)$$

where \dot{P} is the asymptotic covariance approximation of the banded filter (which assumes that the process noise covariance is block diagonal, $\mathcal{D}\{Q\}$),

$$\dot{P} = \mathcal{D}\{\dot{P}\} = \left(\left(F \dot{P} F^T + \mathcal{D}\{Q\} \right)^{-1} + H^T R^{-1} H \right)^{-1}.$$

Note that $\dot{P} \geq P_0$ since $\mathcal{D}\{Q\} = \mathcal{D}\{G^V V (G^V)^T + W\} \geq W$. The intuitive meaning of C is the first covariance of $\Delta \hat{\mathbf{v}}_k - \Delta \mathbf{v}_k$ if the coupling, $\Delta \mathbf{v}_k$, was suddenly to appear after the “decoupled” filter has converged.

Careful inspection of Eq. (3.14) shows that C “decreases” with the addition of new subsystems since \dot{P} and W are block diagonal and positive definite. In particular, if the subsystems are identical as in the above example, the norm of C is roughly

inversely proportional to their number,

$$\|C\| = \Theta(n^{-1}).$$

This implies that the effect of coupling on each subsystems diminishes with n .

Formally, let

$$\epsilon_i = \left\| G_i^V C (G_i^V)^T \right\|_F$$

be a measure of the influence of the coupling on the i th system, and note that ϵ_i decreases with n (unlike η in Eq. (3.11)). For small ϵ_i , the difference between the i th block of P_0 and \tilde{P} is linearly bounded by ϵ_i , i.e.,

$$\left\| (P_0)_i - \tilde{P}_i \right\| = O(\epsilon_i)$$

when all other parameters are held constant[79].

The block-diagonal filter is therefore a good approximation of the full KF if the latter becomes decoupled. This is likely to happen for large systems with mixed low-order and uncorrelated process noise such as the WFIRST-CGI model in Sec. 3.2. In those cases, the block-diagonal filter is significantly more accurate than the banded filter and vastly more efficient than the full KF.

Chapter 4

Reduced-order Approach to Post Processing

The standard post-processing pipeline begins with raw history of detector measurements that needs to be parsed, calibrated and reduced to a small set of images at different telescope orientations and wavelengths[83]. These intensity images are mostly of speckles (residual starlight) with other objects (e.g. planets or debris disks) mixed within. Most such objects are not immediately distinguishable from speckles in any single image, and there exist many approaches and several software libraries which allow extracting these dim signals[84, 85].

The simplest post-processing method is PSF subtraction and, as its name suggests, it involves subtracting the presumably known PSF of the speckles from the images. The subtracted PSF can be obtained, for example, by making another observation of a reference star with a spectrum similar to that of the target[86]. When multiple reference observations are available, they can be linearly combined into the to-be-subtracted PSF that best matches a given image. The basis for this procedure is known as Karhunen-Loève Image Projection (KLIP)[87] or Principal Component

Analysis (PCA)[88, 89], and more sophisticated versions include the robust PCA [90] and the Non-negative Matrix Factorization(NMF)[91].

It is also possible to discern the incoherent signal from speckles in a single set of images based on their different physical behavior in terms of spectrum[92, 93, 94], polarization[95], phase apodization[96, 97], etc. However, most of these phenomena provide little phase diversity at small angles of separation[98] (for example, the locations of the peaks of the perturbed speckles in Fig. 2.3 do not move significantly as a function of wavelength and, in that sense, behave similarly to planets).

A source of diversity that *will* be exploited by WFIRST-CGI is image orientation. Similarly to ground-based telescopes which rotate with Earth, WFIRST will be rolled along its observation axis by 26deg every two hours[47]. While the diffraction pattern of the starlight and measurement noise are internal to the instrument and therefore do not depend on the roll angle, PSF from other objects will change their location in the image. The Angular Differential Imaging (ADI)[99, 100] and Locally Optimized Combination of Images (LOCI)[101] are methods that subtract and recombine images at different orientation to retrieve those incoherent signals.

Yet, none of the above approaches take DM actuations into account. Most use “co-added” images of thousands of short exposures with different DM setting each, although it is also possible to incorporate short exposures of “frozen speckles” and data from wavefront sensors used by AO[102]. In space-based telescopes, the slowly varying nature of the speckles implies that even minute-long exposures are well modeled by just one electric field per wavelength. It might therefore be beneficial to formulate the post-processing problem in terms of the electric field rather than intensity to allow taking the history of DM controls into account. Moreover, the dithering of the DM which is necessary for dark hole maintenance (Sec. 2.2), itself constitutes a source of phase diversity.

4.1 An Electric Field Formulation of the Post-processing Problem

If the observation of a target star was performed in a closed-loop manner as discussed in chapter 2, then its data would consist of $n \cdot T$ photon counts, $\{y_{i,k}\}$ across n pixels and T frames, the history of the DM controls $\{\mathbf{u}_k\}$, and the controls Jacobian, G^U . The goal of the probabilistic algorithms below is to estimate the total intensity of all incoherent sources, $I_{i,k}^I = I_{i,k}^D + I_{i,k}^P$, by separating it from the time varying intensity of the speckles, $I_{i,k}^S$.

4.1.1 A Maximum Likelihood (ML) Formulation

The number of detected photons has a known distribution parameterized by the incoherent sum of the intensities, (1.1) or

$$y_{i,k} \sim \text{Pois}(I_{i,k}) = \text{Pois}(I_{i,k}^S + I_{i,k}^I).$$

In other words, the probability of observing $y_{i,k}$ given $I_{i,k}$ (or the likelihood of $I_{i,k}$ given the measured $y_{i,k}$) is

$$p(y_{i,k} | I_{i,k}) = \frac{I_{i,k}^{y_{i,k}} \exp(-I_{i,k})}{y_{i,k}!}. \quad (4.1)$$

The effects of the controls, $\mathbf{g}_i^U \cdot \Delta \mathbf{u}_k$, are assumed to be known (see (1.4)),

$$I_{i,k}^S = |E_{i,k}^{OL} + \mathbf{g}_i^U \cdot \mathbf{u}_k|^2,$$

and thus can be incorporated into the above likelihood function,

$$p(y_{i,k} | E_{i,k}^{OL}, I_{i,k}^I) = p(y_{i,k} | |E_{i,k}^{OL} + \mathbf{g}_i^U \cdot \mathbf{u}_k|^2 + I_{i,k}^I). \quad (4.2)$$

Unfortunately, $p(y_{i,k} | E_{i,k}^{OL}, I_{i,k}^I)$ now contains three real-valued parameters for each measurement, and the unconstrained maximum likelihood (ML) formulation,

$$\max_{E_{i,k}^{OL}, I_{i,k}^I} p(y_{i,k} | E_{i,k}^{OL}, I_{i,k}^I), \quad (4.3)$$

is ill-posed. However, arbitrary changes in electric fields and intensities between adjacent time frames are unphysical, and the problem can be regularized by introducing suitable constraints on these quantities.

The first such constraint prevents the incoherent intensity from fluctuating (except if the telescope changes its orientation amidst the observation, see Eq. (4.13) and surrounding discussion). This is a good assumption for the main PSF lobes of the exoplanet as they do not vary significantly with wavefront and DM perturbations (see Fig. 2.3). In other words, the intensity satisfies $I_{i,k}^I = I_i^I = \text{const}, \forall k$, and furthermore it can be approximated as the average discrepancy between the measurements and the speckles,

$$I_i^I \approx \frac{1}{T} \sum_{k=1}^T (y_{i,k} - I_{i,k}^S).$$

Formally, the incoherent intensity estimate, $\hat{I}_{i,k}^I$, is defined as the following function of the estimates of the open-loop electric field, $\{\hat{E}_{i,k}^{OL}\}_{k=1}^T$,

$$\hat{I}_i^I \left(\{\hat{E}_{i,k}^{OL}\}_{k=1}^T \right) = \text{ramp} \left\{ \frac{1}{T} \sum_{k=1}^T \left(y_{i,k} - \left| \hat{E}_{i,k}^{OL} + \mathbf{g}_i^U \cdot \mathbf{u}_k \right|^2 - I_i^D \right) \right\} + I_i^D, \quad (4.4)$$

where, I_i^D , is the constant intensity of the dark current which must be above \hat{I}_i^I . This is ensured by the ramp function,

$$\text{ramp}\{\xi\} = \begin{cases} \xi & \xi > 0 \\ 0 & \xi \leq 0 \end{cases}.$$

Although the estimation problem is now formulated in terms of just the electric field, it still remains ill-posed. The two techniques below regularize the problem by imposing either temporal or spatial smoothness on the speckles.

4.1.2 Regularization via Drift Increments Prior

The assumption that the electric field performs a random walk, (2.2) and (2.12), provides a probability distribution for the electric field increments,

$$p(E_{i,k+1}^{OL} - E_{i,k}^{OL}) = \exp \left(-\frac{1}{2} \begin{bmatrix} \text{Re} \{E_{i,k+1}^{OL} - E_{i,k}^{OL}\} \\ \text{Im} \{E_{i,k+1}^{OL} - E_{i,k}^{OL}\} \end{bmatrix}^T \Sigma_i^{-1} \begin{bmatrix} \text{Re} \{E_{i,k+1}^{OL} - E_{i,k}^{OL}\} \\ \text{Im} \{E_{i,k+1}^{OL} - E_{i,k}^{OL}\} \end{bmatrix} \right), \quad (4.5)$$

$$p(E_{i,0}^{OL}) = \exp \left(-\frac{1}{2} \begin{bmatrix} \text{Re} \{E_{i,0}^{OL} - \mu_{i,0}\} \\ \text{Im} \{E_{i,0}^{OL} - \mu_{i,0}\} \end{bmatrix}^T \Sigma_{i,0}^{-1} \begin{bmatrix} \text{Re} \{E_{i,0}^{OL} - \mu_{i,0}\} \\ \text{Im} \{E_{i,0}^{OL} - \mu_{i,0}\} \end{bmatrix} \right), \quad (4.6)$$

which act as a prior for the trajectory $\{E_{i,k}^{OL}\}_{k=1}^T$. This prior penalizes solutions of (4.3) that overfit the noisy measurements, $\{y_{i,k}\}_{k=1}^T$, with large temporal “jumps” in the electric field.

This soft constraint transforms the goal into finding the maximum of the *a-posteriori* (MAP) probability of the measurements given Eqs. (4.5) and (4.6), i.e.,

$$\max_{\{E_{i,k}^{OL}\}_{k=1}^T, \{I_{i,k}^I\}_{k=1}^T} \left[\prod_{k=1}^T p(y_{i,k} | E_{i,k}^{OL}, I_{i,k}^I) \right] \cdot \left[\prod_{k=1}^{T-1} p(E_{i,k+1}^{OL} - E_{i,k}^{OL}) \right] \cdot p(E_{i,0}^{OL}). \quad (4.7)$$

Numerically, a local maximum of the above posterior is found by optimizing a log-likelihood cost function, J_i . It is defined below in terms of the “intermediate” estimate

of the incoherent intensity, Eq. (4.4)[49],

$$\begin{aligned}
J_i \left(\left\{ \hat{E}_{i,k}^{OL} \right\}_{k=1}^T \right) &= \sum_{k=1}^T \left(\left| \hat{E}_{i,k}^{OL} + \mathbf{g}_i^U \cdot \mathbf{u}_k \right|^2 + \hat{I}_i^I - y_{i,k} \log \left(\left| \hat{E}_{i,k}^{OL} + \mathbf{g}_i^U \cdot \mathbf{u}_k \right|^2 + \hat{I}_i^I \right) \right) \\
&\quad + \frac{1}{2} \sum_{k=1}^{T-1} \left(\begin{bmatrix} \text{Re} \left\{ \hat{E}_{i,k+1}^{OL} - \hat{E}_{i,k}^{OL} \right\} \\ \text{Im} \left\{ \hat{E}_{i,k+1}^{OL} - \hat{E}_{i,k}^{OL} \right\} \end{bmatrix}^T \Sigma_i^{-1} \begin{bmatrix} \text{Re} \left\{ \hat{E}_{i,k+1}^{OL} - \hat{E}_{i,k}^{OL} \right\} \\ \text{Im} \left\{ \hat{E}_{i,k+1}^{OL} - \hat{E}_{i,k}^{OL} \right\} \end{bmatrix} \right) \\
&\quad + \frac{1}{2} \begin{bmatrix} \text{Re} \left\{ E_{i,0}^{OL} - \mu_{i,0} \right\} \\ \text{Im} \left\{ E_{i,0}^{OL} - \mu_{i,0} \right\} \end{bmatrix}^T \Sigma_{i,0}^{-1} \begin{bmatrix} \text{Re} \left\{ E_{i,0}^{OL} - \mu_{i,0} \right\} \\ \text{Im} \left\{ E_{i,0}^{OL} - \mu_{i,0} \right\} \end{bmatrix}, \\
\hat{I}_i^I \left(\left\{ \hat{E}_{i,k}^{OL} \right\}_{k=1}^T \right) &= \text{ramp} \left\{ \frac{1}{T} \sum_{k=1}^T \left(y_{i,k} - \left| \hat{E}_{i,k}^{OL} + \mathbf{g}_i^U \cdot \mathbf{u}_k \right|^2 - I_i^D \right) \right\} + I_i^D,
\end{aligned} \tag{4.8}$$

which is obtained by combining Eqs. (4.1), (4.2) and (4.4)-(4.6) and dropping constant terms. The resulting incoherent intensity estimate is given by

$$\hat{I}_i^I \left(\underset{\left\{ \hat{E}_{i,k}^{OL} \right\}_{k=1}^T}{\text{argmax}} J_i \left(\left\{ \hat{E}_{i,k}^{OL} \right\}_{k=1}^T \right) \right).$$

4.1.3 Regularization via Electric Field Order Reduction (EFOR)¹

Another way to make the ML estimation problem well posed is to reduce the number of free parameters by constraining speckles to a low-dimensional subspace,

$$\mathbf{E}_k^{OL} = G^V \mathbf{v}_k. \tag{4.9}$$

Here, the modes, $G^V \in \mathbb{R}^{2n \times r}$, have the same meaning as in chapter 3 expect that they also contain the initial electric field, since their coefficients, $\{\mathbf{v}_k\}_{k=1}^T$, will be estimated

¹ A code example is available at <https://github.com/leonidprinceton/EFOR>

rather than their increments, $\{\Delta \mathbf{v}_k\}_{k=1}^{T-1}$. If r is small enough, the total number of free parameters, $2nr + Tr$, is small compared to the number of measurements, nT , hence the estimation problem is over-constrained. Note that it is still ill-posed in the absence of DM dithering, $\mathbf{u}_k = 0, \forall k$, due to phase ambiguity.

Since the increments are now coupled between pixels, the optimization will be formulated in global terms, i.e.,

$$p(\mathbf{y}_k | \mathbf{E}_k^{OL}, \mathbf{I}_k^I) = p\left(\mathbf{y}_k | B \cdot (\mathbf{E}_k^{OL} + G^U \cdot \mathbf{u}_k)^{\circ 2} + \mathbf{I}_k^I\right) \quad (4.10)$$

instead of (4.2), with p now standing for the multivariate Poisson distribution, \circ denoting an elementwise squaring, and B – the linear operation of adding the contributions of the real and imaginary parts of the electric field (in the monochromatic case, $B = I_{n \times n} \otimes \begin{bmatrix} 1 & 1 \end{bmatrix} \in \mathbb{R}^{n \times 2n}$, where \otimes stands for the Kronecker product). Similarly, the incoherent intensity estimate in Eq. (4.4) becomes,

$$\hat{\mathbf{I}}^I(\hat{G}^V, \{\hat{\mathbf{v}}_k\}_{k=1}^T) = \text{ramp} \left\{ \frac{1}{T} \sum_{k=1}^T \left(\mathbf{y}_k - B \cdot (\hat{G}^V \hat{\mathbf{v}}_k + G^U \mathbf{u}_k)^{\circ 2} - \mathbf{I}^D \right) \right\} + \mathbf{I}^D, \quad (4.11)$$

where ramp is the elementwise ramp function and \mathbf{I}^D is the vector of dark currents at all pixels.

The cost function used to maximize the log-likelihood with EFOR is [103]

$$J(\hat{G}^V, \{\hat{\mathbf{v}}_k\}_{k=1}^T) = \mathbf{1} \cdot \sum_{k=1}^T \left(B \cdot (\hat{G}^V \hat{\mathbf{v}}_k + G^U \mathbf{u}_k)^{\circ 2} + \hat{\mathbf{I}}^I - \mathbf{y}_k \circ \log \left(B \cdot (\hat{G}^V \hat{\mathbf{v}}_k + G^U \mathbf{u}_k)^{\circ 2} + \hat{\mathbf{I}}^I \right) \right), \quad (4.12)$$

where $\mathbf{1} \cdot$ stands for summing over the log-likelihoods of all pixels. Note that this is a ML estimator and therefore requires no prior knowledge about the distribution of the

electric field trajectories. The only hyper-parameter of this method is the number of electric field modes, r , the choice of which is discussed in Sec. 4.2.

4.1.4 Numerical Optimization Procedure

Expressing the incoherent intensity, $\hat{\mathbf{I}}^I$, in terms of either $\left\{\hat{\mathbf{E}}_k^{OL}\right\}_{k=1}^T$ or $\left\{\hat{G}^V \hat{\mathbf{v}}_k\right\}_{k=1}^T$ facilitates the optimization in two ways: first, the expression in Eqs. (4.4) and (4.11) ensure that the logarithms of the intensities always act on positive numbers; second, the optimization is not sensitive to the relative scaling and domain of the optimized parameters. Since the expressions for J_i and J above are differentiable with respect to their parameters given that

$$\frac{d}{d\xi} \text{ramp}\{\xi\} = \begin{cases} 1 & \xi > 0 \\ 0 & \xi \leq 0 \end{cases},$$

a straightforward approach to optimization is gradient descent.

One particularly convenient optimization library for this purpose is Tensorflow[104]. Given the definitions of J and $\hat{\mathbf{I}}^I$, it computes their analytical gradients with respect to $\hat{G}^V, \{\hat{\mathbf{v}}_k\}_{k=1}^T$ (or $\left\{\hat{E}_{i,k}^{OL}\right\}_{k=1}^T$) and employs the Adam Optimizer[105] as its algorithm for gradient descent (it requires choosing a single order-of-magnitude hyper parameter – the “learning rate”). Moreover, Tensorflow was designed to run on Graphics Processing Units (GPU) which accelerate the optimization by at least an order of magnitude.

Conveniently, since $\lim J = \infty$ as $\left\|\hat{\mathbf{E}}_k^{OL}\right\| \rightarrow \infty$, the optimization always converges. However, the minima it finds are mostly local and not isolated since \hat{G}^V and $\{\hat{\mathbf{v}}_k\}_{k=1}^T$ can be mutually rescaled without affecting J . Additionally, since $J\left(0, \{\mathbf{0}\}_{k=1}^T\right)$ is a saddle point, it is preferable to initialize \hat{G}^V as a random orthogonal matrix and pick a random initial guess for $\{\hat{\mathbf{v}}_k\}_{k=1}^T$ with a reasonable magnitude.

4.1.5 Realistic Telescope Effects with EFOR

EFOR also allows accounting for residual jitter from FSM[62] and telescope rolls[57] (but not DM actuators drift[61] since its effects are not low-order). This is done by modifying, Eq. (4.12) to incorporate a model of these effects as described below.

Telescope rolls may be employed during the observation to increase image diversity since they only affect the constant incoherent sources external to the instrument, \mathbf{I}^P . Formally this can be described as,

$$\mathbf{I}_k^P = R_k \{ \mathbf{I}_0^P \}, \quad (4.13)$$

where R_k denotes the rotation of the nominal incoherent image, \mathbf{I}_0^P , due to the difference in orientation between frames k and 0. R_k is assumed to be invertible and differentiable and for numerical purposes it will be approximated by a permutation matrix acting on the pixels in the way best resembling image rotation.

Line-of-sight jitter requires a more detailed approach. In general, if there is a clear separation between the fast and slow variations of the electric field, it is reasonable to split the electric field modes accordingly,

$$\mathbf{v}(t) = \mathbf{v}_k + \delta\mathbf{v}(t), \quad t \in [t_k, t_{k+1}).$$

where the average of $\delta\mathbf{v}(t)$ over frame k is defined to be zero, $\langle \delta\mathbf{v}(t) \rangle_k = 0$. Consequently, \mathbf{v}_k is the average of $\mathbf{v}(t)$ over $[t_k, t_{k+1})$, i.e., $\mathbf{v}_k = \langle \mathbf{v}(t) \rangle_k$, and it is assumed to vary slowly between frames.

The intensity of the speckles *corresponding to* frame k is now redefined as the average intensity *during* that frame,

$$\langle B \cdot \mathbf{E}^{\circ 2}(t) \rangle_k = \left\langle B \cdot (G^V \mathbf{v}(t) + G^U \mathbf{u}_k)^{\circ 2} \right\rangle_k, \quad (4.14)$$

and can be split into two parts, $\langle B \cdot \mathbf{E}^{\circ 2}(t) \rangle_k = \mathbf{I}_k^S + \mathbf{I}_k^J$, with

$$\begin{aligned}\mathbf{I}_k^S &= B \cdot (G^V \mathbf{v}_k + G^U \mathbf{u}_k)^{\circ 2}, \\ \mathbf{I}_k^J &= B \cdot \left\langle (G^V \delta \mathbf{v}(t))^{\circ 2} \right\rangle_k.\end{aligned}$$

Note that the static speckles term, \mathbf{I}_k^S , has been already accounted for by both online, Eq. (2.8), and offline, Eqs. (4.2) and (4.10), estimators. The additional jitter term, \mathbf{I}_k^J , is not affected by controls and therefore appears as an incoherent source. Since the fast electric field variations cannot be directly sampled, one has to make additional assumptions on $\delta \mathbf{v}(t)$ in order to distinguish between \mathbf{I}_k^J and external incoherent sources, \mathbf{I}_k^P .

To keep the problem tractable, the number of modes that exhibit fast time variations will be constrained to $m \sim \mathcal{O}(1)$, without loss of generality

$$\delta v_j(t) = 0, \forall j > m.$$

Given that the jitter is mostly the residual of the FSM (tip/tilt modes), a reasonable value for m is 2, and the specific choice of which columns of G^V correspond to tip/tilt modes is inconsequential since they are found during the optimization.

Even though the time history of $\delta v_j(t)$, $1 \leq j \leq m$ cannot be recovered, the jitter intensity can be expressed as a linear sum of $\binom{m}{2}$ products of jitter modes,

$$\mathbf{I}_k^J = B \cdot \sum_{i=1}^m \left(\sum_{j=1}^i w_{i,j,k} \mathbf{g}_j^V \right)^{\circ 2} \quad (4.15)$$

where \mathbf{g}_j^V , $1 \leq j \leq m$ are columns of G^V and $w_{i,j,k}$, $1 \leq j \leq i \leq m$ are some coefficients. This allows estimating and distinguishing the contribution due to jitter from external signals.

Combining Eqs. (4.4), (4.12), (4.13), and (4.15) together, the more general cost function, $J \left(\hat{G}^V, \{\hat{\mathbf{v}}_k\}_{k=1}^T, \{\hat{\mathbf{w}}_k\}_{k=1}^T \right)$, becomes

$$J = \mathbf{1} \cdot \sum_{k=1}^T \left(B \cdot \left(\hat{G}^V \hat{\mathbf{v}}_k + G^U \mathbf{u}_k \right)^{\circ 2} + \hat{\mathbf{I}}_k^I - \mathbf{y}_k \circ \log \left(B \cdot \left(\hat{G}^V \hat{\mathbf{v}}_k + G^U \mathbf{u}_k \right)^{\circ 2} + \hat{\mathbf{I}}_k^I \right) \right), \quad (4.16)$$

$$\begin{aligned} \hat{\mathbf{I}}_k^I &= R_k \hat{\mathbf{I}}_0^P + \hat{\mathbf{I}}_k^J + \mathbf{I}^D, \\ \hat{\mathbf{I}}_0^P &= \text{ramp} \left\{ \frac{1}{T} \sum_{k=1}^T R_k^{-1} \left(\mathbf{y}_k - B \cdot \left(\hat{G}^V \hat{\mathbf{v}}_k + G^U \mathbf{u}_k \right)^{\circ 2} - \hat{\mathbf{I}}_k^J - \mathbf{I}^D \right) \right\}, \\ \hat{\mathbf{I}}_k^J &= B \cdot \sum_{i=1}^m \left(\sum_{j=1}^i \hat{w}_{i,j,k} \mathbf{g}_j^V \right)^{\circ 2}, \end{aligned}$$

where $\hat{\mathbf{I}}_0^P$ is the estimate of the external signal (corresponding to the first telescope orientation) and $\hat{\mathbf{w}}_k = [\hat{w}_{i,j,k}]_{1 \leq j \leq i \leq m}$ are the estimates of the jitter coefficients. The additional parameters, $\hat{\mathbf{w}}_k$, behave similarly to $\hat{\mathbf{v}}_k$ and, as long as their number is relatively small ($\binom{m}{2} = \frac{m(m+1)}{2} < r$), they have a minor effect on the numerical optimization.

4.1.6 Reference Observations

As demonstrated with KLIP[87], reference observations allow building a library of speckle intensity modes thus facilitating their subtraction. This is analogous to having a good guess for the electric field modes, G^V , which can be incorporated in several ways. If the modes, G^V , are perfectly known, the optimization of the cost function, J , can be done in terms of just $\{\hat{\mathbf{v}}_k\}_{k=1}^T$. Although, since the system is not perfectly linear, it might be beneficial to perform a follow-up optimization in terms of both G^V and $\{\hat{\mathbf{v}}_k\}_{k=1}^T$.

If an accurate G^V is not available, multiple (or reference) observations from the same instrument can still be combined to achieve more accurate estimates. This is

illustrated below with a single additional observation of a reference star, but can be extended to any number of observations. The reference data consists of a set of measurements, $\{\mathbf{y}_k^r\}_{k=1}^{T_r}$, the corresponding history of controls, $\{\mathbf{u}_k^r\}_{k=1}^{T_r}$, and the Jacobian, G_r^U , which is some scaling of G^U since the reference star may have a different spectrum and intensity than the target. The electric field modes, on the other hand, are properties of the instrument, i.e., shared between observations. It is therefore possible to define a joint cost function based on Eq. (4.16) for both target and reference observations,

$$\begin{aligned}
J + J_r &= J_{joint} \left(G^V, \{\hat{\mathbf{v}}_k\}_{k=1}^T, \{\hat{\mathbf{w}}_k\}_{k=1}^T, \{\hat{\mathbf{v}}_k^r\}_{k=1}^{T_r}, \{\hat{\mathbf{w}}_k^r\}_{k=1}^{T_r} \right), \\
J &= \mathbf{1} \cdot \sum_{k=1}^T \left(B \cdot \left(\hat{G}^V \hat{\mathbf{v}}_k + G^U \mathbf{u}_k \right)^{\circ 2} + \hat{\mathbf{I}}_k^I - \mathbf{y}_k \circ \log \left(B \cdot \left(\hat{G}^V \hat{\mathbf{v}}_k + G^U \mathbf{u}_k \right)^{\circ 2} + \hat{\mathbf{I}}_k^I \right) \right), \\
J_r &= \mathbf{1} \cdot \sum_{k=1}^{T_r} \left(B \cdot \left(\hat{G}^V \hat{\mathbf{v}}_k^r + G_r^U \mathbf{u}_k^r \right)^{\circ 2} + \hat{\mathbf{I}}_k^{r,I} - \mathbf{y}_k^r \circ \log \left(B \cdot \left(\hat{G}^V \hat{\mathbf{v}}_k^r + G_r^U \mathbf{u}_k^r \right)^{\circ 2} + \hat{\mathbf{I}}_k^{r,I} \right) \right).
\end{aligned} \tag{4.17}$$

J_{joint} is then optimized with respect to all of its parameters which include the coefficients history for both observations.

The term, J_r , can be thought of as an additional regularization that prevents G^V from overfitting the data (see Fig. 4.2). Interestingly, the nature of Poisson distribution makes the measurements from a brighter reference star, $\{\mathbf{y}_k^r\}_{k=1}^{T_r}$, have a proportionally larger weight in J_{joint} which is consistent with them having a higher S/N.

4.2 Numerical Results

Below, the observation scenario described in Sec. 2.3 is used to compare intensity-based post-processing (PCA/KLIP and ADI) with a probabilistic formulation in terms of the electric field (previous section). To this end, a planet with an intensity of $1 \frac{\text{photon}}{\text{frame}}$ was added at $4.5 \lambda/D$ and the errors in its estimate were computed inside its half-




Figure 4.1: Post-processing results for the observation scenario described in Sec. 2.3 (only the $\lambda = 546$ nm channel). The added planet at $4.5 \lambda/D$ is best visible in the top three images on the left, corresponding to a closed-loop observation scenario with no simulated jitter. The highest accuracy was achieved by EFOR with[57] and without[103] telescope rolls, followed by maximum *a-posteriori* (MAP)[49] then ADI[100] and PCA[87]. Open-loop wavefront drift or line-of-sight jitter, each had a negative impact on accuracy of all methods.

max region (i.e., at pixels where the intensity of the planet was at least half of its maximum intensity). Scaling this error by the half-max error of a PSF subtraction in a “perfect” observation scenario (a stable dark hole), gives the relative post-processing factor (PPF),

$$\text{rel. PPF} = \frac{\text{avg. post-processing error}}{\text{avg. error of perfect PSF subtraction}}. \quad (4.18)$$

Since the ideal PSF subtraction is limited by the variance of photon arrival events (shot-noise), the relative PPF has a theoretical lower limit of 1.

The effects of telescope rolls were simulated by rotating the PSF of the planet about the center of the image, once per observation for EFOR, and once every 2

hours for ADI. PCA was given 216 reference images from a simulated magnitude 2 star. The jitter was simulated as sinusoidal perturbations in the tip/tilt Zernike modes,

$$z^{1,-1}(t) - \langle z^{1,-1} \rangle_k = a_k^{-1} \sin\left(\frac{2\pi}{\Delta t} t\right), \quad t \in [t_k, t_{k+1}), \quad (4.19)$$

$$z^{1,+1}(t) - \langle z^{1,+1} \rangle_k = a_k^{+1} \sin\left(\frac{2\pi}{\Delta t} t + \phi_k\right), \quad t \in [t_k, t_{k+1}), \quad (4.20)$$

with $a_k^{\pm 1}$ slowly varying between 0 and 1.4nm (the residual error budget for WFIRST-CGI[62]) and ϕ_k between 0 and 2π .

Figure 4.1 shows post-processing results based on the $\lambda = 546$ nm channel of the images simulated in the open-loop and closed-loop scenarios in Sec. 2.3. The best relative post-processing factor of about 1.3 was achieved with EFOR with telescope rolls, Eq. (4.16), in the absence of line-of-sight jitter. The errors of maximum *a-posteriori*, Eq. (4.8), EFOR without rolls, Eq. (4.12), and ADI were about 50% higher while that of PCA was almost 3 times higher.

When jitter was included in the simulation, MAP was the most severely affected since it cannot distinguish between jitter and other incoherent sources. EFOR and ADI became at least 4 times less accurate and PCA was not significantly affected by jitter. EFOR with telescope rolls remained the best performing methods with a relative PPF of about 3.

In the open-loop scenario (without jitter), both PCA and ADI had a relative PPF of about 6. However, in a realistic scenario, the dark hole might not be recreated with as high contrast as depicted in Fig. 2.1. Additionally, changing the telescope's orientation and pointing every 10 and 2 hours respectively, might excite wavefront errors and increase jitter[48]. These effects were not simulated here but would presumably increase the post-processing error even further.




Figure 4.2: (a) Relative post-processing factors, Eq. 4.18, of EFOR with and without reference observations and intensity-based PCA. The PPF of EFOR are plotted against the number of electric field modes, r , and those of PCA are plotted against $\frac{r(r+1)}{2}$. (b) The cost function of EFOR and the truncation singular value, σ , of the speckle basis used for PCA. Both EFOR with reference images and PCA suggest that there are $r \geq 12$ dominant electric field modes present in the data. *adapted from [103]*

4.2.1 Effects of Dimensionality

While the speckle electric field may lie in a low-dimensional subspace, the dimension itself might not be known. Figure 4.2(a) shows that without reference images, EFOR may overfit the data if r is assumed too large. Its error is smaller than that of PCA-based post-processing but has a similar growth rate with a caveat: EFOR's error corresponding to r modes should be compared to PCA's error corresponding to $\frac{r(r+1)}{2}$ modes.

Indeed, if the electric field, \mathbf{E}^{OL} lies in an r dimensional subspace of \mathbb{R}^{2n} , then the intensity $\mathbf{I}^S = B \cdot \mathbf{E}^{o2}$ lies in an $\binom{r}{2} = \frac{r(r+1)}{2}$ dimensional subspace of \mathbb{R}^n (a worst case). This implies that the number of free parameters required to fit the same speckle patterns is significantly lower when formulating the problem in terms of the electric field, and helps explain the lower error of EFOR. Overall, EFOR was between 1.8 and 2.5 times more accurate than PCA depending on whether reference data was used.

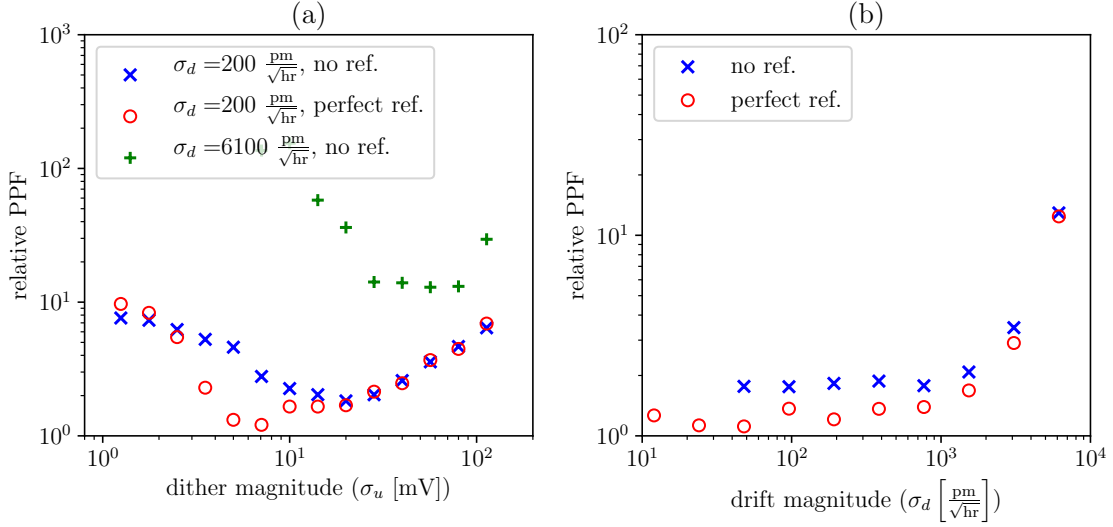


Figure 4.3: (a) Relative post-processing factors, Eq. 4.18, of EFOR with and without a “perfect” guess of the electric field modes, G^V , as a function of the dither magnitude, σ_u , for two drift magnitudes, σ_d . The higher the drift, the more dithering is necessary to reach an optimal PPF. (b) The relative PPF for an optimal choice of dithering magnitude as a function of drift magnitude, σ_d . *from* [60]

Reference data can also provide a guess for the number of significant electric field modes. Figure 4.2(b) shows that the decrease in the cost function, J_{joint} in Eq. (4.17), slows down after about $r = 12$, which also corresponds to the most accurate estimate of EFOR without reference images.

4.2.2 Effects of Drift Magnitude

Finally, the relation between drift magnitude and post-processing factor of EFOR is depicted in Fig. 4.3. In this case, the presence of reference images was simulated by having a “perfect” initial guess for the electric field modes, G^V , before the optimization procedure began.

Similarly to Fig. 2.4(a), the optimal dithering magnitude, this time in terms of the PPF, depends on the magnitude of the drift – the faster the drift, the more dithering was required to accurately estimate the state of the speckles. Also, similarly

to Fig. 2.4(b), closing the loop with optimal dithering provides steady contrast and post-processing accuracy for a wide range of drift values.

It should be noted however, that EFOR breaks down when DM actuator drift becomes dominant since its impact is not low-dimensional as required per Eq. (4.9). In that case a closed-loop observation with at least one telescope roll combined with ADI gives the best PPF.

Chapter 5

Future Directions

5.1 EFOR for Ground-based Coronagraphy

During ground-based observations, the interaction of atmospheric turbulence with Adaptive Optics (AO) results in small unknown speckle perturbation of high spatial variability. It has been previously speculated that this variability increases the S/N in post-processing [106], although it was later proved not to be the case in practice[107]. As shown in the previous chapter for the case of space-based coronagraphs, closing the loop does reduce post-processing error, but the history of DM controls needs to be taken into account to get the full advantage out of it.

Recently, methods which incorporate wavefront sensor data[108] and known DM ripples[102] have been proposed. Although suggested to improve the post-processing factor of ground-based telescopes, the history of DM controls does not explicitly appear in their estimates. In fact, the control history is usually not even stored due to the extreme rapidity of AO cycles.

However, if images with short exposures, corresponding AO controls and their influence function (the Jacobian) are available, the data resembles the inputs of the algorithms presented in chapter 4. Incorporating all of the available information might

bring ground based observations somewhat closer to their theoretical shot noise limit without modifying currently operating instruments. This approach will be entertained in the near future using numerical models of ground based coronagraphs[109, 110].

5.2 Low-order Coordinates Extraction via Debiased Isomaps¹

The emphasis in chapter 4 was on finding electric field modes, $G^V \in \mathbb{R}^{2n \times r}$, which described the evolution of some measurements *given* a non-linear model of how to construct images. The coefficients of these modes, $\mathbf{v}_k \in \mathbb{R}^r$, were of little interest as they were assumed to arise from a random walk of wavefront errors. In a broader context of non-linear dimensionality reduction[111], these coefficients can be thought of as “coordinates” or “latent variables” of an r -dimensional manifold embedded in an n -dimensional space of speckle intensities. Theoretically, one could learn this low-dimensional representation (in the above case - the description of intensities in terms of electric fields) to gain physical insight[112] into and allow controlling of the underlying system[113].

As an example, Fig. 5.1(a) shows frames from a video (i.e., a sequence of noisy images) of a physical pendulum. A natural choice of a latent variable, $v(t)$, to describe its frames is the angle of the pendulum whose evolution is shown in Fig. 5.1(c). However, if the trajectory of v was not given, would it be possible to “learn” it from just those video frames?

In this particular example, the task of extracting the angle of the pendulum yields to standard image-processing techniques[115] since the object of interest is unobscured and has clear boundaries. If the measurements were spatially “smooth”, as is the

¹ A code example is available at
<https://github.com/leonidprinceton/debiased-Isomaps>

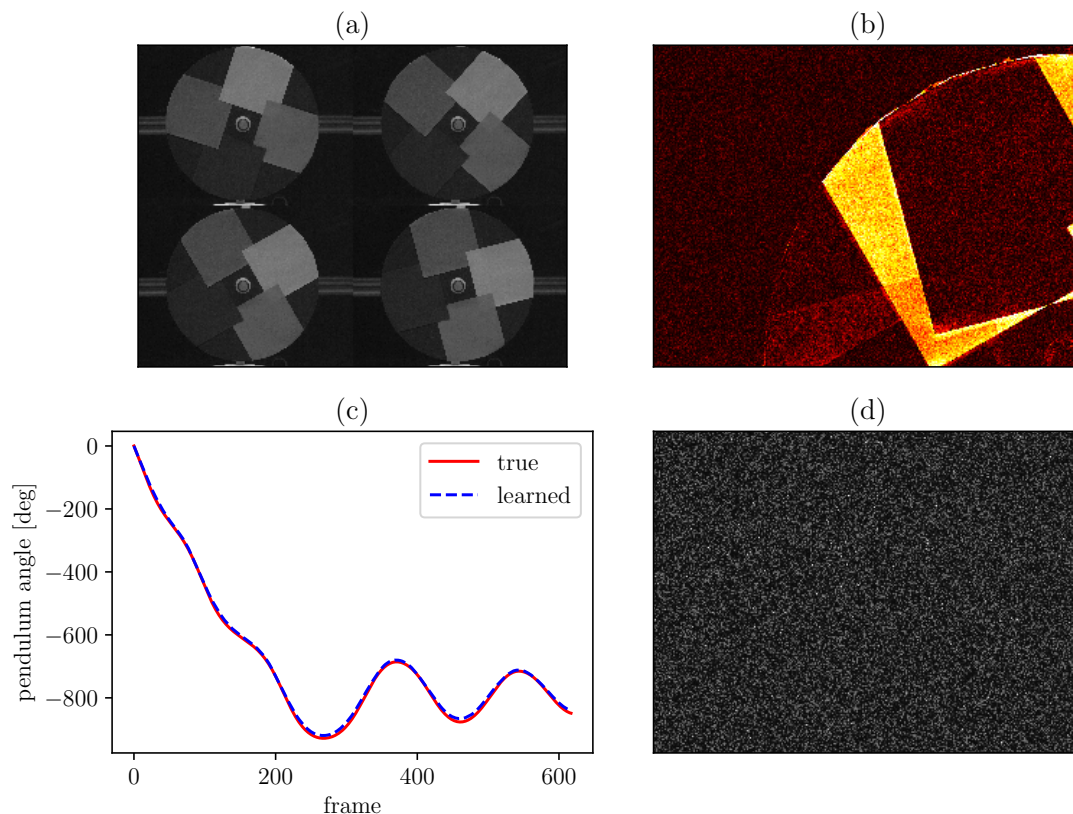


Figure 5.1: (a) Four frames from a slow-motion video of a physical pendulum. (b) Absolute values of pixel-wise differences between two frames with a clearly visible grain noise. (c) The angle of the pendulum (solid red line) as a function of the video frame number. The method proposed in this section “learns” coordinates which (almost) lie on a circle in \mathbb{R}^2 . Their natural parameterization by an angle is shown in dashed blue line. (d) A frame from the pendulum video with pixels randomly shuffled around. The method is equally applicable to this case. *adapted from [114]*
; the video of the pendulum was taken by D. Feng

case with speckles in high contrast imaging, the more natural approach would be projection-based methods[116]. However, both approaches would fail in more general cases, for example, if the pixels of the video were not given in order (i.e., randomly shuffled as in Fig. 5.1(d)). In that case, one would have to resort to unsupervised non-linear dimensionality reduction methods[111], e.g. Locally Linear Embedding[117] or Isomaps[118].

Unsupervised learning of coordinates can be done either with the dynamics[119, 120, 121] (and even controls[122, 123]) in mind from the very beginning, or by extending existing methods[124, 125]. It has been tested on complex datasets such as bio-medical data[126, 127], facial expressions[128, 129] and human motion[130]. Yet, it is difficult to find an example of physically meaningful coordinates extracted from a real video of a simple system, e.g. the pendulum in Fig. 5.1.

It is worth noting that a major motivation for “learning” a system is to autonomously control it. In many cases it is possible to do so by directly interacting with the system (known as “reinforcement learning”) and without explicitly identifying the underlying latent variables[131, 113, 132, 133] (usually with the help of neural networks[134]). Still, it is possible to incorporate low-order constraints, e.g. via autoencoders[135], to enhance the learning rate and performance for such tasks[136, 137, 138].

In particular, autoencoders were successfully used to identify physical coordinates and dynamics from videos[139, 140]. Yet, they require large amounts of training data which, so far, has been simulated rather than real (interestingly, the objects in these simulated videos have “soft” edges, unlike in Fig. 5.1). Below, Isomaps[118] is used to identify the coordinate(s) describing the aforementioned video of a pendulum with very little assumptions about the “measurements” and no external information or training data.

Isomaps is an unsupervised metric-based dimensionality reduction method. As an input, it accepts a matrix of pairwise distances,

$$D = [d(\mathbf{y}_k, \mathbf{y}_l)]_{1 \leq k, l \leq T}, \quad (5.1)$$

between some general data points $\{\mathbf{y}_k\}_{k=1}^T$. Only a fraction of the smallest distances is deemed to be “related to” or “representative of” the latent coordinates (there are numerous criteria for selecting these “representative” distances, some discussed below). The data points that are considered to be close to each other in that sense, are called “connected”. The discarded distances between “disconnected” points are then recomputed as the distances along the shortest paths through the connected points. These are known as the “geodesic distances” (or simply “distances” in graph theory[141]),

$$D_{geo} = [d_{geo}(\mathbf{y}_k, \mathbf{y}_l)]_{1 \leq k, l \leq T}.$$

The hope is that the geodesic distances, D_{geo} , are not significantly affected by the non-linearities of the space in which the data (\mathbf{y}) resides, and are therefore more representative of distances along the submanifold corresponding to the latent variables (\mathbf{v}). The final step of Isomaps is to assign the data points coordinates in a low-dimensional Euclidean space via Multidimensional scaling (MDS)[142] of D_{geo} . The subtleties that arise when applying this procedure to data from a real video are discussed next.

5.2.1 Choice of the Distance Metric

Even if the intensity measurements (images) are precise, the metric used by Isomaps may not be “consistent” with the latent dynamics. To illustrate this with a simple analytical example, consider the following dependency of the “infinite-resolution” image,




Figure 5.2: A “bowl” shaped object described by Eq. (5.2). The latent variable, v , describes the position of the object while the measurements are the images of the object, $I(x)$. Small changes in the position, Δv , are proportional to the shaded area representing the difference between two images shifted by Δv . Because of the sharp discontinuities in the object’s shape, this area can be computed using the L^1 but not the L^2 norm. *adapted from [114]*

$I : \mathbb{R} \rightarrow \mathbb{R}$, on the unidimensional latent variable, v ,

$$I(x) = \begin{cases} (x - v)^2 & x \in [v - 1, v + 1] \\ 0 & \text{otherwise} \end{cases}. \quad (5.2)$$

This image corresponds to a moving object with sharp edges as illustrated in Fig. 5.2.

Intuitively, a consistent metric on the space of images would manifest small changes in the latent variable in a fashion that is *independent of the rate of change of the variable*. Put differently, the metric d should be approximately proportional to small changes in v ,

$$d(I(v), I(v + \Delta v)) \approx \kappa |\Delta v|, \quad (5.3)$$

for small Δv , where $\kappa > 0$ may depend on v .

An example of a metric that is “inconsistent” with Eq. (5.2) is the one induced by the L^2 norm, since

$$\begin{aligned}\|I(v) - I(v + \Delta v)\|_2 &= \left(\int_{-\infty}^{\infty} (I(v - x) - I(v + \Delta v - x))^2 dx \right)^{\frac{1}{2}} = \\ &= \left(2|\Delta v| + \frac{8}{3}\Delta v^2 + \dots \right)^{\frac{1}{2}} = \Theta(|\Delta v|^{\frac{1}{2}}).\end{aligned}$$

This $\sqrt{|\Delta v|}$ dependency will result in a non-diminishing local curvature “computed” by Isomaps even as the data points get closer to each other. Indeed, let $\delta v \ll 1$ and consider the three images, $\{I(-\delta v), I(0), I(\delta v)\}$. The L^2 distances between these images are

$$\begin{aligned}\|I(-\delta v) - I(0)\|_2 &\approx \sqrt{2}\sqrt{\delta v}, \\ \|I(-\delta v) - I(\delta v)\|_2 &\approx 2\sqrt{\delta v}, \\ \|I(0) - I(\delta v)\|_2 &\approx \sqrt{2}\sqrt{\delta v}.\end{aligned}$$

Yet, any three points with the above distances cannot lie on a straight line no matter how small δv is (how fine the data is sampled). The output of Isomaps will therefore suggest that the latent variable lies in a Euclidean space of dimension two or higher, even though $v \in \mathbb{R}^1$ by construction.

On the other hand, the L^1 norm gives

$$\|I(v) - I(v + \Delta v)\|_1 = \int_{-\infty}^{\infty} |I(v - x) - I(v + \Delta v - x)| dx = \Theta(|\Delta v|),$$

and Isomaps will correctly identify some affine transformation of v as the latent variable, as long as the sampling is fine enough and no noise was added to I . The

technique described next, provides a way for estimating the impact of noise as well as testing whether a given metric is consistent.

5.2.2 Estimation of Metric Bias

When measurement noise is present (e.g. Fig. 5.1(b)), two identical states, $\mathbf{v}_k = \mathbf{v}_l$, correspond to the same images, $\mathbf{I}(\mathbf{v}_k) = \mathbf{I}(\mathbf{v}_l)$, but different measurements, $\mathbf{y}_k \neq \mathbf{y}_l$. Therefore, regardless of the choice of metric, d ,

$$d(\mathbf{y}_k, \mathbf{y}_l) > d(\mathbf{I}(\mathbf{v}_k), \mathbf{I}(\mathbf{v}_l)),$$

or

$$\mathbb{E}\{d(\mathbf{y}_k, \mathbf{y}_l) - d(\mathbf{I}(\mathbf{v}_k), \mathbf{I}(\mathbf{v}_l))\} = b > 0, \quad (5.4)$$

where $\mathbb{E}\{\cdot\}$ denotes expectation (across multiple realizations of the measurement). This bias in all the distances causes Isomaps to overestimate the local curvature (which manifests itself similarly to the metric “inconsistency” in the previous example).

Assuming that the bias, b , does not depend on the coordinate, \mathbf{v} , is it possible to estimate it from measurements, $\{\mathbf{y}_k\}$, alone? In the particular case of $\{\mathbf{y}_k\}$ being generated by a smoothly varying latent variable, the answer depends on the metric d and the sampling of $\{\mathbf{v}_k\}$. If the trajectory of \mathbf{v} is smooth and the sampling is “fine” enough, nearby points on the trajectory “almost” lie on a line, in which case

$$\|\mathbf{v}_{k+m} - \mathbf{v}_k\| \approx \sum_{i=1}^m \|\mathbf{v}_{k+i} - \mathbf{v}_{k+i-1}\|.$$

Additionally, if the metric d is consistent in the sense of Eq. (5.3), then

$$d(\mathbf{I}(\mathbf{v}_{k+m}), \mathbf{I}(\mathbf{v}_k)) \approx \sum_{i=1}^m d(\mathbf{I}(\mathbf{v}_{k+i}), \mathbf{I}(\mathbf{v}_{k+i-1})),$$




Figure 5.3: An empirical method for testing the consistency a metric and computing its bias, b . It is based on a matrix of distances $D = [D_{i,j}]$ and the assumption that adjacent images (i close to j) belong to adjacent points on a smooth trajectory of a latent variable. If true, Eq. (5.5) suggests that $\frac{1}{k-1} \sum_{j=i}^{i+k-1} D_{j+1,j}$ should be proportional to $\frac{1}{k-1} D_{i+k,i} + b$ for most i and small k . (a) This is indeed the case for the pendulum video images from Fig. 5.1 when the metric is induced by the l^1 norm. It allows estimating the bias via linear regression (in this case using Random Sample Consensus[143] to reject the visible outliers). (b) The l^2 norm is inconsistent with the dynamics of the pendulum, in the sense of (5.3), and Isomaps will not reproduce the local curvature of the latent variable (angle) correctly. *from* [114]

which, in light of Eq. (5.4), gives

$$\begin{aligned}
 b - \mathbb{E} \{d(\mathbf{y}_{k+m}, \mathbf{y}_k)\} &= 0 = \\
 &= \sum_{i=1}^m (b - \mathbb{E} \{d(\mathbf{y}_{k+i}, \mathbf{y}_{k+i-1})\}) = mb - \sum_{i=1}^m \mathbb{E} \{d(\mathbf{y}_{k+i}, \mathbf{y}_{k+i-1})\}. \quad (5.5)
 \end{aligned}$$

Although one does not have access to multiple realizations of \mathbf{y}_k for any single k , it is possible to estimate b based on all available measurements. That is, one needs to find b such that

$$b \approx \frac{\sum_{i=1}^m \mathbb{E} \{d(\mathbf{y}_{k+i}, \mathbf{y}_{k+i-1})\} - \mathbb{E} \{d(\mathbf{y}_{k+m}, \mathbf{y}_k)\}}{m-1}$$

is accurate for most k .

This is illustrated in Fig. 5.3 for various choices of m and two metrics on the pendulum video images induced by the l^1 and l^2 norms. Eq. (5.5) also acts as an empirical consistency test, which holds well for small m and just the l^1 norm. In light of the example in Fig. 5.2, a possible explanation for the failure of l^2 is the existence of sharp discontinuities between adjacent sections of the pendulum.

5.2.3 Choice of Neighborhood Criterion

It remains to describe a mechanism for choosing what distances in D in Eq. (5.1) that are considered “representative” of the latent variables. One common approach is to pick a fixed number of nearest neighbors of data points \mathbf{y}_k . However, if the data comes from a finely sampled trajectory (which is a prerequisite for bias estimation), most of the nearest neighbors are also likely to be *temporally* close to each other.

In the worst case, which occurs when picking just two nearest neighbors, the pendulum data is mapped by Isomaps into a trajectory in \mathbb{R}^1 . This suggests that in order for Isomaps to deduce the correct topology from smooth-trajectories data, the choice of neighbors must include points on “returning trajectories” (sometimes also called “nearest nontrivial neighbors”[124]).

The ε -neighborhood criterion which keeps all distances below some threshold, ε , is arguably more fitting for this purpose. In the context of smoothly varying latent variables, ε must be large enough to include “returning trajectories” but small enough such that the consistency criterion, (5.3), remains valid (otherwise, second order terms in the image function, \mathbf{I} , will be attributed to the coordinates, \mathbf{v}).

Figure 5.4(a) provides a possible way for choosing ε . The MDS invoked at the last step of Isomaps provides eigenvalues, $\{\lambda_1, \lambda_2, \dots\}$, which correspond to the relative importance of the “learned” modes that describe the data. The ratios of these eigenvalues can be used to infer the dimensionality of the latent space and, if a low-dimensional representation of the data exists, should not be sensitive to other parameters. Indeed,

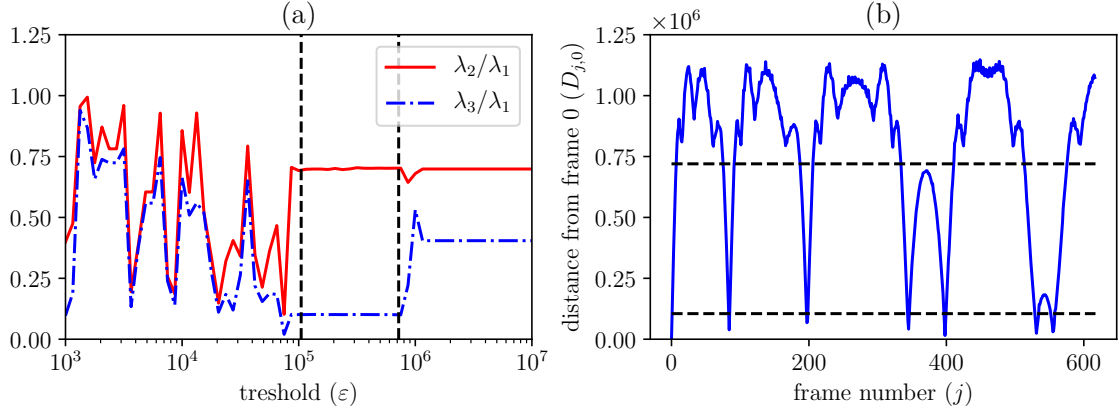


Figure 5.4: A technique for choosing the threshold, ε , for the nearest-neighbors criterion of Isomaps (a) The ratios of the first three eigenvalues computed by MDS (the last step of Isomaps) for various values of ε . The ratios are insensitive to the exact threshold value in a region which spans almost an order of magnitude beginning at 10^5 . (b) This “insensitivity” region corresponds to thresholds which include all “returning” trajectories and excludes far away regions of the latent space where the metric depends non-linearly on the angle. This is illustrated by the distance, $D_{j,0}$, between each j th and the zeroth frames (using the l^1 norm). *from* [114]

for a wide range of thresholds in Fig. 5.4(a), the ratios $\lambda_1/\lambda_{2,3}$ remain almost fixed. This is a consequence of the wide range of angle differences which are proportional to the distances between corresponding images (Fig. 5.4(b)).

5.2.4 Results

With a “consistent” and debiased metric choice and a reasonable neighbors criterion, Isomaps yields a smooth trajectory on what appears to be a circle in \mathbb{R}^2 (Fig. 5.5(a)). Assigning a phase to each point on that circle, gives a shifted approximation of the true angle of the pendulum (dashed blue line in Fig. 5.1(c)). On the other hand, using the l^2 norm or a significantly different ε , results in spurious oscillation or a wrong topology altogether ((Fig. 5.5(b,c,d)).




Figure 5.5: (c) The output of Isomaps with the debiased distance based on the l^1 norm and $\varepsilon = 3 \cdot 10^5$. The coordinates lie close to a circle in \mathbb{R}^2 which is a “natural” embedding of the pendulum angle in a Euclidean space. (d) When the threshold is too small, some “returning” trajectories are not “captured” by Isomaps resulting in a wrong topology of the coordinates. When the threshold is too large, (e), or the norm is inconsistent, (f), Isomaps’ approximation of the local curvature is inaccurate and gives rise to oscillatory artifacts. *from* [114]

It the future, this method will be used to find the equations of motion of the pendulum and possibly other systems. This will be done with the help of algorithms for automated discovery of dynamics from trajectories[144, 145, 146].

Chapter 6

Conclusions

This thesis presents a theoretical approach to mitigating the effects of wavefront instabilities in space-based high-contrast imaging. Chapter 2 describes the potential benefits of continuous Focal Plane Wavefront Sensing and Control (FPWSC) during lengthy observations. Specifically, it allows maintaining a fixed contrast at a significantly higher level than the obtained by periodically “re-creating” a dark hole. As a consequence, it reduces the number of necessary telescope maneuvers, avoiding potential structural and thermal instabilities as well as increased line-of-sight jitter.

Closed loop maintenance of the contrast relies on non-linear control and estimation of the electric field in the dark hole. The coupling between the two manifests itself in the necessity to dither the deformable mirrors to generate enough phase diversity for the estimator to remain stable. This dither slightly decreases the contrast in the dark hole but the closed loop approach appears to be worthwhile for a wide range of wavefront drift rates. It will be subject to numerical optimization and lab experiments in the near future[147].

The accuracy of the electric field estimator, and therefore the stability of the contrast, can be increased by incorporating a more detailed model of the evolution of the electric field. This can be achieved by accounting for the “smoothness” of the

speckles, interpreted as cross-pixel correlation between the increments of the electric field. Even simply bundling adjacent pixels together into subsystems, significantly increases the accuracy of the estimates at each individual pixel.

However, the computation time of the optimal estimator (the Kalman Filter) grows proportionally to the cube of the number of pixels per subsystem. The sub-optimal filter described in chapter 3 takes advantage of the potentially low-order wavefront errors while doing so in linear time complexity. It keeps track of the error covariances only within each subsystem (the main block-diagonal of the full covariance matrix), but advances each estimate based on all available measurements. Although slower than treating each subsystem separately, its accuracy becomes close to that of the full Kalman Filter under certain conditions. These conditions seem to hold for a space-based coronagraph model whose wavefront drift is well described by variations in just tens of Zernike modes.

A low dimensional model of wavefront drift, if true, becomes especially advantageous in post-processing. The number of measurements per observation is proportional to the number of detectors (thousands) and exposures (hundreds), and the number of unknown system parameters is at least three times as large. Chapter 4 formulates the post-processing problem in terms of electric fields and then constrains them to a low-dimension subspace. This greatly reduces the number of free parameters and, under favorable conditions, results in a numerical scheme which gets close to the theoretical (shot-noise) error limit. Unlike intensity-based order-reduction methods such as KLIP, a formulation in terms of electric fields makes the problem non-linear but also allows incorporating the history of controls. This is especially helpful in a closed-loop observation scenario since the effects of deformable mirrors on the dark hole intensity are usually high dimensional.

Although the results presented here are based on a simulation of the WFIRST-CGI, they are applicable to all future telescopes capable of FPWSC, e.g. LUVOIR.

In fact, as the contrast of space coronagraphs increases, they will become more susceptible to wavefront drift, to the point where stability issues may become a major driving factor behind their cost. The ideas presented in this thesis will hopefully help alleviate some of those issues.

Bibliography

- [1] B. Lyot. The study of the solar corona and prominences without eclipses. *Monthly Notices of the Royal Astronomical Society*, 99:580, 1939.
- [2] P. Kalas, J. R. Graham, E. Chiang, M. P. Fitzgerald, M. Clampin, E. S. Kite, K. Stapelfeldt, C. Marois, and J. Krist. Optical images of an exosolar planet 25 light-years from earth. *Science*, 322(5906):1345–1348, 2008.
- [3] C. Thalmann, J. Carson, M. Janson, M. Goto, M. McElwain, S. Egner, M. Feldt, J. Hashimoto, Y. Hayano, T. Henning, K. W. Hodapp, R. Kandori, H. Klahr, T. Kudo, N. Kusakabe, C. Mordasini, J.-I. Morino, H. Suto, R. Suzuki, and M. Tamura. Discovery of the coldest imaged companion of a sun-like star. *The Astrophysical Journal*, 707(2):L123–L127, 2009.
- [4] B. Macintosh, J. R. Graham, P. Ingraham, Q. Konopacky, C. Marois, M. Perrin, L. Poyneer, B. Bauman, T. Barman, A. S. Burrows, A. Cardwell, J. Chilcote, R. J. De Rosa, D. Dillon, R. Doyon, J. Dunn, D. Erikson, M. P. Fitzgerald, D. Gavel, S. Goodsell, M. Hartung, P. Hibon, P. Kalas, J. Larkin, J. Maire, F. Marchis, M. S. Marley, J. McBride, M. Millar-Blanchaer, K. Morzinski, A. Norton, B. R. Oppenheimer, D. Palmer, J. Patience, L. Pueyo, F. Rantakyro, N. Sadakuni, L. Saddlemyer, D. Savransky, A. Serio, R. Soummer, A. Sivaramakrishnan, I. Song, S. Thomas, J. K. Wallace, S. Wiktorowicz, and S. Wolff. First light of the Gemini Planet Imager. *Proceedings of the National Academy of Sciences*, 111(35):12661–12666, 2014.
- [5] G. Chauvin, S. Desidera, A.-M. Lagrange, A. Vigan, R. Gratton, M. Langlois, M. Bonnefoy, J.-L. Beuzit, M. Feldt, D. Mouillet, M. Meyer, A. Cheetham, B. Biller, A. Boccaletti, V. D’Orazi, R. Galicher, J. Hagelberg, A.-L. Maire, D. Mesa, J. Olofsson, M. Samland, T. O. B. Schmidt, E. Sissa, M. Bonavita, B. Charnay, M. Cudel, S. Daemgen, P. Delorme, P. Janin-Potiron, M. Janson, M. Keppler, H. L. Coroller, R. Ligi, G. D. Marleau, S. Messina, P. Mollière, C. Mordasini, A. Müller, S. Peretti, C. Perrot, L. Rodet, D. Rouan, A. Zurlo, C. Dominik, T. Henning, F. Menard, H.-M. Schmid, M. Turatto, S. Udry, F. Vakili, L. Abe, J. Antichi, A. Baruffolo, P. Baudoz, J. Baudrand, P. Blanchard, A. Bazzon, T. Buey, M. Carillet, M. Carle, J. Charton, E. Cascone, R. Claudi, A. Costille, A. Deboulbe, V. D. Caprio, K. Dohlen, D. Fantinel, P. Feautrier, T. Fusco, P. Gigan, E. Giro, D. Gisler, L. Gluck, N. Hubin, E. Hugot, M. Jaquet, M. Kasper, F. Madec, Y. Magnard, P. Martinez,

- D. Maurel, D. L. Mignant, O. Möller-Nilsson, M. Llored, T. Moulin, A. Origné, A. Pavlov, D. Perret, C. Petit, J. Pragt, P. Puget, P. Rabou, J. Ramos, R. Rigal, S. Rochat, R. Roelfsema, G. Rousset, A. Roux, B. Salasnich, J.-F. Sauvage, A. Sevin, C. Soenke, E. Stadler, M. Suarez, L. Weber, F. Wildi, S. Antonucci, J.-C. Augereau, J.-L. Baudino, W. Brandner, N. Engler, J. Girard, C. Gry, Q. Kral, T. Kopytova, E. Lagadec, J. Milli, C. Moutou, J. Schlieder, J. Szulágyi, C. Thalmann, and Z. Wahhaj. Discovery of a warm, dusty giant planet around HIP 65426. *Astronomy & Astrophysics*, 605:L9, 2017.
- [6] NASA Exoplanet Archive. Confirmed planets table, 2019.
- [7] D. Charbonneau, T. M. Brown, D. W. Latham, and M. Mayor. Detection of planetary transits across a sun-like star. *The Astrophysical Journal*, 529(1):L45–L48, 2000.
- [8] A. Udalski, B. Paczynski, K. Zebrun, M. Szymanski, M. Kubiak, I. Soszynski, O. Szewczyk, L. Wyrzykowski, and G. Pietrzynski. The optical gravitational lensing experiment. search for planetary and low-luminosity object transits in the galactic disk. results of 2001 campaign. *arXiv preprint astro-ph/0202320*, 2002.
- [9] M. Mayor and D. Queloz. A Jupiter-mass companion to a solar-type star. *Nature*, 378(6555):355–359, 1995.
- [10] D. Charbonneau, T. M. Brown, R. W. Noyes, and R. L. Gilliland. Detection of an extrasolar planet atmosphere. *The Astrophysical Journal*, 568(1):377–384, 2002.
- [11] N. Madhusudhan. Exoplanetary atmospheres: Key insights, challenges, and prospects. *Annual Review of Astronomy and Astrophysics*, 57(1):617–663, 2019.
- [12] A. Tsiaras, I. P. Waldmann, G. Tinetti, J. Tennyson, and S. N. Yurchenko. Water vapour in the atmosphere of the habitable-zone eight-earth-mass planet K2-18 b. *Nature Astronomy*, 2019.
- [13] B. Macintosh, J. Graham, D. Palmer, R. Doyon, D. Gavel, J. Larkin, B. Oppenheimer, L. Saddlemyer, J. K. Wallace, B. Bauman, J. Evans, D. Erikson, K. Morzinski, D. Phillion, L. Poyneer, A. Sivaramakrishnan, R. Soummer, S. Thibault, and J.-P. Veran. The Gemini Planet Imager. In *Advances in Adaptive Optics II*, volume 6272, pages 177 – 188. International Society for Optics and Photonics, SPIE, 2006.
- [14] B. R. Oppenheimer and S. Hinkley. High-contrast observations in optical and infrared astronomy. *Annual Review of Astronomy and Astrophysics*, 47(1):253–289, 2009.
- [15] G. Chauvin, A.-M. Lagrange, C. Dumas, B. Zuckerman, D. Mouillet, I. Song, J.-L. Beuzit, and P. Lowrance. A giant planet candidate near a young brown dwarf. *Astronomy & Astrophysics*, 425(2):L29–L32, 2004.

- [16] A.-M. Lagrange, D. Gratadour, G. Chauvin, T. Fusco, D. Ehrenreich, D. Mouillet, G. Rousset, D. Rouan, F. Allard, É. Gendron, J. Charton, L. Mugnier, P. Rabou, J. Montri, and F. Lacombe. A probable giant planet imaged in the β pictoris disk. *Astronomy & Astrophysics*, 493(2):L21–L25, 2008.
- [17] R. Belikov, A. Give'on, B. D. Kern, E. Cady, M. Carr, S. B. Shaklan, K. Balasubramanian, V. White, P. Echternach, M. Dickie, J. Trauger, A. Kuhnert, and N. J. Kasdin. Demonstration of high contrast in 10% broadband light with the shaped pupil coronagraph. In *Techniques and Instrumentation for Detection of Exoplanets III*, volume 6693, page 66930Y. International Society for Optics and Photonics, 2007.
- [18] L. Pueyo, J. Kay, N. J. Kasdin, T. Groff, M. McElwain, A. Give'on, and R. Belikov. Optimal dark hole generation via two deformable mirrors with stroke minimization. *Applied Optics*, 48(32):6296, 2009.
- [19] C. C. Stark, A. Roberge, A. Mandell, M. Clampin, S. D. Domagal-Goldman, M. W. McElwain, and K. R. Stapelfeldt. Lower limits on aperture size for an exoearth detecting coronagraphic mission. *The Astrophysical Journal*, 808(2):149, 2015.
- [20] J. W. Hardy. *Adaptive optics for astronomical telescopes*, volume 16. Oxford University Press on Demand, 1998.
- [21] O. Guyon. Extreme adaptive optics. *Annual Review of Astronomy and Astrophysics*, 56(1):315–355, 2018.
- [22] O. Guyon. Limits of adaptive optics for high-contrast imaging. *The Astrophysical Journal*, 629(1):592–614, 2005.
- [23] C. Cavarroc, A. Boccaletti, P. Baudoz, T. Fusco, and D. Rouan. Fundamental limitations on earth-like planet detection with extremely large telescopes. *Astronomy & Astrophysics*, 447(1):397–403, 2006.
- [24] L. Pueyo. *Direct Imaging as a Detection Technique for Exoplanets*, pages 705–765. Springer International Publishing AG, 6330 Cham, Switzerland, 2018.
- [25] N. Jovanovic, O. Absil, P. Baudoz, M. Beaulieu, M. Bottom, E. Cady, B. Carlomagno, A. Carlotti, D. Doelman, K. Fogarty, R. Galicher, O. Guyon, S. Haffert, E. Huby, J. Jewell, C. Keller, M. A. Kenworthy, J. Knight, J. KÅĖhn, K. Miller, J. Mazoyer, M. N'Diaye, E. Por, L. Pueyo, A. J. E. Riggs, G. Ruane, D. Sirbu, F. Snik, J. K. Wallace, M. Wilby, and M. Ygouf. Review of high-contrast imaging systems for current and future ground-based and space-based telescopes: Part II. common path wavefront sensing/control and coherent differential imaging. In *Adaptive Optics Systems VI*, volume 10703, 2018.

- [26] B. Nemati, J. E. Krist, and B. Mennesson. Sensitivity of the WFIRST coronagraph performance to key instrument parameters. In *Techniques and Instrumentation for Detection of Exoplanets VIII*, volume 10400, page 1040007. International Society for Optics and Photonics, 2017.
- [27] R. T. Demers, F. Dekens, R. Calvet, Z. Chang, R. Effinger, E. Ek, L. Hovland, L. Jones, A. Loc, B. Nemati, C. Noecker, T. Neville, H. Pham, M. Rud, H. Tang, and V. Juan. Requirements and design reference mission for the WFIRST/AFTA coronagraph instrument. In *Techniques and Instrumentation for Detection of Exoplanets VII*, volume 9605, page 960502. International Society for Optics and Photonics, 2015.
- [28] M. R. Bolcar, S. Aliez, V. T. Bly, C. Collins, J. Crooke, C. D. Dressing, L. Fantano, L. D. Feinberg, K. France, G. Gochar, Q. Gong, J. E. Hylan, A. Jones, I. Linares, M. Postman, L. Pueyo, A. Roberge, L. Sacks, S. Tompkins, and G. West. The Large UV/Optical/Infrared Surveyor (LUVOIR): Decadal mission concept design update. In *UV/Optical/IR Space Telescopes and Instruments: Innovative Technologies and Concepts VIII*, volume 10398, pages 79 – 102. International Society for Optics and Photonics, SPIE, 2017.
- [29] B. Mennesson, S. Gaudi, S. Seager, K. Cahoy, S. Domagal-Goldman, L. Feinberg, O. Guyon, J. Kasdin, C. Marois, D. Mawet, M. Tamura, D. Mouillet, T. Prusti, A. Quirrenbach, T. Robinson, L. Rogers, P. Scowen, R. Somerville, K. Stapelfeldt, D. Stern, M. Still, M. Turnbull, J. Booth, A. Kiessling, G. Kuan, and K. Warfield. The habitable exoplanet (HabEx) imaging mission: preliminary science drivers and technical requirements. In *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*, volume 9904, pages 212 – 221. International Society for Optics and Photonics, SPIE, 2016.
- [30] C. C. Stark, R. Belikov, M. R. Bolcar, E. Cady, B. P. Crill, S. Ertel, T. Groff, S. Hildebrandt, J. Krist, P. D. Lisman, J. Mazoyer, B. Mennesson, B. Nemati, L. Pueyo, B. J. Rauscher, A. J. Riggs, G. Ruane, S. B. Shaklan, D. Sirbu, R. Soummer, K. S. Laurent, and N. Zimmerman. Exoearth yield landscape for future direct imaging space telescopes. *Journal of Astronomical Telescopes, Instruments, and Systems*, 5(2):1 – 20, 2019.
- [31] K. Balasubramanian, V. E. White, K. Y. Yee, P. M. Echternach, R. E. Muller, M. R. Dickie, E. J. Cady, C. M. Prada, D. J. Ryan, I. Poberezhskiy, B. D. Kern, H. Zhou, J. E. Krist, B. Nemati, A. J. E. Riggs, N. T. Zimmerman, and N. J. Kasdin. WFIRST-AFTA coronagraph shaped pupil masks: design, fabrication, and characterization. *Journal of Astronomical Telescopes, Instruments, and Systems*, 2(1):1 – 16, 2015.
- [32] K. Patterson, J. Shields, X. Wang, H. Tang, A. Azizi, P. Brugarolas, M. Mandic, and F. Shi. Control design for momentum-compensated fast steering mirror for WFIRST-AFTA coronagraph instrument. In *Techniques and Instrumentation*

for *Detection of Exoplanets VII*, volume 9605, page 96052C. International Society for Optics and Photonics, 2015.

- [33] F. Shi, E. Cady, B.-J. Seo, X. An, K. Balasubramanian, B. Kern, R. Lam, D. Marx, D. Moody, C. M. Prada, K. Patterson, I. Poberezhskiy, J. Shields, E. Sidick, H. Tang, J. Trauger, T. Truong, V. White, D. Wilson, and H. Zhou. Dynamic testbed demonstration of WFIRST coronagraph low order wavefront sensing and control (LOWFS/C). In *Techniques and Instrumentation for Detection of Exoplanets VIII*, volume 10400, pages 74 – 90. International Society for Optics and Photonics, SPIE, 2017.
- [34] J. E. Krist, B. Nemati, and B. P. Mennesson. Numerical modeling of the proposed WFIRST-AFTA coronagraphs and their predicted performances. *Journal of Astronomical Telescopes, Instruments, and Systems*, 2(1):1 – 26, 2015.
- [35] R. E. Morgan, E. S. Douglas, G. W. Allan, P. Bierden, S. Chakrabarti, T. Cook, M. Egan, G. Furesz, J. N. Gubner, T. D. Groff, C. A. Haughwout, B. G. Holden, C. B. Mendillo, M. Ouellet, P. d. Vale Pereira, A. J. Stein, S. Thibault, X. Wu, Y. Xin, and K. L. Cahoy. MEMS deformable mirrors for space-based high-contrast imaging. *Micromachines*, 10(6), 2019.
- [36] A. Give'on, B. D. Kern, S. Shaklan, D. C. Moody, and L. Pueyo. Electric field conjugation-a broadband wavefront correction algorithm for high-contrast imaging systems. In *Bulletin of the American Astronomical Society*, volume 39, page 975, 2007.
- [37] B.-J. Seo, K. Patterson, K. Balasubramanian, B. Crill, T. Chui, D. Echeverri, B. Kern, D. Marx, D. Moody, C. M. Prada, G. Ruane, F. Shi, J. Shaw, N. Siegler, H. Tang, J. Trauger, D. Wilson, and R. Zimmer. Testbed demonstration of high-contrast coronagraph imaging in search for earth-like exoplanets. In *Techniques and Instrumentation for Detection of Exoplanets IX*, volume 11117, pages 599 – 609. International Society for Optics and Photonics, SPIE, 2019.
- [38] S. B. Shaklan, L. Marchen, J. E. Krist, and M. Rud. Stability error budget for an aggressive coronagraph on a 3.8 m telescope. In *Techniques and Instrumentation for Detection of Exoplanets V*, volume 8151, page 815109. International Society for Optics and Photonics, 2011.
- [39] H. P. Stahl, M. Postman, and W. S. Smith. Engineering specifications for large aperture UVO space telescopes derived from science requirements. In *UV/Optical/IR Space Telescopes and Instruments: Innovative Technologies and Concepts VI*, volume 8860, page 886006. International Society for Optics and Photonics, 2013.
- [40] H. Zhou, J. Krist, E. Cady, and I. Poberezhskiy. High accuracy coronagraph flight WFC model for WFIRST-CGI raw contrast sensitivity analysis. In *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*,

- volume 10698, pages 811 – 824. International Society for Optics and Photonics, SPIE, 2018.
- [41] G. Ruane, D. Mawet, J. Jewell, and S. Shaklan. Performance and sensitivity of vortex coronagraphs on segmented space telescopes. In *Techniques and Instrumentation for Detection of Exoplanets VIII*, volume 10400, pages 140 – 155. International Society for Optics and Photonics, SPIE, 2017.
 - [42] R. G. Paxman, T. J. Schulz, and J. R. Fienup. Joint estimation of object and aberrations by using phase diversity. *Journal of the Optical Society of America A*, 9(7):1072, 1992.
 - [43] T. D. Groff, A. J. E. Riggs, B. Kern, and N. J. Kasdin. Methods and limitations of focal plane sensing, estimation, and control in high-contrast imaging. *Journal of Astronomical Telescopes, Instruments, and Systems*, 2(1):011009, 2015.
 - [44] J. W. Goodman. *Introduction to Fourier optics*. Roberts and Company Publishers, 2005.
 - [45] M. Fox. *Quantum optics: an introduction*, volume 15. OUP Oxford, 2006.
 - [46] A. Give'on, B. D. Kern, and S. B. Shaklan. Pair-wise, deformable mirror, image plane-based diversity electric field estimation for high contrast coronagraphy. In *Techniques and Instrumentation for Detection of Exoplanets V*, volume 8151, page 815110. International Society for Optics and Photonics, 2011.
 - [47] V. P. Bailey, M. Bottom, E. Cady, F. Cantalloube, J. Boer, T. Groff, J. Krist, M. A. Millar-Blanchaer, A. Vigan, J. Chilcote, E. Choquet, R. J. De-Rosa, J. H. Girard, O. Guyon, B. Kern, A.-M. Lagrange, B. Macintosh, J. R. Males, C. Marois, T. Meshkat, J. Milli, M. N'Diaye, H. Ngo, E. L. Nielsen, J. Rhodes, G. Ruane, R. G. v. Holstein, J. J. Wang, and W. Xuan. Lessons for WFIRST CGI from ground-based high-contrast systems. In *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*, volume 10698, page 106986P. International Society for Optics and Photonics, 2018.
 - [48] J. Krist, R. Effinger, B. Kern, M. Mandic, J. McGuire, D. Moody, P. Morrissey, I. Poberezhskiy, A. J. Riggs, N. Saini, E. Sidick, H. Tang, and J. Trauger. WFIRST coronagraph flight performance modeling. In *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*, volume 10698, pages 788 – 810. International Society for Optics and Photonics, SPIE, 2018.
 - [49] L. Pogorelyuk and N. J. 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, 2019.
 - [50] K. Miller, O. Guyon, and J. Males. Spatial linear dark field control: stabilizing deep contrast for exoplanet imaging using bright speckles. *Journal of Astronomical Telescopes, Instruments, and Systems*, 3(04):1, 2017.

- [51] T. Currie, E. Pluzhnik, R. Belikov, and O. Guyon. Developing linear dark-field control for exoplanet direct imaging in the laboratory and on ground-based telescopes. In *Techniques and Instrumentation for Detection of Exoplanets IX*, volume 11117, pages 639 – 648. International Society for Optics and Photonics, SPIE, 2019.
- [52] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35, 1960.
- [53] R. F. Stengel. *Optimal control and estimation*. Dover Publications, New York, US, 1994.
- [54] T. D. Groff and N. J. Kasdin. Kalman filtering techniques for focal plane electric field estimation. *Journal of the Optical Society of America A*, 30(1):128, 2013.
- [55] A. E. Riggs, N. J. Kasdin, and T. D. Groff. Optimal wavefront estimation of incoherent sources. In *Space Telescopes and Instrumentation 2014: Optical, Infrared, and Millimeter Wave*, volume 9143, page 914324. International Society for Optics and Photonics, 2014.
- [56] J. P. Hespanha. *Linear systems theory*. Princeton university press, 2009.
- [57] L. Pogorelyuk, L. Pueyo, and N. J. Kasdin. On the effects of pointing jitter, actuators drift, telescope rolls and broadband detectors in dark hole maintenance and electric field model reduction. *submitted*.
- [58] A. E. Riggs, G. Ruane, C. T. Coker, S. B. Shaklan, B. D. Kern, and E. Sidick. Fast linearized coronagraph optimizer (FALCO) i: a software toolbox for rapid coronagraphic design and wavefront correction. In *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*, volume 10698. SPIE, 2018.
- [59] L. K. Harding, R. Demers, M. E. Hoenk, P. Peddada, B. Nemati, M. Cherng, D. Michaels, L. S. Neat, A. Loc, N. L. Bush, D. J. Hall, N. J. Murray, J. P. D. Gow, R. Burgon, A. D. Holland, A. L. Reinheimer, P. R. Jorden, and D. Jordan. Technology advancement of the CCD201-20 EMCCD for the WFIRST coronagraph instrument: sensor characterization and radiation damage. *Journal of Astronomical Telescopes, Instruments, and Systems*, 2(1):1 – 31 – 31, 2015.
- [60] L. Pogorelyuk and N. J. Kasdin. Maintaining a dark hole in a high contrast coronagraph and the effects of speckles drift on contrast and post processing factor. In *Techniques and Instrumentation for Detection of Exoplanets IX*, volume 11117, pages 397 – 403. International Society for Optics and Photonics, SPIE, 2019.
- [61] C. M. Prada, E. Serabyn, and F. Shi. High-contrast imaging stability using MEMS deformable mirror. In *Techniques and Instrumentation for Detection of Exoplanets IX*, volume 11117, pages 112 – 118. International Society for Optics and Photonics, SPIE, 2019.

- [62] F. Shi, J. Shields, T. Truong, B.-J. Seo, F. Fregoso, B. Kern, D. Marx, K. Patterson, C. M. Prada, J. Shaw, and C. Shelton. WFIRST low order wavefront sensing and control (LOWFS/C) performance on line-of-sight disturbances from multiple reaction wheels. In *Techniques and Instrumentation for Detection of Exoplanets IX*, volume 11117, pages 170 – 178. International Society for Optics and Photonics, SPIE, 2019.
- [63] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge, M. Donahue, A. Dressler, B. Gaudi, T. Greene, O. Guyon, et al. Wide-field infrared survey telescope-astronomy focused telescope assets wfirst-afta 2015 report. *arXiv preprint arXiv:1503.03757*, 2015.
- [64] D. F. Parrish. A Kalman filter for a two-dimensional shallow-water model: Formulation and preliminary experiments. *Office note 304*, 1985.
- [65] S.-H. Wang, E. Wang, and P. Dorato. Observing the states of systems with unmeasurable disturbances. *IEEE Transactions on Automatic Control*, 20(5):716–717, 1975.
- [66] M. Hou and R. Patton. Optimal filtering for systems with unknown inputs. *IEEE Transactions on Automatic Control*, 43(3):445–449, 1998.
- [67] K. M. Nagpal, R. E. Helmick, and C. S. Sims. Reduced-order estimation part 1. filtering. *International Journal of Control*, 45(6):1867–1888, 1987.
- [68] J. Keller and M. Darouach. Reduced-order Kalman filter with unknown inputs. *Automatica*, 34(11):1463–1468, 1998.
- [69] R. Todling and S. E. Cohn. Suboptimal schemes for atmospheric data assimilation based on the Kalman filter. *Monthly Weather Review*, 122(11):2530–2557, 1994.
- [70] I. Fukumori and P. Malanotte-Rizzoli. An approximate Kalman filter for ocean data assimilation: An example with an idealized Gulf Stream model. *Journal of Geophysical Research*, 100(C4):6777, 1995.
- [71] M. A. Cane, A. Kaplan, R. N. Miller, B. Tang, E. C. Hackert, and A. J. Busalacchi. Mapping tropical Pacific sea level: Data assimilation via a reduced state space Kalman filter. *Journal of Geophysical Research: Oceans*, 101(C10):22599–22617, 1996.
- [72] M. Verlaan and A. W. Heemink. Tidal flow forecasting using reduced rank square root filters. *Stochastic Hydrology and Hydraulics*, 11(5):349–368, 1997.
- [73] S. E. Cohn and R. Todling. Approximate Kalman filters for unstable dynamics. In *Second International Symposium on the Assimilation of Observations in Meteorology and Oceanography*, pages 241–246. World Meteorological Organization, Geneva, Switzerland, 1995.

- [74] P. F. J. Lermusiaux and A. R. Robinson. Data assimilation via error subspace statistical estimation. part I: Theory and schemes. *Monthly Weather Review*, 127(7):1385–1407, 1999.
- [75] G. Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99(C5):10143, 1994.
- [76] L. Glielmo, R. Setola, and F. Vasca. An interlaced extended Kalman filter. *IEEE Transactions on Automatic Control*, 44(8):1546–1549, 1999.
- [77] G. V. Puskorius and L. A. Feldkamp. Decoupled extended Kalman filter training of feedforward layered networks. In *IJCNN-91-Seattle International Joint Conference on Neural Networks*, volume i, pages 771–777 vol.1. IEEE, 1991.
- [78] S.-L. Sun and Z.-L. Deng. Multi-sensor optimal information fusion Kalman filter. *Automatica*, 40(6):1017–1023, 2004.
- [79] L. Pogorelyuk, C. W. Rowley, and N. J. Kasdin. An efficient approximation of the Kalman filter for multiple systems coupled via low-dimensional stochastic input. *submitted*.
- [80] M. A. Woodbury. Inverting modified matrices. *Statistical Research Group, Memo. Rep. no. 42*, 42(106):336, 1950.
- [81] P. Lancaster and L. Rodman. *Algebraic riccati equations*. Clarendon Press, Oxford, 1995.
- [82] F. L. Bauer and C. T. Fike. Norms and exclusion theorems. *Numerische Mathematik*, 2(1):137–141, 1960.
- [83] M. D. Perrin, J. Maire, P. Ingraham, D. Savransky, M. Millar-Blanchaer, S. G. Wolff, J.-B. Ruffio, J. J. Wang, Z. H. Draper, N. Sadakuni, C. Marois, A. Rajan, M. P. Fitzgerald, B. Macintosh, J. R. Graham, R. Doyon, J. E. Larkin, J. K. Chilcote, S. J. Goodsell, D. W. Palmer, K. Labrie, M. Beaulieu, R. J. D. Rosa, A. Z. Greenbaum, M. Hartung, P. Hibon, Q. Konopacky, D. Lafreniere, J.-F. Lavigne, F. Marchis, J. Patience, L. Pueyo, F. T. Rantakyro, R. Soummer, A. Sivaramakrishnan, S. Thomas, K. Ward-Duong, and S. Wiktorowicz. Gemini Planet Imager observational calibrations I: Overview of the GPI data reduction pipeline. In *Ground-based and Airborne Instrumentation for Astronomy V*, volume 9147, pages 1168 – 1180. International Society for Optics and Photonics, SPIE, 2014.
- [84] J. J. Wang, J.-B. Ruffio, R. J. De Rosa, J. Aguilar, S. G. Wolff, and L. Pueyo. pyKLIP: PSF subtraction for exoplanets and disks. *Astrophysics Source Code Library*, 2015.

- [85] C. A. G. Gonzalez, O. Wertz, O. Absil, V. Christiaens, D. Defrère, D. Mawet, J. Milli, P.-A. Absil, M. V. Droogenbroeck, F. Cantalloube, P. M. Hinz, A. J. Skemer, M. Karlsson, and J. Surdej. VIP: Vortex image processing package for high-contrast direct imaging. *The Astronomical Journal*, 154(1):7, 2017.
- [86] D. Mawet, B. Mennesson, E. Serabyn, K. Stapelfeldt, and O. Absil. A dim candidate companion to Epsilon Cephei. *The Astrophysical Journal Letters*, 738(1):L12, 2011.
- [87] R. Soummer, L. Pueyo, and J. Larkin. Detection and characterization of exoplanets and disks using projections on Karhunen-Loève eigenimages. *The Astrophysical Journal Letters*, 755(2):L28, 2012.
- [88] A. Amara and S. P. Quanz. PYNPOINT: An image processing package for finding exoplanets. *Monthly Notices of the Royal Astronomical Society*, 427(2):948–955, 2012.
- [89] R. Fergus, D. W. Hogg, R. Oppenheimer, D. Brenner, and L. Pueyo. S4: A spatial-spectral model for speckle suppression. *The Astrophysical Journal*, 794(2):161, 2014.
- [90] C. A. G. Gonzalez, O. Absil, P.-A. Absil, M. V. Droogenbroeck, D. Mawet, and J. Surdej. Low-rank plus sparse decomposition for exoplanet detection in direct-imaging ADI sequences. *Astronomy & Astrophysics*, 589:A54, 2016.
- [91] B. Ren, L. Pueyo, G. B. Zhu, J. Debes, and G. Duchêne. Non-negative matrix factorization: Robust extraction of extended structures. *The Astrophysical Journal*, 852(2):104, 2018.
- [92] C. Marois, R. Doyon, R. Racine, and D. Nadeau. Efficient speckle noise attenuation in faint companion imaging. *Publications of the Astronomical Society of the Pacific*, 112(767):91–96, 2000.
- [93] W. B. Sparks and H. C. Ford. Imaging spectroscopy for extrasolar planet detection. *The Astrophysical Journal*, 578(1):543–564, 2002.
- [94] F. Martinache. Kernel phase in Fizeau interferometry. *The Astrophysical Journal*, 724(1):464–469, 2010.
- [95] N. Baba and N. Murakami. A method to image extrasolar planets with polarized light. *Publications of the Astronomical Society of the Pacific*, 115(814):1363–1366, 2003.
- [96] J. L. Codona, M. A. Kenworthy, P. M. Hinz, J. R. P. Angel, and N. J. Woolf. A high-contrast coronagraph for the MMT using phase apodization: design and observations at 5 microns and 2 λ /d radius. In *Ground-based and Airborne Instrumentation for Astronomy*, volume 6269, 2006.

- [97] M. A. Kenworthy, J. L. Codona, P. M. Hinz, J. R. P. Angel, A. Heinze, and S. Sivanandam. First on-sky high-contrast imaging with an apodizing phase plate. *The Astrophysical Journal*, 660(1):762–769, 2007.
- [98] D. Mawet, L. Pueyo, P. Lawson, L. Mugnier, W. Traub, A. Boccaletti, J. T. Trauger, S. Gladysz, E. Serabyn, J. Milli, R. Belikov, M. Kasper, P. Baudoz, B. Macintosh, C. Marois, B. Oppenheimer, H. Barrett, J.-L. Beuzit, N. Devaney, J. Girard, O. Guyon, J. Krist, B. Mennesson, D. Mouillet, N. Murakami, L. Poyneer, D. Savransky, C. Verinaud, and J. K. Wallace. Review of small-angle coronagraphic techniques in the wake of ground-based second-generation adaptive optics systems. In *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*, volume 8442. SPIE, 2012.
- [99] P. J. Lowrance, E. E. Becklin, G. Schneider, D. Hines, J. D. Kirkpatrick, D. Koerner, F. Low, D. McCarthy, R. Meier, M. Rieke, B. A. Smith, R. Terile, R. Thompson, and B. Zuckerman. A coronagraphic search for substellar companions to young stars. In *NICMOS and the VLT: A New Era of High Resolution Near Infrared Imaging and Spectroscopy*, volume 55 of *European Southern Observatory Conference and Workshop Proceedings*, page 96. European Southern Observatory, 1998.
- [100] C. Marois, D. Lafreniere, R. Doyon, B. Macintosh, and D. Nadeau. Angular differential imaging: A powerful high-contrast imaging technique. *The Astrophysical Journal*, 641(1):556–564, 2006.
- [101] D. Lafreniere, C. Marois, R. Doyon, D. Nadeau, and E. Artigau. A new algorithm for point-spread function subtraction in high-contrast imaging: A demonstration with angular differential imaging. *The Astrophysical Journal*, 660(1):770–780, 2007.
- [102] J. L. Codona and M. Kenworthy. Focal plane wavefront sensing using residual adaptive optics speckles. *The Astrophysical Journal*, 767(2):100, 2013.
- [103] L. Pogorelyuk, N. J. Kasdin, and C. W. Rowley. Reduced order estimation of the speckle electric field history for space-based coronagraphs. *The Astrophysical Journal*, 881(2):126, 2019.
- [104] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, Savannah, GA, 2016. USENIX Association.
- [105] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- [106] J. R. P. Angel. Ground-based imaging of extrasolar planets using adaptive optics. *Nature*, 368(6468):203, 1994.
- [107] A. Sivaramakrishnan, J. P. Lloyd, P. E. Hodge, and B. A. Macintosh. Speckle decorrelation and dynamic range in speckle noise-limited imaging. *The Astrophysical Journal Letters*, 581(1):L59, 2002.
- [108] R. A. Frazin. Statistical framework for the utilization of simultaneous pupil plane and focal plane telemetry for exoplanet imaging. I. accounting for aberrations in multiple planes. *J. Opt. Soc. Am. A*, 33(4):712–725, 2016.
- [109] B. Carlomagno, O. Absil, M. Kenworthy, G. Ruane, C. U. Keller, G. Otten, M. Feldt, S. Hippler, E. Huby, D. Mawet, C. Delacroix, J. Surdej, S. Habraken, P. Forsberg, M. Karlsson, E. V. Catalan, and B. R. Brandl. End-to-end simulations of the E-ELT/METIS coronagraphs. In *Adaptive Optics Systems V*, volume 9909, pages 2070 – 2079. International Society for Optics and Photonics, SPIE, 2016.
- [110] E. H. Por, S. Y. Haffert, V. M. Radhakrishnan, D. S. Doelman, M. v. 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, pages 1112 – 1125. International Society for Optics and Photonics, SPIE, 2018.
- [111] J. A. Lee and M. Verleysen. *Nonlinear Dimensionality Reduction*. Springer New York, 2007.
- [112] M. Schmidt and H. Lipson. Distilling free-form natural laws from experimental data. *Science*, 324(5923):81–85, 2009.
- [113] F. L. Lewis and D. Vrabie. Reinforcement learning and adaptive dynamic programming for feedback control. *IEEE Circuits and Systems Magazine*, 9(3):32–50, 2009.
- [114] L. Pogorelyuk, S. E. Otto, and C. W. Rowley. *in preparation*.
- [115] M. Sonka, V. Hlavac, and R. Boyle. *Image processing, analysis, and machine vision*. Cengage Learning, 2014.
- [116] P. Benner, S. Gugercin, and K. Willcox. A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Review*, 57(4):483–531, 2015.
- [117] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.
- [118] J. B. Tenenbaum. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

- [119] H. Valpola and J. Karhunen. An unsupervised ensemble learning method for nonlinear dynamic state-space models. *Neural Computation*, 14(11):2647–2692, 2002.
- [120] L. Wiskott and T. J. Sejnowski. Slow feature analysis: Unsupervised learning of invariances. *Neural Computation*, 14(4):715–770, 2002.
- [121] G. Doretto, A. Chiuso, Y. N. Wu, and S. Soatto. Dynamic textures. *International Journal of Computer Vision*, 51(2):91–109, 2003.
- [122] L. Yuan, F. Wen, C. Liu, and H.-Y. Shum. Synthesizing dynamic texture with closed-loop linear dynamic system. In *Computer Vision - ECCV 2004*, pages 603–616, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.
- [123] S. Mahadevan and M. Maggioni. Proto-value functions: A Laplacian framework for learning representation and control in Markov decision processes. *Journal of Machine Learning Research*, 8(Oct):2169–2231, 2007.
- [124] O. C. Jenkins and M. J. Matarić. A spatio-temporal extension to isomap nonlinear dimension reduction. In *Twenty-first international conference on Machine learning - ICML '04*. ACM Press, 2004.
- [125] M. Thida, H. Eng, and P. Remagnino. Laplacian eigenmap with temporal constraints for local abnormality detection in crowded scenes. *IEEE Transactions on Cybernetics*, 43(6):2147–2156, 2013.
- [126] G. Lee, C. Rodriguez, and A. Madabhushi. Investigating the efficacy of nonlinear dimensionality reduction schemes in classifying gene and protein expression studies. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 5(3):368–384, 2008.
- [127] R. Talmon, S. Mallat, H. Zaveri, and R. R. Coifman. Manifold learning for latent variable inference in dynamical systems. *IEEE Transactions on Signal Processing*, 63(15):3843–3856, 2015.
- [128] A. Rahimi, B. Recht, and T. Darrell. Learning appearance manifolds from video. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*. IEEE.
- [129] G. Zhao and M. Pietikainen. Dynamic texture recognition using local binary patterns with an application to facial expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(6):915–928, 2007.
- [130] P. Turaga, R. Chellappa, V. S. Subrahmanian, and O. Udrea. Machine recognition of human activities: A survey. *IEEE Transactions on Circuits and Systems for Video technology*, 18(11):1473, 2008.
- [131] C. J. C. H. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8(3-4):279–292, May 1992.

- [132] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [133] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- [134] W. Böhmer, J. T. Springenberg, J. Boedecker, M. Riedmiller, and K. Obermayer. Autonomous learning of state representations for control: An emerging field aims to autonomously learn state representations for reinforcement learning agents from their real-world sensor observations. *KI - Künstliche Intelligenz*, 29(4):353–362, 2015.
- [135] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- [136] M. Watter, J. Springenberg, J. Boedecker, and M. Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in Neural Information Processing Systems 28*, pages 2746–2754. Curran Associates, Inc., 2015.
- [137] N. Wahlström, T. B. Schön, and M. P. Deisenroth. From pixels to torques: Policy learning with deep dynamical models. *arXiv preprint arXiv:1502.02251*, 2015.
- [138] T. Lesort, N. Diaz-Rodriguez, J.-F. Goudou, and D. Filliat. State representation learning for control: An overview. *Neural Networks*, 108:379–392, 2018.
- [139] M. Karl, M. Soelch, J. Bayer, and P. v. d. Smagt. Deep variational Bayes filters: Unsupervised learning of state space models from raw data. *arXiv preprint arXiv:1605.06432*, 2016.
- [140] K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116(45):22445–22451, 2019.
- [141] D. B. West et al. *Introduction to graph theory*, volume 2. Prentice hall Upper Saddle River, NJ, 1996.
- [142] K. V. Mardia and J. T. K. J. M. Bibby. *Multivariate Analysis*. Academic Press, London, 1979.
- [143] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.

- [144] J. Bongard and H. Lipson. Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 104(24):9943–9948, 2007.
- [145] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley. A data-driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015.
- [146] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.
- [147] S. M. Redmond, L. Pogorelyuk, and N. J. Kasdin. *in preparation*.