UNIVERSITY OF CALIFORNIA
SANTA CRUZ

EVERY PART OF THE PIG FOR HIGH CONTRAST IMAGING AND EXTREME ADAPTIVE OPTICS

A dissertation submitted in partial satisfaction
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

ASTRONOMY & ASTROPHYSICS

by

Jules Fowler

June 2026

The Dissertation of Jules Fowler is approved:

Rebecca Jensen-Clem

Phil Hinz

Maxwell Millar-Blanchaer

Bruce Macintosh

Lucie Leboulleux

Peter Biehl
Vice Provost and Dean of Graduate Studies

Copyright © by

Jules Fowler

2026

Table of Contents

List of Figures viii

List of Tables xii

Abstract xiii

1 Introduction 1
- 1.1 Finding an exo-Earth . . . 1
- 1.2 High contrast imaging . . . 7
  - 1.2.1 Direct and indirect methods of exoplanet characterization. 7
  - 1.2.2 Coronagraphy and differential imaging techniques . . . 8
- 1.3 Polarimetry and polarimetric differential imaging . . . 13
  - 1.3.1 Physical properties of polarized light . . . 13
  - 1.3.2 Finding Earth-like planets with reflected light polarimetry 16
  - 1.3.3 Disk integrated polarimetry of giant self-luminous companions . . . 17
  - 1.3.4 Polarimetric differential imaging in practice . . . 19
- 1.4 Extreme adaptive optics and visible light adaptive optics . . . 21
  - 1.4.1 Speckles and atmospheric turbulence . . . 21
  - 1.4.2 Anatomy of an adaptive optics system . . . 24
  - 1.4.3 Error terms in an adaptive optics system . . . 27
  - 1.4.4 Keck-II adaptive optics error budget . . . 33
  - 1.4.5 Translating error terms to performance metrics . . . 34
  - 1.4.6 Adaptive optics for visible light astronomy . . . 35
- 1.5 Predictive wavefront control . . . 38
  - 1.5.1 Kalman filtering and linear quadratic Gaussian control 40
  - 1.5.2 Empirical orthogonal functions . . . 43
  - 1.5.3 Neural networks . . . 47
  - 1.5.4 Reinforcement learning . . . 50
  - 1.5.5 Overlaps between methods . . . 51

1.6 Summary 53

2 When PIGSS Fly: the Polarimetric InvestiGation of Self-luminous Substellar companions 55
2.1 Abstract 55
2.2 Introduction 56
2.3 Target selection 58
2.3.1 HIP 21152 B 62
2.3.2 AF Lep b 62
2.4 Subaru/SCExAO observations 63
2.4.1 Observing and adaptive optics conditions 63
2.4.2 CHARIS spectropolarimetric imaging 64
2.4.3 Polarimetric differential imaging 65
2.5 Data reduction, contrast curves, and upper limits 68
2.5.1 Data reduction and contrast curves 68
2.6 Conclusions and future work 73

3 Don't Heckle my (visible light) Speckle: Characterizing Keck-II adaptive optics performance in visible light with the ORKID instrument 76
3.1 Abstract 77
3.2 Introduction 78
3.3 ORCAS Keck Instrument Demonstrator 80
3.3.1 ORKID instrument 80
3.3.2 Keck-II adaptive optics system 81
3.4 ORKID observations 82
3.4.1 Point spread function sharpening for ORKID during daytime testing 85
3.4.2 State of the segmented primary after a segment exchange 86
3.5 Performance and comparison to analytic predictions 88
3.5.1 Strehl ratio and encircled energy 89
3.5.2 Full-width half-maximum and jitter of the point spread function 93
3.5.3 Predictions from the Keck AO error budget 99
3.5.4 Empirically tuned analytic models of strehl ratio and encircled energy 101
3.5.5 Comparison between data and models 104
3.6 Identifying adverse impacts to ORKID performance 106
3.6.1 Impact of the sodium dichroic on visible light wavefront sensing in natural guide star mode 107

3.6.2 Recurring vibrations in the AO system 110
3.6.3 PSF splitting in fast ORKID exposures 115
3.7 Comparisons between ORKID PSF jitter and Keck-II AO tip and tilt residuals 119
3.7.1 X/Y asymmetry in PSF jitter 119
3.7.2 Large drifts over long timescales 120
3.8 Conclusion 122
3.8.1 Looking forward to HAKA science 123

4 Predictions vs. Reality: Towards facilitizing data-driven predictive wavefront control on-sky at Keck Observatory 127
4.1 Abstract 127
4.2 Introduction 128
4.3 Summary of empirical orthogonal functions and tunable hyper-parameters in the predictive filter matrix 130
4.4 Optimizing filter hyper-parameters in simulation 132
4.5 On-the-fly optimization with simulated annealing 134
4.6 Predictive control update to the Keck real time controller . . . 138
4.7 Performance verification with the Keck-II AO bench 139
4.8 On-sky results with Keck-II AO and NIRC2 142
4.8.1 Initial on-sky performance verification 142
4.8.2 Optimizing predictive filter hyper-parameters on-sky . 144
4.8.3 Residual wavefront error and Strehl ratio as performance metrics 148
4.8.4 High contrast imaging performance metrics 151
4.9 HAKA + predictive wavefront control 153
4.10 Conclusion 155

5 Ground Control to Major Time-lag: Comparing predictive wavefront control methods for the next generation of extreme adaptive optics 159
5.1 Abstract 160
5.2 Introduction 161
5.3 Brief summary of empirical orthogonal functions 164
5.4 Comparing open and closed-loop predictive wavefront control 166
5.4.1 Our development of a closed-loop empirical orthogonal functions update 166
5.4.2 Verifying the time flow of information through a control system 169
5.4.3 Selecting a truth condition for training data 173

5.4.4 Preliminary implementation of closed-loop predictive wavefront control on simulated data 174
5.5 Comparing data and model-driven predictive wavefront control 177
5.5.1 Overview of predictive Fourier control as a model-driven predictor 177
5.5.2 Predictive Fourier control math in detail 181
5.5.3 Optimized implementation of predictive wavefront controllers on simulated data and on-sky telemetry 191
5.6 Conclusion 202

6 Conclusions **204**
6.1 Polarimetry makes use of parts of light we normally neglect to understand atmospheric asymmetries that are hard to probe with other observational techniques 204
6.2 Visible light adaptive optics is extremely challenging but current performance is very well matched to error budgets and models 206
6.3 Data-driven predictive wavefront control improves adaptive optics corrections for a ten meter telescope 208
6.4 There are many different ways to do predictive wavefront control and finding the "best" one will require a direct comparison on a telescope 209
6.5 Future outlook and machinations 214

A Conceivably Useful Appendices **304**
A.1 Examples of predictive control in the literature, broken down by method 304
A.1.1 Empirical orthogonal functions 304
A.1.2 Kalman filtering and linear quadratic Gaussian controllers 306
A.1.3 Neural networks 308
A.2 Center-finding strategies for ORKID PSF jitter analysis 311
A.3 Math mistakes in fundamental papers 313
A.3.1 Paschall & Anderson, 1993 313
A.3.2 Dessenne et al., 1998 314
A.3.3 Poyneer et al., 2007 315

B Foolish Musings **316**
B.1 In violation of the prime directive: simulating detriments to Delta-Quadrant civilizations from the starship Voyager's impact on planetary rings 316
B.1.1 Abstract 316

B.1.2 Introduction . . . 317
B.1.3 Initial conditions of a voyager-planet system . . . 319
B.1.4 Simulated Voyager impacts . . . 325
B.1.5 Violations of the prime directive . . . 329
B.1.6 Conclusions . . . 332
B.2 Wavefront sensing and control with the many headed hydra modal basis . . . 333
B.2.1 Abstract . . . 333
B.2.2 Introduction . . . 334
B.2.3 Formulation of the many headed hydra modal basis . . . 336
B.2.4 Efficacy amongst the hydra heads . . . 336
B.2.5 Control applications . . . 341
B.2.6 Conclusions . . . 343
B.3 Full suite of hydra heads . . . 344
B.3.1 Hirschauer basis . . . 344
B.3.2 Fowler basis . . . 345
B.3.3 Hinz basis . . . 346
B.3.4 Soummer basis . . . 347
B.3.5 Macintosh basis . . . 348
B.3.6 Leboulleux basis . . . 349
B.3.7 Pearce basis . . . 350
B.3.8 Stelter basis . . . 351
B.3.9 Jensen-Clem basis . . . 352
B.3.10 Millar-Blanchaer basis . . . 353
B.3.11 Steiger basis . . . 354
B.3.12 Lewis basis . . . 355
B.3.13 Balmer basis . . . 356

List of Figures

1.1 Small angle approximation. 3
1.2 Habitable zones for M through G type stars. 4
1.3 Exoplanet detection methods. 6
1.4 Mass period diagram of different detection methods. 10
1.5 Diagram of a classic Lyot coronagraph. 11
1.6 Polarization states. 14
1.7 Polarization ellipse. 15
1.8 Chrysina gloriosa beetle, a natural polarizer. 16
1.9 Water glint feature from an extremely Earth-like planet. . . . 18
1.10 Disk integrated polarization. 19
1.11 Diagram of polarimetric differential imaging. 20
1.12 AO on/off on Neptune. 22
1.13 AO on/off on Neptune. 23
1.14 Diagram of an adaptive optics sytem. 24
1.15 Flow of information through an adaptive optics system. 25
1.16 Physical scale of IR and visible light wavelengths. 37
1.17 Kalman filter diagram 41
1.18 Linear Quadratic Gaussian Controller diagram. 42
1.19 Empirical Orthogonal Functions history vector. 44
1.20 Empirical Orthogonal Functions calculation in real time. . . . 45
1.21 Diagrams of reinforcement learning and neural nets. 49
1.22 Showing connections and overlaps between different machine learning methods. 51
1.23 Notorious PIG 54

2.1 Target selection considerations for HIP 21152 B and AF Lep b. 60
2.2 Example HIP 21152 B images at different pipeline stages. . . . 69
2.3 HIP 21152 B images. 70
2.4 AF Lep b images. 71
2.5 Impact of removing 50 vs 100 KL modes on contrast curves. . 72
2.6 Preliminary 5-σ contrast curves for HIP 21152 B and AF Lep b. 73

3.1 Natural guide star wavefront sensing modes for ORKID. . . . 82
3.2 Seeing during the ORKID night according the DIMM monitor. 85
3.3 Daytime PSF sharpening for ORKID. 86
3.4 Keck primary mirror measurements. 88
3.5 Strehl calculations. 91
3.6 Point spread fitting splitting in the ORKID PSF. 93
3.7 ORKID data with and without residual tip/tilt errors. 94
3.8 ORKID full-width half-max. 95
3.9 ORKID point spread function jitter. 97
3.10 Closed-loop adaptive optics on ORKID. 98
3.11 ORKID performance with varying guide star magnitude. . . . 105
3.12 Sodium dichroic transmission and ORKID wavefront sensing detector quantum efficiency. 110
3.13 Full estimated transmission of the sodium dichroic wavefront sensing mode. 111
3.14 Keck-II Strehl performance with guide star magnitude for ORKID and NIRC2. 112
3.15 Recurring vibrations in the ORKID data. 114
3.16 ORKID data with and without the ring down effect. 116
3.17 Point spread function splitting in ORKID data and in simulation. 117
3.18 ORKID point spread function jitter compared to tip/tilt mirror telemetry. 120
3.19 Long duration drifts in the ORKID point spread function. . . 121
3.20 Predictions of future ORKID performance with the HAKA deformable mirror. 125

4.1 Performance of Empirical Orthogonal Functions with varying hyper-parameters: r_0 varies with wind speed constant. 133
4.2 Performance of Empirical Orthogonal Functions with varying hyper-parameters: speed varies with r_0 constant. 134
4.3 Simulated annealing performance on simulated data: grid perspective. 136
4.4 Simulated annealing performance on simulated data: n iterations to optimal solution. 137
4.5 PWFC BU block diagram. 139
4.6 Predictor vs integrator performance on Keck-II AO bench. . . 141
4.7 Predictor vs integrator performance during first on-sky test: residual wavefront on the SHWFS. 143
4.8 Predictor vs integrator performance during first on-sky test: NRC2 images. 144

4.9 On-sky hyper-parameter optimzation for predictive control. . . 147
4.10 Predictor vs integrator performance during engineering half-night: residual wavefront on the SHWFS. 150
4.11 Predictor vs integrator performance during engineering half-night: NIRC2 images. 151
4.12 Predictor vs integrator performance during engineering half-night: SRs from imaging data. 152
4.13 Predictor vs integrator performance during engineering half-night: high contrast imaging metrics. 153
4.14 Project HAKA + ORKID performance using predictive control. 155

5.1 Wavefront control lop diagram. 170
5.2 Simulated turbulence and parameters. 174
5.3 Comparing open and closed-loop control methods in simulation. 175
5.4 Peak identification for Predictive Fourier Control. 180
5.5 Reconstruction of a turbulent scene with Fourier modes. . . . 183
5.6 Peak identification of injected wind layers. 184
5.7 Model vs. data-driven predictors applied to simulated data: single layer turbulence. 192
5.8 Model vs. data-driven predictors applied to simulated data: multi-layer turbulence. 193
5.9 Recovery of a known wind-layer in Keck telemetry via peak identification. 196
5.10 Model vs. data-driven predictors applied to Keck telemetry: RMS residuals in time-series. 197
5.11 Model vs. data-driven predictors applied to Keck telemetry: residuals aperture on. 198
5.12 Comparing model and data-driven predictors: residual wavefront aperture on. 199

6.1 GAIA Data Release 4 exoplanet yeild after GRAVITY+ AO improvements. 213

A.1 Estimating error in point spread function jitter recovery for ORKID data. 313

B.1 Initial conditions from Voyager opening credits. 322
B.2 Initial conditions in the REBOUND simulation. 323
B.3 Ring particle height evolution. 326
B.4 Ring particle heights radially. 327
B.5 Ring particle density estimate. 328

B.6 Evidence of tiny civilizations. 330
B.7 Example modal basis using a Debes head. 337
B.8 Reconstruction of turbulence using Fowler modal basis. 339
B.9 Difference between reconstruction and original image using Fowler modal basis. 341
B.10 Hirschauer basis. 344
B.11 Fowler basis. 345
B.12 Hinz basis. 346
B.13 Soummer basis. 347
B.14 Macintosh basis. 348
B.15 Macintosh basis. 349
B.16 Pearce basis. 350
B.17 Stelter Basis. 351
B.18 Jensen-Clem Basis. 352
B.19 Millar-Blanchaer basis. 353
B.20 Steiger Basis. 354
B.21 Lewis Basis. 355
B.22 Balmer Basis. 356

List of Tables

1.1 Keck-II Adaptive Optics Error Budget summary post HAKA upgrade. . . . 34
1.2 Necessary and optimal AO criteria. . . . 39

2.1 System Properties . . . 61
2.2 Subaru/SCExAO observing data. . . . 61

3.1 Engineering targets from the 9/19/2024 ORKID run. . . . 83
3.2 Keck-II Adaptive Optics Error Budget summary. . . . 100
3.3 Strehl Ratio estimate from the Keck-II Adaptive Optics Error Budget. . . . 100
3.4 Impact of the sodium dichroic wavefront sensing mode. . . . 109

4.1 Summary of controller performance. . . . 149

5.1 Flow of information through controller. . . . 172
5.2 Success rate of peak identification. . . . 185
5.3 Multi-layer atmospheric profile simulation paramers. . . . 194
5.4 Strehl Ratios with predictor-estimated wavefront error. . . . 200

B.1 Initial conditions of the n-body simulation. . . . 324
B.2 Parameters for simulation of atmospheric turbulence. . . . 338
B.3 Eigenface efficacy. . . . 340

Abstract

Every Part of the PIG for High Contrast Imaging and Extreme Adaptive Optics

by

Jules Fowler

More than 6000 exoplanets have been discovered to date, but how well do we really understand them? Exoplanets give off their own thermal glow and by directly imaging this signal we gain key insight into a planet's size, rotational dynamics, and atmosphere. These characteristics are harder to study with other exoplanet observation techniques that indirectly infer planetary properties from the way a planet affects the star it orbits. The signatures of these planets are at a high contrast to their host star (i.e., much fainter), giving off less than 10^{-6} times the amount of light for giant planets (for earth-like planets 10^{-10}); this problem is often likened to searching for a firefly next to a lighthouse. Telescopes on the ground have the advantage of larger primary mirrors (i.e., collecting diameters for light) because they are not constrained by the Herculean efforts of space launches. However, observing these signatures from the ground means contending with Earth's atmosphere, which is made of up tiny thermally varying pockets that impart different indices of refraction,

varying on millisecond timescales; this blurs the incoming light (or wavefront). We combat these effects with adaptive optics (AO); by placing instrumentation that acts as a constantly evolving pair of glasses on our telescope systems, we can correct for these errors in real time. In particular, extreme adaptive optics (XAO) achieves exquisite correction over a small field of view, prioritizing the correction that enables high contrast imaging and reveals a planetary companion next door to a bright host star. Only with current and future XAO systems can we understand these 6000 exoplanets, and more to come. I look at every part of the pig for high contrast imaging and extreme adaptive optics in four parts: (1) probing the cloud properties of giant planets with polarized light observations from XAO-fed telescopes, (2) understanding the current state of AO at Keck Observatory, (3) exploring the practical implementations of novel control strategies at current large-class observatories, and (4) improving future XAO systems with novel control strategies.

I discuss results from 2 half nights of Subaru telescope time on the brown dwarf HIP 21152 B and giant planet AF Lep b, seeking the first polarimetric detection of a giant planet and direct evidence for clouds on giant companions. Due to poor data quality I do not detect either companion in total intensity or reach informative upper limits for polarized intensity. Data are limited by adaptive optics performance, nicely motivating the rest of this work, to

improve adaptive optics systems. I present the results of my work to commission the visible light ORCAS Keck Instrument Demonstrator (ORKID) at Keck Observatory, including improved ORKID image quality with PSF-sharpening daytime experiments, and AO characterization from the perspective of a visible light imager with a night of engineering time at Keck Observatory. I isolate strategies to improve Keck-II AO performance in visible light with updated wavefront sensing and vibration mitigation strategies. ORKID performance demonstrates how challenging visible light AO is, and the need for advanced control techniques to deliver high quality PSFs. I discuss the process of commissioning a data-driven predictive wavefront controller into the new Keck-II real time controller. I present results from bench and on-sky time that show a predictor consistently out-performing an integrator in wavefront residuals and in some imaging performance metrics. This implementation of predictive wavefront control currently represents the only predictive controller running on an 8-10 meter telescope. Finally, I summarize my work to make the first comparison of predictive wavefront control (PWFC) methods side-by-side, focusing on data and model-driven predictors. I find that model-driven PWFC is better at learning information from data collected during on-sky observations. In total, my thesis work explores cutting edge exoplanetary science achievable with today's instrumentation, characterizes and improves current state-of-the-

art AO instrumentation, and contributes to the next generation of extreme AO.

In the humbler works of my salad days I gathered up my poor researches in heroic strophes, but here I have sifted them into a series of dramatic scenes and avoided through omission the pernicious voluptuousness of pagan writers.

— *Preface to the Plays of Hrotswitha, German Religious and Virgin of the Saxon Race, 1923*

I have a lot of people who, through contributions both large and larger, have helped me make it here today. Thank you to Frank Mikan, the first teacher who got me excited about astronomy, who let me and my weird nerd friends hang around our little high school observatory, and would lock me in the lab so I could skip morning chapel and nobody could come check up on me. Thank you to Anna Sajina and Tim Atherton, who first taught me what kind of scientist and academic I want to be, and were endlessly patient with me when I was a clueless undergrad. Thank you to Remi Soummer and Laurent Pueyo, who first helped me discover why I love high contrast imaging, helped me find my way into grad school, and have continued to support my career even though I haven't really thought about a space telescope in years. Thank you to Marshall Perrin, the first person who taught me senior academics don't want to do paperwork they just want to solve little puzzles with you, whose encouraging words right as I was starting grad school have stuck with me for the past 6 years. Thank you to my postdoc mentor Kyle Franson, who talked me through a million postdoc applications and was incredibly understanding when I needed to freak out about nothing and everything. Thank you to Nour Skaf, Jaren Ashcraft, Saavi Perara, Emiel Por, and Maissa Salama for being my astronomical and emotional support (at-some-point-was-a-)postdocs. Thank you to Sebastiaan Haffert, who I totally did not use as high contrast

imaging wikipedia while I was writing my fellowship applications, for becoming a staple of my conference community, and who whether or not it was because I made him promise when he would hire me when he got his fancy faculty job, offered me my dream job. Thank you especially to Maaike van Kooten, who held my hand when I was first trying to understand AO, and made my first few years of grad school feel possible. Thank you to Don Gavel for taking time out of his retirement to help me understand control theory, and Claire Max for lending me a physical copy of Hardy that has been my constant companion throughout grad school. Thank you to all of the members of my thesis committee, Lucie Leboulleux, Bruce Macintosh, and especially to Phil Hinz and Max Millar-Blanchaer for your support these last few months getting something that walks and quacks like a thesis together despite all odds. Thank you a hundred times to my amazing advisor Rebecca Jensen-Clem. You are the first successful academic I know who has modeled work-life balance, you take academic game-playing as seriously as I do and always have good advice, and you give amazing quality feedback at lightning speed. I have learned a lot from you and I expect I still have a lot more.

Thank you to the astoundingly excellent grad community at UC Santa Cruz. We are so lucky to have a weird gay little enclave and I have been so lucky to have been a part of it. Thank you to my carpool buddy Diego, my

wrangler jeep wrangler Courtney, and Isabel for taking over union organizing and showing me everything I could and should have been doing. Thank you to my absolutely perfect pal and office mate Kyle. It is tragic that we were unable to solve our two-body office mate problem and I will never tell you if my new office mate is funnier and cooler than you. Thank you to Alex, Anne, Arcelia, Dom, Nicholas, and Sierra. Even though I will never forgive you for graduating first and leaving me behind this last year our amazing years prior almost make up for it. Thank to Ethan for learning how to scuba dive with me even if your ear has never been the same. Thank you to Tyler for roasting a pig while I defend my thesis just because the bit was too good to pass up. Thank you to Maddy, Diana, and Sagnick, my cryptic cohort. We may be the most peculiar cohort in the department but we are also the most powerful. I would not have made it through covid without you guys.

Thank you to my amazing friends in Santa Cruz. Ken and Erin I have loved getting to know you, getting up to low-impact hijinks, and seeing you become parents. McKenna and Martha, you have both been so kind to me as I wax and wane through being a contactable and real person and I always know a little hike is with you guys is waiting for me whenever I hold still long enough to remember to ask you about it. Lara, I have loved being you gym/tan/laundry buddy. You have been so kind to me in the past year when we have often felt

like we were the last two people left in Santa Cruz. To the wonderful Space Telescope crew, even if our friendships were mostly forged in being in our early 20s and making bad decisions despite nominally being working professionals, those relationships have crystallized into something I depend on. Thank you to Clare and Varun for letting me be your third roommate, to Matthew and Keira for trusting me to officiate your wedding even though I made you talk about a lobster during your wedding vows. Thank you to Shannon for always yes-anding a silly trip and providing consistently excellent vibes. Iva, our friendship has taken us all over the world together. You are someone who I look up to as a friend and as an academic, and I can't wait to only live a few hours via train away from you. I hope we trade off tourists in South America again soon.

To my wonderful college/vacation/you're-stuck-with-me-now pals Andy, Danny, Ronny (it's a new thing I'm trying?) it has been so fun to go adventures with you and convert our college relationships into adult ones. I'm so lucky to get to tell you every year where I have been whisked off to and you dutifully appear there to keep me company. Monica and Maya, despite the fact that we have been friends since the start of the paleolithic era, I can only hope we will get to continue that for twice as long. For 20 years you guys have been there and I can't imagine what it would be like without you. I'm sorry I am

leaving the bay area now that we finally all live in the same place again – I have to assume the universe knew we were too powerful when concentrated in a 100 mile radius. Thank you to my mom for being endlessly supportive as I traipse all over the world and don't visit home nearly as often as I should. Thank you to my dad, who if he were here to see this would probably think my little diagram of a pig is just as silly as I do.

Chapter 1

Introduction

1.1 Finding an exo-Earth

There are planets where it rains rubies. Specifically, some planets orbiting alien suns show evidence of clouds made of corundum, which when it condenses on Earth becomes rubies and sapphires [287]. Since the discovery of the first exoplanet, 51 Pegasi b in 1995 [169], more than six thousand planets have been observed. Now our field's most pressing priority is not just detecting planets, but to characterize these planets, understanding their compositions, atmospheres, formation pathways, and ultimately, whether or not they resemble Earth. For years, the gold standard of exoplanet astronomy has been searching for the signals of planets with liquid water, plant life, methane, and generally

signs of disequilibrium chemistry that point to something organic, or the potential for something organic to thrive.

While previous scientists theorized that life could exist in unusual forms, like floating creatures in the hot gas atmosphere of Jupiter [235], now it is generally accepted that for planets to host life they will likely resemble the Earth, and therefore be a particular size and a particular distance from their host star: these two characteristics allow a planet to be rocky (as opposed to made of gas or ice or other inhospitable things), retain an atmosphere, and be a temperature that allows clouds and organic life to form. For example, for a G type star (that resembles our sun), broadest estimates suggest a habitable planet with a biosphere could exist within 0.38-10 astronomical units (AU) around its host star [309, 212]. The habitable zones of M-type stars (fainter, redder stars with big puffy envelopes), due to their cooler temperatures, are even smaller, on the order of 0.03-1.1 AU [136, 137]. Within that zone, a planet must be $\sim \frac{1}{2}$ to 2 times the size of the Earth to be habitable; any larger and it begins to accrete too much gas, leading to high surface pressures that make planets inhospitable [201]. These physical properties lead to stringent technical requirements for the future instruments that will be able to detect faint planets very near their suns.

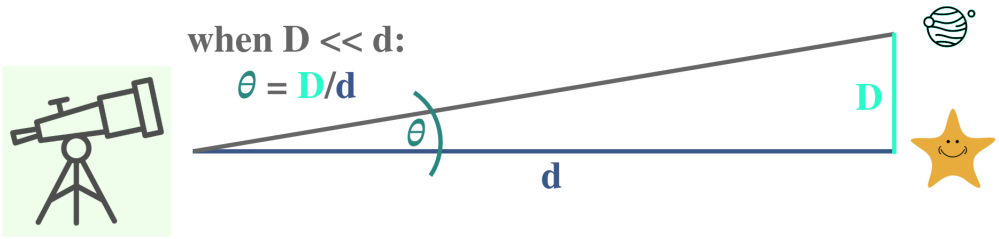


Figure 1.1: Habitable zones determine feasible separations between host star and planet for exo-Earths. Given a potential planet's distance from its host star and the system's distance from (our) Earth we can estimate how close we must be able to observe next to a star in arcseconds (often cited as miliarcseconds or mas), which sets the inner working angle (or nearness to star) requirements for our telescopes.

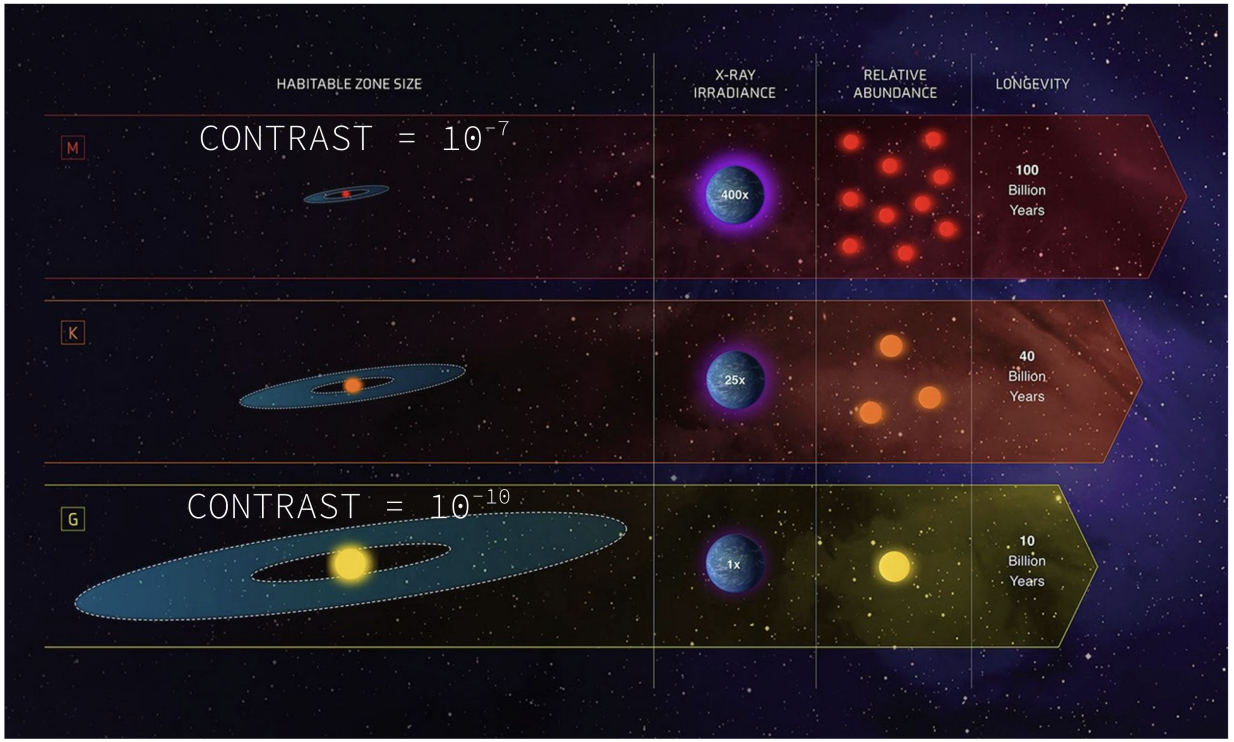


Figure 1.2: The habitable zone is dependent on the temperature of the star. Hotter G (solar-type) stars could sufficiently heat planets, even if they are further away, as compared to cooler M stars with much smaller habitable zones. But brighter stars are at a higher contrast, requiring better nulling of stray starlight to resolve a planet. Figure adapted from NASA, ESA, and Z. Levy (STScI).

Using the small angle approximation (shown in Figure 1.1) we can estimate how close potential planets would be to their host star in seconds of arc in astronomical observations. The nearest G type star to earth, Alpha Centauri, is 4.3 light years (or 1.3 pc) away, meaning its habitable zone is about 7700 mas from its star. However, more targets will exist at closer to 10pc from Earth, with habitable zones of 38-1000 mas for sun-like stars, and 3-110 mas for M-type stars. Given the amount of light a small Earth-like planet will reflect, searching for Earth-like planets around M or G-type stars will mean searching for a planet that is at a contrast of 10^{-7} to 10^{-10} to its host star, as shown in Figure 1.2. These constraints are closely related, as the further a planet is from its host star, the less stellar light it will reflect, making it at a higher contrast. For example, an exo-Earth 8000 mas (10 AU) from Alpha Centauri would have a contrast of $\sim 4 \times 10^{-12}$, whereas an exo-Earth only 300 mas (0.4 AU) would have a contrast of $\sim 2.5 \times 10^{-9}$. These two constraints and their tradeoff specify our need to see extremely faint objects extremely close to their host stars. This problem is colloquially likened to searching for a firefly next to a lighthouse, which is why it is often referred to as high contrast imaging.

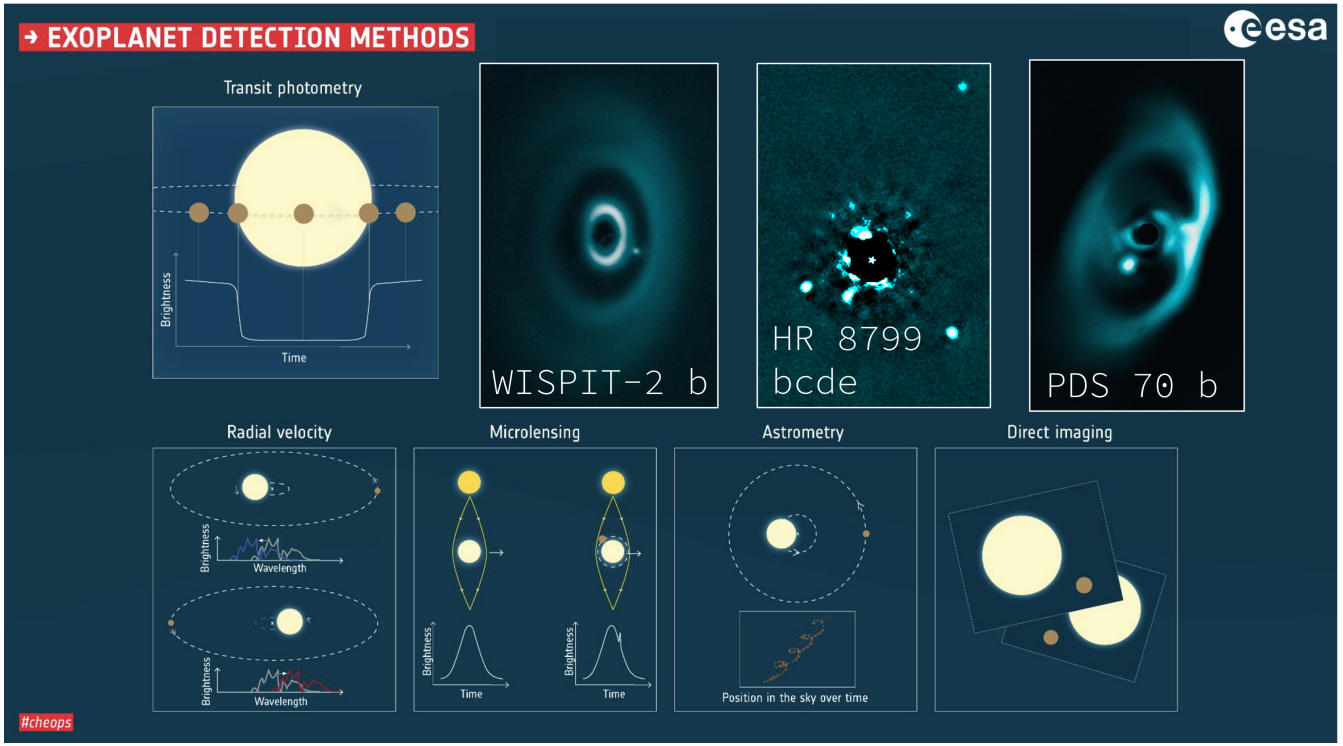


Figure 1.3: A summary of exoplanet detection methods. Indirect methods include transit photometry, radial velocity, microlensing, and astrometry. Direct imaging is unique in that it directly resolves the light of the planet from its host star, instead of looking for variations in the star to infer an invisible planet. Figure adapted from ESA. Added to this diagram are three famous systems with directly imaged planets: WISPIT 2 b [274], HR 8799 bcde (image courtesy of Jason Wang), and PDS 70 b [185].

1.2 High contrast imaging

1.2.1 Direct and indirect methods of exoplanet characterization.

In particular, this thesis will comment on imaging, or direct imaging, of substellar companions at a high contrast to their host star, as a detection and characterization method for exoplanets. Direct imaging refers to resolving the light of a planet uniquely from the light of its host star. This is distinct from other methods of exoplanet detection and characterization, which infer the presence of a planet and its properties by watching its host star do something unexpected.

For example, the radial velocity method uses subtle shifts in the spectrum of a star to estimate the presence of an unseen planet tugging on the star, which induces a Doppler shift in the spectrum as the star moves subtly towards or away from us. This gravitational pull on the star by a second body is not only observable in shifted spectra; high precision astrometry searches for this same effect by seeking subtle deviations in the star's path over time. Transit photometry monitors for variations in stellar brightness, which point to a transiting planet blocking out star light. Observing spectroscopically during this transit (or transit spectroscopy) provides information on the absorption

spectrum of the planet, as it blocks out the emission spectrum of the star. Finally, microlensing waits for the perfect geometry such that a planet acts as a micro gravitational lens, creating an extra bump in brightness in a star. Figure 1.3 diagrams these methods and includes some famous examples of directly imaged planets.

However indirect methods require detailed knowledge of a host star. A particularly active star could mimic a radial velocity wiggle, stellar variability could mimic a planet transit, etc, as shown in Figure 1.4. Similarly, while transit photometry and spectroscopy is the most prolific of exoplanet detection methods (as shown in Figure 1.4), it requires a planet to geometrically transit in front of a star. In contrast, direct imaging is limited not by stellar activity or geometry, but by technology, making it excellent fodder for a PhD thesis.

1.2.2 Coronagraphy and differential imaging techniques

The challenging part of direct imaging comes in extinguishing enough light from a host star to reveal the presence of a planet. This is typically performed with coronagraphy, originally designed to observe solar corona, but used today to block out as much light from a star as possible. A classic coronagraph includes two masks, one placed at the focal (or image) plane of an optical system to null the stellar core, and one placed in the pupil plane (or plane conjugate to

the telescope aperture, also called the Fourier plane), to mask light that is diffracted to the edges of the pupil. Figure 1.5 shows an example of a classic Lyot coronagraph and how light is suppressed at each stage.

In tandem with optical suppression via coronagraphy, differential imaging methods are typically used to improve contrast even further. Chapter 2 will discuss data taken with angular, spectral, and polarimetric differential imaging (ADI, SDI, and PDI), so I will briefly summarize these methods here. For a full description of differential imaging methods see Ref. [69].

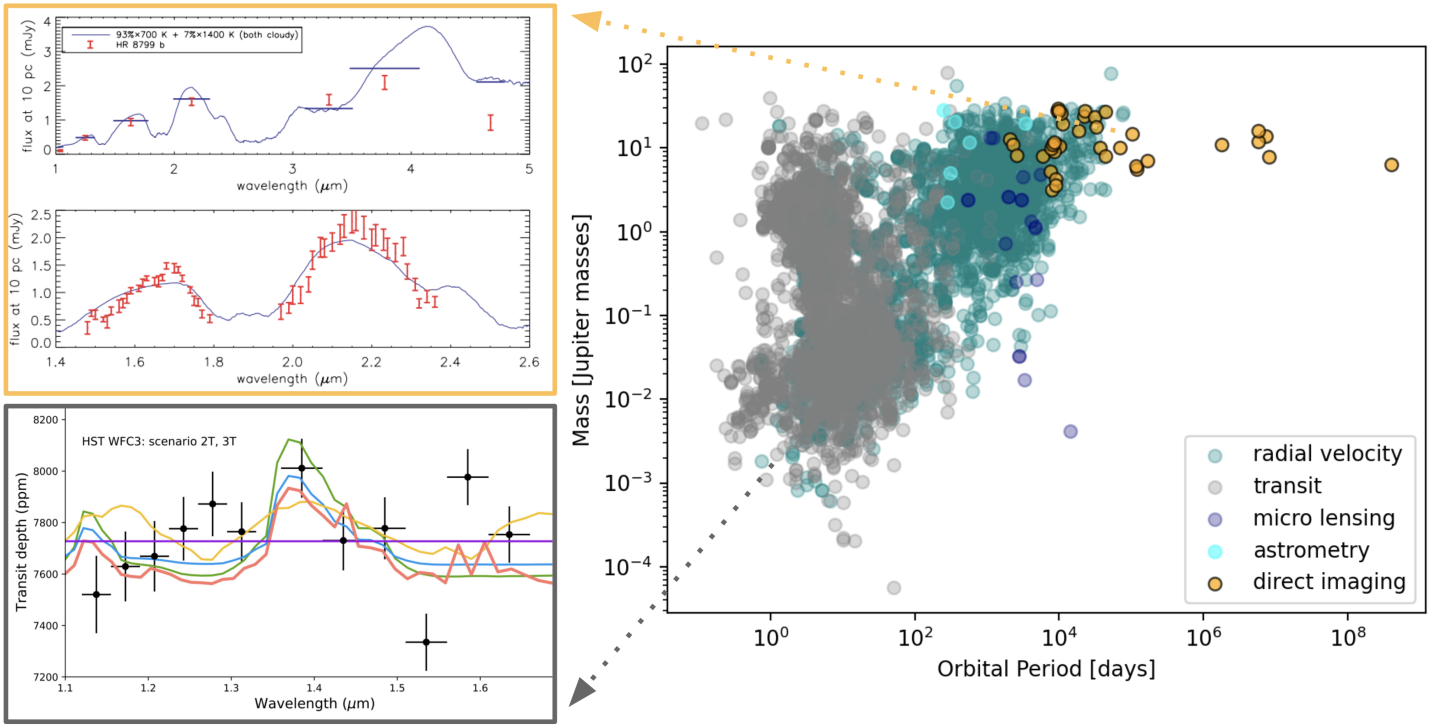


Figure 1.4: Top Left: Directly imaged spectroscopy of HR 8799 b [257], successfully fit to an atmospheric profile. Bottom Left: Transmission spectroscopy of Trappist-1 e [287]; these observations were dominated by stellar fluctuations of Trappist-1 and did not result in a successful fit to an atmospheric profile. Right: Mass period diagram of every detected planet, with data from the NASA Exoplanet Archive as of May 29th [39]. Direct imaging (shown in orange) is not nearly as productive as indirect methods, but offers one of the few pathways to directly detect planetary atmospheres. To directly image Earths we need to develop technology that will let us detect planets with lower masses and shorter orbital periods.

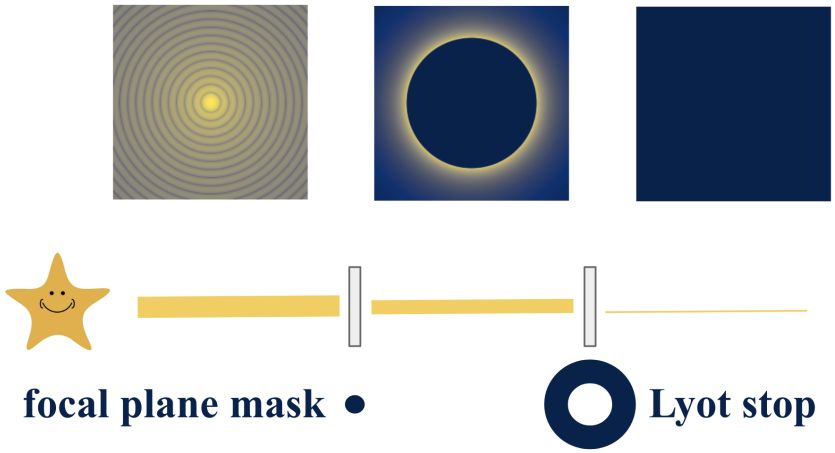


Figure 1.5: Light from a star is nulled with two optical masks, one in the focal plane to null the central light, and one in the pupil plane to null light that diffracts around the edge of the pupil, named a Lyot stop after Bernard Lyot, inventor of the coronagraph. The images at the top are simulated examples of the state of the wavefront while using a coronagraph, starting with left) a perfect unobscured PSF with airy rings, middle) most of the central light is obscured by the focal plane mask, but light has diffracted around the edges of the pupil, after the Lyot stop blocks that light right) a dark zone appears. A coronagraph makes a dark region where we could observe the faint signal of an exoplanet. Simulated images made with `hcipy` [213].

Angular differential imaging (ADI [167]) allows the sky background to rotate during an observation, as opposed to a more classic strategy where the target is held fixed. Then, principal component analysis uses these many realizations of the same point-spread function with varying angular positions to subtract off noise without risking subtracting out the planet signal (as it has been dispersed amongst many frames, unlike the telescope PSF).

Spectral differential imaging (SDI, [259]) often uses an integral field spec-

trograph (IFS), where cubes of spectra across every resolution element of an image are taken simultaneously. Through wavelength space the planet will remain at a constant location, while the shape of the point-spread function and location of speckles of aberrated light will change. This technique can also employ knowledge of where a planet least likely to have signal, creating a planet-free reference image for speckle subtraction.

Finally, polarimetric differential imaging (PDI [138]), exploits the fact that astrophysical objects may have a subtle polarized signal. If a planet has its own self-luminous polarized signature, is embedded within certain kinds of disks or dusty clouds, or reflects light, that induces a polarized signal. However, stars are less likely to be polarized, which allows for a gain in contrast in polarized light where so much stellar light is naturally extinguished. Furthermore, (as I will describe in Chapter 1.3), PDI uses a double differencing method to subtract speckles out of polarimetric images at orthogonal polarization states. Many of these techniques can be used together, for example, a common combination is ADI and SDI by observing with an IFS and allowing your sky background to rotate.

With different observatories, observational techniques, and data reduction strategies, many different contrasts have technically been achieved (with performance estimates down to 10^{-8} for infrared (IR) imaging with the James Webb

Space Telescope [153]), however, the faintest exoplanet imaged to date is 51 Eri b, with a contrast of 2×10^{-6} in J band (at λ =1250nm), at a distance of ∼475 mas from its host star [51, 161]. This is an order of magnitude brighter and 4 times further from its host star than we need to achieve with ground-based high contrast imaging to find an exo-Earth around an M-dwarf.

1.3 Polarimetry and polarimetric differential imaging

1.3.1 Physical properties of polarized light

Light in a basic sense is an electromagnetic wave with an electric (E) and magnetic (B) component. Our detectors are typically only sensitive to the E field, so I will focus on how the E field propagates through space. The electric field component is complex, with an x and y direction in space. The polarization state of light describes how the x and y components of the E field waves become misaligned with each other, for example if the x component experiences a different amount of optical path delay by some scattering feature. Figures 1.6 and 1.7 demonstrate how misalignment between x and y components of the electric field create polarized light. We often summarize the state of

polarized light with the Stokes vector:

$$\mathbf{S} = \begin{bmatrix} I \\ Q \\ U \\ V \end{bmatrix} \tag{1.1}$$

where I represents total intensity, Q and U represent the linear polarization states along the x and y directions, and V represents circular polarization, as diagrammd in Figure 1.6.

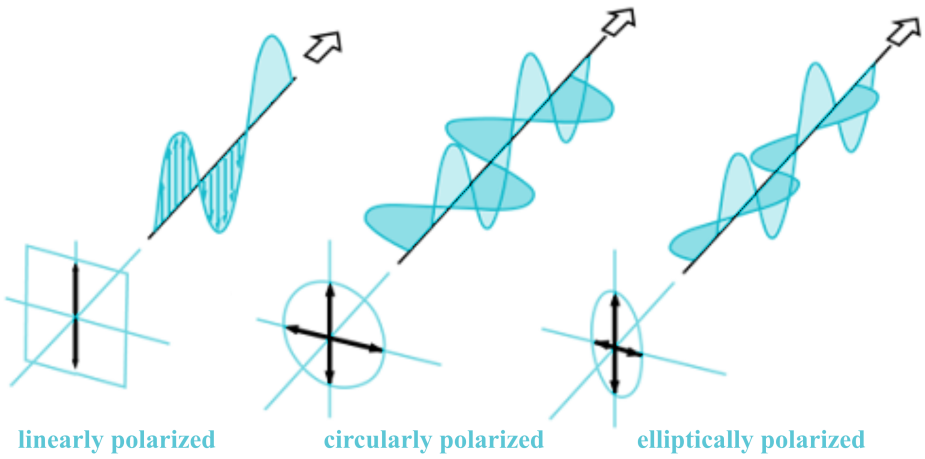


Figure 1.6: As the x and y components of the electric field diverge diverge, this asymmetry creates a polarized signal. This creates linearly, circularly, or elliptically polarized light. Figure adapted from GSU Astronomy HyperPhysics.

Most light is a mix of polarization states, however, certain scattering and reflection events create strong polarized signatures, for example light reflecting

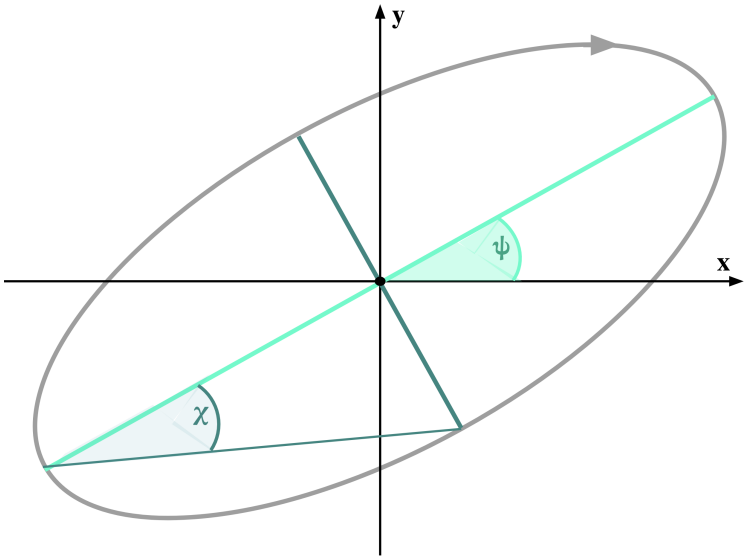


Figure 1.7: Polarization states are often described using the polarization ellipse. The amplitude of the x and y components (which can be summarized with χ), as well as their directional offset ψ determine if it is linear or circularly polarized light.

off of water, or the perfect natural polarizer the Chrysina gloriosa scarab beetle, as shown in Figure 1.8. Polarizers filter light that consists of these mixed states and allow only one state through. If a strong density of signal exists in a particular polarization state, viewing the object through a polarizer has the effect of highlighting it from its surroundings, e.g., the scarab beetle likely has vision optimized for polarized light and can spot its fellow beetles more effectively, even compared to the glare from water or shiny leaves.

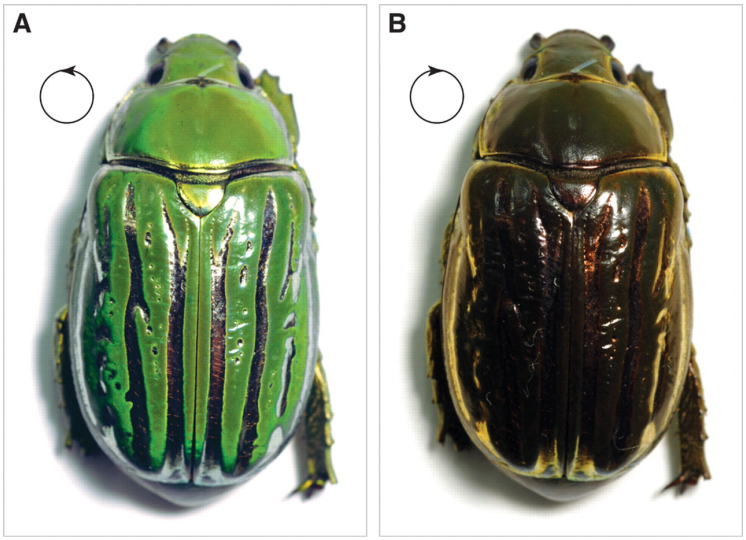


Figure 1.8: The Chrysina gloriosa scarab beetle is a perfect natural polarizer. The micro-structure of its exoskeleton absorbs all right hand circularly polarized light, and leaves only left hand circularly polarized light reflected. Left: Standard imaging of the beetle. Right: Using a quarter wave plate and a polarizer, left hand circularly polarized light is blocked, leaving the beetle to appear dark and colorless. Figure from Ref. [249].

1.3.2 Finding Earth-like planets with reflected light polarimetry

Within the wealth differential imaging techniques, my observing work explores polarimetric differential imaging (PDI). PDI has three major advantages. First, as stars are largely unpolarized, you often gain up to an order of magnitude in contrast as compared to total intensity imaging (this does rely on some polarized signal from the companion which is not always the case). Current polarimeters (for example ZIMPOL at the Very Large Telescope) have shown that PDI can

achieve the photon noise limit (e.g., limited by the noise of the photons hitting the detector) [114]. Additionally, total intensity information is still available within a PDI dataset; PDI data can be further reduced and used as total intensity direct imaging data. And finally, polarized light has complimentary information to standard full intensity data. This includes spectropolarimetric data that describes cloud morphology, grain sizes, and cloud depths [264], as well as surface properties like foliage and oceans [231, 282].

In particular, polarimetric imaging will be vital to resolve Earth-like features in future habitable planets. While direct spectroscopy will tell us what is in the atmosphere of a planet, polarimetry is one of the few techniques that will tell us what is on the surface of a planet, and will, for example, be one of the only ways to determine if a planet has liquid oceans [282].

For example Figure 1.9 shows a real world example of a reflected light water feature emitting a visible polarized signal.

1.3.3 Disk integrated polarimetry of giant self-luminous companions

Unlike total intensity, polarized light has a vector describing the direction of the signal. When combined with other polarized signatures of other directions, this signal can cancel. This is relevant in the context of imaging self-luminous

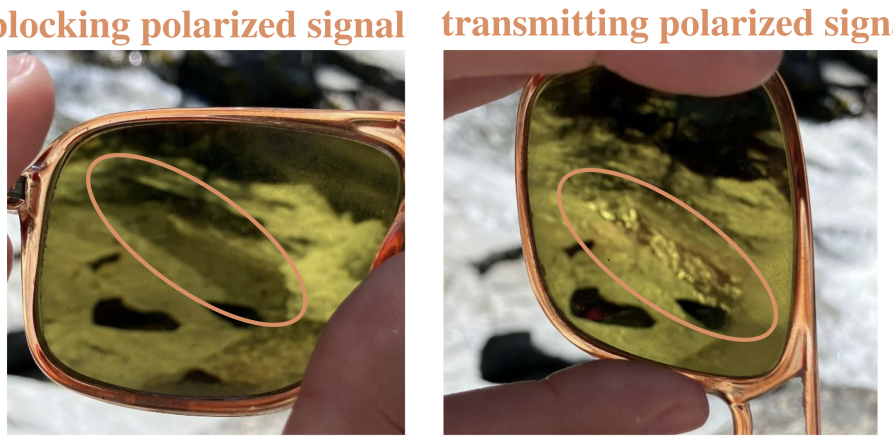


Figure 1.9: Light reflecting off of water will have a polarized flint feature in visible light. Only in the right-most frame, where polarized signal is transmitted, is a visible light water glint feature visible, due to water running over the falls. This is a visible light image of Earth-like water reflecting the light of a Sun-like star taken from Carlon Falls, California in the Sierras.

point-like sources (e.g., planets), as the single point will represent the sum of information across the planetary disk, or disk-integrated signal. This means that self-luminous polarized signatures will be indicative not just of scattering, but of strong asymmetric scattering features. For example, a completely uniform atmosphere scattering a planets thermal emission would perfectly cancel as you sum over the signal from the entire planet, but asymmetries in the cloud coverage would create a non-uniform net signal that could be detected. This is visualized in Figure 1.10.

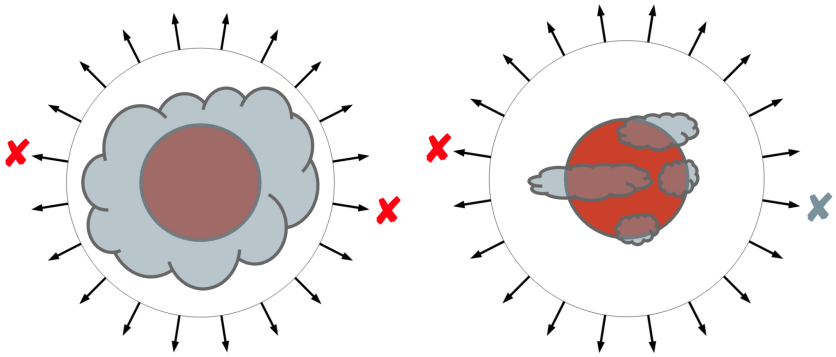


Figure 1.10: Visual representation of disk-integrated polarized signal for a planet with (left) uniform cloud coverage and (right) asymmetric cloud coverage. As bulk cloud features scatter light asymmetrically, polarized signals do not perfectly cancel, and result in a net signal that is one of the only observational methods to detect cloud bands or hot spots. Note, polarization vectors are pseudo-vectors and would not directionally point outward, this has been simplified for demonstration.

1.3.4 Polarimetric differential imaging in practice

Polarimetric Differential Imaging (PDI) uses a series of optics to split light into orthogonal polarized states. This is often done with a Wollaston prism and a modulating half-wave plate (HWP). The Wollaston prism splits light in two, creating a left and right side copy of the images at orthogonal polarization states, and a half wave plate modulates the signal, allowing us to see the Q and U states. In practice a full HWP cycle rotates through four angles: 0, 45, 22, 67°. Information on linear polarization (Q) is contained in the HWP=0: 0-90° state and information on the linear polarization (U) is contained in the HWP=22: 45-135° state. Using the additional states (90-0 and 135-45

respectively) creates an inverse image, which allows for subtraction of errors due to instrumental polarization (if they are downstream of the half-wave plate) and aberrated stellar light. this subtraction process is called double differencing. Using an optic other than a half-wave plate along with additional calibration could allow most dual-channel polarimeters to resolve circular polarization information, but it has yet to be prioritized as an observing mode. Because the Wollaston prism evenly splits incoming light, the total intensity can be recovered by adding the left and right hand sides of each image together. The double differencing method used to recovered Q, U, and I polarization is shown in Figure 1.11.

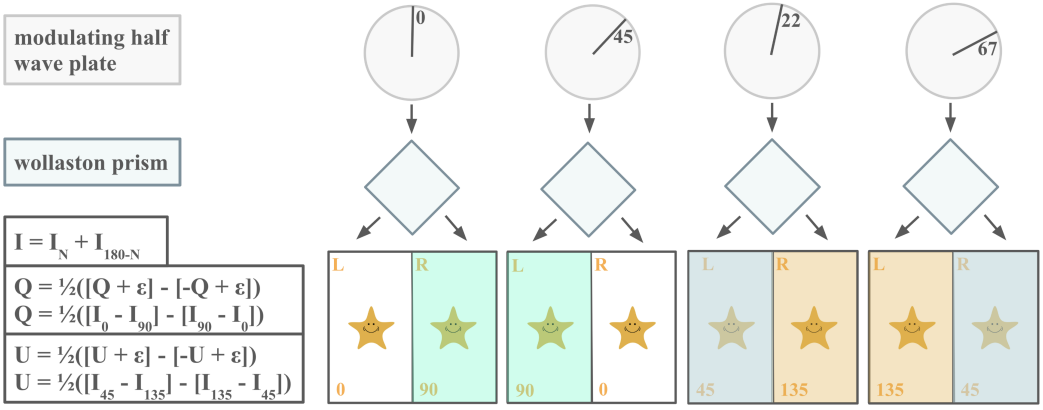


Figure 1.11: Diagram outlining the method and components of polarimetric differential imaging in an ideal system. A Wollaston prism splits light into two orthogonal polarization states, and a half-wave plate induces the delay (of N degrees) that allows for the extraction of I, Q, and U polarization states. Instrumental polarization effects ϵ cancel when double differencing is applied.

1.4 Extreme adaptive optics and visible light adaptive optics

1.4.1 Speckles and atmospheric turbulence

To find the ultra-faint signatures (in polarized light or otherwise) of planets at a high contrast from their host stars, we need technology that will null starlight extremely well. The tool of choice, coronagraphy, often depends on symmetric wavefronts, or waveforms of incoming light to the telescope. Any aberrations, be they from the instrument (imperfectly shaped optics, vibrations from cryo-coolers, etc) or external conditions like atmospheric turbulence, will impact the ability of coronagraphy to effectively null starlight, and create speckles of starlight in the final science image that are often as bright as planetary mass companions. Figure 1.12 shows an example of the performance of a Vortex coronagraph on a perfect PSF as opposed to a PSF impacted by Earth-atmosphere. In particular, this thesis will discuss wavefront errors induced by atmospheric turbulence and ground-based optical and infrared (IR) instrumentation.

The Earth's atmosphere is made up of small pockets of varying temperature (and therefore index of refraction), which evolve on sub millisecond timescales. One way to combat this is with active control systems, or adaptive optics, which

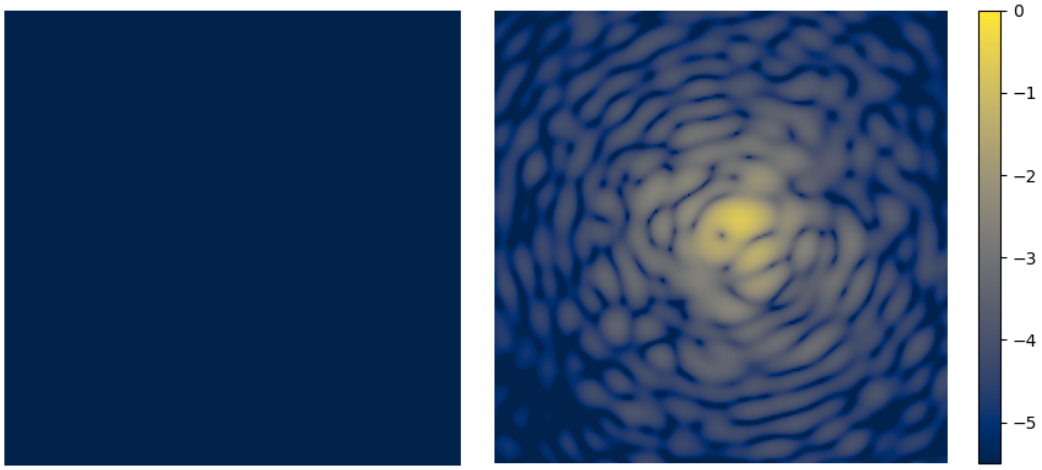


Figure 1.12: Coronagraphy is excellent at extinguishing starlight from a perfect source, but its performance degrades for more complicated wavefronts. Left: Simulated image of nulled light from a vortex coronagraph. Right: the same system with wavefront errors included creates speckles in the dark zone that obscure potential planet observations. The left-most scale bar, which indicates coronagraphic contrast (with arbitrary unites), applies to both images. Figure generated using `hcipy` [213].

sense and correct for the atmosphere (or other aberrations) in real time and correct them before they can get to our final science images, akin to putting a constantly evolving pair of glasses on our telescope systems. Figure 1.13 shows an example of imaging the planet Neptune with and without adaptive optics.

We expect the predictable elements of atmospheric physics are governed by linear effects [216]. Basic atmospheric modeling for adaptive optics is done with Kolmogorov turbulence [135, 267] that evolves according to the Taylor frozen flow hypothesis [268]. Together these models describe velocities imparted by wind layers, often at multiple levels in the height of the atmosphere, which blow

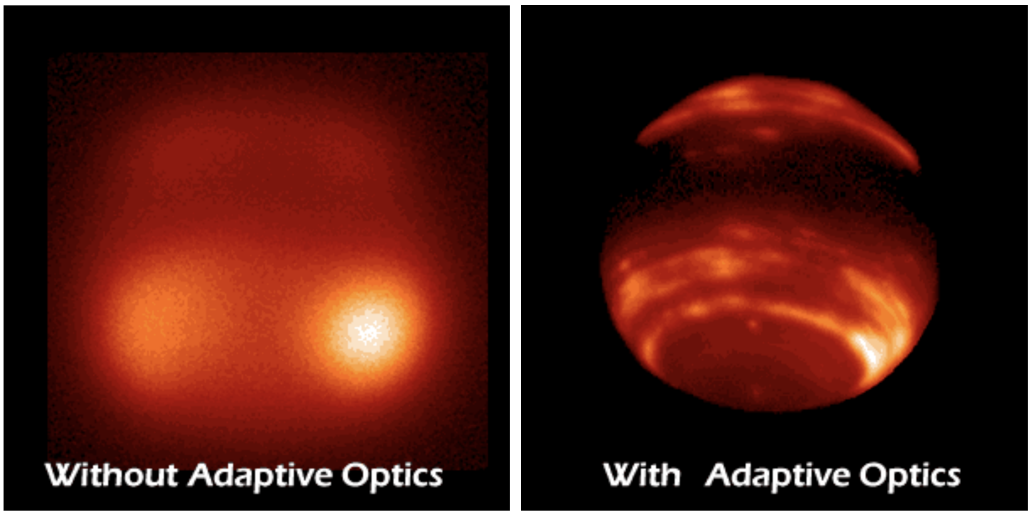


Figure 1.13: Example observation of Neptune with Keck Observatory before (left) and after (right) adaptive optics (AO) correction. Before AO, the observation is limited by the atmospheric seeing, or the spatial scale over which the atmosphere has coherence. With the correction we can see the structures of the planet, and see scales down to the diffraction limit of the telescope. Figure adapted from Claire Max and W. M. Keck Observatory.

static scenes of Kolmogorov spectrum turbulence over the telescope. While Kolmogorov turbulence does not perfectly encapsulate the full atmosphere, it does to within the sensitivity of modern AO systems. [216] Additional atmospheric turbulence (i.e., boiling [60], dome seeing [298]) adds random features that could still be statically encapsulated with linear noise propagation (i.e., Kalman filtering and autoregressive models.)

1.4.2 Anatomy of an adaptive optics system

An adaptive optics system at its simplest is made of 3 fundamental components: a wavefront sensor (which records information on the state of the wavefront after it has been aberrated by Earth-atmosphere), a deformable mirror (which attempts to correct for those aberrations), and a computer that synthesizes information from the wavefront sensor and determines what correction the deformable mirror (DM) should apply. Figure 1.14 shows a control diagram of this system. Figure 1.15 shows the flow of information through a control system and the impact of noise and digitization at each stage.

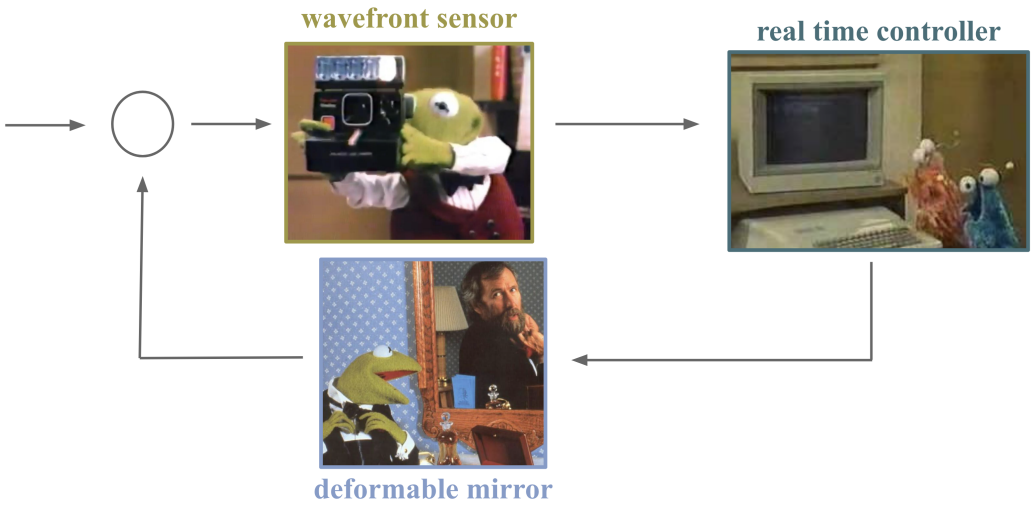


Figure 1.14: A simple adaptive optics system consists of a camera that senses the shape of the incoming light waves (a wavefront sensor), a computer that synthesizes information and calculates a correction out (a real time controller), and a corrective optic that applies a correction out (a deformable mirror). This diagram includes content from Jim Henderson and the Muppets.

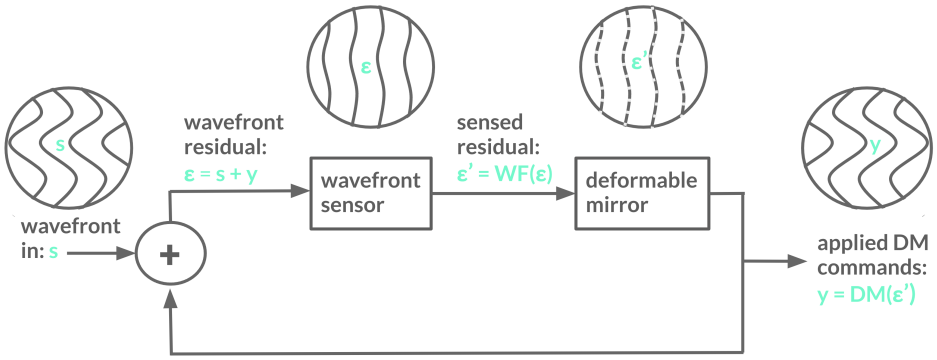


Figure 1.15: The full state of turbulence $\mathbf{s}$ enters the system and immediately encounters the wavefront correction $\mathbf{y}$, noted by the sum junction. The wavefront sensor senses the residual wavefront $\epsilon = \mathbf{s} + \mathbf{y}$, and reports a digitized and imperfect realization ϵ'. Finally, a computer uses this information to calculate the command update to the deformable mirror, which feeds the next iteration.

To date, every extreme AO system runs their standard control for mid and high order modes as an integrator. This refers specifically to a closed-loop integral style controller, where the wavefront sensor is downstream of the DM, senses residual wavefront after the correction, estimates the error with a simple minimization between that residual and a flat reference, and applies that command as an addition to the current shape. For example, for Keck Observatory, the integral control law for a DM command C_{int} at some time t takes the form:

$$C_{int}(t) = k \times C_{int}(t-1) - g(\mathbf{M}_C[W(t) - W_{ref}]) \tag{1.2}$$

where the current state of the atmosphere is estimated with wavefront residuals W as compared to a static reference W_{ref}. $\mathbf{M}_C$ is a command matrix that converts wavefront sensor residuals to the control space of the DM, with some gain that is typically between $0.2 < g < 0.6$ (for some advanced controllers gain is optimized per mode [80, 163]). Finally, as a closed-loop system, the estimate of the residual wavefront error is added to the previous command $C_{int}(t-1)$, with a leak k (usually set to 0.99) that provides further stability to the integral control law [280]. This correction is applied as quickly as possible, with the hope that the speed of the system will keep up with the speed of the evolving atmosphere. In practice this lag-time ranges from 1-5 ms.

Adaptive optics (AO) is a broad collection of methods that is best used when it is tuned for a specific set of requirements, often driven by a kind of science problem. For example, some AO systems work to cover a large (many arcminute) field of view fairly consistently (e.g., optimizing for repeatable pictures of a crowded star field to find subtle star motions to find the galactic center [41]). This thesis focuses on extreme adaptive optics, concerned with getting the highest quality correction (often measured with the highest Strehl ratio or deepest contrast achievabl as compared to a host star) over a small (few arcsecond) field of view where stars may have planetary companions.

1.4.3 Error terms in an adaptive optics system

A few major error terms are common to extreme adaptive optics. Here I briefly outline these terms and discuss estimates at Keck Observatory as a practical example.

Fitting error

Fitting error describes errors induced by lack of correction by the corrective optics, typically a deformable mirror, in the adaptive optics system. Commercially available stroke (or maximum amplitude a DM can apply) is such that this problem is rarely limited by available stroke of the DM, and we typically consider this a degrees of freedom problem. This means that the limiting factor in this correction is actuators at too rough a spacing to adequately encapsulate the variations in turbulence. The atmospheric spatial scale r_0 describes the coherence length of an atmosphere, or the spatial scale over which the correlation of the atmospheric phase is $\leq \frac{1}{e}$. Matching the deformable mirror actuator spacing to r_0 is typically considered the lower limit on fitting error, though practically this may not be sufficient for extreme AO systems [70]. We estimate fitting error analytically with:

$$\sigma_{fitting} = \sqrt{\mu(d_a/r_0)^{5/3}} \tag{1.3}$$

where μ is a constant that describes actuator influence and projection (see more in Ref. [104]) and d_a describes the actuator spacing projected onto the primary mirror.

Measurement error

Measurement error concerns errors in measuring the wavefront accurately, typically considered a function of signal-to-noise (SNR) on the wavefront sensor detector. AO observations use a reference star (or guide star) to act as perfect calibrator. For high contrast imaging, a bright central star to act as a calibrator is usually built in (i.e., a host star), but as we search for planets around fainter stars this problem becomes more complicated. For all but the brightest stars, measurement error is a trade off with fitting error, as increasing degrees of freedom on the DM requires splitting guide star light into more samples (or subapertures), so there is unique information for each DM actuator. For a Shack-Hartmann wavefront sensor (the current WFS at Keck) we calculate measurement error with:

$$\sigma_{measurement} = \eta \chi d_a \frac{\sigma_{spot}}{SNR} \tag{1.4}$$

with the noise propagator η, closed-loop average factor χ, actuator spacing d_a, and Hartmann spot size σ_{spot}. For more details on these constants and system

specifications see Ref. [104]. Finally SNR is a function of photons and detector readnoise

$$SNR = \frac{N_{photons}}{\sqrt{N_{photons} + N_{pix}\sigma^2_{readnoise}}} \tag{1.5}$$

where $N_{pix} = 4$ for a detector that uses 2x2 pixels for each subaperture.

Equation 1.4 assumes a bright star is close enough to the observation to act as a natural guide star (NGS); when an NGS will no longer suffice, an artificial laser beacon (laser guide star or LGS) is used to create an artificial star, often by exciting sodium atoms high in the Earth atmosphere to create a bright calibration source. LGS has a different error profile and does not perfectly sample all of atmospheric turbulence, but does avoid issues of low SNR to the wavefront sensor. For more details on LGS see Ref. [104].

Temporal error

Temporal errors describe errors that occur when the speed of the control loop falls behind the evolution of the atmosphere. With the spatial scale r_0, and a given wind layer blowing scenes of atmospheric turbulence over the telescope, we can estimate a temporal scale $\tau_0 = r_0/v$, where v is the velocity of a wind layer. For all but the brightest stars, temporal error is a tradeoff with measurement error, as running AO quickly requires minimal sampling times,

which reduce SNR to the wavefront sensor. We estimate temporal error with:

$$\sigma_{temporal} = \sqrt{\left(\frac{f_g}{f_c}\right)^{5/3}} \tag{1.6}$$

where f_c is the control frequency (typically around 1/10th of the system timelag), and f_g is the turbulence-specific Greenwood frequency:

$$f_g = 0.426\frac{v}{r_0} = \frac{0.426}{\tau_0} \tag{1.7}$$

We quote r_0 at a specific wavelength of light (often 500nm), and r_0 varies with wavelength to the $\frac{5}{3}$ power. Modern AO systems can run at 1-2kHz, which is enough to keep up with longer wavelength (IR at $\sim \lambda = 2000$ nm, but this becomes challenging for visible light observing, where these timescales are faster than 1kHz. (This will be discussed at length in Chapter 3). Temporal errors can be mitigated with predictive methods that predict where the atmosphere will be in one system timelag's time. (This will be discussed further in Chapters 4 and 5).

Static and quasi-static errors induced by the telescope system

Telescope apertures are far from perfect idealized circles. Most telescopes have an on-axis second mirror, meaning there is a large central obscuration. On the

ground telescopes are heavy and require support structures, which include struts to hold a secondary mirror (often called spiders). Finally, Keck Observatory is fairly unique in that it has a series of abutting hexagonal segments that make up its primary mirror. While Keck is currently the only segmented telescope with facility AO, this will be (and already is) the standard for space and extremely large telescopes, meaning the Extremely Large Telescope (ELT), the Giant Magellan Telescope (GMT), the Thirty Meter Telescope (TMT), the James Webb Space Telescope (JWST), and the Habitable Worlds Observatory (HWO) will all contend with these issues.

Light bending around a sharp edge creates a complex aperture, complicating the the job of coronagraphs that will valiantly try to extinguish as much diffracted star light as possible. For ground-based telescopes, AO residuals will be the limiting factor, not complex apertures. However, from the ground, these complex apertures interact with air flow, turbulence, etc, and create more complex effects. For example, at the Subaru telescope, harsh thermal variations around the telescope spiders along with missing information (where the spider blocks light there is no light to sense) leads to large discontinuous DM solutions across a spider, called petal modes. [149].

As light moves from the telescope primary through the optical system it only encounters more imperfect optics, as well as motors, coolers, materials

that are sensitive to temperature, electronics that vibrate, and the telescope dome itself interacting with air flow over the course of the night. All of these further contribute errors, though these are often sensible and correctable by the AO system (with the caveat that quickly evolving and high spatial frequency features are always tricky).

Segment co-phasing errors

At Keck, the segmented aperture not only creates segment gaps that create a complex aperture, but each segment moves independently, requiring regular alignment with the other segments (or co-phasing) to a common global shape, as well as regular tuning of the individual shape of a segment. Segment co-phasing errors are difficult to sense with many slope wavefront sensors (e.g., the Shack-Hartmann or modulated Pyramid wavefront sensor), and require additional sensing strategies (e.g., using more sensitive wavefront sensors [236], and point by point measurements during the daytime [81]). While these difficulties in sensing make it hard to directly estimate the impact, segment co-phasing errors are suspected to dominate the unknown sources of error (often described as "margin" error) in the Keck AO system [27, 236].

Non common path aberrations

Non common path aberrations (NCPA) occur when the wavefront sensor does not sense the full path of light to the final science instrument. This is the case for most adaptive optics systems, where wavefront sensing is done with sensing optics upstream of the final science instrument. Advanced control strategies circumvent this issue with separate sensing and control based solely on science instrument data [21], incorporating instrument information directly into the controller [146], or by using lucky imaging strategies to update a wavefront sensor reference based on image quality [256]. Practically, NCPA are often addressed with daytime calibrations on an internal source. Most NCPA are static (e.g., optics in the non common path), however some NCPA (thermal drifts, insensible vibrations) evolve with time, making a static calibration routine insufficient.

1.4.4 Keck-II adaptive optics error budget

To demonstrate the impact of all of these errors, here I summarize the estimate of these terms at Keck Observatory, in the context of a natural guide star system on the Keck-II telescope. Note that for brevity I skip some of the smaller terms, so the total high order error and total effective error will not exactly match the terms listed added in quadrature.

Error Term	Cause	WFE
Fitting Error	too few degrees of correction	51 nm
Temporal Error	controller lag time	60 nm
Measurement Error	low wavefront sensor SNR	23 nm
Margin Error	segment misalignment, unknown errors	130 nm
Total High Order Error	error terms added in quadrature	185 nm
Tip/Tilt Errors	low order sensing errors	58 nm †
Total Effective Error	all error terms added in quadrature	198 nm

Table 1.1: Summary of measurements and analytic predictions for error terms of note in the current Keck-II AO system. † Tip/tilt errors are measured from the focal plane; this term can also be quoted as 3.5 mas.

1.4.5 Translating error terms to performance metrics

A common performance metric in an imaging system is the Strehl ratio (SR), which is a metric that describes image quality. The simplest description of SR, is that it is a ratio of maximum intensity in a real image as compared to a perfect version of that image. There is much debate about how exactly you perform that ratio, what is the best way to estimate a perfect image, and how to best reduce and compare data, as described in Ref. [230]. However, a simple way to relate SR and errors in an AO system is the Marechal approximation, sometimes used in tandem with an empirical correction and the impact of a lower order error term. For example, we can estimate the Keck SR given the

above high order error term:

$$\mathrm{SR} = \mathrm{SR}_M \times \mathcal{E} \times \mathrm{SR}_{TT} \tag{1.8}$$

with an empirical correction $\mathcal{E}$ based on [83], the Marechal approximation

$$\mathrm{SR}_M = \exp - \left(\frac{2\pi\sigma_{HO}}{\lambda} \right)^2 \tag{1.9}$$

where σ_{HO} is the high order error term in meters and λ is the wavelength at which we calculate SR. SR_{TT} is the impact to the SR given low order aberrations, which typcially desginate just the impact of the lowest order zernike modes tip and tilt. Strehl Ratio due to residual tip/tilt is calculated given some tip/tilt error in radians σ_{TT}

$$SR_{TT} = \frac{\sigma_D^2}{\sigma_D^2 + \sigma_{TT}^2} \tag{1.10}$$

where $\sigma_D = 0.44\lambda/D$ and D is telescope diameter.

1.4.6 Adaptive optics for visible light astronomy

A root-mean-square (RMS) error term (i.e., all of the AO error budget terms added in quadrature) can be described as a delay in optical path in units of

physical distance (e.g., some nanometer delay). However, this delay has a proportionally varying impact on varying wavelengths. For example 200 nm of error at 656nm (Hydrogen α) represents a third of a wave of delay, whereas it would only be $\frac{1}{10}$ of a wave of delay at 2190nm (K band). Similarly, the characteristic timescale of the atmosphere (driven by a quasi-static wind-layer) is shorter and requires faster AO correction at shorter wavelengths. Assuming the lower performance end of nominal Maunakea-like atmospheric conditions (a characteristic turbulence spatial scale of r_0 =10 cm, reported for 500nm that varies with wavelength to the 6/5 power) and an average wind velocity of $v = 15$ m/s, the characteristic timescale ($\tau_0 = \frac{r_0}{v}$) for at 2190 nm would be 39 ms (i.e., require a sampling time of $\frac{1}{10}$ of that, or 3.9 ms), whereas a visible light AO correction at 656 nm would require a sampling time on the order of 0.9 ms, or an AO system that runs faster than 1 kHz. Figure 1.16 shows an example of the physical wavelength of light compared to the Keck-II error term.

For this reason, AO systems have traditionally focused on correcting near IR light, where AO systems can still provide good performance while allowing more margin for error. Practically, these systems can operate more slowly, with fewer degrees of freedom in the correction, less precise calibration, and with higher tolerances for optical quality while still delivering a PSF that is limited only by

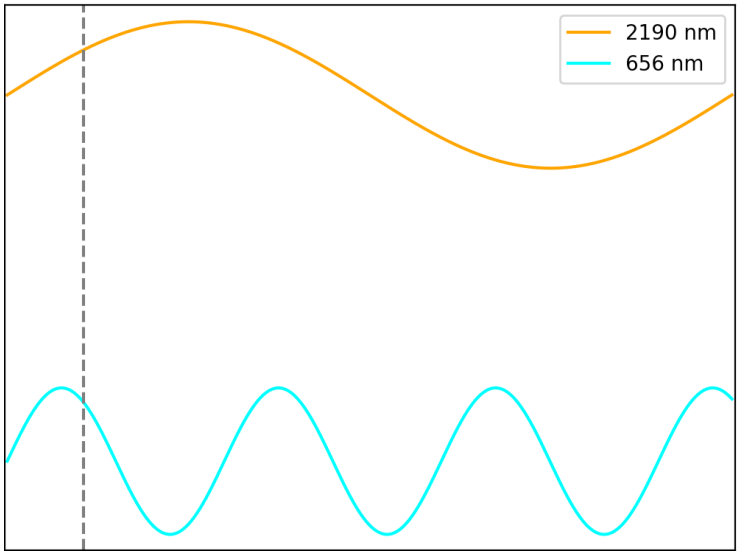


Figure 1.16: Visible light is a physically shorter wavelength than in th IR. This means a Keck-length error term (indicated with the grey dashed line) is a much larger fraction of the optical path in the visible.

the diffraction limit of the telescope. For example, Keck Observatory routinely provides K Band Strehl ratios (SR) of 50-60% [236]. Empirically adjusted analytic estimates suggest the high order error term (or RMS wavefront error across the telescope pupil caused by all effects past residual tip/tilt errors) that provides that SR at 2190nm is 245nm. Translating the 245 nm error term to 656nm in visible light results in an SR of less than 1%.

Better performance in visible light is achievable with extreme AO systems, and some large (6-8 meter) class observatories have extreme AO-fed visible light instruments, including Subaru/VAMPIRES, VLT/ZIMPOL and MagAO-X. These AO systems with higher degrees of freedom in their AO corrections

have seen 20-60 % Strehl Ratio through visible light filters [242, 156]. Even these 20-60% SRs will not be sufficient to image Earths in visible light, and advanced control strategies will be necessary to reduce sources of error limiting AO performance.

1.5 Predictive wavefront control

Many approaches to reducing errors in AO systems exist, and indeed all of them will need to work in tandem to achieve observations of Earth-like planets. Much of this thesis will focus on how to extinguish temporal errors, as described in Section 1.3. By using predictive methods to estimate where a correction will be when an AO loop completes, we effectively overcome the temporal lag of the system. This has the added bonus of offloading measurement error, as the AO system can be slowed, allowing for longer wavefront sensor exposures which increase SNR to the wavefront sensor.

There are many ways to do prediction, and here I summarize the most common methods. In order to comment on and evaluate the performance of control methods, I first set priorities for current and future instruments. I define mandatory criteria and ideal exploration space for new methods within an XAO system in Table 1.5. As I summarize these methods, I examine which methods achieve some of these ideal features and whether or not they sacrifice

mandatory AO requirements.

mandatory criteria	
	achieves at least 60% Strehl Ratio in the IR runs in real time robust over a full night accounts for static NCPAs
ideal criteria	
	accounts for non-linearities does not require high fidelity instrument model high sky coverage at high performance tunable and flexible accounts for *dynamic* NCPAs does not require consistent operator attention

Table 1.2: Here I define criteria deemed mandatory for extreme AO systems (i.e., any method in development cannot sacrifice these), and criteria that we could potentially gain with future technology developments.

In this thesis I will often refer to machine learning methods as a category. I define machine learning as any problem that uses one set of data to make rules for a second set of data [107]. From this definition, machine learning not only includes complicated many-layer neural networks and artificial intelligence engines, but also classic linear regression problems, extrapolation problems, and many basic control engineering and statistical concepts (i.e., Kalman filtering and linear quadratic Gaussian controllers.) In this section, I will overview the most common machine learning methods applied to wavefront control, as well as comment on the overlap in methods.

1.5.1 Kalman filtering and linear quadratic Gaussian control

Kalman filtering [125] is a statistical prediction tool, common to engineering and mathematics alike. Simply stated, a Kalman filter is a robust way to predict the evolution of a property according to a physical law, a Gaussian uncertainty distribution, and data fed back into the system to verify and improve that estimate. Figure 1.17 shows a visual representation of how some property (for our purposes, the measurement from a wavefront sensor) evolves according to some physical law (for our purposes, the wind velocity affecting the atmospheric turbulence, any dynamics in the system, and the control law we are applying), and a cloud of uncertainty that accompanies and is altered by that prediction.

A Kalman filter starts with a state vector $\hat{\mathbf{x}}$, which encapsulates all the variables that we predict at each iteration, and a covariance matrix $\mathbf{P}$, which uses the covariance to estimate the uncertainty of those measurements. Given a physical law $\mathbf{F}$, a control law $\mathbf{B}$, and a noise term $\mathbf{Q}$, we predict the next state vector and covariance matrix with

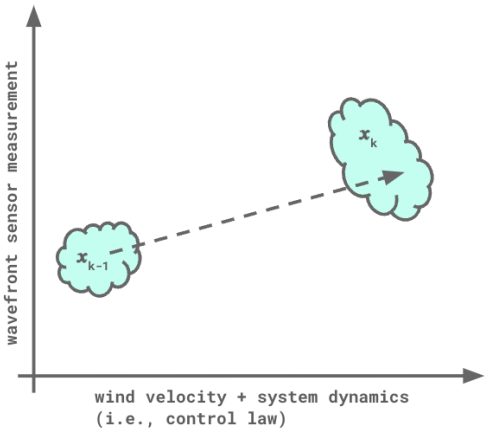


Figure 1.17: "A Kalman filter is used to generate estimates of the states, given the noise-corrupted measurements." [202] A Kalman filter projects a state vector and a Gaussian uncertainty cloud forward according to a physical law.

$$\hat{\mathbf{x}}_{\mathbf{k}} = \mathbf{F}_{\mathbf{k}}\hat{\mathbf{x}}_{\mathbf{k-1}} + \mathbf{B}_{\mathbf{k}}\vec{\mathbf{u}_k} \tag{1.11}$$

$$\mathbf{P}_{\mathbf{k}} = \mathbf{F}_{\mathbf{k}}\mathbf{P}_{\mathbf{k-1}}\mathbf{F}_{\mathbf{k}}^{\mathbf{T}} + \mathbf{Q}_{\mathbf{k}} \tag{1.12}$$

Despite a simple step-by-step update formula, a Kalman filter is not typically updated at each iteration in practice. Instead, a steady-state solution is found after training from previous system data.

A common pair in control engineering is a Kalman filter and a linear quadratic regulator, or a controller that aims to minimize a quadratic cost function at each iteration. This combination is referred to as a linear quadratic

Gaussian (LQG) controller; a controller that assumes Gaussian errors, predicts forward with a linear physical law, and regulates the system with a quadratic cost function. Figure 1.18 shows an example of an LQG controller in the context of wavefront control.

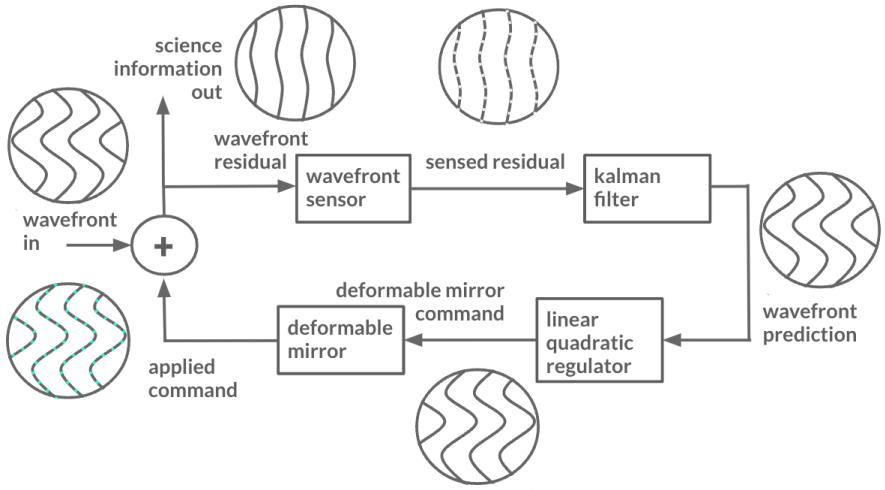


Figure 1.18: Control diagram of a linear quadratic Gaussian (LQG) controller. In this system, a sensed wavefront residual is fed to a Kalman filter, which predicts the future state of the wavefront, calculates the control law with a linear quadratic regulator, and sends that control command to the deformable mirror to correct the system.

$$\mathbf{J} = \mathbf{E}\left\langle \sum_{i=0}^{N} \frac{1}{2}\left[\mathbf{x}^{\mathbf{T}}(t_i)\mathcal{X}\mathbf{x}(t_i) + \mathbf{u}^{\mathbf{T}}(t_i)\mathcal{U}\mathbf{u}(t_i)\right]\right\rangle \tag{1.13}$$

$\mathbf{J}$ is an example cost function [202] used for a linear quadratic regulator, given some state vector $\mathbf{x}$, deformable mirror commands $\mathbf{u}$, and expectation value $\mathbf{E}$. $\mathbf{J}$ minimizes the error in the state as well as DM commands to estimate

the next DM command.

1.5.2 Empirical orthogonal functions

Empirical Orthogonal Functions (EOF) [93] looks for linear relationships in the evolution of the wavefront over time and space. EOF has been successfully run both in simulation [93], during day-time testing at Keck Observatory [116], as well as on-sky at the Keck Observatory [281] and the Subaru Telescope [95]. Here I will briefly summarize the math behind building and applying the EOF predictive filter; the mathematical formalism is taken from Guyon & Males (2017) [93] unless specified.

Empirical orthogonal functions predicts a future state of a wavefront by building a predictive filter that linearly combines previous states of the wavefront. Starting with Guyon, 2017 [93] and a follow up from Jensen-Clem in 2019 [116], I outline our implementation of EOF in (pseudo-)open loop. Given some state of the wavefront (i.e., full scale of the uncorrected turbulence as considered in pseudo open loop), w, with m variables representing wavefront sensor measurements (which for a zonal approach corresponds to the number of deformable mirror actuators) and n associated frames, we build a history

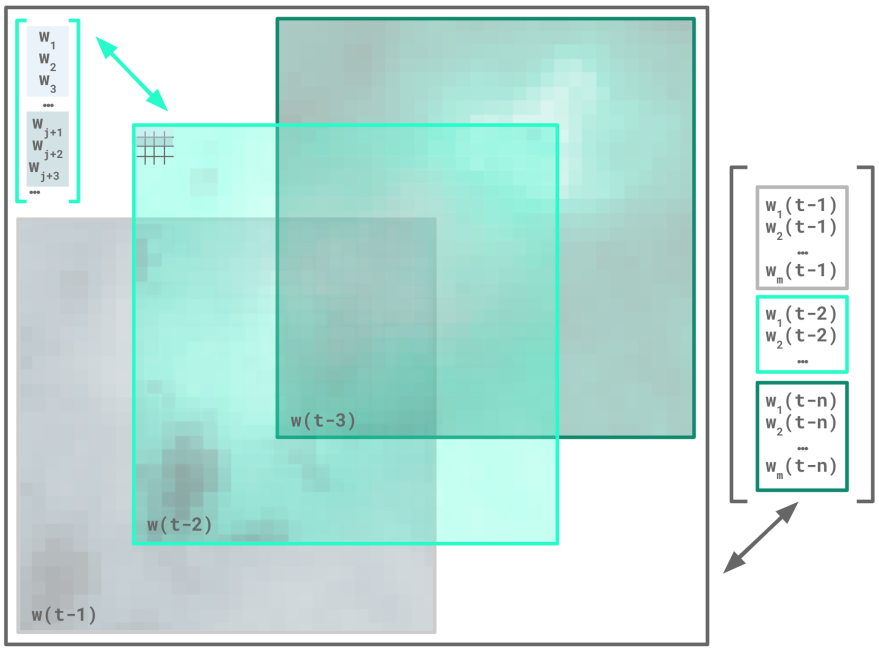


Figure 1.19: The future state of the wavefront is predicted as a linear combination of previous states, w at some time n. The collection of previous states is called the history vector, $\mathbf{h}$.

vector for each subaperture:

$$\mathbf{h}_m(t) = \begin{bmatrix} w(t) \\ w(t-dt) \\ w(t-(n-1)dt) \end{bmatrix} \tag{1.14}$$

A visual representation of the history vector is shown in Figure 1.19.

We build a predictive filter for each mode $\mathbf{F}_m$ that will predict the full

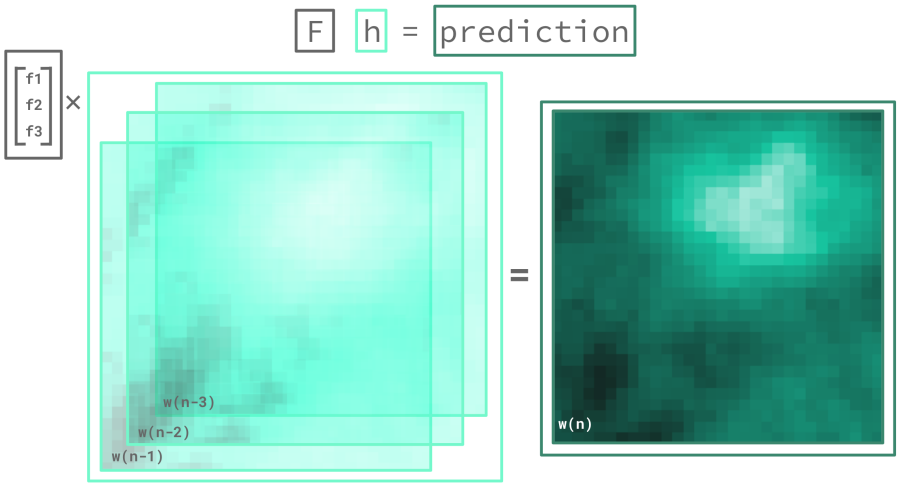


Figure 1.20: Phase screens evolve in time (in sea foam) and a weight is applied to each one (the predictive filter $\mathbf{F} = [f_1, f_2, ...]$, where each f_n contains m weights for m modes) to estimate the final wavefront prediction in teal. Note that mathematically a history vector is a single flat vector that contains the appended information of all 3 screens to predict a final vector that contains a flattened version of the predicted wavefront.

phase at a given point with:

$$\mathbf{F}_m \mathbf{h}_m(t) = w(t + dt) \tag{1.15}$$

This is shown in Figure 1.20. Both the filter and history vector can be written to include information across multiple modes, i.e., $\mathbf{h} = [w_0(t), w_1(t)...w_0(t - dt), w_1(t - dt)...]$, but for ease of notation we leave these as single wavefront sensor measurement/DM actuator filters and predictions.

To find the predictive filter $\mathbf{F}$ we minimize an error term that consists of the difference between the output predicted wavefront (in DM space) and the

true phase at that time. We collect training data $\mathbf{D}$, which contains history vectors and $\mathbf{P}$ which holds the true state (or "future") for each history vector. (I.e., mapping t to its future state at time $t + dt$.)

$$\min_{\mathbf{F}} ||\mathbf{D}^T\mathbf{F}^T - \mathbf{P}^T||^2 \tag{1.16}$$

Solving equation (1.16) requires a pseudo-inverse; Ref. [93] solved this problem with an SVD (singular value decomposition) inversion, but we use a least-squares inversion [116], with regularization constant α (α may be set to 1 for simulations, but is found empirically on-sky [280].)

$$\mathbf{F} = ((\mathbf{D}^T)^\dagger \mathbf{P}^T)^T \tag{1.17}$$

$$\mathbf{F} = \mathbf{P}\mathbf{D}^T(\mathbf{D}\mathbf{D}^T + \alpha\mathbf{I})^{-1} \tag{1.18}$$

Finally, our predictive filter $\mathbf{F}$ holds a coefficient for each previous state, as expressed by a pseudo open loop (POL) reconstruction of the telemetry projected into DM space.

EOF as a predictor works so simply because we consider some wavefront $\mathbf{w}_m$(t) as the full state of turbulence that the atmosphere imparts. In practice, this is not the measurement we have in a typical AO system, which runs in closed-loop, where our wavefront sensor will measure a residual after some DM

correction has been applied. In this implementation, EOF uses pseudo-open loop (POL) reconstructed values of the full wavefront, where information on the applied DM commands and measured wavefront sensors residuals are used to reconstruct what the full state of the turbulence would have been. This is how this is run on-sky both at Subaru/SCExAO [95] and Keck/NIRC2 [280].

However, the math represented by EOF is arguably comparable to that of Dessenne, 1998 [58], which outlined a similar method, but builds a modified history vector using DM commands and wavefront sensor slopes, instead of reconstructing the full state of the turbulence for each input mode. This implementation is more computationally expensive, but encapsulates the real system dynamics and runs in closed-loop.

1.5.3 Neural networks

The previous methods all rely on a linear state space model; the predicted future state can be expressed as a linear combination of previous states and commands. The advantage of this linear model is that one can often find analytical solutions that minimize a quadratic cost function, as discussed in the previous section. However, this assumption of linearity is not valid in all cases. An alternative is to use a nonlinear state-space model. This can be done using Neural Networks (NN), which can act as arbitrary function

approximators. These NN's utilize multiple consecutive matrix multiplications with a so-called activation function in between. This activation function is nonlinear, allowing the NN to model nonlinear functions. These NN's are often visualized by having multiple "hidden" layers between the input and output, where the activation $\mathbf{y}$ in layer i can be mathematically expressed as:

$$\mathbf{y}_i = f(W\mathbf{y}_{i-1} + \mathbf{b}). \tag{1.19}$$

Here, W are the weights of that layer, $\mathbf{b}$ the bias, and $f()$ the activation function. A visualization of a NN is shown in Fig. 1.22. Note that a single-layer NN without an activation function reduces to a familiar linear model, such as in EOF. The amount of hidden layers and number of neurons in each layer is tweaked based on the complexity of the problem. The weights and biases can be optimized by minimizing a chosen loss function on a training dataset. This is commonly done using stochastic gradient descent based optimizers in combination with the backpropagation algorithm. Libraries such as `keras` and `pytorch` hide most of these complexities to end users, which makes building and training basic NN's much easier.

For complex nonlinear problems basic NN's quickly run into the curse of dimensionality, where the number of free parameters explodes. Luckily, the

parameterization of these NN's can be adapted to specific problems by using domain knowledge. This can be seen as a strong form of regularization on the complexity of the model. For example, Convolutional Neural Networks (CNN's) are popular for image processing tasks. They use the knowledge that features are local and translational invariant to drastically reduce the number of free parameters that have to be optimized. For our purposes, CNN's can be used to capture spatial correlation in the wavefront. Alternatively, Recurrent Neural Networks (RNN's) are commonly used for problems involving sequences. They are well suited to deal with the temporal aspect of predictive control. Popular variants of RNN's are Long-short term memory (LSTM) networks and Gated Recurrent Units (GRU).

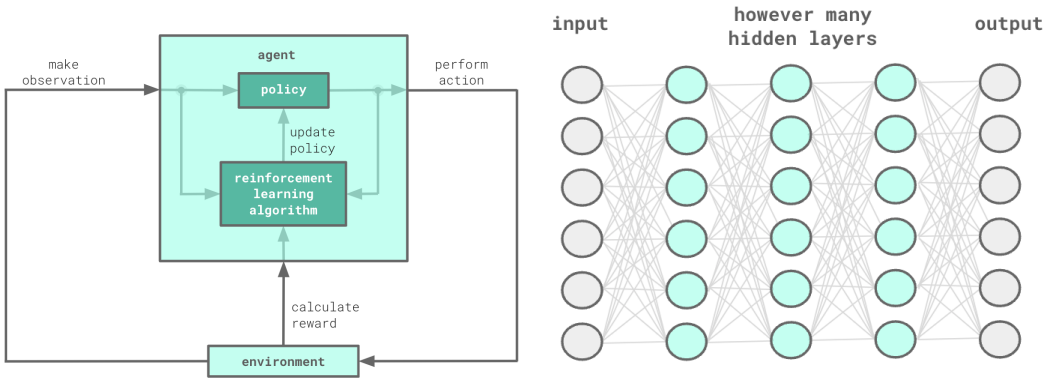


Figure 1.21: Left: Diagram of a reinforcement learning (RL) algorithm. An RL algorithm enacts some policy (control law) on some environment, makes observations, optimizes itself according to a reward. Right: Diagram of a neural network (NN). A neural network finds non-linear connections between layers of nodes.

1.5.4 Reinforcement learning

Reinforcement Learning (RL) is a fast growing field that deals with learning optimal control through interaction with environment. The goal of these algorithms is to find the optimal commands that maximize some arbitrary optimization objective. RL algorithms can be split up into two main branches: model-free and model-based. Note that models are used in both cases.

Model-based RL algorithms explicitly learn a state-space or dynamics model. Using this state-space model and the currently measured state, they aim to find the control command which maximizes the optimization objective at each iteration. This can be done using planning algorithms, which can be quite computionally expensive to run at kHz frequencies. In the case of a linear dynamics model with a quadratic cost function, model-based RL reduces to the LQG problem with online identification of the state-space model.

Instead of learning a dynamics model, model-free RL algorithms aim to directly learn an optimal parameterized control law. Popular choices of model-free RL are variants of the policy gradient algorithm, such as Soft Actor-Critic (SAC) and Deep Deterministic Policy Gradient (DDPG). All these algorithms try to learn a differentiable model of the optimization metric as a function of the state and control command. One can then obtain the gradient of the cost function with respect to the control commands. This can subsequently be used

to optimize the parameters of the control law. Model-free RL methods provide a parameterized control law, with which inference speeds are typically fast enough to run at kHz frequencies. A disadvantage of model-free algorithms are that they are typically less data efficient and would thus more slowly adapt to changing conditions. Promising hybrid methods, that combine the advantages of model-free and model-based algorithms, are also being developed.

1.5.5 Overlaps between methods

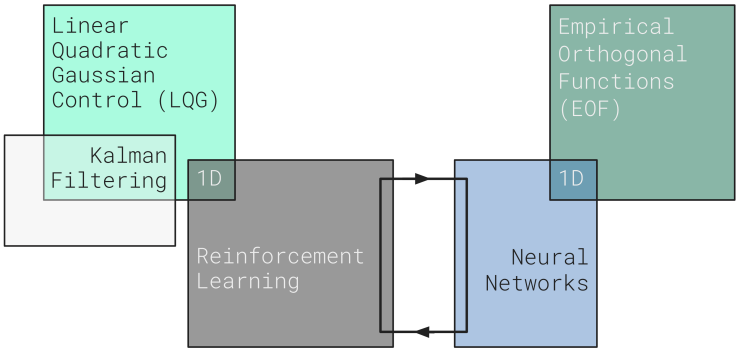


Figure 1.22: Reinforcement learning is a particular way to frame a problem; a linear example of reinforcement learning is a linear quadratic Gaussian controller (which requires a Kalman filter as its estimator.) Neural networks are often used to solve reinforcement learning problems, and empirical orthogonal functions is equivalent to a neural net that has been constrained to a single layer and temporal evolution.

Reinforcement learning and neural networks are two common elements of machine learning that often go hand-in-hand: reinforcement learning as a particular problem type, and neural networks as a common solver for that

problem type. Narrowing the scope of these problems and making them linear shows that an LQG controller is simply a reinforcement learning problem with more conditions about its error propagation, physical law, and cost function. In particular, a mathematical analysis of the overlap between a reinforcement learning problem and an LQG controller is available in Ref. [225]. Kalman filtering is often utilized as an element of an LQG controller, though one of the methods I describe is closer to a pure Kalman filter approach. Finally, EOF acts as a one-layer neural network, and works that have compared EOF to neural networks [297, 97] find that they converge to the EOF solution for linear problems. For a review of the examples of implementation of predictive control, organized by method, see Appendix A.1.

Are non-linear methods justified

Despite pointing to non-linearities as a potential goal for new wavefront control algorithms, some interesting work in the field has explored whether or not temporally evolving non-linear controllers are necessary. Multiple works [97, 297] have compared classic linear solvers (i.e., empirical orthogonal functions or a neural net with a single layer) to multi-layer approaches, only to find that when considering temporally evolving statistics the multi-layer systems converge to the same answers, sometimes with added noise in the multi-layer

networks.

In Ref. [99], they argue that running in closed-loop maintains the condition of linearity for both Pyramid wavefront sensors (PWFS) and DM actuators, meaning non-linear reconstructions and system dynamics are not worth encapsulating in temporally evolving models. However, other works [192, 146] demonstrate that a non-linear solver can learn the entire system model, including wavefront sensing response, deformable mirror dynamics, and DM to wavefront sensor misregistration, which leads to a more painless and flexible system implementation. Furthermore, Ref. [297] found that despite converging to the same solution, the neural network found the answer with ∼5 seconds of training data, as compared to a linear predictor which needed ∼40 seconds of training data.

1.6 Summary

Finding and characterizing an exo-Earth is not within the current capability of ground or space-based telescopes. My thesis uses 8-10 meter class telescopes as a laboratory for future exo-Earth characterization. I present instrumentation work that bridges the technology gap between current telescopes and our future exo-Earth finders, while understanding and improving our current telescopes. Additionally I describe observational work that seeks to push forward future

techniques we will need to find exo-Earths, while pursuing real physics-driven questions about larger and brighter planets that are currently in reach of 8-10 meter class observatories. I explore high contrast imaging and extreme adaptive optics from snout to tail in four parts by (1) exploring the science we can do with current ground based telescopes with a focus on polarimetric detections of giant planets and brown dwarfs, (2) characterizing the performance of current ground based AO-fed instruments in the visible, (3) commissioning a data-driven predictive wavefront controller on-sky at Keck Observatory, and (4) comparing predictive methods in preparation for the next generation of extremely large telescopes.

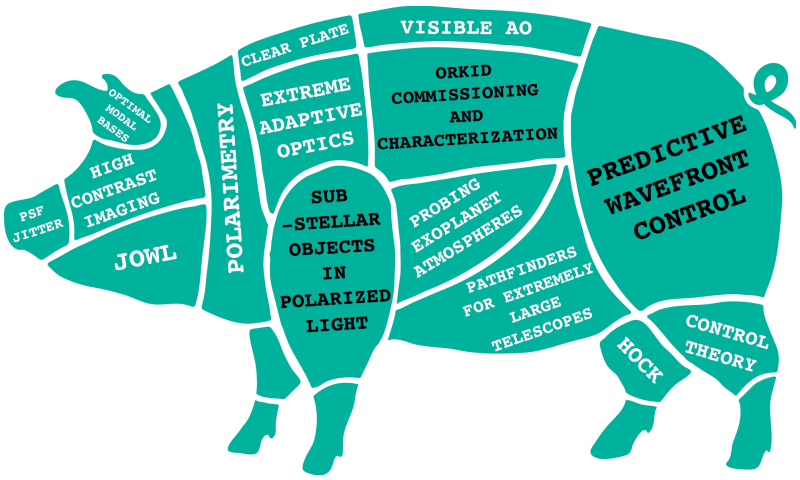


Figure 1.23: Every part of the PIG for high contrast imaging and extreme adaptive optics: Predictive wavefront control, commissioning and characterizing the ORKID Instrument, and Ground-based polarimetry of substellar companions.

Chapter 2

When PIGSS Fly: the Polarimetric InvestiGation of Self-luminous Substellar companions

2.1 Abstract

The convective physics that produces large scale cloud features (e.g., Jupiter's bands and spots) while thought to depend on temperature, is still largely unknown. In this chapter I present two new half nights of Subaru/SCExAO

spectropolarimetric CHARIS observations of two substellar companions (AF Lep b and HIP 21152 B) in attempts to build a sample of polarimetric observations of substellar companions at a range of temperatures. We do not reach our intended upper limits at the level of 0.5 - 0.1% disk-integrated degree of linear polarization on these objects, which would strongly rule out large scale cloud structures. We build total intensity images from these data, but do not reach sufficient contrast to detect either companion at the 5σ level. We note that our polarimetric upper limits were limited by performance of the adaptive optics system, not polarimetric precision. Finally, we comment on polarimetric surveys as a low risk/high reward strategy for current ground-based telescopes and future extremely large and space-based telescopes.

2.2 Introduction

Large scale cloud features (for example, Jupiter's bands and spots) are theorized to be driven by convective regions deep in planetary atmospheres and only possible within certain temperature regimes [307]. However, with only a handful of examples of these cloud features in our solar system, it is challenging to confirm this relationship. To fully investigate these processes and confirm a temperature dependence requires more examples of substellar companions confirmed to have large scale cloud features. Exoplanets and brown dwarfs

have a wealth of indirect observational evidence confirming the presence of clouds [253], but directly constraining cloud morphology is challenging outside of our solar system. Doppler mapping shows promise to resolve large scale cloud features, e.g. maps of the brown dwarf binary Luhman 16 [44, 126, 37], but these techniques are currently limited to bright (K$\leq$10 magnitude) targets. Similarly, photometric monitoring has found 1-10% variability in the brightness of substellar companions that implies the existence of asymmetric cloud features [166, 18], but this technique is limited to bright objects with long periods (i.e., objects that can be detected in enough time to resolve multiple data points in time along an object's rotation). Future techniques will use kilometer baseline array interferometry to spatially resolve cloud features on substellar objects [23], but we do not currently have the technology to resolve a substellar object at more than 2 pixels across. As a solution, we present direct imaging of self-luminous companions in polarized light. The disk-integrated self-scattering signatures required for a polarimetric detection of a substellar companion are indicative of atmospheric asymmetry, and one of the few ways to cause this asymmetry is large scale cloud features like bands and spots.

Young exoplanets and brown dwarfs have self-luminous emission that peaks in the infrared. As that emission moves through clouds it is scattered by aerosols (driven by Rayleigh scattering in the visible and Mie scattering in

the IR), which polarizes the object's self-luminous emission [50]. However the positive and negative polarized signal from a symmetrical source is perfectly canceled when integrated over a planet, i.e., a disk-integrated polarimetric observation is representative only of asymmetries in the polarized signature. These asymmetries are caused by non-uniform cloud distributions and hot-spots, oblateness of the object, and circum-planetary disks [264]. Studies of free floating planetary mass companions and brown dwarfs have detected these signals at levels of 0.4- 1.5% [175, 176], but without a host star it is difficult to estimate the age and mass of an object with accuracy, making these polarized signals more difficult to interpret physically. Two studies have detected the polarized signature of a brown dwarf, finding a signal of 0.031 $\pm$ 0.004 and 0.010 $\pm$ 0.004 % signal in fractional polarized intensity from Luhman 16 A and B respectively (due to cloud vortex features [177, 183]) and a 0.48 $\pm$ 0.03 % signal from DH Tau B (due to a circum-planetary disk [276]). This survey intended to grow this population of directly imaged point-sources detected in polarized light.

2.3 Target selection

Our study targeted 3 companions: AF Lep b, HIP 21152 B, and HD 984 B, which were chosen because they (1) span a range of mass, temperature, and

L/T space, (2) are all bright enough that 0.5-0.1 % of their signal could be resolvable in ∼6 hours of observing time (a half night), and (3) were within the field-of-view of the CHARIS spectropolarimetry mode (< 1") in the northern sky. All three objects were selected to be high fidelity detections in total intensity in a half night of observing time (for example, HIP 21152 B had been observed by SCExAO/CHARIS before in ∼2 hours [140]), so that even without informative upper limits these data would still provide a detection in total intensity, and therefore updated astrometry and spectroscopy. The target upper limits of 0.5-0.1% in fractional polarized intensity are driven by detailed models of AF Lep b [79] and model grids [34], that project the presence of asymmetric clouds and hot spots or oblate atmospheres would create a polarized fraction of $\geq$ 0.5 or 0.1 % respectively. These three objects are summarized in Table 2.1 and Figure 2.1 shows how the three objects compare in mass, age, temperature, and L/T space. AF Lep b was particularly timely; due to its edge-on inclination, 2024 was one of the last opportunities to observe it before it got too close to its host star for conventional coronagraphy. Due to poor weather we did not collect data on HD 984 B. The rest of this chapter will focus on data on the two companions HIP 21152 B and AF Lep b.

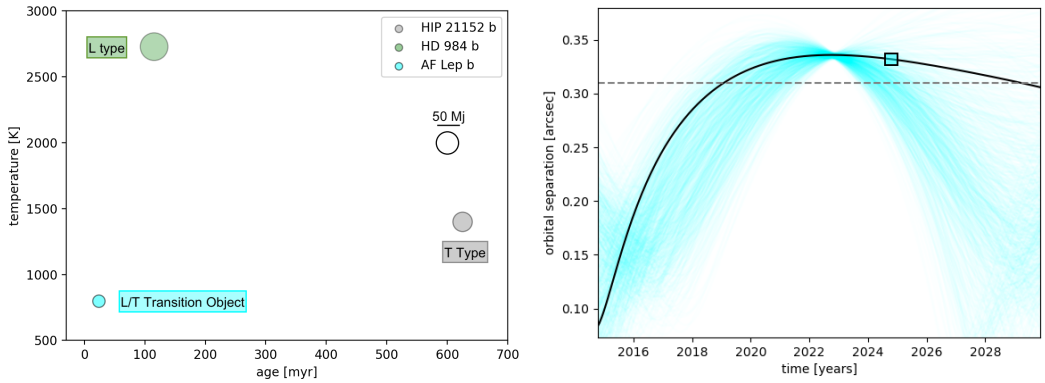


Figure 2.1: Left: Our three proposed targets span the temperature, age, mass, and L/T space. HD 984 B is a large, hot, L dwarf. HIP 21152 B is a cooler, smaller T dwarf. AF Lep b is a very small and very cool L/T Transition object. Right: AF Lep b's estimated orbital parameters. Based on data from Ref. [308] we estimate the orbital solution of AF Lep b. Due to its highly edge-on inclination, it is quickly approaching its host star. While it may remain visible until 2028, with the existing data there is a density of orbital solutions that predict AF Lep b is no longer observable with conventional coronagraphy (where a focal plane mask $\geq$ 300 mas).

Target	Contrast (Δ L' mag)	Companion (L' mag)	Mass (M_J)	Spectral Type	Age (Myr)
HD 984 B	6.0 $\pm$ 0.02	12.0 $\pm$ 0.2	61 $\pm$ 4	L Dwarf	115 $\pm$ 85
HIP 21152 B	9.31 $\pm$ 0.09	14.66 $\pm$ 0.11	24 $\pm$ 6	T Dwarf	650 $\pm$ 100
AF Lep b	12.81 $\pm$ 0.26	14.95 $\pm$ 0.26	3.75 $\pm$ 0.5	L/T Transition	24 $\pm$ 3

Table 2.1: Properties of each system in our intended survey. Information is compiled from existing imaging data in Refs. [174, 73, 173].

Target	Exp Time (seconds)	HWP Cycles	Integration (minutes)	Wall Clock (UTC)	Date (UTC)
HIP 21152 B	60.5	46	185.53	11:06-15:57	10/13/2024
AF Lep b	60.5	19	76.63	07:16-09:49	01/16/2025

Table 2.2: Observing data on the two companions HIP 21152 B and AF Lep b. Here we include only the frames usable in the final analysis. Due to poor data quality and inconsistencies in data collection (including frames that were mostly NaN values and frames taken with varying satellite spots that could not be processed with other data) we were unable to reduce and process all of the data from the 01/16/2025 night on AF Lep b.

2.3.1 HIP 21152 B

HIP 21152 B is a hot (1300 K) late (650 Myr) T dwarf, with a dynamical mass of 24±6 M_J [73]. HIP 21152 B has a magnitude of L'= 14.66 ± 0.11 at a separation of 370 mas [73] and a measured vsini of 12.6 km/s [112], which combined with its large size and projected surface gravity [73] make it unlikely to be oblate. HIP 21152 B has been the subject of recent modeling efforts which have been unable to determine the difference between cloudless, grey, and cloud types of varying opacity [307]. An observation of polarized emission would be evidence of a cloudy atmosphere, as this object is too old for a disk and there is no spectral evidence for a disk; the clouds would have to be asymmetric because HIP 21152 B's oblateness would not cause a detectable amount of polarization for an 8-10 meter telescope (e.g., SCExAO).

2.3.2 AF Lep b

AF Lep b is a young (24 Myr), cool (∼800 K) giant planet with a dynamical mass of 3.75± 0.5 M_J [173, 74]. Recent photometry estimates its magnitude at L'=14.95 ± 0.26 at a separation of 335 mas [173]. Recent modeling efforts have predicted AF Lep b has Silicate clouds [308], but are unable to place constraints on the Silicate species (i.e., the grain size of the clouds akin to Ref. [183]) or comment on cloud morphologies with existing spectra. AF Lep b is an

L/T transition object [74], and previous data point to extreme variability [308]; both of these characteristics make AF Lep b likely to have asymmetric cloud features that would lead to a net disk-integrated polarization. Furthermore, due to AF Lep b's small mass and low surface gravity, it is likely to be very oblate (more than twice the oblateness of any object in previous studies like that of Ref. [120]); as oblateness is a strong function of rotational velocity, we also planned to use a detection or upper limit in polarized light to constrain AF Lep b's velocity.

2.4 Subaru/SCExAO observations

2.4.1 Observing and adaptive optics conditions

We received three half nights of Subaru/SCExAO time, each timed for the visibility of a single companion. The first night was completely weathered out due to Hurricane Hone; flood conditions made the summit road impassable and we collected no data on our first intended companion HD 984 B. We collected data as planned for the following nights 10/13/2024 UTC (on HIP 21152 B) and 01/16/2025 UTC (AF Lep b).

At the time of our observations, Subaru was in the process of commissioning their AO-3k upgrade. For the first (HIP 21152 B) night, Subaru was corrected

by a first-stage facility adaptive optics (AO) system, AO-188, which included a curvature wavefront sensor [291] and a CILAS bimorph deformable mirror with 188 degrees of correctional freedom [198]. For the the second (AF Lep b) night, Subaru was running with the 188-subaperture wavefront sensor powering an ALPAO deformable mirror with 3000 degrees of freedom [154], while using influence functions that mimic the previous 188 element DM (i.e., only correcting 188 degrees of freedom). SCExAO further includes a second stage AO system with a visible light pyramid wavefront senor and Boston Micromachines micro-electro-mechanical systems (MEMS) deformable mirror with 1000 degrees of freedom [2]. Anecdotes from AO operators suggest that despite median seeing conditions, the state of AO correction was not optimal for either night. Approximately half of the night on AF Lep was lost to AO calibration and startup issues, likely related to AO-3k commissioning. The total wall clock time of observations and integration time after rejecting bad frames on each target is described in Table 2.2.

2.4.2 CHARIS spectropolarimetric imaging

We took integral field spectrograph (IFS) data of the substellar companions with the Coronagraphic High Angular Resolution Imaging Spectrograph (CHARIS [87]) instrument in the low resolution mode, which spans $\lambda = 1150 - 2370$ nm

(J-H-K) with a resolution $R \sim 30$. Due to a lack of detection we do not present spectra, but we do use these data for spectral differential imaging (SDI [259]). Data were taken allowing the sky to rotate, inducing angular diversity for angular differential imaging (ADI [167]). Finally, we used the dual beam polarization mode to take spectropolarimetric information, allowing us to (1) seek potential polarized signatures from our companions and (2) further use polarimetric differential imaging (PDI [138]) to improve contrast.

2.4.3 Polarimetric differential imaging

Polarimetric imaging aims to recover light in linear and circularly polarized states. This is often described by the stokes vector

$$\mathbf{S} = \begin{bmatrix} I \\ Q \\ U \\ V \end{bmatrix} \tag{2.1}$$

where I represents total intensity, Q and U represent the linear polarized intensity along two directions, and V represents circular polarized intensity. Most high contrast imaging polarimeters target linear (Q and U) polarization and do not have sufficient instrumentation (e.g., a less than half wave plate) to

measure circular polarization (V), which some studies suggest may be valuable for future studies of grain sizes on Earth-like planets [184].

Polarimetric Differential Imaging (PDI) splits light into orthogonal polarization states. In particular, Subaru/SCExAO uses a Wollaston prism and a modulating half-wave plate (HWP). The Wollaston prism splits light in two, creating left and right side records of the target at linear orthogonal polarization states [87]. A half wave plate modulates the signal, allowing for the reconstruction of Q and U polarization states. In practice a full HWP cycle rotates through four angles: 0, 45, 22.5, 67.5°. Information on the first linear polarization state (Q) is contained in the HWP rotated to 0°: 0-90° state and information on the second linear polarization state (U) is contained in the HWP rotated to 22.5° : 45-135° state. Using the additional states creates an inverse image (90-0 and 135-45 respectively), also known as double differencing, which allows for subtraction of errors between the HWP and Wollaston prism. Additionally, because PDI uses closely associated states in time to build a full HWP cycle, it is one of the most effective techniques to remove speckle noise. Finally, because the Wollaston prism evenly splits incoming light, the total intensity can be recovered by adding the left and right hand sides of each image together.

Double differencing builds I, Q, and U stokes vectors with total intensity:

$$I = I_N + I_{N+90} \tag{2.2}$$

the first linear polarization state:

$$\begin{aligned} Q &= \frac{1}{2}\left([Q - \mathcal{E}] - [-Q + \mathcal{E}]\right) && (2.3)\\ Q &= \frac{1}{2}\left([I_0 - I_{90}] - [I_{90} - I_0]\right) && (2.4)\\ & && (2.5) \end{aligned}$$

and the second linear polarization state:

$$\begin{aligned} U &= \frac{1}{2}\left([U + \mathcal{E}] - [-U + \mathcal{E}]\right) && (2.6)\\ U &= \frac{1}{2}\left([I_{45} - I_{135}] - [I_{135} - I_{45}]\right) && (2.7) \end{aligned}$$

where $\mathcal{E}$ represents polarimetric errors that are removed through double differencing, and I_N denotes the image created via HWP at the rotation N. Previous studies have found CHARIS spectropolarimetric differential imaging to have an average polarimetric precision of $\leq 0.1\%$, as described in Ref.[277].

2.5 Data reduction, contrast curves, and upper limits

2.5.1 Data reduction and contrast curves

Our data consist of CHARIS low-resolution integral field spectrograph (IFS) data taken in the previously described sets of half-wave plate angles. Data reduction consists of three major steps: (1) reduction of IFS traces into multi-spectral image cubes, (2) reduction of multi-spectral image cubes into stokes Q, U, and total fractional polarized intensity ($I_P = \sqrt{Q^2 + U^2}$) images, and (3) conversion of Wollaston split images into total intensity images and ADI and SDI analysis on these images. We used the CHARIS pipeline [26, 239] to reduce the IFS traces into multi-spectral image cubes. We used the CHARIS Data Processing Pipeline (DPP) [148] to do sky subtraction, image registration, spectrophotometric calibration, half wave plate angle matching, polarimetric correction, and finally to estimate the polarized intensity, and Q and U stokes vectors. Additionally, using intermediary products we build an effective total intensity (or I) image by adding each half of the images split by the Wollaston prism. Figure 2.2 shows examples of less processed data of the target HIP 21152 B, including the raw CHARIS IFS data product before any wavelength calibration is applied and a pre-PDI image with the Wollaston prism split.

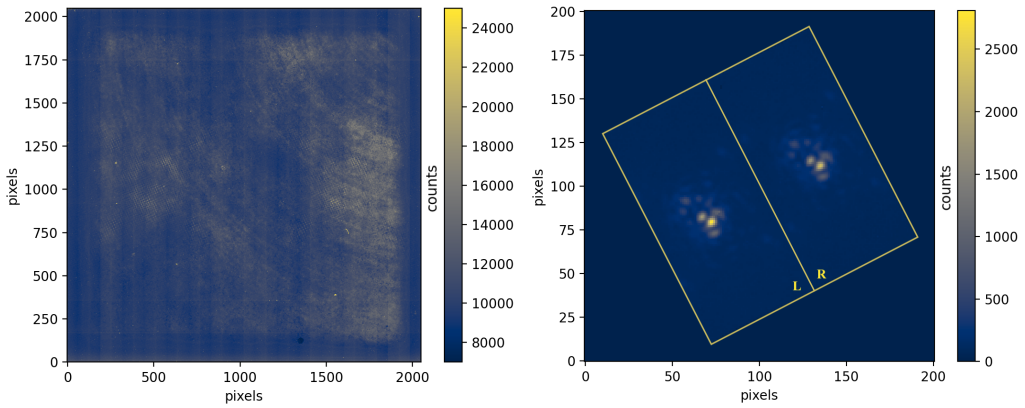


Figure 2.2: Examples of HIP 21152 B data at varying stages of the data processing pipelines. Left: Raw CHARIS data frame before any processing is applied. This contains a spectra for every pixel, hence the higher dimensionality than the final images. Right: CHARIS data slice through a single wavelength after CHARIS pipeline (the wavelength calibration has been applied), but before the CHARIS DPP (any PDI data processing has taken place). This contains the left and right sides of the image split by the Wollaston, as outlined by the yellow detector foot print and the L/R annotation.

The Q, U, and polarized intensity pipeline products have already been de-rotated and had instrumental polarization correction via Mueller matrix applied [209]. Because this correction is pixel/detector orientation dependent, doing further SDI/ADI analysis is non-trivial, and is left for future work. We show the polarized intensity image and the total intensity image of the bright brown dwarf companion HIP 21152 B and AF Lep b in Figures 2.3 and 2.4 respectively. We do not detect either companion in polarized light.

We estimate the contrast for polarized intensity using `pyklip.klip. meas_contrast` function with an aperture set by the CHARIS low resolution IFS on the wave-

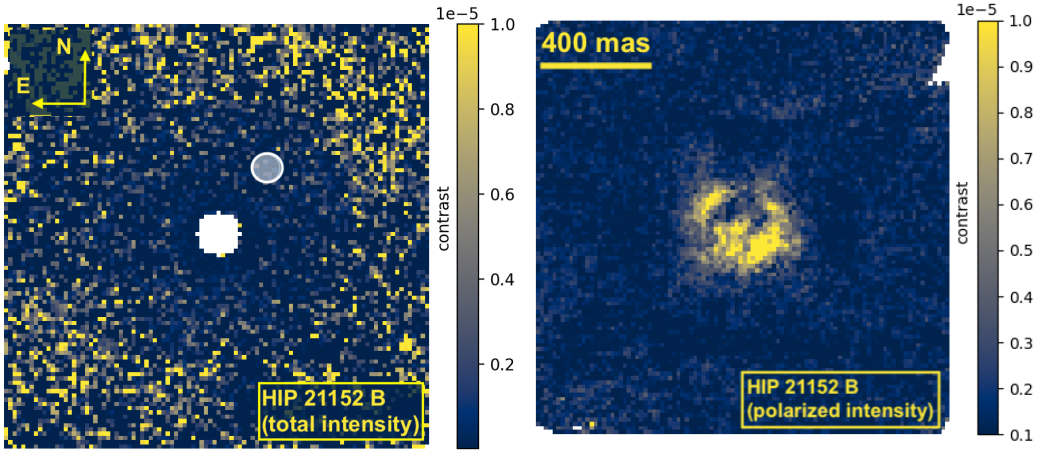


Figure 2.3: HIP 21152 B. Left: Total intensity image after post-processing with `pyklip`. The estimated companion position [289] is marked in white. Right: Polarized intensity image. We do not detect the companion in total or polarized intensity. Cardinal directions and spatial scale are shared between the total and polarized intensity images.

length averaged polarized intensity image output of the CHARIS DPP pipeline to calculate the 5-σ contrast curve, as shown in Figure 2.6. The green marker indicates a generous upper limit of companion contrast from L band data on the two companions. Future work would compare averaged contrast estimated from model spectra of HIP 21152 B and AF Lep b averaged through J/H/K, for a more accurate comparison to CHARIS J/H/K wavelength collapsed cubes. We do not reach 0.5% of the L band (upper limit) in polarized intensity. We note that for AF Lep b, the polarized contrast should not be worse than the contrast in total intensity. This is likely due to an unknown error in the PDI reduction (possibly due to missing data during HWP cycles), and future work would

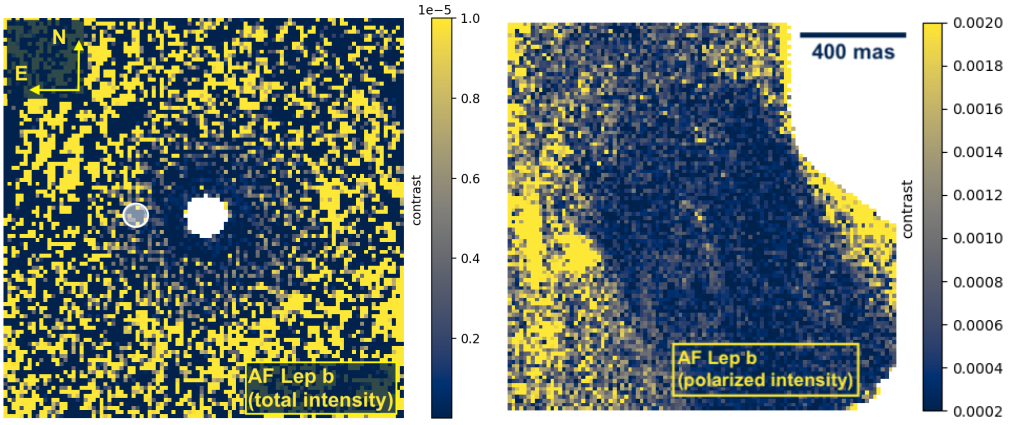


Figure 2.4: AF Lep b. Left: Total intensity image after post-processing with `pyklip`. The estimated companion position [289] is marked in white. Right: Polarized intensity image. We do not detect the companion in standard or polarized intensity. Note that missing data in the top right hand corner of the polarized image comes from the shorter observing sequence. When the Wollaston splits the detector frame images are long rectangles, which are then combined in rotation. Without $\sim$ half of the planned data there is not enough to rotate through the entire sky and cover the standard detector footprint. Cardinal directions and spatial scale are shared between the total and polarized intensity images.

explore if better reduction can improve the contrast in polarized intensity.

After rebuilding the total intensity image products, we perform ADI and SDI with `pyklip`'s built in ADI and SDI `klip` tool [290]. After comparing the contrast before algorithmic throughput correction of data reduced while removing between 0 and 150 modes via principal component analysis (PCA), we select 100 modes as our optimal contrast curves. We note that optimizing for contrast before algorithmic throughput correction may have optimized contrast without preserving planet signal. Future work would explore reduction

strategies using fewer modes to ensure planet signal is not being spuriously subtracted. However, a comparison of removing 50 and 100 modes shows similar loss due to algorithmic throughput, while the 100 modes showed better contrast, as shown in Figure 2.5.

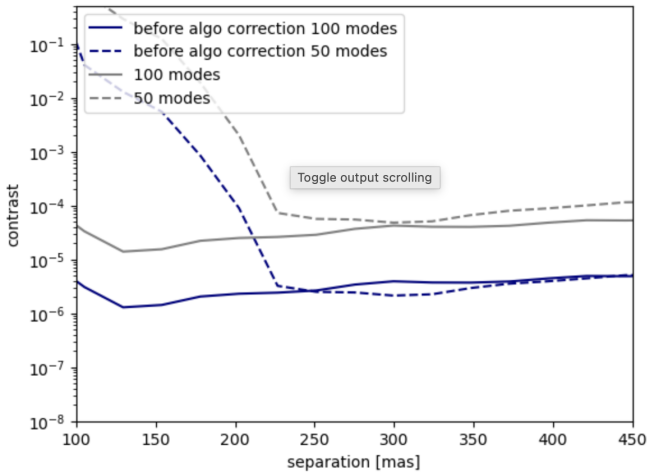


Figure 2.5: Comparing the impact of of removing 50 and 100 modes during ADI and SDI on the algorithmic throughput correction for data of the target HIP 21152 B. The dashed lines indicate 50 modes, while the solid lines indicate 100. Navy curves are contrast before algorithmic throughput correction is applied, and grey indicates the final algorithmic throughput correction we provide in Figure 2.6. Note that the difference before and after algorithmic throughput correction is similar between the 50 and 100-mode removed curves, but the contrast improves after removing 100 modes, especially at close inner working angles.

The final wavelength-averaged images of each companion are shown in Figures 2.3 (of HIP 21152 B) and 2.4 (of AF Lep b). We do not detect the companion in this wavelength-stacked image or any individual wavelength slice. We then inject fake companions into our data to estimate the impact

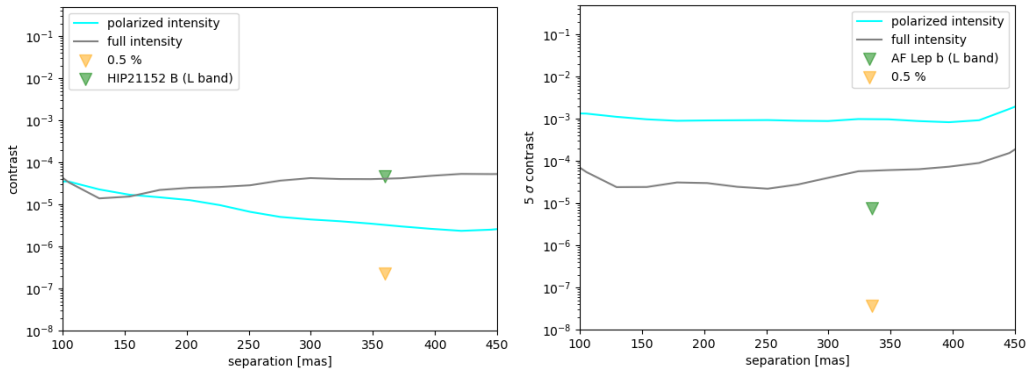


Figure 2.6: 5-σ contrast curves for the two companions. Left: HIP 21152 B. Right: AF Lep b. After accounting for algorithmic throughput, it is clear that we do not reach a sufficient contrast to detect the either companion. The green markers are L band contrast, which is a generous upper limit of contrast, which would be much fainter than the average of J/H/K in these data. Using L band contrast as a rough starting estimate we also include a 0.5 % upper limit, indicated by the orange marker. Despite seeing better contrast with the polarized intensity for HIP 21152 B, we do not reach the target upper limit to successfully rule out asymmetric cloud features. We note that the poor PDI contrast (especially in comparison to total intensity) for AF Lep b is unexpected, and this points to potential issues in the data reduction.

of algorithmic throughput loss. This algorithmically corrected contrast curve shows that we are not sufficiently able to resolve the contrast needed to detect either companion, as shown in Figure 2.6.

2.6 Conclusions and future work

In conclusion, we present two half nights of polarimetric data on the substellar companions HIP 21152 B and AF Lep b. At the time of this analysis, we do not reach sufficient contrast in total intensity to detect either companion.

We do not reach our target upper limits of 0.5 and 0.1 % fractional polarized intensity on either companion. Future work would include:

- Exploration of the polarimetric differential imaging data reduction, to better understand the limited contrast of the AF Lep data after PDI.
- Better optimization of the angular differential imaging and spectral differential imaging steps, to prioritize companion signal over contrast, which may result in removing fewer modes.
- Using a better proxy for companion contrast than existing L band detections, by using models of objects at similar spectral types to estimate the flux through J/H/K bands and average over the broadband CHARIS filter.
- Comparing other ADI and SDI reduction methods like Vortex Image Processing (`VIP`, [85, 38]) or Signal Safe Speckle Subtraction (`4S`, [20])

While we did not successfully complete the goals of this study, we note that the limiting factor was issues with adaptive optics performance; previous studies with CHARIS under better conditions have resolved HIP 21152 B in $\sim$ 2 hour of observing. Despite sub-optimal AO performance, we came close to resolving both companions in the reconstructed full intensity image cubes, showing that polarimetric imaging can be a low risk/high reward strategy for direct imaging

of exoplanets. If there is any chance that the companion may have a polarized signature (either self-scattering or in reflected light), a polarimetric detection will provide up to a magnitude deeper contrast than other differential imaging methods, and the original data products allow photometry, spectroscopy, and astrometry from the reconstructed I images. This study also motivates future work in extreme adaptive optics systems, specifically stability and repeatability, to ensure the efficacy of follow up observations on known targets.

Chapter 3

Don't Heckle my (visible light) Speckle: Characterizing Keck-II adaptive optics performance in visible light with the ORKID instrument

3.1 Abstract

This chapter uses the ORCAS Keck Instrument Demonstrator (ORKID), a fast (up to 500 Hz) visible light camera on the Keck-II AO bench, as a tool to analyze current and predict future Keck AO performance at visible light wavelengths. We perform the first comparison between empirical data and analytic predictions of Keck-II AO in visible light. We report a diffraction limited core of 13.5±0.4 mas at 656nm, and a Strehl Ratio of 2.4% on a bright target for an image with residual tip/tilt errors removed (via post-facto image median). We find a point spread function (PSF) jitter of 3.2 and 4.4 ± 0.7 mas in two axes respectively over sub-second timescales, which is comparable to the magnitude of analytic predictions, but shows consistent asymmetry between x and y that diverges from analytic estimates. We note a frequent split in the ORKID PSF, likely due to optical coma that was not calibrated out of the ORKID path. We find recurring tip/tilt vibrations at 20 Hz and 300 Hz and explore their impact on current and future performance. We identify a PSF drift on the order of 40 mas over longer (∼ hour timescales). We discuss the limited ORKID sky coverage and wavefront sensing strategies. Finally, we note that current ORKID performance is limited by the low degree of freedom (DOF) deformable mirror on the Keck-II AO bench at the time of these observations, and predict that with the high DOF deformable mirror

upgrade for Keck-II AO (which recently completed commissioning) ORKID will provide Strehl Ratios of 6-7% at 656nm.

3.2 Introduction

As discussed in Chapter 1.3.6, high contrast visible light science fed by adaptive optics system represents an unsolved technical challenge. Despite the complexities, AO-fed visible light instruments (especially for the 6-10 meter class telescopes) provide higher angular resolution than IR instruments; this technology is vital for resolving closely associated binaries, circumstellar disk imaging, searching for forming protoplanets, and other galactic and extra-galactic science cases, as described in the Keck Observatory Strategic Plan. Furthermore, high angular resolution visible light imaging will be key for confirming a habitable planet around nearby M-stars with future extremely large telescopes [260].

At the W. M. Keck Observatory, the largest facility AO-fed ground based observatory (with two segmented 10 meter telescopes), the current generation of AO has been developed exclusively for the near IR [292]. However, with the introduction of the visible light ORCAS Keck Instrument Demonstrator camera [208, 178] (ORKID) fed by the Keck-II AO bench, we have an opportunity to evaluate the Keck-II AO performance (and that of the twin Keck-I AO system) from the perspective of a visible light instrument and the highest

angular resolution diffraction limited PSF on-sky. By comparing ORKID data to analytic models and the AO error budget for Keck-II [27, 121] we perform a health and wellness check on Keck-II AO. This is important precursor work for the characterization of the High order Advanced Keck Adaptive optics (HAKA) system, a deformable mirror upgrade to Keck-II AO that increases degrees of freedom in correction to 2844 actuators, projected to deliver higher Strehl Ratios by a factor of up to 6, in visible light [293]. In this chapter we analyze a night of ORKID engineering data, with experiments designed to match ORKID data to analytic models and the Keck-II AO error budget. After finding good agreement with these models we use them to comment on the performance of future visible light astronomy at Keck, with a focus on the HAKA upgrade.

In this section we have motivated and contextualized visible light adaptive optics. In Section 3.3 we describe the ORKID instrument. In Section 3.4 we outline the observational engineering data we collected for these experiments. In Section 3.5 we present our estimated performance metrics for ORKID, including Strehl, Encircled Energy, Point Spread Function (PSF) jitter, and how these compare to analytic predictions based on the Keck-II AO error budget. In Section 3.6 we explore current factors impacting ORKID performance, including poor wavefront sensor performance on faint guide stars, recurring vibration events, and PSF splitting. In Section 3.7 we compare ORKID PSF jitter to

Keck-II AO telemetry, with a focus on tip and tilt residuals, and identify an asymmetry between x and y motion, as well as a large drift over long timescales unique to the ORKID PSF. Finally in Section 3.8 we make final conclusions, look forward to HAKA performance, and comment on the future of AO-fed visible light imaging for Keck Observatory and the broader visible-light science community.

3.3 ORCAS Keck Instrument Demonstrator

3.3.1 ORKID instrument

The ORCAS Keck Instrument Demonstrator (ORKID) was originally commissioned to demonstrate AO-fed visible light capabilities for the Orbiting Configurable Artificial Star (ORCAS) project [207]. The goal of ORKID is to unlock visible wavelength science via a bright space-based laser beacon that expands its sky coverage. ORKID is a fast (subarray readout at $\sim$ 500Hz) Andor Marana CCD camera, Nyquist sampled at 650 nm, with a plate scale of 6.7 mas/pix. ORKID is fed by the Keck-II adaptive optics system, making it the only AO-fed visible light instrument on Keck-II. It has broadband filters covering Bessell R, Sloan i and z, as well as narrowband filters for OI (620-640nm), Hα and Hα continuum (651-661nm, 634-646nm), Jovian Methane

(871-907nm), and CT (920-960nm). For more details on the ORKID instrument, its commissioning, and first light results see Ref. [208, 178].

3.3.2 Keck-II adaptive optics system

Keck-II AO includes a tip/tilt correction mirror, a visible light Shack-Hartmann wavefront sensor (SHWFS), and, at the time of these observations, a high order Xinetics DM with 349 actuators, each mapped to one SHWFS subaperture [27]. In natural guide star mode, the tip/tilt slopes are estimated from the SHWFS. Before 2025 (and at the time of the data discussed in this chapter), the wavefront sensor detector was a CCD39 from e2v. The current wavefront sensor detector for Keck-II is an OCAM2K EMCCD from First Light Inc./Oxford Instruments. In this work we will discuss the CCD39 performance and extrapolate forward to the OCAM2K performance where relevant. The SHWFS typically uses most of the visible light signal we would ideally send to the ORKID detector, so for ORKID nights we operate in a non-standard wavefront sensing mode, using the sodium dichroic designed to transmit the 580-595 nm band including the 589 nm sodium-wavelength light generated by a laser guide star. All remaining light, including light below this band, is reflected, e.g., sent to the ORKID camera. ORKID can also run with standard R SHWFS bandpass, which due to its anti-reflective coating transmits 97% of light to the wavefront sensor and

reflects 3% to the science camera. Practically the use of R band SHWFS with ORKID requires longer science exposures for sufficient signal to the science instrument, but does improve AO performance for fainter guide stars. These two wavefront sensing modes are diagrammed in Figure 3.1.

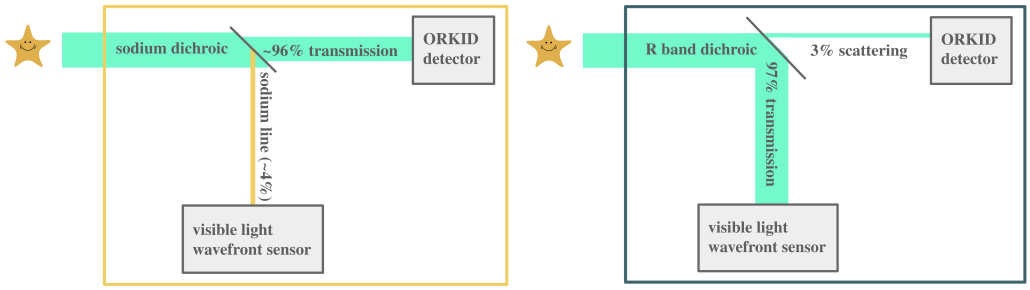


Figure 3.1: The two natural guide star wavefront sensing modes for ORKID. Left: The sodium dichroic transmits light through the sodium line between 580-595 nm (∼4% for a flat bandpass) to the visible light Shack-Hartmann wavefront sensor. All remaining light between 400-900 nm is sent to the ORKID instrument. This is the wavefront sensing mode used in this work. Right: The standard R filter transmits 97% of light between 400-900 nm to the wavefront sensor. Due to an anti-reflective coating, 3% of light scatters to the ORKID instrument.

3.4 ORKID observations

Target	V Magnitude	R Magnitude	Exposure Time (ms)	N Exposures	Total Wall Clock (UTC)	AO Rate (kHz)	Use
Alpha And	2.06	2.09	1	2155	08:58 - 09:07	1.0	4†
			1	400001	13:10 - 13:41	1.0	
BS 458	4.1	3.64	10	12500	10:32 - 10:35	1.0	3
			5	12501	12:18 - 12:22	1.0	
BS 603	2.1	1.16	10	12502	10:24 - 10:28	1.0	3
			1	62501	12:11 - 12:15	1.0	
FS 104	11.06	> 10.5	1000	126	13:03 - 13:06	0.15	3
HD 984	7.32	> 6.4	50	50891	09:18 - 10:13	1.0	1, 2, 4
HR 249	6.06	>4.4	5	392102	11:02 - 12:03	1.0	1, 2
HR 476	5.63	> 5.2	10	12500	10:39 - 10:42	1.0	3
			20	7505	12:31 - 12:38	1.0, 0.5	
HR 876	6.03	> 3.8	5	362504	13:46-14:55	1.0	1, 2
SAO 038218	9.04	> 8.7	500	250	10:53 - 10:57	1.0, 0.5	3
			500	501	12:53 - 12:58	0.5	
SAO54443	7.78	> 6.6	10	12500	10:46 - 10:48	1.0	3
			100	2502	12:44 - 12:49	0.5	

Table 3.1: Targets from engineering night of 09/19/2024:08:25:00 UTC. † Number indicates experiment from the list of experiments in Section 3.4.

ORKID observations were performed during the second half of the September 18th, 2024 HST engineering night, starting at 09/19/2024:08:25:00 UTC. We designed four experiments: (1) PSF characterization and stability over ∼hour long timescales, (2) impact of high and low order control loops on tip/tilt jitter and AO performance in general, (3) impact of faint guide star magnitudes on AO performance, and (4) a demonstration of the system angular resolution. For most of the night we have accompanying telemetry, specifically lenslet spot movements (or slopes) from the Shack-Hartmann wavefront sensor, and commands sent to the tip/tilt mirror and the high order deformable mirror. Table 3.1 lists the targets we observed and their use within these experiments. For the AO bandpass we used the sodium dichroic in natural guide star mode, to maximize ORKID camera performance and short science exposures, and most of these data were taken with a 1kHz control loop. For all of these experiments, ORKID exposures were as short as possible while preserving signal in each exposure (set interactively for each target). Targets were observed very close to zenith and taken in the narrowband filter $H\alpha, \lambda$ =656nm, to control for PSF-smearing due to atmospheric dispersion. Over the course of the night, seeing conditions were standard for Maunakea, on average sub-arcsecond seeing ranging from 0.4 to 1.2" according to the Canada-France-Hawaii Telescope Differential Image Motion Monitor, as shown in Figure 3.2.

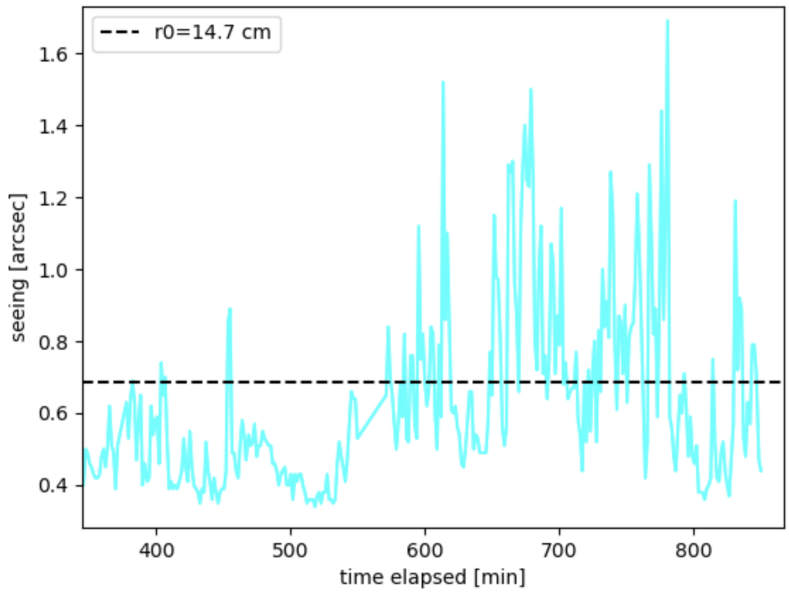


Figure 3.2: Seing during the ORKID night, according to the CFHT DIMM seeing monitor. Dashed line indicates the seeing used by the semi-analytic Keck AO error budget used later to fit ORKID performance.

3.4.1 Point spread function sharpening for ORKID during daytime testing

The ORKID instrument sits at a different optical path than that of other Keck-II instruments (e.g., NIRC2), and therefore experiences a different set of non common path aberrations (error induced by varying paths between a wavefront sensor and focal plane imager, NCPA). Before our engineering night of September 18, 2024, we performed daytime image sharpening by looking for peak PSF flux over a scan of zernike modes [13] applied to the DM. We calculated an offset to the SHWFS slopes of 79nm RMS, as shown in Figure

3.3 and included a defocus offset via motorized stage to the ORKID camera of -0.5mm. Figure 3.3 shows the PSF before and after image sharpening was applied.

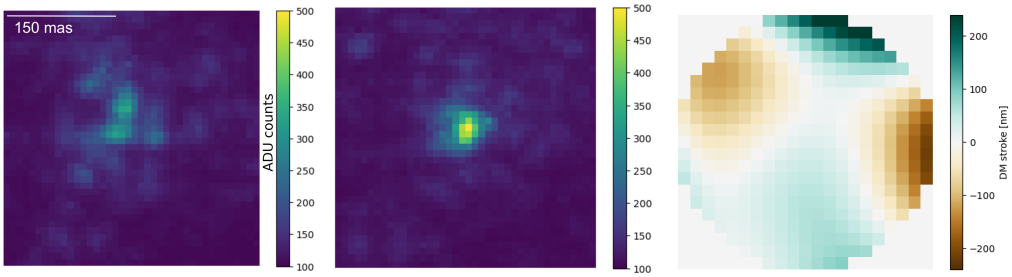


Figure 3.3: ORKID PSF before (left) and after (center) daytime PSF sharpening experiments. Most of the offset was included as a defocus movement 0.5 mm to the ORKID camera. The rest was applied as SHWFS slope offsets, a correction of 79nm RMS (left).

3.4.2 State of the segmented primary after a segment exchange

The first half of the 9/19/2024 night was a segment exchange and phasing night. Keck segments are periodically swapped; during the daytime proceeding this night three telescope segments were exchanged and the telescope was rephased during the first half of the night. Following rephasing, segments are warped (to manually change the shape of each segment) but the segments are not able to be warped until the following day. This means that the state of the segmented primary during these observations was at a recently updated phasing solution

but with three individual segments that have uncorrected optical aberrations.

Data taken from the Phasing Camera System Super Ultra Fine Screen (PCS SUFS) measurement system prior to handoff, analyzed with BAMM! [81] gives estimates of wavefront error at grid points across the entire Keck primary, as shown in Figure 3.4. Three segments, 2, 9, and 23 were exchanged and indeed show 2.7 σ more wavefront error than nominal segments, as shown by the distinct population in the histogram in Figure 3.4. The total RMS wavefront error across the primary from this measurement is 649 nm, as compared to 517 nm when masking affected segments. Warping error could in part be sensed by the wavefront sensor and corrected by the high order DM. However, we assume any remaining error is reflected in the margin error term of the Keck AO error budget, shown in Table 3.2. The margin error term is a non-physical term required to match empirical Keck data, and may be due to primary errors which are otherwise unaccounted for in the error budget. However, even with PSF splitting, ORKID SRs are in line with analytic predictions for a high order error term that does not include additional margin error (that would reflect major uncorrected segment warping). Therefore, we do not expect the segment error had a major effect on the ORKID performance on this night.

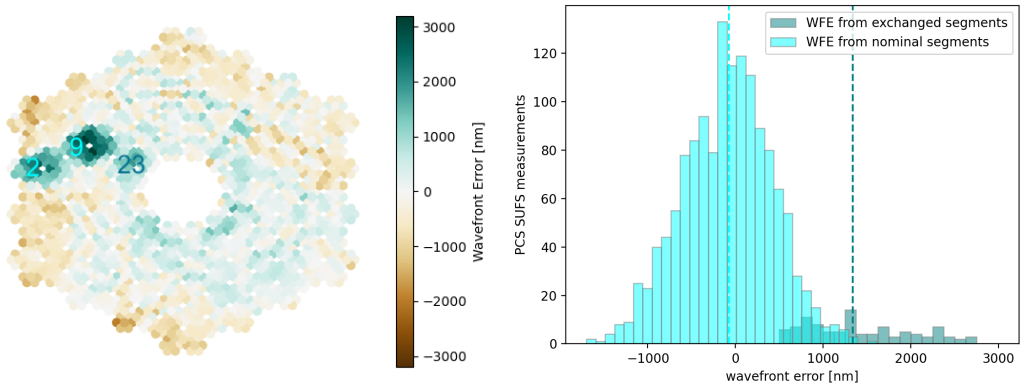


Figure 3.4: Left: Measurements from the PCS SUFS system showing wavefront error across the Keck-II primary. Three segments (2, 9, and 23) were swapped and not yet de-warped, which show higher error that the other segments. Right: Histogram of these data, divided into the nominal segment data and data from the three swapped segments. Dashed lines indicate the median of these populations. The swapped segments peak at 2.7 σ from the nominal segment population.

3.5 Performance and comparison to analytic predictions

In this section we describe ORKID performance according to a few key metrics, focusing on Strehl Ratio and encircled energy to report the imaging efficiency, as well as full-width half-maximum (FWHM) and point spread function (PSF) jitter to report the quality of image morphology. We directly compare these performance metrics to predictions from the AO error budget and similar analytic models to isolate where Keck predictions agree with or deviate from our reported performance.

3.5.1 Strehl ratio and encircled energy

We calculate Strehl Ratio and encircled energy for the ORKID data. For this analysis the ORKID data are dark subtracted, and short consecutive frames are added to effective exposure times of 250ms such that every directly reported value is conducted on data that is of equivalent exposure time, unless otherwise stated.

Strehl ratio calculation

We calculate Strehl Ratio with a procedure most similar to method (7) outlined in Ref. [230]. We start by drawing a 48 pixel (321.6 mas) radius aperture around the center of the source in the science image. We note that Ref. [230] suggests an aperture of 1" radius, but 48 pixels is the largest full circle that fits centered on the source within every image on which we wanted to run comparable Strehl Ratio analysis.

We draw a 17 pixel (113.9 mas) radius aperture immediately abutting the edge of the aperture around the central source, as shown in Figure 3.5. We use `photutils.aperture_photometry` [25] to estimate the photometry through this region, and calculate a per-pixel sky-background, which is subtracted from the science image. While Ref. [230] uses a ring around the aperture for sky background estimates, drawing the largest possible aperture around the source

does not leave room for a ring aperture.

We simulate a perfect Keck-like PSF, using the narrowband $H\alpha$=656nm and the Keck segment geometry, with the same pixel scale as the ORKID detector (Nyquist sampled at 656 nm). We draw a 48 pixel radius aperture around the center of the perfect simulated image. We perform aperture photometry through the aperture on the sky-subtracted science and the perfect simulated image, and use these to normalize the images, by dividing each image its aperture sum.

We then take $4\lambda/D$ ($\sim$ 8 pixel) box around the source for both the science and perfect simulated image, and upsample them by a factor of 8 using a fast Fourier transform upsampling method (built using `np.fft` methods) [106]. We take the maximum value of each upsampled region, and report their ratio as Strehl Ratio for this work.

Strehl ratio and encircled energy values

We report a median of SR of $1.3 \pm 0.25\%$ using bright (V magnitude = 2.1-6.1) natural guide stars at Hα, λ=656nm. We note that Strehl Ratios this low are not expected to be a robust benchmark for analytic predictions, but for visible light this performance is still above where we expect the Marechal approximation to break down, as discussed in Section 3.5.4.

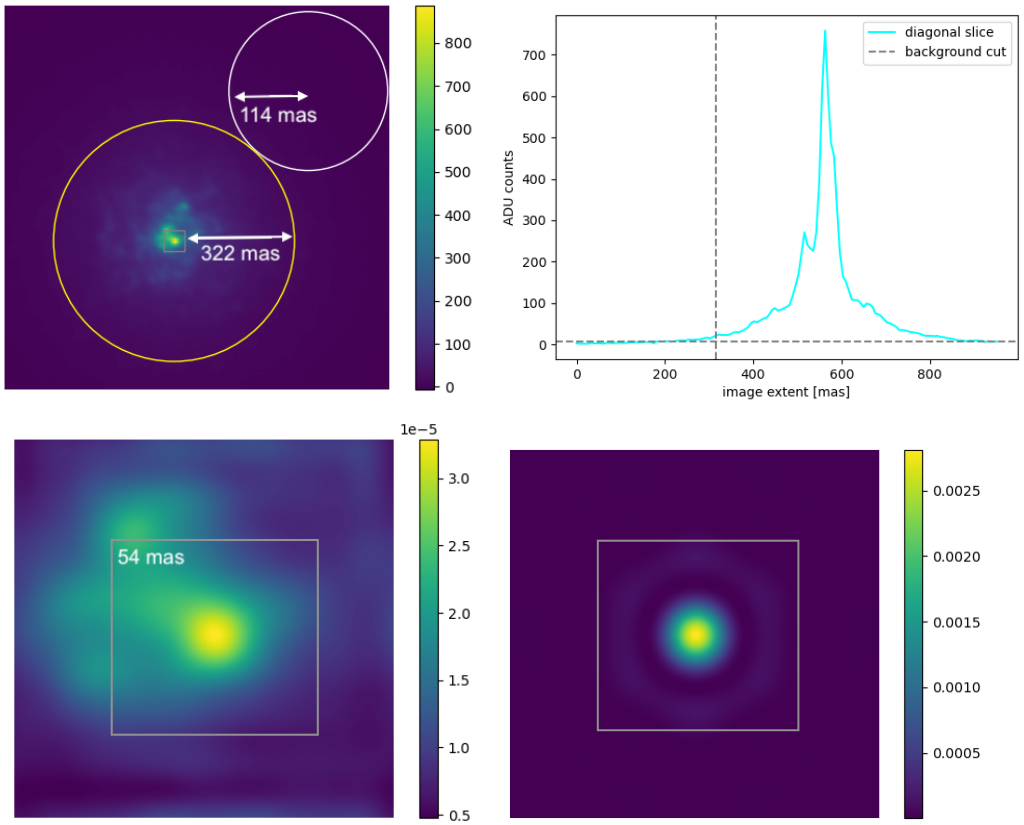


Figure 3.5: Top left: Example of the 48 and 17 pixel aperture (shown in yellow and white) to estimate photometry for normalization and background. This image is a median combination of 5 ms exposures of the bright target BS 603 through the Hα=656 nm filter. The grey box indicates the 8 pixel across region that is selected for upsampling. Top right: A diagonal cut through the top left image, showing that this cut clears the seeing limited halo for an appropriate background estimate. Bottom left: Upsampled 8 pixel region (by a factor of 8) to estimate maximum value of the science image. Bottom right: Upsampled 8 pixel region (by a factor of 8) to estimate maximum value of the perfect simulated image. The maximum pixel value is compared between the real image and perfect simulated PSF to estimate Strehl.

The greatest unforeseen impact to Strehl Ratio is a splitting effect in the PSF, which reduces the achievable Strehl Ratio by spreading light from the core into a bright second lobe or speckle, shown in Figure 3.6. For this reason we also present encircled energy as a performance metric. We do not see PSF splitting across every image, so encircled energy acts as a more stable metric than Strehl Ratio. Because the PSFs are Nyquist sampled (at 2 pixels across the PSF core) we use an aperture of 6 pixels to encapsulate the full intensity of a would-be core even when the PSF splits. Figure 3.6 shows an example of the PSF split as it presents in 3 separate ORKID frames, and how the 6 pixel radius aperture encapsulates core flux from both lobes from all three. We normalize the encircled energy to the total counts in a 48 pixel radius around the star. We report a median encircled energy of 7.5$\pm$0.3% through this 6 pixel aperture for bright stars.

While ORKID cannot readout at 1kHz, ORKID can consistntly take 1ms exposures, which is short enough to freeze system vibrations and atmospheric tip and tilt. Therefore, we also report the Strehl Ratio after recentering postprocessing is applied. Specifically, we estimate PSF jitter between each 1ms exposure, and then recenter and median-combine images with a common origin to remove the impact of residual tip/tilt errors. In the example shown in Figure 3.7 we report a Strehl Ratio of 2.4% on the bright 1.1 magnitude

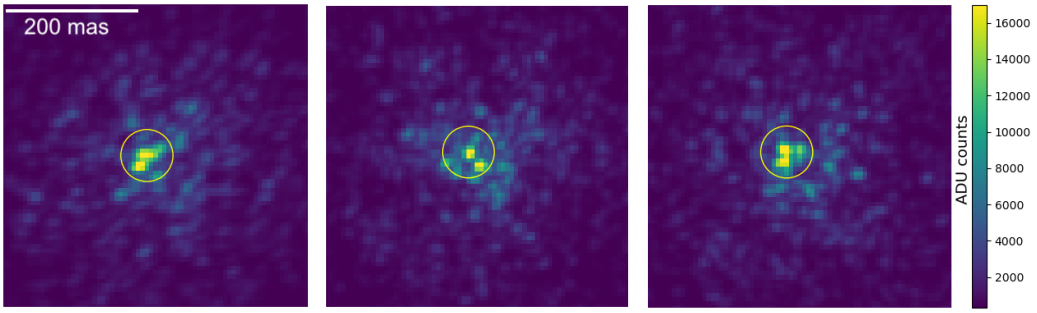


Figure 3.6: Three examples of PSF splitting in ORKID frames. All images are 5ms exposures on the bright target BS 603 through the Hα filter. The yellow circle shows the 6 pixel (3λ/D) aperture we use for encircled energy calculations to incorporate the flux of both bright lobes.

target BS 603, as compared to 1.2% on the same data median-combined but not recentered, i.e., still including residual tip/tilt errors.

3.5.2 Full-width half-maximum and jitter of the point spread function

We fit the ORKID PSF with the sum of two 2-dimensional Moffat profiles to model the core and halo. We report the core width of the fit as our core width for these data. Figure 3.8 shows the model fit as compared to a horizontal slice through a recentered image. After recentering one minute of 1ms Hα frames on the bright star BS 603 we resolve the diffraction limited core at 13.5 $\pm$ 0.4 mas, or 2 pixels for our Nyquist sampled system.

We track PSF jitter, or measured motion of the PSF center in consecutive

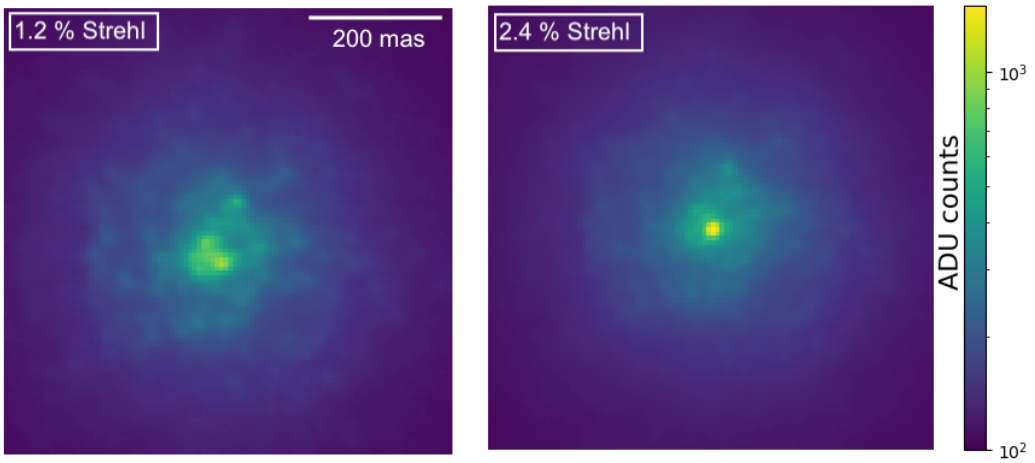


Figure 3.7: Left: One minute of median combined 1 ms exposures of the R=1.1 magnitude star BS 603 through the Hα filter at 656nm. Right: The same data with recentering to remove PSF jitter and the impact of recurring vibrations. Strehl Ratio improves notably in the (right-most) recentered cube, where we have effectively removed all residual tip and tilt errors and accounted for systematic vibrations.

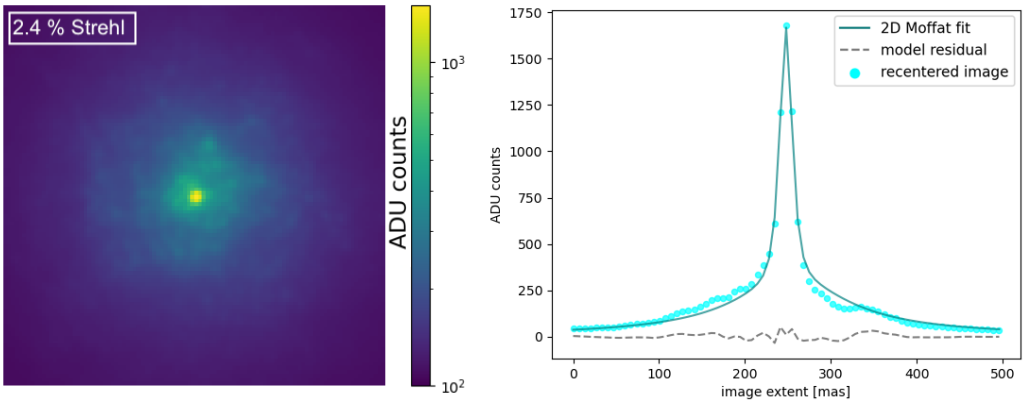


Figure 3.8: Left: One minute of median combined and recentered 1 ms exposures of the R=1.1 magnitude star BS 603 through the Hα filter at 656nm. Right: Horizontal central slice of the image as compared to a model built from the sum of two Moffat profiles. The profiles were free to fit the core width, the center in x and y, and exponential profile as free variables, but the peak of one profile was forced to match the peak of the science image. We report a diffraction limited core with a FWHM=13.5 mas. The standard deviation derived from the covariance of the model fit for the core width parameter sets an error of 0.4mas. The residuals between the model fit and data are shown in the dashed grey line.

ORKID image frames. With two hour long data sets of bright targets near zenith (HR 249 and HR 876), we track the motion of one frame compared the previous and mask outliers (a move more than 20 mas between frames). We report a root mean square (RMS) PSF jitter of (3.2, 4.4) $\pm$ 0.7 mas in (x,y) over sub-second timescales (as opposed to larger global moves discussed in Section 3.7.2. (Note that x and y are reported from the perspective of the ORKID camera, where y is aligned with elevationa nd x is alignd with azimuth.) For a discussion of how we report jitter in these images and estimate this accuracy see Appendix A.2. We note a consistent asymmetry with more jitter in the y direction from the perspective of the ORKID camera. Figure 3.9 shows a histogram of x and y jitter estimated from ORKID images and an example of jitter evolving in time. The spike, as well as other recurring features in these data will be discussed in Section 3.6.

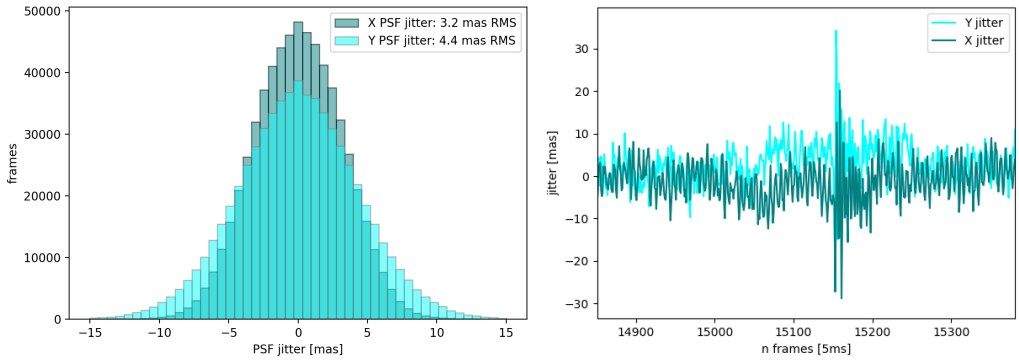


Figure 3.9: Left: Two hours of ORKID data for the two bright stars HR 249 and HR 876. The data consist of 5 ms exposures taken through the narrowband filter λ=656nm and summed to an effective exposure time of 250ms. The RMS PSF jitter collected for both targets is 3.2 and 4.4 mas $\pm$ 0.7 for the two camera x/y axes respectively. Right: Short selection of the HR 249 data to show the evolution of jitter in time series. Over short timescales, the data show a periodic pattern, as well as notable recurring spikes, termed "ring down events" [251].

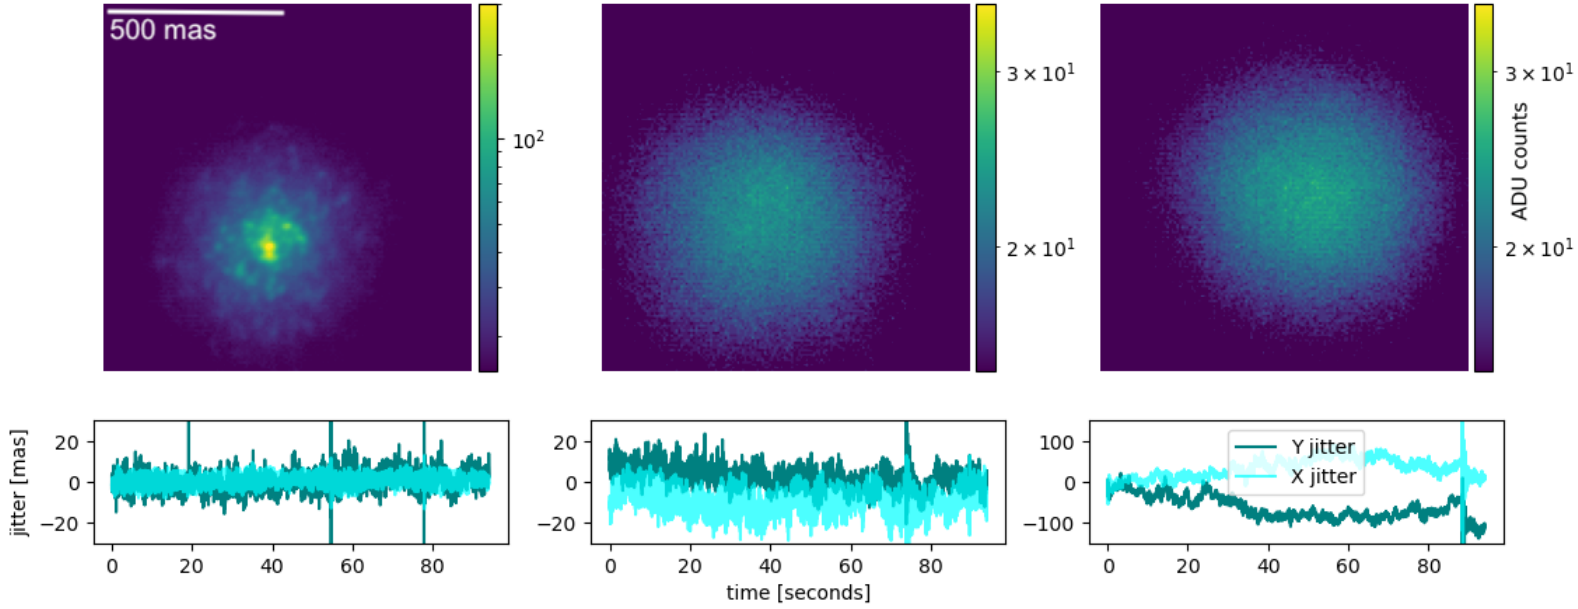


Figure 3.10: AO performance on HR 249 for (left) nominal closed loop data, (middle) data where only the tip/tilt loop is closed, and (right) fully open loop data, as shown by a median combination of 1.5 minutes of ORKID data. The bottom panels track the PSF jitter for each image. Note that while there is a near-diffraction limited core in the closed loop image, there is a PSF blurring effect reducing performance. The closed loop provides a substantial performance improvement over the seeing limited data.

The effect of closing the AO loops and delivering a near-diffraction limited core is shown in Figure 3.10 with the ORKID PSFs for the fully closed loop, just lower order tip and tilt loop closed, and open loop seeing limited data on HR 249. While the performance of a visible light instrument fed by a system with relatively few (349) degrees of freedom is notably poor performance compared to standard IR imaging, this demonstrates how much is being done by the Keck AO system to deliver this near-diffraction limited PSF. We will further discuss features of this PSF and limits to performance in Section 3.6.

3.5.3 Predictions from the Keck AO error budget

The Keck AO development team has put considerable work over the years into measuring and estimating various error terms in the Keck-II AO system [27]. In Table 3.2 we reproduce the main error terms we will discuss and a brief description of the causes of these error terms. In Table 3.3 we report the resulting Strehl Ratios for photometric bands of note. The values are reported for optimal performance for the current Keck-II AO system, which is a natural guide star of ≤ 5 magnitudes through the standard R bandpass on the visible light Shack-Hartmann wavefront sensor, a deformable mirror with 349 degrees of freedom, and a 1kHz control loop. The atmosphere assumed for these estimates has a velocity layer of 10.97 m/s and an r_0 of 14.7 cm (estimated at

a wavelength of 500 nm), as compared the to equivalent r_0 ranging from 8 to 25 cm according to the CFHT DIMM seeing monitor from the night as shown in Figure 3.2. See Section 3.5.4 for further description of each error term in Table 3.2.

Error Term	Cause	WFE
Fitting Error	too few degrees of correction	121 nm
Temporal Error	limited control frequency	60 nm
Measurement Error	low wavefront sensor SNR	20 nm
Margin Error	unknown errors, (segment misalignment)	130 nm
Total High Order Error	error terms added in quadrature	225 nm
Tip/Tilt Errors	low order sensing errors	58 nm †
Total Effective Error	all error terms added in quadrature	231 nm

Table 3.2: Summary of measurements and analytic predictions for error terms of note in the Keck-II AO system. † Tip/tilt errors are measured from the focal plane; this term can also be quoted as 3.5 mas.

Band	r'	i'	K
Central Wavelength (nm)	620	750	2190
Strehl Ratio (%)	2	7	65

Table 3.3: Strehl Ratio estimates through various photometric bands given the predicted error term using empirically adjusted analytic models as described in Section 3.5.4. We note that this is similar to the AO error budget outlined in Chapter 1.3, but here we present the Keck-II AO system at the time that these data were taken, with a different DM and WFS than the current Keck-II AO system.

3.5.4 Empirically tuned analytic models of strehl ratio and encircled energy

To build our empirically adjusted analytic model of Strehl Ratio and Encircled Energy, we start with the assumption that Strehl Ratio and encircled energy through small (in our analysis a $6\lambda/D$ diameter) apertures can be related with a freely fit linear factor, such that

$$\text{encircled energy} = A \times \text{Strehl Ratio} \tag{3.1}$$

where for the ORKID data, $A = 3.8$. Following the convention of the AO error budget [27], we calculate Strehl Ratios with three parts, by (1) using the Marechal approximation (SR_M) based on the high order error terms, (2) using an empirical correction ($\mathcal{E}$) mapped to real to data from the Advanced Electronic Optical System Telescope [83], and (3) multiplying this by a Strehl Ratio estimate based purely on the impact of PSF broadening induced by residual tip/tilt errors (SR_{TT}). We estimate each SR with guidance from Ref. [104]. Our final Strehl Ratio is

$$\text{SR} = \text{SR}_M \times \mathcal{E} \times \text{SR}_{TT} \tag{3.2}$$

with the Marechal approximation

$$\mathrm{SR}_M = \exp - \left(\frac{2\pi\sigma_{HO}}{\lambda} \right)^2 \tag{3.3}$$

where σ_{HO} is the high order error term in meters and λ is the wavelength at which we calculate Strehl.

The high order error term is a combination of all error terms (except for tip and tilt) in quadrature

$$\sigma_{HO} = \sqrt{\sigma^2_{measurement} + \sigma^2_{fitting} + \sigma^2_{temporal} + \sigma^2_{margin} + ...} \tag{3.4}$$

For the Keck AO system fitting error ($\sigma_{fitting}$) remains constant, as it is set by the number of degrees of freedom (or actuators) in the AO system. We use the AO error budget estimate for the high order Xinetics DM such that $\sigma_{fitting} = 121$nm. For HAKA, $\sigma_{H,fitting} = 51$nm. All of the ORKID data compared to models in this analysis were taken with an AO loop speed of 1kHz, which sets a constant temporal error of $\sigma_{temporal} = 60$nm. Margin error represents unknown remaining error terms needed to balance the AO error budget and match empirical data. The only feasible deviation between the AO error budget and ORKID come in the form of (1) higher than usual segment misalignment, discussed in Section 3.4.2, and (2) non common path aberrations

unique to the ORKID path, discussed in Section 3.4.1.

Measurement error varies with guide star magnitude and AO loop speed. With AO loop speed constant across increasing guide star magnitudes, we see an increase in measurement error, leading to a reduction in performance as tracked by encircled energy, as shown in Figure 3.14. Measurement error is calculated as a function of signal-to-noise ratio on the wavefront sensor detector

$$\sigma_{measurement} = \eta \chi d_a \frac{\sigma_{spot}}{SNR} \tag{3.5}$$

with the noise propagator η, closed-loop average factor χ, actuator spacing d_a, nd Hartmann spot size σ_{spot}. For more details on these constants and system specifications see Ref. [104]. Finally SNR is a function of photons and detector readnoise

$$SNR = \frac{N_{photons}}{\sqrt{N_{photons} + N_{pix}\sigma^2_{readnoise}}} \tag{3.6}$$

where $N_{pix} = 4$ for a detector binned 2x2, and $\sigma_{readnoise}$ is approximately $4e^-$ per pixel [122]. The number of photons can be estimated from the target flux through the detector zero point over the given exposure time, though with concurrent telemetry data on each target we empirically know the ADU counts associated with each guide star.

Strehl Ratio due to tip/tilt is calculated given some tip/tilt error in radians

σ_{TT}

$$SR_{TT} = \frac{\sigma_D^2}{\sigma_D^2 + \sigma_{TT}^2} \tag{3.7}$$

where $\sigma_D = 0.44\lambda/D$

We note that Marechal approximation breaks down with increasing error. For a fitting error dominated system, this can be estimated based on the number of degrees of freedom. For the Keck-II AO system with 349 DOF, this occurs at 2.7 [78] radians of phase, or 276nm at a wavelength of 656nm. This amounts to $< 0.1\%$ Strehl, which is an order of magnitude less than the values we compare to analytic predictions.

3.5.5 Comparison between data and models

To put ORKID SRs in the context of performance expectations, we build an analytical model of Strehl Ratio as a function of wavelength and AO error terms based on Ref. [104], following the method outlined in the previous Section 3.5.4. To explore how performance degrades with fainter guide star magnitude (i.e., how increased measurement error impacts performance) we explore the encircled energy of ORKID images with increasingly faint guide star magnitudes. We make the assumption that encircled energy and Strehl Ratio have a linear relation. We fit analytic expressions of Strehl Ratio degradation with increased fitting error to encircled energy to constrain the total high order error and

what fraction is made up of measurement error. Studying performance with guide star magnitude is a useful tool for empirical and analytic comparisons because performance falls off sharply as measurement error increases, creating a "knee" with guide star magnitude. Figure 3.11 shows the model matches to ORKID data.

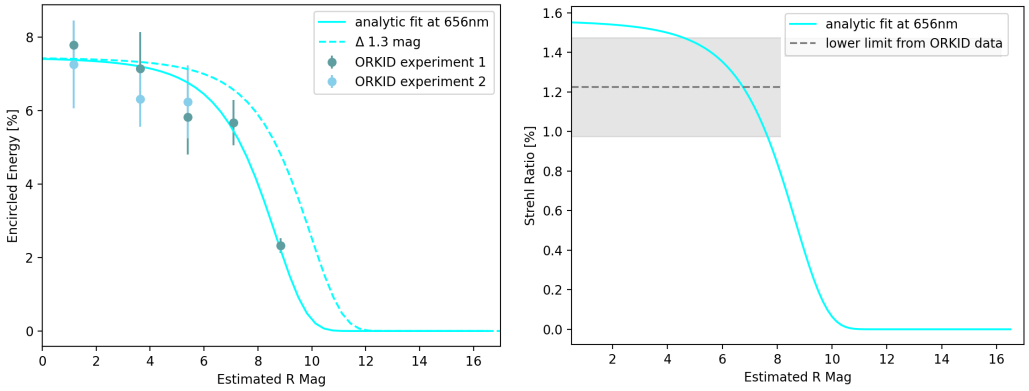


Figure 3.11: Left: Median encircled energy of ORKID data within a 6 pixel radius aperture vs. guide star magnitude while running at a constant 1 kHz AO control loop. Performance drops off, creating a "knee", as measurement error dominates the error budget and performance degrades. Error bars indicate the standard deviation of the encircled energy over minutes of data for each target, with a peak median performance of 7.5 $\pm$0.3% through a 6 pixel aperture. Experiment 1 and 2 refer to the two times over the course of the night this experiment was conducted. Targets that were too faint for good guiding at 1kHz are not included in this experiment. The dashed line indicates 1.3 magnitudes better performance, which would be in line with previous empirically-derived expectations for the impact of the sodium dichroic on SHWFS performance. Right: Comparing the original model's Strehl Ratio prediction to the Strehl Ratio values from ORKID data on bright targets with a median of 1.3$\pm$0.25%. We present these Strehl Ratio values as a lower limit at bright magnitudes, as Strehl Ratio values are severely impacted by PSF splitting. Note that even at 1.3 % Strehl, the Marechal approximation is still useful for fitting error dominated systems, see more discussion in Section 3.5.4.

With this model we estimate a high order error term of 228nm for bright ($\leq$ 5 magnitude guide stars) as compared to the AO error budget prediction of 225nm. Our model fit to the two guide star magnitude experiments has a median χ^2 value of 1.7 including errors. The difference of 3nm between the Keck AO error budget prediction and our model fit amounts to $\sim$ 1% difference. We also tie the comparison of Strehl Ratio before and after recentering to analytic predictions. With a measured low order residual tip/tilt error of 3.2-4.4 mas, we expect (from the analytic model described in 3.5.4) to see reduction in Strehl Ratio by a factor of up to 0.69, which under-predicts the difference between SRs with the recentered and median combined SRs (2.4 and 1.2, with an empirical loss of 0.5). Figure 3.7 shows these images side-by-side. We expect the PSF splitting accounts for the deviation between SR loss of 0.69 vs. 0.5.

3.6 Identifying adverse impacts to ORKID performance

In this section we identify unusual features in the ORKID PSF and limits to ORKID performance, explore potential causes, and discuss future mitigation strategies.

3.6.1 Impact of the sodium dichroic on visible light wavefront sensing in natural guide star mode

We examine ORKID performance (as measured by Strehl Ratio and Encircled Energy) as guide star magnitude increases, as shown in Figures 3.11 and 3.14. With stars on the order of 1-3 magnitudes in R band, ORKID performance is limited by fitting error and additional system error terms (e.g., optical errors and non common path aberrations), and not impacted by measurement error. Fainter than 6-7 magnitudes, if AO control frequency is held constant, the performance begins to degrade rapidly. For reference, nominal wavefront sensor performance in the near-IR would see a similar decline around 13th magnitude, as shown in Figure 3.14. The difference is the result of two effects: (1) RMS wavefront error is a larger fraction of visible light wavelengths, and as we will discuss in this section (2) the impact of the sodium dichroic during natural guide star wavefront sensing.

When using the sodium dichroic, light to the wavefront sensor undergoes a difference in brightness due to the narrower 15 nm bandpass of the sodium dichroic as compared to the broadband R 400 nm bandpass. Previous Keck estimates based on empirical data suggested this difference should amount to 2.2 magnitudes [151], but our analysis of the ORKID data suggest this gap is larger. We measure this gap directly by analyzing wavefront sensor telemetry,

which includes the counts per subaperture that hit the SHWFS for each target. However, a priori this can be estimated from ancillary data, by including the impact of the spectral energy distribution (SED) of a target, the different bandpasses between the R band and the sodium dichroic, and the Quantum Efficiency (QE) of the SHWFS detector.

Figures 3.12 and 3.13 show the process of estimating this performance difference. We fit the QE of the SHWFS detector used when this data was taken (the CCD39, as shown in Figure 3.12) and estimate the impact to (1) a flat spectrum and (2) spectral energy distributions compiled by SIMBAD/Vizier [64, 187, 66, 234, 299, 115, 35, 164, 96, 7, 102, 108, 272, 241, 304, 63, 62, 45, 157, 3, 199, 6, 49, 189, 188, 246, 228, 59, 8, 11, 303, 182, 31, 165, 224, 129, 53, 210, 147, 245, 36, 128, 5, 306, 61, 181, 263, 273, 88, 302, 200, 103, 221, 162, 123, 24, 172, 269, 40, 270, 127, 222, 258, 196, 171, 124, 238, 229, 48, 305, 301] for each target used in the guide star magnitude experiment, to find a magnitude hit from the sodium dichroic as compared to the standard wavefront sensor R band. These results are shown in Table 3.4 and shown in Figure 3.13. Since the time of these observations, a new detector has been commissioned for the SHWFS (an OCAM2K), so we also present estimates for this detector. We find this gap is closer to 3.5 ± 0.4, depending on the spectral type of the source and the detector. We compare this analysis to a similar experiment conducted with

NIRC2 data, as shown in Figure 3.14. We use the analytic model fit to ORKID data and map it to the IR, specifically K band, to compare with NIRC2. While we cannot use the NIRC2 data to further constrain our model, we show that our model is consistent with IR performance and independently collected data. (These data were collected as part of a wavefront sensor gain-tuning experiment conducted by the Keck AO Development Team.)

Target	CCD39 Δ mag	OCAM2k Δ mag	Spectral Type
flat	3.8	3.5	N/A
BS 603	3.6	3.9	K2
BS 458	3.5	3.7	F9
HR 476	3.4	3.5	A7
SAO 054443	3.4	3.6	G1
SAO 038218	3.4	3.5	A1

Table 3.4: Estimates of flux differences between the sodium dichroic and the R band during natural guide star mode for the CCD39 (SHWFS detector used during this experiment) and the OCAM2k (SHWFS detector used currently).

We note the impact of the sodium dichroic on natural guide star observations for future ORKID users, who should carefully consider the impact of this unusual wavefront sensing mode on their AO performance for faint targets. In practice, when allowing the control loop to slow down to account for fainter targets, we were able to control on targets down to $\sim$10.5 in R magnitude – fainter targets may not be possible with the sodium dichroic in natural guide

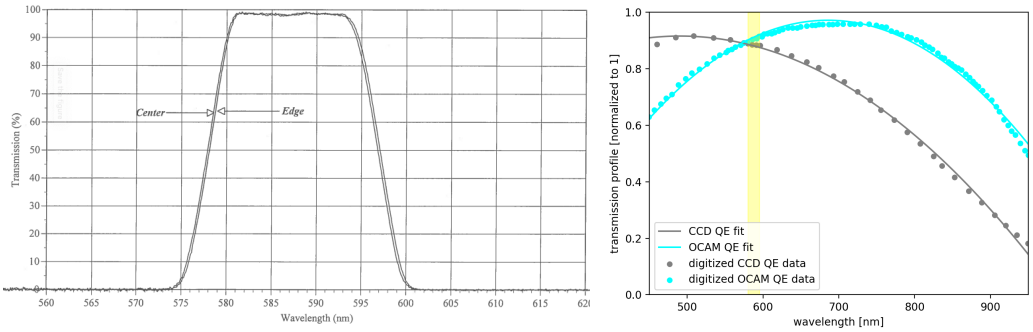


Figure 3.12: Left: Transmission profile of the sodium dichroic from Custom Scientific, Inc. We model the Sodium transmission as a flat 580-595 nm bandpass. Right: Using documentation from e2v and Oxford Instruments we provide digitized versions of the CCD39 (the SHWFS detector while these data were collected) and the OCAM2k (the current SHWFS detector.) We fit these data with a 2D polynomial and use this to estimate the full transmission profile in Figure 3.13. The sodium dichroic profile is shaded in yellow.

star mode. An alternative strategy is to use the full R bandpass for the SHWFS, to maximize AO performance on fainter targets, in which case 97% of light is transmitted to the wavefront sensor. This leaves only 3% for the ORKID camera, requiring longer exposures, and removing the ability to compensate for residual tip/tilt errors with recentering. More exploration of this in the context of the HAKA upgrade is discussed in Section 3.8.

3.6.2 Recurring vibrations in the AO system

When median combining ORKID images without recentering, ORKID has a notable image blur, for example the median combined image with closed-loop control in Figure 3.10. This is likely due to recurring vibrations, which we

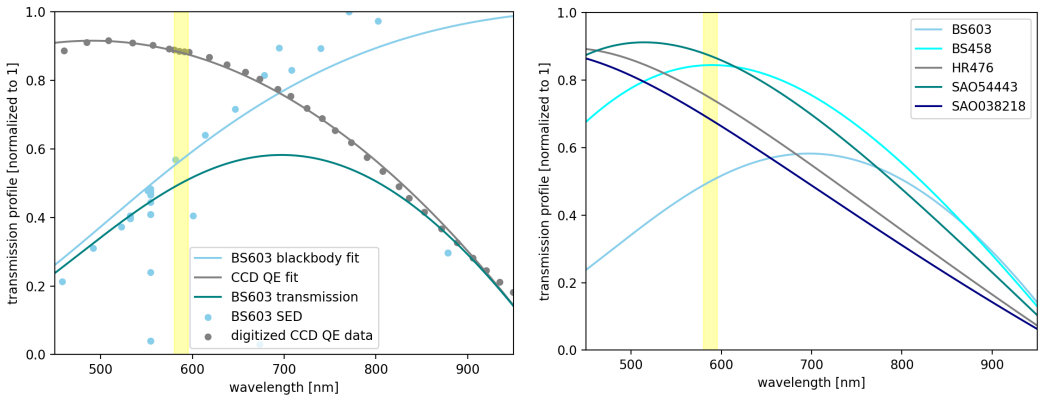


Figure 3.13: Left: We estimate the full transmission profile of a star to the SHWFS by combining the QE of the detector (in grey) and the SED of the target (in blue). Using data from Vizier collected via SIMBAD [64, 187, 66, 234, 299, 115, 35, 164, 96, 7, 102, 108, 272, 241, 304, 63, 62, 45, 157, 3, 199, 6, 49, 189, 188, 246, 228, 59, 8, 11, 303, 182, 31, 165, 224, 129, 53, 210, 147, 245, 36, 128, 5, 306, 61, 181, 263, 273, 88, 302, 200, 103, 221, 162, 123, 24, 172, 269, 40, 270, 127, 222, 258, 196, 171, 124, 238, 229, 48, 305, 301], we model the SED of each source and fit it with a simple blackbody. (Note that while a blackbody is not a perfect fit for the entire SED of each source, we fit around the data through R band, and see close agreement over this small spectral range for each fit.) We integrate the blackbody curve over the R bandpass of the SHWFS from 450-950 nm, including a 3% flux reduction due to the anti-reflective coating on the R band filter, and through the yellow region (the sodium dichroic bandpass from 580-595nm) to estimate the flux difference, and translate that to an effective loss in magnitude. Note that in this analysis we only need to preserve the ratio of flux between the sodium dichroic and the full bandpass to model an effective magnitude hit from the expected target brightness. Right: The effective profile for each of the 5 sources used in the guide star magnitude experiment through the CCD, the SHWFS detector at the time these data were taken. See Table 3.4 for an analytic estimate of the magnitude difference for each source through the CCD39 and OCAM2K (the current SHWFS detctor).

can analyze using Keck-II AO telemetry. With current Keck AO operations,

telemetry data in the form of Shack-Hartmann wavefront sensor slopes, tip/tilt

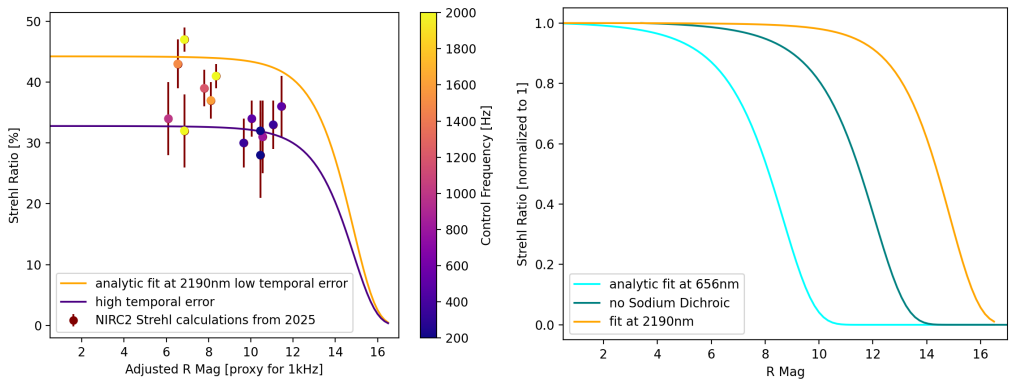


Figure 3.14: Left: NIRC2 K band data from 2025 with an analytic model matched to ORKID data, and then adjusted to NIRC2 parameters and incorporating a 3.4 magnitude hit from the sodium dichroic. Note that these Strehl Ratio measurements were taken at varying control loop frequencies, as indicated by the color bar. For the purposes of comparison with ORKID data we have adjusted the x axis magnitude as a proxy for flux hitting the detector (i.e., we report a magnitude indicative of how many photons hits the detector). We plot profiles including two temporal error regimes (due to slower control loop frequencies), as increased temporal error would impact performance. While our models are empirically tuned with ORKID data, they are consistent with NIRC2 data. Right: The analytic model matched to ORKID data at 656nm as compared to the same model after having removed the impact of the sodium dichroic (i.e., without the 3.4 magnitude hit as compared to an analytic model at 2190nm matched to NIRC2 data (shown on the left for the low temporal error case). The curves here are normalized to 1 to highlight the guide star magnitude at which these fall off. Note the full impact of the switch from visible to IR and the Sodium Dichoric is $\sim$ 6 magnitudes.

mirror commands, and high order (HO) Xinetics deformable mirror commands are recorded during operations and saved for a week after observations. We collected telemetry from the entirety of this night, and are able to match it to observing data with timestamps from the ORKID computer and the Telemetry Record Server. We note that these timestamps are not well-synced enough for direct frame to frame comparisons, but are sufficiently similar to examine trends over the same time periods. Using these data we make power spectral densities (PSD) of associated telemetry to isolate recurring vibrations affecting ORKID images, as shown in Figure 3.15. Among many interesting features in the PSD, we will comment on the 20 Hz and 302 Hz features that we can map to known Keck-II AO effects. For further discussion of recurring vibrations in Keck-II AO see Ref. [251].

We see high frequency vibrations at 302 Hz in tip, tilt, and HO DM commands, meaning that only with exposures under 3ms can we fully correct for these vibrations with post-processing. (Nominally for a perfect correction, we would need 0.3 ms exposures for a factor of 10 sampling.) Indeed, the analysis in Figure 3.8 to reveal the 13.5 mas diffraction limited core was only possible with 1ms exposures, and using 5 ms exposures did not successfully compensate for all residual tip/tilt errors. These vibrations were due to previous camera electronics for the CCD39 [275] that are no longer used for the SHWFS,

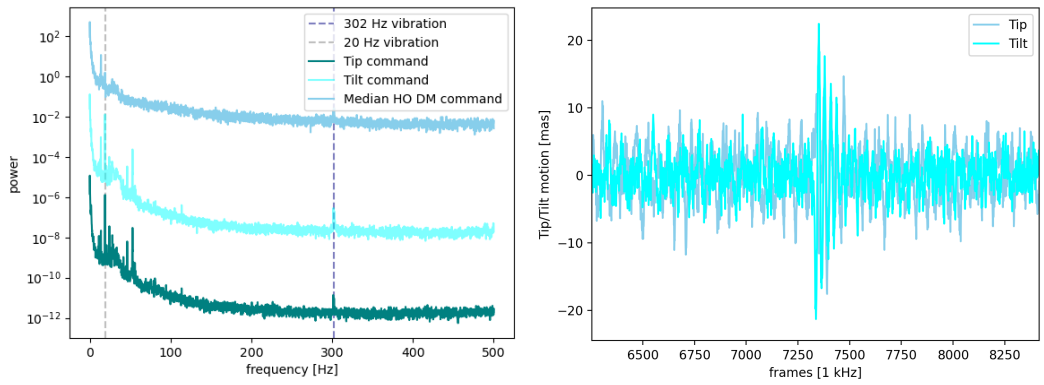


Figure 3.15: Left: Power spectral density functions (PSDs) of Keck-II AO telemetry, associated with the ORKID night. We present tip and tilt deformable mirror commands as well as high order Xinetics deformable mirror commands. These have been artificially separated in y space to display the features of each PSD uniquely. In this work we discuss the 302 Hz vibration due to wavefront sensor camera electronics and the 20 Hz spike corresponding to offloading to the secondary mirror at semi-regular intervals. Right: Ring down event in a time-series plot of Tip and Tilt deformable mirror commands.

and therefore will not be a lasting issue for ORKID images. However, we also see lower frequency vibration around 20 Hz, or the "ring down" effect, as described in Ref. [251]. Figure 3.15 shows these vibrations in both tip and tilt commands as well as high order deformable mirror commands at low and high temporal frequencies. During standard AO observations, when too much of the focus optical aberration has built up on the HO deformable mirror, it is offloaded to the secondary mirror. This occurs regularly with a frequency around 20 Hz. We not only see ring down events in telemetry, but in the ORKID PSF jitter, as shown in Figure 3.9. Offloading will only occur when a certain focus threshold has been met, so ring down events are not present in

all data.

We show the impact of this ring down effect by comparing two summed images. We add 50 consecutive 5 ms exposures including the ring down, and a sum with the ring down affected exposures removed. Figure 3.16 shows the impact of removing the ring down frames from the image sum, with a reduction in Strehl Ratio by a factor of 0.9, from 1.44 to 1.31%. Assuming this loss in Strehl Ratio is primarily due to additional tip and tilt error, this would be a consistent proportional loss in Strehl on affected frames, reducing images of higher quality by 10 %. Ring down events have been noted in previous studies impacting many Keck instruments [54, 110, 191], and can be mitigated by pausing offloading to the secondary during long exposures.

3.6.3 PSF splitting in fast ORKID exposures

Many of the ORKID exposures feature a PSF split into multiple lobes, as shown in Figures 3.6 and 3.17. Here we explore three potential causes and suggest that optical coma, perhaps in combination with errors from the segmented primary, is the most likely cause.

We simulate ORKID images using `HCIPy` [213], simulating two PSFs that approximate closed and open loop. For both we match the ORKID pixel scale, include the Keck primary geometry, atmospheric residuals appropriate for a

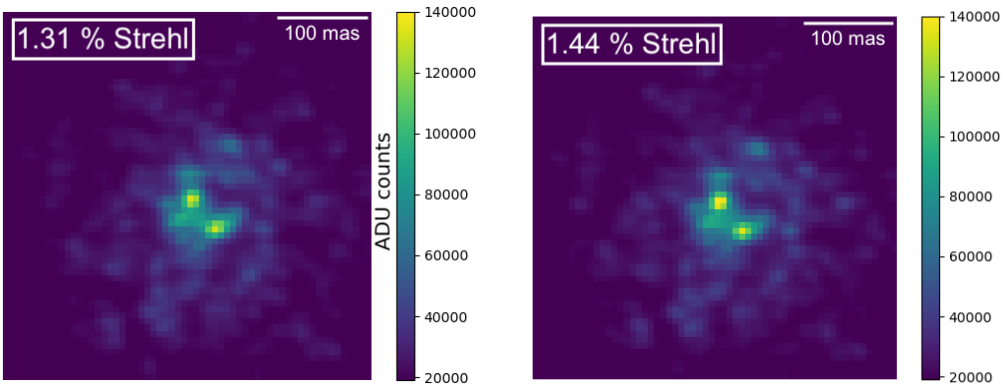


Figure 3.16: Impact of ring down events on longer exposures. Left: The sum of 50 5ms exposures (0.25s) including a ring down event (as isolated from ORKID PSF jitter). Right: A nearly identical set of frames with the ring down event removed. To match a similar data quantity this sum uses a slightly longer time window to match 0.25s worth of combined ORKID data. Note the Strehl Ratio degredation by a factor of 0.9 from the ring down event.

nominal night for Maunakea [190], and include non-common path aberrations (including the first 10 zernike modes) to visually match error in the PSF (using speckle background, image broadening, and PSF shape). For closed loop we then apply a control loop of 21 by 21 actuators (to match the Keck-II Xinetics deformable mirror) with a lag-time of 2 frames at 1KHz, assuming perfect estimation of the phase at every iteration.

We start by simulating AO residuals for Maunakea-like turbulence corrected by the Keck-II AO system, and their impact on the visible light ORKID PSF. These residuals alone are not enough to split the PSF, but we find that they are a required baseline to recreate this splitting effect. We first examine a

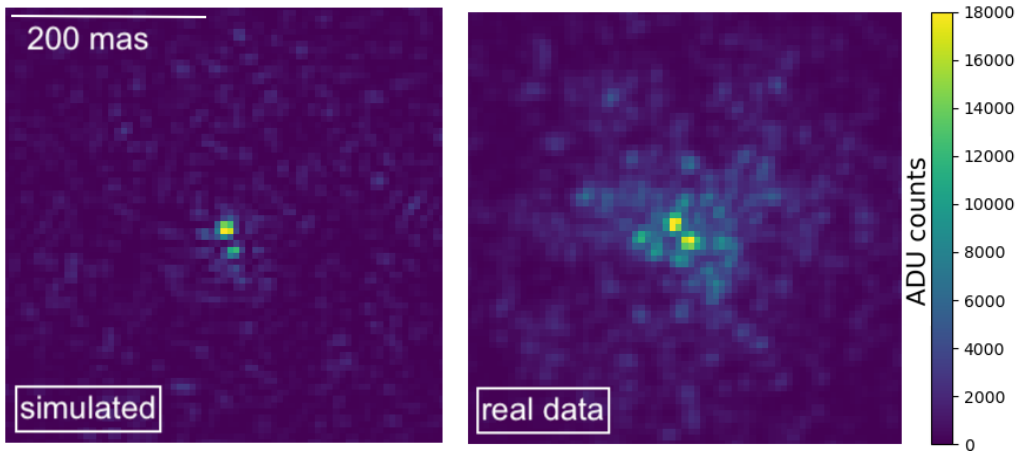


Figure 3.17: PSF splitting in the ORKID images. Left: Simulated ORKID image with AO residuals, segment primary errors, and an injected optical coma (414 nm peak to valley and 65 nm RMS) that mimicks the ORKID PSF split. This image was simulated to match the ORKID pixel scale at 656nm and normalized to the brightest pixel to share a flux scale. Right: ORKID data (a 5 ms exposure on the bright target BS 603) showing a split in the PSF, with two distinct lobes.

recurring vibration integrated over an exposure time that is longer than the vibration frequency, with the hypothesis that the vibration must slow at the peak amplitude where it changes direction and make two bright lobes at these two extremes. We simulate a vibration at 302 Hz, with an amplitude set to match the average spacing of the bright PSF lobes (3 λ/D). While this does indeed double the PSF in simulation, it also doubles the entire speckle pattern in a way that does not resemble the residuals and speckles in the ORKID.

Next we explore the impact of segment primary errors in the context of leftover warping due to the segment exchange night. With three segments

swapped, there was a large phase error on three adjacent segments. This is not unlike low wind effect [149], where a large discontinuous phase error around the telescope struts has been shown to split the PSF into multiple lobes. Section 3.4.2 discusses the state of the segmented primary based on ancillary PCS SUFS data and shows the three segments that were swapped. We matched the geometry of segments that were swapped prior to the ORKID night and injected phase errors starting with negligible (on the order of a few nm) to enough phase error to fully scatter the PSF (in practice ∼500nm of error at 656nm). In simulation this did not split the PSF; the result of including this was to scatter light into mid-spatial frequencies, in agreement with longer term studies of the Keck segmented primary [236].

Finally, we explore the impact of optical coma on the ORKID PSF. Coma shifts light asymmetrically into the side of an airy ring; this effect, along with the AO residuals creates the appearance of a second PSF lobe or bright speckle. We find the best visual match comes from combining this with segment primary errors. The coma injected in this simulation was $\sim$ 0.6 waves peak-to-valley, or 414nm, with an RMS of 65nm. Figure 3.17 shows the impact of injecting coma in the ORKID PSF with AO residuals and segment primary errors (using a phase error of 100nm). If the coma is indeed the cause of ORKID PSF splitting, better daytime calibrations and non common path aberration (NCPA)

correction will improve future ORKID performance. For more information on the daytime calibrations before our ORKID night, see Section 3.4.1.

3.7 Comparisons between ORKID PSF jitter and Keck-II AO tip and tilt residuals

3.7.1 X/Y asymmetry in PSF jitter

Over sub-second durations using two hours of data on the bright targets HR 249 and HR 876 we find the RMS jitter in x and y to be 3.2 and 4.4 $\pm$ 0.7 mas, and see that y consistently experiences more jitter. For comparison, the Keck-II AO error budget predicts 3.5 mas of tip/tilt jitter for both axes [27]. This asymmetry is not unique to ORKID, and is reflected in telemetry data, as shown in Figure 3.18. We directly compare a minute of concurrent ORKID data on HR 249 and corresponding telemetry, and find that across these data the RMS ORKID PSF jitter is (3.5, 4.4) for (x,y) whereas the RMS tilt/tip residuals is (3.3, 4.0). While the telemetry reports different values (on the order of $\sim$ 9% difference) when compared to concurrent PSF jitter estimates, we note that (1) the telemetry information comes from frames on the order of 1ms, as opposed to 5ms exposures used for this comparison and

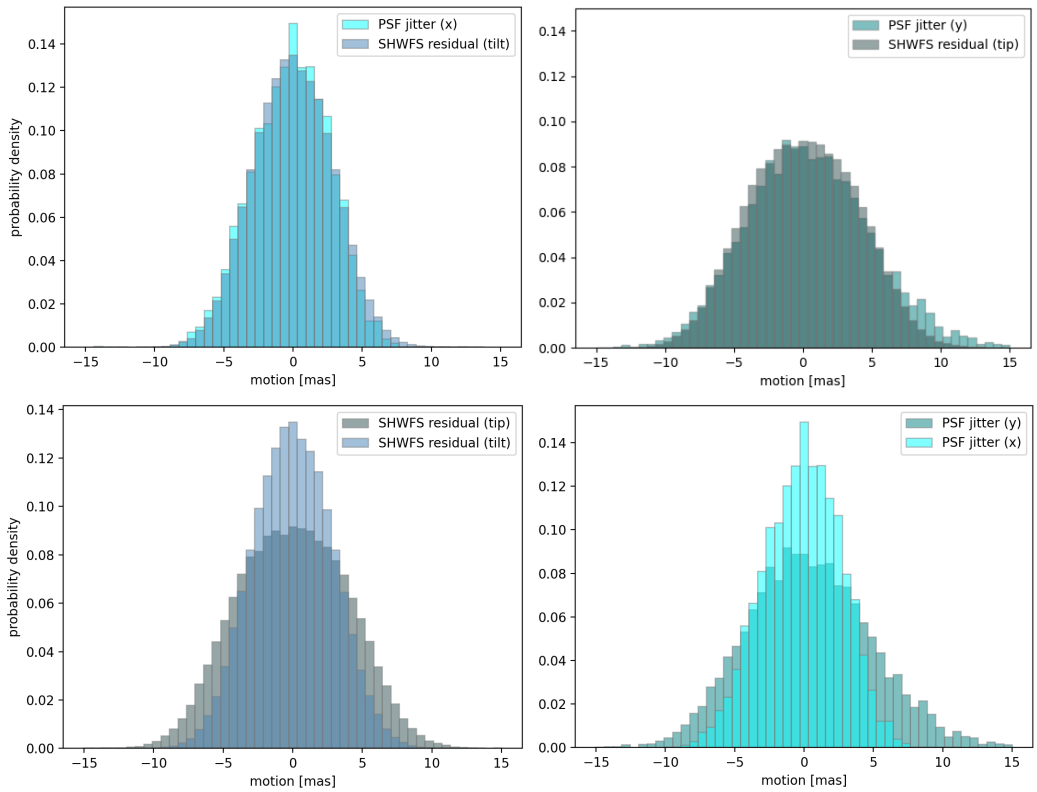


Figure 3.18: Concurrent data from ORKID (tracking PSF jitter) and telemetry (tracking tip/tilt residuals) over 1 minute of data for HR 249. The data consist of 5 ms exposures taken through the narrowband filter $\lambda = 656$nm. The RMS PSF jitter is 3.5 and 4.4 mas $\pm$ 0.7 for the two camera x/y axes respectively. The RMS residual wavefront error in tilt and tip is 3.3 and 4.0 mas.

(2) the discrepancy between the two sets of data is within the accuracy of our center-finding method.

3.7.2 Large drifts over long timescales

We find that while jitter remains fairly stable over short durations, y and to a lesser extent x jitter experiences a notable drift with time. Figure 3.19 shows

the evolution of the PSF drift over these longer ($\sim$ hour timescales), with a total y drift of 40 mas $\pm$ 4. This represents a movement of 6 pixels. For extended targets and smaller subarrays this could mean the object drifts out of frame.

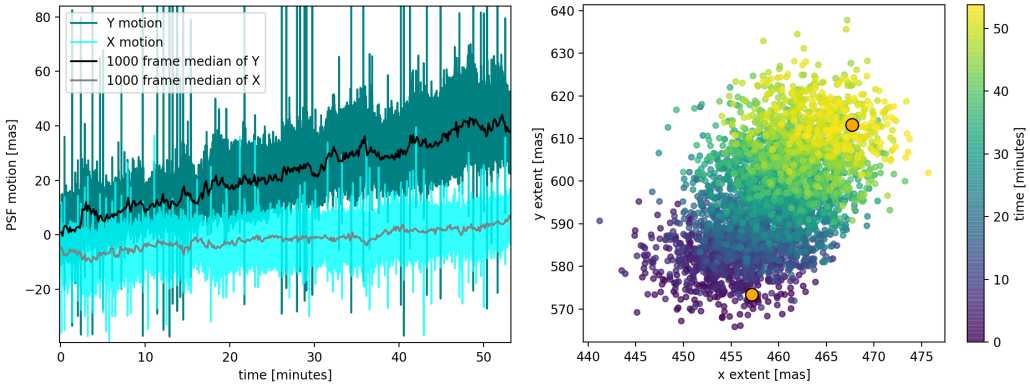


Figure 3.19: Left: Long duration drifts over the course of 30 minutes of data. We plot the x and y drift in positions for the target HR 249, as well as a running median position of 1000 frames, with a total drift of 40 $\pm$ 4 mas. Right: The same set of data from the perspective of the ORKID detector, showing the PSF center drift in Y (and to a lesser extent X) with time. Orange circles mark the start and end point of the PSF over time.

This drift is not reflected in telemetry, and therefore is unique to the ORKID path (non-common to the wavefront sensing optics.) At this time we are not confident of the cause of this effect. We note that the drift is aligned with the elevation direction. The Atmospheric Dispersion Corrector (ADC) is the only moving optic unique to the ORKID path and because our images were taken in a narrowband Hα filter, with improper ADC correction we would not expect to

see atmospheric dispersion but we would see motion. However, this could also be due to thermal expansion, rotation axis misalignment, or pointing origin misalignment.

3.8 Conclusion

In conclusion, we present the first characterization of Keck-II AO performance feeding a visible light instrument. We directly compare empirical visible light data from the ORKID instrument, Keck-II AO telemetry, and predictions from the Keck-II AO error budget. We report a Strehl Ratio of 2.4% and demonstrate a diffraction limited visible light ($\lambda = 656$nm) core (after recentering with post processing) with a FWHM of 13.5 $\pm$ 0.4 mas, and a median Strehl Ratio of 1.3$\pm$0.25% for images that were not recentered.

For future natural guide star AO with ORKID nights, we suggest taking the impact of the sodium dichroic into account, with an average difference of 3.5 ± 0.4 magnitudes in R. This effect varies with target SED and detector QE, and the exact difference can be directly estimated for targets with available SED data through the R band. We explore a persistent PSF splitting effect, likely due to optical coma and suggest more robust daytime calibration strategies be used to calculate a custom non common path aberration solution for ORKID. We note a previously characterized vibration, [191, 54, 110] that occurs in PSF

jitter, AO telemetry and an independent analysis of other Keck instruments [251], and present its impacts on science exposures, with a factor of 0.9 in Strehl Ratio loss. Not including this vibration we find the sub-second PSF jitter to be 3.2 and 4.4 mas $\pm$ 0.7 across the two camera axes, slightly more than AO error budget prediction of 3.5 mas and with an asymmetry echoed by Keck-II AO telemetry. We also find drifts (as much as 40 $\pm$ 4 mas) over durations of an hour that are unique to the ORKID path.

The Keck AO error budget predicts $\sim$ 2% Strehl Ratio should be achievable at 656 nm for a nominal night at Maunakea with the current real time controller and the 349-actuator Xinetics DM for a natural guide star of 5 magnitudes [27]. We report a median of $1.3 \pm 0.25\%$ Strehl Ratio for a natural guide star with an effective magnitude of 5.1 through the sodium dichroic. We expect this difference is due chiefly to PSF splitting, and recentering images recovers the expected 2% Strehl. Otherwise, our empirical data is in excellent agreement with the Keck-II AO error budget, with a similar prediction of the high order error term and overall performance.

3.8.1 Looking forward to HAKA science

The former Keck-II AO system has 349 actuators sampled across the pupil and a visible light Shack-Hartmann wavefront sensor. These two pieces of hardware

were the major limitations to ORKID performance; for bright stars fitting error dominated, with too few degrees of freedom for an accurate visible light correction. For faint stars measurement error will dominate, with performance that falls off steeply after R magnitudes > 10.5.

With the high order HAKA DM, which has just finished commissioning with Keck-II AO, the actuator count increases to 2844. This is predicted to deliver Strehl Ratios of 7-14% for R and i band imaging. With close agreement with the AO error budget, we see no obstacles to achieving this performance with HAKA + ORKID science for bright targets. However, using more actuators to achieve lower fitting error will require more wavefront sensor subapertures which will only further limit the natural guide stars that can deliver good performance for ORKID. Extrapolating forward the analytic models built from current ORKID performance to future performance with HAKA does recover 6-7% SRs in Hα, but extremely bright guide stars (on the order of -1) are needed to use the sodium dichroic mode with natural guide stars, as shown in Figure 3.20. This motivates in part why a space-based artificial guide star like ORCAS could be a powerful part of future visible light extreme adaptive optics. Additionally, 'IWA, the NIR pyramid wavefront sensor will be commissioned in 2029 [294], allowing visible light observations to use all of the light for science images while preserving high signal to the wavefront sensor. These

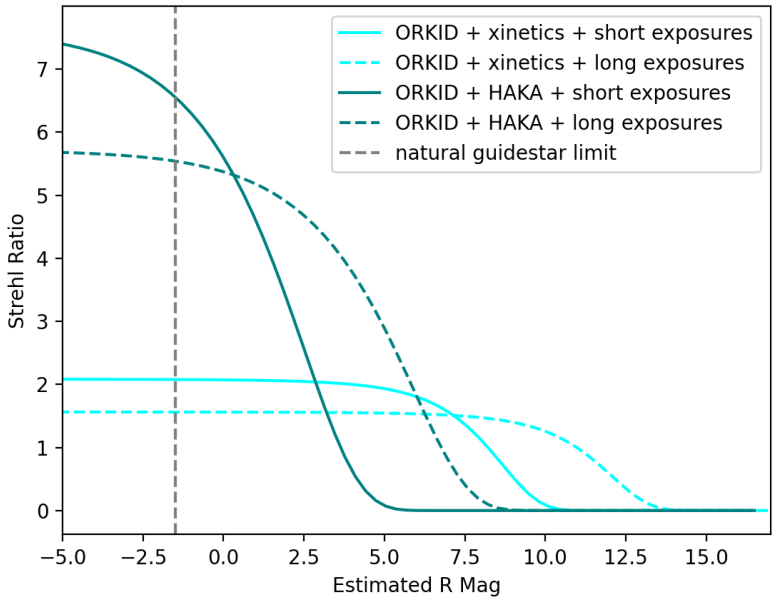


Figure 3.20: Current ORKID + Xinetics performance and future ORKID + HAKA performance with 1 kHz AO. We show an analytic model fit to current ORKID data using the Xinetics DM (in cyan) and extrapolated to HAKA (in teal) using predictions for less fitting error and increased measurement error (with a higher degree of freedom DM providing a better correction but also requiring more subapertures and therefore more light). Here the solid and dashed lines denote short and long exposures, or whether images are short enough to be recentered to remove residual tip/tilt errors and recurring vibrations (in practice 1ms) to resolve the diffraction limited core. For short exposures we plot the AO performance using the sodium dichroic, which would send enough light to the science camera for short exposures and include the impact of removing residual tip and tilt errors from the final Strehl Ratios. For long exposures we assume we can use the standard R band, increasing the AO performance at fainter magnitudes, but requiring long exposures that cannot be recentered. We have chosen to show the full performance profile for each mode, but the dashed grey line indicates feasible natural targets bright enough to meet this limit – past this line an artificial guide star is required.

upgrades will unlock a new era of visible light performance at Keck-II. Improved fitting and measurement error will likely leave segment cophasing errors, NCPA and temporal errors as the dominant error terms, motivating work on better segment primary corrections [236, 28], focal plane wavefront sensing to correct for NCPA [22], and predictive wavefront control to overcome temporal errors [71]. As the only facility AO-fed segmented telescope, Keck Observatory, and the ORKID instrument, are a unique testbed to investigate visible light AO performance for future extremely large telescopes.

Chapter 4

Predictions vs. Reality: Towards facilitizing data-driven predictive wavefront control on-sky at Keck Observatory

4.1 Abstract

Empirical Orthogonal Functions (EOF) has been demonstrated on-sky at Subaru/SCExAO and Keck/NIRC2 as a linear data-driven predictor, but has yet to be optimized. This chapter explores the practical elements of implementing,

optimizing, and facilitizing EOF on-sky at Keck Observatory. In simulation we explore the optimal predictive filter hyper-parameter space, study its stability under varying atmospheric conditions, and present simulated annealing as a method for fast on-sky hyper-parameter tuning. We discuss an update to the Keck real time controller (RTC) to include an EOF implementation and present daytime bench and on-sky results comparing predictive control to a standard integrator. We show improved wavefront correction from EOF as compared to a classic integral controller during on-sky observing, with an improvement of up to 20 % in root mean square wavefront error, and a point spread function stability improvement of 3 % using singular value strength as a metric. This work represents the first demonstration of EOF with a Shack-Hartmann wavefront sensor, and the latest implementation to have run on a large telescope.

4.2 Introduction

Predictive wavefront control, first suggested by Dessenne in 1997 [56], has been explored through many implementations over the years, ranging from classic and predictive Linear Quadratic Gaussian controllers [255, 217], neural network approaches [297, 266], through reinforcement learning [99, 146, 194, 215], and through classic minimization problems [93, 117, 278, 72]. Predictive controllers

consistently show promising performance with on-sky testing [94, 58, 168], and in particular, empirical orthogonal functions (EOF) at Keck has shown up to 3x improvement in contrast at separations of 3-7 λ/D, and 2x for 3 λ/D [280].

However, no fully facilitized implementations of predictive control are running on-sky and little has been published in the way of optimizing its performance for regular usage under varying atmospheric conditions. The previous EOF implementation at Keck [280] explored the optimal mixing between an integrator and predictor and optimal gain for the high order correction, but built a predictive filter with a consistent set of hyper-parameters. Similarly, one work [297] explored the optimal amount of data for an EOF filter to train on with Subaru/SCExAO telemetry, but predictive control at SCExAO was run on-sky with a static set of filter hyper-parameters (private communication, SCExAO team).

In the effort to move predictive control from an on-sky demonstration to a facilitized mode, we must address the the practical challenges of running a predictor on-sky. In this chapter I present a study of simulations of predictive filter performance for varying nightly conditions and make recommendations for a hyper-parameter optimization strategy during observing nights. I also discuss a new implementation of EOF in the Keck RTC, and present results of a predictor running on the Keck datytime AO bench and during on-sky Keck

observing.

4.3 Summary of empirical orthogonal functions and tunable hyper-parameters in the predictive filter matrix

A detailed description of Empirical Orthogonal Functions is available in Chapter 1.5.2. Here I present a summary for the purposes of this chapter. Using wavefront sensor residuals and deformable mirror commands we create pseudo-open loop (POL) records of the wavefront. This is the condition of the wavefront before the the deformable mirror correction is applied, which linearly encodes the trends of the atmosphere. We collect $\mathbf{n}$ POL states in a history vector $\vec{\mathbf{h}}$, in practice $3 \leq \mathbf{n} \leq 10$, which are used in real time to calculate a wavefront correction of the form:

$$\text{prediction} = \mathbf{F} \times \vec{\mathbf{h}} \tag{4.1}$$

where $\mathbf{F}$ is a filter matrix trained using up to a 1.5 minutes of data.

To calculate F we minimize the difference between $\mathbf{l}$ frames of training data

$\mathbf{D}$ and a truth condition $\mathbf{P}$ (data that is one time lag ahead of $\mathbf{D}$)

$$\left|\left|\mathbf{D}^T\mathbf{F}^T - \mathbf{P}^T\right|\right|^2 \tag{4.2}$$

and calculate $\mathbf{F}$ with a least squares inversion

$$\mathbf{F} = \left((\mathbf{D}^T) + \mathbf{P}^T\right)^T = \mathbf{P}\mathbf{D}^T(\mathbf{D}\mathbf{D}^T + \alpha\mathbf{I})^{-1} \tag{4.3}$$

where α is a regularization factor, akin to setting the number of singular values. For simulations and bench experiments we are able to run with $\alpha = 1$ but for on-sky we use higher values, which provides a more stable and less predictive correction.

In practice, on-sky we do not run with a pure predictor, but using a mixing factor $\mathbf{m}$ to combine a predictor with an integrator solution, such that

$$\text{final correction} = \mathbf{g} \times \left((1 - \mathbf{m})C_i + \mathbf{m}C_p\right) \tag{4.4}$$

where $\mathbf{g}$ is the high order integrator gain, and C_i, C_p represent the corrections calculated by the integrator and predictor respectively.

4.4 Optimizing filter hyper-parameters in simulation

We simulate the performance of EOF with varying hyper-parameters **l** and **n** as applied to atmospheres with varying windspeeds and Fried parameters. We simulate a classic single-layer Kolmogorov [135] phase screen with the Taylor frozen flow approximation [268] with a given Fried parameter r_0 and windspeed for a 10 meter telescope at a resolution of 48x48 pixels. We subtract piston, tip, and tilt modes, and apply a circular aperture. We sample and correct the atmosphere at 1kHz, with a lag of two time steps (2 ms), with perfect wavefront sensing and perfect correction – so the error is indicative only of error from the control itself.

For each hyper-parameter grid search, we build a predictive EOF filter with varying numbers of frames of training data (**l**) and varying history vector lengths (**n**), and build a heat map of the the median performance in root mean square (RMS) error. We compute a residual by subtracting the prediction of the full phase-screen at each iteration from the known injected phase at each iteration, and take an RMS of that difference at each point in time, leaving us with the RMS for each of 10000 iterations. We then take the median of that time-series distribution for the final value plotted and used for optimization.

Figures 4.1 and 4.2 show examples of differing r_0 with the same windspeed and differing windspeed with the same r_0 respectively. We note that the optimal history vector length changes based on different turbulence conditions.

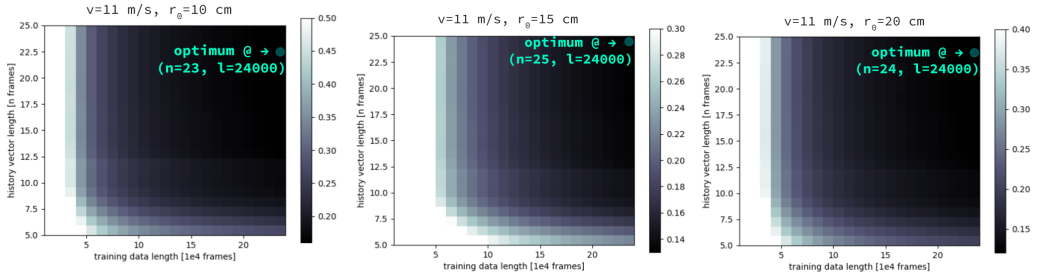


Figure 4.1: Heatmap of EOF predictor performance (in median RMS error in nm) as training data length **l** and history vector length **n** vary. Left: r_0 of 10 cm. Middle: r_0 of 15 cm. Right: r_0 of 20 cm. All of the above simulations have a windspeed of 11 m/s. Note that the optimal history vector length changes based on r_0.

While these simulations show that the optimal history vector length changes, simulations consistently show that the optimal amount of training data is the maximum amount available. We expect that predictor performance due to different lengths of training data in the filter matrix would be more impactful on real data, where including more frames of data introduces more noise into the training (our simulations are lacking measurement noise), and when the velocity layer will vary over some timescale.

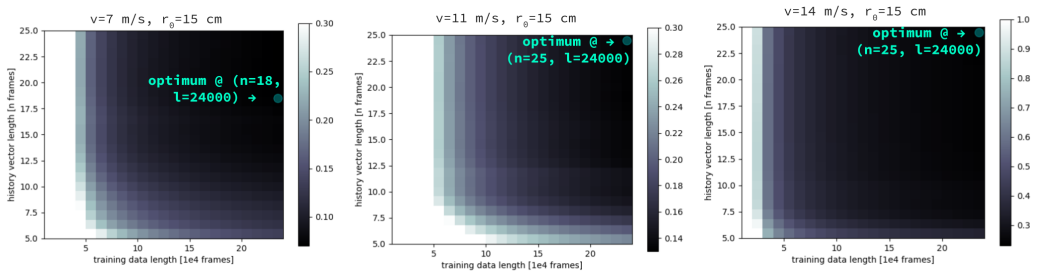


Figure 4.2: Heatmap of EOF predictor performance (in median RMS error in nm) as training data length **l** and history vector length **n** vary. Left: windspeed of 7 m/s. Middle: windspeed of 11 m/s. Right: windspeed of 14 m/s. All of the above simulations have an r_0 of 15 cm. Note that the optimal history vector length is different for the slowest atmosphere. Given the two right-most plots find an optimal in the corner of our grid, they may also find different optimal values given an expanded search space.

4.5 On-the-fly optimization with simulated annealing

If the desired strategy for EOF implementation is to do on-sky hyper-parameter tuning, building a full grid of filter hyper-parameter runs is likely not feasible. For these initial grids, we used 10 seconds of data and built a filter based on up to 24 seconds of training data; even for a system optimized to run in real time, to do this 600 times would take $\sim$ 6 hours. We present an alternative way to explore this hyper-parameter space more optimally with simulated annealing, a probabilistic method for finding the global optimum of a non-differentiable function.

Simulated annealing [141] is a probabilistic way to explore a discrete pa-

rameter space. The algorithm picks a random point of a given step size away from the current point and accepts it if the point provides a better solution, or if a Metropolis criterion is met with a given temperature. This process is named after metal annealing, where metal is heated and cooled strategically for optimal performance.

For each iteration, simulated annealing picks a new point, evaluates if that point is a better solution, and if not decides to keep the new point only if it meets the Metropolis criterion, which is met by comparing a randomly generated number to a function built from a cooling input temperature and how far the evaluation of the new point is from the previous. We note that simulated annealing and the Metropolis criterion make this algorithm vulnerable to getting trapped in a deep local minimum. For a given iteration i:

$$\text{cooled temperature} = \frac{\text{temperature}}{i+1} \tag{4.5}$$

$$\text{difference} = \text{evaluation}(i-1) - \text{evaluation}(i) \tag{4.6}$$

$$\text{metropolis} = \exp\left(\frac{\text{difference}}{\text{cooled temperature}}\right) \tag{4.7}$$

This logic runs for a given number of steps, but on-sky this could also be a cutoff based on how much the algorithm is still improving the solution or some minimum allowable wavefront error. Figure 4.3 shows how the algorithm walks

through one of the hyper-parameter search grids shown above.

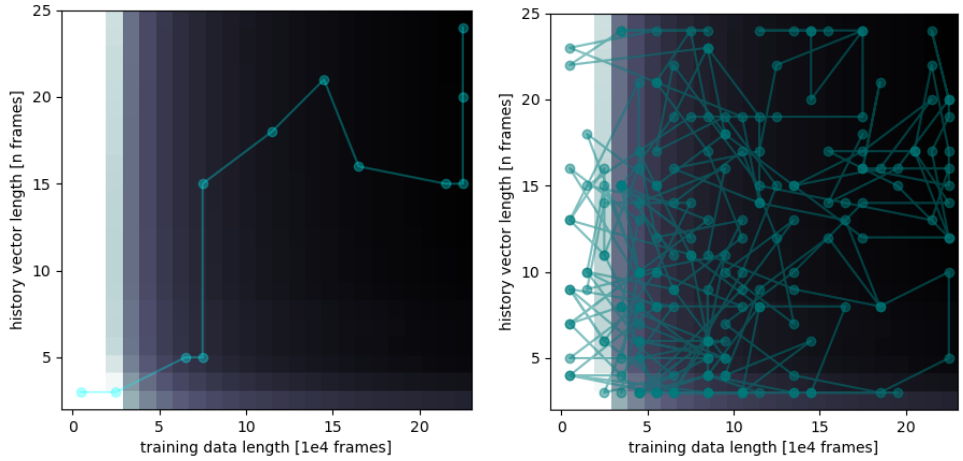


Figure 4.3: Simulated annealing samples many values, but it does not accept every step. This is the run of a simulated annealing algorithm applied to a grid with a windspeed of 14 m/s and an r_0 of 15 cm, chosen because the grid had the most variation. Left: Entire walk of the simulated annealing to find the optimal solution. Right: Every point sampled (whether or not it was accepted) over the 400 steps.

We applied simulated annealing to the previously shown filter hyper-parameter grid for a windspeed of 14 m/s and an r_0 of 15 cm, and found that the algorithm could repeatably find within 0.1 nm of the optimal value within 10 steps and repeatably recover the exact solution within 400 steps. For an on-sky strategy, we expect we could use closer to 10 steps, to find a very good if not exactly optimal solution. However, how quickly the simulated annealing algorithm converges to near its minimum will be dependent on the performance of each filter over a set of conditions, so testing with lab-data and on-sky telemetry will be needed before specific strategies can be recommended.

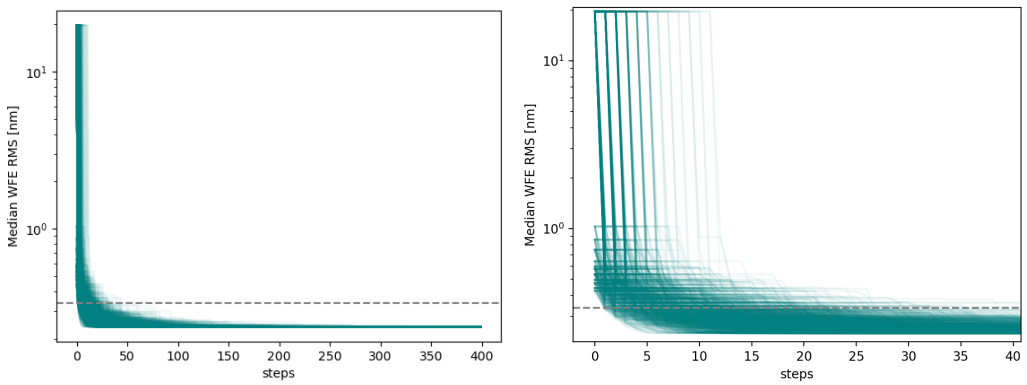


Figure 4.4: Simulated annealing is applied to the previously shown dataset, with a windspeed of 14 m/s and an r_0 of 15 cm. The input temperature is 119 and the step size is 5. The dashed line indicates when the algorithm reaches within 0.1 nm of the optimal solution. Left: full run. Right: closer detail on the first few iterations.

Figure 4.4 shows 1000 runs of simulated annealing on the same data set. This example was run with a step size of 5 (which maps to a single history vector frame and 50000 training data frames), and a temperature of 119. Initial testing found that higher temperatures (100-200) worked better for this data set but no quantitative optimization was done for temperature and step size in this case, due to how well the algorithm performed with these parameters. A strength of simulated annealing is its effective application to problems without significant optimization or previous knowledge of a parameter space.

4.6 Predictive control update to the Keck real time controller

Predictive control requires the use of multiple time-series frames of AO telemetry in real time – in practice up to 10. Because we need access to more frames of wavefront sensor information in real time than were available in the existing Keck real time controller (RTC) this required an RTC update. In 2024 we began discussions with Microgate and Keck to commission a new predictive wavefront control business unit in the Keck RTC. Those specifications and design documentation are described in an updated version of Ref. [293].

Figure 4.5 shows how this update interacted with the existing RTC architecture. We specify new computations and data products required, as well as the required access to these products in soft and hard real time. We define hard real time as less than the AO frame rate, such that any computations can successfully run without slowing down the AO loop (expected to run at 1-2 kHz). We define soft real time as the 1-3 minute timescale on which we want to be able to successfully build and update a filter matrix. In particular, the predictive control business unit requires a new data type built into the Telemetry Recording Server (TRS) which contains the psuedo-open loop (POL) reconstruction for every real time iteration.

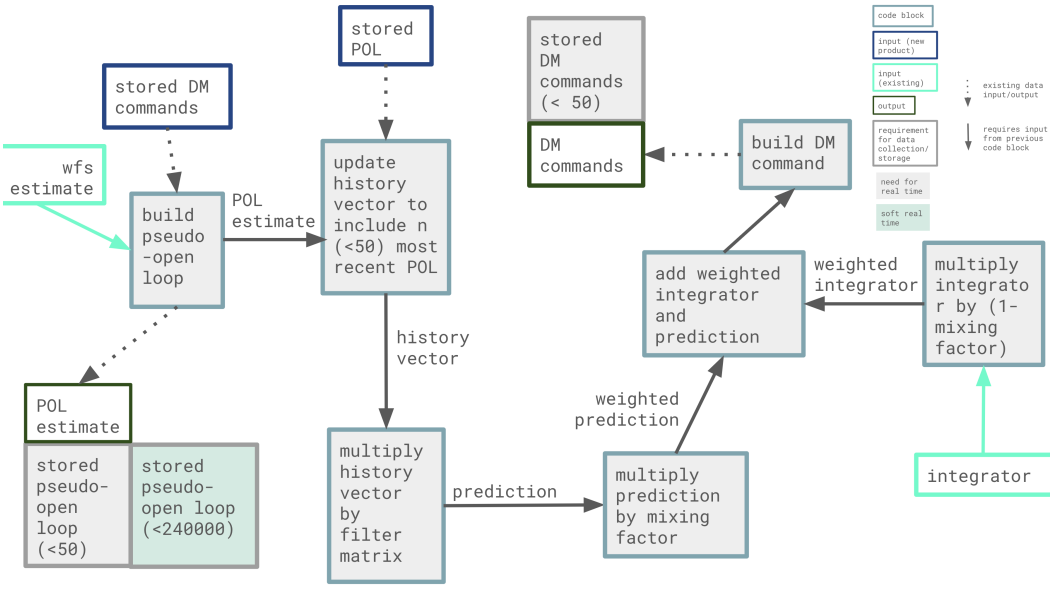


Figure 4.5: Block diagram of the predictive control update to the Keck RTC.

4.7 Performance verification with the Keck-II AO bench

After the completion of this upgrade in May of 2025, we tested the implementation on the Keck-II AO bench. During the day with the dome closed, Keck-II operates as a testbed, fed by a broadband fiber source, using the wavefront sensor, high order deformable mirror, tip/tilt mirror, and RTC that are used for nighttime on-sky operations. For the data collected on bench and on-sky described in this chapter, Keck was operating with a 349 degree of freedom high order Xinetics deformable mirror (DM) with each aperture mapped to one

Shack-Hartmann wavefront sensor (SHWFS) subaperture, with an OCAM2k as the SHWFS detector. During natural guide star mode (which also applies to the fiber source), tip and tilt slopes are estimated from the average tip and tilt of the SHWFS, and applied to fast tip/tilt mirror. We test the predictive controller by injecting two atmospheric turbulence profiles on the DM, one with a single injected wind layer, and one with the seven layer turbulence profile based on atmospheric profiling and modeling from Ref. [190]. We inject each profile by applying the shape on the high order DM. We note that this creates turbulence that the SHWFS could perfectly sense (i.e., there are no strutures at finer spatial features) and the DM could perfectly correct (i.e., there is no amplitude outside the linear range) which is not realistic to true atmospheric turbulence.

We compare the root mean square (RMS) performance of the integrator with a high order gain of 0.5 and the predictor with a mixing factor and α of 1, which is only possible in simulation or in optimal conditions with lab data. Both the predictor and integrator were run at 1kHz. The data are compared in a collapsed time series (i.e., 0 is the start of each run for all three populations, but they are not the same 0) to show the comparable evolution in time of each controller and the open loop turbulence. The turbulence is limited to two minutes by the DM control software, and loops at the two minute mark, which

creates a spike in wavefront residuals, as neither integrator nor predictor expect a fully disconnected scene of turbulence. We find promising bench results on both turbulence profiles, as shown in Figure 4.6. The predictor consistently out-performs the integrator in this lab verification experiment by a factor of 1.7 for the higher turbulence case.

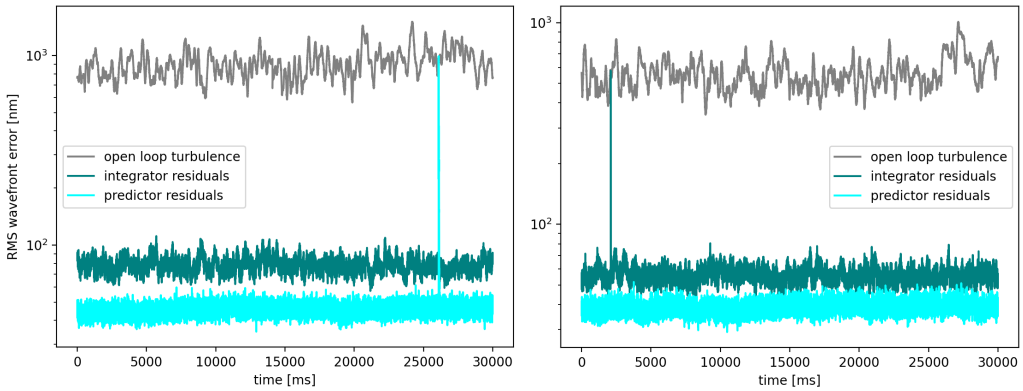


Figure 4.6: Performance of the integrator and predictor as compared to the full state of injected turbulence on the Keck-II daytime bench. We show the RMS of the wavefront residual error as measured by the SHWFS with iterations of the controller. Left: Seven-layer atmospheric profile from Ref. [190]. On this high turbulence profile the gap between the improvement of the predictor over the integrator is more noticeable. Right: For this single profile with one simulated wind layer, there is still notable performance improvement offered by the predictor.

4.8 On-sky results with Keck-II AO and NIRC2

4.8.1 Initial on-sky performance verification

We performed initial verification in a half hour of observing on the target BS 8088, V magnitude = 6.26 on the night of 10/5/2025, where we confirmed that we could successfully switch to the predictive control mode and close AO loops while maintaining nominal performance. In this initial test the α parameter was set to 2600, which results in a predictive filter that closely resembles an integrator. This was meant only as a first operational test.

We performed our first quantitative test in about an hour of observing on the target HD 189337, V magnitude = 6.5 the night of 10/19/2025, where we briefly explored possible filter parameters by manually scanning through filters, and ran consecutive predictor on/off experiments (predictor interleaved with integrator data) for a direct comparison between the predictor and integrator. For this experiment we used a high order integrator gain of 0.5, a mixing factor of 0.5, a low α=250, a history vector ($\mathbf{n}$) of 10 frames, and 60000 frames of training data ($\mathbf{l}$), with both the integrator and predictor running at 1 kHz. In $\sim$20 minutes of data we see a shift of 19% in the peak of the RMS wavefront error residuals sensed by the SHWFS towards a lower value. However, with these initial filter hyper-parameters we see a high noise tail. This is likely

due to the higher mixing factor and a low regularization parameter α, which creates a correction that is more aggressive and therefore more sensitive to data quality, as compared to optimal parameters from the second night, as shown in Figure 4.7.

Along with measuring real time wavefront residuals, we include NIRC2 imaging through the Bracket γ filter at 2190 nm. Figure 4.8 shows the Strehl ratios (SR) on the median combined data with and without the predictor running, which show comparable performance around 56% SR. Data were minimally reduced, with bad pixel masking and dark subtraction. Strehl ratios are calculated using a modified version of method 7 from Ref. [230].

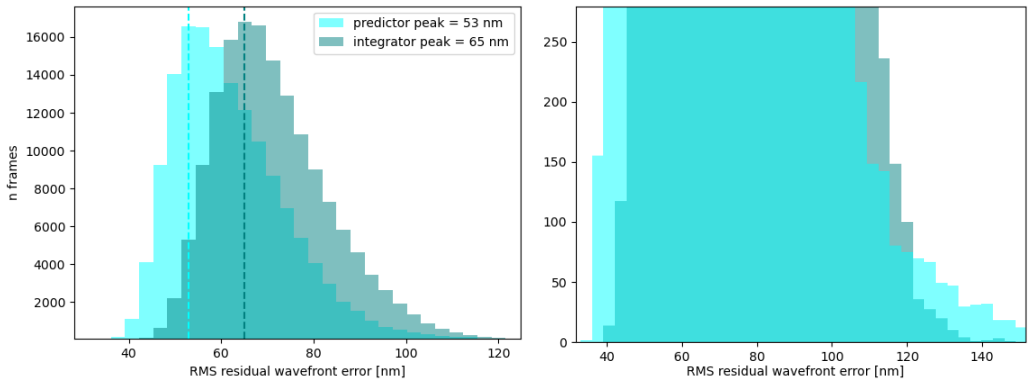


Figure 4.7: Histogram of RMS wavefront residuals during the initial predictor on/off experiment. Left: The mode of the two datasets are offset by 19%, with the predictor outperforming the integrator by the measure of the SHWFS wavefront residuals. Right: A zoomed in view of the left figure; in this experiment we see a high noise tail from the prediction population, likely due to lower regularization parameter α and lower mixing factor making the correction less stable that optimal filter hyper-parameters.

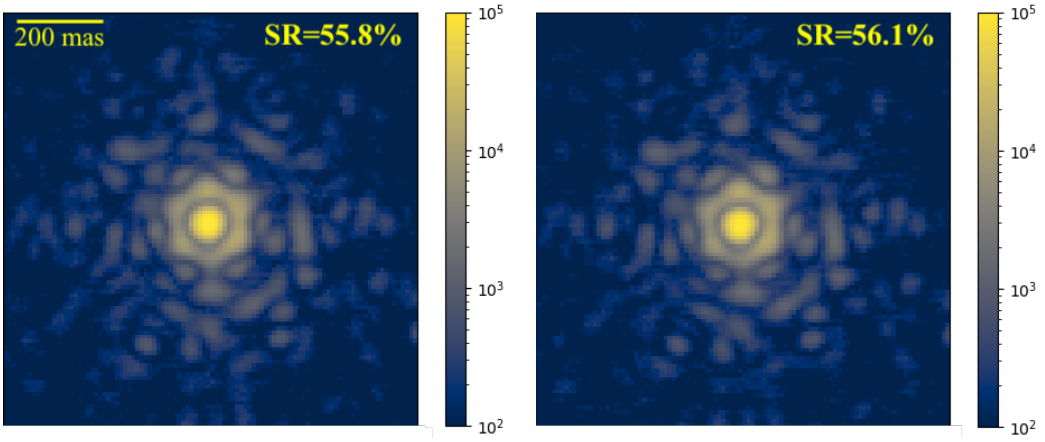


Figure 4.8: NIRC2 Bracket γ images of the target HD 189337. We show a median combination of 5 0.2 ms frames. Left: Images while the integrator was running. Right: Images while the predictor was running. The performance in SR is comparable between the two median combined images.

4.8.2 Optimizing predictive filter hyper-parameters on-sky

On the night of 12/03/2025, we conducted a half night of Keck engineering time to conduct parameter scan experiments to quantify predictor performance and explore optimal hyper-parameters for the predictor on-sky. For these experiments we used a high order gain of 0.5 and ran the AO system at 1kHz. We used three targets, BS 8936, V magnitude = 6.64, BS 327, V magnitude $\sim$ 5.8, and BS 1191, V magnitude = 5.77. These targets were chosen to be bright natural guide stars that would experience low airmass during the duration of the observations. We designed the night with 4 experiments in mind:

1. Conduct predictor on/off experiments for a quantitative comparison between integrator and predictor performance.

2. Scan through mixing factors, to gauge what weighted combination of integrator and predictor performed best on-sky.

3. Scan through regularization parameter α, to gauge how aggressive the predictive filter should be for optimal performance.

4. Scan through history vector and training data length, to gauge how much previous information produced the best predictive filter.

In our mixing factor experiment (2), we found that any mixing factor higher than 0.4 resulted in an unstable loop, which was prone to high noise or loosing closed-loop performance; similarly, mixing factors less than 0.4 were not notably different than the integrator population. Varying α had a similar effect: any lower than an α of 300 caused the loop to go unstable (or be highly prone to going unstable), and higher values had the effect of washing out the predictor improvement, e.g., moving it more towards a integrator correction. For these experiments where the correction was unstable we did not collect consistent data, as intentionally collecting minutes of data on an unstable correction risks building up large DM shapes that impact quality the rest of night. For this reason we focus the rest of this analysis on experiments (1) and (4).

We compare the performance of various filter hyper-parameters by comparing the residual wavefront error sensed by the SHWFS. We take the RMS of the residual at each frame, and report that value as wavefront error. We note that on-sky the SHWFS is not sensitive to the entire wavefront error term (estimated to be an RMS error of 225 nm [27]). The SHWFS is blind to fitting error (error due to limited deformable mirror degrees of freedom) due to its sampling to a matching number of subapertures. Similarly, SHWFS signal should be noisier as measurement error (error due to low SNR on the SHWFS) increases, but may not reflect the exact predicted error term. Finally, the SHWFS has been shown to be insensitive to segment co-phasing errors in the primary mirror, and will not sense non-common path aberrations (NCPA) that occur in a path other than the SHWFS path.

Figure 4.9 shows the results from experiment (4), exploring how much data are used to train the predictive filter and how much data are used in real time to calculate the correction for the next step. We try history vector lengths of 5, 8, and 10, and see a clear preference for a history vector of 10 frames. We try a training data length of 30000, 60000, 90000 frames; with a frame rate of 1kHz this amounts to a 0.5, 1, and 1.5 minutes of data. We see a clear preference for a training data length of 60000 frames. During this experiment more training data introduced noise into the correction, which as expected

does diverge from simulated performance. We use these filter parameters for the predictor on/off experiment (1), to compare integrator performance with optimal predictor performance.

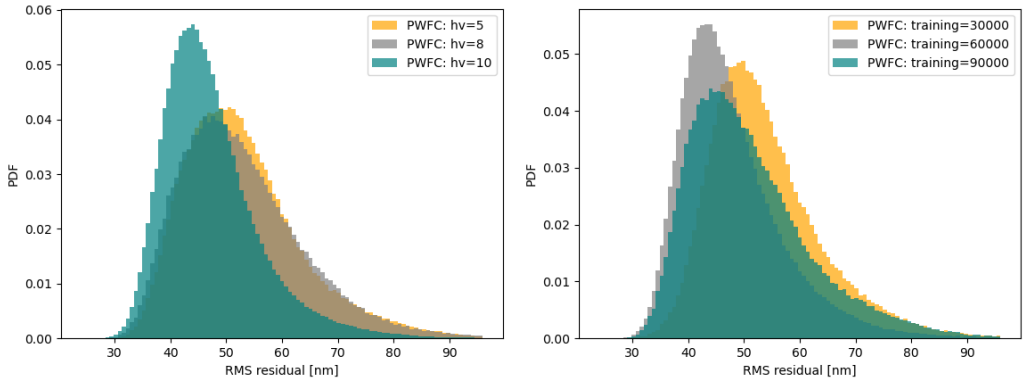


Figure 4.9: Selecting optimal history vector (hv) and training data lengths for the predictive control filter. Using SHWFS wavefront sensor residuals during predictor periods running with different filter hyper-parameters, a clear set of optimal hyper-parameters emerge. Left: Data are sorted by history vector (number of frames used by the controller at every real time iteration) length. The longest history vector of 10 frames is preferred. Right: Data are sorted by training data (number of frames used to train the filter matrix) length. The median training value of 60000 frames is preferred, as more training data introduces noise into the controller. This difference speaks to how real noise in data impacts optimal training, but not history vector length. For both sets of histograms, the training data and history vector respectively vary while the other is held constant. Data on all three targets (BS 8936, BS 327, and BS 1191) are included.

4.8.3 Residual wavefront error and Strehl ratio as performance metrics

We compare performance between the integrator with a high order gain of 0.5 and predictor performance, with AO running at 1kHz. We discuss three populations of predictor performance: (1) with a mixing factor of 0.4, a history vector length of 10, a training data length of 60000, and an α of 300, which were found to be the optimal filter hyper-parameters, as discussed in Section 4.8.2, which we refer to as"static parameters", (2) with a mixing factor of 0.4, varying history vector between 6-10, varying training data between 60000-1200000 frames, and an α of 300, which we refer to as "varying data length", and (3) a mixing factor of 0.4, a history vector length of 10, a training data length of 60000 frames, and a varying α between 1-1000, which we refer to as "varying α or alpha".

We compare how the SHWFS wavefront residuals vary in time series during experiment (1) with fast switches between integrator and the predictor with static parameters. Figure 4.10 shows the notable improvement in wavefront error, as well as reduced noise while the predictor runs. Figure 4.10 also explores the same data, along with the predictors with varying data and varying alpha in the form of a histogram to examine where performance peaks. While both of the varying parameter populations do not perform as well as the static

population, all of the predictor populations have a median performance in RMS wavefront error that is comparable to, if not less than the integrator. We note that extreme ($\alpha < 200$) values can induce noise. These results are summarized in Table 4.1. These results show that hyper-parameter optimization could be used to get the most out of predictor performance, but sub-optimal parameters are unlikely to degrade performance worse than a classic integrator.

We note that the data from experiment (1) were collected while running with the same filter matrix during real on-sky turbulence. This implies that for the conditions on this night a single filter matrix was sufficient for up to ∼20 minutes, meaning that replacing a filter matrix every 5-10 minutes will likely be sufficient for standard on-sky operations.

Table 4.1: Numerical summary of controller performance according to SHWFS residuals.

Controller	Frames	m	n	l	α	Median (nm)	STD (nm)
predictor	470000	0.4	10	60000	300	46.1	6.3
predictor	4520000	0.4	6-10	30000-90000	300	51.4	10.2
predictor	3090000	0.4	10	60000	1-1000	59.0	62.3
integrator	484000	N/A	N/A	N/A	N/A	57.8	8.4

We also present NIRC2 imaging data corresponding to predictor on/off periods shown in Figure 4.10. Figure 4.11 shows a median combination of

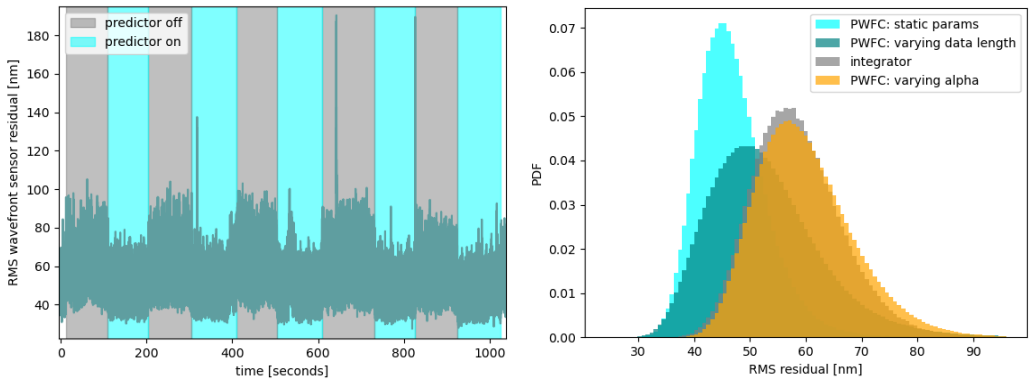


Figure 4.10: Left: RMS residual wavefront error from the SHWFS over time during the predictor on/off experiment. Cyan and grey regions indicate where the predictor and integrator were running respectively. Switching on the predictor consistently reduces the RMS wavefront error. Right: Histogram of the same predictor on and off population, as well as the performance of the predictor while exploring the filter hyper-parameter space. The optimal filter hyper-parameter population show a clear offset from the integrator population, (with a 20% improvement in RMS wavefront error) but even the populations with varying parameters show comparable to if not less wavefront error as compared to the integrator.

Bracket γ NIRC2 exposures, and an SR for each predictor on/off image, calculated with a modified version of method 7 from Ref. [230]. The SRs between the integrator and predictor are comparable, around 46 %. A histogram of the SR for each individual frame and a Gaussian fit to the data is shown in 4.12. The peak in population for each individual frame is comparable between predictor and integrator, with a difference in population that is less half of a standard deviation. The difference as recorded by the SHWFS between the predictor on and off is $\sim$15 nm, which using an empirically adjusted Marechal approximation akin to [27] results in a difference of $\leq$ 1.5%Strehl ratio, which

would be hard to resolve in a seeing averaged night, especially when other factors are likely limiting Strehl ratio (e.g., 10% variation between the October and December night despite comparable wavefront sensor error measured from the SHWFS).

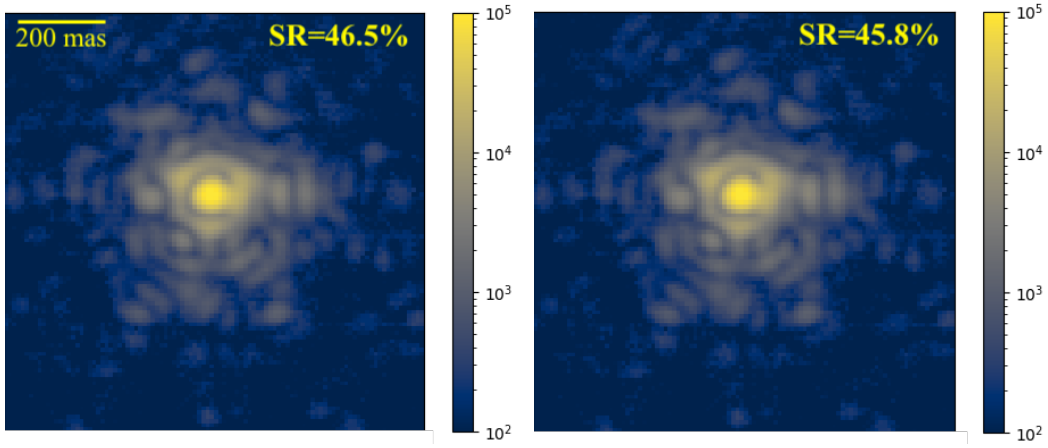


Figure 4.11: Keck/NIRC2 data through the Bracket γ filter from the night of 12/03/2025. These data consist of a median combination of 40 0.1 second frames on the target BS1191 while (left) a classic integrator was running and (right) a predictor with optimal filter hyper-parameters was running.

4.8.4 High contrast imaging performance metrics

In the context of future high contrast imaging applications, we explore two metrics that (1) demonstrate coronagraph-free contrast (due to the time-consuming nature of coronagraphic observing in the context of a half night) and (2) speak to image stability for better PCA-style data reduction methods. In Figure 4.13 we show a coronagraph-free contrast curve, by intentionally over-

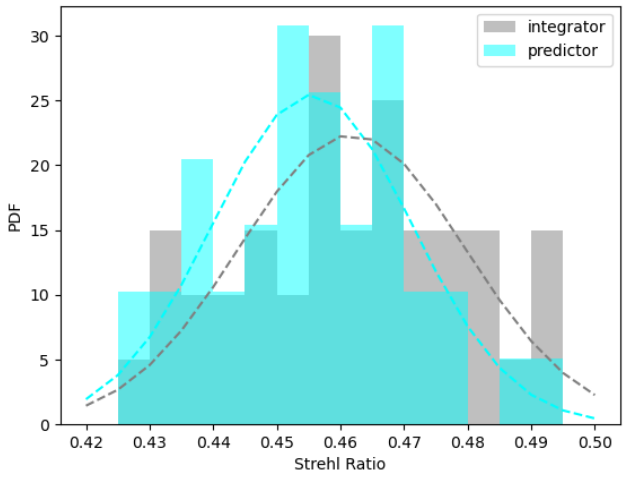


Figure 4.12: Strehl ratios of the Keck/NIRC2 images taken during the predictor on/off experiment. Each NIRC2 frame is 0.1 seconds long on the target BS1191. The integrator shows $< 1\%$ improvement in SR over the predictor, which is less than the population standard deviation and likely negligible.

saturating the NIRC2 detector to explore detector response with separation from the bright star at the center. We estimate contrast on this over-saturated image using the `pyklip.klip.meas_contrast` [290] function applied to a median combination of 5 0.5 s frames. This contrast proxy also shows comparable performance between the predictor and the integrator.

We also calculate the principal components of the 40 predictor and integrator images using `scipy.linalg.svd` [284] to explore PSF stability in the context of standard high contrast imaging reduction strategies. Here we see the predictor outperforming the integrator, with less power in most modes, with a net improvement of 3%, also shown in 4.13. High contrast imaging metrics show a

subtle performance improvement as compared to a standard integrator, and future work will explore issues of K2-AO bench noise floor or model mismatch that may be limiting predictor performance.

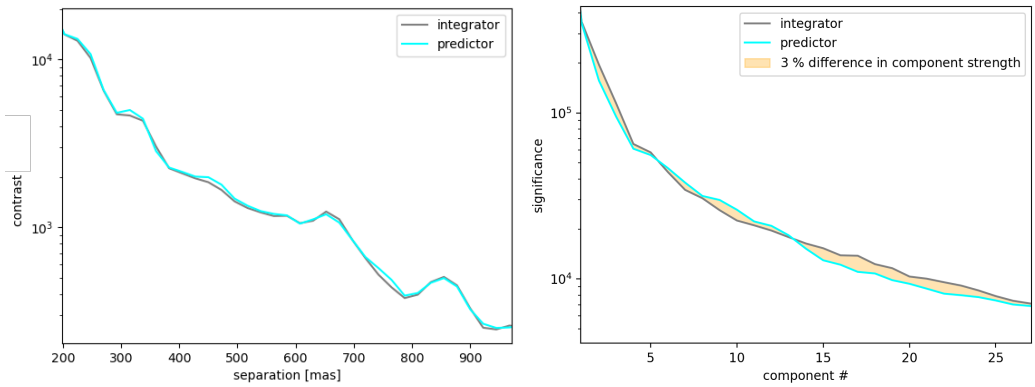


Figure 4.13: Left: Coronagraph-free contrast curves taken with Keck/NIRC2. By intentionally over-saturating the NIRC2 detector (0.5 seconds on the target BS1191) we can roughly estimate comparative contrast between the two controllers. The integrator and predictor show comparable contrast with this proxy method. Right: Measuring PSF stability with the strength of singular values. We compare the singular values (i.e., principal components) of the two populations of images from the predictor and integrator. The data taken while the predictor was running shows 3% less strength than the integrator in singular values, implying that there is less noise and easier paths to data reduction for the predictor images.

4.9 HAKA + predictive wavefront control

At the time of this publication, Keck will have underdone an AO upgrade to use HAKA, a higher count actuator deformable mirror that will increase the current degrees of freedom in the Keck correction by up to a factor of 10.

With fitting error currently dominating the Keck AO error budget, after this upgrade temporal errors induced by control lag time will remain as one of the larger error terms. This work presents an opportunity to mitigate that error term with predictive control and ultimately deliver better performance for the next generation of Keck AO. Figure 4.14 demonstrates this by simulating the impact on SR of the measured reduction in RMS wavefront error for HAKA feeding the visible light ORKID instrument. Using empirically tuned analytic models from Chapter 3.3 we predict a performance improvement of 2 % SR on bright stars. Note that this only explores a consistent AO rate of 1kHz.

However, HAKA's many actuators increase the size of the filter matrix, and threfore the training data needed for EOF. Work is still needed to ensure filter computation for a HAKA-sized filter matrix can happen in the few-minute timescale on which we expect atmosphere will stay consistent. Potential use of multiprocessing or GPU enabled filter computation, as well as strategic use of most significant singular values are potential solutions, but further work is needed to investigate potential performance loss from matrix inversion tools with less precision.

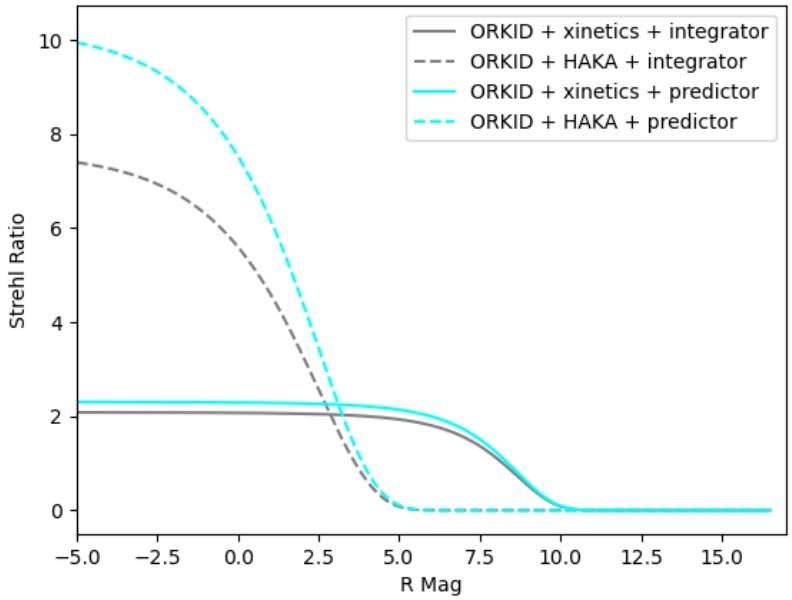


Figure 4.14: In visible light, a few nm of wavefront has a larger impact. Using performance models of Keck/ORKID (the visible light Keck demonstrator instrument) we project the impact of the measured wavefront error improvement. Solid lines indicate the estimated performance using the xinetics deformable mirror and dashed lines show the predicted HAKA performance.

4.10 Conclusion

In conclusion, we present practical considerations of implementing predictive wavefront control on sky. In simulation we explore optimal hyper-parameters for the training data and history vector length of the predictive filter. We find that in simulation, optimal filter hyper-parameters vary based on the atmospheric turbulence, and expect that to reach the optimal predictor performance we will need to optimize on-sky. We explore simulated annealing as a way to probe filter hyper-parameters more efficiently, and find it can repeatably converge to

within 0.1 nm RMS of the optimal solution within 10 iterations.

We also present lab verification and on-sky results of empirical orthogonal functions (EOF) as a linear data-driven predictive wavefront controller built into the Keck-II RTC. On-sky engineering tests of the predictive controller show a 20% performance improvement over a classic integrator according to wavefront residuals from the Shack-Hartmann Wavefront Sensor (SHWFS). Parameter optimzation studies show that there is a clear improvement based on varying predictive filter hyper-parameters, but that within a reasonable regime, varying filter parameters does not degrade performance to notably worse than an integrator. Practically, this means that while hyper-parameter optimzation does improve performance, parameters could be "set and forget" for ease of operation and this would not have an adverse effect on control. We were not able to test simulated annealing as a optimal filter-finding strategy on-sky, but this method still shows promise for future start of night operations.

NIRC2 images through the Brackett γ=2190 nm filter show comparable performance between an integrator and predictor, both comparing Strehl Ratio (SR) and a proxy measurement for contrast. We note that this is not similar to the up to 3x improvement in contrast in mid-spatial frequencies seen in Ref. [280], however the proxy we present for contrast does not include use of the vortex coronagraph, which is highly sensitive to wavefront aberrations [92]. We

also explore power in principal components, and find a modest improvement (on the order of 3% less area under the curve of component strength) from the predictor over the integrator. We expect that the Keck-II AO system at the time of these data was dominated by fitting error, as well as segment cophasing errors and non-common path aberrations. The impact of the latter two terms is demonstrated by the 10% difference in SRs between the October 19 night and the December 3 night, despite similar RMS wavefront error recorded by the SHWFS.

With similar performance in wavefront error between the two nights (both seeing ~ 50 nm RMS), it is also important to understand the full noise dynamics of the system and the true state of atmospheric turbulence. For example, predictive control performance gains may be limited by model mismatch (as discussed in Chapter 5), or we may be hitting the stability floor of the Keck-II AO system for a star of this magnitude. Future work will make a direct comparison to the noise floor of (1) the pseduo-open loop reconstruction of atmospheric turbulence and (2) the noise floor of the SHWFS to put these results in context.

With HAKA on-sky at the time of this publication, and fitting error no longer the dominant term, temporal errors are predicted to be the dominant understood term (after the margin error term) in the Keck AO system, meaning

that this improvement will be easier to estimate. Furthermore, visible light instruments are significantly more sensitive to wavefront errors, and empirically tuned analytic models of the visible light ORKID instrument predict an improvement in SR from 7% (performance without prediction) to 9% (performance with prediction) for bright natural guide stars given the wavefront error improvement demonstrated here.

Future work will explore the improvement of prediction at slower control frequencies. For faint guide stars, long exposures will be beneficial for reducing measurement error, but will incur greater temporal errors as the system lag time grows. A predictor could overcome this gap by predicting the state of the atmosphere after longer time lags, and will likely be the most beneficial for fainter guide stars.

More on-sky demonstrations of novel controllers on large telescopes are necessary to bridge the gap between simulations, injected turbulence in a lab setting, and true on-sky turbulence and full telescope system complexity at a realistic outerscale. SAXO+, the upgrade to the SAXO AO system for the Very Large Telescope/SPHERE will contribute to this goal by testing 6 predictors on the 8 meter VLT [68]. Until that time, Keck remains a one-of-a-kind testbed for the impact of predictive control on key science metrics at 8-10 meter class telescopes.

Chapter 5

Ground Control to Major Time-lag: Comparing predictive wavefront control methods for the next generation of extreme adaptive optics

5.1 Abstract

This chapter explores predictive wavefront control methods side-by-side-by-side. I start with empirical orthogonal functions (EOF). This is a linear data-driven controller that runs in pseudo-open loop, and the only method to have been implemented for full control of an 8-10 meter telescope. I compare EOF to an implementation that (1) uses the same mathematical methods but runs in closed loop (CLEOF) and (2) uses a physics-based equation of state to inform a model-driven controller, predictive Fourier control (PFC). I compare these predictors by applying them as post-facto methods to simulated atmospheres and on-sky telemetry. I find that (a not-yet fully optimized) CLEOF does not match the performance of EOF on perfect simulated data, but is more robust to poor internal models on data simulated with a model error in the deformable mirror command to wavefront sensor slopes. I find that EOF out-performs PFC on perfect simulated data, but that PFC shows improvements over EOF when applied to telemetry. All of the predictive methods show a notable improvement over an integral controller.

5.2 Introduction

Dessenne [58] wrote one of the first papers outlining predictive control for adaptive optics (AO) systems. This predictive method was demonstrated on-sky in 1999 [57] at the 1.52 meter observatory in Haute Provence; they cite a relative Strehl Ratio increase over a classic integrator of 30% (a max Strehl Ratio performance of $\sim$ 14%) at their visible central wavelength of 650 nm. Since then, a plethora of predictive control methods have come into the literature, including linear estimators [93, 95, 116, 280], linear quadratic Gaussian controllers [202, 255, 219, 211], model-based updates to a linear quadratic Gaussian [218, 217, 72], subspace control methods [99], and non-linear neural network solvers [152, 297, 146, 266, 192, 97]. Despite nearly 30 years of predictive methods, Empirical Orthogonal Functions (EOF) [93] is the only method to have been run on-sky as the controller for all spatial frequencies on an 8-10 meter class telescope; demonstrated on Subaru/SCExAO [95] and on Keck/NIRC2 [280].

As opposed to more classic methods (e.g., an integral controller or a linear quadratic Gaussian controller), EOF [93] learns a linear relationship in the evolution of the full scale of input turbulence, which means that EOF trains on (and predicts) the open loop wavefront. This is in contrast to the way a an AO system typically runs, where a deformable mirror corrects turbulence before it

is sensed, and a wavefront sensor never sees the full state of turbulence, only the residuals after a correction, or a closed-loop system. Because open loop systems are so rare, especially in the context of single conjugate AO, in practice these methods train on pseudo-open loop (POL) data, where deformable mirror (DM) commands are added to wavefront sensor (WFS) measurements to reconstruct the full state of turbulence. However, this leaves room for error in the POL reconstruction. Any spatially or temporally evolving mismatches in the calibration (e.g., non-linearities in the sensing and actuator pokes or DM to WFS misregistrations) that are not as impactful in an integral closed-loop controller could lead to poor prediction in open loop.

However, upon closer inspection of Dessenne's adaptive predictor [58] and Guyon's EOF [93], it becomes clear that they are both methods that predict from the same pieces of information. The major departure between the two methods is that Guyon's EOF runs in open loop, and reconstructs full turbulence states, and Dessenne's predictive controller runs in closed-loop, and encapsulates the evolution of both the DM commands and WFS measurements.

Similarly, one could argue that while the simplicity of a purely data-driven method like EOF makes the mathematical formalism extremely easy to understand and implement, but leaves EOF vulnerable to issues of overtraining, potentially over-correcting errors that are unphysical. The alternative is a

model-based controller, for example, predictive Fourier control (PFC). Based on a starting assumption of a frozen flow atmosphere, PFC uses a state space model of an atmosphere with an unknown number of wind layers, uses data to assign these layers a speed and direction across each spatial frequency, and ultimately only corrects for errors that fit the form of this model.

With this work, we focus on comparing predictive control methods across two axes: (1) whether they use closed or (pseudo-)open loop data to train their prediction and (2) whether they explicitly use a physical model to constrain the controller. In particular we explore empirical orthogonal functions [93] (EOF), as a pseduo-open loop purely data-driven method, a closed-loop implementation of empirical orthogonal functions (CLEOF) [56], and predictive Fourier control [218] (PFC), a model-driven method. As three solutions to the same problem, the performance of EOF, CLEOF, and PFC have yet to be directly compared. With this chapter we will study their performance over simulated atmospheric data and saved on-sky telemetry from the W.M. Keck Observatory.

In Section 5.3, we describe the pseudo-open loop EOF method. In Section 5.4 we describe a closed loop update and present the results of preliminary performance simulations and the impact of model-mismatch errors on that performance on the two methods side-by-side. In Section 5.5, we introduce predictive Fourier control and examine the comparative performance of the

two predictors.

5.3 Brief summary of empirical orthogonal functions

A detailed description of Empirical Orthogonal Functions (EOF) is available in Chapter 1.5.2. Chapter 4 explores on-sky integration of EOF at Keck Observatory, and also includes a summary of the math and methods. Here I present an even briefer summary for the purposes of this chapter.

By adding wavefront sensor residuals and deformable mirror commands, we create pseudo-open loop (POL) records of the wavefront and collect **n** POL states in a history vector $\vec{\mathbf{h}}$:

$$\mathbf{h}_m(t) = \begin{bmatrix} w(t) \\ w(t-dt) \\ w(t-(n-1)dt) \end{bmatrix} \tag{5.1}$$

which are used in real time to calculate a wavefront correction of the form:

$$\text{prediction} = \mathbf{F} \times \vec{\mathbf{h}} \tag{5.2}$$

where $\mathbf{F}$ is a filter matrix trained using up to a 1.5 minutes of data.

To calculate F we minimize the difference between $\mathbf{l}$ frames of training data $\mathbf{D}$ and a truth condition $\mathbf{P}$ (data that is one time lag ahead of $\mathbf{D}$)

$$\left|\left|\mathbf{D}^T\mathbf{F}^T - \mathbf{P}^T\right|\right|^2 \tag{5.3}$$

and calculate $\mathbf{F}$ with a least squares inversion

$$\mathbf{F} = ((\mathbf{D}^T) + \mathbf{P}^T)^T = \mathbf{P}\mathbf{D}^T(\mathbf{D}\mathbf{D}^T + \alpha\mathbf{I})^{-1} \tag{5.4}$$

where α is a regularization factor, akin to setting the number of singular values. For simulations we run with $\alpha = 1$, 60000 frames of training data, and history vectors made up of 3 exposures (apart from the 7 m/s layer, see further discussion in Section B.1.5), a result also found by Jensen-Clem (2019) [116]. We use these parameters for our prediction performance estimates in the rest of this chapter.

5.4 Comparing open and closed-loop predictive wavefront control

5.4.1 Our development of a closed-loop empirical orthogonal functions update

Inspired by Dessenne, 1998 [58] and Haffert, 2021 [99], we explore a closed-loop reformulation of the classic empirical orthogonal functions [93]. The fullest realization of this method could:

1. Improve controller stability by running in closed-loop.

2. Avoid error introduced by non-linear wavefront sensing or DM model mismatch when wavefront sensor residuals are converted to psudeo open loop.

3. Track drifts that may impact the DM and wavefront sensor separately by allowing each to evolve with a different set of coefficients.

4. Do robust timekeeping with information into the system.

Following [58], an updated history vector (we refer to this as ϕ to distinguish from the open loop history vector, as is standard in control theory conventions), contains both wavefront sensor residuals $\epsilon(t)$ and DM commands $y(t)$. (While

the original work opted to use KL-modes, we work in the DM zonal basis, i.e., one point of information per DM actuator/WFS subaperture.) We consider only a single mode at a time:

$$\phi(n) = \begin{bmatrix} y(n-1) \\ y(n-2) \\ \dots \\ y(n-p) \\ \epsilon(n-2) \\ \epsilon(n-3) \\ \dots \\ \epsilon(n-p-1) \end{bmatrix} \tag{5.5}$$

In this way, the new history vector and the corresponding predictive filter have twice as many values per mode pair as previous open loop derivations. However, the output of the filter applied to the history vector $\theta \cdot \phi$ predicts the same piece of information as an open loop implementation: the full turbulence in DM space, which is the DM command needed at a given iteration, in our notation $y(n)$.

Dessenne's original formulation suggests reconstructing open loop turbulence from the full transfer function of the AO control loop to train the predictive

filter, and solving for a steady state recursive least squares solution [58]. In this simulation, we instead apply the same minimization technique used in the open loop implementation. Future work, inspired by that of van Kooten, 2019 [279] will explore a recursive least square implementation that updates the filter with each control iteration. We build a new minimization problem:

$$\min_\theta ||\mathbf{D}^T\theta^T - \mathbf{P}^T||^2 \tag{5.6}$$

where $\mathbf{D}$ contains collections of the history vector ϕ, and $\mathbf{P}$ contains the future state of the full turbulence.

In practice, the distinction between the open and closed-loop implementation is twofold: (1) the ability of the correction to apply different coefficients for the DM commands and the wavefront sensor information (whereas POL is one command applied to a summed value) and (2) the ability to robustly encapsulate a time delay into the reconstruction. POL reconstruction adds wavefront sensor measurements to DM commands with a single static delay, while closed-loop methods can account for when each piece of information enters the system (with two-step delays or delays with non-uniform steps). See Section 5.4.2 for additional information on how timekeeping impacts the construction of ϕ and $\mathbf{P}$.

Assuming a classic two-step delay and p associated frames of information in a history vector, the weighted estimate for a single coefficient for the closed-loop formulation (y_{CL}) vs. the original open loop formulation (y_{OL}) takes the form:

$$y_{CL}(n) = b_1 y(n-1) + ... + b_p y(n-p) + a_0 \epsilon(n-2) + ... + a_{p-1} \epsilon(n-p) \quad (5.7)$$

$$y_{OL}(n) = c_0[y(n-1) + \epsilon(n-1)] + ... + c_p[y(n-p-1) + \epsilon(n-p-1)] \quad (5.8)$$

5.4.2 Verifying the time flow of information through a control system

One intriguing issue when comparing and reworking control methods is a robust understanding of time delays over the course of a control loop and how that plays out with different methods. Pseudo-open loop implementations gloss over this issue, by recreating a single point of information at a single point in time, but closed-loop implementations provide the opportunity to represent each piece of information at the time it enters the system, some work even accounts for fractional time-delays [217]. The following is a minimal proof of information flow in a control loop, switching between time reference frames.

We consider this system in two reference frames: (1) the time represented by the physical information flowing through the system (i.e., what time the

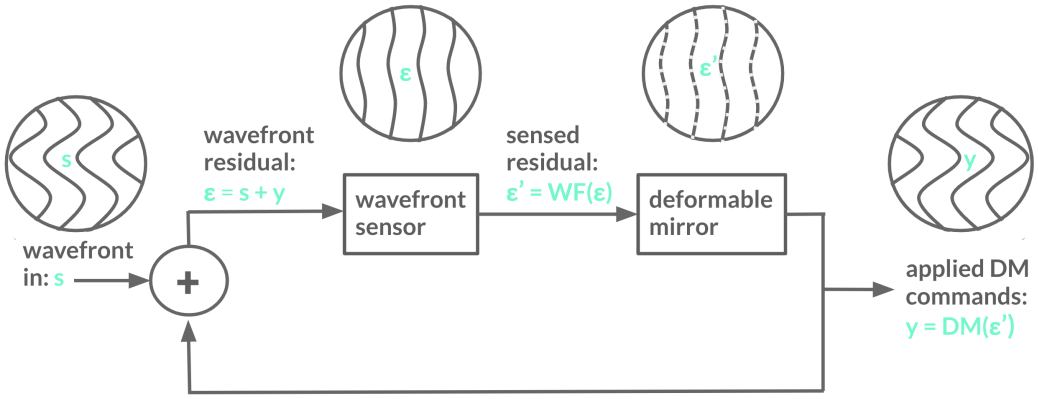


Figure 5.1: Control diagram of a classic two-step delay. The wavefront (s) comes into the system, and first meets a sum junction, where we apply deformable mirror (DM) commands (y), making it a closed-loop system. At this point, our residual is $\epsilon = s + y$. That signal is sensed by a wavefront sensor; that sensed signal is $\epsilon' = \mathrm{WF}(\epsilon)$, where WF represents a functional form of how the wavefront sensor interprets the signal. Finally, this is fed into a computer that will calculate and apply a correction $y = DM(\epsilon')$, where DM represents a functional form of how that correction is applied and calculated based on ϵ'.

turbulence the system is analyzing occurred); this is physically intuitive for timekeeping and residual comparison and we call it a physical clock and (2) the time represented by the control system and when information could be sampled from various sensors or correctors, which we call a control clock.

Dessenne [58], uses a control clock framework. If, for example, we wanted to associate 3 states in time, we would build a history vector (used to predict

the state at iteration n) of the form:

$$\phi(n) = \begin{bmatrix} y(n-1) \\ y(n-2) \\ y(n-3) \\ \epsilon'(n-2) \\ \epsilon'(n-3) \\ \epsilon'(n-4) \end{bmatrix} \tag{5.9}$$

We show that these time steps from the perspective of the control clock associate logical pieces of information from the perspective of a physical clock. If we consider information moving through a system with a two step delay, we would see the chain of control events outlined in Table 5.1

If we now consider the history vector ϕ, we can rewrite it to prediction

Table 5.1: How information moves through the control loop, starting from the loop turning on. With a step of delay between each component, we at first have no information sensed and turned into a DM command.

s	ϵ	ϵ'	y
s_1	$\epsilon_1 = s_1$	-	-
s_2	$\epsilon_2 = s_2$	$\epsilon'_2 = \mathrm{WF}(\epsilon_1)$	-
s_3	$\epsilon_3 = s_3 + y_3$	$\epsilon'_3 = \mathrm{WF}(\epsilon_2)$	$y_3 = \mathrm{DM}(\epsilon'_2)$
s_4	$\epsilon_4 = s_4 + y_4$	$\epsilon'_4 = \mathrm{WF}(\epsilon_3)$	$y_4 = \mathrm{DM}(\epsilon'_3)$
s_5	$\epsilon_5 = s_5 + y_5$	$\epsilon'_5 = \mathrm{WF}(\epsilon_4)$	$y_5 = \mathrm{DM}(\epsilon'_4)$
s_6	$\epsilon_6 = s_6 + y_6$	$\epsilon'_6 = \mathrm{WF}(\epsilon_5)$	$y_6 = \mathrm{DM}(\epsilon'_5)$
s_7	$\epsilon_7 = s_7 + y_7$	$\epsilon'_7 = \mathrm{WF}(\epsilon_6)$	$y_7 = \mathrm{DM}(\epsilon'_6)$
s_8	$\epsilon_8 = s_8 + y_8$	$\epsilon'_8 = \mathrm{WF}(\epsilon_7)$	$y_8 = \mathrm{DM}(\epsilon'_7)$

some iteration $n = 7$

$$\phi(7) = \begin{bmatrix} y(n-1) = y_6 = \mathrm{DM}(\epsilon'_5) = \mathrm{DM}(\mathrm{WF}(\epsilon_4)) \\ y(n-2) = y_5 = \mathrm{DM}(\epsilon'_4) = \mathrm{DM}(\mathrm{WF}(\epsilon_3)) \\ y(n-3) = y_4 = \mathrm{DM}(\epsilon'_3) = \mathrm{DM}(\mathrm{WF}(\epsilon_2)) \\ \epsilon'(n-2) = \epsilon'_5 = \mathrm{WF}(\epsilon_4) \\ \epsilon'(n-3) = \epsilon'_4 = \mathrm{WF}(\epsilon_3) \\ \epsilon'(n-4) = \epsilon'_3 = \mathrm{WF}(\epsilon_2) \end{bmatrix} \tag{5.10}$$

Notice that though the control clock indexing appears to introduce a one step offset, the most expanded part of the expression for each iteration shows that the elements on the history vector depend on the the same physically timed

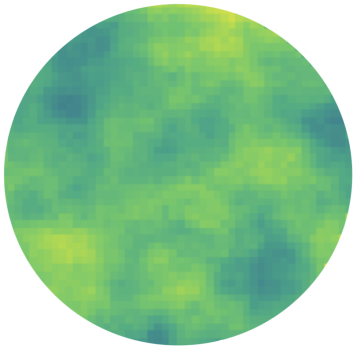
pieces of information.

It should be noted that a classic two-step delay is not actually physically reminiscent of most systems. In Ref. [32] they found that for the Keck pyramid WFS RTC, the time delay over the entire correction (including wavefront sensing time, calculations, and DM latency after correction is applied) takes ∼1.5 ms, where ∼1 ms is the wavefront sensing and readout time, and ∼0.5 ms is the calculation and hardware latency time. For this simulation work we have opted to only update the DM once per new piece of wavefront sensor information, essentially forcing this to be a classic two-step delay. However, future work will explore the efficacy of updating the DM more frequently, as we could project forward the correction easily and provide two DM updates per wavefront sensor readout.

5.4.3 Selecting a truth condition for training data

The truth condition $\mathbf{P}$, against which the predictive filter trains, must be the full uncorrected state of the wavefront (i.e., open loop turbulence) at a given iteration. The original work from Dessenne [58] estimated the open loop wavefront from the transfer function of the AO control loop, essentially a higher fidelity pseudo open loop reconstruction than that of Guyon, 2017 [93]. For these simulations, we give the predictor perfect knowledge, using the full

scale of turbulence from simulation. We note that this training condition is not realistic to on-sky operation, but acts as a laboratory to test the perfect performance of a closed-loop implementation, and future work will explore more realistic $\mathbf{P}$ generation. Figure 5.2 shows the example data used as the truth condition for training the closed-loop implementation.



simulation parameter	value
r_0	20 cm
wavelength	658 nm
wind speed	7 m/s
telescope diameter	8 m
modes / DM actuators	400
loop speed	1 kHz
correction lag-time	2 ms
associated frames in the history vector	3
training data length	60000

Figure 5.2: Left: Piston, tip, and tilt subtracted simulated phase screens from an 8 meter telescope with a single wind layer. Right: Table of parameters for generating the turbulence and AO simulations.

5.4.4 Preliminary implementation of closed-loop predictive wavefront control on simulated data

We simulate atmospheric phase screens and an idealized AO control system with `HCIPy` [213]. We estimate the root-mean-square (RMS) residual error across the pupil for 1) the full state of the uncorrected turbulence, 2) a pseduo-integrator, in which we apply a perfect correction with a 2 ms time delay

(for Keck, the time-delay is ∼1.5 ms [32]), 3) pseduo-open loop EOF (well optimized in previous work [116, 72]), and 4) a preliminary implementation of CLEOF. These simulations use perfect wavefront sensing and correction, and are meant to provide comparative estimates only for the bandwidth error. The parameters for the simulation are shown in Figure 5.2.

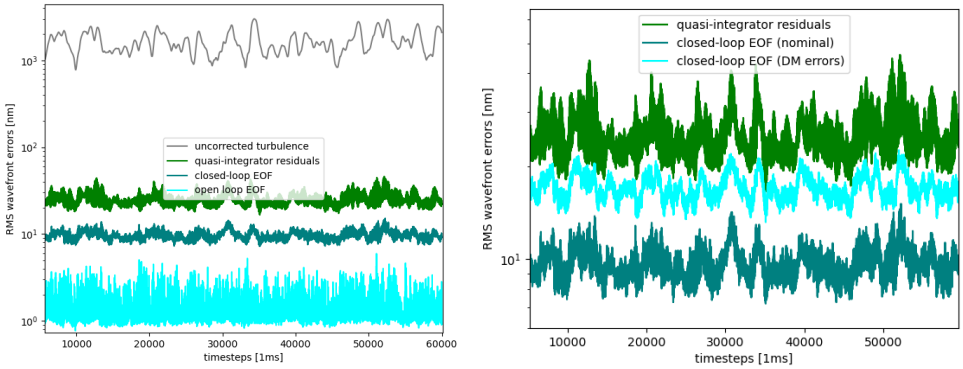


Figure 5.3: Left: Two predictive methods alongside a psuedo-integrator (a perfect correction 2 timesteps behind), all applied with no measurement or fitting error (i.e., perfect wavefront sensing and control) running at 1 kHz. Full uncorrected turbulence RMS median: 1534.21 nm, quasi-integrator: 24.30 nm, closed-loop predictor: 9.49 nm, open loop predictor: 1.03 nm. Right: Impacts of a DM-model mismatch. Closed-loop EOF correction goes from a median of 9.49 (for the nominal run) to 16.66 nm RMS when a DM model error is introduced, but still shows better performance than the pseudo-integrator. The open loop EOF error increases by an order of magnitude and is not displayed on the figure for scale purposes.

Performance in an idealized AO system

Figure 5.3 (left) shows the results of our initial implementation; we find a factor of improvement in RMS error of ~ 2.5 over a standard integrator from CLEOF,

under-performing compared to the standard psuedo open loop. In practice, we should be able to recreate similar (if not better) performance between a closed and open loop implementation, however closed-loop minimization problem contains twice as many variables (treating the WFS and DM independently), twice as many regressors, and will require more training data to converge to an optimal solution, as shown in van Kooten, 2019 [279]. Future work will explore optimizing filter length and training length for a closed-loop implementation for a more fair comparison with open loop.

Robustness to model errors

Figure 5.3 (right) shows the performance of both predictors on a second set of data, in which we simulate a DM model mismatch in our control system – adding a static factor of 2 between what the DM expects to apply and actually applies. We see that the performance of closed-loop EOF is less sensitive to errors in the system model. While the closed-loop performance drops when a DM model issue is introduced, the closed-loop predictor still outperforms a typical integrator. However, introducing the same issue into a open loop EOF increases the residual error by more than an order of magnitude. We note that in theory a closed-loop EOF should exactly learn and reconstruct the model-error, and future work to optimize the filter and training length will

likely make a closed-loop EOF even more robust to model errors.

5.5 Comparing data and model-driven predictive wavefront control

5.5.1 Overview of predictive Fourier control as a model-driven predictor

Predictive Fourier control (PFC) [218] uses data to extract information about wind layers from the system, and then builds a predictive filter informed a state space model that encapsulates those wind layers, as an update to a traditional Linear Quadratic Gaussian controller [203, 139]. Here we will briefly summarize the math behind building and applying the PFC predictive filter; the mathematical formalism is taken from Poyneer (2007) [218] unless specified.

In short, PFC consists of 4 main steps:

1. Decompose turbulent phase screens into complex spatial Fourier modes. Each Fourier mode becomes a separate control problem and a filter is calculated individually for each coefficient. For each coefficient:

2. Make a power spectral density (PSD) function by looking at how a single coefficient evolves in temporal frequency space.

3. Isolate and fit peaks in the PSDs (which map to wind layers in an atmosphere).

4. Feed the coefficients of each fit to a Kalman filter.

From the perspective of a telescope system (under a Frozen Flow assumption [268]), wind layers will present as static turbulent scenes crossing over the telescope primary mirror at the speed of the wind at that height. Given enough data in time to adequately sample a wind velocity, these pupil crossing times will have high signal to noise peaks across spatial Fourier modes, as different spatial scales will cross with different frequencies. Therefore, when turbulent phase screens are decomposed into Fourier modes, each coefficient contains information on a spatial frequency of a particular size that maps to a series of non-conflicting controllers. A PSD for each coefficient reveals the peak frequency, height, and width for each wind layer, which provides input to a controller specific to that coefficient.

At a given Fourier mode indexed (k, l) (over wavefront sensor subapertures), a peak at some ν indicates a velocity (v_x, v_y) component of:

$$\nu = -\frac{(kv_x + lv_y)}{D} \tag{5.11}$$

where D is the telescope diameter. While in practice, the algorithm is not

meant to deliver identifications of wind-layers in a system, testing the wind identification on simulated known atmospheres shows recovery of both single and multi-layer wind layers. Figure 5.4 shows examples of successful recovery of rejected wind layers. See Figure 5.6 for some exploration of less optimal peak recovery and a discussion of potential causes.

Peaks can be identified and fit using built-in `scipy` tools. Also present in each PSD is a peak that occurs at $\nu = 0$, due to a 0 Hz frequency peak (DC peak) from the Fourier transform that builds the PSD [218]. DC noise encapsulates the static offset in a system, which may be caused by consistent noise from electric components in a realistic system, or mean offset from zero for any distribution. Wind velocity peaks that occur near zero (in frequency space) are difficult to resolve, as they blend into the static DC peak.

For each peak, a height σ^2, and placement (i.e., peak width and location in angular frequency space) α, are fit with

$$P(\omega) = \frac{\sigma^2}{|1 - \alpha \exp(-i\omega)|^2} \tag{5.12}$$

The complex term α is expressed $\alpha = |\alpha| e^{i\omega}$, where ω encapsulates the actual wind velocity of each layer. (Without the complex phase, α acts as a gain for a typical Linear Quadratic Gaussian controller without extracting and predicting

based on wind layers.)

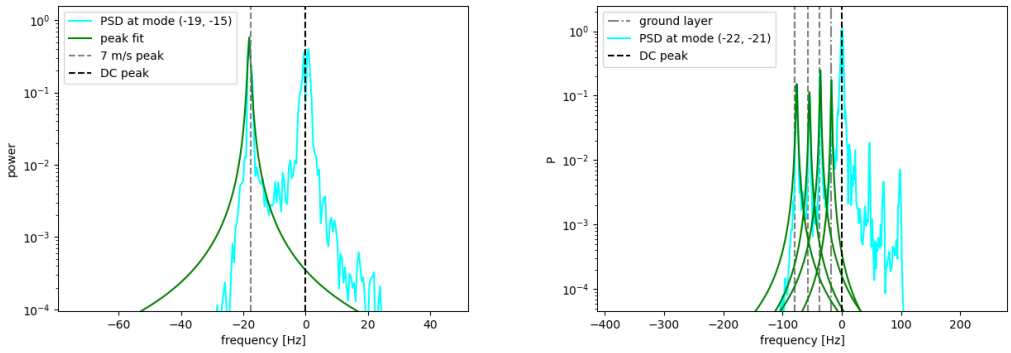


Figure 5.4: Automatic peak identification and fitting with `scipy` tools. Left: Peak identification (for a single layer atmosphere) at the (-19, -15) mode. The 7 m/s peak is correctly identified, along with the DC peak at 0. Right: Peak identification (for a multi-layer atmosphere) at the (-22, -21) mode. All four wind layers are recovered.

The fits for each α and σ inform a Kalman filter, which in turn provides a prediction of each Fourier coefficient. We predict the state vector of our system

$$\mathbf{x}[t] = \begin{bmatrix} \mathbf{a}(t) & \phi(t+1) & \phi(t) & \phi(t-1) & d(t-1) & d(t-2) \end{bmatrix} \tag{5.13}$$

where $\mathbf{a}$ encapsulates the fit parameters, ϕ are the wavefront states, and d are commands to the deformable mirror. We predict the next iteration with:

$$\mathbf{x}(t) = (\mathbf{I} - \mathbf{KC})\mathbf{A}\mathbf{x}(t-1) + (\mathbf{I} - \mathbf{KC})\mathbf{G}d(t-1) + \mathbf{K}y(t) \tag{5.14}$$

where **K** are the Kalman gains, **C** is a control matrix, **A** is a covariance matrix, **G** is a matrix to update DM commands, $d(t)$ holds DM commands, and $y(t)$ holds noisy wavefront sensor measurements. Note that in our implementation, for ease of comparison with EOF, we rewrite the original control law so that we can apply it in open loop.

We found that using 10000 frames of data to build our PSDs and windows between 1024-2048 exposures provided a consistent correction (see Section 5.5.2 for further description of the windowing and its impacts), and used these parameters for our prediction performance estimates.

5.5.2 Predictive Fourier control math in detail

While in the last section we provided a digestible overview, here we fully describe the mathematical formalism for Predictive Fourier Control [218], as well as how we have implemented it, and any deviations from the original method. For consistency across EOF and PFC, all matrices are defined in the convention of m rows by n columns, in departure from typical control engineering conventions and the original formulation from Poyneer (2007) [218].

Decomposition into complex fourier modes

We decompose each turbulent scene into complex Fourier modes. Given some mode (k, l), over some x by y pixel grid, each mode is defined by

$$f_{k,l}[x, y] = \cos\left(2\pi\frac{kx + ly}{N}\right) + i\sin\left(2\pi\frac{kx + ly}{N}\right) \tag{5.15}$$

where N is the total number of modes. (Note that the original paper from Poyneer (2007) [218] contains a typo in its casting of each cos/sin term as complex, which we do not duplicate.)

Fourier modes are meant to operate over square regions, and many telescope pupils are circular; both our Keck telemetry and simulations are over a circular pupil. Starting with a circular pupil embedded within a square grid, and applying a mask for the non-used elements to both the telescope pupil and the Fourier mode allows for a decomposition of the original turbulence image at near machine precision as mathematically proved in Poyneer (2005) [220], and experimentally reproduced in this work.

Figure 5.5 shows an example of a residual turbulent screen, a reconstruction of the same turbulent screen built with complex Fourier modes, the real component of a complex Fourier mode projected onto the same pupil, and the error introduced during this decomposition, on the order of 1×10^{-12},

as compared to the phase of the original turbulence which has a median of ~ 850nm (the ratio of which approaches the machine precision of Python.)

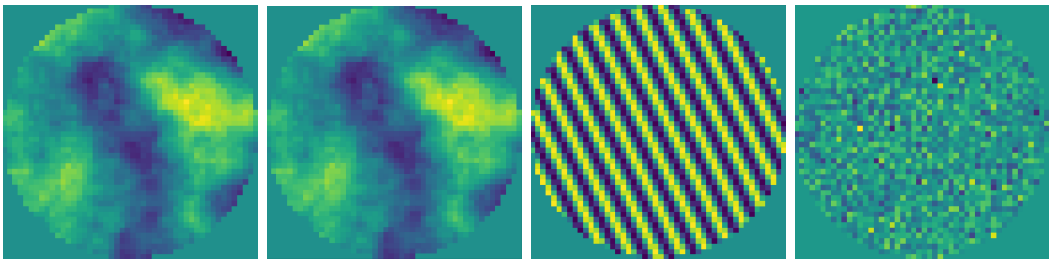


Figure 5.5: Left: Simulated turbulent screen. We start with a circular pupil, and mask the edges within the square. Center left: Reconstruction of the original turbulent scene with the complex Fourier modes, projected into the circular pupil. Center right: Real values of an example complex Fourier mode projected over the same pupil. Right: Residuals between the reconstructed turbulence and the original scene. Residuals are on the order of 1×10^{-12}.

Creating power spectral densities and identifying peaks

To examine the coefficient for each Fourier mode, we create a periodogram in temporal frequency space. We use a Welch periodogram [197] with a Hamming window, using `scipy.signal.welch` and `scipy.signal.get_window`.

We use windows of $n = 1024$ or $n = 2048$, depending on the noise present in the data. Longer windows provide the ability to resolve peaks further from each other and from the 0Hz DC peak, but increase noise in the periodogram. For our smallest window, $n = 1024$ frames, our minimum resolvable frequency would be $\nu_{min} = 1024 \times t_{int} = 2, 0.6$Hz, for our 0.5,1.7 ms resolutions in

simulated data and telemetry respectively. We do not record a peak as input to the Kalman filter unless it is ν_{min} from 0, as well as other identified peaks.

We find the location of each peak in the data with `scipy.signal.find_peaks`, and fit it with `scipy.optimize.curve_fit`. Peak identification appears repeatable over single layer atmospheres (as shown in Table 5.2), though Figure 5.6 shows some examples of sub-optimal peak identification and fitting given known wind layers, which may be due to spectral leakage or scalloping [252].

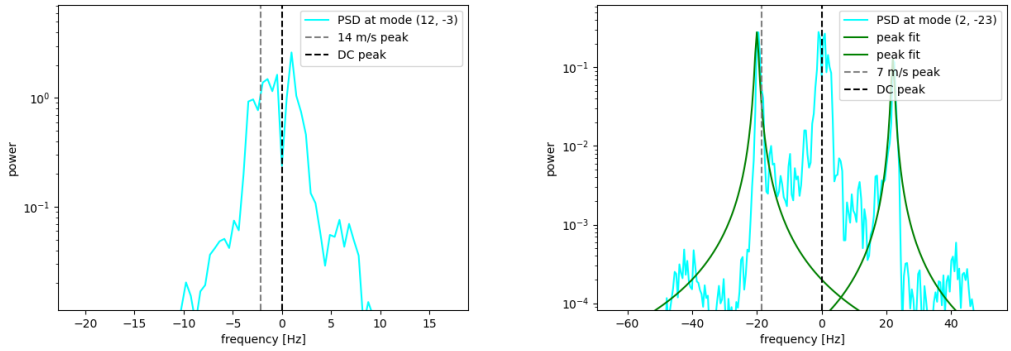


Figure 5.6: Left: With the expected wind layer for the (12, -3) coefficient quite close to the DC layer at 0, we are unable to recover the peak from this mode. Right: Extraneous peak identification at the (-2, -23) mode, as well as the expected 7 m/s peak, likely due to aliasing. Dashed lines in gray and black represent the expected peak placement from the 7 m/s layer and the DC layer respectively.

Fitting α and σ

Each peak can be defined with the power law:

Table 5.2: Successful (within 2 Hz) peak identifications for varying atmospheric profiles.

Identified Layer	Correct Peak IDs	Total Modes
7 m/s	1909	2304
14 m/s	1989	2304
telemetry ground layer	179	484

$$P(\omega) = \frac{\sigma^2}{|1 - \alpha e^{(-i\omega)}|^2} \tag{5.16}$$

where σ^2 describes the peak height for a given coefficient and complex $\alpha = |\alpha|e^{i\omega}$ describes the placement (i.e. peak width and velocity) of a given wind layer. ω, or phase, describes angular frequency that goes with $\omega = -2\pi\nu t_{int}$, where t_{int} is the time between each frame of data and ν is a natural frequency space in Hz.

Building the state vector and covariance matrix

In keeping with the traditional Linear Quadratic Gaussian (LGQ) control, we use a Kalman filter to predict forward the state of the system. While the original Poyneer (2007) [218] paper describes a close-loop control law, we reformulate it to run in open loop, adjusting the covariance matrix **A**, as well as control matrix **C**. The original close loop law includes terms in **A** and **C**

to subtract off a deformable mirror (DM) command. Our update renders the DM commands in the state vector unnecessary for the Kalman filter (which could impact computational performance), but in this initial implementation we keep them for simplicity in following the original math. We denote these changes in each matrix with an*.

The state vector $\mathbf{x}$ is a vector of (L+6) elements that consists of

$$\mathbf{x}(t) = \begin{bmatrix} \mathbf{a}(t) & \phi(t+1) & \phi(t) & \phi(t-1) & d(t-1) & d(t-2) \end{bmatrix} \tag{5.17}$$

where $\phi(t)$ and $d(t)$ map to the wavefront coefficient at each Fourier mode, and the command to the deformable mirror respectively. The system operates with a two-step time delay, and is centered one step off from the deformable mirror and one step off from the wavefront sensor in opposite directions. At the time of some system measurement $y(t)$ taken at $t = t_0$, the wavefront is at time $t = t_0 + dt$, and the deformable mirror has information from time $t = t_0 - dt$; this is why we examine $\phi(t+1), \phi(t)$ as compared to $d(t-1), d(t-2)$.

$\mathbf{a}$ is a vector of (L+1) elements that encapsulates information on every identified wind layer at that coefficient.

$$\mathbf{a} = (a_0, a_1, ..., a_L) \tag{5.18}$$

$$a(t)_L = \alpha_L a(t-1) + w(t) \tag{5.19}$$

Each a_L evolves according to a first order auto-regressive process, AR(1) [252], driven by the α_L for a given wind layer and complex white noise $w(t)$. The periodogram for each Fourier mode includes a direct current (DC) layer peak at $\nu = 0$, for which we include $\alpha_0 = 0.999$ and $\sigma^2 = \max[P(\nu)]$, where $P(\nu)$ is the PSD. (Future implementations could fit the σ^2_{DC} term in a way that is more physically relevant.)

The state vector $\mathbf{x}$ evolves according to the state equation:

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{G}d(t) + \mathbf{B}\mathbf{w}(t) \tag{5.20}$$

with the (L+6) by (L+6) covariance matrix

$$\mathbf{A} = \begin{bmatrix} \mathbf{R} & \mathbf{0}_{L+1x1} & \mathbf{0}_{L+1x1} & \mathbf{0}_{L+1x1} & \mathbf{0}_{L+1x1} & \mathbf{0}_{L+1x1} \\ \mathbf{1}_{1xL+1} & 0 & 0 & 0 & 0 & 0 \\ \mathbf{0}_{1xL+1} & 1 & 0 & 0 & 0 & 0 \\ \mathbf{0}_{1xL+1} & 0 & 1 & 0 & 0 & 0 \\ \mathbf{0}_{1xL+1} & 0 & 0 & 0 & 0 & 0 \\ \mathbf{0}_{1xL+1} & 0 & 0 & 0 & 0^{*} & 0 \end{bmatrix} \tag{5.21}$$

where (L+1) by (L+1) $\mathbf{R}$ holds the α_L values on the diagonals such that

$$\mathbf{R} = \text{Diag}(\alpha_0, \alpha_1, ..., \alpha_L) \tag{5.22}$$

The 1 by (L+6) DM update matrix is defined as

$$\mathbf{G} = \begin{bmatrix} \mathbf{0}_{1xL+1} & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \tag{5.23}$$

Finally, an (L+6) by (L+1) noise propagator matrix

$$\mathbf{B} = \begin{bmatrix} \mathbf{I}_{L+1xL+1} \\ \\ \mathbf{0}_{5xL+1} \end{bmatrix} \tag{5.24}$$

The measurement in the system is defined with

$$y[t] = \mathbf{C}\mathbf{x}[t] + v[t] \tag{5.25}$$

where $v[t]$ is the measurement noise, which is assumed to be zero-mean Gaussian white noise, and 1 by (L+6) $\mathbf{C}$ is the control matrix

$$\mathbf{C} = \begin{bmatrix} \mathbf{0}_{1xL+1} & 0 & 0 & 1 & 0^* & 0 \end{bmatrix} \tag{5.26}$$

Steady state solutions with the algebraic Ricatti equation

Instead of recalculating the Kalman filter at each step – which is quite computationally expensive, we opt to use the Algebraic Ricatti Equation (ARE) [144], specifically with a discrete solver, as the covariance matrices will reach a steady-state solution. We use `scipy.linalg.solve_discrete_are` to calculate $\mathbf{P}_s$. Poyneer (2007) [218] outlines this equation as:

$$\mathbf{P_s} = \mathbf{A}\mathbf{P_s}\mathbf{A^H} + \mathbf{B}\mathbf{P_w}\mathbf{B^H} - \mathbf{A}\mathbf{P_w}\mathbf{C^H}(\mathbf{C}\mathbf{P_s}\mathbf{C^H} + \mathbf{P_v})^{-1}\mathbf{C}\mathbf{P_s}\mathbf{A^H} \tag{5.27}$$

with $\mathbf{P_v}$ as a scalar variance of the white noise distribution and (L+1) by (L+1)

$$\mathbf{P_w} = \text{Diag}(\sigma_0^2, \sigma_1^2, ..., \sigma_L^2) \tag{5.28}$$

We note that this use of the discrete ARE differs from the typical form in that it uses Hermitian conjugates of the known coefficient matrices $\mathbf{A} \rightarrow \mathbf{A}^H$ and $\mathbf{B} \rightarrow \mathbf{C}^H$, and builds the noise term with $\mathbf{Q} \rightarrow \mathbf{B}\mathbf{P_w}\mathbf{B}^H$ [144, 284].

Finally we calculate the Kalman gains (a vector of L+6 elements) for a given coefficient with:

$$\mathbf{K} = \mathbf{P_s}\mathbf{C^H}(\mathbf{C}\mathbf{P_s}\mathbf{C^H} + \mathbf{P_v})^{-1} \tag{5.29}$$

which builds all of the pieces for the update equation to predict the next step:

$$\mathbf{x}(t) = (\mathbf{I} - \mathbf{K}\mathbf{C})\mathbf{A}\mathbf{x}(t-1) + (\mathbf{I} - \mathbf{K}\mathbf{C})\mathbf{G}d(t-1) + \mathbf{K}y(t) \tag{5.30}$$

From the newly calculated state vector we pull the $L+1$st element as the prediction of the wavefront $\phi(t+1)$.

5.5.3 Optimized implementation of predictive wavefront controllers on simulated data and on-sky telemetry

Simulated single and multi-layer turbulence

Using `hcipy`, the High Contrast Imaging package for Python [213], we simulated three test case atmospheres: two with single layer atmospheres with a wind speed of 7 m/s (split into $v_x, v_y = 2.13, 6.67$) and 14 m/s (split into $v_x, v_y = 2, 13.5$), and one that included 8 layers according to a C_N^2 and wind velocity study of Maunakea weather [190]. The layers and heights from this profile are described in Table 5.3; all layers in the multi-layer simulation had a random direction assigned. For the purposes of this simulation we used an 8 meter primary mirror diameter, an $r_0 = 15$cm at 500nm, simulated at a resolution of 48 by 48 pixels (which maps to an idealized wavefront sensor of 48 by 48 pixels), and evolved with a 0.5ms time step. This simulates a Gemini Planet Imager-like system, which was the instrument PFC was originally designed for [218]. We consider an idealized wavefront sensor whose output is exactly the simulated 48 by 48 pixel phase screen at a given time step. I.e., we use a perfect wavefront sensor (there is no difference between the phase screen and the phase screen we sense) and a perfect deformable mirror (there is no difference between the phase screen we predict and the correction we apply)

and focus only on how faithfully we can predict the wavefront given a choice of time delay.

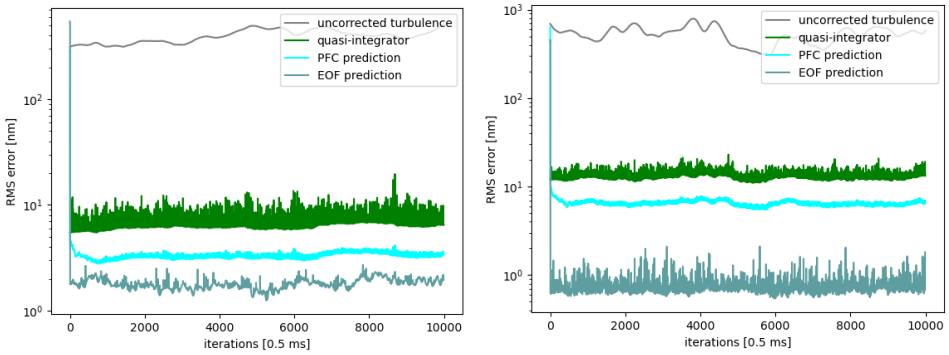


Figure 5.7: Root mean square residuals for the uncorrected turbulence, a quasi-integrator (the known state of the wavefront applied as a correction 2 steps behind), and the two prediction methods. Left: Single 7 m/s wind layer. Right: Single 14 m/s wind layer. Both simulated atmospheres show a performance improvement over a classic integrator from prediction, with EOF performing the best across both atmospheres.

We calculated residuals by subtracting the original simulated wavefront from the predicted wavefront in phase, and then calculate the root mean square (RMS) error at 500nm. For the sake of comparison, we also estimate the performance of a quasi-integrator, where we subtract the simulated phase from itself with a two-step delay (i.e., if the system could apply a perfect correction two steps behind.) For this initial implementation we found that edge effects could have a major impact on the residual error, so as an interim measure we calculated the RMS error only over the inner 7 meter diameter of phase correction. Figure 5.7 shows the results for the two predictors. For the 7 and

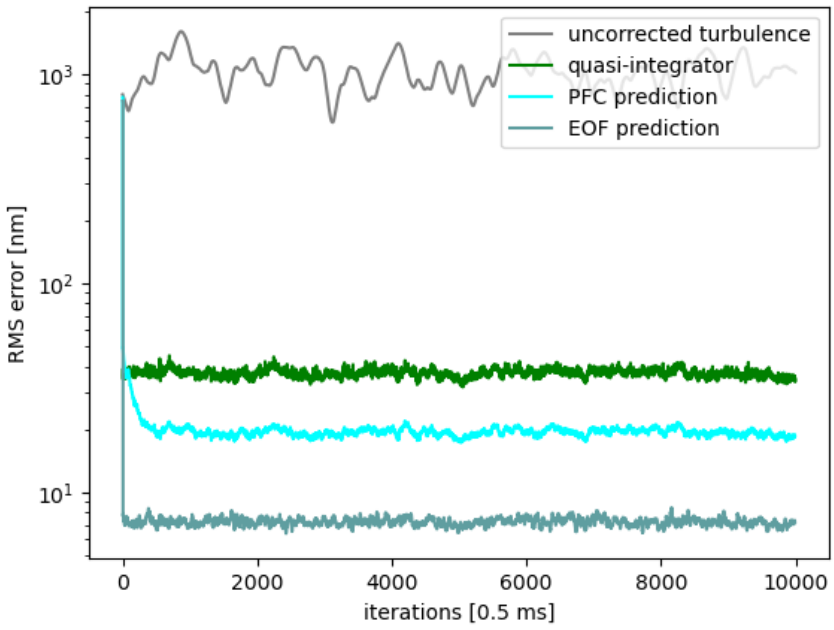


Figure 5.8: Uncorrected turbulence, quasi-integrator with a two-step lag, and two predictors applied to a simulated 8 layer atmosphere from Neyman (2004) [190]. Both prediction methods show improvement over a classic integrator, and EOF shows the most dramatic performance improvement.

14 m/s single-layer turbulence examples, both predictors show improvement over the quasi-integrator, and EOF shows an advantage between the two predictors. Figure 5.12 shows an example of turbulence and residuals of the 14 m/s atmosphere at one step in time. Figure 5.8 shows the results for the multi-layer example, which again shows that both predictors out-perform the integrator, and an advantage to EOF between the two predictors (see Section B.1.5 for further elaboration.)

Table 5.3: Wind layers simulated based on Neyman (2004) [190], to emulate nominal performance at Maunakea.

Velocity	Height	C_N^2
6.7 m/s	0.0 km	0.369×10^{-12}
13.9 m/s	2.1 km	0.219×10^{-12}
20.8 m/s	4.1 km	0.127×10^{-12}
29.0 m/s	6.5 km	0.101×10^{-12}
29.0 m/s	9.0 km	0.046×10^{-12}
29.0 m/s	12.0 km	0.111×10^{-12}
29.0 m/s	14.8 km	0.027×10^{-12}

On-sky telemetry

Our telemetry data is from the Keck II AO system running in closed loop on the night of 04-20-2019, previously published in Jensen-Clem (2019) [116]. It consists of 120000 time steps of data (taken at $\sim$ 1.7ms intervals [32]) in the form of DM commands and near-IR pyramid wavefront sensor [19] (installed in Keck II as part of the KPIC upgrade [288]) phase measurements in units of volts. Keck II is a 10 meter telescope, with a 21 by 21 actuator deformable mirror (we use the pyramid wavefront sensor's 21 by 21 reconstructed residuals to match.) As we see only the residuals and applied commands in this system,

we must first reconstruct the pseudo-open loop turbulence in microns with

$$\text{open loop phase} = (-1 \times \text{dm commands} + \text{wf residuals}) \times 0.6 \tag{5.31}$$

$$\text{integrator residuals} = (\text{wf residuals}) \times 0.6 \tag{5.32}$$

The 0.6 factor accounts for system gain and volt-to-micron conversion [116].

Using data from the Maunakea Weather Center [33], we use the Canada-France-Hawaii Telescope (CFHT) to estimate the wind velocity at the ground layer from the two minutes of Keck telemetry that we use for this analysis. Figure 5.9 shows the median x and y velocity components from our portion of the night, alongside a peak identification of the ground layer wind velocity from telemetry data. We estimate the ground layer to be $(v_x, v_y) = 1.07, 5.17\text{m/s}$. DIMM-reported seeing on that night was ~ 0.7", as opposed to MASS-reported seeing of ~ 0.3" [33]. Because MASS measures free atmosphere, the discrepancy between the two values implies the turbulence that night was dominated by the ground-layer, making it a good candidate for improvement with predictive control. Using the DIMM value for seeing that night, we find $r_0 \sim 15\text{cm}$ at 500nm.

We applied both predictors to reconstructed open-loop phase screens, calculated residuals by subtracting the predicted screen via each method from the

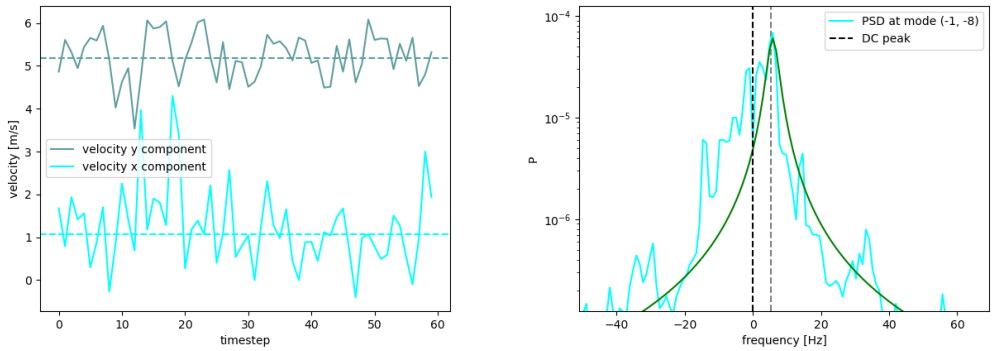


Figure 5.9: Left: Canada-France-Hawaii Telescope data from 04-20-19 used to estimate a median x and y velocity component. Right: Peak identification at the (-1, -8) mode. We identify the ground-layer velocity peak expected from the CFHT data (grey dashed line).

reconstructed screen, and then calculated the root mean square error. Figure 5.10 shows the results for the two predictors. Both predictors show a ≥ 2 factor improvement over the on-sky integrator, and in a departure from simulated results, PFC shows a modest advantage between the two predictors. Figure 5.11 shows an example of turbulence residuals.

Discussion

Table 5.5.3 compiles the residual wavefront error from each prediction method across each atmosphere and estimates (using the Marechal approximation [104]) the Strehl ratio achieved by the integrator as well as each predictor. Note that simulations are highly idealized, with no measurement error from the wavefront sensor or fitting error from the deformable mirror.

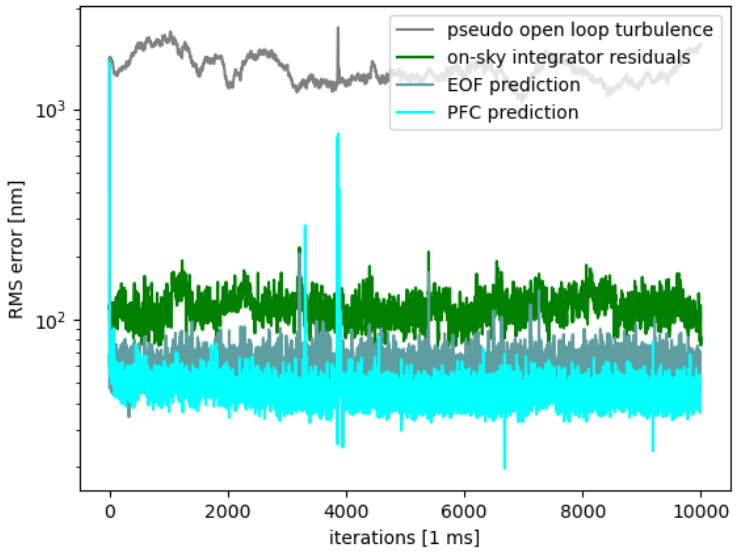


Figure 5.10: Root mean square residuals for the pseudo-open loop reconstructed turbulence, the true on-sky residuals, and the two prediction methods. Both predictors show improvement over a classic integrator, but PFC shows a modest improvement between the two prediction methods.

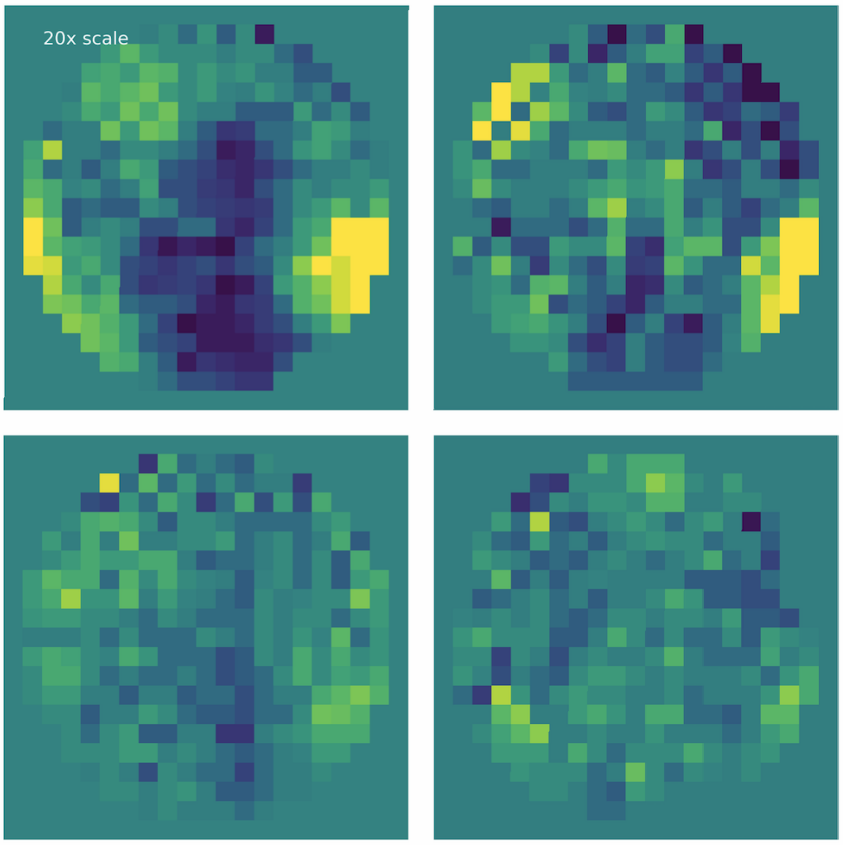


Figure 5.11: Example of turbulence and residuals from on-sky telemetry. Top left: Pseudo-open loop residuals. Top right: on-sky integrator residuals. Bottom left: PFC predictor residuals. Bottom right: EOF predictor residuals. The open-loop turbulence is at 20x the scale of the other three residuals.

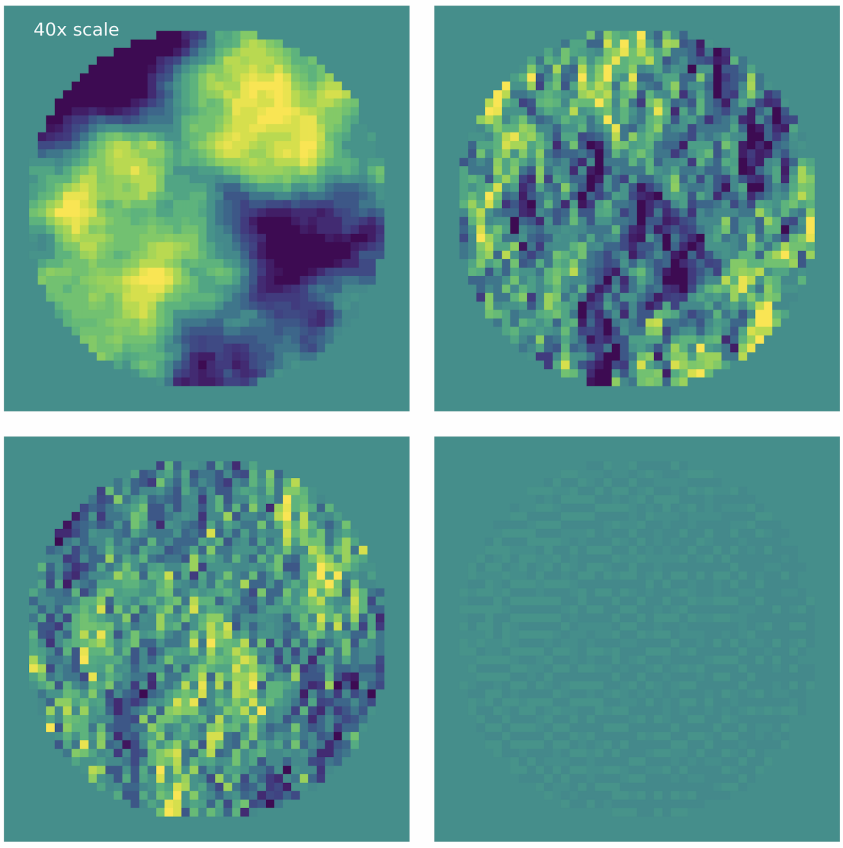


Figure 5.12: Example of turbulence and residuals evolving in real time for the 14 m/s wind simulated atmosphere. Top left: Uncorrected atmospheric turbulence. Top right: Quasi-integrator residuals. Bottom left: PFC predictor residuals. Bottom right: EOF predictor residuals. The open-loop turbulence is at 40x the scale of the other three residuals. Each is one demonstrative step in time.

Atmosphere	wavelength	Predictor	WFE	Strehl Ratio
7 m/s	500 nm	quasi-integrator	6.45 +/- 0.9 nm	92.2 %
7 m/s	500 nm	PFC predictor	3.31 +/- 0.2 nm	95.9 %
7 m/s	500 nm	EOF predictor	1.75 +/- 0.2 nm	97.8 %
14 m/s	500 nm	quasi-integrator	12.84 +/- 1.1 nm	85.1 %
14 m/s	500 nm	PFC predictor	6.48 +/- 0.46 nm	92.2 %
14 m/s	500 nm	EOF predictor	0.67 +/- 0.11 nm	99.2 %
multi-layer	500 nm	quasi-integrator	36.84 +/- 1.54 nm	62.9 %
multi-layer	500 nm	PFC predictor	19.39 +/- 2.37 nm	78.4 %
multi-layer	500 nm	EOF predictor	7.19 +/- 0.30 nm	91.4 %
telemetry	1633 nm	on-sky integrator	112.35 +/- 17.27 nm	64.9 %
telemetry	1633 nm	PFC predictor	44.96 +/- 18.56 nm	84.1 %
telemetry	1633 nm	EOF predictor	58.43 +/- 8.02 nm	79.9 %

Table 5.4: Estimation of Strehl Ratio improvements from residual WFE of each predictor.

For every simulated atmospheric profile and on-sky telemetry we find that both predictors provide an improvement over a classic integrator. PFC provides a consistent improvement of 2-3 over the integrator, while EOF shows more variety. EOF succeeds more notably in picking out one extremely fast wind layer (the single-layer 14 m/s atmosphere), but requires a longer training set (90000 frames of data) to find and remove the 7 m/s wind layer. In contrast, PFC performs consistently with only 10000 frames of training data. Both methods of predictive control also out-perform the integrator in the on-sky telemetry example, which is consistent with the results of Jensen-Clem (2019) [116]. For the on-sky telemetry, we find that PFC shows a greater improvement than EOF, in contrast with its performance over the simulated atmospheres. This may be indicative of the way a Kalman filter handles Gaussian noise (a Kalman filter carries the variance of a distribution in its covariance matrix through each iteration [252]), as opposed to the treatment of EOF as a purely linear filter. While EOF shows greater performance consistently over atmospheric simulation (which includes no measurement noise), the comparative performance over telemetry may better forecast lab and on-sky performance of the two predictors. However, between the two methods, EOF is less complex to implement, as evinced by the length of Section 5.5.2.

5.6 Conclusion

We compare two aspects of predictive wavefront control: whether it runs in closed or pseduo-open loop, and whether it is built purely from data, or from a physics-based model. A preliminary simulation of closed-loop EOF does not perform as well as an optimized implementation of open loop EOF, but still provides improvement over a classic integrator. A closed-loop implementation also proves to be more robust to the introduction of DM model errors. We speculate that future optimization of the closed-loop predictive filter (e.g., exploring history vector length, training data length, etc) will likely close the gap between the open and closed-loop implementations and provide a controller that is even more robust to model errors. We also note that for our preliminary simulations we used perfect knowledge of the system to train our closed-loop predictive filter, which is not realistic for on-sky implementation; if performance comparisons prove promising, a more realistic training method could be devised, for example using the transfer function as in Desenne's original work [58] or training on open and closed-loop data at the beginning of the night. We also compare the purely data driven EOF to the model-driven PFC. We find that both the data-driven predictor using empirical orthogonal functions (EOF) and the model-driven predictor using predictive Fourier control (PFC) provide a consistent improvement across every simulated atmospheric

profile and the telemetry compared with the integrator. While EOF shows greater improvement than PFC over our three simulated atmospheres, PFC shows greater improvement than EOF over on-sky telemetry data, which may better correlate with future on-sky performance. The major difference between these methods is how noise is handled by the Kalman filter, implying that this treatment of error propagation is better for a more realistic system.

Chapter 6

Conclusions

6.1 Polarimetry makes use of parts of light we normally neglect to understand atmospheric asymmetries that are hard to probe with other observational techniques

Finding scattering signatures from exoplanets and brown dwarfs is a promising technique for understanding atmospheric asymmetry that provides complementary information to direct spectroscopy. Typically, these data should also contain all of the information available within direct imaging data (e.g., photometry, updated astrometry, and total intensity spectroscopy). With a sample

size of two targets we were not able to get data of high enough quality to directly detect the known companions in total intensity (i.e., do standard direct imaging science) or place informative upper limits. This work presents the first attempt to use Subuaru/SCExAO for point source polarimetry and within our data is information that can be used to better calibration future polarimetric imaging modes for SCExAO. While this study selected companions for their likelihood to have asymmetric cloud features, signatures that are $\leq 0.5\%$ fractional polarized intensity, future work will target companions that show evidence of circum-planetary disks, a signal more on the order of $0.5 - 15\%$ fractional polarized intensity.

AF Lep b continues to hold promise, as a very unusal and active object, but due to its edge-on inclination it will not be observable again for a few years. A target that has entranced and eluded me for years is the bright, nearby brown dwarf binary, Eps Indi Ba/Bb; I think this system has the potential to be the next Luhman 16, but only when we have an IR wavefront sensor and a polarimeter on a Southern telescope. This will be true with the installation of SAXO+ in $\sim$2028 [68].

While there is exciting science to be done on 8-10 meter class telescopes with point-source polarimetry, I also think it is important precursor work for detecting the scattering features that confirm the presence of an exo-Earth

with the future ELT/HWO. We need to have experience with and develop techniques for point source polarimetric detections if we want to resolve the extremely faint signals on extremely faint planets with our future telescopes.

6.2 Visible light adaptive optics is extremely challenging but current performance is very well matched to error budgets and models

Visible light AO is extremely challenging, and the first generation Keck-II AO system (with its low degree of freedom deformable mirror) will not be able to provide Strehls of more than 2%. Keck AO error budget predictions and empirically adjusted analytic models are in good agreement with real ORKID data, but it is important to regularly verify our models and predictions with empirical data. While data and models provide a global agreement, there are some interesting morphological features in the ORKID PSF that all show ways that Keck AO could be improved today with minimal and non-invasive changes.

Future work will explore how ORKID's performance changes after the commissioning of the HAKA deformable mirror. However, more actuators are not a solution without further complications, as more actuators require more

light, and without it wavefront sensor errors will dominate the error budget. This motivates future technologies like predictive control that can predict around longer exposures and slower control rates or projects like ORCAS that use artificial spaced-based laser beacons as "natural" guide stars.

To push towards higher Strehl Ratios in the visible there is a lot of work left to do. Even with HAKA, Keck is projected to reach $\sim$ 7% SR in the visible. After removing fitting error (due to deformable mirror degrees of freedom), and assuming we have a solution to reduce measurement error (with an extremely bright guide star requirement), other error terms like margin error, temporal errors, known but uncorrectable instrument errors are left to impact performance [27]. Temporal errors will be mitigated with predictive wavefront control, but margin error, the lump term added to balance the AO error budget needs to be demystified. It is speculated that margin error may be due to segment cophasing errors; if so, segment cophasing is only going get more complex in the ELT era, and solving this problem on Keck is a vital precursor to future telescopes.

6.3 Data-driven predictive wavefront control improves adaptive optics corrections for a ten meter telescope

We successfully demonstrated a data-driven predictive wavefront controller on-sky at Keck Observatory. A stable implementation of empirical Orthogonal Functions is now built into the Keck RTC. In a half night of engineering time we saw an improvement of 20 % over a standard integrator in RMS wavefront residuals and 3% in point-spread function stability in principal component strength. Performance of the predictor was dependent on filter hyper-parameters, but most parameter options provide better performance than an integrator. Practically this lends itself to setting some default parameters with the option to tweak when observers are interested in best possible performance.

Facilitizing predictive control as a mode for all future observers will require (1) a strategy for building HAKA-sized filters in soft real time ($\sim$ minute timescales) and (2) a more streamlined method for building and updating filters. Both are problems that are limited by human hours, not by lack of clear solutions.

EOF at Keck is currently the most recent implmentation of a predictor on an 8-10 meter class telescope and the only implementation on a segmented

telescope. It has the potential to be the first facilitized controller other than a standard integrator on an 8-10 meter telescope. Predictive control at Keck improves AO performance and observations the Keck community, but it also makes Keck an important testbed for future advanced control strategies for ELTs and the HWO.

6.4 There are many different ways to do predictive wavefront control and finding the "best" one will require a direct comparison on a telescope

We explored three different flavors of predictive control, comparing open loop data-driven control, closed-loop data-driven control, and model-driven control. While all of these controllers out-perform an integrator, model-driven control provides the best performance on Keck telemetry data. This is likely due to either (1) some kind of model-mismatch when building pseduo-open loop slopes for the data-driven method or (2) ways that a Kalman filter carries noise through a system in a more realistic way. To address (1) we looked a preliminary implementation of running data-driven PWFC without building

POL slopes. This implementation was ultimately not complete enough to make strong comments on performance but in a contrived example does show that a closed-loop implementation has the potential to overcome errors due to model-mismatch in POL slopes.

Working with controllers is a uniquely satisfying way to improve AO performance; no new hardware is required, predictive methods can overcome (to some extent) a requirement for longer integration times in sensing, and next generation controllers can incorporate all useful data. With the 1993 Paschall & Anderson paper [202] on LQG controllers for astronomical telescopes, we celebrate more than 30 years of machine learning for wavefront control, over which machine learning approaches to wavefront control have been suggested, simulated, tested in the lab, and rarely tested on sky. Few [95, 280] works show proven results on-sky (for a 8-10 meter class telescope controlling all modes with a predictive method), making it difficult to promise or benchmark performance for future systems. Furthermore, most of these efforts focus on a single algorithm and its development and progression on a single system. Few efforts [72, 297, 97] have attempted to compare the performance of these methods side-by-side, to determine how these methods may improve performance, stability, or make for simpler problems across a uniform simulation, lab demonstration, or telescope system. We need the opportunity to compare state of the

art algorithms, ideally with on-sky performance on a large telescope. This will be possible in the next few years with the SAXO+ upgrade, scheduled for 2028 [68], to the AO system feeding VLT/SPHERE. This new real time controller will include 6 predictive controllers, with VLT engineering time earmarked to demonstrate their comparative performance on an 8 meter telescope.

For my own future work I am interested in two open questions: (1) when are nonlinear controllers necessary for astronomical AO and (2) how should we incorporate multiple data streams in the next genration of controllers. Wavefront sensing is an interesting potential source of non-linearities, covered at length in Wong, 2022 [296]. In short, wavefront sensing often operates as a tradeoff between sensitivity and linear sensible range (the more sensitive the wavefront sensor is, the lower the usable linear range becomes.) While we attempt to design our wavefront sensor in such a way that this assumption of linearity is valid (i.e. by introducing dynamic modulation), the Pyramid and Zernike wavefront sensors are known to have nonlinear responses for large aberrations. However, running wavefront control in closed-loop leaves residual turbulence sensed by the wavefront sensor within the linear range of most wavefront sensors used for first-stage AO. Similarly, deformable mirrors (DMs) operate with actuator pokes that have a limited linear range; in practice, the linear range of the DM is much larger than that of the wavefront sensor. For

the current generation of large telescopes we do not expect to need nonlinear control. However, this may no longer be true with the sensing strategies for ELTs and HWO. There is an attractive simplicity to neural networks that learn a full system model and optimal controller, removing the need for high fidelity models and optimization studies. However, more work is needed to understand the computation and energy cost of this, especially in the context of space missions. Physically informed neural networks constrain some of these degrees freedom, akin to a model-driven predictive controller, and may be a way to balance flexibility and energy concerns.

Additionally, future extremely large telescopes (and to some extent even current large telescopes) are limited by non-common path aberrations, or errors that are not sensed by a typical upstream wavefront sensor. Controllers that fuse data from the sensor and final science image plane are an interesting way to address this, and there are a lot of techniques to fuse information, ranging from neural network architectures to linear regressors to linear quadratic Gaussian controllers with gussied up equations of state.

Finally, in collaboration with members of the GRAVITY+ AO team, I discovered that the GRAVITY+ AO (GPAO) controller saves a closed loop predictive control state vector. This means that with a non-invasive upgrade, we could run data-driven predictve control on GRAVITY. GRAVITY is an

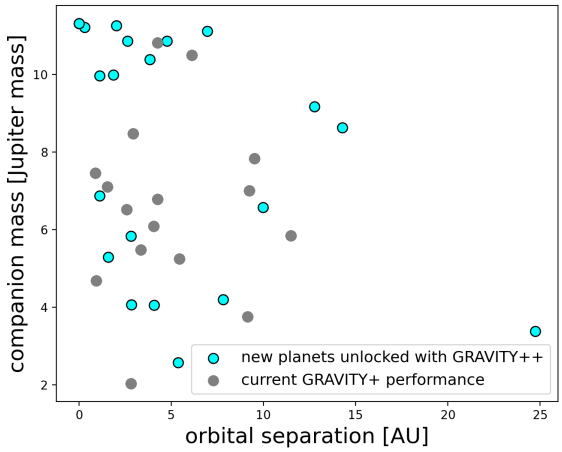


Figure 6.1: Improving the companion sensitivity by 1 magnitude and the guide star limit by 1 magnitude will double the projected planet yield from Gaia DR4. Using a simulated Gaia DR4 population [91] I demonstrate the population of targets available with the current GRAVITY+ sensitivity and AO performance (shown in grey) and the 21 additional targets available with improved GRAVITY performance (shown in cyan). Target populations are simulated with a random age and with 66 % sky coverage to account for only Southern hemisphere targets.

interferometry mode on the Very Large Telescope Interferometer, meaning it contends with the fast jet-stream at Paranal, one of the strongest drivers of temporal errors. GRAVITY will also be a core component of synergistic observing of Gaia driven target selection when the fourth data release (DR4), scheduled for December of 2026. Generous estimates of improvement due to advanced controllers show that GRAVITY could more than double the potential planet detections from Gaia DR4.

6.5 Future outlook and machinations

We are in a liminal time between telescope classes, where we continue to push forward and refine the 8-10 meter class, but also work towards first light of the 20-40 meter class. What I think is unique about my work is that I have used a current 10 meter telescope as a testbed for future extremely large telescopes. Large telescopes like Keck are not only the closest facsimile we have to ELTs, but any work done on large telescopes is an investment in our current architecture and making it better for current observers.

As I have done this work I see tension between how to share resources on our 8-10 meter class telescopes between current observations and future instrumentation. Keck, Gemini, Subaru, and VLT all have varying philosophies of how open they are or should be to experimentation and how that impacts facility use. Keck is the only facility AO-fed segmented large telescope, and will be invaluable for testing control for ELTs and for HWO, but only if we can marshal enough human power do experiments safely without impacting regular observations. SCExAO is so popular because of its flexibility, but sometimes sacrifices consistent performance for the sake of constant experimentation. Gemini and VLT are almost entirely frozen, without the flexibility to adapt and upgrade AO apart from large multi-year projects, but those large multi-year projects are some of the most successful high contrast imaging instruments.

However, GRAVITY+ AO holds promise for a more flexible corner of the VLT. We need to be able to use these telescopes as demonstrators but I'm not sure how we balance that with longer term performance and sustainability for the telescope and its staff.

I wanted to include in my conclusion any available work on AO error budgeting for the Planetary Camera and Spectrograph (PCS), the second generation high contrast imaging instrument for the ELT, and comment on the technology gaps left to overcome, and how the work in this thesis tackles those gaps. But it turns out this error budgeting work has not yet been done, and that's part of what I will be working on as a postdoctoral researcher at Leiden Observatory.

In the next 20 years (if not sooner) we are going to find an Earth like planet. Part of this confirmation will be via finding the reflected light scattering feature of its Ocean, rainbow, or water-laden clouds. I have observed using the same polarimetric techniques we will use to isolate the scattering signature, characterized the performance of visible light AO where that scattering feature will fall, and tested on-sky the control algorithms that will lead to accurate enough AO corrections to find those few photons. This is how I have addressed the problem of finding an exo-Earth, from snout to tail, using every part of the PIG.

Bibliography

[1] Abdurro'uf, Katherine Accetta, Conny Aerts, Víctor Silva Aguirre, Romina Ahumada, Nikhil Ajgaonkar, N. Filiz Ak, Shadab Alam, Carlos Allende Prieto, Andrés Almeida, Friedrich Anders, Scott F. Anderson, Brett H. Andrews, Borja Anguiano, Erik Aquino-Ortíz, Alfonso Aragón-Salamanca, Maria Argudo-Fernández, Metin Ata, Marie Aubert, Vladimir Avila-Reese, Carles Badenes, Rodolfo H. Barbá, Kat Barger, Jorge K. Barrera-Ballesteros, Rachael L. Beaton, Timothy C. Beers, Francesco Belfiore, Chad F. Bender, Mariangela Bernardi, Matthew A. Bershady, Florian Beutler, Christian Moni Bidin, Jonathan C. Bird, Dmitry Bizyaev, Guillermo A. Blanc, Michael R. Blanton, Nicholas Fraser Boardman, Adam S. Bolton, Médéric Boquien, Jura Borissova, Jo Bovy, W. N. Brandt, Jordan Brown, Joel R. Brownstein, Marcella Brusa, Johannes Buchner, Kevin Bundy, Joseph N. Burchett, Martin Bureau, Adam Burgasser, Tuesday K. Cabang, Stephanie Campbell, Michele Cappellari,

Joleen K. Carlberg, Fábio Carneiro Wanderley, Ricardo Carrera, Jennifer Cash, Yan-Ping Chen, Wei-Huai Chen, Brian Cherinka, Cristina Chiappini, Peter Doohyun Choi, S. Drew Chojnowski, Haeun Chung, Nicolas Clerc, Roger E. Cohen, Julia M. Comerford, Johan Comparat, Luiz da Costa, Kevin Covey, Jeffrey D. Crane, Irene Cruz-Gonzalez, Connor Culhane, Katia Cunha, Y. Sophia Dai, Guillermo Damke, Jeremy Darling, James W. Davidson, Jr., Roger Davies, Kyle Dawson, Nathan De Lee, Aleksandar M. Diamond-Stanic, Mariana Cano-Díaz, Helena Domínguez Sánchez, John Donor, Chris Duckworth, Tom Dwelly, Daniel J. Eisenstein, Yvonne P. Elsworth, Eric Emsellem, Mike Eracleous, Stephanie Escoffier, Xiaohui Fan, Emily Farr, Shuai Feng, José G. Fernández-Trincado, Diane Feuillet, Andreas Filipp, Sean P. Fillingham, Peter M. Frinchaboy, Sebastien Fromenteau, Lluís Galbany, Rafael A. García, D. A. García-Hernández, Junqiang Ge, Doug Geisler, Joseph Gelfand, Tobias Géron, Benjamin J. Gibson, Julian Goddy, Diego Godoy-Rivera, Kathleen Grabowski, Paul J. Green, Michael Greener, Catherine J. Grier, Emily Griffith, Hong Guo, Julien Guy, Massinissa Hadjara, Paul Harding, Sten Hasselquist, Christian R. Hayes, Fred Hearty, Jesús Hernández, Lewis Hill, David W. Hogg, Jon A. Holtzman, Danny Horta, Bau-Ching Hsieh, Chin-Hao Hsu, Yun-Hsin Hsu, Daniel Huber, Marc Huertas-Company,

Brian Hutchinson, Ho Seong Hwang, Héctor J. Ibarra-Medel, Jacob Ider Chitham, Gabriele S. Ilha, Julie Imig, Will Jaekle, Tharindu Jayasinghe, Xihan Ji, Jennifer A. Johnson, Amy Jones, Henrik Jönsson, Ivan Katkov, Arman Khalatyan, Dr., Karen Kinemuchi, Shobhit Kisku, Johan H. Knapen, Jean-Paul Kneib, Juna A. Kollmeier, Miranda Kong, Marina Kounkel, Kathryn Kreckel, Dhanesh Krishnarao, Ivan Lacerna, Richard R. Lane, Rachel Langgin, Ramon Lavender, David R. Law, Daniel Lazarz, Henry W. Leung, Ho-Hin Leung, Hannah M. Lewis, Cheng Li, Ran Li, Jianhui Lian, Fu-Heng Liang, Lihwai Lin, Yen-Ting Lin, Sicheng Lin, Chris Lintott, Dan Long, Penélope Longa-Peña, Carlos López-Cobá, Shengdong Lu, Britt F. Lundgren, Yuanze Luo, J. Ted Mackereth, Axel de la Macorra, Suvrath Mahadevan, Steven R. Majewski, Arturo Manchado, Travis Mandeville, Claudia Maraston, Berta Margalef-Bentabol, Thomas Masseron, Karen L. Masters, Savita Mathur, Richard M. McDermid, Myles Mckay, Andrea Merloni, Michael Merrifield, Szabolcs Meszaros, Andrea Miglio, Francesco Di Mille, Dante Minniti, Rebecca Minsley, and Antonela Monachesi. The Seventeenth Data Release of the Sloan Digital Sky Surveys: Complete Release of MaNGA, MaStar, and APOGEE-2 Data. , 259(2):35, April 2022.

[2] Kyohoon Ahn, Olivier Guyon, Julien Lozi, Sébastien Vievard, Vin-

cent Deo, Nour Skaf, Ruslan Belikov, Steven P. Bos, Michael Bottom, Thayne Currie, Richard Frazin, Kyle V. Gorkom, Tyler D. Groff, Sebastiaan Y. Haffert, Nemanja Jovanovic, Hajime Kawahara, Takayuki Kotani, Jared R. Males, Frantz Martinache, Ben Mazin, Kelsey Miller, Barnaby Norris, Alexander Rodack, and Alison Wong. SCExAO: a testbed for developing high-contrast imaging technologies for ELTs. In Stuart B. Shaklan and Garreth J. Ruane, editors, *Techniques and Instrumentation for Detection of Exoplanets X*, volume 11823 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1182303, September 2021.

[3] Romina Ahumada, Carlos Allende Prieto, Andrés Almeida, Friedrich Anders, Scott F. Anderson, Brett H. Andrews, Borja Anguiano, Riccardo Arcodia, Eric Armengaud, Marie Aubert, Santiago Avila, Vladimir Avila-Reese, Carles Badenes, Christophe Balland, Kat Barger, Jorge K. Barrera-Ballesteros, Sarbani Basu, Julian Bautista, Rachael L. Beaton, Timothy C. Beers, B. Izamar T. Benavides, Chad F. Bender, Mariangela Bernardi, Matthew Bershady, Florian Beutler, Christian Moni Bidin, Jonathan Bird, Dmitry Bizyaev, Guillermo A. Blanc, Michael R. Blanton, Médéric Boquien, Jura Borissova, Jo Bovy, W. N. Brandt, Jonathan Brinkmann, Joel R. Brownstein, Kevin Bundy, Martin Bureau,

Adam Burgasser, Etienne Burtin, Mariana Cano-Díaz, Raffaella Capasso, Michele Cappellari, Ricardo Carrera, Solène Chabanier, William Chaplin, Michael Chapman, Brian Cherinka, Cristina Chiappini, Peter Doohyun Choi, S. Drew Chojnowski, Haeun Chung, Nicolas Clerc, Damien Coffey, Julia M. Comerford, Johan Comparat, Luiz da Costa, Marie-Claude Cousinou, Kevin Covey, Jeffrey D. Crane, Katia Cunha, Gabriele da Silva Ilha, Yu Sophia Dai, Sanna B. Damsted, Jeremy Darling, James W. Davidson, Jr., Roger Davies, Kyle Dawson, Nikhil De, Axel de la Macorra, Nathan De Lee, Anna Bárbara de Andrade Queiroz, Alice Deconto Machado, Sylvain de la Torre, Flavia Dell'Agli, Hélion du Mas des Bourboux, Aleksandar M. Diamond-Stanic, Sean Dillon, John Donor, Niv Drory, Chris Duckworth, Tom Dwelly, Garrett Ebelke, Sarah Eftekharzadeh, Arthur Davis Eigenbrot, Yvonne P. Elsworth, Mike Eracleous, Ghazaleh Erfanianfar, Stephanie Escoffier, Xiaohui Fan, Emily Farr, José G. Fernández-Trincado, Diane Feuillet, Alexis Finoguenov, Patricia Fofie, Amelia Fraser-McKelvie, Peter M. Frinchaboy, Sebastien Fromenteau, Hai Fu, Lluís Galbany, Rafael A. Garcia, D. A. García-Hernández, Luis Alberto Garma Oehmichen, Junqiang Ge, Marcio Antonio Geimba Maia, Doug Geisler, Joseph Gelfand, Julian Goddy, Violeta Gonzalez-Perez, Kathleen Grabowski, Paul Green, Catherine J. Grier,

Hong Guo, Julien Guy, Paul Harding, Sten Hasselquist, Adam James Hawken, Christian R. Hayes, Fred Hearty, S. Hekker, David W. Hogg, Jon A. Holtzman, Danny Horta, Jiamin Hou, Bau-Ching Hsieh, Daniel Huber, Jason A. S. Hunt, J. Ider Chitham, Julie Imig, Mariana Jaber, Camilo Eduardo Jimenez Angel, Jennifer A. Johnson, Amy M. Jones, Henrik Jönsson, Eric Jullo, Yerim Kim, Karen Kinemuchi, Charles C. Kirkpatrick, IV, George W. Kite, Mark Klaene, Jean-Paul Kneib, Juna A. Kollmeier, Hui Kong, Marina Kounkel, Dhanesh Krishnarao, Ivan Lacerna, Ting-Wen Lan, Richard R. Lane, David R. Law, Jean-Marc Le Goff, Henry W. Leung, Hannah Lewis, Cheng Li, Jianhui Lian, Lihwai Lin, Dan Long, Penélope Longa-Peña, Britt Lundgren, Brad W. Lyke, J. Ted Mackereth, Chelsea L. MacLeod, Steven R. Majewski, Arturo Manchado, Claudia Maraston, Paul Martini, Thomas Masseron, Karen L. Masters, Savita Mathur, Richard M. McDermid, Andrea Merloni, Michael Merrifield, Szabolcs Mészáros, Andrea Miglio, Dante Minniti, Rebecca Minsley, Takamitsu Miyaji, Faizan Gohar Mohammad, Benoit Mosser, Eva-Maria Mueller, Demitri Muna, Andrea Muñoz-Gutiérrez, Adam D. Myers, Seshadri Nadathur, Preethi Nair, Kirpal Nandra, Janaina Correa do Nascimento, Rebecca Jean Nevin, Jeffrey A. Newman, David L. Nidever, Christian Nitschelm, Pasquier Noterdaeme, Julia E. O'Connell,

Matthew D. Olmstead, Daniel Oravetz, Audrey Oravetz, Yeisson Osorio, Zachary J. Pace, Nelson Padilla, Nathalie Palanque-Delabrouille, and Pedro A. Palicio. The 16th Data Release of the Sloan Digital Sky Surveys: First Release from the APOGEE-2 Southern Survey and Full Release of eBOSS Spectra. , 249(1):3, July 2020.

[4] G. J. Alred, C. T. Brusaw, and W. E. Oliu. *Handbook of Technical Writing, 7th ed.* St. Martin's, New York, 2003.

[5] S. Mark Ammons, Sarah E. Robinson, Jay Strader, Gregory Laughlin, Debra Fischer, and Aaron Wolf. The N2K Consortium. IV. New Temperatures and Metallicities for More than 100,000 FGK Dwarfs. , 638(2):1004–1017, February 2006.

[6] E. Anderson and Ch. Francis. XHIP: An extended hipparcos compilation. *Astronomy Letters*, 38(5):331–346, May 2012.

[7] V. M. Andruk, L. K. Pakuliak, V. V. Golovnia, G. O. Ivanov, A. I. Yatsenko, S. V. Shatokhina, and O. M. Yizhakevych. Catalog of equatorial coordinates and B-magnitudes of stars of the FON project. *Kinematics and Physics of Celestial Bodies*, 32(5):260–263, September 2016.

[8] T. Aparicio Villegas, E. J. Alfaro, J. Cabrera-Caño, M. Moles, N. Benítez, J. Perea, A. del Olmo, A. Fernández-Soto, D. Cristóbal-Hornillos,

C. Husillos, J. A. L. Aguerri, T. Broadhurst, F. J. Castander, J. Cepa, M. Cerviño, R. M. González Delgado, L. Infante, I. Márquez, J. Masegosa, V. J. Martínez, F. Prada, J. M. Quintana, and S. F. Sánchez. The ALHAMBRA Photometric System. , 139(3):1242–1253, March 2010.

[9] Astropy Collaboration, A. M. Price-Whelan, B. M. Sipőcz, H. M. Günther, P. L. Lim, S. M. Crawford, S. Conseil, D. L. Shupe, M. W. Craig, N. Dencheva, A. Ginsburg, J. T. VanderPlas, L. D. Bradley, D. Pérez-Suárez, M. de Val-Borro, T. L. Aldcroft, K. L. Cruz, T. P. Robitaille, E. J. Tollerud, C. Ardelean, T. Babej, Y. P. Bach, M. Bachetti, A. V. Bakanov, S. P. Bamford, G. Barentsen, P. Barmby, A. Baumbach, K. L. Berry, F. Biscani, M. Boquien, K. A. Bostroem, L. G. Bouma, G. B. Brammer, E. M. Bray, H. Breytenbach, H. Buddelmeijer, D. J. Burke, G. Calderone, J. L. Cano Rodríguez, M. Cara, J. V. M. Cardoso, S. Cheedella, Y. Copin, L. Corrales, D. Crichton, D. D'Avella, C. Deil, É. Depagne, J. P. Dietrich, A. Donath, M. Droettboom, N. Earl, T. Erben, S. Fabbro, L. A. Ferreira, T. Finethy, R. T. Fox, L. H. Garrison, S. L. J. Gibbons, D. A. Goldstein, R. Gommers, J. P. Greco, P. Greenfield, A. M. Groener, F. Grollier, A. Hagen, P. Hirst, D. Homeier, A. J. Horton, G. Hosseinzadeh, L. Hu, J. S. Hunkeler, Ž. Ivezić, A. Jain, T. Jenness, G. Kanarek, S. Kendrew, N. S. Kern, W. E. Kerzendorf, A. Khvalko, J. King, D. Kirkby, A. M. Kulkarni,

A. Kumar, A. Lee, D. Lenz, S. P. Littlefair, Z. Ma, D. M. Macleod, M. Mastropietro, C. McCully, S. Montagnac, B. M. Morris, M. Mueller, S. J. Mumford, D. Muna, N. A. Murphy, S. Nelson, G. H. Nguyen, J. P. Ninan, M. Nöthe, S. Ogaz, S. Oh, J. K. Parejko, N. Parley, S. Pascual, R. Patil, A. A. Patil, A. L. Plunkett, J. X. Prochaska, T. Rastogi, V. Reddy Janga, J. Sabater, P. Sakurikar, M. Seifert, L. E. Sherbert, H. Sherwood-Taylor, A. Y. Shih, J. Sick, M. T. Silbiger, S. Singanamalla, L. P. Singer, P. H. Sladen, K. A. Sooley, S. Sornarajah, O. Streicher, P. Teuben, S. W. Thomas, G. R. Tremblay, J. E. H. Turner, V. Terrón, M. H. van Kerkwijk, A. de la Vega, L. L. Watkins, B. A. Weaver, J. B. Whitmore, J. Woillez, V. Zabalza, and Astropy Contributors. The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package. , 156(3):123, September 2018.

[10] Astropy Collaboration, Thomas P. Robitaille, Erik J. Tollerud, Perry Greenfield, Michael Droettboom, Erik Bray, Tom Aldcroft, Matt Davis, Adam Ginsburg, Adrian M. Price-Whelan, Wolfgang E. Kerzendorf, Alexander Conley, Neil Crighton, Kyle Barbary, Demitri Muna, Henry Ferguson, Frédéric Grollier, Madhura M. Parikh, Prasanth H. Nair, Hans M. Unther, Christoph Deil, Julien Woillez, Simon Conseil, Roban Kramer, James E. H. Turner, Leo Singer, Ryan Fox, Benjamin A. Weaver,

Victor Zabalza, Zachary I. Edwards, K. Azalee Bostroem, D. J. Burke, Andrew R. Casey, Steven M. Crawford, Nadia Dencheva, Justin Ely, Tim Jenness, Kathleen Labrie, Pey Lian Lim, Francesco Pierfederici, Andrew Pontzen, Andy Ptak, Brian Refsdal, Mathieu Servillat, and Ole Streicher. Astropy: A community Python package for astronomy. , 558:A33, Oct 2013.

[11] Thomas Ayres. In the Trenches of the Solar-Stellar Connection. VII. Wilson-Bappu 2022. , 266(1):6, May 2023.

[12] S. Bagnulo, M. Landolfi, J. D. Landstreet, E. Landi Degl'Innocenti, L. Fossati, and M. Sterzik. Stellar Spectropolarimetry with Retarder Waveplate and Beam Splitter Devices. , 121(883):993, September 2009.

[13] Vanessa P. Bailey, Philip M. Hinz, Alfio T. Puglisi, Simone Esposito, Vidhya Vaitheeswaran, Andrew J. Skemer, Denis Defrère, Amali Vaz, and Jarron M. Leisenring. Large binocular telescope interferometer adaptive optics: on-sky performance and lessons learned. In Enrico Marchetti, Laird M. Close, and Jean-Pierre Vran, editors, *Adaptive Optics Systems IV*, volume 9148 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 914803, July 2014.

[14] Vanessa P. Bailey, Lisa A. Poyneer, Bruce A. Macintosh, Dmitry Savran-

sky, Jason J. Wang, Robert J. De Rosa, Katherine B. Follette, S. Mark Ammons, Thomas Hayward, Patrick Ingraham, Jérôme Maire, David W. Palmer, Marshall D. Perrin, Abhijith Rajan, Fredrik T. Rantakyrö, Sandrine Thomas, and Jean-Pierre Véran. Status and performance of the Gemini Planet Imager adaptive optics system. In Enrico Marchetti, Laird M. Close, and Jean-Pierre Véran, editors, *Adaptive Optics Systems V*, volume 9909 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 99090V, July 2016.

[15] Rick Berman. Star trek: Deep space nine. Paramount Domestic Television, 1993. Accessed: 2024-3-28.

[16] Rick Berman. Star trek: Voyager. Paramount Domestic Television, 1995. Accessed: 2024-3-28.

[17] E. Bertin and S. Arnouts. SExtractor: Software for source extraction. , 117:393–404, Jun 1996.

[18] Beth A. Biller, Johanna M. Vos, Yifan Zhou, Allison M. McCarthy, Xianyu Tan, Ian J. M. Crossfield, Niall Whiteford, Genaro Suarez, Jacqueline Faherty, Elena Manjavacas, Xueqing Chen, Pengyu Liu, Ben J. Sutlieff, Mary Anne Limbach, Paul Molliere, Trent J. Dupuy, Natalia Oliveros-Gomez, Philip S. Muirhead, Thomas Henning, Gregory Mace,

Nicolas Crouzet, Theodora Karalidi, Caroline V. Morley, Pascal Tremblin, and Tiffany Kataria. The JWST weather report from the nearest brown dwarfs I: multiperiod JWST NIRSpec + MIRI monitoring of the benchmark binary brown dwarf WISE 1049AB. , 532(2):2207–2233, August 2024.

[19] Charlotte Z. Bond, Peter Wizinowich, Mark Chun, Dimitri Mawet, Scott Lilley, Sylvain Cetre, Nemanja Jovanovic, Jacques-Robert Delorme, Edward Wetherell, Shane Jacobson, Charles Lockhart, Eric Warmbier, James K. Wallace, Donald N. Hall, Sean Goebel, Olivier Guyon, Cedric Plantet, Guido Agapito, Christophe Giordano, Simone Esposito, and Bruno Femenia-Castella. Adaptive optics with an infrared pyramid wavefront sensor. In Laird M. Close, Laura Schreiber, and Dirk Schmidt, editors, *Adaptive Optics Systems VI*, volume 10703 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 107031Z, July 2018.

[20] Markus J. Bonse, Timothy D. Gebhard, Felix A. Dannert, Olivier Absil, Faustine Cantalloube, Valentin Christiaens, Gabriele Cugno, Emily O. Garvin, Jean Hayoz, Markus Kasper, Elisabeth Matthews, Bernhard Schölkopf, and Sascha P. Quanz. Use the 4S (Signal-Safe Speckle Sub-

traction): Explainable Machine Learning Reveals the Giant Exoplanet AF Lep b in High-contrast Imaging Data from 2011. , 169(4):194, April 2025.

[21] Steven P. Bos, Michael Bottom, Sam Ragland, Jacques-Robert Delorme, Sylvain Cetre, and Laurent Pueyo. Fast and furious focal-plane wavefront sensing at W.M. Keck Observatory. In Stuart B. Shaklan and Garreth J. Ruane, editors, *Techniques and Instrumentation for Detection of Exoplanets X*, volume 11823 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 118231E, September 2021.

[22] Michael Bottom, Samuel A. U. Walker, Ian Cunnyngham, Charlotte E. Guthery, and Jacques-Robert Delorme. Sequential coronagraphic low-order wavefront control. In *Adaptive Optics for Extremely Large Telescopes (AO4ELT7)*, page 125, June 2023.

[23] G. Bourdarot and F. Eisenhauer. Kilometer-baseline interferometry: science drivers for the next generation instrument. In M. Béthermin, K. Baillié, N. Lagarde, J. Malzac, R. M. Ouazzani, J. Richard, O. Venot, and A. Siebert, editors, *SF2A-2024: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, pages 183–186, December 2024.

[24] L. Bourgés, S. Lafrasse, G. Mella, O. Chesneau, J. L. Bouquin, G. Duvert, A. Chelli, and X. Delfosse. The JMMC Stellar Diameters Catalog v2 (JSDC): A New Release Based on SearchCal Improvements. In N. Manset and P. Forshay, editors, *Astronomical Data Analysis Software and Systems XXIII*, volume 485 of *Astronomical Society of the Pacific Conference Series*, page 223, May 2014.

[25] Larry Bradley, Brigitta Sipőcz, Thomas Robitaille, Erik Tollerud, Zè Vinícius, Christoph Deil, Kyle Barbary, Tom J Wilson, Ivo Busko, Axel Donath, Hans Moritz Günther, Mihai Cara, P. L. Lim, Sebastian Meßlinger, Zach Burnett, Simon Conseil, Michael Droettboom, Azalee Bostroem, E. M. Bray, Lars Andersen Bratholm, William Jamieson, Adam Ginsburg, Geert Barentsen, Matt Craig, Sergio Pascual, Shivangee Rathi, Marshall Perrin, and Brett M. Morris. astropy/photutils: 2.2.0, February 2025.

[26] Timothy D. Brandt, Maxime Rizzo, Tyler Groff, Jeffrey Chilcote, Johnny P. Greco, N. Jeremy Kasdin, Mary Anne Limbach, Michael Galvin, Craig Loomis, Gillian Knapp, Michael W. McElwain, Nemanja Jovanovic, Thayne Currie, Kyle Mede, Motohide Tamura, Naruhisa Takato, and Masahiko Hayashi. Data reduction pipeline for the CHARIS integral-field spectrograph I: detector readout calibration and data cube

extraction. *Journal of Astronomical Telescopes, Instruments, and Systems*, 3:048002, October 2017.

[27] Corria C., Wizonowich P., and Ragland S. Real time controller upgrade performance analysis. Technical report, W. M. Keck Observatory, March 2020.

[28] Benjamin Calvin and Michael P. Fitzgerald. Simultaneous operation of a controllable segmented primary mirror and single conjugate adaptive optics system part 1: design concept and sensitivity analysis. *Journal of Astronomical Telescopes, Instruments, and Systems*, 11:019003, January 2025.

[29] Benjamin Calvin, Michael P. Fitzgerald, and Sam Ragland. Simultaneous operation of a controllable segmented primary mirror and single conjugate adaptive optics system, part 2: simulated operation. *Journal of Astronomical Telescopes, Instruments, and Systems*, 11:019004, January 2025.

[30] F. Cantalloube, E. H. Por, K. Dohlen, J. F. Sauvage, A. Vigan, M. Kasper, N. Bharmal, T. Henning, W. Brandner, J. Milli, C. Correia, and T. Fusco. Origin of the asymmetry of the wind driven halo observed in high-contrast images. , 620:L10, December 2018.

[31] Nicolás Cardiel, Jaime Zamorano, Josep Manel Carrasco, Eduard Masana, Salvador Bará, Rafael González, Jaime Izquierdo, Sergio Pascual, and Alejandro Sánchez de Miguel. RGB photometric calibration of 15 million Gaia stars. , 507(1):318–329, October 2021.

[32] Sylvain Cetre, Olivier Guyon, Charlotte Bond, Mark Chun, Dimitri Mawet, Peter Wizinowich, Charles Lockhart, Sean Goebel, and Ed Wetherell. A near-infrared pyramid wavefront sensor for Keck adaptive optics: real-time controller. In Laird M. Close, Laura Schreiber, and Dirk Schmidt, editors, *Adaptive Optics Systems VI*, volume 10703 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1070339, July 2018.

[33] CFHT. Mauna kea weather center. Technical report, http://mkwc.ifa.hawaii.edu/current/seeing/, 2021.

[34] Aritra Chakrabarty, Sujan Sengupta, and Mark S. Marley. Polarization of Rotationally Oblate Self-luminous Exoplanets with Anisotropic Atmospheres. , 927(1):51, March 2022.

[35] K. C. Chambers, E. A. Magnier, N. Metcalfe, H. A. Flewelling, M. E. Huber, C. Z. Waters, L. Denneau, P. W. Draper, D. Farrow, D. P. Finkbeiner, C. Holmberg, J. Koppenhoefer, P. A. Price, A. Rest, R. P.

Saglia, E. F. Schlafly, S. J. Smartt, W. Sweeney, R. J. Wainscoat, W. S. Burgett, S. Chastel, T. Grav, J. N. Heasley, K. W. Hodapp, R. Jedicke, N. Kaiser, R.-P. Kudritzki, G. A. Luppino, R. H. Lupton, D. G. Monet, J. S. Morgan, P. M. Onaka, B. Shiao, C. W. Stubbs, J. L. Tonry, R. White, E. Bañados, E. F. Bell, R. Bender, E. J. Bernard, M. Boegner, F. Boffi, M. T. Botticella, A. Calamida, S. Casertano, W.-P. Chen, X. Chen, S. Cole, N. Deacon, C. Frenk, A. Fitzsimmons, S. Gezari, V. Gibbs, C. Goessl, T. Goggia, R. Gourgue, B. Goldman, P. Grant, E. K. Grebel, N. C. Hambly, G. Hasinger, A. F. Heavens, T. M. Heckman, R. Henderson, T. Henning, M. Holman, U. Hopp, W.-H. Ip, S. Isani, M. Jackson, C. D. Keyes, A. M. Koekemoer, R. Kotak, D. Le, D. Liska, K. S. Long, J. R. Lucey, M. Liu, N. F. Martin, G. Masci, B. McLean, E. Mindel, P. Misra, E. Morganson, D. N. A. Murphy, A. Obaika, G. Narayan, M. A. Nieto-Santisteban, P. Norberg, J. A. Peacock, E. A. Pier, M. Postman, N. Primak, C. Rae, A. Rai, A. Riess, A. Riffeser, H. W. Rix, S. Röser, R. Russel, L. Rutz, E. Schilbach, A. S. B. Schultz, D. Scolnic, L. Strolger, A. Szalay, S. Seitz, E. Small, K. W. Smith, D. R. Soderblom, P. Taylor, R. Thomson, A. N. Taylor, A. R. Thakar, J. Thiel, D. Thilker, D. Unger, Y. Urata, J. Valenti, J. Wagner, T. Walder, F. Walter, S. P. Watters, S. Werner, W. M. Wood-Vasey, and R. Wyse. The Pan-STARRS1 Surveys.

arXiv e-prints, page arXiv:1612.05560, December 2016.

[36] Colin Orion Chandler, Iain McDonald, and Stephen R. Kane. The Catalog of Earth-Like Exoplanet Survey Targets (CELESTA): A Database of Habitable Zones Around Nearby Stars. , 151(3):59, March 2016.

[37] Xueqing Chen, Beth A. Biller, Johanna M. Vos, Ian J. M. Crossfield, Gregory N. Mace, Callie E. Hood, Xianyu Tan, Katelyn N. Allers, Emily C. Martin, Emma Bubb, Jonathan J. Fortney, Caroline V. Morley, and Mark Hammond. Global weather map reveals persistent top-of-atmosphere features on the nearest brown dwarfs. , 533(3):3114–3143, September 2024.

[38] Valentin Christiaens, Carlos Gonzalez, Ralf Farkas, Carl-Henrik Dahlqvist, Evert Nasedkin, Julien Milli, Olivier Absil, Henry Ngo, Carles Cantero, Alan Rainot, Iain Hammond, Markus Bonse, Faustine Cantalloube, Arthur Vigan, Vijay Kompella, and Paul Hancock. VIP: A Python package for high-contrast imaging. *The Journal of Open Source Software*, 8(81):4774, January 2023.

[39] Jessie L. Christiansen, Douglas L. McElroy, Marcy Harbut, David R. Ciardi, Megan Crane, John Good, Kevin K. Hardegree-Ullman, Aurora Y. Kesseli, Michael B. Lund, Meca Lynn, Ananda Muthiar, Ricky Nilsson,

Toba Oluyide, Michael Papin, Amalia Rivera, Melanie Swain, Nicholas D. Susemiehl, Raymond Tam, Julian van Eyken, and Charles Beichman. The NASA Exoplanet Archive and Exoplanet Follow-up Observing Program: Data, Tools, and Usage. , 6(8):186, August 2025.

[40] Dmitry Chulkov and Oleg Malkov. Visual binary stars with known orbits in Gaia EDR3. , 517(2):2925–2941, December 2022.

[41] Anna Ciurlo, Randall D. Campbell, Mark R. Morris, Tuan Do, Andrea M. Ghez, Eric E. Becklin, Rory O. Bentley, Devin S. Chu, Abhimat K. Gautam, Yash A. Gursahani, Aurélien Hees, Kelly Kosmo O'Neil, Jessica R. Lu, Gregory D. Martinez, Smadar Naoz, Shoko Sakai, and Rainer Schödel. The Swansong of the Galactic Center Source X7: An Extreme Example of Tidal Evolution near the Supermassive Black Hole. , 944(2):136, February 2023.

[42] Ryan Cloutier, René Doyon, Francois Bouchy, and Guillaume Hébrard. Quantifying the Observational Effort Required for the Radial Velocity Characterization of TESS Planets. , 156(2):82, Aug 2018.

[43] Lia Corrales. X-Ray Scattering Echoes and Ghost Halos from the Intergalactic Medium: Relation to the Nature of AGN Variability. , 805(1):23, May 2015.

[44] I. J. M. Crossfield, B. Biller, J. E. Schlieder, N. R. Deacon, M. Bonnefoy, D. Homeier, F. Allard, E. Buenzli, Th. Henning, W. Brandner, B. Goldman, and T. Kopytova. A global cloud map of the nearest known brown dwarf. , 505(7485):654–656, January 2014.

[45] P. Cruzalèbes, R. G. Petrov, S. Robbe-Dubois, J. Varga, L. Burtscher, F. Allouche, P. Berio, K.-H. Hofmann, J. Hron, W. Jaffe, S. Lagarde, B. Lopez, A. Matter, A. Meilland, K. Meisenheimer, F. Millour, and D. Schertl. A catalogue of stellar diameters and fluxes for mid-infrared interferometry. , 490(3):3158–3176, December 2019.

[46] Thayne Currie, Olivier Guyon, Julien Lozi, Tyler Groff, N. Jeremy Kasdin, Frantz Martinache, Timothy D. Brandt, Jeffrey Chilcote, Christian Marois, Benjamin Gerard, Nemanja Jovanovic, and Sébastien Vievard. Performance and early science with the Subaru Coronagraphic Extreme Adaptive Optics project. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 11117 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 111170X, September 2019.

[47] Thayne Currie, Olivier Guyon, Julien Lozi, Ananya Sahoo, Sebastien Vievard, Vincent Deo, Jeffrey Chilcote, Tyler Groff, Timothy D. Brandt,

Kellen Lawson, Nour Skaf, Frantz Martinache, and N. Jeremy Kasdin. On-sky performance and recent results from the Subaru coronagraphic extreme adaptive optics system. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 11448 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 114487H, December 2020.

[48] R. M. Cutri, M. F. Skrutskie, S. van Dyk, C. A. Beichman, J. M. Carpenter, T. Chester, L. Cambresy, T. Evans, J. Fowler, J. Gizis, E. Howard, J. Huchra, T. Jarrett, E. L. Kopan, J. D. Kirkpatrick, R. M. Light, K. A. Marsh, H. McCallon, S. Schneider, R. Stiening, M. Sykes, M. Weinberg, W. A. Wheaton, S. Wheelock, and N. Zacarias. VizieR Online Data Catalog: 2MASS All-Sky Catalog of Point Sources (Cutri+ 2003). VizieR On-line Data Catalog: II/246. Originally published in: University of Massachusetts and Infrared Processing and Analysis Center, (IPAC/California Institute of Technology) (2003), June 2003.

[49] R. M. Cutri, E. L. Wright, T. Conrow, J. W. Fowler, P. R. M. Eisenhardt, C. Grillmair, J. D. Kirkpatrick, F. Masci, H. L. McCallon, S. L. Wheelock, S. Fajardo-Acosta, L. Yan, D. Benford, M. Harbut, T. Jarrett, S. Lake, D. Leisawitz, M. E. Ressler, S. A. Stanford, C.-W. Tsai, F. Liu, G. Helou,

A. Mainzer, D. Gettngs, A. Gonzalez, D. Hoffman, K. A. Marsh, D. Padgett, M. F. Skrutskie, R. Beck, M. Papin, and M. Wittman. VizieR Online Data Catalog: AllWISE Data Release (Cutri+ 2013). VizieR On-line Data Catalog: II/328. Originally published in: IPAC/Caltech (2013), February 2021.

[50] R. J. de Kok, D. M. Stam, and T. Karalidi. Characterizing Exoplanetary Atmospheres through Infrared Polarimetry. , 741(1):59, November 2011.

[51] Robert J. De Rosa, Eric L. Nielsen, Sarah C. Blunt, James R. Graham, Quinn M. Konopacky, Christian Marois, Laurent Pueyo, Julien Rameau, Dominic M. Ryan, Jason J. Wang, Vanessa Bailey, Ashley Chontos, Daniel C. Fabrycky, Katherine B. Follette, Bruce Macintosh, Franck Marchis, S. Mark Ammons, Pauline Arriaga, Jeffrey K. Chilcote, Tara H. Cotten, René Doyon, Gaspard Duchêne, Thomas M. Esposito, Michael P. Fitzgerald, Benjamin Gerard, Stephen J. Goodsell, Alexandra Z. Greenbaum, Pascale Hibon, Patrick Ingraham, Mara Johnson-Groh, Paul G. Kalas, David Lafrenière, Jerome Maire, Stanimir Metchev, Maxwell A. Millar-Blanchaer, Katie M. Morzinski, Rebecca Oppenheimer, Rahul I. Patel, Jennifer L. Patience, Marshall D. Perrin, Abhijith Rajan, Fredrik T. Rantakyrö, Jean-Baptiste Ruffio, Adam C. Schneider, Anand

Sivaramakrishnan, Inseok Song, Debby Tran, Gautam Vasisht, Kimberly Ward-Duong, and Schuyler G. Wolff. Astrometric Confirmation and Preliminary Orbital Parameters of the Young Exoplanet 51 Eridani b with the Gemini Planet Imager. , 814(1):L3, November 2015.

[52] Rafael Ribeiro de Sousa, Alessandro Morbidelli, Sean N. Raymond, Andre Izidoro, Rodney Gomes, and Ernesto Vieira Neto. Dynamical evidence for an early giant planet instability. , 339:113605, March 2020.

[53] N. R. Deacon, Th Henning, and D. E. Kossakowski. Data-driven stellar parameters for southern TESS FGK targets. , 486(1):251–265, June 2019.

[54] Frank Dekens, David Kirkman, Gary A. Chanan, Terry S. Mast, Jerry E. Nelson, Garth Illingworth, and Peter L. Wizinowich. High-speed seeing measurements at the Keck Telescope. In Mark A. Ealey and Fritz Merkle, editors, *Adaptive Optics in Astronomy*, volume 2201 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 310–313, May 1994.

[55] Jacques Désarménien. How to run TEX in french. Technical Report SATN-CS-1013, Computer Science Department, Stanford University, Stanford, California, August 1984.

[56] C. Dessenne, P. Y. Madec, and G. Rousset. Modal prediction forclosed-loop adaptive optics. *Optics Letters*, 22(20):1535–1537, October 1997.

[57] C. Dessenne, P. Y. Madec, and G. Rousset. Sky implementationof modal predictive control in adaptive optics. *Optics Letters*, 24(5):339–341, March 1999.

[58] Caroline Dessenne, Pierre-Yves Madec, and Gérard Rousset. Optimization of a Predictive Controller for Closed-Loop Adaptive Optics. , 37(21):4623–4633, July 1998.

[59] Lauren Doyle, David J. Armstrong, Daniel Bayliss, Toby Rodel, and Vedad Kunovac. The TESS-SPOC FFI target sample explored with Gaia. , 529(2):1802–1813, April 2024.

[60] Dainis Dravins, Lennart Lindegren, Eva Mezey, and Andrew T. Young. Atmospheric Intensity Scintillation of Stars. III. Effects for Different Telescope Apertures. , 110(747):610–633, May 1998.

[61] Thomas F. Droege, Michael W. Richmond, Michael P. Sallman, and Robert P. Creager. TASS Mark IV Photometric Survey of the Northern Sky. , 118(850):1666–1678, December 2006.

[62] Michael P. Egan and Stephan D. Price. The MSX Infrared Astometric Catalog. , 112:2862, December 1996.

[63] C. W. Engelbracht, M. Blaylock, K. Y. L. Su, J. Rho, G. H. Rieke, J. Muzerolle, D. L. Padgett, D. C. Hines, K. D. Gordon, D. Fadda, A. Noriega-Crespo, D. M. Kelly, W. B. Latter, J. L. Hinz, K. A. Misselt, J. E. Morrison, J. A. Stansberry, D. L. Shupe, S. Stolovy, Wm. A. Wheaton, E. T. Young, G. Neugebauer, S. Wachter, P. G. Pérez-González, D. T. Frayer, and F. R. Marleau. Absolute Calibration and Characterization of the Multiband Imaging Photometer for Spitzer. I. The Stellar Calibrator Sample and the 24 μm Calibration. , 119(859):994–1018, September 2007.

[64] Ochsenbein F. et. al. The VizieR database of astronomical catalogues .

[65] Olivier Fauvarque, Benoit Neichel, Thierry Fusco, Jean-Francois Sauvage, and Orion Girault. General formalism for Fourier-based wave front sensing. *Optica*, 3(12):1440, December 2016.

[66] P. N. Fedorov, V. S. Akhmetov, and V. V. Bobylev. Residual rotation of the Hipparcos/Tycho-2 system as determined from the data of the XPM catalogue. , 416(1):403–408, September 2011.

[67] G. J. Ferland, R. L. Porter, P. A. M. van Hoof, R. J. R. Williams, N. P.

Abel, M. L. Lykins, G. Shaw, W. J. Henney, and P. C. Stancil. The 2013 Release of Cloudy. , 49:137–163, Apr 2013.

[68] F. Ferreira, A. Sevin, D. Gratadour, G. Chauvin, E. Diolaiti, R. Gratton, M. Loupias, J. Milli, F. Wildi, and A. Boccaletti. Upgrading SPHERE with the second stage AO system SAXO+: RTC preliminary design. In Kathryn J. Jackson, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems IX*, volume 13097 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 130974M, August 2024.

[69] Katherine B. Follette. An Introduction to High Contrast Differential Imaging of Exoplanets and Disks. , 135(1051):093001, September 2023.

[70] J. Fowler, Sebastiaan Y. Haffert, Maaike A. M. van Kooten, Rico Landman, Alexis Bidot, Adrien Hours, Mamadou N'Diaye, Olivier Absil, Lisa Altinier, Pierre Baudoz, Ruslan Belikov, Markus Johannes Bonse, Kimberly Bott, Bernhard Brandl, Alexis Carlotti, Sarah L. Casewell, Elodie Choquet, Nicolas B. Cowan, Niyati Desai, David Doelman, Kevin Fogarty, Timothy D. Gebhard, Yann Gutierrez, Olivier Guyon, Olivier Herscovici-Schiller, Roser Juanola-Parramon, Matthew Kenworthy, Elina Kleisioti, Lorenzo König, Mariya Krasteva, Iva Laginja, Lucie Leboulleux,

Johan Mazoyer, Maxwell A. Millar-Blanchaer, David Mouillet, Emiel Por, Laurent Pueyo, Frans Snik, Dirk van Dam, Kyle van Gorkom, and Sophia R. Vaughan. Visible extreme adaptive optics on extremely large telescopes: towards detecting oxygen in Proxima Centauri b and analogs. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 12680 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 126801U, October 2023.

[71] J. Fowler, Rebecca Jensen-Clem, Maaike A. M. van Kooten, Vincent Chambouleyron, and Sylvain Cetre. The future looks dark: improving high contrast imaging with hyper-parameter optimization for data-driven predictive wavefront control. In Kathryn J. Jackson, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems IX*, volume 13097 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 130977N, August 2024.

[72] J. Fowler, M. A. M. Van Kooten, and R. Jensen-Clem. Battle of the predictive wavefront controls: comparing data and model-driven predictive control for high contrast imaging. In Laura Schreiber, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems VIII*, volume 12185 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*

Series, page 1218582, August 2022.

[73] Kyle Franson, Brendan P. Bowler, Mariangela Bonavita, Timothy D. Brandt, Minghan Chen, Matthias Samland, Zhoujian Zhang, Anna Lueber, Kevin Heng, Daniel Kitzmann, Trevor Wolf, Brandon A. Jones, Quang H. Tran, Daniella C. Bardalez Gagliuffi, Beth Biller, Jeffrey Chilcote, Justin R. Crepp, Trent J. Dupuy, Jacqueline Faherty, Clémence Fontanive, Tyler D. Groff, Raffaele Gratton, Olivier Guyon, Rebecca Jensen-Clem, Nemanja Jovanovic, N. Jeremy Kasdin, Julien Lozi, Eugene A. Magnier, Koraljka Mužić, Aniket Sanghi, and Christopher A. Theissen. Astrometric Accelerations as Dynamical Beacons: Discovery and Characterization of HIP 21152 B, the First T-dwarf Companion in the Hyades. , 165(2):39, February 2023.

[74] Kyle Franson, Brendan P. Bowler, Yifan Zhou, Tim D. Pearce, Daniella C. Bardalez Gagliuffi, Lauren I. Biddle, Timothy D. Brandt, Justin R. Crepp, Trent J. Dupuy, Jacqueline Faherty, Rebecca Jensen-Clem, Marvin Morgan, Aniket Sanghi, Christopher A. Theissen, Quang H. Tran, and Trevor N. Wolf. Astrometric Accelerations as Dynamical Beacons: A Giant Planet Imaged inside the Debris Disk of the Young Star AF Lep. , 950(2):L19, June 2023.

[75] David Fuchs. The format of TEX's DVI files version 1. *TUGboat*, 2(2):12–16, July 1981.

[76] David Fuchs. Device independent file format. *TUGboat*, 3(2):14–19, October 1982.

[77] Richard K. Furuta and Pierre A. MacKay. Two TEX implementations for the IBM PC. *Dr. Dobb's Journal*, 10(9):80–91, September 1985.

[78] Chanan G., Djorgovski G., Glecker A., Kulkarni S., Mast T., Max C., Nelson J., and Wizonowich P. Keck observatory report no. 208. Technical report, W. M. Keck Observatory, January 1996.

[79] Maxwell Galloway, Theodora Karalidi, Sagnick Mukherjee, Max Millar-Blanchaer, Connor Vancil, and Joseph Harrington. Expected Thermal and Reflected Light Polarization of Three Imaged Exoplanets: Model Spectra and Light Curves for VHS 1256 b, HIP 99770 b and AF Lep b. In *American Astronomical Society Meeting Abstracts #243*, volume 243 of *American Astronomical Society Meeting Abstracts*, page 230.04, February 2024.

[80] E Gendron and P Léna. Astronomical adaptive optics. ii. experimental results of an optimized modal control. *Astronomy and Astrophysics Supplement, v. 111, p. 153*, 111:153, 1995.

[81] Luke Gers. BAMM! An improved warping solver for segmented mirrors. In Heather K. Marshall, Jason Spyromilio, and Tomonori Usuda, editors, *Ground-based and Airborne Telescopes X*, volume 13094 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 130943F, August 2024.

[82] C. Ginski, M. Benisty, R. G. van Holstein, A. Juhász, T. O. B. Schmidt, G. Chauvin, J. de Boer, M. Wilby, C. F. Manara, P. Delorme, F. Ménard, P. Pinilla, T. Birnstiel, M. Flock, C. Keller, M. Kenworthy, J. Milli, J. Olofsson, L. Pérez, F. Snik, and N. Vogt. First direct detection of a polarized companion outside a resolved circumbinary disk around CS Chamaeleonis. , 616:A79, August 2018.

[83] Szymon Gladysz, Julian C. Christou, L. William Bradford, and Lewis C. Roberts. Temporal Variability and Statistics of the Strehl Ratio in Adaptive-Optics Images. , 120(872):1132, October 2008.

[84] P. Goldreich and S. Tremaine. The dynamics of planetary rings. , 20:249–283, January 1982.

[85] Carlos Alberto Gomez Gonzalez, Olivier Wertz, Olivier Absil, Valentin Christiaens, Denis Defrère, Dimitri Mawet, Julien Milli, Pierre-Antoine Absil, Marc Van Droogenbroeck, Faustine Cantalloube, Philip M. Hinz,

Andrew J. Skemer, Mikael Karlsson, and Jean Surdej. VIP: Vortex Image Processing Package for High-contrast Direct Imaging. , 154(1):7, July 2017.

[86] John Green. *Turtles All The Way Down.* Penguin Books, 2019.

[87] Tyler Groff, Jeffrey Chilcote, Timothy Brandt, N. Jeremy Kasdin, Michael Galvin, Craig Loomis, Maxime Rizzo, Gillian Knapp, Olivier Guyon, Nemanja Jovanovic, Julien Lozi, Thayne Currie, Naruhisa Takato, and Masahiko Hayashi. First light of the CHARIS high-contrast integral-field spectrograph. In Stuart Shaklan, editor, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 10400 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1040016, September 2017.

[88] S. B. Gudennavar, S. G. Bubbly, K. Preethi, and Jayant Murthy. A Compilation of Interstellar Column Densities. , 199(1):8, March 2012.

[89] S. F. Gull. Developments in maximum-entropy data analysis. In J. Skilling, editor, *Maximum Entropy and Bayesian Methods*, pages 53–71. Kluwer Academic, Dordrecht, 1989.

[90] Charlotte Guthery, Michael Bottom, Antonin Bouchez, Jacques Delorme, Sonia Karkar, Paul Richards, and Peter Wizinowich. Optimizing Keck

adaptive optics: correcting residual errors with focal-plane wavefront sensing. In Kathryn J. Jackson, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems IX*, volume 13097, page 130975Y. International Society for Optics and Photonics, SPIE, 2024.

[91] Mayra Gutierrez, Steph Sallum, Rebecca Jensen-Clem, Maaike van Kooten, and Charlotte Bond. High-contrast exoplanet imaging at Keck Observatory with SCALES and next generation adaptive optics. In Garreth J. Ruane and Maxwell A. Millar-Blanchaer, editors, *Techniques and Instrumentation for Detection of Exoplanets XII*, volume 13627, page 1362705. International Society for Optics and Photonics, SPIE, 2025.

[92] O. Guyon, E. A. Pluzhnik, M. J. Kuchner, B. Collins, and S. T. Ridgway. Theoretical Limits on Extrasolar Terrestrial Planet Detection with Coronagraphs. , 167(1):81–99, November 2006.

[93] Olivier Guyon and Jared Males. Adaptive Optics Predictive Control with Empirical Orthogonal Functions (EOFs). *arXiv e-prints*, page arXiv:1707.00570, July 2017.

[94] Olivier Guyon, Arnaud Sevin, Florian Ferreira, Hatem Ltaief, Jared Males, Vincent Deo, Damien Gratadour, Sylvain Cetre, Frantz Martinache, Julien Lozi, Sebastien Vievard, Neelay Fruitwala, Steven Bos, and Nour

Skaf. Adaptive optics real-time control with the compute and control for adaptive optics (Cacao) software framework. In Laura Schreiber, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems VII*, volume 11448 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 114482N, December 2020.

[95] Olivier Guyon, Arnaud Sevin, Damien Gratadour, Julien Bernard, Hatem Ltaief, Dalal Sukkari, Sylvain Cetre, Nour Skaf, Julien Lozi, Frantz Martinache, Christophe Clergeon, Barnaby Norris, Alison Wong, and Jared Males. The compute and control for adaptive optics (CACAO) real-time control software package. In Laird M. Close, Laura Schreiber, and Dirk Schmidt, editors, *Adaptive Optics Systems VI*, volume 10703 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 107031E, July 2018.

[96] Christian Bernt Haakonsen and Robert E. Rutledge. XID II: Statistical Cross-Association of ROSAT Bright Source Catalog X-ray Sources with 2MASS Point Source Catalog Near-Infrared Sources. , 184(1):138–151, September 2009.

[97] Rehan Hafeez, Finn Archinuk, Sébastien Fabbro, Hossen Teimoorinia, and Jean-Pierre Véran. Forecasting wavefront corrections in an adaptive

optics system. *Journal of Astronomical Telescopes, Instruments, and Systems*, 8:029003, April 2022.

[98] S. Y. Haffert, R. G. van Holstein, C. Ginski, J. Brinchmann, I. A. G. Snellen, J. Milli, T. Stolker, C. U. Keller, and J. Girard. CS Cha B: A disc-obscured M-type star mimicking a polarised planetary companion. , 640:L12, August 2020.

[99] Sebastiaan Y. Haffert, Jared R. Males, Laird M. Close, Kyle van Gorkom, Joseph D. Long, Alexander D. Hedglen, Lauren Schatz, Jennifer Lumbres, Alexander Rodack, Justin M. Knight, He Sun, Kevin Fogarty, and Logan Pearce. Data-driven subspace predictive control: lab demonstration and future outlook. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 11823 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 118231C, September 2021.

[100] R. J. Hanisch and C. D. Biemesderfer. T$_E$X and LAT$_E$X Macro Definition Files for Astronomical Publications. In , page 780, Mar 1989.

[101] K. M. Hanson. Introduction to Bayesian image analysis. In M.Ħ. Loew, editor, *Medical Imaging: Image Processing*, volume 1898 of *Proc. SPIE*, pages 716–731, Feb. 1993. [doi:10.1117/12.154577].

[102] Caleb K. Harada, Courtney D. Dressing, Stephen R. Kane, and Bahareh Adami Ardestani. Setting the Stage for the Search for Life with the Habitable Worlds Observatory: Properties of 164 Promising Planet-survey Targets. , 272(2):30, June 2024.

[103] Kevin K. Hardegree-Ullman, Dániel Apai, Galen J. Bergsten, Ilaria Pascucci, and Mercedes López-Morales. Bioverse: A Comprehensive Assessment of the Capabilities of Extremely Large Telescopes to Probe Earth-like O_2 Levels in Nearby Transiting Habitable-zone Exoplanets. , 165(6):267, June 2023.

[104] John W. Hardy. *Adaptive Optics for Astronomical Telescopes*. 1998.

[105] A. Harris, J. J. Sluss, Jr., H. H. Refai, and P. G. LoPresti. Free-space optical wavelength diversity scheme for fog migration in a ground-to-unmanned-aerial-vehicle communications link. *Opt. Eng.*, 45:086001, 2006. [doi:10.1117/1.2338565].

[106] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-

Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020.

[107] Tibshirani Robert Hastie, Trevor and Jerome Friedman. *The Elements of Statistical Learning, Data Mining, Inference, and Prediction, second edition.* 01 2017.

[108] Arne A. Henden, Stephen Levine, Dirk Terrell, and Douglas L. Welch. APASS - The Latest Data Release. In *American Astronomical Society Meeting Abstracts #225*, volume 225 of *American Astronomical Society Meeting Abstracts*, page 336.16, January 2015.

[109] Herodotus. *The Histories.* Harvard University Press, 1920.

[110] Michael Hess, Craig E. Nance, John W. Vause, Michael Hrynevych, Mark R. Swain, and Mark Colavita. Strategy for identifying and mitigating facility vibrations to improve optical performance at the W.M. Keck Observatory. In Jacobus M. Oschmann and Larry M. Stepp, editors, *Large Ground-based Telescopes*, volume 4837 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 342–351, February 2003.

[111] Naoyuki Hirata, Ryuji Morishima, Keiji Ohtsuki, and Akiko M. Nakamura.

Disruption of Saturn's ring particles by thermal stress. , 378:114919, May 2022.

[112] Chih-Chun Hsu, Jason J. Wang, Jerry W. Xuan, Yapeng Zhang, Jean-Baptiste Ruffio, Dimitri Mawet, Luke Finnerty, Katelyn Horstman, Julianne Cronin, Yinzi Xin, Ben Sappey, Daniel Echeverri, Nemanja Jovanovic, Ashley Baker, Randall Bartos, Geoffrey A. Blake, Benjamin Calvin, Sylvain Cetre, Jacques-Robert Delorme, Gregory W. Doppmann, Michael P. Fitzgerald, Quinn M. Konopacky, Joshua Liberman, Ronald A. López, Evan Morris, Jacklyn Pezzato, Tobias Schofield, Andrew Skemer, J. Kent Wallace, and Ji Wang. Distinct Rotational Evolution of Giant Planets and Brown Dwarf Companions. , 171(4):224, April 2026.

[113] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.

[114] S. Hunziker, H. M. Schmid, D. Mouillet, J. Milli, A. Zurlo, P. Delorme, L. Abe, H. Avenhaus, A. Baruffolo, A. Bazzon, A. Boccaletti, P. Baudoz, J. L. Beuzit, M. Carbillet, G. Chauvin, R. Claudi, A. Costille, J.-B. Daban, S. Desidera, K. Dohlen, C. Dominik, M. Downing, N. Engler, M. Feldt, T. Fusco, C. Ginski, D. Gisler, J. H. Girard, R. Gratton, Th. Henning, N. Hubin, M. Kasper, C. U. Keller, M. Langlois, E. Lagadec,

P. Martinez, A. L. Maire, F. Menard, M. R. Meyer, A. Pavlov, J. Pragt, P. Puget, S. P. Quanz, E. Rickman, R. Roelfsema, B. Salasnich, J.-F. Sauvage, R. Siebenmorgen, E. Sissa, F. Snik, M. Suarez, J. Szulágyi, Ch. Thalmann, M. Turatto, S. Udry, R. G. van Holstein, A. Vigan, and F. Wildi. RefPlanets: Search for reflected light from extrasolar planets with SPHERE/ZIMPOL. , 634:A69, February 2020.

[115] T. Jayasinghe, C. S. Kochanek, K. Z. Stanek, B. J. Shappee, T. W.-S. Holoien, Todd A. Thompson, J. L. Prieto, Subo Dong, M. Pawlak, J. V. Shields, G. Pojmanski, S. Otero, C. A. Britt, and D. Will. The ASAS-SN catalogue of variable stars I: The Serendipitous Survey. , 477(3):3145–3163, July 2018.

[116] Rebecca Jensen-Clem, Charlotte Z. Bond, Sylvain Cetre, Eden McEwen, Peter Wizinowich, Sam Ragland, Dimitri Mawet, and James Graham. Demonstrating predictive wavefront control with the Keck II near-infrared pyramid wavefront sensor. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 11117 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 111170W, September 2019.

[117] Rebecca Jensen-Clem, Charlotte Z. Bond, Sylvain Cetre, Eden McEwen,

Peter Wizinowich, Sam Ragland, Dimitri Mawet, and James Graham. Demonstrating predictive wavefront control with the Keck II near-infrared pyramid wavefront sensor. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 11117 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 111170W, September 2019.

[118] Rebecca Jensen-Clem, Daren Dillon, Benjamin Gerard, M. A. M. van Kooten, J. Fowler, Renate Kupke, Sylvain Cetre, Dominic Sanchez, Philip Hinz, Cesar Laguna, David Doelman, and Frans Snik. The Santa Cruz Extreme AO Lab (SEAL): design and first light. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 11823 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 118231D, September 2021.

[119] Rebecca Jensen-Clem, Daren Dillon, Benjamin Gerard, M. A. M. van Kooten, J. Fowler, Renate Kupke, Sylvain Cetre, Dominic Sanchez, Philip Hinz, Cesar Laguna, David Doelman, and Frans Snik. The Santa Cruz Extreme AO Lab (SEAL): design and first light. In Stuart B. Shaklan and Garreth J. Ruane, editors, *Techniques and Instrumentation for Detection of Exoplanets X*, volume 11823 of *Society of Photo-Optical*

Instrumentation Engineers (SPIE) Conference Series, page 118231D, September 2021.

[120] Rebecca Jensen-Clem, Maxwell A. Millar-Blanchaer, Rob G. van Holstein, Dimitri Mawet, James Graham, Sujan Sengupta, Mark S. Marley, Frans Snik, Arthur Vigan, Sasha Hinkley, Jos de Boer, Julien H. Girard, Robert J. De Rosa, Brendan P. Bowler, Sloane J. Wiktorowicz, Marshall D. Perrin, Justin R. Crepp, and Bruce Macintosh. A Search for Polarized Thermal Emission from Directly Imaged Exoplanets and Brown Dwarf Companions to Nearby Stars. , 160(6):286, December 2020.

[121] Erik M. Johansson, D. Scott Acton, Jong R. An, Kenneth Avicola, Bart V. Beeman, James M. Brase, Carmen J. Carrano, John Gathright, Donald T. Gavel, Randall L. Hurd, Olivier Lai, William Lupton, Bruce A. Macintosh, Claire E. Max, Scot S. Olivier, J. Christopher Shelton, Paul J. Stomski, Kevin Tsubota, Kenneth E. Waltjen, James A. Watson, and Peter L. Wizinowich. Initial performance of the Keck AO wavefront controller system. In Peter L. Wizinowich, editor, *Adaptive Optical Systems Technology*, volume 4007 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 600–607, July 2000.

[122] Erik M. Johansson, Marcos A. van Dam, Paul J. Stomski, James M. Bell,

Jason C. Chin, Roger C. Sumner, Peter L. Wizinowich, Roberto Biasi, Mario Andrighettoni, and Dietrich Pescoller. Upgrading the Keck AO wavefront controllers. In Norbert Hubin, Claire E. Max, and Peter L. Wizinowich, editors, *Adaptive Optics Systems*, volume 7015 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 70153E, July 2008.

[123] H. L. Johnson, R. I. Mitchell, B. Iriarte, and W. Z. Wisniewski. UBVRIJKL Photometry of the Bright Stars. *Communications of the Lunar and Planetary Laboratory*, 4:99–110, January 1966.

[124] Henrik Jönsson, Jon A. Holtzman, Carlos Allende Prieto, Katia Cunha, D. A. García-Hernández, Sten Hasselquist, Thomas Masseron, Yeisson Osorio, Matthew Shetrone, Verne Smith, Guy S. Stringfellow, Dmitry Bizyaev, Bengt Edvardsson, Steven R. Majewski, Szabolcs Mészáros, Diogo Souto, Olga Zamora, Rachael L. Beaton, Jo Bovy, John Donor, Marc H. Pinsonneault, Vijith Jacob Poovelil, and Jennifer Sobeck. APOGEE Data and Spectral Analysis from SDSS Data Release 16: Seven Years of Observations Including First Results from APOGEE-South. , 160(3):120, September 2020.

[125] Rudolph Emil Kalman. A new approach to linear filtering and prediction

problems. *Transactions of the ASME–Journal of Basic Engineering*, 82(Series D):35–45, 1960.

[126] Theodora Karalidi, Dániel Apai, Mark S. Marley, and Esther Buenzli. Maps of Evolving Cloud Structures in Luhman 16AB from HST Time-resolved Spectroscopy. , 825(2):90, July 2016.

[127] N. V. Kharchenko. All-sky compiled catalogue of 2.5 million stars. *Kinematika i Fizika Nebesnykh Tel*, 17(5):409–423, October 2001.

[128] J. Davy Kirkpatrick, Kendra Kellogg, Adam C. Schneider, Sergio Fajardo-Acosta, Michael C. Cushing, Jennifer Greco, Gregory N. Mace, Christopher R. Gelino, Edward L. Wright, Peter R. M. Eisenhardt, Daniel Stern, Jacqueline K. Faherty, Scott S. Sheppard, George B. Lansbury, Sarah E. Logsdon, Emily C. Martin, Ian S. McLean, Steven D. Schurr, Roc M. Cutri, and Tim Conrow. The AllWISE Motion Survey, Part 2. , 224(2):36, June 2016.

[129] J. Davy Kirkpatrick, Federico Marocco, Christopher R. Gelino, Yadukrishna Raghu, Jacqueline K. Faherty, Daniella C. Bardalez Gagliuffi, Steven D. Schurr, Kevin Apps, Adam C. Schneider, Aaron M. Meisner, Marc J. Kuchner, Dan Caselden, R. L. Smart, S. L. Casewell, Roberto Raddi, Aurora Kesseli, Nikolaj Stevnbak Andersen, Edoardo Antonini,

Paul Beaulieu, Thomas P. Bickle, Martin Bilsing, Raymond Chieng, Guillaume Colin, Sam Deen, Alexandru Dereveanco, Katharina Doll, Hugo A. Durantini Luca, Anya Frazer, Jean Marc Gantier, Léopold Gramaize, Kristin Grant, Leslie K. Hamlet, Hiro Higashimura, Michiharu Hyogo, Peter A. Jałowiczor, Alexander Jonkeren, Martin Kabatnik, Frank Kiwy, David W. Martin, Marianne N. Michaels, William Pendrill, Celso Pessanha Machado, Benjamin Pumphrey, Austin Rothermich, Rebekah Russwurm, Arttu Sainio, John Sanchez, Fyodor Theo Sapelkin-Tambling, Jörg Schümann, Karl Selg-Mann, Harshdeep Singh, Andres Stenner, Guoyou Sun, Christopher Tanner, Melina Thévenot, Maurizio Ventura, Nikita V. Voloshin, Jim Walla, Zbigniew Wędracki, Jose I. Adorno, Christian Aganze, Katelyn N. Allers, Hunter Brooks, Adam J. Burgasser, Emily Calamari, Thomas Connor, Edgardo Costa, Peter R. Eisenhardt, Jonathan Gagné, Roman Gerasimov, Eileen C. Gonzales, Chih-Chun Hsu, Rocio Kiman, Guodong Li, Ryan Low, Eric Mamajek, Blake M. Pantoja, Mark Popinchalk, Jon M. Rees, Daniel Stern, Genaro Suárez, Christopher Theissen, Chao-Wei Tsai, Johanna M. Vos, David Zurek, and The Backyard Worlds: Planet 9 Collaboration. The Initial Mass Function Based on the Full-sky 20 pc Census of ∼3600 Stars and Brown Dwarfs. , 271(2):55, April 2024.

[130] Donald E. Knuth. The WEB system for structured documentation, version 2.3. Technical Report STAN-CS-83-980, Computer Science Department, Stanford University, Stanford, California, September 1983.

[131] Donald E. Knuth. *The TEX Book*. Addison-Wesley, Reading, Massachusetts, 1984. Reprinted as Vol. A of *Computers & Typesetting*, 1986.

[132] Donald E. Knuth. Literate programming. *The Computer Journal*, 27(2):97–111, May 1984.

[133] Donald E. Knuth. A torture test for TEX, version 1.3. Technical Report STAN-CS-84-1027, Computer Science Department, Stanford University, Stanford, California, November 1984.

[134] Donald E. Knuth. *TEX: The Program*, volume B of *Computers & Typesetting*. Addison-Wesley, Reading, Massachusetts, 1986.

[135] A. Kolmogorov. The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds' Numbers. *Akademiia Nauk SSSR Doklady*, 30:301–305, January 1941.

[136] Ravi Kumar Kopparapu, Ramses Ramirez, James F. Kasting, Vincent Eymet, Tyler D. Robinson, Suvrath Mahadevan, Ryan C. Terrien, Shawn

Domagal-Goldman, Victoria Meadows, and Rohit Deshpande. Habitable Zones around Main-sequence Stars: New Estimates. , 765(2):131, March 2013.

[137] Ravi kumar Kopparapu, Eric T. Wolf, Jacob Haqq-Misra, Jun Yang, James F. Kasting, Victoria Meadows, Ryan Terrien, and Suvrath Mahadevan. The Inner Edge of the Habitable Zone for Synchronously Rotating Planets around Low-mass Stars Using General Circulation Models. , 819(1):84, March 2016.

[138] J. R. Kuhn, D. Potter, and B. Parise. Imaging Polarimetric Observations of a New Circumstellar Disk System. , 553(2):L189–L191, June 2001.

[139] Caroline Kulcsár, Henri-François Raynaud, Cyril Petit, Jean-Marc Conan, and Patrick Viaris de Lesegno. Optimal control, observers and integrators in adaptive optics. *Optics Express*, 14(17):7464, August 2006.

[140] Masayuki Kuzuhara, Thayne Currie, Takuya Takarada, Timothy D. Brandt, Bun'ei Sato, Taichi Uyama, Markus Janson, Jeffrey Chilcote, Taylor Tobin, Kellen Lawson, Yasunori Hori, Olivier Guyon, Tyler D. Groff, Julien Lozi, Sebastien Vievard, Ananya Sahoo, Vincent Deo, Nemanja Jovanovic, Kyohoon Ahn, Frantz Martinache, Nour Skaf, Eiji Akiyama, Barnaby R. Norris, Mickaël Bonnefoy, Krzysztof G. Hełminiak,

Tomoyuki Kudo, Michael W. McElwain, Matthias Samland, Kevin Wagner, John Wisniewski, Gillian R. Knapp, Jungmi Kwon, Jun Nishikawa, Eugene Serabyn, Masahiko Hayashi, and Motohide Tamura. Direct-imaging Discovery and Dynamical Mass of a Substellar Companion Orbiting an Accelerating Hyades Sun-like Star with SCExAO/CHARIS. , 934(2):L18, August 2022.

[141] P. J. M. Laarhoven. *Simulated annealing: Theory and applications*. 1987.

[142] L. Lamport. *LaTeX: A Document Preparation System*. Addison-Wesley Professional, 2 edition, 1994.

[143] Leslie Lamport. *LaTeX: A Document Preparation System. User's Guide and Reference Manual*. Addison-Wesley, Reading, Massachusetts, 1986.

[144] Peter Lancaster and Leiba Rodman. *Algebraic Riccati Equations*. 01 2002.

[145] R. Landman, S. Y. Haffert, V. M. Radhakrishnan, and C. U. Keller. Self-optimizing adaptive optics control with reinforcement learning. In Laura Schreiber, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems VII*, volume 11448 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1144849, December 2020.

[146] Rico Landman, Sebastiaan Y. Haffert, Vikram M. Radhakrishnan, and Christoph U. Keller. Self-optimizing adaptive optics control with reinforcement learning for high-contrast imaging. *Journal of Astronomical Telescopes, Instruments, and Systems*, 7:039002, July 2021.

[147] Barry M. Lasker, Mario G. Lattanzi, Brian J. McLean, Beatrice Bucciarelli, Ronald Drimmel, Jorge Garcia, Gretchen Greene, Fabrizia Guglielmetti, Christopher Hanley, George Hawkins, Victoria G. Laidler, Charles Loomis, Michael Meakes, Roberto Mignani, Roberto Morbidelli, Jane Morrison, Renato Pannunzio, Amy Rosenberg, Maria Sarasso, Richard L. Smart, Alessandro Spagna, Conrad R. Sturch, Antonio Volpicelli, Richard L. White, David Wolfe, and Andrea Zacchei. The Second-Generation Guide Star Catalog: Description and Properties. , 136(2):735–766, August 2008.

[148] Kellen Lawson, Thayne Currie, John P. Wisniewski, Jun Hashimoto, Olivier Guyon, Jeremy Kasdin N, Tyler D. Groff, Julien Lozi, Timothy D. Brandt, Jeffrey Chilcote, Vincent Deo, Taichi Uyama, and Sebastien Vievard. High-contrast integral field spectropolarimetry of planet-forming disks with SCExAO/CHARIS. In Stuart B. Shaklan and Garreth J. Ruane, editors, *Techniques and Instrumentation for Detection*

of Exoplanets X, volume 11823, page 118230D. International Society for Optics and Photonics, SPIE, 2021.

[149] Lucie Leboulleux, Olivier Guyon, Julien Lozi, Miles Lucas, Sébastien Vievard, Vincent Deo, Kyohoon Ahn, Alexis Carlotti, Alain Delboulbé, Mamadou N'Diaye, Marie-Hélène Sztefek, and Philippe Delorme. Coronagraphy robust to segmented aberrations: low-wind effect mitigation on Subaru/SCExAO. In Kathryn J. Jackson, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems IX*, volume 13097 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 130973G, August 2024.

[150] Leping Li, Jun Zhang, Hardi Peter, Lakshmi Pradeep Chitta, Jiangtao Su, Hongqiang Song, Chun Xia, and Yijun Hou. Quasi-periodic Fast Propagating Magnetoacoustic Waves during the Magnetic Reconnection Between Solar Coronal Loops. , 868(2):L33, Dec 2018.

[151] S. Lilley. Keck adaptive optics note 1463 ao sodium notch filter and beamsplitter performance. Technical report, W. M. Keck Observatory, March 2024.

[152] Xuewen Liu, Tim Morris, Chris Saunter, Francisco Javier de Cos Juez, Carlos González-Gutiérrez, and Lisa Bardou. Wavefront prediction using

artificial neural networks for open-loop adaptive optics. , 496(1):456–464, July 2020.

[153] Jorge Llop-Sayson, Charles Beichman, Geoffrey Bryden, Marie Ygouf, Andras Gaspar, William Thompson, Aniket Sanghi, Dimitri Mawet, Tiffany Meshkat, Alexandra Z. Greenbaum, Jarron Leisenring, Schuyler Wolff, Marcia Rieke, and George Rieke. Searching for Planets Orbiting Eridani with JWST/NIRCam. , 170(4):229, October 2025.

[154] Julien Lozi, Kyohoon Ahn, Hannah Blue, Alicia Chun, Christophe Clergeon, Vincent Deo, Olivier Guyon, Takashi Hattori, Yosuke Minowa, Shogo Nishiyama, Yoshito Ono, Shin Oya, Yuhei Takagi, Sebastien Vievard, and Maria Vincent. AO3k at Subaru: First on-sky results of the facility extreme-AO. *arXiv e-prints*, page arXiv:2407.19188, July 2024.

[155] Miles Lucas, Michael Bottom, Olivier Guyon, Julien Lozi, Barnaby Norris, Vincent Deo, Sébastien Vievard, Kyohoon Ahn, Nour Skaf, and Peter Tuthill. A visible-light Lyot coronagraph for SCExAO/VAMPIRES. In Christopher J. Evans, Julia J. Bryant, and Kentaro Motohara, editors, *Ground-based and Airborne Instrumentation for Astronomy IX*, volume 12184 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 121844E, August 2022.

[156] Miles Lucas, Barnaby Norris, Olivier Guyon, Michael Bottom, Vincent Deo, Sébastien Vievard, Julien Lozi, Kyohoon Ahn, Jaren Ashcraft, Thayne Currie, David Doelman, Tomoyuki Kudo, Lucie Leboulleux, Lucinda Lilley, Maxwell Millar-Blanchaer, Boris Safonov, Peter Tuthill, Taichi Uyama, Aidan Walk, and Manxuan Zhang. Visible-light High-contrast Imaging and Polarimetry with SCExAO/VAMPIRES. , 136(11):114504, November 2024.

[157] R. Earle Luck. Abundances in the Local Region II: F, G, and K Dwarfs and Subgiants. , 153(1):21, January 2017.

[158] Michael B. Lund. Astrology in the era of exoplanets. 2016.

[159] Michael B. Lund. Detecting the ultimate power in the universe with lsst. 2017.

[160] Michael B. Lund, Robert J. Siverd, and Ponder Stibbons. Independent discovery of a sub-earth in the habitable zone around a very close solar-mass star. 2018.

[161] B. Macintosh, J. R. Graham, T. Barman, R. J. De Rosa, Q. Konopacky, M. S. Marley, C. Marois, E. L. Nielsen, L. Pueyo, A. Rajan, J. Rameau, D. Saumon, J. J. Wang, J. Patience, M. Ammons, P. Arriaga, E. Artigau, S. Beckwith, J. Brewster, S. Bruzzone, J. Bulger, B. Burningham, A. S.

Burrows, C. Chen, E. Chiang, J. K. Chilcote, R. I. Dawson, R. Dong, R. Doyon, Z. H. Draper, G. Duchêne, T. M. Esposito, D. Fabrycky, M. P. Fitzgerald, K. B. Follette, J. J. Fortney, B. Gerard, S. Goodsell, A. Z. Greenbaum, P. Hibon, S. Hinkley, T. H. Cotten, L.-W. Hung, P. Ingraham, M. Johnson-Groh, P. Kalas, D. Lafreniere, J. E. Larkin, J. Lee, M. Line, D. Long, J. Maire, F. Marchis, B. C. Matthews, C. E. Max, S. Metchev, M. A. Millar-Blanchaer, T. Mittal, C. V. Morley, K. M. Morzinski, R. Murray-Clay, R. Oppenheimer, D. W. Palmer, R. Patel, M. D. Perrin, L. A. Poyneer, R. R. Rafikov, F. T. Rantakyrö, E. L. Rice, P. Rojo, A. R. Rudy, J.-B. Ruffio, M. T. Ruiz, N. Sadakuni, L. Saddlemyer, M. Salama, D. Savransky, A. C. Schneider, A. Sivaramakrishnan, I. Song, R. Soummer, S. Thomas, G. Vasisht, J. K. Wallace, K. Ward-Duong, S. J. Wiktorowicz, S. G. Wolff, and B. Zuckerman. Discovery and spectroscopy of the young jovian planet 51 Eri b with the Gemini Planet Imager. *Science*, 350(6256):64–67, October 2015.

[162] Eugene. A. Magnier, Edward. F. Schlafly, Douglas P. Finkbeiner, J. L. Tonry, B. Goldman, S. Röser, E. Schilbach, S. Casertano, K. C. Chambers, H. A. Flewelling, M. E. Huber, P. A. Price, W. E. Sweeney, C. Z. Waters, L. Denneau, P. W. Draper, K. W. Hodapp, R. Jedicke, N. Kaiser, R.-P. Kudritzki, N. Metcalfe, C. W. Stubbs, and R. J. Wainscoat. Pan-STARRS

Photometric and Astrometric Calibration. , 251(1):6, November 2020.

[163] Jared R Males and Olivier Guyon. Ground-based adaptive optics coronagraphic performance under closed-loop predictive control. *Journal of Astronomical Telescopes, Instruments, and Systems*, 4(1):019001–019001, 2018.

[164] A. Mallama. Sloan Magnitudes for the Brightest Stars. , 42(2):443, December 2014.

[165] A. Mallama. Erratum: Sloan Magnitudes for the Brightest Stars. , 46(1):90, June 2018.

[166] Elena Manjavacas, Theodora Karalidi, Xianyu Tan, Johanna M. Vos, Ben W. P. Lew, Beth A. Biller, and Natalia Oliveros-Gómez. Top-of-the-atmosphere and Vertical Cloud Structure of a Fast-rotating Late T Dwarf. , 164(2):65, August 2022.

[167] Christian Marois, David Lafrenière, René Doyon, Bruce Macintosh, and Daniel Nadeau. Angular Differential Imaging: A Powerful High-Contrast Imaging Technique. , 641(1):556–564, April 2006.

[168] Lucas Marquis, Henri-François Raynaud, Nicolas Galland, José Marco de La Rosa, Icíar Montilla, Óscar Tubío Araújo, Marcos Reyes García-

Talavera, Gianluca Lombardi, Manuel Huertas Lopez, Daniel Reverte, and Caroline Kulcsár. First on-sky tests of LQG control for a 10m-class telescope: prelude on the Gran Telescopio Canarias adaptive optics system. In Kathryn J. Jackson, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems IX*, volume 13097 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 130977Y, August 2024.

[169] Michel Mayor and Didier Queloz. A Jupiter-mass companion to a solar-type star. , 378(6555):355–359, November 1995.

[170] L. C. Mayorga, E. M. May, J. Lustig-Yaeger, and S. E Moran. Detection of rotational variability in floofy objects at optical wavelengths. 2021.

[171] I. McDonald, A. A. Zijlstra, and R. A. Watson. Fundamental parameters and infrared excesses of Tycho-Gaia stars. , 471(1):770–791, October 2017.

[172] Ilija Medan, Sébastien Lépine, and Zachary Hartman. Bayesian Cross-matching of High Proper-motion Stars in Gaia DR2 and Photometric Metallicities for 1.7 million K and M Dwarfs. , 161(5):234, May 2021.

[173] D. Mesa, R. Gratton, P. Kervella, M. Bonavita, S. Desidera, V. D'Orazi, S. Marino, A. Zurlo, and E. Rigliaco. AF Lep b: The lowest-mass planet

detected by coupling astrometric and direct imaging data. , 672:A93, April 2023.

[174] T. Meshkat, M. Bonnefoy, E. E. Mamajek, S. P. Quanz, G. Chauvin, M. A. Kenworthy, J. Rameau, M. R. Meyer, A.-M. Lagrange, J. Lannier, and P. Delorme. Discovery of a low-mass companion to the F7V star HD 984. , 453(3):2378–2386, November 2015.

[175] P. A. Miles-Páez, M. R. Zapatero Osorio, E. Pallé, and K. Peña Ramírez. Linear polarization of rapidly rotating ultracool dwarfs. , 556:A125, August 2013.

[176] P. A. Miles-Páez, M. R. Zapatero Osorio, E. Pallé, and K. Peña Ramírez. Optical and near-infrared linear polarization of low and intermediate-gravity ultracool dwarfs. , 466(3):3184–3198, April 2017.

[177] Maxwell A. Millar-Blanchaer, Julien H. Girard, Theodora Karalidi, Mark S. Marley, Rob G. van Holstein, Sujan Sengupta, Dimitri Mawet, Tiffany Kataria, Frans Snik, Jos de Boer, Rebecca Jensen-Clem, Arthur Vigan, and Sasha Hinkley. Detection of Polarization due to Cloud Bands in the Nearby Luhman 16 Brown Dwarf Binary. , 894(1):42, May 2020.

[178] Maxwell A. Millar-Blanchaer, Eliad Peretz, Peter Wizinowich, Eduardo Marin, J. Fowler, Jack Grossman, Bert Pasquale, Sam Ragland, Shui

Kwok, Scott Lilley, Jacques-Robert Delorme, Kevin Tsubota, and Edward Wetherell. First light with ORKID, a pathfinder diffraction-limited visible-light camera at the WM Keck Observatory. In Julia J. Bryant, Kentaro Motohara, and Joël. R. D. Vernet, editors, *Ground-based and Airborne Instrumentation for Astronomy X*, volume 13096 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 130961D, July 2024.

[179] J. Milli, D. Mouillet, T. Fusco, J. H. Girard, E. Masciadri, E. Pena, J. F. Sauvage, C. Reyes, K. Dohlen, J. L. Beuzit, M. Kasper, M. Sarazin, and F. Cantalloube. Performance of the extreme-AO instrument VLT/-SPHERE and dependence on the atmospheric conditions. *arXiv e-prints*, page arXiv:1710.05417, October 2017.

[180] Ashai Moreno, Vincent Chambouleyron, Rebecca M. Jensen-Clem, Daren Dillon, Phil M. Hinz, and Bruce Macintosh. Implementation and characterization of the vector vortex coronagraph on the SEAL testbed. In Kathryn J. Jackson, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems IX*, volume 13097 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1309746, August 2024.

[181] A. Moro-Martín, J. P. Marshall, G. Kennedy, B. Sibthorpe, B. C. Matthews, C. Eiroa, M. C. Wyatt, J.-F. Lestrade, J. Maldonado, D. Rodriguez, J. S. Greaves, B. Montesinos, A. Mora, M. Booth, G. Duchêne, D. Wilner, and J. Horner. Does the Presence of Planets Affect the Frequency and Properties of Extrasolar Kuiper Belts? Results from the Herschel Debris and Dunes Surveys. , 801(2):143, March 2015.

[182] M. Mugrauer. Search for stellar companions of exoplanet host stars by exploring the second ESA-Gaia data release. , 490(4):5088–5102, December 2019.

[183] Sagnick Mukherjee, Jonathan J. Fortney, Rebecca Jensen-Clem, Xianyu Tan, Mark S. Marley, and Natasha E. Batalha. Modeling Polarization Signals from Cloudy Brown Dwarfs Luhman 16 A and B in Three Dimensions. , 923(1):113, December 2021.

[184] Elijah Mullens and Nikole K. Lewis. Silicate Sundogs: Probing the Effects of Grain Directionality in Exoplanet Observations. , 988(2):L43, August 2025.

[185] A. Müller, M. Keppler, Th. Henning, M. Samland, G. Chauvin, H. Beust, A.-L. Maire, K. Molaverdikhani, R. van Boekel, M. Benisty, A. Boccaletti, M. Bonnefoy, F. Cantalloube, B. Charnay, J.-L. Baudino, M. Gennaro,

Z. C. Long, A. Cheetham, S. Desidera, M. Feldt, T. Fusco, J. Girard, R. Gratton, J. Hagelberg, M. Janson, A.-M. Lagrange, M. Langlois, C. Lazzoni, R. Ligi, F. Ménard, D. Mesa, M. Meyer, P. Mollière, C. Mordasini, T. Moulin, A. Pavlov, N. Pawellek, S. P. Quanz, J. Ramos, D. Rouan, E. Sissa, E. Stadler, A. Vigan, Z. Wahhaj, L. Weber, and A. Zurlo. Orbital and atmospheric characterization of the planet within the gap of the PDS 70 transition disk. , 617:L2, September 2018.

[186] Ruth A. Murray-Clay and Eugene I. Chiang. Brownian Motion in Planetary Migration. , 651(2):1194–1208, November 2006.

[187] J. R. Myers, C. B. Sande, A. C. Miller, W. H. Warren, Jr., and D. A. Tracewell. VizieR Online Data Catalog: SKY2000 Master Catalog, Version 5 (Myers+ 2006). VizieR On-line Data Catalog: V/145. Originally published in: Goddard Space Flight Center, Flight Dynamics Division (2006), February 2015.

[188] V. Nascimbeni, G. Piotto, S. Ortolani, G. Giuffrida, P. M. Marrese, D. Magrin, R. Ragazzoni, I. Pagano, H. Rauer, J. Cabrera, D. Pollacco, A. M. Heras, M. Deleuil, L. Gizon, and V. Granata. An all-sky catalogue of solar-type dwarfs for exoplanetary transit surveys. , 463(4):4210–4222, December 2016.

[189] G. Neugebauer and R. B. Leighton. *Two-micron sky survey. A preliminary catalogue*. 1969.

[190] Christopher Neyman. Keck adaptive optics note 303: Atmospheric parameters for mauna kea. Technical report, W. M. Keck Observatory, December 2004.

[191] Christopher Neyman and Ralph Flicker. Keck adapative optics note 482: Keck telescope wavefront errors: Implications for ngao. Technical report, W. M. Keck Observatory, March 2007.

[192] J. Nousiainen, C. Rajani, M. Kasper, T. Helin, S. Y. Haffert, C. Vérinaud, J. R. Males, K. Van Gorkom, L. M. Close, J. D. Long, A. D. Hedglen, O. Guyon, L. Schatz, M. Kautz, J. Lumbres, A. Rodack, J. M. Knight, and K. Miller. Toward on-sky adaptive optics control using reinforcement learning. Model-based policy optimization for adaptive optics. , 664:A71, August 2022.

[193] Jalo Nousiainen, Byron Engler, Markus Kasper, Tapio Helin, Cédric T. Heritier, and Chang Rajani. Advances in model-based reinforcement learning for adaptive optics control. In Laura Schreiber, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems VIII*, volume 12185 of

Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, page 121852O, August 2022.

[194] Jalo Nousiainen, Byron Engler, Markus Kasper, Chang Rajani, Tapio Helin, Cédric T. Heritier, Sascha P. Quanz, and Adrian M. Glauser. Laboratory experiments of model-based reinforcement learning for adaptive optics control. *Journal of Astronomical Telescopes, Instruments, and Systems*, 10:019001, January 2024.

[195] Jalo Nousiainen, Chang Rajani, Markus Kasper, and Tapio Helin. Adaptive optics control using model-based reinforcement learning. *Optics Express*, 29(10):15327, May 2021.

[196] Eran O. Ofek. Calibrated griz Magnitudes of Tycho Stars: All-Sky Photometric Calibration Using Bright Stars. , 120(872):1128, October 2008.

[197] A. V. Oppenheim and R. W. Schafer. Digital signal processing. *NASA STI/Recon Technical Report A*, 76:11184, January 1975.

[198] Shin Oya, Aurelien Bouvier, Olivier Guyon, Makoto Watanabe, Yutaka Hayano, Hideki Takami, Masanori Iye, Masayuki Hattori, Yoshihiko Saito, Meguru Itoh, Stephen Colley, Matthew Dinkins, Michael Eldred, and Taras Golota. Performance of the deformable mirror for Subaru

LGSAO. In Brent L. Ellerbroek and Domenico Bonaccini Calia, editors, *Advances in Adaptive Optics II*, volume 6272 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 62724S, June 2006.

[199] Martin Paegert, Keivan G. Stassun, Karen A. Collins, Joshua Pepper, Guillermo Torres, Jon Jenkins, Joseph D. Twicken, and David W. Latham. TESS Input Catalog versions 8.1 and 8.2: Phantoms in the 8.0 Catalog and How to Handle Them. *arXiv e-prints*, page arXiv:2108.04778, August 2021.

[200] G. V. Panopoulou, L. Markopoulioti, F. Bouzelou, M. A. Millar-Blanchaer, S. Tinyanont, D. Blinov, V. Pelgrims, S. Johnson, R. Skalidis, and A. Soam. A Compilation of Optical Starlight Polarization Catalogs. , 276(1):15, January 2025.

[201] Léna Parc, François Bouchy, Julia Venturini, Caroline Dorn, and Ravit Helled. From super-Earths to sub-Neptunes: Observational constraints and connections to theoretical models. , 688:A59, August 2024.

[202] Randall N. Paschall and David J. Anderson. Linear quadratic Gaussian control of a deformable mirror adaptive optics system with time-delayed measurements. , 32(31):6347–6358, November 1993.

[203] Randall N. Paschall and David J. Anderson. Linear quadratic Gaussian control of a deformable mirror adaptive optics system with time-delayed measurements. *Applied Optics*, 32(31):6347–6358, November 1993.

[204] Oren Patashnik. *BibTEXing*. Computer Science Department, Stanford University, Stanford, California, January 1988. Available in the BibTEX release.

[205] Oren Patashnik. *Designing BibTEX Styles*. Computer Science Department, Stanford University, January 1988.

[206] Marven F. Pedbost, Trillean Pomalgu, Chris Lintott, Nora Eisner, and Belinda Nicholson. Defining the really habitable zone. 2020.

[207] Eliad Peretz, Peter Wizinowich, Eduardo Marin, Richard Butler, Bert Pasquale, Maxwell A. Millar-Blanchaer, Scott Lilley, Luke Gers, Shui Hung Kwok, Jason Chin, Sam Ragland, Ed Wetherell, Brett Smith, John O'Meara, Marc Kassis, Greg Aldering, Susana Deustua, Rogier Windhorst, Christian Marois, Saul Perlmutter, Sara Seager, Breann Sitarski, Guillaume Filion, Jean Thomas Landry, Étienne Gauvin, J. Fowler, Rebecca Jensen-Clem, Eric L. Nielsen, Imke de Pater, Peter Plavchan, Steph Sallum, Shobita Satyapal, John Mather, Peter Kurczynski, Kayla Carmical, Jack Grossman, Andrew Lewis, Maya Wertheim,

Vivian Palmer, Keshet Shavit, and Kevin Hall. ORCAS Keck mission and instrument development. In Julia J. Bryant, Kentaro Motohara, and Joël. R. D. Vernet, editors, *Ground-based and Airborne Instrumentation for Astronomy X*, volume 13096 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 130960M, July 2024.

[208] Eliad Peretz, Peter Wizinowich, Maxwell A. Millar-Blanchaer, Bert Pasquale, Jason Chin, Guillaume Filion, Luke Gers, Shui Hung Kwok, Jean Thomas Landry, Scott Lilley, Eduardo Marin, Sam Ragland, Ed Wetherell, Kevin Hall, Jules Fowler, Étienne Gauvin, Rebecca Jensen-Clem, Marc Kassis, Peter Kurczynski, John Mather, Eric L. Nielsen, John O'Meara, Imke de Pater, Peter Plavchan, Steph Sallum, Shobita Satyapal, Brett Smith, Kayla Carmical, Jack Grossman, Andrew Lewis, Maya Wertheim, and Vivian Palmer. ORCAS Keck instrument demonstrator. *Journal of Astronomical Telescopes, Instruments, and Systems*, 11:014004, January 2025.

[209] Marshall D. Perrin, Gaspard Duchene, Max Millar-Blanchaer, Michael P. Fitzgerald, James R. Graham, Sloane J. Wiktorowicz, Paul G. Kalas, Bruce Macintosh, Brian Bauman, Andrew Cardwell, Jeffrey Chilcote, Robert J. De Rosa, Daren Dillon, René Doyon, Jennifer Dunn, Darren

Erikson, Donald Gavel, Stephen Goodsell, Markus Hartung, Pascale Hibon, Patrick Ingraham, Daniel Kerley, Quinn Konapacky, James E. Larkin, Jérôme Maire, Franck Marchis, Christian Marois, Tushar Mittal, Katie M. Morzinski, B. R. Oppenheimer, David W. Palmer, Jennifer Patience, Lisa Poyneer, Laurent Pueyo, Fredrik T. Rantakyrö, Naru Sadakuni, Leslie Saddlemyer, Dmitry Savransky, Rémi Soummer, Anand Sivaramakrishnan, Inseok Song, Sandrine Thomas, J. Kent Wallace, Jason J. Wang, and Schuyler G. Wolff. Polarimetry with the Gemini Planet Imager: Methods, Performance at First Light, and the Circumstellar Ring around HR 4796A. , 799(2):182, February 2015.

[210] Erik A. Petigura and Geoffrey W. Marcy. Carbon and Oxygen in Nearby Stars: Keys to Protoplanetary Disk Chemistry. , 735(1):41, July 2011.

[211] C. Petit, T. Fusco, J. F. Sauvage, A. Sevin, M. Suarez, J. Charton, P. Baudoz, J. L. Beuzit, T. Buey, K. Dohlen, P. Feautrier, E. Fedrigo, B. Fleuzry, J. L. Gach, N. Hubin, M. Kasper, D. Mouillet, D. Perret, P. Puget, J. C. Sinquin, C. Soenke, and F. Wildi. Integration, tests and laboratory performance of SAXO, the VLT-SPHERE extreme AO system. In *Second International Conference on Adaptive Optics for Extremely Large Telescopes. Online at <A href="http://ao4elt2.lesia.ob-*

spm.fr">http://ao4elt2.lesia.obspm.fr</A, page 16, September 2011.

[212] Raymond Pierrehumbert and Eric Gaidos. Hydrogen Greenhouse Planets Beyond the Habitable Zone. , 734(1):L13, June 2011.

[213] Emiel H. Por, Sebastiaan Y. Haffert, Vikram M. Radhakrishnan, David S. Doelman, Maaike van Kooten, and Steven P. Bos. High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator. In Laird M. Close, Laura Schreiber, and Dirk Schmidt, editors, *Adaptive Optics Systems VI*, volume 10703 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1070342, July 2018.

[214] B. Pou, F. Ferreira, E. Quinones, D. Gratadour, and M. Martin. Adaptive optics control with multi-agent model-free reinforcement learning. *Optics Express*, 30(2):2991, January 2022.

[215] B. Pou, J. Smith, E. Quinones, M. Martin, and D. Gratadour. Model-free reinforcement learning with a non-linear reconstructor for closed-loop adaptive optics control with a pyramid wavefront sensor. In Laura Schreiber, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems VIII*, volume 12185 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 121852U, August 2022.

[216] Lisa Poyneer, Marcos van Dam, and Jean-Pierre Véran. Experimental verification of the frozen flow atmospheric turbulence assumption with use of astronomical adaptive optics telemetry. *Journal of the Optical Society of America A*, 26(4):833, March 2009.

[217] Lisa A. Poyneer, S. Mark Ammons, Mike K. Kim, Brian Bauman, Jesse Terrel-Perez, Aaron J. Lemmer, and Jayke Nguyen. Laboratory demonstration of the prediction of wind-blown turbulence by adaptive optics at 8 kHz with use of LQG control. , 62(8):1871, March 2023.

[218] Lisa A. Poyneer, Bruce A. Macintosh, and Jean-Pierre Véran. Fourier transform wavefront control with adaptive prediction of the atmosphere. *Journal of the Optical Society of America A*, 24(9):2645, January 2007.

[219] Lisa A. Poyneer, David W. Palmer, Bruce Macintosh, Dmitry Savransky, Naru Sadakuni, Sandrine Thomas, Jean-Pierre Véran, Katherine B. Follette, Alexandra Z. Greenbaum, S. Mark Ammons, Vanessa P. Bailey, Brian Bauman, Andrew Cardwell, Daren Dillon, Donald Gavel, Markus Hartung, Pascale Hibon, Marshall D. Perrin, Fredrik T. Rantakyrö, Anand Sivaramakrishnan, and Jason J. Wang. Performance of the Gemini Planet Imager's adaptive optics system. , 55(2):323, January 2016.

[220] Lisa A. Poyneer and Jean-Pierre Véran. Optimal modal Fourier-

transform wavefront control. *Journal of the Optical Society of America A*, 22(8):1515–1526, August 2005.

[221] Stephan D. Price, Beverly J. Smith, Thomas A. Kuchar, Donald R. Mizuno, and Kathleen E. Kraemer. 3.6 Years of DIRBE Near-infrared Stellar Light Curves. , 190(2):203–219, October 2010.

[222] Philipp Protte and Susanne M. Hoffmann. Accuracy of magnitudes in pre−telescopic star catalogs. *Astronomische Nachrichten*, 341(8):827–840, October 2020.

[223] Andrej Prša, Petr Harmanec, Guillermo Torres, Eric Mamajek, Martin Asplund, Nicole Capitaine, Jørgen Christensen-Dalsgaard, Éric Depagne, Margit Haberreiter, and Saskia Hekker. Nominal Values for Selected Solar and Planetary Quantities: IAU 2015 Resolution B3. , 152(2):41, Aug 2016.

[224] I. Ramírez, C. Allende Prieto, and D. L. Lambert. Oxygen Abundances in Nearby FGK Stars and the Galactic Chemical Evolution of the Local Disk and Halo. , 764(1):78, February 2013.

[225] Benjamin Recht. A Tour of Reinforcement Learning: The View from Continuous Control. *arXiv e-prints*, page arXiv:1806.09460, June 2018.

[226] H. Rein and S. F. Liu. REBOUND: an open-source multi-purpose N-body code for collisional dynamics. , 537:A128, January 2012.

[227] Hanno Rein and David S. Spiegel. IAS15: a fast, adaptive, high-order integrator for gravitational dynamics, accurate to machine precision over a billion orbits. , 446(2):1424–1437, January 2015.

[228] Reed L. Riddle, Andrei Tokovinin, Brian D. Mason, William I. Hartkopf, Lewis C. Roberts, Jr., Christoph Baranec, Nicholas M. Law, Khanh Bui, Mahesh P. Burse, H. K. Das, Richard G. Dekany, Shrinivas Kulkarni, Sujit Punnadi, A. N. Ramaprakash, and Shriharsh P. Tendulkar. A Survey of the High Order Multiplicity of Nearby Solar-type Binary Stars with Robo-AO. , 799(1):4, January 2015.

[229] Adric R. Riedel, Sarah C. Blunt, Erini L. Lambrides, Emily L. Rice, Kelle L. Cruz, and Jacqueline K. Faherty. LACEwING: A New Moving Group Analysis Code. , 153(3):95, March 2017.

[230] Lewis C. Roberts, Jr., Marshall D. Perrin, Franck Marchis, Anand Sivaramakrishnan, Russell B. Makidon, Julian C. Christou, Bruce A. Macintosh, Lisa A. Poyneer, Marcos A. van Dam, and Mitchell Troy. Is that really your Strehl ratio? In Domenico Bonaccini Calia, Brent L. Ellerbroek, and Roberto Ragazzoni, editors, *Advancements in Adaptive Optics*, vol-

ume 5490 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 504–515, October 2004.

[231] Giulia Roccetti, Claudia Emde, Michael F. Sterzik, Mihail Manev, Julia V. Seidel, and Stefano Bagnulo. Planet Earth in reflected and polarized light: I. Three-dimensional radiative transfer simulations of realistic surface-atmosphere systems. , 697:A170, May 2025.

[232] Gene Roddenberry. Star trek. NBC, 1966. Accessed: 2024-3-28.

[233] Gene Roddenberry. Star trek: The next generation. Paramount Domestic Television, 1987. Accessed: 2024-3-28.

[234] S. Roeser, M. Demleitner, and E. Schilbach. The PPMXL Catalog of Positions and Proper Motions on the ICRS. Combining USNO-B1.0 and the Two Micron All Sky Survey (2MASS). , 139(6):2440–2447, June 2010.

[235] Carl Sagan. *Cosmos*. 1980.

[236] Maïssa Salama, Charlotte Guthery, Vincent Chambouleyron, Rebecca Jensen-Clem, J. Kent Wallace, Jacques-Robert Delorme, Mitchell Troy, Tobias Wenger, Daniel Echeverri, Luke Finnerty, Nemanja Jovanovic, Joshua Liberman, Ronald A. López, Dimitri Mawet, Evan C. Morris,

Maaike van Kooten, Jason J. Wang, Peter Wizinowich, Yinzi Xin, and Jerry Xuan. Keck Primary Mirror Closed-loop Segment Control Using a Vector-Zernike Wavefront Sensor. , 967(2):171, June 2024.

[237] Maïssa Salama, Rebecca Jensen-Clem, Mahawa Cissé, J. Kent Wallace, Mitchell Troy, Laurent Pueyo, Charlotte Guthery, Antonin Bouchez, and Vincent Chambouleyron. Keck Observatory as an HWO Testbed: validating wavefront sensing and control schemes on a large segmented aperture in parallel with high-contrast science. *arXiv e-prints*, page arXiv:2511.01585, November 2025.

[238] M. Salvato, J. Buchner, T. Budavári, T. Dwelly, A. Merloni, M. Brusa, A. Rau, S. Fotopoulou, and K. Nandra. Finding counterparts for all-sky X-ray surveys with NWAY: a Bayesian algorithm for cross-matching multiple catalogues. , 473(4):4937–4955, February 2018.

[239] M. Samland, T. D. Brandt, J. Milli, P. Delorme, and A. Vigan. Spectral cube extraction for the VLT/SPHERE IFS. Open-source pipeline with full forward modeling and improved sensitivity. , 668:A84, December 2022.

[240] Arthur L. Samuel. First grade TEX: A beginner's TEX manual. Techni-

cal Report SATN-CS-83-985, Computer Science Department, Stanford University, Stanford, California, November 1983.

[241] Edward F. Schlafly, Aaron M. Meisner, and Gregory M. Green. The unWISE Catalog: Two Billion Infrared Sources from Five Years of WISE Imaging. , 240(2):30, February 2019.

[242] H. M. Schmid, A. Bazzon, R. Roelfsema, D. Mouillet, J. Milli, F. Menard, D. Gisler, S. Hunziker, J. Pragt, C. Dominik, A. Boccaletti, C. Ginski, L. Abe, S. Antoniucci, H. Avenhaus, A. Baruffolo, P. Baudoz, J. L. Beuzit, M. Carbillet, G. Chauvin, R. Claudi, A. Costille, J. B. Daban, M. de Haan, S. Desidera, K. Dohlen, M. Downing, E. Elswijk, N. Engler, M. Feldt, T. Fusco, J. H. Girard, R. Gratton, H. Hanenburg, Th. Henning, N. Hubin, F. Joos, M. Kasper, C. U. Keller, M. Langlois, E. Lagadec, P. Martinez, E. Mulder, A. Pavlov, L. Podio, P. Puget, S. P. Quanz, F. Rigal, B. Salasnich, J. F. Sauvage, M. Schuil, R. Siebenmorgen, E. Sissa, F. Snik, M. Suarez, Ch. Thalmann, M. Turatto, S. Udry, A. van Duin, R. G. van Holstein, A. Vigan, and F. Wildi. SPHERE/ZIMPOL high resolution polarimetric imager. I. System overview, PSF parameters, coronagraphy, and polarimetry. , 619:A9, November 2018.

[243] H. M. Schmid, A. Bazzon, R. Roelfsema, D. Mouillet, J. Milli, F. Menard,

D. Gisler, S. Hunziker, J. Pragt, C. Dominik, A. Boccaletti, C. Ginski, L. Abe, S. Antoniucci, H. Avenhaus, A. Baruffolo, P. Baudoz, J. L. Beuzit, M. Carbillet, G. Chauvin, R. Claudi, A. Costille, J. B. Daban, M. de Haan, S. Desidera, K. Dohlen, M. Downing, E. Elswijk, N. Engler, M. Feldt, T. Fusco, J. H. Girard, R. Gratton, H. Hanenburg, Th. Henning, N. Hubin, F. Joos, M. Kasper, C. U. Keller, M. Langlois, E. Lagadec, P. Martinez, E. Mulder, A. Pavlov, L. Podio, P. Puget, S. P. Quanz, F. Rigal, B. Salasnich, J. F. Sauvage, M. Schuil, R. Siebenmorgen, E. Sissa, F. Snik, M. Suarez, Ch. Thalmann, M. Turatto, S. Udry, A. van Duin, R. G. van Holstein, A. Vigan, and F. Wildi. SPHERE/ZIMPOL high resolution polarimetric imager. I. System overview, PSF parameters, coronagraphy, and polarimetry. , 619:A9, November 2018.

[244] JÃ¼rgen Schmidt, Keiji Ohtsuki, Nicole Rappaport, Heikki Salo, and Frank Spahn. Dynamics of Saturn's Dense Rings. In Michele K. Dougherty, Larry W. Esposito, and Stamatios M. Krimigis, editors, *Saturn from Cassini-Huygens*, page 413. 2009.

[245] Adam C. Schneider, Jennifer Greco, Michael C. Cushing, J. Davy Kirkpatrick, Amy Mainzer, Christopher R. Gelino, Sergio B. Fajardo-Acosta, and James Bauer. A Proper Motion Survey Using the First Sky Pass of

NEOWISE-reactivation Data. , 817(2):112, February 2016.

[246] Mathew Schofield, William J. Chaplin, Daniel Huber, Tiago L. Campante, Guy R. Davies, Andrea Miglio, Warrick H. Ball, Thierry Appourchaux, Sarbani Basu, Timothy R. Bedding, Jørgen Christensen-Dalsgaard, Orlagh Creevey, Rafael A. García, Rasmus Handberg, Steven D. Kawaler, Hans Kjeldsen, David W. Latham, Mikkel N. Lund, Travis S. Metcalfe, George R. Ricker, Aldo Serenelli, Victor Silva Aguirre, Dennis Stello, and Roland Vanderspek. The Asteroseismic Target List for Solar-like Oscillators Observed in 2 minute Cadence with the Transiting Exoplanet Survey Satellite. , 241(1):12, March 2019.

[247] Greg J. Schwarz, Jan-Uwe Ness, J. P. Osborne, K. L. Page, P. A. Evans, A. P. Beardmore, Frederick M. Walter, L. Andrew Helton, Charles E. Woodward, Mike Bode, Sumner Starrfield, and Jeremy J. Drake. Swift X-Ray Observations of Classical Novae. II. The Super Soft Source Sample. , 197(2):31, Dec 2011.

[248] Ridley Scott. Blade runner. Warner Bros, 1982. Accessed: 2024-3-28.

[249] Vivek Sharma, Matija Crne, Jung Ok Park, and Mohan Srinivasarao. Structural Origin of Circularly Polarized Iridescence in Jeweled Beetles. *Science*, 325(5939):449, July 2009.

[250] Iuliia Shatokhina, Victoria Hutterer, and Ronny Ramlau. Review on methods for wavefront reconstruction from pyramid wavefront sensor data. *Journal of Astronomical Telescopes, Instruments, and Systems*, 6:010901, January 2020.

[251] Joel Sheilds. keck vibration spie proceeding – fill out with spie epub. In Stuart B. Shaklan and Garreth J. Ruane, editors, *Techniques and Instrumentation for Detection of Exoplanets X*, volume 11824 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, September 2025.

[252] Robert H. Shumway. *Applied statistical time series analysis*. 1988.

[253] David K. Sing, Jonathan J. Fortney, Nikolay Nikolov, Hannah R. Wakeford, Tiffany Kataria, Thomas M. Evans, Suzanne Aigrain, Gilda E. Ballester, Adam S. Burrows, Drake Deming, Jean-Michel Désert, Neale P. Gibson, Gregory W. Henry, Catherine M. Huitson, Heather A. Knutson, Alain Lecavelier Des Etangs, Frederic Pont, Adam P. Showman, Alfred Vidal-Madjar, Michael H. Williamson, and Paul A. Wilson. A continuum from clear to cloudy hot-Jupiter exoplanets without primordial water depletion. , 529(7584):59–62, January 2016.

[254] Baptiste Sinquin, Léonard Prengère, Caroline Kulcsár, Henri-François

Raynaud, Eric Gendron, James Osborn, Alastair Basden, Jean-Marc Conan, Nazim Bharmal, Lisa Bardou, Lazar Staykov, Tim Morris, Tristan Buey, Fanny Chemla, and Matthieu Cohen. On-sky results for adaptive optics control with data-driven models on low-order modes. , 498(3):3228–3240, November 2020.

[255] Gaetano Sivo, Caroline Kulcsár, Jean-Marc Conan, Henri-François Raynaud, Éric Gendron, Alastair Basden, Fabrice Vidal, Tim Morris, Serge Meimon, Cyril Petit, Damien Gratadour, Olivier Martin, Zoltan Hubert, Arnaud Sevin, Denis Perret, Fanny Chemla, Gérard Rousset, Nigel Dipper, Gordon Talbot, Eddy Younger, Richard Myers, David Henry, Stephen Todd, David Atkinson, Colin Dickson, and Andy Longmore. First on-sky SCAO validation of full LQG control with vibration mitigation on the CANARY pathfinder. *Optics Express*, 22(19):23565, September 2014.

[256] Nour Skaf, Olivier Guyon, Éric Gendron, Kyohoon Ahn, Arielle Bertrou-Cantou, Anthony Boccaletti, Jesse Cranney, Thayne Currie, Vincent Deo, Billy Edwards, Florian Ferreira, Damien Gratadour, Julien Lozi, Barnaby Norris, Arnaud Sevin, Fabrice Vidal, and Sébastien Vievard. On-sky validation of image-based adaptive optics wavefront sensor referencing. ,

659:A170, March 2022.

[257] Andrew J. Skemer, Philip M. Hinz, Simone Esposito, Adam Burrows, Jarron Leisenring, Michael Skrutskie, Silvano Desidera, Dino Mesa, Carmelo Arcidiacono, Filippo Mannucci, Timothy J. Rodigas, Laird Close, Don McCarthy, Craig Kulesa, Guido Agapito, Daniel Apai, Javier Argomedo, Vanessa Bailey, Konstantina Boutsia, Runa Briguglio, Guido Brusa, Lorenzo Busoni, Riccardo Claudi, Joshua Eisner, Luca Fini, Katherine B. Follette, Peter Garnavich, Raffaele Gratton, Juan Carlos Guerra, John M. Hill, William F. Hoffmann, Terry Jones, Megan Krejny, Jared Males, Elena Masciadri, Michael R. Meyer, Douglas L. Miller, Katie Morzinski, Matthew Nelson, Enrico Pinna, Alfio Puglisi, Sascha P. Quanz, Fernando Quiros-Pacheco, Armando Riccardi, Paolo Stefanini, Vidhya Vaitheeswaran, John C. Wilson, and Marco Xompero. First Light LBT AO Images of HR 8799 bcde at 1.6 and 3.3 μm: New Discrepancies between Young Planets and Old Brown Dwarfs. , 753(1):14, July 2012.

[258] Beverly J. Smith, Stephan D. Price, and Rachel I. Baker. The COBE DIRBE Point Source Catalog. , 154(2):673–704, October 2004.

[259] W. Hayden Smith. Spectral differential imaging detection of planets about nearby stars. , 99:1344–1353, December 1987.

[260] Ignas Snellen, Sebastiaan Haffert, Matthew Kenworthy, and Tomas Stolker. The next frontier in exoplanet science: Imaging our neighbouring planetary systems. *arXiv e-prints*, page arXiv:2512.13756, December 2025.

[261] Michael D. Spivak. *The Joy of TEX*. American Mathematical Society, 1985.

[262] Marco Stangalini, Dario Del Moro, Francesco Berrilli, and Oskar von der Lühe. Zernike basis optimization for solar adaptive optics by using information theory. , 49(11):2090, April 2010.

[263] John Stauffer, Angelle M. Tanner, Geoffrey Bryden, Solange Ramirez, Bruce Berriman, David R. Ciardi, Stephen R. Kane, Trisha Mizusawa, Alan Payne, Peter Plavchan, Kaspar von Braun, Pamela Wyatt, and J. Davy Kirkpatrick. Accurate Coordinates and 2MASS Cross Identifications for (Almost) All Gliese Catalog Star. , 122(894):885, August 2010.

[264] T. Stolker, M. Min, D. M. Stam, P. Mollière, C. Dominik, and L. B. F. M. Waters. Polarized scattered light from self-luminous exoplanets. Three-dimensional scattering radiative transfer with ARTES. , 607:A42, November 2017.

[265] Richard Sutton and Andrew Barto. *Reinforcement Learning, An Introduction, second edition.* 2020.

[266] Robin Swanson, Masen Lamb, Carlos M. Correia, Suresh Sivanandam, and Kiriakos Kutulakos. Closed loop predictive control of adaptive optics systems with convolutional neural networks. , 503(2):2944–2954, May 2021.

[267] Valerian Ilich Tatarskii. *Wave Propagation in Turbulent Medium.* 1961.

[268] G. I. Taylor. The Spectrum of Turbulence. *Proceedings of the Royal Society of London Series A*, 164(919):476–490, February 1938.

[269] Andrei Tokovinin. From Binaries to Multiples. I. Data on F and G Dwarfs within 67 pc of the Sun. , 147(4):86, April 2014.

[270] J. L. Tonry, L. Denneau, H. Flewelling, A. N. Heinze, C. A. Onken, S. J. Smartt, B. Stalder, H. J. Weiland, and C. Wolf. The ATLAS All-Sky Stellar Reference Catalog. , 867(2):105, November 2018.

[271] P Umesh. Image processing in python. *CSI Communications*, 23, 2012.

[272] S. E. Urban, T. E. Corbin, G. L. Wycoff, J. C. Martin, E. S. Jackson, M. I. Zacharias, and D. M. Hall. The AC 2000: The Astrographic Catalogue on

the System Defined by the HIPPARCOS Catalogue. , 115(3):1212–1223, March 1998.

[273] S. E. Urban, N. Zacharias, and G. L. Wycoff. VizieR Online Data Catalog: The UCAC2 Bright Star Supplement (Urban+, 2006). VizieR On-line Data Catalog: I/294A. Originally published in: 2009yCat.1294....0U, October 2009.

[274] Richelle F. van Capelleveen, Christian Ginski, Matthew A. Kenworthy, Jake Byrne, Chloe Lawlor, Dan McLachlan, Eric E. Mamajek, Tomas Stolker, Myriam Benisty, Alexander J. Bohn, Laird M. Close, Carsten Dominik, Sebastiaan Haffert, Rico Landman, Jie Ma, Ignas Snellen, Ryo Tazaki, Nienke van der Marel, Lukas Welzel, and Yapeng Zhang. WIde Separation Planets In Time (WISPIT): A Gap-clearing Planet in a Multi-ringed Disk around the Young Solar-type Star WISPIT 2. , 990(1):L8, September 2025.

[275] M. van Dam. Keck adaptive optics note 1480 compute delay for the keck ao systems. Technical report, W. M. Keck Observatory, May 2024.

[276] R. G. van Holstein, T. Stolker, R. Jensen-Clem, C. Ginski, J. Milli, J. de Boer, J. H. Girard, Z. Wahhaj, A. J. Bohn, M. A. Millar-Blanchaer, M. Benisty, M. Bonnefoy, G. Chauvin, C. Dominik, S. Hinkley, C. U.

Keller, M. Keppler, M. Langlois, S. Marino, F. Ménard, C. Perrot, T. O. B. Schmidt, A. Vigan, A. Zurlo, and F. Snik. A survey of the linear polarization of directly imaged exoplanets and brown dwarf companions with SPHERE-IRDIS. First polarimetric detections revealing disks around DH Tau B and GSC 6214-210 B. , 647:A21, March 2021.

[277] Rob G. van Holstein, Steven P. Bos, Jasper Ruigrok, Julien Lozi, Olivier Guyon, Barnaby Norris, Frans Snik, Jeffrey Chilcote, Thayne Currie, Tyler D. Groff, Joost 't Hart, Nemanja Jovanovic, Jeremy Kasdin, Tomoyuki Kudo, Frantz Martinache, Ben Mazin, Ananya Sahoo, Motohide Tamura, Sébastien Vievard, Alex Walter, and Jin Zhang. Calibration of the instrumental polarization effects of SCExAO-CHARIS' spectropolarimetric mode. In Christopher J. Evans, Julia J. Bryant, and Kentaro Motohara, editors, *Ground-based and Airborne Instrumentation for Astronomy VIII*, volume 11447 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 114475B, December 2020.

[278] M. A. M. van Kooten, N. Doelman, and M. Kenworthy. Robustness of prediction for extreme adaptive optics systems under various observing conditions. An analysis using VLT/SPHERE adaptive optics data. , 636:A81, April 2020.

[279] Maaike van Kooten, Niek Doelman, and Matthew Kenworthy. Impact of time-variant turbulence behavior on prediction for adaptive optics systems. *Journal of the Optical Society of America A*, 36(5):731, May 2019.

[280] Maaike A. M. van Kooten, Rebecca Jensen-Clem, Sylvain Cetre, Sam Ragland, Charlotte Z. Bond, J. Fowler, and Peter Wizinowich. Predictive wavefront control on Keck II adaptive optics bench: on-sky coronagraphic results. *Journal of Astronomical Telescopes, Instruments, and Systems*, 8:029006, April 2022.

[281] Maaike A. M. van Kooten, Rebecca Jensen-Clem, Sylvain Cetre, Sam Ragland, Charlotte Z. Bond, J. Fowler, and Peter Wizinowich. Predictive wavefront control on Keck II adaptive optics bench: on-sky coronagraphic results. *Journal of Astronomical Telescopes, Instruments, and Systems*, 8:029006, April 2022.

[282] Sophia R. Vaughan, Timothy D. Gebhard, Kimberly Bott, Sarah L. Casewell, Nicolas B. Cowan, David S. Doelman, Matthew Kenworthy, Johan Mazoyer, Maxwell A. Millar-Blanchaer, Victor J. H. Trees, Daphne M. Stam, Olivier Absil, Lisa Altinier, Pierre Baudoz, Ruslan Belikov, Alexis Bidot, Jayne L. Birkby, Markus J. Bonse, Bernhard

Brandl, Alexis Carlotti, Elodie Choquet, Dirk van Dam, Niyati Desai, Kevin Fogarty, J. Fowler, Kyle van Gorkom, Yann Gutierrez, Olivier Guyon, Sebastiaan Y. Haffert, Olivier Herscovici-Schiller, Adrien Hours, Roser Juanola-Parramon, Evangelia Kleisioti, Lorenzo König, Maaike van Kooten, Mariya Krasteva, Iva Laginja, Rico Landman, Lucie Leboulleux, David Mouillet, Mamadou N'Diaye, Emiel H. Por, Laurent Pueyo, and Frans Snik. Chasing rainbows and ocean glints: Inner working angle constraints for the Habitable Worlds Observatory. , 524(4):5477–5485, October 2023.

[283] A. Vigan, M. Langlois, K. Dohlen, A. Zurlo, C. Moutou, A. Costille, C. Gry, F. Madec, D. Le Mignant, Laurence Gluck, and J. F. Sauvage. SPHERE/IRDIS: final performance assessment of the dual-band imaging and long slit spectroscopy modes. In Suzanne K. Ramsay, Ian S. McLean, and Hideki Takami, editors, *Ground-based and Airborne Instrumentation for Astronomy V*, volume 9147 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 91474T, July 2014.

[284] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett,

Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

[285] Frédéric P. A. Vogt, Michael A. Dopita, Lisa J. Kewley, Ralph S. Sutherland, Julia Scharwächter, Hassan M. Basurah, Alaa Ali, and Morsi A. Amer. Galaxy Emission Line Classification Using Three-dimensional Line Ratio Diagrams. , 793(2):127, Oct 2014.

[286] H. R. Wakeford, N. K. Lewis, J. Fowler, G. Bruno, T. J. Wilson, S. E. Moran, J. Valenti, N. E. Batalha, J. Filippazzo, V. Bourrier, S. M. Hörst, S. M. Lederer, and J. de Wit. Disentangling the Planet from the Star in Late-Type M Dwarfs: A Case Study of TRAPPIST-1g. , 157(1):11, January 2019.

[287] H. R. Wakeford, C. Visscher, N. K. Lewis, T. Kataria, M. S. Marley, J. J.

Fortney, and A. M. Mandell. High-temperature condensate clouds in super-hot Jupiter atmospheres. , 464(4):4247–4254, February 2017.

[288] Jason Wang, Dimitri Mawet, Nemanja Jovanovic, Jacques Robert Delorme, Charlotte Bond, Jacklyn Pezzato, Sylvain Cetre, Daniel Echeverri, J. Kent Wallace, Randall Bartos, Scott Lilley, Sam Ragland, Garreth Ruane, Peter Wizinowich, Mark Chun, Ji Wang, Michael P. Fitzgerald, Andy Skemer, Emily Martin, Ed Wetherell, Eric Wang, Shane Jacobson, Eric Warmbier, Charles Lockhart, and Donald Hall. First Light of the Keck Planet Imager and Characterizer. In *AAS/Division for Extreme Solar Systems Abstracts*, volume 51 of *AAS/Division for Extreme Solar Systems Abstracts*, page 328.01, August 2019.

[289] Jason J. Wang, Matas Kulikauskas, and Sarah Blunt. whereistheplanet: Predicting positions of directly imaged companions. Astrophysics Source Code Library, record ascl:2101.003, January 2021.

[290] Jason J. Wang, Jean-Baptise Ruffio, Robert J. De Rosa, Jonathan Aguilar, Schuyler G. Wolff, and Laurent Pueyo. pyKLIP: PSF Subtraction for Exoplanets and Disks. Astrophysics Source Code Library, record ascl:1506.001, June 2015.

[291] Makoto Watanabe, Shin Oya, Yutaka Hayano, Hideki Takami, Masayuki

Hattori, Yosuke Minowa, Yoshihiko Saito, Meguru Ito, Naoshi Murakami, Masanori Iye, Olivier Guyon, Stephen Colley, Michael Eldred, Taras Golota, and Matthew Dinkins. Implementation of 188-element curvature-based wavefront sensor and calibration source unit for the Subaru LGSAO system. In Norbert Hubin, Claire E. Max, and Peter L. Wizinowich, editors, *Adaptive Optics Systems*, volume 7015 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 701564, July 2008.

[292] P. Wizinowich, D. S. Acton, C. Shelton, P. Stomski, J. Gathright, K. Ho, W. Lupton, K. Tsubota, O. Lai, C. Max, J. Brase, J. An, K. Avicola, S. Olivier, D. Gavel, B. Macintosh, A. Ghez, and J. Larkin. First Light Adaptive Optics Images from the Keck II Telescope: A New Era of High Angular Resolution Imagery. , 112(769):315–319, March 2000.

[293] P. Wizinowich, A. Bouchez, S. Cetre, M. van Dam, L. Gers, C. Guthery, R. Jensen-Clem, S. Lilley, S. Sallum, J. Steiner, A. Surendran, M. Vincent, M. Wetherell, and with ALPAO and Microgate. Keck adaptive optics note 1456 high order keck adaptive optics (haka) project design review ',;lpl '. Technical report, W. M. Keck Observatory, January 2025.

[294] P. Wizinowich, A. Bouchez, E. Marin, S. Cetre, J. Chin, C. Correia,

M. van Dam, J.-R. Delorme, L. Gers, C. Guthery, S. Karkar, S. Kwok, S. Lilley, J. Lyke, P. Richards, M. Service, J. Steiner, A. Surendran, K. Tsubota, E. Wetherell, M. Bottom, R. Dekany, A. Ghez, P. Hinz, M. Liu, J. Lu, R. Jensen-Clem, M. Millar-Blanchaer, E. Peretz, S. Sallum, T. Treu, and S. Wright. Adaptive optics at W. M. Keck Observatory. In Kathryn J. Jackson, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems IX*, volume 13097 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 130970K, August 2024.

[295] P. Wizinowich, J. R. Lu, S. Cetre, J. Chin, C. Correia, J. R. Delorme, L. Gers, S. Lilley, J. Lyke, E. Marin, S. Ragland, P. Richards, A. Surendran, E. Wetherell, C. F. Chen, D. Chu, T. Do, C. Fassnacht, M. Freeman, A. Gautam, A. Ghez, L. Hunter, T. Jones, M. C. Liu, D. Mawet, C. Max, M. Morris, M. Phillips, J. B. Ruffio, N. E. Rundquist, S. Sabhlok, S. Terry, T. Treu, and S. Wright. Keck All sky Precision Adaptive optics program overview. In Laura Schreiber, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems VIII*, volume 12185 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 121850Q, August 2022.

[296] Alison P. Wong, Barnaby R. M. Norris, Vincent Deo, Olivier Guyon, Peter G. Tuthill, Julien Lozi, Sébastien Vievard, and Kyohoon Ahn. Machine learning for wavefront sensing. In Laura Schreiber, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems VIII*, volume 12185 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 121852I, August 2022.

[297] Alison P. Wong, Barnaby R. M. Norris, Peter G. Tuthill, Richard Scalzo, Julien Lozi, Sébastien Vievard, and Olivier Guyon. Predictive control for adaptive optics using neural networks. *Journal of Astronomical Telescopes, Instruments, and Systems*, 7:019001, January 2021.

[298] N. Woolf. Dome seeing. , 91:523–529, August 1979.

[299] J. T. Wright, S. Upadhyay, G. W. Marcy, D. A. Fischer, Eric B. Ford, and John Asher Johnson. Ten New and Updated Multiplanet Systems and a Survey of Exoplanetary Systems. , 693(2):1084–1099, March 2009.

[300] W. Jerry Xuan, Dimitri Mawet, Henry Ngo, Garreth Ruane, Vanessa P. Bailey, Élodie Choquet, Olivier Absil, Carlos Alvarez, Marta Bryan, Therese Cook, Bruno Femenía Castellá, Carlos Gomez Gonzalez, Elsa Huby, Heather A. Knutson, Keith Matthews, Sam Ragland, Eugene Serabyn, and Zoë Zawol. Characterizing the Performance of the NIRC2

Vortex Coronagraph at W. M. Keck Observatory. , 156(4):156, October 2018.

[301] I. Yamamura, S. Makiuti, N. Ikeda, Y. Fukuda, S. Oyabu, T. Koga, and G. J. White. VizieR Online Data Catalog: AKARI/FIS All-Sky Survey Point Source Catalogues (ISAS/JAXA, 2010). VizieR On-line Data Catalog: II/298. Originally published in: ISAS/JAXA (2010), April 2010.

[302] Ben Yelverton, Grant M. Kennedy, and Kate Y. L. Su. No significant correlation between radial velocity planet presence and debris disc properties. , 495(2):1943–1957, June 2020.

[303] N. Zacharias, C. Finch, and J. Frouard. UCAC5: New Proper Motions Using Gaia DR1. , 153(4):166, April 2017.

[304] N. Zacharias, C. T. Finch, T. M. Girard, A. Henden, J. L. Bartlett, D. G. Monet, and M. I. Zacharias. The Fourth US Naval Observatory CCD Astrograph Catalog (UCAC4). , 145(2):44, February 2013.

[305] N. Zacharias, D. G. Monet, S. E. Levine, S. E. Urban, R. Gaume, and G. L. Wycoff. The Naval Observatory Merged Astrometric Dataset (NOMAD). In *American Astronomical Society Meeting Abstracts*, volume

205 of *American Astronomical Society Meeting Abstracts*, page 48.15, December 2004.

[306] Norbert Zacharias, Valeri V. Makarov, Charles T. Finch, Hugh C. Harris, Jeffrey A. Munn, and John P. Subasavage. USNO Bright Star Catalog, Version 1. , 164(2):36, August 2022.

[307] Xi Zhang and Adam P. Showman. Atmospheric Circulation of Brown Dwarfs: Jets, Vortices, and Time Variability. , 788(1):L6, June 2014.

[308] Zhoujian Zhang, Paul Mollière, Keith Hawkins, Catherine Manea, Jonathan J. Fortney, Caroline V. Morley, Andrew Skemer, Mark S. Marley, Brendan P. Bowler, Aarynn L. Carter, Kyle Franson, Zachary G. Maas, and Christopher Sneden. ELemental abundances of Planets and brown dwarfs Imaged around Stars (ELPIS). I. Potential Metal Enrichment of the Exoplanet AF Lep b and a Novel Retrieval Approach for Cloudy Self-luminous Atmospheres. , 166(5):198, November 2023.

[309] Andras Zsom, Sara Seager, Julien de Wit, and Vlada Stamenković. Toward the Minimum Inner Edge Distance of the Habitable Zone. , 778(2):109, December 2013.

Appendix A

Conceivably Useful Appendices

A.1 Examples of predictive control in the literature, broken down by method

A.1.1 Empirical orthogonal functions

Arguably the first predictive controller for astronomical telescopes is from a work by Dessenne, 1998 [58]. This uses a method quite similar to empirical orthogonal functions, except that it runs in closed-loop; instead of associating the full state of the wavefront in time it builds a similar history vector that includes both wavefront sensor residuals and DM commands from every iteration. The truth condition that the filter is optimized against uses the rejection transfer function

(i.e., a high fidelity pseudo-open loop reconstruction) to reconstruct the full wavefront from the DM commands and wavefront sensor slopes. This method was implemented on-sky in 1999 [57], at the 1.52 meter visible telescope at Haute Provence. They cite a 30% relative Strehl Ratio increase over a classic integrator, which maps to an increase of $\sim$ 11% to 14% Strehl Ratio over varying r_0 conditions at a central wavelength of 650 nm.

After Guyon proposed EOF in 2017 [93], it was implemented on the 8 meter Subaru/SCExAO in 2018 [95]. An example of predictive control on/off shows visual improvement when comparing two on-sky images, and provides a factor of up to 3 improvement in contrast at close-in inner working angles [46].

Shortly thereafter, EOF was also applied in simulation to the 8 meter VLT/SPHERE telemetry [278]. This application used telemetry data (in the form of wavefront sensor slopes and DM commands) to apply predictive control as if it were a post-processing method, and compare this to the actual on-sky performance of the integrator. van Kooten, 2020 found a 5.1 factor of improvement in RMS error over the residuals of the SAXO integrator, and estimated an additional 2.1 increase over an idealized AO system for VLT/SPHERE when applied to data taken at a central wavelength of 658 nm.

Finally, work started in 2019 [116] to apply this on-sky at the 10 meter Keck observatory. Performance improvements have been shown with a similar

post-processing telemetry application [116, 72], and van Kooten, 2022 [280] demonstrates an on-sky application. van Kooten, 2022 found a factor of 3 contrast improvement over the classic integrator at an inner working angle of 3-7 λ/D, and a 1.5 factor decrease in standard deviation when applied to L band coronagraphic observations, making the predictive controller consistently more robust than the integrator. A similar test also looked at Strehl Ratio from the Bracket γ narrowband filter, and found a Strehl Ratio increase from $\sim$ 19% with the integrator to $\sim$ 23% with the predictive EOF method.

A.1.2 Kalman filtering and linear quadratic Gaussian controllers

Haffert, 2021 [99] takes a Kalman filtering-like approach, with a subspace predictive controller. Subspace predictive control builds and updates a predictive Kalman filter, with a state vector built of wavefront sensor slopes and DM commands; this work could also be considered an extension of the Dessenne, 1998 [58] work, but without the need for rejection transfer function (or pseudo-open loop reconstructed) data to train against. This method, when applied to the MagAO-X daytime bench (an instrument for the 8 meter Magellan-Clay telescope) shows a contrast improvement by a factor of of up to 10-20 for low wind speeds and up to two orders of magnitude for higher wind speeds at a

central wavelength of 760 nm [99].

Linear quadratic Gaussian control was suggested as a controller for astronomical telescopes in 1993 [202], making more than 30 years of machine learning methods for wavefront control. The classic implementation of this method quickly gained popularity for the control of lower order modes. LQG runs as the standard low order controller for the VLT/SPHERE, where low order LQG performed with an RMS error of 0.87 mas on tip and tilt, as compared to a classic integrator at 12.8 mas [211]. The Gemini Planet Imager (GPI) also uses low order LQG control, in particular to control consistent and specific frequencies; for example one 27 Hz frequency induced a 19.5 mas tilt, and with the LQG 2.3 mas of tilt remained at a central wavelength of 1570 nm [219].

The only application of this method as a full controller on-sky was on the 4.2 meter CANARY AO demonstrator at the William Herschell Telescope, which reports a Strehl Ratio of 50.9% for a classic integrator, and 61.9% for the LQG controller in the H band [255]. A reformulated data-driven LQG [254] was also tested on CANARY/WHT, finding up to 7% Strehl Ratio improvement (from 0.067 to 0.051 radians RMS) when run on-sky in the H band.

Predictive Fourier control (PFC) [218] is another LQG alternative, which learns wind speeds from telemetry data (often called a wind-driven predictor). Its original implementation was for GPI-like simulations, which showed an

integrator residual RMS error of 37.2 nm, as compared to 21 nm for the predictive Fourier control [218]. Fowler, 2022 applied this method to Keck telemetry, finding an improvement in RMS error on by a factor of 2.5 when comparing PFC to on-sky integrator residuals [72]. Most recently lab results from the LLAMAS (Low Latency Adaptive Optical Mirror System) testbed show a decrease in temporal error power by up to 3, with varying wind speeds and spatial frequencies [217]. The predictive Fourier control method also allows for non-integer lags in the predictor, which makes for a highly tunable predictive corrector.

A.1.3 Neural networks

Some works have explored the use of neural networks (NN) for predicting the evolution of atmospheric turbulence or wavefront sensor telemetry [152, 297, 97]. These mainly focused on the prediction of (pseudo-)open loop measurements, and all show promising improvements over classical, non-predictive controllers (i.e., an integrator). We note that these methods build a nonlinear state-space model that predicts the evolution of turbulence, but do not build an optimal nonlinear controller; a classic integrator is often used as controller. From these works, Wong, 2021 [297] used a simple feed-forward NN, while Hafeez and Liu used Long-Short Term Memory networks to deal with the temporal aspect of

the data [97, 152].

Swanson, 2022 [266] used NN prediction in a closed-loop AO system, showing that in closed-loop there can be a mismatch between the training distribution and observed distribution, which can be mitigated by using an adverserial prior. Swanson, in particular, explored the role of guide star magnitude and increased sky coverage, showing that the NN method was able to show major performance improvements for a 16th magnitude guide star, from a Strehl Ratio of 40% with an integrator to 70% with the NN, for a simulated 8 meter telescope in the R band.

The challenge with nonlinear state-space models is that it is not trivial to find an analytical solution to the control problem with a quadratic cost function (unlike linear examples). To that end, reinforcement learning (RL) aims to derive optimal nonlinear controllers by e.g. using planning algorithms (model-based) or through direct optimization of a parameterized control law (model-free). An additional feature of RL algorithms is that one is not limited to a quadratic cost function of the state variables but is free to choose the optimization objective (i.e., reward), allowing one to directly optimize for e.g. Strehl or contrast.

Nousiainen, 2021 [195] used model-based RL in a simulated 24x24 AO system with a Shack-Hartmann and demonstrated an improvement in Strehl and

contrast compared to an integrator, showing that their algorithm is also able to deal with DM to wavefront sensor misregistrations, but notes that the algorithm may be too computationally expensive to run at kHz frequencies. A consecutive work (Nousiainen, 2022 [192]) solved this issue by using a parameterized control law; they not only show significant contrast improvements in a simulated VLT-like system with a PWFS, and decreased speckle variance in the lab with MagAO-X, but also run a simulations of a 40 meter telescope in real time with this controller, increasing performance from 88% Strehl with a classic integrator to 90% Strehl for the RL method, calculated at 1650 nm. Most recently, Nousiainen demonstrated these results the GHOST testbench and showed promising initial results in a cascaded AO system. [193]

On the other hand, in the work of Landman, 2020, and 2021 [145, 146], a model-free algorithm is chosen based on the Deep Deterministic Policy Gradient algorithm. Landman shows improvements in vibration mitigation in simulations and lab experiments and a large improvement in contrast in a simulated high-contrast imaging instrument on an 8 meter telescope (from 10^{-4} for the integrator to 2×10^{-5} for the RL method, calculated at 1000 nm). By utilizing a combination of an LSTM and CNN, the RL algorithm built a policy that adjusts to changing atmospheric conditions without online learning, making the controller highly adaptable to changing on-sky conditions.

Pou, 2022 [214] also uses a model-free approach using the Soft Actor-Critic algorithm and considers a distributed approach using multi-agent RL. Pou shows improvements in Strehl for a simulated 8-meter telescope with Shack-Hartmann WFS and PWFS, from 85% Strehl to 91% Strehl calculated at 1650 nm. [215] However, for the Pyramid, this performance required a nonlinear wavefront reconstructor before the RL controller.

A.2 Center-finding strategies for ORKID PSF jitter analysis

Center-finding strategies were used regularly in this analysis. PSF jitter estimates and recentering all require a reliable centering method. Here we explore the various centering strategies and explain our rational for the choice in this work.

Our goal for these methods is to recover the impact of tip and tilt zernike modes as they map to motion of the PSF via image jitter, so we estimate its performance by simulating an ORKID PSF at 656 nm and injecting the zernike modes tip and tilt to see how well an algorithm matches this recovery.

We then propagate a zernike mode through an open or closed loop PSF and through a "perfect" PSF with only the Keck primary geometry. Using

the perfect PSF as a truth condition on the injected motion, we compare the recovery of that motion via every `photutils` [25] center-finding method. Figure A.1 shows this recovery for the closed-loop simulations.

From the first analysis we rule out center of mass algorithms, as they are prone to error in the case of high background or halo. While brightest-pixel provides good performance, ideally we wish to recover sub-pixel accuracy, which is not possible with a single max pixel value. (We tested brightest pixel super-sampling methods but this increased the computational time beyond feasibility for hours of 5 ms exposures.) While `photutils.centroid_quadratic` showed the best recovery in our experiment, run over many frames it was prone to outliers, and for this reason we use `photutils.centroid_1dg` for this analysis. This method fits a Gaussian to a 2D array by fitting the x and y values each with a 1D Gaussian.

When tracking jitter in this work, two regimes are important: (1) tracking jitter on the order of $\sim$15 mas to track the nominal performance (which we ideally wish to do with high precision), and understanding large moves. We therefore consider the small and large error independently, as `centroid_1dg` performance worsens as the move increases. We show that with this method we can track tip/tilt motion down to 0.7 mas for jitter under 20 mas and to 4 mas for moves under 50 mas, as shown in Figure A.1.

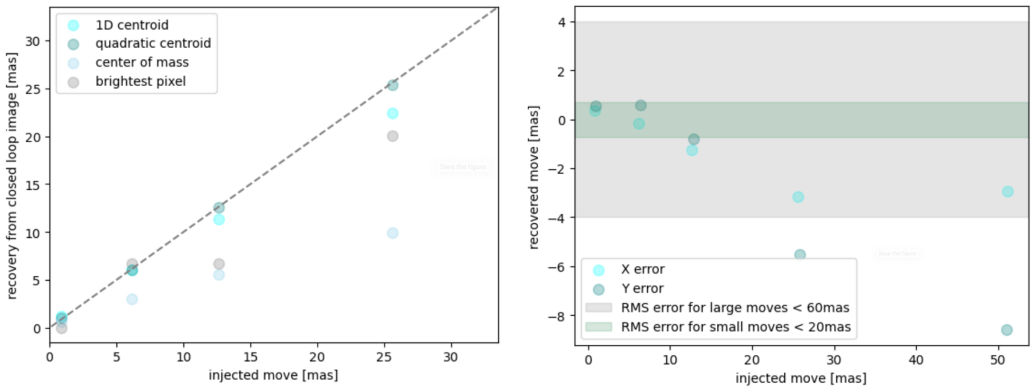


Figure A.1: Left: Injection and comparison for different center-finding methods. Right: Residuals of the `centroid_1dg` method. This puts limits at 4 mas for moves under 50 mas and 0.7 for moves under 20 mas.

A.3 Math mistakes in fundamental papers

A.3.1 Paschall & Anderson, 1993

The first paper to describe applying an LQG controller to an astronomical telescope system. The quadratic cost function described in Eq (17) of the original paper [202], written as:

$$\mathbf{J} = \mathbf{E}\left\langle {}^{N}_{i=0}\frac{1}{2}\left[\mathbf{x}^{\mathbf{T}}(t_i)\mathcal{X}\mathbf{x}(t_i + \mathbf{u}^{\mathbf{T}}(t_i)\mathcal{U}\mathbf{u}(t_i)\right]\right\rangle \tag{A.1}$$

is missing a close paren, it should read:

$$\mathbf{J} = \mathbf{E}\left\langle {}^{N}_{i=0}\frac{1}{2}\left[\mathbf{x}^{\mathbf{T}}(t_i)\mathcal{X}\mathbf{x}(t_i) + \mathbf{u}^{\mathbf{T}}(t_i)\mathcal{U}\mathbf{u}(t_i)\right]\right\rangle \tag{A.2}$$

This is an arguably obvious mistake, as the vector $\mathbf{x}$ must apply to some discrete timestep.

A.3.2 Dessenne et al., 1998

Many credit Dessenne, 1998 [58] as the first paper to describe a fully predictive controller. It builds a state vector ϕ, which holds both wavefront sensor residuals and deformable mirror commands. Equation (20) of the original work lays out the state vector and the predictive filter as:

$$\theta = [b_1 ... b_{p-1}, a_0 ... a_{q-1}]^T$$

$$\phi(i) = [y(i-1) ... y(i-p+1)]$$

$$(y - e'_{meas})(i-2) ... (y - e'_{meas})(i-1-1)]$$

This expression has an extra bracket inserted. It should read:

$$\theta = [b_1 ... b_{p-1}, a_0 ... a_{q-1}]^T$$

$$\phi(i) = [y(i-1) ... y(i-p+1), (y - e'_{meas})(i-2) ... (y - e'_{meas})(i-1-1)]$$

The entirety of the second two lines of the originaly Equation (20) make up a single expression for the state vector. This error has high potential for confusion

as one could interpret it as a third expression, missing its definition, or two series of terms that should be multiplied as cross terms.

A.3.3 Poyneer et al., 2007

Poyneer, 2007 [218] is one of two major predictors explored today. It makes use of a particular modal basis, complex Fourier modes. The expression for complex Fourier modes in the original work from Equation (1) is written:

$$f_{k,l}[x,y] = \cos\left(i2\pi\frac{kx+ly}{N}\right) + i\sin\left(i2\pi\frac{kx+ly}{N}\right) \tag{A.3}$$

which has two extra imaginary terms inserted. It should read:

$$f_{k,l}[x,y] = \cos\left(2\pi\frac{kx+ly}{N}\right) + i\sin\left(2\pi\frac{kx+ly}{N}\right) \tag{A.4}$$

Those familiar with complex modes would likely spot this mistake immediately, but due to the waning use of complex modes it may be easy to forget where the imaginary terms lie. It should also be noted that the cos and sin of imaginary numbers provides a real answer, so it is not immediately obvious in debugging results mathematically that this is incorrect.

Appendix B

Foolish Musings

B.1 In violation of the prime directive: simulating detriments to Delta-Quadrant civilizations from the starship Voyager's impact on planetary rings

B.1.1 Abstract

In the seven years that the starship Voyager spent in the Delta Quadrant, it used many questionable techniques to engage with alien civilizations and ultimately find its way home. From detailed studies of their logs and opening

credits, we simulate Voyager's practice of orbiting a planet, to examine the effect on planetary rings. We outline a feasible planetary system and simulate the extent to which its rings would be disrupted. We find that Voyager's orbit could inflate the height of the rings in the vicinity of the spacecraft by a factor of 2, as well as increase the relative speeds of neighboring planetesimals within the rings. This increase in ring thickness has the potential to alter shadows on any moons of this planet, impacting ring-shadow based religions. Additionally, the acceleration of these planetesimals could rival their gravity, bucking any alien inhabitants and their tiny civilizations off of their planetesimal homeworlds. Finally, we posit that due to increased collisions amongst the planetesimals (which may harbor tiny intelligent life) the trajectory of these civilizations may be forever altered, violating the prime directive.

B.1.2 Introduction

The formation and evolution of planetary rings has been highly studied and modeled over the years (for a review, see Refs. [84, 244]). Nbody simulations are a common tool to model the effect of rings given gravitational forces over time, including due to disruptions resulting from the changing orbits of moons. However, to date no published literature has tackled the impact of intelligently-driven foreign bodies on planetary dynamics and ring evolution. Here, we

consider the problem of of how planetary rings may change after encountering the Starship Voyager.

Star Trek: Voyager's [16] opening credits feature a collection of Voyager's exploits, presumably representative of their seven year journey in the Delta Quadrant of the Milky Way Galaxy. Much of the opening credits feature Voyager interacting with gas and dust, carving paths through nebulae and zooming around planetary rings. We aim to carefully treat a Voyager-planet interaction from the opening credits and examine impacts to the ring system. The Planetary Federation's Prime Directive states that no member of Star Fleet may interfere with the evolution and natural development of alien civilizations. **We examine to what extent simulations show that Voyager in this opening credit scenario could violate the Prime Directive.**

In Section B.1.3 we motivate our problem given a reconstruction of the opening credits and describe our initial assumptions to model a Voyager-planet interaction. In Section B.1.4 we describe numerical integrations of planetary ring dynamics conducted using the Nbody code `REBOUND` [226, 227]. In Section B.1.5 we discuss potential impacts of ring-disruption and consider violations of the Prime Directive. Finally, in Section B.2.6 we make our conclusions and summarize recommendations for the protection of Delta-or-any-Quadrant aliens.

B.1.3 Initial conditions of a voyager-planet system

To estimate Voyager's impact on the rings of a given planet, we use a scene from the opening credits (shown in Figure B.1) to benchmark a typical Voyager-planet interaction. Assuming Voyager is the size of a middle school (we find this to be feasible based on the crew compliment it could support before the catastrophic run-ins of the first few episodes) we size Voyager at ~ 200 meters long. Following standard theoretical procedure, we model Voyager as a sphere of radius $R_{\rm Voy} = 100m$. As illustrated using Voyager's reflection on the ring surface (Figure B.1, right panel), the spacecraft is approximately three Voyager-lenghths ($6R_{\rm Voy}$) above the ring surface. Using Voyager's size relative to the planet in Figure B.1, we rule out gas giant planets. The planet is blue in the opening credits, so we use a planet mass, M_p, and radius, R_p, equal to those of Neptune for our estimates.

The ring geometry in our simulations is illustrated in Figure B.2. For rings extending from 1.35–2.28R_p (left panel), comparable to the rings of Saturn, we identify a viewing geometry comparable to that shown in Figure B.1. However, given this geometry, Voyager itself must be enhanced in size by a factor of 1000 to have an appropriate size. We expect that image enhancement [248] was used to magnify Voyager in the opening credits sequence. An alternative that does not require image enhancement is provided in the right-hand panel.

Here, Voyager is appropriately sized, but the rings are quite thin, extending from 1.3999 to 1.4001 planetary radii. In both viewing geometries, Voyager orbits at a distance of $1.4R_p$. We prefer the wider rings due to a better visual match, but we note that our results for the dynamics of ring particles in the vicinity of Voyager are independent of which geometry we choose.

We assume Voyager's bulk composition resembles the density of Titanium at 4.5 g/cm^3. While Aluminum is more commonly used in spacecraft because it's light and sturdy, we assume that 4800-era Titanium alloys could provide the stability needed for warp-speed travel. With improved space mining and energy sourced from Dilithium crystals, we project an enriched Titanium alloy would be a suitable material. Voyager is not entirely metal, but likely contains $\sim$ two thirds empty space. Given this porosity, $\phi_{\rm Voy} = 0.66$, the spacecraft's density is $\rho_{\rm Voy} = 1.53$ g/cm^3, and its mass $M_{\rm Voy} = 6.41 \times 10^{12}$ g $= 1.07 \times 10^{-15}$ Earth masses. We note that the Hubble Space Telescope has a density of only 0.06 g/cm^3, but Voyager requires greater robustness because it travels faster, further, and has to survive more hull impacts.

While the size of particles in Saturn's rings can be as large as 1km, the typical size is 2-12 m (made mostly of water ice). **(author?)** [111] argue that this maximum size is determined by thermal stresses and suggest that ring particles in the Neptunian region should also have similar maximum sizes. We

therefore assume that the largest ring particle sizes evident as the opening credits sweep through the rings is 1–10 m. However, unlike Saturn's rings, which have a typical thickness of 10m (outside of radially-thin regions where perturbations by shepherding moons produce "walls" that can reach 1km in height), these Delta-quadrant rings appear to be ∼30–100 particle-sizes thick. This puzzling difference may come from a distinct coefficient of restitution in Delta quadrant rings. Indeed, [244] find that while rings modeled using a "frosty" elasticity model are expected to be 10 particle-sizes thick, their preferred "smooth" elasticity model is consistent with a 30 particle-size-thick ring that is marginally optically thick.

For a full-width ring thickness of 60m, our estimate that Voyager is 3 Voyager lengths above the rings (Fig B.1), puts the starship 10 ring thicknesses above the rings, consistent with the sweeping view provided in the opening credits. These starting assumptions are listed in Table B.1.

We simulate the impact of Voyager orbiting for 1 Earth-week on 2000 nearby planetary ring particles. The ring particles (which may be 1–12 meters across at most) are small with respect to Voyager, so they are simulated as test particles For computational efficiency, we restrict our simulated ring particles to orbits with semi-major axis randomly drawn from a uniform distribution between -2.5 and 2.5 $R_{\rm Voy}$ from Voyager's orbital semi-major axis at $1.4R_p$.



Figure B.1: Left: Video evidence from the opening credits allow us to estimate initial conditions. Right: Estimated height of Voyager above the rings of 3 Voyager lengths, based on reflection. Figure adapted from Star Trek: Voyager [16]

The particles are assigned random inclinations from a uniform distribution 0 to H/a_{Voy}, where $H = 30$m is the half-thickness of the rings and $a_{\text{Voy}} = 1.4R_p$ is Voyager's orbital distance. Longitudes of pericenter, ω, and longitudes of ascending node, Ω, are drawn uniformly between 0 and 2π radians, and ring particle true anomalies, f, are chosen such that the initial angular position of the particles ranges uniformly between ± 0.0002 radians from the angular position of Voyager. Voyager is initialized to be at its maximum vertical separation from the rings, with $\omega = 0$, $\Omega = \pi$, $f = \pi/2$, and inclination $H_{\text{Voy}}/a_{\text{Voy}}$ where $H_{\text{Voy}} = 600$ m.

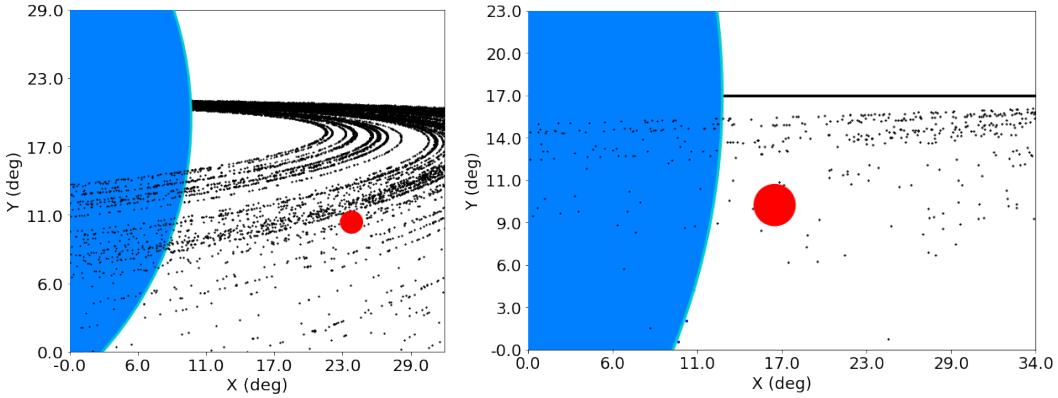


Figure B.2: Initial conditions of the simulations for the `REBOUND` setup. Left: Note that Voyager is 1000x too big in this plot for visual comparison with the original image. We expect that image enhancement [248] was used to magnify the Voyager and the region around it – leaving our use of Voyager as a measuring stick credible but allowing us to pick on other things we didn't like. Right: Here Voyager is a realistic size with respect to the planet, but the rings appear much thinner than in visual comparison. We expect this may be due to artificial projections of larger size, like that of the alien society in Catspaw [232].

Parameter	Estimate	Justification	Notes
planet mass	1.02×10^{29}g	Neptune mass	planet is blue
average ring radius	4.50×10^{9}cm	2 Neptune radii	Saturn-like rings
Voyager diameter	200 m	$\sim$ middle school sized	bigger than a beach motel but smaller than a community college
Voyager height	600 m	3 Voyagers above the rings	see Fig B.1
Voyager density	1.53 g/cm^3	4.5 g/cm^3 at a porosity of 0.66	Dilithium strengthened Titanium alloy

Table B.1: Starting assumptions of the n-body simulation.

B.1.4 Simulated Voyager impacts

We integrated the simulation forward using the orbital dynamics code `REBOUND` with the `ias15` integrator [226, 227]. We allow Voyager to orbit the planet for 7 days, a typical visit duration, and then examine impacts to the rings. Results are shown in Figures B.3 and B.4.

Ring puffing

The major impact of the Voyager-planet interaction takes the form of increasing the thickness of the rings, or ring-puffing. After Voyagers 7-day orbit, the starship carves a small gap in nearby particles, both azimuthally (Figure B.3) and radially (Figure B.4. Just beyond the gap, we see the nearby ring thickness increase by a factor of 2, as shown in Figure B.4.

The ability of Voyager to scatter nearby ring particles may be understood using the impulse approximation. The initial "random velocities" of ring particles (velocities relative to Keplerian) are approximately

$$v_{\rm rand} = H\Omega \sim 1{\rm cm/s} \ , \tag{B.1}$$

where $\Omega = (GM_p/a_{\rm Voy}^3)^{1/2}$ is the orbital angular velocity. The vertical velocity

kick induced by Voyager over one 4.27-hour orbital period is approximately

$$v_{\rm kick} \sim \frac{GM_{\rm Voy}}{H_{\rm Voy} v_{\rm rel}} \sim \frac{2GM_{\rm Voy}}{H_{\rm Voy}^2 \Omega} \sim 0.6 {\rm cm/s} \ , \tag{B.2}$$

where we have taken $v_{\rm rel} \sim H_{\rm Voy}\Omega$.

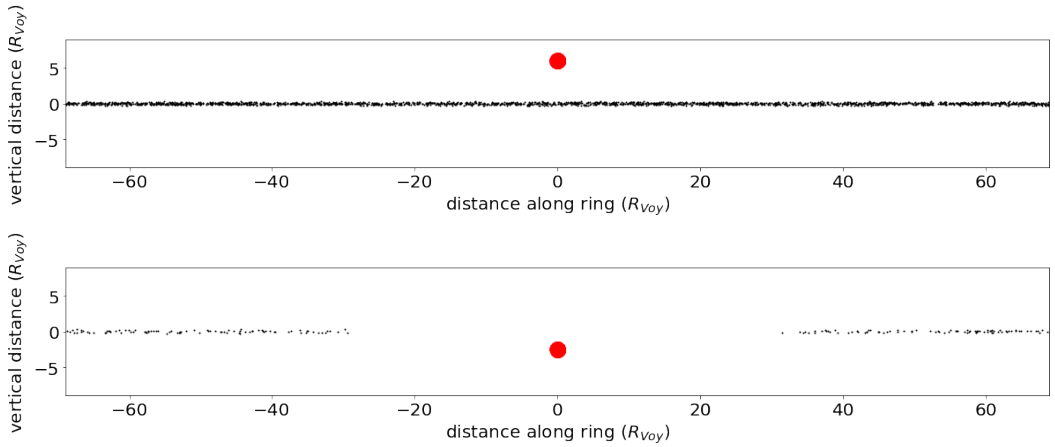


Figure B.3: Ring particle heights as a function of azimuthal distance along the ring (black points). Initially (top), Voyager (red circle) lies above the ring. After 7 days (bottom), Voyager's orbit has cleared an azimuthal gap extending a few tens of starship radii in both directions. Voyager size is to scale.

Increased velocity and collision rate of ring particles

The random velocities of the particles will increase with ring height (as described in Equation (B.1) ; with a 2x increase in ring height, we project the particle speed will increase by approximately a factor of two from its initial value of ∼1 cm/s. This increase in speed over 7 days will create a puff-induced acceleration

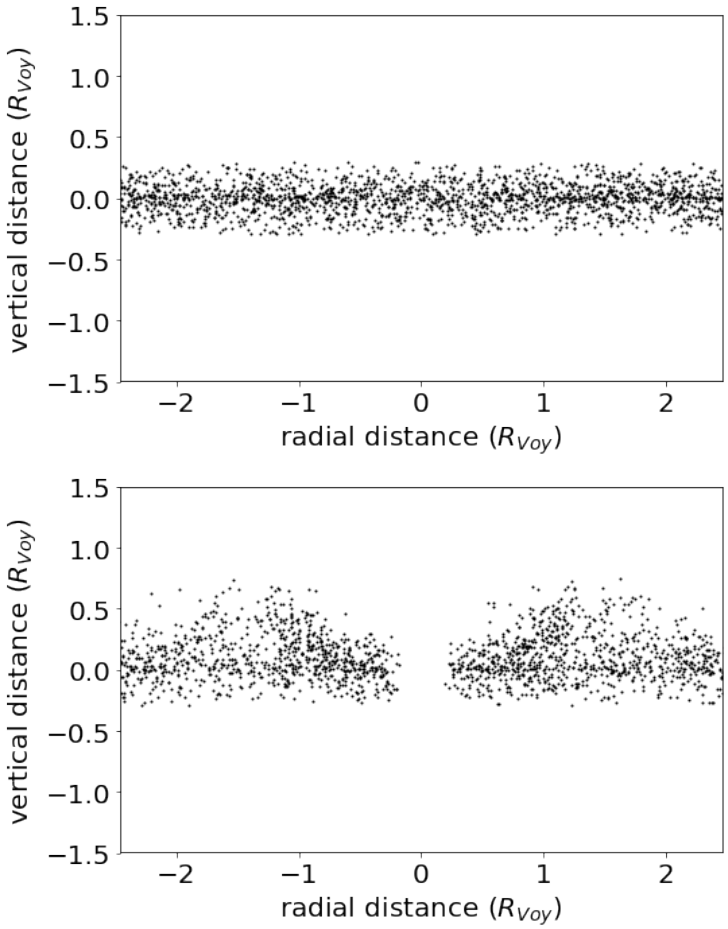


Figure B.4: Ring particle heights as a function of radial distance from the center of the planet. Distances are measured with respect to the orbital separation of Voyager. The smooth initial ring distribution (left) is vertically excited in the neighborhood of Voyager by approximately a factor of 2 (right) due to scattering by the starship. A small radial gap is also evident at the orbital distance of Voyager.

of $\sim 10^{-5}$ cm s^{-2}. We note that the acceleration of some particles likely occurs on less than a 7-day timescale, so this estimate is a minimum acceleration.

In Figure B.5 we estimate a number density of particles in the rings. We

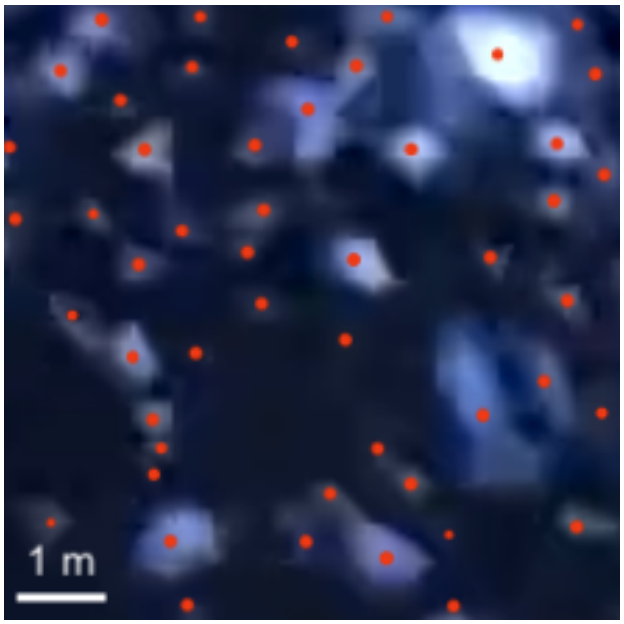


Figure B.5: Estimate of particle number density in the ring. We count 50 particles in this 5 meter (or 5 median-particle-size) box.

assume the number density of particles stays constant across the puffing as particles ejected from the gap pile up at the gap edges, making up for the increased height.[1] Using a median particle as 10 meters across we count 50 particles per 5x5m box (we count the box size as 5 median particles across), and take 30 meters (the ring height) as the depth, giving us a final number density of $n = 6.67 \times 10^{-8}$ cm^{-3}. Using a geometric cross-section for collisions σ and relative velocities, $v_{\rm rel}$ comparable to the random velocities of the disk

[1]This choice was entirely physically motivated and has nothing to do with the fact that we only have one screengrab of the opening credits from which to estimate number densities.

particles, we calculate the collision rate:

$$\text{collision rate} = n\sigma v_{\text{rel}} = 0.42 \text{ s}^{-1} \tag{B.3}$$

meaning we expect a collision of ring particles approximately every 2.5 seconds (where previously, it was only every 5 seconds). This results in 1.5×10^6 more collisions due to Voyager during its 7 day orbit.

B.1.5 Violations of the prime directive

Seeing the notable effects that Voyager has on this ring system, we fear Voyager is likely to have a major impact on any life that exists within this ring system. Below we outline some cases in which the puffing of the rings, the increased velocity of the particles, and the increased collisions of the particles could violate the Prime Directive.

Creation of a super-space virus

As evinced by numerous existing studies on space viruses [232, 233, 15, 16], one of the greatest dangers to space travel are viruses that have ranged from intoxication so powerful it leads to death (e.g., PSI 2000 of [232]), viruses that have disrupted communication (e.g., Aphasia Virus of [15]), and viruses so large you can see them (e.g., Macroscopic virus of [16]). With increased collision

of ring particles, foreign matter that would have otherwise remained dormant will be more likely to interact with other matter – laying the groundwork for a super-space virus built from many ring-particle's worth of viral space matter. This could not only impact the crew of Voyager and other travelers to the region, but bring disease to any aliens that lived on moons or on the small ring particles.



Figure B.6: Left: Morphology of tiny aliens after their humanoid projection is undone Figure adapted from Star Trek: The Original Series [232]. Right: Humanoids and the Defiant shrunken down to tiny size. Figure adapted from Star Trek: Deep Space Nine [15].

Disruptions of tiny intelligent life living on ring particles

Given previous surveys of alien civilizations (e.g., Catspaw aliens in [232]), as well as the feasibility of making humanoid-size aliens tiny (e.g., circumstance of One Little Ship in [15]) it is not unlikely that the small planitesimals could

harbor tiny intelligent life, see Figure B.6 for examples.

Creatures living on planetesimals and ring particles would have extremely low surface gravity; for the case of a 1 meter diameter particle with a density of 1 g/cm^3, their gravity would be $g_p = 1.4 \times 10^{-5}$ cm/s^2. This means that any of these creatures would be particularly vulnerable to acceleration of their planetesimals by increased particle velocity. The minimum Voyager-induced acceleration of 10^{-5} cm/s^2 estimated in Section B.1.4 which is comparable to the planetesimal gravity.

Furthermore, due to the aforementioned increased collisions, the likelihood of two alien species who otherwise may not interact being thrust into each other increases greatly – forever altering the course of their civilization. This could be notably traumatic for tiny aliens species, as one species could be **especially** tiny compared to another that was only somewhat tiny.

Interruptions to speculative shadow-based religions on ring-embedded moons

If a moon were embedded in this planet's rings, the rings would cast a prominent band of shadow on the moon. We expect that for intelligent forms of life this shadow would inspire religious worship or festivals, and be closely monitored for seasonal variability. A two-times increase in the size of this shadow would

certainly be noticeable. Saturn's rings contain abundant small gaps carved by moons, and similar gaps are clearly evident in Figure B.1. Consider, for example, a moon located in the middle of a 30km-wide gap (similar to Saturn's Russell or Jeffreys gap). We estimate the shadow cast by the 60m-thick ring would take up ~ 0.1 degrees in the sky. However if Voyager is orbiting near the gap edge, then for the duration of the puffed rings, that would increase to ~ 0.2 degrees, nearly half the size of our sun in the sky. This changing ring shadow could be comparable to the Ferengi False Profits [16], or the Earth Battle of Hayls [109].

B.1.6 Conclusions

In conclusion, we predict that Voyager could irreparably alter the course of many civilizations with its habit of orbiting planets. By simulating Voygaer as a middle-school sized object for 7 days in orbit around an opening-credits-like planetary system, we find that the system would undergo a ring puff that increases both the ring height and the random velocities of ring particles by up to a factor of 2 in the vicinity of the starship. The increased collision rate of the particles in this system could lead to a greater confluence of space material, potentially creating a super space virus. In terms of intelligent life, we predict that any tiny intelligent creatures on planetestimals within the rings

would suddenly be thrust into each other, artificially disrupting the natural flow of their society. For smaller planetesimals, the puff-induced acceleration of these particles will be comparable to if not greater than their gravity, with high potential of bucking civilizations clear off their home-particle. These all represent clear violations of the Prime Directive, and should be addressed by future Star Fleet regulations. We recommend clear policies that avoid orbiting ringed planets with any signs of life—intelligent or otherwise.

B.2 Wavefront sensing and control with the many headed hydra modal basis

B.2.1 Abstract

The future of space and ground based telescopes is intimately tied to technology and algorithm development surrounding wavefront sensing and control. Only with cutting edge developments and unusual ideas will we be able to build diffraction limited observatories on the ground that contend with earth-atmosphere, as well as space-based observatories that sense and control for optical aberrations and telescope jitter. There exist a variety of mathematical bases for decomposing wavefront images, including the zonal modal basis,

Zernike polynomials, and Fourier modes. However, previous bases have neglected the most vital element of astronomical optics: our field is only as good as the people in it. In this paper we propose a new Many Headed Hydra Modal Basis (hydra heads) for wavefront decomposition, using physical representations of the Adaptive Optics community. We find that the first author makes the best wavefront decomposition, on the order of a correct reconstruction within $\sim 1\%$ of the original turbulence image. We discuss engineering implications of the Many Headed Hydra Modal Basis, as applied to deformable mirror technology and within active and adaptive optics control loops. Finally, we explore predictive control avenues, by positing that the first author being the most effective wavefront predicts a successful future for them as a high contrast imaging scientist.

B.2.2 Introduction

Future advances in astrophysics depend on improving performance our our telescopes and instruments, ranging from verifying sub-Earth exoplanets around surprisingly nearby stars [160], to identifying floofs at optical wavelengths [170], to discovering exigent threats in galaxies far, far away [159]. The performance of these instruments is inextricably entwined with how well we are able to sense and correct for aberrations in our incoming light – be it from earth-atmosphere,

errors built into the telescope system, or often both. Given we start with some representation of phase from an incoming wavefront, and hope to correct it, the next step in this sensing and control problem is a choice of basis. A zonal modal basis [250] is intuitively simple – breaking down each wavefront sensor subaperture into its own mode and controlling along zones. Whereas, a Zernike modal basis [262] builds up an intuition of optical aberrations, framing incoming light in terms of coma, astigmatism, trefoil, etc, that maps to optical defects within our system. Finally, a Fourier modal basis [65] decomposes light into sine and cosine waves, and has a natural connection to Fourier transforms into frequency space.

We present the Many Headed Hydra Modal Basis (hydra heads), as a basis that encapsulates the true character of adaptive optics, by encoding the likeness of the very scientists who do it a basis sets. With Section B.2.3, we work through the formalism of the hydra heads. In Section B.2.4 we examine the effectiveness of the hydra heads as used to decompose a sample turbulence image. In Section B.2.5, we discuss wavefront control applications of the hydra heads. Finally, in Section B.2.6 we make our final conclusions about why the first author makes an extremely good basis set.

B.2.3 Formulation of the many headed hydra modal basis

Given 16 starting candidate heads, we build the modes by:

1. Making the head into a square image.
2. Converting the image from color to black and white.
3. Rescaling the image from 0-1.
4. Appending a negative valued copy of the image.
5. Increasing the spatial frequency (number of heads) and altering the orientation.

We display one sample hydra head in Figure B.7, built with the Debes head.

The `hydra-sim` software used to create each hydra head is publicly available to turn any input head into a modal basis. Section B.3 scans through a few examples of modes for each input head.

B.2.4 Efficacy amongst the hydra heads

Given an input set of 16 heads, we built a Many Headed Hydra Modal Basis for each head, as described in Section B.2.3. With this suite of hydrae, we can

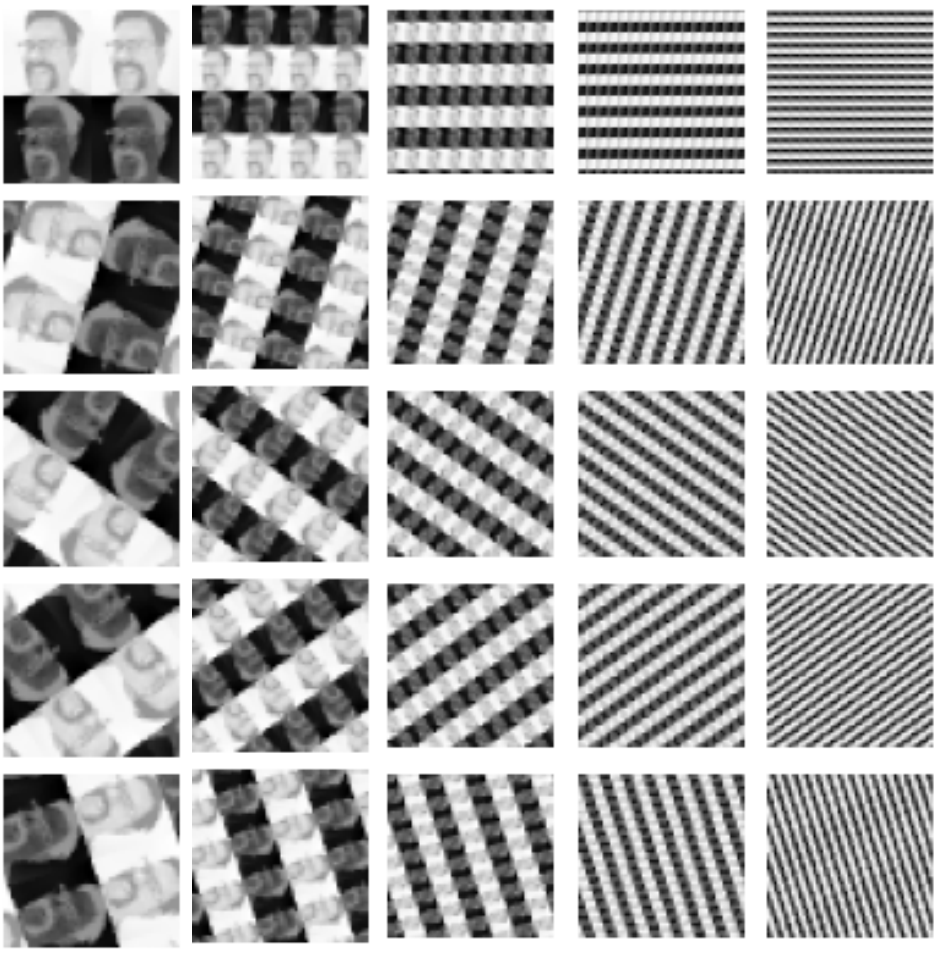


Figure B.7: Example Many Headed Hydra Modal Basis set using a Debes head as starting input. While this full set has 2500 modes (spanning tighter degree increments as the pattern rotates and more repetitions of heads) we display a handful of modes to get an idea of the modal basis images.

then test them practically, by decomposing a sample turbulence image. Using `HCIPy` [213], we simulate a turbulence profile with 5 wind layers, each with their own velocity and r_0. Table B.2 outlines the atmospheric parameters used to simulate this atmosphere.

Layer	Velocity m/s	Direction degrees	r_0 cm
1	22.7	246	38.9
2	3.28	71	44.7
3	16.6	294	45.5
4	5.89	150	38.8
5	19.8	14	43.6

Table B.2: These conditions were used to simulate an atmospheric turbulence profile as a practical test of the hydrae. We used conditions from [218], because the first author happened to have simulated this atmosphere at high resolution recently. Our final images are simulated at a wavelength of 1630 nm over an 8 meter primary telescope mirror.

Given each hydra head mode starts as a square m by m matrix, which could map to a representation of a turbulence profile binned down to a resolution of m x m, we make the Many Headed Hydra Modal Basis set $\mathbf{A}$, by flattening each mode to be a column of $\mathbf{A}$. From there, decomposing a turbulence image becomes a linear algebra problem of the form:

$$\mathbf{A}\vec{x} = \vec{b} \tag{B.4}$$

where $\vec{b}$ is the flattened image we aim to reconstruct, and $\vec{x}$ contains the coefficients of each mode to reconstruct the given image. We solve for $\vec{x}$ using a pseudo-inversion of $\mathbf{A}$, using `numpy`'s singular value decomposition pseduo-inverse method.

For this test we generated 2500 modes for each head, each a 50x50 scene,

which was then used to decompose an image simulated at 100x100 resolution and binned down to match the 50x50 (to preserve Nyquist sampling). Table B.3 shows the relative performance of each hydra head. While the Hirschauer head displays the best performance, a closer inspection of the original input images reveals that the Hirschauer image benefits from a flat background. Considering the unreliability of dome flats, we rule this out as a potential hydra head, and find that the Fowler hydra head is the most effective modal basis in this set. Figure B.8 shows an example turbulence image and the reconstruction using the Fowler hydra head. Figure B.9 shows the difference image between the reconstruction and the original turbulence.

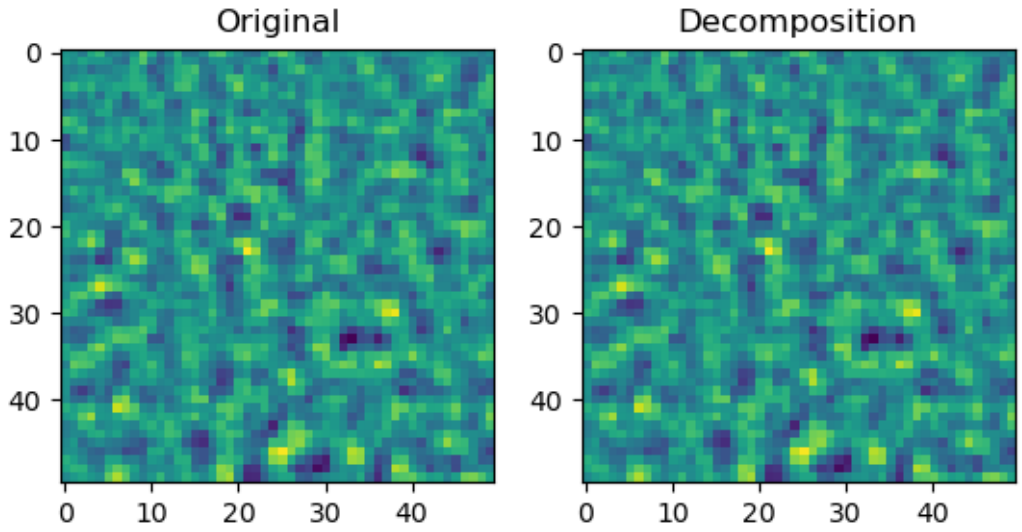


Figure B.8: Left: Original scene of turbulence, originally generated at a resolution of 100 pixels under the conditions described in B.2, and binned down to 50. Right: Reconstruction using the Fowler hydra head.

Head	% Error
Hirschauer	0.09
Fowler	1.07
Hinz	1.54
Soummer	2.47
Macintosh	3.89
Leboulleux	4.34
Pearce	4.93
Stelter	5.02
Jensen-Clem	5.33
Debes	5.75
Millar-Blanchaer	10.40
Steiger	12.34
Lewis	13.92
Balmer	15.66

Table B.3: The relative performance of the sample heads. We decompose a sample turbulent wavefront with a hydra head, and calculate the percent difference in the reconstructed image and the original simulated image.

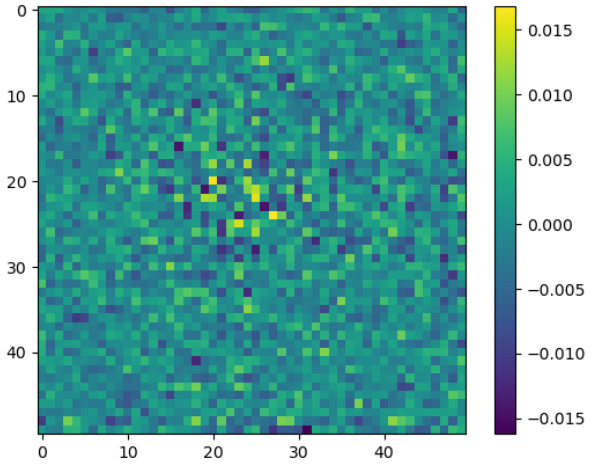


Figure B.9: Difference image between the Fowler hydra head reconstruction and the original turbulence. Note the scale on the order of 0.01 radians in the infrared regime.

B.2.5 Control applications

The hydrae are well-suited for pupils produced by off-axis mirrors (such as OAPs or OAEs) because the pupil planes of off-axis mirrors are tilted with respect to the beam direction resulting in an elongated pupil. The pupil, being non-square, is thus well-matched to a craftily-cropped head shot, although care must be taken to ensure that the head shot and the pupil share the same axial length ratio. Deviations can be systematically addressed by adding more fractal layers, which reduces the problem to one where the solution is known.

While head shots of prominent scientists[2] may not, at first, seem to be a sensible basis for wavefront control, we note that, from the perspective of a deformable mirror (DM) actuator, following a leader is second nature. After all, a 'follower/leader' actuator scheme is a common choice in designing wavefront control loops. We have simplified this scheme by finding the best head shot and designating it the leader, allowing the DM actuators to act in accordance with their natural predilections.

Edge effect artifacts of the head shots can affect the final result. We have implemented a simple fractal-headed approach described below, simply allowing the higher frequency features to be fit with (spatially) copies of smaller head shots as needed.[3] Determining the lower end cut-off frequency scale when using the hydra heads requires the use of a carefully-timed application of an integrating sphere fed by a high-powered incandescent lamp or laser, or cryogenic liquid (we note that open flame light sources are considered more traditional). With some trial and error, the suitable cut-off frequency can be found experimentally with only mild burns.

[2]Or at least the ones who responded to a late-night call on Twitter for head shots.
[3]We expect that snapshots of stacks of turtles would also be effective. [86]

B.2.6 Conclusions

In conclusion, we find that decomposing turbulence with pictures of astronomers is surprisingly feasible, with a shockingly high variability. While it seemed likely that this project would essentially test how good heads are at being sin waves, further inspection of the individual bases (see Appendix B.3) shows a striking similarity at higher order modes, as only a handful of pixels can be used to represent an image that requires many features to resemble a head. This potentially implies that the lower order modes are having a large impact on this decomposition.

Finally, we stand by the conclusion that how well your head can decompose turbulence is strongly predictive of scientific output, with a focus on how well the first author can reconstruct the sample scene.[4] No formal statistics have been conducted on this sample to verify this, and we do not recommend this for future works.

[4]This effect has no relation to Fowler sampling (a commonly-used scheme for reading out infrared detectors), as far as we are aware.

B.3 Full suite of hydra heads

B.3.1 Hirschauer basis

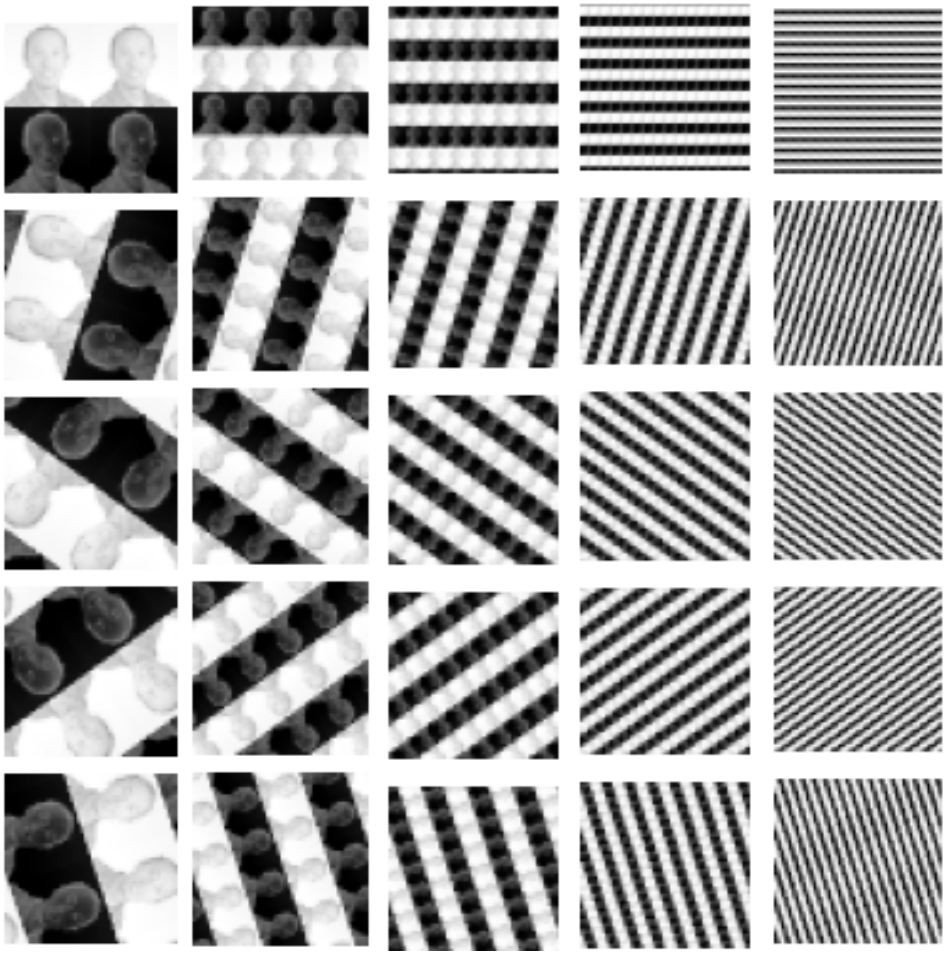


Figure B.10: Hirschauer basis.

B.3.2 Fowler basis

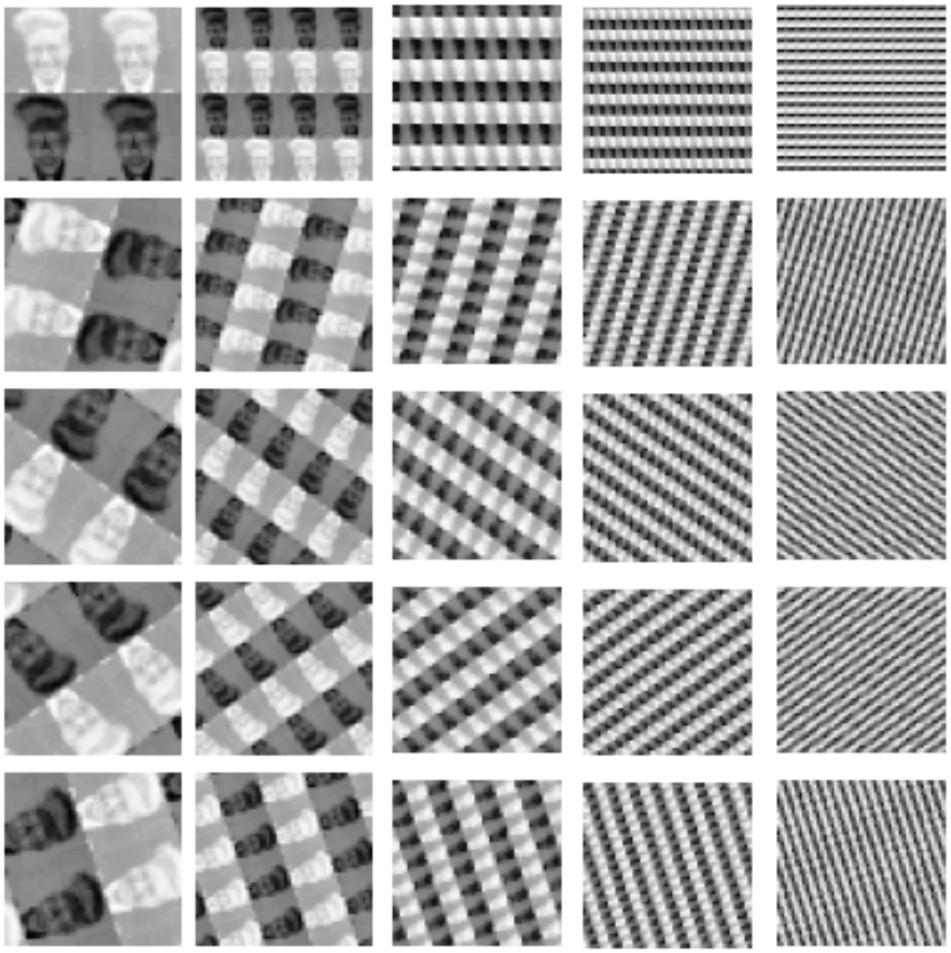


Figure B.11: Fowler basis.

B.3.3 Hinz basis

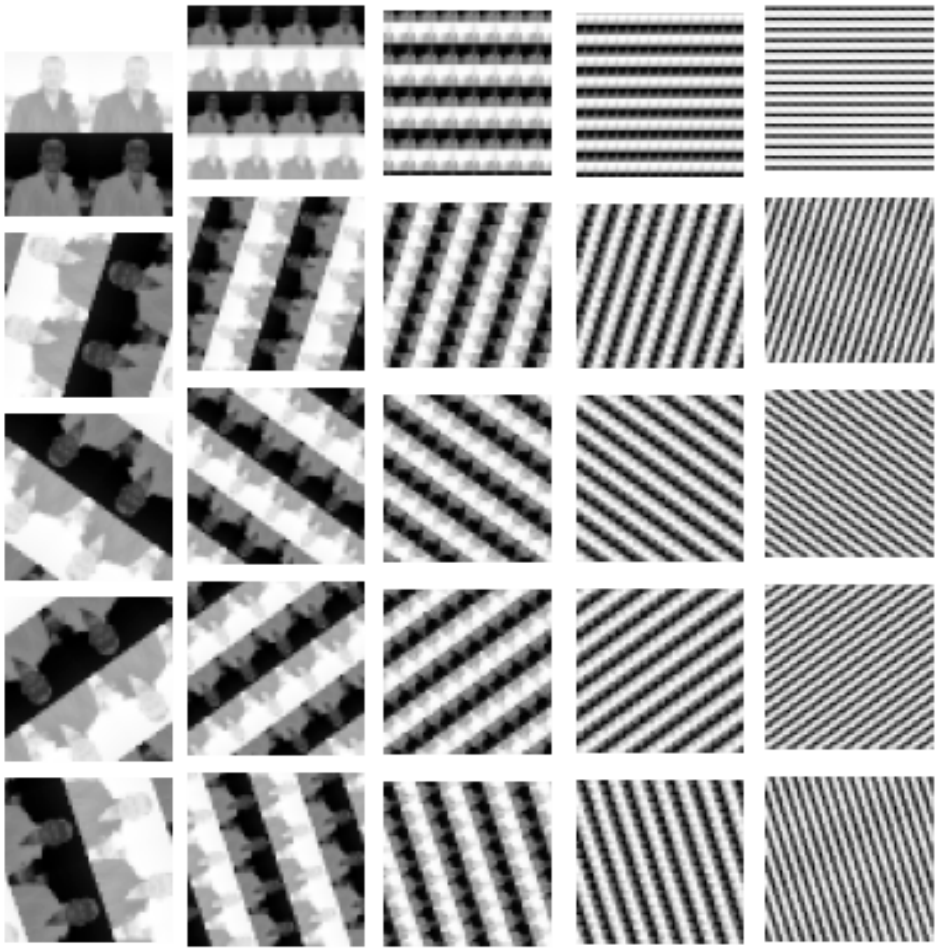


Figure B.12: Φllip Hinz.

B.3.4 Soummer basis

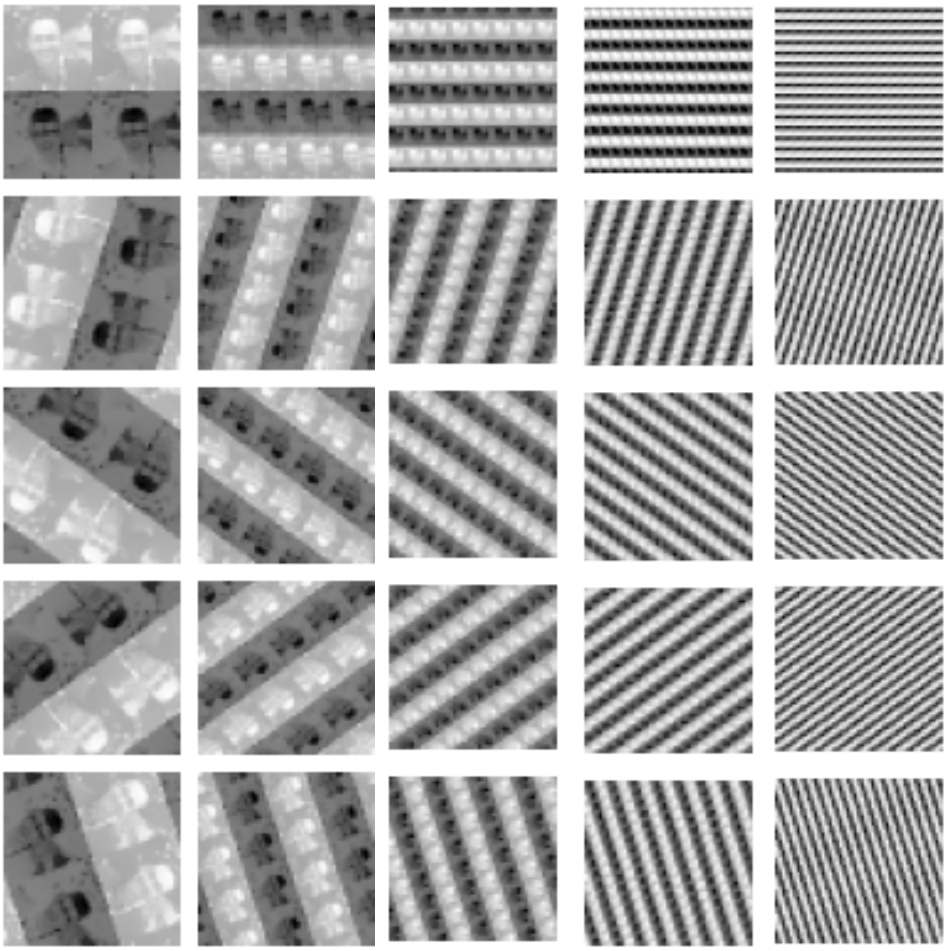


Figure B.13: Soummer basis.

B.3.5 Macintosh basis

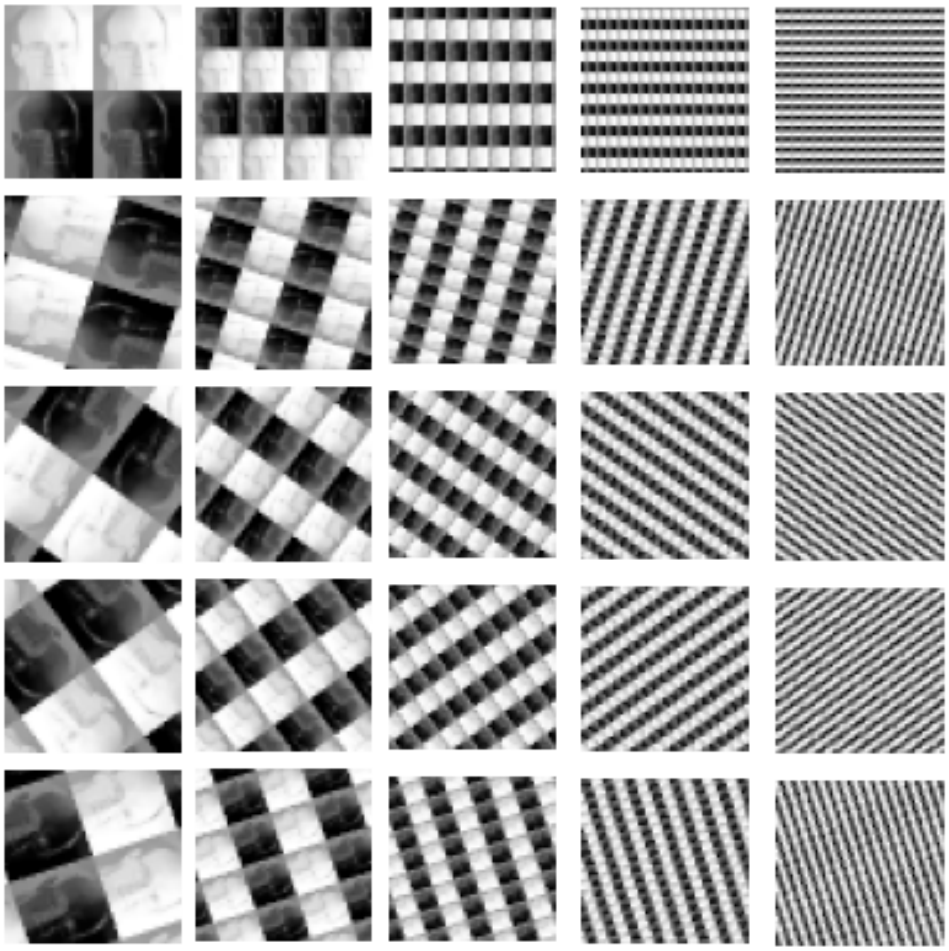


Figure B.14: Bruce DMacintosh.

B.3.6 Leboulleux basis

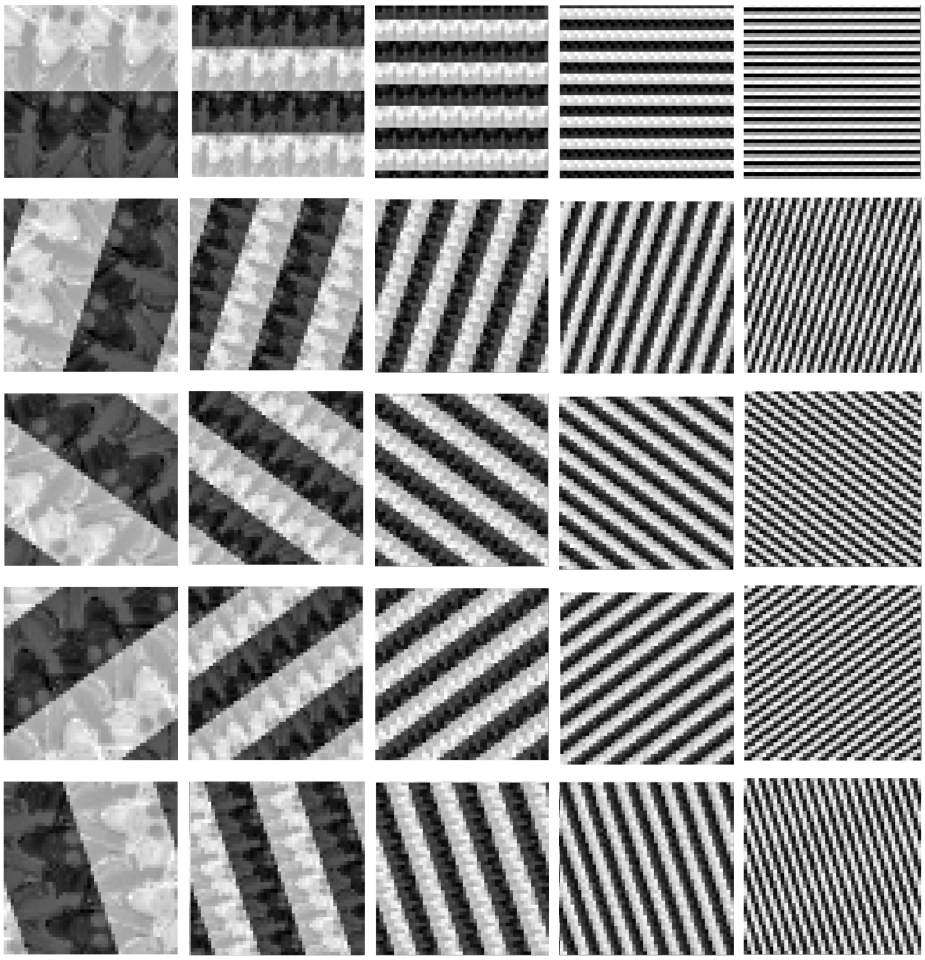


Figure B.15: Lucie Lebou||eu⊗.

B.3.7 Pearce basis

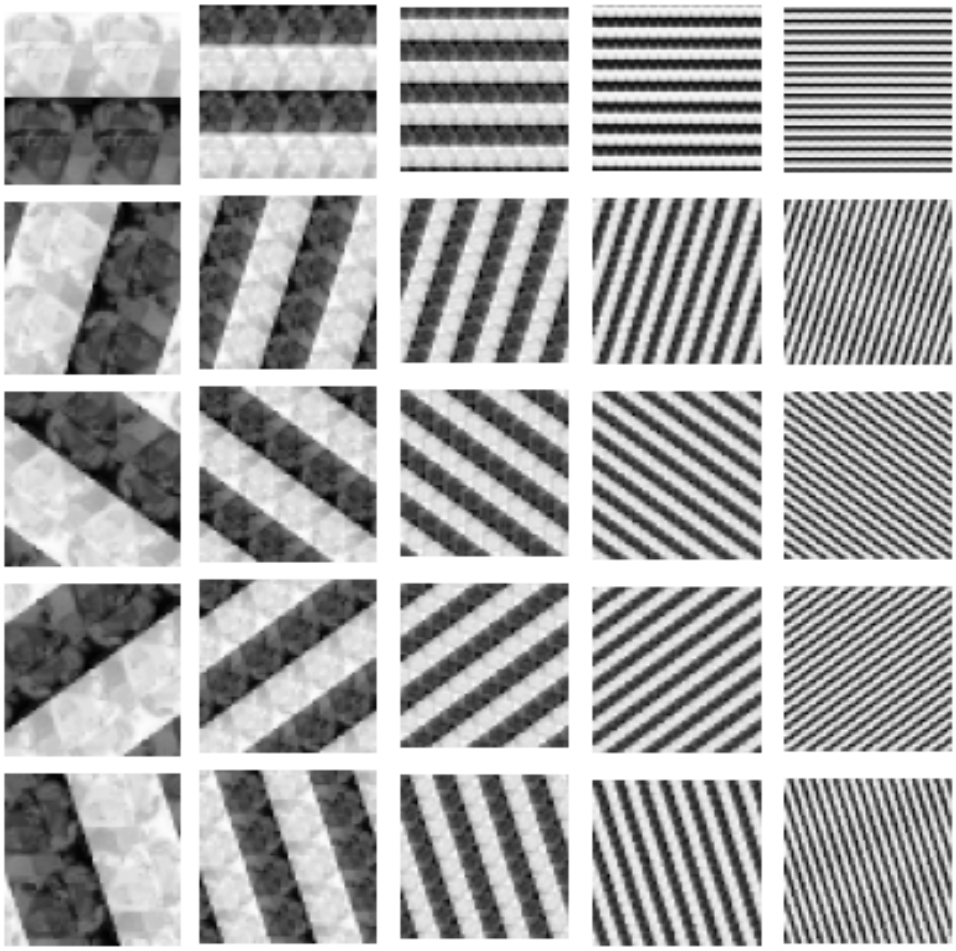


Figure B.16: Pearce basis.

B.3.8 Stelter basis

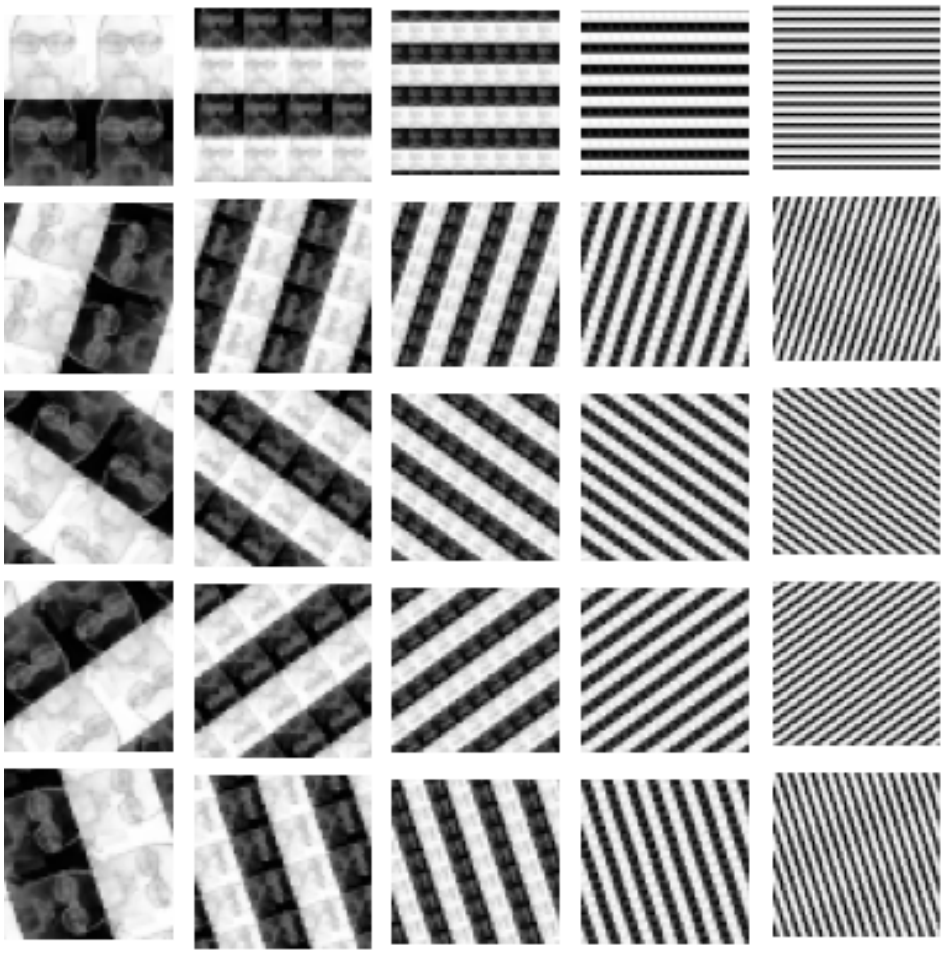


Figure B.17: Stelter Basis.

B.3.9 Jensen-Clem basis

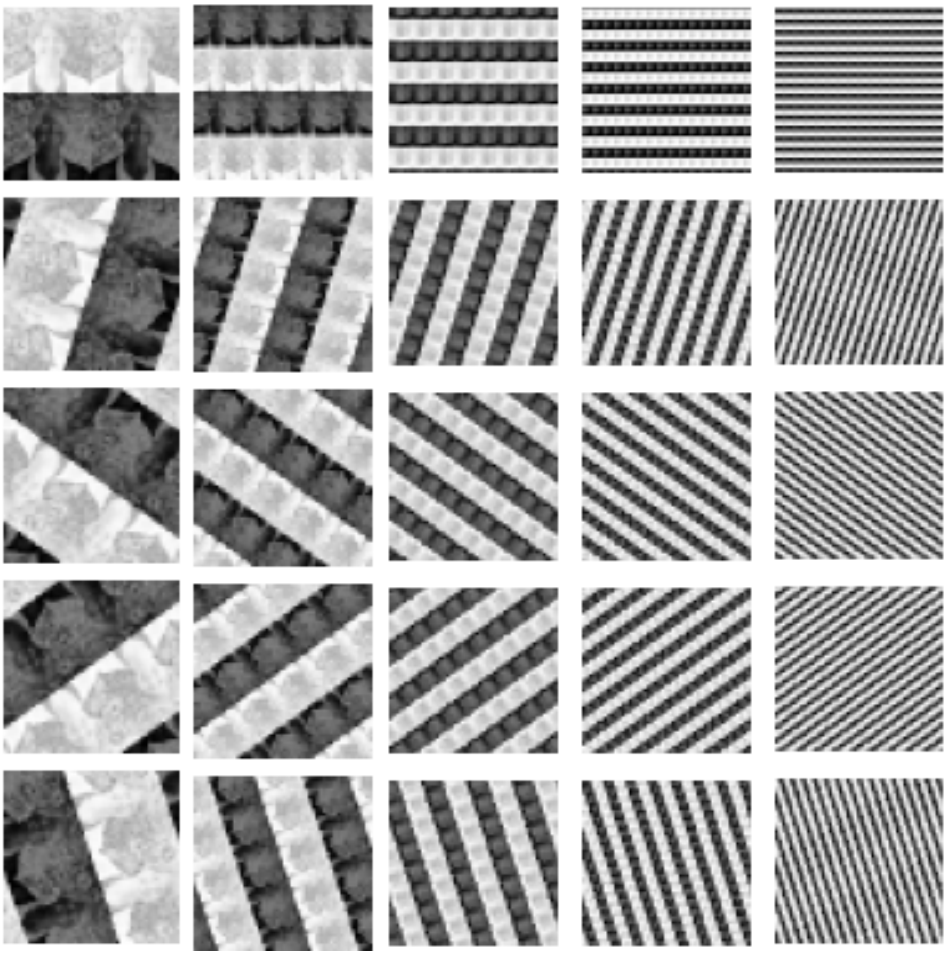


Figure B.18: Rebecca Jensen-KLem.

B.3.10 Millar-Blanchaer basis

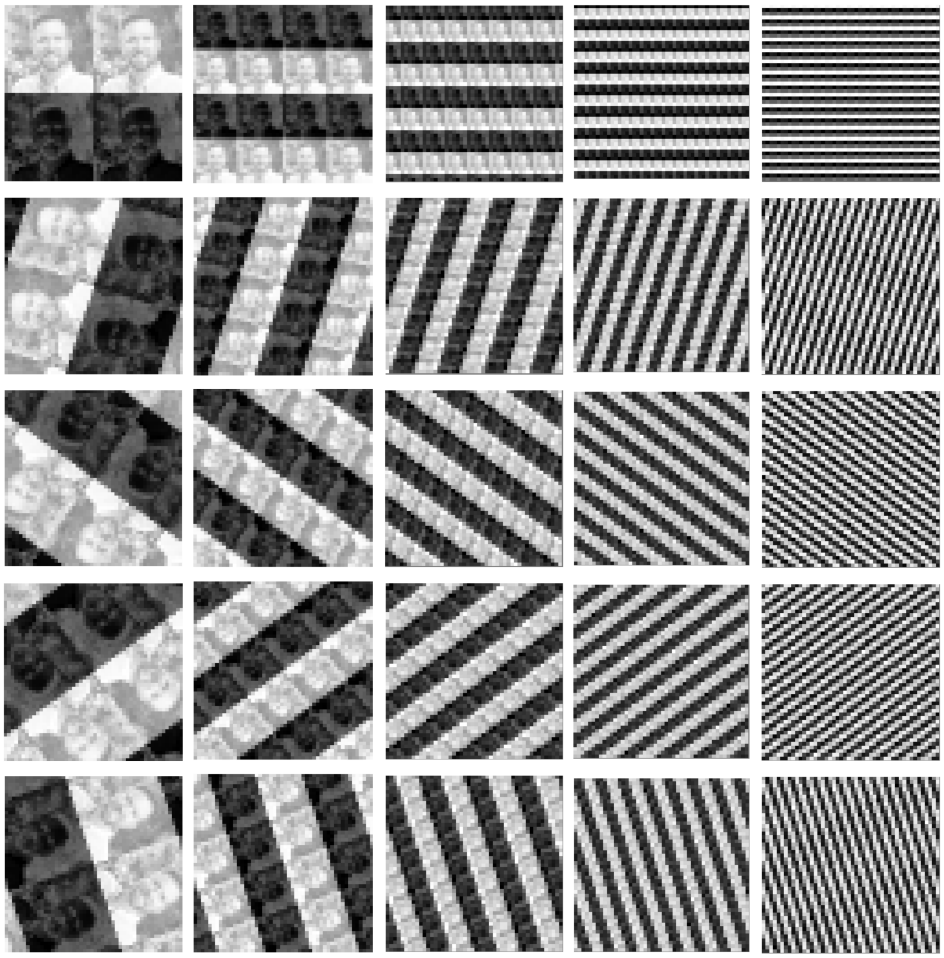


Figure B.19: Maxwell Modal-Blanchaer.

B.3.11 Steiger basis

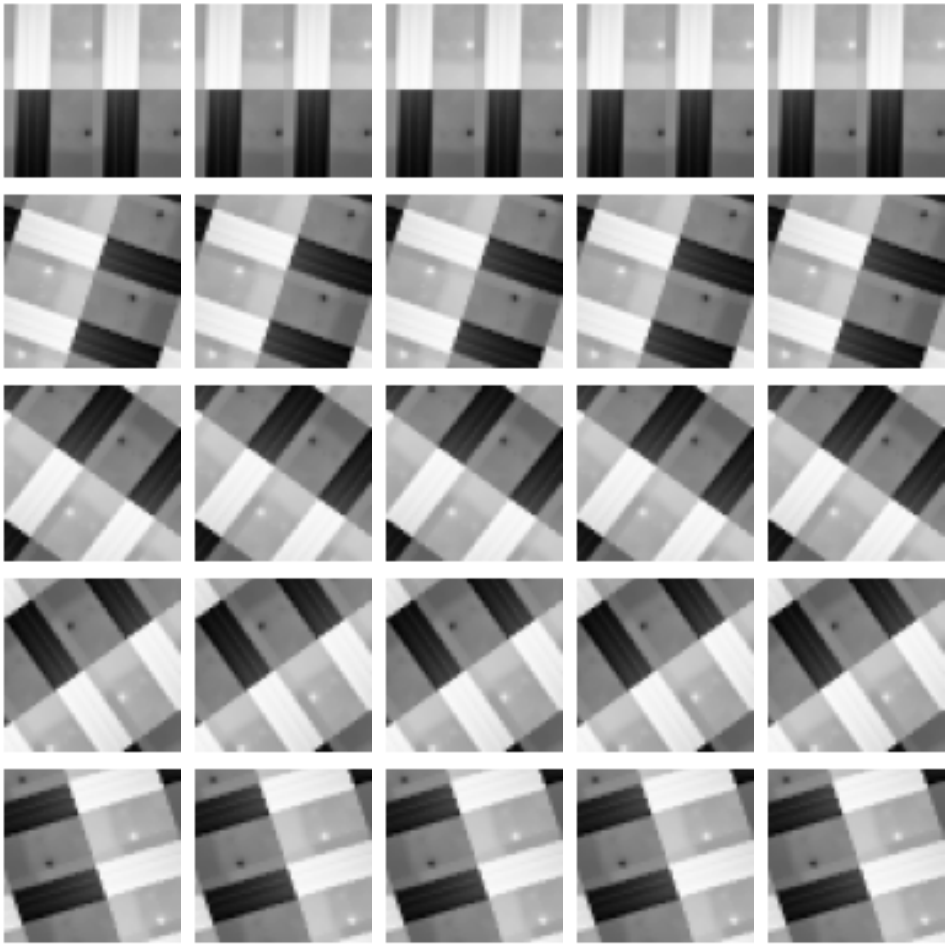


Figure B.20: Steiger Basis.

B.3.12 Lewis basis

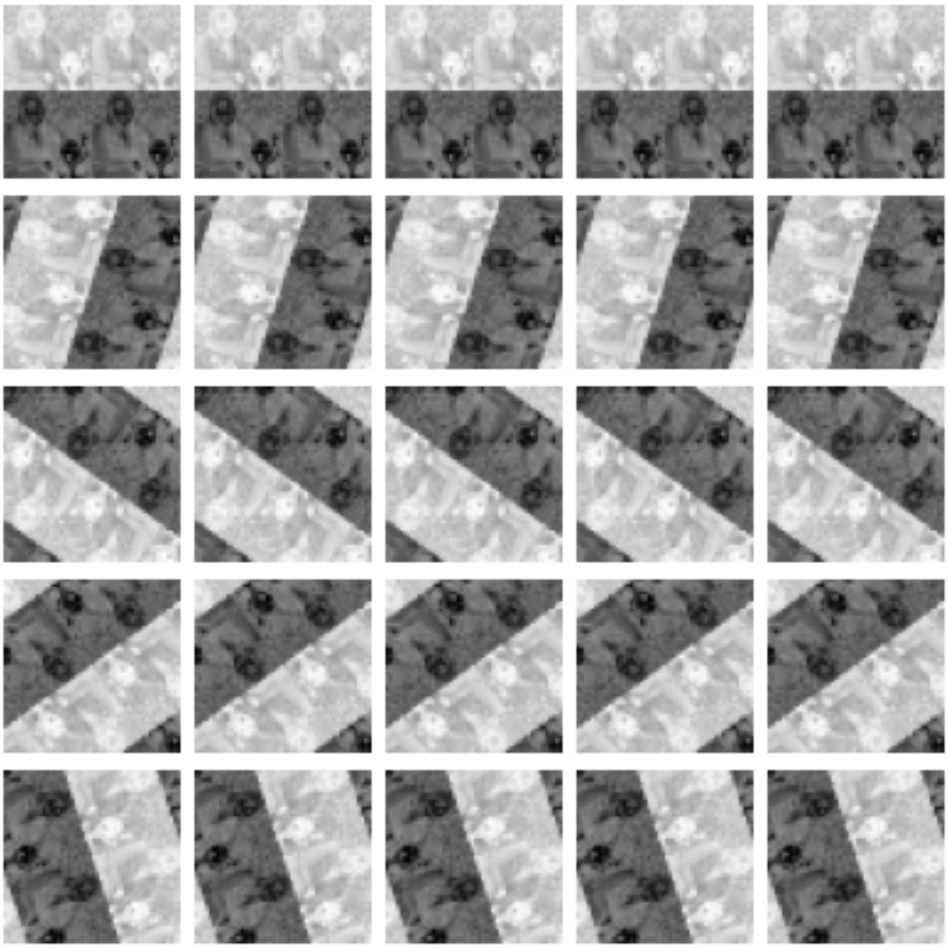


Figure B.21: Lewis Basis.

B.3.13 Balmer basis

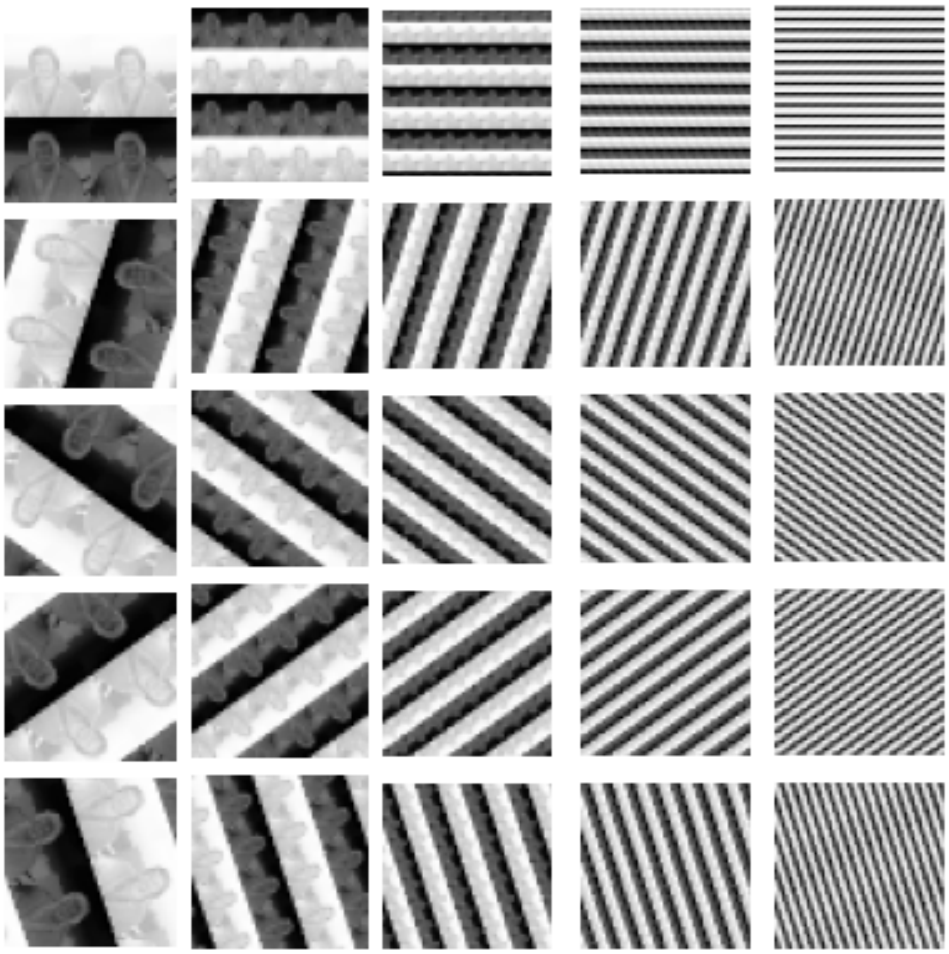


Figure B.22: Balmer Basis.