

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

High-Contrast Imaging with the MKID Exoplanet Camera

Permalink

<https://escholarship.org/uc/item/3c1866gs>

Author

Steiger, Sarah

Publication Date

2023

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

High-Contrast Imaging with the MKID Exoplanet Camera

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

Doctor of Philosophy
in
Physics

by

Sarah Steiger

Committee in charge:

Professor Benjamin A. Mazin, Chair
Professor Timothy D. Brandt
Professor Omer Blaes

June 2023

The Dissertation of Sarah Steiger is approved.

Professor Timothy D. Brandt

Professor Omer Blaes

Professor Benjamin A. Mazin, Committee Chair

June 2023

High-Contrast Imaging with the MKID Exoplanet Camera

Copyright © 2023

by

Sarah Steiger

This work is dedicated to my Mom and Dad, for their unending support.

Acknowledgements

The work presented in this thesis would not have been possible without the support, both professional and personal, of countless people.

First I would like to thank my advisor, Dr. Ben Mazin for both your seemingly endless knowledge of MKID systems and for your support and trust throughout these past six years. I would also like to thank Dr. Tim Brandt for your guidance and insight in particular with the work presented in Chapters 3 and 4 of this thesis. Special recognition also needs to go to Dr. John I. Bailey II (Jeb), for sharing your expertise and for your contributions to the work in Chapter 2 (The MKID Pipeline) which was a joint-effort and could not have been completed without your vision and dedication.

I would also like to thank my labmates of the past and present, Miguel Daal, Gregoire Coiffard, Rupert Dodkins, Kristina Davis, Joe Redford, Seth Meeker, Clint Bockstiegel, Giulia Collura, Alex Walter, Paul Szypryt, Isabel Lipartito, Neelay Fruitwala, Nick Zorbrist, Jenny Smith, Hawkins Clay, Majid Mohammad, Crystal Kim, Josh Breckenridge, and Aled Cuda (Noah, you get your own section later). You are all not only all brilliant scientists but, more importantly, kind and genuine people and it has been an honor to work with and learn from you all.

A special acknowledgment also needs to go to my cohort and lab mate Noah Swimmer. I couldn't imagine going through this journey without you and I am endlessly grateful for you not only as a fantastic collaborator and scientist, but also as a life-long friend.

I would also like to acknowledge the hard work of the incredible scientists at the Subaru Telescope, and SCEXAO in particular, including Olivier Guyon, Thayne Currie, Julien Lozi, Vincent Deo, Sebastien Vievard, Nour Skaf, and Kyohoon Ahn. Thank you for sharing not only your knowledge, but also your time and support over these many years.

I would also like to acknowledge the Subaru Telescope Operators and Day Crew for answering my endless calls and power cycling countless electrical components throughout the course of this Ph.D.. Your hard work and dedication ensures that all of the complex facilities on the mountain are able to run and this work certainly could not have been completed without you.

I would also like to acknowledge my community at both Santa Barbara City College and at the Santa Barbara Museum of Natural History including Krissie Cook, Sean Kelly, Sean Fox, Raphael Cottom, and Erin O'Connor. Your passions for education are infectious and I thank you for welcoming me with open arms, sharing your knowledge and experience, and providing necessary reminders that being able to study astronomy for a living is very, very cool – especially when my code hasn't been working for weeks and I desperately need that reminder.

I also want to take the time to recognize the indigenous Hawaiian community on whose land the Subaru Telescope is located. Maunakea holds a sacred place in the hearts of these communities and I have been incredibly privileged to use the facilities on this mountain over the course of my research.

On a personal note, I would also like to thank all of my friends and family without whose emotional support this thesis would not have been possible. This includes both my Santa Barbara community – Seamus, Demi, Joe, Aleksei, Mark, Remi, Serena, Jon, and Paige – as well as my friends across the country – Morgan, Ilektra, Sam, Erica, and Emma.

To my roommates Josh and Iman, thank you for keeping me sane and helping our house feel like a home with your warmth and friendship, I will miss you both immensely.

I would also like to thank all of my family and in particular Mom, Dad, Grandma, Poppy, Kyle, Nicole, Jared, Maddie, and Jaxon - thank you for your support, your love, and for always reminding me about what is truly important.

Last (but certainly not least) I would like to thank my fiancé Will Schultz. Thank you for your unwavering faith in me and while this thesis represents the end of a journey, you make me nothing but excited to start the next one – I love you.

Curriculum Vitæ

Sarah Steiger

Education

- 2023 Ph.D. in Physics (Expected), University of California, Santa Barbara.
- 2020 M.S. in Physics, University of California, Santa Barbara.
- 2017 B.S. in Physics, Boston College

Honors and Awards

- 2023 STScI Postdoctoral Fellowship
- 2023 51 Pegasi b Postdoctoral Fellowship (declined)
- 2020 Mananya Tantiwiwat Fellowship
- 2019 Worster Fellowship
- 2019 UCSB Physics Department Service Award
- 2018 UCSB Physics Chair's Outstanding Service Award
- 2018 UCSB Physics Outstanding Teaching Assistant Award
- 2017 Boston College George J. Goldsmith Award

Publications

Steiger, S., Brandt, T., Guyon, O., et. al. (2022). Probing Photon Statistics in Adaptive Optics Images with SCEXAO/MEC. *AJ*, 164(5), 186. doi: 10.3847/1538-3881/ac922f

Swimmer, N., Currie, T., **Steiger, S.**, et. al. (2022). SCEXAO and Keck Direct Imaging Discovery of a Low-Mass Companion Around the Accelerating F5 Star HIP 5319. *AJ*, 164(4), 152. doi: 10.3847/1538-3881/ac85a8

Steiger, S., Bailey, J., Zobrist, N., et. al. (2022). The MKID Pipeline: A Data Reduction and Analysis Pipeline for UVOIR MKID Data. *AJ*, 163(5), 193. doi: 10.3847/1538-3881/ac5833

Steiger, S., Currie, T., Brandt, T., et. al. (2021). SCEXAO/MEC and CHARIS Discovery of a Low-mass, 6 au Separation Companion to HIP 109427 Using Stochastic Speckle Discrimination and High-contrast Spectroscopy. *AJ*, 162(2), 44. doi: 10.3847/1538-3881/ac02cc

Zobrist, N., Klimovich, N., Ho Eom, B., Coiffard, G., Daal, M., Swimmer, N., **Steiger, S.**, et. al. (2021). Improving the dynamic range of single photon counting kinetic inductance detectors. *Journal of Astronomical Telescopes, Instruments, and Systems*, 7, 010501. doi: 10.1117/1.JATIS.7.1.010501

Smith, J., Mazin, B., Walter, A., Daal, M., Bailey, I., Bockstiegel, C., Zobrist, N., Swimmer, N., **Steiger, S.**, and Fruitwala, N. (2021). Flexible Coaxial Ribbon Cable for High-Density Superconducting Microwave Device Arrays. *IEEE Transactions on Applied Superconductivity*, 31(1), 3008591. doi: 10.1109/TASC.2020.3008591

Lozi, J., Guyon, O., Vievard, S., [and 49 others, including **Steiger, S.**] (2020). Status of the SCEXAO instrument: recent technology upgrades and path to a system-level demonstrator for PSI. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (pp. 114480N). doi: 10.1117/12.2562832

Swimmer, N., Mazin, B., Bockstiegel, C., [and 14 others, including **Steiger, S.**] (2020). The PICTURE-C MKID camera. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (pp. 114479B). doi: 10.1117/12.2561770

Walter, A., Fruitwala, N., **Steiger, S.**, et. al. (2020). The MKID Exoplanet Camera for Subaru SCEXAO. *PASP*, 132(1018), 125005. doi: 10.1088/1538-3873/abc60f

Coiffard, G., Daal, M., Zobrist, N., Swimmer, N., **Steiger, S.**, Bumble, B., and Mazin, B. (2020). Characterization of sputtered hafnium thin films for high quality factor microwave kinetic inductance detectors. *Superconductor Science Technology*, 33(7), 07LT02. doi: 10.1088/1361-6668/ab8d99

Zobrist, N., Coiffard, G., Bumble, B., Swimmer, N., **Steiger, S.**, Daal, M., Collura, G., Walter, A., Bockstiegel, C., Fruitwala, N., Lipartito, I., and Mazin, B. (2019). Design and performance of hafnium optical and near-IR kinetic inductance detectors. *Applied Physics Letters*, 115(21), 213503. doi: 10.1063/1.5127768

Mazin, B., Bailey, J., Bartlett, J., [and 22 others, including **Steiger, S.**] (2019). MKIDs in the 2020s. In *Bulletin of the American Astronomical Society* (pp. 17).

Abstract

High-Contrast Imaging with the MKID Exoplanet Camera

by

Sarah Steiger

The Microwave Kinetic Inductance Detector (MKID) Exoplanet Camera (MEC) is a $Y-J$ band integral field unit located behind the Subaru Coronagraphic Extreme Adaptive Optics system (SCEXAO) at the Subaru Telescope on Maunakea. The detector inside of MEC is a 20 kilo-pixel photon-counting MKID array that yields not only the energy of each photon ($\mathcal{R} \sim 5$), but also its arrival time (to within a microsecond) with no read noise or dark current. This temporal resolution allows us to perform post-processing techniques that leverage differences in the photon arrival time statistics between stars and their faint companions such as Stochastic Speckle Discrimination (SSD). With SSD, we have demonstrated the ability to uncover low mass stellar companions with the discovery of HIP 109427 B and also obtained the first SSD detection of a diffuse source using MEC observations of the known disk AB Aurigae. Here we can resolve structures in the disk within $0.3''$ without the use of any PSF subtraction or polarization. These analyses are made possible through the use of the MKID Pipeline, a new open-source data reduction and analysis pipeline developed for MKID instruments that takes raw MKID data as its input and can return not only unique MKID data products for specialized analysis, but also images that can easily interface with existing post-processing techniques (e.g. ADI) for more general science.

This thesis first presents an introduction to high-contrast imaging, the MEC instrument, and photon statistics in millisecond images behind an Adaptive Optics system. We then introduce the MKID Pipeline which is the open source MKID Data Reduction and

Analysis Software developed for MEC with extensibility to future MKID instruments. Using the MKID Pipeline, we then present the first science results with SCE_xAO/MEC with the discovery of HIP 109427 B using SSD with MEC as well as high resolution spectroscopy with the CHARIS instrument. We end with a broader discussion of how MKID instruments, like MEC, are uniquely positioned to explore photon statistics at very fast frame rates and provide a path forward for future researchers to understand and utilize the unique capabilities of MKID cameras.

Contents

Curriculum Vitae	viii
Abstract	x
1 Introduction	1
1.1 High-Contrast Imaging	1
1.1.1 Exoplanet, Brown Dwarf, and Disk Science	3
1.1.2 High-Contrast Imaging from the Ground: Adaptive Optics and Coronagraphy	10
1.1.3 Speckle Suppression Techniques	14
1.2 The MKID Exoplanet Camera (MEC)	17
1.2.1 Microwave Kinetic Inductance Detectors (MKIDs)	17
1.2.2 MEC at Subaru/SCEXAO	20
1.3 Photon Statistics in Adaptive Optics Images	20
1.3.1 Speckles	21
1.3.2 Companions	22
1.3.3 Satellite Spots	26
2 The MKID Pipeline	28
2.1 MKID Observing Datasets	29
2.2 Data Processing	30
2.2.1 Data Format	32
2.2.2 Data Calibration	33
2.2.3 Data Products	42
2.3 <code>mkidpipeline</code> : The MKID Pipeline Package	44
2.3.1 Concept	45
2.3.2 Initialization and Configuration	45
2.3.3 Data Specification	47
2.3.4 Output Specification	49
2.3.5 Execution	50
2.3.6 Core Modules and Libraries	51

2.3.7	Interactive Use	51
2.3.8	Extending Functionality to Future MKID Instruments	51
2.4	Summary	52
3	First Science Results with MEC	54
3.1	System Properties and Observations	55
3.2	Data	57
3.2.1	Image Processing: MEC	57
3.2.2	Image Processing: CHARIS and NIRC2	61
3.3	Analysis	63
3.3.1	HIP 109427 B Spectroscopy and Photometry	63
3.3.2	HIP 109427 B Spectral Type, Temperature, and Luminosity	63
3.3.3	HIP 109427 B Astrometry and Dynamical Mass	68
3.4	Conclusion	71
4	Probing Photon Statistics in Adaptive Optics Images with SCExAO/MEC	74
4.1	Utilizing Photon Arrival Time Information with MEC	75
4.2	Photon Arrival Time Based Post-Processing Techniques	76
4.2.1	Stochastic Speckle Discrimination (SSD)	76
4.2.2	Photon-Counting SSD (PCSSD)	78
4.3	Discussion	86
A	Transition Region Between Gamma and MR Statistics	88
	Bibliography	90

Chapter 1

Introduction

1.1 High-Contrast Imaging

At the broadest level, high-contrast imaging is the imaging of a very dim object (a “companion”) around a much brighter object (a “host”). The greater the difference in brightness between the companion and host, the higher the contrast, and the more challenging it is to detect the companion. Hosts can span a wide range of masses and properties and include stars, quasars, pulsars, and even black holes. The companions that surround these hosts can be equally as varied including planets, brown dwarfs (intermediate objects with masses between those of planets and stars), low mass stars, and expansive disks of gas and dust.

In this work, the term “hosts” will exclusively refer to stars and a scientific focus will be placed on exoplanet, brown-dwarf, and disk companions. In order to spatially resolve the companion from the host star, these systems are generally close-by – within a few hundred parsecs – and typical companion separations range from 10–150 au, angular separations (ρ) of $\sim 0.4 - 2$ arcseconds (”) on the sky (e.g. Marois et al., 2008a; Lagrange et al., 2009; Rameau et al., 2013; Kuzuhara et al., 2013; Currie et al., 2014; Macintosh

et al., 2015; Chauvin et al., 2017). Contrasts in these cases can easily range from 10^{-5} to 10^{-10} making this an incredibly technologically challenging method of companion detection. Despite these technological challenges, high-contrast imaging is important as it can explore a unique region of parameter space compared to other detection methods as will be described later on in this section.

Besides high-contrast imaging, the two most common methods for companion detection are the transit and radial velocity (RV) methods. These are each “indirect” methods of detection in that they look at changes in the properties of the host star to infer the presence of a companion. High-contrast imaging is often referred to as “direct” imaging to distinguish it from these indirect methods and these terms will be used interchangeably throughout this work.

In the transit method, companions are found by identifying periodic dips in the brightness of the host star that are caused by companion(s) passing in front of it along an observers line-of-sight (Charbonneau et al., 2000; Henry et al., 2000). The RV method detects companions by looking at the Doppler shift of host star spectral lines due to the gravitational pull of a companion causing it to orbit about the system’s common center of mass (Mayor and Queloz, 1995). Since the companion needs to pass in front of the host star for a transit detection, and there needs to be a line-of-sight component to the velocity of the host star for an RV signal to be detectable, these methods are only sensitive to companions in at least partially edge-on orbits. They are also more sensitive to companions with close-in orbits since the probability of a transit $\sim R_*/P^{\frac{2}{3}}$ where R_* is the radius of the star and P is the orbital period. The RV signal also goes like

$$K \sim \left(\frac{1}{P}\right)^{\frac{1}{3}} \frac{M_P \sin i}{M_T^{\frac{2}{3}}} \frac{1}{(1-e^2)^{\frac{1}{2}}} \quad (1.1)$$

where K is the RV semi-amplitude, P is the orbital period, M_P is the mass of the

companion, i is the inclination of the orbit, M_T is the total mass of the system and e is the orbital eccentricity. This means that for a given host star mass, shorter period massive companions have the highest likelihood of detection.

In high-contrast imaging, we are sensitive to the opposite region of parameter space. Companions in wide separation, face-on orbits are less likely to be swamped by the signal from the host star and are therefore more likely to be detected. Since we are also getting light directly from the companion, we can simultaneously perform spectroscopy to determine the companion's composition as well as directly measure orbital parameters.

Importantly, high-contrast imaging is also currently the only method we have that allows us to detect diffuse sources such as disks. As will be explored in §1.1.1 these objects are incredibly important for understanding the formation and evolution of star and planetary systems and are still very poorly understood.

Figure 1.1 shows all of the currently discovered planetary mass companions as a function of their mass and orbital separation with the color and shape of the points indicating their discovery method. Here it is clear to see the detection biases inherent to each technique with the directly imaged planets at large masses and wide separations and transiting planets covering a wide range of masses but with orbital periods of less than a few hundred days. Currently there are very few Solar System like exoplanets that have been discovered in large part due to these detection biases and current limits.

1.1.1 Exoplanet, Brown Dwarf, and Disk Science

The questions that astronomers who study exoplanetary and related systems are seeking to answer can be summarized by the following “big” questions:

- How unique is our Solar System?
- How did our Solar System come to be?

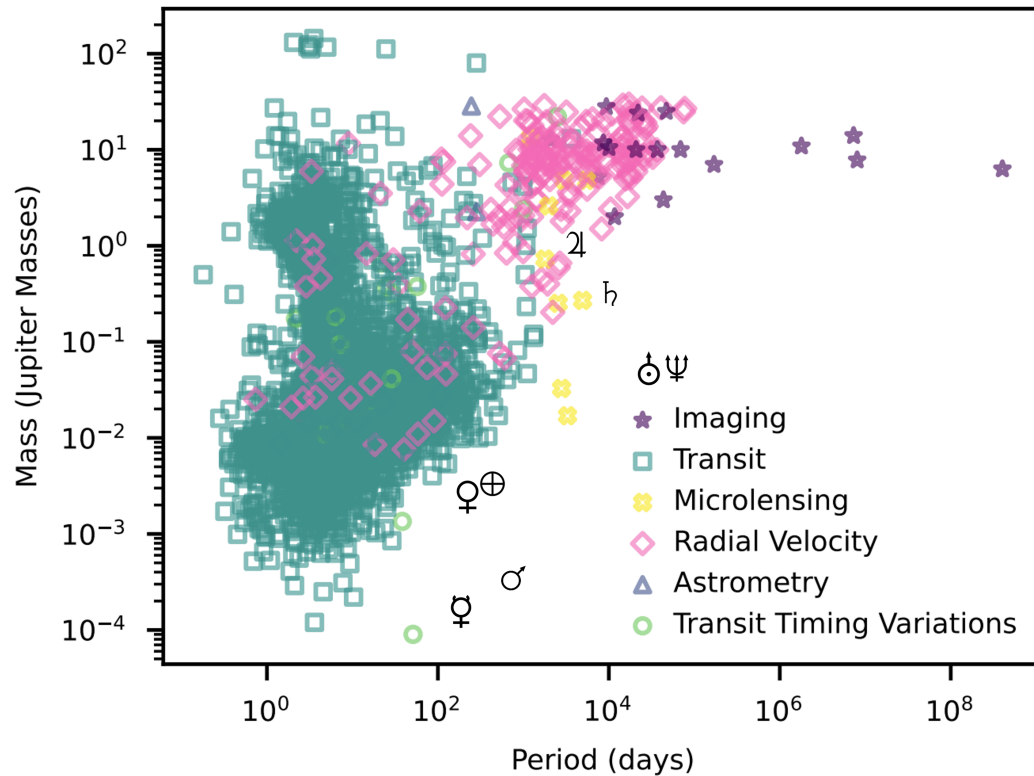


Figure 1.1 All currently known exoplanets plot as a function of their mass (in Jupiter masses) and orbital periods (in days) where the different colors and shapes denote the method of detection. The symbols for each of the Solar System planets demonstrate where our Solar System fits into these broader demographics.

- Is there life elsewhere in the universe?

These will each be explored briefly in the following sections and the citations therein. This review is not meant to be comprehensive, but to give an overview of the largest unanswered questions in the field, why they are important, and what tools and techniques are used to attempt to answer them.

How Unique is Our Solar System?

The first exoplanet was discovered in 1992 and was found to have a mass a few times larger than that of Jupiter and is orbiting a highly magnetized rapidly-rotating neutron star called a pulsar (Wolszczan and Frail, 1992). Before this time, it was not known a) whether any planets besides our Solar System planets existed at all or b) whether our Solar System planets are representative of the entire range of possible planetary systems. This first discovery was already indicative that the template we were familiar with by observing the types of planets that exist in our own Solar System does not come close to describing the full range of systems that can exist in the universe. The discovery of a much more familiar configuration of a planet and host came a few years later with discovery of the first planet orbiting a main sequence star, 51 Pegasi b (Mayor and Queloz, 1995). This massive planet (roughly half the mass of Jupiter) however was found to orbit incredibly close-in to its host, with an orbital period of just ~ 4 days. This showed that even for main-sequence stars, our own Solar System with short period rocky planets and wide separation massive gas giants is not necessarily representative of all systems.

After these initial discoveries, more systems were needed to be able to make broader statements about exoplanet demographics, find out the full range of possible planetary configurations and, in the process, learn how unique our own Solar System is.

In the intervening ~ 30 years since the discovery of the first exoplanet, a few key

missions have been critical to expand the handful of known systems in the 1990's to now over 4000 confirmed systems today. The Kepler Space Telescope (Koch et al., 2010; Howell et al., 2014a) was launched in 2009 with the goal of finding the occurrence rate of Earth-like planets (rocky planets in the habitable zone of a variety of main-sequence stars) via the transit method. In 2010, it made its first discovery with Kepler-9 – the first confirmed planetary system with more than one planet transiting the same star (Holman et al., 2010). Kepler was ultimately unable to probe Earth-analogs due to the fact that these planets are expected to reside at wide separations ($\sim > 100$ days) and Kepler had a truncated mission lifetime due to the failure of two of the telescopes reaction wheels (which are responsible for pointing). Since the transit probability decreases with period and many transits need to be observed for a reliable detection, very few of these Earth-like planets were confirmed (Thompson et al., 2018). Throughout its limited mission lifetime, however, Kepler did discover over 2,600 planets and showed that not only do planets exist elsewhere in the galaxy, but that they are common. In fact, most stars contain at least one planet and the number of planets likely outnumber the amount of stars in our own galaxy (Cassan et al., 2012). It also further solidified that the variety of planets in the universe extends far beyond what is present in our own Solar System. The most common size of planet found by Kepler is between the size of Earth and Neptune for which our Solar System has no analog.

Kepler reached the end of its mission lifetime in 2018 after the K2 extension (Howell et al., 2014b), but that same year the Transiting Exoplanet Survey Satellite (TESS; Ricker et al., 2015) was launched. Like Kepler, TESS was also designed to look for transiting exoplanets, but targeted stars 30-100 times brighter than Kepler (allowing for easier follow-up) and surveyed the entire night sky over the course of a two-year period with a search area 400x the size of the Kepler survey area. In March 2021, NASA announced that TESS found 2200 exoplanet candidates and by the end of 2021 TESS

had discovered over 5000 candidates.

Despite probing only the shortest period planets, these missions are still invaluable for helping understand what broader exoplanet demographics look like. They are especially powerful when combined with RV data as transits can break the inclination degeneracy on the planetary mass from RVs while also yielding planet radii. Some key findings include the presence of a bi-modality in the size distribution of small planets with a statistically significant lack of planets with radii of around $\sim 1.7R_{\oplus}$ (Fulton and Petigura, 2018), a lack of short-period planets with Neptune masses (called the “hot Neptune desert”; Mazeh et al., 2016), and that planets within a single multi-planet system have correlated sizes (Weiss et al., 2018). With this plethora of discovered short period transiting exoplanets, increasing the number of long period companions which can be accessed with direct imaging surveys will help us fill in the parameter space seen in Figure 1.1 and draw further conclusions about what types of planetary systems are possible.

How Did Our Solar System Come to Be?

Since most of the currently discovered exoplanets have masses, orbits, and compositions which are very different than any of the planets in our own Solar System, these observations all inform our understanding of planet formation theory and raise many questions with regards to how these different systems formed and, importantly, if our own Solar System formed in the same way 4.5 billion years ago.

The two most prevalent theories of planet formation are core-accretion (the “bottom-up” approach; Pollack, 1984) and gravitational instability (the “top-down” approach; Boss, 1997). In the core-accretion theory of planet formation, small planetesimals in a protoplanetary disk collide until they generate a rocky core large enough for a gaseous atmosphere to accrete onto it. In the outer regions of a disk, beyond the snow line, temperatures are low enough that ices and rocky material can condense forming large

solid cores. Once these rocky cores reach $\sim 10M_{\oplus}$ there is still enough material left in the outer regions of the disk that a process of runaway accretion starts to occur, eventually forming planets with masses and orbits similar to the gas giants in our own Solar System. Close-in to the star, there is not enough material in the disks to form a planetesimal large enough to undergo runaway accretion which was thought to explain the presence of the smaller terrestrial planets in our own Solar System closer-in to the host star.

When 51 Pegasi b was discovered in 1995, the presence of a huge world so close to its star was not compatible with this formation theory and was initially considered an anomaly. However, since then, numerous other “hot Jupiters” (massive planets located at very tight orbital separations) have been discovered prompting the need for alternate theories. In the gravitational instability theory of planet formation, giant self-gravitating clumps of disk material are formed when the protoplanetary disk fragments. This was thought to be able to form giant planets at these close separations, but it was later shown that there needs to be very tight constraints on the properties of the gas in the disk at these locations for this to be possible (Rafikov, 2005). Since hot Jupiters are so common, this mechanism is unlikely to explain the presence of these objects further prompting the need for alternate formation mechanisms.

Planet migration is one such explanation for how hot Jupiters like 51 Pegasi b came to be. In this case, the planet is not formed *in situ*, but instead forms farther away in the disk (where both gravitational instability and core-accretion are more likely) and migrates inward to its currently observed location either via disk interactions (Goldreich and Tremaine, 1980; Lin et al., 1996) or through a process called high-eccentricity migration (Rasio and Ford, 1996) at much later times after the disk has already dissipated.

In the context of hot Jupiters it is also interesting to look at another class of object called brown dwarfs which are objects with masses that are too small to burn hydrogen in their cores, but large enough to burn deuterium (Kumar, 1963; Nakajima et al., 1995;

Oppenheimer et al., 1995). These objects span the transitional mass region between stars and planets and so whether they form via the same mechanisms as planets or stars is unknown making these systems an interesting testing ground for formation theories.

How planets form and arrive at their stable late-time configurations are still largely unanswered questions in the field, but a very important tool provided by high-contrast imaging is the ability to look at young planetary systems which are still in the process of formation to see (in “real time”) how different stages of the planetary formation process unfold. These protoplanetary disks of dense gas and dust around young stars are objects that can only be observed through direct imaging and in recent years a handful of protoplanetary candidates have been discovered including PDS 70 bc (Keppler et al., 2018; Haffert et al., 2019) and AB Aurigae b (Currie et al., 2022; Zhou et al., 2022) providing valuable insights into these systems.

In addition to attempting to image protoplanets directly, there is evidence that as young planets form, they can carve out openings in the surrounding disk material forming “gaps” (Lin and Papaloizou, 1986; Bate et al., 2003) which are much easier to detect than young protoplanets themselves. Directly imaging planets in these gaps places constraints on these systems and provides evidence that other gaps such as those viewed by the Atacama Large Millimeter/submillimeter Array (ALMA) in sub mm wavelengths may also harbor planetary companions.

Is the Life Elsewhere in the Universe?

Despite the plethora of planets currently discovered, Earth is still the only one that we know to host life. Significant work, however, has been put into understanding where else in the universe can be hospitable for life as we know it and how to identify life if we find such places.

The first indicator for life is the idea of a “Habitable Zone” (HZ) which is traditionally

defined as the circumstellar region in which a terrestrial-mass planet with a $CO_2 - H_2O - N_2$ atmosphere can sustain liquid water on its surface (Huang, 1959; Kasting et al., 1993). Depending on the size of the host star, this region can vary from 0.1 to 10s of au and is also highly dependent on things like stellar age, planetary albedo, climate models, etc. (Kopparapu et al., 2013).

Due to the close separations and advanced ages of planets that can meet such criteria, these types of systems are, as of now, beyond the reach of direct imaging which is currently limited to probing young giant planets still hot with the heat of their own formation. Future space missions such as the Habitable Worlds Observatory – a combination of the HabEx (Gaudi et al., 2020) and LUVOIR (The LUVOIR Team, 2019) mission concepts – will, for the first time, have the sensitivity to directly image the first Earth-like planet around a Sun-like star. In order to achieve this feat, however, there are still significant technological advances that need to be achieved in both the detectors used to take the science images (which will be the subject of this thesis), and in the upstream optics used to block as much of the light from the host star as possible which are discussed in §1.1.2.

1.1.2 High-Contrast Imaging from the Ground: Adaptive Optics and Coronagraphy

High-contrast imaging plays a unique role in answering these big questions by detecting wide separation planets in face-on orbits, imaging and characterizing disks, directly determining companions orbital parameters, and simultaneously performing spectroscopy on these objects to learn about their chemical compositions and abundances. The main challenge of high-contrast imaging, however, lies in the fact that companions are both very close to, and millions to billions of times fainter than, the stars they orbit.

Despite recent developments in space based instrumentation to perform direct imaging

with JWST (Rigby et al., 2023) and the upcoming Nancy Grace Roman Space Telescope (Spergel et al., 2015), currently much effort is focused on ground-based imaging. Here large 8-10 m class telescopes capable of achieving the contrasts and IWAs needed to see planetary companions are already deployed and research and development cycles can happen on much more rapid timescales.

A major limitation from the ground, however, is Earth’s turbulent atmosphere which imparts aberrations on the wavefronts of light hitting a telescope. In order to correct for these aberrations, adaptive optics (AO) systems are used which apply mechanical corrections to the wavefront in an attempt to recover the image as it would appear to a space telescope located above the atmosphere, see Figure 1.2.

Briefly, the aberrated wavefront is first sent to a deformable mirror (DM) which is typically a flexible sheet of glass controlled by anywhere from hundreds to thousands of actuators. These actuators can be controlled individually and alter the shape of the mirror surface. The light from this mirror is then sent to a beamsplitter which simultaneously directs the light to a science camera and a wavefront sensor. The wavefront sensor is connected to a control system which calculates how the actuators on the DM need to be moved to recover an (ideally) perfectly flat, un-aberrated wavefront. This wavefront sensing and control (WFS/C) loop runs many of thousands of times per second (kilohertz speeds) to be able to keep up with the rate at which the atmosphere is changing and maintain a stable corrected wavefront at the science camera.

In recent years, next-generation extreme AO instruments (AO systems with > 1000 actuators), such as the Gemini Planet Imager (GPI; Macintosh et al., 2014), the Spectro-Polarimetric High-contrast Exoplanet REsearch at VLT (SPHERE; Beuzit et al., 2019), and the Subaru Coronagraphic Extreme Adaptive Optics System (SCEXAO; Jovanovic et al., 2015a) have achieved factors of 100 improvement in contrast at sub-arcsecond separations over conventional AO systems. It is these systems which are now enabling

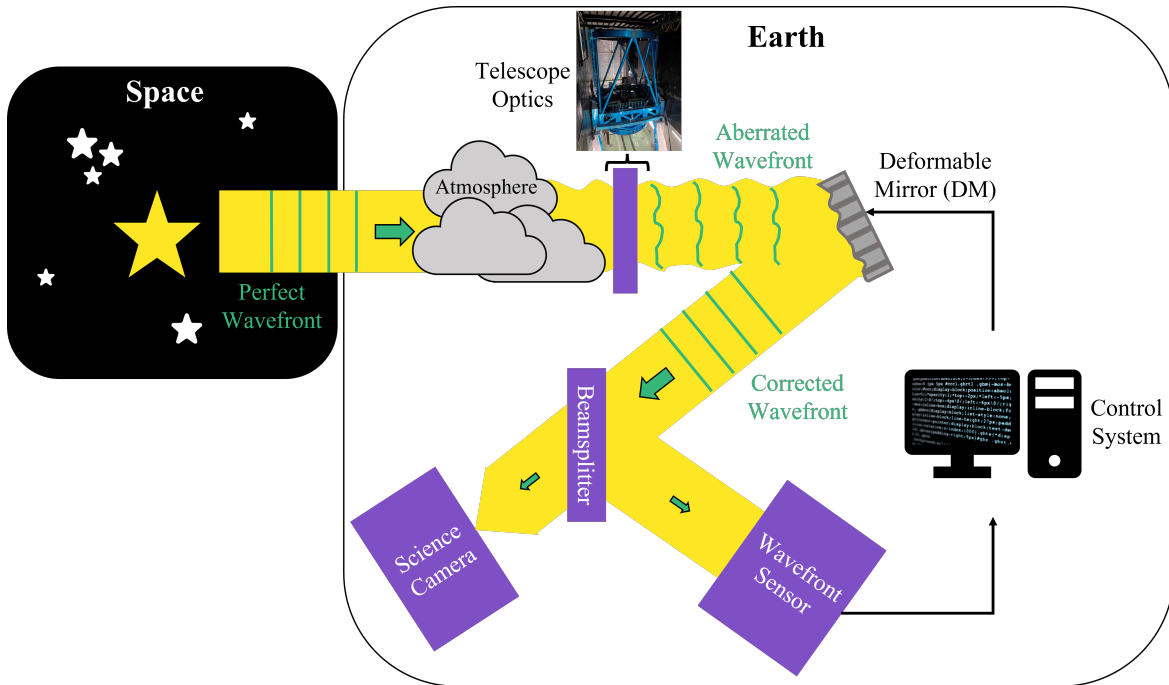


Figure 1.2 Schematic diagram of an Adaptive Optics System.

the next-generation of high-contrast imaging though current limits still restrict detectable companions to jovian-sized exoplanets beyond ~ 10 au.

The extreme contrasts and small angular separations that we need to consider for direct imaging mean that even behind an AO system if one were to simply take an image of the two-body system, the light from the bright host would completely overwhelm any signal from the faint companion of interest. In order to successfully extract this companion signal we must therefore block as much of the light from the host star as possible. This is done physically by using a set of specialized optics called a coronagraph. The simplest type of coronagraph is the Lyot coronagraph and is named for the French astronomer of the same name, Bernard Lyot, who invented it in 1931. It is comprised of two main optics, the first is an occulting spot which is placed over the bright host star in the image plane to block as much on-axis light as possible. The second is the Lyot

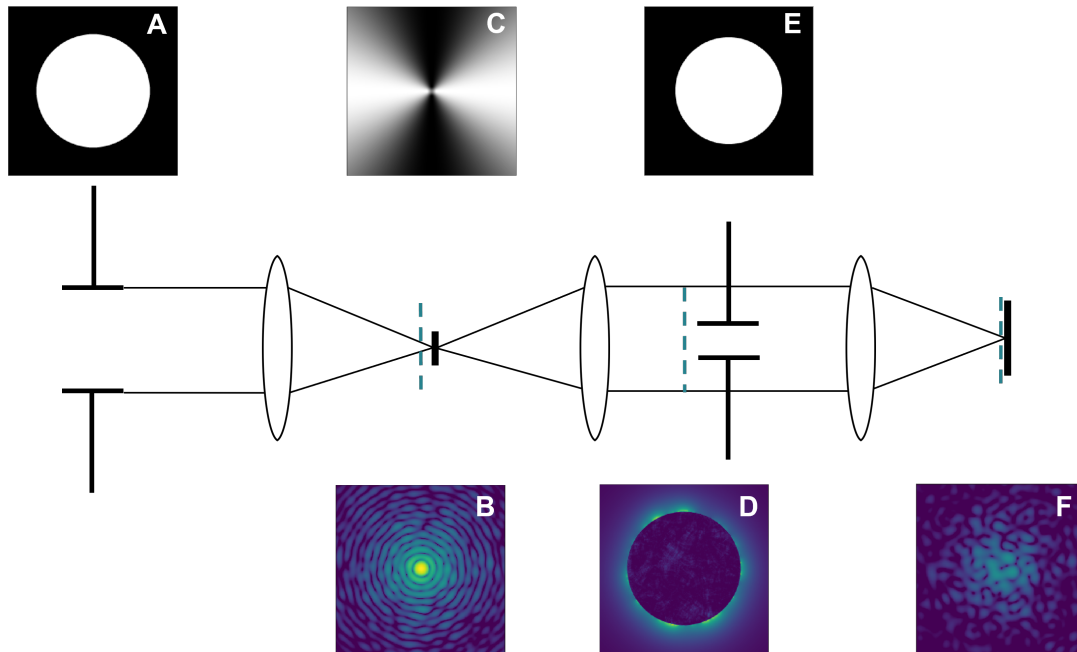


Figure 1.3 Schematic diagram of an optical vortex coronagraph with simulated data showing relevant optical components and intensities. Dashed lines show the locations at which panels B, D, and F are sampled. From left to right we have A) The telescope pupil B) the telescope PSF with an aberrated wavefront C) The optical vortex of charge 2. D) Lyot plane intensity . E) Lyot stop F) Intensity in the image plane at the detector where much of the original starlight has been suppressed. Simulated data was generated using the HCIPy package (Por et al., 2018) with a Vortex focal plane mask and Lyot stop acting on an aberrated wavefront.

stop which is placed in the pupil plane to block any diffracted light that makes its way around the occulting spot.

Modern-day coronagraphs have made many improvements on the classical Lyot coronagraph and the design and fabrication of increasingly novel and complex architectures to remove as much starlight as possible is its own expansive research landscape (and beyond the scope of this thesis), but comprehensive reviews can be found in Ruane et al. (2018) and Galicher and Mazoyer (2023). One improvement on the design of the traditional Lyot coronagraph is to replace the simple occulting mask in the focal plane with a more advanced optic. In particular, masks with no radial features, that are robust to



Figure 1.4 Picture of Lyot stops in the SCEXAO filter wheel that we fabricated in 2018. Picture credit: J. Lozi

tip-tilt errors and that can work in broadband are desired. One such design is called the vortex coronagraph (Foo et al., 2005) which uses an azimuthal phase ramp as shown in Figure 1.3. These designs importantly have been pushing for not only increased starlight suppression, but also smaller inner working angles (IWAs) where companions are more likely to be located.

1.1.3 Speckle Suppression Techniques

After the AO system and coronagraph, the limiting noise source for high-contrast imaging from the ground are point spread function (PSF) sized “speckles” of leftover diffracted starlight with a range of correlation timescales (τ) and sources – see Figure 1.3, panel F. Rapidly-evolving atmospheric speckles ($\tau \sim 1\text{-}20$ ms) result from aberrations left uncorrected by an AO system and average out over the course of long-exposure images,

forming a smooth halo (e.g. Perrin et al., 2003; Soummer et al., 2007a). These “fast” speckles can be corrected by improved AO control loops which will mitigate temporal bandwidth error and measurement (photon noise) error (e.g. Guyon, 2005).

Alternatively, quasi-static speckles result from imperfections in the instrument such as non-common path errors, telescope vibrations, etc. (Guyon, 2005; Lozi et al., 2018). These speckles interfere with atmospheric speckles and can be pinned to the diffraction rings (Soummer et al., 2007a). Quasi-static speckle noise follows a highly non-Gaussian (modified Rician distribution) and is temporally well correlated ($\tau \sim 10\text{-}60$ minutes), presenting a fundamental obstacle in exoplanet direct imaging (e.g. Marois et al., 2008b). In order to remove these speckles, additional techniques are needed and can be separated into two categories – “real-time” and “post-processing” speckle suppression techniques.

Real-time Techniques

As the name suggests, the term “real-time” techniques here will be used to broadly describe a collection of algorithms that attempt to correct for, or remove, speckles while data is being taken on-sky. These include techniques which apply patterns to the DM to remove speckles such as speckle-nulling (Bordé and Traub, 2006; Martinache et al., 2014) and Electric Field Conjugation (EFC; Give’on et al., 2007; Pueyo et al., 2009; Riggs et al., 2016) as well as techniques which use the science camera as a simultaneous focal plane wavefront sensor – see Jovanovic et al. (2018) for a review. In the latter class of techniques, non-common path errors can be eliminated since the location at which the wavefront is being corrected is the same location that the science image is taken. Unfortunately for this to be effective however, fast low noise cameras capable of keeping up with the kilohertz speeds of the AO control loop are needed.

Post-processing Techniques

While focal-plane wavefront control methods can conceivably suppress quasi-static speckles and have achieved contrasts of up to 10^{-9} (Haffert et al., 2023), post-processing methods currently provide the most common way of removing them on-sky. Unfortunately, common post-processing techniques utilizing advanced PSF subtraction methods (e.g. Lafrenière et al., 2007; Soummer et al., 2012) become less effective at small IWAs where direct detections are most challenging. Angular Differential Imaging (ADI; Marois et al., 2006) exploits parallactic angle (PA) rotation to distinguish speckles, which will rotate with the telescope field of view, from companions, which are at a fixed location on-sky. The magnitude of this rotation, however, scales proportionally with angular separation for a given unit time, resulting in less rotation at smaller IWAs. Additionally, the rotation in λ/D units is smaller within a few diffraction beamwidths, resulting in severe self-subtraction of a planet signal (Mawet et al., 2012). Similarly, Spectral Differential Imaging (SDI; Marois et al., 2000) utilizes the wavelength-independent nature of phase-induced speckle noise to rescale (magnify) slices of polychromatic images. However, SDI requires broad spectral coverage close to the primary otherwise it also suffers from self-subtraction effects. Reference Star Differential Imaging (RDI/RSDI; Soummer et al., 2012) uses an image of a reference star without a companion to subtract off the quasi-static speckles, which should be common, and preserve the differential companion signal. This technique does not inherently suffer at small IWA, but requires careful magnitude and color matching between the target of interest and the reference star. Mismatches in the direction of the gravity vector with respect to the primary mirror and in the position of the telescope rotator between reference observations and target observations can also degrade RDI performance, placing even tighter constraints on the choice of reference star (Ruane et al., 2019).

A method to suppress quasi-static speckles that is free of the limitations of ADI, SDI, and RDI would significantly improve our ability to detect jovian planets at Jupiter-to-Saturn like separations. One such technique called Stochastic Speckle Discrimination (SSD; Gladysz and Christou, 2008a) will be discussed later in this work and exploits differences in the intensity distributions of companions and speckles at millisecond timescales to differentiate the two signals.

1.2 The MKID Exoplanet Camera (MEC)

1.2.1 Microwave Kinetic Inductance Detectors (MKIDs)

In order to perform many of the techniques highlighted in the previous section, fast, low-noise cameras are essential. Typical astronomical detectors are semiconductor based sensors such as charge-coupled devices (CCDs) and complementary metal oxide semiconductor (CMOS) detectors, but in the past decade a push has been made towards developing large format superconducting arrays for astronomy.

These superconducting detectors have the benefit of essentially no read noise or dark current and also have inherent energy resolving capabilities allowing for spectra to be obtained without the addition of throughput reducing optics. Microwave Kinetic Inductance Detectors (MKIDs) developed to operate in the ultraviolet, optical, and near-infrared (IR) are one such detector technology being explored for this purpose and currently show much promise for both interfacing with AO systems in real time, and in imaging and obtaining simultaneous spectroscopy of objects.

MKIDs are superconducting photodetectors capable of measuring the arrival time (to within a microsecond) and energy ($R \sim 20$) of incident photons with no read noise or dark current (Day et al., 2003; Mazin et al., 2012; Szypryt et al., 2017; Zobrist et al., 2022).




Figure 1.5 Left: Subsection of an MKID array showing many pixels connected to a single microwave readout feedline. Right: Single MKID pixel showing the photosensitive meandered inductor (yellow, top) and interdigitated capacitor (purple, bottom).

Each MKID pixel is a lumped element superconducting LC resonant circuit composed of a photosensitive inductor and a tunable interdigitated capacitor, see Figure 1.5. With the tunable capacitor, each MKID pixel can be fabricated with a unique identifying resonant frequency, currently selected to fall between 4-8 GHz (in the microwave).

When a photon strikes the inductor of an MKID pixel (cooled below its superconducting transition temperature), quasiparticles are generated through the breaking of Cooper pairs which are the charge carriers in a superconductor. This increases the inductance of the material and lowers the resonant frequency of the circuit which can analogously be measured as a change in phase by room temperature readout electronics – see Fruitwala et al. (2020). Since the number of quasiparticles generated by the photon event is proportional to the energy of the incident photon, MKIDs have an inherent energy resolution without the use of filters or gratings.

Since each MKID pixel has its own unique resonant frequency, many thousands of individual MKID pixels can be placed onto a single microwave readout feedline to generate




Figure 1.6 Image of a full 20 kpix MEC array in its microwave readout packaging.

large format arrays as is shown in Figure 1.6. This is especially important for astronomical imaging where megapixel arrays are the standard.

Another advantage of MKIDs is that even in these large arrays, each MKID pixel can be sampled at a rate of up to 1 MHz, yielding microsecond timing resolution (Fruitwala et al., 2020). Unlike traditional cameras, the output of an MKID camera is therefore a photon list containing the registered arrival time and energy of each photon which can then be reduced to perform science, see Chapter 2.

One advantage of this timing resolution is MKIDs ability to identify and cut out false counts due to cosmic rays. Cosmic rays are high energy particles that move through space at near light speeds and can damage the silicon lattice in CCDs through a phenomenon called charge transfer deterioration. With MKID instruments, we can identify spurious counts due to a cosmic ray event and remove them from further analysis without any lasting effects on the detector itself (see §2.2.2). This feature is particularly advantageous for space-based applications where expected count rates are very low and cosmic ray events more likely.

1.2.2 MEC at Subaru/SCEXAO

The MKID Exoplanet Camera (MEC; Walter et al., 2020) is the first permanently deployed MKID instrument for high-contrast imaging and is located behind the Subaru Coronagraphic Extreme Adaptive Optics system (SCEXAO; Jovanovic et al., 2015a) at the Subaru Telescope on Maunakea. The enabling technology for MEC is a 20,440 pixel MKID array optimized for detecting photons in the near-IR ($Y - J$ band) to perform high-contrast imaging of exoplanets, brown dwarfs, and disks. Due to the fast timing information available to MKIDs it was also designed to serve as a focal plane wavefront sensor and perform real time speckle suppression techniques on-sky including speckle nulling (Fruitwala, 2021).

The MEC cryostat (see Figure 1.7) is a pulse tube cooled Adiabatic Demagnetization Refrigerator (ADR) with an operating temperature of 90 mK for typical observing. It gets light directly from the SCEXAO bench via a fold mirror that also allows for simultaneous observing with other instruments behind SCEXAO in different observing bands. MEC is equipped with an internal filter wheel to apply neutral density (ND) to the system to prevent saturation of the detector for bright objects. It also has a 2.54 mm gold mirror mounted on a Newport CONEX-AG-M100D controller to be able to perform dithering observing sequences to help mitigate detector effects, especially the presence of dead pixels – see §2.2.3. More detailed specifications for the MEC instrument can be found in Walter et al. (2020).

1.3 Photon Statistics in Adaptive Optics Images

As was discussed in §1.1.2, in order to perform high-contrast imaging from the ground, AO systems are vital for removing the effects of Earth’s turbulent atmosphere and make direct imaging possible. This process of correcting the aberrated wavefronts however




Figure 1.7 MEC located behind SCEXAO on the Nasmyth platform at the Subaru Telescope.

noticeably changes the properties of the light in the final image at very fast frame rates which can be easily picked up by MKID detectors. In particular, the intensity distributions of on-axis companions and off-axis diffracted light speckles in millisecond exposures will be different. These distributions can have fundamental impacts on the noise in these images (see Bonse et al. (2023)) and are explored in depth in the following sections. These differences can also be exploited as a post-processing technique to separate diffracted light speckles from true companions as will be discussed in Chapters 3 and 4.

1.3.1 Speckles

Originally derived by Goodman (1975) and experimentally verified by Cagigal and Canales (2001) and Fitzgerald and Graham (2006), the underlying probability density function (PDF) that estimates the intensity distribution of off-axis stellar speckles in the image plane is given by a modified Rician (MR)

$$p_{MR}(I) = \frac{1}{I_S} \exp\left(-\frac{I + I_C}{I_S}\right) I_0\left(\frac{2\sqrt{II_C}}{I_S}\right) \quad (1.2)$$

where $I_0(x)$ denotes the zero-order modified Bessel function of the first kind, I_C describes the coherent intensity component attributed to the unaberrated PSF of the primary, and I_S is the time variable component of the total intensity that describes the speckle field (see also Marois et al., 2008b). It is important to note that the shape of this distribution is always positively skewed (i.e. the distribution tail falls to the right hand side of the mean) and that the skewness of the MR can be parameterized by the ratio of I_C/I_S (the larger the I_C/I_S ratio, the less skewed the distribution).

For a sequence of exposures shorter than the decorrelation time of atmospheric speckles (~ 10 - 100 ms), a histogram of the image plane intensity follows a MR and I_C and I_S can be determined for each pixel in an image (Fitzgerald and Graham, 2006). The optimal exposure time should also be long enough that each bin contains many photons. If too large of an exposure time is chosen, many realizations of the speckle intensity will be averaged over. Conversely, if too small of an exposure time is selected, then not enough photons will arrive per exposure and the distribution will tend towards Poisson statistics.

1.3.2 Companions

On-axis sources (non-diffracted light, i.e. astronomical objects in the image) behind an Extreme AO (ExAO) system do not follow MR statistics because of the ExAO system itself which can be thought of as a high-pass spatial filter acting on the phase (Sivaramakrishnan et al., 2001). At the center of an image, the complex amplitude is simply the integral of the pupil complex amplitude as given by Equation 1.3 (Soummer et al., 2007a).

$$\Psi(0) = \int P(x) e^{i\phi(x)} dx \quad (1.3)$$

Here $P(x)$ denotes the pupil function and $\phi(x)$ is the phase of the wavefront in the pupil plane.

With no (or low) AO correction, $\phi(x)$ is large and so the phase vectors, $e^{i\phi(x)}$, can take any orientation in the complex plane. Summing a large number of these vectors will produce a random walk resulting in circular Gaussian statistics by the Central Limit Theorem. This adherence to circular Gaussian statistics in phase will result in a MR distribution in intensity and be indistinguishable from off-axis speckles.

At high correction levels (i.e. behind an ExAO system), $\phi(x)$ is small and so the vectors are not oriented randomly in the complex plane. Summing them will not produce a random walk, the corresponding distribution is not a circular Gaussian, and the resulting on-axis intensity distribution will not follow a MR. The spatial extent where the transition between the on-axis and off-axis intensity distributions occurs is at $\ll 1\lambda/D$ (Soummer et al., 2007a) which can be seen qualitatively in Figure 1.8 – see also Appendix A. For reference, at Subaru $1\lambda/D$ is equal to 27.6 mas (2.75 MEC pixels) at $1.1 \mu\text{m}$.

Instead of following the MR, at these high correction levels the Strehl Ratio (SR) distribution (which is proportional to the intensity) instead follows the PDF described by Equation 1.4 which was derived independently by Soummer et al. (2007a) and Gladysz and Christou (2008b) :

$$p_{SR}(sr) = \frac{p_{\hat{\sigma}^2}(-\ln(sr))}{sr} \quad (1.4)$$

where sr is the instantaneous SR, and $p_{\hat{\sigma}^2}$ is given by

$$p(x; k, \theta, \mu) = \frac{\left(\frac{x-\mu}{\theta}\right)^{k-1} \exp\left(-\frac{x-\mu}{\theta}\right)}{\Gamma(k)\theta} \quad (1.5)$$




Figure 1.8 Histograms of the binned satellite spot intensity moving from three pixels away from the center of the spot (purple, top) to the center (yellow, bottom). Histograms shift from positive to negative skewness as they approach the center of the satellite spot as is shown by the best fit MR I_c/I_s ratio (dashed lines). Here a higher value indicates less skew. Since this data was taken in $Y+J$ band, the satellite spots are elongated in the image and the plotted intensities were found by adding the intensities for six pixels at the same specified distance from the spot center (three on each side). Insets: satellite spot image where the pixels used to generate the histograms are color-coded by their distance from the spot center. Dead pixels were purposefully avoided.

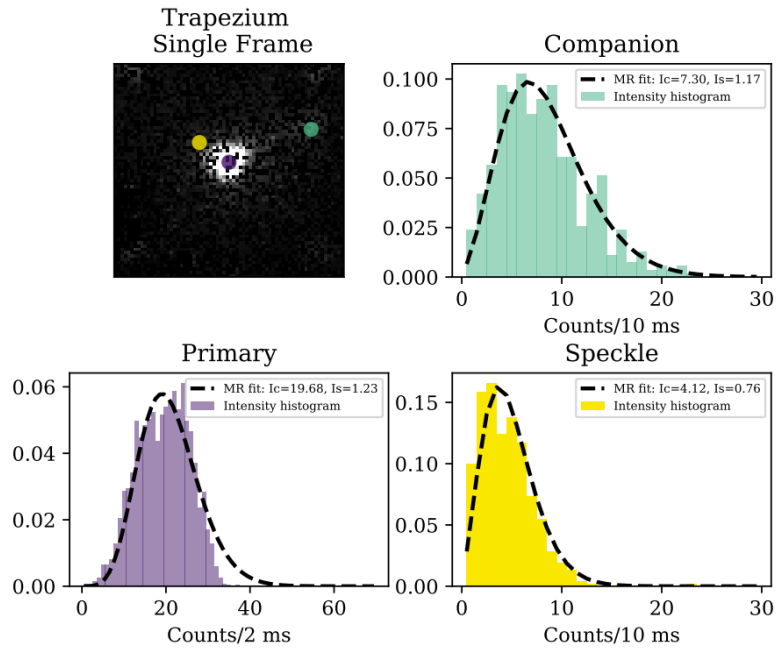


Figure 1.9 Single-pixel arrival time histograms taken on-sky with SCEXAO/MEC of the unocculted star Θ 1 Orionis (purple, bottom left), a faint secondary companion (teal, upper right), and at a random location in the field with comparable brightness to the faint secondary companion (yellow, bottom right). The best fit MR was plotted for each distribution and is shown with the dashed line. Note that due to the brightness of the primary, 2 ms time bins were used to generate that intensity histogram instead of the 10 ms time bins used for the secondary and field locations. It is clear to see that the MR adequately describes the field location, but breaks down at the location of the primary. The companion is seen to be a convolution of the primary and field PDFs.

Here $k > 0$ is the shape parameter, $\theta > 0$ is the scale parameter, and μ is the location parameter, which shifts the PDF left and right. $\Gamma(k)$ is a Gamma function. For ease of discussion we will refer to the entirety of Equation 1.4 as the ‘Gamma’ distribution for the remainder of this work.

In contrast to the MR distribution, this Gamma distribution is negatively skewed (i.e. the distribution tail falls to the left hand side of the mean). Figure 1.9 illustrates these differences in skewness for various on and off-axis sources with a SCEXAO/MEC observation of Θ 1 Orionis (The Trapezium Cluster).

1.3.3 Satellite Spots

Astrometric and spectrophotometric calibrations are very difficult when the target star (typically the only reference in the field-of-view) is obscured by a coronagraph. In order to perform these calibrations, faint copies of the obscured stellar PSF called “satellite spots” may be intentionally placed into the image plane. For SCExAO, these are generated by placing a “waffle” pattern (two orthogonal sine waves) on the AO system’s deformable mirror.

In SCExAO, the satellite spots are additionally modulated at a rate of 2 kHz by flipping the sign of the two sine waves (equivalent to phase shifting them by π). This is done to make the speckles incoherent with the underlying background and improve photometric performance (Jovanovic et al., 2015b). If there is a coherent speckle underneath the spots then these two polarities will not have the same brightness as one will interact constructively and the other destructively. In the regime of these millisecond images, this phase switching is unlikely to affect the statistics at the spot locations as 10s - 100s of these transitions are being averaged over in a single time bin. It is then expected that the satellite spots should follow the same on-axis statistics of the primary.

Figure 1.8 shows intensity histograms moving from the speckle field towards the satellite spot center for a single 25 s MEC observation of HIP 109427. The statistics shift from positively to negatively skewed, showing that the satellite spot statistics are not only distinct from that of the speckle field, but also qualitatively follow the same distribution expected for the primary. Additionally, this transition between the on and off-axis distributions occurs at the expected location of $\ll 1\lambda/D$ (< 2.75 MEC pixels).

The satellite spot statistics themselves are important because if the distribution of the satellite spots matches that of the on-axis source, then they can be used to help fit the free parameters of the on-axis Gamma distribution (Equation 1.4) when observing with

a coronagraph. Since the brightness of the satellite spots can be controlled by changing the amplitude of the sine waves on the DM, this would be the simplest way to determine the shape of the companion distribution using on-sky data at high signal-to-noise. Once the shape of this distribution is known it could then enable the use of post-processing techniques that try to explicitly separate the companion and speckle PDFs such as PDF deconvolution (see Gladysz and Yaitskova (2010)).

Chapter 2

The MKID Pipeline

MKID instruments such as MEC produce a raw data stream that differs from typical semiconductor based astronomical detectors and requires significant post-processing before it is effectively accessible to the broader astronomical community. To this end we have created the MKID Pipeline¹, a Python package to provide an open-source extensible data reduction pipeline for MKID data. This pipeline takes raw MKID data as input and processes it into either a traditional form (i.e. FITS cubes) to be used with existing astronomical analysis packages, or a unique MKID data product suitable for advanced analysis tailored to the detector’s unique abilities.

The MKID Pipeline is based on the development and use of the only three optical/near-infrared astronomical MKID instruments to date: ARCONS (Mazin et al., 2013), DARKNESS (Meeker et al., 2018), and the MKID Exoplanet Camera (MEC, Walter et al., 2020) with recent extensions implemented to analyze data from the newly deployed XKID instrument located behind MagAO-X (Males et al., 2018; Swimmer et al., 2022).

In this chapter, we first briefly describe the contents of a typical MKID observing dataset (§2.1). We then discuss how data is processed in §2.2, beginning with a de-

¹<https://github.com/MazinLab/MKIDPipeline>

scription of the contents of raw MKID data before diving into the specific calibration algorithms in depth. Finally, in §2.3 we end with a discussion of how these steps are implemented in software and how a user would perform basic data reduction.

2.1 MKID Observing Datasets

MKID detectors take data by recording the time, location, and phase response for each detected photon. For this reason, all time binning is performed in post processing. An MKID observation or ‘exposure’ therefore refers not to specific exposures determined during a night of observing, but to time ranges where the object of interest is on the detector at an intended position. The resulting total observational dataset consists of some number of science observations, associated observatory and instrumental metadata, and necessary calibration data.

Science observations consist of a single time range, target, sky position, and associated calibration datasets. Due to the current high level of detector defects (e.g. cold/dead pixels), it is common to take dithered data suitable for reconstruction of a sky mosaic (Hook and Fruchter, 2000). In MEC, a tip/tilt mirror is used for this purpose. This dither then consist of a series of science observations and corresponding tip/tilt mirror positions that are combined in post processing to generate a single output image (see §2.2.3).

Calibration data consists of a series of uniform monochromatic laser exposures relatively evenly spaced across the wavelength coverage of the detector. These exposures are used for wavelength calibration (§2.2.2) and can also be used for flat-fielding (§2.2.2), though sky flats may also be taken and used instead. Dark observations (intervals obtained with a closed shutter or on blank sky) may be included to remove instrumental or astrophysical backgrounds. Finally, support for observations of an astrometric refer-

ence is provided to calibrate the final output products to reflect real on-sky coordinates (§2.2.2).

Both science and calibration data have associated observatory and instrumental meta-data (e.g. observatory and telescope status information, detector temperatures). The instrument control software records all of this data periodically in a machine readable format for later use by the pipeline. A subset of this data must be provided either via these logs or specified by the user when defining data to ensure proper reduction.

2.2 Data Processing

Raw MKID data consists of per-resonant-frequency (an analog to pixel) time series of photon-induced phase shifts. These are associated with individual pixels, converted to tabulated photon event data for each observation, and calibrated via the pipeline diagrammed in Figure 2.2.

In brief, the telescope and instrument logs (along with user overrides) are first used to create an associated metadata time series for for each observation to properly carry out later steps and determine eventual FITS and output header keys. Cosmic rays are then identified within the photon list. A linearity calibration may be performed which calculates a weight for each photon to statistically correct for missing photons caused by a detector-imposed dead time inherent to the MKID readout. This dead time prevents the recording of a photon that arrives too close the the tail of the preceding photon and causes non linear responses at high count rates ($\gtrsim 5000$ photon pixel⁻¹, exceeding current instrument limitations). A series of monochromatic exposures is next used to determine the relationship between phase shift and wavelength for each pixel. Pixels that exhibit too strong (hot), too weak (cold), or no (dead) response to incoming photons are then masked and ignored in further analysis. Inter-pixel variations are next corrected by

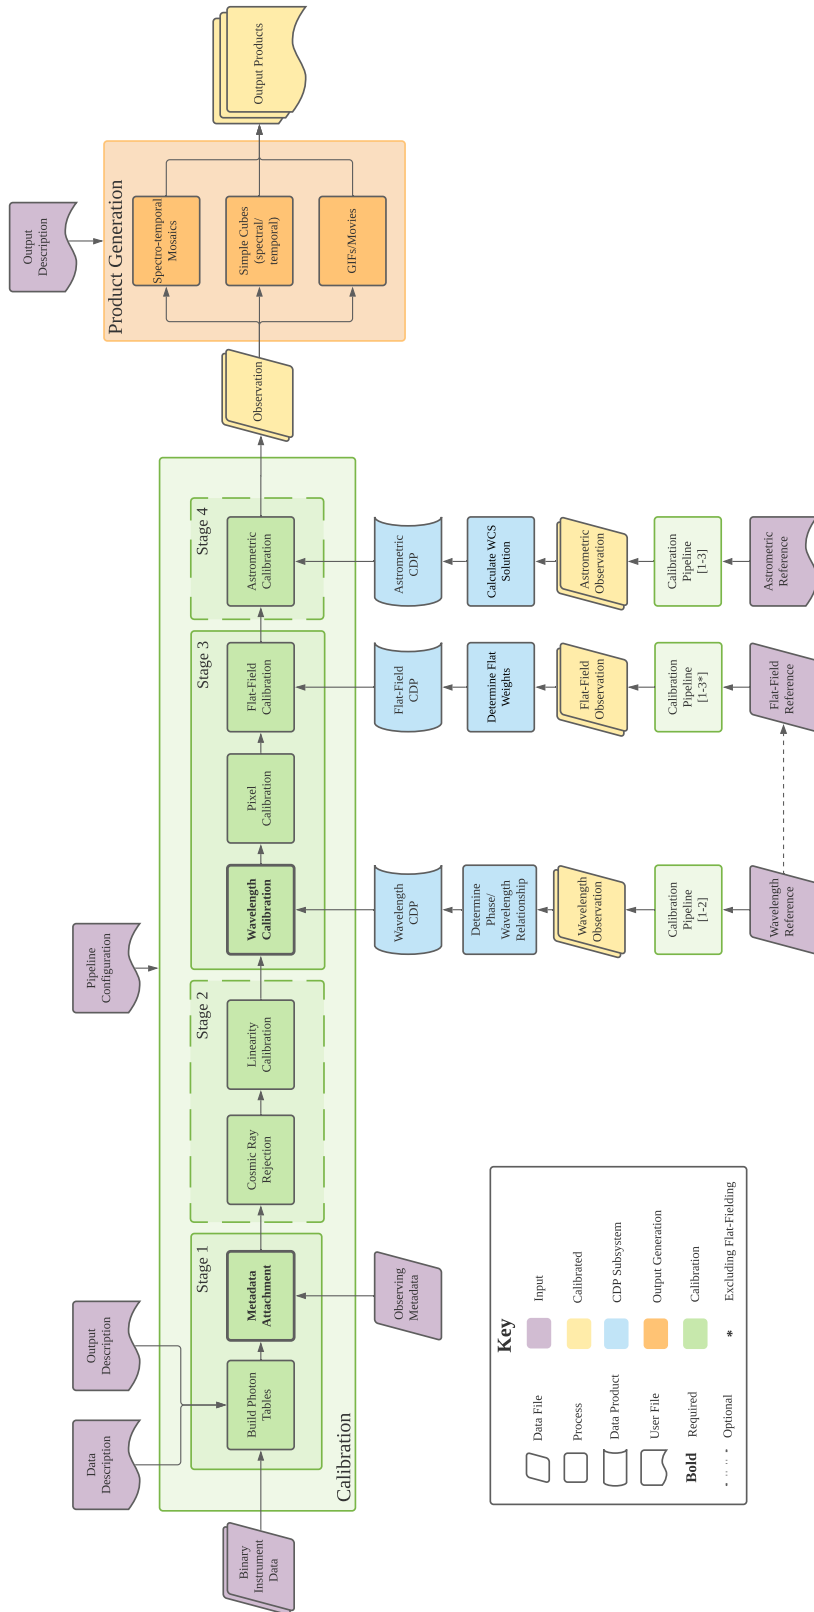


Figure 2.1 A block diagram of the MKID Pipeline. Processing steps are shown as rectangles with calibration steps in green, product generation steps in orange, and processes that are a part of the calibration data product (CDP) subsystem in blue. Files containing observing data (in various formats) are depicted by purple (input) and yellow (processed) rhombuses. User inputs are in purple. Calibration data products themselves are shown as blue cylinder faces. Rectangles with the wavy bottom edge represent files where the purple are user input configuration files and the yellow are the files representing calibrated output data products. All calibration steps are optional with the exception of the bolded 'Metadata Attachment' and 'Wavelength Calibration' steps without which product generation cannot occur. The dashed stages of the pipeline may be completely omitted. The bracketed numbers in the 'Calibration Pipeline' boxes denote pipeline stages that must be completed to generate the respective CDPs. Flat-fielding can be accomplished via reference to wavelength calibration data, denoted by the dashed arrow.

using a uniform polychromatic exposure or set of monochromatic exposures to determine a spectrally dependent flat weight for each pixel. Finally, an astrometric reference can be used to determine the pixel to world coordinate system (WCS) mapping for the instrument to yield physical output units for both the spatial and spectral dimensions of the output.

The resulting calibrated data is then used to create output products such as spectro-temporal FITS cubes, calibrated tables of photons, and movies. This section describes the algorithmic details of each calibration step outlined above. For details on the implementation of the pipeline itself, see §2.3.

2.2.1 Data Format

MKID detectors are read out via frequency multiplexing sets of pixels that share a microwave feedline. Photon arrival locations are therefore discriminated by frequency rather than detector position. This means the resulting raw MKID data is a per-resonant-frequency time-series of photon-induced phase shifts. Due to the potential for data rates up to $40 \text{ MB s}^{-1} \text{ kpix}^{-1}$ the data is recorded in a packed binary format (Fruitwala et al., 2020). This, coupled with the environmental sensitivity of MKIDs, necessitates the occasional determination of an optical beam position to pixel frequency mapping. At the start of pipeline processing this mapping, or “beammap”—which also contains information about malformed, inoperative pixels—is used to ingest this packed binary data and produce a tabulated photon list for further processing.

2.2.2 Data Calibration

Metadata Attachment

During observing, the instrument captures a record of telescope and instrument status information in addition to the photon data from the detector. After photon table construction, this data is parsed for records within the observing interval as well as the record immediately prior, forming a metadata time-series for each. These series, supplemented with any user specified values, are attached to the photon table.

Cosmic Ray Rejection

Cosmic rays incident on an MKID detector excite phonons in the detector substrate causing the majority of pixels to register photon events near-simultaneously for a brief duration. The cosmic ray rejection step identifies intervals where these false photons are recorded for use in later analysis. This is done by splitting observations into $\sim 10 \mu\text{s}$ time bins and using one of two techniques to compute a count rate above which a cosmic ray event is flagged.

The first approach assumes that count rates should follow Poisson statistics and employs `scipy.stats` to generate a count rate threshold (Virtanen et al., 2020). First, a cumulative density function (CDF) is determined which is defined by the number of standard deviations away from the mean that a given count rate needs to be for that time bin to be classified as containing a cosmic ray. A percent point function is then evaluated on that CDF at the average count rate to generate the threshold value. The second method calculates the standard deviation of the count rates using the total binned time stream, excluding data that falls outside of three standard deviations from the mean. The threshold is then defined as a user input number of those standard deviations above the mean value.




Figure 2.2 Photon time stream where the red arrows denote locations of identified cosmic rays. Excluding all photon data obtained in a $10 \mu\text{s}$ window around each event would eliminate a total of 0.00195 s , $< 0.01\%$ of the exposure. Any missed cosmic rays would contribute no more than a single photon per pixel in an astrophysical source.

In both cases, bins that exceed the computed threshold are flagged as cosmic ray events and their time intervals, total and average counts, and peak count rates recorded in the photon table's header. Due to the microsecond timing resolution of MKIDs the total time lost due to cosmic rays in a typical dataset is less than 0.01% of the total observation time. In contrast to a CCD detector, missed events would only add a single count to each pixel. For this reason, cosmic ray rejection is presently implemented in a way that does not alter the original photon time stream and removal is not merited unless a particular analysis is sensitive to false counts at the 10's of photons level. Figure 2.2 shows an example MKID photon time stream with cosmic rays identified.

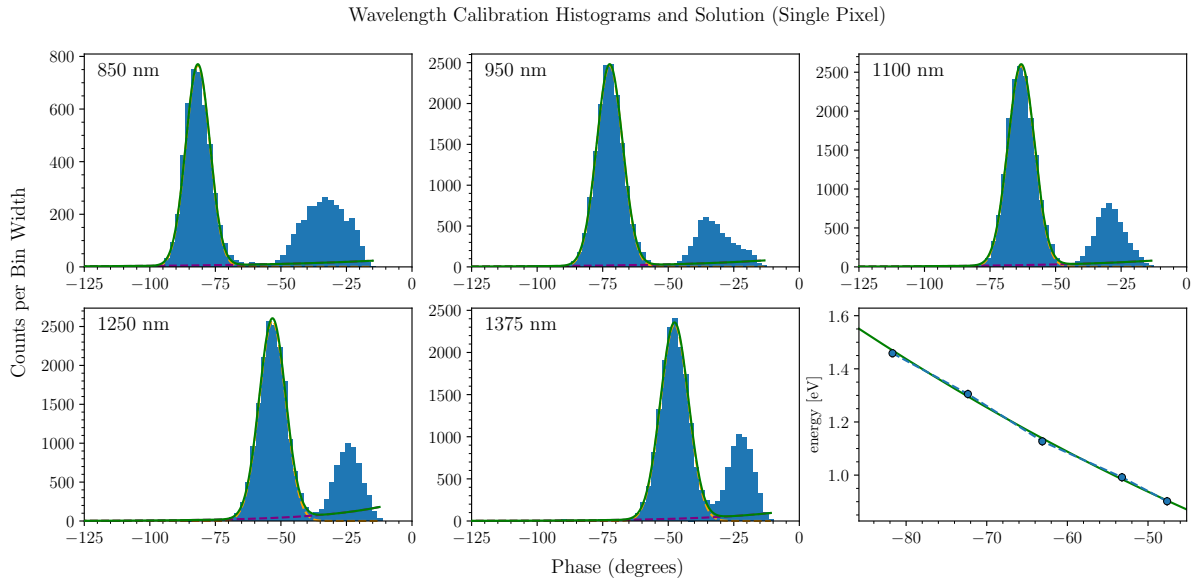


Figure 2.3 Single pixel count rate histograms for each laser wavelength as well as the calibration solution fit (bottom right). The small Gaussian bump at low phases is likely due to an IR leak around $2.7 \mu\text{m}$ in the filter stack of the instrument (MEC) used to take this data.

Wavelength Calibration

The wavelength calibration calculates the relationship between the phase response of each pixel and the wavelength of each incident photon via phase pulse-height histograms generated from a series of monochromatic laser exposures. These exposures are typically generated by using a series of lasers spanning the wavelength sensitivity range of the particular instrument coupled with an integrating sphere to ensure a uniform illumination on the array.

The phase histograms are fit using one or more of a series of models. Current supported models are a Gaussian signal plus a Gaussian background, and a Gaussian signal plus an exponential background. If more than one model is specified then all are attempted and the best fit one used. When provided, a dark observation is used to subtract a background count rate from the phase histograms to yield a better fit.




Figure 2.4 Resolution image at $1.1 \mu m$ for the detector in MEC as of January 2022. The median energy resolution (\mathcal{R}) across the array at this wavelength is 3.93 excluding dead pixels. This particular detector has three defective feed lines (each containing 2000 pixels) which results in the large strip of dead ($\mathcal{R} = 0$) pixels seen to the left of the image

Once the phase histograms are fit, the centers of each histogram are determined and fit as a function of laser wavelength with a linear or quadratic function to determine a final phase-wavelength calibration for each pixel (Fig. 2.3). The resulting fits constitute a wavelength calibration data product that consists of a per-pixel mapping of phase to wavelength, a set of associated calibration quality flags, and general solution metadata. A sample resolution map at $1.1 \mu m$ is shown in Figure 2.4

Individual observations are then calibrated using the appropriate (e.g. user-specified, temporally proximate) solution for a given observation by loading each pixels' phases and feeding them through the associated mapping. The resulting wavelengths, associated flags, and wavelength calibration metadata are then stored in the observation.

Pixel Calibration

The pixel calibration identifies ‘hot’, ‘cold’, and ‘dead’ pixels to be removed from further analysis. Pixels that register counts a specified number of intervals above a threshold are flagged hot, below a threshold cold. Dead pixels are first determined based on the detector’s beammap and the array image is then passed through a filter which iteratively replaces the dead pixels with the mean value of pixels in a surrounding box until none remain. This is done before the determination of hot and cold pixels so as to not skew the algorithms. Three algorithms are provided for determining the hot and cold thresholds and associated interval for each pixel.

Threshold: This method compares the ratio of the image to a moving-box median that excludes both the central pixel and any defective pixels. Ratios greater than some tolerance above/below the peak-to-median of a Gaussian PSF are flagged as hot/cold respectively. See Figure 2.5 for a sample of this algorithm used on a dataset. Care must be exercised to ensure the moving box is sufficiently large so as to not be biased by clusters of hot or cold pixels.

Median: The detector array’s median count value is used as the global threshold. The tolerance interval is determined by applying a standard deviation moving box filter to the counts image.

Laplacian: A Laplace filter (`scipy.ndimage.filters.laplace`) is applied to the image and the result adopted as the count threshold. The standard deviation of the filtered image is used for the tolerance interval.

Linearity Calibration

Each pixel has a finite dead time, imposed in firmware, that precludes detection of photons arriving within a small time interval following the preceding photon. The exact




Figure 2.5 Subset of an MKID array with hot, cold, and dead pixels labeled. The threshold method was used in the determination of the pixel flags with default settings.

interval value depends on the quasi-particle recombination time of the superconducting film and the LC time-constant of the resonator. For MEC, this dead time is set in firmware to be $\sim 10 \mu s$. As a result, MKID detectors exhibit a nonlinear response that requires correction at high count rates (see Fig. 15 of van Eyken et al. (2015)). This correction is equal to $(1 - N \cdot \tau / T)^{-1}$ where N is the number of detected photons in time T for a pixel with dead time τ . The time T is set by the user and should be small so as to effectively determine the instantaneous count rate for each photon.

The need to compute and operate this calibration on per-pixel inter-photon arrival times can result in expensive computation, especially as single exposures may easily contain $> 10^9$ photons. As the effect is less than one part in 1000 for typical count rates, the use of this step is generally discouraged.

Flat-field Calibration

Flat-field calibration has two modes: laser and white light. In both modes, a spectro-temporal cube is generated and used to determine the per-pixel wavelength response weights necessary to achieve a uniform response across the detector array. To calculate this weight, the cube is normalized by the integrated average flux at each wavelength and then a user-specified number of the highest and lowest flux temporal bins are excluded to control for time-dependent contamination of the flat, e.g. radio frequency interference. The average of the remaining temporal bins is then fit as a polynomial function of wavelength and the fit saved as the flat-field calibration data product for later application. Data is flat-fielded by evaluating the polynomial at the wavelength of each photon and incorporating the resulting spectral weight into the photon table.

White Light Mode: Uses an observation of a uniform continuum source (e.g. twilight, dome) to generate the spectral cube. In this mode, the spectral sampling is determined by the nominal wavelength resolution set by the associated wavelength calibration.

Laser Mode: Generates the spectral cube using a series of monochromatic laser exposures such as the ones used for the wavelength calibration (see 2.2.2). This can be done by either positing that the laser frames are truly monochromatic (i.e. not imposing any wavelength cut on the exposures), or by using the wavelength calibration solution to only include photons within a small window around each laser wavelength. An example of the flat-field calibration using the laser mode applied to a real dataset is shown in Figure 2.6.

Astrometric Calibration

The astrometric calibration determines the World Coordinate System (WCS) transformation parameters to convert an image from pixel (x, y) to on-sky (RA, Dec.) co-

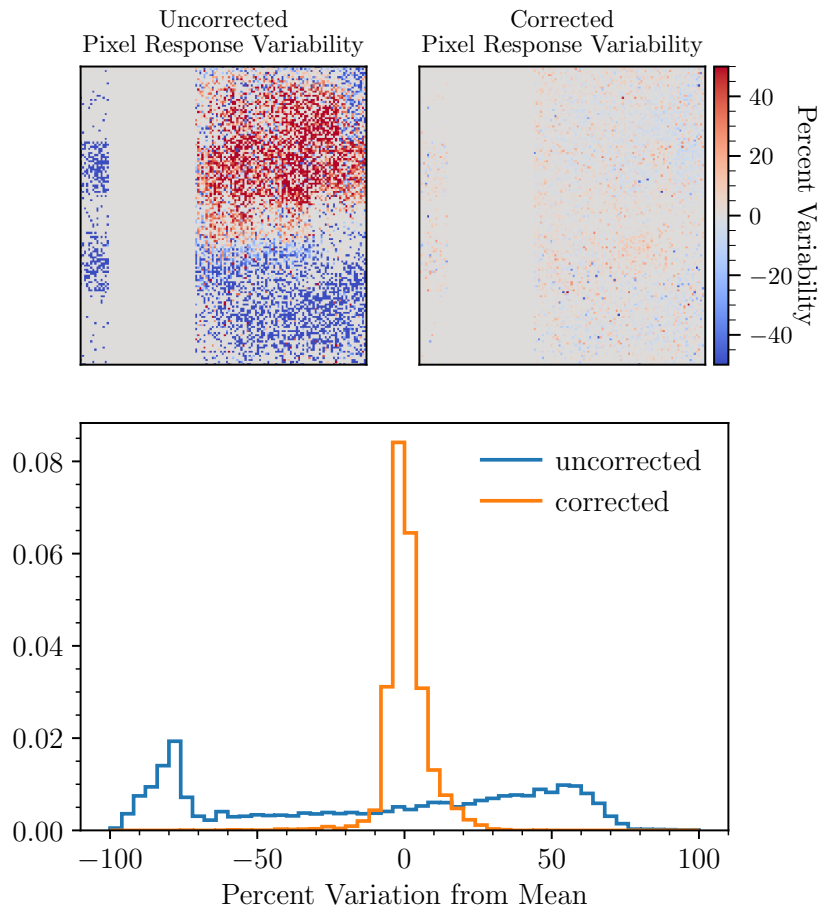


Figure 2.6 Top: Percent variability in pixel response before and after applying the flat-field calibration. This is calculated by subtracting and then dividing the median counts registered on the detector from each pixel. The structure seen in this MEC data is dominated by vignetting from the optical system. Bottom: Histograms of the percent variability with the uncorrected pixel response shown in blue and the corrected pixel response shown in orange.

ordinates. First, a point spread function (PSF) fit is performed to determine the pixel location of each source in each image of the observation. Here an ‘image’ is defined by any single exposure where the pixel and sky locations of the sources are expected to remain constant (e.g. the telescope pointing does not change, the tip/tilt mirror is in the same position, etc.).

Each fit PSF location is then assigned an RA and Dec. through the use of an interactive tool where the user selects the approximate pixel location of the PSF for each source coordinate. The fit position of the nearest PSF to the selected coordinates is then assigned to the corresponding sky coordinate to generate a dictionary of pixel-sky coordinate pairs. When complete, the transformation between pixel and sky coordinate is then determined by solving for the WCS parameters by performing the following.

First, the tip/tilt mirror to pixel mapping is determined by fitting a linear model to the PSF centers (p_x, p_y) and corresponding mirror positions (c_x, c_y) .

$$p_{x,y} = \mu_{x,y}c_{x,y} + a_{x,y} \quad (2.1)$$

Here, the slopes $\mu_{x,y}$ give the number of pixels moved for a given tip/tilt mirror position change in either x or y, and $a_{x,y}$, is the pixel location corresponding to tip/tilt mirror position $(0, 0)$.

Next, the x and y platescales (η_x, η_y) are found using the known separation and pixel displacement of the sources. The platescale is calculated for each image and the mean value saved.

Finally, an affine transform is applied to the pixel coordinate point consisting of the following steps:

1. Rotation by an angle Φ to account for the detector’s rotation with respect to the telescope beam

2. Translation by an amount $(\mu_x c_x, \mu_y c_y)$ where c_x and c_y are the tip/tilt mirror positions.
3. Scaling by the platescale (η_x, η_y)
4. Rotation by the telescope position angle (Θ)

The (RA, Dec.) telescope offset is then added to the transformed pixel coordinate to complete the mapping. This results in two equations for each image (n_{im}) and each coordinate pair (n_s) giving a total of $n_{im} \cdot n_s \cdot 2$ equations. Each equation is solved for the last unknown WCS variable, the detector rotation Φ , using `scipy.optimize.fsolve` and the mean value saved. Values of μ_x , μ_y , η_x , η_y , and Φ are all saved within the photon table metadata.

2.2.3 Data Products

The calibrated photon tables output by the calibration stage of the pipeline consist of rows of individual photons with columns of time, resonator ID, wavelength, and weight. The resonator ID is a unique five to six digit identifying number given to each pixel to determine its location on the array in conjunction with the beammap. The weights are the multiplicative combination of the linearity and flat-field calibration steps. These photon tables may be used directly for analyses that rely on photon arrival time information, such as stochastic speckle discrimination (see Chapters 3 and 4).

The pipeline is also able to produce traditional astronomical outputs in the form of spectro-temporal cubes from individual observations or dithered mosaics and movies as are described below. Spectral and temporal FITS cubes with arbitrary wavelength and time bin widths may also be generated from individual exposures.




Figure 2.7 Left: MEC Image of the HIP 109427 system with each dither position color-coded by its order in the sequence. The image of the HIP 109427 (behind a coronagraph), satellite spots, and stellar companion are shown in grey scale to be able to better see the frame boundaries. Right: Exposure coverage footprint for the same dataset. Here bright regions have more effective exposure time than darker regions. The dithering script used to generate this dataset enforced a rectangular dithering pattern leading to a non-uniform footprint, but recent improvements have optimized this pattern for maximal uniform coverage.

Spectro-temporal Mosaics

A common observing strategy with MKID instruments is to dither using a tip/tilt mirror to fill in regions of dead pixels and increase the field of view. A mosaic from these dithered observations may be formed into a spectro-temporal FITS cube by combining each frame onto a common on-sky grid using the `DrizzlePac` implementation (Gonzaga et al., 2012) of the `Drizzle` algorithm (Fruchter and Hook, 2002). Each frame is mapped onto a sub-sampled output image to generate a single combined image, a spectral cube, a temporal cube, or a spectro-temporal cube with arbitrary wavelength and temporal axes. This allows for the generation of contiguous outputs even with pixel yields of $\sim 75\%$ on active feedlines (Walter et al., 2020), see Figure 2.7.

As all presently supported MKID instruments operate without an image derotator

the sky rotation is generally removed from each frame, resulting in an output where every frame is North aligned. An ADI mode is offered to facilitate interfacing with the Vortex Image Processing package for high-contrast direct imaging (VIP; Gomez Gonzalez et al. (2017)) in which each frame in the sequence is rotated so that the first frame is North aligned, but the parallactic angle rotation between frames is preserved.

Movies

Movies may be output in GIF or MPEG-4 format and come in two types. The first shows subsequent frames with the desired temporal resolution and run time and is well suited to show rapidly changing features, such as diffracted speckle patterns that vary on millisecond timescales (Goebel et al., 2018). The second format integrates the series of frames and is helpful to illustrate how increasing exposure time affects the final output image.

2.3 `mkidpipeline`: The MKID Pipeline Package

The MKID Pipeline is implemented as the Python 3 package `mkidpipeline`² that includes a corresponding conda environment definition file. The package provides a command-line program, `mkidpipe`, to process observational data and is configured via an instrument name and three YAML files: `pipe.yaml`, for general and step specific settings; `data.yaml`, which defines the data; and `out.yaml`, which specifies output products. Instructions for basic pipeline setup and execution of a sample dataset are provided in the package `README`. Complex data processing is expected to require direct use of pipeline methods in a user script. The following subsections describe the pipeline implementation. Additional details may be found in the source code.

²<https://github.com/MazinLab/MKIDPipeline>

`mkidpipeline` is composed of the modules `pipeline`, `photontable`, `definitions`, `config`, and `samples`, along with the sub-packages `steps`, `utils`, `data`, and `legacy`. Example data and default configuration files are stored in the `data` and `config` directories, respectively.

2.3.1 Concept

The pipeline steps as outlined in §2.2 are implemented as modules in `steps`, with the requirement that each define a `FlagSet` at `FLAGS`, a `StepConfig` (see §2.3.2), and an `apply()` method. Steps may also implement `fetch()` when there is a need to compute a persistent calibration data product (CDP), e.g. a wavelength or flat-field calibration solution file. If implemented, `fetch()` will be provided a path that is guaranteed to be unique for the input data and step configuration used to generate the CDP. This allows multiple users to use these files from a shared location without duplication of effort.

2.3.2 Initialization and Configuration

Each step module with settings is required to implement a subclass of `config.BaseStepConfig` named `StepConfig`. In its simplest form, this merely consists of a class-member listing of setting names, default values, descriptions and a YAML tag, though support is provided for additional verification of parameters that may have complex inter-dependencies or depend on other settings from other steps.

The pipeline places a configuration object for programmatic and interactive use at `config.config` after initial configuration (e.g. by loading a `pipe.yaml`). Access to a fully populated, isolated configuration object is available via the `PipelineConfigFactory`. This allows individual steps to not worry about whether or not the pipeline has been configured via a file and ensures that required step defaults are present. It also means

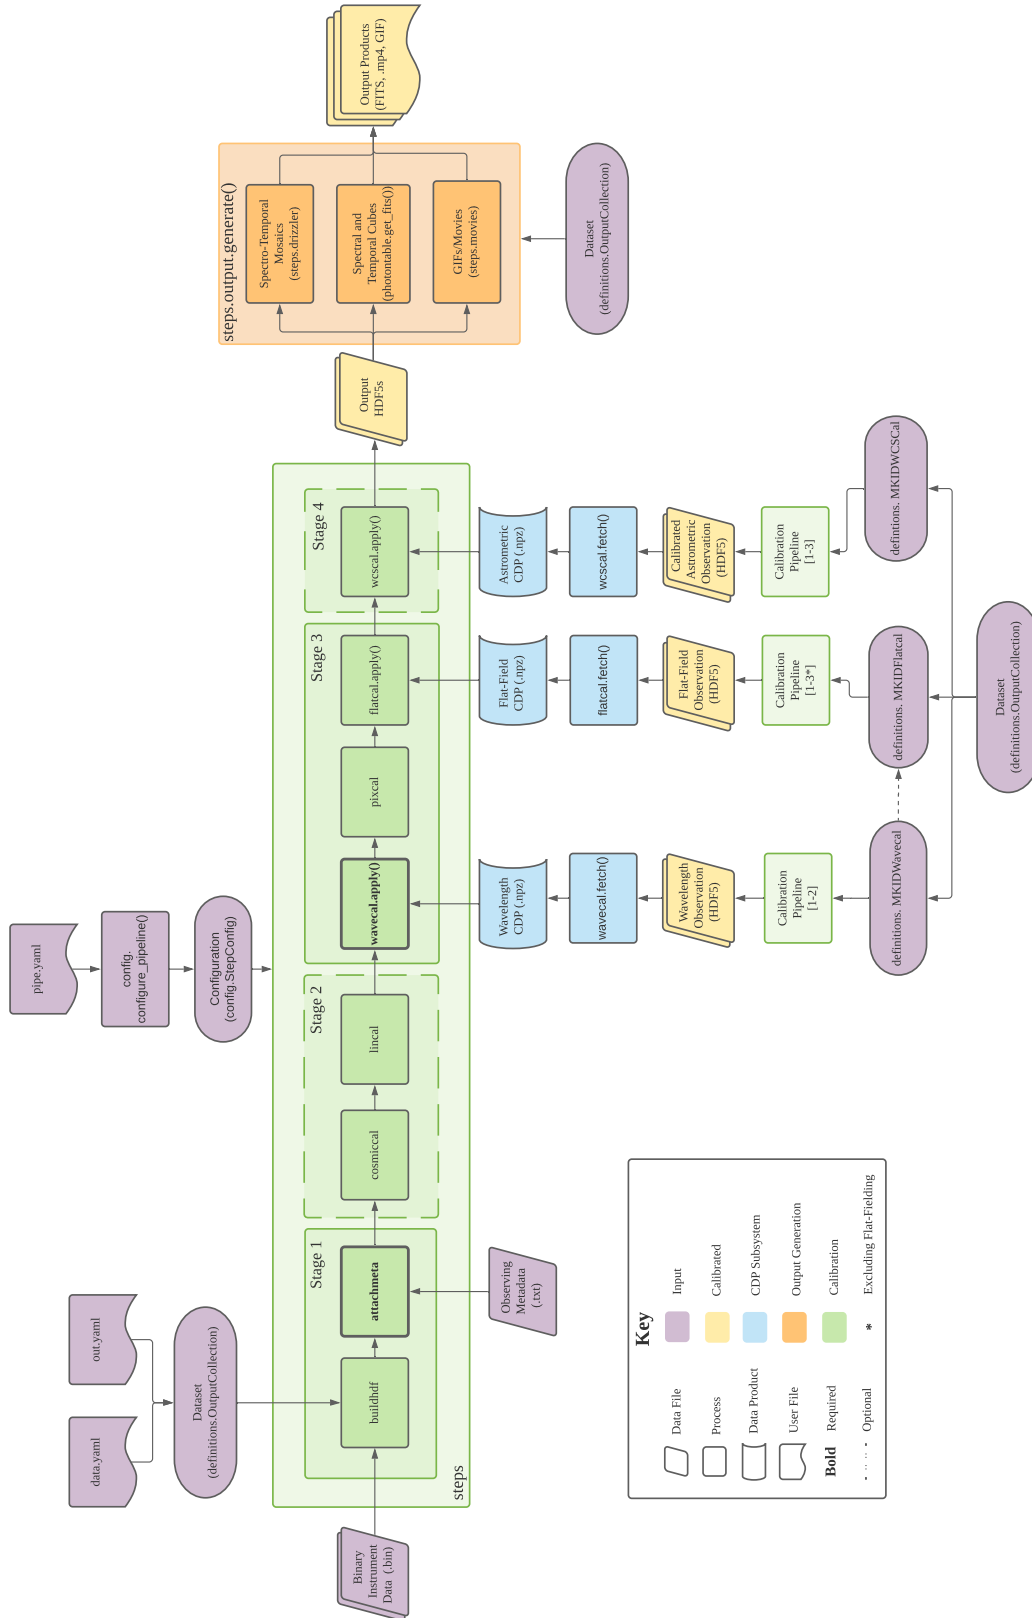


Figure 2.8 Block diagram of the MKID Pipeline depicting implementation details.

that any accidental mutations of the configuration do not propagate to other steps or processes. The configuration object supports parameter inheritance, however default values for individual steps can result in unexpected behaviour as the existence of a child default will take precedence over an explicitly set parent setting.

When imported, the `pipeline` module loads steps from the `steps` sub-package, registering any defined configuration classes with the pipeline YAML loader. These are then available to build a `config.PipeConfig` of top-level and step specific settings via `pipeline.generate_default_config()` or by loading a config file `config.configure_pipeline()`.

In addition to configuration options, the pipeline maintains a set of named flags that may be associated with individual pixels. Flag support is achieved by requiring steps to list any flags they would like to set as a tuple of strings named `FLAGS`. These are parsed when the steps are loaded and used to build a `FlagSet` object at `config.PIPELINE_FLAGS` that is capable of converting between flag names and bitmasks. The `FlagSet` is implemented in such a manner to ensure forward compatibility with pipeline data as new flags are added.

`mkidpipeline.samples` provides sample datasets and outputs for both programmatic reference and use by the pipeline to generate default `data.yaml` and `out.yaml` configuration files during initialization. The resulting files provide comprehensive samples with sensible defaults that may be used to test the pipeline. The raw data is not included due to its extremely large size.

2.3.3 Data Specification

`definitions` provides classes to manage the description and use of calibration and science data. Data definitions may be created either via class instantiation or via YAML,

where support is provided for linking unnested data descriptions by name. For example, an observation may specify a `wavecal` to use via the name of a top-level `wavecal` (i.e. not defined explicitly within a different observation) within the same `data.yaml`. Though possible, it is not generally advised to nest definitions. The `MKIDObservingDataset` is used to represent collections of data definitions and defines properties to access key groupings of data: `<stepname>able` (e.g. `wavecalable`) are definitions that can have the step applied and `<thing>s` (e.g. `wavecal`s) are definitions of that thing.

All observational data are sub-classes or collections of `MKIDTimerange` objects. This object is defined by a name, a UTC start time (as a Unix timestamp), a stop time or duration, an optional nested `MKIDTimerange` for a dark exposure, and an optional set of header key overrides. It provides support for metadata retrieval from instrument logs, accessing the associated detector beammap, HDF5 path, and convenience methods for accessing the table of photon data (see §2.3.6).

Scientific observations are instances of `MKIDObservation`, which requires the specification of a `wavecal`, `flatcal`, and `wscal`. Dithered observations are represented by `MKIDDither` which has similar calibration requirements to `MKIDObservation`. The dither, however, takes a single `data` specification which may be either a list of `MKIDObservations`, a timestamp within a dither log, or the fully qualified path to a dither log. In the latter two cases, the list of `MKIDObservations` is built from the dither, specified calibrations, and any extra header information.

All calibration datasets (including `MKIDWavecal`, `MKIDFlatcal`, and `MKIDWCSCal`) include the `CalibMixin` mix-in. This provides support for accessing the input time ranges as well as the creation of unique hash strings to identify calibrations made with specific data and settings. Wavelength calibration datasets are represented by an `MKIDWavecal` and take a list of `MKIDTimeranges` named by laser wavelength (e.g. ‘1000 nm’) as its data. Flat-field calibrations (`MKIDFlatcal`) take either a list of `MKIDObservations` or

the `name` of an `MKIDWavecal` as data input. If a `MKIDWavecal` name is provided then a `wavecal_duration` and `wavecal_offset` must be given. These specify the duration and starting offset relative to the wavecal’s photon tables are used to create new, wavelength calibrated tables for the flat-field calibration.

Astrometric calibration data is represented by an `MKIDWCSCal` and takes either a `platescale`, an `MKIDDither`, or an `MKIDObservation` as data. A `pixel_ref` and `conex_ref` are also required that define a tip/tilt mirror home position and a corresponding pixel location, if applicable. If a `dither` or `observation` is used, `source_locs` must list the sky coordinates of the targets.

2.3.4 Output Specification

Individual outputs are defined by a named `MKIDOutput`. This class is defined by a `name`, a `data` string specifying a `MKIDDither` or `MKIDObservation`, and a `kind` which specifies the type of output (e.g. `movie`, `drizzle`). Optional keys include minimum and maximum wavelength bounds (`min_wave`, `max_wave`), `exclude_flags`, a `duration`, a `filename` which specifies the desired name of the output file, `units` (`photons` or `photons/s`), `use_weights` which weights photons by their pipeline weights, `adi_mode` which preserves parallactic rotation between drizzled frames (see §2.2.3), a `timestep` which will yield a temporal cube if non-zero, a `wavestep` which will yield a spectral cube if non-zero, and fields that determine which calibration steps will be applied to the output (e.g. `wavecal`). If a `movie` is requested, a `movie_runtime` is also required. `MKIDOutput` provides the pipeline with the properties `wants_<outputtype>` and `output_settings` to help determine what output types are needed and what settings need to be used with `output.generate()` (see §2.3.5).

The `MKIDOutputCollection` manages the outputs and defines relevant properties to

be used by `outputs.generate()`. These include `to_<stepname>` (e.g. `to_wavecals`) which gather all of the data definitions needing a particular step given the current configuration, data, and outputs requested. It also provides properties similar `MKIDObservingDataset` that filter a potentially large dataset down to the subset needed for a particular set of outputs.

2.3.5 Execution

The command-line program `mkidpipe` provides arguments for help, initialization, input verification, and pipeline execution. On initialization, it creates a commented set of pipeline YAML files in the working directory populated with all available settings and a set of default data and output definitions. Re-invoking `mkidpipe` will validate the files and begin the data reduction process.

On execution, `mkidpipe` configures the pipeline via `config.configure_pipeline()` and then loads the data and output YAMLs by instantiating an `MKIDOutputCollection`. The dataset and outputs are then validated via `outputs.validation_summary` and any issues presented to the user for correction. The program then proceeds to call first `fetch()` and then `apply()` for each applicable step required for each output. This can be seen diagrammatically in Figure 2.3.1. Finally, the entire `MKIDOutputCollection` is fed to `outputs.generate()`. This function executes `photontable.Photontable.get_fits()` for a spectral or temporal FITS cube from an observation, `movies.fetch()` for a GIF or MPEG-4 output, or `drizzler.form()` for a combined spectro-temporal mosaic FITS cube as needed. Existing outputs are not, by default, overwritten.

2.3.6 Core Modules and Libraries

Much of this functionality is mediated by the `Photontable` class described below. The pipeline also depends on AstroPy (Astropy Collaboration, 2018), PyTables (Team, 2002–), and the python2/3 compatible library `mkidcore`³. `mkidcore` is used for tasks such as logging, flagging, parsing instrument readout information, and managing instrument specific settings. This package ensures compatibility may be maintained across a number of instrument readout systems without editing the pipeline.

The `photontable` module implements `Photontable` which handles all interaction with underlying photon data, loading data from and manipulating the underlying HDF5 file representation. Key functionality is provided to (un)flag pixels, interact with observing metadata, select subsets of photons by wavelength range, time range, and pixel, and form FITS images and cubes (with associated WCS information, if available). Functionality is generally dependent upon what pipeline processing has been completed.

2.3.7 Interactive Use

Users are able to import the `mkidpipeline` package, create data and output definitions programmatically in a similar manner to that done for pipeline initialization in `mkidpipeline.samples`. Step operations and numerous utility functions are then available to be used interactively on the data from a terminal.

2.3.8 Extending Functionality to Future MKID Instruments

To utilize the pipeline with a new MKID instrument a CSV file containing all of the desired FITS keys should be added to `mkidcore`. FITS keys should comply with the FITS standard format and any non-standard keys should provide a mapping so that the

³<https://github.com/MazinLab/MKIDCore>

corresponding standard is populated with the correct values. All `PIPELINE_KEYS` (keys that the pipeline uses to save and store necessary values) must be common for each instrument.

The new instrument should also be added to the `INSTRUMENT_KEY_MAP` in `mkidcore.metadata` which will take the name of the instrument from the `pipe.yaml` and automatically pull the specified keys to appropriately populate the HDF5 files during metadata attachment. Additional instrument defaults and specifications (such as feedline and readout board information) should be added to `mkidcore.instruments`.

All of the packages in `mkidpipeline` are instrument agnostic and so once these mappings are added, and the configuration YAML files updated, the pipeline may then be run as previously described.

2.4 Summary

The MKID Pipeline is an open-source extensible pipeline for the reduction and calibration of MKID data. It takes binary per-pixel time-series of photon-induced phase shifts as its input and can perform cosmic ray rejection, linearity calibration, wavelength calibration, flat-fielding, bad pixel masking, and astrometric calibration. This results in calibrated spectro-temporal FITS cubes which can be integrated with traditional astronomical tool chains for scientific analysis. Additionally, unique MKID specific data products, such as time tagged photon lists, can be easily accessed and manipulated for the use and development of new post-processing techniques that utilize photon arrival time statistics – see Chapter 4.

The pipeline is designed with automation in mind to allow users to run basic reductions from the command line with unique reductions requiring only the editing of a few configuration files. It also allows future developers to add new algorithms and calibration

steps in a modular framework to serve as a base for future MKID instruments and mixed instrument reductions.

Chapter 3

First Science Results with MEC

In this chapter, we demonstrate the use of a post-processing technique called Stochastic Speckle Discrimination (SSD; Fitzgerald and Graham, 2006; Gladysz and Christou, 2008b; Meeker et al., 2018) for detecting new low mass companions using SCEXAO and the Microwave Kinetic Inductance Detector (MKID) Exoplanet Camera (MEC; Walter et al., 2020). SSD works by utilizing the timing resolution of MKID detectors to break up an observation into a series of short exposures in post-processing. These short exposure images are then used to generate intensity histograms for each pixel in an image. If the time binning is short enough, we can adequately sample the underlying probability density function (PDF) that describes the off-axis intensity in the image (light from a speckle) which can be written analytically as a modified Rician distribution. Fitting this distribution to the intensity histograms then allows us to diagnose whether a bright point in an image is a quasi-static speckle, or a true companion, see Section 3.2.1.

We also report the discovery of a low mass stellar companion to HIP 109427 using SSD with SCEXAO/MEC as well as SCEXAO/MEC photometry, SCEXAO/CHARIS spectroscopy, and Keck/NIRC2 photometry. This companion has a best fit dynamical mass of $\sim 0.25 M_{\odot}$ consistent with a spectral type of M4–M5.5 from spectral analysis.

Table 3.1. HIP 109427 Observing Log

UT Date	Instrument	Seeing (")	Passband	t_{exp} (s)	N_{exp}	ΔPA ($^{\circ}$)	Post-Processing
New Data							
20200731	SCExAO/CHARIS	0.6	<i>JHK</i>	10.32	43	5.4	RDI/KLIP
20201007	SCExAO/MEC	0.35	<i>YJ</i>	25.0	36	2.3	SSD
–	SCExAO/CHARIS	0.35	<i>H</i>	16.2-20.7 ^a	78	5.4	none
20201225	Keck/NIRC2+PyWFS	0.7	<i>Lp</i>	22.5	49	3.5	RDI/KLIP
Archival Data							
20151028	Keck/NIRC2	0.7	<i>Lp</i>	25	25	11.6	RDI/ALOC1

Note. — a) Total integration time is 1524 s.

This discovery serves as an important proof-of-concept for the use of time-domain information in addition to standard PSF subtraction methods exploiting spectral and spatial information to remove quasi-static speckles in high-contrast images.

3.1 System Properties and Observations

HIP 109427 (tet Peg) is a nearby ($d = 28.3$ pc) λ Boo star with a spectral type of A1V (Gray et al., 2006; van Leeuwen, 2007). David and Hillenbrand (2015) and Stone et al. (2018) derive system ages of $t \sim 400$ – 700 Myr; Banyan- Σ does not reveal evidence that the star’s kinematics are consistent with younger moving groups (Gagné et al., 2018). While HIP 109427 lacks a published detected radial-velocity trend indicative of a companion (Lagrange et al., 2009; Howard and Fulton, 2016), Makarov and Kaplan (2005) suggest evidence for a potential companion at a 5.7σ level from Hipparcos astrometry. Previous direct imaging observations taken as a part of the thermal infrared LEECH survey conducted with the Large Binocular Telescope failed to image any companions (Stone et al., 2018). Searches through public archives show that the star has not been targeted as a part of the Gemini Planet Imager campaign planet search, but it has been observed with VLT/NaCo and SPHERE without a reported companion.

Astrometry derived from the *Hipparcos-Gaia Catalogue of Accelerations* (HGCA; Brandt, 2018) reveals a substantial deviation from simple linear kinematic motion ($\chi^2 = 108.83$) consistent with a $\sim 11\text{-}\sigma$ -significant acceleration. We therefore targeted this star as a part of our survey to discover low-mass companions to accelerating stars (e.g. Currie et al., 2020a).

In three epochs between July and December 2020, we observed HIP 109427 with the Subaru Telescope using SCEXAO coupled to CHARIS and MEC and with the Keck II telescope using the NIRC2 camera. (Jovanovic et al., 2015a; Groff et al., 2016; Currie et al., 2020b; Walter et al., 2020) (Table 3.1). Conditions were photometric each night with average to excellent optical seeing ($\theta_V = 0''.35 - 0''.7$).

The SCEXAO Pyramid wavefront sensor (PyWFS) ran at 2 kHz, correcting for 1080 spatial modes and delivering a diffraction-limited PSF core. MEC data (7 October 2020) covers wavelengths over the *Y* and *J* passbands (0.95 - 1.4 μm) at a spectral resolution of $\mathcal{R} \sim 3.3$. We obtained CHARIS data in broadband (1.1–2.4 μm ; 31 July 2020) at a resolution of $\mathcal{R} \sim 18$ or in *H* band at a higher resolution ($\mathcal{R} \sim 70$).

The Keck near-IR PyWFS (Bond et al., 2020) corrected the wavefront at 1 kHz, correcting for 349 spatial modes and NIRC2 data (25 December 2020) was taken in the L_p broadband filter ($\lambda_o = 3.78 \mu\text{m}$).

All observations were conducted in “vertical angle”/pupil tracking mode, enabling angular differential imaging (ADI; Marois et al., 2006). The CHARIS data also enables spectral differential imaging (SDI; Marois et al., 2000). CHARIS and MEC data utilized the Lyot coronagraph ($\sim 0''.23$ diameter) to suppress the stellar halo, as well as satellite spots for precise astrometric and spectrophotometric calibration (e.g. Jovanovic et al., 2015b; Currie et al., 2018a; Sahoo et al., 2020). NIRC2 exposures left the HIP 109427 primary unocculted and unsaturated. Parallax angle rotation for all datasets was small to negligible; however, we obtained reference star observations for the CHARIS

broadband and NIRC2 data (HIP 105819 and HIP 112029, respectively).

To complement these new data, we analyzed Keck/NIRC2 L_p data for HIP 109427 taken on 28 October 2015 from the Keck Observatory Archive (Program ID C197NI). These data were obtained with Keck II’s facility (Shack-Hartmann) adaptive optics system and the vector vortex coronagraph (Serabyn et al., 2017). We used HD 212061, observed immediately after HIP 109427, for reference star subtraction.

3.2 Data

3.2.1 Image Processing: MEC

Basic Processing

MEC data was reduced using the MKID Data Reduction and Analysis Pipeline¹ (Steiger et al. (2022) – see also Chapter 2). As with the CHARIS data, satellite spots were used for the spectrophotometric calibration reference. We adopted the scaling between modulation amplitude and contrast from Currie et al. (2018b) to generate the expected satellite spot flux values per passband. A stellar spectrum from the PHOENIX stellar library appropriate for an A1V star was used and the data normalized to match HIP 109427’s reported J band flux (Ducati, 2002). Given MEC’s low energy resolution, we focused on broadband MEC photometry (not spectra). Additionally, due to the wavelength scaling of the spots, the satellite spots are extended out into elongated streaks instead of appearing as copies of an unocculted stellar PSF. This is similar to the case for GPI’s polarimetry mode.

To derive photometry for the satellite spots, we therefore follow similar methods to those outlined for GPI’s polarimetry mode from Millar-Blanchaer et al. (2016). Briefly,

¹GitHub: <https://github.com/MazinLab/MKIDPipeline>




Figure 3.1 Left: Total intensity image of HIP 109427 B taken with SCEXAO/MEC at Y and J band where the location of the companion has been circled in red. Right: SSD I_C/I_S image of HIP 109427 B. Here the companion is plainly visible as well as dark regions at the edge of the coronagraph showing the removal of pinned speckles from the total intensity image.

we subtract off a plane fitted background from a region surrounding each of the four satellite spots. We then use a “racetrack aperture” to extract satellite photometry, where the aperture radius (width perpendicular to the line connecting the spot and the star) equals that for the diffraction limit at the center wavelength for each wavelength bin (i.e. for Y or J band). The aperture radial elongation is determined empirically using the start and stop wavelengths of the bin. Photometric errors consider the intrinsic SNR of the detection, the SNR of the satellite spots, and flat-fielding errors.

Stochastic Speckle Discrimination (SSD) Analysis

Stochastic Speckle Discrimination (SSD) is a post-processing technique first demonstrated by Gladysz and Christou (2008b) that relies solely on photon arrival time statistics

to distinguish between speckles and faint companions in coronagraphic images.

Recall that the underlying probability density function that estimates the intensity distribution of off-axis stellar speckles in the image plane can be given by a modified Rician (MR)

$$p_{MR}(I) = \frac{1}{I_S} \exp\left(-\frac{I + I_C}{I_S}\right) I_0\left(\frac{2\sqrt{II_C}}{I_S}\right) \quad (3.1)$$

where $I_0(x)$ denotes the zero-order modified Bessel function of the first kind, I_C describes the coherent intensity component attributed to the unaberrated PSF of the primary, and I_S is the time variable component of the total intensity that describes the speckle field.

For a sequence of exposures shorter than the decorrelation time of atmospheric speckles (~ 10 ms), a histogram of the image plane intensity follows a MR: I_C and I_S determined for each pixel in an image (Fitzgerald and Graham, 2006). Because MEC stores the arrival time information of every photon, all time binning can be done in post-processing, which is important since the bin size that ideally samples the MR distribution is difficult to determine a priori and may vary across the image.

In order to ideally sample the MR, a bin size should be chosen that is shorter than the decorrelation timescale of the speckles in the image. If too large of a bin size is chosen, many realizations of the speckle intensity will be averaged over. Conversely, if too small of a bin size is selected, then not enough photons will arrive per bin and the distribution will tend towards Poisson statistics.

While the individual components of the MR distribution themselves do not inherently describe the signal from a faint companion, the *ratio* of the coherent component to time variable component, I_C/I_S , may reveal faint companions from a comparably bright speckle field (Gladysz and Christou, 2009; Meeker et al., 2018). This is because the

addition of light from a companion will make the MR more negatively skewed or, analogously, increase the best fit I_C . This results in a larger I_C/I_S ratio at the location of the companion compared to other pixels at the same angular separation from the primary.

We wrote an SSD analysis code to interface with the MKID Pipeline, which breaks up a MEC observation into a series of short-exposure images. Given a user-defined bin size, we then fit a MR distribution to the histogram of the intensities for each pixel using a maximum likelihood approach. Detector dithers mitigated the large number of dead pixels in the current (engineering grade) MEC array. The SSD code is run on a single dither position at a time, and the resulting I_C and I_S images are drizzled together into a combined image using an adaptation of the STScI DrizzlePac software package (Gonzaga et al., 2012).

We used this SSD code to process our 15 minute observation of HIP 109427 taken on 7 October 2020 to generate the image in Figure 3.1. The companion is clearly visible. Dark circular regions close to the edge of the coronagraph represent pinned speckles that have been suppressed by SSD due to their large I_S component.

For this analysis, a conservative bin size of 10 ms was chosen. Macintosh et al. (2005) found that speckles evolve on timescales similar to the aperture clearing time of the telescope which is given by $\tau_0 = (0.6 * D)/\bar{v}$. Here, D is the diameter of the telescope and \bar{v} is the mean wind speed for the observation. During the MEC observations of HIP 109427 B, we had quite slow wind speeds of ~ 5 m/s which, combined with a telescope diameter of 8.2 m for Subaru, yields a τ_0 of ~ 1 s. 10 ms is therefore a conservative choice since we are unlikely to be sampling over more than one realization of the speckle intensity while still having enough photons per bin to not become Poissonian.

To quantify the power of this technique, we calculated the SNR by performing aperture photometry on the companion and at a series of sky apertures located in a ring at the same angular separation from the host star. These apertures all had a diameter equal




Figure 3.2 Detections of HIP 109427 B from SCE_xAO/CHARIS in broadband (JHK) and H band and Keck/NIRC2 in L_p . For the CHARIS broadband data (NIRC2 L_p data), we retained 5 (3) KL modes for PSF subtraction but obtain similar results for other settings.

to the diffraction limit at the center of the MEC bandwidth. Since the satellite spots are at a sufficient distance away from the close-in companion, all apertures were able to be used. The noise was calculated by taking the standard deviation of the sums of the sky subtracted flux for each non-companion containing aperture (see also Currie et al., 2011; Mawet et al., 2014). This procedure was performed for both the total intensity and SSD I_C/I_S image of HIP 109427 B. The SNR of the I_C/I_S image is 21.2, about a factor of 3 higher than the SNR of 6.9 found for the total intensity image.

3.2.2 Image Processing: CHARIS and NIRC2

We extracted CHARIS data cubes from the raw data using the standard CHARIS pipeline (Brandt et al., 2017) to perform basic reduction steps – sky subtraction, image registration, and spectrophotometric calibration. For spectrophotometric calibration, we adopted a Kurucz stellar atmosphere model appropriate for an A1V star. For NIRC2 data, a well-tested general purpose high-contrast ADI broadband imaging pipeline (Currie et al., 2011) performed basic processing. To subtract the PSF for CHARIS broadband data and December 2020 NIRC2 L_p data, we used a full-frame implementation of reference star differential imaging (RDI) using the *Karhunen-Loe'Ve Image Projection* (KLIP;

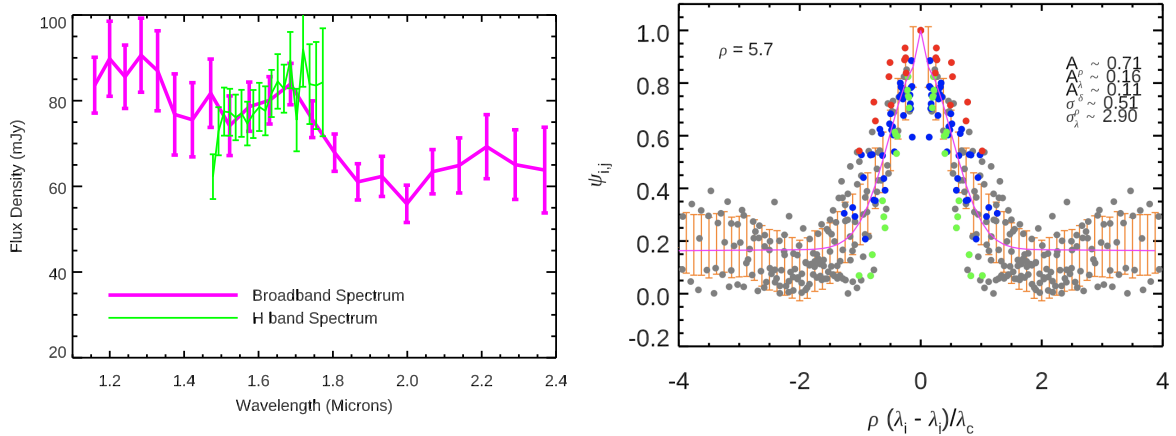


Figure 3.3 Left: SCEXAO/CHARIS spectra for HIP 109427 B extracted from broadband data (magenta) and in H band (green). Right: Spectral covariance for the CHARIS broadband data. The magenta line shows our fit to the spectral covariance as a function of scaled separation $-\rho(\lambda_i - \lambda_j)/\lambda_c$ – where ρ is the separation in λ/D units for the central wavelength λ_c (see Greco and Brandt, 2016). Blue, red, and green circles denote individual measurements between channels within the same major near-IR filter (J , H , or K_s) while grey circles denote other individual measurements. Orange points with error bars denote binned averages with 68% confidence intervals. The broadband and H band SCEXAO/CHARIS data are available in the machine-readable format as data behind the Figure.

Soummer et al., 2012) algorithm as in Currie et al. (2019), although results obtained with A-LOCI were similar (Currie et al., 2012, 2015). For the 2015 NIRC2 data, we used a full-frame version of A-LOCI.

Figure 3.2 shows detections of HIP 109427 B in each 2020 dataset. The SNRs of HIP 109427 B in the CHARIS wavelength-collapsed broadband and H band images and 2020 NIRC2 image are ~ 19 , 15, and 12, respectively. HIP 109427 B is easily visible in each CHARIS channel. We failed to obtain a decisive detection of HIP 109427 B in the 2015 NIRC2 data. No other companions are seen in the field-of-view for any dataset.

3.3 Analysis

3.3.1 HIP 109427 B Spectroscopy and Photometry

For the CHARIS broadband data, we corrected for algorithm signal loss induced by KLIP using forward-modeling as described in Pueyo (2016). Because we subtracted the PSF using a reference star, only oversubtraction (not self-subtraction terms) attenuates the companion signal flux and throughput is high ($\sim 95\text{--}97\%$). No throughput correction is applied for the H band data since we simply subtracted a median radial profile in each channel. The longest wavelength channel for the H band spectrum was deemed unreliable due to extremely poor throughput and a large dispersion (a factor of 3) in the satellite spot flux densities used to map between counts and physical units (mJy).

Figure 3.3 (top panel) shows the CHARIS spectrum. The broadband and H band flux densities agree to within $1\text{-}\sigma$ except at $\sim 1.45\ \mu\text{m}$, where telluric absorption is strongest. The CHARIS spectra show clear local minima at $1.4\ \mu\text{m}$ and $1.8\text{--}2.0\ \mu\text{m}$, consistent with absorption from water opacity (e.g. Currie et al., 2020a). In the standard Mauna Kea Observatory bandpasses, HIP 109427 B photometry drawn from the CHARIS broadband spectrum and NIRC2 imaging data is $J = 10.62 \pm 0.10$, $H = 10.30 \pm 0.07$, $K_s = 10.02 \pm 0.11$, and $L_p = 9.58 \pm 0.13$. The MEC Y and J band photometry is consistent with CHARIS-driven values: $Y = 10.73 \pm 0.24$ and $J = 10.67 \pm 0.23$.

3.3.2 HIP 109427 B Spectral Type, Temperature, and Luminosity

Following recent work (Currie et al., 2020a), we compared the CHARIS spectra for HIP 109427 B to entries in the Montreal Spectral Library² (e.g. Gagné et al., 2015),

²<https://jgagneastro.com/the-montreal-spectral-library/>

Table 3.2. HIP 109427 B Detection Significance and Photometry

UT Date	Instrument	Passband	SNR ^a	Photometry
20200731	SCExAO/CHARIS	<i>JHK</i>	19	$J = 10.62 \pm 0.10$, $H = 10.31 \pm 0.08$, $K_s = 10.02 \pm 0.10$
20201007	SCExAO/MEC	<i>YJ</i>	7.0, 21.4 ^b	$Y = 10.73 \pm 0.23$, $J = 10.67 \pm 0.24$
–	SCExAO/CHARIS	<i>H</i>	15	$H = 10.28 \pm 0.09$
20201225	Keck/NIRC2	<i>L_p</i>	12	$L_p = 9.58 \pm 0.13$

Note. — a) All HIP 109427 B SNR estimates were drawn from reductions used to calculate astrometry. b) The higher SNR SSD image can be used to determine MEC astrometry only: MEC photometry is performed using the simple sequence-combined image without post-processing (SNR = 7.0).

Table 3.3. HIP 109427 B Astrometry

UT Date	Instrument	[E,N]
20200731	SCExAO/CHARIS	$[0.229, 0.100] \pm [0.004, 0.004]$
20201007	SCExAO/MEC	$[0.228, 0.092] \pm [0.010, 0.010]$
–	SCExAO/CHARIS	$[0.229, 0.086] \pm [0.004, 0.004]$
20201225	Keck/NIRC2	$[0.222, 0.077] \pm [0.003, 0.003]$

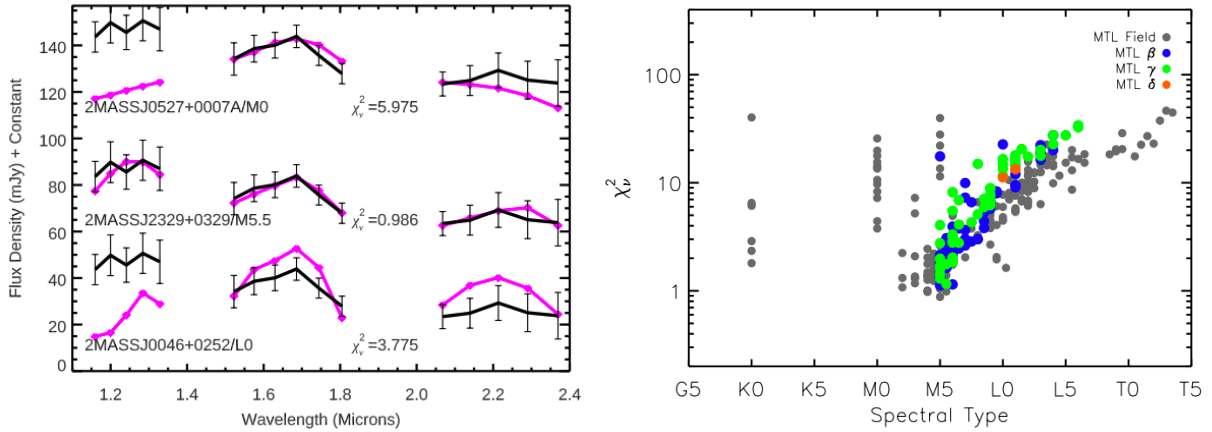


Figure 3.4 Left: The CHARIS HIP 109427 B spectrum (black) compared to field brown dwarf spectra (magenta) with M0, M5.5, and L0 spectral types from the Montreal Spectral Library (binned to CHARIS’s resolution). Right: The χ^2_ν distribution comparing HIP 109427 B’s spectrum to objects in the Montreal Spectral Library. The blue, green, and orange symbols refer to dwarfs with gravity classifications of intermediate, low, and very low which tend to be intermediate aged (e.g. ~ 100 Myr), young (10–100 Myr), and very young (< 10 Myr), whereas field (older) dwarfs are shown as grey circles (see Currie et al., 2018a).

considering the impact of spatially and spectrally correlated noise (Greco and Brandt, 2016)³. The CHARIS data reveal highly correlated errors (Figure 3.3, right panel). The spectral covariance at HD 109427 B’s location includes substantial off-diagonal terms, especially for spatially-correlated noise ($A_\rho \sim 0.71$) and (to a lesser extent) residuals speckles well correlated as a function of wavelength ($A_\lambda \sim 0.16$).

As shown in Figure 3.4, HIP 109427 B’s CHARIS spectrum is best matched by M4–M5.5 field objects (left panel). Three objects in the Montreal library yield $\chi^2_\nu \leq 1$, even with the full spectral covariance included: 2MASSJ0326-0617 (M5), 2MASSJ0854-3051 (M4), and 2MASSJ2329+032 (M5.5). Using the mapping between spectral type and effective temperature from Pecaut and Mamajek (2013), empirical comparisons to the CHARIS spectra then favor a temperature of 3000-3200 K for HIP 109427 B. Adopting

³We do not also compare the MEC or NIRC2 photometry due to sparse coverage of the library outside of the *JHK* passbands

the relationship from Casagrande et al. (2008) and assuming a distance of 28.3 *pc*, HIP 109427 B’s luminosity is $\log_{10}(L/L_{\odot}) = -2.28^{+0.04}_{-0.04}$.

We compared the MEC *YJ* band photometry, CHARIS *JHK* spectra, and NIRC2 L_p photometry to the BT-Settl atmosphere models (Allard et al., 2012) with the Asplund et al. (2009) abundances downloaded from the Theoretical Spectra Web Server⁴. The grid covers temperatures of 2500–4000 *K*, surface gravities of $\log(g) = 3.5$ –5.5, and metallicities of $[\text{Fe}/\text{H}] = -1$ to 0.5. Following Currie et al. (2018b), we focus only on the CHARIS channels unaffected by telluric absorption, resulting in 21 photometric/spectrophotometric points fit. We define the fit quality for the *k*th model using the χ^2 statistic, considering the spectral covariance:

$$\chi^2 = R_k^T C^{-1} R_k + \sum_i (f_{phot,i} - \alpha_k F_{phot,ik})^2 / \sigma_{phot,i}^2. \quad (3.2)$$

Here, the vector R_k is the difference between measured and predicted CHARIS data points ($f_{spec} - \alpha_k F_{spec}$) and C is the covariance for the CHARIS spectra. The vectors $f_{phot,i}$, $F_{phot,ik}$, and $\sigma_{phot,i}$ are measured photometry, model predicted photometry, and photometric uncertainty; α_k is the scaling factor for the model that minimizes χ^2 (see also De Rosa et al., 2016).

Figure 3.5 shows the best-fit solar and non-solar metallicity models (top panels) and the associated χ^2 contours (bottom panels). An atmosphere with a temperature of $T_{\text{eff}} = 3200$ *K* and a high gravity ($\log(g) = 5.5$) fits the data the best in both cases. The 1- σ contour for temperature and gravity is narrowly defined about this peak for both metallicities: $T_{\text{eff}} = 3100$ –3300 *K* and $\log(g) = 5.25$ –5.5. At the 2- σ level, the best-fit temperature and gravity ranges widen to 3000–3400 *K* and $\log(g) = 5$ –5.5. The radii that minimize χ^2 are ~ 2.1 –2.6 Jupiter radii.

⁴<http://svo2.cab.inta-csic.es/theory/newov2/>

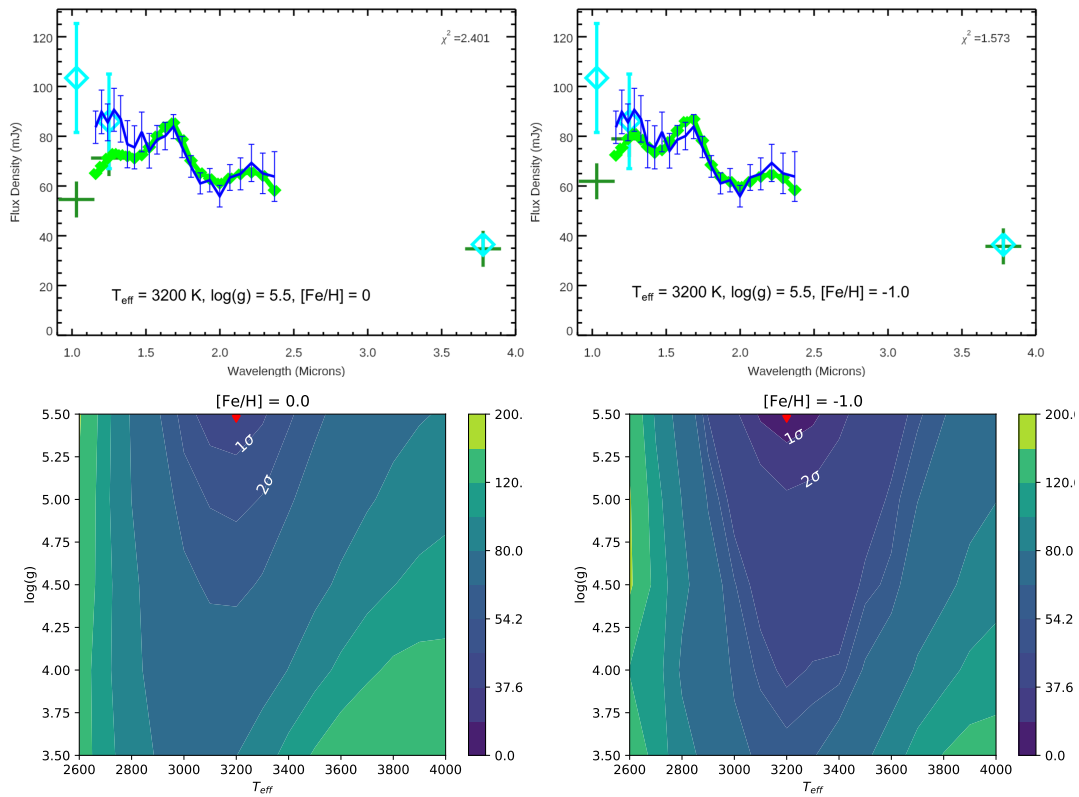


Figure 3.5 Top: Best fit BT-Settl models for a solar and non-solar metallicity. Bottom: corresponding contour plots of χ^2 as a function of temperature and surface gravity. The 1σ and 2σ contours are labeled in white and the best fit solution denoted with a red diamond. The χ^2_{ν} value shown is for 20 degrees of freedom. CHARIS spectra is shown in blue, MEC and NIRC2 photometry in cyan, model-predicted CHARIS spectrophotometry in light green, and predicted MEC/NIRC2 photometry in dark green.

The best-fit solar metallicity model accurately reproduces the H and K portions of CHARIS spectrum and the NIRC2 L_p photometry; however, it underpredicts the brightness of HIP 109427 B in Y and J band by 85% and 25%, respectively. Subsolar metallicity models systematically produce a rough match in J band and show less severe disagreement at Y band. Future MEC calibration work, such as improving the wavelength dependent flat-fielding, may yield better agreement with expected Y band photometry.

The $2\text{-}\sigma$ ranges for temperature correspond to M3–M5.5 dwarfs, a range that overlaps with the spectral types of best-matching objects in the Montreal Spectral Library, although the best-fit is skewed towards earlier, hotter objects by ~ 1 subclass. For M3–M5.5 objects with the HIP 109427 system’s estimated age of $\sim 0.4\text{--}0.7$ *Gyr*, the expected surface gravities are $\log(g) \sim 5\text{--}5.1$ (Baraffe et al., 2003), or about 0.25–0.5 dex lower than the best-fit values considered by our grid. Expected radii are 2–3 Jupiter radii: consistent with our best-fit values.

3.3.3 HIP 109427 B Astrometry and Dynamical Mass

Evidence for Common Proper Motion

To rule out the possibility that HIP 109427 B is a background object, we analyzed archival 2015 Keck/NIRC2 data shown in Figure 3.6. The data do not reveal a statistically significant detection of any signal that could be HIP 109427 B. Using the small sample statistics correction from Mawet et al. (2014), we estimate a $5\text{-}\sigma$ contrast of $\Delta L_p \sim 5, 5.75, 11.3,$ and 12 magnitudes at $0''.15, 0''.225, 1''.0,$ and $1''.5,$ respectively. Companions at HIP 109427 B’s current angular separation would be just undetectable at $5\text{-}\sigma$. Those with contrasts like HIP 109427 B near $2 \lambda/D$ would be well below the detection limit and those at arcsecond or wider separations would be easily detected.

HIP 109427 has an extremely high proper motion of $\mu_\alpha \cos(\delta), \mu_\delta \sim 282.18, 30.46$

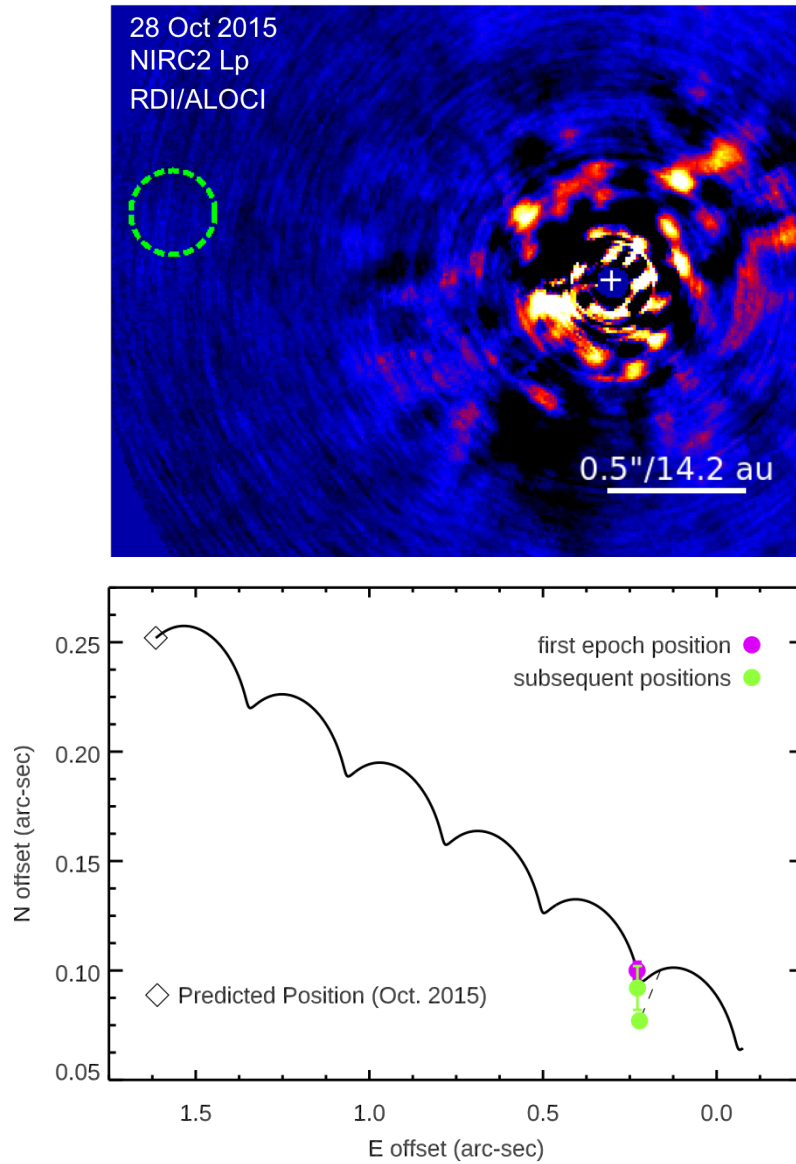


Figure 3.6 Top: Keck/NIRC2 data taken in L_p showing a non-detection at the expected location of the companion, which is circled in green. Bottom: Expected track for a background object showing its predicted location in October 2015. The dashed line connects the measured Dec 2020 position to the predicted position for a background object.

mas yr⁻¹ (van Leeuwen, 2007). If HIP 109427 B were a background star, it would appear at an angular separation of $\sim 1''.6$ in October 2015 data with an expected SNR of ~ 1000 . However, no signal is present at its expected location (dashed circle). HIP 109427 B's position in December 2020 also deviates by ~ 65 mas compared to the expected location of a background star, far larger than our astrometric precision (bottom panel). This implies HIP 109427 B is a common proper motion companion to the primary.

Orbit and Dynamical Mass

We used the open source code `orvara`, Brandt et al. (2021), to fit the mass and orbit of HIP 109427 B. `orvara` is an MCMC orbit fitting code for exoplanetary systems that uses a combination of absolute astrometry, relative astrometry, and radial velocities (RVs) to fit one or more Keplerian orbits to a system. For this dataset, we used HGCA absolute astrometry measurements for the star and the three measured epochs of relative astrometry for the companion from CHARIS and MEC. We do not consider RV limits since previous data has had a limited time baseline and poor precision. A Gaussian prior of $2.1 \pm 0.15 M_{\odot}$ was chosen for the primary in concordance with literature values derived from isochrone fitting (De Rosa et al., 2014; David and Hillenbrand, 2015; Stone et al., 2018).

Figure 3.7 shows the posterior distributions of select orbital parameters as well as the primary and secondary mass. A summary of the fit parameters can also be found in Table 3.4. The mass of the primary is nearly identical to the adopted prior with a value of $2.09^{+0.16}_{-0.16} M_{\odot}$ and the fit secondary mass is $0.280^{+0.18}_{-0.059} M_{\odot}$. The best fit eccentricity is $0.54^{+0.28}_{-0.15}$ with an inclination of $66.7^{+8.5}_{-14}$ degrees. The best fit semimajor axis is $6.55^{+3.0}_{-0.48}$ au, although the distribution is bimodal with HIP 109427 B's mass with one family of solutions favoring a ~ 6 au separation with a mass of $\sim 0.25 \pm 0.05 M_{\odot}$ and another favoring a mass of $0.5 M_{\odot}$ and semimajor axis of 9 au. Main-sequence stars with masses

Table 3.4. HIP 109427 B Orbit Fitting Results and Priors

Parameter	Fitted Value	Prior
$M_{pri} (M_{\odot})$	2.09 ± 0.16	Gaussian, 2.1 ± 0.15
$M_{sec} (M_{\odot})$	$0.280^{+0.18}_{-0.059}$	$1/M_{sec}$
Semimajor axis a (au)	$6.55^{+3.0}_{-0.48}$	$1/a$
Eccentricity e	$0.54^{+0.28}_{-0.15}$	uniform
Inclination i ($^{\circ}$)	$66.7^{+8.5}_{-14}$	$\sin(i)$

Note. — Posterior distributions for the secondary mass and semimajor axis are bimodal with a favored solution of $\sim 0.25 M_{\odot}$ and ~ 6 au - see Figure 3.7 and text for more details.

of $0.5 M_{\odot}$ have early M spectral types (e.g. Pecaut and Mamajek, 2013), which are excluded from our spectral analysis. In contrast, the lower-mass solution is consistent with M4 V object allowed by the CHARIS spectral comparisons.

A mass of $\sim 0.25 M_{\odot}$ is broadly consistent with inferred masses based on luminosity evolution models, given HIP 109427 B’s likely age. From the Baraffe et al. (2003) models, an M3–M5.5 object with an age of 400–700 Myr is predicted to have a mass of 0.15–0.3 M_{\odot} . Modeling absolute astrometry of the primary and relative astrometry of the star likely then yields much more precise (20%) constraints on the companion mass than available from luminosity evolution models alone (50%).

3.4 Conclusion

With SCEXAO/MEC photometry, SCEXAO/CHARIS spectroscopy, and Keck/NIRC2 photometry, we have identified a low mass stellar companion at a near-Jupiter-like separation around the nearby A1V star HIP 109427. Comparison of this target’s spectrum with entries in the Montreal Spectral Library indicates a spectral type of M4–M5.5. This is consistent with a best fit a dynamical mass of $\sim 0.25 M_{\odot}$ with a semimajor axis of ~ 6

au from orbital fitting using measurements from both *Hipparcos* and *Gaia* DR2 as well as MEC, CHARIS, and NIRC2 relative astrometry. There is a degeneracy in the orbital fit with another favored solution of $\sim 0.5 M_{\odot}$ with a semimajor axis of ~ 9 au that is excluded by our spectral analysis. Future RV measurements, Gaia astrometry, and relative astrometry from high-contrast imaging will help to better constrain this orbit.

This result demonstrates the efficacy of Stochastic Speckle Discrimination (SSD) in identifying faint companions. SSD increases the SNR of HIP 109427 B by about a factor of 3 versus the total intensity image (comparable to the CHARIS SNR of this target) without the use of any additional PSF subtraction techniques. This technique is especially effective at small angular separations (inside $10 \lambda/D$) where algorithms exploiting traditional observing strategies like ADI and SDI suffer.

Work expanding the SSD framework to be agnostic to bin size and to directly fit an off-axis Poisson source has been shown to be effective on simulated data and is currently being adapted for use on real datasets (Walter et al., 2019). This will be explored in-depth in the following Chapter.

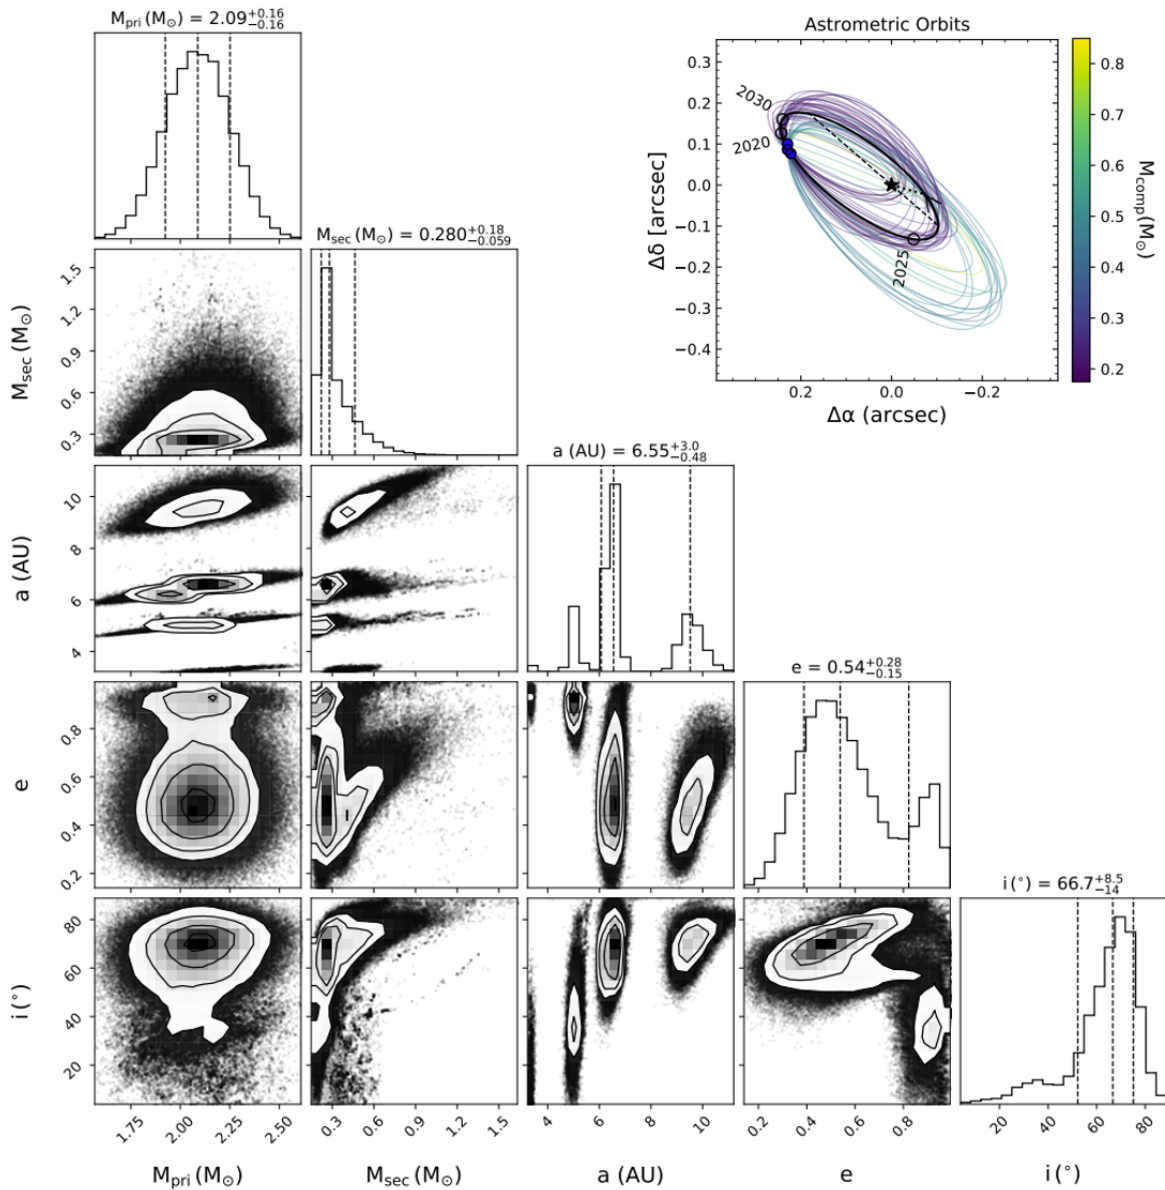


Figure 3.7 Corner plot displaying select posterior orbital parameters. The orbit fits were performed using HGCA data and relative astrometry points from SCExAO/CHARIS and MEC data. The mass of the primary is nearly identical to the chosen prior of $2.1^{+0.15}_{-0.15} M_{\odot}$. (Inset) The best fit orbit of HIP 109427 B in black with 50 randomly selected orbits from the MCMC fit color-coded by HIP 109427 B’s mass. The blue circles represent the measured relative astrometry points and the unfilled black circles are the predicted locations of the companion at different epochs. The arrow indicates that HIP 109427 B is orbiting counter-clockwise

Chapter 4

Probing Photon Statistics in Adaptive Optics Images with SCE_xAO/MEC

The work of this thesis is focused on MEC and, specifically, the development and use of software tools to be able to harness some of the unique abilities of MKID detectors to tackle fundamental problems in high-contrast imaging. In addition to containing the first SSD detection of a diffuse source (the protoplanetary disk AB Aurigae), this chapter will be focused on how others can best utilize MEC data for their own analyses, the limitations of current techniques, and a path forward to build off of current proofs-of-concept to push current speckle noise limits.

4.1 Utilizing Photon Arrival Time Information with MEC

Arrival time information can be easily accessed using The MKID Pipeline¹ which is open-source and can be downloaded via GitHub (Steiger et al., 2022). Since MKID detectors record the arrival time and energy of each incident photon, the format of raw MKID data is a time and energy-tagged photon list that can be queried using the `Photontable` class on pixel location, time range, photon wavelength, or any combination thereof. A result of this is that all spectral and temporal binning is done in post-processing and MEC has no set ‘exposure time’ for its observations. This is especially beneficial for post-processing techniques that leverage differences in arrival time statistics, like those that are discussed in the following section, as many different timescales can be probed from a single observation.

MKIDs temporal resolution limit is determined by the readout speed of the detector (~ 1 MHz) and the firmware-imposed dead time. This dead time is set by material properties of the MKID array and for MEC has a value of $10 \mu\text{s}$. During this time, no additional photons are able to be recorded for that pixel to allow it time to return to its idle state (Fruitwala et al., 2020).

Immediately after this dead time, a pile-up of photon events has been observed which is likely insignificant for total integrated observations, but can cause unintended effects when using photon arrival time information directly. Empirically this effect decays rapidly after $40 \mu\text{s}$ (Figure 4.1) and any work done with MEC that uses arrival time information should take this into account so as to not contaminate results.

¹<https://github.com/MazinLab/MKIDPipeline>




Figure 4.1 Interarrival time histogram for a single pixel. An excess of interarrival times between 10-40 μs can be clearly seen that are inconsistent with Poisson statistics.

4.2 Photon Arrival Time Based Post-Processing Techniques

4.2.1 Stochastic Speckle Discrimination (SSD)

As was explained in prior chapters, SSD is achieved with MEC by fitting a MR to every pixel in an image with a user-specified temporal bin size. While the components of the MR distribution do not inherently describe the signal from a faint companion, the *ratio* of the coherent component to time variable component - I_C/I_S - may reveal faint companions from a comparably bright speckle field (Gladysz and Christou, 2009; Meeker et al., 2018; Steiger et al., 2021). This is because the addition of light from a companion (whose statistics follow a negatively skewed Gamma distribution – Equation 1.5) will make the best-fit MR more negatively skewed at that location. This is analogous to increasing the best fit I_C and results in a larger I_C/I_S ratio at the location of the companion compared to other pixels at the same angular separation from the primary.




Figure 4.2 Left: Total Y band intensity dither combined image of AB Aurigae taken with SCEXAO/MEC. Disk features here are largely obscured. The satellite spots can be seen as radially smeared bright patches on the edge of this image due to the corrected sky rotation. Right: I_C/I_S image of the disk around AB Aurigae clearly showing some of the inner disk features. These include the two main spirals that have roots to the North and South as well as the extended point-like source to the South at $\rho = 0''.16$ as found in Boccaletti et al. (2020). A Gaussian filter has been applied over this image to smooth over small scale inter-pixel variations.

This can be used to identify companions in coronagraphic images as was shown with the discovery of HIP 109427 B (Chapter 3).

SSD on AB Aur

Some post-processing techniques, such as ADI, struggle to reveal structures with azimuthal symmetry. SSD does not suffer the same limitations for these sources because the light from the disk will still be on-axis (and thus follow a Gamma distribution) even if it is spread over a region instead of contained within a single PSF.

Here we tested the performance of SSD on extended sources using a SCEXAO/MEC Y band observation of AB Aurigae (AB Aur) taken on 14 October 2021. AB Aur is a well

studied system with a known protoplanetary disk that also harbors a recently discovered protoplanet (Boccaletti et al., 2020; Currie et al., 2022; Zhou et al., 2022). AB Aur was observed using SCEXAO/MEC for 16 minutes in exceptional seeing conditions of $\sim 0''.3$. The results from the SSD reduction of the protoplanetary disk surrounding AB Aur can be seen in Figure 4.2 where inner disk structures are revealed in the I_C/I_S map (right) not seen in the Y band total intensity image (left). This bears a strong resemblance to images taken of this system in polarized intensity by the Spectro-Polarimetric High-contrast Exoplanet REsearch instrument at the VLT (SPHERE; Beuzit et al., 2019) as described in Boccaletti et al. (2020) – see their Figure 4. In contrast with the SPHERE observations, here only the millisecond intensity distributions were used to generate these images with no polarization information or PSF subtraction techniques employed. It is important to note that this result was in large part facilitated by the exceptional seeing conditions since better seeing leads to a less intense speckle halo and allows the disk’s Gamma distributed intensity to significantly modulate the underlying MR distribution of the speckles.

4.2.2 Photon-Counting SSD (PCSSD)

PCSSD is an extension of the SSD formalism where contributions from an incoherent source of constant intensity (I_P) are accounted for in addition to I_C and I_S which define the shape of the MR. Given a list of photon inter-arrival times, the maximum likelihood value of I_C , I_S and I_P are determined. Since all inter-photon arrival times are used in this technique, no temporal binning is done and it has been shown to perform twice as well as perfect PSF subtraction on simulated data where the companions were modeled as constant, incoherent sources (Walter et al., 2019).

One of the main motivations for expanding the SSD formalism is that the SSD I_C/I_S

maps – while helpful for extracting companion astrometry and disk morphology – only quantify the skewness of the fit MR to the data. This output is therefore not easily converted to physically meaningful units and the images are unhelpful for performing spectroscopy or photometry as would typically be desired to determine key companion properties such as temperature, composition, and mass. PCSSD attempts to solve for this by calculating a likelihood for each inter-photon arrival time so that the output units can be reported in counts or counts/s. Additionally, PCSSD is able to leverage the photon counting nature of MKID detectors to beat the long exposure noise limit by not temporally binning.

PCSSD on HIP 109427 B

In the form described by Walter et al. (2019), PCSSD makes the following assumptions:

1. The speckle halo intensity is entirely described by the MR distribution.
2. I_C , I_S , and I_P (the intensity of a companion) remain constant over an observation.
3. Chromaticity is ignored.

The assumptions that are perhaps the most problematic are that the off-axis intensity is entirely described by the MR distribution and that the companion intensity (I_P) remains constant over the course of an observation. MEC has a known infrared (IR) background that can cause the intensity at any pixel to not be fully described by the MR PDF. Additionally, for realistic observing conditions the on-axis companion intensity (which is proportional to the SR) varies quite considerably and we know is described by the Gamma PDF in Equation 1.4.




Figure 4.3 Output of PCSSD on HIP 109427 B (see also Figure 3.1). The red arrows in each image point to the same location of the low mass stellar companion. Here the companion light clearly shows up in I_C with corresponding minima in the intensity in the I_P and I_S images

The fact that the on-axis light follows this negatively skewed distribution is in fact one of the main bases that allows traditional SSD to work. Companions stand out in I_C/I_S images due to the addition of the negatively skewed companion PDF which shifts the whole distribution or, analogously, increases the best-fit MR I_C (Gladysz and Christou, 2008a; Meeker et al., 2018; Steiger et al., 2021).

For this reason, using the PCSSD technique on real data does not accurately separate the companion light into I_P , but instead attributes that light largely to I_C . See Figure 4.3 for an example of PCSSD run on the same HIP 109427 B dataset as Figure 3.1.

PCSSD on More Realistic Simulated Data

To verify that a limitation for performing PCSSD on on-sky data is the assumption that the companion intensity (I_P) is constant, we generated new mock photon lists following the procedure as described in Walter et al. (2019), but with the notable exception that the companion intensities were sampled from the Gamma PDF. The exact distribu-




Figure 4.4 Gamma distribution (Equation 1.4) from which the companion intensities were sampled in §4.2.2 and Figures 4.5 & 4.6. Here $k = 5$, and $\Theta = 30$, and μ corresponds to the median strehl ratio used (0.7).

tion from which the companion intensities were sampled can be seen in Figure 4.4 – here sensible values of k , Θ , and μ were chosen to roughly match observed values. These companion photons were also assumed to be correlated in time with a decorrelation timescale (τ) of 0.1 s - the same as that of the MR. A suite of companion separations and contrasts was tested ranging from 3.5 to $12.5 \lambda/D$ and $5 \cdot 10^{-5}$ to $4 \cdot 10^{-4}$ respectively. A suite of brighter companions was also tested (with contrasts ranging from $7.5 \cdot 10^{-4}$ to $6 \cdot 10^{-3}$) to more closely replicate the on-sky PCSSD results for binary stellar companions like HIP 109427 B which has a J band contrast of $1.27 \cdot 10^{-3}$.

The results are summarized in Figure 4.5. The addition of the Gamma distributed companion flux causes more companion light to be misattributed to I_C over I_P in the case of bright companions ($< 1 \cdot 10^{-3}$ contrast - see right most column) matching observations.

Interestingly, in the case of higher contrast sources (left column) the shape of the companion distribution doesn't seem to make much of a difference in the performance of the PCSSD. We believe that in this regime there are not enough companion photons to

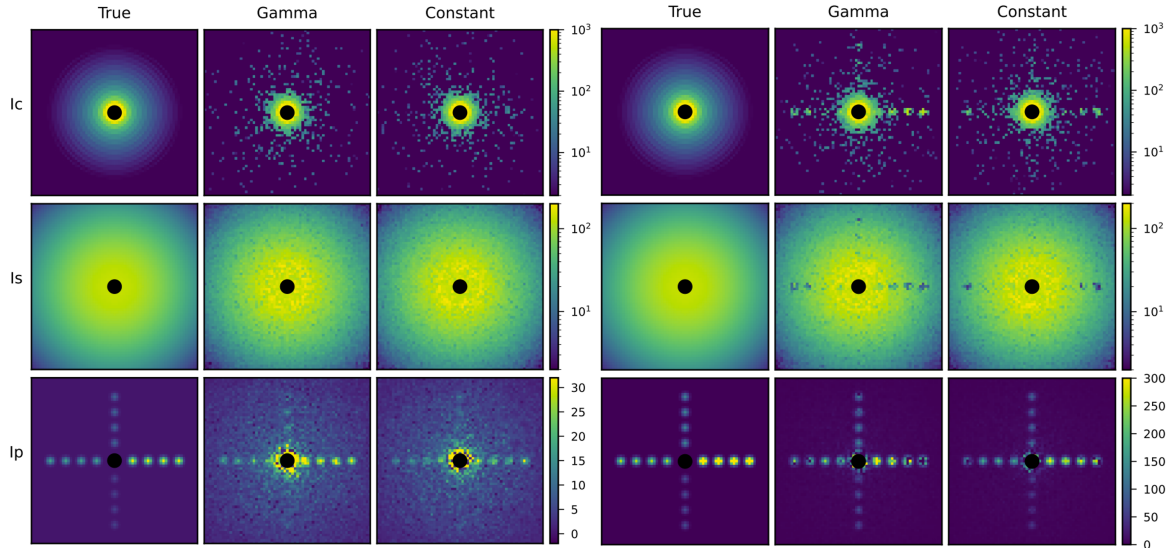


Figure 4.5 Simulated I_C , I_S and I_P images generated by running PCSSD on constant and Gamma distributed companion intensities for two different suites of companion contrasts. Left: $4 \cdot 10^{-4}$, $2 \cdot 10^{-4}$, $1 \cdot 10^{-4}$, $5 \cdot 10^{-5}$ Right: $6 \cdot 10^{-3}$, $3 \cdot 10^{-3}$, $1.5 \cdot 10^{-3}$, $7.5 \cdot 10^{-4}$. The ‘True’ columns are the input I_C , I_S , and I_P images.

shift the skewness of the underlying MR distribution at the location of companion pixels and so the light does not incorrectly end up in I_C . In other words, at these low count rates it becomes hard to distinguish between the Poisson and Gamma distributions.

At these higher contrasts however, other factors such as background noise sources are likely to become more significant. The effect of background counts on the performance of the PCSSD is summarized in Figure 4.6. Here the PCSSD was run on mock photon lists using the suite of higher contrast companions (see Figure 4.5, left panels) with an added constant and uniform 50 photons/s/pixel background with uncorrelated Poisson distributed arrival times. This count rate approximately matches the current IR background count rate for MEC and significantly degrades PCSSD performance (center column).

We also wanted to explore the effect of performing a wavelength cut on MEC data

before running PCSSD. Since MEC stores the arrival time and wavelength of each incident photon, a wavelength range can be specified and only the arrival times of photons with wavelengths within that range returned. This is typically done in part to remove background counts since MEC’s thermal background lies outside of our science bands ($Y - J$). While these out-of-band photons themselves can be removed from analysis, their effect on temporally proximate photons still remains. Even if it is not used, an out-of-band photon incident on the detector still activates the $10 \mu\text{s}$ dead time of that pixel meaning that it is no longer photosensitive over that time range in a way that is unaccounted for by the PCSSD code. For the typical background count rates observed with MEC however this effect doesn’t appear significant enough to impact PCSSD performance. This can be seen in the right column of Figure 4.6 as the background removed I_P images are consistent with the case of no thermal background present (left column).

For all of these tests, the resulting I_P signal-to-noise ratios (SNR) are summarized in Table 4.2.2. SNR was calculated using 10^3 independently generated 30s photon lists and is given by $\text{SNR} = (\langle I_P \rangle - \langle \text{Background} \rangle) / (\text{std.dev.} \langle I_P \rangle)$. Here the ‘Background’ is determined using photon lists without any injected companions and ‘std.dev. $\langle I_P \rangle$ ’ is the standard deviation of the mean companion intensity.

In summary, A constant 50 photons/s/pixel background significantly degrades the I_P SNR of the faint companions at all separations when not removed. The perfect removal of these photons in post-processing recovers the results as if there had never been a background present, but this type of perfect subtraction can be challenging with real data due to the current energy resolution of MEC ($\mathcal{R} \sim 5$). An exploration of how imperfect background subtraction effects PCSSD performance will be left for future work.




Figure 4.6 Simulated I_C , I_S and I_P images generated by running PCSSD on mock photonlists that sampled a Gamma PDF for the companion intensities. Left: No background included (same as the second column of Figure 4.5, left). Middle: Adding an additional constant and uncorrelated 50 photons/s/pixel background. Right: The 50 photons/s/pixel background is inserted, but then removed before running PCSSD. Any photons that were excluded from the analysis due to their proximity to a background photon (i.e. falling within the $10 \mu\text{s}$ deadtime) are not accounted for which simulates performing a wavelength cut on MEC data in post-processing.

Table 4.1. Figure 4.6 Companion Signal-to-Noise

Contrast	Separation=3.5 λ/D			Separation=6.5 λ/D		
	Control	50 cps	Removed	Control	50 cps	Removed
$4 \cdot 10^{-4}$	4.5	3.4	4.2	6.3	4.4	6.8
$2 \cdot 10^{-4}$	2.6	1.9	2.6	3.5	2.6	3.6
$1 \cdot 10^{-4}$	1.6	1.0	1.7	2.2	1.4	2.2
$5 \cdot 10^{-5}$	0.8	0.6	0.8	1.2	0.7	1.4

Table 4.2. Figure 4.6 Companion Signal-to-Noise Continued

Contrast	Separation=9.5 λ/D			Separation=12.5 λ/D		
	Control	50 cps	Removed	Control	50 cps	Removed
$4 \cdot 10^{-4}$	7.4	3.0	7.0	5.8	1.3	6.4
$2 \cdot 10^{-4}$	4.7	2.1	4.8	5.5	1.1	5.4
$1 \cdot 10^{-4}$	2.6	1.5	2.7	3.4	0.8	3.6
$5 \cdot 10^{-5}$	1.6	0.8	1.6	1.9	0.3	1.9

4.3 Discussion

To use photon arrival time statistics as a metric for differentiating photons from the bright speckle halo from those that originate from a faint companion, understanding and incorporating information about the distributions from which they originate is essential. This is especially true in regimes where the companion intensity is comparable to that of the underlying speckle field and deconvolving those two distributions becomes more important.

Additionally, identifying and removing any possible sources of light outside of these distributions, or incorporating them explicitly into the model being used, is important in very high contrast regimes where background photons can easily outnumber photons coming from the faint source of interest.

Thankfully hardware fixes can remove out-of-band background light. Future work will identify and mitigate MEC's known IR background at wavelengths longer than its science band ($Y - J$). This will significantly reduce the number of background counts and should aid the performance of even the current PCSSD algorithm on faint sources.

PCSSD has been derived only for a linear combination of a MR intensity distribution and a constant intensity component (Walter et al., 2019). Section 4.2.2 shows that a companion whose intensity instead follows a Gamma distribution can significantly weaken the sensitivity of PCSSD. PCSSD could be expanded within its current analytical framework to account for a Gamma-distributed intensity component. Doing so could require fitting an additional three parameters, rather than one for a constant-intensity companion, unless some of the Gamma distribution's parameters are independently known (e.g. from AO telemetry). This could make performing the fit untenable for even typical MKID datasets that can easily run in excess of 10^9 photons.

Non-analytic approaches can be computationally less expensive than PCSSD and

assume no knowledge about the underlying intensity distributions except for the fact that they are measurably different. This work, and the works cited herein, have already demonstrated that intensity distributions differ for speckles and incoherent on-axis sources. The most prevalent non-analytic SSD techniques overlap with machine learning based approaches for detecting patterns in time series data such as Long Short-Term Memory recurrent neural networks (LSTM; Hochreiter and Schmidhuber, 1997). To facilitate further development of machine-learning-based SSD, the MKID Exoplanet Direct Imaging Simulator (MEDIS) can be used to generate more realistic simulated MKID datasets. (Dodkins et al., 2020).

In high-contrast imaging there is unique information available at millisecond and microsecond timescales that is lost when taking long exposure images. New instruments deploying fast, noiseless detectors (like MEC) are now able to access this information and use it to start suppressing/differentiating quasi-static speckles and companions. Though not discussed here, work using millisecond images (sometimes combined with wavefront sensor telemetry) to directly measure and remove quasi-static speckles on-sky is another promising path forward. However, these techniques have additional challenges needing to run in conjunction with the AO loop during observing instead of in post-processing (Martinache et al., 2014; Gerard et al., 2018; Rodack et al., 2021). Moving forward, the advancement of both real-time and post-processing techniques that leverage the information available at these fast timescales will be necessary for reaching the best achievable contrasts with current and next generation telescopes.

Appendix A

Transition Region Between Gamma and MR Statistics

The complex amplitude in the focal plane is the Fourier transform of the pupil plane complex amplitude:

$$\Psi(\vec{r}) = \int P(\vec{x}) e^{i\phi(\vec{x})} e^{-2i\pi\vec{x}\cdot\vec{r}} dx \quad (\text{A.1})$$

Where \vec{x} is the pupil plane coordinate vector ($\vec{x} \in [0, D]$, where D is the telescope diameter), \vec{r} is the focal plane coordinate vector, $P(\vec{x})$ is the pupil function, and $\phi(\vec{x})$ is the phase of the wavefront in the pupil plane. At the center of an image ($r = 0$), this reduces to

$$\Psi(0) = \int P(x) e^{i\phi(x)} dx \quad (\text{A.2})$$

In the case of a high adaptive optics (AO) correction, $\phi(x)$ will be small and so $e^{i\phi(x)}$ can be written as

$$e^{i\phi(x)} = \cos(\phi(x)) + i \sin(\phi(x)) \quad (\text{A.3})$$

$$e^{i\phi(x)} \approx 1 + i\phi(x) \quad (\text{A.4})$$

Here it can be seen that at the center of the image, the phasors, $\phi(\vec{x})$, are not oriented randomly in the complex plane and are instead constrained to a single real value for all \vec{x} . The statistics will therefore not be circular Gaussian and the resulting PDF of the intensity will not follow a modified Rician distribution (Equation A.5, see Soummer et al. (2007b); Goodman (1975)).

$$p_{MR}(I) = \frac{1}{I_S} \exp\left(-\frac{I + I_C}{I_S}\right) I_0\left(\frac{2\sqrt{II_C}}{I_S}\right) \quad (\text{A.5})$$

It will instead follow a Gamma distribution as given below. The full derivation has not been copied here, but can be found in Gladysz and Christou (2008b).

$$p_{SR}(sr) = \frac{p_{\hat{\sigma}^2}(-\ln(sr))}{sr} \quad (\text{A.6})$$

where sr is the instantaneous SR, and $p_{\hat{\sigma}^2}$ is given by

$$p(x; k, \theta, \mu) = \frac{\left(\frac{x-\mu}{\theta}\right)^{k-1} \exp\left(-\frac{x-\mu}{\theta}\right)}{\Gamma(k) \theta} \quad (\text{A.7})$$

We now know that the statistics which govern the intensities at the center of an image will follow a Gamma distribution and that sufficiently far away (at large \vec{r}), or for a sufficiently small AO correction ($\phi(\vec{x})$ large), they will follow a modified Rician. The exact point at which this transition occurs will be explored now.

For $\vec{r} \neq 0$, the Fourier component of Equation A.1 does not disappear and will have

the effect of rotating the phase vectors $\phi(\vec{x})$. Combining the two exponential terms from Equation A.1, we get

$$e^{i\phi(\vec{x})-2i\pi\vec{x}\cdot\vec{r}} \quad (\text{A.8})$$

which can be rewritten as

$$e^{i(\phi(\vec{x})-2\pi\vec{x}\cdot\vec{r})} = \cos(\phi(\vec{x}) - 2\pi\vec{x}\cdot\vec{r}) + i \sin(\phi(\vec{x}) - 2\pi\vec{x}\cdot\vec{r}) \quad (\text{A.9})$$

to fully circularize $\phi(\vec{x})$, we need the term $2\pi\vec{x}\cdot\vec{r}$ over the range of \vec{x} to go from $[0, 2\pi]$.

Plugging in $\vec{x} = 0$, we get a shift of 0. For $x = D$ we therefore need $2\pi\vec{x}\cdot\vec{r} = 2\pi$. It is important to note that up until now we have only been considering the monochromatic case. Inserting wavelength dependence we also know that $\vec{x} \propto \frac{1}{\lambda}$. Combining all of this, we can derive the value of \vec{r} needed to fully circularize $\phi(\vec{x})$:

$$2\pi\vec{x}\cdot\vec{r} = 2\pi \quad (\text{A.10})$$

$$\frac{D\cdot\vec{r}}{\lambda} \propto 1 \quad (\text{A.11})$$

$$\vec{r} \propto \frac{\lambda}{D} \quad (\text{A.12})$$

So for $\vec{r} > \lambda/D$ the Fourier term in Equation A.1 will completely circularize the phase vectors even at high correction levels and the resulting intensities will follow a modified Rician. For $0 < r < \frac{\lambda}{D}$ the vectors will be partially, but not completely, circularized due to the effect of the Fourier term in Equation A.1

Bibliography

- C. Marois, B. Macintosh, T. Barman, et al. Direct Imaging of Multiple Planets Orbiting the Star HR 8799. *Science*, 322(5906):1348, November 2008a. doi: 10.1126/science.1166585.
- A. M. Lagrange, M. Desort, F. Galland, et al. Extrasolar planets and brown dwarfs around A-F type stars. VI. High precision RV survey of early type dwarfs with HARPS. *A&A*, 495(1):335–352, February 2009. doi: 10.1051/0004-6361:200810105.
- J. Rameau, G. Chauvin, A. M. Lagrange, et al. Confirmation of the Planet around HD 95086 by Direct Imaging. *ApJ*, 779(2):L26, December 2013. doi: 10.1088/2041-8205/779/2/L26.
- M. Kuzuhara, M. Tamura, T. Kudo, et al. Direct Imaging of a Cold Jovian Exoplanet in Orbit around the Sun-like Star GJ 504. *ApJ*, 774(1):11, September 2013. doi: 10.1088/0004-637X/774/1/11.
- T. Currie, S. Daemgen, J. Debes, et al. Direct Imaging and Spectroscopy of a Candidate Companion Below/Near the Deuterium-burning Limit in the Young Binary Star System, ROXs 42B. *ApJL*, 780(2):L30, January 2014. doi: 10.1088/2041-8205/780/2/L30.
- B. Macintosh, J. R. Graham, T. Barman, et al. Discovery and spectroscopy of the young jovian planet 51 Eri b with the Gemini Planet Imager. *Science*, 350(6256):64–67, October 2015. doi: 10.1126/science.aac5891.
- G. Chauvin, S. Desidera, A. M. Lagrange, et al. Discovery of a warm, dusty giant planet around HIP 65426. *A&A*, 605:L9, September 2017. doi: 10.1051/0004-6361/201731152.
- D. Charbonneau, T. M. Brown, D. W. Latham, et al. Detection of Planetary Transits Across a Sun-like Star. *ApJ*, 529(1):L45–L48, January 2000. doi: 10.1086/312457.
- G. W. Henry, G. W. Marcy, R. P. Butler, et al. A Transiting “51 Peg-like” Planet. *ApJ*, 529(1):L41–L44, January 2000. doi: 10.1086/312458.
- M. Mayor and D. Queloz. A Jupiter-mass companion to a solar-type star. *Nature*, 378(6555):355–359, November 1995. doi: 10.1038/378355a0.

- A. Wolszczan and D. A. Frail. A planetary system around the millisecond pulsar PSR1257 + 12. *Nature*, 355(6356):145–147, January 1992. doi: 10.1038/355145a0.
- D. G. Koch, W. J. Borucki, G. Basri, et al. Kepler Mission Design, Realized Photometric Performance, and Early Science. *ApJL*, 713(2):L79–L86, April 2010. doi: 10.1088/2041-8205/713/2/L79.
- S. B. Howell, C. Sobeck, M. Haas, et al. The K2 Mission: Characterization and Early Results. *PASP*, 126(938):398, April 2014a. doi: 10.1086/676406.
- M. J. Holman, D. C. Fabrycky, D. Ragozzine, et al. Kepler-9: A System of Multiple Planets Transiting a Sun-Like Star, Confirmed by Timing Variations. *Science*, 330(6000):51, October 2010. doi: 10.1126/science.1195778.
- S. E. Thompson, J. L. Coughlin, K. Hoffman, et al. Planetary Candidates Observed by Kepler. VIII. A Fully Automated Catalog with Measured Completeness and Reliability Based on Data Release 25. *ApJS*, 235(2):38, April 2018. doi: 10.3847/1538-4365/aab4f9.
- A. Cassan, D. Kubas, J. P. Beaulieu, et al. One or more bound planets per Milky Way star from microlensing observations. *Nature*, 481(7380):167–169, January 2012. doi: 10.1038/nature10684.
- S. B. Howell, C. Sobeck, M. Haas, et al. The K2 Mission: Characterization and Early Results. *PASP*, 126(938):398, April 2014b. doi: 10.1086/676406.
- G. R. Ricker, J. N. Winn, R. Vanderspek, et al. Transiting Exoplanet Survey Satellite (TESS). *Journal of Astronomical Telescopes, Instruments, and Systems*, 1:014003, January 2015. doi: 10.1117/1.JATIS.1.1.014003.
- B. J. Fulton and E. A. Petigura. The California-Kepler Survey. VII. Precise Planet Radii Leveraging Gaia DR2 Reveal the Stellar Mass Dependence of the Planet Radius Gap. *AJ*, 156(6):264, December 2018. doi: 10.3847/1538-3881/aae828.
- T. Mazeh, T. Holczer, and S. Faigler. Dearth of short-period Neptunian exoplanets: A desert in period-mass and period-radius planes. *Astronomy & Astrophysics*, 589:A75, May 2016. doi: 10.1051/0004-6361/201528065.
- L. M. Weiss, G. W. Marcy, E. A. Petigura, et al. The California-Kepler Survey. V. Peas in a Pod: Planets in a Kepler Multi-planet System Are Similar in Size and Regularly Spaced. *AJ*, 155(1):48, January 2018. doi: 10.3847/1538-3881/aa9ff6.
- J. B. Pollack. Origin and History of the Outer Planets: Theoretical Models and Observations L Constraints. *Annual review of astronomy and astrophysics*, 22:389–424, January 1984. doi: 10.1146/annurev.aa.22.090184.002133.

- A. P. Boss. Giant planet formation by gravitational instability. *Science*, 276:1836–1839, January 1997. doi: 10.1126/science.276.5320.1836.
- R. R. Rafikov. Can Giant Planets Form by Direct Gravitational Instability? *ApJL*, 621(1):L69–L72, March 2005. doi: 10.1086/428899.
- P. Goldreich and S. Tremaine. Disk-satellite interactions. *ApJ*, 241:425–441, October 1980. doi: 10.1086/158356.
- D. N. C. Lin, P. Bodenheimer, and D. C. Richardson. Orbital migration of the planetary companion of 51 Pegasi to its present location. *Nature*, 380(6575):606–607, April 1996. doi: 10.1038/380606a0.
- F. A. Rasio and E. B. Ford. Dynamical instabilities and the formation of extrasolar planetary systems. *Science*, 274:954–956, November 1996. doi: 10.1126/science.274.5289.954.
- S. S. Kumar. The Structure of Stars of Very Low Mass. *ApJ*, 137:1121, May 1963. doi: 10.1086/147589.
- T. Nakajima, B. R. Oppenheimer, S. R. Kulkarni, et al. Discovery of a cool brown dwarf. *Nature*, 378(6556):463–465, November 1995. doi: 10.1038/378463a0.
- B. R. Oppenheimer, S. R. Kulkarni, K. Matthews, et al. Infrared Spectrum of the Cool Brown Dwarf Gl 229B. *Science*, 270(5241):1478–1479, December 1995. doi: 10.1126/science.270.5241.1478.
- M. Keppler, M. Benisty, A. Müller, et al. Discovery of a planetary-mass companion within the gap of the transition disk around PDS 70. *A&A*, 617:A44, September 2018. doi: 10.1051/0004-6361/201832957.
- S. Y. Haffert, A. J. Bohn, J. de Boer, et al. Two accreting protoplanets around the young star PDS 70. *Nature Astronomy*, 3:749–754, June 2019. doi: 10.1038/s41550-019-0780-5.
- T. Currie, K. Lawson, G. Schneider, et al. Images of embedded Jovian planet formation at a wide separation around AB Aurigae. *Nature Astronomy*, April 2022. doi: 10.1038/s41550-022-01634-x.
- Y. Zhou, A. Sanghi, B. P. Bowler, et al. HST/WFC3 H α Direct-imaging Detection of a Pointlike Source in the Disk Cavity of AB Aur. *ApJ*, 934(1):L13, July 2022. doi: 10.3847/2041-8213/ac7fef.
- D. N. C. Lin and John Papaloizou. On the Tidal Interaction between Protoplanets and the Protoplanetary Disk. III. Orbital Migration of Protoplanets. *ApJ*, 309:846, October 1986. doi: 10.1086/164653.

- M. R. Bate, S. H. Lubow, G. I. Ogilvie, et al. Three-dimensional calculations of high- and low-mass planets embedded in protoplanetary discs. *MNRAS*, 341(1):213–229, May 2003. doi: 10.1046/j.1365-8711.2003.06406.x.
- S. Huang. Occurrence of Life in the Universe. *American Scientist*, 47(3):397–402, September 1959.
- J. F. Kasting, D. P. Whitmire, and R. T. Reynolds. Habitable Zones around Main Sequence Stars. *Icarus*, 101(1):108–128, January 1993. doi: 10.1006/icar.1993.1010.
- R. K. Kopparapu, R. Ramirez, J. F. Kasting, et al. Habitable Zones around Main-sequence Stars: New Estimates. *ApJ*, 765(2):131, March 2013. doi: 10.1088/0004-637X/765/2/131.
- B. S. Gaudi, S. Seager, B. Mennesson, et al. The Habitable Exoplanet Observatory (HabEx) Mission Concept Study Final Report. *arXiv e-prints*, art. arXiv:2001.06683, January 2020.
- The LUVOIR Team. The LUVOIR Mission Concept Study Final Report. *arXiv e-prints*, art. arXiv:1912.06219, December 2019.
- J. Rigby, M. Perrin, M. McElwain, et al. The Science Performance of JWST as Characterized in Commissioning. *PASP*, 135(1046):048001, April 2023. doi: 10.1088/1538-3873/acb293.
- D. Spergel, N. Gehrels, C. Baltay, et al. Wide-Field Infrared Survey Telescope—Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report. *arXiv e-prints*, art. arXiv:1503.03757, March 2015. doi: 10.48550/arXiv.1503.03757.
- B. Macintosh, J. R. Graham, P. Ingraham, et al. First light of the Gemini Planet Imager. *Proceedings of the National Academy of Science*, 111(35):12661–12666, September 2014. doi: 10.1073/pnas.1304215111.
- J. L. Beuzit, A. Vigan, D. Mouillet, et al. SPHERE: the exoplanet imager for the Very Large Telescope. *A&A*, 631:A155, November 2019. doi: 10.1051/0004-6361/201935251.
- N. Jovanovic, F. Martinache, O. Guyon, et al. The Subaru Coronagraphic Extreme Adaptive Optics System: Enabling High-Contrast Imaging on Solar-System Scales. *PASP*, 127(955):890, September 2015a. doi: 10.1086/682989.
- E. H. Por, S. Y. Haffert, V. M. Radhakrishnan, and et. al. High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator. In *Adaptive Optics Systems VI*, volume 10703 of *Proc. SPIE*, 2018. doi: 10.1117/12.2314407. URL <https://doi.org/10.1117/12.2314407>.

- G. Ruane, A. Riggs, J. Mazoyer, et al. Review of high-contrast imaging systems for current and future ground- and space-based telescopes I: coronagraph design methods and optical performance metrics. In *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*, volume 10698 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 106982S, August 2018. doi: 10.1117/12.2312948.
- R. Galicher and J. Mazoyer. Imaging exoplanets with coronagraphic instruments. *arXiv e-prints*, art. arXiv:2302.10833, February 2023. doi: 10.48550/arXiv.2302.10833.
- G. Foo, D. M. Palacios, and Jr. Swartzlander, G. A. Optical vortex coronagraph. *Optics Letters*, 30(24):3308–3310, December 2005. doi: 10.1364/OL.30.003308.
- M. D. Perrin, A. Sivaramakrishnan, R. B. Makidon, et al. The Structure of High Strehl Ratio Point-Spread Functions. *ApJ*, 596(1):702–712, October 2003. doi: 10.1086/377689.
- R. Soummer, A. Ferrari, C. Aime, et al. Speckle Noise and Dynamic Range in Coronagraphic Images. *ApJ*, 669(1):642–656, November 2007a. doi: 10.1086/520913.
- O. Guyon. Limits of Adaptive Optics for High-Contrast Imaging. *ApJ*, 629(1):592–614, August 2005. doi: 10.1086/431209.
- J. Lozi, O. Guyon, N. Jovanovic, et al. Characterizing vibrations at the Subaru Telescope for the Subaru coronagraphic extreme adaptive optics instrument. *Journal of Astronomical Telescopes, Instruments, and Systems*, 4:049001, October 2018. doi: 10.1117/1.JATIS.4.4.049001.
- C. Marois, D. Lafrenière, B. Macintosh, et al. Confidence Level and Sensitivity Limits in High-Contrast Imaging. *ApJ*, 673(1):647–656, January 2008b. doi: 10.1086/523839.
- P. J. Bordé and W. A. Traub. High-Contrast Imaging from Space: Speckle Nulling in a Low-Aberration Regime. *ApJ*, 638(1):488–498, February 2006. doi: 10.1086/498669.
- F. Martinache, O. Guyon, N. Jovanovic, et al. On-Sky Speckle Nulling Demonstration at Small Angular Separation with SCEXAO. *PASP*, 126(940):565, June 2014. doi: 10.1086/677141.
- A. Give'on, B. Kern, S. Shaklan, et al. Broadband wavefront correction algorithm for high-contrast imaging systems. In *Astronomical Adaptive Optics Systems and Applications III*, volume 6691 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 66910A, September 2007. doi: 10.1117/12.733122.
- L. Pueyo, J. Kay, N. J. Kasdin, et al. Optimal dark hole generation via two deformable mirrors with stroke minimization. *Appl. Opt.*, 48(32):6296, November 2009. doi: 10.1364/AO.48.006296.

- A. J. Eldorado Riggs, N. J. Kasdin, and T. D. Groff. Recursive starlight and bias estimation for high-contrast imaging with an extended Kalman filter. *Journal of Astronomical Telescopes, Instruments, and Systems*, 2:011017, January 2016. doi: 10.1117/1.JATIS.2.1.011017.
- N. Jovanovic, O. Absil, P. Baudoz, et al. Review of high-contrast imaging systems for current and future ground-based and space-based telescopes: Part II. Common path wavefront sensing/control and coherent differential imaging. In *Adaptive Optics Systems VI*, volume 10703 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 107031U, July 2018. doi: 10.1117/12.2314260.
- S. Y. Haffert, J. R. Males, K. Ahn, et al. Implicit electric field Conjugation: Data-driven focal plane control. *arXiv e-prints*, art. arXiv:2303.13719, March 2023. doi: 10.48550/arXiv.2303.13719.
- D. Lafrenière, C. Marois, R. Doyon, et al. A New Algorithm for Point-Spread Function Subtraction in High-Contrast Imaging: A Demonstration with Angular Differential Imaging. *ApJ*, 660(1):770–780, May 2007. doi: 10.1086/513180.
- R. Soummer, L. Pueyo, and J. Larkin. Detection and Characterization of Exoplanets and Disks Using Projections on Karhunen-Loève Eigenimages. *ApJ*, 755(2):L28, August 2012. doi: 10.1088/2041-8205/755/2/L28.
- C. Marois, D. Lafrenière, R. Doyon, et al. Angular Differential Imaging: A Powerful High-Contrast Imaging Technique. *ApJ*, 641(1):556–564, April 2006. doi: 10.1086/500401.
- D. Mawet, L. Pueyo, P. Lawson, et al. Review of small-angle coronagraphic techniques in the wake of ground-based second-generation adaptive optics systems. In *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*, volume 8442 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 844204, September 2012. doi: 10.1117/12.927245.
- C. Marois, R. Doyon, R. Racine, et al. Efficient Speckle Noise Attenuation in Faint Companion Imaging. *PASP*, 112(767):91–96, January 2000. doi: 10.1086/316492.
- G. Ruane, H. Ngo, D. Mawet, et al. Reference Star Differential Imaging of Close-in Companions and Circumstellar Disks with the NIRC2 Vortex Coronagraph at the W. M. Keck Observatory. *AJ*, 157(3):118, March 2019. doi: 10.3847/1538-3881/aafec2.
- S. Gladysz and J. C. Christou. Detection of Faint Companions through Stochastic Speckle Discrimination. *ApJ*, 684(2):1486–1495, September 2008a. doi: 10.1086/589679.
- P. K. Day, H. G. LeDuc, B. A. Mazin, et al. A broadband superconducting detector suitable for use in large arrays. *Nature*, 425(6960):817–821, October 2003. doi: 10.1038/nature02037.

- B. A. Mazin, B. Bumble, S. R. Meeker, et al. A superconducting focal plane array for ultraviolet, optical, and near-infrared astrophysics. *Optics Express*, 20(2):1503, January 2012. doi: 10.1364/OE.20.001503.
- P. Szypryt, S. R. Meeker, G. Coiffard, et al. Large-format platinum silicide microwave kinetic inductance detectors for optical to near-IR astronomy. *Optics Express*, 25(21):25894, October 2017. doi: 10.1364/OE.25.025894.
- N. Zobrist, W. H. Clay, G. Coiffard, et al. Membraneless Phonon Trapping and Resolution Enhancement in Optical Microwave Kinetic Inductance Detectors. *Phys. Rev. Lett.*, 129(1):017701, July 2022. doi: 10.1103/PhysRevLett.129.017701.
- N. Fruitwala, P. Strader, G. Cancelo, et al. Second generation readout for large format photon counting microwave kinetic inductance detectors. *Review of Scientific Instruments*, 91(12):124705, December 2020. doi: 10.1063/5.0029457.
- A. B. Walter, N. Fruitwala, S. Steiger, et al. The MKID Exoplanet Camera for Subaru SCEXAO. *PASP*, 132(1018):125005, December 2020. doi: 10.1088/1538-3873/abc60f.
- N. Fruitwala. *Readout and Calibration of Large Format Optical/IR MKID Arrays and Applications to Focal Plane Wavefront Control*. PhD thesis, University of California, Santa Barbara, 2021.
- M. J. Bonse, E. O. Garvin, T. D. Gebhard, et al. Comparing Apples with Apples: Robust Detection Limits for Exoplanet High-Contrast Imaging in the Presence of non-Gaussian Noise. *arXiv e-prints*, art. arXiv:2303.12030, March 2023. doi: 10.48550/arXiv.2303.12030.
- J. W. Goodman. *Statistical properties of laser speckle patterns*, volume 9, page 9. 1975. doi: 10.1007/BFb0111436.
- M. P. Cagigal and V. F. Canales. Experimental checking of the Rician statistics in partially compensated wave fronts. *Optical Engineering*, 40:2690–2697, December 2001. doi: 10.1117/1.1417495.
- M. P. Fitzgerald and J. R. Graham. Speckle Statistics in Adaptively Corrected Images. *ApJ*, 637(1):541–547, January 2006. doi: 10.1086/498339.
- A. Sivaramakrishnan, C. D. Koresko, R. B. Makidon, et al. Ground-based Coronagraphy with High-order Adaptive Optics. *ApJ*, 552(1):397–408, May 2001. doi: 10.1086/320444.
- S. Gladysz and J. C. Christou. Detection of Faint Companions through Stochastic Speckle Discrimination. *ApJ*, 684(2):1486–1495, September 2008b. doi: 10.1086/589679.

- N. Jovanovic, O. Guyon, F. Martinache, et al. Artificial Incoherent Speckles Enable Precision Astrometry and Photometry in High-contrast Imaging. *ApJ*, 813(2):L24, November 2015b. doi: 10.1088/2041-8205/813/2/L24.
- S. Gladysz and N. Yaitskova. Finding and measuring extrasolar planets using speckle statistics. In *Speckle 2010: Optical Metrology*, volume 7387 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 73870Z, September 2010. doi: 10.1117/12.870757.
- B. A. Mazin, S. R. Meeker, M. J. Strader, et al. ARCONS: A 2024 Pixel Optical through Near-IR Cryogenic Imaging Spectrophotometer. *PASP*, 125(933):1348, November 2013. doi: 10.1086/674013.
- S. R. Meeker, B. A. Mazin, A. B. Walter, et al. DARKNESS: A Microwave Kinetic Inductance Detector Integral Field Spectrograph for High-contrast Astronomy. *PASP*, 130(988):065001, June 2018. doi: 10.1088/1538-3873/aab5e7.
- J. R. Males, L. M. Close, K. Miller, et al. MagAO-X: project status and first laboratory results. 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 1070309, July 2018. doi: 10.1117/12.2312992.
- N. J. Swimmer, B. Mazin, J. I Bailey III, et al. An mkid camera for use behind magao-x. In *Ground-based and Airborne Instrumentation for Astronomy IX*, volume 12184, page 121843Z. SPIE, 2022.
- R. N. Hook and A. S. Fruchter. Dithering, Sampling and Image Reconstruction. In *Astronomical Data Analysis Software and Systems IX*, volume 216 of *Astronomical Society of the Pacific Conference Series*, page 521, January 2000.
- P. Virtanen, R. Gommers, T. E. Oliphant, et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- J. C. van Eyken, M. J. Strader, A. B. Walter, et al. The ARCONS Pipeline: Data Reduction for MKID Arrays. *ApJS*, 219(1):14, July 2015. doi: 10.1088/0067-0049/219/1/14.
- S. Gonzaga, W. Hack, A. Fruchter, et al. *The DrizzlePac Handbook*. 2012.
- A. S. Fruchter and R. N. Hook. Drizzle: A Method for the Linear Reconstruction of Undersampled Images. *PASP*, 114(792):144–152, February 2002. doi: 10.1086/338393.
- C. A. Gomez Gonzalez, O. Wertz, O. Absil, et al. VIP: Vortex Image Processing Package for High-contrast Direct Imaging. *AJ*, 154(1):7, July 2017. doi: 10.3847/1538-3881/aa73d7.

- S. B. Goebel, O. Guyon, D. N. B. Hall, et al. Measurements of Speckle Lifetimes in Near-infrared Extreme Adaptive Optics Images for Optimizing Focal Plane Wavefront Control. *PASP*, 130(992):104502, October 2018. doi: 10.1088/1538-3873/aad8ed.
- Astropy Collaboration. The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package. *AJ*, 156(3):123, September 2018. doi: 10.3847/1538-3881/aabc4f.
- PyTables Developers Team. PyTables: Hierarchical datasets in Python, 2002–. URL <http://www.pytables.org/>.
- R. O. Gray, C. J. Corbally, R. F. Garrison, et al. Contributions to the Nearby Stars (NStars) Project: Spectroscopy of Stars Earlier than M0 within 40 pc-The Southern Sample. *AJ*, 132(1):161–170, July 2006. doi: 10.1086/504637.
- F. van Leeuwen. Validation of the new Hipparcos reduction. *A&A*, 474(2):653–664, November 2007. doi: 10.1051/0004-6361:20078357.
- T. J. David and L. A. Hillenbrand. The Ages of Early-type Stars: Strömgren Photometric Methods Calibrated, Validated, Tested, and Applied to Hosts and Prospective Hosts of Directly Imaged Exoplanets. *ApJ*, 804(2):146, May 2015. doi: 10.1088/0004-637X/804/2/146.
- J. M. Stone, A. J. Skemer, P. M. Hinz, et al. The LEECH Exoplanet Imaging Survey: Limits on Planet Occurrence Rates under Conservative Assumptions. *AJ*, 156(6):286, December 2018. doi: 10.3847/1538-3881/aaec00.
- J. Gagné, E. E. Mamajek, L. Malo, et al. BANYAN. XI. The BANYAN Σ Multivariate Bayesian Algorithm to Identify Members of Young Associations with 150 pc. *ApJ*, 856(1):23, March 2018. doi: 10.3847/1538-4357/aaae09.
- A. W. Howard and B. J. Fulton. Limits on Planetary Companions from Doppler Surveys of Nearby Stars. *PASP*, 128(969):114401, November 2016. doi: 10.1088/1538-3873/128/969/114401.
- V. V. Makarov and G. H. Kaplan. Statistical Constraints for Astrometric Binaries with Nonlinear Motion. *AJ*, 129(5):2420–2427, May 2005. doi: 10.1086/429590.
- T. D. Brandt. The Hipparcos-Gaia Catalog of Accelerations. *ApJS*, 239(2):31, December 2018. doi: 10.3847/1538-4365/aaec06.
- T. Currie, T. D. Brandt, M. Kuzuhara, et al. SCEXAO/CHARIS Direct Imaging Discovery of a 20 au Separation, Low-mass Ratio Brown Dwarf Companion to an Accelerating Sun-like Star. *ApJL*, 904(2):L25, December 2020a. doi: 10.3847/2041-8213/abc631.

- T. D. Groff, J. Chilcote, N. J. Kasdin, et al. Laboratory testing and performance verification of the CHARIS integral field spectrograph. In *Ground-based and Airborne Instrumentation for Astronomy VI*, volume 9908 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 99080O, Aug 2016. doi: 10.1117/12.2233447.
- T. Currie, O. Guyon, J. Lozi, et al. On-sky performance and recent results from the Subaru coronagraphic extreme adaptive optics system. *arXiv e-prints*, art. arXiv:2012.05241, December 2020b.
- C. Z. Bond, S. Cetre, S. Lilley, et al. Adaptive optics with an infrared pyramid wavefront sensor at Keck. *Journal of Astronomical Telescopes, Instruments, and Systems*, 6: 039003, July 2020. doi: 10.1117/1.JATIS.6.3.039003.
- T. Currie, T. D. Brandt, T. Uyama, et al. SCEXAO/CHARIS Near-infrared Direct Imaging, Spectroscopy, and Forward-Modeling of κ And b: A Likely Young, Low-gravity Superjovian Companion. *AJ*, 156(6):291, December 2018a. doi: 10.3847/1538-3881/aae9ea.
- A. Sahoo, O. Guyon, J. Lozi, et al. Precision Photometric and Astrometric Calibration Using Alternating Satellite Speckles. *AJ*, 159(6):250, June 2020. doi: 10.3847/1538-3881/ab88cd.
- E. Serabyn, E. Huby, K. Matthews, et al. The W. M. Keck Observatory Infrared Vortex Coronagraph and a First Image of HIP 79124 B. *AJ*, 153(1):43, January 2017. doi: 10.3847/1538-3881/153/1/43.
- S. Steiger, J. I. Bailey, N. Zobrist, et al. The MKID Pipeline: A Data Reduction and Analysis Pipeline for UVOIR MKID Data. *AJ*, 163(5):193, May 2022. doi: 10.3847/1538-3881/ac5833.
- T. Currie, N. J. Kasdin, T. D. Groff, et al. Laboratory and On-sky Validation of the Shaped Pupil Coronagraph’s Sensitivity to Low-order Aberrations With Active Wavefront Control. *PASP*, 130(986):044505, April 2018b. doi: 10.1088/1538-3873/aaab41.
- J. R. Ducati. VizieR Online Data Catalog: Catalogue of Stellar Photometry in Johnson’s 11-color system. *VizieR Online Data Catalog*, January 2002.
- M. A. Millar-Blanchaer, M. D. Perrin, L. Hung, et al. GPI observational calibrations XIV: polarimetric contrasts and new data reduction techniques. In *Ground-based and Airborne Instrumentation for Astronomy VI*, volume 9908 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 990836, August 2016. doi: 10.1117/12.2233071.

- S. Gladysz and J. C. Christou. Reference-Less Detection, Astrometry, and Photometry of Faint Companions with Adaptive Optics. *ApJ*, 698(1):28–42, June 2009. doi: 10.1088/0004-637X/698/1/28.
- B. Macintosh, L. Poyneer, A. Sivaramakrishnan, et al. Speckle lifetimes in high-contrast adaptive optics. In *Astronomical Adaptive Optics Systems and Applications II*, volume 5903 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 170–177, August 2005. doi: 10.1117/12.627854.
- T. Currie, A. Burrows, Y. Itoh, et al. A Combined Subaru/VLT/MMT 1-5 μm Study of Planets Orbiting HR 8799: Implications for Atmospheric Properties, Masses, and Formation. *ApJ*, 729(2):128, March 2011. doi: 10.1088/0004-637X/729/2/128.
- D. Mawet, J. Milli, Z. Wahhaj, et al. Fundamental Limitations of High Contrast Imaging Set by Small Sample Statistics. *ApJ*, 792(2):97, September 2014. doi: 10.1088/0004-637X/792/2/97.
- T. D. Brandt, M. Rizzo, T. Groff, et al. 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. doi: 10.1117/1.JATIS.3.4.048002.
- T. Currie, C. Marois, L. Cieza, et al. No Clear, Direct Evidence for Multiple Protoplanets Orbiting LkCa 15: LkCa 15 bcd are Likely Inner Disk Signals. *ApJL*, 877(1):L3, May 2019. doi: 10.3847/2041-8213/ab1b42.
- T. Currie, J. Debes, T. J. Rodigas, et al. Direct Imaging Confirmation and Characterization of a Dust-enshrouded Candidate Exoplanet Orbiting Fomalhaut. *ApJL*, 760(2):L32, December 2012. doi: 10.1088/2041-8205/760/2/L32.
- T. Currie, R. Cloutier, S. Brittain, et al. Resolving the HD 100546 Protoplanetary System with the Gemini Planet Imager: Evidence for Multiple Forming, Accreting Planets. *ApJL*, 814(2):L27, December 2015. doi: 10.1088/2041-8205/814/2/L27.
- J. P. Greco and T. D. Brandt. The Measurement, Treatment, and Impact of Spectral Covariance and Bayesian Priors in Integral-field Spectroscopy of Exoplanets. *ApJ*, 833(2):134, December 2016. doi: 10.3847/1538-4357/833/2/134.
- L. Pueyo. Detection and Characterization of Exoplanets using Projections on Karhunen Loeve Eigenimages: Forward Modeling. *ApJ*, 824(2):117, June 2016. doi: 10.3847/0004-637X/824/2/117.
- J. Gagné, J. K. Faherty, K. L. Cruz, et al. BANYAN. VII. A New Population of Young Substellar Candidate Members of Nearby Moving Groups from the BASS Survey. *ApJS*, 219(2):33, August 2015. doi: 10.1088/0067-0049/219/2/33.

- M. J. Pecaut and E. E. Mamajek. Intrinsic Colors, Temperatures, and Bolometric Corrections of Pre-main-sequence Stars. *ApJS*, 208(1):9, September 2013. doi: 10.1088/0067-0049/208/1/9.
- L. Casagrande, C. Flynn, and M. Bessell. M dwarfs: effective temperatures, radii and metallicities. *MNRAS*, 389(2):585–607, September 2008. doi: 10.1111/j.1365-2966.2008.13573.x.
- F. Allard, D. Homeier, and B. Freytag. Models of very-low-mass stars, brown dwarfs and exoplanets. *Philosophical Transactions of the Royal Society of London Series A*, 370 (1968):2765–2777, June 2012. doi: 10.1098/rsta.2011.0269.
- M. Asplund, N. Grevesse, A. J. Sauval, et al. The Chemical Composition of the Sun. *ARA&A*, 47(1):481–522, September 2009. doi: 10.1146/annurev.astro.46.060407.145222.
- R. J. De Rosa, J. Rameau, J. Patience, et al. Spectroscopic Characterization of HD 95086 b with the Gemini Planet Imager. *ApJ*, 824(2):121, June 2016. doi: 10.3847/0004-637X/824/2/121.
- I. Baraffe, G. Chabrier, T. S. Barman, et al. Evolutionary models for cool brown dwarfs and extrasolar giant planets. The case of HD 209458. *A&A*, 402:701–712, May 2003. doi: 10.1051/0004-6361:20030252.
- T. D. Brandt, T. J. Dupuy, Y. Li, et al. orvara: An Efficient Code to Fit Orbits Using Radial Velocity, Absolute, and/or Relative Astrometry. *AJ*, 162(5):186, November 2021. doi: 10.3847/1538-3881/ac042e.
- R. J. De Rosa, J. Patience, P. A. Wilson, et al. The VAST Survey - III. The multiplicity of A-type stars within 75 pc. *MNRAS*, 437(2):1216–1240, January 2014. doi: 10.1093/mnras/stt1932.
- A. B. Walter, C. Bockstiegel, T. D. Brandt, et al. Stochastic Speckle Discrimination with Time-tagged Photon Lists: Digging below the Speckle Noise Floor. *PASP*, 131(1005):114506, November 2019. doi: 10.1088/1538-3873/ab389a.
- S. Steiger, T. Currie, T. D. Brandt, et al. SCEXAO/MEC and CHARIS Discovery of a Low-mass, 6 au Separation Companion to HIP 109427 Using Stochastic Speckle Discrimination and High-contrast Spectroscopy. *AJ*, 162(2):44, August 2021. doi: 10.3847/1538-3881/ac02cc.
- A. Boccaletti, E. Di Folco, E. Pantin, et al. Possible evidence of ongoing planet formation in AB Aurigae. A showcase of the SPHERE/ALMA synergy. *A&A*, 637:L5, May 2020. doi: 10.1051/0004-6361/202038008.

- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- R. H. Dodkins, K. K. Davis, B. Lewis, et al. First Principle Simulator of a Stochastically Varying Image Plane for Photon-counting High Contrast Applications. *PASP*, 132(1016):104503, October 2020. doi: 10.1088/1538-3873/aba9e4.
- B. L. Gerard, C. Marois, and R. Galicher. Fast Coherent Differential Imaging on Ground-based Telescopes Using the Self-coherent Camera. *AJ*, 156(3):106, September 2018. doi: 10.3847/1538-3881/aad23e.
- A. T. Rodack, R. A. Frazin, J. R. Males, et al. Millisecond exoplanet imaging: I method and simulation results. *Journal of the Optical Society of America A*, 38(10):1541, October 2021. doi: 10.1364/JOSAA.426046.
- R. Soummer, A. Ferrari, C. Aime, et al. Speckle Noise and Dynamic Range in Coronagraphic Images. *ApJ*, 669(1):642–656, November 2007b. doi: 10.1086/520913.