Title

Applications of photonic lanterns in exoplanet astronomy

Author

Lin, Jonathan

Publication Date

2025

UNIVERSITY OF CALIFORNIA
Los Angeles

Applications of photonic lanterns in exoplanet astronomy

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Physics & Astronomy

by

Jonathan Lin

2025

ABSTRACT OF THE DISSERTATION

Applications of photonic lanterns in exoplanet astronomy

by

Jonathan Lin
Doctor of Philosophy in Physics & Astronomy
University of California, Los Angeles, 2025
Professor Michael P. Fitzgerald, Chair

Though almost 6,000 confirmed exoplanets discovered to date, only ~ 200 have measured atmospheric spectra, none of which resemble the Earth. Even though such planets, and by extension solar-system-like exoplanetary systems may be rare, at least given current planet formation theories, the sheer number of exoplanetary systems suggests that our non-detection is mostly technological in nature. In particular, indirect detection methods such as the radial velocity and transit methods are biased towards higher masses, larger sizes, and/or smaller orbits, often yielding close-in gas giants or barren terrestrial planets scorched by stellar radiation. On the other hand, direct imaging techniques are biased towards larger orbits and larger planets, yielding cold gas giants often several times the mass of Jupiter. Planet formation remains similarly uncertain due to a lack of observables, with confirmed protoplanetary systems numbering only in the low single digits.

Nevertheless, the search for an Earth analog will be a priority for astronomy in the next decades. These efforts will be supported by the upcoming 30-m-class ground-based telescopes such as ESO's Extremely Large Telescope, as well as NASA's planned spaceborne flagship, the 6–8 m Habitable Worlds Observatory. These observatories will rely on new instruments which must overcome a suite of technical limits holding back the current state-of-the-art. At the highest level, our primary limit is one of contrast: an Earth-mass exoplanet orbiting a sun-like star will be outshone by 10 orders of magnitude in visible wavelengths, so that any exoplanetary light is obscured by a stellar glare too intense for current instruments to adequately block. The solution to this challenge will in turn require developments in the suppression of starlight, as well as wavefront control, while the subsequent characterization of the faint exoplanetary signal will require improvements to astronomical spectrometers and ultra-low-noise detectors. Of course, this list is non-exhaustive.

My dissertation considers how photonic devices – primarily the photonic lantern – can be applied to meet these technical challenges, and how such devices fit in the wider contexts of spectroscopy, wavefront sensing, and imaging. In the following chapters, I quantify the improvements in high-spectral-resolution diffraction-limited spectroscopy afforded by photonic lanterns, and then develop the photonic lantern as a focal-plane wavefront sensor. Taken together, these developments yield a combined spectrometer and wavefront sensor which achieves greater light efficiency and stability than conventional spectrometers, and which can simultaneously assist in coronagraphic wavefront control. Such a device would be well-suited for the future search of an exo-Earth, and more generally, the study of faint circumstellar environments. In tandem, my dissertation considers the tools and methods which enable the development of new photonic technologies for astronomy.

The dissertation of Jonathan Lin is approved.

John Miao

Smadar Naoz

Erik Petigura

Michael P. Fitzgerald, Committee Chair

University of California, Los Angeles

2025

TABLE OF CONTENTS

V Theory, models & methods 232

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I acknowledge the results presented in this thesis are based upon published works with additional coauthors, and are reprinted here with permission. Chapter 3 is adapted from publication [9], below. Chapter 4 is adapted from [8]. Chapter 5 is adapted from publications [6] & [7]. Chapter 6 is adapted from [2]. Chapter 7 is adapted from [5]. Chapter 8 is adapted from [1]. Thank you to all my coauthors, who made my research possible (and more importantly, enjoyable!)

It should go without saying that I'd like to thank my advisor, Mike Fitzgerald, who made all this work possible in the first place, and then put up with me for the 6 years it took me to do it, as well as the rest of the Fitzgerald group (both current and past) for all the support they've given me.

I'd also like to thank the team at Subaru/SCExAO, both for hosting me as a visitor multiple times and also for somehow keeping SCExAO functional, as well as Nem Jovanovic, who at times was like my second advisor. Going back even further, special thanks to Eve Lee and Eugene Chiang for putting me on an astronomy track in the first place.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2034835. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. This work was also supported by the National Science Foundation under Grant No. 2109232.

Last but not least, I wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Maunakea has always had within the indigenous Hawaiian community, and that I am fortunate to have the opportunity to conduct observations from this mountain.

VITA

2019	B.S. (Engineering Physics), University of California Berkeley
2022	M.S. (Astronomy), UCLA
2019–2025	Doctoral student, UCLA

PUBLICATIONS

[1] Lin, J. W., Fitzgerald, M. P., Xin, Y., Kim, Y. J., Guyon, O., Norris, B., Betters, C., Leon-Saval, S., Ahn, K., Deo, V., Lozi, J., Vievard, S., Levinstein, D., Sallum, S., & Jovanovic, N. (2025). Experimental and on-sky demonstration of spectrally dispersed wavefront sensing using a photonic lantern. Opt. Lett., 50(8), 2780–2783. doi:10.1364/OL.551624

[2] Lin, Jonathan, & Fitzgerald, M. P. (2024). Nonlinear techniques for few-mode wavefront sensors. Applied Optics, 63(34), 8748. doi:10.1364/ao.537925

[3] Jovanovic, N., Gatkine, et al. (2023). 2023 Astrophotonics Roadmap: pathways to realizing multi-functional integrated astrophotonic instruments. Journal of Physics: Photonics, 5(4), 042501. doi:10.1088/2515-7647/ace869

[4] Lin, J. W. (2024). cbeam: Coupled-mode propagator for slowly-varying waveguides (p. eascl:2404.001). Astrophysics Source Code Library, record ascl:2404.001

[5] Lin, J. W., Fitzgerald, M. P., Xin, Y., Kim, Y. J., Guyon, O., Norris, B., Betters, C., Leon-Saval, S., Ahn, K., Deo, V., Lozi, J., Vievard, S., Levinstein, D., Sallum, S., & Jovanovic, N. (2023). Real-time Experimental Demonstrations of a Photonic Lantern Wave-front Sensor. The Astrophysical Journal Letters, 959(2), L34. doi:10.3847/2041-8213/ad12a4

[6] Lin, J., Fitzgerald, M. P., Xin, Y., Kim, Y. J., Guyon, O., Leon-Saval, S. G., Norris, B., & Jovanovic, N. (2023). Focal-plane wavefront sensing with photonic lanterns II: numerical characterization and optimization. Journal of the Optical Society of America B, 40(12), 3196. doi:10.1364/josab.502962

[7] Lin, Jonathan, Fitzgerald, M., Xin, Y., Guyon, O., Leon-Saval, S., Norris, B., & Jovanovic, N. (2022). Focal-plane wavefront sensing with photonic lanterns I: theoretical framework. Journal of the Optical Society of America B, 39, 2643-2656. doi:10.1364/josab.466227

[8] Lin, J., Vievard, S., Jovanovic, N., Norris, B., Fitzgerald, M. P., Betters, C., … Xin, Y. (2022, August). Experimental measurements of AO-fed photonic lantern coupling efficiencies. Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation, 12188, 121882E. doi:10.1117/12.2630608

[9] Lin, J., Jovanovic, N., & Fitzgerald, M. P. (2021). Design considerations of photonic lanterns for diffraction-limited spectrometry. J. Opt. Soc. Am. B, 38(7), A51–A63. doi:10.1364/JOSAB.423664

[10] Lin, J. (2021). Lightbeam: Simulate light through weakly-guiding waveguides (p. eascl:2102.006). Astrophysics Source Code Library, record ascl:2102.006.

Part I

Background

This dissertation is divided into five parts. This part, **Part I**, reviews motivations and technical challenges from exoplanet science (Chapter §1) and provides an overview of astrophotonics (Chapter §2). **Part II** considers the improvements to diffraction-limited spectroscopy offered by photonic lanterns, including numerical studies (Chapter §3) as well as experimental validations (Chapter §4). **Part III** contains analytic and numerical studies on wavefront sensing with photonic lanterns (Chapter §5) and other few-moded wavefront sensors which are highly sensitive but also highly nonlinear (Chapter §6). **Part IV** reviews ongoing developments to demonstrate the PL wavefront sensor concept both in a laboratory setting (Chapter §7) and on-sky (Chapter §8), as well as future prospects (Chapter §9). Lastly, **Part V** provides an introduction to the underlying physics, mathematics, and numerical methods which support this dissertation, including a review of waveguide eigenmode solvers and propagators (Chapter §10), coupled-mode theory (Chapter §11), atmospheric turbulence theory and modelling (Chapters §12 & §13), adaptive optics control theory (Chapter §14), and optimal linear phase retrieval (Chapter §15).

CHAPTER 1

Challenges in exoplanet science

In 1995, Mayor & Queloz confirmed detection of the first known exoplanet, 51 Peg b [MQ95]; since then, astronomers have detected over 5,000 more exoplanets, primarily through radial velocity (RV) and transit techniques. One of the next major steps in exoplanet science is to directly image and characterize Earth-like exoplanets, a goal which in turn is now driving a two-fold push in the field of instrumentation to increase our sensitivity for smaller exoplanets at wider separations, and to better isolate exoplanet light from obscuring starlight.

The necessary advancements in instrumentation will be met, in part, by the upcoming 30 m-class ground-based telescopes, as well as NASA's upcoming Habitable Worlds Observatory (HWO), all of which will be fitted with a new generation of adaptive optics (AO) systems and science instruments. I briefly summarize some of the challenges to the direct characterization of Earth-like exoplanets that this new generation of instrumentation must resolve. First: the wavefronts collected by ground-based telescopes are distorted by atmospheric turbulence. These distortions must corrected with adaptive optics to preserve the precision and sensitivity of any downstream instruments. Second: a sun-like host star will outshine any Earth-like companion by a factor of 10^{10} in the optical [TO10]. This starlight must be separated from the reflected exoplanetary light, or otherwise rejected, in order to detect the planet - but this is complicated by the small angular separations between stars and their planetary companions. Third: planetlight is faint, and must be coupled in an efficient manner to any downstream instruments to preserve signal.

1.1 Science motivators

One of next big pushes in astronomy is "direct exoplanet characterization": the isolation and subsequent spectral analysis of exoplanet light. Below, I list some of the science that precise exoplanet spectroscopy might enable.

1. Molecular abundances for exoplanet atmospheres, which can help constrain planet formation models and inform us of ongoing chemical processes [KBM13]. Ultimately such measurements will enable future searches for the biochemical markers of life outside our solar system.

2. Planetary radial velocities, which in conjunction with astrometric measurements of the planet's position can constrain orbital parameters [RMK19]. If planetary radial velocities can be measured to a precision of a few m/s, we can also search for exomoons.

3. Exoplanet rotation rates, measured indirectly through Doppler-broadening of spectral lines. Measurements of exoplanet spin in turn informs of us of its accretion history [LK91].

4. The mapping of a planet's weather patterns and the thermal structure of its atmosphere, through the technique of Doppler mapping [Cro14]. This technique leverages the fact that atmospheric absorption lines will be modulated by local temperature variations on the planet's surface.

5. "Super-resolution" spectroastrometry [Bai98] and more generally spectrally-dispersed imaging of sub-diffraction-limit structures, like unresolved binaries or debris disks.

The need for direct characterization - especially of Earth-like exoplanets - was reinforced by the Astro2020 decadal survey, which listed direct characterization as one of the three major focuses for astronomy in the next decade [Nat21]. However, direct characterization is extremely challenging: out of the 5,000+ exoplanets detected so far, only some 30 are directly imaged [CBL22], and of those few, none are Earth-like. There are several major challenges, as follows.

1.2 Technical challenges

1.2.1 Wavefront error

The size of the upcoming 30 m-class telescopes relegates them as ground-based facilities: as such, these telescopes must contend with atmospheric speckle noise, which will distort incoming wavefronts and limit the performance of all downstream scientific instruments (space telescopes also must deal with speckles, albeit to a lesser extent). I give a brief overview of speckle noise, following [Rod81]. Suppose the phase of an incoming wavefront varies by some characteristic scale r_0, where $r_0 < D$, the telescope diameter. We can then treat the aperture as a collection of subapertures, r_0 in size, such that the phase over each subaperture is approximately constant. Each pair of apertures constitutes a two-beam interferometer, producing an Airy pattern of scale λ/r_0 (the size of the seeing disk) and modulated by a fringe pattern of scale λ/D. An aperture filled with subapertures that are in phase constructively interferes at the center of the focal plane and destructively interferes elsewhere, leaving a single centrally-located diffraction pattern of size λ/D, but if the subapertures are out of phase, the larger Airy pattern λ/r_0 is "cut up" by fringes into multiple "speckles" of size λ/D, whose locations and intensities will vary with time.

The issue of speckle noise for ground-based telescopes can be mitigated through the use of adaptive optics (AO). These are typically closed-loop wavefront control systems composed of a wavefront sensor (WFS) and deformable mirror (DM); measurements of phase errors by the sensor drive corrections via the DM, which uses actuators to alter the optical path length over the aperture. Performance of these sensors will be limited to an extent by DM spatial resolution and actuator speed, but also by the quality of sensing provided by the WFS. This is by no means an easy task. The crux of the issue is one of phase retrieval: how to recover phase information from a wavefront when the only observable is electromagnetic intensity. In astronomy, phase retrieval has historically been done by splitting off a portion of telescope light, placing an auxiliary optic (e.g. a microlens array or pyramid prism) in a conjugate pupil plane within the new optical arm, and then inverting a linearized model relating pupil phase to the sensor image. In this respect, one of the current roadblocks in improving correction are non-common-path aberrations (NCPAs, [MLA12]): quasi-static aberrations evolving on the timescale of minutes to hours that occur due to differences between WFS arm and science arm of an instrument, induced by humidity, temperature, and gravity vector changes. Also problematic are so-called "petal" aberrations, which arise due aperture segmentation and appear as phase distortions which are discontinuous in the telescope pupil plane at subaperture boundaries. Petal aberrations have several sources, including radiative cooling of supporting spiders (the low-wind effect [NMJ18]), cophasing and alignment errors for telescopes with segmented primary mirrors, and erroneous phase retrieval when the supporting spiders are thicker than the atmosphere's coherence length.

Adaptive optics systems are critical even for space-borne telescopes which avoid atmospheric wavefront errors, since such telescopes still must contend with NCPAs, optomechanical drifting, and segment cophasing errors. In fact, NASA's upcoming flagship mission, the Habitable Worlds Observatory (HWO), is currently one of the main drivers for AO and DM development: HWO is expected to require wavefront control at a precision of ~ 10 pm and a wavefront stability of ~ 10 pm/hr in order to meet its goal of exo-Earth direct imaging. Much of this requirement is driven by coronagraphy, covered in §1.2.3.

1.2.2 Angular resolution

An Earth-Sun analog will have an angular separation of 100 mas at distance of 10 pc, a typical distance to the nearest stars. Considering our goal of Earth-analog direct imaging, a back-of-the-envelope calculation shows that a binary with an Earth-like separation will be unresolved for current state-of-the-art 10-m telescopes (e.g. in NIR, $\lambda = 1$ μm) if the system is more than $\sim$ 50 pcs away. Outside this distance, any signal from an Earth-like exoplanet overlaps the core of the stellar PSF, and will almost assuredly be dominated by stellar photon noise. Even for exoplanets at separations of a few λ/D, isolation of planetlight is still extremely challenging due to photon noise from the extended diffraction pattern of the host star, which can still dominate the overall signal due to the relative faintness of the planet. On the ground, planetlight and starlight are further mixed by speckle noise. One solution is to build bigger telescopes (e.g. the upcoming ground-based 30-m class telescopes, or the 6-8 m HWO), which tightens the diffractive spread of starlight and reduces stellar photon noise, while also lowering the necessary integration times for the faint exoplanetary signal. In turn, larger telescopes come with their own associated challenges in manufacturing, mirror segment alignment, wavefront sensing, and DM technology.

1.2.3 Starlight rejection

The third challenge is the inherent difficulty in disentangling planetlight from starlight when planets are outshone by their host stars by many orders of magnitude (10^{10} for the Earth-Sun-like system in the visible), and are so closely located next to their host stars on-sky. In conjunction with increasing aperture size, the obscuring starlight can also be suppressed through coronagraphy: a general term for a class of techniques that block incoming light from an on-axis point source. Most coronagraphic methods involve placing an occulting mask (the coronagraphic mask, which in its simplest form is a small disk that blocks the star) in the focal plane and a secondary mask (the Lyot stop) in a conjugate pupil plane. Of course, coronagraphy is not a complete solution, either: the presence of wavefront error or even slight misalignment can allow on-axis light to leak through the system. The minimum separation a star and companion can have before becoming undetectable is termed the inner working angle (IWA) of the coronagraph. Typical IWAs of modern coronagraphs are typically around 2 λ/D or larger. For HWO, potential detections of Earth-like companions are expected to occur at the several λ/D scales in visible wavelengths, which in turn drives a need for targeted coronagraphy (i.e. coronagraphic suppression of starlight only around the known exoplanet position) at 1-2 λ/D scales in the near-infrared.

While coronographic techniques using focal-plane occulters or phase masks are most common among current observing facilities, they are not the only method. Other techniques include interferometric nulling, in which starlight is destructively interfered with itself (e.g. [MNT21]), and vortex fiber nulling, in which a phase corkscrew produced by a "vortex mask" prevents the coupling of on-axis starlight into a fiber (e.g. [REJ19]). Note that perfect suppression of on-axis starlight also removes a portion of the off-axis planetlight[1] [GPK06], giving a fundamental limit to coronagraphy. While no coronagraphs operating at this limit have yet been demonstrated, active interferometric nulling using photonic devices provides a plausible pathway [BSJ21].

1.2.4 Precision spectroscopy

The fourth major challenge in direct characterization is the precise extraction of spectral information from the planetary signal - in other words, spectrometer design. The precision of a spectrometer hinges on its optomechanical stability, as well as the stability of the input signal (e.g. planetlight). The latter, in turn, is limited primarily by wavefront errors, even with AO correction. These instabilities can be compounded when aberrated light is transferred to an optical fiber if the core diameter of the fiber is large compared to the wavelength of the light. Since large-diameter fibers let light propagate in multiple eigenmodes, the power traveling down these so-called "multimode fibers" (MMFs) will continually shift between the available modes over the course of propagation, due to bends, imperfections, stresses, and strains along the fiber. This is modal noise. While the issue of modal noise for MMF-fed spectrometers has been partially circumvented in the past by scrambling the propagating light with fiber agitators, in many cases SMF-fed spectrometers are becoming

---

[1] This holds even for the ideal case of a filled circular aperture, except at specific angular separations corresponding to nulls in the Airy pattern.

the more attractive option. SMFs have only a single, "fundamental" mode of propagation which prevents modal noise and provides a static beam profile for the spectrometer even in the presence of time-varying wavefront errors, greatly improving instrument precision. Diffraction-limited SMF-fed spectrometers are also significantly more compact than MMF-fed "seeing-limited" spectrometers, reducing cost and improving stability against optomechanical deflections [JSG17]. These gains come at the expense of instrument throughput. Because the fundamental mode is flat in phase and Gaussian-like in amplitude (at least for circular step-index optical fibers), efficient coupling of telescope light into SMFs is difficult for ground-based observatories, resulting in greater light losses in SMF-fed spectrometers over MMF-fed spectrometers. The above discussion outlines a tension in spectrometer design between instrument *stability* and *throughput*. For precise spectra of a faint exoplanetary signal, we require both.

For radial velocity detections of exo-Earths with the upcoming 30-m telescopes, another challenge will be the suppression of so-called "stellar rotation bias". The spin of a star gives it a bluer limb and redder limb. If such a star is partially resolved, then the spectrum of the star may change depending on the alignment of the star relative to the spectrograph slit (or fiber feed), as well as any time-varying wavefront errors. This spectral noise must be suppressed if the next generation of radial velocity instruments is to reach the 10 cm/s precision required to detect an Earth analog.

REFERENCES

[Bai98] Jeremy A. Bailey. "Spectroastrometry: a new approach to astronomy on small spatial scales." In Sandro D'Odorico, editor, *Optical Astronomical Instrumentation*, volume 3355, pp. 932 – 939. International Society for Optics and Photonics, SPIE, 1998.

[BSJ21] Ruslan Belikov, Dan Sirbu, Jeffrey B. Jewell, Olivier Guyon, and Christopher C. Stark. "Theoretical performance limits for coronagraphs on obstructed and unobstructed apertures: how much can current designs be improved?" In Stuart B. Shaklan and Garreth J. Ruane, editors, *Techniques and Instrumentation for Detection of Exoplanets X*, volume 11823 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 118230W, September 2021.

[CBL22] Thayne Currie, Beth Biller, Anne-Marie Lagrange, Christian Marois, Olivier Guyon, Eric Nielsen, Mickael Bonnefoy, and Robert De Rosa. "Direct Imaging and Spectroscopy of Extrasolar Planets." *arXiv e-prints*, p. arXiv:2205.05696, May 2022.

[Cro14] Crossfield, Ian J. M. "Doppler imaging of exoplanets and brown dwarfs." *A&A*, **566**:A130, 2014.

[GPK06] O. Guyon, E. A. Pluzhnik, M. J. Kuchner, B. Collins, and S. T. Ridgway. "Theoretical Limits on Extrasolar Terrestrial Planet Detection with Coronagraphs." *ApJS*, **167**(1):81–99, November 2006.

[JSG17] Jovanovic, N., Schwab, C., Guyon, O., Lozi, J., Cvetojevic, N., Martinache, F., Leon-Saval, S., Norris, B., Gross, S., Doughty, D., Currie, T., and Takato, N. "Efficient injection from large telescopes into single-mode fibres: Enabling the era of ultra-precision astronomy." *A & A*, **604**:A122, 2017.

[KBM13] Quinn M. Konopacky, Travis S. Barman, Bruce A. Macintosh, and Christian Marois. "Detection of Carbon Monoxide and Water Absorption Lines in an Exoplanet Atmosphere." *Science*, **339**(6126):1398–1401, 2013.

[LK91] Jack J. Lissauer and David M. Kary. "The origin of the systematic component of planetary rotation I. Planet on a circular orbit." *Icarus*, **94**(1):126–159, November 1991.

[MLA12] Martinez, P., Loose, C., Aller Carpentier, E., and Kasper, M. "Speckle temporal stability in XAO coronagraphic images." *A&A*, **541**:A136, 2012.

[MNT21] Marc-Antoine Martinod, Barnaby Norris, Peter Tuthill, Tiphaine Lagadec, Nemanja Jovanovic, Nick Cvetojevic, Simon Gross, Alexander Arriola, Thomas Gretzinger, Michael J. Withford, Olivier Guyon, Julien Lozi, Sébastien Vievard, Vincent Deo, Jon S. Lawrence, and Sergio Leon-Saval. "Scalable photonic-based nulling interferometry with the dispersed multi-baseline GLINT instrument." *Nature Communications*, **12**(1):2465, Apr 2021.

[MQ95] Michel Mayor and Didier Queloz. "A Jupiter-mass companion to a solar-type star." *Nature*, **378**(6555):355–359, Nov 1995.

[Nat21] National Academies of Sciences, Engineering, and Medicine. *Pathways to Discovery in Astronomy and Astrophysics for the 2020s.* The National Academies Press, Washington, DC, 2021.

[NMJ18] N´Diaye, M., Martinache, F., Jovanovic, N., Lozi, J., Guyon, O., Norris, B., Ceau, A., and Mary, D. "Calibration of the island effect: Experimental validation of closed-loop focal plane wavefront control on Subaru/SCExAO." *A&A*, **610**:A18, 2018.

[REJ19] Garreth Ruane, Daniel Echeverri, Nemanja Jovanovic, Dimitri Mawet, Eugene Serabyn, J. Kent Wallace, Jason Wang, and Natasha Batalha. "Vortex fiber nulling for exoplanet observations: conceptual design, theoretical performance, and initial scientific yield predictions." In Stuart B. Shaklan, editor, *Techniques and Instrumentation for Detection of Exoplanets IX*, volume 11117, pp. 366 – 381. International Society for Optics and Photonics, SPIE, 2019.

[RMK19] Jean-Baptiste Ruffio, Bruce Macintosh, Quinn M. Konopacky, Travis Barman, Robert J. De Rosa, Jason J. Wang, Kielan K. Wilcomb, Ian Czekala, and Christian Marois. "Radial Velocity Measurements of HR 8799 b and c with Medium Resolution Spectroscopy." *AJ*, **158**(5):200, November 2019.

[Rod81] F. Roddier. "V The Effects of Atmospheric Turbulence in Optical Astronomy." In E. Wolf, editor, *Progress in Optics*, volume 19 of *Progress in Optics*, pp. 281–376. Elsevier, 1981.

[TO10] W. A. Traub and B. R. Oppenheimer. "Direct Imaging of Exoplanets." In S. Seager, editor, *Exoplanets*, pp. 111–156. University of Arizona Press, 2010.

CHAPTER 2

Astrophotonics

The field of astrophotonics crosses photonic devices - which guide and control photons in a manner analogous to how electronic devices guide and control electrons - with astronomical applications. The capabilities of astrophotonics ranges from the basic routing of light between telescopes and scientific instruments, to more complex applications involving active control (i.e. interferometry) of confined light; the 2023 Astrophotonics Roadmap [JGA23] provides an excellent overview. As we will see, astrophotonic light processing is one pathway to expand the technical capability of scientific instruments. Astrophotonics can also presents a more compact, optomechanically stable, and cost-efficient alternative to bulk optical instruments. Below, I give a non-exhaustive list of astrophotonic devices.

2.1 Example devices and applications

2.1.1 Optical fibers

Arguably the first astrophotonic device, optical fibers are primarily distinguished by the number of propagating modes they support. Multi-mode fibers (MMFs) support many optical modes and are often used to both collect as much telescope light as possible and route it to a downstream science instrument, whereas single-mode fibers (SMFs) allow only one propagating mode[1] and are often used when control over the beam's shape and flatten the

---

[1] Technically, two modes with the same spatial distribution but orthogonal polarizations. See §10 for more details.

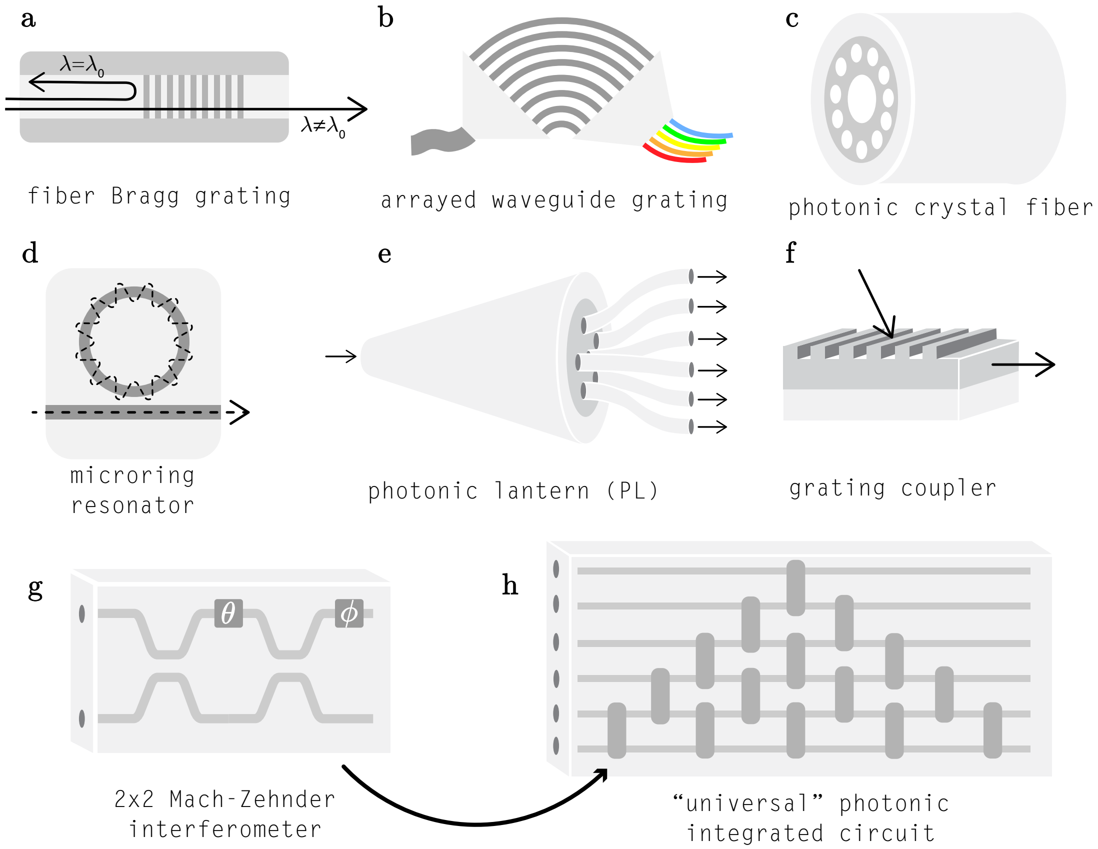

Figure 2.1: Cartoons of several astrophotonic devices. **a**: the fiber Bragg grating, which reflects light in a narrow wavelength band. **b**: the arrayed waveguide grating, a compact photonic disperser. **c**: a photonic crystal fiber, which uses refractive index patterns to guide light. The fiber shown here uses small circular voids to create a photonic bandgap. **d**: a microring resonator. The ring supports whispering-gallery modes which can be used to produce a frequency comb. **e**: the photonic lantern, which can couple light from an MMF into multiple SMFs. **f**: a grating coupler, which couples free-space light into a slab waveguide. **g**: a Mach-Zehnder interferometer formed from 2 single-mode waveguide channels. When the channels are "pinched" together, they act as a beamsplitter. In-line phase shifters, denoted by θ and ϕ, complete the interferometer. **h**: a 6×6 circuit composed of Mach-Zehnder units. This "mesh" can apply an arbitrary unitary operation to the 6 input signals.

beam's phase. SMFs can be used to improve the stability of high-R spectroscopy, albeit at the cost of greater light losses. Another application is vortex fiber nulling, where the symmetric shape of an SMF's fundamental mode is exploited to null on-axis starlight [REJ19]. In between these two extremes are few-mode fibers (FMFs), which support on the order of 10 propagating modes.

2.1.2 Fiber Bragg gratings

One challenge in obtaining precise NIR spectra is that for λ = 1-2 μm, the night sky background is dominated by the bright spectral lines of OH radicals in the upper atmosphere. Fiber Bragg gratings (FBGs, Figure 2.1a) are a compact solution to reject this unwanted light at extremely high spectral resolution ($R \sim 1$ million). An FBG is an an optical fiber whose core refractive index modulates along the propagation direction; this modulation acts as a wavelength-specific mirror (like a dielectric stack in an anti-reflective optical coating) which can be tuned to reflect OH spectral lines. FBGs were among the first astrophotonic technologies to appear in the field of astronomy, initially used by the GNOSIS instrument and its successor PRAXIS [TEB13, EBB18].

2.1.3 Arrayed waveguide gratings

As shown in Figure 2.1b, the arrayed waveguide grating (AWG; [TSK90]) is a planar device that takes in one input fiber for broadband light, and passes that light through an array of waveguides with different path lengths (the 'phased array', which applies a wavelength-dependent phase ramp to the propagating light). The phased array spectrally disperses light, which is then split into multiple output fibers. Such devices can provide a more compact, optomechanically stable, and cost-effective solution for spectral dispersion than conventional designs that use gratings and prisms.

2.1.4 Microring resonators

Microring resonators (Figure 2.1d) are small ring-shaped waveguides. When evanescently coupled to another waveguide channel, these resonators produce sharp, periodic spectral features similar to a Dirac delta comb, which are of particular interest in the wavelength calibration high-R spectrometers. Another astronomical application is dispersion-less cross-correlation spectroscopy, where the microresonator ring is tuned to produce spectral features similar to that of a molecule of interest (e.g. mimicking ro-vib lines); the periodic spectral filter produced by the resonator is then optically cross-correlated with the signal (e.g. exoplanetary light) to perform molecular detection.

2.1.5 Photonic lanterns

The photonic lantern (PL; [BGY15]) is a tapered waveguide that gradually transitions from a few-mode optical fiber (FMF) geometry to an array of single-mode cores, similar to a multicore fiber (MCF), which can then be fanned out to a series of SMFs. Figure 2.1e shows the geometry of a typical 3-port PL. When the FMF end is placed in the focal plane, the PL can efficiently couple telescope light into multiple SMFs, even in the presence of aberrations. These devices were originally developed as spatial multiplexers by the telecommunications industry. While the design space for PLs is large, the space can be mostly divided into three categories. In what we call the "standard" PL, embedded single-mode cores are uniform in structure and refractive index. At the other extreme, "mode-selective" PLs use differing core radii or index contrasts, so that each fiber mode at the FMF-like lantern entrance routes to a unique output [LFS14]. Lastly, we term lanterns that operate between these two extremes "hybrid lanterns." These lanterns funnel light from the fundamental fiber mode into a specific output while mixing the remaining light in the remaining outputs.

PLs have several astronomical applications. In general, these devices are a method of efficiently coupling focal plane light collected by a telescope into other astrophotonic or fiber-fed instruments, including photonic circuits (e.g. for active nulling of on-axis starlight) or SMF-fed diffraction-limited spectrometers. The PL itself can also be exploited to sample electric field, enabling techniques such as wavefront sensing and imaging. Because the outputs of the PL are relatively straightforwards to disperse, these techniques can incorporate spectrally dispersed signals.

2.1.6 Grating couplers

One of the key technical challenges in astrophotonics is how to efficiently couple light into the waveguides which comprise astrophotonic devices, a process that is often light-inefficient. Grating couplers (Figure 2.1f) are one potential solution. These devices use a grating pattern (typically periodic, though aperiodic grating designs exist) to couple normally incident light into a waveguide oriented in the transverse direction. These devices are being considered in astronomy as a method to individual subapertures of a telescope pupil into channels within a PIC, e.g. for coronagraphy.

2.1.7 Photonic integrated circuits

Photonic integrated circuits (PICs, [NB19]) use "circuits" of waveguides and inline phase-shifters to coherently split, interfere, and recombine guided light. These circuits can be designed to apply an arbitrary linear operation to the circuit inputs. One promising architecture is a "Mach-Zehnder mesh": a circuit composed of multiple interlocking Mach-Zehnder interferometers [Mil17]. Figure 2.1g & h show a Mach-Zehnder interferometer and mesh, respectively. With tunable phase-shifters, such circuits might provide a secondary level of wavefront control in addition to what is provided by the deformable mirror (DM) of a conventional AO system. PICs can also be used in nulling applications, where the goal is to actively and destructively interfere starlight (through electro-optical phase shifters within the circuit) to reveal fainter planetlight. This approach is being pursued for instance in the

AstroPIC project [SBF24].

2.1.8 Pupil reformatters

From Fourier optics, the pupil and focal plane of an imaging system are Fourier pairs, such that every pair of baselines measured in the pupil plane correspond to a spatial frequency (a fringe pattern) in the image plane. For a seeing-limited system, wavefront phase varies significantly over the pupil, and thus power measured in the focal plane is added incoherently from multiple redundant baselines, introducing speckle noise. Pupil reformatters divide the pupil into sub-pupils and inject light from each (e.g. using a microlens array or photonic slicer) into a photonic chip that physically rearranges the light from the sub-pupils into a non-redundant configuration, which can be imaged, thereby eliminating speckle noise while also preserving most of the telescope light (as opposed to approaches like sparse aperture masking). In essence, this approach converts a telescope to an array of interferometers, and is used in the FIRST instrument [VDH22] at Subaru telescope.

2.2 Platforms and fabrication

This section introduces fabrication techniques for astrophotonic devices, and briefly discusses both the current manufacturing landscape as will as current limitations.

2.2.1 Lithography

For astrophotonic devices with a flat geometry (e.g. "planar" photonic circuits, sometimes called planar lightwave circuits, which have waveguide channels all embedded in the same spatial plane), the most well-developed fabrication methods use lithographic techniques originally developed in the electronics industry for complementary metal-oxide semiconductor (CMOS) circuits. Lithographic fabrication is also the standard for directional couplers, multimode interferometers, grating couplers, microresonators, arrayed waveguide gratings, and some photonic crystal waveguides. Within lithography, there are several combinations of optical substrates, or "platforms", that astrophotonic waveguides can be built from. The most mature platform is silicon-on-insulator (SOI), which confines light within a small silicon core (refractive index $n \sim 3.4$ at $\lambda = 1.55$ μm) a few hundred nms across, surrounded by a fused silica cladding or air. The waveguides constructed from this platform have very high index contrast and thus can be bent tightly without loss; however, the small mode field diameters of SOI waveguide channels, compared to that of the typical weak-index-contrast optical fibers used in astronomy, makes the coupling of light from astronomical telescopes into SOI devices very lossy. Light signals are also attenuated as they propagate through the silicon cores. Silicon nitride (SiN) is an alternate platform with intermediate index contrast and lower optical losses than SOI, but suffers from polarization-dependent losses because the shapes of SiN waveguides are often highly asymmetric. The third silicon-based platform is silica-on-silicon, in which waveguides are constructed from doped silica, like a typical weakly guiding optical fiber. Such devices have low loss and can be more easily coupled with optical fibers, suiting astronomical use, but are also typically more expensive to manufacture and much less compact than devices built on higher contrast SOI and SiN platforms. There are vendors for all three platforms who currently offer custom fabrication.

In general, silicon-based platforms are attractive due to their compatibility with CMOS fabrication, meaning that electrical and photonic circuitry can be stacked on the same device for electrical control of the photonic circuit.[2] For instance, a thermoelectric heating pad can be placed over a section of optical waveguide, enabling active control over the waveguide's refractive index, and therefore the amount of phase shift applied. Non-silicon-based materials such as indium phosphide (InP), lithium niobate (LiN), and polymers are also of interest, particularly in nonlinear optics and medical applications.

---

[2] Other integration methods for electrical and photonic circuits, such as microwire bonding, are also possible. This step is often referred to as "packaging."

2.2.2 Ultrafast laser inscription

Waveguides can also be written by moving the point of a focused laser beam through an optical substrate. This technique, called ultrafast laser inscription (ULI), relies on the local modification of refractive index caused by the high-power laser focus, and enables the production of fully 3D waveguides. ULI has been applied to many types of glasses, but is limited to weak index contrast (up to $\sim 10^{-2}$). ULI waveguides are also more difficult to integrate with electronics, and are more often used in passive applications. Examples include the GLINT nulling chip [MNT21] and FIRST pupil remapper [VDH22].

2.2.3 Fiber tapering

Optical fibers or fiber-like waveguides such as photonic lanterns can be produced in an optical fiber draw tower, in which a piece of glass stock resembling a scaled-up version of the desired waveguide is heated and pulled, tapering the stock down to fiber size. In PL fabrication, the glass stock is formed by inserting a bundle of SMFs into a lower-index glass capillary. Then, one end is drawn down to the size of a few-moded optical fiber, while maintaining the other end's geometry. PLs can also be made in this manner using multicore optical fibers as the glass stock. In the latter approach, the PL can be used as-is, or a multicore fiber fanout can be attached to provide SMF outputs.

REFERENCES

[BGY15] T. A. Birks, I. Gris-Sánchez, S. Yerolatsitis, S. G. Leon-Saval, and R. R. Thomson. "The photonic lantern." *Adv. Opt. Photon.*, **7**(2):107–167, Jun 2015.

[EBB18] S. C. Ellis, S. Bauer, C. Bacigalupo, J. Bland -Hawthorn, J. J. Bryant, S. Case, R. Content, T. Fechner, D. Giannone, R. Haynes, E. Hernandez, A. J. Horton, U. Klauser, J. S. Lawrence, S. G. Leon-Saval, E. Lindley, H. G. Löhmannsröben, S. S. Min, N. Pai, M. Roth, K. Shortridge, L. Waller, Pascal Xavier, and Ross Zhelem. "PRAXIS: an OH suppression optimised near infrared spectrograph." In *Proceedings of the SPIE*, volume 10702 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 107020P, Jul 2018.

[JGA23] Nemanja Jovanovic, Pradip Gatkine, Narsireddy Anugu, Rodrigo Amezcua-Correa, Ritoban Basu Thakur, Charles Beichman, Chad Bender, Jean-Philippe Berger, Azzurra Bigioli, Joss Bland-Hawthorn, Guillaume Bourdarot, Charles M. Bradford, Ronald Broeke, Julia Bryant, Kevin Bundy, Ross Cheriton, Nick Cvetojevic, Momen Diab, Scott A. Diddams, Aline N. Dinkelaker, Jeroen Duis, Stephen Eikenberry, Simon Ellis, Akira Endo, Donald F. Figer, Michael Fitzgerald, Itandehui Gris-Sanchez, Simon Gross, Ludovic Grossard, Olivier Guyon, Sebastiaan Y. Haffert, Samuel Halverson, Robert J. Harris, Jinping He, Tobias Herr, Philipp Hottinger, Elsa Huby, Michael Ireland, Rebecca Jenson-Clem, Jeffrey Jewell, Laurent Jocou, Stefan Kraus, Lucas Labadie, Sylvestre Lacour, Romain Laugier, Katarzyna Ławniczuk, Jonathan Lin, Stephanie Leifer, Sergio Leon-Saval, Guillermo Martin, Frantz Martinache, Marc-Antoine Martinod, Ben Mazin, Stefano Minardi, John D Monnier, Reinan Moreira, Denis Mourard, Abani Shankar Shankar Nayak, Barnaby Norris, Ewelina Obrzud, Karine Perraut, François Reynaud, Steph Sallum, David Schiminovich, Christian Schwab, Eugene Serbayn, Sherif Soliman, Andreas Stoll, Liang Tang, Peter Tuthill, Kerry Vahala, Gautam Vasisht, Sylvain Veilleux, Alexander B. Walter, Edward J Wollack, Yinzi Xin, Zongyin Yang, Stephanos Yerolatsitis, Yang Zhang, and Chang-Ling Zou. "2023 Astrophotonics Roadmap: pathways to realizing multi-functional integrated astrophotonic instruments." *Journal of Physics: Photonics*, 2023.

[LFS14] Sergio G. Leon-Saval, Nicolas K. Fontaine, Joel R. Salazar-Gil, Burcu Ercan, Roland Ryf, and Joss Bland-Hawthorn. "Mode-selective photonic lanterns for space-division multiplexing." *Opt. Express*, **22**(1):1036–1044, Jan 2014.

[Mil17] David A. B. Miller. "Setting up meshes of interferometers – reversed local light interference method." *Opt. Express*, **25**(23):29233–29248, Nov 2017.

[MNT21] Marc-Antoine Martinod, Barnaby Norris, Peter Tuthill, Tiphaine Lagadec, Nemanja Jovanovic, Nick Cvetojevic, Simon Gross, Alexander Arriola, Thomas Gretzinger, Michael J. Withford, Olivier Guyon, Julien Lozi, Sébastien Vievard, Vincent Deo, Jon S. Lawrence, and Sergio Leon-Saval. "Scalable photonic-based nulling interferometry with the dispersed multi-baseline GLINT instrument." *Nature Communications*, **12**(1):2465, Apr 2021.

[NB19] Barnaby Norris and Joss Bland-Hawthorn. "Astrophotonics: The Rise of Integrated Photonics in Astronomy." *Opt. Photon. News*, **30**(5):26–33, May 2019.

[REJ19] Garreth Ruane, Daniel Echeverri, Nemanja Jovanovic, Dimitri Mawet, Eugene Serabyn, J. Kent Wallace, Jason Wang, and Natasha Batalha. "Vortex fiber nulling for exoplanet observations: conceptual design, theoretical performance, and initial scientific yield predictions." In Stuart B. Shaklan, editor, *Techniques and Instrumentation for Detection of Exoplanets IX*, volume 11117, pp. 366 – 381. International Society for Optics and Photonics, SPIE, 2019.

[SBF24] Dan Sirbu, Ruslan Belikov, Kevin Fogarty, Carson Valdez, Zhanghao Sun, Annie Kroo, Olav Solgaard, David A. B. Miller, and Olivier Guyon. "AstroPIC: near-infrared photonic integrated circuit coronagraph architecture for the Habitable Worlds Observatory." In Laura E. Coyle, Shuji Matsuura, and Marshall D. Perrin, editors, *Space Telescopes and Instrumentation 2024: Optical, Infrared, and Millimeter Wave*, volume 13092, p. 130921T. International Society for Optics and Photonics, SPIE, 2024.

[TEB13] Christopher Q. Trinh, Simon C. Ellis, Joss Bland-Hawthorn, Jon S. Lawrence, Anthony J. Horton, Sergio G. Leon-Saval, Keith Shortridge, Julia Bryant, Scott Case, Matthew Colless, Warrick Couch, Kenneth Freeman, Hans-Gerd Löhmannsröben, Luke Gers, Karl Glazebrook, Roger Haynes, Steve Lee, John O'Byrne, Stan Miziarski, Martin M. Roth, Brian Schmidt, Christopher G. Tinney, and Jessica Zheng. "GNOSIS: The first instrument to use fiber Bragg gratings for OH suppression." *The Astronomical Journal*, **145**(2):51, jan 2013.

[TSK90] H. Takahashi, S. Suzuki, K. Kato, and I. Nishi. "Arrayed-waveguide grating for wavelength division multi/demultiplexer with nanometre resolution." *Electronics Letters*, **26**(2):87, January 1990.

[VDH22] Sébastien B. Vievard, Vincent Deo, Elsa Huby, Sylvestre Lacour, Olivier Guyon, Nick Cvetojevic, Kevin Barjot, Guillermo Martin, Manon Lallement, Julien Lozi, Takayuki Kotani, Franck Marchis, Daniel Rouan, Kyohoon Ahn, Nour Skaf, and Guy Perrin. "Interferometric wavefront sensing at FIRST/SCExAO : self-calibrated fibered pupil-remapping spectrocopy using a metrology laser source." In Antoine Mérand, Stephanie Sallum, and Joel Sanchez-Bermudez, editors, *Optical and Infrared Interferometry and Imaging VIII*, volume 12183, p. 121830R. International Society for Optics and Photonics, SPIE, 2022.

Part II

Applications to spectroscopy

CHAPTER 3

Design considerations of photonic lanterns for diffraction-limited spectrometry

The wavelength axis of an astronomical signal holds a wealth of information. Atoms and molecules make themselves known through the characteristic emission and absorption features they imprint on passing light, and variations in feature shapes can reveal the motion of those particles through the Doppler effect. At a more macroscopic level, a periodic Doppler shift of spectral features indicates an astronomical source which may be wobbling due to an orbiting companion – this is the well-known radial velocity method for exoplanet detection. Improvements to the wavelength resolution of our spectroscopic measurements could make possible the radial velocity detection of an Earth analog, which would apply a reflex motion of only 10 cm/s to its host star. Highly resolved spectra can also be used for Doppler tomography and spectroastrometry, both of which exploit wavelength information to do a basic form of high-angular-resolution imaging. Precision spectroscopy is challenging; but astrophotonics may provide solutions to some of the technical challenges.

3.1 Introduction

The stability of a ground-based spectrometer is primarily limited by atmospheric turbulence, which distorts incoming wavefronts and makes the point-spread function (PSF) of an optical system temporally unstable. These instabilities can be compounded when aberrated light is transferred to an optical fiber if the core diameter of the fiber is large compared to the wavelength of the light. Since large-diameter fibers let light propagate in multiple eigenmodes (specifically, the linearly polarized or LP modes, which apply for weakly guiding and radially symmetric waveguides [Mit16]), the power travelling down these "multi-mode fibers" (MMFs) will continually shift between the available modes over the course of propagation, due to bends, imperfections, stresses, and strains along the fiber. This is termed modal noise [BW01, JSC16a]. While the issue of modal noise has been partially circumvented in the past by scrambling the propagating light with fiber agitators [PBG13], in many cases spectrometers fed with single-mode fibers (SMFs) are becoming the more attractive option. SMFs have only a single mode of propagation (LP_{01}) which automatically eliminates modal noise and provides a static beam profile for the spectrometer even in the presence of a temporally fluctuating input field distribution, greatly improving instrument precision. Diffraction-limited SMF-fed spectrometers are also significantly more compact than traditional seeing-limited spectrometers, reducing cost and improving stability against thermal and mechanical deflections [BLR10, JSC16b]. These gains come at the expense of instrument throughput. Because the LP_{01} mode is flat in phase and Gaussian-like in amplitude, efficient coupling into SMFs is difficult at ground-based observatories, resulting in greater light losses in SMF-based fiber injectors over MMF-based fiber injectors.

There are several pathways to increasing throughput while maintaining instrument stability. In SMF coupling, the first hurdle to overcome is the illumination mismatch between the Airy pattern of the telescope beam and the Gaussian-like distribution of the fiber mode. One method to overcome this is by using phase-induced amplitude apodization (PIAA) optics [Guy03]. PIAA optics use a pair of lenses (or mirrors) in collimated space to reshape the flat-topped aperture illumination profile of incoming waves into a distribution more amenable for applications such as fiber injection and coronagraphy. In the context of SMF injection, incoming waves are typically reshaped to be Gaussian in amplitude, thereby preventing the formation of diffraction rings in the focal plane. This, in turn, promotes a PSF that better matches the fundamental mode of an SMF, increasing coupling efficiency. Such

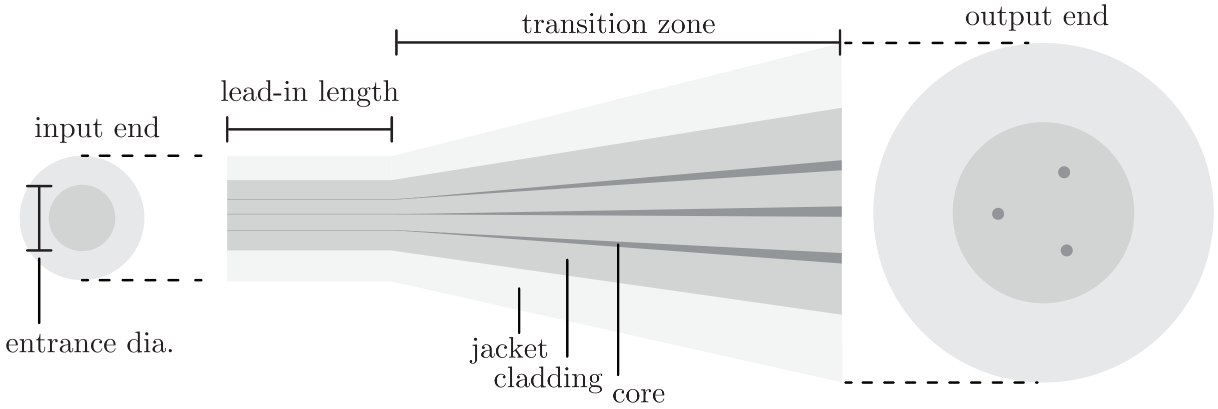

Figure 3.1: The photonic lantern, a tapered waveguide that can adiabatically transfer light from a multi-mode fiber-like input end into multiple single-mode cores. In particular, the lanterns considered in this work have output geometries similar to multi-core fibers. We choose this architecture for simplicity, and neglect propagation in the single-mode cores after the lantern transition. For the purpose of this work we additionally assume a linear tapering profile between the "lead-in" portion of the waveguide and the lantern's output, along with a core-cladding index contrast of 8.8×10^{-3} and a cladding-jacket contrast of 5.5×10^{-3}. Important geometrical considerations in lantern design include the lead-in length and taper factor (the degree by which SMF cores shrink from output to input).

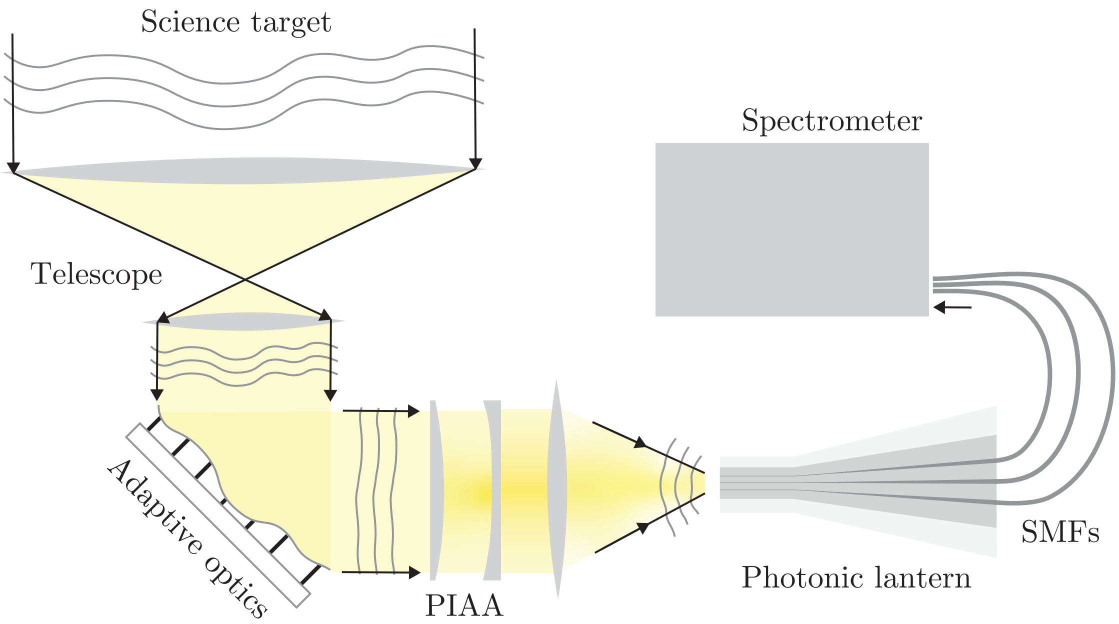

Figure 3.2: A diagram of a fiber-fed diffraction-limited spectrometer coupled to a telescope with an AO system and a PL fiber injector. The amplitude-remapping effect of PIAA optics is shown through gradients, in yellow.

"beam-shaping" optics have been shown to work well (e.g. [JSG17, CJR21], on SCExAO and KPIC respectively). However, PIAA optics only correct for errors in field amplitude, not phase. Phase errors (formed from both atmospheric and instrumental effects) are the second hurdle in efficient SMF coupling. Such errors can be corrected with adaptive optics (AO), but these systems are less effective at shorter wavelengths. For example, at the short end of the near-infrared (NIR) y and J bands (0.97–1.35 μm), where the Strehl ratio (the ratio between the peak intensity of an aberrated PSF and the peak intensity of an unaberrated PSF) is low ($< 50\%$), residual phase errors will limit the coupling to $< 50\%$. These wavelength regimes will be the focus for the next push in direct exoplanet spectroscopy with upcoming instruments such as HISPEC and MODHIS [MFK19]. An alternative approach is to initially couple aberrated telescope light into a MMF, which can more efficiently accept light from the complex speckle patterns characteristic of low Strehl beams, and "stabilize" the light before injection into the instrument. The field of astrophotonics makes such a pathway possible via the photonic lantern (PL) [LFS14, BGY15]. The PL is a waveguide that transitions adiabatically from an MMF-like geometry to a multi-core fiber (MCF) geometry, allowing light to move between the MMF-like and MCF-like ends with low loss. Such devices can be formed by heating and tapering one end of either an MCF or a bundle of SMFs (though methods such as laser inscription can also be used [TBL11], for a review of construction methods see [BGY15]). A diagram of a PL formed from a tapered MCF is shown in Fig. 3.1. Such devices are already in use with fiber Bragg gratings (e.g. GNOSIS [BEL11], PRAXIS [EBB18]) and are currently being considered for use in focal plane wavefront sensing [CMH18, NWB20]. Figure 3.2 shows how a PL might be used in the context of high-resolution spectrometry. AO-corrected (and potentially beam-shaped) telescope light is injected into the MMF-like end of a PL; over the course of propagation through the lantern, this light becomes efficiently confined within the single-mode lantern cores, which can then be routed into the science instrument through SMFs. As such, lantern-based fiber injectors retain both the gain in input coupling efficiency offered by MMF-based injectors as well as the gain in instrument stability offered by SMF-based injectors.

In recent years, significant headway has been made in our understanding of PLs and their potential uses in astronomy [HMC15, CJG17, AHC18, CMH18]. PLs are also currently being considered for a number of applications outside astronomy, including mode-selective fiber lasers [WZA18] and amplifiers [MHM17], spatial mode control in optical fibers [MAH16], LIDAR systems [OTW15], optical ground receivers for laser satellite communications [TVL19], and free space optical communications [OTW15]. The latter three investigate how PLs can couple light through the atmosphere, similar to what we consider in astronomy. This work aims at extending existing studies of PLs to the realm of high-resolution spectrometry. In particular, we focus on characterizing the chromatic behavior of PLs, their coupling as function of Strehl, and their performance in the few-port regime (especially relevant for high-resolution spectrometry which requires many pixels): all properties that must be better understood before lanterns can make their way into the next generation of spectrometers. To this end, we present a numerical model for PL performance, as outlined in §3.2. We then apply this model in §3.3 to showcase the potential gains offered by PLs across a range of geometries, wavelengths, and wavefront errors (WFEs), with the aim of informing the next generation of instrument designs. In the process, we include a first look at the interaction between PLs and beam-shaping PIAA optics. Results are discussed in §3.4.

3.2 Methodology

On-sky testing with an experimental setup is the ultimate way to determine the true performance of a lantern-based injector. However, before investing in the development of complex infrastructure, it is important to understand the expected performance with a detailed model; such models will also help interpret future experimental results. We present a model for lantern performance based on the statistical assessment of the scalar electric fields at the lantern's input and output. From these fields, we compute our prioritized metric for lantern performance: the lantern throughput, defined as the ratio between total power in the single-mode cores at the lantern's output and the total incident power at the lantern's input (though other metrics, such as the degree of mutual coherence between the lantern output cores, may also be applicable in other contexts). The electric field at the lantern's input is a result of partially compensated aberrations as well as potential beam shaping, while the output electric field depends on the propagation through the lantern device.

There are multiple approaches to computing lantern throughput. The electric field distribution at the lantern output after beam propagation can be readily converted to a lantern throughput, via overlap integrals (note that this is different than taking the total power of the output field, which may include extra power in the cladding, not guided by the SMF cores). However, under certain circumstances, throughput can also be computed from a "modal analysis" perspective through careful comparison between the input field and the

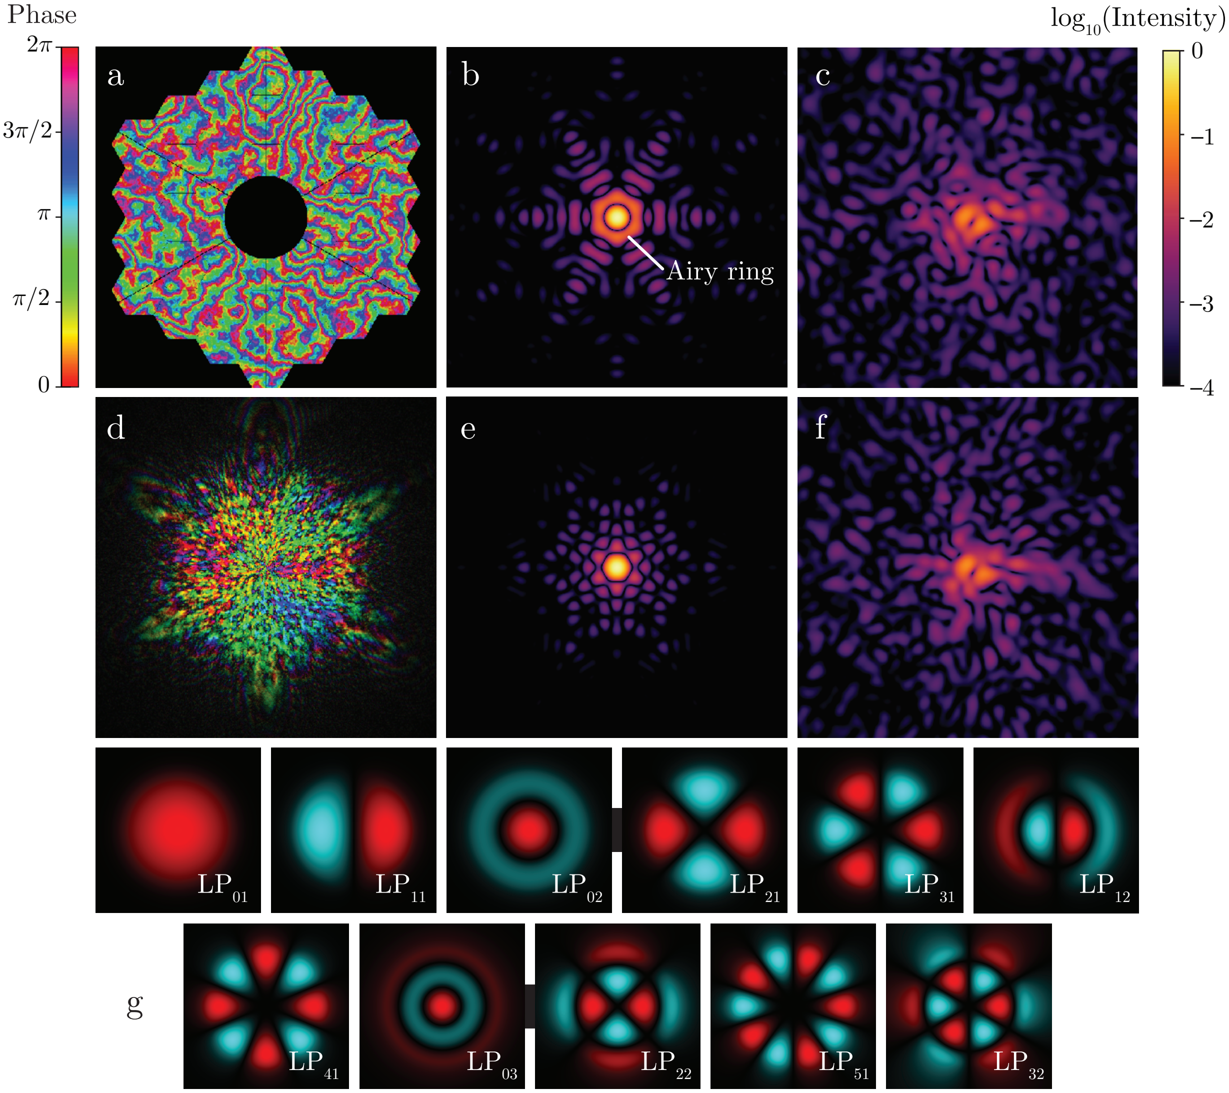

Figure 3.3: **Panel a**: the Keck pupil, with turbulence-induced phase errors plotted in color. **b**: the PSF in the absence of WFE. **c**: same as c but with WFE (~10% Strehl). **d**: same thing as a, after beam-shaping to a Gaussian illumination profile. **e**: same thing as b, but with beam-shaping. Note the disappearance of the Airy ring. **f**: same as c, but with beam-shaping. **g**: the first 19 LP_{lm} (linearly polarized) modes, the basis for propagation down a circular, weakly-guiding step-index fiber. Modes are presented in the order they appear as core diameter scales up. Note that all modes with $l \neq 0$ can be rotated by 90° to yield another linearly independent mode; these rotations are not shown. Modes that appear simultaneously are shown connected. Panels a, d, and g display phase information according to the left color bar while panels b, c, e, and f display intensity information according to the right color bar.

propagating modes of the lantern. We take both approaches in this work. In this section, we detail our models for WFE (§3.2.1), AO correction (§3.2.2), beam shaping (§3.2.3), and injection into and propagation through a PL (§3.2.4).

3.2.1 Wavefront error

To compute a coupling efficiency for any sort of fiber injector, we first need a statistical model describing the WFE of light entering our optical system. To this end we assume the standard Kolmogorov turbulence model for a single thin layer of frozen-flow turbulence translating at a velocity of 10 m/s. Realizations of turbulent phase screens are generated over a period of 10 seconds, corresponding to 10 pupil crossings over the course of a simulation for an assumed 10 meter aperture. The atmospheric coherence length is a parameter of the simulation, which we use to adjust the average Strehl ratio of the PSF delivered by the AO correction. Our particular implementation of Kolmogorov turbulence employs the Fourier-based autoregressive algorithm of [SPR15] (see also §13 for a derivation of the method); an example aberrated wavefront in the pupil plane generated with this method is shown in Fig. 3.3a. However, Fourier-based turbulence models typically under-represent low-order WFEs [JG94]. Additionally, the Kolmogorov power spectrum does not account for other sources of low-order error such as non-common path aberrations (NCPAs) and mechanical vibrations. As such, extra low-order phase errors can be added to make WFE more realistic. In this work, we allow for the injection of extra tip-tilt error by applying a Gaussian-distributed random displacement to the PSF in the waveguide input focal plane. The magnitude of this displacement is a parameter of the simulation that can be adjusted. Static WFEs, such as those caused by misalignment or imperfect optical prescriptions (giving rise to, say, coma or astigmatism), are neglected.

3.2.2 Adaptive optics

The turbulence-aberrated wavefronts as simulated above are partially corrected using an AO model built with HCIPy [PHR18], a high-contrast imaging package written for Python. The general trends in our results will be independent of pupil geometry, but for the purpose of this work we assume the Keck pupil. We assume a deformable mirror (DM) with 30 × 30 actuators across the pupil, in line with current "extreme AO" systems and similar to the $\sim$1,000-element DM which will be introduced to Keck as part of the phase II of the KPIC project [JCP20]. The control law is handled using a basic leaky-integrator scheme. Propagation from the pupil to focal plane is accomplished using HCIPy's Fraunhofer propagator. Example PSFs from our model in the presence and absence of turbulence-induced WFE are shown in Figs. 3.3b and 3.3c. Finally, for each tested lantern entrance diameter, we adjust the focal ratio of the optical system in order to maximize the spatial overlap between the PSF and the lantern modes, ensuring optimal coupling efficiency.

3.2.3 Beam-shaping (PIAA) optics

PIAA optics remap the amplitude distribution of light in the pupil plane, in order to promote a more amenable amplitude distribution (e.g. for fiber injection or coronagraphy) in the focal plane. This remapping is usually accomplished by a pair of lenses (or mirrors) placed after the wavefront corrector. In the context of SMF injection, the target amplitude distribution is usually selected to follow a truncated, circularly symmetric Gaussian profile. This profile removes the Airy ring, created by diffraction from the pupil mask, and promotes a Gaussian-like PSF that matches well with the Gaussian-like fundamental mode of SMFs ([JSG17]; also see the first panel of Fig. 3.3g for an example of the fundamental mode). The action of such beam-shaping optics is shown in the second row of Fig. 3.3: Fig. 3.3d shows the Keck pupil post-apodization, while Figs. 3.3e and 3.3f show apodized PSFs with and without turbulence-induced aberrations. Without beam-shaping optics, coupling efficiencies into

SMFs for optical systems with Keck-like pupils (secondary obstruction ratio of ~30%) would be limited to ~65%; with such optics, coupling efficiencies can reach upwards of 80% [JSG17, CJR21]. In this work we consider how these optics may also be applied to lantern-based fiber injectors. We computed the required lens profiles corresponding to the desired Gaussian amplitude distribution by numerically solving a differential equation for the lens surfaces, obtained in the geometric optics limit [Guy03]. When we use beam-shaping optics in our simulations, we propagate wavefronts through lenses with the derived lens profiles, after correction from the DM but before propagation to the focal plane. For free-space propagation between the two lenses, we use HCIPy's angular spectrum propagator.

3.2.4 Photonic lantern

After propagation to the focal plane, light is injected into a PL. In this work, we only consider simple lanterns of the type shown in Fig. 3.1, where the transverse structure of the transition zone at any point along its length is just a linearly scaled-down version of the lantern's output (i.e. isolated single-mode cores in a common cladding). The scale difference between the output and input ends we term the "taper factor". Such lanterns deviate slightly from real lanterns in that the claddings of real lanterns do not necessarily maintain circularity as they taper down; however, they are a good first approximation. Due to the construction method, lanterns formed by tapering fibers or fiber bundles may also have a "lead-in" length of MMF-like waveguide at their input end, which we simulate in our model by extending the tapered end of the transition zone.

The most straightforward way to simulate lantern throughput is to numerically propagate injected light through the lantern and compute the power in the output lantern modes. However, numerical propagation is computationally expensive. This expense is multiplied in the presence of stochastic processes such as turbulence, where numerical propagation must be repeated for multiple realizations of the input electric field in order to obtain a statistical average, and additionally compounded during parameter space explorations (e.g.

over WFE or wavelength, as done in this work). To simplify and accelerate our simulations, we split throughput into two parts: coupling efficiency, which reflects the initial loss of light during coupling into the lantern entrance, and transition efficiency, which reflects losses that occur during propagation through the lantern. Previous works have shown that transition efficiencies can be as high as ∼97% [TEB13, LFS14]; a necessary but insufficient condition for this to occur is to have at least as many single-mode outputs (or "ports") at the lantern exit as guided modes at the lantern entrance [BGY15]. Assuming negligible transition effects, the input coupling efficiency of the lantern can be taken as a proxy for overall lantern throughput, bypassing the need to simulate the lantern transition and saving on computation time. In this work, lantern coupling efficiencies are generally computed under the assumption that the number of ports equals the number of modes and transition losses are negligible.

Coupling efficiencies on their own are quick to simulate. We only need to compute a basis of "lantern modes": a set of orthogonal, linearly independent field distributions at the lantern entrance which span the space of bound electric field distributions when propagated to the lantern output. Then, the coupling efficiency into a lantern entrance for a given input electric field is simply the total amount of power coupled into all of the lantern modes, which we can compute via overlap integrals between the input field and the lantern modes. The overlap integral η for two scalar complex fields E_1 and E_2 is given by

$$\eta = \frac{\left|\int E_1^* E_2 dA\right|^2}{\int |E_1|^2 \, dA \int |E_2|^2 \, dA}. \tag{3.1}$$

For large lantern taper factors (small residual SMF cores at the lantern's entrance), the lantern modes approach the LP$_{lm}$ modes: analytically known eigenmodes of propagating waves in a circularly symmetric step-index fiber. Example amplitude distributions for the first 19 LP modes are shown in Fig. 3.3g (rotated variants of modes $l \neq 0$ are not shown).

Later in this work we compute lantern throughputs (rather than input coupling efficiencies), and extend lantern characterization over regimes where lantern operation has been considered "inefficient" and the above approach breaks down. In these regimes, we need to apply full numerical beam propagation to connect the electric field distributions at the lantern input and output. In order to simulate the propagation of light through PLs, we have built `lightbeam` [Lin20]: a numerical package which uses the "finite-differences beam propagation method" (FD BPM) algorithm to numerically and implicitly solve the paraxial Helmholtz equation in the weakly guiding regime. Our particular variant of FD BPM, which is based on the algorithm presented in [YSS97] and is reviewed in §10, is accurate to fourth order transversely, to second order longitudinally, and works on an adaptive mesh. We have tested this code against RSoft BeamProp, a commercial package for optical propagation, with consistent results. We have also verified that our code reproduces the expected behavior for single and few-mode circularly symmetric step-index fibers.

3.2.5 Tested configurations

All simulated lanterns are assumed to have a fluorine-based glass jacket and a fused silica cladding with an index contrast of 5.5×10^{-3}, materials typically used to construct lanterns (private communication with S. Leon-Saval). We consider lantern operation across the astronomical y and J bands (0.97–1.35 μm), where the Strehl ratio of typical AO systems is lower and there is a stronger need for PLs. The embedded SMF cores are assumed to have similar specifications to OFS ClearLite 980 16 fiber, which has a core-cladding index contrast of 8.8×10^{-3} and a nominal core diameter of 4.4 μm. This fiber was chosen due to its single-mode cutoff at 970 nm, consistent with our chosen wavelength range. We additionally assume linear lantern taper profiles throughout this work; other geometric parameters are adjustable.

We use this model to assess lantern performance as a function of entrance diameter, wavelength, and WFE, as well as the length of the lantern's lead-in waveguide and taper factor. We also compare the performance of SMFs and 3-, 6-, and 19-port lanterns with and without beam-shaping optics. We have chosen to consider 3- and 6- port lanterns because their low port counts make them more applicable for use in high-resolution spectrometers, which require many pixels and lack room for many fibers. We additionally consider the 19-port lantern because such devices are already in use (e.g. GNOSIS [TEB13], PRAXIS [EBB18]). Additionally, these port numbers were selected because they reflect how the number of guided modes increases as core diameter is scaled up and/or wavelength is scaled down.

3.3 Results

In this section, we characterize lantern performance as a function of Strehl ratio, operating wavelength, and geometric factors including lantern entrance diameter, taper factor, and lead-in waveguide length. Sections 3.3.1 through 3.3.5 characterize lantern performances in terms of coupling efficiency into the lantern entrance, which serves under most circumstances as a good proxy for lantern throughput. In Section 3.3.6, we use full numerical beam propagation to compute proper throughputs and extend the results of prior sections.

3.3.1 Coupling efficiency vs. Strehl ratio

With the index contrast of the lantern fixed, the primary factor that governs lantern coupling efficiency is the diameter of the lantern's input waveguide, which sets the number of supported modes at the lantern's entrance and thus determines the lantern's ability to accept aberrated light. We present coupling efficiencies at $\lambda = 1$ μm into the input ends of lanterns for various entrance diameters (and hence mode counts) as a function of Strehl ratio in Fig. 3.4a. The Strehl ratio is computed from the intensity average of 100 turbulence-aberrated PSF realizations. We additionally adjust focal ratio to maximize the overall coupling efficiency at each core diameter (re-optimizing the focal ratio when beam-shaping is included in the optical system). We assume that the embedded SMF cores are small after tapering (large taper factor) and that the lantern entrance is circular, so that the lantern modes can be taken as the LP modes. Other geometric parameters such as lead-in length are for now irrelevant; we consider them again in Section 3.3.6. WFE was solely generated from Fourier-based Kolmogorov turbulence (no extra tip-tilt was added; we consider extra tip-tilt later in §3.3.5). Figure 3.4a indicates that coupling efficiencies increase linearly with Strehl ratio independent of mode count, with the largest mode count lantern consistently outperforming all other configurations. Performance gains from PLs over SMFs are also evident across the entire range of tested Strehl ratios: even at 10% Strehl, the simplest three-mode lantern outperforms direct injection into an SMF by a factor of 2, netting a throughput of ∼10%. In the absence of beam-shaping optics, six-mode lanterns offer even greater gains, reaching a ∼15% increase over both SMFs and three-port lanterns at 70% Strehl ratio.

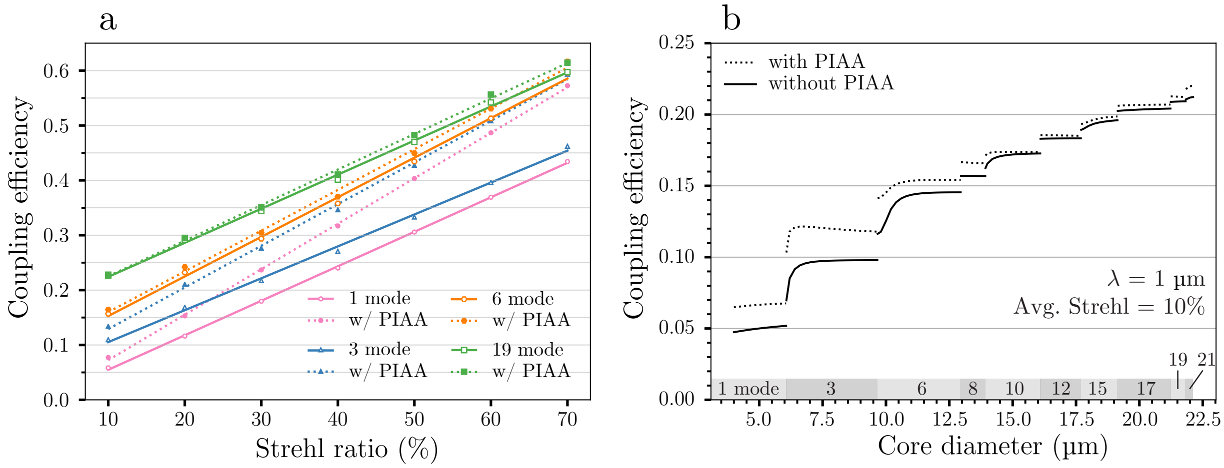

Figure 3.4: a: coupling efficiency comparison for 1-, 3-, 6-, and 19-mode lanterns against Strehl ratio at $\lambda = 1$ μm, with and without beam-shaping (PIAA) optics. Lantern entrances were approximated as step-index fibers. The specified mode counts were selected because they reflect how the number of LP modes increases with core diameter. PSF degradation was simulated using standard Kolmogorov turbulence. b: Mean fiber coupling efficiency in the presence and absence of beam-shaping optics, at a wavelength of 1 μm for an average Strehl ratio of 10%. The number of available LP modes is annotated at the bottom.

3.3.2 Effect of lantern size

In Fig. 3.4b, we extend the previous calculation over a range of core diameters, this time fixing turbulence to give an average 10% Strehl at the operating wavelength of 1 μm. As its diameter increases, the number of guided modes at the lantern's entrance increases in a stepwise fashion, resulting in discontinuous jumps in the lantern coupling efficiency. Importantly, the number of guided modes (annotated at the bottom of the plot) does not increase one at a time: modes usually appear in groups of 2 or 3. Figure 3.4b shows that the most significant gains in coupling efficiency occur in the 1- to 3-mode transition and the 3- to 6-mode transition at 10% Strehl, beyond which increases in lantern size and modality yield diminishing returns. However, if the increases in instrument complexity and exposure times are not a concern, it may still be beneficial to adopt a high-mode count lanterns: as seen in Fig. 3.4, the coupling efficiency of a 19-port lantern is around double that of a 3-port lantern at 10% Strehl.

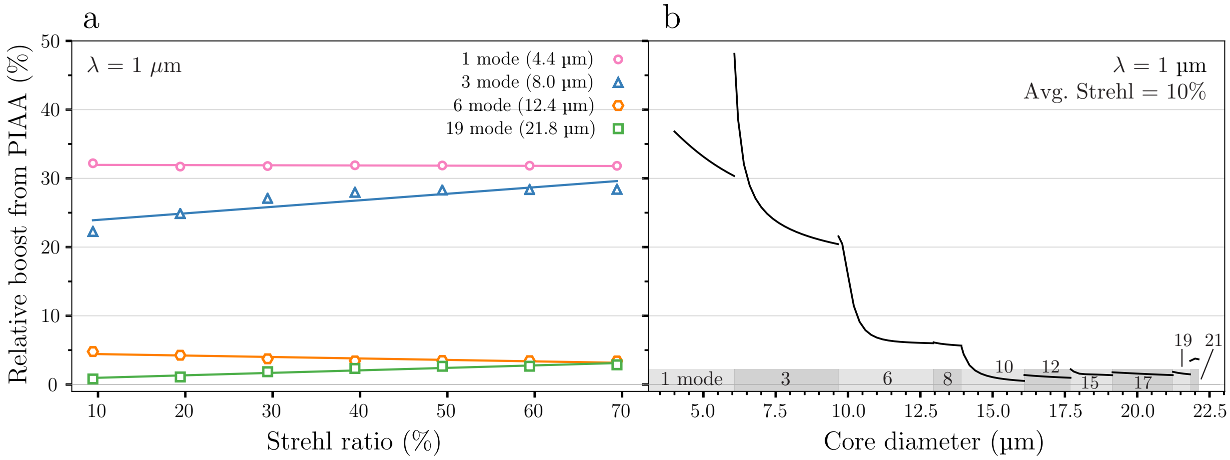

Figure 3.5: **Panel a**: relative coupling efficiency boost offered by Gaussian beam shaping optics as a function of Strehl ratio, for various lantern entrance diameters. Large lanterns show less gain over all tested Strehl ratios. **b**: relative coupling efficiency boost from beam shaping optics as a function of lantern entrance diameter. We approximate the lantern entrance as a step-index MMF. As diameter increases, the boost offered by beam shaping optics decreases, dropping below 10% in the 6-mode region.

3.3.3 Effect of beam-shaping (PIAA) optics

Figure 3.4 shows the absolute coupling gains from beam-shaping optics as a function of Strehl ratio and lantern entrance diameter. We recast these results in terms of the relative gains offered by beam-shaping optics in Fig. 3.5. As indicated in panel a, we see that beam shaping optics offer a significant ∼30% boost to SMF coupling, consistent with results from [JSG17, CJR21]. We additionally see that these beam-shaping optics also offer a significant ∼25% boost to coupling into a 3-port lantern, but offer negligible ∼5% gain for a 6-port lantern. Consequently, we find that the 3-port PL with beam-shaping performs almost as well as a 6-port lantern over the entire range of Strehl ratios (see Fig. 3.4a). Figure 3.5b makes a similar argument, showing that as lantern entrance diameter (and the number of supported modes) increases, the relative boost afforded by beam-shaping optics drops, declining below 5% beyond the 6-mode region, with negligible gains for lanterns with 10+ ports. The reason for this will be described in Section 3.4.

The bump in relative boost in the 21-mode region shown in Fig. 3.5b can be traced to the LP_{13} mode, which in the 21-mode region is very close to its cutoff wavelength and thus extends further into the cladding. This mode shape matches well with the amplitude distribution produced by our beam-shaping optics; a similar effect can be seen at the beginning of the 3 mode region in the same figure, where the relative boost initially spikes because the LP_{11} mode is close to its cutoff wavelength and thus has a larger spatial extent. Such modes can be more sensitive to losses (e.g. due to bending [Mit16]) so it is unclear if these coupling gains are realizable or if they are an artifact of our idealized simulations.

3.3.4 Chromatic behavior

We extend the results from Fig. 3.4a over wavelength in Fig. 3.6, again fixing atmospheric turbulence to give PSFs with an average 10% Strehl at the nominal operating wavelength of 1 μm. Here, we fix the number of PL ports to be equal to the number of modes at the PL entrance, at the reference wavelength $\lambda = 1$ μm. With the exception of the 19-port lantern, this assumption ensures that the lanterns will have at least as many output ports as entrance modes over the y and J bands. Unsurprisingly, we see from Fig. 3.6 that configurations with larger lanterns consistently outperform those with smaller lanterns across the entire range of tested wavelengths. Additionally, as wavelength increases, phase error decreases, and coupling efficiency into lanterns typically increases for all entrance diameters, though there are dips and discontinuous drops as the guided modes at the lantern entrance change shape and eventually disappear. The discontinuous drops in coupling efficiency seen in Fig. 3.6 indicate transition wavelengths across which the number of guided modes at the lantern entrance changes. Importantly, near such wavelengths, we no longer expect coupling efficiency to be an accurate proxy for lantern throughput. The reasoning is as follows: as wavelength increases past a transition wavelength, lantern modes at the waveguide entrance that were originally guiding now become weakly radiative. Normally, we expect that such modes would not be able to carry power; however, if the characteristic length scale

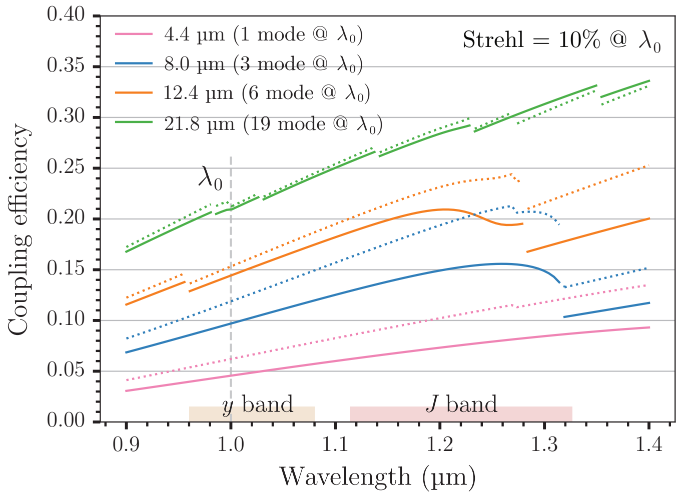

Figure 3.6: Throughput comparisons for the lantern configurations in Fig. 3.4a, now as a function of wavelength. Simulated atmospheric conditions were fixed for an average of 10% Strehl at $\lambda = 1$ μm. Dotted lines show the coupling efficiencies with beam-shaping (PIAA) optics. For each curve, the focal ratio is optimized at the nominal core diameter assuming a wavelength of 1 μm.

for radiative attenuation is small compared to the length of the lantern lead-in, losses are negligible and these radiative modes can still transfer power into the lantern cores. This effect is relatively unimportant for large lanterns (e.g. a few percent difference in coupling efficiency for the 19-port lantern in Fig. 3.6) since such lanterns will distribute power over many modes. However the effect can be significant for few-mode lanterns. In Section 3.3.6, we apply numerical beam propagation to compute throughputs of few-port lanterns in these transition regimes.

3.3.5 Tip-tilt wavefront error

To make our WFE more realistic, we can inject extra tip-tilt in addition to the existing WFE from Kolmogorov turbulence. Similar to Fig. 3.4b, Fig. 3.7 plots the mean coupling efficiency for a PL as a function of entrance diameter, though now under varying amounts of injected tip-tilt error. The focal ratio of the optical system was optimized to match the size of the lantern modes at each core diameter prior to tip-tilt injection. Scenarios without beam-shaping optics (panel a) and with such optics (panel b) are considered. The atmospheric coherence length is tuned to give an average 70% Strehl ratio at a wavelength of 1 μm (actual Strehl of the intensity-averaged image, after tip-tilt injection, will be lower). Figure 3.7 shows that coupling degrades as both the amount of injected tip-tilt increases and as entrance diameter decreases, suggesting that the reduction in coupling performance due to the injected tip-tilt ultimately depends on the relative amount of tip-tilt displacement compared to the size of the lantern entrance. Coupling appears mostly unaffected so long as the RMS tip-tilt displacement is smaller than the lantern entrance radius by more than a factor of 2.

We present an alternative perspective on lantern performance in the presence of tip-tilt in Fig. 3.8, comparing coupling efficiency against Strehl ratio for a lantern entrance diameter of 8 μm in two distinct cases: either tuning Strehl by adjusting atmospheric coherence length or by adjusting the RMS displacement of injected tip-tilt error. Figure 3.8 indicates that

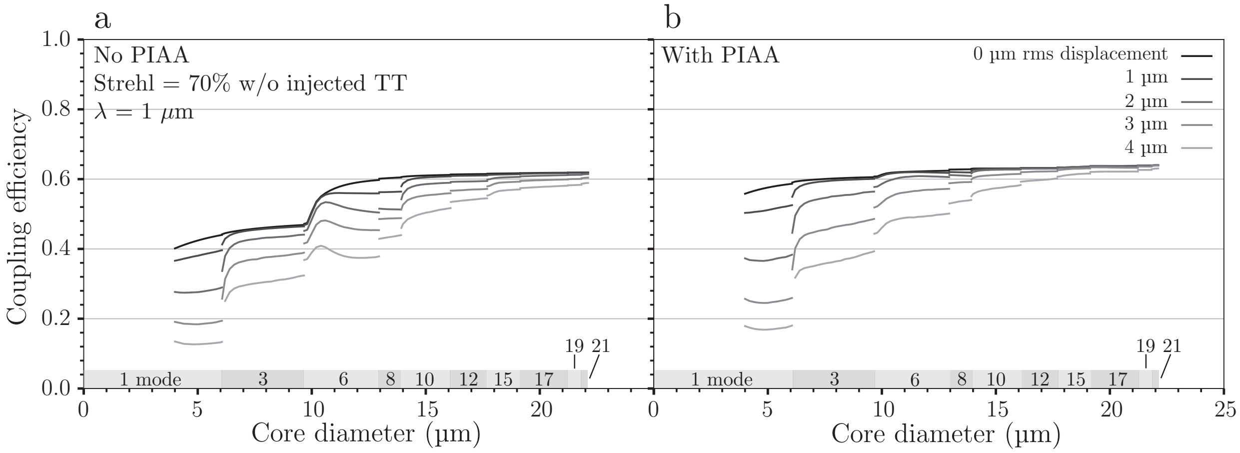

Figure 3.7: Similar to Fig. 3.4b, but the atmospheric coherence length is fixed to give an average Strehl of 70%, after which tip-tilt WFE in the form of random displacements in the focal plane is injected to wavefront realizations. This tip-tilt is added with the intention to approximate instrumental tip-tilt and compensate for potential underestimation of low-order error in our turbulence model.

coupling efficiencies are comparatively better in the latter case, implying that few-mode fibers and PLs are resilient against tip and tilt: for a given amount of WFE, coupling efficiency of the beam into the fiber or lantern improves when there is more power in tip-tilt aberrations vs pure Kolmogorov turbulence. This holds whether or not beam-shaping optics are used. With this in mind, we see that the results of §3.3.1-§3.3.4 and later §3.3.6, which do not include additional tip-tilt, are lower bounds on the coupling gains that can be realized by lanterns in real optical systems. The degree of this effect will depend on lantern size and the specific characteristics of instrumental tip-tilt.

3.3.6 Full beam propagation

Because beam propagation is more computationally expensive than the prior coupling calculations, we apply this method only to study the simplest 3-port lantern; the general trends from this section will be extensible to larger lanterns. Figure 3.9 presents the throughput (not input coupling efficiency) for a 3-port lantern with an 8 μm entrance diameter with

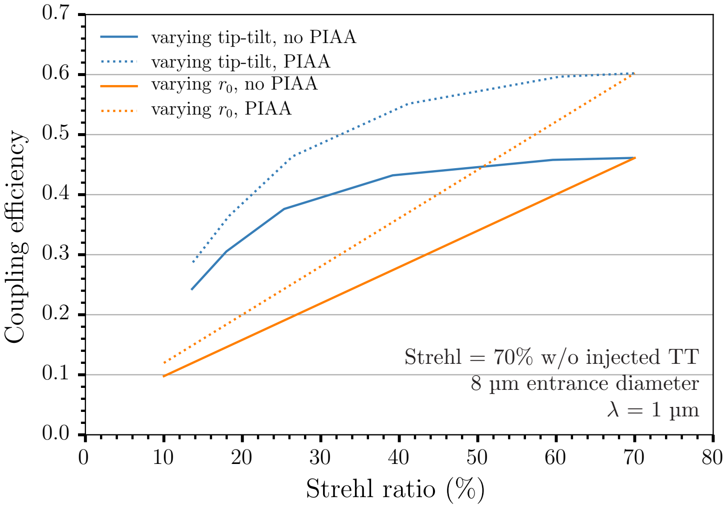

Figure 3.8: Coupling efficiency vs Strehl ratio into an 8 μm diameter (3 mode) fiber for aberrated PSFs. Aberrations can be introduced in multiple ways; we consider two. In the first we assume only single-layer atmospheric turbulence and vary the coherence length r_0. In the second we degrade the PSFs by fixing some r_0 and then injecting additional tip-tilt error in the form of random displacements in the focal plane. In this case, we fixed the r_0 to give 70% average Strehl before injecting tip-tilt. Both with and without beam-shaping (PIAA) optics, MMFs are better able to capture light degraded via the second method.

beam-shaping optics (black curve). In throughput calculations, geometric factors such as lead-in length and taper factor now become relevant; we assume a 4× taper factor, and no lead-in waveguide. To generate the throughput curve in Fig. 3.9, we have numerically computed the lantern modes at the waveguide entrance by backpropagating the fundamental modes of the SMF cores from the lantern output and orthogonalizing the resulting basis. These modes are LP-like, but will no longer necessarily be guiding within the lantern entrance if there are more lantern cores than propagating modes at the lantern's MMF-like input. Calculation of the lantern modes allows us to retain our modal analysis approach and compute lantern throughput from the input electric field distribution at the lantern entrance. As a result, we bypass the need to forward-propagate multiple realizations of the input field through the guide, saving on computation time. As seen in Fig. 3.9, lantern throughput agrees well with the coupling efficiency into a similarly sized step-index fiber until the propagating wavelength approaches the $3 \rightarrow 1$ mode transition wavelength at ~ 1.32 μm. Beyond 1.32 μm, the lantern throughput exceeds step-index fiber coupling efficiency by almost a factor of 2. This is because an 8 μm diameter step-index fiber with our assumed core and cladding materials becomes single-mode beyond $\lambda = 1.3$ μm; however, an 8 μm diameter lantern entrance retains weakly attenuating radiative modes which can help transfer power to the lantern cores.

As such, over transition wavelengths where radiative lantern modes are active, the length of the lead-in waveguide becomes a relevant factor in lantern design. To show this effect, Fig. 3.9 plots throughput curves for lanterns that are similar to our 3-port configuration with the exception of differing lead-in lengths (note these throughputs were only computed for $\lambda > 1.26$ μm to save on computation time). Throughputs are now computed using the "brute-force" method of numerically propagating all input field realizations to the lantern's output. Note that we can no longer use backpropagation when attenuation through the lantern is non-negligible (as can be the case for lead-in waveguides): the departure of light from the simulation zone makes the operation of numerical propagation non-invertible. Figure 3.9

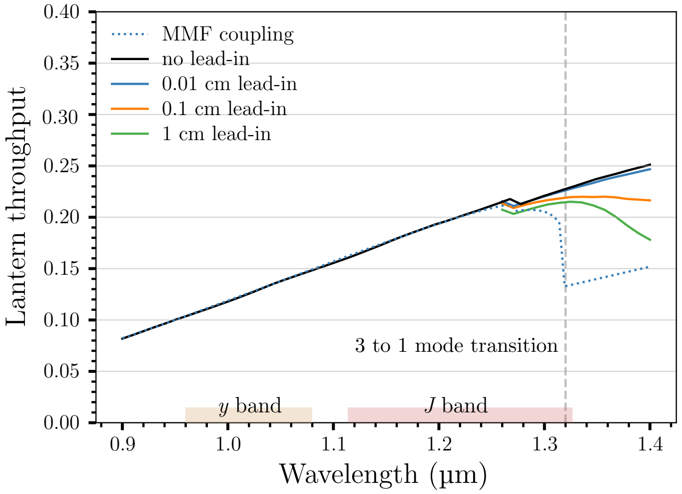

Figure 3.9: Throughput of an 8 μm diameter lantern against wavelength, for various lantern lead-in lengths. WFE here is wholly generated from Kolmogorov turbulence with an average 10% Strehl. The dotted curve denotes coupling efficiency into an 8 μm diameter circular step-index fiber; all other curves denote lantern coupling efficiencies and were generated using beam propagation. Beyond 1.32 μm, the step-index fiber becomes single-mode, and coupling efficiency drops precipitously. However lanterns with short lead-in lengths do not experience such a drop due to radiative modes. Note that throughputs for non-zero lead-in lengths were only computed and plotted for $\lambda > 1.26$ μm to save on computation time.

verifies that longer lead-in lengths result in greater dampening in lantern throughput at wavelengths beyond the transition wavelength of 1.32 μm, due to increased attenuation of the radiative modes (note that this attenuation is being captured by the numerical beam propagation itself, not some assumed fiber model). The magnitude of attenuation also increases as the propagation wavelength increases beyond the transition wavelength.

The reduction in throughput due to attenuation of the radiative modes will also depend on the lantern taper factor, which sets the degree by which the embedded SMF cores in the lantern shrink from output to input. Figure 3.10 plots the degree of attenuation for the LP_{11}-like radiative mode in the lead-in waveguide of a 3-mode (8 μm entrance diameter) lantern, as a function of taper factor t; since the SMF cores at the output end of the lantern are 4.4 μm in diameter, the diameter of the tapered-down cores embedded in this lead-in waveguide can be computed as 4.4 μm / t. The propagation wavelength is taken to be 1.4 μm. Attenuation is quantified in terms of the attenuation length d, the distance over which power in the radiative mode drops by a factor of e. The attenuation length d can be computed from

$$d = \frac{\lambda}{4\pi \mathrm{Im}\left(n_{\mathrm{eff}}\right)} \tag{3.2}$$

where $n_{\text{ eff}}$ is the (complex) effective index of refraction for the radiative mode and Im denotes taking the imaginary part. In turn, we compute the effective index using FemSIM [Syn], a finite-element mode solver that can compute the effective indices of radiative modes. From Fig. 3.10, we see that the attenuation of this radiative mode increases as taper factor decreases and the embedded SMF cores at the lantern entrance become larger. The dependence between the attenuation length and taper factor is asymptotic, climbing rapidly to infinity as the embedded SMF cores at the lantern's entrance approach the minimum required size to become guiding. Finally, we note from Fig. 3.10 that the attenuation length for a 3-port lantern with a taper factor of 4 is ~3 cm. This is order-of-magnitude consistent with Fig. 3.9, computed using our numerical beam propagator, where we see that the 3-port, 4× tapered lantern begins to show significant attenuation at $\lambda = 1.4$ μm when the lead-in length reaches 1 cm.

3.4 Discussion/Future directions

Our results from Fig. 3.4a recover the linearly increasing trend of coupling efficiency with Strehl ratio, consistent with prior works [JSG17]. This is sensible: the Strehl ratio is directly proportional to peak PSF intensity, which in turn is directly proportional to the amount of PSF power that couples into the lantern modes. The coupling boost provided by the beam-shaping optics also increases with Strehl ratio, since beam-shaping optics only correct for average pupil illumination and not WFE. The boost is greatest for SMFs, which in conjunction with beam-shaping actually approach the performance of 6-port lanterns at high

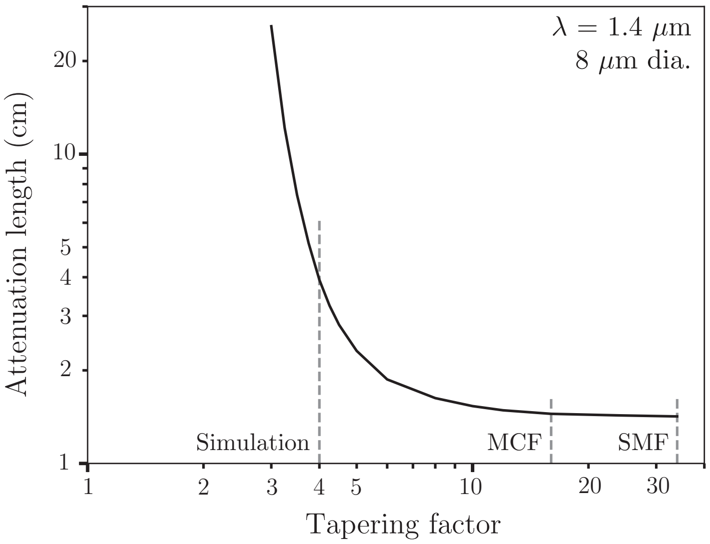

Figure 3.10: Attenuation strength vs. taper factor for one of the radiative LP_{11}-like modes ($\lambda = 1.4\ \mu$m) in a lead-in waveguide for a 3-port, 8 μm diameter lantern, computed with FEMSIM [Syn]. Attenuation strength is quantified as the length over which the total power of the radiative LP_{11}-like mode drops by a factor of e. The required taper factors to construct the assumed lantern from a bundle of SMFs (125 μm diameter cladding) and a MCF (125 μm diameter cladding overall), as well as the taper factor of the lantern simulated in Fig. 3.9, are annotated with vertical dashed lines.

Strehl. This behavior is expected: at high Strehl, beam-shaping moves the bulk of light at the focal plane into the PSF core, promoting coupling into the single Gaussian-like mode of SMFs. In contrast, larger lanterns have access to other higher-order modes which can couple light away from the PSF core, such as light in the Airy rings, reducing the utility of beam-shaping. In the case of the 6-port lantern, the presence of the LP_{02} mode allows such lanterns to efficiently couple light from the Airy ring. However, in the low-Strehl regime PLs are still preferable. This result is reaffirmed in the presence of increased tip-tilt error: the extra size and modality of lanterns makes their coupling efficiencies less sensitive to displacement of the PSF core. In fact, Fig. 3.8 shows that for tip-tilt-varied WFE, lantern coupling efficiency scales roughly with the square root of the Strehl ratio. In contrast, coupling scales linearly with Strehl when WFE is purely generated from AO-filtered Kolmogorov turbulence.

Unsurprisingly, we also find that larger lanterns couple more light than smaller lanterns at a given Strehl ratio, because the number of supported modes increases with lantern entrance diameter. However, in the absence of beam-shaping, we show in Fig. 3.6 that gains from increasing lantern size begin to diminish beyond the 6-port lantern: the absolute coupling gain from a 6- to a 19-port lantern is roughly the same as the gain from a 3- to a 6-port lantern. This result is reiterated by Fig. 3.4b, and stems from the fact that low-order modes such as LP_{01}, LP_{11}, and LP_{02} play a larger role in accepting injected light than higher-order modes, especially when WFE is dominated by low-order aberrations. The LP_{11} modes are well-suited for dealing with tip-tilt aberrations, and account for the increase in coupling efficiency from a SMF to a 3-port lantern. Similarly, in the absence of beam-shaping, the increase in coupling efficiency from a 3 to 6-port lantern is primarily due to the inclusion of the LP_{02} mode (Fig. 3.3g), which can accept light from the Airy ring.

In the presence of Gaussian beam-shaping optics, diminishing returns begin even sooner: as seen in Fig. 3.4a, a 3-port lantern with beam-shaping performs almost as well as a 6-port lantern with or without beam-shaping. As mentioned above, such lenses have a significant effect for three-port lanterns (and SMFs) since such devices have no other way to efficiently accept light from the Airy ring; in contrast, such lenses are less useful for 6-port and larger lanterns, which have access to the LP_{02} mode. However, it is important to recognize that these results are specific to our simple Gaussian-remapping implementation of beam-shaping. Lenses with more advanced remapping functions, tailored to individual lantern mode geometries and imaging systems, may boost coupling more and warrant further exploration. Also of interest are optical systems with two DMs, allowing for simultaneous phase and amplitude modulation without any requirement for radial symmetry. Careful configuration of such setups may boost coupling performance even further.

3.4.1 Signal-to-noise considerations

While lantern coupling efficiencies are useful for measuring the raw performance of PLs, the signal-to-noise ratio (SNR) is often a more practical metric. Full SNR calculations will

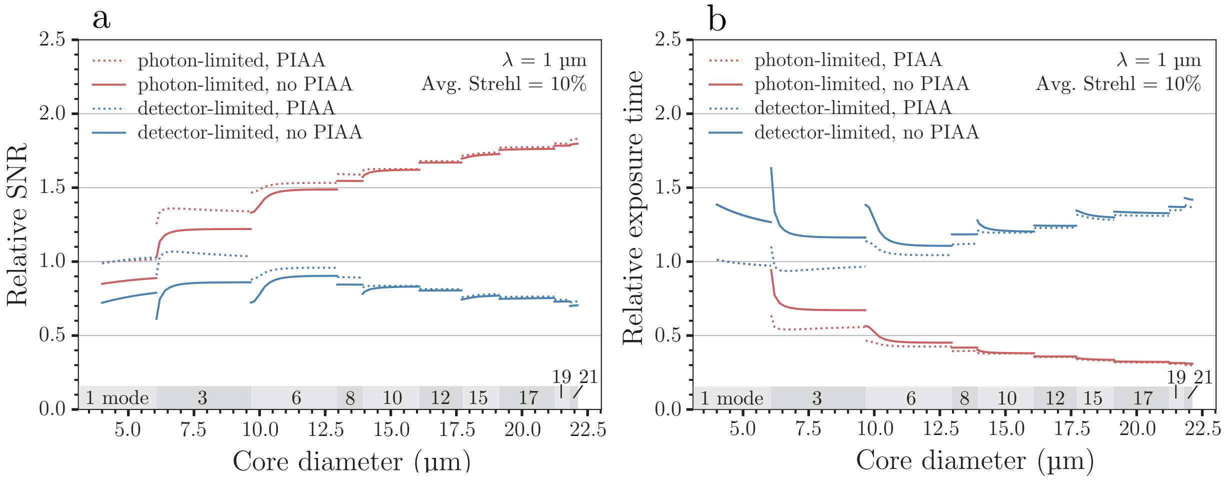

Figure 3.11: a: SNR of a lantern-fed instrument as a function of entrance diameter at $\lambda = 1$ μm, in the photon-limited (red) and detector-limited (blue) noise regimes. Turbulence was fixed to give an average 10% Strehl. All lantern ports are assumed to be in the same noise regime, and noise is assumed to be spread evenly among the lantern ports. SNR values are relative to the SNR of an SMF with beam-shaping optics. b: the required exposure time for a lantern-fed instrument to attain some reference SNR, relative to the time it would take an SMF with beam-shaping optics to attain that same SNR, as a function of lantern entrance diameter. Wavelength and Strehl are the same as in a.

depend on instrument specifics and the observing context, so we only consider two limiting cases: the fully photon-limited case, where all lantern ports are dominated by photon noise, and the fully detector-limited case, where all lantern ports are dominated by detector (read, dark current) noise. Figure 3.11a presents the SNR as a function of lantern entrance diameter in these two limits. Atmospheric turbulence is fixed to provide a 10% Strehl ratio. Here, lantern SNRs are normalized against the SNR of an SMF with beam-shaping optics. In the fully photon-limited (Poissonian) regime, SNR scales as the square root of the lantern coupling efficiency μ: $\mathrm{SNR}_{\mathrm{photon}} \propto \sqrt{\mu}$. In the fully detector-limited regime, the noise variance will be linearly proportional to the number of lantern ports n_{c}, penalizing larger lanterns: $\mathrm{SNR}_{\mathrm{detector}} \propto \mu/\sqrt{n_{\mathrm{c}}}$. As a result, we see that significant gains in SNR can be made in adopting larger lanterns only if we are in the fully photon-limited regime, with 19-port lanterns reaching an $\sim 80\%$ gain over an SMF. Conversely, in the fully detector-limited regime, lantern SNRs are about the same as or worse than that of an SMF. Fig. 3.11b presents the same argument from the perspective of exposure times. Here, we plot the time required by a lantern-fed instrument to hit some reference SNR, relative to the time it would take an SMF with beam-shaping optics to hit that same SNR. Figure 3.11b shows that in the photon-limited regime, exposure times for a 3-port lantern with beam-shaping optics are around half that of a similarly configured SMF, and that exposure time halves again going from 3 to 19 ports. In the detector-limited regime, exposure times are increased by a factor of $\sqrt{n_c}$ over the photon-limited regime, so that lanterns generally require more time than an SMF with beam-shaping optics. As such, at 10% Strehl ratio, if the signal from the lantern ports will be in the detector-limited regime, then an SMF should be used instead. If the output ports will be photon-noise limited, then any lantern can be used and higher port counts will result in greater overall SNR, but only if the signal in each port will remain dominated by photon-noise. In reality, lantern behavior will likely be in between the two limiting cases presented above. If most but not all lantern outputs are in the photon-limited regime, SNR gains from larger lanterns will drop off more quickly than in the fully photon-limited case, strengthening the argument for few-port lanterns. The impact of detector noise may also be lessened with mode-selective lanterns [LFS14], which can direct light so that the bulk of the signal is concentrated within a few high-flux ports, as opposed to being split evenly among many ports.

A system level consideration that arises when selecting the number of ports to use is the availability of pixels. For a given set of spectrometer parameters (resolving power and bandwidth), the number of pixels required will increase linearly with the number of SMF ports. For very high resolving power applications (>100,000), which typically desire large bandwidths (y-K band for example) as well, port counts using even the largest format NIR detectors will be limited to ∼6. Going beyond the footprint of a single detector and requiring a secondary detector or a larger version is typically a costly upgrade that most instruments cannot absorb. In addition, splitting the light across more ports lowers the overall flux per port, requiring longer integration times to regain SNR. As large telescope time is costly, this can also be an undesirable choice. For these reasons, few-port lanterns can be more attractive. Regardless, as Figure 3.11 shows, if both pixel constraints and detector noise are not an issue, high mode-count lanterns can provide significant gains in SNR. The cost of introducing adding more ports may also be offset due the increased tip-tilt resilience of larger lanterns. In this case, a large mode-count PL could used instead of high-performance tip-tilt correction, resulting in a trade in system complexity, rather than a strict increase.

Finally, it is clear that the beam-shaping optics can work in parallel to the PL to boost coupling. Given that beam-shaping optics can be manufactured for <$5k, such optics can be readily used in conjunction with a lantern to help reduce the required number of pixels. For instance, as seen in Fig. 3.6, performance for a 3-port lantern in conjunction with beam-shaping is similar to that of a 6-port lantern with or without beam-shaping optics. By using a 3-port+beam-shaping optics configuration instead of a 6-port lantern, the number of detector pixels can be reduced by a factor of 2! This means that the wavelength coverage could be extended, or large gaps between the traces could be used to minimize cross-talk to improve overall precision. Although the relative cost of the beam-shaping optics is similar to the PL, the addition of this optic can have a tremendous impact compared to the much more expensive detector pixels.

3.4.2 Optimizing coupling efficiency

After deciding on a targeted number of lantern modes and beam-shaping configuration, the next focus in lantern design is the maximization of input coupling efficiency. This maximization requires the fulfillment of two separate conditions: firstly, the physical extents of the lantern modes and PSF must match; secondly, the beam profile shapes should also match. Condition 1 sets the optimal camera focal ratio (and hence PSF scale) for a given entrance diameter. On the other hand, condition 2 sets the lantern entrance diameter itself. As seen in Fig. 3.4b (computed assuming optimal focal ratios for each entrance diameter), over diameter ranges where the number of available modes is constant, coupling efficiency is still variable, implying in turn that the field structure of the lantern modes depends on entrance diameter. This is an analytically known effect, and is most clearly displayed by the fundamental LP_{01} mode, which appears more Lorentzian-like at small entrance diameters, parabolic at large entrance diameters, and Gaussian-like in between. For the same reason, coupling efficiencies can actually drop as core diameter increases. This effect appears to be the strongest in the presence of Gaussian beam-shaping optics (see the 3-mode lantern w/ PIAA in Figure 3.4b), presumably because increasing core diameter causes mode shapes to deviate further from a Gaussian profile. As such, lantern entrance diameters should be selected to promote lantern modes that are most similar in shape to the PSF. The optimal entrance diameter can be determined via a brute-force search over a range of entrance diameters, like in Fig. 3.4b, ensuring that focal ratio is optimized at each tested diameter (otherwise coupling effects due to mismatches in size and shape will be entangled). Ultimately, while effects involving lantern mode shape are not large, they are worth considering during system design and manufacture.

The last parameters to consider in lantern design are the length of the lead-in waveguide and the taper factor, relevant in regimes where the number of lantern ports strictly exceeds the number of propagating modes at the lantern's entrance (e.g. Fig. 3.9 for wavelengths beyond 1.3 μm). In such regimes, weakly attenuating radiative modes can still help transfer light to the lantern cores, promoting lantern throughputs that exceed the coupling efficiency into the guided modes at the lantern's entrance. A similar experimental effect is mentioned in Birks et al. [BGY15], which notes that prior lantern experiments have found throughputs that were better than expected in wavelength regimes where lantern operation was no longer strictly efficient. As a potential explanation Birks et al. [BGY15] proposed that experimental setups may not have excited all available modes in their lanterns; we offer the alternative explanation that instead the presence of weakly radiative modes at a lantern's entrance may be boosting the effective modality of the lanterns. In turn, lantern designs in wavelength regimes near and beyond transition wavelengths should carefully consider lead-in length and taper factor: the tuning of these parameters presents an avenue through which the bandwidth for efficient lantern operation can be maximized. Optimal performance occurs when there is no lead-in, corresponding to direct injection of light into the lantern's transition zone: this approach minimizes the attenuation of any light travelling in the radiative modes. From a manufacturing perspective, our result implies that lanterns should be cleaved as close to the transition zone as possible. For the particular case of the 3-port lantern in Fig. 3.9, we have shown that such a lantern can maintain its modality over the entirety of the y and J bands if the lead-in waveguide length is less than 0.01 cm.

As seen in Fig. 3.10, attenuation of the radiative lantern modes can also be reduced by lowering the lantern taper factor, which in turn increases the size of the embedded SMF cores at the lantern entrance and hence lowers the imaginary effective refractive index at that location. In practice, obtaining smaller taper factors with SMF-based lanterns often involves etching down the claddings of SMFs or pre-tapering the individual SMFs before stacking and tapering the bundle. Otherwise, taper factors will tend to be quite large: for instance, a 3-core lantern constructed from a bundle of 3 SMFs with cladding diameters of 125 μm requires a taper factor of ∼34 to yield a lantern entrance diameter of 8 μm. Taper factors can also be lowered somewhat by instead constructing lanterns from MCFs. For reference, commercially available 4-core MCFs such as Chiral Photonics MCF-004_1 or Fibercore SM-4C1500 have typical cladding diameters of ~ 125 μm, requiring only a ∼16× taper to produce an 8 μm entrance diameter. The required taper factors to construct a 3-port lantern from a 3-core MCF will likely be similar. Beyond radiative attenuation, taper factor will also affect lantern throughput in another more indirect way: the presence of large residual SMF cores in the lead-in waveguide (corresponding to a small taper factor) will change the shape of the lantern modes to diverge from that of the LP modes, impacting the initial coupling efficiency into the lantern. We do not treat this higher-order effect, and leave it for future work.

Finally, radiative mode attenuation may also be reduced by shortening a lantern's transition zone, though this effect is minor since light only has to travel through a fraction of the transition zone before the lantern cladding becomes large enough to support the previously radiative LP modes. However, care must also be taken not to make the lantern too short, otherwise the transformation of light from input to output will not be adiabatic, and performance will drop [BGY15]. The tension between adiabaticity and radiative mode attenuation implies that lanterns have an optimal taper length (see [CMH18]), which should be solved for during lantern design.

3.4.3 Future directions

The next step will be to verify the results of our simulations on-sky. While we have conducted an initial investigation on the interaction between tip-tilt errors and PLs, real wavefronts have extra low-order error components that are not captured by our simulations. On-sky verification will also allow us to extend our results beyond the idealized, linearly tapered and perfectly circular lanterns assumed in this work. While we believe our simulations accurately capture the first-order behavior of PLs, these higher-order effects will need to be measured before PLs can be fully integrated into the next generation of instruments.

3.5 Conclusion

In this work we have presented numerical simulations characterizing the potential performance gains of lantern-based fiber injectors and their interaction with beam-shaping optics over a range of Strehl ratios, wavelengths, and lantern geometries. Lantern performance was first characterized in terms of coupling efficiency into the lantern input. This proxy metric sets the first-order behavior of the PL and serves as a good approximation for overall lantern throughput over wavelengths where the number of guided modes at the lantern's entrance remains constant. We find that PLs show a linear scaling in coupling efficiency with Strehl ratio when WFE is solely generated from Kolmogorov turbulence. At low Strehl ($\sim$10%), and in the absence of beam-shaping, coupling efficiency into the simplest 3-port (8 μm entrance diameter) lantern is around twice that of an SMF. A 6-port lantern (12.4 μm entrance diameter) and 19-port lantern (21.8 μm) triple and quadruple the performance of an SMF, respectively. When lanterns are combined with Gaussian beam-shaping (PIAA) optics, we find that the relative boost offered by such optics diminishes rapidly, such that 3-port lanterns and 6-port lanterns now perform similarly at low Strehl. This drop-off occurs because lanterns with 6 modes or more have access to the LP_{02} mode, which can accept light from the Airy ring, in turn lessening the utility of beam-shaping. At high Strehl ($\sim$ 70%), in the presence of beam-shaping, lanterns and SMFs perform similarly. We also show that PLs are resilient to tip-tilt aberrations, scaling with the square root of Strehl when WFE is modulated via injected tip-tilt. As such, in real optical systems, where WFE has additional low-order components from NCPAs and mechanical vibrations, the gains from PLs may be even greater. From our analysis of lantern coupling efficiencies, we find that for operation in the astronomical y and J bands when Strehl is low ($\sim$ 10%), 3-port lanterns in conjunction with beam-shaping strike a good balance between performance and pixel count, offering nearly the same coupling efficiency as a 6-port configuration with half the lantern outputs. If pixels are not a constraint, 19-port lanterns can provide even greater coupling efficiencies and better resilience to tip-tilt jitter, doubling that of a 3-port lantern with beam-shaping optics at $\lambda = 1$ μm, as long as all ports have sufficient flux to be dominated by photon noise.

Over wavelengths where the number of guided modes at the lantern entrance drops, we find that the performance of PLs may not necessarily suffer. In fact, past such wavelengths, lanterns with sufficiently short lead-in waveguide lengths will act as to retain their modality due to the presence of weakly radiative modes at the lantern entrance. To promote this behavior and ensure optimal lantern performance across large wavelength bands where the number of guided modes at the lantern entrance changes, we find that lantern lead-in lengths should be minimized in order to keep radiative losses low. Maximal performance occurs when telescope light is directly injected into the transition zone of the lantern, and when lantern taper factors are small. Future instruments that adopt PLs should carefully consider these geometric factors to ensure optimal performance over the widest range of wavelengths.

REFERENCES

[AHC18] Th Anagnos, R J Harris, M K Corrigan, A P Reeves, M J Townson, D G MacLachlan, R R Thomson, T J Morris, C Schwab, and A Quirrenbach. "Simulation and optimisation of an astrophotonic reformatter." *Monthly Notices of the Royal Astronomical Society*, **478**(4):4881–4889, 05 2018.

[BEL11] J. Bland-Hawthorn, S. C. Ellis, S. G. Leon-Saval, R. Haynes, M. M. Roth, H. G. Löhmannsröben, A. J. Horton, J. G. Cuby, T. A. Birks, and J. S. Lawrence. "A complex multi-notch astronomical filter to suppress the bright infrared sky." *Nature Communications*, **2**:581, Dec 2011.

[BGY15] T. A. Birks, I. Gris-Sánchez, S. Yerolatsitis, S. G. Leon-Saval, and R. R. Thomson. "The photonic lantern." *Adv. Opt. Photon.*, **7**(2):107–167, Jun 2015.

[BLR10] Joss Bland-Hawthorn, Jon Lawrence, Gordon Robertson, Sam Campbell, Ben Pope, Chris Betters, Sergio Leon-Saval, Tim Birks, Roger Haynes, Nick Cvetojevic, and Nem Jovanovic. "PIMMS: photonic integrated multimode microspectrograph." In Ian S. McLean, Suzanne K. Ramsay, and Hideki Takami, editors, *Ground-based and Airborne Instrumentation for Astronomy III*, volume 7735 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 77350N, July 2010.

[BW01] Jacques Baudrand and Gordon A. H. Walker. "Modal Noise in High-Resolution, Fiber-fed Spectra: A Study and Simple Cure." *Publications of the Astronomical Society of the Pacific*, **113**(785):851–858, jul 2001.

[CJG17] N. Cvetojevic, N. Jovanovic, S. Gross, B. Norris, I. Spaleniak, C. Schwab, M. J. Withford, M. Ireland, P. Tuthill, O. Guyon, F. Martinache, and J. S. Lawrence. "Modal noise in an integrated photonic lantern fed diffraction-limited spectrograph." *Opt. Express*, **25**(21):25546–25565, Oct 2017.

[CJR21] Benjamin Calvin, Nemanja Jovanovic, Garreth Ruane, Jacklyn Pezzato, Jennah Colborn, Daniel Echeverri, Tobias Schofield, Michael Porter, J. Kent Wallace, Jacques-Robert Delorme, and Dimitri Mawet. "Enhancing Direct Exoplanet Spectroscopy with Apodizing and Beam Shaping Optics." *Publications of the Astronomical Society of the Pacific*, **133**(1020):024503, feb 2021.

[CMH18] Mark K. Corrigan, Timothy J. Morris, Robert J. Harris, and Theodoros Anagnos. "Demonstration of a photonic lantern low order wavefront sensor using an adaptive optics testbed." In Laird M. Close, Laura Schreiber, and Dirk Schmidt, editors, *Adaptive Optics Systems VI*, volume 10703, pp. 1313 – 1320. International Society for Optics and Photonics, SPIE, 2018.

[EBB18] S. C. Ellis, S. Bauer, C. Bacigalupo, J. Bland -Hawthorn, J. J. Bryant, S. Case, R. Content, T. Fechner, D. Giannone, R. Haynes, E. Hernandez, A. J. Horton, U. Klauser, J. S. Lawrence, S. G. Leon-Saval, E. Lindley, H. G. Löhmannsröben, S. S. Min, N. Pai, M. Roth, K. Shortridge, L. Waller, Pascal Xavier, and Ross Zhelem. "PRAXIS: an OH suppression optimised near infrared spectrograph." In *Proceedings of the SPIE*, volume 10702 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 107020P, Jul 2018.

[Guy03] Guyon, O. "Phase-induced amplitude apodization of telescope pupils for extrasolar terrestrial planet imaging." *A & A*, **404**(1):379–387, 2003.

[HMC15] R. J. Harris, D. G. MacLachlan, D. Choudhury, T. J. Morris, E. Gendron, A. G. Basden, G. Brown, J. R. Allington-Smith, and R. R. Thomson. "Photonic spatial reformatting of stellar light for diffraction-limited spectroscopy." *Monthly Notices of the Royal Astronomical Society*, **450**(1):428–434, 04 2015.

[JCP20] N. Jovanovic, B. Calvin, M. Porter, T. Schofield, J. Wang, M. Roberts, G. Ruane, J. K. Wallace, R. Bartos, J. Pezzato, J. Colborn, J. R. Delorme, D. Echeverri, D. Mawet, C. Z. Bond, S. Cetre, S. Lilley, S. Ragland, P. Wizinowich, and R. Jensen-Clem. "Enhanced high-dispersion coronagraphy with KPIC phase II: design, assembly and status of sub-modules." In Christopher J. Evans, Julia J. Bryant, and Kentaro Motohara, editors, *Ground-based and Airborne Instrumentation for Astronomy VIII*, volume 11447, pp. 970 – 981. International Society for Optics and Photonics, SPIE, 2020.

[JG94] Erik M. Johansson and Donald T. Gavel. "Simulation of stellar speckle imaging." In James B. Breckinridge, editor, *Amplitude and Intensity Spatial Interferometry II*, volume 2200, pp. 372 – 383. International Society for Optics and Photonics, SPIE, 1994.

[JSC16a] N. Jovanovic, C. Schwab, N. Cvetojevic, O. Guyon, and F. Martinache. "Enhancing Stellar Spectroscopy with Extreme Adaptive Optics and Photonics." *Publications of the Astronomical Society of the Pacific*, **128**(970):121001, nov 2016.

[JSC16b] N. Jovanovic, C. Schwab, N. Cvetojevic, O. Guyon, and F. Martinache. "Enhancing Stellar Spectroscopy with Extreme Adaptive Optics and Photonics." *Publications of the Astronomical Society of the Pacific*, **128**(970):121001, December 2016.

[JSG17] Jovanovic, N., Schwab, C., Guyon, O., Lozi, J., Cvetojevic, N., Martinache, F., Leon-Saval, S., Norris, B., Gross, S., Doughty, D., Currie, T., and Takato, N. "Efficient injection from large telescopes into single-mode fibres: Enabling the era of ultra-precision astronomy." *A & A*, **604**:A122, 2017.

[LFS14] Sergio G. Leon-Saval, Nicolas K. Fontaine, Joel R. Salazar-Gil, Burcu Ercan, Roland Ryf, and Joss Bland-Hawthorn. "Mode-selective photonic lanterns for space-division multiplexing." *Opt. Express*, **22**(1):1036–1044, Jan 2014.

[Lin20] J. Lin. "Lightbeam." Astrophysics Source Code Library, February 2020. ascl:2102.006.

[MAH16] Juan Montoya, Chris Aleshire, Christopher Hwang, Nicolas K. Fontaine, Amado Velázquez-Benítez, Dale H. Martz, T.Y. Fan, and Dan Ripin. "Photonic lantern adaptive spatial mode control in LMA fiber amplifiers." *Opt. Express*, **24**(4):3405–3413, Feb 2016.

[MFK19] Dimitri Mawet, Michael Fitzgerald, Quinn Konopacky, Charles Beichman, Nemanja Jovanovic, Richard Dekany, David Hover, Eric Chisholm, David Ciardi, Etienne Artigau, Ravinder Banyal, Thomas Beatty, Bjorn Benneke, Geoffrey A. Blake, Adam Burgasser, Gabriela Canalizo, Guo Chen, Tuan Do, Greg Doppmann, Rene Doyon, Courtney Dressing, Min Fang, Thomas Greene, Lynne Hillenbrand, Andrew Howard, Stephen Kane, Tiffany Kataria, Eliza Kempton, Heather Knutson, Takayuki Kotani, David Lafreniere, Chao Liu, Shogo Nishiyama, Gajendra Pand ey, Peter Plavchan, Lisa Prato, S. P. Rajaguru, Paul Robertson, Colette Salyk, Bunei Sato, Everett Schlawin, Sujan Sengupta, Thirupathi Sivarani, Warren Skidmore, Motohide Tamura, Hiroshi Terada, Gautam Vasisht, Ji Wang, and Hui Zhang. "High-resolution Infrared Spectrograph for Exoplanet Characterization with the Keck and Thirty Meter Telescopes." *arXiv e-prints*, p. arXiv:1908.03623, Aug 2019.

[MHM17] Juan Montoya, Christopher Hwang, Dale Martz, Christopher Aleshire, T. Y. Fan, and Daniel J. Ripin. "Photonic lantern kW-class fiber amplifier." *Opt. Express*, **25**(22):27543–27550, Oct 2017.

[Mit16] Fedor Manuel Mitschke. *Fiber optics physics and technology.* Springer., 2 edition, 2016.

[NWB20] B.R.M. Norris, J. Wei, C.H. Betters, A. Wong, and S.G. Leon-Saval. "All-fibre focal-plane wavefront sensor." In *14th Pacific Rim Conference on Lasers and Electro-Optics (CLEO PR 2020)*, p. PDP_1. Optical Society of America, 2020.

[OTW15] Ibrahim Ozdur, Paul Toliver, and T. K. Woodward. "Photonic-lantern-based coherent LIDAR system." *Opt. Express*, **23**(4):5312–5316, Feb 2015.

[PBG13] Peter P. Plavchan, M. Bottom, P. Gao, J. K. Wallace, B. Mennesson, D. Ciardi, S. Crawford, S. Lin, C. Beichman, C. Brinkworth, J. Johnson, C. Davison, R. White, G. Anglada-Escude, K. von Braun, G. Vasisht, L. Prato, S. Kane, A. Tanner, B. Walp, and S. Mills. "Precision near-infrared radial velocity instrumentation II: noncircular core fiber scrambler." In Stuart Shaklan, editor, *Techniques and Instrumentation for Detection of Exoplanets VI*, volume 8864 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 88640G, September 2013.

[PHR18] Emiel H. Por, Sebastiaan Y. Haffert, Vikram M. Radhakrishnan, David S. Doelman, Maaike van Kooten, and Steven P. Bos. "High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator." In Laird M. Close, Laura Schreiber, and Dirk Schmidt, editors, *Adaptive Optics Systems VI*, volume 10703, pp. 1112 – 1125. International Society for Optics and Photonics, SPIE, 2018.

[SPR15] Srikar Srinath, Lisa A. Poyneer, Alexander R. Rudy, and S. Mark Ammons. "Computationally efficient autoregressive method for generating phase screens with frozen flow and turbulence in optical simulations." *Opt. Express*, **23**(26):33335–33349, Dec 2015.

[Syn] Synopsys. "FEMSIM." https://www.synopsys.com/photonic-solutions/rsoft-photonic-device-tools.

[TBL11] R. R. Thomson, T. A. Birks, S. G. Leon-Saval, A. K. Kar, and J. Bland-Hawthorn. "Ultrafast laser inscription of an integrated photonic lantern." *Opt. Express*, **19**(6):5698–5705, Mar 2011.

[TEB13] Christopher Q. Trinh, Simon C. Ellis, Joss Bland-Hawthorn, Jon S. Lawrence, Anthony J. Horton, Sergio G. Leon-Saval, Keith Shortridge, Julia Bryant, Scott Case, Matthew Colless, Warrick Couch, Kenneth Freeman, Hans-Gerd Löhmannsröben, Luke Gers, Karl Glazebrook, Roger Haynes, Steve Lee, John O'Byrne, Stan Miziarski, Martin M. Roth, Brian Schmidt, Christopher G. Tinney, and Jessica Zheng. "GNOSIS: The first instrument to use fiber Bragg gratings for OH suppression." *The Astronomical Journal*, **145**(2):51, jan 2013.

[TVL19] Sarah A. Tedder, Brian E. Vyhnalek, Sergio Leon-Saval, Christopher Betters, Bert Floyd, Jeremy Staffa, and Robert Lafon. "Single-mode fiber and few-mode fiber photonic lanterns performance evaluated for use in a scalable real-time photon counting ground receiver." In Hamid Hemmati and Don M. Boroson, editors, *Free-Space Laser Communications XXXI*, volume 10910, pp. 69 – 78. International Society for Optics and Photonics, SPIE, 2019.

[WZA18] Ning Wang, J. C. Alvarado Zacarias, J. Enrique Antonio-Lopez, Z. Sanjabi Eznaveh, Cedric Gonnet, Pierre Sillard, Sergio Leon-Saval, Axel Schülzgen, Guifang Li, and Rodrigo Amezcua-Correa. "Transverse mode-switchable fiber laser based on a photonic lantern." *Opt. Express*, **26**(25):32777–32787, Dec 2018.

[YSS97] J. Yamauchi, J. Shibayama, M. Sekiguchi, and H. Nakano. "Finite-difference beam propagation method based on the generalized Douglas scheme for a nonuniform grid." *IEEE Photonics Technology Letters*, **9**(1):67–69, 1997.

CHAPTER 4

Experimental measurements of AO-fed photonic lantern coupling efficiencies

This chapter builds on Chapter 3 with experimental validations of the predicted coupling gains for PL-fed spectrographs, using the photonic lantern injection platform on the SCExAO testbed at Subaru telescope. These efforts were part of SCExAO's initial astrophotonics development, particularly in working with the photonic lantern, and helped make clear some of the underappreciated challenges in working with real astrophotonic devices. Beyond challenges photonics fabrication, a running theme in this chapter is the difficulty in aligning a lantern, which is inside an extreme AO testbed, to the level of precision required for high coupling efficiency. Later in Chapter 5, we see that this curse is in some sense a blessing: sensitivity to misalignments allows the photonic lantern to act as a wavefront sensor.

4.1 Introduction

Diffraction-limited spectrographs fed by single-mode fibers (SMFs) offer advantages over seeing-limited spectrographs in both compactness and thermo-mechanical stability[BLR10, JSC16], and are quickly becoming the preferred architecture in spectrograph design. However, the coupling of large telescopes into SMF-fed spectrographs has historically been challenging due to the mismatch between the Gaussian-like fundamental mode of the fiber, and the telescope point-spread function (PSF), which is unstable due to time-varying wavefront error (WFE) originating from both atmospheric turbulence and the instrument itself.

The photonic lantern (PL) is a tapered waveguide which can efficiently couple light from a few-mode fiber (FMF) into multiple single-mode outputs [LBB05, BGY15]. This device represents an alternate way to inject telescope light into SMF-fed spectrographs. Notably, initial simulations [LJF21] have shown that lantern-based injection systems may offer significant benefits in coupling efficiency over traditional SMF-fed injection. However, while simulations are undoubtedly a crucial first step in characterizing the coupling performance of the PL, they should not be the only step. The challenge in performing a simulation that is truly representative of real physics is two-fold. For one, real lanterns, due to quirks in the fabrication process, may look significantly different (and thus behave differently) than the ideal lanterns assumed in simulations. This is especially the case for lanterns made by stacking and tapering bundles of SMFs: claddings for such lanterns at the FMF end may look more flower-shaped than circular (e.g. Figure 6 from [LFA17]). Additionally, realistic WFE, which includes the effect of atmospheric turbulence, non-common path aberrations (NCPAs; [MLA12, MKC13]), tip-tilt jitter, and the island effect [NMJ18, VBC19], is difficult to simulate. These challenges motivate the need for experimental validation and characterization of the PL.

In this work we present initial results from our ongoing experiments in the near-infrared (NIR) to characterize the 19-port PL on the SCExAO instrument [AGL21] at the Subaru Telescope. An overview of SCExAO's IR bench is shown in Figure 4.1. We now go over the relevant optical components for the experiments covered in this work. Light is first emitted from a supercontinuum calibration source, and passes through a 1550 ± 25 nm filter placed in the source box. This light is collimated by an off-axis parabolic (OAP) mirror, and then sent to a 2,000 actuator deformable mirror (DM). Light then passes through a set of phase-induced amplitude apodization (PIAA) lenses, which reformat the beam to be Gaussian in amplitude, thereby removing Airy rings in the focal plane [Guy03]. After the PIAA optics, the bulk of the light is picked off, and then sent to the PL via another OAP and a focusing lens. The PL is mounted on a translation stage (we will call this the "injection

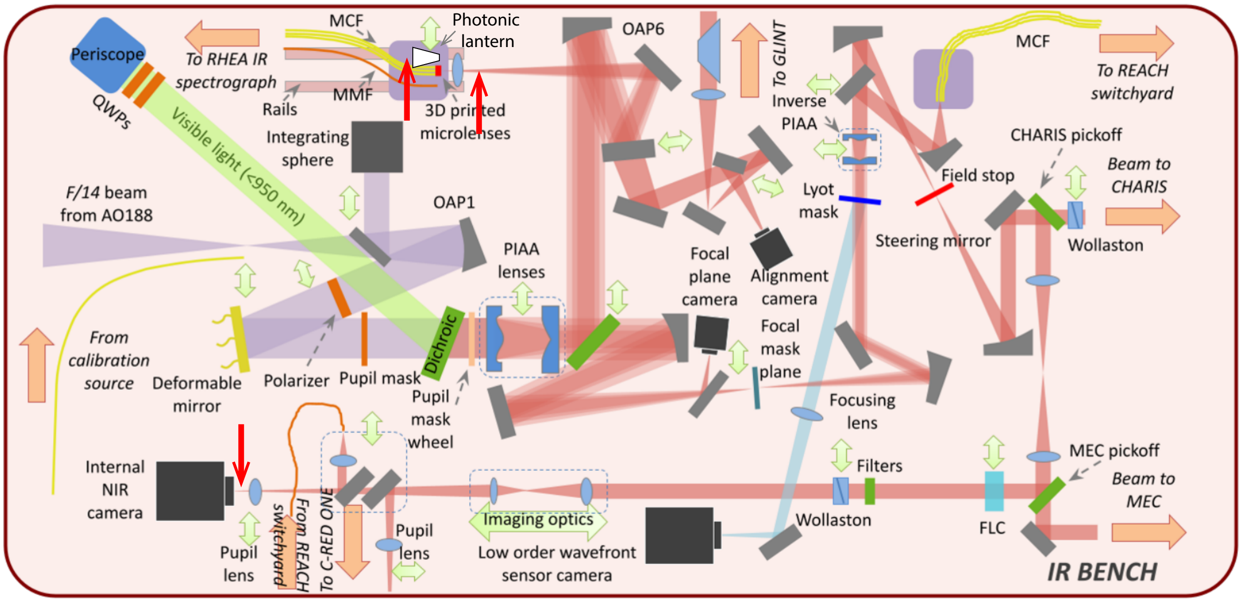

Figure 4.1: Diagram of the SCExAO IR test bench. Adapted from [AGL21]. The photonic lantern is located in the top left, mounted on a 5-axis translation stage. Also on the stage is an SMF, which was used for NCPA calibrations. The three red arrows denote locations where power meter measurements were taken.

stage") which can move in x, y, f, and z. The first two coordinates control the alignment of the stage transverse to the optical beam, while the latter two coordinates can be adjusted to change the focal ratio of the injection (for more details see [JSG17]). The output of the lantern, which looks like a multicore fiber (MCF), is imaged by the GLINT[MNT21] detector (a FirstLight Imaging C-red2, InGaAs detector). Meanwhile, the light not sent to the lantern is ultimately directed to SCExAO's internal NIR camera (bottom left of Figure 4.1). In Section §4.2, we provide an overview of the experimental procedures used to collect data from the 19-port PL on SCExAO. In §4.3, we present the results of these tests, and compare them to numerical models from [LJF21]. Finally, in §4.4 we discuss our results, and lay a course for future experiments.

4.2 Method

In this work, we present three different types of characterizations in relative throughput and one preliminary measurement of absolute throughput for the 19-port PL on SCExAO. All measurements were taken off-sky. We first go over some initial preparations we took before beginning our main data collection, and then detail the experimental setup for each set of measurements.

4.2.1 Preparations

Before taking measurements with the PL, we first aligned the stage in x, y, f, and z for optimal coupling into the Corning SMF-28 fiber mounted on the same translation stage as the PL. Our reason for this is two-fold: first, because SMFs only accept light that is flat in phase, we can modulate the DM and monitor SMF coupling in order to construct an NCPA compensation map. Second, coupling efficiencies of the same SMF were experimentally measured previously by [JSG17]; repeating these measurements will improve confidence in our subsequent measurements with the PL.

We first correct for non-common-path aberrations (NCPAs) between the GLINT lightpath and the SCExAO internal NIR camera lightpath. Using the DM, we then scan over the first 10 non-piston Zernike modes, monitoring the effect on coupling efficiency. For each mode, we record the mode amplitude which maximized SMF coupling; this list of mode amplitudes is then converted to a phase map which we apply to the DM for NCPA compensation.

With NCPA compensation applied, we then used a power meter to measure fluxes before and after the injection stage, in order to estimate the absolute coupling efficiency into the SMF. This efficiency was measured before by [JSG17] to be $\sim 87\%$; we recover a value of 83 ±3%, which is lower than the previous measurement but still close. Note that both estimates are solely of coupling efficiencies: losses due to Fresnel reflection and lens transmission are accounted for and divided out.

After our tests with the SMF, we realigned the system solely in x and y to maximize coupling into the PL, then moved on to PL characterization, detailed below.

4.2.2 Throughput vs. xy

For this test, we simply put the PL into focus and then shift the stage in x and y while recording the flux measured by the GLINT detector and the internal NIR camera detector. We then take the ratio of these two measurements: this proxy metric should be proportional to absolute lantern throughput, with the added benefit of being immune to variability in the supercontinuum source. After collecting a set of flux ratios, we divide by the largest value to scale our data relative to the maximum measured throughput.

We also maintain a focal ratio of 4.3: this was the optimal focal ratio for injection into the SMF mounted next to the PL, for our nominal wavelength of 1550 nm. According to simulations, $f/4.3$ should also be close to optimal for the 19-port PL, even though the resulting PSF will be undersized compared to the lantern's mode field diameter (MFD). This claim is also corroborated by our throughput vs. focal ratio measurements, covered in the next subsection. Note that converting the f and z coordinates of the injection stage position into a focal ratio is non-trivial: our approach is outlined in Appendix 4.A. Note that the NCPA map mentioned in 4.2.1 was not applied; however, because the PSF is undersized compared to the lantern MFD, it is unlikely that NCPA corrections will make a significant difference. We apply NCPA correction for later tests.

4.2.3 Throughput vs. focal ratio

For this test, apply our NCPA correction on the FM and scan the injection stage over z. For each z we align the PL in x, y, and f, then measured a flux ratio in the same manner as the previous subsection. We also scale these flux ratios relatively, as in the previous subsection.

4.2.4 Throughput vs. Strehl ratio

In this test we apply our NCPA compensation map to the DM and align the PL, setting the focal ratio to 4.3. We then artificially introduce WFE to the system. This WFE, which is produced by the turbulence simulator on the SCExAO real-time computer (RTC) and played back by the DM, is primarily Kolmogorov in nature, but with power in the low spatial frequencies suppressed in order to emulate partial correction of the wavefront by the AO188 system, Subaru Telescope's facility AO system. To characterize how relative throughput depends on Strehl ratio, we ramp up this simulated WFE, and for each Strehl, record 100 flux ratios. PSF images from the internal NIR camera are used to manually calculate Strehl ratio.

We later repeat this characterization, but increase the relative proportion of low-order wavefront error. This is done by changing the value of the "low-order coefficient" parameter in the turbulence simulator from 0.01 (the default value) to 1.

4.2.5 Absolute throughput measurement

In this subsection we detail our procedure for estimating the absolute throughput of the 19-port PL. After applying NPCA compensation, we aligned the injection stage in x and y for the PL (keeping the same 4.3 focal ratio) and then took four sets of three power meter measurements, the locations of which are indicated by the three red arrows in Figure 4.1. Our measurements and the corresponding estimate of the PL throughput are summarized in subsection §4.3.4.

4.3 Results

4.3.1 Throughput vs. xy

Figure 4.2 (left) shows the throughput map produced from the experimental procedure outlined in subsection §4.2.2. Interestingly, we find that the relative throughput varies by a factor of up to 3 over the face of the lantern, and note that the throughput pattern is stable, at least on the order of several days. This is already a deviation from simulation: for comparison, Figure 4.2 (right) shows the simulated coupling efficiency of the SCExAO PSF (accounting for the effect of PIAA optics) into an FMF matching the dimensions of the PL entrance. This coupling map is flat over the entrance of the fiber. Note that coupling efficiency is not the same as throughput: coupling efficiency accounts only for light loss that occurs at the FMF-like entrance of the lantern, while throughput additionally accounts for loss during propagation through the lantern transition. However, the former is often taken as a proxy for the latter because internal losses for PLs are usually assumed to be low. In fact, PLs with internal losses of only a few % have experimentally demonstrated[TEB13, LFS14].

4.3.2 Throughput vs. focal ratio

Figure 4.3 compares measured relative throughput of the 19-port PL (black points) against the simulated coupling efficiency into an FMF whose core diameter matches the size of the lantern entrance (dotted blue), over a range of focal ratios. We note decent agreement between experimental results and simulation, except at around $f/10$. Interestingly, at this focal ratio, the size of the PSF should roughly match the MFD of the PL; however, experiment and simulation both show that throughput/coupling efficiency at $\sim f/10$ is slightly lower than the maximal value at $\sim f/5$, suggesting that optimal coupling occurs when the PSF is undersized compared to the lantern MFD.

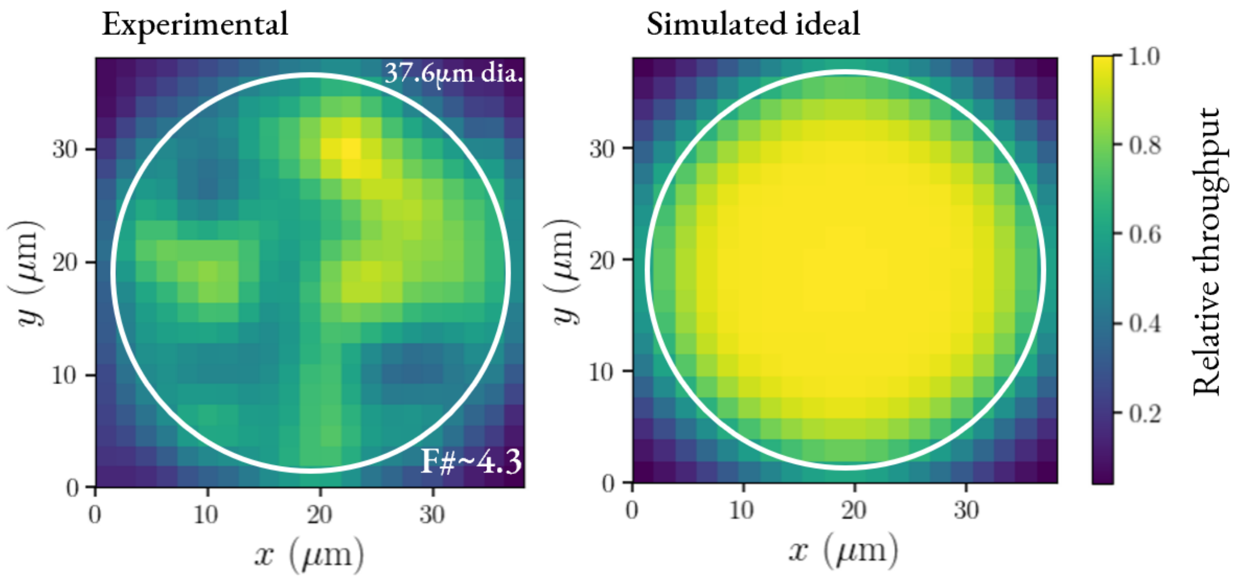

Figure 4.2: Left: total measured throughput of the 19-port PL, scaled against the maximum measured throughput, as a function of xy position. The cladding-jacket interface (37.6 μm diameter) is outlined in white. The focal ratio was set to 4.3. Right: simulated coupling efficiency, at the same focal ratio, between the SCExAO PSF and a few-mode step-index optical fiber that matches the entrance geometry of the PL.

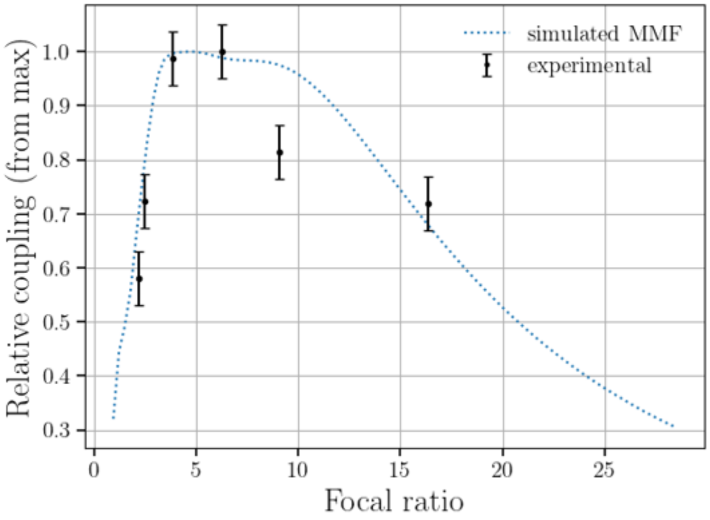

Figure 4.3: Black: net throughput of the 19-port PL, for various focal ratios. Blue: simulated coupling efficiency between the SCExAO PSF and an FMF matching the dimensions of the lantern entrance. Agreement is good except for around $f/10$.

4.3.3 Throughput vs. Strehl ratio

Figure 4.4 shows the results of our tests described in subsection §4.2.4. Blue points show experimentally measured relative throughputs of the PL, as a function of Strehl, while the black dashed line shows the simulated relative coupling efficiency of an aberrated beam into a similarly sized FMF. For simulations, WFE was produced by simulating a closed-loop AO system which partially corrects a single layer of Kolmogorov turbulence; Strehl is tuned by adjusting the Fried parameter, r_0. For more detail, refer to the method for WFE simulation described in [LJF21]. Our experiment recovers a linear relation between throughput and Strehl ratio, as expected from simulation; however, the slopes of the measured and simulated relations differ slightly. At our current stage of experimental testing with the PL, it is hard to determine the precise cause of this discrepancy, which could be due to differences in how the WFE was generated, or due to some intrinsic property of the lantern.

The red points in Figure 4.4 show how relative throughput of the lantern changes with Strehl, when the proportion of low spatial frequency to high spatial frequency WFE is increased. In the presence of increased low-order error, the throughput no longer behaves linearly, and instead approaches a dependence on Strehl that is more square-root in nature. Qualitatively, this result matches simulations from [LJF21](e.g. Figure 8), though the experimental setup in this work is not completely analogous to the numerical modelling done there. Regardless, both experiment and simulation indicate that higher order aberrations degrade throughput more than lower order aberrations. This is a property of PLs and FMFs in general, which can accept light through a variety of low-order fiber propagation modes, and thus are well-suited to accepting low-order aberrated light.

4.3.4 Absolute throughput measurement

On July 10, 2022, we took four sets of power meter measurements around the SCExAO IR bench, as described in subsection §4.2.5. Our results are summarized in Table 4.1. Conversion

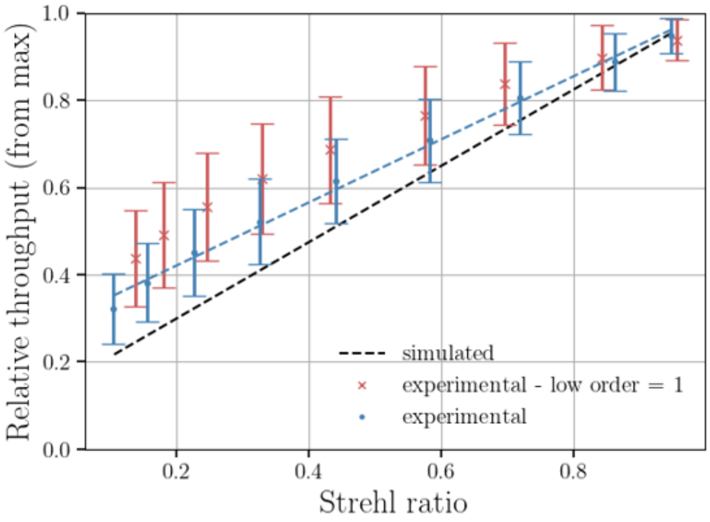

Figure 4.4: Blue: measured (relative) net throughput of the 19-port PL as a function of Strehl ratio. Strehl is altered by adjusting the amplitude of the WFE produced by the turbulence simulator on the SCExAO RTC. For each Strehl we measure relative throughput 100 times: points and bars show the mean and standard deviation of these measurements, respectively. Black: relative coupling efficiency of an aberrated beam into an FMF matching the entrance geometry of the lantern. Red: similar to the blue points, but the relative proportion of low to high spatial frequency is increased.

of power meter measurements to an absolute throughput value relies on only the first two columns; the third column of data was taken as a consistency check, and for comparison in the event of future power meter measurements. Note that between the two measurement locations corresponding to the first two columns of Table 4.1, there are also three lenses, one in front and two behind the lantern. The front lens was measured to have a transmission of 0.964; the back lenses are Thorlabs AC080-010-C and AC254-100-C, which together have a theoretical transmission of 0.981. Finally, because the lantern is uncoated, we also must account for Fresnel reflection at the front and back of the lantern. For an interface between air and fused silica, transmission should be 0.96 per surface. This gives the following formula for absolute throughput:

$$\text{absolute throughput} = \frac{\text{power after PL}}{\text{power before lens}} \times \frac{1}{0.964 \times 0.981 \times 0.96^2}. \tag{4.1}$$

Applying this equation to Table 4.1 gives an estimated net throughput for the PL of 59±4%. To compare with simulations, numerical models from [LJF21] predict a coupling efficiency into a similarly sized FMF of around 92%: this sets an upper bound for the coupling efficiency into the PL. Assuming that losses in the PL are predominately internal (occurring as light propagates through the body of the lantern, rather than during coupling into the lantern's entrance) we estimate the internal loss of the lantern to be $1 - 0.59/0.92 = 36\%$. This is significantly higher than our expectation of 5-10%. However, this result is preliminary: future work needs to be done to verify this figure, to separately quantify coupling and internal losses.

4.4 Discussion

While our preliminary characterizations of the 19-port PL on the SCExAO bench agree with previous numerical modelling in some areas, they deviate in others. Our next immediate goal will be to determine why these deviations occur. In this section we discuss some possible

Table 4.1: Power meter measurements (μW).

Before injection lens	After PL	In front of NIR camera
2.00	1.10	0.14
1.97	0.95	0.15
1.78	0.93	0.12
1.61	0.79	0.12

sources for the discrepancies seen between our experimental results and previous simulations.

The first source to consider is error in the experimental process. As an example, one contributor to the observed discrepancies is our NCPA compensation for the GLINT arm of the SCExAO IR bench. Anecdotally, we note that the impact of NCPA compensation is relatively minor, and often improves relative throughput by 5% or less; additionally our NCPA compensation cannot be too far off since we recover an SMF coupling efficiency similar to previous measurements[JSG17]. As such, imperfect NCPA compensation alone cannot explain the the factor of 3 variations in relative throughput shown in Figure 4.2, or the mismatch between measured throughput and simulated coupling efficiency at $f/10$ in Figure 4.3. However, in the future, it would still be prudent to refine our somewhat simplistic NCPA compensation scheme, as outlined in §4.2.3. Another source of error we are considering is potential angular misalignment between the PL and the optical path: because our injection setup is more sensitive to misalignment when injecting light into waveguides with larger MFDs, misalignment may explain why our coupling efficiency into the SMF is comparatively high, while our measured PL throughput is comparatively low.

Besides experimental error, the other main source for our observed discrepancies between experiment and simulation is that the photonic lantern on the bench is significantly non-ideal. Here, it is important to determine whether this non-ideal-ness is internal, or if it occurs when light enters/exits the lantern. One potential way to disambiguate between the two is to mount a FMF on the injection stage and repeat the characterizations covered in Section

§4.3. Significant differences between the two sets of characterizations might indicate internal imperfections within the lantern. Another useful test might be to reverse-inject light into the PL. This would involve flipping the PL on the injection stage and coupling light into the individual SMF cores at the lantern's MCF end, then imaging the FMF end using the GLINT detector. Measuring the throughput in this manner can help us better characterize PL losses, and can tell us if any of the lantern's cores are underperforming.

The last potential source for our observed discrepancies between experiment and simulation that we will discuss is inaccuracy in previous simulations (e.g. [LJF21]). However, we believe this to to be unlikely, since such simulations are based on the physics of coupling into FMFs, which is well-understood. Here, repeating our experimental characterizations with an FMF will again be useful: presuming no issues in the experimental setup, comparisons with the simulated results will confirm or deny the accuracy of the previous simulations.

Finally, we briefly discuss future experiments, after the discrepancies between experiment and simulation have been resolved. The most important test will be to take PLs on sky, so that we can gauge their throughputs in truly realistic conditions. We also plan repeating these characterizations for 3-, 6-, 10-, and 12-port PLs, which will be fabricated at the Sydney Astrophotonic Instrumentation Laboratory (SAIL).

4.5 Conclusion

We presented initial characterizations of the 19-port PL on the SCExAO testbench at Subaru Telescope, and compared these characterizations with simulated results for an FMF that matches the entrance geometry of the lantern. We noted several discrepancies: for one, we find that the experimentally measured throughput of the lantern varies by a factor of up to 3 over the entrance of the lantern, while simulated results suggest coupling efficiency (and hence throughput, for an ideal lantern) should remain relatively flat. Next, we measured PL throughput as a function of focal ratio, and found good agreement with simulations except around a focal ratio of $f/10$. Using the SCExAO turbulence simulator, we then injected WFE into the optical system in order to characterize lantern throughput as a function of Strehl ratio. Our results agree with simulations. Finally, we presented initial results on our measurement of absolute PL throughput. We found a net throughput of 59±4%; assuming a coupling efficiency of 92%, the upper limit set by simulations, we compute an estimated internal loss for the lantern of 36%. In the future, we plan to extend our characterizations to 3-,6-,10-,and 12-port PLs, and to measure PL throughput on-sky.

Appendix - Chapter 4

4.A Conversion of injection stage position to focal ratio

For two converging lenses separated by some distance d, the formula for equivalent focal length (and hence focal ratio F) will take the form

$$F(d) = \frac{1}{a + bd} \tag{4.2}$$

for real-valued constants a and b. In the context of SCExAO, we instead have an OAP (see OAP6, Figure 4.1) and a convex lens, but the physics is the same. Separation between the OAP and the lens is linearly related to the optical carriage position z, so that focal ratio F also has the dependence $F(z)$ as described by equation 4.2. To determine the best values for a and b, we measured the coupling efficiency into the SMF on the injection stage, denoted $C_e(z)$ as a function of z. This data was provided by Sébastien Vievard. We also numerically computed the theoretical SMF coupling efficiency as a function of focal ratio F, denoted $C_s(F)$ and plotted in blue in Figure 4.A.1. Finally, we used a Nelder-Meade optimizer to find the values of a and b which minimized the squared residual between $C_e(z)$ and $C_s(F(z))$ over an array of z values. This gives the best-fit relation $F(z)$, plotted in the right panel of Figure 4.A.1. The experimentally measured values of SMF coupling, plotted against focal

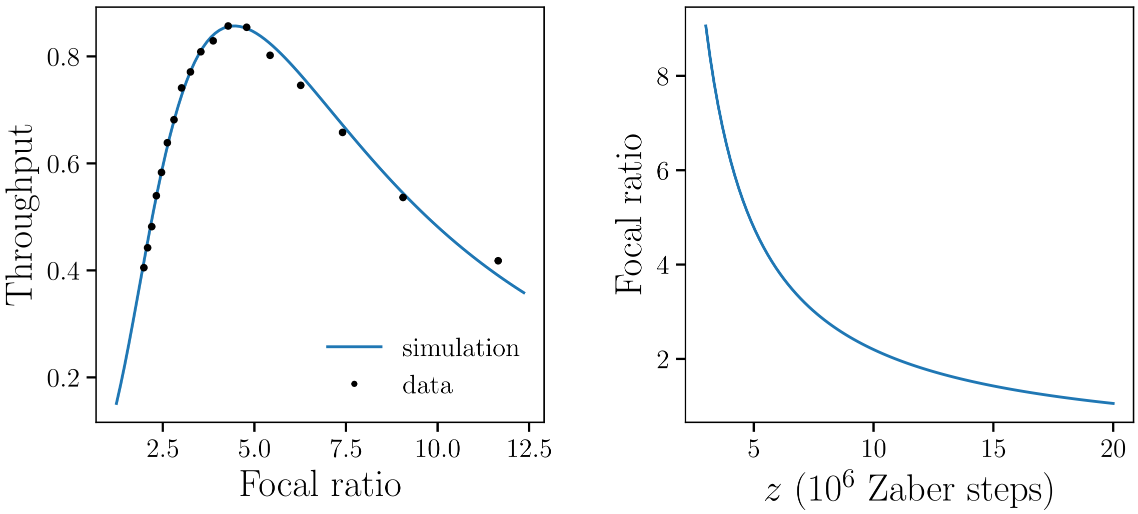

Figure 4.A.1: Left: comparison of simulated SMF coupling and experimentally measured coupling, after applying our best-fit transformation $F(z)$ to the carriage position coordinate z. Right: best-fit relation for $F(z)$.

ratio instead of z using the best-fit transformation $F(z)$, is shown by the black points in Figure 4.A.1, left. These points agree well with simulation.

REFERENCES

[AGL21] Kyohoon Ahn, Olivier Guyon, Julien Lozi, Sébastien Vievard, Vincent Deo, Nour Skaf, Ruslan Belikov, Steven P. Bos, Michael Bottom, Thayne Currie, Richard Frazin, Kyle V. Gorkom, Tyler D. Groff, Sebastiaan Y. Haffert, Nemanja Jovanovic, Hajime Kawahara, Takayuki Kotani, Jared R. Males, Frantz Martinache, Ben Mazin, Kelsey Miller, Barnaby Norris, Alexander Rodack, and Alison Wong. "SCExAO: a testbed for developing high-contrast imaging technologies for ELTs." In Stuart B. Shaklan and Garreth J. Ruane, editors, *Techniques and Instrumentation for Detection of Exoplanets X*, volume 11823, pp. 9 – 21. International Society for Optics and Photonics, SPIE, 2021.

[BGY15] T. A. Birks, I. Gris-Sánchez, S. Yerolatsitis, S. G. Leon-Saval, and R. R. Thomson. "The photonic lantern." *Adv. Opt. Photon.*, **7**(2):107–167, Jun 2015.

[BLR10] Joss Bland-Hawthorn, Jon Lawrence, Gordon Robertson, Sam Campbell, Ben Pope, Chris Betters, Sergio Leon-Saval, Tim Birks, Roger Haynes, Nick Cvetojevic, and Nem Jovanovic. "PIMMS: photonic integrated multimode microspectrograph." In Ian S. McLean, Suzanne K. Ramsay, and Hideki Takami, editors, *Ground-based and Airborne Instrumentation for Astronomy III*, volume 7735 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 77350N, July 2010.

[Guy03] Guyon, O. "Phase-induced amplitude apodization of telescope pupils for extrasolar terrestrial planet imaging." *A & A*, **404**(1):379–387, 2003.

[JSC16] N. Jovanovic, C. Schwab, N. Cvetojevic, O. Guyon, and F. Martinache. "Enhancing Stellar Spectroscopy with Extreme Adaptive Optics and Photonics." *Publications of the Astronomical Society of the Pacific*, **128**(970):121001, nov 2016.

[JSG17] Jovanovic, N., Schwab, C., Guyon, O., Lozi, J., Cvetojevic, N., Martinache, F., Leon-Saval, S., Norris, B., Gross, S., Doughty, D., Currie, T., and Takato, N. "Efficient injection from large telescopes into single-mode fibres: Enabling the era of ultra-precision astronomy." *A & A*, **604**:A122, 2017.

[LBB05] S. G. Leon-Saval, T. A. Birks, J. Bland-Hawthorn, and M. Englund. "Multimode fiber devices with single-mode performance." *Opt. Lett.*, **30**(19):2545–2547, Oct 2005.

[LFA17] Sergio G. Leon-Saval, Nicolas K. Fontaine, and Rodrigo Amezcua-Correa. "Photonic lantern as mode multiplexer for multimode optical communications." *Optical Fiber Technology*, **35**:46–55, 2017. Next Generation Multiplexing Schemes in Fiber-based Systems.

[LFS14] Sergio G. Leon-Saval, Nicolas K. Fontaine, Joel R. Salazar-Gil, Burcu Ercan, Roland Ryf, and Joss Bland-Hawthorn. "Mode-selective photonic lanterns for space-division multiplexing." *Opt. Express*, **22**(1):1036–1044, Jan 2014.

[LJF21] Jonathan Lin, Nemanja Jovanovic, and Michael P. Fitzgerald. "Design considerations of photonic lanterns for diffraction-limited spectrometry." *J. Opt. Soc. Am. B*, **38**(7):A51–A63, Jul 2021.

[MKC13] Martinez, P., Kasper, M., Costille, A., Sauvage, J. F., Dohlen, K., Puget, P., and Beuzit, J. L. "Speckle temporal stability in XAO coronagraphic images - II. Refine model for quasi-static speckle temporal evolution for VLT/SPHERE." *A&A*, **554**:A41, 2013.

[MLA12] Martinez, P., Loose, C., Aller Carpentier, E., and Kasper, M. "Speckle temporal stability in XAO coronagraphic images." *A&A*, **541**:A136, 2012.

[MNT21] Marc-Antoine Martinod, Barnaby Norris, Peter Tuthill, Tiphaine Lagadec, Nemanja Jovanovic, Nick Cvetojevic, Simon Gross, Alexander Arriola, Thomas Gretzinger, Michael J. Withford, Olivier Guyon, Julien Lozi, Sébastien Vievard, Vincent Deo, Jon S. Lawrence, and Sergio Leon-Saval. "Scalable photonic-based nulling interferometry with the dispersed multi-baseline GLINT instrument." *Nature Communications*, **12**(1):2465, Apr 2021.

[NMJ18] N´Diaye, M., Martinache, F., Jovanovic, N., Lozi, J., Guyon, O., Norris, B., Ceau, A., and Mary, D. "Calibration of the island effect: Experimental validation of closed-loop focal plane wavefront control on Subaru/SCExAO." *A&A*, **610**:A18, 2018.

[TEB13] Christopher Q. Trinh, Simon C. Ellis, Joss Bland-Hawthorn, Jon S. Lawrence, Anthony J. Horton, Sergio G. Leon-Saval, Keith Shortridge, Julia Bryant, Scott Case, Matthew Colless, Warrick Couch, Kenneth Freeman, Hans-Gerd Löhmannsröben, Luke Gers, Karl Glazebrook, Roger Haynes, Steve Lee, John O'Byrne, Stan Miziarski, Martin M. Roth, Brian Schmidt, Christopher G. Tinney, and Jessica Zheng. "GNOSIS: The first instrument to use fiber Bragg gratings for OH suppression." *The Astronomical Journal*, **145**(2):51, jan 2013.

[VBC19] Sebastien Vievard, Steven Bos, Frederic Cassaing, Alban Ceau, Olivier Guyon, Nemanja Jovanovic, Christoph U. Keller, Julien Lozi, Frantz Martinache, Aurelie Montmerle-Bonnefois, Laurent Mugnier, Mamadou NDiaye, Barnaby Norris, Ananya Sahoo, Jean-Francois Sauvage, Frans Snik, Michael J. Wilby, and Alisson Wong. "Overview of focal plane wavefront sensors to correct for the Low Wind Effect on SUBARU/SCExAO.", 2019.

Part III

Wavefront sensors: analysis and modelling

CHAPTER 5

Wavefront sensing with photonic lanterns

One of the key draws of photonic lanterns and other astrophotonic devices is their potential to simultaneously perform scientific observations (e.g. high-R spectroscopy, as in the prior chapter) along with calibration measurements which can be used to compensate for measurement-corrupting noise. This self-calibrating capability will be necessary for future scientific instruments to meet the extreme technical requirements posed by direct exo-Earth imaging. The next sections, adapted from two journal papers, show how photonic lanterns may be leveraged as sensors which measure and drive the correction of one of the dominant types of astronomical noise: wavefront error. In tandem with spectroscopy, this technique improves the precision of retrieved spectra; with imaging, it sharpens angular resolution; with coronagraphy, it deepens contrast. Note that while the techniques derived below were motivated by astrophotonic wavefront sensing, the results can be generalized to any few-moded wavefront sensor, at least in astronomical contexts. Some ideas from this work are continued in this more general sense, in Chapter 6.

5.1 Introduction

High-contrast imaging is becoming one of the primary tools for the direct detection and characterization of exoplanets. This class of techniques combines ground-based extreme adaptive optics (AO), which corrects for wavefront aberrations induced by passage of light through the atmosphere and the instrument, and coronagraphy, which suppresses on-axis starlight to reveal the circumstellar environment, as well as contrast-boosting post-processing techniques such as angular differential imaging [MLD06] and spectral differential imaging [SF02]. Together, these techniques enable contrasts down to $\sim 10^{-6}$ and angular separations down to 200 mas. So far, some 30 exoplanets have been detected through high-contrast imaging techniques [Bow16]; however, almost all are widely separated gas giants with masses several times that of Jupiter. One of the main roadblocks in increasing current sensitivity are non-common-path aberrations (NCPAs): quasi-static aberrations evolving on the timescale of minutes to hours that occur due to instrument instabilities induced by humidity, temperature, and gravity vector changes [MLA12, MKC13]. Because these aberrations appear downstream from the wavefront sensor, they cannot be removed via typical pupil-plane wavefront control systems. As a result, wavefront control must be improved before instruments can attain the necessary contrasts and angular separations typical for systems similar to the Sun and Earth: $\sim 10^{-10}$ and $\sim$ 100 mas, at a distance of 10 pc, in visible light [TO10]. One way forward is to sense wavefront aberrations in the final focal plane with the science camera, so that sensor and science light travel down the same optical path. This approach, known as focal-plane wavefront sensing (FPWFS), removes NCPAs.

In parallel, a number of new ideas and techniques are being proposed to further advance direct exoplanet characterization. One development is in short-exposure exoplanet imaging, which leverages statistical differences in planet and star speckle behavior at millisecond timescales to distinguish between planet light from starlight [RFM21, GBD19]. This technique is distinct from ADI and SDI. Coherent detection, which exploits the incoherence of planet light, presents an alternative pathway for separating planet light and starlight. A related technique is nulling interferometry, an alternative to conventional coronagraphy that can achieve smaller inner working angles, and which works by destructively interfering starlight collected from different subapertures or telescopes. Other advances in direct characterization will need to be made not in the isolation of planet light, but the spectral analysis of that light. The high-resolution spectral analysis of faint objects like exoplanets will require methods for both the efficient coupling of light into the science instrument, and stabilization of that same light, which will vary with time due to passage through the atmosphere and instrument. These two requirements are typically in tension [LJF21], and thus hard to achieve simultaneously.

The photonic lantern (PL; [BGY15]) provides a capable platform for the above applications; other notable applications include OH line suppression through fiber Bragg gratings [TEB13, EBB18], and spectroastrometry [GVD19]. As in Figure 3.1, the PL is a tapered waveguide that gradually transitions from a few-mode optical fiber (FMF) geometry to multiple widely-spaced single-mode cores, similar to a multi-core fiber (MCF), which can then be fanned out to an array of single-mode fibers (SMFs). When the FMF end is placed in the focal plane, the PL can efficiently couple multi-modal telescope light into multiple SMFs. While PLs come in a wide array of port counts and geometries, they can be largely classified into three groups. In what we call the "standard" PL, embedded cores are uniform in structure and refractive index. At the other extreme, "mode-selective" PLs use differing single-mode core radii or index contrasts, so that each fiber mode at the FMF-like lantern entrance routes to a distinct output port [LFS14]. Lastly, we term lanterns that operate between these two extremes "hybrid lanterns." These lanterns have one core mismatched from the rest, thereby funnelling light from the fundamental fiber mode into a single output port while mixing the remaining light in the rest of the ports. This concept is similar to the "mode-group selective" lantern, introduced in [VAA18].

Critically, in the process of coupling light into an array of SMFs, PLs map phase aberrations into intensity variations in a one-to-one manner, at least for small aberrations. This behavior enables the PL to additionally act as a 100% duty cycle focal-plane WFS [CMH18, NWB20, WYH22]. Because PLs have a limited number of outputs (set by the manufacturing process, though PLs with up to 511 modes have been reported [BGY15]), these devices as of now can only give low-order wavefront information. Therefore, while PLs are well-suited to sense low-order aberrations like NCPAs [SFR07] and island modes [NMJ18], they are not a standalone WFS solution in XAO systems, which correct upwards of 1000 modes. In such applications, PLs will likely need to work in tandem with pupil-plane sensors like the Shack-Hartmann or pyramid WFS.

We show an example of this phase-to-intensity mapping in Figure 5.1.1, which plots the non-degenerate intensity responses of a 6-port PL in the presence of positive and negative astigmatism. The focus of this work is to assess the performance of the photonic lantern wavefront sensor (PL WFS), in contexts like instrument coupling or coherent detection where PLs are already being considered for use. In these scenarios, the utility of the PL is doubled, enabling both the aforementioned non-WFS applications as well as focal-plane wavefront sensing. We focus on two contexts the first being fiber-fed, high-resolution spectrometry, mentioned above; and vortex-fiber nulling (VFN), a high contrast imaging technique which exploits symmetries in optical fiber modes to separate star and planet light [REJ19]. In turn, we restrict our analysis to the infrared, since this wavelength regime will be the staging ground for the next push in direct exoplanet spectrometry, with upcoming instruments such as HISPEC and MODHIS [MFK19].

Research in PL wavefront sensing is ongoing. For instance, [NWB20] recently combined a 19-port PL with a neural net to enable nonlinear wavefront reconstruction of the first 9 non-piston Zernike modes. In comparison, we take a broader, but less in-depth approach: our goal is to provide a general baseline overview of the capabilities of the PL WFS, as well as the methods through which the sensing properties of these devices might be controlled. We place added emphasis on the linear analysis of the PL WFS, in order assess the limits of the PL WFS under more standard and simplistic linear AO control schemes. In Section §5.2, we establish the math that will enable wavefront reconstruction with the PL WFS. To begin, we present power series expansions for the PL WFS intensity response to first and second order in phase (§5.2.1-§5.2.3). We also consider methods through which these models can be inverted, thereby enabling wavefront sensing. Next, we expand our models to arbitrary modal basis (§5.2.4): this both increases computational efficiency of the reconstruction models and allows them to be expressed in terms of common phase aberration bases such as the Zernike

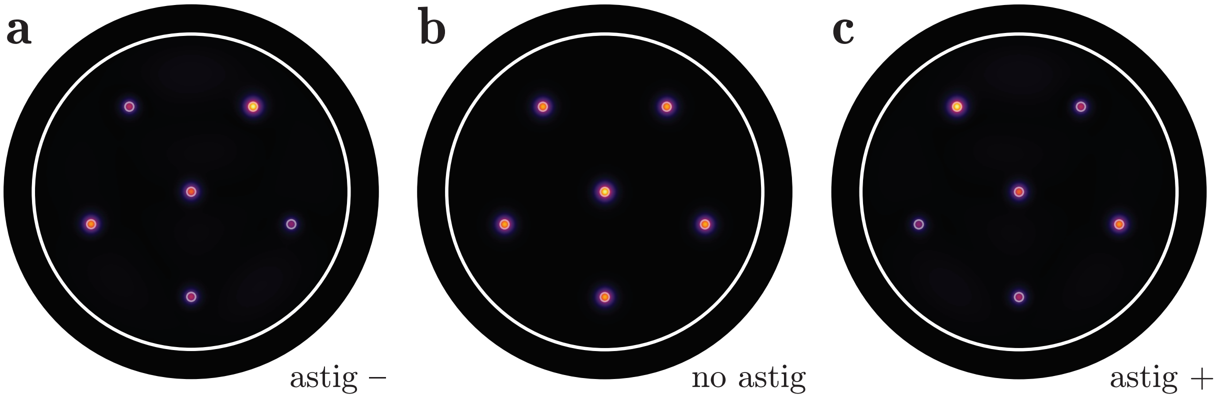

Figure 5.1.1: Simulated response of a 6 port lantern in the presence a: -1 rad rms astigmatism; b: 0 rad rms astigmatism; and c: +1 rad rms astigmatism. The photonic lantern converts phase variations into unique intensity variations among the output cores. Circles show the jacket-cladding interface and the cladding-core interfaces. Optical propagation is simulated using the Python packages HCIPy and Lightbeam.

polynomials. In Section §5.3, we apply our models to quantify the behavior of the PL WFS. This analysis includes deriving conditions for WFS linearity (§5.3.1-§5.3.3), and estimating maximum amount of WFE that can be handled by these sensors (§5.3.4).

Finally, we combine our models with numerical simulations, to provide a first look at the wavefront-sensing abilities of a standard, hybrid, and mode-selective 6-port PL. Our aim in this work is to develop an initial understanding of the capabilities of the PL WFS, and in doing so we assume "perfect" lanterns and neglect noise (though we provide some reference to noise propagation in the linear regime in §5.2.2). We present an overview of our numerical method in §5.4, and the corresponding results in §5.5. In a companion paper [LF22], we extend these simulations to cover a range of PL WFS configurations beyond the 6-port geometries considered in this paper, in order to establish a rough baseline of the sensing abilities of PL WFSs. There, we also investigate potential strategies through which PL WFS performance can be further controlled and optimized.

5.2 Propagation analysis and phase reconstruction

5.2.1 General model

Consider the following general setup for a backend device to an AO-equipped telescope. AO-corrected light passes into an instrument backend, which may contain components such as beam-shaping (PIAA) optics [Guy03] and additional phase and/or amplitude optics (e.g. vortex fiber nuller mask). After light passes through some number of upstream components, it is focused onto the FMF end of a PL, ultimately propagating into the SMF ports at the PL output. These output ports may also optionally be inteferometrically combined. Because optical propagation is linear in complex electric field, the action of all backend optical components can be lumped into a single complex-valued transfer matrix, which we denote A. This matrix connects the input electric field $\mathbf{u}_{\text{in}}$ and the output electric field $\mathbf{u}_{\text{out}}$ of the backend device:

$$\mathbf{u}_{\text{out}} = A\mathbf{u}_{\text{in}}. \tag{5.1}$$

In the case of the PL WFS, the transfer matrix A will contain a projection component, since an N-port lantern will support only N complex-valued electric field modes, meaning that the vector $\mathbf{u}_{\text{out}}$ is N-dimensional. Note that, unlike the modes of a standard optical fiber, the modes of a PL are three-dimensional, encompassing the full propagation of light from the FMF-like input to the MCF-like output of the lantern. Here, we have a choice of mode basis. The modes we use in this work, which we term “lantern modes,” look like individual SMF modes at the lantern exit, and complex linear combinations of fiber modes at the lantern entrance. These modes can be computed by illuminating a single output core at the lantern exit and numerically back-propagating light to the lantern entrance. Simulated cross-sections of lantern modes at the PL entrance, computed in this manner, are shown in Figure 5.2.1 for a standard 6-port lantern. The A matrix accounts for optical propagation through the telescope and any subsequent beam-shaping to the PL entrance, and then projects the focal plane electric field onto these lantern modes. Accordingly, A has dimensions $N \times M$, for an

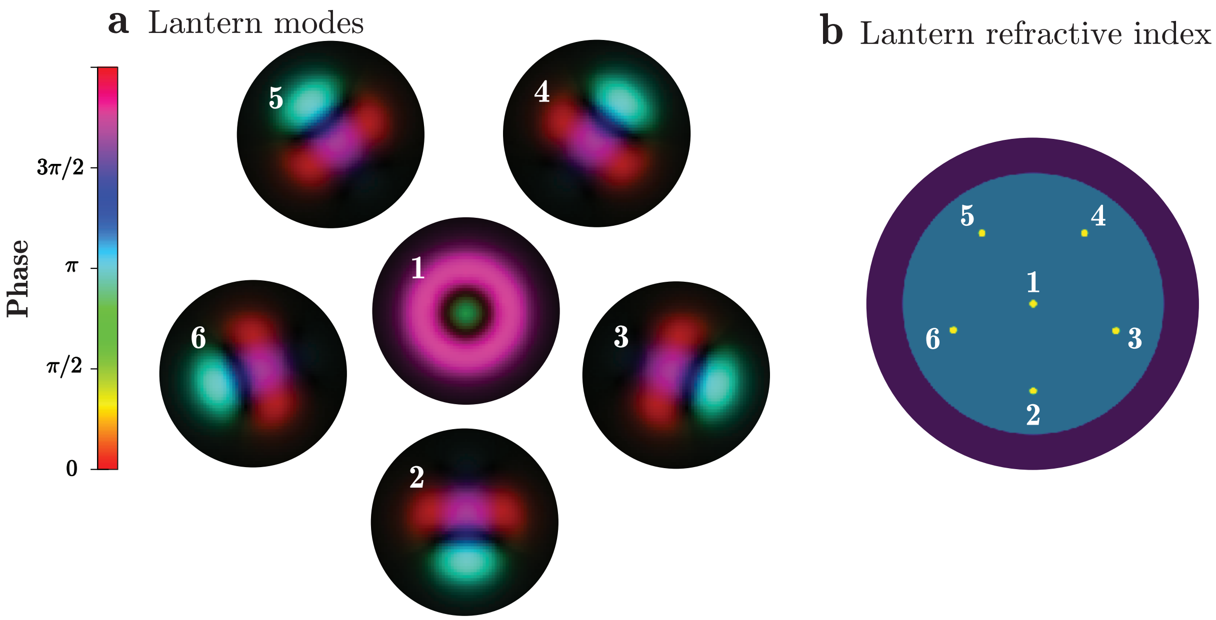

Figure 5.2.1: Panel a: Lantern modes for the same 6-port lantern as in Figure 5.1.1, evaluated at the lantern entrance. Phase is plotted in color, while amplitude is plotted in opacity. The 6 lantern modes are oriented to reflect the location of their corresponding lantern ports, shown on the right. To identify the ports and lantern modes, we index them according to the numerical labels. Panel b: the refractive index profile of the output (MCF-like) end of the PL. Embedded SMF cores are shown in yellow. Numerical labels connect each core to its corresponding lantern mode in panel a.

N-port lantern and M pupil samples.

Since we ultimately measure intensity, not complex amplitude, we recast equation 5.1 in terms of the intensity response $\mathbf{p}_{\rm out}$:

$$\mathbf{p}_{\rm out} = |A\mathbf{u}_{\rm in}|^2. \tag{5.2}$$

For phase-only aberrations, the goal of wavefront sensing is to invert equation 5.2 and recover the phase of $\mathbf{u}_{\rm in}$. We go over methods to do so in the following subsections.

5.2.2 Linearizing intensity response

In this subsection, we provide a review of wavefront sensing in the linear regime. While optical propagation is linear in complex amplitude, it is nonlinear in intensity. However, for small changes in aberration amplitude, the intensity response will vary in a near-linear manner. Consider a phase-only aberration ϕ in an electric field with assumed uniform intensity $I_{\text{in}} = 1$. We can approximate the intensity response of the system about some arbitrary reference phase ϕ_0 as

$$\mathbf{u}_{\text{in}} = \exp(i\phi) \approx e^{i\phi_0} \odot [\mathbf{1} + i(\phi - \phi_0)] \tag{5.3}$$

where the vector $\mathbf{1}$ represents the electric field of a flat wavefront, and $(\odot)$ represents element-wise (Hadamard) vector-vector multiplication. For clarity, we denote $\boldsymbol{\Delta}\phi \equiv \phi - \phi_0$, and modify the transfer matrix as $A_{ij} \to A_{ij}e^{i\phi_{0,j}}$; for a flat reference wavefront, $\phi_{0,j} = 0$ and A_{ij} is unchanged. The intensity resulting from the phase aberration ϕ is

$$\begin{aligned}
\mathbf{p}_{\text{out}} &= |A\mathbf{u}_{\text{in}}|^2 \\
&\approx \left|A\left[\mathbf{1} + i\boldsymbol{\Delta}\phi\right]\right|^2 \\
&\approx |A\mathbf{1}|^2 + 2\,\text{Im}\left[(A\mathbf{1}) \odot (A^*\boldsymbol{\Delta}\phi)\right]
\end{aligned} \tag{5.4}$$

where the squaring and $(||)$ operators are element-wise, and Im denotes taking the imaginary part. We can define the matrix B, having the same dimensions as A, as

$$B_{ij} \equiv 2\,\text{Im}\left[A^*_{ij}\sum_k A_{ik}\right] \tag{5.5}$$

and recover

$$\mathbf{p}_{\text{out}} \approx |A\mathbf{1}|^2 + B\boldsymbol{\Delta}\phi. \tag{5.6}$$

We see that the first quantity represents the bias intensity when there is no phase error, while the matrix B (often called the "interaction matrix" in the context of adaptive optics) describes the linear response of the intensities to phase perturbations from the reference wave – in other words, B is the Jacobian of the PL's intensity response, evaluated at the reference wavefront determined by ϕ_0. Equation 5.6 can be inverted (e.g. via Moore-Penrose pseudo-inverse), enabling the reconstruction of phase errors from intensity responses. The phase aberration modes which this backend device can sense in the linear regime will be determined by the B matrix; un-sensed aberration modes will lie in the null space of B. Alternatively, A and B can be used to compute gradients and Hessians for cost functions, enabling iterative nonlinear estimation for phase aberrations (see [Fra18] for an example of this sort of analysis, with the pyramid WFS).

Finally, to understand how error and noise propagates through the linear reconstruction process, we follow the analysis of [Fra18]. We write the intensity response of an N-port PL WFS as

$$\mathbf{p}_{\text{out}} = |A\mathbf{1}|^2 + B\boldsymbol{\Delta}\phi + \mathbf{n}(\boldsymbol{\Delta}\phi) + \nu \tag{5.7}$$

where $\mathbf{n}$ is an N-length vector of functions that accounts for the error caused by linearization, and ν is the noise, assumed to be composed of N independent zero-mean random processes. The least-squares estimate $\widetilde{\boldsymbol{\Delta}\phi}$ for the original phase aberration $\boldsymbol{\Delta}\phi$ is obtained through B^+, the pseudo-inverse of B, as follows:

$$\begin{aligned}\widetilde{\boldsymbol{\Delta}\phi} &= B^+\left[\mathbf{p}_{\text{out}} - |A\mathbf{1}|^2\right] \\ &= B^+B\boldsymbol{\Delta}\phi + B^+\left[\mathbf{n}(\boldsymbol{\Delta}\phi) + \nu\right].\end{aligned} \tag{5.8}$$

Small singular values of B will amplify both the error incurred by linearization, as well as random noise. Such amplification can be partially mitigated through regularization of the singular values. Future analysis of the reconstruction properties of the PL WFS, particularly for nonlinear reconstruction and closed-loop operation, will require more detailed considerations of noise and error propagation — we leave this for later work.

5.2.3 Second-order analysis of intensity response

Under perfect knowledge of the system transfer matrix A, we may obtain greater accuracy by expanding WFS response to second order. Express the incident electric field as

$$\mathbf{u}_{\text{in}} \approx e^{i\phi_0}\left(1 + i\boldsymbol{\Delta}\phi - \frac{1}{2}\boldsymbol{\Delta}\phi^2\right). \tag{5.9}$$

Repeating the analysis of the previous subsection, again making the substitution $A_{ij} \rightarrow A_{ij}e^{i\phi_{0,j}}$, leads to the following:

$$\mathbf{p}_{\text{out}} = |A\mathbf{1}|^2 + 2\,\mathrm{Im}\left[(A\mathbf{1}) \odot (A^*\boldsymbol{\Delta}\phi)\right] - \mathrm{Re}\,\left[(A\mathbf{1}) \odot (A^*\boldsymbol{\Delta}\phi^2)\right] + |A\boldsymbol{\Delta}\phi|^2. \tag{5.10}$$

We define the matrix C as

$$C_{ij} \equiv 2\,\mathrm{Re}\left(A^*_{ij}\sum_k A_{ik}\right) \tag{5.11}$$

where Re denotes taking the real part. This yields the following formula for how phase errors up to second order affect intensity:

$$\mathbf{p}_{\text{out}} \approx |A\mathbf{1}|^2 + B\boldsymbol{\Delta}\phi - \frac{1}{2}C\boldsymbol{\Delta}\phi^2 + |A\boldsymbol{\Delta}\phi|^2. \tag{5.12}$$

Inversion of equation 5.12 can be accomplished using iterative techniques like Landweber iteration, the Levenberg-Marquardt algorithm, or gradient descent. Such methods often benefit from knowledge of the Jacobian, which can be derived from equation 5.12:

$$J_{ij} = \frac{\partial I_{\text{out},i}}{\partial \Delta\phi_j} = B_{ij} + (|A_{ij}|^2 - C_{ij})\Delta\phi_j + \sum_k A_{ij}A^*_{ik}\Delta\phi_k. \tag{5.13}$$

Our reliance on numerical solving techniques begs the question: prior to inversion, why approximate the intensity response of the PL WFS at all? An alternative strategy is to numerically solve equation 5.2 directly. However, we note two benefits of making the initial approximation. First, doing so simplifies the inverse problem, which improves numerical stability and mitigates issues where the numerical solver becomes stuck in local minima (similar to the phenomenon observed by [Fra18], for the pyramid WFS). This issue is exacerbated as the nonlinearity of a PL increases. Second, as we will later see in Section §5.2.4, truncation of the power series enables the use of a modal basis for phase aberrations, which increases computational efficiency, especially for a low-order sensors like the PL WFS. As an added note, the preliminary analysis presented in this work may one day enable or accelerate non-iterative nonlinear reconstruction akin to neural net methods [NWB20], as well as more detailed analytic or semi-analytic characterization of the PL WFS.

However, it is important to emphasize that inversion of quadratic and higher-order models is more complicated than their linear counterpart, primarily because nonlinear models can admit multiple solutions, at least for large WFE. This multiplicity may be a fundamental property of the WFS system, or an artifact due to truncation of the power series. We briefly discuss how some of these issues may be mitigated in §5.5.3.

A cubic expansion is presented in Appendix 5.B.

5.2.4 Modal basis

The matrices A, B, and C will each have N by M entries, where N is the number of output ports and M is the number of sample points in the pupil plane. This is computationally inefficient — the number of sample points will almost always greatly exceed the number of lantern ports, making the above matrices unnecessarily large. It is more efficient to represent phase aberrations in terms of some modal basis (e.g. the Zernike modes, the Karhunen-Loève modes derived from second-order phase aberration statistics, or the singular vectors of the A matrix, projected onto pupil phase). To do so, write the phase aberration displacement vector $\mathbf{\Delta}\phi$ as

$$\mathbf{\Delta}\phi = R\mathbf{a}, \tag{5.14}$$

where $\mathbf{a}$ is the real-valued vector of modal coefficients and R is the change-of-basis-matrix, whose columns correspond to the basis vectors. Defining $B' \equiv BR$, the linear model given by equation 5.6 is easily extended to modal basis as follows:

$$\mathbf{p}_{\text{out}} \approx |A\mathbf{1}|^2 + B'\mathbf{a}. \tag{5.15}$$

Extension of the quadratic model to modal basis is more involved. Inserting equation 5.14 into equation 5.12 results in the following:

$$\mathbf{p}_{\text{out},i} \approx |A\mathbf{1}|_i^2 + (B'\mathbf{a})_i - \frac{1}{2}\sum_{jk} C'_{ijk} a_j a_k + |A'\mathbf{a}|_i^2 \tag{5.16}$$

where the tensor C' is defined as

$$C'_{imn} \equiv \sum_j C_{ij} R_{jm} R_{jn} \tag{5.17}$$

and the $A' \equiv AR$. Differentiating equation 5.16 yields the Jacobian, under the quadratic approximation, in terms of modal basis:

$$J'_{ij} = B'_{ij} + \sum_k \left(\operatorname{Re}\left[A'_{ij} A'^{*}_{ik}\right] - \frac{1}{2} C'_{ijk} \right) a_k + \left(|A'_{ij}|^2 - \frac{1}{2} C'_{ijj} \right) a_j. \tag{5.18}$$

5.3 PL WFS properties

In this section, we provide an initial analysis into the wavefront-sensing properties of the PL WFS. Denote $\mathbf{u}_{\text{in}}$ and $\mathbf{u}_{\text{out}}$ as the input electric field (located in the pupil-plane) and output electric field (located at the backend of lantern), respectively, of the overall telescope-PL WFS system. The number of PL outputs is N. Following the analysis of the previous section, $\mathbf{u}_{\text{in}}$ and $\mathbf{u}_{\text{out}}$ are related by the complex-valued transfer matrix A. Additionally, assume that there is no flux loss during propagation through the PL. We expand the A

matrix as a product of constituent matrices U, P, and F such that

$$\mathbf{u}_{\text{out}} = UPF\mathbf{u}_{\text{in}}. \tag{5.19}$$

Here, $F \propto -i\mathcal{F}$ is the Fraunhofer propagator, where $\mathcal{F}$ is the Fourier transform. The P matrix determines how the electric field couples into the lantern entrance, which resembles an FMF. More specifically, P projects $\mathbf{u}_{\text{out}}$ onto the basis of the N first guided fiber modes for an FMF matching the geometry of the lantern entrance. For this work, our fiber modes are assumed to be the linearly polarized/LP modes, relevant for weakly guiding, circular, step-index optical fibers, which implies that P is real-valued. Lastly, U is the unitary matrix representing propagation through the lantern. In other words, U transforms a focal-plane electric field, expressed in terms of LP mode amplitudes, into a set of complex-valued SMF amplitudes. Let us further assume phase-only aberrations. Expanding the complex exponential $e^{i\phi}$ with Euler's identity yields

$$\mathbf{u}_{\text{in}} = \mathbf{t} \odot \cos\phi + i\mathbf{t} \odot \sin\phi \tag{5.20}$$

where $\mathbf{t}$ is the real-valued transmission mask of the pupil. We now derive some results.

5.3.1 Impact of perfect mode selectivity

In this section, we show that for an even pupil transmission $\mathbf{t}$, a "perfect" mode-selective lantern (i.e. one free of manufacturing imperfections, which can separate the LP modes with zero crosstalk) maps $\pm\phi$ to the same intensity response. This symmetry in the intensity response makes wavefront reconstruction impossible, preventing mode-selective lanterns from performing effectively as wavefront sensors. First, note that for such a lantern, the propagation matrix U is the identity matrix. Therefore, the complex response of the system for a positive and negative phase aberration is

$$\begin{aligned}\mathbf{u}_{\text{out}}(\pm\phi) &= -iP\mathcal{F}\left[\mathbf{t}\odot\cos\phi \pm i\mathbf{t}\odot\sin\phi\right]\\ &= -iP\left[\mathbf{a}\pm i\mathbf{b}\right]\end{aligned} \tag{5.21}$$

where we have defined

$$\begin{aligned}\mathbf{a} &\equiv \mathcal{F}\left[\mathbf{t}\odot\cos\phi\right],\\ \mathbf{b} &\equiv \mathcal{F}\left[\mathbf{t}\odot\sin\phi\right].\end{aligned} \tag{5.22}$$

We now make use of the following properties of the Fourier transform:

1. The Fourier transform of a real, even function is real and even.
2. The Fourier transform of a real, odd function is imaginary and odd.

First, consider ϕ even. In this case, due to the Fourier transform properties, the real-ness of ϕ, and the symmetry properties of composite functions, both **a** and **b** are real and even. Therefore, the intensity response is

$$\mathbf{p}_{\text{out}}(\pm\phi_{\text{even}}) = |\mathbf{u}_{\text{out}}(\pm\phi_{\text{even}})|^2 = (P\mathbf{a})^2 + (P\mathbf{b})^2. \tag{5.23}$$

For even phase aberrations, the intensity response of a mode-selective PL WFS is even. Next, consider odd phase aberrations. Repeating a similar analysis, we now find that while **a** is still real and even, **b** is now odd and imaginary. Therefore,

$$\mathbf{p}_{\text{out}}(\pm\phi_{\text{odd}}) = |\mathbf{u}_{\text{out}}(\pm\phi_{\text{odd}})|^2 = (P\mathbf{a})^2 + (iP\mathbf{b})^2 \pm 2(P\mathbf{a})\odot(iP\mathbf{b}). \tag{5.24}$$

While an even phase aberration produces a real and imaginary field component, an odd phase aberration produces two real field components that interfere with each other. Under certain circumstances, this interference can break sign ambiguity. However, for the PL WFS, the vectors $\mathbf{a}$ and $\mathbf{b}$ are ultimately projected by P onto the LP mode basis: a basis of real-valued, even and odd electric field distributions. As a result, the last term in equation 5.24 is always 0. This is because $\mathbf{a}$ is even, and only has non-zero overlap with even modes, while $\mathbf{b}$ is odd, and only has non-zero overlap with odd modes. Finally, since any field can be decomposed into an even and odd component, the intensity response of the mode-selective PL WFS is even for all ϕ, at least in the vicinity of the origin.

As a corollary, the above implies that mode-selective lanterns have a linear response matrix $B = 0$.

5.3.2 Non-mode-selectivity can break sign ambiguity

For a non-mode-selective lantern, the matrix U is not the identity matrix; the rows of the matrix UP are the (complex-conjugated) lantern modes. We repeat the analysis from the prior section. The intensity response is

$$\mathbf{p}_{\text{out}}(\pm\phi) = |\mathbf{u}_{\text{out}}(\pm\phi)|^2 = |UP(\mathbf{a} \pm i\mathbf{b})|^2. \tag{5.25}$$

From the above, we see sign ambiguity is broken. The matrix U applies a “rotation” to the vector $P\mathbf{a} + iP\mathbf{b}$. While this rotation preserves the overall norm of the vector, it alters the the modulus of the individual elements, and hence, the powers in the individual ports of the PL WFS.

In other words, switching the sign of a phase aberration is equivalent to conjugating the complex response of the telescope. If we immediately measure the focal plane electric field in the LP mode basis, this conjugation cannot be detected. However, if we apply a unitary transformation (e.g. a PL) after this conjugation, and then measure, the conjugation can be detected.

5.3.3 Conditions for linearity

In this section we derive criteria that the PL WFS must meet to maximize linear sensitivity to a given mode. We will restrict ourselves to the second-order expansion of intensity response for the PL WFS, equation 5.12.

To maximize the linear response of the PL WFS for a particular aberration mode, denoted by unit vector $\hat{\mathbf{z}}_i$, we require that the linear term in equation 5.12 is maximized and the quadratic terms are minimized. We can encourage this behavior by demanding that the quantity

$$Q \equiv [(A\mathbf{1}) \odot (A\hat{\mathbf{z}}_i)^*] \tag{5.26}$$

is purely imaginary. Repeating the same expansion of A from the prior subsections, we equivalently require that

$$Q \equiv [UP\mathcal{F}\mathbf{1} \odot (UP\mathcal{F}\hat{\mathbf{z}}_i)^*] \tag{5.27}$$

is purely imaginary. To connect with the analysis of §5.2.2, note that $B\hat{\mathbf{z}}_i = 2\,\mathrm{Im}\,Q$. Ultimately, linearity imposes a phase restriction on Q: linear response is maximized when Q is purely imaginary, and minimized when Q is purely real. Note that this maximization only enforces that intensity response of the PL WFS is predominately linear in the vicinity of the reference wavefront $e^{i\phi_0}$; this is not a maximization of linear range, although it is likely the first step in an analytically-informed optimization of the latter.

Optimization for the above metric entails designing a PL such that its corresponding propagation matrix U satisfies equation 5.27. This is tricky, but can be simplified in certain cases. In Appendix 5.A, we simplify the above linearity condition for a standard 6-port lantern in the presence of defocus.

5.3.4 WFS limitations

Even with a perfect nonlinear reconstruction model, wavefront sensing breaks down when two distinct phase aberrations can map to the same WFS response. These "degenerate" aberrations are not a concern when the WFS is operating in the linear regime and the mapping of aberrations to sensor intensity responses is one-to-one, but become increasingly problematic as the amplitude of phase aberrations increases. A way to estimate when degenerate aberrations may become problematic is to find input phase aberrations for which a column of the Jacobian becomes zero-valued. This estimate may be conservative, as the response of the PL in this regime can still carry useful information about the input WFE for a subset of the sensed modes. Mathematically we look for an aberration vector $\mathbf{a}_0$ (defined, for instance, in Zernike basis) such that

$$\left.\frac{\partial \mathbf{p}_{\text{out}}}{\partial a_j}\right|_{\mathbf{a}_0} = 0. \tag{5.28}$$

To motivate this criterion, suppose we find some aberration $\mathbf{a}_0$ where the above criterion is fulfilled. In turn, the WFS response about $\mathbf{a}_0$, in the a_j direction, may behave quadratically:

$$\mathbf{p}_{\text{out}}(a_{0,k} + a_j) = \mathbf{p}_{\text{out}}(a_{0,k}) + \frac{\mathbf{p}''_{\text{out}}(a_{0,j})}{2} a_j^2 + o\left(a_j^3\right). \tag{5.29}$$

Here, $a_{0,j}$ is the j_{th} element of $\mathbf{a}_0$. We immediately see that for small a_j, aberrations $a_{0,j} \pm a_j$ map to the same intensity response. More widely separated pairs of degenerate aberrations may also occur around $\mathbf{a}_0$, although they most likely will not be positioned symmetrically about $a_{0,j}$. For an alternative perspective, consider the modal-basis representation of the Jacobian, which has dimensions N rows by M columns for N lantern ports and M aberration modes, with $N \geq M$. The zeroing of a column makes the Jacobian rank-deficient, implying that locally about $\mathbf{a}_0$, the mapping of phase aberrations to PL intensity outputs can no longer be injective. In other words, we are guaranteed scenarios where two or more distinct phase aberrations map to the same intensity response.

The norm (or total RMS WFE) of the smallest aberration vector $\mathbf{a}_0$ which satisfies 5.28 sets the scale in phase aberration space beyond which degeneracy can occur. We term this scale the "degenerate radius." To actually compute the degenerate radius, we take a numerical approach: feeding a standard root-solving algorithm (e.g. Levenberg-Marquardt) a series of random initial guesses in the vicinity of the origin, repeatedly solving 5.28, and then taking the solution with the smallest norm from the returned set. In this approach, we require the full form of the Jacobian for the WFS, without any power series approximations. We derive the following form for the Jacobian:

$$\frac{\partial(p_{\mathrm{out},i}/p_{\mathrm{in}})}{\partial a_k} = -2\,\mathrm{Im}\left[\sum_j A_{ij}e^{i\phi_j}R_{jk} \times \sum_{j'} A^*_{ij'}e^{-i\phi_{j'}}\right]. \tag{5.30}$$

Here, $\phi \equiv R\mathbf{a}$, similar to section §5.2.4, with the exception that we are no longer expanding about some reference phase ϕ_0. A rougher but simpler approximation for the degenerate radius can be made by expanding wavefront response only to second-order: essentially, we set equation 5.18 equal to 0, for fixed aberration index j. This conveniently gives an ordinary matrix-vector equation which can be solved quickly and directly using the Moore-Penrose pseudo-inverse, giving exactly one solution $\mathbf{a}_0$ per aberration. However, this approach can be inaccurate if the WFS response contains little quadratic component.

Lastly, we consider the maximum number of modes that an N-port lantern can sense. In the linear model, it is clear that such a lantern at most can sense N aberration modes. However, this limit holds for nonlinear models as well. This is because our optical system, while nonlinear in intensity, is linear in complex amplitude. A lantern attempting to sense more aberration modes than it has ports is guaranteed to map two distinct phase aberrations to the same complex-valued lantern response, and in turn, the same real-valued intensity. Topological theorems, such as invariance of domain, lead to the same conclusion.

5.4 Simulations

In order to provide the initial steps for general characterization of the PL WFS, we simulate these devices using a numerical model in Python. This model has three primary components: a telescope model, which takes in an incident wavefront and returns a focal plane electric field; a PL propagator, which takes both a focal plane electric field and a lantern geometry, and returns the resulting power distribution of the output ports; and wavefront reconstructer, based on the analysis in Section §5.2. Sections §5.4.1, §5.4.2, and §5.4.3 expand upon these components, respectively. Finally Section §5.4.4 goes over the specific 6-port PL geometries which we simulate with our numerical model.

5.4.1 Telescope simulation

Propagation through telescope optics is handled using the HCIPy package [PHR18]. Simulations are monochromatic, at a wavelength of 1.55 μm. While pure monochromatic light in unrealistic, we anecdotally expect PLs to perform similarly out to bandwidths of perhaps 100 nm: this is roughly the scale of variation we observe in the spectral traces of the dispersed 3- and 19-port PLs installed in the SCExAO/Subaru testbed [LGV20], when injecting a broadband source. We additionally assume a 10 m circular, unobstructed aperture; the focal ratio of the system is optimized to ultimately maximize coupling of an unaberrated wavefront into the PL. Pupil-to-focal plane propagation is handled via HCIPy's Fraunhofer propagator.

5.4.2 Lantern propagation

After computing the focal-plane electric field distribution, the next step is to determine the corresponding electric field at the output of the lantern. To do so, we multiply the electric field vector by the lantern's propagation matrix, UP, which can be computed in pixel basis by discretizing the input plane of the PL and repeatedly propagating single-pixel electric fields.

Alternatively, we can compute the lantern modes for a given PL design by illuminating each single-mode port at the lantern output with its fundamental mode and backpropagating light to the lantern entrance; the complex conjugate of the lantern modes form the rows of the propagation matrix. When the number of PL outputs is less than the number of pixels in the input plane, the backpropagation approach is more efficient; in this work, we use the latter. Numerical propagations through PLs are handled with the Lightbeam Python package [Lin20].

5.4.3 Wavefront reconstruction

Given some PL WFS intensity response, we may now attempt to reconstruct the original phase aberration. Critically, to simplify our models, we neglect the impact of noise; the treatment of noise, and related analyses of PL WFS sensitivity and closed-loop performance, are left for future work. In the meantime, our noiseless model will still be useful for an initial characterization of PL WFS capabilities. We also set our reference wavefront to be flat (i.e. in equation 5.3 we set $\phi_0 = 0$) Our reconstruction model is as follows.

First, we expand phase aberrations in terms of the Zernike basis. To implement linear reconstruction, we compute the matrix B' from equation 5.15; this is done by numerically measuring the matrix of slopes $\partial I_i / \partial a_j$ about the origin. Here, I_i denotes the intensity of the i_{th} output port and a_j denotes the amplitude of the j_{th} Zernike mode, in radians RMS. We then calculate the Moore-Penrose pseudo-inverse of B', which enables inversion of equation 5.15. Note that this reconstruction method neglects any sort of flux normalization, which is unnecessary in the context of simulations but may be more desirable in a more practical implementation with real optics.

In contrast, quadratic reconstruction requires knowledge of the A matrix, equation 5.12, which in turn determines the modal-basis matrices A' and B', and the tensor C'. The A matrix can be computed by probing the pupil-plane electric field (resolved into a 128 by 128 grid of samples) pixel-by-pixel, and measuring the complex-valued response of the PL

WFS, or alternatively through a backpropagation technique like in Section §5.4.2. This is straightforward in the case of simulations, since the complex-valued electric field is known. In contrast, experimental measurement of the A matrix will likely require some phase-diversity method. Inversion of the quadratic model, equation 5.16, is handled using the Levenberg-Marquardt root-finding algorithm, as implemented by the Python package SciPy. We set the starting point of the root-finding routine to the linearly-reconstructed wavefront aberration.

In later numerical explorations, we compare the sensing properties of different kinds of PLs using two standardized prescriptions. In the first, we fix the aberration basis to the first five non-piston Zernike modes; in the second, we use each PL's "control modes": the right-singular vectors of the linear response matrix measured about zero WFE. While sensing is often done with the latter type of basis in conventional AO, the former is useful because it provides a common reference for comparisons.

5.4.4 Simulated lanterns

We choose a 6-port PL geometry as the starting point for our simulated analysis. To demonstrate the validity of the above approach, we simulate wavefront reconstruction with two types of 6-port PL: "standard" and "hybrid" – formal definitions of these designations are in §5.C. Both PLs obey the following assumptions. Firstly, we assume that PLs taper uniformly and linearly so that cross sections of the cores and overall cladding of a PL remain perfectly circular throughout the transition zone. While this is an idealization, it remains a useful starting point for a first-order analysis of the PL WFS, especially since it is unclear whether PL imperfections (such the non-circular claddings exhibited by PLs formed via the tapering of SMF bundles) will help or hurt sensing performance.

Beyond the above idealization, we assume that all PLs taper by a factor of 8 from entrance to exit, with cores spaced in the cladding in such a way that is consistent with the geometries produced when constructing lanterns from a bundle of uniformly sized SMFs. Cladding index is set to 1.444, corresponding to fused silica at 1.55 μm wavelength, while jacket-cladding contrast is set to 5.5×10^{-3}; these parameters are typical for lantern construction (private communication with S. Leon-Saval). Core index is set so that the mode field diameter is ~7.5 μm, matching OFS ClearLite 980 16 fiber. The main difference between our simulated standard and hybrid PLs is in lantern core diameter. In the standard non-selective variant, all SMF cores have the same diameter (4.4 μm), while in the hybrid variant one SMF core is made 2 μm micron larger in diameter to accept the LP_{01} mode. In either case, entrance diameter (i.e. the diameter of the cladding at the input FMF end of the lantern) is set to 20 μm. Additionally, both lanterns have their lengths set by an optimization routine that maximizes for linearity in the lantern's intensity response to the first five non-piston Zernike aberrations. For more details on this procedure, see our companion paper [LJF21].

Lastly, as a sanity check, we also simulate a fully mode-selective variant of the 6-port lantern, to verify our result from §5.3.1 that such lanterns are insensitive to all aberration modes. For simplicity, we assume that the modes of this lantern are exactly the first 6 LP modes, bypassing the need for numerical beam propagation.

5.5 Results for 6-port lanterns

In this section, we apply our numerical model to a standard, hybrid, and mode-selective 6-port lantern. In §5.5.1, we look at the intensity response of these PLs, in the presence of single aberrations. These response curves can be thought of as 1D slices of the PL WFS response "surface" in the presence of many aberration modes. Subsections §5.5.2 and §5.5.3 compare the performances of the linear and quadratic reconstruction models in the presence of the first five non-piston Zernike aberrations. While this basis cannot fully describe "realistic" seeing conditions (and neglects any sort of cross-talk in the reconstruction process from higher order aberration modes) we leave analysis of low-order wavefront reconstruction in the presence of higher-order error for future work. Nevertheless, because the spatial-frequency spectrum of real WFE is typically very bottom-heavy [SFR07, NMJ18], and because PLs are primarily sensitive to low-order modes, we believe that our simplified analysis is still useful.

5.5.1 Intensity response

Figure 5.5.1a shows the intensity response of a standard 6-port lantern as a function of mode amplitude for the first 5 (non-piston) Zernike modes. Empirically, we find that this is the maximum number of modes a 6-port lantern can sense. In [LF22], we find the more general result that an N-port PL can sense at most $N-1$ Zernikes, without additional optics. Our heuristic explanation is that the complex-valued response of an N-port PL is sensitive to piston, which takes up one degree of freedom out of the N total degrees in the system. This piston sensitivity is typically useless for wavefront sensing, and is lost in the conversion of complex amplitude to intensity.

We additionally mark the regions where the linear approximation holds. This "linear range" is defined as the interval in Zernike mode amplitude within which the linear model reconstructs the original phase aberration with less than 0.1 radians RMS of error. Intensity responses to the tilt and astigmatism modes exhibit good linearity in the interval around $[-0.4, +0.4]$ radians, while defocus exhibits linearity over a larger but more asymmetric range: around $[-0.4, 0.8]$ radians. Note that the large linear range for defocus is primarily due to the taper length optimization outlined in §5.4.4. Conversely, certain values of taper length can lead to a lantern that is almost completely insensitive to defocus. We consider this and similar effects in more detail in our companion paper [LJF21].

Figure 5.5.1b shows intensity responses for a 6-port hybrid lantern against the same modes. The introduction of a single, larger lantern core changes the lantern mode structure, both by replacing one of the modes with the LP_{01} mode and by breaking the rotational symmetry of the lantern. We find that the 6-port hybrid lantern begins to behave nonlinearly more quickly than its non-selective counterpart. Additionally, as seen in Figure 5.5.1c, a fully mode-selective 6-port lantern has completely symmetric intensity responses, and therefore is not useful for wavefront sensing. This simulated result corroborates our analytic result

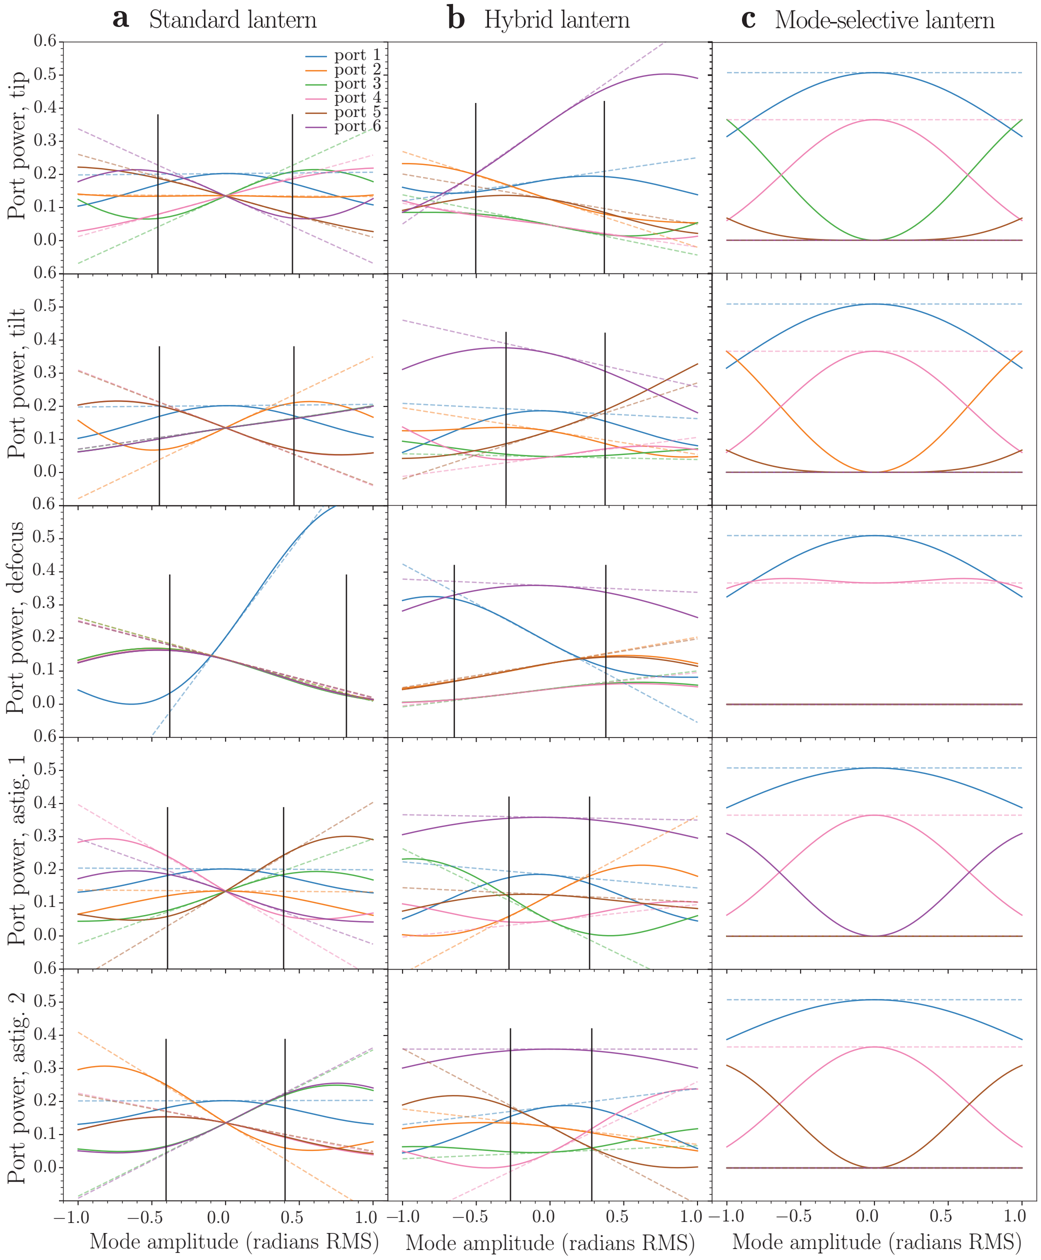

Figure 5.5.1: Column a: intensity response (solid, colored lines) of the 6 SMF outputs for a 6-port, standard lantern, as a function of aberration mode amplitude for Zernike modes 2-5 (tilt, defocus, and astigmatism). Vertical black lines denote the range where the linear model reconstructs the original aberration within 0.1 radians RMS. Dashed lines show the linear approximation for each port's response. Columns b, c: same as column a, but for a hybrid and mode-selective 6-port lantern, respectively.

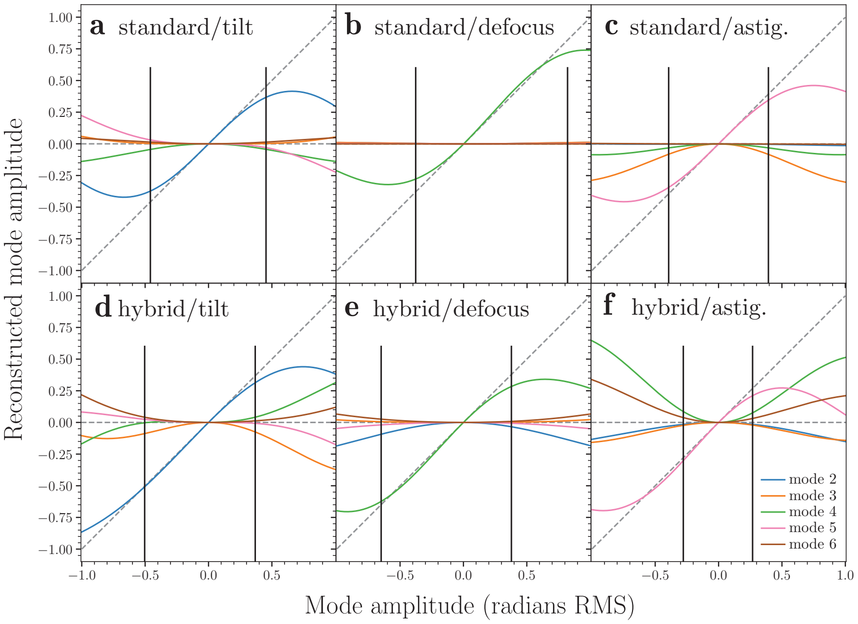

Figure 5.5.2: Panels a,b,c: Aberration reconstruction with the linear model, for various Zernike modes. Aberration amplitude is varied along the horizontal axis; this aberration is propagated into a lantern intensity response through our numerical model. We then attempt to do a linear reconstruction of the aberration amplitude, which we plot along the vertical axis. Under perfect reconstruction, the trace corresponding to the scanned mode should follow the line $y = x$ (marked by the diagonal, dashed gray line) while all other traces should follow $y = 0$ (flat, dashed gray line). As before, vertical black lines mark the region where linear reconstruction is accurate within 0.1 radians RMS. Panels d,e,f: the same as previous panels, but for a 6-port hybrid lantern.

from Section §5.3.1. Finally, we find that our tested 6-port hybrid lanterns outperformed its standard counterpart in terms of degenerate radius (1.3 vs. 0.86 radians). This suggests that hybrid lanterns may exceed standard lanterns when using nonlinear reconstruction methods. Crucially, we emphasize that the above results are for a specific subset of 6-port lantern geometries. Beginning from §5.6, we extend these results to a wider range of PL designs.

5.5.2 Linear reconstruction

Given the intensity responses in Figures 5.5.1, computed over a range of Zernike mode amplitudes, we now apply the linear model (equation 5.15) in an attempt to reconstruct the original mode amplitude. Figure 5.5.2 plots reconstructed aberration mode amplitude against true mode amplitude for Zernike tilt, defocus, and astigmatism, both for a standard and hybrid 6-port lantern. From the Figure, we see that in terms of reconstruction range, the hybrid lantern performs worse than the standard lantern in all modes, particularly in astigmatism. This is in line with results from §5.5.1.

In order to test reconstruction performance in the presence of multiple aberrations, we use a Monte-Carlo approach. We first randomly draw 10,000 aberrated wavefronts (each composed of a random linear combination of Zernike modes 2-6), then pass each wavefront through our PL WFS model to obtain the corresponding intensity response. Given the intensity response, we attempt linear reconstruction. The root-mean-square of the difference between the "true" wavefront and the reconstructed wavefront gives an estimate of the overall accuracy of the reconstruction scheme. Figures 5.5.3a and b plots this accuracy against total aberration, for the standard and hybrid lantern, respectively. From the Figure, we see that reconstruction accuracy for the standard lantern remains under 0.1 radians for wavefront aberrations with up to ~ 0.35 radians of total RMS WFE; the hybrid lantern remains similarly accurate up to a lesser ~ 0.25 radians of total RMS WFE. While this result — that hybrid lanterns behave more nonlinearly than standard lanterns — is specific to 6-port PLs, we find in [LF22] that it also applies for PLs of other sizes.

5.5.3 Quadratic reconstruction

In this subsection we present simulated results for the simplest nonlinear reconstruction method: quadratic reconstruction. This method is based off equation 5.16, which we invert using the Levenberg-Marquardt root-finding algorithm as implemented by the Python package SciPy. For the initial guess required by the root-finder, we use the linearly reconstructed phase aberration vector.

We use the same Monte-Carlo approach outlined in the previous section to test the reconstruction performance of the quadratic model. Our results – reconstruction accuracy against total RMS WFE for 10,000 randomly sampled aberrations – are shown in Figure 5.5.3d and e, for the standard and hybrid 6-port lanterns, respectively. Comparing with panels a and b, which were generated using the linear reconstruction model, we see that the quadratic model lowers the overall error in wavefront reconstruction, as expected. Specifically, for the standard lantern, quadratic reconstruction allows aberrations with up to ~ 0.45 radians of total RMS WFE to be reconstructed to an accuracy of 0.1 radians RMS. The hybrid lantern is similarly accurate up to ~ 0.35 radians of total RMS WFE. These results represents a $\sim 30-40\%$ increase in reconstruction range over the linear model. Notably, the hybrid PL benefits more from quadratic reconstruction than the standard PL, which reinforces the notion that the hybrid PL behaves more nonlinearly.

The quadratic model has the potential to provide even greater gains in reconstruction range when applied to PLs that behave more nonlinearly than the 6-port lanterns tested in this work, whose lengths were specifically optimized to maximize linearity. To show this, we apply the linear and quadratic reconstruction models to a 6-port hybrid lantern without any linearity optimization. Results are shown in Figures 5.5.3c and f, respectively. The large spread and diverging pattern of points in panel c clearly shows the highly nonlinear nature of this particular PL; nevertheless, when switching to quadratic reconstruction model in panel f, the reconstruction error for most aberrations drops dramatically. In scenarios where linearity optimization is infeasible, quadratic reconstruction may provide an alternate path to improving WFS performance.

However, the quadratic model is not without downsides. The major issue is that quadratic reconstruction tends to become increasingly numerically unstable as total RMS WFE increases. We see this behavior reflected in Figure 5.5.3 particularly in panels e and f, where the scatter of points increases substantially with increasing RMS WFE. These instabilities can occur when the root-finder used to invert equation 5.16 gets stuck in a local minimum; a similar phenomenon was observed for the pyramid WFS in [Fra18]). We discuss how this instability may be circumvented in Section §5.8.

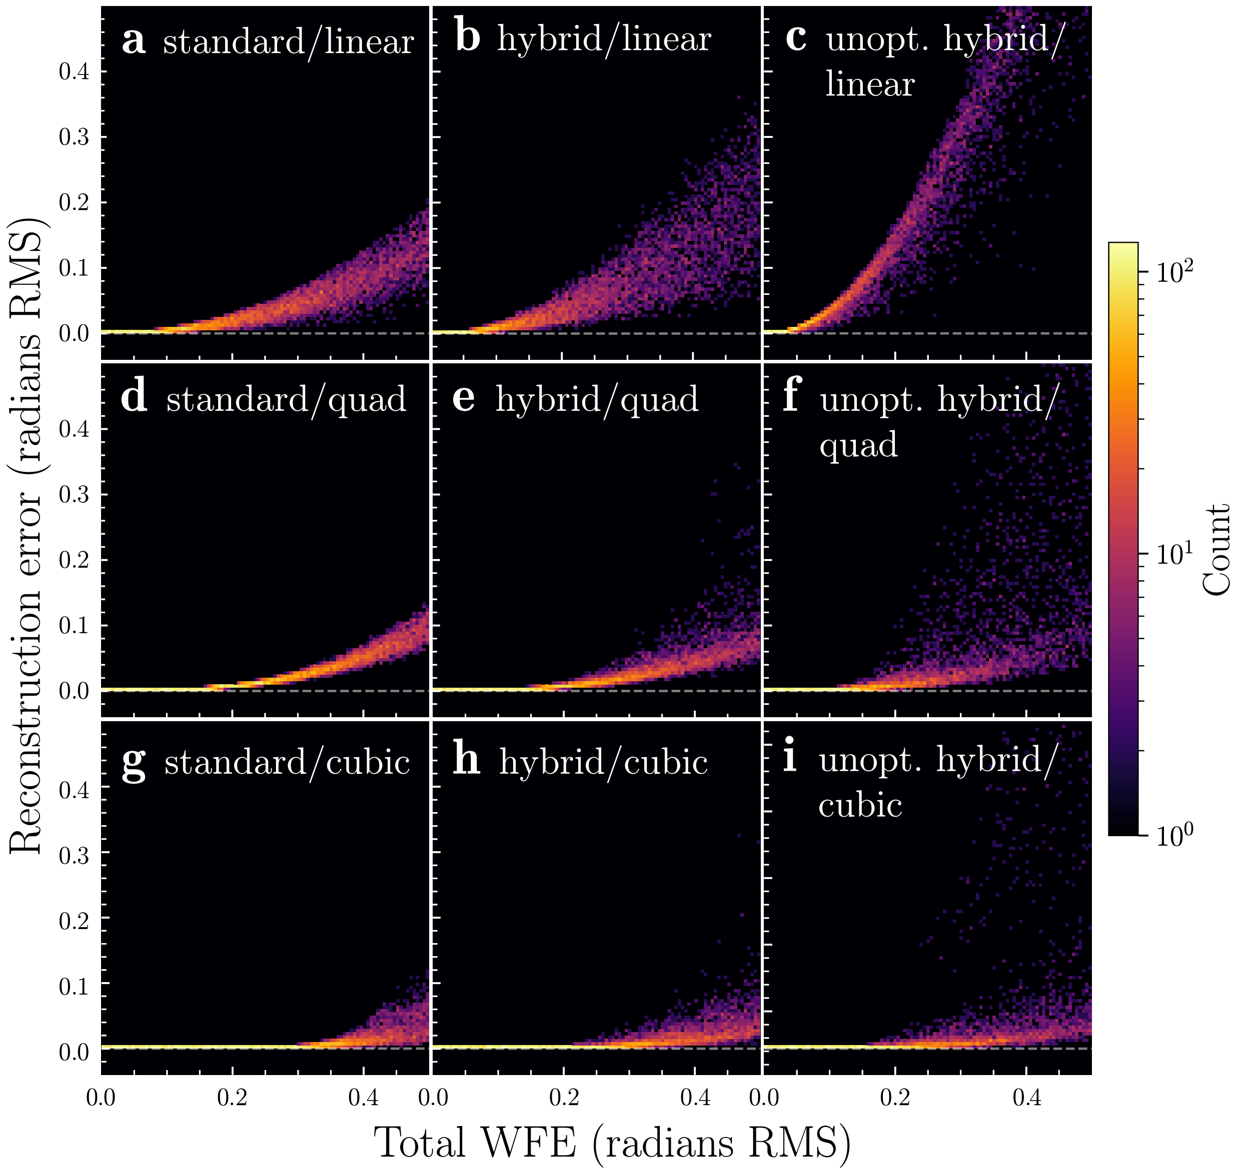

Figure 5.5.3: Panels a,d,g : heatmap of wavefront reconstruction accuracy against total RMS WFE for the standard 6-port lantern in the presence Zernike modes 2-5, using the linear, quadratic, and cubic reconstruction methods, respectively. Reconstruction accuracy was computed for 10,000 randomly sampled phase aberration vectors. A perfect wavefront sensor would have all reconstruction error along the line $y = 0$, shown by the dashed gray line. Panels b,e,h : same as the previous panels but for the hybrid 6-port lantern. Panels c,f,i: same as the previous panels but for a hybrid 6-port lantern without linearity-maximizing taper length optimization.

5.6 Numerical explorations

This and subsequent sections use the same numerical models as above to characterize and compare the WFS properties of a variety of few-moded "standard" and "hybrid" PLs (see §5.C) in the NIR. For simplicity, we neglect noise. While our models are undoubtedly idealizations, they will nonetheless help to determine manufacturing requirements for future PL WFS systems.

Our secondary goal is to find ways to improve the PL WFS. Strategies are divided into three categories: through physical modification of the lantern itself, via geometric properties such as taper length or core arrangement; through "frontend" modification of light upstream of the lantern, via pupil-plane masks and DM control; and through "backend" modification of the light downstream of the lantern, via coherent beam recombination. Frontend and backend optimizations are particularly important because they are independent of the PL, giving pathways through which the PL WFS can be tuned even after the PL has been fabricated. In this work, we present initial investigations into all three strategies. Overall, the questions we address in this work can be divided into two areas — characterization and optimization — and are organized in Table 5.6.1. In §5.6.1, we cover the different PL WFS configurations explored, and in §5.6.2, we define the metrics used to benchmark the sensing performance of those configurations. In §5.6.3, we show the results of our analysis, applied to an ensemble of PL WFS configurations, while in §5.7 we discuss further pathways through which the sensing properties of the PL WFS can be optimized.

5.6.1 PL parameter space

All simulated PLs obey the following assumptions. First, we assume that every PL tapers uniformly and linearly, and that cross-sections of the cores and overall cladding remain perfectly circular through the taper. Additionally, all PLs taper by a factor of 8 from

Table 5.6.1: PL WFS questions

Area	Question	Section(s)
Characterization	What are the useful metrics for benchmarking PL WFSs?	§5.6.2
	What are the impacts of beam-shaping (PIAA) optics?	§5.6.3
	What are the impacts of vortex masks?	§5.6.3
	What approximate wavefront sensing performance might we expect from 3-, 6-, 10-, 12-, and 19-port PLs?	§5.6.3
Design optimization	Can we control sensing properties by tuning PL design parameters such as taper length and core spacing?	§5.7.1-§5.7.2
	What *pre-lantern* optics can we add to improve the PL WFS's sensing properties?	§5.7.3
	What *post-lantern* optics can we add to improve the PL WFS's sensing properties?	§5.7.4

entrance to exit (i.e. the geometry of the small end is the same as the large end, but scaled down by 8×), with cores spaced in the cladding to reflect the geometries produced when constructing lanterns from bundles of uniformly-sized SMFs; for an initial investigation into different core spacings, see §5.7.2. The cladding index is set to 1.444, corresponding to fused silica at $\lambda = 1.55$ μm, while the jacket-cladding contrast is set to 5.5×10^{-3}, corresponding to a fluorine-based glass. The core index is set so that the mode field diameter is $\sim$7.5 μm, matching OFS ClearLite 980 16 fiber.

Under the above assumptions, we consider five sizes of PL: 3-, 6-, 10-, 12- and 19-port, whose output-end geometries are shown in Figure 5.6.1. Each lantern size has two variants: a standard, non-selective variant where all SMF cores have the same diameter (4.4 μm), and a hybrid variant where one core is made 2 μm micron larger in diameter to select out the LP_{01} mode. For all hybrid PLs, the selective core is placed as far off-center as possible, in order to break mode structure symmetries and minimize cross-talk. This choice is especially important for the 6-port hybrid lantern, which would be insensitive to defocus if the central core was made selective. Notably, the hybrid 3-port lantern is technically "mode-group-selective" (see 5.C). By a similar argument to what was laid out for fully mode-selective lanterns in §5.3.1, we also expect mode-group-selective lanterns to have

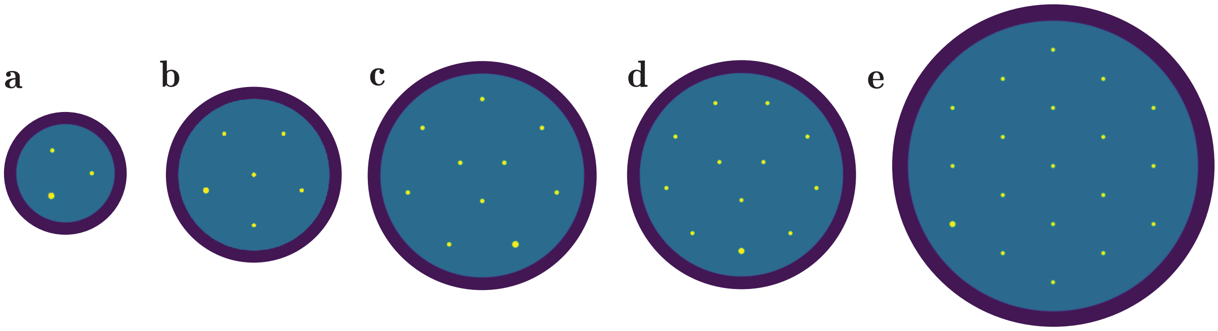

Figure 5.6.1: Refractive index cross-sections for tested PLs. Panels a, b, c, d, and e show geometries for hybrid 3-, 6-, 10-, 12- and 19-port hybrid lanterns, respectively. Non-selective variants are not shown for brevity; they are similar in structure to the hybrid lanterns but with uniform core size.

poor WFS performance without help from devices like backend photonic integrated circuits (PICs). Regardless, we include this lantern in our simulations for completeness.

We additionally optimize the length of the transition region for standard 3- and 6-port lanterns, and the hybrid 6-port lantern, to improve the linearity of their intensity response to phase aberrations. This optimization is done by setting the PL to an initial length (40 mm for the 6-port hybrid lantern and 20 mm for the rest) and entrance diameter (12 μm for the 3-port, 20 μm for the 6-port), and then gradually shortening the lantern, as if material is being cleaved from the lantern entrance, while computing the range in aberration amplitude for which linear reconstruction is accurate to within 0.1 radians RMS, for each aberration mode. The product of these ranges gives a rough metric for linearity (for a more precise but computationally expensive metric, refer to "linear radius", §5.6.2). As the lantern is shortened, and the entrance diameter increases, we also repeatedly apply focal ratio optimizations to keep the coupling of the PSF into the PL maximal. For simplicity, these optimizations were done in the absence of beam-shaping optics and vortex masks; a more in-depth treatment on the interaction between length optimization and different lantern optics is outside the scope of this work. Similarly, we opted not to apply this optimization to the larger 10-, 12-, and 19-port lanterns since the added complexity of these lanterns makes optimization less straightforward: at least anecdotally, gains in reconstruction range in one phase aberration mode typically come at the cost of another, and overall the effect of length optimization on WFS properties seems to be lesser for larger lanterns with more outputs. The 10- and 12-port standard lanterns have fixed lengths of 20 mm, while the hybrid variants have lengths of 60 mm. The standard 19-port lantern has a fixed length of 60 mm while the hybrid has a length of 100 mm. The hybrid lanterns are made longer to reduce cross-talk and improve their mode-selective properties. Entrance diameters are 24 μm, 26 μm, and 37.6 μm for the 10-, 12-, and 19-port lanterns, respectively.

Finally, for each lantern geometry mentioned above, we additionally test the impact of a charge-1 vortex mask of the system, implemented as a helical phase ramp in the pupil plane. As a caveat, the 3-, 10-, and 12-port lanterns with the vortex mask do not function as nullers, but we include their results anyways to assess the impact the vortex mask has on wavefront sensing. We also consider the impact of beam-shaping, but only for the 6-port hybrid PL: we expect beam-shaping optics (which apply an amplitude apodization) to impact WFS properties in a less drastic and more consistent manner than vortex masks, so that any specific results will be qualitatively generalizable to other lanterns.

5.6.2 Performance metrics

We consider the following five metrics:

Throughput We define throughput as the fraction of power of the incident electric field that makes it through the entire telescope + PL WFS system. All optical losses are assumed to occur during coupling into or propagation through the PL; other losses (e.g. from imperfect telescope optics) are neglected.

Sensed low-order Zernikes We consider, for each PL WFS configuration, which subset of the first five non-piston Zernikes (tip/tilt, defocus, and astigmatisms) can be sensed simultaneously - these low-spatial-frequency modes will likely contain the bulk of the power in both atmospheric and instrumental WFE, and form a common basis for later comparisons.

Rank The maximum number of modes that can be sensed using a linear reconstructor is equal to the rank of the interaction matrix (e.g. B in equatioin 5.6). To measure rank numerically for each PL WFS configuration, we form the interaction matrix for the first 30 non-piston Zernikes, and compute the singular value decomposition, which produces the control modes of the sensor. We then count the number of modes with singular value larger than 0.05 simulated flux unit/radian (flux units are normalized to set the power of incoming wavefronts to one).

Linear radius r_L When considering only a single aberration mode, we define the linear range as the interval in mode amplitude within which linear reconstruction is accurate to 0.1 radians RMS. When sensing N modes simultaneously, phase aberrations become an N-length vector of mode amplitudes, and the region where linear reconstruction is accurate becomes N-dimensional. We estimate the volume of this region using Monte Carlo integration. Approximating this ND region as a ball then gives the “linear radius” r_L in units of radians RMS. Aberrations with total RMS WFE less than this radius can generally be reconstructed with good accuracy using the linear model. In principle, non-linear models and closed-loop control should be able to perform wavefront reconstruction beyond this radius.
Because r_L depends on the phase aberration modes being sensed, we compute it for both the low-order Zernike basis and control mode basis.

Degenerate radius r_D Following §5.3.4, the degenerate radius r_D estimates the maximum amount of RMS WFE an aberration can have before the process of phase retrieval becomes ill-posed. Beyond this radius, a WFS may map distinct mode amplitude vectors (each corresponding to a different phase aberration) to the same intensity response, making the propagation process non-invertible without additional regularization. This metric is an approximation, designed to be easily calculable numerically, and the presence of degeneracies do not necessarily impose a fundamental barrier to WFS if the volume of phase space within which degeneracies occur is small. Degeneracy might also be lifted by spectral dispersion of the lantern outputs. Like for r_L, we compute r_D for both the low-order Zernike basis and for control mode basis. The concept of degenerate radius is expanded upon in Chapter §6.

5.6.3 Results

We now use our numerical model to compute the above metrics for various PL WFS configurations. We first consider performance in two contexts: injection for diffraction-limited spectrometers, and VFN. Finally, we review the performance of all tested configurations.

Beam-shaping: Here, we simulate the interaction between beam-shaping optics and PLs. These tests are motivated by SMF-fed, high-resolution spectrometry, where it is desirable from a detector-usage perspective to couple as much light as possible into a single fiber. Accordingly, this section only considers hybrid lanterns, which can couple the bulk of light into their selective port, at least in the absence of wavefront aberrations. The combination of Gaussian beam-shaping optics and hybrid PLs should boost single-port coupling even further, since such optics promote a Gaussian amplitude distribution for incoming light, which matches the approximately Gaussian-like principal mode of hybrid lanterns.

Simulations recover these expectations. Panel a of Figure 5.6.2 shows the intensity response of a 6-port hybrid lantern without beam-shaping optics, in the presence of tilt (Zernike mode 2), while Panel c plots the same for a hybrid lantern with beam-shaping optics. Comparing the panels, we see that beam-shaping more than doubles coupling into port 6 (the selective port) from 35% to 80%. This increase in coupling to the selective port comes at the cost of sensing range: with less light overall in the non-selective lantern ports, the PL WFS behaves more non-linearly. Our results are reiterated in Figures 5.6.2b and d, which show linearly-reconstructed tilt amplitude for the sensing configurations without and with beam-shaping, respectively. Beam-shaping reduces the linear reconstruction range by around 40%. We obtain similar results when scanning over defocus and astigmatism.

Vortex masks: The impact of a charge-1 vortex mask on sensing performance depends strongly on the azimuthal symmetry of the principal lantern modes (as defined in §5.2.1), which in turn depends on factors like core arrangement and lantern length. One of the main impacts of a vortex mask is to alter the sensed set of aberration modes. The clearest example is the hybrid 3-port PL, which cannot sense any modes without the vortex mask, but can sense Zernike mode 6 with the mask in. This effect is shown in Figure 5.6.3, which plots intensity responses in the presence of Zernike mode 6 for the 3-port hybrid lantern, in the absence and presence of a pupil-plane vortex mask: the vortex mask makes the intensity response of the lantern asymmetric, so that the response of the system is no longer degenerate in the linear regime.

While the vortex mask is beneficial in this particular scenario, the effect is not consistent. In fact, the vortex mask can occasionally remove sensitivity to a particular mode (for instance, defocus with the 6-port PL); vortex masks can also impact linear and degenerate radii both positively and negatively. In the next section we present tabulated results that show the effect of the vortex mask on case-by-case basis.

Characterization summary: Here, we evaluate the five metrics detailed in Section §5.6.2 for all tested PL WFS configurations. Results are summarized in Table 5.6.2 for standard PLs, and Table 5.6.3 for hybrid PLs. We also outline some qualitative trends that can be seen in the results, below.

- Consistent with [LJF21], we find that throughput increases with the number of ports.
- Rank also increases with number of ports, although no configuration is able to sense as many modes as it has ports. We believe this is because one of the PL WFS's

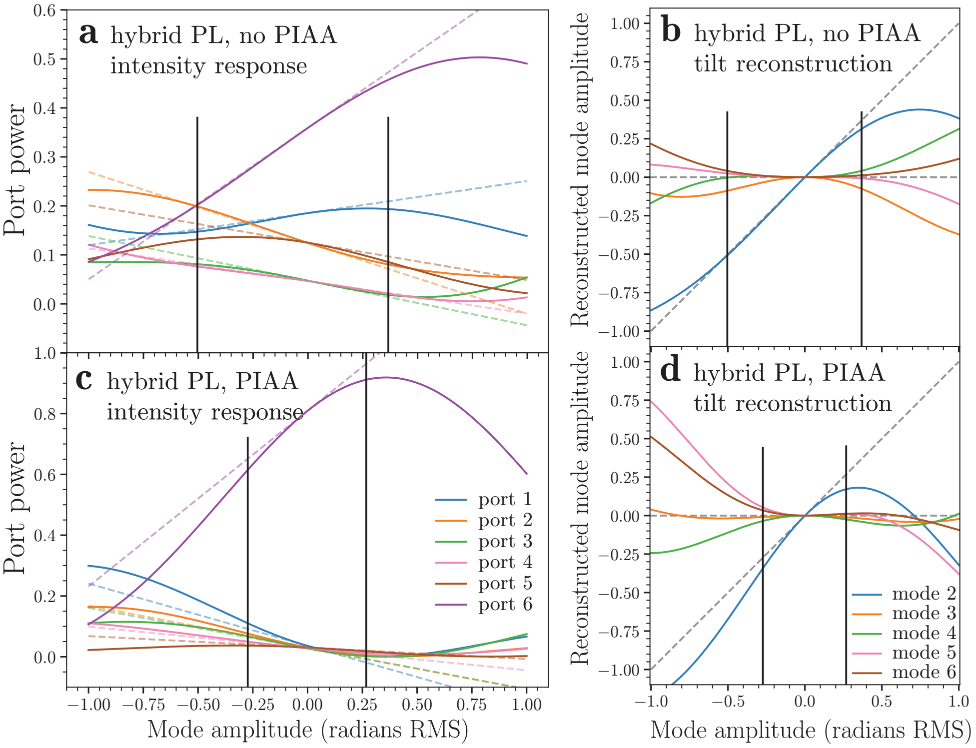

Figure 5.6.2: Panel a: intensity response of the hybrid 6-port lantern in the presence of tilt, Zernike mode 2, without beam-shaping optics. Dashed, colored lines show the linear approximation to the WFS response, while vertical black bars denote the linear range. Panel b: linear reconstruction of the tilt mode for the same lantern. Under perfect reconstruction, the trace corresponding to the scanned mode should follow the line $y = x$ (diagonal dashed line) while all other traces should follow $y = 0$ (horizontal dashed line). Panels a and b are reproduced from Figures 5.5.1 and 5.5.2, respectively. Panels c, d: same as Panels a and b but for a hybrid 6-port lantern with beam-shaping (PIAA optics). The presence of said optics boost coupling into port 6, the science port. As a tradeoff, we see that beam-shaping reduces the linear range of the sensor by around 30–40%.

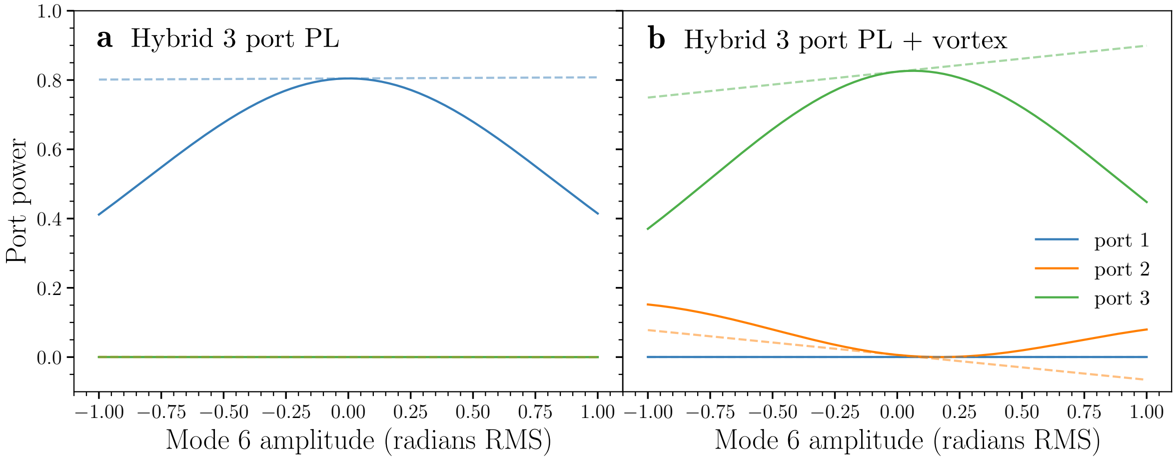

Figure 5.6.3: Intensity response of a hybrid 3-port PL (a) without a pupil-plane vortex mask and (b) with a pupil-plane vortex mask, in the presence of Zernike mode 6. Dashed lines show the slopes of the intensity response in each port, about the origin. Without the vortex mask, only port 1 (the selective port) has non-zero flux, and the response in this port is completely symmetric, so sensing cannot be done. With the vortex, the selective port is nulled, and the symmetry in the intensity response is ultimately broken (albeit only slightly).

degrees of freedom is taken by a piston-like mode, corresponding to the LP01 mode backpropagated to the pupil plane. See further discussion in 5.8.2.

- r_D drops significantly when switching from the low-order Zernike basis to the control mode basis, when the number of control modes exceed the number of sensed Zernikes. In other words, the ability to sense more modes comes at the cost of sensor range.
- For standard lanterns, the degenerate radius in control mode basis is typically improved by the vortex mask; there is no such trend for hybrid lanterns. One explanation is that the vortex mask introduces an asymmetry, where the intensity response of a standard lantern was before symmetric.
- Considering the 10-, 12-, and 19-port PLs, the linear radius in the low-order Zernike basis increases as the number of ports increases. This is expected due to the increased degrees of freedom. The 3- and 6-port lanterns break this trend because they were already optimized for linearity.
- The vortex mask often can add sensitivity to a certain Zernike mode (see the hybrid 3-port PL, and both 10-port PLs), but occasionally removes sensitivity (see both 6-port PLs).
- Configurations using hybrid PLs typically have a smaller linear radius compared to their standard PL counterparts; this is consistent with our result from §5.5 that hybrid lanterns become non-linear more quickly than their standard counterparts. Notable exceptions are the the 10- and 12-port lanterns, without the vortex mask.
- Hybrid lanterns also typically lose rank compared to their standard counterparts. This is unsurprising since the selective port is not useful for wavefront sensing. As an extreme case, the hybrid 3-port lantern has a completely symmetric intensity response for all aberration modes, and thus has both 0 rank and 0 degenerate radius.

- Degenerate radius is typically large by a factor of $\gtrsim 2$, motivating non-linear reconstruction techniques.

Finally, we note that in a few cases, the degenerate radius is actually less than the linear radius, when intuitively we might expect the opposite to always hold. This unintuitive behavior is clearest with the standard 19-port lantern, which in control mode basis has $r_L = 0.30$ and $r_D = 0.21$. The discrepancy is an artifact of how r_L and r_D are defined: recall that r_L is an averaged quantity, but r_D is not. Scenarios where $r_D < r_L$ occur when linearity is poor in one mode but not the rest: then that mode's nonlinearity will negatively impact r_D more than r_L.

Our results establish a rough estimate of PL WFS capability, with an upper limit of around 1–2 radians in RMS WFE before the onset of degeneracies, and an upper limit for linear reconstruction of up to 0.25-0.5 radians, depending on the choice of mode basis. Therefore, in cases where input aberrations are larger, the PL WFS may better suited in a secondary control loop behind a conventional pupil-plane wavefront control system, which can have a larger reconstruction range at the expense of non-common path. However, it remains to be seen if real photonic lanterns, which deviate from the ideal lanterns simulated in this work, will perform better or worse, though the qualitative trends observed in our modelling should be relatively robust. We also emphasize that our numbers are not yet fundamental: increasing the number of ports, exploiting polarization, and dispersing the lantern outputs are all future pathways to improving wavefront sensing capability. Beyond increasing the number of degrees of freedom, we can also make progress by altering the lantern structure and using additional optics. We discuss such prospects in the next section.

Config.	Throughput	Sensed ZMs	ZM r_L	ZM r_D	CM Rank	CM r_L	CM r_D
3	0.815	2-3	0.52	1.0	2	0.54	2.8
6	0.879	2-6	0.46	0.86	5	0.47	0.92
10	0.890	2-5	0.19	2.9	8	0.18	0.41
12	0.895	2-6	0.27	2.0	11	0.28	0.25
19	0.921	2-6	0.49	2.3	15	0.30	0.21
3V	0.843	2-3	0.46	1.9	2	0.51	0.88
6V	0.877	2-3,5-6	0.44	1.0	4	0.45	1.5
10V	0.883	2-6	0.40	2.1	9	0.39	0.90
12V	0.883	2-6	0.47	2.3	11	0.43	0.97
19V	0.894	2-6	0.52	2.0	15	0.30	0.43

Table 5.6.2: **Standard PL performance.** Throughput, rank, sensed Zernikes (out of modes 2-6), linear radius r_L, and degenerate radius r_D for standard 3-, 6-, 10-, 12-, and 19-port lanterns, with and without a vortex mask in the pupil plane. Lantern configurations are identified in the first column by their number of ports; configurations with vortex masks are additionally marked with a capital "V". Linear and degenerate radius are presented for both a low-order Zernike mode ("ZM") basis, formed from whichever of the first five non-piston Zernikes the lantern configuration can sense, and control mode ("CM") basis.

Config.	Throughput	Sensed ZMs	ZM r_L	ZM r_D	CM Rank	CM r_L	CM r_D
3	0.805	None	0	0	0	0	0
6	0.888	2-6	0.43	0.84	5	0.41	0.44
10	0.893	2-3,5-6	0.29	2.7	7	0.25	0.56
12	0.898	2-6	0.39	2.3	10	0.24	0.45
19	0.921	2-6	0.39	3.1	13	0.30	0.87
3V	0.831	6	0.13	-	1	0.13	-
6V	0.884	2-3,5-6	0.36	1.3	4	0.37	1.5
10V	0.887	2-6	0.33	1.6	9	0.26	0.71
12V	0.888	2-6	0.37	2.1	10	0.29	0.54
19V	0.889	2-6	0.50	2.9	15	0.27	0.77

Table 5.6.3: **Hybrid PL performance.** Throughput, rank, sensed Zernikes (out of modes 2-6), linear radius r_L, and degenerate radius r_D for hybrid 3-, 6-, 10-, 12-, and 19- port lanterns, with and without a vortex mask in the pupil plane, analogous to Table 5.6.2. The empty cells for the 3V configuration indicate that no degeneracies could be found, and that the system only becomes degenerate as aberration amplitude goes to infinity and flux in all outputs goes to 0.

5.7 Design optimization

In this section we consider how the wavefront sensing properties of the PL WFS can be improved. One strategy is to modify the structure of the lantern itself; however, optimizations of the optics before and after the PL are also important. Understanding of latter methods will enable more ways to fine-tune the properties of the PL WFS, some of which might be impossible to do through modification of the lantern alone. As a secondary motivation, "optimal" PL designs may not always be feasible to manufacture, but deviations from design specifications may be corrected by pre- and post-lantern signal processing.

We first provide an initial investigation in optimization of the lantern itself, by changing a lantern's taper length and the radial spacing of its cores (§5.7.1–S5.7.2). We then consider how PL WFS performance may be further improved through pre-lantern phase masks (§5.7.3), and in post-lantern beam recombination (§5.7.4).

5.7.1 Lantern length

One way to optimize a PL for wavefront sensing is to tune the lantern's length ("taper length") to maximize a sensing metric such as linear range. As an example, we simulate the effect lantern length has on linear reconstruction range for the standard 3-port PL specified in §5.6.1. Figure 5.7.1a tracks the linear reconstruction range for the tip and tilt modes of the PL as a function of length. Note that we alter length while holding the tapering angle constant, as if cleaving material from the lantern entrance. While the range of lengths that we tested is narrow enough to not alter the number of modes supported at the PL entrance, length also impacts the size of the entrance, requiring re-optimization of our telescope model's F#. For our chosen PL, we find that the linear reconstruction ranges for the tip and tilt modes oscillate with length, at a period of around 300 μm, likely due to beating between modes with different propagation constants. This oscillation period has real-word ramifications: in particular, it sets restrictions in the manufacturing tolerances for taper length, which must be comparatively small in order to enable control over the lantern's sensing properties.

Figure 5.7.1b (blue curve) shows similar oscillations in linear reconstruction range for the standard 6-port lantern, specifically for the defocus mode. These oscillations are reminiscent of a phenomenon discussed in §5.A, which showed that defocus *sensitivity* (not reconstruction range, though the two metrics are related) peaks when the LP_{01} and LP_{02} components of the central lantern mode (i.e. the mode corresponding to the central port of the lantern) are 90° out of phase. In turn, oscillations in relative phase (which would lead to alternation in defocus reconstruction range, as seen in Figure 5.7.1b) might occur as a function of lantern length, since different modes travel with a different effective refractive index through the lantern.

While oscillations in reconstruction range also occur for higher mode-count lanterns, the addition of more lantern modes complicates the optimization process. Different sets of modes within may oscillate at different periods, so that gains in for one mode are more likely to come at the cost of another. We leave this for future work.

5.7.2 Core arrangement

The geometric arrangement of single-mode cores within a standard PL can also strongly impact wavefront sensing ability [FRB12]. For example, prior work in the telecommunications field has shown that a PL's "rings" (sets of cores that are the same radial distance from the PL's center) should contain odd numbers of cores in order to break symmetries in the lantern mode structure [DDM21]. In this section, we show that the radial spacing between rings also impacts wavefront sensing, using the standard 6-port lantern (assumed to be fabricated by tapering a bundle of SMFs) as an example. For a PL constructed from uniformly-sized SMFs, the outer ring of cores will be located at a distance of $2/3\, r_{\mathrm{clad}}$ from the lantern's geometric center, where r_{clad} is the radius of the lantern's cladding-jacket interface. However, by changing the placement of the outer ring (which can be achieved by etching down SMFs to different outer diameters, or inserting glass spacers into the fiber bundle before tapering), we

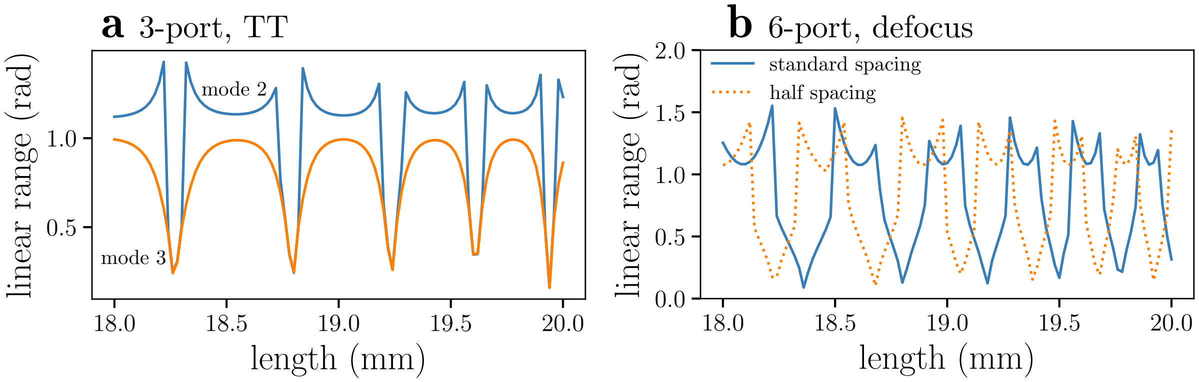

Figure 5.7.1: Panel a: Oscillatory behavior of tip (Zernike mode 2) and tilt (Zernike mode 3) reconstruction range for a standard 3-port lantern, as a function of lantern length. These oscillations occur over a period of around 300 μm. Panel b: Oscillatory behavior of defocus (Zernike mode 4) reconstruction range for a non-mode-selective 6-port lantern, as a function of lantern length. Curves for two lanterns are shown: the first has “standard” radial spacing, where the outer single-mode cores are placed two-thirds along the way from the lantern’s center to the cladding-jacket interface. The second lantern places the outer lantern cores halfway between the lantern’s center and the cladding-jacket interface. This change in radial spacing displaces the defocus reconstruction oscillations.

can alter the 6-port lantern's response to modes such as defocus. Figure 5.7.1b plots defocus reconstruction range against lantern length for two kinds of non-selective 6-port lanterns: one where the outer cores are placed at a distance of $2/3\, r_{\rm clad}$ from the center, and one where they are placed $1/2\, r_{\rm clad}$ from the center. Changing the radial spacing of the lantern cores causes the oscillations in defocus reconstruction range to shift. Therefore, a 6-port lantern's linear range in the defocus mode can be tuned either by altering taper length or core spacing (though the former is likely easier to control).

Core arrangement also matters for hybrid lanterns. One trend we expect from radial symmetry arguments, and have confirmed through simulation, is that the selective core should not be placed in the lantern's center: doing so can remove sensitivity to circularly symmetric aberration modes.

5.7.3 Improved sensing with pre-lantern phase offsets

In this section, we consider how the manipulation of light before the lantern can change a lantern's sensing characteristics. One such method is to apply a phase offset (either using a dedicated mask in the pupil plane, or a DM). As a proof-of-concept, we show that phase offsets can increase the number of sensed modes. Our fiducial example is the standard 6-port lantern, which alone can sense the first 5 non-piston Zernike modes. We find that we can add a sixth sensed mode simply by adding a phase mask that resembles the desired mode; this static wavefront error biases the intensity response (e.g. Figure 5.5.1) of the PL WFS, changing the slopes and enabling the mode to be sensed in the linear regime. As a tradeoff, the increase in rank from the phase mask will come at the cost of throughput. We numerically verify this method by applying a spherical aberration phase offset (mode amplitude -0.5 radians) in the pupil plane of our 6-port PL WFS model; its intensity response in the presence of spherical aberration, as well as the linearly reconstructed mode amplitudes, are shown in Figure 5.7.2. Comparing the throughput indicated by Figure 5.7.2a (∼0.8) to that of a standard 6-port lantern without a phase mask (∼0.88, as per Table 5.6.2), we see

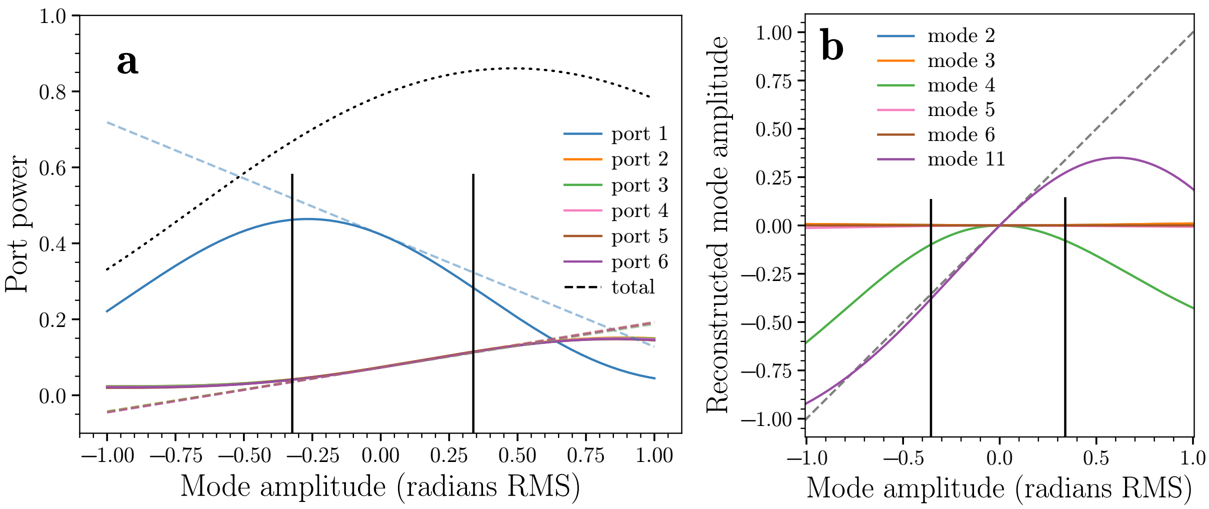

Figure 5.7.2: Panel a: intensity response of a standard 6-port lantern; a pupil-plane phase mask applies a static amount of spherical aberration (mode amplitude -0.5 radians) before injection into the PL. The dotted black curve tracks the total flux in all lantern outputs. Like in Figure 5.6.2, the dashed, colored lines show the linearly-approximated WFS response, while vertical black bars denote linear range. Panel b: linearly reconstructed spherical aberration mode amplitude as a function of true mode amplitude.

that the cost of this added sensitivity to spherical aberration is a ∼10% drop in throughput. We also note that the linear range of the sensor in this setup is limited by cross-talk between the defocus and spherical aberration modes.

5.7.4 Improved sensing with post-lantern beam-combination

Because PLs coherently couple light into single-mode cores, they naturally serve as a platform for the coherent manipulation of light through systems of beam splitters/combiners [DM19]. These beam recombination systems have been produced within monolithic glass blocks using ultra-fast laser inscription (e.g. GLINT [MNT21]). Such devices can apply linear transformations to the complex-valued vector of PL outputs; as a result, compared to the phase mask approach of the previous section, we anticipate that beam recombiners will be able to more drastically change sensing characteristics, theoretically without loss of light. These transformations could be used to improve WFS capabilities, or to correct imperfections in lantern mode structure left by the manufacturing process. As a downside, this method requires more complicated, dedicated optics for complex beam-recombination. Such optics will likely be required for coherence-based exoplanet detection methods [KFL24].

As a proof-of-concept, we consider a simple beam recombination setup for PIAA-equipped hybrid lanterns, where the bulk of telescope light is coupled into a single "selective" output to be used for science. In this case, the PL WFS behaves more non-linearly, since all ports besides the science port now have near-zero intensity. To improve linearity, we can divert a small fraction f of the light from the selective port, split it into $N-1$ equal parts, and pairwise-interfere each part with one of the $N-1$ non-selective ports, where N is the number of PL outputs. For illustration purposes, we assume a unitary transfer matrix $A_{\,50:50}$ for each beam pair of the form

$$A_{\,50:50} = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}. \tag{5.31}$$

The above matrix corresponds to a "50:50" beam splitter; the overall beam recombination system turns N PL outputs into $2N-1$ system outputs. Figure 5.7.3 shows the effect of such a beam recombiner for a 6-port hybrid PL (the same as the one simulated in §5.5, with beam-shaping optics), reconstructing tilt with the linear model. We find that this method trades throughput in the science port for linearity; the beam recombiner lowers flux in the selective port by 10% but improves reconstruction range by $\sim 30\% - 40\%$.

In the future, we envision recombiner designs that improve sensing properties while preserving the number of PL outputs. Recombiners can also be constructed from more advanced building blocks such as tricouplers or photonic Mach-Zehnder interferometers; the latter architecture can enable the application of an arbitrary unitary transformation to the PL outputs [RZB94]. These possibilities are further discussed in §5.8.

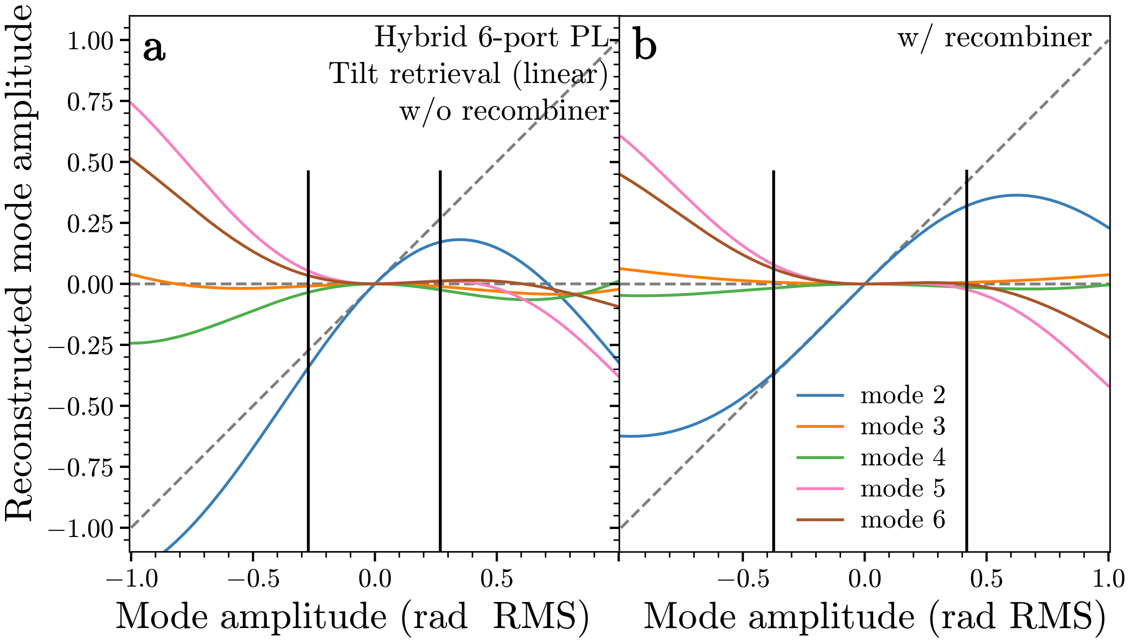

Figure 5.7.3: A beam recombiner, applied to the outputs of a 6-port hybrid PL, can increase the PL's linear reconstruction range. In this example, the simulated beam recombiner siphons 10% of the light from the selective output and interferes it with each of the remaining ports, so that the 6 original outputs of the hybrid PL become 11 outputs. Panel a: reconstruction plot for tilt (Zernike mode 2), using a hybrid PL without the backend beam recombiner. Panel b: the same as a, but with the recombiner. The recombiner increases the reconstruction range for tilt by $\sim 30\% - 40\%$

5.8 Discussion

5.8.1 A framework for the PL WFS

In Section §5.2, we laid out a general mathematical framework, in arbitrary modal basis, for the intensity response of a WFS to errors in phase. While we recover the usual linear model in our first-order expansion, we additionally derive a quadratic reconstruction model; however, it remains to be seen whether the increased accuracy afforded by the quadratic and higher-order models outweigh the penalties in numerical stability and computation time, and if these techniques can be applied to closed-loop operation. We expect additional complications when moving to wavefront reconstruction with real PL WFSs. For one, we will have to contend with detector and photon noise, which will degrade both sensitivity and reconstruction range. Noise will likely be particularly problematic at the kHz framerates used for atmospheric compensation, but may be less of an issue when sensing slower NCPAs. An additional complication is that, in practice, the complex-valued A matrix must be experimentally determined (e.g. through phase diversity methods), and hence will be prone to the effects of random and systematic uncertainties. While linear reconstruction, which requires only intensity knowledge, will be largely unaffected, uncertainties in A may make nonlinear reconstruction even more numerically unstable. These uncertainties may be mitigated if we can constrain the A matrix (for instance, through its modulus, or through B).

We imagine several potential next steps in our mathematical analysis. One continuation is to extend our phase-only aberration analysis to amplitude aberrations as well. Another is the expansion of WFS intensity response to third order, which we begin in Appendix 5.B; see also Figure 5.5.3g-i, which plots the similar reconstruction heatmaps as panels a-f but for a cubic reconstruction model. Figure 5.5.3 confirms that this expansion can offer a significant boost in reconstruction accuracy, especially for the 6-port standard lantern. However, the drawbacks are similar to the quadratic model. Each increase in order is accompanied by an increase in the model degeneracy, as well as an increase in the rank of the tensors required by the model.

More advanced reconstruction models may overcome these drawbacks. For one, stochastic optimization algorithms like simulated annealing, while computationally expensive, are one potential way to avoid local minima in the inversion process. Another idea is to use wavelength diversity, leveraging the chromatic dependence of the PL WFS response: extra measurements at multiple wavelengths may make the reconstruction process for our nonlinear models significantly easier. These measurements can be made though spectral dispersion of the PL WFS outputs, as in the so-called photonic "TIGER" configuration [LBB12]. Lastly, we emphasize that while going to higher order may amplify numerical instability, it does not amplify experimental uncertainties in the A matrix; this is because intensity will always have a second-order dependence on complex amplitude.

Besides enabling wavefront reconstruction, mathematical models have a second, important use: they allow us to derive certain WFS properties and metrics through which the WFS can be optimized. For instance, in §5.3.1, we derived that a fully mode-selective lantern is insensitive to phase aberrations, for even pupil transmission. It remains to be shown whether or not this limitation can be practically overcome with pupil masks or other additional optics. In contrast, there are no such restrictions with standard and hybrid lanterns. As a corollary, we found that the linearity of the PL WFS, at least for small aberrations, depends on the phase of what we call the Q metric (equation 5.27). We also show how this linearity condition simplifies for certain cases, such as the 6-port standard PL in the presence of defocus (Appendix 5.A). In the future, it may be desirable to optimize the PL WFS for this linearity condition. However, if nonlinear reconstruction methods, such as the quadratic or cubic methods in this work or the neural-net approach from [NWB20], can be developed that are fast and stable enough to compete with linear reconstruction, it may instead be desirable to optimize lanterns according to degenerate radius (equation 5.29). Both the linear Q metric and the degenerate radius are only the first steps in analytically defining the sensing properties of the PL WFS.

5.8.2 Numerical comparisons of PL WFS designs

We observed several trends in the WFS metrics computed by our numerical model. As expected for both standard and hybrid lanterns, throughput and the number of sensed modes increase with number of ports. We also found that hybrid lanterns exhibit a slightly higher degree of non-linearity over their standard counterparts, perhaps due to their more complicated mode structures. Linear reconstruction ranges of $\lesssim 0.5$ radians RMS imply that PL WFSs will likely need to work in tandem with a first-stage pupil-plane WFS: these pupil-plane control schemes can remove larger WFEs (such as those from atmospheric turbulence) while leaving the correction of lower-amplitude residuals, produced by NCPAs and potentially even petaling modes, to the PL WFS. However, in the future these limitations may be lifted by exploiting spectral dispersion and polarization, employing more advanced non-linear reconstruction methods, and increasing mode counts.

We also note that for all sensor configurations tested in §5.6.3, the number of sensed modes was strictly less than the number of outputs; we propose that this deficiency is due to a piston-like mode corresponding to the LP_{01}, backpropagated to the pupil plane, which takes up one degree of freedom but is not useful for wavefront sensing. This explanation is corroborated by §5.7.3, where we showed that a 6-port PL could sense a full 6 Zernike modes though use of a pupil-plane phase mask, in essence replacing piston with a more useful mode. Beyond the aforementioned effect, we also expect to lose an additional degree of freedom in real systems because the sum of all sensor outputs can carry no WFS signature: there are many high-spatial-frequency aberration modes which will induce a drop in overall throughput. Conversely, in simulations where we inject and retrieve exactly the same number of modes, and ignore confusion with higher-order modes, the sum of all ports can have WFS signature.

We additionally test the impact beam-shaping optics and vortex masks have on the sensing performance of the PL WFS. Tests with the former are motivated by high-resolution fiber-fed spectrometry, where it is desirable to inject as much light as possible into the single diffraction-limited fiber that feeds the spectrometer, and then use the remaining light for wavefront sensing. We found that beam-shaping can route upwards of 80% of total light into the selective core, at the cost of sensor linearity: the lack of light in the non-selective cores forces responses in those cores to behave more quadratically. When accounting for noise, the low fluxes in the non-selective ports will hinder wavefront sensing even more. We envision several ways to circumvent this issue, for instance by altering the PIAA lens design to leave some minimum amount of light in the non-selective ports, applying phase offsets with DMs as in §5.7.3, or applying beam recombiners as in §5.7.4.

Our tests with vortex masks are motivated by VFN. In this setup, we envision using a vortex mask and a PL to simultaneously perform VFN in one port and wavefront sensing/control with the other ports (for an alternative VFN setup using fully mode-selective lanterns, see [XJR22]). We find that the interaction between vortex masks and PLs is complex, alternately removing sensitivity to aberration modes, swapping sensitivity between modes, or adding sensitivity to a completely new mode. Interestingly, we also find that vortex masks can significantly increase the degenerate radius for standard PLs; heuristically, we believe this is due to the vortex mask introducing an asymmetry into the system. Further optimizations in this sort of setup will need to be made on a case-by-case basis, owing to the complex interaction between vortex masks and the PL WFS. In cases where the lantern modes have certain azimuthal symmetries (for instance, partially mode-selective lanterns), the impact of vortex masks may be more analytically tractable.

5.8.3 Improving the PL WFS

Sensing quality can be improved in two areas: linear reconstruction range, and the number of sensed modes. One strategy to improve reconstruction range is to adjust lantern length: as in Figure 5.7.1, we found that reconstruction ranges for certain modes can oscillate from 0, marking complete insensitivity, to some maximal value as length is varied. Note that such a strategy is only possible if the manufacturing tolerances in PL length are small compared to the oscillation period (~ 300 μm for our simulated 3-port and 6-port lanterns); alternatively, lantern length can be controlled to a lesser extent by etching back the lantern entrance post-fabrication, although care should be taken not to widen the lantern entrance so much that the shape or number of modes supported at the lantern entrance changes drastically. The radial spacing between lantern cores also impacts wavefront sensing, as was shown in section §5.7.2. In the future, we might envision even more exotic core arrangements with asymmetrically-placed cores. Because the parameter space of PL geometries is so large, the identification of "good" PL designs will likely require the use of waveguide mode-solvers working in tandem with numerical optimization techniques.

But improvement to PL WFS performance does not necessarily require alteration of the PL itself. Instead, we can manipulate light before and after the PL, through optics such as phase masks and beam recombiners. The first technique, considered in section §5.7.3, involves applying a static phase offset in the pupil plane of the telescope-PL WFS system. More refined phase masks, designed through global numerical optimization methods, may be able to boost sensing quality while preserving throughput. Another natural continuation would be to consider masks (or dual-DM setups) that alter both amplitude and phase. We leave both ideas for future work. Phase offsets applied by DMs have the added benefit that they can be dynamically controlled, perhaps ramping down in amplitude as the overall wavefront error decreases, or switching on the fly to modify the set of sensed aberration mode as needed.

As shown in section §5.7.4, we can also improve PL WFS performance by using chains of beam splitters to repeatedly interfere the PL's outputs. Mathematically, this process applies a matrix transformation to the vector of complex SMF mode amplitudes at the lantern output, and can thereby alter the linearity of the system through the Q metric (§5.3). As a proof-of-concept, we used a system of 50:50 beam splitters to improve the linear reconstruction range of a hybrid lantern; the next step will be to design more complex beam-combining systems. One promising architecture is the "Mach-Zehnder" mesh [ZDC22] — a photonic circuit composed of multiple stages of Mach-Zehnder interferometers which repeatedly interfere pairs of beams. Such a mesh, combined with piezoelectrically or thermally tunable phase delays [DCL22], could enable active control of the PL outputs. The photonic circuit could be modulated to extend sensitivity to more aberration modes, increase the overall reconstruction range of the sensor, correct for PL manufacturing imperfections, and even function as a secondary stage of wavefront control.

In all, there are four general classes of methods through which we can improve PL WFS performance: "frontend" methods like phase masks, which manipulate light before the PL; "intrinsic" methods which directly alter the lantern geometry; "backend" methods like beam-combining, which manipulate confined light after the PL; and finally, post-processing methods (e.g. the reconstruction model). All four are important in achieving greater control over the sensing properties of the PL WFS.

5.9 Conclusion

In this work, we provide an end-to-end mathematical analysis of the PL WFS. In Sections §5.2 and §5.3, we developed linear and higher-order mathematical models for the intensity response of the PL WFS. These models enable the reconstruction of wavefront aberrations from intensity responses, and enable the derivation of certain metrics, such as the degenerate radius, which estimates the maximum amount of RMS WFE an aberration can have before the mapping of aberrations to intensities is no longer one-to-one.

We also applied our numerical model to benchmark the sensing performance of a variety of PL WFS configurations, covering both standard PLs and hybrid PLs, as well as the impact of beam-shaping optics in §5.6.3 and vortex masks in §5.6.3. We find that, in conjunction with a hybrid lantern, beam-shaping can couple the bulk of incident light into the selective core, at the cost of a significant decrease of $\sim 40\%$ in linear reconstruction range. The interaction between vortex masks and PL WFSs is less clear: alternatively adding to, changing, or subtracting from the set of sensed modes. Results are tabulated in Tables 5.6.2 and 5.6.3. In general, linear reconstruction ranges are found to be $\lesssim 0.5$ radians RMS, and degeneracies in the sensor response typically appear when WFE reaches 1-2 radians RMS. Our results do not mark a fundamental upper limit to the performance of the PL WFS: future pathways for improvement include spectral dispersion of lantern outputs, exploiting polarization, and switching from linear to more advanced non-linear and/or neural net reconstruction methods. We also consider how sensing performance could be improved by altering aspects of the PL WFS design in §5.7, through application of phase masks in §5.7.3, and through backend beam recombination in §5.7.4.

Appendix - Chapter 5

5.A Defocus performance for standard 6-port lantern

The linearity criterion from §5.3.3 can be simplified for a standard 6-port lantern, located in the focal plane of a telescope with a filled circular aperture, in the presence of defocus. We order our basis of LP modes as (LP_{01}, LP_{02}, rest of the LP modes) and our output ports as (central port, rest of the ports). For simplicity, we also assume a reference phase $\phi_0 = 0$. Due to symmetry, both an unaberrated wavefront and a defocused wavefront will only couple into LP_{01} and LP_{02}. Furthermore, the coupling coefficients will be real. Therefore, we can set

$$P\mathcal{F}\mathbf{1} \equiv \begin{pmatrix} a \\ b \\ 0 \\ \vdots \end{pmatrix}, P\mathcal{F}\mathbf{z} \equiv \begin{pmatrix} d \\ f \\ 0 \\ \vdots \end{pmatrix} \tag{5.32}$$

where a, b, c, d are real numbers and $\mathbf{z}$ is the vector corresponding to the defocus mode. Denoting the columns of the lantern propagation matrix U as $\mathbf{c}_i$, we find that equation 5.27

becomes

$$Q = ad|\mathbf{c}_1|^2 + af\mathbf{c}_1 \odot \mathbf{c}_2^* + bd\mathbf{c}_2 \odot \mathbf{c}_1^* + bf|\mathbf{c}_2|^2. \tag{5.33}$$

We want Q to be "as imaginary as possible." Clearly, the first and last terms are real, so a lantern that satisfies

$$ad|\mathbf{c}_1|^2 + bf|\mathbf{c}_2|^2 = 0 \tag{5.34}$$

will behave "more linearly" than one that doesn't. The middle terms apply another condition: the each element in $\mathbf{c}_1$ should be 90° out of phase with its corresponding element in $\mathbf{c}_2$. In turn, this condition implies that the LP_{01} and LP_{02} components for each lantern mode must be 90° out of phase. We have verified this behavior numerically.

It is also useful to consider the converse of the above conclusion. Suppose that the LP_{01} and LP_{02} mode coefficients are in phase. Then, $\mathbf{c}_1 \odot \mathbf{c}_2^*$ will be real, and Q will be purely real. Consequently, the linear B matrix will be 0 - lantern response is locally quadratic.

5.B Cubic expansion

Expand the incident electric field with a phase ϕ about a reference phase ϕ_0:

$$\mathbf{u}_{\text{in}} = e^{i\phi_0} \odot \left[\mathbf{1} + i\boldsymbol{\Delta}\phi - \frac{1}{2}\boldsymbol{\Delta}\phi^2 - \frac{i}{6}\boldsymbol{\Delta}\phi^3 + o\left(\boldsymbol{\Delta}\phi^4\right)\right], \tag{5.35}$$

where $\boldsymbol{\Delta}\phi \equiv \phi - \phi_0$. As before, the intensity response of the WFS is

$$\mathbf{p}_{\text{out}} = \left|A\mathbf{u}_{\text{in}}\right|^2 \tag{5.36}$$

where A is the complex-valued transfer matrix of the overall optical system. Modifying $A_{ij} \rightarrow A_{ij}e^{i\phi_{0,j}}$ and combining the above two equations, keeping only terms up to third order, yields

$$\mathbf{p}_{\text{out}} \approx \mathbf{p}_{\text{out, quad}} - \frac{1}{3}\text{Im}\left[A\mathbf{1} \odot A^* \boldsymbol{\Delta}\phi^3\right] + \text{Im}\left[A\Delta\phi \odot A^* \boldsymbol{\Delta}\phi^2\right] \tag{5.37}$$

where $\mathbf{p}_{\text{out, quad}}$ is the quadratic approximation for output intensity, as per equation 5.12. We now expand the above model to an arbitrary modal basis. Let R be a change-of-basis matrix, such that $\boldsymbol{\Delta}\phi = R\mathbf{a}$. The additional terms from the cubic expansion can be expressed as a single tensor multiplication of the form

$$\sum_{lmn} D'_{ilmn}\, a_l a_m a_n \tag{5.38}$$

where the tensor D'_{ilmn} is defined as

$$D'_{ilmn} = \text{Im}\left[-\frac{1}{3}\sum_j A_{ij}\sum_k A^*_{ik}R_{kl}R_{km}R_{kn} + \sum_{jk} A_{ij}A^*_{ik}R_{jl}R_{km}R_{kn}\right]. \tag{5.39}$$

The D' tensor has dimensions $N \times M \times M \times M$ for an N-port lantern sensing M aberration modes. The full cubic model, in modal basis, is

$$\mathbf{p}_{\text{out},i} \approx |A\mathbf{1}|_i^2 + (B'\mathbf{a})_i - \frac{1}{2}\sum_{jk} C'_{ijk}a_j a_k + |A'\mathbf{a}|_i^2 + \sum_{lmn} D'_{ilmn}\, a_l a_m a_n \tag{5.40}$$

Brief empirical testing with this model shows that it can provide a significant increase in reconstruction accuracy, especially for PLs that have already been optimized for linearity. Heatmaps of reconstruction error against total RMS WFE for 10,000 randomly sampled aberrations are shown in Figure 5.5.3g, h, and i, for various 6-port lantern designs. Notably, going to higher order consistently extends the reconstruction range of the sensor, suggesting that the main downside of going to a higher-order model is additional computational complexity rather than numerical instability, at least for the first few orders.

5.C Types of photonic lanterns

There are four types of PLs referenced throughout this work; to distinguish between them, we will consider the lantern's principal modes, the N complex-valued distributions of light at the lantern entrance, each of which an N-port lantern couples into a single one of its outputs. Alternatively, the principal modes can be thought of as the front-end fields produced when reverse-injecting the PL through each of its SMF outputs. The prinicpal modes are occasionally referred to simply as "lantern modes", e.g. [LFS14, LJF21].

Differences can also be expressed through each lantern's complex-valued transfer matrix U, the unitary matrix that transforms the input vector expressing the electric field at the lantern entrance (in the basis of "entrance modes", the eigenmodes supported by the fiber-like lantern entrance) into a complex-valued vector of single-mode output amplitudes. In the limit where the single-mode cores taper to negligible size at the lantern entrance, and assuming circular symmetry, the entrance modes are the LP modes.

Non-selective or "standard": A non-selective, non-mode-selective, or "standard" photonic lantern contains single-mode cores of uniform size and refractive index. As such, modes in any basis propagating through the lantern become degenerate in their propagation constants, leading to cross-coupling. The principal modes of a standard lantern are complex linear combinations of the entrance modes, and the transfer matrix U has no clear properties other than being unitary.

Partially-selective or "hybrid": In general, a "hybrid" or partially-selective PL "selects" a subset of the entrance modes, exclusively routing each to its own single-moded output; light in the remaining entrance modes is coupled non-selectively into the remaining outputs. The hybrid PL considered throughout this work and in [NBW22] is specifically "fundamental-mode-selective", with only one core larger (or of higher index) than the rest. The output corresponding to this core exclusively couples all light from the fundamental entrance mode. The principal modes for the smaller "non-selective" cores are linear combinations of entrance modes, excluding the fundamental. The transfer matrix for an N-port hybrid PL that selects out only the fundamental mode has the block diagonal form of

$$U = \begin{bmatrix} 1 & 0 \\ 0 & U_{\text{N-1}} \end{bmatrix} \tag{5.41}$$

where $U_{\text{N-1}}$ is an $N-1 \times N-1$ unitary matrix. There can also be a phase variation along the diagonal, which we have omitted for clarity.

Mode-group-selective: The entrance modes can be divided into groups sharing the same propagation constant. For instance, LP_{01} is its own group, while the two LP_{11} modes correspond to another. A mode-group-selective lantern uses differing core geometries to route each mode group to its own subset of single-mode outputs. The transfer matrix has the block diagonal form

$$U = \begin{bmatrix} 1 & & 0 \\ & U_1 & \\ 0 & & \ddots \end{bmatrix} \tag{5.42}$$

where each U_i is an $N_i \times N_i$ unitary matrix corresponding to mode group i with N_i modes.

Mode-selective: A mode-selective or "fully-mode-selective" lantern uses differing core geometries *and* asymmetries in the waveguide structure to break all degeneracies between entrance modes, and then route each entrance mode to a separate output. The principal modes of a mode-selective lantern are the entrance modes. The transfer matrix of an ideal lantern of this type is identity:

$$U = \begin{bmatrix} 1 & 0 \\ 0 & \ddots \end{bmatrix}. \tag{5.43}$$

The 3-port lantern from [LFS14], which has one large core and two small cores, and an asymmetric cladding shape, is fully mode-selective. The same lantern with a circular cladding, like the one simulated in this work, is mode-group-selective, and also hybrid.

REFERENCES

[BGY15] T. A. Birks, I. Gris-Sánchez, S. Yerolatsitis, S. G. Leon-Saval, and R. R. Thomson. "The photonic lantern." *Adv. Opt. Photon.*, **7**(2):107–167, Jun 2015.

[Bow16] Brendan P. Bowler. "Imaging Extrasolar Giant Planets." *Publications of the Astronomical Society of the Pacific*, **128**(968):102001, aug 2016.

[CMH18] Mark K. Corrigan, Timothy J. Morris, Robert J. Harris, and Theodoros Anagnos. "Demonstration of a photonic lantern low order wavefront sensor using an adaptive optics testbed." In Laird M. Close, Laura Schreiber, and Dirk Schmidt, editors, *Adaptive Optics Systems VI*, volume 10703, pp. 1313 – 1320. International Society for Optics and Photonics, SPIE, 2018.

[DCL22] Mark Dong, Genevieve Clark, Andrew J. Leenheer, Matthew Zimmermann, Daniel Dominguez, Adrian J. Menssen, David Heim, Gerald Gilbert, Dirk Englund, and Matt Eichenfield. "High-speed programmable photonic circuits in a cryogenically compatible, visible–near-infrared 200 mm CMOS architecture." *Nature Photonics*, **16**(1):59–65, Jan 2022.

[DDM21] John J. Davenport, Momen Diab, Kalaga Madhav, and Martin M. Roth. "Optimal SMF packing in photonic lanterns: comparing theoretical topology to practical packing arrangements." *J. Opt. Soc. Am. B*, **38**(7):A7–A14, Jul 2021.

[DM19] Momen Diab and Stefano Minardi. "Modal analysis using photonic lanterns coupled to arrays of waveguides." *Opt. Lett.*, **44**(7):1718–1721, Apr 2019.

[EBB18] S. C. Ellis, S. Bauer, C. Bacigalupo, J. Bland -Hawthorn, J. J. Bryant, S. Case, R. Content, T. Fechner, D. Giannone, R. Haynes, E. Hernandez, A. J. Horton, U. Klauser, J. S. Lawrence, S. G. Leon-Saval, E. Lindley, H. G. Löhmannsröben, S. S. Min, N. Pai, M. Roth, K. Shortridge, L. Waller, Pascal Xavier, and Ross Zhelem. "PRAXIS: an OH suppression optimised near infrared spectrograph." In *Proceedings of the SPIE*, volume 10702 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 107020P, Jul 2018.

[Fra18] Richard A. Frazin. "Efficient, nonlinear phase estimation with the nonmodulated pyramid wavefront sensor." *J. Opt. Soc. Am. A*, **35**(4):594–607, Apr 2018.

[FRB12] Nicolas K. Fontaine, Roland Ryf, Joss Bland-Hawthorn, and Sergio G. Leon-Saval. "Geometric requirements for photonic lanterns in space division multiplexing." *Opt. Express*, **20**(24):27123–27132, Nov 2012.

[GBD19] Galicher, R., Baudoz, P., Delorme, J.-R., Mawet, D., Bottom, M., Wallace, J. K., Serabyn, E., and Shelton, C. "Minimization of non-common path aberrations at the Palomar telescope using a self-coherent camera." *A&A*, **631**:A143, 2019.

[Guy03] Guyon, O. "Phase-induced amplitude apodization of telescope pupils for extrasolar terrestrial planet imaging." *A & A*, **404**(1):379–387, 2003.

[GVD19] Pradip Gatkine, Sylvain Veilleux, and Mario Dagenais. "Astrophotonic Spectrographs." *Applied Sciences*, **9**(2), 2019.

[KFL24] Yoo Jung Kim, Michael P. Fitzgerald, Jonathan Lin, Steph Sallum, Yinzi Xin, Nemanja Jovanovic, and Sergio Leon-Saval. "Coherent Imaging with Photonic Lanterns." *ApJ*, **964**(2):113, April 2024.

[LBB12] Sergio G. Leon-Saval, Christopher H. Betters, and Joss Bland-Hawthorn. "The Photonic TIGER: a multicore fiber-fed spectrograph ." In Ramón Navarro, Colin R. Cunningham, and Eric Prieto, editors, *Modern Technologies in Space- and Ground-based Telescopes and Instrumentation II*, volume 8450, pp. 550 – 557. International Society for Optics and Photonics, SPIE, 2012.

[LF22] Jonathan Lin and Michael Fitzgerald. "Focal-plane wavefront sensing with photonic lanterns II: characterization and optimization." *in prep.*, 2022.

[LFS14] Sergio G. Leon-Saval, Nicolas K. Fontaine, Joel R. Salazar-Gil, Burcu Ercan, Roland Ryf, and Joss Bland-Hawthorn. "Mode-selective photonic lanterns for space-division multiplexing." *Opt. Express*, **22**(1):1036–1044, Jan 2014.

[LGV20] Julien Lozi, Olivier Guyon, Sébastien Vievard, Ananya Sahoo, Vincent Deo, Nemanja Jovanovic, Barnaby Norris, Marc-Antoine Martinod, Ben Mazin, Alex Walter, Neelay Fruitwala, Sarah Steiger, Kristina Davis, Peter Tuthill, Tomoyuki Kudo, Hajime Kawahara, Takayuki Kotani, Michael Ireland, Theodoros Anagnos, Chrstian Schwab, Nick Cvetojevic, Elsa Huby, Sylvestre Lacour, Kevin Barjot, Tyler D. Groff, Jeffrey Chilcote, Jeremy Kasdin, Frantz Martinache, Romain Laugier, Mamadou N'Diaye, Justin Knight, Jared Males, Steven Bos, Frans Snik, David Doelman, Kelsey Miller, Eduardo Bendek, Ruslan Belikov, Eugene Pluzhnik, Thayne Currie, Masayuki Kuzuhara, Taichi Uyama, Jun Nishikawa, Naoshi Murakami, Jun Hashimoto, Yosuke Minowa, Christophe Clergeon, Yoshito Ono, Naruhisa Takato, Motohide Tamura, Hideki Takami, and Masa Hayashi. "Status of the SCExAO instrument: recent technology upgrades and path to a system-level demonstrator for PSI." In Laura Schreiber, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems VII*, volume 11448, p. 114480N. International Society for Optics and Photonics, SPIE, 2020.

[Lin20] Jonathan W. Lin. "Lightbeam." Astrophysics Source Code Library, February 2020. ascl:2102.006.

[LJF21] Jonathan Lin, Nemanja Jovanovic, and Michael P. Fitzgerald. "Design considerations of photonic lanterns for diffraction-limited spectrometry." *J. Opt. Soc. Am. B*, **38**(7):A51–A63, Jul 2021.

[MFK19] Dimitri Mawet, Michael Fitzgerald, Quinn Konopacky, Charles Beichman, Nemanja Jovanovic, Richard Dekany, David Hover, Eric Chisholm, David Ciardi, Etienne Artigau, Ravinder Banyal, Thomas Beatty, Bjorn Benneke, Geoffrey A. Blake, Adam Burgasser, Gabriela Canalizo, Guo Chen, Tuan Do, Greg Doppmann, Rene Doyon, Courtney Dressing, Min Fang, Thomas Greene, Lynne Hillenbrand, Andrew Howard, Stephen Kane, Tiffany Kataria, Eliza Kempton, Heather Knutson, Takayuki Kotani, David Lafreniere, Chao Liu, Shogo Nishiyama, Gajendra Pand ey, Peter Plavchan, Lisa Prato, S. P. Rajaguru, Paul Robertson, Colette Salyk, Bunei Sato, Everett Schlawin, Sujan Sengupta, Thirupathi Sivarani, Warren Skidmore, Motohide Tamura, Hiroshi Terada, Gautam Vasisht, Ji Wang, and Hui Zhang. “High-resolution Infrared Spectrograph for Exoplanet Characterization with the Keck and Thirty Meter Telescopes.” *arXiv e-prints*, p. arXiv:1908.03623, Aug 2019.

[MKC13] Martinez, P., Kasper, M., Costille, A., Sauvage, J. F., Dohlen, K., Puget, P., and Beuzit, J. L. “Speckle temporal stability in XAO coronagraphic images - II. Refine model for quasi-static speckle temporal evolution for VLT/SPHERE.” *A&A*, **554**:A41, 2013.

[MLA12] Martinez, P., Loose, C., Aller Carpentier, E., and Kasper, M. “Speckle temporal stability in XAO coronagraphic images.” *A&A*, **541**:A136, 2012.

[MLD06] Christian Marois, David Lafreniere, Rene Doyon, Bruce Macintosh, and Daniel Nadeau. “Angular Differential Imaging: A Powerful High-Contrast Imaging Technique.” *The Astrophysical Journal*, **641**(1):556–564, apr 2006.

[MNT21] Marc-Antoine Martinod, Barnaby Norris, Peter Tuthill, Tiphaine Lagadec, Nemanja Jovanovic, Nick Cvetojevic, Simon Gross, Alexander Arriola, Thomas Gretzinger, Michael J. Withford, Olivier Guyon, Julien Lozi, Sébastien Vievard, Vincent Deo, Jon S. Lawrence, and Sergio Leon-Saval. “Scalable photonic-based nulling interferometry with the dispersed multi-baseline GLINT instrument.” *Nature Communications*, **12**(1):2465, Apr 2021.

[NBW22] Barnaby Norris, Christopher Betters, Jin Wei, Stephanos Yerolatsitis, Rodrigo Amezcua-Correa, and Sergio Leon-Saval. “Optimal broadband starlight injection into a single-mode fibre with integrated photonic wavefront sensing.” *Opt. Express*, **30**(19):34908–34917, Sep 2022.

[NMJ18] N´Diaye, M., Martinache, F., Jovanovic, N., Lozi, J., Guyon, O., Norris, B., Ceau, A., and Mary, D. “Calibration of the island effect: Experimental validation of closed-loop focal plane wavefront control on Subaru/SCExAO.” *A&A*, **610**:A18, 2018.

[NWB20] Barnaby R. M. Norris, Jin Wei, Christopher H. Betters, Alison Wong, and Sergio G. Leon-Saval. “An all-photonic focal-plane wavefront sensor.” *Nature Communications*, **11**(1):5335, Oct 2020.

[PHR18] Emiel H. Por, Sebastiaan Y. Haffert, Vikram M. Radhakrishnan, David S. Doelman, Maaike van Kooten, and Steven P. Bos. “High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator.” In Laird M. Close, Laura Schreiber, and Dirk Schmidt, editors, *Adaptive Optics Systems VI*, volume 10703, pp. 1112 – 1125. International Society for Optics and Photonics, SPIE, 2018.

[REJ19] Garreth Ruane, Daniel Echeverri, Nemanja Jovanovic, Dimitri Mawet, Eugene Serabyn, J. Kent Wallace, Jason Wang, and Natasha Batalha. “Vortex fiber nulling for exoplanet observations: conceptual design, theoretical performance, and initial scientific yield predictions.” In Stuart B. Shaklan, editor, *Techniques and Instrumentation for Detection of Exoplanets IX*, volume 11117, pp. 366 – 381. International Society for Optics and Photonics, SPIE, 2019.

[RFM21] Alexander T. Rodack, Richard A. Frazin, Jared R. Males, and Olivier Guyon. “Millisecond exoplanet imaging: I. method and simulation results.” *J. Opt. Soc. Am. A*, **38**(10):1541–1556, Oct 2021.

[RZB94] Michael Reck, Anton Zeilinger, Herbert J. Bernstein, and Philip Bertani. “Experimental realization of any discrete unitary operator.” *Phys. Rev. Lett.*, **73**:58–61, Jul 1994.

[SF02] William B. Sparks and Holland Cole Ford. “Imaging Spectroscopy for Extrasolar Planet Detection.” *The Astrophysical Journal*, **578**:543–564, 2002.

[SFR07] Jean-François Sauvage, Thierry Fusco, Gérard Rousset, and Cyril Petit. “Calibration and precompensation of noncommon path aberrations for extreme adaptive optics.” *J. Opt. Soc. Am. A*, **24**(8):2334–2346, Aug 2007.

[TEB13] Christopher Q. Trinh, Simon C. Ellis, Joss Bland-Hawthorn, Jon S. Lawrence, Anthony J. Horton, Sergio G. Leon-Saval, Keith Shortridge, Julia Bryant, Scott Case, Matthew Colless, Warrick Couch, Kenneth Freeman, Hans-Gerd Löhmannsröben, Luke Gers, Karl Glazebrook, Roger Haynes, Steve Lee, John O’Byrne, Stan Miziarski, Martin M. Roth, Brian Schmidt, Christopher G. Tinney, and Jessica Zheng. “GNOSIS: THE FIRST INSTRUMENT TO USE FIBER BRAGG GRATINGS FOR OH SUPPRESSION.” *The Astronomical Journal*, **145**(2):51, jan 2013.

[TO10] W. A. Traub and B. R. Oppenheimer. “Direct Imaging of Exoplanets.” In S. Seager, editor, *Exoplanets*, pp. 111–156. 2010.

[VAA18] Amado M. Velázquez-Benítez, J. Enrique Antonio-López, Juan C. Alvarado-Zacarías, Nicolas K. Fontaine, Roland Ryf, Haoshuo Chen, Juan Hernández-Cordero, Pierre Sillard, Chigo Okonkwo, Sergio G. Leon-Saval, and Rodrigo Amezcua-Correa. "Scaling photonic lanterns for space-division multiplexing." *Scientific Reports*, **8**(1):8897, Jun 2018.

[WYH22] Thomas A. Wright, Stephanos Yerolatsitis, Kerrianne Harrington, Robert J. Harris, and Tim A. Birks. "All-fibre wavefront sensor.", 2022.

[XJR22] Yinzi Xin, Nemanja Jovanovic, Garreth Ruane, Dimitri Mawet, Michael P. Fitzgerald, Daniel Echeverri, Jonathan Lin, Sergio Leon-Saval, Pradip Gatkine, Yoo Jung Kim, Barnaby Norris, and Steph Sallum. "Efficient Detection and Characterization of Exoplanets within the Diffraction Limit: Nulling with a Mode-selective Photonic Lantern." *The Astrophysical Journal*, **938**(2):140, oct 2022.

[ZDC22] Hailong Zhou, Jianji Dong, Junwei Cheng, Wenchan Dong, Chaoran Huang, Yichen Shen, Qiming Zhang, Min Gu, Chao Qian, Hongsheng Chen, Zhichao Ruan, and Xinliang Zhang. "Photonic matrix multiplication lights up photonic accelerator and beyond." *Light: Science & Applications*, **11**(1):30, Feb 2022.

CHAPTER 6

Nonlinear phase retrieval for few-moded and multiwavelength wavefront sensors

Adaptive optics systems both within astronomy and without rely on phase retrieval algorithms, which infer the shape of incoming wavefronts from wavefront sensor measurements. The standard phase retrieval method used in astronomical observatories is also the simplest, and assumes a linear relation between wavefront shape and sensor output — but this approximation becomes worse as wavefronts become more severely distorted. Concurrently, linear retrieval methods have a fundamental tradeoff between sensitivity and accuracy when wavefront distortion is high. This tradeoff can be problematic for highly sensitive and therefore highly nonlinear wavefront sensors such as the unmodulated pyramid, Zernike sensor, and photonic lantern, all of which are candidates to perform second-stage correction of non-common-path and petal aberrations on next generation of observatories. In this chapter, I consider how the accuracy of wavefront sensors in this context can be improved by extending phase retrieval algorithms beyond the typical linear approach. I also develop a numerical approach which quantifies the fundamental performance limit of a wavefront sensors, and consider the potential improvements offered by spectral dispersion.

6.1 Introduction

Linear wavefront sensing methods have long been the standard for astronomical adaptive optics (AO) due to their speed and simplicity. In exchange, such methods lose dynamic range as the intensity response of the wavefront (WFS) deviates from linearity, which can happen rapidly for sensors such as the unmodulated pyramid [Rag96, ER01], photonic lantern (PL; [LBB05, LFX22, LFX23]), and Zernike WFS [BW03]. As such, there has been a growing movement in the community to investigate and implement "nonlinear" sensing methods, an umbrella term applied to any phase retrieval method that uses more information than just the Jacobian and reference intensity image of the WFS. The need for such techniques has been underscored by the recent pushes for wavefront sensing in the final focal plane, in order to mitigate non-common-path aberrations (NCPAs; [MLA12]) and low-wind-effect aberrations [NMJ18], as well as the possibility of wavelength-dispersed WFSs, potentially enabled by microwave kinetic inductance detectors (MKIDs; [DLM03]) and/or astrophotonic devices [JGA23]. A final motivation for nonlinear sensing is that it will expedite alignment and NCPA compensation of AO systems, especially for WFSs with small dynamic range.

Existing non-linear phase retrieval methods include phase diversity [SMP12, KKD14, SGG22], Gerchberg-Saxton algorithms [GS72, AGL23], and most recently, neural networks [LH20, NWB20, PFQ22]. Some types of WFSs, such as the pyramid and Zernike sensors, also have sensor-specific strategies (e.g. [Fra18, HR18, CFJ20] for the pyramid and [Haf24] for the Zernike WFS). Broadly speaking, methods are often divided into "model-based" and "model-free" methods, though in the authors' opinion both classes of methods use models, with the main difference being whether the model is analytically developed or empirically calibrated; the classic linear sensing scheme is an example of the latter. The purpose of this article is to discuss alternative, empirically calibrated or "data-driven" retrieval methods, and to present numerical demonstrations. To offset the additional complexity of nonlinear sensing, we first apply our methods to few-moded ($\lesssim$ 10 mode) WFSs: this includes sensors which physically cannot sense large numbers of modes (e.g. the PL) as well as sensors whose outputs are projected to a few-mode subspace (e.g. a pyramid sensor restricted to the first few control modes) and any sensors composed of many smaller few-mode sensors (e.g. Shack-Hartmann). We leave the extension of the concepts presented here to higher-mode-count sensors for future work.

This article is divided into four subsequent parts. In §6.2, we establish the notation that will be used throughout the rest of the work. In §6.3 we discuss the different methods by which the nonlinear response of a WFS can be calibrated, including power-series expansion as well as interpolation using radial basis functions or polyharmonic splines. In §6.4 we discuss how models can be solved, as well as the possibility of “backwards” models where the calibration procedure takes WFS intensities as inputs and phase aberrations as outputs. In §6.5 we present numerical demonstrations of these methods, and additionally show how the technique of numerical continuation can be used to better understand WFS nonlinearity. We also briefly discuss polychromatic wavefront sensing, leaving a more detailed analysis for future work.

6.2 Framework

We represent a telescope + WFS system as a vector-valued nonlinear function which maps pupil-plane wavefront phases to intensity values. Extending the “response matrix” terminology from linear phase retrieval, we will refer to this nonlinear function as the “response map” of the WFS; the response map may also be loosely thought of as the WFS’s “transfer function” in the sense of control theory. Denote the monochromatic response map for an arbitrary WFS as $\mathcal{A}_\lambda$, which takes both a real scalar field ϕ and a wavelength λ, and returns a vector corresponding to the N intensity outputs of the WFS:

$$\mathcal{A}_\lambda \left[\phi(x, y, \lambda)\right] = \mathbf{p}_\lambda. \tag{6.1}$$

Here, $\phi(x, y, \lambda)$ is the phase at wavelength λ measured over the pupil plane. Note that λ appears twice: once representing the chromaticity of the phase, and once representing the chromaticity of the response map. We can approximate a polychromatic WFS as a series of such equations for each λ. In the case where the phase chromaticity dominates, the polychromatic WFS acts as phase-diverse monochromatic sensor that observes the same phase aberration "shape" at different amplitudes. However, if the "intrinsic" WFS chromaticity dominates, the WFS may see orthogonal sets of phase modes at different wavelengths. For brevity, we will drop λ going forwards, and only re-add it when discussing polychromatic sensors.

The goal of wavefront sensing is to robustly solve the above equation on a timescale faster than that of the wavefront aberrations. Solutions are often approximated by expanding the aberration ϕ in a finite-dimensional and incomplete modal basis

$$\phi(x,y) = \sum_j a_j m_j(x,y) \tag{6.2}$$

where the modes $m_j(x,y)$ are usually the "most important" modes ("control modes"), as determined by a singular value decomposition, and the a_j are the mode amplitudes. The output vector $\mathbf{p}$ may also be projected to the subspace corresponding to the image of the selected phase modes. We denote this modal-basis, vector-valued function as $\mathcal{T}$, whose scalar components are

$$\mathcal{T}_i\left[\mathbf{a}\right] = p_i. \tag{6.3}$$

Assuming an incomplete modal basis, $\mathcal{T}$ is an approximation of $\mathcal{A}$. For M modes and N WFS outputs, $\mathcal{T}$ maps $\mathbb{R}^M$ onto some subset of $\mathbb{R}^N$ corresponding to the range of intensities that the WFS can produce. From here, we would ideally obtain a relation of the form

$$a_j = \mathcal{T}_j^{-1}\left[\mathbf{p}\right]. \tag{6.4}$$

However, because there is no guarantee that $\mathcal{T}$ is bijective, there is usually no *single-valued* inverse function $\mathcal{T}^{-1}$ which satisfies the above over the entire range of $\mathcal{T}$; the notation $\mathcal{T}^{-1}\left[\mathbf{p}\right]$ instead denotes the *preimage* of the intensity $\mathbf{p}$. A typical function which shows such behavior is the complex exponential, e^z for $z \in \mathbb{C}$, which is single-valued but whose inverse is the multi-valued complex logarithm: $\log(z) = \{\ \ln|z| + i(\ \mathrm{Arg}\, z + 2\pi n) — n \in \mathbb{Z}\}$. Here, $\mathrm{Arg}(z)$ denotes the principal value complex argument of z, which is restricted to take values in the interval $(-\pi, \pi]$. In the context of interferometry, the failure for e^z to admit a single-valued inverse is referred to as "phase wrapping".

In §6.5.4 we will explicitly investigate how $\mathcal{T}$ fails to admit a single-valued inverse. To do so, it will be useful to embed $\mathcal{T}$ into the function $\mathcal{F} : \mathbb{R}^{M+N} \to \mathbb{R}^N$ with

$$\mathcal{F}[\mathbf{a}, \mathbf{p}] = \mathcal{T}[\mathbf{a}] - \mathbf{p}. \tag{6.5}$$

Pairs of corresponding intensities and phases fall along the surfaces with $\mathcal{F}[\mathbf{a}, \mathbf{p}] = \mathbf{0}$.

6.3 WFS models

In this section, we discuss some methods which can be used to empirically model the nonlinear WFS response represented by equation 6.1: the first step in a sensor-agnostic nonlinear phase retrieval scheme.

6.3.1 Power series representation

6.3.1.1 Linear approximation

To begin, we review the standard linear approach to wavefront sensing. A Taylor expansion to first order about a reference aberration $\mathbf{r}$ yields

$$p_i(\mathbf{r} + \mathbf{x}) - p_i(\mathbf{r}) = \partial_j p_i\big|_{\mathbf{r}} x_j \tag{6.6}$$

where $\partial_i \equiv \partial/\partial x_i$ and the vertical bar denotes the point of evaluation. The derivative term is typically measured using centered finite difference; thus, an M-moded sensor requires $2M$ probes to fully calibrate the linear model. Denoting the slope matrix (the Jacobian) as

$T_{ij} \equiv \partial_j p_i \big|_{\mathbf{r}}$, we can write the above relation as

$$\mathbf{p}(\mathbf{r}+\mathbf{x}) - \mathbf{p}(\mathbf{r}) \approx T\mathbf{x} \tag{6.7}$$

which is inverted as

$$\mathbf{x} = T^{+}\left[\mathbf{p}(\mathbf{r}+\mathbf{x}) - \mathbf{p}(\mathbf{r})\right] \tag{6.8}$$

where T^{+} is the inverse or pseudo-inverse of T. The map T is the linearization of $\mathcal{T}$.

6.3.1.2 Higher-order Taylor expansion

The natural extension of the linear technique is a power series expansion to higher order:

$$p_i(\mathbf{r}+\mathbf{x}) = p_i(\mathbf{r}) + \partial_j p_i \big|_{\mathbf{r}} x_j + \frac{1}{2}\partial_j \partial_k p_i \big|_{\mathbf{r}} x_j x_k + \frac{1}{6}\partial_j \partial_k \partial_l p_i \big|_{\mathbf{r}} x_j x_k x_l + ... \tag{6.9}$$

Repeated indices within the same term imply summation. For a few-moded WFS, we can truncate the power series beyond linear order. For instance, calibration up to second order requires an additional $2M^2$ probes: off-diagonal terms in the Hessian each involve four phase probes corresponding to the vertices of a small square centered on the reference phase, while the diagonal terms each require three measurements, one of which is supplied by the reference response of the WFS. In principle, truncation can occur at arbitrary order, though eventually calibration will take an unfeasible amount of time. Also, depending on the convergence of the series, higher order terms may be detrimental to dynamic range.

6.3.1.3 Higher-order expansion from complex response map

While the response map $\mathcal{T}$ is nonlinear, we note that the response map relating the complex-valued electric field in the pupil plane to that at the WFS output is linear. If this map can be measured (e.g. with phase diversity, off-axis holography, or a spatial light modulator), then it can be used to determine all higher-order tensors of the Taylor expansion. We direct the reader to [LFX22] for more details.

6.3.2 Radial basis function interpolation

Instead of expanding the nonlinear WFS response in terms of polynomials, another approach to build a model of the nonlinear response, often used in unstructured data analysis, is to expand in terms of radial basis functions (RBFs; [Har71]). An RBF is a function $\theta : \mathbb{R}^M \to \mathbb{R}$ whose value depends only on the distance to some reference point. For example, the Gaussian RBF is

$$\theta(\mathbf{x}, \mathbf{r}) = e^{-\epsilon^2 ||\mathbf{x}-\mathbf{r}||^2} \tag{6.10}$$

where ϵ is called the “shape parameter” of the RBF, $\mathbf{r}$ is the reference point, and $||\mathbf{x}||$ is the norm of $\mathbf{x}$. A nonlinear function $f : \mathbb{R}^M \to \mathbb{R}$ can be expanded as

$$f(\mathbf{x}) = \sum_i w_i \theta(\mathbf{x}, \mathbf{r}_i) \tag{6.11}$$

where the w_i are the weights and the $\mathbf{r}_i$ are the “control points” of the interpolation. Given a set of P control points, one can determine the weights by solving the $P \times P$ linear system

$$\begin{bmatrix} \theta(\mathbf{r}_1, \mathbf{r}_1) & \theta(\mathbf{r}_1, \mathbf{r}_2) & \dots \\ \theta(\mathbf{r}_2, \mathbf{r}_1) & \theta(\mathbf{r}_2, \mathbf{r}_2) & \dots \\ \vdots & \vdots & \ddots \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ \vdots \end{bmatrix} = \begin{bmatrix} f(\mathbf{r}_1) \\ f(\mathbf{r}_2) \\ \vdots \end{bmatrix}. \tag{6.12}$$

To extend the technique to vector-valued nonlinear functions, the simplest approach (and the approach adopted later in this work) is to apply the scalar interpolation technique to each output dimension independently; more advanced extensions are discussed in §6.3.3. Other RBFs besides the Gaussian are also used; notably, functions with compact support will yield sparse linear systems. A related technique is polyharmonic spline interpolation [HD72], which augments the RBFs with a polynomial term. For instance, we may add a linear term as

$$f(\mathbf{x}) = \sum_i w_i \theta(\mathbf{r}_i, \mathbf{x}) + \mathbf{v}^T \cdot \begin{bmatrix} 1 \\ \mathbf{x} \end{bmatrix}. \tag{6.13}$$

Here, $\mathbf{v}$ is a column vector of linear polynomial weights, with dimension $M + 1$. Like with pure RBF interpolation, the weights for the polyharmonic splines are determined by solving a linear system. The default interpolation scheme implemented by `scipy`'s `RBFinterpolator`, which we use for numerical tests in §6.5, is a polyharmonic spline interpolation using the RBF $\theta(r) = r^2 \log r$ and a linear polynomial term.

The simplest way to select the control points for an RBF model is to randomly sample from a subset of the phase aberration space — for example, a ball with maximum radius corresponding to the sensor's expected dynamic range. We would then experimentally measure the WFS response at these points and determine the interpolation weights as above. If first-order derivative information is also required, we can additionally develop an RBF model of the Jacobian. In §6.5, we present some initial numerical results on how many control points are required to obtain a good fit. The number of control points can be reduced through greedy algorithms, in which new sample points are repeatedly inserted where the model error is highest [RA10].

6.3.3 Other interpolation methods

There are many other options for data-driven interpolation models other than the RBF method mentioned above, which we do not consider but could be explored in future works. For instance, a common method used to interpolate univariate curves in computer graphics is basis spline (b-spline; [Uns99]) interpolation. B-splines have finite support and can be extended to higher-dimensional spaces by forming multi-dimensional basis functions from the tensor products of the univariate b-splines. Another technique often used in statistics and machine learning is Gaussian process modeling [SSK18], a non-parametric Bayesian regression technique which considers the space of functions as an "infinitely" multivariate normal distribution whose covariance is determined by a kernel function. This distribution is then conditioned on observed data, and the mean of the conditioned distribution is taken as the interpolant. Similar to RBF interpolation, the standard formulation of Gaussian process modeling treats univariate data, but the technique can be extended to the multivariate case by applying univariate method to each output independently. Alternatively, more advanced "multi-output" Gaussian processes, which involve prescriptions to model the output covariances, are possible [LCO18].

Notably, there is a deep connection between RBF interpolation, spline interpolation, and Gaussian process modeling, all of which are a type of non-parametric regression. For instance, spline interpolation can be obtained as a specific case of Gaussian process modeling through the choice of kernel [KW70]; similarly, RBF interpolation and Gaussian process modeling are equivalent if the RBF is taken as the kernel of a Gaussian process, though this may not always be warranted (for instance, not all RBFs are valid covariances). All of these techniques are a form of functional regression in Reproducing Kernel Hilbert Spaces [AL11]. We leave further investigations of alternate interpolation methods beyond RBF interpolation to future work.

6.3.4 Neural network models

Similar to a non-parametric interpolation method, neural networks provide a way to model an arbitrary nonlinear function. The calibration process is also similar: we must provide a training data set of input phase aberrations and output intensities. The weights of the neural network are then determined via gradient descent. The amount of time required for RBF interpolation models and neural network models seems comparable: [NWB20] implements a neural network reconstructor for a few-mode PL sensor which is trained using calibration datasets containing $10^3 - 10^4$ points, similar to what we use in our simulations in §6.5. However, we note that as of now RBF interpolation is simpler to implement. We direct the reader to [NWB20, LH20, PFQ22] for more details.

6.4 Inversion strategies

Given a model of the WFS response in the neighborhood of some reference phase $\mathbf{r}$, we seek a way to find the preimages of the observed intensity patterns. We discuss options in the next sections.

6.4.1 Standard methods

Briefly, we mention that there are a number of well-developed nonlinear equation solvers including conjugate gradient methods (e.g. Newton-CG), nonlinear least-squares methods such as Levenberg-Marquardt [Lev44], and more general trust region methods. For few modes, these methods may run quickly enough for real-time AO.

6.4.2 Solution to the Taylor model — successive approximations

In the case of Taylor expansion, there is another method based on successive approximations [Kan39]. We seek to solve the following equation for $\mathbf{x}$ given a vector ρ and tensors $(B, C, D, ...)$:

$$\rho_i(\mathbf{r}, \mathbf{x}) = B_{ij}x_j + C_{ijk}x_jx_k + D_{ijkl}x_jx_kx_l + ... \tag{6.14}$$

In the context of WFSs, ρ_i is the reference-subtracted intensity response. To first order, we have

$$\mathbf{x}^{(1)} = B^{+}\rho + o(\mathbf{x}^2). \tag{6.15}$$

The second order equation is approximated by a partial substitution:

$$\rho_i \approx \left[B_{ij} + C_{ijk}x_k^{(1)}\right] x_j = \left\{B_{ij} + C_{ijk}\left[B^{+}_{km}\rho_m + o(\mathbf{x}^2)\right]\right\} x_j \equiv B_{ij}^{(2)}\left[x_j + o(\mathbf{x}^3)\right] \tag{6.16}$$

and so the second-order approximant to $\mathbf{x}$ is

$$\mathbf{x}^{(2)} = B^{(2)+}\rho + o(\mathbf{x}^3), \tag{6.17}$$

and so on. Note that if $\mathbf{x}^{(1)}$ is 0, all $x^{(n)}$ are 0: we cannot sense modes in the nullspace of B.

6.4.3 Radial basis function models

There are several possibilities in solving interpolated models such as those constructed from RBFs. For one, the model of the WFS response map $\mathcal{T}$ can be combined with one of the nonlinear solving methods mentioned in §6.4.1, e.g. by minimizing the scalar quantity

$$||\mathcal{T}[\mathbf{a}] - \mathbf{p}|| \tag{6.18}$$

over aberration amplitude $\mathbf{a}$, given WFS image $\mathbf{p}$. Models of the Jacobian of $\mathcal{T}$ may also enable gradient descent methods as well as phase retrieval models based on local linearization of the WFS [KVL07], and can be used at the very least to speed up nonlinear solvers. However, the simplest scheme, akin to what is often done with neural network WFS models, is to construct the model "in reverse", swapping the inputs and outputs during the calibration process so that the RBF model returns phase aberrations as a "function" of WFS intensities. As mentioned in §6.2 this will not work over arbitrarily large regions of phase aberration space, so control points must be placed in smart manner. We discuss this more in §6.5.4, but leave the development of an optimal method for WFS control point placement to future work.

6.4.4 Numerical continuation

The class of numerical continuation methods [AG90] aim to solve a system of nonlinear equations by starting with a different system admitting known solutions, and then continuously transforming the starting system into the one in question. As the system is transformed, the solutions, which will also vary continuously, are tracked, so that when the transformation is complete we are left with the solutions to the target system. One example is homotopy continuation for polynomial root-finding: suppose we want the roots of a polynomial $P(\mathbf{x})$. We begin with a polynomial $P_0(\mathbf{x})$ whose roots are known; preferably, the number of roots in the starting polynomial should also match the number of expected roots in the target polynomial, otherwise we must augment our numerical continuation methods with *bifurcation theory* [Cra91], which is outside the scope of this paper. We then embed both systems in a homotopy class, for instance

$$H(\mathbf{x}, s) = (1 - s)P_0(\mathbf{x}) + sP(\mathbf{x}), \tag{6.19}$$

where $s \in [0, 1]$. Note that the choice of homotopy is not unique, and can strongly impact the quality and speed of the solving process. Finally, the solutions $\mathbf{x}_0$ for which $H(\mathbf{x}_0, s) = 0$ are tracked as s is increased to 1. Typically, this tracking is done with "predictor-corrector" methods such as pseudo-arclength continuation [DKI06], which is briefly introduced in Appedix 6.A.

Continuation methods for polynomials are well-developed with multiple existing software packages, making homotopy continuation a plausible method to solve the Taylor polynomial equation, at least in the case of few-mode sensors. Numerical continuation is also applicable for non-polynomial systems. Consider the embedded WFS response map in equation 6.5. Any curve in "intensity space", which we define as the image of the phase aberration space with respect to $\mathcal{T}$, implicitly defines a homotopy class of functions. By constructing a curve from some reference intensity $\mathbf{p}_0$ with known reference phase $\mathbf{r}$ to the observed intensity $\mathbf{p}$ and applying numerical continuation, we may determine the phase aberration solution $\mathbf{r}+\mathbf{x}$. If $\mathbf{p}_0$ is associated with multiple phase aberrations $\mathbf{r}_i$, then numerical continuation from each $\mathbf{r}_i$ will yield a different phase solution for $\mathbf{p}$.

6.5 Numerical results

6.5.1 Method

We simulate a 10-m telescope with an unobstructed circular aperture using the Python package `hcipy` [PHR18]. We fix the wavelength to 1.55 μm and set a focal ratio of 4. As an initial demonstration, we then couple the optical system to a quad-cell tip-tilt sensor composed of four square pixels each with side length 10 μm. The four outputs of the quad cell are combined into two values in the usual bi-cell configuration, i.e. by taking the difference in flux between the left and right halves of the sensor, and the difference between the top and bottom halves, and then normalizing both differences by the summed flux of all four pixels. We chose a quad-cell as a first test because it is simple and demonstrates known non-linearity, and because techniques applied to quad-cells may be extended to Shack-Hartmann sensors. Later, we simulate a few-mode sensor based on a 6-port PL, which was modeled using the package `cbeam` [Lin24]. The phase aberration space of this sensor was restricted to the first 5 non-piston Zernike modes (Noll indices 2-6, corresponding to tilt, defocus, and astigmatism). The lantern has an output single-mode core size of 2.2 μm, and the geometry of the input is similar to that of a step-index few-mode fiber with a core radius of 10 μm. The lantern tapers by a factor of 8 over a length of 4 cm.

We tested two classes of phase retrieval methods. The first class is based on solving the Taylor polynomial equation using successive approximations; for reference, calibration up to third order for the 5-mode PL sensor required 285 deformable mirror probes. We also tested methods based on RBF interpolation. The first variant combines a forward model of the WFS response constructed from RBFs (and a linear term) with a nonlinear least-squares solver; we call this method “RBF-LS” for “RBF least-squares”. In the second, we evaluate an inverse model of the WFS response, constructed from the same type of RBFs; we call this method “RBF-I” for “RBF inverse”. In the last method, we solve the RBF-LS model using numerical continuation.

Phase retrieval schemes were implemented in Python using `numpy` and `scipy`. In particular, for RBF interpolation and nonlinear solving, respectively, we used the functions `RBFinterpolator` and `least_squares` as implemented by `scipy v1.14.0`. For the quad-cell, both the RBF-LS and RBF-I methods used the same set of 100 control points in phase space, which were randomly sampled in a ball of radius 2 radians RMS. Samples were selected uniformly in radius, so that there are more points near the origin. For the lantern sensor, we used 2000 control points sampled in a ball of radius 1 radians RMS. We later found that it may be possible to reduce the number of control points significantly while maintaining reasonable accuracy; see §6.5.3. A simple code was written for numerical continuation, based off the method of pseudo-arclength continuation (Appendix 6.A).

6.5.2 Phase retrieval comparisons

In this section we present performance comparisons for the different phase retrieval methods mentioned in §6.3 and §6.4. We first injected varying amounts of a single aberration into the quad-cell system and attempted phase retrieval. Figure 6.5.1 (left) shows the residual of the phase retrieval estimate for each method, as a function of injected tilt, while Figure 6.5.1 (right) plots the retrieved tilt amplitude as a function of the injected tilt; the dashed grey line $y = x$ represents a perfect WFS. Comparing the Taylor expansion methods, we note that the quadratic method offers no benefit, and that the largest gains in accuracy come from odd orders; this is expected because the highest order term of an even Taylor expansion is symmetric about the expansion point, and is likely not as useful for sensing. The best performers were the numerical continuation and RBF-I methods. Defining the dynamic range as the region where the absolute value of the phase retrieval error is less than 0.1 rad, the numerical continuation method more than doubles the dynamic range of the linear method. However, numerical continuation is quite slow, which we discuss at the end of this section. Next, we applied the phase retrieval methods to the 5-mode PL sensor. Figure 6.5.2 (left) shows the error in the phase retrieval estimate for each method, specifically as

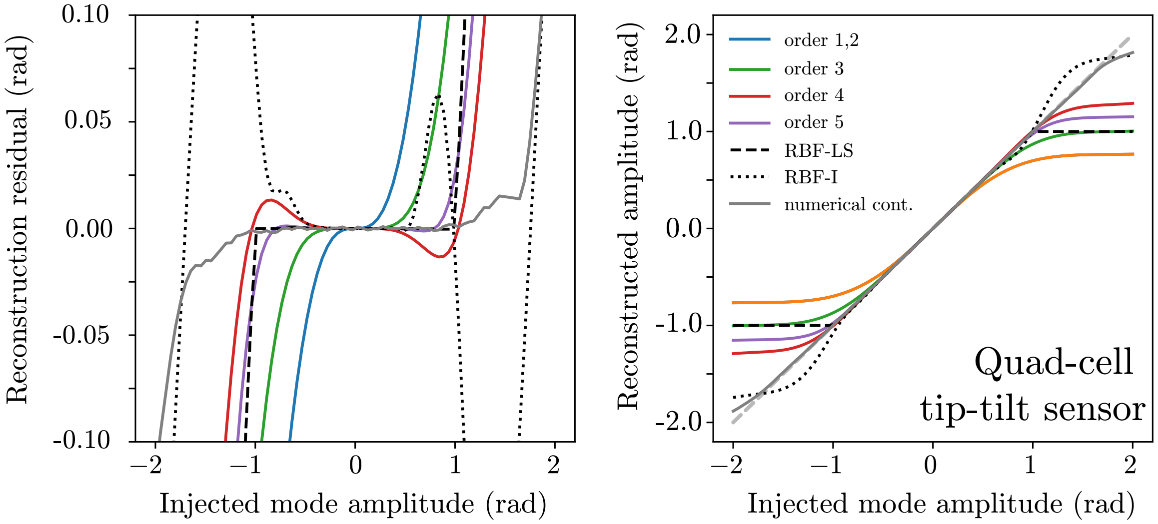

Figure 6.5.1: Left: residual of different phase retrieval schemes as a function of injected tilt for our fiducial quad-cell sensor. Order 1, 2, and so on correspond to Taylor expansion phase retrieval methods of the specified order. The “RBF-LS” curve corresponds to solving a forward model of the WFS built using RBF interpolation, while the “RBF-I” curve corresponds to direct evaluation of a backward model of the WFS, also built with RBFs. Right: reconstructed tilt amplitude as a function of injected amplitude. The dashed grey line $y = x$ represents the behavior of a perfect of sensor.

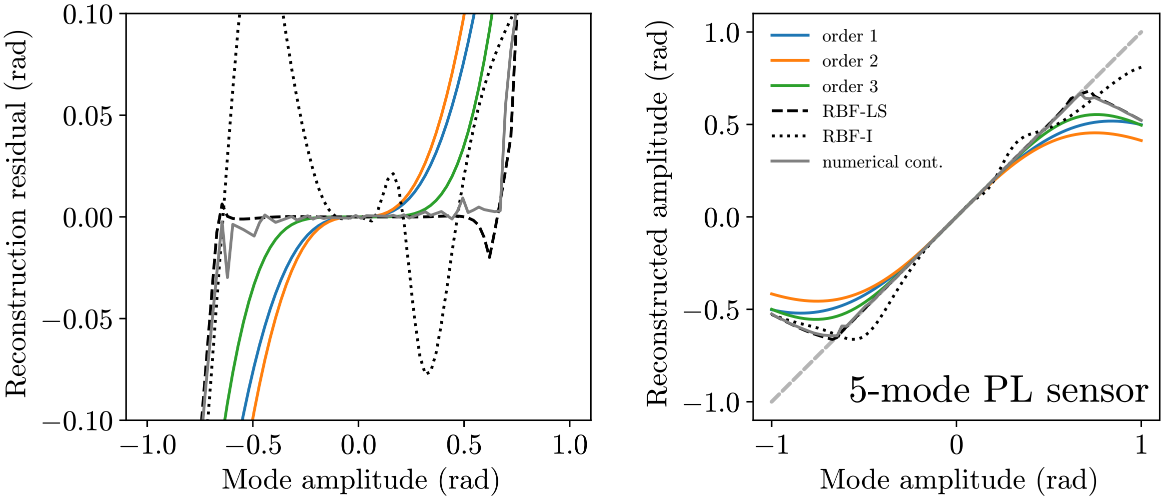

Figure 6.5.2: Left: residual of different phase retrieval schemes as a function of injected aberration mode amplitude, for our fiducual 5-mode PL sensor. In this particular plot, we chose to vary only the amplitude in Zernike mode 2 (tilt). Orders 1, 2, and 3 correspond to linear, quadratic, and cubic phase retrieval using Taylor approximation methods. The "RBF-LS" curve corresponds to solving a forward model of the WFS built using RBF interpolation, while the "RBF-I" curve corresponds to direct evaluation of a backward model of the WFS, also built with RBFs. Right: reconstructed tilt amplitude as a function of injected amplitude.

a function of injected tilt, while Figure 6.5.2 (right) plots the retrieved tilt amplitude as a function of injected tilt. Similar to the quad-cell, the largest gains in dynamic range from the Taylor methods come from the cubic model. The best performers were the RBF-LS and the numerical continuation methods. Interestingly, the RBF-I method performed worse than even the linear reconstructor for small amounts of aberration, but outperformed it for larger amounts. The increased error at small aberrations is likely due to our calibration data including pairs of degenerate or nearly degenerate phase aberrations, which make the model inconsistent; we discuss this more in §6.5.4.

We next sampled 10,000 random phase aberrations with at most 1 radian of total RMS wavefront error (WFE), and compared retrieval errors at each sample point for the 5-mode PL sensor. Our results are shown in Figure 6.5.3, which plots a 2D histogram for each retrieval method with total input WFE on the horizontal axis and retrieval error on the vertical axis. The numerical continuation method is omitted because in this case it performs similarly to RBF-LS, but takes longer to run. The histograms show that the RBF-LS triples the dynamic range of the linear method, and that the RBF-I method is outperformed by all methods until around 0.5 radians of total RMS WFE. We also recorded each method's solving times for all 10,000 sampled aberrations. The linear, quadratic, and cubic methods took a total of 0.85, 1.85, and 3.12 seconds for all 10,000 solves. The RBF-LS and RBF-I methods took 39.49 and 0.59 seconds, respectively. We suppose that the RBF-I method only outperforms the linear method due to some additional overhead incurred by our Python code. Nevertheless, a better-calibrated RBF-I method with a smarter allocation of control points could enable extremely high-speed wavefront sensing at dynamic ranges much larger than the linear range of the sensor. Even the RBF-LS method could be used as-is for live correction of quasistatic NCPAs. For reference, our simple numerical continuation code takes ~ 0.1 s for a single solve. However, the continuation code was written in pure Python and is completely unoptimized.

6.5.3 Number of control points

We next consider how many control points are required to obtain a good fit of WFS response using RBFs; this number directly correlates to the amount of time needed for calibration and determines how far the technique can be scaled in terms of the number of sensed modes. We repeat the random sample process used to generate Figure 6.5.3 for the 5-mode PL sensor, this time attempting phase retrieval using models calibrated with different amounts of control points. Our results are shown in Figure 6.5.4. The phase retrieval error scales relatively weakly with the number of control points: even at only a 100 points, the RBF-LS method clearly outperforms the linear method, and beyond 1000 points there is almost no gain in accuracy. The RBF-I method shows an even weaker scaling. For the quad-cell, we find that this same threshold happens at around 100 control points. These two results suggest

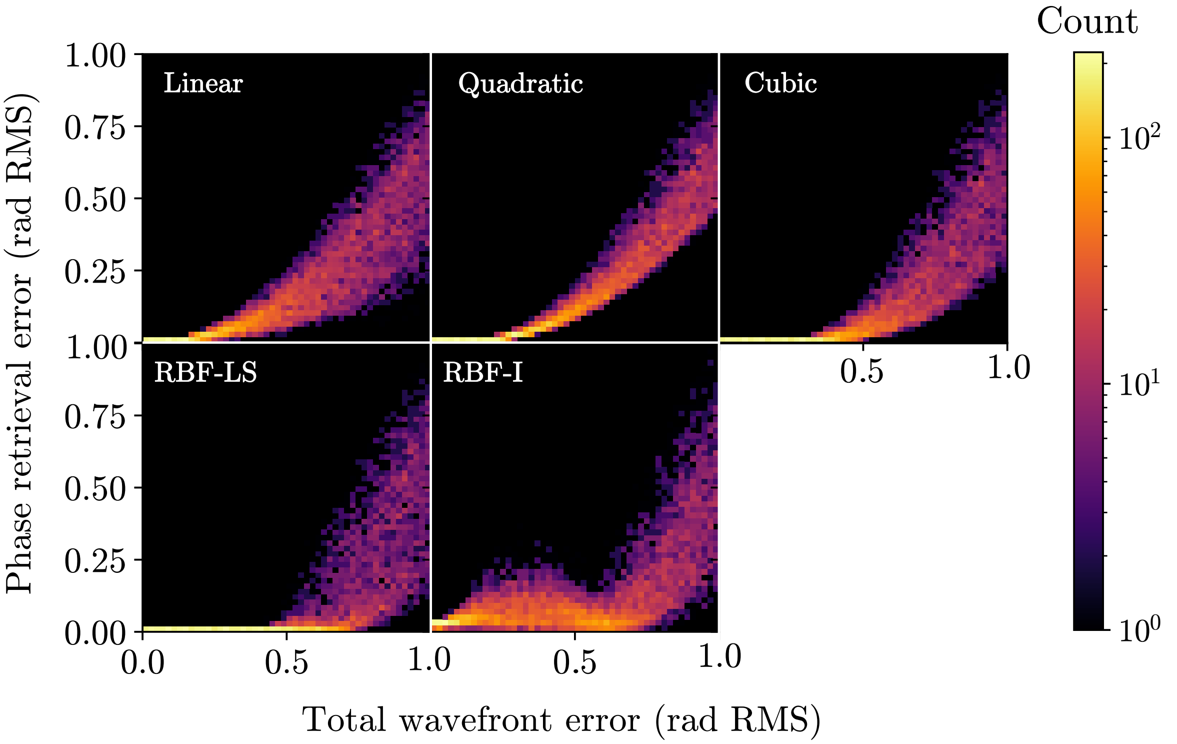

Figure 6.5.3: 2D histograms plotting total phase retrieval error on the vertical axis and input RMS WFE on the horizontal axis, for different retrieval methods combined with the 5-mode PL sensor. Each plot was produced using 10000 randomly sampled phase aberrations within a ball of radius 1 radian RMS.

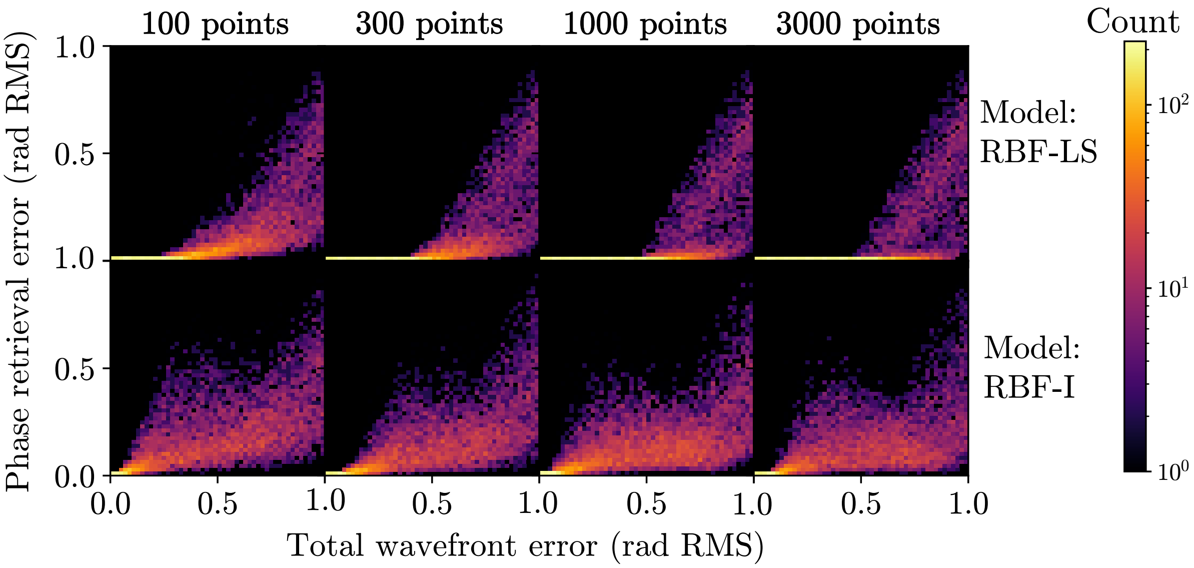

Figure 6.5.4: Comparison of RBF-based phase retrieval methods for different numbers of control points. We note the RBF-LS model begins seeing diminishing returns beyond 1000 control points, and still outperforms the linear model using only 100 control points. The RBF-I model appears almost insensitive to the number of control points.

that the number of control points required for an RBF model scales sub-exponentially with dimensionality/mode count.

6.5.4 Characterization through numerical continuation

In this section we demonstrate that numerical continuation can be used to better understand WFS nonlinearity. The idea is as follows: we embed the nonlinear response map $\mathcal{T}$ in the function $\mathcal{F}$, as in equation 6.5. This function is $\mathbb{R}^N$-valued, with a domain that is some subset of $\mathbb{R}^N \times \mathbb{R}^M$ corresponding to the direct product of the space of WFS intensities and the space of phase aberrations. Then, we set $\mathcal{F} = 0$, defining an implicit relation between the two spaces. Finally, we probe the intensity space along different paths, while computing the corresponding solution curves in phase space which together satisfy $\mathcal{F} = 0$. Note that we may also do the reverse, and probe phase space along different paths while computing the corresponding WFS intensity. This is often done in WFS characterization and does not require numerical continuation. However, if the goal is to search for features such as phase degeneracies, or to otherwise understand how the intensity-to-phase map fails to be a single-valued function, this simpler method is less informative: for instance, we do not know *a priori* how phase space should be traversed to find degeneracies. In comparison, degeneracies are naturally identified using numerical continuation, as we will later see.

Because the curves have the potential to be quite complicated, we first use our method to display the nonlinearity of the simple two-mode quad-cell sensor from §6.5.2. Since this sensor maps aberrations in a 2D phase space to a 2D intensity space, we can simultaneously visualize paths in intensity space and their corresponding solution curves in phase space with side-by-side plots. This method can also be extended to higher dimensions by projecting onto different 2D or 3D subspaces; see §6.B for an example with the 5-mode PL sensor.

To begin our study, we start with a reference pair of points in intensity space and phase space corresponding to the reference phase and intensity; in our particular example the reference in each is [0,0], because a flat phase yields a zero measurement for both bi-cell summations. We then attempt to move in intensity space along a ray emanating from the reference intensity, and simultaneously use pseudo-arclength continuation to compute the corresponding phase aberration solutions. By repeating this process for a number of rays, we obtain the plots in Figure 6.5.6. Each line segment in the left panel represents a path in intensity space which was constrained to move along a ray, and defines a corresponding curve of solutions in phase space through the implicit relation $\mathcal{F} = 0$. Black circles in either space mark the location of simple folds, where the intensity response "turns over" as a function of arclength (i.e. goes from increasing to decreasing or vice versa). Such locations mark where the Jacobian loses rank. In phase aberration space, each aberration near a simple fold has a "twin" aberration on the other side of the fold which is mapped to the same intensity. Note that we only plot the first simple fold found during each numerical continuation to reduce clutter. From this plot, we make the following conclusions:

1. Simple folds appear at around 2.5 radians RMS of tilt. Beyond this threshold, the mapping from intensity space to phase space is no longer bijective. The simple folds divide phase space into two branches. The region bounded by simple folds and containing the origin is the principal branch, which roughly corresponds to the dynamic range of the sensor under closed-loop control and linear phase retrieval.

2. Looking at the density of solution curves in phase space, we note that when increasing the total WFE, the quad-cell is less sensitive to diagonal tilts.

3. The reference phase $[0, 0]$ is unique: there is no other phase aberration that yields the same WFS response as the reference. Such aberrations, if they existed, would manifest as curve crossings in the full M-dimensional phase space.

As a more complex example, we repeat the above analysis for a two-mode tip-tilt sensor which uses a 3-port PL; the 3-dimensional output intensity vector is reference-subtracted and projected onto the two control modes with the two highest singular values so that our WFS maps $\mathbb{R}^2$ onto some subset of $\mathbb{R}^2$, like the quad-cell. In our case, this projection is also equivalent to a normalization of the WFS image. Figure 6.5.6 shows our results, this time plotting all simple folds. We make several conclusions:

1. There is a triangular region in phase space, bounded by simple folds, where the straight lines in intensity space are mapped to (mostly) straight lines. Outside this region we may find additional phase aberrations which are degenerate to the ones in this region.

2. Both plots have three-fold rotational symmetry, like the 3-port PL.

3. We can explicitly see the 6 phase aberration solutions that are degenerate to the reference phase $[0, 0]$. These points are represented by curve crossings in phase space.

4. The simple folds divide the space of phase solutions into 7 branches, with the principal branch corresponding to the central triangular region. Each branch maps to a different

but overlapping region of intensity space. As a corollary, an intensity may have a different number of associated phase aberrations, depending on its location in intensity space.

5. Regions devoid of curves in phase space indicate where our WFS is insensitive; this may change with the choice of control modes.

We envision that similar analyses may be done for other WFSs. Such a method provides a means of computing the bounded region within which it is possible to formulate a one-to-one mapping between intensity and phase. Additionally, for any point in intensity space, it only takes two numerical continuations to compute the set of aberrations which map to that point. Finally, the most powerful aspect of this technique is the characterization of the WFS response into branches. If a one-to-one map can be constructed for each branch of the WFS response (i.e. by constructing a model at the reference phase and every point degenerate to it), and there is a regularizing method of predicting which solution branch we are in at any given time (e.g. using auxiliary AO telemetry), then accurate phase retrieval outside the limits set by degeneracy is possible.

6.6 Discussion

6.6.1 Feasibility

The primary barrier to a real-time demonstration of the methods presented in this work is in calibration. For one, the calibration for some methods may take a prohibitively long time. Taylor expansion methods in particular show poor scaling with the number of sensed modes: an M-mode sensor calibrated to order p will require $\lesssim M^p$ probes on the deformable mirror. As such the realistic upper-bound for calibration order with ∼10-mode sensors is probably 3 or 5 (as mentioned earlier, it is not efficient to truncate the power series at even orders), and the calibration process will take tens of thousands of probes on the deformable mirror. The

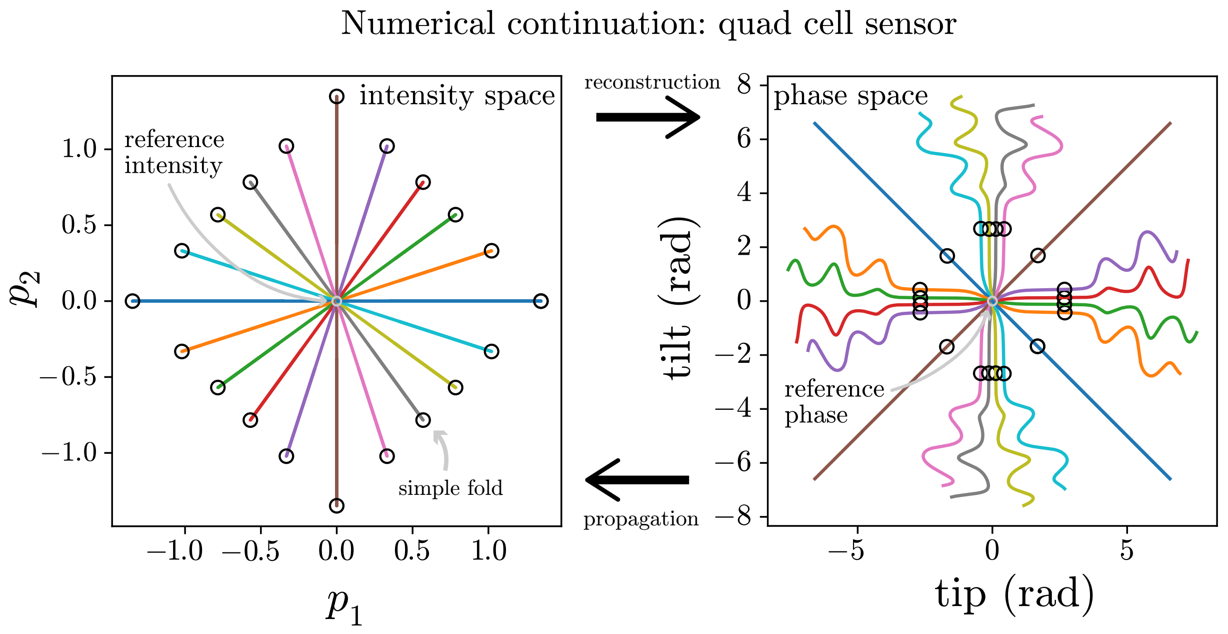

Figure 6.5.5: Numerical continuation plot for the quad-cell sensor from §6.5.2. Each colored line segment in intensity space (left) is mapped to a curve of the same color in phase space (right) by the nonlinear WFS response map. Black circles in either space mark the location of simple folds, where the intensity response "turns over" (i.e. goes from increasing to decreasing or vice versa). For clarity, only the first simple fold in each pair of curves is shown. Such locations indicate the presence of pairs of phase aberrations located on either side of the fold which are mapped to the same WFS intensity.

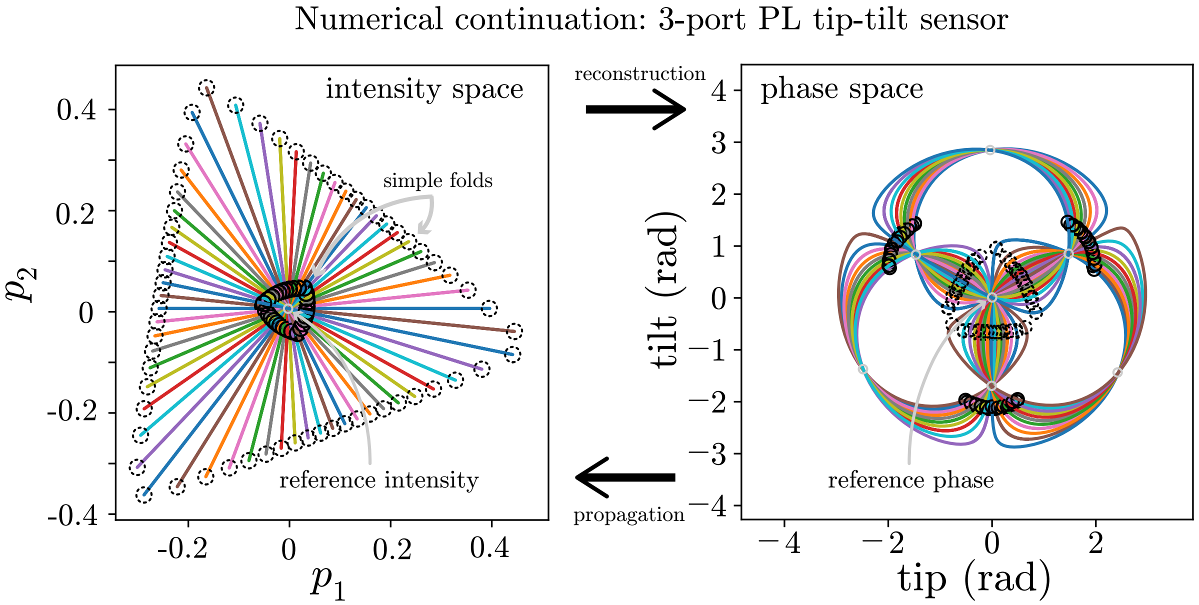

Figure 6.5.6: Similar to Figure 6.5.5 but for the 3-port lantern tip-tilt sensor. Dashed and solid black circles indicate corresponding sets of simple folds, as they appear in intensity and phase space. Line segments in intensity space map to closed curves, and the reference intensity (gray circle in left panel) is mapped to 7 distinct phases (gray circles in right panel), each of which are marked by curve crossings. The presence of phase degeneracies indicate that the function taking intensity to phase is multi-valued; branches of this function are partitioned by simple folds.

scaling between the number of required calibration probes and the number of sensed modes is less clear for phase retrieval strategies based on RBF interpolation, since we only have two datapoints for two specific sensors: around 100 control points for a 2-mode quad-cell sensor and 1000 for a 5-mode PL sensor. Future modelling of higher-mode sensors would help determine the exact scaling. For now, under the assumption of a power law scaling, the number of calibration points required for a 10 mode sensor is around 5,000 points, similar to the Taylor expansion models and some neural network schemes [NWB20]. In any case, the longer calibration times for nonlinear phase retrieval methods likely necessitates stabilization of the telescope and instrument, since any uncorrected and time-varying aberrations will add inconsistency to the nonlinear model; alternatively, we may account for noise using smoothing interpolators or Gaussian process models.

6.6.2 Composite sensing strategies

The phase retrieval methods presented in this paper may be combined with each other, as well existing methods, to produce more robust and performant sensing schemes. For instance, we might construct a dual-stage sensing scheme which uses RBF-I when the phase error is anticipated to be large and a linear or cubic Taylor method when the phase error is small, which would have an effectively nonlinear dynamic range while maintaining fast computation times and a high degree of accuracy. Such a composite method is similar to the combined neural network + linear methods being tested on the SCExAO testbed at Subaru telescope [AGL23].

For another example, in §6.5 we found that the performance of the RBF-I method may suffer if the calibration dataset happens to contain phase aberrations from different branches of the WFS response. Because it may be difficult to exclude problematic control points *a priori*, one workaround is to instead calibrate a backwards RBF model in a very small neighborhood of the reference phase, and use this model to calibrate a higher-order Taylor expansion which directly takes intensity to phase. Lastly, simple nonlinear phase retrieval schemes for quad-cells might be chained together to improve the performance of Shack-Hartmann sensors. Analogously, because the response matrices for pyramid WFSs typically have a banded diagonal structure, few-mode nonlinear sensing strategies might also be composited together for the pyramid as well.

6.6.3 Forwards or backwards?

The relative performance between forward methods like RBF-LS and backward methods like RBF-I will depend on the properties of the nonlinear response map $\mathcal{T}$. For instance, RBF-I outperforms the RBF-LS method for the quad-cell sensor, but is outperformed in the case of the PL sensor. This discrepancy is likely due to a difference in where degenerate aberrations occur in the phase space for each sensor. In the case of the quad-cell, we found via numerical continuation (Figure 6.5.5) that degeneracies begin to occur when the total RMS WFE exceeds around 2.5 radians RMS. The phase aberrations randomly sampled for calibration all fall within 2 radians RMS, so the RBF-I model is well-behaved. In contrast, the calibration probes for the 5-mode PL sensor extend out to 1 rad RMS total wavefront error, while numerical continuation (see §6.B and Figure 6.A.1) shows that degeneracies may occur within 1 rad RMS. If the calibration dataset includes aberrations from different branches of the WFS response, then the inverse model will be inconsistent. Even so, it may be beneficial to include degenerate aberrations during the calibration process, because in sampling a larger area of phase space we may gain dynamic range in exchange for accuracy. Still, the optimal strategy is to map each branch of the WFS response separately.

6.6.4 Prospects for dispersion

We can extend monochromatic phase retrieval to polychromatic sensors in the simplest way. Suppose we have wavefront aberrations that are achromatic in optical path difference (λ^{-1} chromaticity in phase), and a polychromatic sensor which observes at L distinct wavelengths

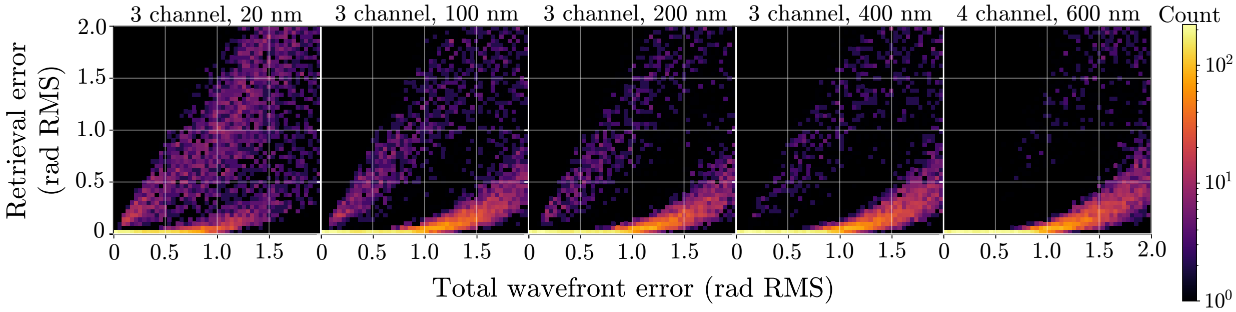

Figure 6.6.1: Phase retrieval error for various polychromatic PL sensors, each with different spectral bandwidths. The modal space for each sensor is restricted to the first 7 non-piston Zernike modes. The presence of data points that fall in the vicinity of the line $y = x$ for the 3-channel sensors indicate that they cannot sense a subset of the modal space; this subset becomes smaller when the spectral bandwidth is increased, and mostly disappears for the four-channel sensor.

over the sensing band. We treat this sensor as a function $\mathbb{R}^M \to \mathbb{R}^{NL}$, and attempt to apply the same techniques developed for monochromatic sensing to this higher-dimensional function. As an example, Figure 6.6.1 shows retrieval error plots for polychromatic sensors using different numbers of wavelength channels and spectral bandwidths; each configuration is composed of several monochromatic 5-mode PL sensors, and uses the RBF-LS reconstruction method. We restrict the modal phase space of these sensors to the first 7 non-piston Zernike modes. All 3-channel sensors observe at a central wavelength of 1.55 μm while the 4-channel sensor augments the 3-channel, 400 nm bandwidth sensor with an additional channel at 1.15 μm. We note that at the smallest bandwidths, the residual does not go to zero faster than linear, even for small amounts of WFE, signalling that the sensor is unable to distinguish between certain phase modes. However, as the bandwidth and number of channels increases, so too does correction quality, with the 4-channel configuration showing good performance out to 1 rad RMS. Thus, we have at least shown in this rudimentary test that chromaticity can improve both dynamic range and spatial resolution. We leave more detailed analysis, experimental verification, and the sensing of chromatic aberrations to a followup work.

Lastly, we mention one caveat to polychromatic sensing. In the simple construction above we will run into calibration issues when the guide star spectrum changes, since the response in each spectral channel depends on the intensity at each wavelength. If the spectrum changes only slightly, we can correct for the spectral offset in a similar way to how sensitivity loss of the pyramid WFS in the presence of phase offsets is corrected ("optical gain compensation", e.g. [CFJ20]). However, there is no guarantee that the target spectrum will change only slightly.

6.7 Conclusion

In §6.3-§6.4, we consider two methods for empirically modelling the nonlinear response of few-mode wavefront senors: Taylor expansion and radial basis function interpolation. In conjunction with these models, we also discuss how such models can be solved. In particular we show that higher-order Taylor expansions can be solved via successive approximations, and that forward models of the WFS built with radial basis functions can be solved with standard least-squares techniques as well as numerical continuation. Furthermore, we show that backwards models, which map intensity directly to phase, can be built from radial basis functions as well. §6.5 presents numerical demonstrations of the aforementioned modelling techniques and solving strategies for both a tip-tilt quad-cell sensor as well as a PL WFS which is sensitive to the first 5 non-piston Zernike modes, with additional discussion on how the method of numerical continuation may be used to better understand wavefront sensor nonlinearity and phase aberration degeneracy. Finally, in §6.6 we discuss future outlooks for both nonlinear and dispersed wavefront sensing.

Appendix - Chapter 6

6.A Pseudo-arclength continuation

Define a potentially vector-valued nonlinear function $\mathcal{F}$ which takes in M algebraic variables $\mathbf{a}$ and N parameters $\mathbf{p}$; technically, there is no difference between the two and that some of the variables in $\mathbf{a}$ may at any time be taken as parameters, or vice versa. The goal is find the set of solutions $\mathbf{a}$ which satisfy

$$\mathcal{F}\left[\mathbf{a}, \mathbf{p}\right] = 0 \tag{6.20}$$

as $\mathbf{p}$ moves along some predetermined curve in the parameter space, beginning at some known start point $\mathbf{a}_0, \mathbf{p}_0$ for which the above is satisfied. For simplicity, we will consider only the one-parameter case, since any curve can be parameterized in terms of arclength. Therefore, the reduced problem at hand is to numerically continue the solutions $\mathbf{a}$ which solve

$$\mathcal{F}\left[\mathbf{a}, p\right] = 0 \tag{6.21}$$

as p varies through some interval $[p_0, p_f]$. We denote the (Fréchet) derivatives as

$$\begin{aligned} \mathcal{F}_{\mathbf{a}} &= \frac{\partial \mathcal{F}}{\partial \mathbf{a}} \\ \mathcal{F}_p &= \frac{\partial \mathcal{F}}{\partial p} \end{aligned} \tag{6.22}$$

For a vector-valued $\mathcal{F}$, the first line is the Jacobian with respect to the variables $\mathbf{a}$. In the simplest numerical continuation schemes, we move p through the interval and repeatedly use a nonlinear solver (e.g. Newton's method) to obtain the solutions $\mathbf{a}$ as a function of p. Each value of $\mathbf{a}$ is taken as an initial guess for the next value of p. However, this process will fail if $\mathcal{F}_{\mathbf{a}}$ is singular anywhere along the the path, because the nonlinear solver will not converge.

Therefore, we will parameterize all arguments of $\mathcal{F}$ in terms of a path length s:

$$\mathcal{F}\left[\mathbf{a}(s), p(s)\right] = 0. \tag{6.23}$$

By differentiating the above with respect to s we obtain

$$\mathcal{F}_{\mathbf{a}}\dot{\mathbf{a}} + \mathcal{F}_p\dot{p} = 0 \tag{6.24}$$

where dots denote derivative with respect to s. We can additionally apply a normalization condition

$$||\dot{\mathbf{a}}||^2 + ||\dot{p}||^2 = 1. \tag{6.25}$$

The combination of the above two formulas determines the tangent vector as

$$\begin{aligned} \dot{\mathbf{a}} &= -\mathcal{F}_{\mathbf{a}}^{-1}\mathcal{F}_p\dot{p} \\ \dot{p} &= \pm\frac{1}{\sqrt{1 + ||\mathcal{F}_{\mathbf{a}}^{-1}\mathcal{F}_p||^2}} \end{aligned} \tag{6.26}$$

which we can evaluate at some starting $\mathbf{a}_0$ and p_0 to obtain the starting tangent vector $[\dot{\mathbf{a}}_0^T, \dot{p}_0]^T$. Finally, we augment our nonlinear system as

$$\begin{bmatrix} \mathcal{F}\left[\mathbf{a}(s), p(s)\right] \\ \dot{\mathbf{a}}_0 \cdot (\mathbf{a} - \mathbf{a}_0) + \dot{p}_0(p - p_0) \end{bmatrix} = \begin{bmatrix} 0 \\ s - s_0 \end{bmatrix} \tag{6.27}$$

where the additional scalar equation comes from the normalization condition. The system is solved with initial guess

$$\begin{aligned} \tilde{\mathbf{a}} &= \mathbf{a}_0 + (s - s_0)\dot{\mathbf{a}}_0 \\ \tilde{p} &= p_0 + (s - s_0)\dot{p}_0 \end{aligned} \tag{6.28}$$

to obtain a solution **a** located a distance $ds = s - s_0$ down the path. This process is then iterated to generate the entire solution curve.

Pseudo-arclength continuation allows us to numerically continue through points where $\mathcal{F}_\mathbf{a}$ is singular so long as we also have $\dot{p} = 0$; these points are called simple folds or limit points and correspond to when the path in parameter space folds back on itself. If we do not have $\dot{p} = 0$, the solutions may bifurcate along multiple paths, and leads into *bifurcation theory*.

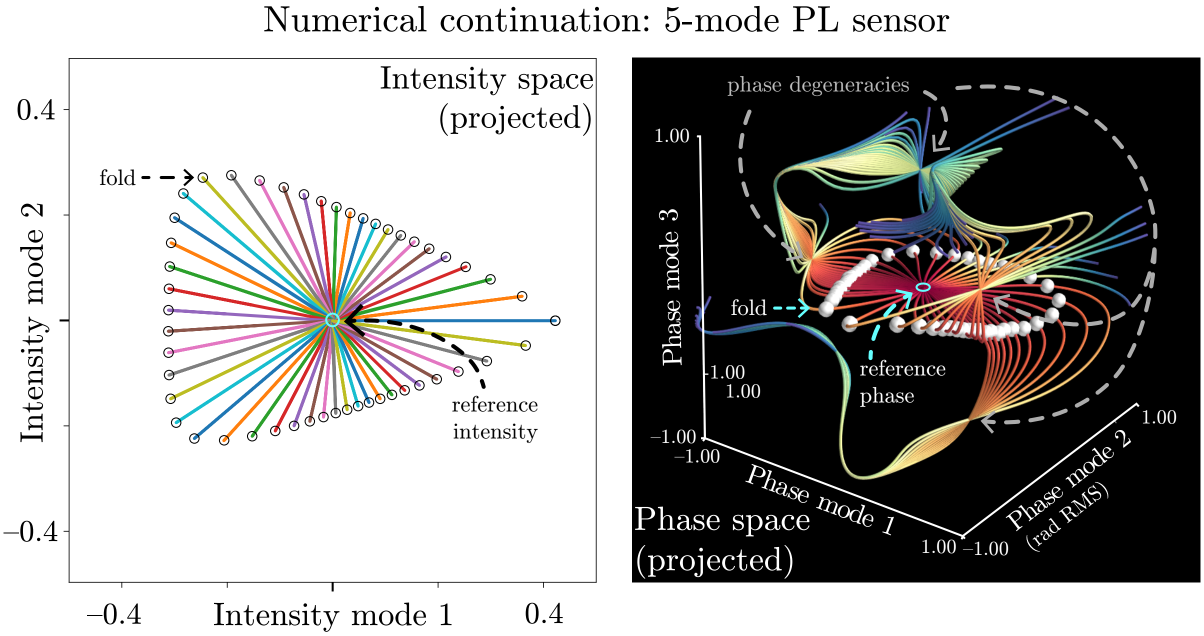

Figure 6.A.1: Numerical continuation plot for the 5-mode PL sensor. Intensity space and phase space are represented in the bases corresponding to the left and right singular vectors of the WFS Jacobian; the left and right panels show projections of intensity and phase space. Intensity space is probed along the plane corresponding to the first two modes. In the vicinity of the reference phase, corresponding phase space solution curves fall onto the plane corresponding to the first two phase space modes. As the RMS WFE increases, curves deviate from the plane, indicating nonlinearity. Phase curves are colored by arclength. We note that the reference phase is degenerate to at least 4 other phase aberrations. Simple folds are marked on the left with black circles and on the right by white dots.

6.B Nonlinear characterization by numerical continuation in >3 dimensions

Numerical continuation plots for 2- and 3-mode sensors are straightforward because we can view the entirety of phase space without projection. However, beyond 3 dimensions, locating branches, folds, and degenerate phase aberrations becomes challenging. One way forward is to work in both the left and right singular vectors (the control modes) of the Jacobian. In this basis, the WFS applies a scaling transformation in the neighborhood of the reference phase and intensity. Thus, if we probe intensity space along a plane, solution curves in phase space will also be embedded in a plane, at least near the reference; in other words, with this this approach numerically-continued curves will locally behave like those for 2-mode sensors. As the RMS WFE increases, the phase space surface containing the solutions will bend out of the plane, indicating nonlinearity. As an example, we apply this method to the 5-mode PL sensor from §6.5, probing intensity space in the plane of the first two intensity modes. These paths are shown in Figure 6.A.1 (left); note that some paths were truncated to save on computation time. As before, simple folds in intensity space are marked by open black circles. The WFS relates each path in this 2D intensity space to a solution curve in the 5D phase space. We plot these curves, as projected onto the subspace of the first 3 phase modes, in Figure 6.A.1 (right); curves are colored by arclength and simple folds are marked by white dots. Near the origin (corresponding to the reference phase), all curves lie on the plane corresponding to the first two phase modes, because in our chosen basis the WFS acts locally like a scaling transformation. Moving away from the origin, the surface containing the phase solutions bends away from this plane. These curves intersect again at four other locations in phase space, indicating aberrations that yield the same intensity response as the reference. Finally, we note that simple folds begin to appear around 0.5 rad RMS from the reference phase, consistent with the maximum dynamic range of this 5-mode sensor, as shown in Figure 6.5.4 (RBF-LS method). Simple folds divide the 2D solution manifold into at least 5 distinct branches, though we might find more if we extend the numerical continuation.

REFERENCES

[AG90] Eugene L. Allgower and Kurt Georg. *The Basic Principles of Continuation Methods*, pp. 7–16. Springer Berlin Heidelberg, Berlin, Heidelberg, 1990.

[AGL23] K. Ahn, O. Guyon, J. Lozi, S. Vievard, V. Deo, M. Lallement, and J. C. Bragg. "A non-linear curvature wavefront sensor for the Subaru telescope's AO3k system." In Garreth J. Ruane, editor, *Techniques and Instrumentation for Detection of Exoplanets XI*, volume 12680, p. 126800B. International Society for Optics and Photonics, SPIE, 2023.

[AL11] Ken Anjyo and J.P. Lewis. "RBF Interpolation and Gaussian Process Regression Through an RKHS Formulation." *Journal of Math-for-Industry*, **3**:63–71, 04 2011.

[BW03] Eric E Bloemhof and J. Kent Wallace. "Phase contrast techniques for wavefront sensing and calibration in adaptive optics." In Robert K. Tyson and Michael Lloyd-Hart, editors, *Astronomical Adaptive Optics Systems and Applications*, volume 5169, pp. 309 – 320. International Society for Optics and Photonics, SPIE, 2003.

[CFJ20] Chambouleyron, V., Fauvarque, O., Janin-Potiron, P., Correia, C., Sauvage, J-F., Schwartz, N., Neichel, B., and Fusco, T. "Pyramid wavefront sensor optical gains compensation using a convolutional model." *A&A*, **644**:A6, 2020.

[Cra91] John David Crawford. "Introduction to bifurcation theory." *Rev. Mod. Phys.*, **63**:991–1037, Oct 1991.

[DKI06] K. I. Dickson, C. T. Kelley, I. C. F. Ipsen, and I. G. Kevrekidis. "Condition Estimates for Pseudo-Arclength Continuation.", 2006.

[DLM03] Peter K. Day, Henry G. LeDuc, Benjamin A. Mazin, Anastasios Vayonakis, and Jonas Zmuidzinas. "A broadband superconducting detector suitable for use in large arrays." *Nature*, **425**(6960):817–821, Oct 2003.

[ER01] Esposito, S. and Riccardi, A. "Pyramid Wavefront Sensor behavior in partial correction Adaptive Optic systems." *A&A*, **369**(2):L9–L12, 2001.

[Fra18] Richard A. Frazin. "Efficient, nonlinear phase estimation with the nonmodulated pyramid wavefront sensor." *J. Opt. Soc. Am. A*, **35**(4):594–607, Apr 2018.

[GS72] R. W. GERCHBERG and W. O. SAXTON. "PRACTICAL ALGORITHM FOR DETERMINATION OF PHASE FROM IMAGE AND DIFFRACTION PLANE PICTURES." *OPTIK*, **35**(2):237–&, 1972.

[Haf24] Haffert, S. Y. "Into nonlinearity and beyond for Zernike-like wavefront sensors." *A&A*, **683**:A113, 2024.

[Har71] Rolland L. Hardy. “Multiquadric equations of topography and other irregular surfaces.” *Journal of Geophysical Research (1896-1977)*, **76**(8):1905–1915, 1971.

[HD72] ROBERT L. HARDER and ROBERT N. DESMARAIS. “Interpolation using surface splines.” *Journal of Aircraft*, **9**(2):189–191, 1972.

[HR18] Victoria Hutterer and Ronny Ramlau. “Nonlinear wavefront reconstruction methods for pyramid sensors using Landweber and Landweber–Kaczmarz iterations.” *Appl. Opt.*, **57**(30):8790–8804, Oct 2018.

[JGA23] Nemanja Jovanovic, Pradip Gatkine, Narsireddy Anugu, Rodrigo Amezcua-Correa, Ritoban Basu Thakur, Charles Beichman, Chad Bender, Jean-Philippe Berger, Azzurra Bigioli, Joss Bland-Hawthorn, Guillaume Bourdarot, Charles M. Bradford, Ronald Broeke, Julia Bryant, Kevin Bundy, Ross Cheriton, Nick Cvetojevic, Momen Diab, Scott A. Diddams, Aline N. Dinkelaker, Jeroen Duis, Stephen Eikenberry, Simon Ellis, Akira Endo, Donald F. Figer, Michael Fitzgerald, Itandehui Gris-Sanchez, Simon Gross, Ludovic Grossard, Olivier Guyon, Sebastiaan Y. Haffert, Samuel Halverson, Robert J. Harris, Jinping He, Tobias Herr, Philipp Hottinger, Elsa Huby, Michael Ireland, Rebecca Jenson-Clem, Jeffrey Jewell, Laurent Jocou, Stefan Kraus, Lucas Labadie, Sylvestre Lacour, Romain Laugier, Katarzyna Ławniczuk, Jonathan Lin, Stephanie Leifer, Sergio Leon-Saval, Guillermo Martin, Frantz Martinache, Marc-Antoine Martinod, Ben Mazin, Stefano Minardi, John D Monnier, Reinan Moreira, Denis Mourard, Abani Shankar Shankar Nayak, Barnaby Norris, Ewelina Obrzud, Karine Perraut, François Reynaud, Steph Sallum, David Schiminovich, Christian Schwab, Eugene Serbayn, Sherif Soliman, Andreas Stoll, Liang Tang, Peter Tuthill, Kerry Vahala, Gautam Vasisht, Sylvain Veilleux, Alexander B. Walter, Edward J Wollack, Yinzi Xin, Zongyin Yang, Stephanos Yerolatsitis, Yang Zhang, and Chang-Ling Zou. “2023 Astrophotonics Roadmap: pathways to realizing multi-functional integrated astrophotonic instruments.” *Journal of Physics: Photonics*, 2023.

[Kan39] L. Kantorovitch. “The method of successive approximation for functional equations.” *Acta Mathematica*, **71**(1):63–97, Dec 1939.

[KKD14] Visa Korkiakoski, Christoph U. Keller, Niek Doelman, Matthew Kenworthy, Gilles Otten, and Michel Verhaegen. “Fast & Furious focal-plane wavefront sensing.” *Appl. Opt.*, **53**(20):4565–4579, Jul 2014.

[KVL07] Visa Korkiakoski, Christophe Vérinaud, Miska Le Louarn, and Rodolphe Conan. “Comparison between a model-based and a conventional pyramid sensor reconstructor.” *Appl. Opt.*, **46**(24):6176–6184, Aug 2007.

[KW70] George S. Kimeldorf and Grace Wahba. “A Correspondence Between Bayesian Estimation on Stochastic Processes and Smoothing by Splines.” *The Annals of Mathematical Statistics*, **41**(2):495 – 502, 1970.

[LBB05] S. G. Leon-Saval, T. A. Birks, J. Bland-Hawthorn, and M. Englund. "Multimode fiber devices with single-mode performance." *Opt. Lett.*, **30**(19):2545–2547, Oct 2005.

[LCO18] Haitao Liu, Jianfei Cai, and Yew-Soon Ong. "Remarks on multi-output Gaussian process regression." *Knowledge-Based Systems*, **144**:102–121, 2018.

[Lev44] K. Levenberg. "A method for the solution of certain problems in least squares." *Quaterly Journal on Applied Mathematics*, (2):164–168, 1944.

[LFX22] Jonathan Lin, Michael P. Fitzgerald, Yinzi Xin, Olivier Guyon, Sergio Leon-Saval, Barnaby Norris, and Nemanja Jovanovic. "Focal-plane wavefront sensing with photonic lanterns: theoretical framework." *J. Opt. Soc. Am. B*, **39**(10):2643–2656, Oct 2022.

[LFX23] Jonathan W. Lin, Michael P. Fitzgerald, Yinzi Xin, Yoo Jung Kim, Olivier Guyon, Barnaby Norris, Christopher Betters, Sergio Leon-Saval, Kyohoon Ahn, Vincent Deo, Julien Lozi, Sébastien Vievard, Daniel Levinstein, Steph Sallum, and Nemanja Jovanovic. "Real-time Experimental Demonstrations of a Photonic Lantern Wave-front Sensor." *The Astrophysical Journal Letters*, **959**(2):L34, dec 2023.

[LH20] R. Landman and S. Y. Haffert. "Nonlinear wavefront reconstruction with convolutional neural networks for Fourier-based wavefront sensors." *Opt. Express*, **28**(11):16644–16657, May 2020.

[Lin24] Jonathan Lin. "cbeam." Astrophysics Source Code Library, record ascl:2404.001, April 2024.

[MLA12] Martinez, P., Loose, C., Aller Carpentier, E., and Kasper, M. "Speckle temporal stability in XAO coronagraphic images." *A&A*, **541**:A136, 2012.

[NMJ18] N´Diaye, M., Martinache, F., Jovanovic, N., Lozi, J., Guyon, O., Norris, B., Ceau, A., and Mary, D. "Calibration of the island effect: Experimental validation of closed-loop focal plane wavefront control on Subaru/SCExAO." *A&A*, **610**:A18, 2018.

[NWB20] Barnaby R. M. Norris, Jin Wei, Christopher H. Betters, Alison Wong, and Sergio G. Leon-Saval. "An all-photonic focal-plane wavefront sensor." *Nature Communications*, **11**(1):5335, Oct 2020.

[PFQ22] B. Pou, F. Ferreira, E. Quinones, D. Gratadour, and M. Martin. "Adaptive optics control with multi-agent model-free reinforcement learning." *Opt. Express*, **30**(2):2991–3015, Jan 2022.

[PHR18] E. H. Por, S. Y. Haffert, V. M. Radhakrishnan, D. S. Doelman, M. Van Kooten, and S. P. Bos. “High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator.” In *Adaptive Optics Systems VI*, volume 10703 of *Proc. SPIE*, 2018.

[RA10] T.C.S. Rendall and C.B. Allen. “Reduced surface point selection options for efficient mesh deformation using radial basis functions.” *Journal of Computational Physics*, **229**(8):2810–2820, 2010.

[Rag96] Roberto Ragazzoni. “Pupil plane wavefront sensing with an oscillating prism.” *Journal of Modern Optics*, **43**(2):289–293, 1996.

[SGG22] Skaf, Nour, Guyon, Olivier, Gendron, Éric, Ahn, Kyohoon, Bertrou-Cantou, Arielle, Boccaletti, Anthony, Cranney, Jesse, Currie, Thayne, Deo, Vincent, Edwards, Billy, Ferreira, Florian, Gratadour, Damien, Lozi, Julien, Norris, Barnaby, Sevin, Arnaud, Vidal, Fabrice, and Vievard, Sébastien. “On-sky validation of image-based adaptive optics wavefront sensor referencing.” *A&A*, **659**:A170, 2022.

[SMP12] J.-F. Sauvage, L. Mugnier, B. Paul, and R. Villecroze. “Coronagraphic phase diversity: a simple focal plane sensor for high-contrast imaging.” *Opt. Lett.*, **37**(23):4808–4810, Dec 2012.

[SSK18] Eric Schulz, Maarten Speekenbrink, and Andreas Krause. “A tutorial on Gaussian process regression: Modelling, exploring, and exploiting functions.” *Journal of Mathematical Psychology*, **85**:1–16, 2018.

[Uns99] M. Unser. “Splines: a perfect fit for signal and image processing.” *IEEE Signal Processing Magazine*, **16**(6):22–38, 1999.

Part IV

Demonstrating the photonic lantern wavefront sensor

CHAPTER 7

Real-time experimental demonstrations of a photonic lantern wavefront sensor

7.1 Introduction

From the ground, planar wavefronts from distant stars are warped due to Earth's turbulent atmosphere, and can become further distorted by the imperfect and unstable optics of astronomical instruments. In astronomy, these challenges motivate adaptive optics (AO), a technique which combines wavefront sensors (WFSs) and deformable mirrors (DMs) to actively flatten incoming wavefronts. Using AO, modern observing facilities have imaged exoplanets, uncovered circumstellar environments, examined the galactic center, and probed the structure of active galactic nuclei. Beyond astronomy, AO finds applications in free-space optical communications, microscopy, and remote sensing, all of which must contend with the propagation of light through inhomogeneous and dynamic media.

One of the primary goals for astronomy in the upcoming decade will be the direct imaging of an Earth-like exoplanet, and the characterization of any potential biological signatures: a scientific and technical feat whose importance was reiterated by the Astro 2020 decadal survey [Nat21], and whose challenges are now driving the development of new technologies. One of the primary hurdles is the suppression of light from the host star, which will outshine an Earth-like companion by a factor of 10^{10} in the visible [TO10]. To do so, we require coronagraphs or external starshades to blot out starlight, paired with wavefront control which must flatten incoming wavefronts at sub-nanometer precision. Such stringent demands cannot be met by conventional AO, which use sensors (e.g. Shack-Hartmann, pyramid) located in a conjugate pupil separate from the science focal plane and therefore suffer from non-common-path aberrations (NCPAs): instrumental aberrations that appear exclusively either in the science or sensing arms of the instrument [MLA12]. Additional complications include the low-wind effect (LWE), where the temperature gradients in a segmented aperture lead to aberrations that are discontinuous at segment boundaries [NMJ18], and petaling, in which the same kind of aberrations arise due to drifting misalignments in telescopes with fragmented primaries. Both are challenging to correct with general purpose sensors which re-image the pupil; petaling is of particular concern for the upcoming 30 m-class telescopes, which will provide the best chance of directly imaging Earth-like exoplanets from the ground.

We can avoid these difficulties by using wavefront sensors that operate in the science focal plane. This can be enabled in two ways: through phase diversity (e.g. COFFEE [SMP12], and related algorithms such as F&F [KKD14] and DrWho [SGG22]); or through purpose-built sensing optics such as phase holograms, asymmetric pupil masks, and dephased pinhole masks (e.g. cMWFS, APF-WFS, ZELDA, vAPP [WKS17, Mar13, NDC14, BDL19]), sometimes coupled with non-linear phase retrieval algorithms. In this work we consider an alternate architecture that works with standard linear phase retrieval algorithms, using a photonic lantern (PL): a slowly transitioning waveguide which efficiently couples multi-modal light into multiple single-moded outputs [LBB05, BGY15]. Importantly, when used to couple the aberrated telescope beam in the focal plane, PLs already have utility in several high-contrast imaging applications, including spectro-imaging and starlight nulling. PLs are also uniquely suited to couple aberrated telescope light into highly stable diffraction-limited spectrometers [LJF21], enabling direct exoplanet spectroscopy, and can serve as a gateway into a wider ecosystem of astrophotonic devices such as arrayed waveguide gratings and photonic integrated circuits [JGA23]. Simultaneously, low-spatial-frequency aberrations such as NCPAs and LWE modes can be sensed in the true science focal plane by monitoring the fluxes of the lantern's outputs. By construction, the photonic lantern wavefront sensor

(PL WFS), at least in a single monochromatic channel, is a low-order sensor, with a maximum number of sensed modes equal to the number of lantern ports (or twice that if polarizations can be separated) [LFX22, LFX23]. Accordingly, we envision that in practical ground-based applications the PL WFS will be most useful as a second-stage system, correcting the low-order NCPAs and petalling aberrations left over by a first-stage pupil-plane WFS control loop.

Previous numerical modelling [LFX22, LFX23] and experimental results [CMH18, NWB20] have made important progress in developing the PL WFS. However, unknowns, such as the PL WFS's linearity, dynamic range, and stability, especially as part of a real-time AO system, have so far prevented the PL WFS from being realistically considered. The next step — a real-time demonstration of the PL WFS as part of a modern AO system — would verify the PL WFS as a future pathway to wavefront sensing in the true science focal plane, one that is further unique due to the non-WFS capabilities that PLs can simultaneously provide, including starlight nulling, spectro-imaging, and high-resolution spectroscope injection. In this work, we take this step, using the SCExAO testbed at Subaru telescope. Our results place the PL WFS firmly on the path to eventual integration with the next generation of astronomical instruments, which will ultimately enable the imaging and characterization of an Earth-like exoplanet.

7.2 Methodology

7.2.1 Testbed setup

We used the near-infrared testbed on SCExAO [LGV20, JSC16] to demonstrate real-time control with the PL WFS. Microscope images of the 19-port PL used in this work, as well as a picture of the lantern mounted in the SCExAO testbed, are shown in Figure 7.1.1. A simplified diagram of the SCExAO beampath, containing only optics relevant for our tests, is given in the top panel of Figure 7.3.1, while an overview of the closed loop calibration

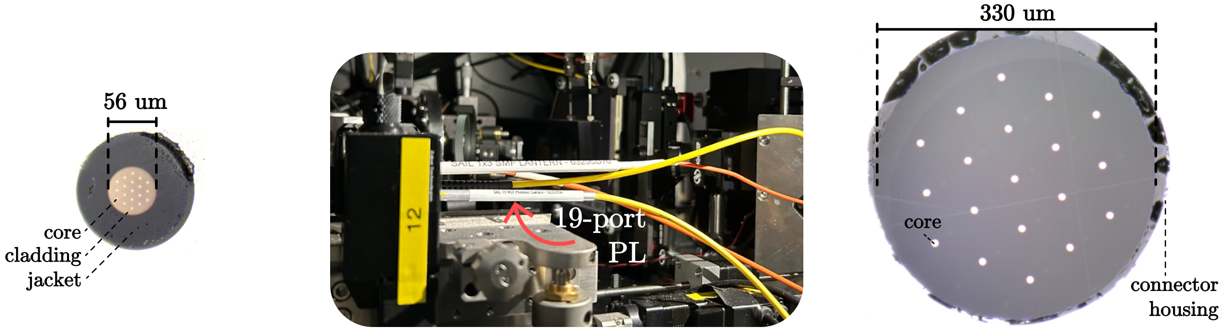

Few-mode fiber end **Side view** **Multicore fiber end**

Figure 7.1.1: End face and side views of the 19-port PL used for our wavefront sensing tests. End views are taken using a microscope at 792× magnification, and are approximately to scale relative to each other. To increase the visibility of the single-mode cores, visible light was injected through the lantern during imaging for the multicore end, and reverse injected during imaging of the few-mode fiber end.

and control process is shown at the bottom. Our light source is a supercontinuum white light laser, with a narrowband 1550 ± 50 nm filter, which is then collimated onto a 50×50 actuator DM and apodized using a set of Gaussian beam-shaping lenses. We use these lenses only because the downstream optics of the fiber injection unit were sized for the smaller apodized beam. The beam is then divided by a 90:10 beam splitter, which sends 10% of the light through to SCExAO's internal NIR camera (FLI C-Red 2) for monitoring of the point-spread function. The remaining 90% of the light is sent into the fiber injection unit: a 4-axis translation stage (shown in Figure 7.1.1, middle) that allows us to change the focal ratio of the injection as well as reposition the lantern in the focal plane. Finally, the 19 single-mode lantern output spots are imaged onto a detector (FLI C-Red 1).

7.2.2 Software

To calibrate and close the PL WFS loop, we use the Compute And Control for Adaptive Optics (CACAO) package. Figure 7.3.1 (bottom) gives an overview of the software steps required to close the loop. The raw frame taken from the CRED1 detector must be processed before it can be used for AO control. This includes dark subtraction, photometry extraction, image normalization, and reference subtraction. Initial data processing converts the full-frame detector image into a 19-dimensional vector, which we then multiply against the control matrix. The output is then fed into a leaky integrator; the output mode values are converted back to a 50×50 displacement map which is then sent out to the SCExAO DM. To compute the control matrix, we measure the response of the PL WFS to a set of aberration modes (Zernike, petaling, Hadamard, etc.) and calculate the pseudo-inverse, setting our regularization parameter to 0.1. The loop was run at 1 kHz, but in principle could be run even faster; because the computations for linear phase retrieval with only a few modes are lightweight, control speed is limited by hardware, not software.

7.2.3 Photonic lantern

The lantern in our demonstration was manufactured at the Sydney Astrophotonics Instrumentation Laboratory (SAIL) using a tapered-fiber process: a custom multicore fiber with 19 single-mode channels (hexagonal array, 6.5 μm core diameter, 60 μm core spacing, numerical aperture 0.14) was inserted into a lower-index fluorine glass capillary, and then one end was heated and drawn to form a tapered structure. The left and right panels of figure 7.1.1 show the endface geometries of the PL.

7.3 Results

7.3.1 Closed-loop correction of static aberrations

We first closed the loop on the first 5 non-piston Zernike modes, at a frequency of 1 kHz using a leaky integrator control law, with leak = 0.99 and gain = 0.2. We choose this few-mode, low-spatial-frequency basis both because the PL WFS is a few-mode sensor, and because most of the power in instrumental aberrations will appear in the first few Zernike modes

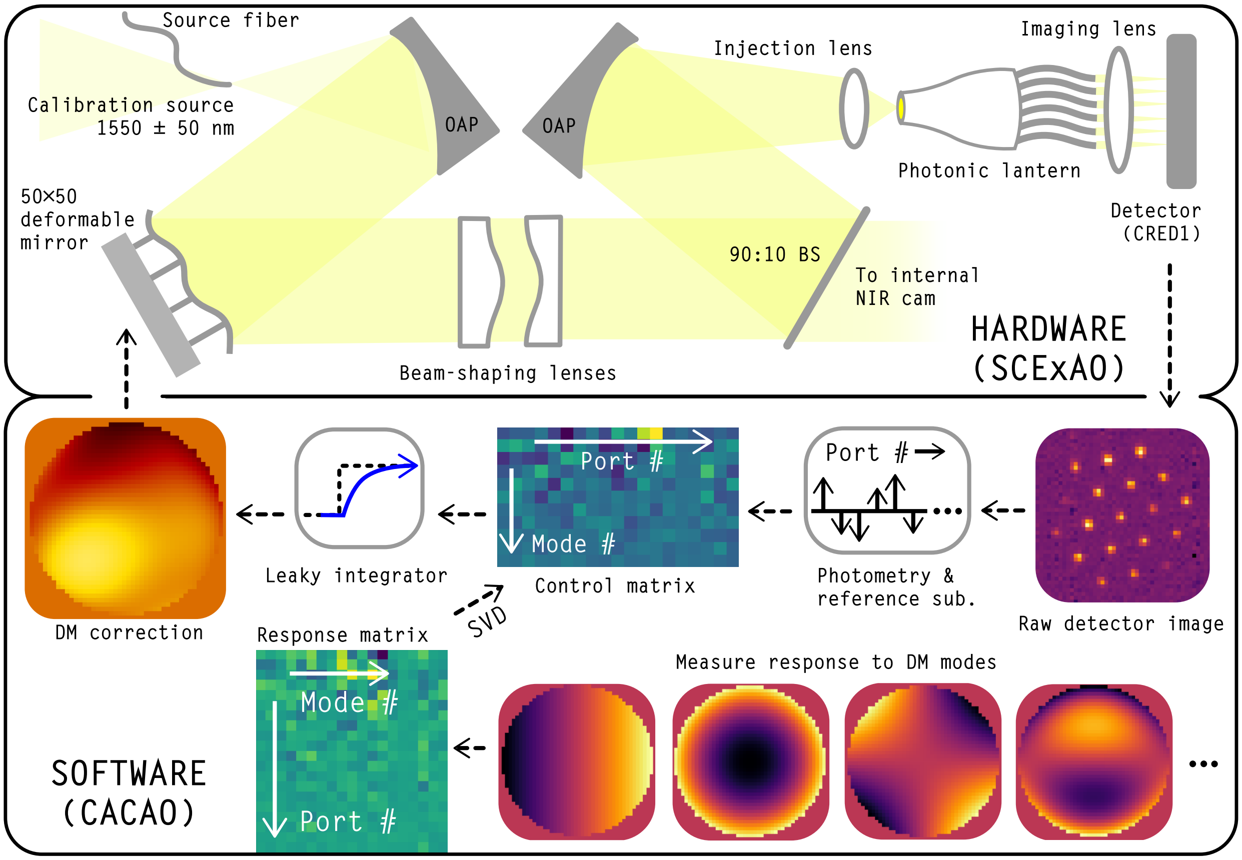

Figure 7.3.1: Top: basic hardware setup for PL wavefront sensing tests on SCExAO. OAP and BS stand for off-axis parabola and beam splitter, respectively. Only the relevant components of the SCExAO test bench are shown. Bottom: overview of closed loop control for the PL WFS, including intermediate steps like response matrix measurement. Loop calibration and control is handled with the CACAO package.

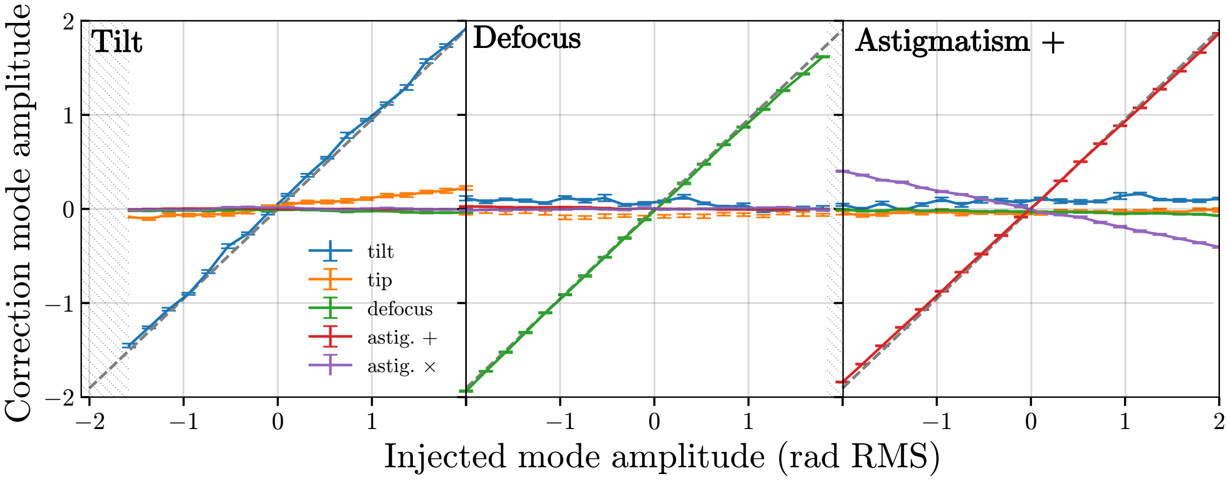

Figure 7.3.2: Correction applied by closed loop, as a function of the amount of static Zernike aberration injected into the system. The sign of the y axis has been flipped for clarity (in a perfect system, the injection and correction mode amplitude would sum to 0). The left, middle, and right panels show the loop behavior when injecting tilt, defocus, and astigmatism, respectively. For each injected mode amplitude, 20 measurements of the closed loop correction were taken; vertical bars show the standard deviation. The diagonal dashed grey lines show the line $y = 0.95x$, following the expected correction of a static aberration for our chosen loop parameters. Hatching shows regions where the loop becomes unstable.

[SFR07]. To test the loop, we injected a fixed amount of a single Zernike mode, scanning in both mode amplitude and mode index. At each point in the scan, after injecting the WFE and letting the loop settle, we sampled the closed-loop correction 20 times over the course of 2 seconds. This procedure is idealized in the sense that it neglects the effect of AO residuals, which will degrade the performance of the sensor in real-world operation; nevertheless, we believe that a comprehensive treatment of WFS performance in the presence of realistic AO residuals (e.g. as was done in [ELV22] for the pyramid) outside the scope of this work. We plan on pursuing such characterizations with future on-sky tests.

Figure 7.3.2 shows the correction of three Zernike modes over a range of static injected mode amplitudes; the other two modes show similar behavior and are omitted for brevity. We note that the loop was able to consistently correct ∼95% of the injected WFE, which is the expected amount of correction when setting leak = 0.99 and gain = 0.2 (see §7.B). Furthermore, the loop remained stable out to around ±1.6 radians of RMS WFE, or about 400 nm at our injection wavelength. The dynamic range was limited by tilt. However, it is important to note that precise loop properties are sensitive to lantern alignment, and that we did not apply a rigorous optimization to align the lantern in the best location for WFSing (which may not be at the lantern's center, or even in focus). Nevertheless, our estimate of dynamic range falls in the middle of the predicted lower and upper limits from prior simulations [LFX23], which suggested that a 19-port lantern sensing the first 5 non-piston Zernikes would have good linearity out to at least 0.5 radians and would be limited by degeneracies, arising from nonlinearity, beyond 2.3 radians. Figure 7.3.2 also shows aliasing between the two astigmatism modes (and to a lesser extent, the two tilt modes), which was not shown in simulations; it is unclear if this result is specific to our particular 19-port lantern, or if it can be reduced by adjusting the alignment.

We also consider the correction of petaling/LWE modes, an additional source of aberration besides NCPAs which will be particularly problematic for the upcoming 30-m telescopes, and are difficult to sense with conventional pupil-plane sensors. The SCExAO pupil is divided by spiders into four aperture segments, giving a total of 12 LWE modes: a local piston and two local tilt modes per aperture segment. These 12 LWE modes correspond to 11 degrees of freedom, since overall piston has no effect. We find that out of these 11 degrees, our 19-port PL WFS is able to sense 8 of them, including all 4 segment pistons; more modes might be recovered by tweaking alignment or by using a higher mode-count lantern. Figure 7.3.3 (left) repeats the laboratory procedure of Figure 7.3.2 for two of the four LWE segment piston modes, while 7.3.3 (right) compares phase maps from the WFE injection and closed-loop correction channels of the DM. We find that the PL WFS is capable of tracking all four piston modes, but with a slight undercorrection of ∼5–10%, perhaps due to cross-talk. Unlike our previous test, we observe no loop instabilities over the entire mode amplitude range of -2 to 2 radians, for any of the segment piston modes.

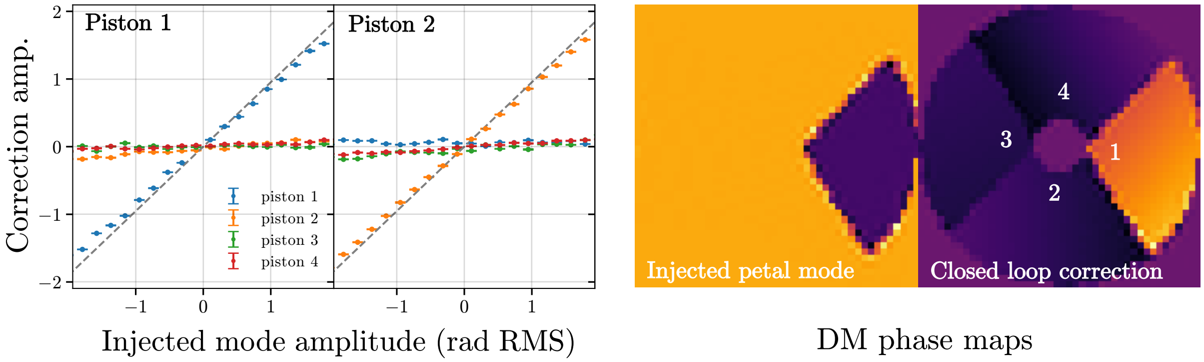

Figure 7.3.3: Left: correction applied by closed loop, as a function the amount of static LWE piston aberration artificially injected into the system. We only show correction of the first two piston modes for brevity; the correction of the other two modes behaves similarly. Right: a comparison between an example LWE piston mode and the associated correction applied by the PL WFS loop. Numbers in the rightmost panel show how the piston segments are indexed.

7.3.2 Dynamic aberrations

While the PL WFS can suppress static aberrations in closed-loop operation, in practice instrumental aberrations evolve temporally due to mechanical drifts, thermal expansion, and gravity vector changes caused by the slewing telescope. Excluding vibrations, the temporal evolution is typically on the timescale of seconds to minutes [MLA12, MKC13]. Accordingly, we tested the loop next by injecting time-evolving low order WFE with a decorrelation timescale of 1 s. This artificial WFE was composed from the first 7 non-piston Zernike modes (tilt, defocus, astigmatism, and coma); each mode amplitude is independently updated at a rate of 1 kHz according to an autoregressive formula, setting a per-mode amplitude of 0.05 radians RMS. For more details, see §7.A. To calibrate the loop, we measured the sensor response matrix against the first 20 non-piston Zernike modes, and used a singular value decomposition to compute the control modes of the system. In all, we found 12 control modes, constructed from independent linear combinations of the first 20 Zernikes. Note that we cannot recover a full 19 modes with a 19-port PL WFS because we lose one degree of

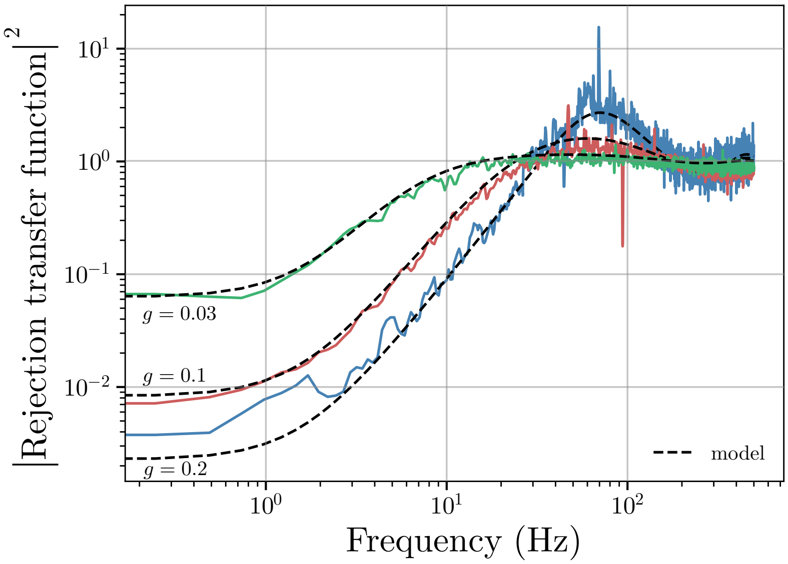

Figure 7.3.4: Colored lines: experimentally measured squared modulus of the AO system's rejection transfer function (equal to the ratio between the closed and open loop power spectral distributions) for the first control mode, at different leaky integrator gains. Transfer functions for other control modes look similar. Dashed black lines: theoretically expected squared modulus of the rejection transfer function, at different gains.

freedom to image normalization and another to a global piston-like mode [LFX22]; beyond that, there is no guarantee *a priori* that the 19 lantern outputs will behave independently for a given aberration basis. We then logged the PL WFS output in both open and closed loop; the closed-loop leak was 0.99. The ratio of the open- and closed-loop power spectral densities (PSDs), which we estimated from the collected time-series data, approximates the squared modulus of the system's rejection transfer function. We plot the experimentally measured transfer functions for the first control mode in Figure 7.3.4 at different gains, along with the transfer functions expected by a theoretical model for our closed-loop system, set entirely by system latency, detector framerate, leak, and gain. Our model transfer function is presented in §7.B. All measured transfer functions in Figure 7.3.4 agree with models, and show successful rejection of the slowly varying components of the injected WFE. We find 0 dB frequencies of 15, 23, and 32 Hz, and rejection at 1 Hz of roughly 4×, 10×, and 12×, for gains of 0.03, 0.1, and 0.2, respectively. We also note that these transfer functions can be recovered by comparing open- and closed-loop PSDs, even without injecting artificial WFE (relying solely on the actual system instabilities), implying that the PL WFS can actively correct the ambient WFE of the SCExAO testbed.

7.3.3 Stability

In a perfectly stable AO system, measurement of the response matrix only needs to happen once; in reality, an accumulation of slight changes in WFS optical properties (e.g. caused by temperature drifts) will require eventual re-calibration. In this subsection, we present a rudimentary estimate of the average drift rate of the PL response matrix, separate from drift inherent to the SCExAO testbed. The basis of our test is as follows. Suppose at some initial time, the response matrix of the PL WFS, A_0, is measured, and the control matrix A_0^+ is computed. Then, $A_0^+ A_0 = I$, where I is the identity matrix. However, later measurements of the response matrix, denoted A_n, will deviate from A_0, and thus the product $A_0^+ A_n$ will drift from I. This accumulated miscalibration error, err_n, can be quantified as

$$\mathrm{err}_n = |A_0^+ A_n - I| \tag{7.1}$$

where $|M|$ denotes the Frobenius norm of matrix M.

We applied this test to the PL WFS, over the course of one week. In practice, we first measured the reference response and control matrices, A_0 and A_0^+, for 11 control modes. At roughly the same time on later days, we then re-aligned the PL injection, closed the loop using the reference control matrix, and used `cacao`'s built-in `selfRM` function to measure $A_0^+ A_n$ directly. `selfRM` works by manually injecting a small amount of each control mode, in sequence, into the system while the loop is closed, and recording the applied correction. After one week, we compute an accumulated miscalibration error of ~ 1.4. On average, this corresponds to an accumulated per-mode error of $1.4/11 \approx 0.13$ (13%, since each column of

$A_0^+ A_n$ should have a norm of 1) over the course of the week, or a per-mode drift of around 2% per day. This measurement of long-term stability (separate from diurnal stability) is a lower bound since we currently cannot completely compensate for the drifting bulk optics on the SCExAO testbed. Thus, the accumulated miscalibration error we measure is at least partially attributable to testbed instability; nevertheless, our results confirm that we can still close the wavefront control loop using a week-old calibration.

7.4 Discussion and Conclusion

The PL WFS offers a novel approach to focal-plane wavefront sensing, enabling both the correction of NCPAs and petaling modes, which currently limit achievable contrast in exoplanet direct imaging, as well as a host of additional high-contrast applications such as starlight nulling and high-resolution spectrometer injection. We present the first real-time demonstration of the photonic lantern wavefront sensor, and verify its ability to correct both low-order Zernike modes and LWE/petaling modes from the science path focal plane. We further confirm the system's ability to track and correct dynamic Zernike WFE which varies on a timescale of $\sim$ 1s, representative of quasistatic NCPAs — currently one of the main limiting factors in achieving higher contrasts. Our control method has the added benefit of simplicity, using conventional linear phase retrieval and thus easily integrable into existing AO systems. However, future work still needs to be done before a PL WFS can be adopted in next-generation high-contrast imaging instruments. For one, we require a better understanding of how PL geometry and the overall system alignment impacts the properties of the AO system. In terms of practicality, we must also measure the PL WFS's WFE sensitivity and limiting magnitude in realistic conditions. In regards to the latter, we expect that the PL WFS should be able to work with fainter guide stars than conventional pupil-plane AO systems, because PLs have relatively high coupling efficiencies and concentrate collected light into fewer pixels. In the near future, we hope to take the PL WFS on-sky, and to leverage spectral dispersion and/or polarization to increase the amount of wavefront information provided by the lantern. Eventually, we believe that the PL WFS, and the wider ecosystem of astrophotonic devices, will ultimately play a critical part in the successful direct imaging and characterization of another Earth.

Appendix - Chapter 7

7.A Simulating temporal WFE

To simulate dynamic WFE, we continuously apply a dynamic phase map composed from the first 7 non-piston Zernike modes (tilt, defocus, astigmatism, and coma) to a DM channel separate from the closed loop channel. Denote z_j^n as the amplitude of mode j at timestep n. We use the following autoregressive formula to update the mode amplitudes:

$$z_j^{n+1} = e^{-1/t_c}\, z_j^n + a_j \sqrt{1 - e^{-2/t_c}}\, m^n. \tag{7.2}$$

Here, n superscripts denote timestep while j subscripts denote mode index. t_c is the decorrelation time in frames, a_j is the desired average mode amplitude for mode j in units of radians RMS, and m^n is the nth draw of a Gaussian distributed random variable with mean $\mu = 0$ and standard deviation $\sigma = 1$. In this particular test, we set $t_c = 1000$ frames and $a_j = 0.05$ radians RMS, for all modes. The overall phase map was updated at a rate of 1 kHz. See Chapter §13 for derivation of the autoregressive method.

7.B Modeled rejection transfer function

A classic AO loop using linear phase retrieval and a leaky integrator control law has the following complex-valued rejection transfer function for mode j:

$$h_{\mathrm{rej},j}(f) = \left[1 + g_j \frac{\mathrm{sinc}(T_e f) e^{-\pi i f (T_e + 2\tau)}}{1 - l_j\, e^{-2\pi i T_e f}}\right]^{-1}. \tag{7.3}$$

Here, $\mathrm{sinc}(x) \equiv \sin(\pi x)/\pi x$. This transfer function model is determined by four parameters: the modal gain g_j, the modal leak l_j, the detector integration time T_e, and the loop latency τ. The leakage term l_j is defined such that $l_j = 1$ corresponds an ordinary (i.e. non-leaky) integrator. For our PL WFS tests on SCExAO, the detector integration time is 1 ms and the loop latency was independently measured using CACAO to be 2.7 ms. The leak was set to 0.99 for all tests. See Chapter 14 for a detailed derivation.

For static aberrations, the rejection becomes

$$h_{\mathrm{rej},j}(f = 0) = \frac{1 - l_j}{1 - l_j + g_j}. \tag{7.4}$$

REFERENCES

[BDL19] Bos, S. P., Doelman, D. S., Lozi, J., Guyon, O., Keller, C. U., Miller, K. L., Jovanovic, N., Martinache, F., and Snik, F. "Focal-plane wavefront sensing with the vector-Apodizing Phase Plate." *A&A*, **632**:A48, 2019.

[BGY15] T. A. Birks, I. Gris-Sánchez, S. Yerolatsitis, S. G. Leon-Saval, and R. R. Thomson. "The photonic lantern." *Adv. Opt. Photon.*, **7**(2):107–167, Jun 2015.

[CMH18] Mark K. Corrigan, Timothy J. Morris, Robert J. Harris, and Theodoros Anagnos. "Demonstration of a photonic lantern low order wavefront sensor using an adaptive optics testbed." In Laird M. Close, Laura Schreiber, and Dirk Schmidt, editors, *Adaptive Optics Systems VI*, volume 10703, p. 107035H. International Society for Optics and Photonics, SPIE, 2018.

[ELV22] Byron Engler, Miska Louarn, Christophe Verinaud, Steve Weddell, and Richard Clare. "Flip-flop modulation method used with a pyramid wavefront sensor to correct piston segmentation on ELTs." *Journal of Astronomical Telescopes, Instruments, and Systems*, **8**, 03 2022.

[JGA23] Nemanja Jovanovic, Pradip Gatkine, Narsireddy Anugu, Rodrigo Amezcua-Correa, Ritoban Basu Thakur, Charles Beichman, Chad Bender, Jean-Philippe Berger, Azzurra Bigioli, Joss Bland-Hawthorn, Guillaume Bourdarot, Charles M. Bradford, Ronald Broeke, Julia Bryant, Kevin Bundy, Ross Cheriton, Nick Cvetojevic, Momen Diab, Scott A. Diddams, Aline N. Dinkelaker, Jeroen Duis, Stephen Eikenberry, Simon Ellis, Akira Endo, Donald F. Figer, Michael Fitzgerald, Itandehui Gris-Sanchez, Simon Gross, Ludovic Grossard, Olivier Guyon, Sebastiaan Y. Haffert, Samuel Halverson, Robert J. Harris, Jinping He, Tobias Herr, Philipp Hottinger, Elsa Huby, Michael Ireland, Rebecca Jenson-Clem, Jeffrey Jewell, Laurent Jocou, Stefan Kraus, Lucas Labadie, Sylvestre Lacour, Romain Laugier, Katarzyna Ławniczuk, Jonathan Lin, Stephanie Leifer, Sergio Leon-Saval, Guillermo Martin, Frantz Martinache, Marc-Antoine Martinod, Ben Mazin, Stefano Minardi, John D Monnier, Reinan Moreira, Denis Mourard, Abani Shankar Shankar Nayak, Barnaby Norris, Ewelina Obrzud, Karine Perraut, François Reynaud, Steph Sallum, David Schiminovich, Christian Schwab, Eugene Serbayn, Sherif Soliman, Andreas Stoll, Liang Tang, Peter Tuthill, Kerry Vahala, Gautam Vasisht, Sylvain Veilleux, Alexander B. Walter, Edward J Wollack, Yinzi Xin, Zongyin Yang, Stephanos Yerolatsitis, Yang Zhang, and Chang-Ling Zou. "2023 Astrophotonics Roadmap: pathways to realizing multi-functional integrated astrophotonic instruments." *Journal of Physics: Photonics*, 2023.

[JSC16] N. Jovanovic, C. Schwab, N. Cvetojevic, O. Guyon, and F. Martinache. "Enhancing Stellar Spectroscopy with Extreme Adaptive Optics and Photonics." *Publications of the Astronomical Society of the Pacific*, **128**(970):121001, nov 2016.

[KKD14] Visa Korkiakoski, Christoph U. Keller, Niek Doelman, Matthew Kenworthy, Gilles Otten, and Michel Verhaegen. “Fast & Furious focal-plane wavefront sensing.” *Appl. Opt.*, **53**(20):4565–4579, Jul 2014.

[LBB05] S. G. Leon-Saval, T. A. Birks, J. Bland-Hawthorn, and M. Englund. “Multimode fiber devices with single-mode performance.” *Opt. Lett.*, **30**(19):2545–2547, Oct 2005.

[LFX22] Jonathan Lin, Michael P. Fitzgerald, Yinzi Xin, Olivier Guyon, Sergio Leon-Saval, Barnaby Norris, and Nemanja Jovanovic. “Focal-plane wavefront sensing with photonic lanterns: theoretical framework.” *J. Opt. Soc. Am. B*, **39**(10):2643–2656, Oct 2022.

[LFX23] Jonathan Lin, Michael P. Fitzgerald, Yinzi Xin, Yoo Jung Kim, Olivier Guyon, Sergio Leon-Saval, Barnaby Norris, and Nemanja Jovanovic. “Focal-plane wavefront sensing with photonic lanterns II: numerical characterization and optimization.” *J. Opt. Soc. Am. B*, Oct 2023.

[LGV20] Julien Lozi, Olivier Guyon, Sébastien Vievard, Ananya Sahoo, Vincent Deo, Nemanja Jovanovic, Barnaby Norris, Marc-Antoine Martinod, Ben Mazin, Alex Walter, Neelay Fruitwala, Sarah Steiger, Kristina Davis, Peter Tuthill, Tomoyuki Kudo, Hajime Kawahara, Takayuki Kotani, Michael Ireland, Theodoros Anagnos, Chrstian Schwab, Nick Cvetojevic, Elsa Huby, Sylvestre Lacour, Kevin Barjot, Tyler D. Groff, Jeffrey Chilcote, Jeremy Kasdin, Frantz Martinache, Romain Laugier, Mamadou N’Diaye, Justin Knight, Jared Males, Steven Bos, Frans Snik, David Doelman, Kelsey Miller, Eduardo Bendek, Ruslan Belikov, Eugene Pluzhnik, Thayne Currie, Masayuki Kuzuhara, Taichi Uyama, Jun Nishikawa, Naoshi Murakami, Jun Hashimoto, Yosuke Minowa, Christophe Clergeon, Yoshito Ono, Naruhisa Takato, Motohide Tamura, Hideki Takami, and Masa Hayashi. “Status of the SCExAO instrument: recent technology upgrades and path to a system-level demonstrator for PSI.” In Laura Schreiber, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems VII*, volume 11448, p. 114480N. International Society for Optics and Photonics, SPIE, 2020.

[LJF21] Jonathan Lin, Nemanja Jovanovic, and Michael P. Fitzgerald. “Design considerations of photonic lanterns for diffraction-limited spectrometry.” *J. Opt. Soc. Am. B*, **38**(7):A51–A63, Jul 2021.

[Mar13] Frantz Martinache. “The Asymmetric Pupil Fourier Wavefront Sensor.” *pasp*, **125**(926):422, April 2013.

[MKC13] Martinez, P., Kasper, M., Costille, A., Sauvage, J. F., Dohlen, K., Puget, P., and Beuzit, J. L. “Speckle temporal stability in XAO coronagraphic images - II. Refine model for quasi-static speckle temporal evolution for VLT/SPHERE.” *A&A*, **554**:A41, 2013.

[MLA12] Martinez, P., Loose, C., Aller Carpentier, E., and Kasper, M. "Speckle temporal stability in XAO coronagraphic images." *A&A*, **541**:A136, 2012.

[Nat21] National Academies of Sciences, Engineering, and Medicine. *Pathways to Discovery in Astronomy and Astrophysics for the 2020s.* The National Academies Press, Washington, DC, 2021.

[NDC14] Mamadou N'Diaye, Kjetil Dohlen, Amandine Caillat, Anne Costille, Thierry Fusco, Aïssa Jolivet, Fabrice Madec, Laurent Mugnier, Baptiste Paul, Jean-François Sauvage, Rémi Soummer, Arthur Vigan, and J. K. Wallace. "Design optimization and lab demonstration of ZELDA: a Zernike sensor for near-coronagraph quasi-static measurements." In Enrico Marchetti, Laird M. Close, and Jean-Pierre Vran, editors, *Adaptive Optics Systems IV*, volume 9148 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 91485H, August 2014.

[NMJ18] N´Diaye, M., Martinache, F., Jovanovic, N., Lozi, J., Guyon, O., Norris, B., Ceau, A., and Mary, D. "Calibration of the island effect: Experimental validation of closed-loop focal plane wavefront control on Subaru/SCExAO." *A&A*, **610**:A18, 2018.

[NWB20] Barnaby R. M. Norris, Jin Wei, Christopher H. Betters, Alison Wong, and Sergio G. Leon-Saval. "An all-photonic focal-plane wavefront sensor." *Nature Communications*, **11**(1):5335, Oct 2020.

[SFR07] Jean-François Sauvage, Thierry Fusco, Gérard Rousset, and Cyril Petit. "Calibration and precompensation of noncommon path aberrations for extreme adaptive optics." *J. Opt. Soc. Am. A*, **24**(8):2334–2346, Aug 2007.

[SGG22] Skaf, Nour, Guyon, Olivier, Gendron, Éric, Ahn, Kyohoon, Bertrou-Cantou, Arielle, Boccaletti, Anthony, Cranney, Jesse, Currie, Thayne, Deo, Vincent, Edwards, Billy, Ferreira, Florian, Gratadour, Damien, Lozi, Julien, Norris, Barnaby, Sevin, Arnaud, Vidal, Fabrice, and Vievard, Sébastien. "On-sky validation of image-based adaptive optics wavefront sensor referencing." *A&A*, **659**:A170, 2022.

[SMP12] J.-F. Sauvage, L. Mugnier, B. Paul, and R. Villecroze. "Coronagraphic phase diversity: a simple focal plane sensor for high-contrast imaging." *Opt. Lett.*, **37**(23):4808–4810, Dec 2012.

[TO10] W. A. Traub and B. R. Oppenheimer. "Direct Imaging of Exoplanets." In S. Seager, editor, *Exoplanets*, pp. 111–156. 2010.

[WKS17] M. J. Wilby, C. U. Keller, F. Snik, V. Korkiakoski, and A. G. M. Pietrow. "The coronagraphic Modal Wavefront Sensor: a hybrid focal-plane sensor for the high-contrast imaging of circumstellar environments." *aap*, **597**:A112, January 2017.

CHAPTER 8

On-sky testing and spectral dispersion

In this chapter, we demonstrate on-sky wavefront control using a photonic lantern wavefront sensor, coupled to a state-of-the-art 8-m telescope and high-order adaptive optics system. This is, to my best knowledge, the first such demonstration. While photonic wavefront sensors have already been tested off-sky, on-sky experiments are critical in confirming if these devices can perform in realistic conditions which are difficult to emulate in the lab, in turn improving the readiness of these technologies for the upcoming ground-based 30-m telescopes and space-based HWO. Furthermore, our on-sky demonstration is simultaneously the first for a wavelength-dispersed wavefront sensor, photonic or otherwise. Wavelength dispersion is currently underexplored within the adaptive optics, even though the wavelength dependence of the sensor carries useful information for wavefront sensing. Spectrally dispersed wavefront sensors will likely gain interest in tandem with photonic wavefront sensors, due to the relative ease at which photonic devices can be dispersed. The possibility of simultaneous dispersion and sensing makes possible future self-calibrating astronomical spectrographs, as well as multiband free-space optical communications systems and compact hyperspectral imagers.

This chapter is adapted from a manuscript currently accepted for publication in Optics Letters, which is divided into a letter-format paper and a supplement. For the dissertation version, I have combined the two components, with §8.1-8.3 constituting the letter component while §8.4-8.5 constitutes the supplement, along with extra material not included in the journal version.

8.1 Introduction

Adaptive optics (AO) systems cancel wavefront distortions in real-time by driving wavefront correctors such as deformable mirrors (DMs) with wavefront sensors (WFSs), typically in closed-loop control. Such systems have found use in a wide range of imaging and telecommunications applications, including ground-to-space optical communications [Tys96], deep tissue microscopy [RCR21], and remote sensing [DKC09]. In astronomy, AO systems correct for wavefront distortions which may originate from the overhead turbulent atmosphere (in the case of ground-based observation) or time-varying imperfections in the scientific instrument. AO has enabled highly-resolved imaging of the galactic center [GSW08], exoplanets [MZK10], and more; [Guy18] provides a review.

One of the primary goals for astronomy for the next decades is the direct imaging of an Earth-like exoplanet and the identification of potential biosignatures [Nat21]. This achievement will require the separation of exoplanetary light from that of a host star which outshines its companion by 10 orders of magnitude in visible wavelengths [TO10], for instance by using coronagraphy [GPK06]. However, coronagraphic contrast is highly sensitive to wavefront distortions, requiring correction at a precision and stability which exceeds the capabilities of current AO systems. Such systems typically use general-purpose WFSs such as the Shack-Hartmann lenslet array or pyramid WFS, which re-image the telescope pupil and split off a portion of the light to a dedicated WFS arm. This gives rise to non-common-path aberrations (NCPAs): quasi-static aberrations arising from opto-mechanical deformations which affect one arm of the instrument but not the other [MLA12]. Other sources of wavefront error include the low-wind effect [SFL16] and petaling [BGR22], both of which arise from fragmentation of the telescope aperture and induce discontinuous phase aberrations over the pupil. While such aberrations are difficult to correct with pupil-reimaging WFSs, they can be sensed from the focal plane: a strategy that also mitigates NCPAs because the instrument now has a single common path. To date, several focal-plane sensing techniques have been proposed

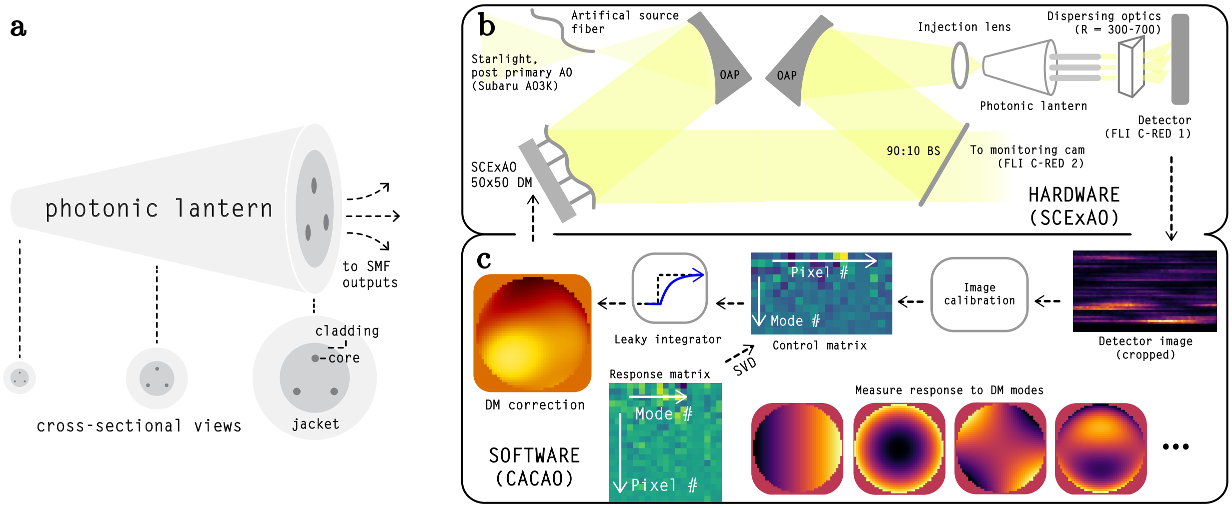

Figure 8.1.1: **a:** A 3-port PL, with 3 outputs. The PL couples light copropagating in multiple fiber modes into multiple single-mode outputs. Note that we SCExAO contains both a 3-port and 19-port PL, both of which we use in our tests; a 3-port PL is drawn for simplicity. **b:** Simplified beam diagram of the astrophotonics platform at SCExAO. Light comes from either the facility AO or a supercontinuum laser, and reflects off a 50 × 50 actuator DM. The light is then divided between an internal camera, used to monitor the Strehl ratio, and the PL, which is dispersed at a resolution of several hundred. **c:** Software steps used to perform wavefront control; we use a linear phase retrieval method.

(e.g. DRWHO [SGG22]), as well as dedicated sensor designs like the self-coherent camera [GBR10]. Recently, focal-plane sensing was also demonstrated with a photonic lantern (PL; [LBB05]). As shown in Figure 8.1.1a, the PL is a slowly varying waveguide similar to a tapered multicore fiber which can efficiently couple aberrated telescope light into single-mode fibers (SMFs); optical throughputs of PLs have reached 97% in laboratory testing [NSS12]. Such devices also have potential astronomical uses beyond AO: for instance, they can stably inject light into spectrographs [LJF21, VLL24] e.g. for high-resolution exoplanet spectroscopy, probe sub-diffraction-limit spatial features with spectroastrometry [KFLed], and null starlight for high-contrast imaging [XJR22]. Simultaneously, the PL can drive a second stage of common-path wavefront correction [LFX22], complementing the telescope's primary AO.

Previous demonstrations of the PL WFS used the non-dispersed spot images of a PL's single-mode outputs to perform phase retrieval [NWB20] and drive wavefront correction [LFX23]. The purpose of this work is to show that the sensing capability of the PL is retained, and even improved, when wavelength dispersion is introduced, and that such devices are suitable for second stage correction of NCPA and low-wind effect. To this end, we use the SCExAO high-contrast imaging testbed at the 8 m Subaru Telescope [GAA22], also used in [LFX23]. This near-infrared testbed contains an astrophotonics platform that was recently upgraded to include a low-resolution diffraction-limited spectrograph, as well as two photonic lanterns: a 19-port PL optimized for high throughput in near-infrared wavelengths ($\lambda = 1 - 1.8\,\mu$m), and a 3-port PL optimized for throughput in the astronomical H-band. Both devices were fabricated at the Sydney Astrophotonics Instrumentation Laboratory by tapering either multicore fibers or SMF bundles. We present both off-sky and on-sky demonstrations of real-time wavefront control.

8.2 Results

8.2.1 Off-sky laboratory demonstration: 19-port PL

We first present off-sky results from SCExAO. Figure 8.1.1b gives an overview of our setup: briefly, this testbed emulates a point source using a supercontinuum laser, which is collimated onto a 2500 actuator DM. 90% of the light is directed to a 4-axis stage which holds the PL in a focal plane and enables tuning of alignment and focal ratio, while the remainder is sent to an FLI C-RED 2 detector for PSF and Strehl ratio monitoring. The output of the PL is dispersed at a resolution of 700 - 300 over $\lambda = 1 - 1.8$ μm with a prism onto an FLI C-RED 1 detector. When on-sky, the laser source is replaced with starlight from the facility AO system (AO3K, [LAB24]), which uses a separate 3228-actuator DM and pyramid WFS. See §8.5 for more details.

Figure 8.1.1c overviews the calibration and control process for the 19-port PL WFS tests. First, we measure the slope of the PL's response to the first 100 non-piston Zernike modes. This gives the response matrix, which approximates the relation between the pupil phase and WFS output; the singular value decomposition (SVD) of this matrix determines the sensed aberration modes, or control modes. To compute the phase correction, the phase retrieved by the PL WFS is fed into a leaky integrator. We then closed the AO loop (i.e. applied the correction to the DM in real-time) at a framerate of 1.5 kHz, without injecting any additional WFE into the system. After running a "self-test" of the loop, where a static amount of each control mode is injected into the system one-by-one on a separate channel of the DM, we found that the first 52 control modes — each a linear combination of Zernike modes — were clearly corrected, with correction of the remainder too slow to be useful. Note that an un-dispersed 19-port PL senses strictly less than 19 aberration modes [LFX22], and that previous experiments from [LFX23] demonstrated control over only 15 modes. We provide an explanation for this increase in §8.3. Next, to test the loop, we generated phase screens with an RMS amplitude of 150 nm and windspeed of 10 m/s, and applied them to an independent channel of the DM. These phase screens were high-pass-filtered Kolmogorov screens to emulate first-stage correction of atmospheric turbulence by Subaru's facility AO. Finally, we recorded the amplitudes for each control mode, as measured through the PL WFS, in both open- and closed-loop operation. The ratio of the closed-loop and open-loop power spectral densities (PSDs) estimates the squared modulus of the rejection transfer function. We find that the correction loop is stable up to a limit of 150 - 200 nm RMS of wavefront error, dependent on alignment and control parameters. This somewhat limited dynamic range is expected by simulations, e.g. [LFX22].

Figure 8.2.1 shows the PSDs and rejection transfer function for an example control mode of the Zernike-calibrated 19-port PL WFS. We find good agreement with an analytic model derived in [LFX23], except at the slowest frequencies where the deviation approaches 2-3×. This suggests that a subset of control modes had lower sensitivity and hence slower correction than expected; one possible cause is a slow optomechanical drift within SCExAO, which would change the alignment of the PL and eventually outdate the sensor calibration.

We also tried a calibrating the 19-port PL in the zonal mode basis, which is more typical for astronomical AO. Zonal modes correspond to the pistoning of each DM element. This calibration yielded 58 control modes; we then closed the AO loop using a leaky integrator controller with a leak of 0.99 and a gain of 0.1. Due to external constraints, the PL was dispersed onto an FLI C-RED 2 camera instead of the C-RED 1 used in previous tests; we ran this camera at a framerate of 1450 Hz and measured a system latency of 2.96 frames. Similar to the earlier demonstration which used Zernike basis, we then tested the system by injecting high-pass-filtered Kolmogorov phase screens with an RMS amplitude of 100 nm and a windspeed of 10 m/s onto a separate channel of the DM. From this test, we obtained the open- and closed-loop PSDs and the rejection transfer function, as shown in the middle row of Figure 8.2.1. We note that the agreement between the experimentally measured average transfer function and the theoretical expectation is somewhat better in this test than the prior test, but it is unclear if this discrepancy is due to the difference in calibration basis,

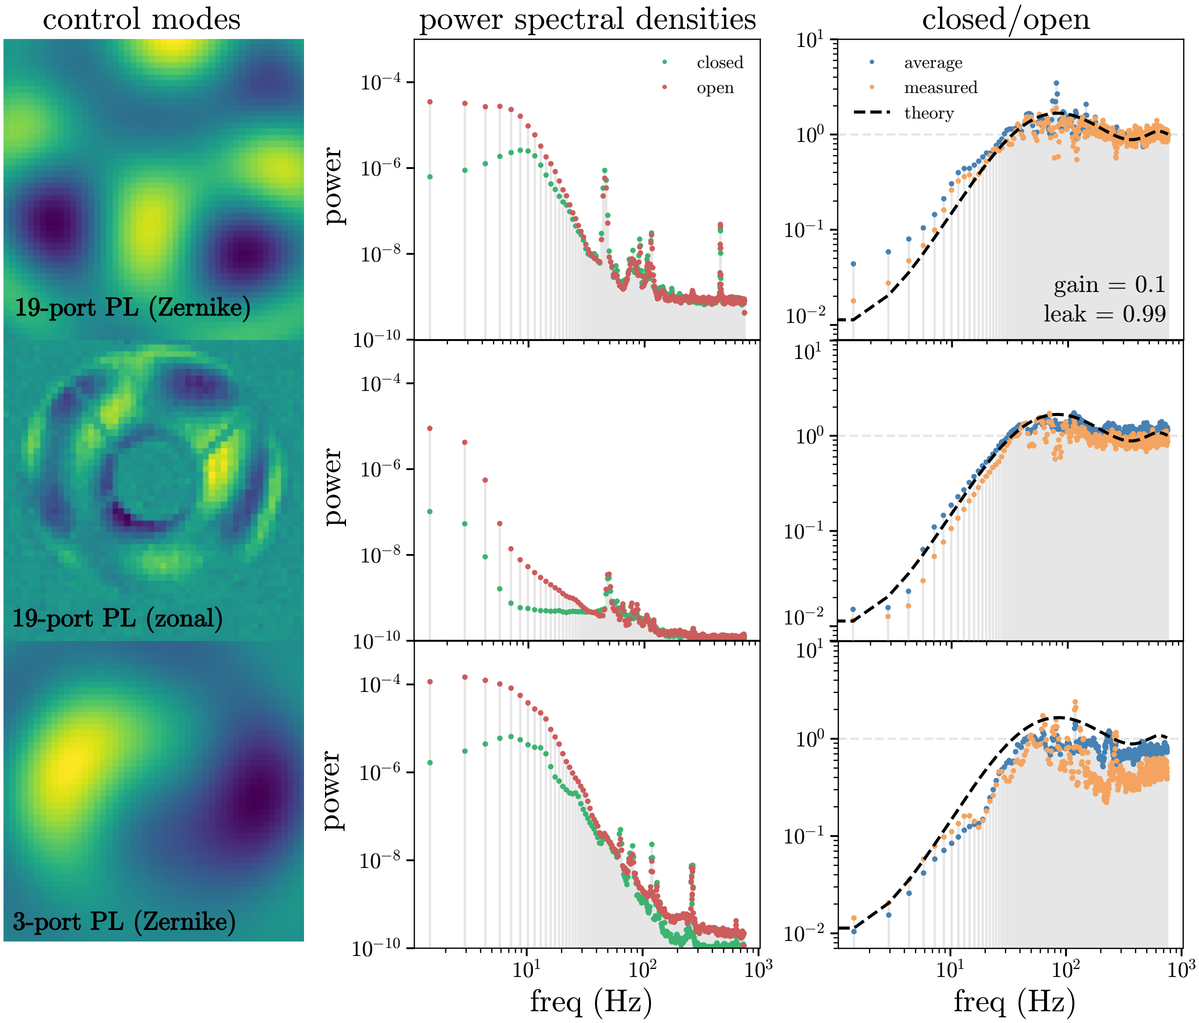

Figure 8.2.1: Top row, from left to right: an example control mode, open- and closed-loop power spectral densities for this mode's amplitude, and the rejection transfer function of this mode, for a 19-port PL calibrated in Zernike basis; orange dots correspond to the given mode; blue dots correspond to the average of the control modes; the dashed black lines give the expected transfer function, derived analytically. Middle row: the same as the top row, but for a 19-port PL calirbated in zonal (DM actuator) basis. Bottom row: the same as the top row, but for a 3-port PL calibrated in Zernike basis. The detector used in the top row was an FLI C-RED 1 while the bottom two rows used an FLI C-RED 2.

the detector swap, optomechanical instability of the SCExAO system, or some mixture of the three.

8.2.2 Off-sky laboratory demonstration: 3-port PL

We also demonstrated closed-loop operation for the dispersed 3-port PL, imaged on an FLI C-RED 2. To explicitly ensure that the outputs of this device were single-moded, we blocked all wavelengths $\lesssim 1260$ nm using a filter. The calibration process was the same as the above, yielding 8 control modes: after closing the loop using the same loop parameters, we confirmed that all modes were clearly corrected by applying the self-test. For reference, a non-dispersed 3-port PL can in principle sense at most 2 modes [LFX23]. Finally, we injected the same type of filtered Kolmogorov WFE onto an independent channel of the DM, at an amplitude of 100 nm RMS. Notably, the loop was unstable for 150 nm of WFE, possibly indicating that this device has a lower dynamic range than the 19-port PL, though we were later able to close on 200 nm of turbulence with some tuning of modal gains. The bottom row of Figure 8.2.1 shows an example control mode, power spectral densities, and the squared rejection transfer function of this system. When increasing the wavelength range to use the entire bandwidth of the spectrograph, we found that the device is capable of controlling 10 modes; further increases might be possible by optimizing the alignment of the system, which we discuss in 3.4.

8.2.3 On-sky demonstration

On 9/17/2024, we demonstrated closed-loop wavefront control using the spectrally dispersed 19-port PL at Subaru/SCExAO, using Humu (Altair) as our natural guide star. This demonstration was performed as part of SCExAO engineering time, under proposal ID S24B-EN03 (PI Julien Lozi). The seeing during the test was $\sim$ 0.5”, as estimated by the Maunakea Weather Center. Our test was also downstream of AO3K, which provided first-stage wavefront correction and yielded ∼50% Strehl for $\lambda \sim 1.5 - 1.8$ μm, as measured from the PSF monitoring camera: as such, this demonstration is also the first to show the PL working as a second-stage focal plane WFS.

Because our PL's spectral traces overlap on the detector, we were unable to account for spectral differences between SCExAO's calibration source and that of our natural guide star, and had to calibrate the WFS on-sky. This overlap is due to hardware constraints, which will be rectified in the future with a new PL; for more details, see §8.4. In the interest of on-sky time, we only calibrated the WFS against the first 30 non-piston Zernikes. The SVD of this matrix obtained 14 control modes, though higher fidelity control should be possible using an on-sky calibration with more basis modes or an off-sky, daytime calibration, as in the previous section. Finally, we closed the loop using a leak of 0.97 and a gain of 0.2 at a rate of 1.7 kHz and latency of 3.77 frames, using the PL and SCExAO DM to perform secondary correction after AO3k. Our somewhat high latency could be reduced with improvements to the (currently Python-based) processing of the dispersed PL image. Figure 8.2.2 gives the average open- and closed-loop PSDs for the 14 control modes, as well as the estimated rejection transfer function; the peak at 100 Hz is a result of noise amplification at high gain, and is captured by the analytic model. Beyond 100 Hz, we find deviations which might be related to the large spikes in the PSDs near 100 and 200 Hz, perhaps corresponding to strong vibrations.

The rightmost panel of 8.2.2 plots the estimated *H*-band Strehl ratio over the course of the test, and includes insets of the *H*-band point spread function (PSF) immediately before and after loop closing, averaged over a two minute timescale. Note that the closed-loop data was taken before the open-loop data, and that after opening the loop, the seeing gradually increased to ∼1", at which point we were not able to close the PL loop again due to the limited dynamic range of the PL. To account for this trend, we performed linear fits of both the closed- and open-loop Strehl data, which indicate that the second-stage wavefront control provided a Strehl improvement of ∼17%. Using the extended Maréchal approximation, this

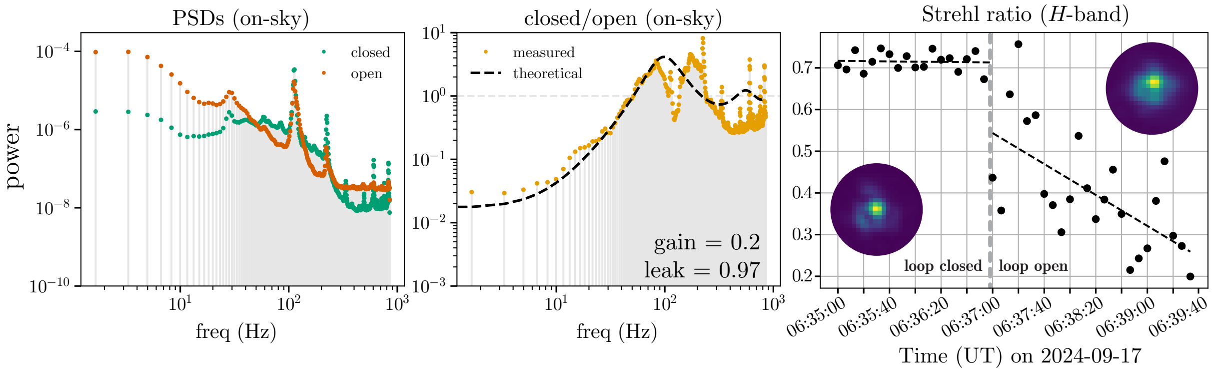

Figure 8.2.2: Left: on-sky open- and closed-loop PSDs for the dispersed 19-port PL WFS. Middle: the ratio of the PSDs, which estimates the squared modulus of the rejection transfer function. Agreement between the experimentally measured transfer function and the theoretical expectation for our control parameters is good up to around 100 Hz. Right: the *H*-band Strehl ratio as a function of time. The wavefront control loop is initially closed and is then opened. Dashed black lines show linear regressions for the open- and closed-loop Strehl data. Circular insets show the averaged *H*-band PSF in under open- and closed-loop operation.

corresponds to a reduction in RMS WFE of ∼190 nm to ∼140 nm, which in turn matches the integral of the difference between the experimentally measured open- and closed-loop PSDs, computed to be 52 nm.

8.3 Discussion

It is encouraging to see that wavelength dispersion already has a beneficial impact on the WFS capability of the PL, even using standard linear phase retrieval. Since PLs concentrate light on a few pixels and have high optical throughputs, a second stage PL WFS may also use light more efficiently compared to other cascaded WFS designs. Considering the benefits in precision spectroscopy offered by PLs, we believe that a combined astrophotonic spectrograph and WFS warrants further study. We note some considerations for a PL WFS/spectrograph, as well as some improvements that could be made in WFSing strategies.

In the case of linear phase retrieval, our results suggest that the primary benefit of wavelength dispersion is to improve the spatial resolution of the correction by increasing the number of sensed aberration modes. This result may be somewhat surprising because the space of focal-plane electric fields that couple into the lantern depends only weakly on wavelength, at least over wavelength ranges where the number of guided modes supported at the PL entrance remains constant. Therefore, we might assume that the space of sensed phase aberrations might have a similarly weak dependence on wavelength. However, this assumption does not hold, primarily because the vector space of focal-plane electric fields is complex and the vector space of phase aberrations is real. For a concrete counterexample, consider a two-mode wavefront sensor that only accepts light in the two orthogonal LP_{11} fiber modes, which we denote LP_{11}a and LP_{11}b. When individually backpropagated to the pupil plane, the phase sampled by each takes the form of a step function that divides the pupil either horizontally or vertically. But the complex linear combination $\mathrm{LP}_{11}\mathrm{a} + i\mathrm{LP}_{11}\mathrm{b}$ forms a charge 1 orbital angular momentum mode, which samples a phase vortex. So, the space of sensed phase aberrations depends not only on the shape of the electric field basis vectors (LP_{11}a/b), but also the particular linear combination(s) sampled by the wavefront sensor. In the case of the wavelength-dispersed N-port PL, different linear combinations are sampled at different wavelengths due to the dispersion of modes as they copropagate through the waveguide. Put another way: at a single, fixed wavelength, a PL senses a small subset of wavefront aberrations (both phase and amplitude) roughly corresponding to a projection of the global wavefront aberration space onto the "photonic lantern principal modes" [KFL24]. As the wavelength is changed, this modal basis rotates to reveal different projections of the wavefront aberration space [LFX22], sweeping across the space like a lighthouse beam.

We also applied the dispersed PL WFS on-sky to drive a second stage AO system, yielding a $\sim$17% increase in H-band Strehl ratio over primary correction provided by a state-of-the-art AO system on an 8-m telescope. There are two practical limitations. First, there is a detector tension between spectroscopy, which favors slower readout, and wavefront sensing, which favors faster readout. For applications such as NCPA correction where errors evolve slowly, this tension may be sidestepped, but for a general dispersed WFS this tension motivates high speed detectors with low read noise (e.g. FLI C-RED 1) and brighter targets. Another solution to the tension is to divide the outputs of the PL, dispersing some at low spectral resolution for fast WFSing and dedicating others for long-exposure spectroscopy.

The second limitation is the dynamic range of the PL WFS, which in our testing could only provide stable correction when the seeing was good ($\sim$ 0.5” or less) and the residual RMS WFE of the primary AO was $\lesssim$ 150–200 nm. While less of an issue for the correction of smaller-amplitude NCPAs, low-wind effect errors can exceed this threshhold, motivating the development of nonlinear phase retrieval methods for wavelength-dispersed, few-moded sensors. Combined with nonlinear methods, wavelength dispersion should lead to an improvement of dynamic range, in the same way that multi-wavelength measurements of a 2π-wrapped wavefront enable phase unwrapping. These nonlinear methods might use neural networks as in [NWB20] or interpolating methods [LF24]. Still, the wavefront sensing benefits offered by spectral dispersion even in the linear regime suggest that dispersed PLs may be well-suited to simultaneously fulfill sensing and spectroscopy functions in future instruments.

8.4 Methods

8.4.1 Photonic lantern

Both the 19-port and 3-port PLs used in this work were manufactured by the Sydney Astrophotonics Instrumentation Lab (SAIL). The 19-port lantern was formed from a custom multicore fiber with 19 single-mode channels (hexagonal array, 6.5 μm core diameter, 60 μm core spacing, numerical aperture 0.14); this fiber was inserted into a lower-index fluorine glass capillary, and then one end was heated and drawn to form a tapered structure with an entrance diameter of 56 μm. According to these parameters, the outputs of this lantern are single-moded for $\lambda > 1200$ nm. Figure 8.4.1 gives microscope images for both endfaces of the lantern, as well as a picture of the lantern as mounted within SCExAO. To make the lantern cores more visible, they were illuminated with visible light during microscope imaging. This PL has a maximum measured coupling efficiency of 80%, which was obtained at $\lambda = 1550$ nm.

The 3-port lantern was formed similarly by inserting 3 single-mode fibers (Corning SMF-28e) into a lower-index capillary and then heating and drawing one end to an entrance diameter 17.8 μm; the diameter was optimized for 3-moded performance over H band, while the outputs are single-moded for $\lambda > 1260$ nm. Because this lantern was formed from individual fibers, and did not use a preform or "skeleton" to arrange the fibers and fill empty space in the fiber stack, the cladding of this lantern is somewhat triangular. Figure 7.3.1 gives microscope images for the endfaces of both lanterns, as well as a picture of the lanterns as mounted within SCExAO. To make the cores of the lanterns more visible, they were illuminated with visible light during microscope imaging.

8.4.2 SCExAO testbed

All experiments in this paper were performed on the infrared bench of the SCExAO testbed[JMG15]. The top panel of Figure 7.3.2 gives a simplified overview of the infrared bench, showing only optical components immediately relevant to our experiments. Light is emitted from a supercontinuum laser source (NKT superK) which is collimated by an off-axis parabolic mirror. The collimated beam reflects off SCExAO's 50×50 actuator MEMS deformable mirror, and then is optionally apodized to a near-Gaussian beam profile using a pair of beam-shaping lenses. These lenses were used for daytime tests because some of the downstream optics were designed for the smaller, shaped beam; for later on-sky testing, we tweaked the alignment of the setup to omit these lenses and simplify the beam path. The lenses do not appear to have a strong impact on the WFSing ability of our system. From here, the beam is intercepted by a 90:10 splitter which directs the majority of the light to the injection platform, a four-axis

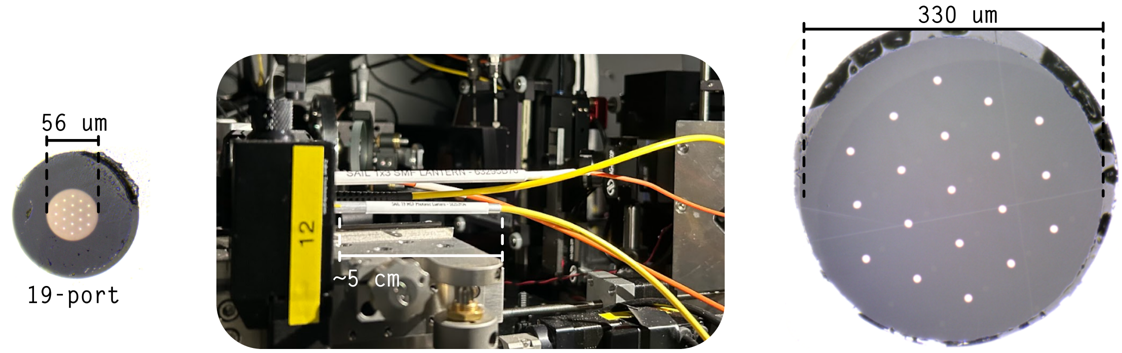

Few-mode fiber end

Side view

Multicore fiber end (19-port)

Figure 8.4.1: The photonic lanterns currently mounted in the infrared bench of SCExAO. Shown are microscope images of the smaller few-mode fiber and larger output end of the 19-port PL. The middle picture shows the lanterns as mounted within the bench; also shown in the middle is a 3-port PL, which was not used for the experiments in this paper. To make the lantern cores more visible in microscope images of the end faces, visible light was injected into the PL during imaging. The wavelength of the visible light is short enough that it can be confined even in the tapered-down cores at the small end of the PL (left).

movement stage holding a lens and the 19-port PL. The movement stage controls the focal ratio of the beam and the transverse alignment relative to the lantern (for more details about the injection see [JSG17]). The remaining 10% of light is directed to an internal InGaAs camera (FLI C-RED 2) for PSF monitoring. Finally, the output of the PL is routed to a diffraction-limited spectrograph with $R \approx 700$, 500, and 300 in the astronomical y, J, and H bands, respectively; the spectral traces of the lanterns are imaged either onto an FLI C-RED 1 (320×256 pixels, 24 μm HgCdTe e-APD) or C-RED 2 detector (640x512 pixels, 10 μm pixel InGaAs) — our setup occasionally changed between experiments due to external constraints. All plots in the main manuscript show data taken when our setup used the C-RED 1 detector.

Because the output of the 19-port PL is a multicore fiber, it is dispersed in a TIGER configuration [LBB12]; the narrow core-to-core spacing of this particular lantern makes the 19 spectral traces partially overlap. As far as we can tell, this does not impact wavefront sensing, though it would impact spectroscopic applications. Another complication is that the shortest wavelength sampled by the spectrograph is ∼1100 nm, shorter than the single-mode cutoff wavelength for the PL outputs. Due to the TIGER dispersion and the overlap of the spectral traces, we were unable to remove wavelength information below this cutoff without sacrificing light in the correct wavelength range from other ports, implying that the PL was not acting exactly as a multi-mode to single-mode converter. That said, barring manufacturing defects, we would still expect the PL outputs at these wavelengths to be close to the LP_{01} mode, which matches our observations. Due to the tapered structure of the PL, any wavefront propagating through the PL will initially couple into cores that are small enough to be single-moded. Additionally, since propagation through a PL is often assumed to be adiabatic, any light initially in the fundamental mode of a core will remain in the fundamental mode. To make this argument more rigorous, we use the `cbeam` [Lin24] package to compute the coupling coefficients of a tapered optical fiber whose parameters match a single tapered core of our lantern, in the region where the cores become multi-moded. Under the simplifying assumption that our lantern has a linear taper profile, we find that the coupling coefficient between the LP_{01} mode and the LP_{11} mode group is of order 10^{-4} cm^{-1}. A comparison of the corresponding length scale with that of our PL implies that cross-coupling is negligible.

In comparison, the 3-port PL has 3 SMF outputs which are mounted on a v-groove array before dispersion. As a result, the spectra of this device are well-separated.

8.4.3 Strehl calculation

Strehl is estimated from PSFs recorded by SCExAO's PSF monitoring camera, which is operated simultaneously with the PL wavefront control loop. We perform a basic Strehl extraction. Given a time series of aberrated PSFs, we first compute the average of the series and locate the point of maximum flux. We then normalize each PSF and compare the flux at this point to the maximum flux of a normalized and "unaberrated" PSF, measured experimentally during the day using SCExAO's calibration source. Of course, due to small optomechanical drifts our calibration PSF is not completely unaberrated, slightly biasing our extracted Strehl ratios upwards. Given the performance of SCExAO, we expect this bias to be $1-2\%$ at most, with negligible impact on relative comparisons between Strehl ratios. Such comparisons are more relevant to gauge the on-sky performance of the PL WFS.

8.4.4 Software

We use the Compute And Control for Adaptive Optics (CACAO; [GSG18]) package to calibrate the wavelength-dispersed PL WFS and run the wavefront correction loop. An overview of the software is given in the bottom panel of Figure 7.3.2. We process the WFS image in a number of ways before passing it to the control loop, involving cropping the image around the lit area of the detector and the usual steps of dark subtraction, masking bad pixels, image normalization, and WFS reference image subtraction. When working with the 3-port PL, we additionally perform a basic spectral extraction by summing across each spectral trace. In some tests, we used a tunable narrowband filter to determine a rudimentary wavelength solution for the spectral extraction, though this is not strictly necessary for phase retrieval. For closed-loop operation, we use a linear reconstructor and a leaky integrator controller. For closed-loop operation, we use a linear reconstructor and a leaky integrator controller, detailed in the next subsection.

8.4.5 WFS calibration

We calibrate and operate dispersed PL WFS in the standard way for astronomical AO, which we review below. First, we measure the sensor response to an incomplete modal basis of pupil-plane phase aberrations — our "calibration basis". Common choices include the Zernike, Fourier, and zonal (DM actuator) modes. We column-stack the measured response to each of the N calibration modes into an $M \times N$ response matrix, denoted A, where M

is the total number of pixels in the WFS intensity image. Next, we apply a singular value decomposition to A. When computing the control matrix (pseudoinverse of A) we choose to only keep singular values that were at least on-tenth the largest value. Finally, we drive the SCExAO DM using a leaky integrator controller: the mirror control signal ϕ at step $n+1$ is set as

$$\phi_{n+1} = l\phi_n - gA^+ \left(\mathbf{I}_n - \mathbf{I}_0\right). \tag{8.1}$$

Here, $\mathbf{I}$ is the intensity image of the WFS, $\mathbf{I}_0$ is the reference intensity, g is the gain, and l is the leak. The same loop parameters are used for all modes.

8.5 Further topics: nonlinear retrieval and impact of bandwidth

8.5.1 Open-loop nonlinear phase retrieval

In comparison with previous studies of non-dispersed PL sensors, we have found that wavelength dispersion combined with linear phase retrieval can increase the number of sensed aberration modes. We also anticipate that wavelength dispersion may be able to improve the dynamic range of PL sensors; however, this benefit will likely only be accessible through nonlinear phase retrieval techniques. To explore this idea, we test open-loop phase retrieval using experimental data taken from the 3-port PL on SCExAO. The particular method we have chosen uses radial basis function interpolation to model the intensity nonlinearity of the PL, and combines this approximate model with a nonlinear least-squares solver[LF24]. We chose this method because it is simple to implement and is generalizable to any few-moded wavefront sensor.

Our experiment is as follows. First, we applied a series of 10,000 phase aberrations composed of random linear combinations of the first 5 non-piston Zernike modes to the SCExAO DM. The maximum amount of RMS WFE applied was 1 radian at 1550 nm, and the average amount of WFE was 0.5 radians. For each random phase, we logged the three PL

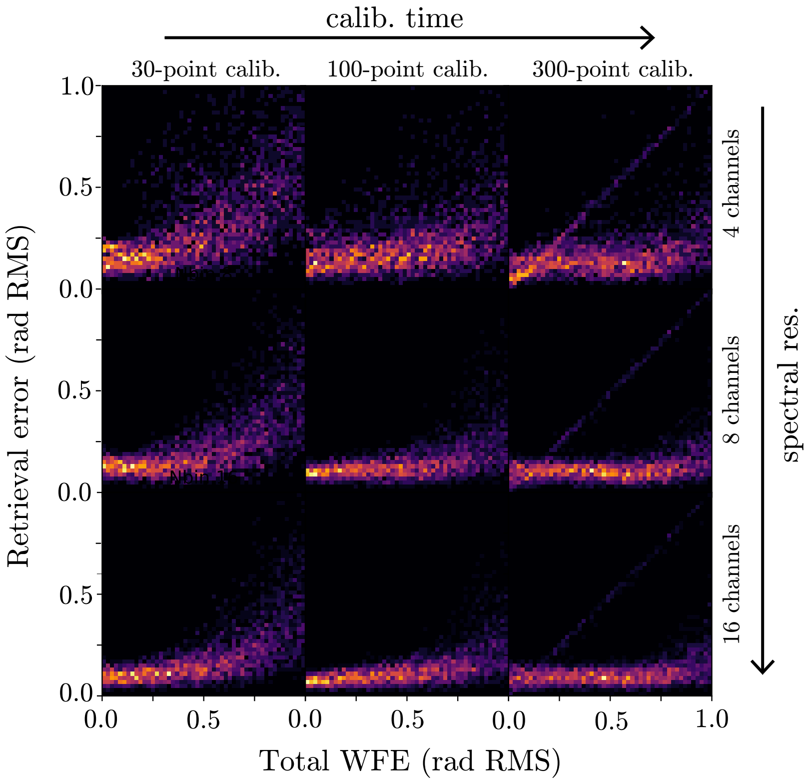

Figure 8.5.1: 2D histograms showing phase retrieval error against total WFE, using a non-linear method based on radial basis function interpolation. Different heatmaps show how the retrieval error depends on both the calibration time and effective spectral resolution of the sensor.

spectra. This calibration step took around 1 hour. Offline, we reduced the dimensionality of the sensor output by binning it down into a few wavelength channels and then used a subset of the binned data to calibrate a radial basis function model. We then combined this model with a Levenburg-Marquardt solver to retrieve phase from the remaining data, and compared our results with the true phase that was originally applied to the DM. To visualize data, we produced 2D histograms which plot the accuracy of the phase retrieval against the total WFE of the aberration. A functional wavefront sensor will give points located below the diagonal of this plot, while a perfect wavefront sensor will place all points along the line $y = 0$. Figure 8.5.1 gives multiple such histograms for nonlinear models calibrated with different amounts of data, binned down to different effective spectral resolutions. As the effective spectral resolution increases, the phase retrieval accuracy also improves, though even with only 4 spectral channels (which should hardly be called wavelength dispersion) the method is still able to sense a portion of the aberration. We also note that as the amount of data used in the calibration process increases, the retrieval accuracy first improves, then worsens for a subset of the phase aberration, which fall on the diagonal. We believe that this behavior is a result of instability in the SCExAO bench: as we include more information in the calibration dataset, the model becomes increasingly inconsistent. This result both emphasizes the sensitivity of the dispersed PL WFS, as well as the need for stabilization of the optical setup during sensor calibration. Nevertheless, we are encouraged by the fact that a 30-probe calibration is enough to produce a usable nonlinear model. By optimizing the placement of the calibration points in phase space[RA10], it may be possible to improve the retrieval accuracy even further, without incurring much overhead in the calibration process. This method should also be fast enough to run at 1 kHz with some optimizations; we leave integration with a closed-loop controller for future work.

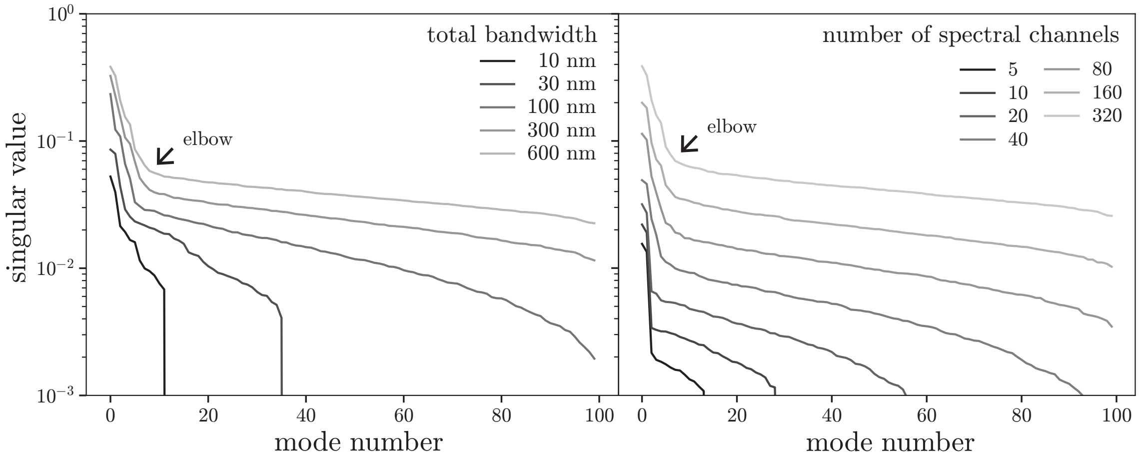

Figure 8.5.2: Left: singular value spectrum of the dispersed 3-port PL WFS for different spectral bandwidths. As the bandwidth increases, the entire spectrum lifts up, indicating greater sensitivity. The “elbow” also shifts right, indicating an increase in the number of sensed modes. Right: similar to the left, except that the entire spectral response of the sensor response is binned down to different effective spectral resolutions, instead of being cropped. The singular values increase as the number of spectral channels increases. At 80 channels and beyond, the singular value spectra are similar in shape.

8.5.2 Effect of spectral bandwidth and resolution

By cropping an experimentally measured response matrix of the dispersed 3-port PL WFS along the wavelength axis, we investigated the effect that bandwidth has on the number of control modes. Our results are shown in the left panel of Figure 8.5.2, which gives the singular value spectrum of the 3-port PL WFS at 5 effective bandwidths ranging from 10 to 600 nm. As the bandwidth increases, so to do the singular values, indicating an improvement in sensitivity. We also note that the "elbow" of the spectra moves out from around mode 3 to around mode 10, indicating a corresponding increase in spatial resolution; the location of the elbow at the largest bandwidth roughly aligns with the number of control modes in our closed-loop test from §8.2.2. The right panel of Figure 8.5.2 presents a similar analysis where we bin down the response matrix over the wavelength axis to investigate the effect of different spectral resolutions. An increase in the number of spectral channels yields a similar overall increase in sensitivity. A similar study of the 19-port PL was not possible due to the overlapping spectral traces.

Appendix - Chapter 8

8.A Reference spectrum compensation

If the calibration of the dispersed WFS is dependent on the reference spectrum used for calibration, then such a sensor would require recalibration every time the calibrator source spectrum changes. We consider how changes in source spectrum may be accounted for in the linear model, in the case of aberrations which are achromatic in terms of optical path difference (OPD). First, note that the power $\mathbf{p}$ measured through the WFS at some phase aberration $\mathbf{r}$ will be proportional to the source spectrum $s(\lambda)$:

$$p_j(\mathbf{r},\lambda) = f_j(\mathbf{r},\lambda)s(\lambda). \tag{8.2}$$

Here, f_j is a scalar-valued function representing the wavelength-dependent throughput of sensor output j, which is intrinsic to the WFS. Next, discretize over wavelength, yielding

$$\begin{aligned} p_{ij} - p_{ij,0} &= \sum_k A_{ijk} x_k \\ p_{ij,0} &= f_j(\mathbf{r}, \lambda_i) s(\lambda_i) \\ p_{ij} &= f_j(\mathbf{r} + \mathbf{x}, \lambda_i) s(\lambda_i) \end{aligned} \tag{8.3}$$

where i iterates over wavelengths and j iterates over output ports; the matrices A_{1jk}, A_{2jk}... are the linear response matrices in each spectral channel. The reference intensity measured using the calibration source is $\mathbf{p}_{ij,0}$, and $\mathbf{x}_k$ now represents OPD. The sensor "slopes" in A_{ijk} are linearly proportional to the absolute power in spectral channel i. Thus, during the calibration process we should factor out the source spectrum, defining the normalized response matrix A' as

$$\frac{p_{ij}}{p_{ij,0}} - 1 = \sum_k \frac{1}{p_{ij,0}} A_{ijk} x_k \equiv \sum_k A'_{ijk} x_k. \tag{8.4}$$

Wavelengths λ_i for which $p_{ij,0}$ are small should be discarded, since they will amplify noise; as a corollary, it may be preferable to choose a calibration source with the flattest possible spectrum over the sensor's wavelength band. During operation, the sensor may see a different reference spectral response, $p_{ij,\star}$. In this case, the control law is obtained by inverting

$$p_{ij} - p_{ij,\star} = \sum_k p_{ij,\star} A'_{ijk} x_k. \tag{8.5}$$

WFS image normalization can be accounted for by a rescaling of $p_{ij,\star}$. The above implies that the control matrix should be recomputed whenever the source spectrum changes, corresponding to the new spectral reference. However, the spectro-normalized response matrix A' does not need to be recomputed unless the instrument is unstable. We have confirmed that this method works by applying different passband filters and closing the loop using the

3-port PL on SCExAO, with the same normalized response matrix. In practice, it might be desirable to live-update both the reference response $p_{ij,\star}$ and the control matrix. Alternatively, one may choose to use the spectrally normalized signal $p_{ij}/p_{ij,\star}$ in conjunction with $(A')^+$ as the control matrix; however, since spectral normalization will re-scale noise, the predicted phase will no longer be the maximum likelihood estimate, as demonstrated in Chapter §15. This is especially problematic if there regions where $p_{ij,\star}$ is near 0 (e.g. atmospheric absorption features). We leave a more rigorous performance comparison between different spectral compensation methods for future work.

If the calibration spectrum is somehow similar to the target guide star spectrum, then we might approximately solve equation 8.5 as

$$x_k \approx \sum_{ij} \frac{p_{ij,0}}{p_{ij,\star}} A^{+}_{ijk} \left(p_{ij} - p_{ij,\star}\right) \tag{8.6}$$

which is reminiscent of how the sensitivity loss of the pyramid WFS in the presence of small phase offsets is compensated ("optical gain compensation", e.g. [CFJ20]). Therefore, this effect might be called "spectral gain compensation", though as a name, "reference spectrum compensation" is probably more informative. Whereas optical gain arises from a phase offset combined with WFS nonlinearity [DGR18], "spectral gain" arises from a spectral offset coupled with the *linear* dependence of the WFS slopes on the source spectrum.

REFERENCES

[BGR22] Bertrou-Cantou, A., Gendron, E., Rousset, G., Deo, V., Ferreira, F., Sevin, A., and Vidal, F. "Confusion in differential piston measurement with the pyramid wavefront sensor." *A & A*, **658**:A49, 2022.

[CFJ20] Chambouleyron, V., Fauvarque, O., Janin-Potiron, P., Correia, C., Sauvage, J-F., Schwartz, N., Neichel, B., and Fusco, T. "Pyramid wavefront sensor optical gains compensation using a convolutional model." *A & A*, **644**:A6, 2020.

[DGR18] V. Deo, É. Gendron, G. Rousset, F. Vidal, and T. Buey. "A modal approach to optical gain compensation for the pyramid wavefront sensor." In Laird M. Close, Laura Schreiber, and Dirk Schmidt, editors, *Adaptive Optics Systems VI*, volume 10703, p. 1070320. International Society for Optics and Photonics, SPIE, 2018.

[DKC09] J.-F. Daigle, Y. Kamali, M. Châteauneuf, G. Tremblay, F. Théberge, J. Dubois, G. Roy, and S. L. Chin. "Remote sensing with intense filaments enhanced by adaptive optics." *Applied Physics B*, **97**(3):701–713, Nov 2009.

[GAA22] Olivier Guyon, Kyohoon Ahn, Masayuki Akiyama, Thayne Currie, Vincent Deo, Takshi Hattori, Tomoyuki Kudo, Julen Lozi, Yosuke Minowa, Yoshito Ono, Nour Skaf, Motohide Tamura, and Sebastien Vievard. "High contrast and high angular imaging at Subaru Telescope." In Laura Schreiber, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems VIII*, volume 12185, p. 121856J. International Society for Optics and Photonics, SPIE, 2022.

[GBR10] Galicher, R., Baudoz, P., Rousset, G., Totems, J., and Mas, M. "Self-coherent camera as a focal plane wavefront sensor: simulations." *A & A*, **509**:A31, 2010.

[GPK06] O. Guyon, E. A. Pluzhnik, M. J. Kuchner, B. Collins, and S. T. Ridgway. "Theoretical Limits on Extrasolar Terrestrial Planet Detection with Coronagraphs." *The Astrophysical Journal Supplement Series*, **167**(1):81–99, November 2006.

[GSG18] Olivier Guyon, Arnaud Sevin, Damien Gratadour, Julien Bernard, Hatem Ltaief, Dalal Sukkari, Sylvain Cetre, Nour Skaf, Julien Lozi, Frantz Martinache, Christophe Clergeon, Barnaby Norris, Alison Wong, and Jared Males. "The compute and control for adaptive optics (CACAO) real-time control software package." In Laird M. Close, Laura Schreiber, and Dirk Schmidt, editors, *Adaptive Optics Systems VI*, volume 10703 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 107031E, July 2018.

[GSW08] A.M. Ghez, S. Salim, N. N. Weinberg, J.R. Lu, T. Do, J.K. Dunn, K. Matthews, M.R. Morris, S. Yelda, E. E. Becklin, T. Kremenek, M. Milosavljevic, and

J. Naiman. “Measuring Distance and Properties of the Milky Way’s Central Supermassive Black Hole with Stellar Orbits.” *ApJ*, **689**(2):1044–1062, December 2008.

[Guy18] Olivier Guyon. “Extreme Adaptive Optics.” *ARAA*, **56**:315–355, September 2018.

[JMG15] N. Jovanovic, F. Martinache, O. Guyon, C. Clergeon, G. Singh, T. Kudo, V. Garrel, K. Newman, D. Doughty, J. Lozi, J. Males, Y. Minowa, Y. Hayano, N. Takato, J. Morino, J. Kuhn, E. Serabyn, B. Norris, P. Tuthill, G. Schworer, P. Stewart, L. Close, E. Huby, G. Perrin, S. Lacour, L. Gauchet, S. Vievard, N. Murakami, F. Oshiyama, N. Baba, T. Matsuo, J. Nishikawa, M. Tamura, O. Lai, F. Marchis, G. Duchene, T. Kotani, and J. Woillez. “The Subaru Coronagraphic Extreme Adaptive Optics System: Enabling High-Contrast Imaging on Solar-System Scales.” *PASP*, **127**(955):890, September 2015.

[JSG17] N. Jovanovic, C. Schwab, O. Guyon, J. Lozi, N. Cvetojevic, F. Martinache, S. Leon-Saval, B. Norris, S. Gross, D. Doughty, T. Currie, and N. Takato. “Efficient injection from large telescopes into single-mode fibres: Enabling the era of ultra-precision astronomy.” *AAP*, **604**:A122, August 2017.

[KFL24] Yoo Jung Kim, Michael P. Fitzgerald, Jonathan Lin, Steph Sallum, Yinzi Xin, Nemanja Jovanovic, and Sergio Leon-Saval. “Coherent Imaging with Photonic Lanterns.” *The Astrophysical Journal*, **964**(2):113, mar 2024.

[KFLed] Yoo Jung Kim, Michael P. Fitzgerald, Jonathan Lin, Daniel Levinstein Yinzi Xin, Nemanja Jovanovic Steph Sallum, and Sergio Leon-Saval. “On the Potential of Spectroastrometry with Photonic Lanterns.” *JATIS*, accepted.

[LAB24] Julien Lozi, Kyohoon Ahn, Hannah Blue, Alicia Chun, Christophe Clergeon, Vincent Deo, Olivier Guyon, Takashi Hattori, Yosuke Minowa, Shogo Nishiyama, Yoshito Ono, Shin Oya, Yuhei Takagi, Sébastien Vievard, and Maria Vincent. “AO3k at Subaru: first on-sky results of the facility extreme-AO.” In Kathryn J. Jackson, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems IX*, volume 13097, p. 1309703. International Society for Optics and Photonics, SPIE, 2024.

[LBB05] S. G. Leon-Saval, T. A. Birks, J. Bland-Hawthorn, and M. Englund. “Multimode fiber devices with single-mode performance.” *Opt. Lett.*, **30**(19):2545–2547, Oct 2005.

[LBB12] Sergio G. Leon-Saval, Christopher H. Betters, and Joss Bland-Hawthorn. “The Photonic TIGER: a multicore fiber-fed spectrograph.” In Ramón Navarro, Colin R. Cunningham, and Eric Prieto, editors, *Modern Technologies in Space- and Ground-based Telescopes and Instrumentation II.* SPIE, September 2012.

[LF24] Jonathan Lin and Michael P. Fitzgerald. "Nonlinear techniques for few-mode wavefront sensors." *Appl. Opt.*, **63**(34):8748–8759, Dec 2024.

[LFX22] Jonathan Lin, Michael P. Fitzgerald, Yinzi Xin, Olivier Guyon, Sergio Leon-Saval, Barnaby Norris, and Nemanja Jovanovic. "Focal-plane wavefront sensing with photonic lanterns: theoretical framework." *J. Opt. Soc. Am. B*, **39**(10):2643–2656, Oct 2022.

[LFX23] Jonathan W. Lin, Michael P. Fitzgerald, Yinzi Xin, Yoo Jung Kim, Olivier Guyon, Barnaby Norris, Christopher Betters, Sergio Leon-Saval, Kyohoon Ahn, Vincent Deo, Julien Lozi, Sébastien Vievard, Daniel Levinstein, Steph Sallum, and Nemanja Jovanovic. "Real-time Experimental Demonstrations of a Photonic Lantern Wave-front Sensor." *The Astrophysical Journal Letters*, **959**(2):L34, dec 2023.

[Lin24] Jonathan Lin. "cbeam." Astrophysics Source Code Library, record ascl:2404.001, April 2024.

[LJF21] Jonathan Lin, Nemanja Jovanovic, and Michael P. Fitzgerald. "Design considerations of photonic lanterns for diffraction-limited spectrometry." *J. Opt. Soc. Am. B*, **38**(7):A51–A63, Jul 2021.

[MLA12] Martinez, P., Loose, C., Aller Carpentier, E., and Kasper, M. "Speckle temporal stability in XAO coronagraphic images." *A & A*, **541**:A136, 2012.

[MZK10] Christian Marois, B. Zuckerman, Quinn M. Konopacky, Bruce Macintosh, and Travis Barman. "Images of a fourth planet orbiting HR 8799." *Nature*, **468**(7327):1080–1083, December 2010.

[Nat21] National Academies of Sciences, Engineering, and Medicine. *Pathways to Discovery in Astronomy and Astrophysics for the 2020s.* The National Academies Press, Washington, DC, 2021.

[NSS12] D. Noordegraaf, P. M. W. Skovgaard, R. H. Sandberg, M. D. Maack, J. Bland-Hawthorn, J. Lawrence, and J. Lægsgaard. "Nineteen-port photonic lantern with multimode delivery fiber." *Optics Letters*, **37**:452, 2012.

[NWB20] Barnaby R. M. Norris, Jin Wei, Christopher H. Betters, Alison Wong, and Sergio G. Leon-Saval. "An all-photonic focal-plane wavefront sensor." *Nature Communications*, **11**(1):5335, Oct 2020.

[RA10] T.C.S. Rendall and C.B. Allen. "Reduced surface point selection options for efficient mesh deformation using radial basis functions." *Journal of Computational Physics*, **229**(8):2810–2820, 2010.

[RCR21] Cristina Rodríguez, Anderson Chen, José A. Rivera, Manuel A. Mohr, Yajie Liang, Ryan G. Natan, Wenzhi Sun, Daniel E. Milkie, Thomas G. Bifano, Xiaoke Chen, and Na Ji. "An adaptive optics module for deep tissue multiphoton imaging in vivo." *Nature Methods*, **18**(10):1259–1264, Oct 2021.

[SFL16] Jean-François Sauvage, Thierry Fusco, Masen Lamb, Julien Girard, Martin Brinkmann, Andres Guesalaga, Peter Wizinowich, Jared O'Neal, Mamadou N'Diaye, Arthur Vigan, David Mouillet, Jean-Luc Beuzit, Markus Kasper, Miska Le Louarn, Julien Milli, Kjetil Dohlen, Benoît Neichel, Pierre Bourget, Pierre Haguenauer, and Dimitri Mawet. "Tackling down the low wind effect on SPHERE instrument." In Enrico Marchetti, Laird M. Close, and Jean-Pierre Véran, editors, *Adaptive Optics Systems V*, volume 9909 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 990916, July 2016.

[SGG22] Skaf, Nour, Guyon, Olivier, Gendron, Éric, Ahn, Kyohoon, Bertrou-Cantou, Arielle, Boccaletti, Anthony, Cranney, Jesse, Currie, Thayne, Deo, Vincent, Edwards, Billy, Ferreira, Florian, Gratadour, Damien, Lozi, Julien, Norris, Barnaby, Sevin, Arnaud, Vidal, Fabrice, and Vievard, Sébastien. "On-sky validation of image-based adaptive optics wavefront sensor referencing." *A & A*, **659**:A170, 2022.

[TO10] W. A. Traub and B. R. Oppenheimer. "Direct Imaging of Exoplanets." In S. Seager, editor, *Exoplanets*, pp. 111–156. 2010.

[Tys96] Robert K. Tyson. "Adaptive optics and ground-to-space laser communications." *Appl. Opt.*, **35**(19):3640–3646, Jul 1996.

[VLL24] Vievard, S., Lallement, M., Leon-Saval, S., Guyon, O., Jovanovic, N., Huby, E., Lacour, S., Lozi, J., Deo, V., Ahn, K., Lucas, M., Sallum, S., Norris, B., Betters, C., Amezcua-Correa, R., Yerolatsitis, S., Fitzgerald, M. P., Lin, J., Kim, Y. J., Gatkine, P., Kotani, T., Tamura, M., Currie, T., Kenchington, H-D., Martin, G., and Perrin, G. "Spectroscopy using a visible photonic lantern at the Subaru Telescope: Laboratory characterization and the first on-sky demonstration on Ikiiki (α Leo) and 'Aua (α Ori)." *A & A*, **691**:A140, 2024.

[XJR22] Yinzi Xin, Nemanja Jovanovic, Garreth Ruane, Dimitri Mawet, Michael P. Fitzgerald, Daniel Echeverri, Jonathan Lin, Sergio Leon-Saval, Pradip Gatkine, Yoo Jung Kim, Barnaby Norris, and Steph Sallum. "Efficient Detection and Characterization of Exoplanets within the Diffraction Limit: Nulling with a Mode-selective Photonic Lantern." *The Astrophysical Journal*, **938**(2):140, oct 2022.

CHAPTER 9

Ongoing work and future prospects

The previous chapters summarize the journal-published output of my dissertation. In this chapter I will summarize some of the ongoing and potential future developments in astrophotonics that are related to my dissertation work and which I plan to support.

9.1 Current areas of research

9.1.1 Correction of LWE with dispersed PLs in visible light

One of the limiting sources of wavefront error that remain downstream of Subaru telescope's AO3k correction is the low-wind effect (LWE). Simultaneously, The SCExAO testbed at Subaru contains the FIRST spectrograph: a visible-light spectropolarimeter with $R \sim 3000$, which can be fed with a 19-port PL to give a total of 38 spectral traces. Building off of developments from Chapter 8, one idea is to use FIRST as a wavefront sensor to drive the correction of LWE aberrations in visible light, during NIR science observations. Working with SCExAO team, we demonstrated off-sky closed-loop control of LWE errors using this technique. However, because phase errors are larger in the visible compared to the NIR, an on-sky demonstration is challenging. We believe the next course of action is to apply a multiwavelength phase-unwrapping technique to the dispersed output of the PL, to increase the dynamic range of this WFS and close the loop on-sky.

9.1.2 PL spectroastrometry

Concurrent with the development of the PL WFS, there is an ongoing push to develop PL spectroastrometry. For context, the technique of astrometry involves precisely determining the spatial locations of sources on the sky, typically through centroid measurements of focal plane images. This technique, when applied to stars, can be used to determine parallaxes (and therefore distances to stars in the galactic neighborhood, as done by Hipparcos and GAIA), as well as proper motions. Importantly, because the primary signal produced by the astrometric method is a centroid position on the detector, astrometric precision is *distinct* from the telescope's classical resolution limit. With sufficient centroid precision, astrometry can be and has been used to identify the presence of companions (e.g. exoplanets) in systems with angular scales below the diffraction limit; in this way, astrometry is a "super-resolution" method.

Spectroastrometry (SA, [Bai98, WG08]) is the astrometry of spectrally dispersed images (both those produced by long-slit and integral field spectroscopy). This class of techniques can be used to probe the spatial structure of an extended but unresolved or barely resolved source, where the spectral characteristics of the source vary across the sky. In long-slit SA, spectroastrometric signals manifest as physical shifts in the centroid of a spectral trace ('wiggles' or 'bumps'), occurring perpendicular to the direction of spectral dispersion.

Finally, the more general technique of spectro-imaging (SI) refers to a class of techniques which try to reconstruct the spatial structure of an unresolved source using its spectrally dispersed image(s). These techniques often use SA in combination with forward models, and/or tackle the problem of ill-posed inversion from the spectral image back to the sub-diffraction-limit source structure via some sort of regularization (e.g. maximum entropy, sparsity).

PLs provide a novel platform for SA and SI. Instead of conventional long-slit spectroscopy, the SMF outputs of a PL can individually dispersed. This has several benefits. For one, PL SA can provide more spatial information regarding the unresolved source: the spectroastrometric signal, which now takes the form of wavelength-dependent intensity differences between each of the PL's individual spectral traces, is no longer limited to a tilt along one axis. This arrangement also is more efficient with telescope light, and can provide more precise spectra, due to the spatial filtering effect of SMFs.

PL SA has been studied in simulation [KFL24b], and has been combined with simulated observations of protoplanetary systems [LSK23]. Ongoing efforts led by UCLA and SCExAO are pursuing an on-sky demonstration of PL SA: for instance, using the technique on $H\alpha$ emission from rapidly rotating stars with compact circumstellar disks, such as βCMi.

9.1.3 Interferometric imaging with astrophotonics

Another idea is to apply a PL combined with a beam recombiner circuit to perform interferometric imaging at potentially sub-diffraction-limited scales. This technique was developed in [KFL24a], which showed in simulation that imaging with a PL and a passive PIC could outperform more standard interferometric techniques such as sparse aperture masking. Astrophotonic interferometric imaging is also employed by the VLTI/GRAVITY instrument, which uses a 4-beam photonic circuit to simultaneously interfere four separate telescopes for long-baseline optical interferometry [Cit].

9.1.4 PL nulling

In an extension to the vortex fiber nulling (VFN) technique [REJ19], a mode-selective PL can be applied to null on-axis starlight for high-contrast imaging applications; this technique is called PL nulling (PLN; [XJR22, XEJ24]). In VFN, a phase corkscrew is applied to incoming wavefronts, preventing on-axis starlight from coupling into the symmetric (and Gaussian-like) fundamental mode of a single-mode fiber. In PLN, a similar effect occurs because a mode-selective PL accepts light in either symmetric modes (e.g. the fundamental mode) or anti-symmetric modes (e.g LP_{11}). An incoming wavefront coming from an on-axis star, assumed to be symmetric, will not couple into the antisymmetric modes of the mode-selective PL: thus, a subset of the PL's outputs are nulled. An off-wavefront, say from an exoplanet, may couple into these nulled outputs. Current efforts led by Caltech and SCExAO are working towards an on-sky demonstration of this technique on a low-contrast binary.

9.2 Looking forward

9.2.1 Photonic spectrometers

Arrayed waveguide gratings, multimode interferometers, and other dispersive photonic components have the potential to reduce the size and complexity of spectrometers which previously relied on bulk optics. Even further reductions in complexity can be had in cross-correlation spectrometers, by using components such as microring resonators [CSD20] or fiber Bragg gratings to detect the presence of absorption or emission features without any spectral dispersion. These size and complexity reductions yield increased optomechanical stability, as well as potential cost and weight savings. For dispersive photonic spectrometers, current challenges include improving light efficiency and spectral resolution, as well as on-sky demonstration.

9.2.2 Self-calibration

As the techniques of PL spectroscopy, SA, and nulling become demonstrated in their respective scientific applications (e.g. the study of circumstellar environments, or exoplanet imaging and spectroscopy), the next step will be to combine these techniques with simultaneous wavefront sensing. An initial step in this direction was demonstrated by [GVL24], which used the visible light, PL-fed FIRST spectrograph at SCExAO to simultaneously measure the spectrum of the star Humu (Altair) and the stellar rotation bias. In this initial demonstration, medium-resolution spectroscopy was performed on a quickly rotating star; the next step will be to extend these results to higher spectral resolutions and smaller rotational biases.

Another possibility for the techniques of PL SA and nulling is to use the PL simultaneously as a wavefront sensor. In the context of PL SA, simultaneous wavefront sensing and control should improve the precision of the retrieved spectra and therefore the spectroastrometric sensitivity, in particular by mitigating systematic errors that may arise from quasi-static aberrations. For PL nulling, active wavefront control will darken nulls and improve achievable contrast. At Subaru telescope, integrations of PL SA and wavefront sensing could be performed on the $R \sim 60,000$ RHEA spectrograph ($\lambda = 600 - 800$ nm; [RIJ16]) and the upcoming $R \sim 4000$ exo-NINJA spectrograph (JHK bands; [ELG24]). A combined PL nulling & wavefront sensing technique would require a novel "partially mode-selective" photonic lantern in which certain outputs are designated for nulling and others for wavefront sensing.

9.2.3 Active photonic circuits

No coronagraph has yet demonstrated starlight suppression at the 10^{10} level demanded by the high-contrast imaging of exo-Earths; however, astrophotonic coronagraphs-on-a-chip are one potential way forward. In particular, [GPK06] and [BSJ21] derived the fundamental performance limits for coronagraphs, and showed that such devices could in principle be realized with only systems of beamsplitters and phase shifters. While a bulk-optics implementation would be complicated to build and hard to stabilize, astrophotonic circuits provide a compact and stable platform for optimal (or at least near-optimal) coronagraphy. This concept is being pursued by the AstroPIC pathfinder instrument [SBF24] which will use grating couplers to feed a telescope pupil into an active Mach-Zehnder mesh, as well as by [MJW24], which uses PLs to inject the focal-plane beam of a telescope into a similar type of circuit. These two projects are being developed in SOI and SiN platforms, respectively. Simultaneously, a UCLA- and Caltech-led effort is pursuing the development of photonic nulling circuits using directional couplers (e.g. nulling tricouplers) and multimode interferometers, in a silica-on-silicon platform.

In tandem with the above developments, efforts should also be made to develop a technique for simultaneous on-chip wavefront sensing. As established before, wavefront errors such as NCPAs are one of the primary limitations in deepening coronagraphic contrast. A hybrid coronagraph-wavefront sensor will not suffer from NCPAs. Furthermore, a PIC-based WFS has the potential to push boundaries in wavefront control in the following ways:

1. Similar to a coronagraph, any WFS may be treated as a complex-valued matrix which maps an input electric field to an output electric field. A PIC has *complete* control over the unitary transformation it applies, given enough beam recombiners. In other words, a PIC can be actively tuned in WFS properties such as linearity, sensitivity, and dynamic range, just by modulating phase shifters. There is no other WFS with this flexibility.

2. The thermo-optic phase shifters in a PIC can hit 100 kHz actuation speeds, opening the door to novel phase retrieval algorithms, for instance exploiting phase diversity, where the wavefront is modulated in real time to better constrain its phase.

3. By combining **1** & **2**, a PIC WFS can actively change its control space (the vector space of sensed phase aberration modes). This enables the same PIC device to serve multiple WFSing functions, for instance by correcting low-spatial-frequency NCPAs and mid-frequency petal modes in alternating fashion. An extension is “WFS diversity”, where the WFS itself is modulated in real time to better constrain the wavefront phase.

I hope to develop this technique further, in the near future.

REFERENCES

[Bai98] Jeremy A. Bailey. "Spectroastrometry: a new approach to astronomy on small spatial scales." In Sandro D'Odorico, editor, *Optical Astronomical Instrumentation*, volume 3355, pp. 932 – 939. International Society for Optics and Photonics, SPIE, 1998.

[BSJ21] Ruslan Belikov, Dan Sirbu, Jeffrey B. Jewell, Olivier Guyon, and Christopher C. Stark. "Theoretical performance limits for coronagraphs on obstructed and unobstructed apertures: how much can current designs be improved?" In Stuart B. Shaklan and Garreth J. Ruane, editors, *Techniques and Instrumentation for Detection of Exoplanets X*, volume 11823 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 118230W, September 2021.

[Cit] Paris Cité. "15— Future of Photonic Beam Combining Technologies for Interferometers." *2023 Astrophotonics Roadmap: pathways to realizing multi-functional integrated astrophotonic instruments.*

[CSD20] Ross Cheriton, Suresh Sivanandam, Adam Densmore, Ernst De Mooij, Daniele Melati, Mohsen Kamandar Dezfouli, Pavel Cheben, Danxia Xu, Jens H. Schmid, Jean Lapointe, Rubin Ma, Shurui Wang, Luc Simard, and Siegfried Janz. "Spectrum-free integrated photonic remote molecular identification and sensing." *Optics Express*, **28**(19):27951, September 2020.

[ELG24] Mona El Morsy, Julien Lozi, Olivier Guyon, Thayne Currie, Sébastien Vievard, Julia Bryant, Chihiro Tokoku, Vincent Deo, Kyohoon Ahn, Fred Crous, Adeline Haobing Wang, and Zinat Mahol Sathi. "Exo-NINJA at Subaru: fiber-fed spectro-imaging of exoplanets and circumstellar disks at R∼4000." *arXiv e-prints*, p. arXiv:2407.16746, July 2024.

[GPK06] O. Guyon, E. A. Pluzhnik, M. J. Kuchner, B. Collins, and S. T. Ridgway. "Theoretical Limits on Extrasolar Terrestrial Planet Detection with Coronagraphs." *Astrophysical Journal, Supplement*, **167**(1):81–99, November 2006.

[GVL24] Olivier Guyon, Sébastien Vievard, Manon Lallement, Sergio Leon-Saval, Takayuki Kotani, Motohide Tamura, Vincent Deo, Julien Lozi, Rodrigo Amezcua-Correa, Stephanos Yerolatsitis, Sylvestre Lacour, and Elsa Huby. "Extreme precision radial velocity with extreme-AO and photonics: challenges and opportunities." In Kathryn J. Jackson, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems IX*, volume 13097, p. 1309719. International Society for Optics and Photonics, SPIE, 2024.

[KFL24a] Yoo Jung Kim, Michael P. Fitzgerald, Jonathan Lin, Steph Sallum, Yinzi Xin, Nemanja Jovanovic, and Sergio Leon-Saval. "Coherent Imaging with Photonic Lanterns." *Astrophysical Journal*, **964**(2):113, April 2024.

[KFL24b] Yoo Jung Kim, Michael P. Fitzgerald, Jonathan Lin, Yinzi Xin, Daniel Levinstein, Steph Sallum, Nemanja Jovanovic, and Sergio Leon-Saval. "On the potential of spectroastrometry with photonic lanterns." *Journal of Astronomical Telescopes, Instruments, and Systems*, **10**:045001, October 2024.

[LSK23] Daniel M. Levinstein, Stephanie Sallum, Yoo Jung Kim, Jonathan Lin, Julien Lozi, Nemanja Jovanovic, Michael P. Fitzgerald, and Sébastien Vievard. "Spectroastrometry of embedded accreting protoplanets using photonic lanterns." In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 12680 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 126800J, October 2023.

[MJW24] Dylan McKeithen, Jeffrey Jewell, and J. Kent Wallace. "Hardware development towards the demonstration of an ideal photonic coronagraph." In Laura E. Coyle, Shuji Matsuura, and Marshall D. Perrin, editors, *Space Telescopes and Instrumentation 2024: Optical, Infrared, and Millimeter Wave*, volume 13092, p. 130925V. International Society for Optics and Photonics, SPIE, 2024.

[REJ19] Garreth Ruane, Daniel Echeverri, Nemanja Jovanovic, Dimitri Mawet, Eugene Serabyn, J. Kent Wallace, Jason Wang, and Natasha Batalha. "Vortex fiber nulling for exoplanet observations: conceptual design, theoretical performance, and initial scientific yield predictions." In Stuart B. Shaklan, editor, *Techniques and Instrumentation for Detection of Exoplanets IX*, volume 11117, pp. 366 – 381. International Society for Optics and Photonics, SPIE, 2019.

[RIJ16] Adam D. Rains, Michael J. Ireland, Nemanja Jovanovic, Tobias Feger, Joao Bento, Christian Schwab, David W. Coutts, Olivier Guyon, Alexander Arriola, and Simon Gross. "Precision single mode fibre integral field spectroscopy with the RHEA spectrograph." In Christopher J. Evans, Luc Simard, and Hideki Takami, editors, *Ground-based and Airborne Instrumentation for Astronomy VI*, volume 9908 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 990876, August 2016.

[SBF24] Dan Sirbu, Ruslan Belikov, Kevin Fogarty, Carson Valdez, Zhanghao Sun, Annie Kroo, Olav Solgaard, David A. B. Miller, and Olivier Guyon. "AstroPIC: near-infrared photonic integrated circuit coronagraph architecture for the Habitable Worlds Observatory." In Laura E. Coyle, Shuji Matsuura, and Marshall D. Perrin, editors, *Space Telescopes and Instrumentation 2024: Optical, Infrared, and Millimeter Wave*, volume 13092, p. 130921T. International Society for Optics and Photonics, SPIE, 2024.

[WG08] E. Whelan and P. Garcia. *Spectro-astrometry: The Method, its Limitations, and Applications*, pp. 123–149. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[XEJ24] Yinzi Xin, Daniel Echeverri, Nemanja Jovanovic, Dimitri Mawet, Sergio Leon-Saval, Rodrigo Amezcua-Correa, Stephanos Yerolatsitis, Michael P. Fitzgerald, Pradip Gatkine, Yoo Jung Kim, Jonathan Lin, Barnaby Norris, Garreth Ruane, and Steph Sallum. "Laboratory demonstration of a Photonic Lantern Nuller in monochromatic and broadband light." *Journal of Astronomical Telescopes, Instruments, and Systems*, **10**:025001, April 2024.

[XJR22] Yinzi Xin, Nemanja Jovanovic, Garreth Ruane, Dimitri Mawet, Michael P. Fitzgerald, Daniel Echeverri, Jonathan Lin, Sergio Leon-Saval, Pradip Gatkine, Yoo Jung Kim, Barnaby Norris, and Steph Sallum. "Efficient Detection and Characterization of Exoplanets within the Diffraction Limit: Nulling with a Mode-selective Photonic Lantern." *Astrophysical Journals*, **938**(2):140, October 2022.

Part V

Theory, models & methods

Part V collates the various notes taken over the course of my dissertation which support the previous Parts. In §10, I review the approximations and discretizations to Maxwell's equations used in numerical simulation codes for photonic waveguides. In §11, I review coupled-mode theory, which can provide a physical intuition for how light propagates through slowly-varying waveguides such as directional couplers and photonic lanterns. In §12 and §13, I review the theory of atmospheric turbulence and its numerical simulation methods, along with some of the fundamental mathematics. Finally, in §14 and §15, I review topics in adaptive optics, including rejection transfer functions and linear phase retrieval.

NOMENCLATURE

This table lists some of the shared mathematical nomenclature for all of Part V. In general, **bolded** letters denote vector-valued quantities. Chapter-specific symbol tables are also given at the beginnings of Chapters 10 – 12.

Symbol	Description
$\vec{\nabla}$	del operator, $[\partial/\partial x, \partial/\partial y, \partial/\partial z]^T$
$\vec{\nabla} f$	gradient of f, $[\partial f/\partial x, \partial f/\partial y, \partial f/\partial z]^T$
$\vec{\nabla} \cdot \mathbf{u}$	divergence of $\mathbf{u}$
$\vec{\nabla} \times \mathbf{u}$	curl of $\mathbf{u}$
∇^2	Laplacian, $\partial^2/\partial x^2 + \partial^2/\partial y^2 + \partial^2/\partial z^2$
$\nabla^2_\perp$	transverse Laplacian, $\partial^2/\partial x^2 + \partial^2/\partial y^2$
$\vec{\nabla}^2\mathbf{u}$	vector Laplacian of $\mathbf{u}$, $[\nabla^2 u_1, \nabla^2 u_2, \dots]^T$
$\mathbf{u} \cdot \mathbf{v}$	dot product (inner product) of $\mathbf{u}$ and $\mathbf{v}$
$\oint_\Gamma$	contour integration around closed curve Γ
$\iint_\Omega$	surface integration over domain Ω
$[\ ]^T$	transpose of matrix or vector
$\lvert A \rvert$	determinant of matrix A
A^+	pseudoinverse of matrix A
$\lvert \mathbf{u} \rvert$	Euclidean norm of $\mathbf{u}$, also implicitly written as u
$\lvert x \rvert$	absolute value of x
$f \circledast g$	convolution of functions f and g
$\mathcal{F}$	Fourier transform, $\mathcal{F}^{-1}$ is the inverse
$\hat{f}$	the Fourier pair of function f, $\hat{f} = \mathcal{F} f$
$f'(x)$	derivative of $f(x)$
$f\big\rvert_{x_0}$	f evaluated at x_0, $f(x_0)$, often applied to derivative expressions
$[\ ]^*$	complex conjugate

i	imaginary unit
ω	temporal angular frequency
λ	wavelength
$\mathbf{r}$	position vector, $[x, y, z]^T$
$\delta(x)$	Dirac delta function
$\delta(\mathbf{r})$	3D Dirac delta function, $\delta(x)\delta(y)\delta(z)$
$J_\nu(x)$	Bessel function of the 1st kind, order ν

CHAPTER 10

Modeling photonic devices

Contents

Simulation tools for photonic devices are a crucial component in astrophotonics research, allowing for the rapid exploration of novel devices as well as the informed manufacture of new astrophotonic instruments. In this chapter, I review several analytic and numerical techniques, some of which I have developed into software packages. All such packages, including `lightbeam` [Lin20] and `wavesolve`[1], introduced in this chapter, and `cbeam` [Lin24] in the next, are open-source.

The starting point is Maxwell's equations in a dielectric medium. When a dielectric substance is placed in an electric field, the field within the dielectric is reduced due to the partial anti-alignment of dipoles. This interaction between the electric field and the intervening dielectric medium, which occurs due to the medium's polarizability, is the operating principle used by (most) waveguides to manipulate and confine electromagnetic waves.

---

[1] https://github.com/jw-lin/wavesolve

NOMENCLATURE - CHAPTER 10

Symbol	Description
$\tilde{u}$	finite element approx. of scalar field u
$\mathcal{D}$	a differential operator
k	free-space longitudinal (z) angular wavenumber
β	effective longitudinal angular wavenumber, β_i is for eigenmode i
n_0	effective index (scalar), $\beta = kn_0$
$\mathbf{r}$	position vector, $[x, y, z]^T$
$n(\mathbf{r}, \omega)$	refractive index
δ_x, δ_y	central finite difference operators in x, y
N_i	nodal shape functions for finite element analysis
Ω_Δ	region corresponding to a triangular finite element
Γ_Δ	contour corresponding to the boundary of a triangular finite element
x', y'	affine coordinates
$\mathbf{N}_i$	edge shape functions for finite element analysis
x_{ij}	shorthand for $x_i - x_j$
l_{ij}	signed length between points i and j
d	determinant of Jacobian of the affine transformation, $x_{21}y_{31} - x_{31}y_{21}$

10.1 Maxwell to Helmholtz

This section adapts material from [Bur09] as well as Feynman's notes on the subject. Bolded quantities are vector-valued. I begin with Maxwell's equations in free space:

$$\vec{\nabla} \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \tag{10.1}$$

$$\vec{\nabla} \cdot \mathbf{B} = 0 \tag{10.2}$$

$$\vec{\nabla} \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \tag{10.3}$$

$$\vec{\nabla} \times \mathbf{B} = \mu_0 \left(\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right). \tag{10.4}$$

Here, $\mathbf{E}$ is vector-valued the electric field; $\mathbf{B}$ is the magnetic field; $\mathbf{J}$ is the current density; ρ is the charge density; and ϵ_0 and μ_0 are the free-space permittivity and permeability, respectively. However, since we are interested in propagation through dielectric media, not free space, we must modify the first and last equations since they have explicit dependence on charge and current. First, denote the electric displacement field D as

$$\mathbf{D} = \epsilon_0 \epsilon_r \mathbf{E} = \epsilon_0 \mathbf{E} + \mathbf{P} \tag{10.5}$$

where ϵ_r is the relative permittivity of the medium and $\mathbf{P}$ is the polarization field. The permittivity can be tensorial in the case of anisotropic media. Similarly, define the auxiliary magnetic field H such that

$$\begin{aligned} \mathbf{B} &= \mu_0 \mu_r \mathbf{H} = \mu_0 \left(\mathbf{H} + \mathbf{M} \right) \\ \mathbf{H} &= \frac{1}{\mu_0} \mathbf{B} - \mathbf{M} \end{aligned} \tag{10.6}$$

where μ_r is the relative permeability and $\mathbf{M}$ is the magnetization field. Substituting the free charge and current for total charge and current (including induced components) the first of

Maxwell's equations becomes becomes

$$\vec{\nabla} \cdot \mathbf{D} = \rho_f \tag{10.7}$$

and the last becomes

$$\begin{aligned} \vec{\nabla} \times \mathbf{B} &= \mu_0 \left(\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right) \\ &= \mu_0 \left(\mathbf{J}_f + \vec{\nabla} \times \mathbf{M} + \frac{\partial \mathbf{P}}{\partial t} \right) + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \\ \vec{\nabla} \times \mathbf{H} &= \mathbf{J}_f + \frac{\partial \mathbf{D}}{\partial t} \end{aligned} \tag{10.8}$$

where the f subscripts denote free charges and currents (as opposed to bound). Thus, to summarize Maxwell's equations in dielectric matter:

$$\vec{\nabla} \cdot \mathbf{D} = \rho_f \tag{10.9}$$

$$\vec{\nabla} \cdot \mathbf{B} = 0 \tag{10.10}$$

$$\vec{\nabla} \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \tag{10.11}$$

$$\vec{\nabla} \times \mathbf{H} = \mathbf{J}_f + \frac{\partial \mathbf{D}}{\partial t}. \tag{10.12}$$

Now make the following assumptions, which hold for most optical waveguides:

1. There are no free charges: $\rho_f = 0$.
2. There are no free currents: $J_f = 0$.
3. All media are non-magnetic: $M = 0$.
4. All media are dielectrically linear, homogeneous, and isotropic, so that $\mathbf{D} = \epsilon \mathbf{E}$, where the permittivity ϵ is a scalar.

In an attempt to solve Maxwell's equations, we assume time-harmonic solutions of the form

$$\mathbf{E} = \boldsymbol{\mathcal{E}}(\mathbf{r})e^{-i\omega t} \tag{10.13}$$

$$\mathbf{H} = \boldsymbol{\mathcal{H}}(\mathbf{r})e^{-i\omega t} \tag{10.14}$$

where ω is the frequency of the electromagnetic wave and $\mathbf{r}$ is the position vector. Since we have assumed non-magnetic, dielectrically linear, and isotropic media,

$$\begin{aligned} \mathbf{B} &= \mu_0 \mathbf{H} \\ \mathbf{D} &= \epsilon_0 \epsilon_r \mathbf{E} = \epsilon_0 n^2(\mathbf{r}, \omega)\mathbf{E} \end{aligned} \tag{10.15}$$

where $n(\mathbf{r}, \omega)$ is the refractive index, which may vary spatially as well as chromatically. Plugging ansatzes in, the first two of Maxwell's equations become

$$\vec{\nabla} \cdot \left(n^2(\mathbf{r}, \omega)\boldsymbol{\mathcal{E}}\right) = 0 \tag{10.16}$$

$$\vec{\nabla} \cdot \boldsymbol{\mathcal{H}} = 0. \tag{10.17}$$

The third of Maxwell's equations becomes

$$\begin{aligned} &\vec{\nabla} \times \mathbf{E} = e^{-i\omega t}\vec{\nabla} \times \boldsymbol{\mathcal{E}} = -\mu_0 \frac{\partial \mathbf{H}}{\partial t} = i\omega\mu_0 \boldsymbol{\mathcal{H}} e^{-i\omega t} \\ &\vec{\nabla} \times \boldsymbol{\mathcal{E}} = i\omega\mu_0 \boldsymbol{\mathcal{H}} \end{aligned} \tag{10.18}$$

and the fourth becomes

$$\begin{aligned} &\vec{\nabla} \times \mathbf{H} = e^{-i\omega t}\vec{\nabla} \times \boldsymbol{\mathcal{H}} = \frac{\partial \mathbf{D}}{\partial t} = -i\omega\epsilon_0 n^2(\mathbf{r}, \omega)\boldsymbol{\mathcal{E}} e^{-i\omega t} \\ &\vec{\nabla} \times \boldsymbol{\mathcal{H}} = -i\omega\epsilon_0 n^2(\mathbf{r}, \omega)\boldsymbol{\mathcal{E}}. \end{aligned} \tag{10.19}$$

To derive the propagation equations for fields traveling down a waveguide we take the curl of equation 10.18, substituting into 10.19:

$$\vec{\nabla} \times \vec{\nabla} \times \boldsymbol{\mathcal{E}} = \vec{\nabla} \times i\omega\mu_0 \boldsymbol{\mathcal{H}} = \omega^2 \epsilon_0 \mu_0 n^2(\mathbf{r}, \omega)\boldsymbol{\mathcal{E}} \equiv k^2 n^2(\mathbf{r}, \omega)\boldsymbol{\mathcal{E}} \tag{10.20}$$

where the k is the free-space wavenumber, $k = \omega/c$. We can also express the LHS as

$$\vec{\nabla} \times \vec{\nabla} \times \boldsymbol{\mathcal{E}} = \vec{\nabla}(\vec{\nabla} \cdot \boldsymbol{\mathcal{E}}) - \vec{\nabla}^2 \boldsymbol{\mathcal{E}} \tag{10.21}$$

where $\vec{\nabla}^2$ is the vector Laplacian (giving a vector consisting of the scalar Laplacians of each field component). The middle term of the above can be rewritten by applying the divergence-product rule to Maxwell's equations in matter 10.9:

$$0 = \vec{\nabla} \cdot \big(n^2(\mathbf{r}, \omega)\boldsymbol{\mathcal{E}}\big) = n^2(\mathbf{r}, \omega)\vec{\nabla} \cdot \boldsymbol{\mathcal{E}} + \boldsymbol{\mathcal{E}} \cdot \vec{\nabla} n^2(\mathbf{r}, \omega). \tag{10.22}$$

Thus,

$$\vec{\nabla}(\vec{\nabla} \cdot \boldsymbol{\mathcal{E}}) = -\vec{\nabla}\left(\frac{1}{n^2(\mathbf{r}, \omega)} \boldsymbol{\mathcal{E}} \cdot \vec{\nabla} n^2(\mathbf{r}, \omega)\right) = -\vec{\nabla}\Big(\boldsymbol{\mathcal{E}} \cdot \vec{\nabla} \ln\ n^2(\mathbf{r}, \omega)\Big) \tag{10.23}$$

A similar calculation can be repeated for the magnetic field, using the curl-product rule. In the end, the vector propagation equations are

$$\begin{aligned} \vec{\nabla}^2 \boldsymbol{\mathcal{E}} + k^2 n^2(\mathbf{r}, \omega)\boldsymbol{\mathcal{E}} &= -\vec{\nabla}\Big(\boldsymbol{\mathcal{E}} \cdot \vec{\nabla} \ln\ n^2(\mathbf{r}, \omega)\Big) \\ \vec{\nabla}^2 \boldsymbol{\mathcal{H}} + k^2 n^2(\mathbf{r}, \omega)\boldsymbol{\mathcal{H}} &= \Big(\vec{\nabla} \times \boldsymbol{\mathcal{H}}\Big) \times \vec{\nabla} \ln\ n^2(\mathbf{r}, \omega), \end{aligned} \tag{10.24}$$

both of which are inhomogeneous Helmholtz equations. The inhomogeneous terms on the right are responsible for polarization mixing.

10.1.1 Weak guidance limit

In the weak guidance limit, we assume refractive index variations are so small that all derivatives of the refractive index can be set to 0. In this case, the RHS of both equations in 10.24 go to 0, giving homogeneous Helmholtz equations, e.g.

$$\vec{\nabla}^2 \boldsymbol{\mathcal{E}} + k^2 n^2(\mathbf{r}, \omega)\boldsymbol{\mathcal{E}} = 0. \tag{10.25}$$

If the refractive index has no longitudinal (z) dependence, then neither does $\boldsymbol{\mathcal{E}}$, and the above reduces to two mathematically identical equations corresponding to each polarization.

10.1.2 Paraxial limit

In the paraxial approximation, we assume waves propagate primarily along one axis, conventionally chosen to be z. Thus, we can make a second ansatz, e.g.

$$\boldsymbol{\mathcal{E}} = \mathbf{u}(x, y, z)e^{i\beta z} \tag{10.26}$$

where the dependence of $\mathbf{u}$ on z is assumed to be weak, which we will soon quantify; β is the "effective" propagation constant (wavenumber) for the propagating field within the waveguide. In conjunction with weak guidance, substituting the above into equation 10.25 gives the scalar-valued relation

$$\frac{\partial^2 u}{\partial z^2} + 2i\beta\frac{\partial u}{\partial z} + \nabla_\perp^2 u + \left[k^2 n^2(\mathbf{r}, \omega) - \beta^2\right] u = 0 \tag{10.27}$$

where $\nabla_\perp \equiv \partial^2/\partial x^2 + \partial^2/\partial y^2$ is the "transverse" Laplacian. The paraxial approximation states

$$\left|\frac{\partial^2 u}{\partial z^2}\right| \ll \left|\beta\frac{\partial u}{\partial z}\right| \tag{10.28}$$

which is equivalent to assuming that the direction of propagation remains close to the z axis.[2] This yields the scalar-valued paraxial Helmholtz equation:

$$2i\beta\frac{\partial u}{\partial z} + \nabla_{\perp}^2 u + \left[k^2 n^2(\mathbf{r},\omega) - \beta^2\right] u = 0. \tag{10.29}$$

If u is assumed not depend on z at all (an eigenmode of the waveguide), the above becomes the eigenvalue problem

$$\left[\nabla_{\perp}^2 + k^2 n^2(\mathbf{r},\omega)\right] u = \beta^2 u \tag{10.30}$$

with eigenvalue β^2.

10.2 Eigenmode solutions by the finite element method

There are several methods for solving equation 10.30; in this section I cover a popular approach using the finite element method (FEM). This approach is implemented by many commercial waveguide mode solvers, as well as `wavesolve`, a software package developed as part of this dissertation. This section adapts notes from [Say08].

10.2.1 Introduction

The finite element method solves differential equations by discretizing space into a mesh of primitive elements (usually triangles). An ansatz is constructed in piecewise continuous fashion: on each element, the ansatz is defined as a linear combination of "shape" or "interpolation" functions. The coefficients of the linear combination are termed "nodal weights", each of which correspond to a point in the primitive element. Next, the ansatz is constrained with the differential equation in question, reducing it to a system of N equations. Some-

---

[2] To motivate the claim, consider an x-polarized free-space ansatz $\boldsymbol{\mathcal{E}} = e^{i\boldsymbol{\beta}\cdot\mathbf{r}}\hat{\mathbf{x}} \equiv e^{i(k_x x + k_y y + k_z z)}\hat{\mathbf{x}}$. We may write $k_z = \sqrt{\beta^2 - k_x^2 - k_y^2}$. For the ansatz $\mathcal{E} = u(x,y,z)e^{i\beta z}$ to obey the paraxial assumption 10.28, we require $k_x \ll \beta$ and $k_y \ll \beta$ from the paraxial approximation. It's also worth noting that this $\boldsymbol{\mathcal{E}}$ solution does not strictly satisfy $\vec{\nabla}\cdot\boldsymbol{\mathcal{E}} = 0$, and only does so approximately under paraxiality.

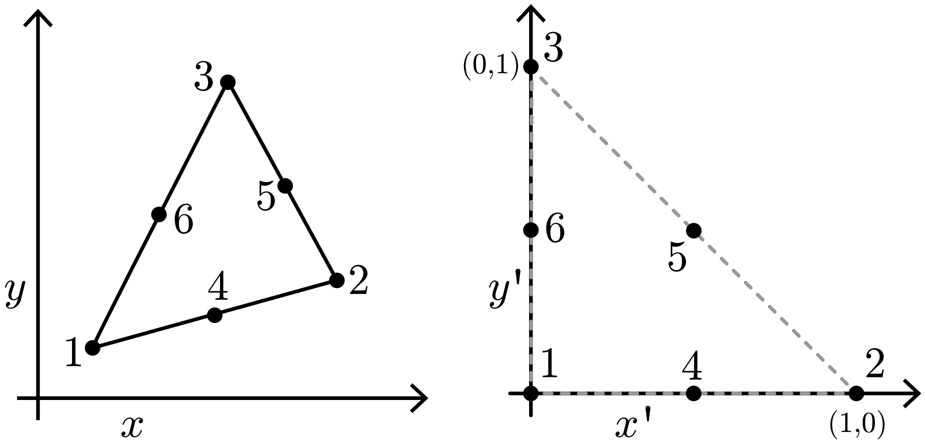

Figure 10.2.1: Left: the quadratic triangle element. Vertices and edge midpoints are indexed in counter-clockwise order. Right: the same element in affine (x', y') coordinates.

times the resulting system is an ordinary linear system, other times a generalized eigenvalue problem. Finally, boundary conditions are applied and the solution is found using some numerical algorithm. I first provide an overview of the quadratic triangle (QT) finite element – the basic primitive element used by `wavesolve`. I then apply FEM to solve equation 10.30 for a single QT element. Our constraint, which allows us to go from a 6-node QT element to a system of 6 equations, will follow Galerkin's method.

10.2.2 Quadratic triangular element

The quadratic triangle (QT; Figure 10.2.1) has 6 nodes: 3 vertices and 3 edge midpoints, conventionally indexed in counter-clockwise (CCW) order. Within the QT, the field u is expanded in terms of 6 shape functions; these functions can be thought of as a partial basis for u over the triangle. Each shape function corresponds to a node, such that it evaluates to 1 at that specific node and 0 at all other nodes. We additionally enforce that each shape function varies like $ax^2+bxy+cy^2$ over the triangle. Below, I present the shape functions for a QT with points (x_j, y_j), where $j = 1, 2, 3$ are vertices and $j = 4, 5, 6$ are edge midpoints. To simplify notation, the shape functions are written in affine coordinates (x', y'), which sets

(x_1, y_1) as the origin and the two remaining vertices to $(0,1)$ and $(1,0)$, as shown in the right panel of Figure 10.2.1. The shape functions $N_j(x', y')$ are:

$$\begin{aligned}
N_1 &= 2(1 - x' - y')(1/2 - x' - y') \\
N_2 &= 2x'(x' - 1/2) \\
N_3 &= 2y'(y' - 1/2) \\
N_4 &= 4x'(1 - x' - y') \\
N_5 &= 4x'y' \\
N_6 &= 4y'(1 - x' - y').
\end{aligned} \tag{10.31}$$

The affine transformation can be written as

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x_{21} & x_{31} \\ y_{21} & y_{31} \end{bmatrix} \begin{bmatrix} x' \\ y' \end{bmatrix} + \begin{bmatrix} x_1 \\ y_1 \end{bmatrix}. \tag{10.32}$$

Here, x_{ij} is shorthand for $x_i - x_j$, and similarly for y_{ij}.

10.2.3 Galerkin's method

Galerkin's method is a special case of a weighted-residual method for approximating the solution to a differential equation. Suppose we have an equation of the form $\mathcal{D}[u] = 0$, where $\mathcal{D}$ is some linear operator that involves differentiation. In the weighted residual method, we multiply the LHS with some weighting function $w(x, y)$ and integrate (average) over some area Ω. Setting this integral to 0 converts the differential equation into an integral equation, known as the "weak form" of the equation:

$$\iint_\Omega w(x, y)\mathcal{D}[u]\, d\Omega = 0. \tag{10.33}$$

If we plug in an approximate solution $\tilde{u}$ to the above, $\mathcal{D}[\tilde{u}]$ will not quite be 0 everywhere in the domain S – instead, we are left with a non-zero residual, multiplied by some weight function, which goes to 0 as the approximation becomes closer to the true solution.

In Galerkin's method, we solve equation 10.33 exactly over some simplified domain Ω. The shape of Ω is usually chosen to be some sort of polygon, parameterized by n nodes. For the case of the QT element, $n = 6$. Over Ω, we expand $\tilde{u}$ in terms of a basis, the shape functions N_i of Ω. We then use each shape function as a weighting function, yielding n integral equations. Each integral can be interpreted as decomposing the residual of $\tilde{u}$ in terms of the shape functions, then finding the values of $\tilde{u}$ so that each of the n components of the residual average to 0 over Ω. In this way, we solve the weak form of the differential equation $\mathcal{D}[u] = 0$ over the element Ω. Global solutions are created by subdividing the global domain into multiple simpler elements, and simultaneously solving for the local solution on each element.

We now explicitly apply Galerkin's method to approximately solve the paraxial wave equation, 10.30, on a single QT element, denoted as Δ. As per the Galerkin method, we multiply equation 10.30 by the shape functions N_i, integrate over the triangle, denoted by Ω_Δ, and set the result to 0.

$$\iint_{\Omega_\Delta} dydx\, N_i \left[\nabla^2 u + \left(k^2 n^2 - \beta^2\right) u\right] = 0. \tag{10.34}$$

Next we integrate the first portion of the above by parts to get rid of the Laplacian. Recall the product rule, and divergence theorem for a domain Ω and closed bounding curve Γ:

$$\begin{aligned} &\int_\Omega \vec{\nabla} \cdot (f \vec{\nabla} g) d\Omega \\ &= \oint_\Gamma f \vec{\nabla} g \cdot \mathbf{n} d\Gamma \\ &= \iint_\Omega \vec{\nabla} f \cdot \vec{\nabla} g d\Omega + \iint_\Omega f \nabla^2 g d\Omega. \end{aligned} \tag{10.35}$$

Here, $\mathbf{n}$ is the unit normal vector for Γ. The second line is the divergence theorem while the third is the product rule. Equating the second and third lines gives us the formula for integration by parts. Applying this to equation 10.34 and replacing u with $\tilde{u}$ gives

$$\oint_{\Gamma_\Delta} N_i \vec{\nabla} u \cdot \mathbf{n} d\Gamma - \iint_{\Omega_\Delta} \vec{\nabla} N_i \cdot \vec{\nabla} \tilde{u} d\Omega + \iint_{\Omega_\Delta} N_i \left[k^2 n^2 - \beta^2\right] \tilde{u} d\Omega = 0. \tag{10.36}$$

The first term in equation 10.36 is our boundary condition, while the latter two terms set the residual of our approximate solution $\tilde{u}$ to 0. Note that in the boundary term, we do *not* approximate u, because boundary conditions are assumed to be known.

We then expand $\tilde{u}$ in terms of the shape functions: $\tilde{u} = \sum_j \tilde{u}_j N_j$, where $\tilde{u}_j$ is the value of $\tilde{u}$ at node j, giving

$$\begin{aligned}
&\oint_{\Gamma_\Delta} N_i \vec{\nabla} u \cdot \mathbf{n} d\Gamma - \iint_{\Omega_\Delta} \vec{\nabla} N_i \cdot \vec{\nabla} \sum_j \tilde{u}_j N_j d\Omega + \iint_{\Omega_\Delta} N_i \left[k^2 n^2 - \beta^2\right] \sum_j \tilde{u}_j N_j d\Omega = 0 \\
&\oint_{\Gamma_\Delta} N_i \vec{\nabla} u \cdot \mathbf{n} d\Gamma = \sum_j \tilde{u}_j \iint_{\Omega_\Delta} \vec{\nabla} N_i \cdot \vec{\nabla} N_j d\Omega - \left[k^2 n^2 - \beta^2\right] \sum_j \tilde{u}_j \iint_{\Omega_\Delta} N_i N_j d\Omega \\
&\oint_{\Gamma_\Delta} N_i \vec{\nabla} u \cdot \mathbf{n} d\Gamma = \sum_j \tilde{u}_j \left[\iint_{\Omega_\Delta} \vec{\nabla} N_i \cdot \vec{\nabla} N_j d\Omega - k^2 n^2 \iint_{\Omega_\Delta} N_i N_j d\Omega\right] \\
&\qquad\qquad\qquad + \beta^2 \sum_j \tilde{u}_j \iint_{\Omega_\Delta} N_i N_j d\Omega.
\end{aligned} \tag{10.37}$$

In the above, we have assumed that the refractive index profile n does not vary over the QT element, which is a good approximation if the element is small enough and/or if the refractive index profile is piecewise-constant. Next, we make the following definitions:

$$\begin{aligned}
A_{ij} &\equiv - \iint_{\Omega_\Delta} \vec{\nabla} N_i \cdot \vec{\nabla} N_j d\Omega + k^2 n^2 \iint_{\Omega_\Delta} N_i N_j d\Omega \\
B_{ij} &\equiv \iint_{\Omega_\Delta} N_i N_j d\Omega \\
c_i &\equiv \oint_{\Gamma_\Delta} N_i \vec{\nabla} u \cdot \mathbf{n} d\Gamma.
\end{aligned} \tag{10.38}$$

Equation 10.37 can equivalently be expressed as the matrix equation

$$A\tilde{\mathbf{u}} + \mathbf{c} = \beta^2 B\tilde{\mathbf{u}} \tag{10.39}$$

where $\tilde{\mathbf{u}} \equiv [\tilde{u}_1, \tilde{u}_2, \ldots]^T$ and similarly for $\mathbf{c}$. The shape function integrals in the A and B matrices can be computed analytically, e.g. using a symbolic math engine; see §10.A for tabulated results. Application of Galerkin's method to each QT element yields a separate system of 6 equations like the above; from here, the systems for each QT element are combined (essentially by addition into a system of N of equations where N is the total number of nodes in the mesh). Combining all the c_i terms yields a contour integral which traverses the exterior edge of the mesh; this term sets a boundary condition on the electric field. `wavesolve` applies a homogeneous Neumann boundary condition, which sets $\mathbf{c} = 0$ and reduces equation 10.39 into a generalized eigenvalue problem that can be solved numerically.

10.2.4 Vectorial solvers

`wavesolve` also implements a mode solving method when the weak guidance criterion is not met – in this case the eigenmodes of the waveguide may have mixed polarization. To solve for the modes, we may apply Galerkin's method to equations 10.18 and 10.19 directly to derive a generalized eigenvalue problem, like before. For brevity, I will provide only an overview of the calculations. First, note that vectorial solving methods based on finite elements (e.g. [LSC91]) typically expand the transverse component of the electric field over element edges, not nodes: this is to ensure that continuity relations in Maxwell's equations[3] are automatically satisfied. The longitudinal field component is expanded over the nodes, as usual, so that a vectorial solver must ultimately solve an eigenvalue problem of dimension $M + N$, where M is the number of edges and N is the number of nodes in the mesh. A common approach using linear triangular (LT) elements is to define the vector-valued "edge

---

[3] In particular, the continuity of the tangential component of the electric field across a material boundary.

shape functions" from the scalar-valued nodal shape functions. Specifically, the nodal shape functions for an LT element are, in affine coordinates:

$$\begin{aligned} N_1 &= 1 - x' - y' \\ N_2 &= x' \\ N_3 &= y' \end{aligned} \tag{10.40}$$

and the vector-valued edge shape functions are

$$\mathbf{N}_{ij} = \left[N_i \vec{\nabla} N_j - N_j \vec{\nabla} N_i\right] l_{ij} \tag{10.41}$$

where l_{ij} is the signed edge length between vertices i and j. Application of Galerkin's method to equations 10.18 and 10.19 ultimately yields the generalized eigenvalue problem

$$\begin{bmatrix} A_{tt} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{u}}_t \\ \tilde{\mathbf{u}}_z \end{bmatrix} = \beta^2 \begin{bmatrix} B_{tt} & B_{tz} \\ B_{zt} & B_{zz} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{u}}_t \\ \tilde{\mathbf{u}}_z \end{bmatrix} \tag{10.42}$$

where the vectors $\mathbf{u}_t$ and $\mathbf{u}_z$ correspond to the transverse and longitudinal components of the electric field, which are defined over the mesh edges and nodes, respectively. The A and B matrices are defined next; for brevity, indices of bold (vector-valued) quantities iterate through mesh edges, while indices of non-bolded quantities iterate through mesh nodes.

$$\begin{aligned} A_{tt,ij} &= \iint \left[k^2 n^2 \mathbf{N}_i \cdot \mathbf{N}_j - \left(\vec{\nabla} \times \mathbf{N}_i\right) \cdot \left(\vec{\nabla} \times \mathbf{N}_j\right)\right] \\ B_{tt,ij} &= \iint \mathbf{N}_i \cdot \mathbf{N}_j \\ B_{tz,ij} &= \iint \mathbf{N}_i \cdot \vec{\nabla} N_j \\ B_{zt,ij} &= B_{tz,ji} \\ B_{zz,ij} &= \iint \left[\vec{\nabla} N_i \cdot \vec{\nabla} N_j - k^2 n^2 N_i N_j\right]. \end{aligned} \tag{10.43}$$

These integrals can be analytically computed as before; see §10.B for tabulated results.

10.3 Propagation methods

For waveguides whose properties depend on z, there are no eigenmodes in the conventional sense and electric fields must be propagated. In the next section, I review a popular propagation method for waveguide modeling.

10.3.1 Finite difference beam propagation method

This section adapts material from [Ped15]. The finite difference beam propagation method (FD BPM) is used in commercial packages such as RSoft BeamProp, as well as the open-source package `lightbeam` [Lin20]. I review FD BPM in the paraxial and weak guidance regime. This method solves the paraxial, scalar Helmholtz equation 10.29. For clarity, we write this equation in a new form using the differential operator $\mathcal{D}$:

$$\begin{aligned} -2ikn_0\frac{\partial u}{\partial z} &= \mathcal{D}u \\ \mathcal{D}u &\equiv \nabla_\perp^2 u + k^2\left[n^2(\mathbf{r},\omega) - n_0^2\right]u. \end{aligned} \tag{10.44}$$

Here, we have set the unknown wavenumber β to kn_0, where n_0 is some reference index that is chosen *a priori* as a simulation parameter, typically following some average value of the waveguide's refractive index values. FD BPM solves the above replacing differential operators with finite difference approximations and then applying the Crank-Nicolson method. First, define the x central difference operator δ_x for function f such that

$$\delta_x u \equiv u(x+\Delta x/2) - u(x-\Delta x/2) \tag{10.45}$$

and likewise for y and z. Thus,

$$\frac{\partial}{\partial x} \approx \frac{\delta_x}{\Delta x} \tag{10.46}$$

and

$$\frac{\partial^2}{\partial x^2} \approx \frac{\delta_x^2}{\Delta x^2}, \tag{10.47}$$

yielding the typical central finite difference approximation to the second order derivative:

$$\frac{\partial^2 u}{\partial x^2} \approx \frac{u(x+\Delta x) - 2u(x) + u(x-\Delta x)}{\Delta x^2}. \tag{10.48}$$

We now apply the Crank-Nicolson scheme to equation 10.44, which discretizes u along the z direction and converts the differential equation into a linear system. Denote u^j as the field u at the jth z-step. Then by Crank-Nicolson:

$$-\frac{2ikn_0}{\Delta z}\left(u^{j+1} - u^j\right) = \frac{1}{2}\left[\mathcal{D}u^{j+1} + \mathcal{D}u^j\right] \tag{10.49}$$

or rearranging based on z-step:

$$\left[-\frac{2ikn_0}{\Delta z} - \frac{1}{2}\mathcal{D}\right]u^{j+1} = \left[-\frac{2ikn_0}{\Delta z} + \frac{1}{2}\mathcal{D}\right]u^j. \tag{10.50}$$

We then split the operator $\mathcal{D}$ into an x and y component (equally splitting the part of $\mathcal{D}$ not specific to x or y) and apply finite differences:

$$\begin{aligned} \mathcal{D} &= \mathcal{D}_x + \mathcal{D}_y \\ \mathcal{D}_x &\equiv \frac{\partial}{\partial x^2} + \frac{1}{2}k^2\left(n^2 - n_0^2\right) \approx \frac{\delta_x^2}{\Delta x^2} + \frac{1}{2}k^2\left(n^2 - n_0^2\right) \\ \mathcal{D}_y &\equiv \frac{\partial}{\partial y^2} + \frac{1}{2}k^2\left(n^2 - n_0^2\right) \approx \frac{\delta_y^2}{\Delta y^2} + \frac{1}{2}k^2\left(n^2 - n_0^2\right). \end{aligned} \tag{10.51}$$

The above is solved by approximately factorizing[4] into an x and y component, giving

$$\left(-\frac{2ikn_0}{\Delta z} - \frac{1}{2}\mathcal{D}_x\right)\left(-\frac{2ikn_0}{\Delta z} - \frac{1}{2}\mathcal{D}_y\right)u^{j+1} = \left(-\frac{2ikn_0}{\Delta z} + \frac{1}{2}\mathcal{D}_x\right)\left(-\frac{2ikn_0}{\Delta z} + \frac{1}{2}\mathcal{D}_y\right)u^j. \tag{10.52}$$

---

[4] Specifically, we introduce a $\mathcal{D}_x\mathcal{D}_y$ term, which assumed to be negligible for small Δz.

Finally, to numerically advance u, we split each jth z-step into halves corresponding to the $\mathcal{D}_x$ and $\mathcal{D}_y$ operators, respectively:

$$\begin{aligned}\left(-\frac{2ikn_0}{\Delta z}-\frac{1}{2}\mathcal{D}_y\right)u^{j+1/2}&=\left(-\frac{2ikn_0}{\Delta z}+\frac{1}{2}\mathcal{D}_y\right)u^j\\\left(-\frac{2ikn_0}{\Delta z}-\frac{1}{2}\mathcal{D}_x\right)u^{j+1}&=\left(-\frac{2ikn_0}{\Delta z}+\frac{1}{2}\mathcal{D}_x\right)u^{j+1/2}.\end{aligned}\tag{10.53}$$

For the above to hold, the operators within the parentheses in 10.52 must commute, at least approximately. This should be the case at least in the weak guidance regime. Ultimately, the above constitutes two sets of tridiagonal linear systems of equations which must be solved at each z-step, e.g. with the Thomas algorithm.

10.3.1.1 Improving transverse accuracy: Douglas method

The second-order derivatives in x and y will have an error on the order of Δx^2 and Δy^2, the where Δx and Δy are the grid spacings of the simulation. This can be improved to fourth-order accuracy using the Douglas method, proposed by [SY93] and which `lightbeam` implements. I present an overview of the Douglas method next. For a function u we have by Taylor expansion

$$u(x+\Delta x)=\sum_n \frac{1}{n!}\frac{\partial^n u}{\partial x^n}\bigg|_x \Delta x^n \tag{10.54}$$

and a similar relation for $u(x-\Delta x)$. We will use this expansion to find the lowest-order error term in the central finite difference approximation to the second-order derivative. Expanding $\delta_x^2 u$, we find

$$\begin{aligned}\delta_x^2 u &= u(x+\Delta x)+u(x-\Delta x)-2u(x)\\&=2u(x)+0+\frac{\partial^2 u}{\partial x^2}\Delta x^2+0+\frac{1}{12}\frac{\partial^4 u}{\partial x^4}\Delta x^4+0+o(\Delta x^6)-2u(x)\\\frac{\partial^2 u}{\partial x^2}&=\frac{1}{\Delta x^2}\left[\delta_x^2-\frac{1}{12}\delta_x^4+o(\Delta x^6)\right]u.\end{aligned}\tag{10.55}$$

In the Douglas method, we apply the operator $(1 + \delta_x^2/12)$ on both sides of the above to eliminate the fourth-order term:

$$\left(1 + \frac{\delta_x^2}{12}\right)\frac{\partial^2 u}{\partial x^2} = \frac{1}{\Delta x^2}\left(\delta_x^2 + o(\Delta x^6)\right)u \tag{10.56}$$

and so we obtain

$$\frac{\partial^2}{\partial x^2} \approx \frac{\delta_x^2}{\Delta x^2(1 + \delta_x^2/12)} \tag{10.57}$$

with fourth-order accuracy. This operator can be substituted into equation 10.50, and will modify the coefficients of the resulting tridiagonal system of equations. Note that `lightbeam` applies a modified version of the above Douglas method which also works on non-uniform grids [YSS97].

Appendix - Chapter 10

10.A Integrals for finite element mode solvers (scalar)

This section tabulates integrals for scalar solvers using QT elements. For brevity, denote d as $|J|$, where J is the Jacobian of the affine transformation 10.32, i.e. $d \equiv x_{21}y_{31} - x_{31}y_{21}$. Vertices and edge midpoints are indexed in CCW order, as in Figure 10.2.1. Note that the results may not be in their simplest form.

Condition	Description	Value
$\iint_{\Omega_\Delta} dydx\, N_i N_j$		
$i=1,\, j=1$	vertex - same vertex	$d/60$
$i=1,\, j=2$	vertex - other vertex	$-d/360$
$i=1,\, j=4$ or 6	vertex - adjacent edge	0
$i=1,\, j=5$	vertex - opposite edge	$-d/90$
$i=4,\, j=4$	edge - same edge	$4d/45$
$i=4,\, j=5$	edge - other edge	$2d/45$
$\iint_{\Omega_\Delta} dydx\, \vec{\nabla} N_i \cdot \vec{\nabla} N_j$		
$i=1,\, j=1$	vertex - same vertex (grad.)	$\frac{y_{32}^2+x_{23}^2}{2d}$
$i=1,\, j=2$	vertex - other vertex (grad.)	$\frac{y_{32}y_{31}+x_{32}x_{31}}{6d}$
$i=1,\, j=4$	vertex - CCW edge (grad.)	$-\frac{2(y_{32}y_{31}+x_{32}x_{31})}{3d}$
$i=1,\, j=5$	vertex - opposite edge (grad.)	0
$i=1,\, j=6$	vertex - CW edge (grad.)	$\frac{2(y_{21}y_{32}+x_{21}x_{32})}{3d}$
$i=4,\, j=4$	edge - same edge (grad.)	$\frac{4(y_{32}^2+y_{31}y_{21}+x_{32}^2+x_{31}x_{21})}{3d}$
$i=4,\, j=5$	edge - other edge (grad.)	$\frac{4(y_{21}y_{32}+x_{21}x_{32})}{3d}$

Table 10.A.1: Integrals for a QT element. For brevity, not all results are shown; however, all required integrals can be obtained by simultaneously permuting vertices $1 \to 2 \to 3 \to 1$ and edges $4 \to 5 \to 6 \to 4$, and/or by swapping indices.

10.B Integrals for finite element mode solvers (vector)

This section tabulates integrals for vectorial solvers using LT elements. Vertices for the LT element are assumed to be indexed in CCW order. l_{ij} is the signed edge length between vertices i and j; the choice of which direction (CW or CCW) is positive does not matter, we just need to be consistent. Similarly, $\mathbf{N}_{ij}$ is the vector-valued shape function corresponding to the edge connecting vertices i and j, while N_k is the scalar-valued shape function corresponding to vertex k. Note that the results may not be in their simplest form.

Condition	Description	Value
$\iint_{\Omega_\Delta} dydx\, \mathbf{N}_{ij} \cdot \mathbf{N}_{mn}$		
$i, j = 1, 2;\ m, n = 1, 2$	edge - same edge	$l_{12}^2 \frac{l_{12}^2 - 3x_{21}x_{31} + 3l_{31}^2 - 3y_{21}y_{31}}{12d}$
$i, j = 1, 2;\ m, n = 2, 3$	edge - other edge	$l_{12}l_{23} \frac{l_{12}^2 - x_{21}x_{31} - l_{31}^2 - y_{21}y_{31}}{12d}$
$\iint_{\Omega_\Delta} dydx\, \left(\vec{\nabla} \times \mathbf{N}_{ij}\right) \cdot \left(\vec{\nabla} \times \mathbf{N}_{mn}\right)$		
any	curl edge - curl edge	$2l_{ij}l_{mn}/d$
$\iint_{\Omega_\Delta} dydx\, \mathbf{N}_{ij} \cdot \vec{\nabla} N_k$		
$i, j = 1, 2;\ k = 1$	edge - CW vertex	$l_{12} \frac{-x_{13}x_{23} - l_{23}^2 - y_{13}y_{23}}{6d}$
$i, j = 1, 2;\ k = 2$	edge - CCW vertex	$l_{12} \frac{-x_{21}x_{31} + 2l_{31}^2 - y_{21}y_{31}}{6d}$
$i, j = 1, 2;\ k = 3$	edge - opposite vertex	$l_{12} \frac{l_{12}^2 - 2x_{21}x_{31} - 2y_{21}y_{31}}{6d}$
$\iint_{\Omega_\Delta} dydx\, N_i N_j$		
$i = j$	vertex - same vertex	$d/12$
$i \neq j$	vertex - other vertex	$d/24$
$\iint_{\Omega_\Delta} dydx\, \vec{\nabla} N_i \cdot \vec{\nabla} N_j$		
$i = 1,\ j = 1$	vertex - same vertex (grad.)	$\frac{l_{23}^2}{2d}$
$i = 1,\ j = 2$	vertex - other vertex (grad.)	$-\frac{x_{13}x_{23} + y_{13}y_{23}}{2d}$

Table 10.B.1: Integrals of the vector-valued edge shape functions and their various derivatives for an LT element. For brevity, not all results are shown; however, all required integrals can be obtained by permuting vertices 1 → 2 → 3 → 1, and/or by swapping indices.

REFERENCES

[Bur09] Jacques Bures. *Guided optics: Optical fibers and all-fiber components.* John Wiley & Sons, 2009.

[Lin20] Jonathan W. Lin. “Lightbeam.” Astrophysics Source Code Library, February 2020. ascl:2102.006.

[Lin24] Jonathan Lin. “cbeam.” Astrophysics Source Code Library, record ascl:2404.001, apr 2024.

[LSC91] J.-F. Lee, D.-K. Sun, and Z.J. Cendes. “Full-wave analysis of dielectric waveguides using tangential vector finite elements.” *IEEE Transactions on Microwave Theory and Techniques*, **39**(8):1262–1271, 1991.

[Ped15] Ginés Lifante Pedrola. *Vectorial and Three-Dimensional Beam Propagation Techniques*, chapter 3, pp. 71–129. John Wiley & Sons, Ltd, 2015.

[Say08] Francisco-Javier Sayas. “A gentle introduction to the Finite Element Method.” *Lecture notes, University of Delaware*, 2008.

[SY93] Lizhong Sun and Gar Lam Yip. “Modified finite-difference beam-propagation method based on the Douglas scheme.” *Opt. Lett.*, **18**(15):1229–1231, Aug 1993.

[YSS97] J. Yamauchi, J. Shibayama, M. Sekiguchi, and H. Nakano. “Finite-difference beam propagation method based on the generalized Douglas scheme for a nonuniform grid.” *IEEE Photonics Technology Letters*, **9**(1):67–69, 1997.

CHAPTER 11

Coupled-mode theory for astrophotonics

Contents

The cross-disciplinary field of astrophotonics applies photonic devices to the technical challenges of astronomy [JGA23], and has since led to novel demonstrations of photonic nullers [MNT21, XJR22], spectrometers [GVD19, LJF21], wavefront sensors [NWB20, LFX23], and wavefront correctors [DCT24]. Such devices offer new ways to process the light collected by astronomical telescopes, in a form factor much smaller than bulk-optical components. With recent improvements in the fabrication capability for such devices, the astrophotonics community is now poised to consider how to tune designs for better performance in astronomical applications. This might include optimizing directional couplers or tricouplers [KMT22] and phase shifters for broadband nulling, or designing PLs which simultaneously maximize wavefront sensing and spectroscopy capability.

To do so requires a physical understanding of how light propagates through astrophotonic devices. At least within the astrophotonics community, this understanding is often heavily reliant on numerical modelling techniques such as the finite-difference beam propagation method, which are accurate but opaque and difficult to intuit. At the same time, design optimizations for photonic devices would benefit from simulation methods that are faster than the standard methods used today.

The goal of this paper is to introduce a mathematical tool that can be used to model and understand the propagation of light through a "slowly-varying" photonic device, through the lens of astronomical instrumentation. The precise meaning of the slowly-varying constraint will be developed later in this work, but for now it suffices to say that this constraint applies to the vast majority of astrophotonic devices, including PLs, photonic integrated circuits, arrayed waveguide gratings, directional couplers, multimode interferometers, fiber Bragg gratings, and Kerr combs; and that photonic devices not satisfying the slowly-varying constraint should probably be redesigned, if possible, to satisfy this constraint and simplify the modelling process. The mathematical tools reviewed in this work all fall under the broad category of "coupled-mode theory" (CMT).

SYMBOLS

Symbol	Description
$\|v\rangle$	vector (ket), from Dirac's bra-ket notation
$\langle u\|$	covector (dual vector, bra)
$\langle u\|v\rangle$	inner product, $\iint dxdy\, u^* v$
$\langle u\|A\|v\rangle$	$\iint dxdy\, u^* Av$, where A is an operator
$\|\phi_j\rangle$	an eigenmode of a z-invariant waveguide
$\|\xi_j^z\rangle$	an instantaneous eigenmode of a z-varying waveguide
$\|\Psi\rangle$	an ansatz for the waveguide propagation equation
D_z	a linear operator involving differentiation against z
$A(z)$	an operator that varies slowly with z; eigenbasis is $\|\xi_j^z\rangle$
$n_0(x,y)$	unperturbed refractive index, a scalar field
$\delta\epsilon$	perturbation to the refractive index (squared)
$n(x,y,z)$	perturbed refractive index, $n^2 = n_0^2 + \delta\epsilon$
a_i	amplitude of CMT mode i
k	free-space longitudinal (z) angular wavenumber
β	effective longitudinal angular wavenumber, $\beta_i \to$ eigenmode i
κ_{ij}	coupling coeff. between modes i and j in fixed-basis CMT
$\tilde{\kappa}_{ij}$	coupling coeff. between modes i and j in varying basis CMT
σ_i	Pauli spin matrices
$\mathbb{C}^n$	complex vector space, dimension n
θ_j	dynamical phase of mode j
ϕ_j	geometric phase of mode j
$\mathcal{A}$	Berry connection
$\|p\rangle$	Jones vector representing a polarization state
$\vec{\nabla}_\perp f$	transverse gradient of f, $[\partial f/\partial x, \partial f/\partial y]^T$

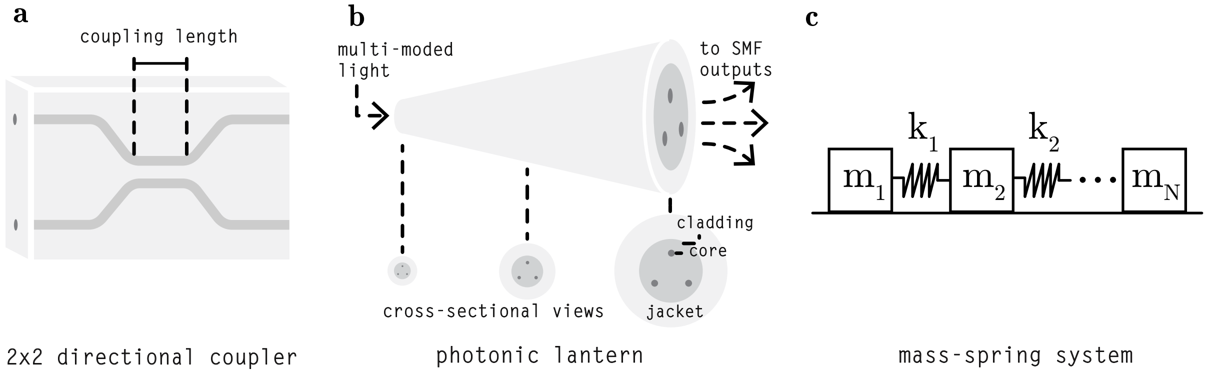

Figure 11.1.1: **a.** A 2 × 2 directional coupler, formed from two single-moded waveguide channels which brought close enough to transfer power and then separated again. **b.** A photonic lantern: an optical waveguide which looks like a few-moded optical fiber on one end and splits off into multiple single-moded outputs at the other. The propagation of light through this device also can be modeled as coupled harmonic oscillators. **c.** A mechanical system of coupled harmonic oscillators composed of masses and springs, constrained to move in one dimension along a frictionless surface.

11.1 Introduction to coupled-mode theory

The CMT is a well-known and well-used mathematical tool. The origins of the CMT are often attributed to studies of radio transmission lines and electron beams by Miller and Pierce, respectively [Mil54, Pie54]; but the underlying ideas appear even earlier in quantum mechanics, for instance in studies of molecular dynamics by Born and Oppenheimer [BO27]. In telecommunications, the CMT has been widely used to study the propagation of light through few-moded optical waveguides, both analytically and numerically [Oka06, Agr21, HHK87]. To motivate the CMT, we will first consider a mechanical analog. Suppose we have a system of N masses connected by springs, as shown in Figure 11.1.1c; our goal is to determine the dynamics of the system given an arbitrary initial condition. We first consider the dynamics of a *single* mass-spring system (where the other end of the spring is fixed), which is a harmonic oscillator obeying

$$\frac{d^2x}{dt^2} = -\frac{k}{m}x \tag{11.1}$$

where k is the spring constant, m the mass, and x denotes position. The corresponding solution for $x(t)$ is the complex exponential. If we chain many mass-spring systems together, the dynamics of the system is that of a many harmonic oscillators, which can also transfer energy between each other. Mathematically, the equation of motion is

$$\frac{d^2\mathbf{x}}{dt^2} = -A\mathbf{x} \tag{11.2}$$

where $\mathbf{x} = [x_1, x_2, ..., x_N]^T$ are the positions of each mass and A is an $N \times N$ matrix whose off-diagonal terms correspond to the linkage between the individual mass-spring systems. To solve this equation, we assume the solution is harmonic in time, i.e. we propose the ansatz $\mathbf{x} \propto \mathbf{u}e^{i\omega t}$. This reduces equation 11.2 to the eigenvalue problem

$$\omega^2\mathbf{u} = A\mathbf{u} \tag{11.3}$$

where $\mathbf{u}$ and ω^2 are identified as an eigenvector-eigenvalue pair. The motion of all the masses in the system may therefore be expanded in terms of the eigenvectors, whose amplitudes oscillate harmonically in time; this is "normal mode analysis".

So far, our analysis has been exact. We now consider the case where the the matrix A is allowed to vary with time, for instance through a gradual weakening of the springs or application of a varying external force. This is the problem that coupled-mode theory (CMT) attempts to solve. The principle of the CMT is as follows: suppose there is a system which supports a basis of harmonic oscillator eigenmodes. When a slowly varying perturbation is added, we assume that perturbed solutions will be "almost" harmonic, and attempt to find solutions formed from a linear combination of the unperturbed eigenmodes. To account for the perturbation, the modes are allowed to cross-couple through varying mode coefficients — hence the name "coupled-mode" theory.

Existing treatments in the context of optical waveguides typically provide the "fixed-basis" formulation of the theory, suitable for optical fibers or composite waveguides composed of individual channels, such as directional couplers [Agr21], but not for slowly varying waveguides whose eigenmodes change incrementally with longitudinal coordinate z. A motivating example for such a waveguide is the PL, which in recent years has been the subject of growing interest for astronomical applications such as telluric line suppression [TEB13], high-resolution spectroscopy [JSG17, LJF21], starlight nulling [MNT21, XJR22], wavefront sensing [LFX23], and spectroastrometry [KFL24].

The main purpose of this paper, comprising §11.2-11.3, is to provide an overview for the fixed-basis CMT as well as for the more broadly applicable "varying-basis" formulation [SL83], in some instances called "coupled local mode theory" (CLMT) [JJJ10, CLW20], and to demonstrate its applicability to slowly varying waveguides. Unlike some previous implementations [CLW20, SS20] which assume non-degenerate modes, the formulation presented here is applicable even in the presence of degeneracy, and accounts for non-adiabatic cross-coupling. The secondary purpose, comprising §11.4, is to show how many of the ideas in the varying-basis CMT have well-studied counterparts in quantum mechanics: concepts including the WKB approximation, perturbation theory, the adiabatic approximation, and geometric phase. While individual connections have been often drawn in previous studies, this work aims at establishing a broader connection, hopefully introducing the reader to techniques and ideas that have normally not been associated with coupled-mode theory and which could be useful in future developments of astrophotonic devices. Throughout, derivations are written using mathematical notation borrowed from quantum mechanics; I briefly review this notation for readers who may be less familiar with it. Finally, in §11.5 I verify my analysis with `cbeam`, a simple numerical implementation of the varying-basis CMT, and use it to simulate a two-channel directional coupler and a 6-port PL.

11.2 Coupled-mode theory for z-invariant waveguides

This section presents a brief derivation of a form of the CMT using a fixed eigenbasis [Agr19], which is similar to time-dependent perturbation theory from quantum mechanics [CHC21]. For simplicity I will also assume weak guidance, though the CMT is equally applicable to vector eigenmodes ([Oka06]; §11.4.3). First, I will introduce the necessary notation, including Dirac's bra-ket notation. Broadly speaking, the electric field within a waveguide is some complex-valued function $u(x, y, z)$, where (x, y) are the transverse coordinates (perpendicular to the axis of the waveguide) and z is the longitudinal coordinate (parallel to the axis of the waveguide). I will denote an inner product over the transverse directions with bra-ket notation:

$$\langle u|v\rangle \equiv \iint u^* v \, dxdy. \tag{11.4}$$

Here, u^* is the complex conjugate of u, and the brackets around the "ket" $|v\rangle$ indicate that the quantity is a vector in a complex vector space.[1] The "bra" $\langle u|$ is a "dual vector"; the inner product $\langle u|v\rangle$ is often referred to as the "overlap integral" between $|u\rangle$ and $|v\rangle$.

Wave propagation under the approximation of weak guidance obeys the Helmholtz equation [Agr19]: a wavefront function $\psi(x, y, z)$ or equivalently a wavefront ket $|\psi\rangle$ obeys

$$\left[\nabla^2 + k^2 n_0^2\right] |\psi\rangle = 0 \tag{11.5}$$

where k is the free space wavenumber and $n_0 = n_0(x, y)$ is the refractive index profile. Since the waveguide structure does not change with z, solutions should also not change with z (modulo a phase factor), leading to the ansatz

$$|\psi\rangle = |\phi_j\rangle e^{i\beta_j z} \tag{11.6}$$

---

[1] Specifically, an element of the Hilbert space of square-integrable functions, L^2. This is because we must have $\langle v|v\rangle$ finite, which is loosely the same as saying that the electric field is transversely "bound" by the waveguide.

where β_j is some constant of propagation, and $|\phi_j\rangle$ denotes a square-integrable function $\phi_j(x, y)$, which represents the electric field of a guided mode: a particular distribution of light which does not change shape as it propagates down the waveguide. Substituting this ansatz into the Helmholtz equation converts it into the following the eigenvalue problem

$$\left[\frac{\partial}{\partial x^2} + \frac{\partial}{\partial y^2} + k^2 n_0^2(x, y)\right] |\phi_j\rangle = \beta_j^2 |\phi_j\rangle \tag{11.7}$$

with eigenvalues β_j^2; note the structural similarity of this equation and the time-independent Schrödinger equation. More general solutions can be expanded in terms of $|\phi_j\rangle$, the eigenmodes.

In an ideal waveguide, power does not transfer between eigenmodes. However, cross-coupling can be induced by a weak perturbation in the refractive index (e.g. a bend in an optical fiber). Define the perturbed refractive index profile as

$$n^2(x, y, z) = n_0^2(x, y) + \delta\epsilon(x, y, z) \tag{11.8}$$

where $\delta\epsilon(x, y, z)$ is a weak function of x,y in accordance with weak guidance, as well as z. The perturbed Helmholtz equation is

$$\left[\nabla^2 + k^2 n_0^2 + k^2 \delta\epsilon\right] |\psi\rangle = 0. \tag{11.9}$$

The CMT uses the following ansatz for the perturbed Helmholtz equation.

$$|\Psi\rangle = \sum_j a_j(z)\, e^{i\beta_j} |\phi_j\rangle. \tag{11.10}$$

The above essentially claims that any propagating wavefront can be expressed as a linear combination of the ideal and z-invariant eigenmodes, modulated by z-varying coefficients, motivating the name "fixed-basis" CMT. Further assume that the mode coefficients $a_j(z)$

vary slowly compared to β_j — this approximation is covered in more detail in §11.3.2. This yields

$$\begin{gathered}\left[\frac{\partial^2}{\partial z^2}+\nabla_\perp^2+k^2n_0^2+k^2\delta\epsilon\right]\sum_j a_j(z)\,e^{i\beta_j z}|\phi_j\rangle=0\\ \sum_j\left[2i\beta_j\frac{da_j}{dz}e^{i\beta_j z}|\phi_j\rangle+k^2\delta\epsilon a_j e^{i\beta_j z}|\phi_j\rangle\right]=0\end{gathered}\tag{11.11}$$

where $\nabla_\perp^2\equiv\partial^2/\partial x^2+\partial^2/\partial y^2$. Complete an inner product by left-multiplying with $\exp(-i\beta_i z)\langle\phi_i|$:

$$2i\beta_i\frac{da_i}{dz}+\sum_j a_j k^2 e^{i(\beta_j-\beta_i)z}\langle\phi_i|\delta\epsilon|\phi_j\rangle=0.\tag{11.12}$$

The coupling coefficients κ_{ij} are defined as

$$\kappa_{ij}=\frac{k^2}{2\beta_i}\langle\phi_i|\delta\epsilon|\phi_j\rangle.\tag{11.13}$$

The definition is conventional and may differ from text to text. The mode coefficients of the propagating wavefront in the perturbed waveguide solve the following coupled system of differential equations:

$$\frac{da_i}{dz}=i\sum_j e^{i(\beta_j-\beta_i)z}\kappa_{ij}a_j.\tag{11.14}$$

Appendix 11.A solves the above for the simple case of a two-mode waveguide (e.g. a 2 × 2 directional coupler).

11.3 Coupled-mode theory for slowly varying waveguides

The fixed-basis CMT fails for waveguides whose eigenbases (i.e. guided modes) vary with z. One such waveguide is the PL, which transitions slowly from a few-mode step-index fiber geometry to a multi-core fiber geometry. Even in the absence of perturbations, cross-coupling can be induced by the changing structure of the waveguide; in other words, the change in the waveguide structure for two nearby values of z is a weak perturbation in the sense of §11.2.

To study such a waveguide, I apply the CMT to solve a differential equation of the form

$$D_z|\Psi\rangle + A(z)|\Psi\rangle = 0 \tag{11.15}$$

where D_z is a linear differential operator with respect to parameter z and the operator $A(z)$ varies slowly with z; for more details on this criterion, see §11.3.1-11.3.3. $A(z)$ admits an "instantaneous" eigenbasis $|\xi_j(z)\rangle$ with eigenvalues $\lambda_j(z)$:

$$A(z)|\xi_j(z)\rangle = \lambda_j(z)|\xi_j(z)\rangle. \tag{11.16}$$

Going forward I will indicate z dependence with a superscript, i.e. $|\xi_j^z\rangle \equiv |\xi_j(z)\rangle$. As before, the above form can be obtained under weak guidance, which reduces Maxwell's equations to a set of decoupled scalar Helmholtz equations. This gives the identifications

$$\begin{aligned} D_z &= \frac{\partial^2}{\partial z^2} \\ A(z) &= \left[\nabla_\perp^2 + k^2 n^2(x,y,z)\right]. \end{aligned} \tag{11.17}$$

This is not the only identification possible; see §11.4.3.1, which applies the CMT to polarized plane waves. The operator $A(z)$ has the eigenvalue equation

$$A(z)|\xi_j^z\rangle = \beta_j^2(z)\,|\xi_j^z\rangle \tag{11.18}$$

where the $\lambda_j(z) \equiv \beta_j^2(z)$ are the eigenvalues. In the varying-basis CMT, solutions are expanded in $|\xi_j^z\rangle$. An ansatz is (c.f. equation 11.10):

$$|\Psi\rangle = \sum_j a_j(z) e^{i\int_0^z \beta_j(z')dz'}|\xi_j^z\rangle. \tag{11.19}$$

Plugging the above back into the Helmholtz equation and completing an inner product with $\langle\Psi_i| \equiv e^{-i\int_0^z \beta_i(z')dz'}\langle\xi_i^z|$ yields:

$$\begin{aligned}
0 = {} & 2i\beta_i\frac{da_i}{dz} + ia_i\frac{d\beta_i}{dz} \\
& + \sum_j e^{i\int_0^z(\beta_j-\beta_i)dz'}\left[2\frac{da_j}{dz} + 2ia_j\beta_j\right]\langle\xi_i^z|\frac{\partial}{\partial z}|\xi_j^z\rangle \\
& + \frac{d^2a_i}{dz^2} + \sum_j a_j e^{i\int_0^z(\beta_j-\beta_i)dz'}\langle\xi_i^z|\frac{\partial^2}{\partial z^2}|\xi_j^z\rangle.
\end{aligned} \tag{11.20}$$

The following sections review common approximations to the above coupled-mode equations.

11.3.1 Paraxial approximation

The paraxial approximation is usually stated as:

$$\left|\frac{\partial^2\xi_j}{\partial z^2}\right| \ll \beta_j\left|\frac{\partial\xi_j}{\partial z}\right|. \tag{11.21}$$

This is equivalent to assuming that the propagation direction of the wavefront is close to the optical axis of the waveguide.[2] This approximation removes all second-order derivatives of $|\xi\rangle$, leaving:

$$\begin{aligned}
0 = {} & 2i\beta_i\frac{da_i}{dz} + ia_i\frac{d\beta_i}{dz} \\
& + \sum_j e^{i\int_0^z(\beta_j-\beta_i)dz'}\left[2\frac{da_j}{dz} + 2ia_j\beta_j\right]\langle\xi_i^z|\frac{\partial}{\partial z}|\xi_j^z\rangle + \frac{d^2a_i}{dz^2}.
\end{aligned} \tag{11.22}$$

---

[2] There are so-called "wide-angle" approximations (e.g. [Had92]) which can be used to relax the assumption of paraxiality.

11.3.2 Slowly varying envelope

The slowly varying envelope approximation (SVEA) claims:

$$\left|\frac{d^2 a_j}{dz^2}\right| \ll \beta_j \left|\frac{da_j}{dz}\right| \ll \beta_j^2 |a_j|. \tag{11.23}$$

In other words, the SVEA[3] claims that coupling between the instantaneous eigenmodes is weak, so that the $a_j(z)$ vary slowly compared to $\exp(i\beta_j z)$. I define the coupling coefficients $\tilde{\kappa}_{ij}$ as

$$\tilde{\kappa}_{ij} \equiv \langle \xi_i^z | \frac{\partial}{\partial z} | \xi_j^z \rangle. \tag{11.24}$$

An analogous version of the coupling coefficients in the case of vector propagation are given in [JJJ10]. Under the SVEA, the coupled-mode equations can be written as

$$\frac{da_i}{dz} = -\frac{1}{2}\frac{d\ln\beta_i}{dz} a_i - \sum_j \frac{\beta_j}{\beta_i} e^{i\int_0^z (\beta_j - \beta_i) dz'} \tilde{\kappa}_{ij} a_j. \tag{11.25}$$

In general, the above form is what should be used when modelling slowly varying waveguides, though the $d\ln\beta_i/dz$ term may be negligible in the case of weak guidance. When $A(z)$ is self-adjoint (or Hermitian, in our case), the eigenbasis can be chosen to be purely real, and $\tilde{\kappa}_{ij}$ is a real antisymmetric matrix.

For power transfer between modes i and j to occur, $\tilde{\kappa}_{ij}$ must be large and $\beta_j - \beta_i$ must be small. If the latter condition is not met, the sign of the coupling will rapidly flip with z and power transfer will be negligible.

---

[3] The definitions of the paraxial approximation and SVEA are somewhat muddled and occasionally used interchangeably. In this treatment I specify the former as applying to the instantaneous eigenmodes and latter to the mode amplitudes, but this may not be conventional.

11.3.3 Adiabatic approximation

In the adiabatic approximation, we assume that the coupling coefficients are negligible, so that the modal evolution is decoupled:

$$\frac{da_i}{dz} = -\frac{1}{2}\frac{d\ln\beta_i}{dz}a_i. \tag{11.26}$$

This approximation is analogous to the adiabatic approximation of quantum mechanics (e.g. [TTB21] who applies this approximation to waveguide couplers); also, see §11.4.3. The general solution is

$$a_i(z) = a_i(0)\sqrt{\frac{\beta_i(0)}{\beta_i(z)}} \tag{11.27}$$

and the full solution for a propagating wavefront is

$$|\Psi(z)\rangle = \sum_j a_j(0)\sqrt{\frac{\beta_j(0)}{\beta_j(z)}}e^{i\int_0^z \beta_j(z')dz'}|\xi_j^z\rangle \tag{11.28}$$

which is reminiscent of the Wentzel–Kramers–Brillouin (WKB) approximation for a particle in a slowly varying potential; see §11.4.1. While the SVEA implicitly assumes that $|\tilde{\kappa}_{ij}| \ll \beta_i$ for all $i \neq j$, the condition for adiabaticity is even more stringent [BF28]:

$$|\tilde{\kappa}_{ij}| \ll |\beta_i - \beta_j|. \tag{11.29}$$

This condition[4] can be loosely motivated by noting that the antiderivative of $\exp\left[i\int_0^z(\beta_j - \beta_i)dz'\right]$, which appears in the coupled-mode equations 11.25, is of order $1/|\beta_i - \beta_j|$. In terms of the refractive index profile, the adiabatic criterion can be recast using perturbation theory (e.g.

---

[4] While not strictly sufficient, this condition holds in "most" cases, and is often used in quantum mechanics [Com09].

§11.4.2) as

$$k^2 \left| \frac{\langle \xi_i^z | \frac{\partial n^2}{\partial z} | \xi_j^z \rangle}{\beta_i^2 - \beta_j^2} \right| \ll |\beta_i - \beta_j| \ ; \ i \neq j. \tag{11.30}$$

In the case of degeneracy, the adiabatic approximation breaks down: *there is no rate of change slow enough to prevent cross-coupling between degenerate modes.*

11.3.4 Example applications of the varying-basis CMT

To reiterate, the prior analysis suggests that light moves through a slowly-varying waveguide in co-propagating modes which cross-couple depending on mode shapes and propagation constants. The evolution of the amplitudes of these modes is controlled by a system of ordinary differential equations. This allows us to build an intuition for how light propagates through some astrophotonic devices. I provide some examples:

2 × 2 **Directional coupler** As shown in Figure 11.1.1a, this device has two single-moded channels embedded in an optical substrate. The two channels are initially well-separated, then are slowly brought closer together, kept close over some "coupling length", and then slowly brought apart. When light is injected into one input, the device acts like a beamsplitter; the power ratio of the split is determined by the coupling length. One explanation of how this splitting works is given in §11.A, using the fixed-basis CMT, but the varying-basis CMT provides a different perspective. When the two channels are very far apart, the instantaneous eigenmodes are indeed just the fundamental modes of each channel. However, when the two channels are close, the instantaneous eigenmodes look like the symmetric and antisymmetric combinations of each channel's fundamental mode. When light is launched in one input of the coupler, it is gradually split between the symmetric and antisymmetric mode. Then, as the light propagates down the coupling length, power is observed to oscillate between the two channels: this is due to the beating of the symmetric and antisymmetric mode, which have different propagation constants. The oscillation slows and eventually stops when the

channels are separated again. The relevant approximations are paraxiality and the SVEA; an example device is numerically simulated in §11.5.3.

Mode-selective PL A mode-selective PL couples light from each mode at its entrance into a separate single-mode output. In other words, a mode-selective PL has no cross-coupling between its modes, and thus should satisfy the adiabatic condition, along with paraxiality and the SVEA. One way to ensure negligible cross-coupling, mentioned in §11.3.2, is to give each mode a sufficiently distinct propagation coefficient β_i. [LFS14] achieves this by forming the PL from single-mode fibers with different core sizes; another option would be differentiate the cores in refractive index. Conversely, a non-mode-selective PL uses cores which are identical in size and refractive index. As a result, there is a region of the PL (particularly towards the larger end of the taper) where the co-propagating PL modes become degenerate and can rapidly cross-couple. In other words, the larger end of an N-port PL acts like an $N \times N$ directional coupler; this is also why the propagation characteristics of fabricated PLs are difficult to control. A non-mode-selective PL is simulated in §11.5.4.

Fiber Bragg gratings Briefly, I mention that the CMT can be extended to structures such as fiber Bragg gratings, which use a rapid variation within the refractive index structure of an optical fiber to reflect narrow spectral bands. To apply the CMT, we must include both the forward and backward propagating waves when expanding the electric field within the waveguide; see [McC00]. Note that despite the rapid variation, a fiber Bragg grating may still be thought of as slowly varying in the sense of the SVEA due to the low index contrast of the variation.

11.4 Connections to quantum mechanics

Several studies have illustrated how certain optical waveguides have quantum mechanical counterparts: for instance, the PL and the Kronig-Penney model for electrons in a periodic potential well [LAB10, BEL11, Kit04]. From another perspective, the fixed-basis CMT is essentially an application of time-dependent perturbation theory [CHC21], with the time t replaced with longitudinal coordinate z. In this section, I demonstrate that the varying-basis CMT is also connected to the quantum mechanics. In §11.4.1, I show that the first-order WKB approximation and the CMT under the adiabatic approximation give the same solution when applied to the Helmholtz equation. In §11.4.2, I review how time-independent perturbation theory can be used to compute the coupling coefficients $\tilde{\kappa}_{ij}$ [SWZ16]. Lastly, in §11.4.3, I show how geometric phase, which appears in quantum mechanics under the adiabatic approximation, is a natural result of the CMT.

11.4.1 WKB approximation

The WKB approximation is closely related to the varying-basis CMT, and in fact motivates the ansatz in equation 11.19. To see this, I present a brief overview of the WKB approximation. Suppose we have a differential equation

$$\frac{dy}{dx} = f(x, y). \tag{11.31}$$

The above can be recast in integral form as

$$\begin{aligned} y(x) &= \int_{x_0}^{x} f(x', y(x'))dx' + C \\ C &= y(x_0). \end{aligned} \tag{11.32}$$

To solve this integral equation we can provide an initial guess for $y(x)$, denoted $y_0(x)$. Substituting this guess back into the integral equation gives a more accurate guess $y_1(x)$. This process can be repeated:

$$y_n(x) = \int_{x_0}^{x} f(x', y_{n-1}(x'))dx' + C. \tag{11.33}$$

In both quantum mechanics and optics (under the scalar approximation), we deal the Helmholtz equation:

$$\frac{d^2\psi}{dx^2} + k^2(x)\psi = 0. \tag{11.34}$$

For $k(x) = k$ constant, solutions are the plane waves: $\psi(x) = \exp(\pm ikx)$. If $k(x)$ is allowed to vary, then we might replace kx with $\int k(x)dx$. This is gives us the 0th order approximation

$$\psi_0(x) = e^{\pm i\int k(x)dx}. \tag{11.35}$$

which is valid for $|dk/dx| \ll k^2$. I now apply the method of successive approximations to solve the wave equation 11.34. Assume a solution of the form $\psi(x) = \exp\left[iS(x)\right]$. Plugging this into equation 11.34 yields

$$iS'' - (S')^2 + k^2 = 0. \tag{11.36}$$

Manipulate the above to derive the series relation for S_n, and insert the 0th order approximation. This yields the first-order solution

$$\psi_1(x) \propto \frac{1}{\sqrt{k(x)}} e^{\pm i\int k(x)\,dx}. \tag{11.37}$$

For more details, see appendix 11.B. This form is the same as what was derived in §11.2 under the adiabatic approximation, and justifies our ansatz 11.19.

11.4.2 Time-independent perturbation theory

Besides using finite differences, the varying-basis coupling coefficients $\tilde{\kappa}_{ij}$ can also be estimated from the fixed-basis coefficients κ_{ij} using time-independent perturbation theory — or in our case, z-independent perturbation theory. Define the refractive index structure of our waveguide as $n(x, y, z)$. We can treat the difference in $n^2(x, y, z)$ between two values of z, separated by some small δz, as a perturbation $\delta\epsilon$:

$$\delta\epsilon = n^2(x, y, z + \delta z) - n^2(x, y, z). \tag{11.38}$$

Following perturbation theory, we treat $k^2\delta\epsilon$ as a perturbation to the operator $A \equiv \nabla_\perp^2 + k^2 n^2(x, y, z)$ in the eigenvalue equation $A|\xi_j^z\rangle = \beta_j^2(z)|\xi_j^z\rangle$. The first-order approximation of the perturbed eigenmode is [SN20]

$$|\xi_j^{z+\delta z}\rangle \approx |\xi_j^z\rangle + \sum_{i \neq j} \frac{\langle\xi_i^z|k^2\delta\epsilon|\xi_j^z\rangle}{\beta_j^2 - \beta_i^2}|\xi_i^z\rangle. \tag{11.39}$$

Therefore, the coupling coefficient matrix can be approximated as

$$\begin{aligned}\langle\xi_k^z|\frac{\partial}{\partial z}|\xi_j^z\rangle &\approx \frac{1}{\delta z}\sum_{i \neq j} \frac{\langle\xi_i^z|k^2\delta\epsilon|\xi_j^z\rangle}{\beta_j^2 - \beta_i^2}\langle\xi_k^z|\xi_i^z\rangle \\ &= \begin{cases} \dfrac{\langle\xi_k^z|k^2\delta\epsilon/\delta z|\xi_j^z\rangle}{\beta_j^2 - \beta_k^2}, & k \neq j \\ 0, & k = j. \end{cases}\end{aligned} \tag{11.40}$$

In the limit $\delta z \to 0$, we have $\delta\epsilon/\delta z \to \partial n^2/\partial z$. This formula links the coupling matrix of the varying-basis CMT in §11.3 to the that of the fixed-basis CMT in §11.2. Degeneracy can be treated if it is lifted by the perturbation, using degenerate perturbation theory. However, a degeneracy enforced by a symmetry in the waveguide design may not be broken by the perturbation. The resulting coupling between degenerate modes can be seen as a result of a non-abelian geometric phase, and is discussed further in §11.4.3.2.

11.4.3 Adiabatic approximation: geometric phase

In this section I show that the varying-basis CMT naturally includes geometric phase, which was first observed in studies of polarized light by Pancharatnam [Pan56] and has been widely studied in the context of quantum mechanics under the adiabatic approximation (e.g. the

Aharanov-Bohm effect [AB59, Ber84]). Several well-known examples of geometric phase in optics include the spin redirection phase, in which the polarization orientation is rotated by out-of-plane propagation [Ros84]; the Gouy phase, which occurs for a Gaussian beam propagating through focus [Sub95]; and the Pancharatnam-Berry (PB) phase, which can occur in birefringent materials. Taking the PB phase as an example, I first show how geometric phase appears in the CMT, and then discuss how non-adiabatic evolution produces the non-abelian Wilczek-Zee phase.

11.4.3.1 Pancharatnam-Berry phase

Following §11.3, the coupled-mode equations are

$$\frac{da_i}{dz} = -\frac{1}{2}\frac{d\ln\beta_i}{dz}a_i - \sum_j \frac{\beta_j}{\beta_i}\langle\zeta_i^z|\frac{\partial}{\partial z}|\zeta_j^z\rangle a_j \tag{11.41}$$

where the eigenbasis has been temporarily modified as

$$|\zeta_j^z\rangle \equiv e^{i\int_0^z \beta_j\, dz'}|\xi_j^z\rangle. \tag{11.42}$$

Using the above form and splitting the summation gives

$$\frac{da_i}{dz} = -\frac{1}{2}\frac{d\ln\beta_i}{dz}a_i - \sum_{j\neq i} \frac{\beta_j}{\beta_i}\langle\zeta_i^z|\frac{\partial}{\partial z}|\zeta_j^z\rangle a_j - \langle\xi_i^z|\frac{\partial}{\partial z}|\xi_i^z\rangle a_i. \tag{11.43}$$

Under the adiabatic approximation, all off-diagonal terms are set to 0, and the coupled equations can be solved via the matrix exponential:

$$\begin{aligned} a_i(z) &= \exp\left[-\int_0^z \left(\frac{1}{2\beta_i}\frac{d\beta_i}{dz'} + \langle\xi_i^{z'}|\frac{\partial}{\partial z'}|\xi_i^{z'}\rangle\right) dz'\right] \\ &\propto \frac{1}{\sqrt{\beta_i(z)}}\exp\left[-\int_0^z \langle\xi_i^{z'}|\frac{\partial}{\partial z'}|\xi_i^{z'}\rangle dz'\right]. \end{aligned} \tag{11.44}$$

The matrix exponential is applicable here because the adiabatic terms in equation 11.43 can be represented as diagonal matrices, which commute. Plugging $a(z)$ back into the original ansatz 11.19, the overall wavefront evolves as

$$|\Psi_j(z)\rangle \propto \frac{1}{\sqrt{\beta_j(z)}} e^{i\phi_j(z)} e^{i\theta_j(z)} |\xi_j^z\rangle \tag{11.45}$$

where $\phi_j(z)$ and $\theta_j(z)$ are defined as

$$\begin{aligned} \phi_j(z) &\equiv i \int_0^z \langle \xi_i^{z'} | \frac{\partial}{\partial z'} | \xi_i^{z'} \rangle dz' \\ \theta_j(z) &\equiv \int_0^z \beta_j(z') dz'. \end{aligned} \tag{11.46}$$

It can be shown that ϕ_j is real.[5] The term $\theta_j(z)$ is the dynamical phase, while $\phi_j(z)$ is identified as the PB phase when the system evolves in a loop, i.e. when $|\xi_j^0\rangle = |\xi_j^z\rangle$. In quantum mechanics, $\phi_j(z)$ is often expressed as a contour integral in some parameter space with parameters λ_μ. Then, $|\xi_i(z)\rangle = |\xi_i(\lambda_\mu(z)\rangle$, and the PB phase is

$$\phi_j(z) = i \oint_\gamma \sum_\mu \langle \xi_j | \frac{\partial}{\partial \lambda_\mu} | \xi_j \rangle d\lambda_\mu \equiv \oint_\gamma \mathcal{A} \cdot d\lambda \tag{11.47}$$

where $\gamma(z) = (\lambda_1(z), \lambda_2(z), ...)$ denotes a closed path in λ_μ parameter space; for the last definition, I have switched to vector notation. The vector $\mathcal{A}$ is the Berry or Aharanov-Anandan connection. The curl of the connection $\mathcal{A}$ is the Berry curvature, denoted $\mathbf{\Omega}$.[6]

When the eigenmodes $|\xi_j^z\rangle$ are real-valued, as is the case under the scalar approximation (neglecting polarization), the Berry connection is 0 and no PB phase can be accrued.[7] Thus,

---

[5] The operator $\partial/\partial z$ is anti-Hermitian, so the diagonal of the operator in matrix representation (the bracketed term in equation 11.46) is purely imaginary. This can also be proven using the product rule.

[6] The connection and curvature can be alternatively written using differential forms. Specifically, we may write the connection as the one-form $\mathcal{A} = \langle \xi_j | d | \xi_j \rangle$, where d is the exterior derivative. The connection is the two-form $\Omega = d\mathcal{A}$.

[7] For a mode ξ, if we have $\xi = \xi^*$ then $\int \xi^* \frac{\partial}{\partial z} \xi \, dxdy = \int \xi \frac{\partial}{\partial z} \xi \, dxdy = \frac{1}{2} \frac{\partial}{\partial z} \int \xi^2 \, dxdy = 0$.

to see PB phase, the eigenvectors must be complex — equivalently, the operator $A(z)$ cannot be self-adjoint. This is the case for the propagation of polarized plane waves through a slowly varying birefringent medium. For instance, represent a polarized plane wave $|P\rangle$ in terms of its normalized Jones vector $|p\rangle$:

$$\begin{aligned} |P\rangle &= e^{i(\beta z - \omega t)} e^{i\gamma} |p\rangle \\ |p\rangle &\equiv \begin{bmatrix} \cos\chi/2 \\ e^{-i\psi}\sin\chi/2 \end{bmatrix}. \end{aligned} \tag{11.48}$$

Here, $\gamma, \psi \in [0, 2\pi]$, $\chi \in [0, \pi]$, and β is the propagation constant. The parametrization for $|p\rangle$ is also used for rank-1 spinors in quantum mechanical two-state systems, and applies to anything that can be represented as a unit vector in $\mathbb{C}^2$. The angle γ represents absolute phase; for visualization, the "azimuthal" angle ψ and "polar" angle χ are often mapped onto the surface of the so-called Poincaré sphere[8], which represents all possible states of the polarization ellipse over its surface ([AD22a] provides a nice explanation).

Following studies of quantum-mechanical two-state systems, the propagation of a polarized plane wave $|P\rangle$ through a homogeneous birefringent material that slowly varies with z can be represented by the differential equation

$$\frac{\partial}{\partial z}|p\rangle = \mathfrak{J}(z)|p\rangle \tag{11.49}$$

where $\mathfrak{J}(z)$ is a z-dependent anti-Hermitian matrix, i.e. a linear combination of the matrices $i\sigma_j$, where the σ_j are the Hermitian Pauli spin matrices. We can motivate the above by noting that the Jones matrices, which represent the optical transfer matrices for homogeneous slabs of birefringent material, are a representation of the special unitary group SU(2), whose Lie algebra is $\mathfrak{su}(2)$, which in turn has the basis $i\sigma_j$. The evolution of $|p\rangle$ is parameterized by the path $\gamma(z) = [\psi(z), \chi(z)]$.

---

[8] Or the Bloch sphere, in the study of two-state quantum mechanical systems.

Comparing the above equation with the form assumed in §11.3, we identify $D(z) = \partial/\partial z$ and $A(z) = -\mathfrak{J}(z)$; the differential equation governing the evolution of the wavefront is first-order, not second-order. This simplifies the application of the CMT, removing the "WKB" term from the coupled-mode equations 11.41; our prior results for geometric phase still hold. The Berry connection for polarization states, $\mathcal{A}_{\text{pol}}$, has components

$$\begin{aligned} \mathcal{A}_{\text{pol},\chi} &= 0 \\ \mathcal{A}_{\text{pol},\psi} &= \frac{\sin^2 \chi/2}{\sin \chi}. \end{aligned} \tag{11.50}$$

The PB phase can be computed using Stokes' theorem:

$$\phi = \oint_{\gamma} \mathcal{A} \cdot d\lambda = \iint_{\mathcal{S}} \left(\vec{\nabla} \times \mathcal{A} \right) \cdot d\mathcal{S} = \frac{1}{2} \iint_{\mathcal{S}} d\mathcal{S}. \tag{11.51}$$

Here, $\mathcal{S}$ denotes the surface of the unit 2-sphere (the Poincaré sphere) bound by the closed curve γ, and $d\mathcal{S}$ is an area element of the sphere; the curvature is $\Omega = 1/2$.[9] Thus, we recover that the PB phase accumulated by a polarized plane wave is half the solid angle enclosed by the path it takes over the Poincaré sphere [AD22b].

The PB phase can be more generally understood as a holonomy in the context of fiber bundles [CGF22]. The full polarization state $\exp[-i\psi]|p\rangle$ "lives" in a unit 3-sphere, denoted S^3, which can be represented as a fiber bundle whose base space is the unit 2-sphere S^2 and whose fiber space is the unit circle S^1 through the Hopf fibration[10] [Urb03]. The base space — the Poincaré sphere — records the orientation and eccentricity of the polarization ellipse, while the fibers record absolute phase. Closed paths on the base space may be lifted by the fiber bundle connection — the Berry connection — to open paths in the total space of polarization states S^3. These open paths start and end on the same fiber, and the fiber-

---

[9] This can be derived using the definition of the connection from equation 11.47 and the formula for the curl in spherical coordinates.

[10] Equivalently, the fiber space is the unitary group $U(1)$.

wise "distance" between the start and end points of such paths encodes the geometric phase accumulated by cyclic evolution about the Poincaré sphere.

11.4.3.2 Degenerate adiabatic approximation: Wilczek-Zee phase

Though wavefronts in the scalar approximation accumulate no PB phase, they can accumulate Wilczek-Zee (WZ) phase so long as the eigenbasis is at least partially degenerate. The WZ phase generalizes the PB phase, which was derived under the adiabatic approximation, to degenerate eigenspaces [RO14]. To see how the WZ phase appears in the CMT, consider solving a simple form of the coupled-mode equations 11.41, assuming degeneracy of all eigenmodes, dropping the WKB term, and discretizing the interval $[0, z]$ at N evenly-spaced values $z_n = n\,\Delta z$, with $n = 0, 1, 2 \ldots N$ and $z_N = z$. Denote the series of matrices $M_{ij}^{z_n}$ as

$$M_{ij}^{z_n} \equiv \langle \zeta_i^z | \frac{\partial}{\partial z} | \zeta_j^z \rangle \Big|_{z=z_n}. \tag{11.52}$$

The M^{z_n} form a non-commuting group. The solution to the coupled-mode equations can be cast as the limit:

$$\begin{aligned} \mathbf{a}(z) &= \lim_{\Delta z \to 0} \left(e^{M^{z_N} \Delta z} \ldots e^{M^{z_1} \Delta z} e^{M^{z_0} \Delta z} \right) \mathbf{a}(0) \\ &\equiv \mathcal{P} \left\{ e^{\int_0^z M(z') dz'} \right\} \mathbf{a}(0) \end{aligned} \tag{11.53}$$

where I have used the notation $\mathcal{P}\{\exp\{\int_0^z M(z')dz'\}\}$ to denote the path-ordered exponential of $M(z)$; in general, we cannot take the usual matrix exponential because the M_{ij}^z do not commute, and so we must take our matrix exponential in short steps of Δz in a particular order — hence why the WZ phase is termed non-abelian. The WZ phase arises from the path-ordered exponential of $M_{ij}(z)$, and controls the variation in the phases and relative weights of the degenerate eigenmodes. In the language of fiber bundles, the WZ phase can be seen as the holonomy of a vector bundle, while abelian geometric phases are specifically holonomies of line bundles [ZKF23].

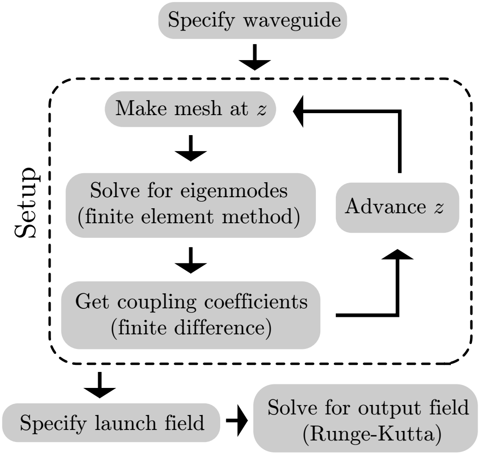

Figure 11.5.1: Block diagram of the `cbeam` package. The loop within the dashed lines takes the bulk of the simulation time but only needs to be run once per waveguide.

Geometric phases both abelian and non-abelian are of particular interest in quantum computing due their anticipated resilience against noise [ZKF23], and have been experimentally realized by manipulating qubit systems [AFJ13]. While not completely analogous, there is a similar need for noise-robust and achromatic optical interferometry, especially for the purpose of astronomical nulling. Achromatic nullers using rotational-shearing interferometers have been implemented [SWH99, TKT05]; it would be interesting to see if a photonic equivalent could also be made by applying geometric phase manipulations within a waveguide. As a monolithic waveguide, such a device would be more stable and compact.

11.5 Numerical implementation: `cbeam`

In this section I present a lightweight numerical implementation of the varying-basis CMT, written in Python and Julia, called `cbeam` [Lin24]. Compared to alternate packages which use the finite-difference beam propagation method (FDBPM), `cbeam` can be more than an order of magnitude faster; see §11.5.4. Figure 11.5.1 gives a graphical overview of `cbeam`'s structure. A user first specifies a waveguide; `cbeam` then adaptively computes the following as a function of z:

1. the instantaneous propagation coefficients $\beta_j(z)$
2. the instantaneous eigenmodes $|\xi_j^z\rangle$
3. the coupling coefficients $\tilde{\kappa}_{ij}$

This setup step is run once per waveguide. Items 1 and 2 are computed using the finite-element method, though care must be taken in the case of degeneracy — see §11.5.1. Item 3 is computed by interpolating through the modes using cubic splines and evaluating the derivative; though a correction must be made — see §11.5.2. With these items and a launch field in hand, propagation through the waveguide becomes an initial value problem, which `cbeam` solves using an adaptive Runge-Kutta method. In §11.5.3 and §11.5.4, I use `cbeam` to simulate a directional coupler and 6-port photonic lantern. For more detailed examples, refer to the Github repository for `cbeam` and online documentation.[11]

11.5.1 Eigenmode continuity

When eigenvalues are degenerate, the eigenbasis is not well-defined. Under such conditions, the instantaneous eigenmodes returned by a finite-element solver may not evolve continuously with z. To enforce continuity, `cbeam` models the z-dependent variation in waveguide geometry by applying a continuous spatial transformation to the nodes of the finite element mesh. This transformation is not specific to a single type of waveguide, and works by expanding or contracting regions of the finite element mesh so that triangles remain approximately similar. For example, this transformation can model a directional coupler by shrinking and expanding the space between a set of single-mode channels. Additionally, if the waveguide admits a degenerate subspace of eigenmodes throughout its entire length, `cbeam` uses least-squares to fix a degenerate eigenbasis and prevent it from rotating. The procedure is as follows: when going from $z \rightarrow z + \delta z$, apply a unitary transformation to the degenerate modes at $z + \delta z$

---

[11] Located here: https://github.com/jw-lin/cbeam/.

to "match" the modes at z. The transformation can be found via least-squares using the singular value decomposition (SVD). Define the degenerate eigenbases at z and $z + \delta z$ as $|\xi_k^z\rangle$ and $|\xi_k^{z+\delta z}\rangle$, respectively. Next, define the overlap matrix C as

$$C_{ij} \equiv \langle \xi_i^{z+\delta z} | \xi_j^z \rangle. \tag{11.54}$$

Construct the unitary matrix Q, which rotates the eigenbasis, from the SVD components of C:

$$\begin{aligned} C &= USV^T \\ Q &\equiv VU^T. \end{aligned} \tag{11.55}$$

Denote M a matrix whose columns are the eigenmodes computed at $z + \delta z$, i.e. $|\xi_j^{z+\delta z}\rangle$. The new basis should be chosen as the columns of M' with

$$M' = MQ. \tag{11.56}$$

Another complication is that eigenvalues of a general operator $A(z)$ may cross as a function of z, requiring a method to track eigenvalues through crossings.[12] In tandem with the above degeneracy correction, `cbeam` also attempts to re-sort modes so that they can be tracked through crossings.

11.5.2 Estimate of the coupling matrix

To compute the coupling matrix κ_{ij}, `cbeam` interpolates the eigenmode profiles through z and evaluates the derivative of the interpolation directly. However, because the eigenmodes are only defined on the mesh nodes, and because the mesh nodes move with the transformation,

---

[12] There are theorems called "avoided crossing" or "anticrossing" theorems from quantum mechanics which establish when crossing may occur [LL76].

the derivative of the interpolation function does not precisely give the partial derivative with respect to z. Denoting the position of a mesh node as $\mathbf{p}(z) = (x(z), y(z))$, the z partial derivative of a mode $|\xi_j\rangle$ at that point is

$$\frac{\partial}{\partial z}|\xi_j\rangle = \frac{d}{dz}|\xi_j\rangle - \vec{\nabla}_\perp|\xi_j\rangle\Big|_{\mathbf{p}} \cdot \frac{d\mathbf{p}}{dz} \tag{11.57}$$

where $\frac{d}{dz}|\xi_j\rangle$ is the derivative of $|\xi_j\rangle$ that "comoves" with the mesh point $\mathbf{p}$; this is what is measured when taking the derivative of the mode's interpolation function. To compute the transverse gradient of $|\xi_j\rangle$, `cbeam` uses the "average gradient on star method" [MLP19].

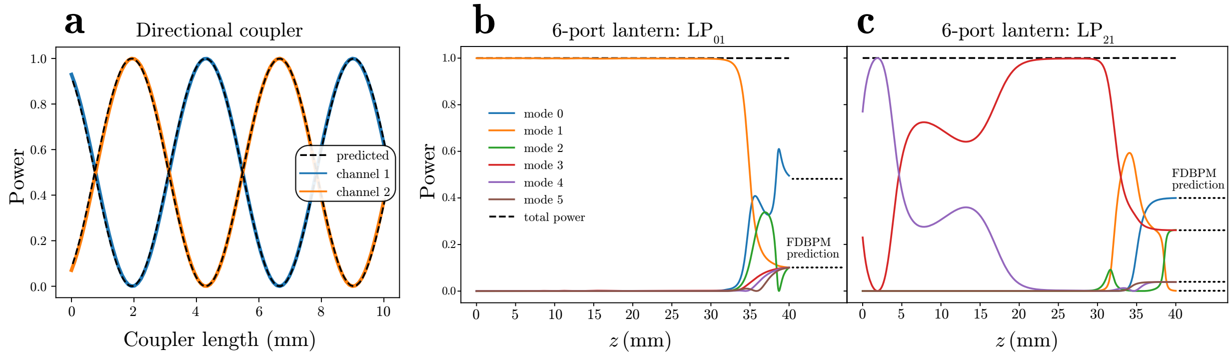

Figure 11.5.2: **a.** Output powers of a 2 × 2 directional coupler as a function of coupling length. The dashed curves show the prediction from theory. **b, c.** The powers of the instantaneous eigenmodes in a standard 6-port PL as a function of z, when light is launched in the LP_{01} mode (**b**) and LP_{21} mode (**c**). The dotted lines past $z = 40$ mm denotes the predicted lantern output from `lightbeam`, a waveguide simulation package which uses the finite difference beam propagation method.

11.5.3 Example: directional coupler

Figure 11.5.2 (left) shows the simulated outputs of a 2 × 2 directional coupler constructed from two single-moded circular channels which bend toward and away from each other along arctangent paths. The output power in each channel is plotted as a function of the length of the coupling region. These powers oscillate like $\sin^2$ and $\cos^2$, as expected, while the oscillation period matches a common empirical approximation [TT86], shown by the dashed curves. I include this approximation in §11.A because some reference texts [Agr21, Kha04] contain a typographical error.

11.5.4 Example: photonic lantern

The right panels of Figure 11.5.2 show the powers of the instantaneous eigenmodes as a function of z, when light is launched in the LP_{01} and LP_{21} modes. The modelled PL has a 5-fold rotational symmetry, with one centrally located core and five outer cores located at the points of a regular pentagon, and thus we expect that light launched in a rotationally symmetric mode should maintain a symmetric distribution at the lantern output. `cbeam` recovers this behavior: launching light into LP_{01} yields equal power in all outer cores at the lantern output. Symmetry is not maintained throughout the entirety of the lantern because of the freedom of eigenbasis choice in presence of degeneracy.

The right panels of Figure 11.5.2 also show the predicted lantern port powers from `lightbeam`, a Python-based numerical package for waveguide modelling analogous to `cbeam`, which uses the FDBPM instead of the CMT. When propagating the LP_{11} modes (not shown) and LP_{21} modes, `cbeam` and `lightbeam` produce output power vectors that are consistent to within 0.2% total error. For the radially symmetric modes LP_{01} and LP_{02}, the error is $\sim 1.5\%$. While far from a rigorous explanation, I provide one possible reason for this discrepancy. The CMT predicts small ($\lesssim 1\%$) amount of oscillatory cross-coupling at the lantern entrance between the LP_{01} and LP_{02} modes, which notably have the largest difference in effective index in the 6-mode group supported by the PL. This oscillatory coupling may be more difficult to model with FDBPM methods, which assign a single reference index for all propagating waves, and which accumulate more error when the chosen index chosen deviates from the effective indices of the eigenmodes [Ped15].

Finally, I note runtime comparisons: for the fiducial 6-port PL, a single propagation in `lightbeam` takes around an hour of single-threaded computation time, and determination of the waveguide transfer matrix takes 6 propagations. With `cbeam`, the entire waveguide is modelled in around 1 minute, and subsequent propagations of any guided wavefront through the lantern take $\sim$ 1s. Therefore, in this example `cbeam` is almost 400× faster.

Appendix - Chapter 11

11.A Two-mode coupling

Consider a waveguide such as a directional coupler supporting only two modes whose eigenvalues are β_1^2 and β_2^2. The two coupled-mode equations governing waveguide propagation are

$$\frac{d}{dz}\begin{bmatrix} a_1 \\ a_2 \end{bmatrix} = i \begin{bmatrix} 0 & \kappa_{12}e^{i(\beta_2-\beta_1)z} \\ \kappa_{21}e^{i(\beta_1-\beta_2)z} & 0 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix}. \tag{11.58}$$

To solve the above, make the following transformation:

$$\begin{aligned} a_1(z) &= b_1(z)e^{-i\delta z} \\ a_2(z) &= b_2(z)e^{+i\delta z}. \end{aligned} \tag{11.59}$$

Here, $\delta \equiv (\beta_1 - \beta_2)/2$. Under this transformation, the coupled system becomes:

$$\begin{aligned} \frac{d}{dz}\begin{bmatrix} b_1 \\ b_2 \end{bmatrix} &= \begin{bmatrix} i\delta & i\kappa_{12} \\ i\kappa_{21} & -i\delta \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \end{bmatrix} \\ &\equiv B \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}. \end{aligned} \tag{11.60}$$

Define $\kappa_e = \sqrt{\delta^2 + \kappa_{12}\kappa_{21}}$. The solution to the coupled system is

$$\mathbf{b}(z) = e^{zB}\mathbf{b}_0 \tag{11.61}$$

for initial condition $\mathbf{b}_0$.

Assuming $\mathbf{b}_0 = [1, 0]$ (light is launched only in mode 1), the power in mode 2 varies as

$$\begin{aligned} b_2(z) &= \frac{\kappa_{21} e^{i\kappa_e z} - \kappa_{21} e^{-i\kappa_e z}}{2\kappa_e} \\ &= \frac{i\kappa_{21}}{\kappa_e} \sin\left(\kappa_e z\right). \end{aligned} \tag{11.62}$$

Notice that up to 100% of the power may be transferred to mode 2 if $\delta = 0$: this is resonant coupling. This limit goes to 0 as the effective index difference increases. The maximum power transfer first occurs at $z = L_c = \pi/2\kappa_e$, defining the so-called coupling length L_c. The above shows that degenerate modes are the most susceptible to an exchange of power.

For a symmetric directional coupler composed of two circular single-mode cores separated by a distance d, we may take as our idealized eigenmodes the fundamental modes supported by each core *when ignoring the other core*. The refractive index of the other core is treated as a perturbation $\delta\epsilon$ in the sense of equation 11.8. The coupling coefficient has the following approximation [TT86]:

$$\begin{aligned} \kappa &\equiv \kappa_{21} = \kappa_{12} \\ &\approx \frac{\pi V}{2kn_{\text{clad}}a^2} \exp\left[-2.3026\left(a + b\tilde{d} + c\tilde{d}^2\right)\right] \end{aligned} \tag{11.63}$$

where V is the normalized frequency parameter

$$V = ka\sqrt{n_{\text{core}}^2 - n_{\text{clad}}^2}. \tag{11.64}$$

Here k is the free-space wavenumber; n_{core} and n_{clad} are the core and cladding indices of the coupler; a is the radius of the single-mode cores; and $\tilde{d} \equiv d/a$ is the normalized center-to-center spacing. The coefficients a, b, and c are:

$$\begin{aligned} a &= 2.2926 - 1.5910V + 0.1668V^2 \\ b &= -0.3374 + 0.5321V - 0.0066V^2 \\ c &= -0.0076 - 0.0028 + 0.0004V^2. \end{aligned} \tag{11.65}$$

The sign of the last term in c is sometimes erroneously flipped in the literature. This approximation is accurate within 1% for $1.5 \leq V \leq 2.5$ and $2.0 \leq \tilde{d} \leq 4.5$.

11.B The WKB method from successive approximations

To derive the first-order WKB solution, begin with the differential equation 11.36. I manipulate this equation to get the first-order derivative on the left-hand side, and then apply successive approximations as follows:

$$\begin{aligned} S'(x) &= \pm\sqrt{iS''(x) + k^2(x)} \\ S_n(x) &= \pm\int \sqrt{iS''_{n-1} + k^2} dx \\ S_0(x) &= \pm\int k(x)\, dx \\ S_1(x) &= \pm\int \sqrt{\pm ik' + k^2} dx \\ &\approx \pm\int k\left[1 \pm \frac{i}{2}\frac{k'}{k^2}\right] dx \\ &= \pm\int k\, dx + \frac{i}{2}\ln(k) + C \end{aligned} \tag{11.66}$$

where going from line 4 to line 5 I made use of the slowly varying condition $|dk/dx| \ll k^2$ to apply a Taylor expansion. Therefore, the first-order solution is

$$\psi_1(x) = e^{iS_1(x)} \propto \frac{1}{\sqrt{k(x)}} e^{\pm i\int k(x)\, dx}. \tag{11.67}$$

REFERENCES

[AB59] Y. Aharonov and D. Bohm. “Significance of Electromagnetic Potentials in the Quantum Theory.” *Phys. Rev.*, **115**:485–491, Aug 1959.

[AD22a] Miguel A. Alonso and Mark R. Dennis. “The geometric phase made simple.” *Photoniques*, (116):58–63, 2022.

[AD22b] Alonso, Miguel A. and Dennis, Mark R. “The geometric phase made simple.” *Photoniques*, (116):58–63, 2022.

[AFJ13] A. A. Abdumalikov Jr, J. M. Fink, K. Juliusson, M. Pechal, S. Berger, A. Wallraff, and S. Filipp. “Experimental realization of non-Abelian non-adiabatic geometric gates.” *Nature*, **496**(7446):482–485, Apr 2013.

[Agr19] Govind P. Agrawal. “Chapter 14 - Multimode fibers.” In Govind P. Agrawal, editor, *Nonlinear Fiber Optics (Sixth Edition)*, pp. 621–683. Academic Press, sixth edition edition, 2019.

[Agr21] Govind P. Agrawal. “Chapter 2 - Directional couplers.” In Govind P. Agrawal, editor, *Applications of Nonlinear Fiber Optics (Third Edition)*, pp. 57–107. Academic Press, third edition edition, 2021.

[BEL11] J. Bland-Hawthorn, S. C. Ellis, S. G. Leon-Saval, R. Haynes, M. M. Roth, H.-G. Löhmannsröben, A. J. Horton, J.-G. Cuby, T. A. Birks, J. S. Lawrence, P. Gillingham, S. D. Ryder, and C. Trinh. “A complex multi-notch astronomical filter to suppress the bright infrared sky.” *Nature Communications*, **2**(1):581, Dec 2011.

[Ber84] Michael Victor Berry. “Quantal phase factors accompanying adiabatic changes.” *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, **392**(1802):45–57, 1984.

[BF28] M. Born and V. Fock. “Beweis des Adiabatensatzes.” *Zeitschrift für Physik*, **51**(3):165–180, Mar 1928.

[BO27] M. Born and R. Oppenheimer. “Zur Quantentheorie der Molekeln.” *Annalen der Physik*, **389**(20):457–484, 1927.

[CGF22] C. Cisowski, J. B. Götte, and S. Franke-Arnold. “Colloquium: Geometric phases of light: Insights from fiber bundle theory.” *Rev. Mod. Phys.*, **94**:031001, Jul 2022.

[CHC21] Luis Cortes-Herrera, Xiaotong He, Jaime Cardenas, and Govind P. Agrawal. “Coupled-mode theory of the polarization dynamics inside a microring resonator with a uniaxial core.” *Phys. Rev. A*, **103**:063517, Jun 2021.

[CLW20] Siyu Chen, Yan ge Liu, Zhi Wang, Huiyi Guo, Hongwei Zhang, and Baiwei Mao. "Mode transmission analysis method for photonic lantern based on FEM and local coupled mode theory." *Opt. Express*, **28**(21):30489–30501, Oct 2020.

[Com09] Daniel Comparat. "General conditions for quantum adiabatic evolution." *Phys. Rev. A*, **80**:012106, Jul 2009.

[DCT24] Momen Diab, Ross Cheriton, Jacob Taylor, Dhwanil Patel, Libertad Rojas, Mark Barnet, Polina Zavyalova, Dan-Xia Xu, Pavel Cheben, Siegfried Janz, Jens H. Schmid, and Suresh Sivanandam. "Experimental demonstration of photonic phase correctors based on grating coupler arrays and thermo-optic shifters.", 2024.

[GVD19] Pradip Gatkine, Sylvain Veilleux, and Mario Dagenais. "Astrophotonic Spectrographs." *Applied Sciences*, **9**(2), 2019.

[Had92] G. Ronald Hadley. "Wide-angle beam propagation using Padé approximant operators." *Opt. Lett.*, **17**(20):1426–1428, Oct 1992.

[HHK87] H. Haus, W. Huang, S. Kawakami, and N. Whitaker. "Coupled-mode theory of optical waveguides." *Journal of Lightwave Technology*, **5**(1):16–23, 1987.

[JGA23] Nemanja Jovanovic, Pradip Gatkine, Narsireddy Anugu, Rodrigo Amezcua-Correa, Ritoban Basu Thakur, Charles Beichman, Chad Bender, Jean-Philippe Berger, Azzurra Bigioli, Joss Bland-Hawthorn, Guillaume Bourdarot, Charles M. Bradford, Ronald Broeke, Julia Bryant, Kevin Bundy, Ross Cheriton, Nick Cvetojevic, Momen Diab, Scott A. Diddams, Aline N. Dinkelaker, Jeroen Duis, Stephen Eikenberry, Simon Ellis, Akira Endo, Donald F. Figer, Michael Fitzgerald, Itandehui Gris-Sanchez, Simon Gross, Ludovic Grossard, Olivier Guyon, Sebastiaan Y. Haffert, Samuel Halverson, Robert J. Harris, Jinping He, Tobias Herr, Philipp Hottinger, Elsa Huby, Michael Ireland, Rebecca Jenson-Clem, Jeffrey Jewell, Laurent Jocou, Stefan Kraus, Lucas Labadie, Sylvestre Lacour, Romain Laugier, Katarzyna Ławniczuk, Jonathan Lin, Stephanie Leifer, Sergio Leon-Saval, Guillermo Martin, Frantz Martinache, Marc-Antoine Martinod, Ben Mazin, Stefano Minardi, John D Monnier, Reinan Moreira, Denis Mourard, Abani Shankar Shankar Nayak, Barnaby Norris, Ewelina Obrzud, Karine Perraut, François Reynaud, Steph Sallum, David Schiminovich, Christian Schwab, Eugene Serbayn, Sherif Soliman, Andreas Stoll, Liang Tang, Peter Tuthill, Kerry Vahala, Gautam Vasisht, Sylvain Veilleux, Alexander B. Walter, Edward J Wollack, Yinzi Xin, Zongyin Yang, Stephanos Yerolatsitis, Yang Zhang, and Chang-Ling Zou. "2023 Astrophotonics Roadmap: pathways to realizing multi-functional integrated astrophotonic instruments." *Journal of Physics: Photonics*, 2023.

[JJJ10] Long Jin, Wei Jin, Jian Ju, and Yiping Wang. "Coupled Local-Mode Theory for Strongly Modulated Long Period Gratings." *J. Lightwave Technol.*, **28**(12):1745–1751, Jun 2010.

[JSG17] Jovanovic, N., Schwab, C., Guyon, O., Lozi, J., Cvetojevic, N., Martinache, F., Leon-Saval, S., Norris, B., Gross, S., Doughty, D., Currie, T., and Takato, N. "Efficient injection from large telescopes into single-mode fibres: Enabling the era of ultra-precision astronomy." *A & A*, **604**:A122, 2017.

[KFL24] Yoo Jung Kim, Michael P. Fitzgerald, Jonathan Lin, Steph Sallum, Yinzi Xin, Nemanja Jovanovic, and Sergio Leon-Saval. "Coherent Imaging with Photonic Lanterns." *arXiv e-prints*, p. arXiv:2402.08158, feb 2024.

[Kha04] RP Khare. *Fiber optics and optoelectronics*. Oxford University Press, 2004.

[Kit04] Charles Kittel. *Introduction to Solid State Physics*. Wiley, 8 edition, 2004.

[KMT22] Teresa Klinner-Teo, Marc-Antoine Martinod, Peter Tuthill, Simon Gross, Barnaby Norris, and Sergio Leon-Saval. "Achromatic design of a photonic tricoupler and phase shifter for broadband nulling interferometry." *Journal of Astronomical Telescopes, Instruments, and Systems*, **8**(04), October 2022.

[LAB10] Sergio G. Leon-Saval, Alexander Argyros, and Joss Bland-Hawthorn. "Photonic lanterns: a study of light propagation in multimode to single-mode converters." *Opt. Express*, **18**(8):8430–8439, Apr 2010.

[LFS14] Sergio G. Leon-Saval, Nicolas K. Fontaine, Joel R. Salazar-Gil, Burcu Ercan, Roland Ryf, and Joss Bland-Hawthorn. "Mode-selective photonic lanterns for space-division multiplexing." *Opt. Express*, **22**(1):1036–1044, Jan 2014.

[LFX23] Jonathan W. Lin, Michael P. Fitzgerald, Yinzi Xin, Yoo Jung Kim, Olivier Guyon, Barnaby Norris, Christopher Betters, Sergio Leon-Saval, Kyohoon Ahn, Vincent Deo, Julien Lozi, Sébastien Vievard, Daniel Levinstein, Steph Sallum, and Nemanja Jovanovic. "Real-time Experimental Demonstrations of a Photonic Lantern Wave-front Sensor." *The Astrophysical Journal Letters*, **959**(2):L34, dec 2023.

[Lin24] Jonathan Lin. "cbeam." Astrophysics Source Code Library, record ascl:2404.001, apr 2024.

[LJF21] Jonathan Lin, Nemanja Jovanovic, and Michael P. Fitzgerald. "Design considerations of photonic lanterns for diffraction-limited spectrometry." *J. Opt. Soc. Am. B*, **38**(7):A51–A63, Jul 2021.

[LL76] L. D. Landau and E. M. Lifshitz. *Mechanics, Third Edition: Volume 1 (Course of Theoretical Physics)*. Butterworth-Heinemann, 3 edition, January 1976.

[McC00] Martin McCall. "On the Application of Coupled Mode Theory for Modeling Fiber Bragg Gratings." *J. Lightwave Technol.*, **18**(2):236, Feb 2000.

[Mil54] S. E. Miller. “Coupled wave theory and waveguide applications.” *The Bell System Technical Journal*, **33**(3):661–719, 1954.

[MLP19] Claudio Mancinelli, Marco Livesu, and Enrico Puppo. “A comparison of methods for gradient field estimation on simplicial meshes.” *Computers & Graphics*, **80**:37–50, 2019.

[MNT21] Marc-Antoine Martinod, Barnaby Norris, Peter Tuthill, Tiphaine Lagadec, Nemanja Jovanovic, Nick Cvetojevic, Simon Gross, Alexander Arriola, Thomas Gretzinger, Michael J. Withford, Olivier Guyon, Julien Lozi, Sebastien Vievard, Vincent Deo, Jon S. Lawrence, and Sergio Leon-Saval. “Scalable photonic-based nulling interferometry with the dispersed multi-baseline GLINT instrument.” *Nature Communications*, **12**(1):2465, Apr 2021.

[NWB20] Barnaby R. M. Norris, Jin Wei, Christopher H. Betters, Alison Wong, and Sergio G. Leon-Saval. “An all-photonic focal-plane wavefront sensor.” *Nature Communications*, **11**(1):5335, Oct 2020.

[Oka06] Katsunari Okamoto. “Chapter 4 - Coupled mode theory.” In Katsunari Okamoto, editor, *Fundamentals of Optical Waveguides (Second Edition)*, pp. 159–207. Academic Press, Burlington, second edition edition, 2006.

[Pan56] S. Pancharatnam. “Generalized theory of interference, and its applications.” *Proceedings of the Indian Academy of Sciences - Section A*, **44**(5):247–262, Nov 1956.

[Ped15] G. L. Pedrola. *Vectorial and Three-Dimensional Beam Propagation Techniques*, chapter 3, pp. 71–129. John Wiley & Sons, Ltd, 2015.

[Pie54] J. R. Pierce. “Coupling of Modes of Propagation.” *Journal of Applied Physics*, **25**(2):179–183, 02 1954.

[RO14] Gustavo Rigolin and Gerardo Ortiz. “Degenerate adiabatic perturbation theory: Foundations and applications.” *Phys. Rev. A*, **90**:022104, Aug 2014.

[Ros84] J. N. Ross. “The rotation of the polarization in low birefringence monomode optical fibres due to geometric effects.” *Optical and Quantum Electronics*, **16**(5):455–461, Sep 1984.

[SL83] A .W. Snyder and J. Love. *Optical Waveguide Theory*. Springer, 1 edition, 1983.

[SN20] J. J. Sakurai and Jim Napolitano. *Modern Quantum Mechanics*. Cambridge University Press, 3 edition, 2020.

[SS20] Sugeet Sunder and Anurag Sharma. “Adiabatic propagation algorithm for photonic lanterns.” *Optical Fiber Technology*, **57**:102219, 2020.

[Sub95] D. Subbarao. “Topological phase in Gaussian beam optics.” *Opt. Lett.*, **20**(21):2162–2164, Nov 1995.

[SWH99] E. Serabyn, J. K. Wallace, G. J. Hardy, E. G. H. Schmidtlin, and H. T. Nguyen. “Deep nulling of visible laser light.” *Appl. Opt.*, **38**(34):7128–7132, Dec 1999.

[SWZ16] Chen Shi, Xiaolin Wang, Pu Zhou, Xiaojun Xu, and Qisheng Lu. “Theoretical study of mode evolution in active long tapered multimode fiber.” *Opt. Express*, **24**(17):19473–19490, Aug 2016.

[TEB13] Christopher Q. Trinh, Simon C. Ellis, Joss Bland-Hawthorn, Jon S. Lawrence, Anthony J. Horton, Sergio G. Leon-Saval, Keith Shortridge, Julia Bryant, Scott Case, Matthew Colless, Warrick Couch, Kenneth Freeman, Hans-Gerd Löhmannsröben, Luke Gers, Karl Glazebrook, Roger Haynes, Steve Lee, John O’Byrne, Stan Miziarski, Martin M. Roth, Brian Schmidt, Christopher G. Tinney, and Jessica Zheng. “GNOSIS: The first instrument to use fiber Bragg gratings for OH suppression.” *The Astronomical Journal*, **145**(2):51, jan 2013.

[TKT05] Alexander V. Tavrov, Yosuke Kobayashi, Yosuke Tanaka, Tatsutoshi Shioda, Yukitoshi Otani, Takashi Kurokawa, and Mitsuo Takeda. “Common-path achromatic interferometer–coronagraph: nulling of polychromatic light.” *Opt. Lett.*, **30**(17):2224–2226, Sep 2005.

[TT86] R. Tewari and K. Thyagarajan. “Analysis of tunable single-mode fiber directional couplers using simple and accurate relations.” *Journal of Lightwave Technology*, **4**(4):386–390, 1986.

[TTB21] Adam K. Taras, Alessandro Tuniz, Musawer A. Bajwa, Judith M. Dawes Vincent Ng, Christopher G. Poulton, and C. Martijn De Sterke. “Shortcuts to adiabaticity in waveguide couplers–theory and implementation.” *Advances in Physics: X*, **6**(1):1894978, 2021.

[Urb03] H.K. Urbantke. “The Hopf fibration—seven times in physics.” *Journal of Geometry and Physics*, **46**(2):125–150, 2003.

[XJR22] Yinzi Xin, Nemanja Jovanovic, Garreth Ruane, Dimitri Mawet, Michael P. Fitzgerald, Daniel Echeverri, Jonathan Lin, Sergio Leon-Saval, Pradip Gatkine, Yoo Jung Kim, Barnaby Norris, and Steph Sallum. “Efficient Detection and Characterization of Exoplanets within the Diffraction Limit: Nulling with a Mode-selective Photonic Lantern.” *The Astrophysical Journal*, **938**(2):140, oct 2022.

[ZKF23] Jiang Zhang, Thi Ha Kyaw, Stefan Filipp, Leong-Chuan Kwek, Erik Sjöqvist, and Dianmin Tong. “Geometric and holonomic quantum computation.” *Physics Reports*, **1027**:1–53, 2023. Geometric and holonomic quantum computation.

CHAPTER 12

Statistical theory of atmospheric turbulence

Contents

Ground-based telescopes must correct for the wavefront errors imprinted by Earth's turbulent atmosphere to recover diffraction-limited imaging performance. This section provides a background on the statistical theory of turbulence, with the end goal of deriving the optical effects caused by turbulent layers in the atmosphere from *almost* first principles. Appendices provide a mathematical background on random processes.[1]

---

[1] Some advice: a lot of stuff goes into turbulence theory, so it's best to go slow. My advisor once warned me that "the study of turbulence has driven people crazy." I now believe he was referring to himself.

SYMBOLS

Symbol	Description
Re	Reynolds number
v	characteristic velocity for a turbulent fluctuation
ν	kinematic viscosity
l	characteristic lengthscale for a turbulent fluctuation
ϵ	viscous dissipation rate
$\mathbf{u}$	fluid velocity vector field, components u_i
$\mathbf{k}$	spatial angular wavenumber for turbulent fluctuations, $[k_x, k_y, k_z]^T$
$\mathbf{f}$	spatial cyclic wavenumber for turbulent fluctuations, $[f_x, f_y, f_z]^T$
k_λ	free-space angular wavenumber for light waves, $2\pi/\lambda$
$\Phi(k)$	3D turbulence power spectrum, 1D and 2D variants are subscripted
$\psi(\mathbf{x})$	electric field wavefront; $\mathbf{x}$ is the pupil coordinate
$\phi(\mathbf{x})$	wavefront phase, $\psi = e^{i\phi}$
$\Phi_n(k)$	the 3D spectral density for refractive index fluctuations
$\Phi_p(k)$	the 2D spectral density for Kolmogorov-aberrated phase
$\boldsymbol{\rho}$	displacement vector
$D_{ij}(\boldsymbol{\rho})$	structure function tensor
$D(\rho)$	structure function
B	correlation
C^2	structure constant
h	altitude of turbulent layer; δh is the thickness
$\Gamma(z)$	Gamma function
$\langle\theta\rangle$	expected value or ensemble average of random variable θ
$\boldsymbol{\alpha}$	on-sky angular coordinate
$\boldsymbol{\xi}$	on-sky frequency coordinate/pupil-plane coordinate, λ units

$S(\boldsymbol{\alpha})$	point-spread function (PSF)
$P(\mathbf{x})$	pupil transmission function
$\mathcal{R}$	resolving power
d	telescope pupil diameter
R	telescope pupil radius
r_0	atmosphere coherence length or Fried parameter
$\mathcal{S}$	Strehl ratio
Z_i	Zernike polynomial, ordered by Noll index
Q_i	Fourier transform of Z_i
a_j	Zernike mode amplitude
Σ_{ij}	covariance matrix for Kolmogorov turbulence in Zernike basis
$\int f(\omega) d\phi(\omega)$	Riemann-Stieltjes integral
$F(x)$	cumulative distribution function (CDF); specifically, $F(\omega)$ is the spectral distribution function with $F'(\omega)$ the spectral density
inf	infimum
sup	supremum
$\wedge$	intersection (of sets)
$\cup$	union (of sets)
$\mathbb{R}$	space of real numbers
$\mathbb{Z}$	space of integers
$\xi(t)$	a random function
$\zeta_\tau(t)$	difference function of $\xi(t)$ with lag τ
δ_{ij}	Kronecker delta
ϵ_{ijk}	Levi-Civita symbol

12.1 Order-of-magnitude derivation of Kolmogorov turbulence

This section and the next follow material from [TSC61]. Consider an initially laminar flow, into which we add a fluctuation of velocity v and size l. The stability of this fluctuation is determined by the Reynolds number Re, where

$$\mathrm{Re} = \frac{vl}{\nu}. \tag{12.1}$$

Here, ν denotes kinematic viscosity. The physical importance of the Reynolds number can be seen using dimensional considerations. The aforementioned velocity fluctuation has a per-mass energy of around v^2 and a characteristic a time scale of around $\tau = l/v$, meaning that the creation of the fluctuation involves a power of around v^3/l. On the other hand, the power dissipated into heat by a laminar flow is of the order $\nu v^2/l^2$, derived again by considering units. The velocity fluctuation can only continue to exist if the power flowing into it exceeds the viscous dissipation rate which gives

$$\frac{v^3}{l} > \nu\frac{v^2}{l^2} \implies \frac{vl}{\nu} > 1. \tag{12.2}$$

Of course, the above derivation is rough, so it would be more accurate to define a critical Reynolds number Re_{cr} above which velocity fluctuations can exist. But the rest of the equations derived in this subsection assume a critical Reynolds number of 1, similar to other texts.

Now we consider what a turbulent flow – a flow with velocity fluctuations – would looks like. Suppose we start with a laminar flow with characteristic velocity v and length L. Now we gradually increase the Reynolds number of this flow. Eventually, the Reynolds number $\mathrm{Re} = Lv/\nu$ of this flow will surpass the critical Reynolds number Re_{cr}, and instabilities of some size l and velocity v_l will develop. These velocity fluctuations will be stable at first. But as we keep increasing the outer Reynolds number Re, the velocity fluctuations will become stronger, forcing the *fluctuation's* Reynolds number above the critical value. When this happens, the velocity fluctuations themselves become unstable to smaller fluctuations. If we make the outer Reynolds number of the flow very high, we get a cascade of larger velocity fluctuations transferring energy to smaller and smaller fluctuations, all the way down to some inner limit set by the critical Reynolds number. This type of turbulence is *Kolmogorov turbulence.*

Since energy dissipation is controlled by the transition to laminar flow at the smallest scale of turbulence, we can assume that, in a steady state scenario, the power being transferred by the fluctuations is independent of the fluctuation size, and equal to ϵ, the viscous dissipation rate at the inner scale:

$$\epsilon = \frac{\nu v_0^2}{l_0^2} = \frac{v_0^3}{l_0}. \tag{12.3}$$

Thus, for Kolmogorov turbulence, fluctuational energy within a turbulent eddy of size l scales as $l^{2/3}$:

$$v^2 = \epsilon^{2/3} l^{2/3}. \tag{12.4}$$

We can additionally derive a relation between the inner and outer scales. At the inner scale we have the following two relations

$$\begin{aligned} v_0 &= (\epsilon l_0)^{1/3} \\ \epsilon &= \nu \frac{v_0^2}{l_0^2}. \end{aligned} \tag{12.5}$$

The above system can be solved, yielding

$$\begin{aligned} l_0 &= \frac{\nu^{3/4}}{\epsilon^{1/4}} \\ v_0 &= (\nu\epsilon)^{1/4}. \end{aligned} \tag{12.6}$$

Substituting in $v_L^3/L = \epsilon$, we find the following relations

$$\begin{aligned} l_0 &= \frac{L}{\mathrm{Re}^{3/4}} \\ v_0 &= \frac{v_L}{\mathrm{Re}^{1/4}}. \end{aligned} \tag{12.7}$$

A larger Reynolds number at the outer scale of the flow leads to the creation of smaller velocity fluctuations.

Lastly, we derive the power spectra $\Phi(k)$ of Kolmogorov turbulence. Define the wavenumber of a turbulent eddy (fluctuation) as $k = 2\pi/l$. Then in the 1D case we have

$$\begin{aligned} \Phi_{\mathrm{1D}}(k)dk &\propto v^2 \propto k^{-2/3} \\ \Phi_{\mathrm{1D}}(k) &\propto k^{-5/3}. \end{aligned} \tag{12.8}$$

In the 3D homogeneous/isotropic case we have

$$\begin{aligned} \Phi_{\mathrm{3D}}(k)k^2dk &\propto v^2 \propto k^{-2/3} \\ \Phi_{\mathrm{3D}}(k) &\propto k^{-11/3}. \end{aligned} \tag{12.9}$$

In future sections I take $\Phi(k) \equiv \Phi_{\mathrm{3D}}(k)$ for brevity.

12.2 Turbulence structure functions

12.2.1 Velocity structure function

Structure functions are useful because they allow us to compute spectral representations for random processes such as turbulence which are only locally homogeneous and isotropic. A review is given in §12.B.7. In the context of turbulent fluids, there are multiple structure functions for the various fluid parameters (velocity, refractivity) that are of interest. The first structure function we will develop is the fluid velocity structure function. Since velocity is a vector, characterization of the velocity field of a turbulent fluid involves not 1 but 9 structure functions, forming a structure function tensor D_{ij}, where

$$D_{ij}(\boldsymbol{\rho}) = \left\langle (u_i - u_i')(u_j - u_j') \right\rangle . \tag{12.10}$$

Here, angled brackets denote ensemble average (after all, in our statistical treatment the velocity field is not deterministic), u_i denotes the Cartesian velocity components of the fluid at some point $\mathbf{r}$, and u_j' denotes the velocity components at another point $\mathbf{r}' = \mathbf{r} + \boldsymbol{\rho}$.

Fortunately, under the assumption of local isotropy and homogeneity, the velocity structure function tensor can be expressed with only two components instead of nine. Under the assumption of incompressibility that reduces to a single component; see §12.C.1 for a proof. This single component is the longitudinal structure function, denoted D_{pp}. The *pp* subscript stands for parallel-parallel, referring to the direction along which we are projecting velocity vectors. Mathematically, if we write u_ρ as the component of the fluid along the separation vector $\boldsymbol{\rho}$, we have

$$D_{pp}(\rho) = \left\langle [u_\rho(\mathbf{r} + \boldsymbol{\rho}) - u_\rho(\mathbf{r})]^2 \right\rangle . \tag{12.11}$$

Physically, D_{pp} measures the intensity of velocity fluctuation along the along the direction of the separation vector $\boldsymbol{\rho}$ over length scales of ρ. The above function fully determines the structure function tensor D_{ij}.

Recalling equation 12.4 and noting that $D_{pp} \approx u_\rho^2$, we find

$$D_{pp}(\rho) = C_v^2 \rho^{2/3} \tag{12.12}$$

where the constant C_v^2 is known as the *velocity structure constant*. Unit considerations lead to the same result. Equation 12.12 is the "two-thirds law" of Kolmogorov turbulence.

Note that the 2/3 law only holds when we are in the regime of Kolmogorov turbulence, i.e. when we are considering eddies with length scales in between the inner and outer scales of the turbulent flow. For length scales much shorter than the inner scale l_0, flow is laminar and velocity changes smoothly, allowing us to approximate $u_\rho(\mathbf{r}+\boldsymbol{\rho})$ with a first-order power expansion in ρ. This shows

$$D_{pp}(\rho) \propto \rho^2; \rho \ll l_0. \tag{12.13}$$

The 2/3 power law also breaks down as ρ approaches the outer scale length L. In this case, the assumptions of isotropy and homogeneity for the difference function $u_\rho(\mathbf{r}+\boldsymbol{\rho}) - u_\rho(\mathbf{r})$ no longer hold. Modifications to the structure function to account for this behavior were proposed by von Kármán, but this treatment is left to other texts. For now, we will assume we are in the *inertial* regime, where the 2/3 law holds.

12.2.2 Other structure functions

In the context of optical propagation, the velocity structure is not of any particular use, at least directly. Optical effects are primarily dependent on the refractive index distribution of the fluid, which in turn depends on fluid parameters such as humidity and temperature.[2] The next step in deriving the optical effects of turbulence is to then characterize the structure functions of other fluid parameters besides velocity.

Obukhov and Yaglom [OY59] showed that as long as the value of a parameter does not affect the physical behavior of turbulence (passive property), and that a fluid parcel when translated conserves the value of that parameter (conservative property), the parameter admits a structure function of the same form as the velocity structure function. Such parameters are termed *conservative passive additives.*

Qualitatively, the similarity between structure functions for velocity and conservative passive additives makes sense. Inhomogeneities in these additives are ultimately created, transported, and mixed by action of the velocity field of the turbulent flow, and thus reflect

---

[2] Equivalently, density; pressure does not affect propagation that much since any fluctuations are smoothed out at the speed of sound. Vertical pressure variation of the atmosphere does create dispersive effects, though.

inhomogeneities in the velocity field. Additionally, there is a minimum physical size for the inhomogeneities of these additives, set by molecular diffusion in an analogous way to how the size of the smallest velocity fluctuations are set by viscous dissipation.

It is important to realize that quantities such as temperature and refractive index are not, strictly speaking, conservative passive additives. The temperature of a fluid parcel displaced vertically changes adiabatically in accordance with external pressure. Regardless, often times the refractive index and temperature are assumed to be conservative passive additives. This approximation becomes even better when we consider that the bulk of atmospheric turbulence in astronomical settings arises due to wind shear between atmospheric layers, concentrating turbulence into relatively thin layers 100-200 m in thickness.

The most important structure function in this analysis is the refractive index structure function. Because of arguments given above, we can express the structure function in a manner similar to equation 12.12:

$$D_n = C_n^2 \rho^{2/3}. \tag{12.14}$$

Similar relations can be defined for quantities such as temperature and refractivity. Lastly, when considering optical propagation it will be useful to do a spectral decomposition of the refractive index field. Using a form of the Wiener-Khinchin theorem, equation 12.185 derived in §12.B, we have

$$C_n^2 \rho^{2/3} = 8\pi \int_0^\infty \left[1 - \frac{\sin(kr)}{kr}\right] \Phi_n(k) k^2 dk \tag{12.15}$$

where C_n is the refractive index structure constant. The above can be solved using a power law ansatz for $\Phi_n(k)$, the spectral density of refractive index fluctuations for Kolmogorov turbulence. The result is

$$\Phi_n(k) = \frac{5}{18\pi\Gamma(\frac{1}{3})} C_n^2 k^{-11/3} \approx 0.033 C_n^2 k^{-11/3}. \tag{12.16}$$

Outside the inertial regime, the above is typically modified to account for an inner and outer lengthscale (bracketing the inertial regime), as was done by von Kármán.

12.3 Effect of turbulent layers on wavefront statistics

This section follows material from [Rod81]. We first consider the effect of a single thin layer of locally homogeneous and isotropic turbulence. First, what do we mean by thin? Well, to make math easier we would like to neglect diffraction effects within this layer; such an assumption is fine as long as diffracted rays will not interfere within the layer. A turbulent cell of characteristic size l will deflect a ray by $\theta \approx \lambda/l$. Thus, if we want to ignore diffraction the maximum propagation length the rays can have is $L = l/\theta = l^2/\lambda$. For severe turbulence where the turbulent cells have a characteristic size of 10 cm, the maximum propagation length is 10 km. On the flip side, we would also like to apply Gaussian statistics to wavefronts passing through this layer, which requires wavefronts to pass through many turbulent cells (so that the wavefront phase is a sum of many independent random variables) before exiting the layer. This means that the minimum thickness of the layers we are considering should be much larger than l, which can top out at around 1 m. Thus, we consider a layer "thin" if the layer thickness δh satisfies the condition $l \ll \delta h \ll l^2/\lambda \implies 1\,\mathrm{m} \ll \delta h \ll 10\,\mathrm{km}$. Note that turbulent layers in the atmosphere have been measured to be on the order of 100 m thick.

Now, suppose we have a plane wave with unit magnitude incident on a thin turbulent layer of thickness δh with a refractive index field $n(\mathbf{x}, z)$; $\mathbf{x}$ is the transverse spatial coordinate. For convenience, we fix a coordinate system so that $z = 0$ right above the turbulent layer with z increasing as we approach the ground. After passing through the layer, the wavefront has the form

$$\psi_0(\mathbf{x}) = e^{i\phi_0(\mathbf{x})} \tag{12.17}$$

where the phase shift ϕ induced by the refractive index fluctuations is

$$\phi_0(\mathbf{x}) = k_\lambda \int_0^{\delta h} n(\mathbf{x}, z) dz. \tag{12.18}$$

Here, $k_\lambda = 2\pi/\lambda$ is the wavenumber of the wavefront; the subscript is to differentiate this wavenumber from the spatial wavenumber of turbulent fluctuations. The "0" subscripts are added to indicate that the wavefront is located at $z = 0$. 12.17 follows because we have neglected diffraction.

We assume $\phi_0(\mathbf{x})$ has Gaussian statistics by the central limit theorem since it is essentially a sum of many independent random variables. The correlation of the wavefront is

$$B_{\psi_0}(\mathbf{x}, \mathbf{x} + \boldsymbol{\rho}) = \left\langle e^{i[\phi(\mathbf{x}) - \phi(\mathbf{x}+\boldsymbol{\rho})]} \right\rangle \tag{12.19}$$

where $\boldsymbol{\rho}$ is a displacement in the xy plane. Note that for a 0 mean Gaussian random variable θ, we have

$$\begin{aligned}
\left\langle e^{i\theta} \right\rangle &= \frac{1}{\sqrt{2\pi \langle \theta^2 \rangle}} \int_{-\infty}^{\infty} e^{i\theta} e^{-\frac{1}{2}\frac{\theta^2}{\langle \theta^2 \rangle}} d\theta \\
&= \frac{1}{\sqrt{2\pi \langle \theta^2 \rangle}} \int_{-\infty}^{\infty} d\theta \; e^{-\frac{1}{2}\left[\frac{\theta}{\sqrt{\langle \theta^2 \rangle}} - i\sqrt{\langle \theta^2 \rangle}\right]^2} e^{-\frac{1}{2}\langle \theta^2 \rangle} \\
&= e^{-\frac{1}{2}\langle \theta^2 \rangle} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} du \, e^{-\frac{1}{2}[u - \text{constant}]^2} \\
&= e^{-\frac{1}{2}\langle \theta^2 \rangle}.
\end{aligned} \tag{12.20}$$

We now replace $\theta = \phi_0(\mathbf{x}) - \phi_0(\mathbf{x} + \boldsymbol{\rho})$, which is also Gaussian distributed and furthermore is homogeneous (ϕ_0 itself is assumed locally homogeneous), meaning we can assume without loss of generality that $\langle \phi_0(\mathbf{x}) - \phi_0(\mathbf{x} + \boldsymbol{\rho}) \rangle = 0$ holds everywhere. The correlation of the wavefront can then be expressed as

$$B_{\psi_0}(\rho) = B_{\psi_0}(\mathbf{x}, \mathbf{x} + \boldsymbol{\rho}) = e^{-\frac{1}{2}\left\langle [\phi_0(\mathbf{x}) - \phi_0(\mathbf{x}+\boldsymbol{\rho})]^2 \right\rangle} = e^{-\frac{1}{2} D_{\phi_0}(\rho)} \tag{12.21}$$

where $D_{\phi_0}(\rho)$ is the phase structure function of the wavefront immediately after exiting the turbulent layer. Interestingly, even though we haven't assumed homogeneous turbulence, the correlation of the wavefront still depends only on the magnitude of the separation ρ.

Our next goal is to connect D_{ϕ_0} to D_n, the refractive index structure function. The following algebraic identity is useful.

$$(a-b)(c-d) = \frac{1}{2}\left[(a-d)^2 + (b-c)^2 - (a-c)^2 - (b-d)^2\right]. \tag{12.22}$$

We begin by using 12.18 to express the phase structure function. Let $z = z_2 - z_1$, $\rho = \sqrt{(x_1-x_2)^2 + (y_1-y_2)^2}$ and $r = \sqrt{(x_1-x_2)^2 + (y_1-y_2)^2 + (z_1-z_2)^2}$.

$$\begin{aligned}
\phi_0(x_1,y_1) - \phi_0(x_2,y_2) &= k_\lambda \int_0^{\delta h} \left[n(x_1,y_1,z) - n(x_2,y_2,z)\right] dz \\
\left\langle[\phi_0(x_1,y_1) - \phi_0(x_2,y_2)]^2\right\rangle &= k_\lambda^2 \int_0^{\delta h} dz_1 \int_0^{\delta h} dz_2 \left\langle[n(x_1,y_1,z_1) - n(x_2,y_2,z_1)]\left[n(x_1,y_1,z_2) - n(x_2,y_2,z_2)\right]\right\rangle \\
D_{\phi_0}(\rho) &= k_\lambda^2 \iint_0^{\delta h} dz_1 dz_2 \, \frac{1}{2}\Big[\left(n(x_1,y_1,z_1) - n(x_2,y_2,z_2)\right)^2 \\
&\quad + \left(n(x_2,y_2,z_1) - n(x_1,y_1,z_2)\right)^2 \\
&\quad - \left(n(x_1,y_1,z_1) - n(x_1,y_1,z_2)\right)^2 \\
&\quad - \left(n(x_2,y_2,z_1) - n(x_2,y_2,z_2)\right)^2\Big] \\
&= k_\lambda^2 \iint_0^{\delta h} dz_1 dz_2 \, \frac{1}{2}\left[D_n(r) + D_n(r) - D_n(z) - D_n(z)\right] \\
&= k_\lambda^2 \iint_0^{\delta h} dz_1 dz_2 \left[D_n(r) - D_n(z)\right].
\end{aligned} \tag{12.23}$$

Local isotropy implies that $D_n(r) - D_n(z)$ is an even function of $z = z_2 - z_1$, meaning the above integral can be written as

$$D_{\phi_0}(\rho) = 2k_\lambda^2 \int_0^{\delta h} dz \, (\delta h - z)\left[D_n(r) - D_n(z)\right]. \tag{12.24}$$

Note that if $z \gg \rho$, then $r \to z$ and the integrand goes to 0. Thus, the bulk of the contribution to the integral occurs when $z \lesssim \rho$. The transverse displacement ρ must be smaller than the size L of the largest turbulent eddies[3] ($\sim 1 - 10$m). Since we have assumed $L \ll \delta h$ in the beginning of this subsection, $z \ll \delta h$ holds in the region where the integrand is non-zero. Furthermore, we can extend the upper bound of the integral to ∞ since the integrand is very close to 0 for $z > \delta h$. The above expression reduces to

$$D_{\phi_0}(\rho) = 2k_\lambda^2 \delta h \int_0^\infty dz \ \left[D_n(r) - D_n(z)\right]. \tag{12.25}$$

Finally, Kolmogorov's 2/3 law implies

$$D_n(r) = C_n^2 r^{2/3}$$

which we can substitute into our expression for D_ϕ to get

$$D_{\phi_0}(\rho) = 2k_\lambda^2 \delta h C_n^2 \int_0^\infty \left[\left(\rho^2 + z^2\right)^{1/3} - z^{2/3}\right] dz. \tag{12.26}$$

The integral term evaluates to

$$\int_0^\infty \left[\left(\rho^2 + z^2\right)^{1/3} - z^{2/3}\right] dz = \frac{1}{5}\frac{\Gamma(\frac{1}{2})\Gamma(\frac{1}{6})}{\Gamma(\frac{2}{3})} r^{5/3} \approx 1.457\rho^{5/3}$$

so the phase structure function of the wavefront immediately after exiting the turbulent layer is

$$D_{\phi_0}(\rho) \approx 2.914 k_\lambda^2 \delta h C_n^2 \rho^{5/3} \tag{12.27}$$

---

[3] Recall that in the study of a non-homogeneous field $\xi(\mathbf{r})$ we construct a difference field $\xi(\mathbf{r}_1) - \xi(\mathbf{r}_2)$ which we assume is homogeneous. Fields that we can do this with are locally homogeneous Intuitively, the assumption of local homogeneity is only valid if $|\mathbf{r}_1 - \mathbf{r}_2|$ is smaller than the size of fluctuations in the mean of the field.

and the correlation function of the wavefront immediately after is

$$B_{\psi_0}(\rho) = e^{-1.457 k_\lambda^2 \delta h C_n^2 \rho^{5/3}}. \tag{12.28}$$

Of course, we are more interested with the what the wavefront looks like on the ground. We can use the Fresnel propagation formula (valid since optical/NIR wavelengths are much smaller than even the smallest physical scales of turbulence, and the distance from the ground to the turbulent layer is much larger than the largest physical scales of turbulence) to propagate the wavefront to the ground. The Fresnel propagation formula can be expressed as a convolution

$$\psi(x, y, z) = \psi(x, y, 0) \circledast \frac{e^{ik_\lambda z}}{i\lambda z} e^{i\frac{k_\lambda}{2z}\left(x^2+y^2\right)}. \tag{12.29}$$

Note that the constant $e^{ik_\lambda z}$ phase factor can be dropped. Assuming our turbulent layer is a height h off the ground, we can express the wavefront on the ground ψ_g, as

$$\psi_g(x, y) = \psi_0(x, y) \circledast \frac{1}{i\lambda h} e^{i\frac{k_\lambda}{2h}\left(x^2+y^2\right)}. \tag{12.30}$$

The correlation of the wavefront on the ground can be written as

$$B_{\psi_g}(\rho) = \langle \psi_0(\mathbf{x}) \psi_0^*(\mathbf{x} + \boldsymbol{\rho}) \rangle \circledast \frac{1}{i\lambda h} e^{i\frac{k_\lambda}{2h}\left(x^2+y^2\right)} \circledast \frac{1}{-i\lambda h} e^{-i\frac{k_\lambda}{2h}\left(x^2+y^2\right)}. \tag{12.31}$$

The second convolution is equivalent to $\delta(\mathbf{x})$. Therefore,

$$B_{\psi_g}(\rho) = B_{\psi_0}(\rho) \tag{12.32}$$

and

$$B_{\psi_g}(\rho) = \exp -\frac{1}{2} D_{\phi_0}(\rho). \tag{12.33}$$

Note that the phase structure function is for the wavefront up at the turbulent layer, not the ground, at which the phase structure function will be different due to diffraction effects.

Regardless, often times the two structure functions are assumed the same. This is supposedly a fine assumption in astronomy, and is termed the "near-field" approximation. The above results for a single thin layer of turbulence generalize to any distribution of turbulence that can be formed by stacking such layers. The phase structure function becomes

$$D_{\phi_0}(\rho) = 2.914 k_\lambda^2 \rho^{5/3} \int C_n^2(z) dz.$$

In the case of observations off of zenith by some angle ζ, layers are thickened by a factor of $1/\cos\zeta$ so the structure function becomes

$$D_{\phi_0}(\rho) = 2.914 k_\lambda^2 \rho^{5/3} \sec\zeta \int C_n^2(z) dz \tag{12.34}$$

and

$$B_{\psi_g}(\rho) = \exp\left[-1.457 k_\lambda^2 \rho^{5/3} \sec\zeta \int C_n^2(z) dz\right]. \tag{12.35}$$

We may use the above structure function along with a form of the Wiener-Khinchin theorem given in equation 12.191 to derive a 2D spectral density for the wavefront phase (see the end of §12.4), which can then be used to simulate phase screens (Chapter 13).

12.4 Imaging through turbulence

This section characterizes the effect of turbulence on optical images, and follows material from [Rod81]. It concludes with an expression for the phase power spectrum, which is useful for simulating Kolmogorov-aberrated wavefronts.

A coherent imaging system is linear and the complex amplitude of the electric field, while an incoherent imaging system is linear in intensity. Since astronomical imaging is broadband and incoherent, it is linear in intensity. Further assuming isoplanaticity (the PSF and the effects of turbulence are the same all over the field of view), the following relation for the on-sky intensity distribution $O(\boldsymbol{\alpha})$, the PSF $S(\boldsymbol{\alpha})$, and the image intensity distribution $I(\boldsymbol{\alpha})$ holds.

$$I(\boldsymbol{\alpha}) = O(\boldsymbol{\alpha}) \circledast S(\boldsymbol{\alpha}) \tag{12.36}$$

i.e., the image intensity is the convolution of the on-sky intensity with the point spread function. Here, $\boldsymbol{\alpha}$ denotes angular direction on sky. $I(\boldsymbol{\alpha})$ is an instantaneous image. In the case where we are imaging through turbulence, it is often more useful to take long exposure (time-averaged) images. Assuming ergodicity, time-averaging is equivalent to ensemble averaging, and we have

$$\langle I(\boldsymbol{\alpha})\rangle = O(\boldsymbol{\alpha}) \circledast \langle S(\boldsymbol{\alpha})\rangle \, . \tag{12.37}$$

Taking the Fourier transform of the above and using convolution theorem, we get

$$\left\langle \hat{I}(\boldsymbol{\xi}) \right\rangle = \hat{O}(\boldsymbol{\xi}) \left\langle \hat{S}(\boldsymbol{\xi}) \right\rangle \tag{12.38}$$

where $\hat{I}(\boldsymbol{\xi})$ is the Fourier pair of $I(\boldsymbol{\alpha})$ and so on. The term $\left\langle \hat{S}(\boldsymbol{\xi}) \right\rangle$ is known as the optical transfer function, or OTF.

Now we derive expressions for the OTF. Suppose we have wavefront $\psi(\mathbf{x})$ at our telescope pupil, described by some pupil transmission function $P(\mathbf{x})$ which is 1 when $\mathbf{x}$ is in the pupil and 0 otherwise; for this section, the pupil coordinate $\mathbf{x}$ is assumed to be in wavelength units. The complex field amplitude distribution $\mathcal{A}$ at the telescope focal plane can be computed as a Fourier transform of the field amplitude distribution in the pupil plane (Fraunhofer propagation):

$$\mathcal{A}(\boldsymbol{\alpha}) \propto \mathcal{F}\left[P(\mathbf{x})\psi(\mathbf{x})\right] . \tag{12.39}$$

Note that the above is a proportionality since we are neglecting some constants. The PSF is the modulus of $\mathcal{A}$, squared.

$$S(\boldsymbol{\alpha}) = |\mathcal{A}(\boldsymbol{\alpha})|^2 \propto |\mathcal{F}\left[P(\mathbf{x})\psi(\mathbf{x})\right]|^2 . \tag{12.40}$$

Finally, the OTF is the Fourier transform of the PSF. Defining the Fourier transform as

$$\hat{f}(\boldsymbol{\xi}) = \mathcal{F}f(\boldsymbol{\alpha}) = \int_{-\infty}^{\infty} d\boldsymbol{\alpha} f(\boldsymbol{\alpha}) \exp\left[-2\pi i \boldsymbol{\alpha} \cdot \boldsymbol{\xi}\right]$$

and the inverse Fourier transform as

$$f(\boldsymbol{\alpha}) = \mathcal{F}^{-1}\hat{f}(\boldsymbol{\xi}) = \int_{-\infty}^{\infty} d\boldsymbol{\xi} \hat{f}(\boldsymbol{\xi}) \exp\left[2\pi i \boldsymbol{\alpha} \cdot \boldsymbol{\xi}\right]$$

we notice

$$\hat{f}(-\boldsymbol{\alpha}) = \mathcal{F}\hat{f}(\boldsymbol{\xi}) = \mathcal{F}\mathcal{F}f(\boldsymbol{\alpha}).$$

This means the OTF can be expressed as

$$\begin{aligned}
\hat{S}(\boldsymbol{\xi}) &= \mathcal{F}S(\boldsymbol{\alpha}) \\
&\propto \mathcal{F}\left|\mathcal{F}\left[P(\mathbf{x})\psi(\mathbf{x})\right]\right|^2 \\
&\propto \mathcal{F}\mathcal{F}P(\mathbf{x})\psi(\mathbf{x}) \circledast \mathcal{F}\left[\mathcal{F}P(\mathbf{x})\psi(\mathbf{x})\right]^* \\
&\propto P(-\mathbf{x})\psi(-\mathbf{x}) \circledast P^*(\mathbf{x})\psi^*(\mathbf{x}) \\
&\propto \int_{-\infty}^{\infty} P^*(\mathbf{x})\psi^*(\mathbf{x})P(\mathbf{x}-\boldsymbol{\xi})\psi(\mathbf{x}-\boldsymbol{\xi})d\mathbf{x} \\
&\propto \int_{-\infty}^{\infty} P(\mathbf{x})\psi(\mathbf{x})P^*(\mathbf{x}+\boldsymbol{\xi})\psi^*(\mathbf{x}+\boldsymbol{\xi})d\mathbf{x}.
\end{aligned} \tag{12.41}$$

If we neglect turbulence and assume that ψ has unit amplitude, we find

$$\hat{S}(\boldsymbol{\xi}) \propto \int_{-\infty}^{\infty} P(\mathbf{x})P^*(\mathbf{x}+\boldsymbol{\xi})d\mathbf{x}. \tag{12.42}$$

In other words, in the absence of turbulence, the OTF is proportional the autocorrelation of the pupil transmittance, denoted T. If we decide to normalize the OTF to be 1 at the origin, then

$$\hat{S}(\boldsymbol{\xi}) = \frac{1}{A_0}\int_{-\infty}^{\infty} P(\mathbf{x})P^*(\mathbf{x}+\boldsymbol{\xi})d\mathbf{x} \equiv T(\boldsymbol{\xi}) \tag{12.43}$$

where $A_0 = \pi d^2/4\lambda^2$ is the pupil area in units of λ^2; d is the pupil diameter.

In the presence of turbulence, the wavefront at the pupil ψ is no longer guaranteed to be planar, and in general will vary with time. For a long-exposure image we have

$$\left\langle \hat{S}(\boldsymbol{\xi}) \right\rangle = \frac{1}{A_0} \int_{-\infty}^{\infty} P(\mathbf{x}) P^*(\mathbf{x} + \boldsymbol{\xi}) \left\langle \psi(\mathbf{x}) \psi^*(\mathbf{x} + \boldsymbol{\xi}) \right\rangle d\mathbf{x}. \tag{12.44}$$

We identify part of the integrand as the correlation of the wavefront, $B(\xi) = B_{\psi_g}(\lambda \xi)$. Since the B only depends on ξ, we can pull it out of the integral to derive

$$\left\langle \hat{S}(\boldsymbol{\xi}) \right\rangle = T(\boldsymbol{\xi}) B(\boldsymbol{\xi}). \tag{12.45}$$

In the presence of turbulence, the OTF for long exposures is the product of the telescope transfer function and the correlation of the wavefront.

The *resolving power* $\mathcal{R}$ of an imaging system is defined as the integral of its OTF. Importantly, the resolving power is connected to the PSF. In fact, since the PSF and the OTF are Fourier pairs, we have

$$\mathcal{R} \equiv \int_{-\infty}^{\infty} d\boldsymbol{\xi} \hat{S}(\boldsymbol{\xi}) = S(0). \tag{12.46}$$

The central PSF intensity $S(0)$ is related to another imaging performance metric, the Strehl ratio, introduced later in §12.5.

We now compute $\mathcal{R}$ in two regimes: the turbulence-free regime, and the turbulence-limited regime. The following derivation is from [Fri66]. Generally, the resolving power in both regimes can be expressed as

$$\mathcal{R} = \int_{-\infty}^{\infty} d\boldsymbol{\xi}\; T(\boldsymbol{\xi}) B(\boldsymbol{\xi}). \tag{12.47}$$

From 12.34, we know that the the structure function of the wavefront has a $\xi^{5/3}$ dependence, where ξ is the separation distance over which the structure function is computed. Using 12.35, we can write the correlation function of the wavefront as $B(\xi) = \exp -K\xi^{5/3}$ with some scaling factor K. Thus,

$$\mathcal{R} = \int_{-\infty}^{\infty} d\boldsymbol{\xi}\ T(\xi) \exp -K\xi^{5/3}.$$

Suppose there is some associated length scale with the turbulence-induced aberrations in the wavefront. When the pupil diameter is much smaller than this length scale, the exponential term in the integrand is approximately 1 for all values of ξ where the integrand is non-zero. Physically, when the pupil diameter is small enough, the wavefront impinging on the aperture can be approximated as planar, meaning we can neglect the turbulent term $B(\xi)$. This is the turbulence-free regime. The resolving power is

$$\mathcal{R}_0 = \int_{-\infty}^{\infty} d\boldsymbol{\xi} T(\boldsymbol{\xi}).$$

Fried assumed a circular aperture in his calculations, so we do the same. Recalling that T is the autocorrelation of the pupil transmission function P, we note that T simply computes the area (in λ^2 units) between two identical intersecting circles as a function of their separation (also in λ^2 units). The area is given by the following formula:

$$A(\rho) = \frac{d^2}{2}\left[\cos^{-1}(\rho/d) - \frac{\rho}{d}\sqrt{1-(\rho/d)^2}\right]$$

where d is the physical diameter of the circles and ρ is the separation. Integrating the above, we find

$$\begin{aligned}\iint \mathbf{d}\boldsymbol{\rho} A(\rho) &= \int_0^{2\pi}\int_0^d A(\rho)\rho d\rho d\theta \\ &= \frac{d^4}{2}\int_0^{2\pi}\int_0^1 \left(\cos^{-1} u - u\sqrt{1-x^2}\right) u du d\theta \\ &= \frac{\pi^2 d^4}{16}.\end{aligned}$$

To do the above integral, we have made the substitution $u = \rho/D$. Note that $\xi = \rho/\lambda$. Using the above result, we can derive the following expression for the resolving power in the turbulence-free regime:

$$
\begin{aligned}
\mathcal{R}_0 &= \int_{-\infty}^{\infty} d\boldsymbol{\xi} T(\boldsymbol{\xi}) = \frac{1}{A_0}\int_{-\infty}^{\infty} d\boldsymbol{\xi} A(\xi) \\
&= \frac{1}{A_0}\int_0^{2\pi}\int_0^{d/\lambda} A(\xi)\xi d\xi d\theta \\
&= \frac{4\lambda^2}{\pi d^2}\frac{\pi^2 d^4}{16\lambda^4} \\
&= \frac{\pi d^2}{4\lambda^2}.
\end{aligned}
\tag{12.48}
$$

In other words, the resolving power for a circular aperture without turbulence is the geometric area of the aperture in λ^2 units.

Conversely, when the pupil diameter is much larger than this length scale, the wavefront correlation will be very small except for when $\xi \approx 0$. Since $T(\xi) \approx 1$ in the immediate neighborhood of $\xi = 0$, we can neglect it. Physically, when the pupil diameter is large, turbulence imposes a smaller "effective" diameter. This is the turbulence-limited regime. We calculate the resolving power in this regime as follows

$$
\begin{aligned}
\mathcal{R}_\infty &= \int_{-\infty}^{\infty} \mathbf{d}\boldsymbol{\xi} B(\xi) \\
&= \int_0^{2\pi}\int_0^{\infty} \xi \exp -K\xi^{5/3} d\xi d\theta \\
&= \frac{6\pi\Gamma(\frac{6}{5})}{5K^{6/5}}.
\end{aligned}
\tag{12.49}
$$

At this point we pose the question: At what diameter is the resolving power in the two regimes equivalent? Let r_0 denote the critical diameter at which this occurs. Then, equating $\mathcal{R}_0 = \mathcal{R}_\infty$ we get

$$
K = \left[\frac{24}{5}\Gamma\left(\frac{6}{5}\right)\frac{\lambda^2}{r_0^2}\right]^{5/6} \approx 3.44\left(\frac{\lambda}{r_0}\right)^{5/3}.
$$

Using the above, the correlation function of the wavefront can be written as

$$B(\xi) = \exp -3.44\left(\xi\lambda/r_0\right)^{5/3}. \tag{12.50}$$

In physical units,

$$B(\rho) = \exp -3.44\left(\rho/r_0\right)^{5/3}. \tag{12.51}$$

Clearly, r_0 is the relevant length scale of turbulence. Loosely speaking, when the aperture diameter is smaller than r_0, our resolving power is limited by the aperture diameter. When the aperture diameter is larger than r_0, our resolving power is limited by turbulence. r_0 is often referred to as the *Fried parameter*.

Equating the above to 12.35, we find the following equivalent definition for r_0:

$$r_0 = \left[0.423 k_\lambda^2 \sec\zeta \int C_n^2(z)dz\right]^{-3/5} \tag{12.52}$$

where $k = 2\pi/\lambda$ is the wavenumber of the wavefront. With the Fried parameter, we can write the phase structure function more simply as

$$D_{\phi_0} = 6.88\left(\frac{\rho}{r_0}\right)^{5/3}. \tag{12.53}$$

Using equation 12.191, we may derive the 2D phase PSD, $\Phi_p(k)$.

$$\begin{aligned} D_{\phi_0} = 6.88\left(\frac{\rho}{r_0}\right)^{5/3} &= 4\pi\int_0^\infty \left[1 - J_0(k\rho)\right]\Phi_p(k)kdk \\ \Phi_p(k) &= \frac{6.88}{4\pi r_0^{5/3}k^{11/3}}\left[\int_0^\infty \left[1 - J_0(x)\right]x^{-8/3}dx\right]^{-1} \\ &\approx 0.490\, r_0^{-5/3}k^{-11/3} \end{aligned} \tag{12.54}$$

which matches the result from [JG94]. To perform the integral, I used the identity

$$\int_0^\infty x^{-p}\left[1 - J_0(x)\right]\, dx = \frac{\pi}{2^p\left[\Gamma\left(\frac{p+1}{2}\right)\right]^2 \sin\left(\pi\frac{p-1}{2}\right)} \tag{12.55}$$

which is listed in [Nol76]. The units of the power spectral density (PSD) are (length/radian)^2. In the case of a discrete numerical model, the above PSD should be multiplied by an area in frequency space corresponding to the discretization. For instance, if the simulation domain has a spatial extent $S \times S$ resolved into an $N \times N$ grid, and we convert to cyclic spatial frequency $f \equiv k/2\pi$, then the phase power spectrum is

$$\begin{aligned} \Phi_p(f)\Delta f^2 &\approx 0.00058\,\frac{4\pi^2 N^2}{S^2} r_0^{-5/3} f^{-11/3} \\ &\approx 0.023\,\frac{N^2}{S^2} r_0^{-5/3} (f_x^2 + f_y^2)^{-11/6} \end{aligned} \tag{12.56}$$

where f_x and f_y are the x and y cyclic spatial frequencies. The above expression is often used to numerically generate turbulence-aberrated wavefronts.

12.5 Modal correction of turbulence

In this section we consider the modal representation of Kolmogorov turbulence, which we use to quantify imaging performance when lower order modes are compensated.

12.5.1 Imaging metrics for corrected wavefronts

We consider aberrated wavefronts that have undergone some sort of partial compensation by an AO system. A useful metric for characterizing partially compensated wavefronts is the Strehl ratio, defined as the ratio between the actual peak intensity of a source image and the peak intensity of a perfectly diffraction-limited PSF produced by a completely planar wavefront. By 12.46, the Strehl ratio is equivalently the ratio between the actual resolving power of an optical system, divided by the resolving power of the system neglecting turbulence. Denote the Strehl ratio $\mathcal{S}$. As before, denote the PSF $S(\alpha)$ and the OTF $\hat{S}(\xi)$. Assuming good turbulence compensation,

$$\begin{aligned}
\mathcal{S} &= \frac{\int \mathbf{d}\boldsymbol{\xi} T(\xi) B(\xi)}{\int \mathbf{d}\boldsymbol{\xi} T(\xi)} \\
&= \frac{\int \mathbf{d}\boldsymbol{\xi} T(\xi) \exp -\frac{1}{2} D_{\phi_0}(\lambda\xi)}{\int \mathbf{d}\boldsymbol{\xi} T(\xi)} \\
&= \frac{\int \mathbf{d}\boldsymbol{\xi} T(\xi) \exp \left[B_{\phi_0}(\lambda\xi) - B_{\phi_0}(0)\right]}{\int \mathbf{d}\boldsymbol{\xi} T(\xi)} \\
&= \exp\left[-B_{\phi_0}(0) \frac{\int \mathbf{d}\boldsymbol{\xi} T(\xi) \exp B_{\phi_0}(\lambda\xi)}{\int \mathbf{d}\boldsymbol{\xi} T(\xi)}\right] \\
&\approx \exp\left[-B_{\phi_0}(0) \frac{\int \mathbf{d}\boldsymbol{\xi} T(\xi)}{\int \mathbf{d}\boldsymbol{\xi} T(\xi)}\right] \\
&= e^{-\left\langle \phi_0^2(\mathbf{x}) \right\rangle} \\
&\approx e^{-\sigma_\phi^2}
\end{aligned}$$

where σ_ϕ^2 is the variance of the phase over the aperture (the last line only follows under the assumption of ergodicity.) Note that the assumption of good turbulence compensation allows us to assume that the phase of the wavefront after compensation is statistically homogeneous (giving line 3) and that the correlation of the phase is $\ll 1$ in the aperture (giving line 5). The approximation

$$\mathcal{S} \approx e^{-\sigma_\phi^2} \tag{12.57}$$

is called the extended Maréchal approximation, and is empirically valid for up to 2 radians of RMS phase error.

Now we develop a mathematical expression for the RMS phase variation σ_ϕ, which will allow us to compute the Strehl of a compensated wavefront. To simulate wavefront compensation, we expand the full aberrated wavefront in terms of Zernike modes (an overview is given in 12.D) and remove the lower order terms.

In general, if we consider perfect compensation of some wavefront up to Zernike order N, the phase of the wavefront after correction ϕ' is

$$\phi'(r,\theta) = \phi(r,\theta) - \sum_{k=1}^{N} a_k Z_k(r,\theta) = \sum_{k=1}^{\infty} a_k Z_k(r,\theta) - \sum_{k=1}^{N} a_k Z_k(r,\theta). \tag{12.58}$$

The squared phase is

$$\phi'^2 = \phi^2 - 2\sum_{k=1}^{N} |a_k|^2 + \sum_{k=1}^{N} |a_k|^2 = \phi^2 - \sum_{k=1}^{N} |a_k|^2 \tag{12.59}$$

so the phase variance is

$$\sigma_N^2 = \left\langle \phi'^2 \right\rangle = \left\langle \phi^2 \right\rangle - \sum_{k=1}^{N} \left\langle |a_k|^2 \right\rangle \tag{12.60}$$

where the N subscript denotes perfect compensation up to the Nth Zernike mode. [Nol76] showed that the above could be expressed as

$$\sigma_N^2 = A_N \left(\frac{d}{r_0}\right)^{5/3}. \tag{12.61}$$

Note that since the piston term doesn't really affect imaging performance, σ_1^2 gives phase variation for uncompensated turbulence. The first few values of A_N are given below:

A_1	A_2	A_3	A_4	A_5	A_6	A_7	A_8	A_9	A_{10}
1.030	0.582	0.134	0.111	0.088	0.065	0.059	0.053	0.046	0.040

12.5.2 Modal representation of Kolmogorov turbulence

In this subsection we express Kolmogorov turbulence as a covariance matrix of the Zernike mode coefficients a_j. This is called the Zernike matrix representation of Kolmogorov turbulence, and allows us to compute the probability distributions of the a_j coefficients.

Suppose a wavefront has some instantaneous phase distribution $\phi(r,\theta)$. We can expand the phase in terms of the Zernike modes

$$\phi(r,\theta)=\sum_j a_j Z_j(r/R,\theta).$$

Noll claims that the coefficients a_j can be considered 0 mean Gaussian random variables. The 0 mean assumption makes sense, and the Gaussian assumption is probably an approximation since we are working in the context of correlation theory. Anyways, defining $\rho = r/R$, the covariance (correlation) matrix of the a_j coefficients is

$$\Sigma_{ij}=\left\langle a_i a_j^*\right\rangle=\int \mathbf{d}\boldsymbol{\rho}\int \mathbf{d}\boldsymbol{\rho}' W(\rho)W(\rho')Z_i(\rho,\theta)Z_j(\rho',\theta')\left\langle \phi(R\rho,\theta)\phi^*(R\rho',\theta')\right\rangle. \tag{12.62}$$

In the following, I use f as the *cyclic* spatial wavenumber to distinguish from k, the *angular* spatial wavenumber from prior sections, and apologize in advanced for any confusion. Using the spectral representations of the Zernike polynomials

$$\begin{aligned}\Sigma_{ij}&=\iint \mathbf{d}\boldsymbol{\rho}\mathbf{d}\boldsymbol{\rho}'\iint \mathbf{df}\mathbf{df}' Q_i(f,\psi)Q_j^*(f',\psi')e^{2\pi i\mathbf{f}\cdot\boldsymbol{\rho}}e^{-2\pi i\mathbf{f}'\cdot\boldsymbol{\rho}'}\left\langle \phi(R\rho,\theta)\phi^*(R\rho',\theta')\right\rangle\\&=\iint \mathbf{df}\mathbf{df}' Q_i(f,\phi)Q_j^*(f',\phi')\Phi_p(f,f')\end{aligned} \tag{12.63}$$

where the Kolmogorov phase PSD in cyclic wavenumber units[4] is

$$\Phi_p(f,f')\approx 0.023 r_0^{-5/3}f^{-11/3}\delta(f-f').$$

Then, the covariance matrix can be expressed as

$$\Sigma_{ij}=\iint \mathbf{df}\mathbf{df}' Q_i(f,\phi)Q_j^*(f',\phi')\Phi_p(f/R,f'/R) \tag{12.64}$$

---

[4] Note that the spectral density carries units which also much change when going between cyclic and angular frequency.

where the spectral density is

$$\begin{aligned}\Phi_p(f/R, f'/R) &= 0.023 R^{-4} r_0^{-5/3} (f/R)^{-11/3} \delta\left(\frac{\mathbf{f}-\mathbf{f}'}{R}\right) \\ &= 0.023 R^{-2} r_0^{-5/3} (f/R)^{-11/3} \delta(\mathbf{f}-\mathbf{f}') \\ &= 0.023 \left(\frac{R}{r_0}\right)^{5/3} f^{-11/3} \delta(\mathbf{f}-\mathbf{f}').\end{aligned} \tag{12.65}$$

Note the extra factor of R^{-4} that was moved from the differential term in the integral into the above definition to make it non-dimensional.

Equation 12.64 is the Zernike matrix representation of Kolmogorov turbulence; the integrals in the matrix can all be evaluated in closed form. It turns out that the piston autocorrelation term, Σ_{11}, is infinite, but since piston does not affect imaging, the term is ignored. Additionally, the Σ_{1j} and Σ_{j1} terms for $j \neq 1$ are all 0, so the Kolmogorov turbulence matrix is usually expressed without the first row and first column.

Notably, Σ_{ij} is not diagonal. The Zernike mode amplitudes for Kolmogorov turbulence are not statistically independent. In the context of wavefront correction, the optimum correction "shapes" are the eigenfunctions of a diagonal covariance matrix, so the Zernike modes are technically not ideal. Another set of functions called the Karhunen-Loève (KL) functions do have a diagonal covariance matrix, but these functions cannot be expressed in closed form for Kolmogorov turbulence. Comparisons between the two bases suggests that the Zernike modes are often "good enough", anyways. This covariance matrix can be used to simulate random draws of Kolmogorov turbulence; see §13 for more details.

12.5.3 PSDs for Zernike-compensated turbulence

In this subsection we derive the form of the wavefront phase PSD when an arbitrary number of low-order Zernike modes are perfectly compensated. First we derive the shape of the wavefront after Zernike modes up to some number n are compensated. This can be useful for numerically simulating wavefronts downstream some primary AO system. Note that we cannot simply fix our aperture of radius R at one particular spot, and compensate the wavefront only within that fixed aperture, since the compensation should hold over the entire wavefront: regardless of where we place our aperture, the wavefront within that aperture should have no low order Zernike modes. In other words, the wavefront should have no low order Zernike modes *anywhere*. The radius R sets the scale over which modal compensation occurs. If we want our wavefront to have no low order Zernike modes over some physical telescope aperture, then R corresponds to the telescope aperture radius. However, we might also consider making R some length scale d smaller than the telescope aperture; as an example, this might correspond to active wavefront correction by a segmented mirror with d-sized segments.

Suppose we have some turbulence aberrated wavefront with phase field ϕ. We imagine placing an aperture of radius R at some location $\mathbf{r}$. The wavefront within this aperture can be expanded in terms of the Zernike modes. Denote $\mathbf{r}_1$ a vector pointing from the center of the aperture some point within the aperture. The phase at that point is

$$\phi(\mathbf{r}+\mathbf{r}_1) = \sum_{j=1}^{\infty} a_j(\mathbf{r}) Z_j(\mathbf{r}_1/R). \tag{12.66}$$

The weighting factor $a_j(\mathbf{r})$ is computed using Fourier's trick (equation 12.209)

$$a_j(\mathbf{r}) = \frac{1}{R^2}\int d^2\boldsymbol{\mathscr{r}}\, W(\boldsymbol{\mathscr{r}}/R) Z_j(\boldsymbol{\mathscr{r}}/R)\phi(\mathbf{r}+\boldsymbol{\mathscr{r}}) \tag{12.67}$$

where the integration variable $\boldsymbol{\mathscr{r}}$ traverses the aperture. In terms of convolution, we can write

$$a_j(\mathbf{r}) = \frac{1}{R^2}\left[W(\boldsymbol{\mathscr{r}}/R) Z_j(\boldsymbol{\mathscr{r}}/R) \;\circledast\; \phi(\boldsymbol{\mathscr{r}})\right](\mathbf{r}). \tag{12.68}$$

Therefore, the phase at the point $\mathbf{r} + \mathbf{r}_1$ can be written as

$$\phi(\mathbf{r} + \mathbf{r}_1) = \sum_{j=1}^{\infty} \frac{1}{R^2} \left[W(\boldsymbol{r}/R) Z_j(\boldsymbol{r}/R) \ \circledast \ \phi(\boldsymbol{r}) \right](\mathbf{r}) \times Z_j(\mathbf{r}_1/R). \tag{12.69}$$

Now suppose that Zernike modes up to order m were compensated within the aperture. We can write the compensated wavefront ϕ_m as

$$\phi_m(\mathbf{r} + \mathbf{r}_1) = \phi(\mathbf{r} + \mathbf{r}_1) - \sum_{j=1}^{m} \frac{1}{R^2} \left[W(\boldsymbol{r}/R) Z_j(\boldsymbol{r}/R) \ \circledast \ \phi(\boldsymbol{r}) \right](\mathbf{r}) \times Z_j(\mathbf{r}_1/R). \tag{12.70}$$

We now apply a Fourier transform with respect to the variable $\mathbf{r}$. Recalling that Fourier transforms have the following "shift" property for a function g

$$\mathcal{F} g(\mathbf{r} + \mathbf{r}_1) = e^{2\pi i \mathbf{f} \cdot \mathbf{r}_1} \hat{g}(\mathbf{f}) \tag{12.71}$$

and applying the convolution theorem, the Fourier transform of the corrected phase field is

$$\begin{aligned} \mathcal{F} \phi_m(\mathbf{r} + \mathbf{r}_1) &= \mathcal{F} \left[\phi(\mathbf{r} + \mathbf{r}_1) - \sum_{j=1}^{m} \frac{1}{R^2} \left[W(\boldsymbol{r}/R) Z_j(\boldsymbol{r}/R) \ \circledast \ \phi(\boldsymbol{r}) \right](\mathbf{r}) \times Z_j(\mathbf{r}_1/R) \right] \\ e^{2\pi i \mathbf{f} \cdot \mathbf{r}_1} \hat{\phi}_m(\mathbf{f}) &= e^{2\pi i \mathbf{f} \cdot \mathbf{r}_1} \hat{\phi}(\mathbf{f}) - \hat{\phi}(\mathbf{f}) \sum_{j=1}^{m} Z_j(\mathbf{r}_1/R) \ \frac{1}{R^2} \mathcal{F} \left[W(\boldsymbol{r}/R) Z_j(\boldsymbol{r}/R) \right]^* . \end{aligned} \tag{12.72}$$

Referring to 12.213, the remaining Fourier transform evaluates to $Q(R\mathbf{f})$. So,

$$\begin{aligned} e^{2\pi i \mathbf{f} \cdot \mathbf{r}_1} \hat{\phi}_m(\mathbf{f}) &= e^{2\pi i \mathbf{f} \cdot \mathbf{r}_1} \hat{\phi}(\mathbf{f}) - \hat{\phi}(\mathbf{f}) \sum_{j=1}^{m} Z_j(\mathbf{r}_1/R) Q_j^*(R\mathbf{f}) \\ e^{2\pi i \mathbf{f} \cdot \mathbf{r}_1} \hat{\phi}_m(\mathbf{f}) &= \hat{\phi}(\mathbf{f}) \left[e^{2\pi i \mathbf{f} \cdot \mathbf{r}_1} - \sum_{j=1}^{m} Z_j(\mathbf{r}_1/R) Q_j^*(R\mathbf{f}) \right] . \end{aligned} \tag{12.73}$$

Noting that $\Phi(\mathbf{f}) = \left\langle \hat{\phi}(\mathbf{f}) \hat{\phi}^*(\mathbf{f}) \right\rangle$, the spectral density Φ_m of the compensated phase field ϕ_m

is

$$\begin{aligned}\Phi_m(\mathbf{f},\mathbf{r}_1) = \Phi(\mathbf{f})\Bigg[&1 - e^{2\pi i\mathbf{f}\cdot\mathbf{r}_1}\sum_{j=1}^{m} Z_j\left(\frac{\mathbf{r}_1}{R}\right)Q_j(R\mathbf{f}) - e^{-2\pi i\mathbf{f}\cdot\mathbf{r}_1}\sum_{j=1}^{m} Z_j\left(\frac{\mathbf{r}_1}{R}\right)Q_j^*(R\mathbf{f})\\ &+\sum_{i=1}^{m}\sum_{j=1}^{m} Z_i\left(\frac{\mathbf{r}_1}{R}\right)Z_j\left(\frac{\mathbf{r}_1}{R}\right)Q_i^*(R\mathbf{f})Q_j(R\mathbf{f})\Bigg]\\ \equiv{}& H_m(\mathbf{f},\mathbf{r}_1)\Phi(\mathbf{f}).\end{aligned} \tag{12.74}$$

The term in the brackets acts as a high-pass filter, which we have denote $H_m(\mathbf{f},\mathbf{r}_1)$. Applying this filter to a normal phase PSD gives the compensated PSD. Notice that the filter has an $\mathbf{r}_1$ dependence. We might consider approximating the filter by averaging out the $\mathbf{r}_1$ dependence. In other words, we can average H_m over the aperture as follows:

$$\bar{H}_m(\mathbf{f}) \equiv \frac{1}{\pi R^2}\int_0^{2\pi}\int_0^R H_m(\mathbf{f},\mathbf{r}_1)d\mathbf{r_1}. \tag{12.75}$$

We can reintroduce the aperture function $W(\mathbf{r}_1/R)$ to extend the integration bounds. Recalling that W as defined in this treatment carries a factor of $1/\pi$,

$$\bar{H}_m(\mathbf{f}) = \frac{1}{R^2}\int_0^{2\pi}\int_0^\infty W(\mathbf{r}_1/R)H_m(\mathbf{f},\mathbf{r}_1)d\mathbf{r_1}. \tag{12.76}$$

Applying the above integration to H_m and identifying the Fourier transform with respect to $\mathbf{r}_1$ as well as the Zernike mode orthogonality relation, we find

$$\begin{aligned}\bar{H}_m(\mathbf{f}) &= 1 - 2\sum_{j=1}^{m}\left|Q_j(R\mathbf{f})\right|^2 + \sum_{i=1}\sum_{j=1}\delta_{ij}Q_i^*(R\mathbf{f})Q_j(R\mathbf{f})\\ &= 1 - \sum_{j=1}^{m}\left|Q_j(R\mathbf{f})\right|^2.\end{aligned} \tag{12.77}$$

For future reference, below are the filters for removing piston and piston/tip/tilt modes.

Note that these filters are isotropic, though the filters generally are not isotropic.

$$\begin{aligned} H_1(f) &= 1 - \left(\frac{J_1(2\pi f R)}{\pi f R}\right)^2 \\ H_3(f) &= 1 - \left(\frac{J_1(2\pi f R)}{\pi f R}\right)^2 - 4\left(\frac{J_2(2\pi f R)}{\pi f R}\right)^2. \end{aligned} \tag{12.78}$$

Also note that we are not limited to removing all modes up to some m. We can pick and choose which modes to remove in the summation term of the filter equation 12.77.

Appendix - Chapter 12

12.A Riemann-Stieltjes integrals

This section follows material from [Ane12].

12.A.1 Introduction

Riemann-Stieltjes integrals are a powerful tool that allow us to treat treat discrete and continuous random variables under a unified mathematical framework, as well as define a spectral decomposition for stochastic processes. Formally, the Riemann-Stieltjes integral is an extension of the Riemann integral. Recall that a Riemann integral is defined as the limit

$$\int g(x)dx = \lim_{\max|x_{i+1}-x_i|\to 0} \sum_{i=1}^{n} g(\xi_i)(x_{i+1} - x_i) \tag{12.79}$$

where ξ_i can be anywhere in the interval (x_{i-1}, x_i). Supposing the limit exists, g is “Riemann-integrable.”

For a strictly increasing function F, the Riemann-Stieltjes (RS) integral is defined as

$$\int g(x)dF(x) = \lim_{\max|x_{i+1}-x_i|\to 0} \sum_{i=1}^{n} g(\xi_i)[F(x_{i+1}) - F(x_i)]. \tag{12.80}$$

Note that if F is differentiable, we have

$$\begin{aligned} \frac{dF}{dx} &= p(x) \\ dF &= p(x)dx \\ \int g(x)dF(x) &= \int g(x)p(x)dx. \end{aligned} \tag{12.81}$$

However, F doesn't need to be differentiable for the Riemann-Stieltjes integral to exist, making it more general than a normal Riemann integral.

In the context of probability, the form of the above suggests interpreting $p(x)$ and $F(x)$ as the probability density function (PDF) and cumulative distribution function (CDF) of random variable x, respectively, leading us to define the expected value of x as

$$\langle x \rangle = \int x dF(x). \tag{12.82}$$

This works in both discrete, continuous, and mixed scenarios.

12.A.2 A more rigorous framework

To prove the various properties of the RS integral, we begin by defining it exclusively for step functions. Let $h(x)$ be a step function whose argument x is within the real interval I, which has been subdivided into multiple intervals I_i such that $I = \cup_{i=1}^{n} I_i$. The subdivisions occur on x_i. The step function can be expressed as

$$h(x) = \sum_{i=1}^{n} c_i 1\{x \in I_i\} \tag{12.83}$$

for some set of constants $c_1 ... c_n$. (The function 1{cond} just returns 1 if cond is true.) For a monotonic real-valued function F, we define the RS integral of $h(x)$ as

$$\int h(x) dF(x) = \sum_{i=1}^{n} c_i \left[F(x_{i+1}) - F(x_i) \right]. \tag{12.84}$$

Now we define RS integration proper. First of all, for an increasing function F defined on interval I, we say a function g (also defined on I) is RS integrable if for every $\epsilon > 0$ there are step functions h_1, h_2 with $h_1 \leq g \leq h_2$ and

$$\int_I h_2(x) dF(x) - \int_I h_1(x) dF(x) < \epsilon. \tag{12.85}$$

Loosely, this means that g is RS integrable if it can be arbitrarily well approximated by step functions (which is, I think most functions of interest in physics).

If g is RS integrable, the actual integral value is formally

Definition 12.A.1

$$\int g(x) dF(x) \equiv \sup\{ \int_I h(x) dF(x) : h \leq g \} \tag{12.86}$$

where h is a step function. sup denotes supremum. This definition leads to useful lemmas, such as:

Lemma 12.A.2 *If g is continuous and bounded on I, and F is increasing and bounded on I, then g is RS integrable on I.*

From the above definition it is also possible to express the RS integral as the limit of a sum, like a Riemann integral. Suppose g is continuous and bounded on I. We construct the step functions

$$\begin{aligned} h_1(x) &= \sum_{i=1}^{n} \inf_{x \in I_i} g(x)\ 1\{x \in I_i\} \\ h_2(x) &= \sum_{i=1}^{n} \sup_{x \in I_i} g(x)\ 1\{x \in I_i\}. \end{aligned} \tag{12.87}$$

By construction, $h_1 \leq g \leq h_2$. Furthermore, we have

$$\int_I h_1(x) dF(x) \leq \sum_{i=1}^{n} g(\xi_i) \left[F(x_{i+1}) - F(x_i)\right] \leq \int_I h_2(x) dF(x) \tag{12.88}$$

also by construction. Now, since g is continuous we can let $\epsilon \to 0$ (with ϵ as defined in equation 12.85) simply by making the subdivisions of the interval I smaller and smaller. Thus, the outer terms in equation 12.88 squeeze the inner term to 12.86, showing

$$\int g(x) dF(x) = \lim_{\max|x_{i+1} - x_i| \to 0} \sum_{i=1}^{n} g(\xi_i) \left[F(x_{i+1}) - F(x_i)\right]. \tag{12.89}$$

This is the form commonly used in most texts.

Another useful result is that RS integrals allow for integration by parts. The following result can be derived by expanding the sum in 12.A.2 (it is useful to fix ξ_i to x_{i+1} in this case).

Lemma 12.A.3

$$\int_a^b g(x) dF(x) = g(b)F(b) - g(a)F(a) - \int_a^b F(x) dg(x). \tag{12.90}$$

12.A.3 Integration in higher dimensions

In normal Riemann integration we repeatedly sum the product of the integrand and the interval size $x_{i+1} - x_i$. In RS integration we do something similar, except that the interval size is replaced by the *transformed* interval size $F(x_{i+1}) - F(x_i)$. The question is, how do we extend this process to higher dimensions? We look first at the two dimensional case. Suppose we have an interval $I = [x_1, x_2] \times [y_1, y_2]$. The area of this interval is

$$|I| = (x_2 - x_1)(y_2 - y_1) = x_2 y_2 - x_2 y_1 - x_1 y_2 + x_1 y_1. \tag{12.91}$$

To expand RS integration to higher dimensions, we define the transformed area as

$$|F(I)| = F(x_2, y_2) - F(x_2, y_1) - F(x_1, y_2) + F(x_1, y_1). \tag{12.92}$$

The above can be further extended to higher dimensions. Suppose our multidimensional interval is now of the form $[a_1, b_1] \times [a_2, b_2] \times ....$ Define the operator

$$\Delta_j F(I) = F(x_1, ...x_{j-1}, b_j, x_{j+1}, ...x_n) - F(x_1, ...x_{j-1}, a_j, x_{j+1}, ...x_n) \tag{12.93}$$

which computes a difference in F according to the bounds of I along the jth axis. Repeatedly applying this operator yields

$$|F(I)| = \Delta_1 ... \Delta_n F(I). \tag{12.94}$$

Then the RS integral is given by

$$\int_I g dF = \lim_{\max |I_i| \to 0} \sum_{i=1}^{n} g(\xi_i) |F(I_i)|. \tag{12.95}$$

You can see for $F(I) = |I|$, RS integration reduces to Riemann integration.

One important property that extends from 1D RS integrals to nD RS integrals is the reduction of an RS integral to a normal integral if F is differentiable. In particular, if we suppose that some continuous f exists such that

$$f(x_1 ... x_n) = \frac{\partial^n}{\partial x_1 ... \partial x_n} F(x_1 ... x_n) \tag{12.96}$$

then the following is true for RS integrable function $g(x)$:

$$\int g(x)dF(x) = \int g(x)f(x)\mathbf{dx} \tag{12.97}$$

where bolding emphasizes that the above is a multidimensional integral. This holds true because the difference operator approximates the partial derivative when F is differentiable.

$$\frac{\partial F}{\partial x_1} \approx \frac{F(b_1, x_2 ... x_n) - F(a_1, x_2 ... x_n)}{b_1 - a_1} = \frac{\Delta_1 F}{b_1 - a_1} \tag{12.98}$$

so

$$\begin{aligned} |F(I)| &= \Delta_1 ... \Delta_n F(I) \\ &= \Delta_1 ... \Delta_{n-1} \frac{\partial F}{\partial x_n}(b_n - a_n) \\ &= \frac{\partial^n F}{\partial x_1 ... \partial x_n}|I| \\ &= f|I|. \end{aligned} \tag{12.99}$$

Substituting the above into 12.95 completes the proof.

12.A.4 Connections to probability

Let's say we have a CDF for a random variable X:

$$F(x) = P(X \leq x). \tag{12.100}$$

Note that F is monotonic. We can also use F in an RS integral. For one, consider the quantity

$$\int x dF(x) \approx \sum_{i=1}^{n} \xi_i \left[F(x_{i+1}) - F(x_i)\right] \tag{12.101}$$

where $\xi_i \in [x_i, x_{i+1})$. The approximation becomes better as we make the partitions of the summation smaller, as per the previous section. Note that critically, for small intervals, we have

$$P(X = \xi_i) \approx F(\xi_{i+1}) - F(\xi_i) \tag{12.102}$$

with equality reached in the limit, so immediately we recognize equation 12.101 as the expected value. Equation 12.102 is the basis for the statistical utility of RS integrals. The expected value is

$$\langle X \rangle = \int x dF(x). \tag{12.103}$$

This allows us to define things like variance, also using RS integrals. Similarly, if we consider a function of a random variable X (denote this $Y \equiv g(X)$ for an RS-integrable function g which maps $\mathbb{R} \to \mathbb{R}$), we can derive a related result:

$$\langle Y \rangle = \int g(x) dF(x). \tag{12.104}$$

12.A.5 Connections to stochastic processes

A definite signal $\zeta(t)$ can be expanded as a sum of harmonic oscillations through the Fourier transform. According to [YN63], it "turns out" that a similar expansion exists for stationary random processes.[5] For a stationary random process $\xi(t)$, we can write the Fourier-RS integral:

$$\xi(t) = \int_{-\infty}^{\infty} e^{i\omega t} d\phi(\omega) \tag{12.105}$$

where $d\phi(\omega)$ are the random complex amplitudes of each frequency component. Note that the above expression only makes sense as a RS integral, since $d\phi$ is random and ϕ will not be differentiable everywhere.

---

[5] But Yaglom doesn't cite any sort of proof, and only mentions that this theorem was proved after the Wiener-Khinchin theorem.

12.B Stochastic processes and fields

This section mainly follows the brief primer on stochastic processes given in [TSC61], with extra material from [YN63]. The purpose is to give a practical background for all the math that goes into the statistical theory of turbulence.

12.B.1 What is a stochastic process?

Generally speaking, a stochastic process is a collection of random variables. Often times this collection takes the form of a time series (real world examples might be a noisy voltage signal, or the amount of rainfall at a location over a sequence of days), but it is important to remember that a stochastic process can also be interpreted as a single multidimensional random variable. A stochastic processes can be discrete or continuous, and finite or infinite dimensional.

To put this into math, consider a stochastic process represented by the random function $\xi(t)$, which takes in an argument $t \in T$ and returns a random variable given t. This random function is fully defined if we can specify the all the multivariable CDFs

$$F_{t_1,t_2...t_N}(x_1, x_2...x_n) = P(\xi(t_1) < x_1 \wedge \xi(t_2) < x_2 \wedge ...) \tag{12.106}$$

for any N elements $\{t_1...t_n\}$ we choose from T. In other words, we need to specify the multivariate CDF as well as all possible joint CDFs.

The above is the most general form of a stochastic process of a single variable. However, the above looks pretty tedious, so stochastic processes are often instead defined as deterministic functions of random variables. You can see how given such a function (and the CDFs of the random variables), we could compute all variations of equation 12.106. An example would be

$$\xi(t) = \cos(t + \phi) \tag{12.107}$$

where ϕ is a random variable uniformly distributed across $[0, 2\pi)$. Much of the mathematical framework of stochastic processes rely on such functions, because they are much easier to work with.

Before proceeding, I note some terminology. A random variable has no definite value, but can be sampled to yield many definite values. These definite values are called *realizations* of the random variable. Similarly, a random (stochastic) process can be sampled to yield a definite collection of values: a realization of the stochastic process.

12.B.2 Moments

Stochastic processes can be complicated, so in studying them it is often useful to look at moments, which reflect the general characteristics of a process. Moments take the general form

$$\mu_{m_1,m_2...m_n}(t_1, t_2...t_n) = \langle \xi^{m_1}(t_1)\xi^{m_2}(t_2)...\xi^{m_n}(t_n)\rangle \tag{12.108}$$

where $\langle\rangle$ denotes the *ensemble average*; the value you would get if you obtained infinitely many realizations of the quantity within the brackets and averaged them (loosely, the expected value). The above can also be written as an n-dimensional RS integral (as alluded to in §12.A)

$$\mu_{m_1,m_2...m_n}(t_1, t_2...t_n) = \int_{-\infty}^{\infty} ... \int_{-\infty}^{\infty} x_1^{m_1}...x_n^{m_n} dF_{t_1...t_n}(x_1...x_n). \tag{12.109}$$

If F admits a probability density

$$f_{t_1...t_n}(x_1...x_n) = \frac{\partial^n F_{t_1...t_n}(x_1...x_n)}{\partial x_1...\partial x_n} \tag{12.110}$$

then as expected we can write the more familiar equation for moments:

$$\mu_{m_1,m_2...m_n}(t_1, t_2...t_n) = \int_{-\infty}^{\infty} ... \int_{-\infty}^{\infty} x_1^{m_1}...x_n^{m_n} f_{t_1...t_n}(x_1...x_n) dx_1...dx_n \tag{12.111}$$

which follows as per 12.97.

Now we go over some examples of commonly used moments. The first is the mean

$$\mu_1(t) = \langle \xi(t) \rangle = \int_{-\infty}^{\infty} x dF_t(x). \tag{12.112}$$

Another moment is the correlation:[6]

$$\mu_{11}(t_1, t_2) = \langle \xi(t_1)\xi(t_2) \rangle = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x_1 x_2 dF_{t_1,t_2}(x_1, x_2). \tag{12.113}$$

Often times you will see the mean written as μ and the correlation as $B(t_1, t_2)$. Note that the correlation for a stochastic process is different than the cross-correlation formula for deterministic signals.

Aside: before, we said a random process is only defined when all the multidimensional distribution functions for the process are defined at all times. Regardless, significant effort has been put into the the development *correlation theory* whereby stationary processes are characterized only by their mean and correlation (first and second moments), because the formal definition of a random process is clunky to work with. Since a random process isn't defined fully by its mean and correlation, correlation theory cannot as be regarded as complete; however, for real world use it's often good enough.[7]

12.B.3 Stationarity

A random function $\xi(t)$ is regarded as stationary if all of its multidimensional distribution functions F remain the same as we translate forwards in time. Mathematically:

$$F_{t_1+\tau \ldots t_n+\tau}(x_1 \ldots x_n) = F_{t_1 \ldots t_n}(x_1 \ldots x_n) \tag{12.114}$$

---

[6] Or autocorrelation or covariance.

[7] Especially since many random processes turn out be Gaussian, and *are* fully defined by mean and correlation. Also, the central limit theorem.

must hold for any choice of n, $t_1...t_n$, and τ.

While 12.114 seems somewhat abstract, this condition applies mathematical restrictions on the distribution functions. For instance, if we look at the 1D distribution functions for the process $\xi(t)$, applying the previous constraint we find

$$F_{t+\tau}(x) = F_t(x) \tag{12.115}$$

for all τ and t, which in turn implies $F_t(x)$ cannot depend on t at all. This implies that the mean μ cannot depend on time t.

$$\begin{aligned} \langle \xi(t) \rangle &= \langle \xi(t+\tau) \rangle \\ \langle \xi(t) \rangle &= \text{constant}. \end{aligned} \tag{12.116}$$

In the 2D case we have

$$F_{t_1+\tau, t_2+\tau}(x_1, x_2) = F_{t_1,t_2}(x_1, x_2). \tag{12.117}$$

Yaglom claims that this means the 2D distribution functions can only depend on the difference in times $t_2 - t_1$.[8] This means that the correlation must have the form $B(\tau)$, $\tau = t_2 - t_1$.

Yaglom makes no claim about what to expect in higher dimensions, but it seems the restrictions on F become less and less useful. For instance, in the 3D case, I think that F must depend on an argument of the form

$$at_1 + bt_2 + ct_3 \tag{12.118}$$

where we require $a + b + c = 0$.

However, the above definition of stationarity is both extremely stringent and hard to verify for a real world process. Thus, in light of correlation theory, we define a process as

---

[8] While the converse is easy to prove. This claim is not immediately obvious to me. You can prove this if F admits a power series expansion.

wide-sense stationary (WSS) if only the 1D and 2D distribution functions look stationary. In other words, a process is WSS if its mean is constant and the correlation between two times t_1 and t_2 depends only on the difference $\tau = t_2 - t_1$. In practice, Yaglom notes that most processes we encounter that are WSS end up fulfilling the "true" stationarity condition anyways.

12.B.4 Ergodicity

In the context of random processes which vary over time, ergodicity refers to the equivalence of time averages and ensemble averages. In the context of correlation theory, the mean and correlation of a ergodic random process can be computed either averaging over times or realizations. For a continuous process:

$$\begin{aligned} \mu = \langle \xi(t) \rangle &= \lim_{T\to\infty} \frac{1}{T} \int_0^T \xi(t) dt \\ B(\tau) = \langle \xi(t+\tau)\xi(t) \rangle &= \lim_{T\to\infty} \frac{1}{T} \int_0^T \xi(t+\tau)\xi(t) dt. \end{aligned} \tag{12.119}$$

Physically, ergodicity is important because it implies that all the information carried by a process (mean, correlation) can be inferred from a single realization. Such a process would be much easier to work with in the real world: you only have to run one long experiment instead of multiple experiments.

As a side note, an ergodic process is necessarily stationary. Intuitively, if the process is not stationary, a time average of a single realization will not be the same as an ensemble average at a single time. However, the converse is not necessarily true: a process can be stationary but not ergodic. An example is the following (physical process): at each timestep, you choose between either a six or four sided dice with equal probability, then repeatedly roll and record the number. This process is stationary. A time average will converge to either 3.5 or 2.5 depending on which dice you chose but the ensemble average will be 3.

12.B.5 Spectral representation

As alluded to in §12.A, a stationary random process can be expressed as a Fourier-RS integral:

$$\xi(t) = \int_{-\infty}^{\infty} e^{i\omega t} d\phi(\omega) \tag{12.120}$$

where $d\phi$'s are random complex amplitudes. While Yaglom never proves this outright, he does motivate it as follows. Suppose we want to find the "simplest" stationary random function. A straightforward way to construct such a process is to multiply an analytic function $f(t)$ by a random variable ξ:

$$\xi(t) = \xi f(t). \tag{12.121}$$

If we now impose stationarity (with 0 mean, for simplicity) we find that

$$\langle \xi \rangle = 0 \tag{12.122}$$

and that the correlation

$$\langle \xi(t+\tau)\xi^*(t) \rangle = f(t+\tau) f^*(t) \left\langle |\xi|^2 \right\rangle \tag{12.123}$$

must be independent of t. This means that the product $f(t+\tau)f^*(t)$ must be independent of t, for all τ. Setting τ to 0 we find

$$\begin{gathered} |f(t)|^2 = \text{constant} \\ f(t) = re^{i\phi(t)} \end{gathered} \tag{12.124}$$

where $r \in \mathbb{R}$ and $\phi(t)$ is some real function. Back-substituting reveals that the quantity $\phi(t+\tau) - \phi(t)$ cannot depend t. Assuming ϕ is a nice, well-behaved function (read: differentiable) then we find

$$\frac{d\phi}{dt} = \text{constant} \equiv \omega. \tag{12.125}$$

In other words, $\phi(t)$ must be linear in t. As a result, stationarity has constrained $f(t)$ to take the form

$$f(t) = re^{i(\omega t+\theta)}. \tag{12.126}$$

Absorbing some constants into the random variable ξ, our "simplest" stationary random process must have the form

$$\xi(t) = \xi e^{i\omega t} \tag{12.127}$$

where ξ is a complex random variable with mean 0 controlling the amplitude and phase of each realization, and $\omega \in \mathbb{R}$ controls the frequency. The correlation function of this random process can be easily verified as

$$B(\tau) = \left\langle |\xi|^2 \right\rangle e^{i\omega\tau}. \tag{12.128}$$

In the context of signal theory, the coefficient, which is isolated by computing $B(0)$, is related to the average power of the oscillation.

Equation 12.127 provides us a "basis" of sorts to construct new stationary stochastic processes. We can construct a more general stochastic process as follows:

$$\xi(t) = \sum_i \xi_i e^{i\omega_i t}. \tag{12.129}$$

Recall that for this process to be stationary in the wide sense we require that the correlation only depend on the time difference τ:

$$\begin{aligned} \langle \xi(t+\tau)\xi^*(t) \rangle &= \left\langle \sum_i \xi_i e^{i\omega_i(t+\tau)} \sum_k \xi_k^* e^{-i\omega_k t} \right\rangle \\ &= \left\langle \sum_i |\xi_i|^2 e^{i\omega_i\tau} \right\rangle + \left\langle \sum_i \sum_{k\neq i} \xi_i \xi_k^* e^{i\omega_i(t+\tau)-i\omega_k t} \right\rangle. \end{aligned} \tag{12.130}$$

For the above to depend only on τ and not t, we require

$$\langle \xi_i \xi_k^* \rangle = 0;\ k \neq i. \tag{12.131}$$

In other words, the random variables ξ_i must be uncorrelated, so

$$B(\tau) = \sum_i \left\langle |\xi_i|^2 \right\rangle e^{i\omega_i \tau}. \tag{12.132}$$

To compute the average power, we again take $B(0)$ and find

$$B(0) = \sum_i \left\langle |\xi_i|^2 \right\rangle \tag{12.133}$$

so we see that the power is simply the sum of the powers of the individual harmonic oscillations.

The above summations are reminiscent of a Fourier series, so it is natural to expand equation 12.132 as a Fourier-like integral. Intuitively, since any function can be expressed as the limit of a Fourier series,[9] a stationary random process formed from summing harmonic oscillations of the form 12.127 can have any possible correlation function. Which is a roundabout way of saying that any stationary random function (in the context of correlation theory) can be expressed as the limit of a sum of harmonic oscillations. This can be shown mathematically using RS integrals. We have

$$B(\tau) = \int_{-\infty}^{\infty} e^{i\omega\tau} dF(\omega) \tag{12.134}$$

where $F(\omega)$ is a real, non-decreasing bounded function. This theorem was originally proved by Khinchin, who showed that any correlation function could be expressed as in 12.134, as well as the converse: that any function of the form 12.134 was the correlation function of

---

[9] at least from the physicist's perspective!

some stationary random process. Note how the RS integral is necessary here, because F might be discontinuous.

In a similar vein, it is also possible to express the random processes themselves as RS integrals; this fact was first proven by Kolmogorov, and a full proof is provided in [YN63]. To use Tatarski's notation, we can expand a stationary random process $\xi(t)$ as

$$\xi(t) = \int_{-\infty}^{\infty} e^{i\omega t} d\phi(\omega) \tag{12.135}$$

where the $d\phi(\omega)$ are random complex amplitudes (taking the place of ξ_i in 12.129). Like before, we can enforce wide-sense stationarity to determine some properties of $d\phi(\omega)$ and $\phi(\omega)$. Again assuming 0 mean for simplicity and expanding the RS integral as the limit of a sum, we find

$$\begin{aligned} \langle \xi(t) \rangle &= \int_{-\infty}^{\infty} e^{i\omega t} \langle d\phi(\omega) \rangle \\ &= \lim_{\max |\omega_{k+1} - \omega_k| \to 0} \sum_k e^{i\omega_k t} \langle \phi(\omega_{k+1}) - \phi(\omega_k) \rangle \\ &= 0 \end{aligned} \tag{12.136}$$

which implies $\langle \phi(\omega) \rangle =$ constant for all ω.[10] Furthermore, we can consider the correlation to impose more constraints. For brevity in the following derivation the dropped limits are

---

[10] Yaglom states $\langle \phi(\omega) \rangle = 0$. I think you can prove it this way. Suppose $\langle \phi(\omega) \rangle = \mu$. Then add and subtract μ within the ensemble average in 12.136 to remove the mean from the random process ϕ. So, without loss of generality we can just set $\langle \phi(\omega) \rangle = 0$ in the first place.

implied.

$$\begin{aligned}\langle\xi(t+\tau)\xi^*(t)\rangle &= \left\langle \int_{-\infty}^{\infty} e^{i\omega_1(t+\tau)} d\phi(\omega_1) \int_{-\infty}^{\infty} e^{-i\omega_2 t} d\phi(\omega_2) \right\rangle \\ &= \left\langle \sum_i e^{i\omega_i(t+\tau)} \left[\phi(\omega_{i+1}) - \phi(\omega_i)\right] \sum_k e^{-i\omega_k t} \left[\phi^*(\omega_{k+1}) - \phi^*(\omega_k)\right] \right\rangle \\ &= \sum_i \sum_k e^{i\omega_i \tau} e^{it(\omega_i - \omega_k)} \left\langle \left[\phi(\omega_{i+1}) - \phi(\omega_i)\right] \left[\phi^*(\omega_{k+1}) - \phi^*(\omega_k)\right] \right\rangle .\end{aligned} \tag{12.137}$$

To cancel out the cross terms and make sure that the correlation depends only on τ, the above requires

$$\left\langle \left[\phi(\omega_{i+1}) - \phi(\omega_i)\right] \left[\phi^*(\omega_{k+1}) - \phi^*(\omega_k)\right] \right\rangle = 0 \tag{12.138}$$

unless the intervals $[\omega_i, \omega_{i+1}]$ and $[\omega_k, \omega_{k+1}]$ are the same.[11] In other words, the $d\phi(\omega)$'s must be uncorrelated (ϕ has "uncorrelated increments").

12.B.6 Correlation and spectral distribution

We now return to the topic of the correlation function. As mentioned in the previous section, Khinchin proved that the correlation of a stationary random process $\xi(t)$ could be expressed as

$$B(\tau) = \int_{-\infty}^{\infty} e^{i\omega\tau} dF(\omega). \tag{12.139}$$

We now show that this is an immediate consequence of the fact that $\xi(t)$ admits a spectral representation. First, going back to the definition of correlation,

$$\begin{aligned}B(\tau) &= \langle\xi(t+\tau)\xi^*(t)\rangle \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{i\omega_1(t+\tau)} e^{-i\omega_2 t} \left\langle d\phi(\omega_1) d\phi^*(\omega_2)\right\rangle .\end{aligned} \tag{12.140}$$

---

[11] To make things simpler, we are assuming that ω_i and ω_k partition the domain of ω in the exact same way.

Now, since $B(\tau)$ must only depend on τ and not t, we necessarily require $\omega_1 = \omega_2$, collapsing the above double integral into a single integral.[12] We then get

$$B(\tau) = \int_{-\infty}^{\infty} e^{i\omega\tau} \left\langle |d\phi(\omega)|^2 \right\rangle . \tag{12.141}$$

Define the function $F(\omega)$ such that

$$\begin{aligned} B(\tau) &= \int_{-\infty}^{\infty} e^{i\omega\tau} dF(\omega) \\ dF(\omega) &= F(\omega + \Delta\omega) - F(\omega) = \left\langle |d\phi(\omega)|^2 \right\rangle \\ &= \left\langle |\phi(\omega + \Delta\omega) - \phi(\omega)|^2 \right\rangle . \end{aligned} \tag{12.142}$$

The above implies that F is nondecreasing. Note that the above equation only defines F up to an arbitrary additive constant. We thus set $F(-\infty) = 0$ for simplicity.

The function F is known as the *spectral distribution function* of the stationary process $\xi(t)$. The differential $dF(\omega)$ tells us how much of each harmonic component $e^{i\omega t}$ is in our process $\xi(t)$, on average, and thus $F(\omega)$ can be seen as a cumulative distribution function, reflecting how much of the $x(t)$ is made from harmonic oscillations with frequency less than ω. If we assume F is differentiable such that $F' = f$, we find

$$B(\tau) = \int_{-\infty}^{\infty} e^{i\omega\tau} f(\omega) d\omega \tag{12.143}$$

where f is known as the *spectral density* of $\xi(t)$. Thus, we see that if f exists, it and the autocorrelation B are Fourier pairs. This is the Wiener-Khinchin theorem. Equivalently, for the above to hold we can require that

$$\int_{-\infty}^{\infty} |B(\tau)| d\tau < \infty. \tag{12.144}$$

---

[12] Mathematically, we are saying the integrand must contain a term $\delta(\omega_1 - \omega_2)$.

This is a sufficient condition[13] for the Fourier representation of $B(\tau)$ to exist. In practice, 12.144 usually holds, meaning we can bypass using RS integrals and instead use the more familiar spectral density. Also in practice, the inverse Fourier transform tends to be more useful:

$$f(\omega) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-i\omega\tau} B(\tau) d\tau. \tag{12.145}$$

One use case is as follows. Suppose we want to simulate a stochastic process. If we can theoretically derive its correlation function, we can get a spectral density. Then, supposing the spectral density is real, simulation of the stochastic process involves Fourier-filtering Gaussian white noise with the square root of that spectral density profile. The validity of this process is easiest to see in the discrete case. A Gaussian white noise process, so named because it has on average equal power in all frequencies, can be written as

$$\nu(t) = \sum_k \nu_k e^{i\omega_k t} \tag{12.146}$$

where the ν_k's are uncorrelated Gaussian random variables with mean 0 and variance 1. Suppose we want to modify this process to have a given (real) spectral density $f(\omega)$. I propose that the following process $\xi(t)$ works:

$$\xi(t) = \sum_k \nu_k \sqrt{f(\omega_k)} e^{i\omega_k t}. \tag{12.147}$$

---

[13] Again, according to Yaglom. Absolute integrability guarantees the existence of the Fourier transform (by dominated convergence theorem), but I am not sure if it guarantees convergence.

We can check that this process has the correct spectral density by computing the correlation.

$$\begin{aligned} B(\tau) = \langle \xi(t+\tau)\xi(t) \rangle &= \left\langle \sum_i \nu_i \sqrt{f(\omega_i)} e^{i\omega_i(t+\tau)} \sum_k \nu_k \sqrt{f(\omega_k)} e^{-i\omega_k t} \right\rangle \\ &= \sum_i \sum_k \sqrt{f(\omega_i) f(\omega_k)} e^{i\omega_i(t+\tau)} e^{-i\omega_k t} < \nu_i \nu_k > \\ &= \sum_i f(\omega_i) e^{i\omega_i \tau} < \nu_i \nu_i > \\ &= \sum_i f(\omega_i) e^{i\omega_i \tau}. \end{aligned} \tag{12.148}$$

We immediately identify $B(\tau)$ and $f(\omega)$ as Fourier pairs, confirming that the process $\xi(\tau)$ does indeed have spectral density $f(\omega)$.

12.B.7 Spectral representation of nonstationary processes

Many phenomena (weather, the stock market, a paper's citation count) are nonstationary stochastic processes. In order to study such process a process $\xi(t)$ using the framework developed above for stationary processes, we might consider instead the difference function

$$\zeta_\tau(t) = \xi(t) - \xi(t-\tau). \tag{12.149}$$

If the mean of the process fluctuates on timescales longer than τ then those fluctuations will not affect the above difference, and $\zeta_\tau(t)$ can be regarded (at least approximately), as a stationary random function. Such processes whose difference functions are stationary are known as random functions with stationary increments.

With the framework of structure functions in place we can extend some of the theorems derived originally for stationary random processes to nonstationary random processes with stationary increments. We take the following as an axiom: just as how stationary random processes can be expressed using a Fourier-RS process, so to can random processes with stationary increments. For such a process $\xi(t)$, these expressions take the form

$$\xi(t) = \xi(0) + \int_{-\infty}^{\infty} (1 - e^{i\omega t}) d\phi(\omega). \tag{12.150}$$

I attempt to motivate this below. We can express $\xi(t)$ as a telescoping series

$$\begin{aligned} \xi(t) &= \xi(t) - \xi(t-\tau) + \xi(t-\tau) - \xi(t-2\tau) + ... + \xi(0) \\ &= \zeta_\tau(t) + \zeta_\tau(t-\tau) + \zeta_\tau(t-2\tau) + ... \\ &= \sum_{k=1}^{n} \zeta_\tau(k\tau) + \xi(0) \end{aligned} \tag{12.151}$$

such that $t = n\tau$.[14] We further know that ζ_τ is stationary, so that it can be expressed as

$$\zeta_\tau(t) = \int_{-\infty}^{\infty} e^{i\omega t} d\phi(\omega). \tag{12.152}$$

Combining the above two equations, we find

$$\xi(t) = \sum_{k=1}^{n} \int_{-\infty}^{\infty} e^{i\omega k\tau} d\phi(\omega) + \xi(0). \tag{12.153}$$

Assuming that we can interchange summation and integration, then

$$\begin{aligned} \xi(t) &= \int_{-\infty}^{\infty} \sum_{k=1}^{n} e^{i\omega k\tau} d\phi(\omega) + \xi(0) \\ &= \int_{-\infty}^{\infty} \frac{(e^{i\omega t} - 1)\, e^{i\omega\tau}}{e^{i\omega\tau} - 1} d\phi(\omega) + \xi(0). \end{aligned} \tag{12.154}$$

The above should be valid as long as τ is shorter than the timescale of fluctuations in the

---

[14] Note that we have implicitly assumed t a multiple of τ, so that the “anchor point” of our series was at time 0 (i.e. the extra term is our series is $\xi(0)$.) This isn’t necessary; we could anchor our series at any time. The point is to show the repeated sum of the difference function ζ_τ builds up the process ξ.

mean, so we can let τ be a constant very close to 0 to derive

$$\xi(t) = \int_{-\infty}^{\infty} \frac{e^{i\omega t} - 1}{i\omega\tau} d\phi(\omega) + \xi(0). \tag{12.155}$$

Absorb the constant τ into $d\phi(\omega)$, since even in doing so the function ϕ still has uncorrelated increments. We can also absorb the $i\omega$ wherever $\omega \neq 0$ for the same reason. The singularity at $\omega = 0$ is trickier to handle. Let's cut out the singularity from the RS integral as follows

$$\xi(t) = \lim_{\epsilon\to 0}\left[\int_{-\infty}^{-\epsilon} \left(e^{i\omega t} - 1\right) d\theta(\omega) + \int_{\epsilon}^{\infty} \left(e^{i\omega t} - 1\right) d\theta(\omega) + \int_{-\epsilon}^{\epsilon} \frac{e^{i\omega t} - 1}{i\omega\tau} d\phi(\omega)\right] + \xi(0) \tag{12.156}$$

where for clarity we have defined $d\theta = d\phi/i\omega\tau$. The last term in the limit might be non-zero; recalling that the RS integral is the limit of a sum, we can write

$$\begin{aligned}\lim_{\epsilon\to 0}\int_{-\epsilon}^{\epsilon} \frac{e^{i\omega t} - 1}{i\omega\tau} d\phi(\omega) &= \lim_{\epsilon\to 0} \frac{e^{i\omega t} - 1}{i\omega\tau}\left[\phi(\epsilon) - \phi(-\epsilon)\right];\ \omega \in [-\epsilon, \epsilon] \\ &= \lim_{\omega\to 0} \frac{e^{i\omega t} - 1}{i\omega\tau} \times \lim_{\epsilon\to 0}\left[\phi(\epsilon) - \phi(-\epsilon)\right] \\ &= \xi_1 t\end{aligned} \tag{12.157}$$

where ξ_1 is a random variable defined as

$$\xi_1 = \frac{1}{\tau}\lim_{\epsilon\to 0}\left[\phi(\epsilon) - \phi(-\epsilon)\right]. \tag{12.158}$$

The term in the brackets is the probability the random process $\xi(t)$ has a component with frequency $\omega \in [-\epsilon, \epsilon]$, i.e. a 0-frequency component.

Thus, we can generally express the random process with stationary increments $\xi(t)$ as

$$\xi(t) = \int_{-\infty}^{\infty} \left(e^{i\omega t} - 1\right) d\theta(\omega) + \xi(0) + \xi_1 t. \tag{12.159}$$

Note that if ϕ is actually continuous about $\omega = 0$, then $\xi_1 = 0$, and the last term in the above can be dropped.

12.B.8 Correlation theory for nonstationary processes

In this section we assume that the random process $\xi(t)$ is real.

In the previous section we analyzed a process with stationary increments $\xi(t)$ by looking at the difference function $\eta_\tau(t) = \xi(t) - \xi(t-\tau)$. In practice, it is also useful to work with the so-called structure function, defined as

$$D(\tau) = \left\langle [\xi(t+\tau) - \xi(t)]^2 \right\rangle \tag{12.160}$$

This is because many theorems involving correlations, which will be functions of t and τ, can be expressed in terms of structure functions, which are only functions of τ. This relation is clearest for a stationary process. Expand $D(\tau)$ as

$$\begin{aligned} D(\tau) &= \left\langle [\xi(t+\tau) - \xi(t)]^2 \right\rangle \\ &= \left\langle \xi^2(t+\tau) \right\rangle + \left\langle \xi^2(t) \right\rangle - 2 \left\langle \xi(t)\xi(t+\tau) \right\rangle \\ &= 2\left[B(0) - B(\tau)\right]. \end{aligned} \tag{12.161}$$

In the case where $B(\infty) = 0$ (most practical cases), we have $D(\infty) = 2B(0)$. The above relation lets us re-express the Wiener-Khinchin theorem in terms of the structure function. Recall that the Wiener-Khinchin theorem states that correlation and spectral density are Fourier pairs:

$$B(\tau) = \int_{-\infty}^{\infty} e^{i\omega\tau} f(\omega) d\omega = \int_{-\infty}^{\infty} \cos(\omega\tau) f(\omega) d\omega \tag{12.162}$$

where the last equality follows if $\xi(t)$ is real.[15] Combine the above two equations to get

$$\begin{aligned} D(\tau) &= 2\left[\int_{-\infty}^{\infty} f(\omega)d\omega - \int_{-\infty}^{\infty} \cos(\omega\tau) f(\omega)d\omega\right] \\ &= 2\int_{-\infty}^{\infty} \left[1-\cos(\omega t)\right] f(\omega)d\omega. \end{aligned} \tag{12.163}$$

Thus, we can also use the structure function to derive spectral density, if the process is stationary. It turns out that 12.163 also holds in the case of nonstationary processes with stationary increments. From the previous subsection we showed that such a process has the spectral representation

$$\xi(t) = \int_{-\infty}^{\infty} \left(1-e^{i\omega t}\right) d\theta(\omega) + \xi_0 + \xi_1 t \tag{12.164}$$

where ξ_0 and ξ_1 are random variables. We now use the expression to show that 12.163 holds even in this nonstationary case.

$$\begin{aligned} D(\tau) &= D(t_1-t_2) = \left\langle \left[\xi(t_1)-\xi(t_2)\right]^2\right\rangle \\ &= \left\langle \left[\xi(t_1)-\xi(t_2)\right]\left[\xi^*(t_1)-\xi^*(t_2)\right]\right\rangle \\ &= \left\langle \left[\int_{-\infty}^{\infty} (e^{i\omega_1 t_2}-e^{i\omega_1 t_1})d\theta(\omega_1) + \xi_1(t_1-t_2)\right]\left[\int_{-\infty}^{\infty} (e^{-i\omega_2 t_2}-e^{-i\omega_2 t_1})d\theta^*(\omega_2) + \xi_1^*(t_1-t_2)\right]\right\rangle \\ &= \int_{-\infty}^{\infty} \left[1-e^{i\omega\tau}-e^{-i\omega\tau}+1\right] \left\langle d\theta(\omega_1)d\theta^*(\omega_2)\right\rangle + \left\langle |\xi_1|^2\right\rangle \tau^2 \\ &= 2\int_{-\infty}^{\infty} \left[1-\cos(\omega\tau)\right] f(\omega)d\omega + \left\langle |\xi_1|^2\right\rangle \tau^2. \end{aligned} \tag{12.165}$$

Note that the last ξ_1 term can usually be neglected.[16]

---

[15] If $\xi(t)$ is real, then $B(\tau)$ is real, meaning that the spectral density $f(\omega)$ is even (property of Fourier transforms). Then the integral can be "folded" about 0, i.e. $B(\tau) = \int_0^{\infty}(e^{-i\omega t}+e^{i\omega t})f(\omega)d\omega$. Note that integrand is $2\cos(\omega\tau)$.

[16] This is just what Tatarski does. I suppose he just assumes that there's no jump discontinuity of $\theta(\omega)$ around 0? Which seems to be physically reasonable...

As a result, we have shown that even for a random process with nonstationary increments, the spectral density of the process can be computed using

$$D(\tau) = 2\int_{-\infty}^{\infty} \left[1 - \cos(\omega\tau)\right] f(\omega) d\omega. \tag{12.166}$$

Note that with $f(\omega)$ for a nonstationary process as derived from the above, we immediately see that spectral densities for stationary and nonstationary processes obey different constraints. Looking at equation 12.162, derived for a stationary process, we find that spectra for such processes must have $f(0) < \infty$. However, looking at 12.166, we see that for a process with stationary increments, $f(\omega)$ can have a singularity at 0 as long as it is of the form $\omega^{-\alpha}$, $\alpha < 3$.[17]. This should make intuitive sense, since nonstationarity is "structure" at the longest physical scales, and thus appears in frequency space at the shortest scales.

12.B.9 Stochastic fields

The framework established by the previous sections can be extended to random functions of multiple variables, termed random fields. In the case of fields over 3D space, we are essentially swapping a single time argument t for a vector argument $\mathbf{r}$. Like stochastic processes, stochastic fields have a mean $\langle \xi(\mathbf{r}) \rangle$ and correlation

$$B(\mathbf{r}_1, \mathbf{r}_2) = \langle \left[\xi(\mathbf{r}_1) - \langle \xi(\mathbf{r}_1) \rangle\right] \left[\xi(\mathbf{r}_2) - \langle \xi(\mathbf{r}_2) \rangle\right] \rangle . \tag{12.167}$$

As before, we are assumed to be working in the context of correlation theory, so that the we only need to consider the mean and correlation of a field when working with it.

---

[17] Get this by noting $1 - \cos(\omega t) \propto \omega^2$ as $\omega \to 0$, considering the convergence of $\int_{-\infty}^{\infty} x^{-p} dx$.

Homogeneity and isotropy: The concept of stationarity extends to the idea of homogeneity for a random field in the following way. For a field to be homogeneous, we require

$$\begin{aligned} \langle \xi(\mathbf{r}) \rangle &= \langle \xi(\mathbf{r}+\mathbf{r}_0) \rangle \\ B(\mathbf{r}_1, \mathbf{r}_2) &= B(\mathbf{r}_1+\mathbf{r}_0, \mathbf{r}_2+\mathbf{r}_0) \end{aligned} \tag{12.168}$$

or equivalently,

$$\begin{aligned} \langle \xi(\mathbf{r}) \rangle &= \text{constant} \\ B(\mathbf{r}_1, \mathbf{r}_2) &= B(\mathbf{r}_1 - \mathbf{r}_2). \end{aligned} \tag{12.169}$$

Furthermore, a homogeneous field is termed isotropic if we can write the correlation as

$$B(\mathbf{r}_1 - \mathbf{r}_2) = B(r) \tag{12.170}$$

where $r = |\mathbf{r}_1 - \mathbf{r}_2|$. For a homogeneous and isotropic field, if we single out a direction and consider values of the field along this direction, we obtain a random process of a single variable, to which all of our previous theorems apply.

The spectral representations of random processes of a single variable can also be expanded to three dimensions. The Fourier-RS integral in 1D naturally generalizes to

$$\xi(\mathbf{r}) = \iiint_{-\infty}^{\infty} e^{i\mathbf{k}\cdot\mathbf{r}} d\phi(\mathbf{k}) \tag{12.171}$$

for three dimensions. Analogous to before, the complex amplitudes $d\phi(\mathbf{k})$ must obey[18] the following equation for any two $\mathbf{k}_1, \mathbf{k}_2$

$$\langle d\phi(\mathbf{k}_1) d\phi^* \mathbf{k}_2 \rangle = \delta(\mathbf{k_1} - \mathbf{k}_2)\Phi(\mathbf{k_1}) d\mathbf{k}_1 \tag{12.172}$$

---

[18] Under conditions analogous to (66), which occur in most practical cases since correlation for physical random processes usually goes to 0 quickly for large separations

so that we can write for a zero-mean field

$$B(\mathbf{r}_1 - \mathbf{r}_2) = \iiint_{-\infty}^{\infty} e^{i\mathbf{k}\cdot(\mathbf{r}_1-\mathbf{r}_2)}\Phi(\mathbf{k})d\mathbf{k} \tag{12.173}$$

which is the Wiener-Khinchin theorem in 3D. For a real process with a real correlation, the above can also be expressed as

$$B(\mathbf{r}) = \iiint_{-\infty}^{\infty} \cos\left(\mathbf{k}\cdot\mathbf{r}\right)\Phi(\mathbf{k})d\mathbf{k}. \tag{12.174}$$

We can isolate the spectral density $\Phi(\mathbf{k})$ with the inverse Fourier transform

$$\Phi(\mathbf{k}) = \frac{1}{(2\pi)^3}\iiint_{-\infty}^{\infty} \cos\left(\mathbf{k}\cdot\mathbf{r}\right)B(\mathbf{r})d\mathbf{r}. \tag{12.175}$$

So as before, the correlation and spectral density are still Fourier pairs in 3D. If we additionally impose isotropy, we can use the above to derive

$$\begin{aligned}
\Phi(\mathbf{k}) &= \frac{1}{(2\pi)^3}\int_0^{2\pi}\int_0^{\pi}\int_0^{\infty} \cos\left(kr\cos\theta\right)B(r)r^2\sin\theta dr d\theta d\phi \\
&= \frac{1}{(2\pi)^3}\times 2\pi \times \int_0^{\infty}\frac{2}{kr}\sin(kr)B(r)r^2 dr \\
&= \frac{1}{2\pi^2 k}\int_0^{\infty} r\sin(kr)B(r)dr
\end{aligned} \tag{12.176}$$

where we have made use of the substitution $u = kr\cos\theta$ to compute the integral over θ. Evidently, the spectral density only depends on $k = |\mathbf{k}|$. We can use that fact to simplify 12.174 for an isotropic field as follows:

$$\begin{aligned}
B(\mathbf{r}) &= \iiint_{-\infty}^{\infty} \cos\left(\mathbf{k}\cdot\mathbf{r}\right)\Phi(\mathbf{k})d\mathbf{k} \\
&= \int_0^{2\pi}\int_0^{\pi}\int_0^{\infty} \cos\left(\mathbf{k}\cdot\mathbf{r}\right)\Phi(k)k^2\sin(\theta)dk d\theta d\phi \\
&= \frac{4\pi}{r}\int_0^{\infty} k\sin(kr)\Phi(k)dk.
\end{aligned} \tag{12.177}$$

Note that for an isotropic field, the full 3D spectral density contains the same information as the 1D spectral density along any direction. To derive the relation between the 3D and 1D spectral densities, denote the 1D density $\Phi_{1\mathrm{D}}(k)$. Then the 1D correlation is

$$B(r) = \int_{-\infty}^{\infty} \cos(kr)\Phi_{1\mathrm{D}}(k)dk. \tag{12.178}$$

Inverting this Fourier transform gives

$$\Phi_{1\mathrm{D}}(k) = \frac{1}{\pi}\int_{0}^{\infty} \cos(kr)B(r)dr. \tag{12.179}$$

Note that differentiating the above gives

$$\frac{d}{dk}\Phi_{1\mathrm{D}}(k) = -\frac{1}{\pi}\int_{0}^{\infty} r\sin(kr)B(r)dr \tag{12.180}$$

which, identifying with 12.176 leads to the relation

$$\Phi(k) = -\frac{1}{2\pi k}\frac{d}{dk}\Phi_{1\mathrm{D}}(k). \tag{12.181}$$

Locally homogeneous/isotropic fields: Just as we expanded the theory of stationary processes to the concept of nonstationary random processes with stationary increments, we can do the same to expand the theory of homogeneous fields to non-homogeneous fields with stationary increments. Firstly we define the structure function of a field as

$$D(\mathbf{r}_1, \mathbf{r}_2) = \left\langle [\xi(\mathbf{r}_1) - \xi(\mathbf{r}_2)]^2 \right\rangle. \tag{12.182}$$

We say that a field is "locally homogeneous" if the distribution functions of the difference $\xi(\mathbf{r}_1) - \xi(\mathbf{r}_2)$ remain unchanged if we translate $\mathbf{r}_1$ and $\mathbf{r}_2$ by some arbitrary vector $\mathbf{r}_0$. This implies the mean and structure function[19] must only depend on the difference $\mathbf{r}_1 - \mathbf{r}_2$.

---

[19] Not correlation! That would mean full homogeneity!

Furthermore, a field $\xi(\mathbf{r})$ is isotropic if the distribution functions for the field $\xi(\mathbf{r}_1)-\xi(\mathbf{r}_2)$ are invariant under rotation and reflection of the difference vector $\mathbf{r}_1-\mathbf{r}_2$. This implies that the structure function of such a field must only depend on the distance $r=|\mathbf{r}_1-\mathbf{r}_2|$.

Locally homogeneous fields also have spectral representations, analogous to the spectral representations of random processes with stationary increments. For instance, we can write for field $\xi(\mathbf{r})$

$$\xi(\mathbf{r})=\xi(0)+\iiint_{-\infty}^{\infty}(1-e^{i\mathbf{k}\cdot\mathbf{r}})d\phi(\mathbf{k}). \tag{12.183}$$

As done previously we can then use the above spectral representation to derive a relation between the structure function and the spectral density of the process. Doing so (following along the 1D derivation), one (should) obtain

$$D(\mathbf{r})=2\iiint_{-\infty}^{\infty}\left[1-\cos(\mathbf{k}\cdot\mathbf{r})\right]\Phi(\mathbf{k})d\mathbf{k}. \tag{12.184}$$

As we did in the previous subsection, we can simplify the above relation (essentially a Wiener-Khinchin theorem for locally homogeneous random fields) if we additionally impose isotropy. Under isotropy, $D(\mathbf{r})=D(r)$ and similarly for $\Phi(\mathbf{k})$. Thus, integrating over the angular components, we can write

$$\begin{aligned} D(r)&=2\int_0^{2\pi}\int_0^{\pi}\int_0^{\infty}\left[1-\cos(kr\cos\theta)\right]\Phi(k)k^2\sin\theta dkd\theta d\phi \\ &=8\pi\int_0^{\infty}\left[1-\frac{\sin(kr)}{kr}\right]\Phi(k)k^2dk. \end{aligned} \tag{12.185}$$

Note that the above integral converges even if $\Phi(k)$ has a singularity at 0, as long as the singularity is of the type $k^{-\alpha}$, $\alpha<5$. This implies that a non-homogeneous field can have infinite energy at the largest physical scales.

In the many situations (such as the study of atmospheric turbulence) it is also useful to consider slices of 3D fields along lines or planes. First we consider the random field along lines. Previously, we derived the relation between the structure function and the spectral density for a process that was a function of time. Replace t with r to get

$$D(r) = 2\int_{-\infty}^{\infty}[1-\cos(kr)]\Phi_{1\mathrm{D}}(k)dk \tag{12.186}$$

We can show that the same relation between 1D and 3D spectral density (12.181)

$$\Phi(k) = -\frac{1}{2\pi k}\frac{d}{dk}\Phi_{1\mathrm{D}}(k).$$

derived in the previous subsection for homogeneous, isotropic fields also holds in the locally homogeneous/isotropic case, by using it to prove the equivalence of the above two expressions for $D(r)$.

$$\begin{aligned} D(r) &= 8\pi\int_0^{\infty}\left[1-\frac{\sin(kr)}{kr}\right]\Phi(k)k^2dk \\ &= 8\pi\int_0^{\infty}\left[1-\frac{\sin(kr)}{kr}\right]\left[-\frac{1}{2\pi k}\frac{d\Phi_{1\mathrm{D}}}{dk}\right]k^2dk \\ &= -4\int_0^{\infty}\left[k-\frac{\sin(kr)}{r}\right]\frac{d\Phi_{1\mathrm{D}}}{dk}dk \\ &= 4\int_0^{\infty}\frac{d}{dk}\left[k-\frac{\sin(kr)}{r}\right]\Phi_{1\mathrm{D}}(k)dk \\ &= 2\int_{-\infty}^{\infty}\left[1-\cos(kr)\right]\Phi_{1\mathrm{D}}(k)dk. \end{aligned} \tag{12.187}$$

It can also be useful to consider 2D slices of locally isotropic fields. The 2D expansion of a field $\xi(x,y,z)$ along some z plane is the RS integral

$$\xi(x,y,z) = \xi(0,0,z) + \iint_{-\infty}^{\infty}\left[1-e^{i(k_1x+k_2y)}\right]d\psi(k_1,k_2,z) \tag{12.188}$$

where the differential term $d\psi$ obeys

$$\langle d\psi(k_1,k_2,z)d\psi^*(k_1',k_2',z')\rangle = \delta(k_1-k_1')\delta(k_2-k_2')\Phi_{2\mathrm{D}}(k_1,k_2,|z-z'|)dk_1dk_2dk_1'dk_2'. \tag{12.189}$$

Φ_{2D} is the 2D spectral density. The above ultimately gives the 2D structure function

$$D(\Delta x, \Delta y, 0) = 2 \iint \left[1 - \cos(k_1 \Delta x + k_2 \Delta y)\right] \Phi_{\text{2D}}(k_1, k_2, 0) dk_1 dk_2. \tag{12.190}$$

If we impose local isotropy along the 2D plane, then Φ_{2D} depends only on $k = \sqrt{k_1^2 + k_2^2}$ and the structure function depends only on $\rho = \sqrt{\Delta x^2 + \Delta y^2}$. Using these facts, the 2D structure function can be written as

$$\begin{aligned} D(\rho) &= 2 \int_0^\infty \int_0^{2\pi} \left[1 - \cos(k\rho \cos\theta)\right] \Phi_{\text{2D}}(k, 0) k d\theta dk \\ &= 4\pi \int_0^\infty \left[1 - J_0(k\rho)\right] \Phi_{\text{2D}}(k, 0) k dk. \end{aligned} \tag{12.191}$$

The 2D spectral density can also be related to the 3D spectral density Φ. We simply state the results here. For a derivation, see [TSC61]. You will need to use the cosine angle addition formula, and assume a real field so that the spectral density Φ is even.

$$\begin{aligned} \Phi_{\text{2D}}(k_1, k_2, z) &= \int_{-\infty}^{\infty} \cos(k_3 z) \Phi(\mathbf{k}) dk_3 \\ \Phi(\mathbf{k}) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \cos(k_3 z) \Phi_{\text{2D}}(k_1, k_2, z) dz. \end{aligned} \tag{12.192}$$

Lastly, in the case where the field is fully isotropic (not just locally isotropic) and homogeneous, we can work with correlation instead of the structure function, where the correlation is given by

$$B(\rho) = 2\pi \int_0^\infty J_0(k\rho) \Phi_{\text{2D}}(k, 0) k dk. \tag{12.193}$$

12.C Tensors in fluid dynamics

The purpose of this section is to expand the idea of correlation and structure functions of scalar fields to vector fields and apply them to fluid dynamics. I will use repeated indices to imply summation, but make a slight abuse of notation by using only lower indices for all components. This is only sensible because this section uses a Cartesian coordinate system, so covariant and contravariant components are equivalent.

12.C.1 Correlation and structure function tensors

A scalar field has one correlation function. A 3D vector field has 9, corresponding to 3 × 3 pairs of vector components. To show this explicitly, let's say we are working with a 3D velocity field. Suppose we want to compute a general correlation function for this field between two points $\mathbf{r}$ and $\mathbf{r}'$. To do so, we need to convert the vector velocity at the two points into scalars; this is done by measuring the velocity at $\mathbf{r}$ along some arbitrary direction $\mathbf{a}$ and measuring the velocity at $\mathbf{r}'$ at along some arbitrary direction $\mathbf{b}$, where $\mathbf{a}$ and $\mathbf{b}$ are unit vectors. Denote the velocity field $\mathbf{v}$; $\mathbf{u} \equiv \mathbf{v}(\mathbf{r})$; $\mathbf{u}' \equiv \mathbf{v}(\mathbf{r}')$; $\boldsymbol{\rho} = \mathbf{r} - \mathbf{r}'$. We can express the velocity correlation generally as

$$\mathcal{B}_{\mathbf{a},\mathbf{b}}(\boldsymbol{\rho}) = \langle [\mathbf{u} \cdot \mathbf{a}] \, [\mathbf{u}' \cdot \mathbf{b}] \rangle \, . \tag{12.194}$$

We can break the vectors into their Cartesian components and write the above as a sum

$$\begin{aligned} \mathcal{B}_{\mathbf{a},\mathbf{b}}(\boldsymbol{\rho}) &= \langle [u_1 a_1 + u_2 a_2 + u_3 a_3] \, [u'_1 b_1 + u'_2 b_2 + u'_3 b_3] \rangle \\ &= \left\langle u_i u'_j \right\rangle a_i b_j . \end{aligned} \tag{12.195}$$

The term $B_{ij}(\boldsymbol{\rho}) \equiv \left\langle u_i u'_j \right\rangle$, which is a function of only the separation vector $\boldsymbol{\rho}$, is the correlation tensor for the velocity field. As we can see it requires 9 components to be fully defined. Given this tensor and any two directions $\mathbf{a}$ and $\mathbf{b}$, a scalar correlation can be computed.

Apparently, working with correlations for 3D vector fields is like working with 9 regular correlation functions simultaneously. This sucks. Fortunately, things become simplified if we impose isotropy. The question becomes, what is the general form of the correlation tensor under isotropic conditions? Of course, for a scalar function isotropy implies $f(\mathbf{r}) = f(r)$, but with tensors it's more complicated. The key insight is that the correlation $B_{\mathbf{a},\mathbf{b}}(\boldsymbol{\rho})$ must be independent of rotation and therefore can only depend on rotationally invariant quantities. In the case of a scalar field the correlation depends only on the separation $\boldsymbol{\rho}$, and the only rotational invariant that can be formed from this vector is the magnitude ρ.

However, when dealing with vector fields the correlation $B_{\mathbf{a},\mathbf{b}}(\boldsymbol{\rho})$ depends on three vectors, $\mathbf{a}, \mathbf{b}$, and $\boldsymbol{\rho}$. The only meaningful rotational invariants that can be constructed are the scalar products $\rho_i \rho_i$, $a_i \rho_i$, $b_i \rho_i$, and $a_i b_i$, which correspond to magnitudes and angles, and the scalar triple product $\epsilon_{ijk} a_i b_j \rho_k$, which corresponds to volume (ϵ_{ijk} is the Levi-Civita symbol). The correlation must be constructed from these terms to remain invariant. Furthermore, regardless of isotropy, the correlation must be expressible in the form given by 12.195 (i.e. the correlation must be linear in a_i and b_i, though it may be nonlinear in $\boldsymbol{\rho}$), which puts strict limitations on the form of the correlation. The following form satisfies all the above constraints and also turns out to be most general:

$$\mathcal{B}_{\mathbf{a},\mathbf{b}}(\boldsymbol{\rho}) = \alpha(\rho)\rho_i \rho_j a_i b_j + \beta(\rho)\delta_{ij} a_i b_j + \gamma(\rho)\epsilon_{ijk} a_i b_j \rho_k. \tag{12.196}$$

If we additionally impose reflection invariance, the last term can be dropped since the scalar triple product flips sign under mirror reflection. Identifying 12.196 and 12.195 we find that the correlation tensor must have the form

$$B_{ij}(\boldsymbol{\rho}) = \alpha(\rho)\rho_i \rho_j + \beta(\rho)\delta_{ij}. \tag{12.197}$$

We could stop here, but often in the literature the functions $\alpha(\rho)$ and $\beta(\rho)$ are replaced by functions with more straightforward physical interpretations. The convention is to express the tensor in terms of components parallel to and transverse to the separation vector ρ. If we fix our coordinate system so that $\boldsymbol{\rho} = (\rho, 0, 0)$, then we can define

$$\begin{aligned} B_{pp} &\equiv B_{11} = \alpha(\rho)\rho^2 + \beta(\rho) \\ B_{tt} &\equiv B_{22} = B_{33} = \beta(\rho) \end{aligned} \tag{12.198}$$

where the pp and tt subscripts stand for "parallel-parallel" and "transverse-transverse", respectively. Therefore, we can write the correlation tensor as

$$B_{ij}(\boldsymbol{\rho}) = \left[B_{pp}(\rho) - B_{tt}(\rho)\right] \frac{\rho_i \rho_j}{\rho^2} + B_{tt}(\rho)\delta_{ij}. \tag{12.199}$$

In the study of locally homogeneous and locally isotropic vector fields, much of the above applies. However, instead of a correlation tensor we have a structure function tensor

$$D_{ij}(\boldsymbol{\rho}) = \left\langle (u_i - u_i') \left(u_j - u_j'\right) \right\rangle . \tag{12.200}$$

The rest of the math is essentially the same though. We can write

$$D_{ij}(\boldsymbol{\rho}) = \left[D_{pp}(\rho) - D_{tt}(\rho)\right] n_i n_k + D_{tt}(\rho)\delta_{ij}. \tag{12.201}$$

In this case we have defined the unit vector $\mathbf{n} = \boldsymbol{\rho}/\rho$ to match Tatarski's notation [TSC61].

Finally, we might ask what additional constraints might arise if the velocity field we were dealing corresponded to a fluid flow. Fluid flows introduce additional constraints. In particular, for flows that are moving much slower than the speed of sound, the flow can be treated as incompressible, $\vec{\nabla} \cdot \mathbf{v} = 0$. This leads to the constraint

$$\frac{\partial D_{ij}}{\partial \rho_i} = 0. \tag{12.202}$$

Substituting in equation 12.201 leads to the following result:

$$D_{tt} = \frac{1}{2\rho}\frac{d}{d\rho}\left(\rho^2 D_{pp}\right). \tag{12.203}$$

So in the case of incompressibility, there really is only one structure function which we need deal with.

12.D Zernike polynomials

This section is primarily based off of [Nol76]. Since aberrated wavefronts can be complicated in shape, we seek a convenient modal basis to represent them. This basis should have several properties. In particular, the basis modes should be:

1. orthogonal over the telescope aperture.
2. expressed in polar coordinates r, θ to accommodate for circularly symmetric imaging systems.
3. separable into $R(r)\Theta(\theta)$ for mathematical convenience.
4. invariant under the transformation $\theta \rightarrow \theta + 2\pi$.
5. maintain the same functional shape under the transformation $\theta \rightarrow \theta + \alpha$ for some angle α (in other words, they should have some rotational symmetry).

The Zernike polynomials satisfy all the above in the case of the simplest aperture, the unit circle. To use the Zernike polynomials over an aperture of arbitrary radius R we use the normalized radius $p \equiv r/R$ as our radial coordinate. The azimuthal dependence of the basis functions is carried by the complex exponential $e^{im\theta}$, $m \in \mathbb{Z}$ and the radial dependence is carried by polynomials related to the Jacobi polynomials. In practice, the azimuthal dependence is broken into sine and cosine terms. [Nol76] writes the Zernike polynomials as

$$Z_m^n(p,\theta) = \sqrt{n+1}\mathcal{R}_m^n(p)\begin{cases}\sqrt{2}\cos(m\theta) & m>0\\ \sqrt{2}\sin(|m|\theta) & m<0\\ 1 & m=0\end{cases} \tag{12.204}$$

where m and n are integers, $m \leq n$, and $n-|m|$ is constrained to be even. Different sources may use different numerical factors than above. The radial polynomials $\mathcal{R}_m^n$ have the following expression:

$$\mathcal{R}_m^n(r) = \sum_{k=0}^{(n-|m|)/2} \frac{(-1)^k(n-k)!}{k!\,[(n+m)/2-k]!\,[(n-m)/2-k]!} r^{n-2k}. \tag{12.205}$$

For convenience, the indices m and n are often remapped to a single index j such that the basis functions with cosine dependence fall on even j and the functions with sine dependence fall on odd j. Table 12.D.1 gives an often used ordering of the first ten Zernike polynomials, along with their expressions and names.

m, n indexing	Noll indexing	Expression	Name
Z_0^0	Z_1	1	piston
Z_1^1	Z_2	$2r\cos(\theta)$	tip (x)
Z_{-1}^1	Z_3	$2r\sin(\theta)$	tilt (y)
Z_0^2	Z_4	$\sqrt{3}\,(2r^2-1)$	defocus
Z_{-2}^2	Z_5	$\sqrt{6}r^2\sin 2\theta$	astigmatism $(+)$
Z_2^2	Z_6	$\sqrt{6}r^2\cos 2\theta$	astigmatism $(\times)$
Z_{-1}^3	Z_7	$\sqrt{8}\,(3r^3-2r)\sin\theta$	coma (y)
Z_1^3	Z_8	$\sqrt{8}\,(3r^3-2r)\cos\theta$	coma (x)
Z_{-3}^3	Z_9	$\sqrt{8}r^2\sin 3\theta$	trefoil (y)
Z_3^3	Z_{10}	$\sqrt{8}r^2\cos 3\theta$	trefoil (x)

Table 12.D.1: The first 10 Zernike modes ordered by Noll index, along with their m, n indices, formulas, and names.

The Zernike polynomials as defined above obey the following orthogonality relation

$$\int d^2p W(p) Z_i Z_j = \delta_{ij} \tag{12.206}$$

where the unit circle aperture function $W(p)$ is defined as

$$W(p) = \begin{cases} \frac{1}{\pi} & p \leq 1 \\ 0 & p > 1. \end{cases} \tag{12.207}$$

The above basis allows us to expand arbitrary wavefronts over a circular region. Suppose we have some function $\phi(r, \theta)$ which we want to expand over a circle of radius R. We can write

$$\phi(r, \theta) = \sum_j a_j Z_j(r/R, \theta). \tag{12.208}$$

The coefficients a_j can be pulled out with Fourier's trick:

$$a_j = \frac{1}{R^2} \int d^2r W(r/R)\phi(r,\theta) Z_j(r/R,\theta). \tag{12.209}$$

Also of use are the spectral representations of the Zernike modes. Denote the normalized cyclic frequency $s \equiv Rf$. The following is modified slightly from Noll:

$$W(p) Z_n^m(p,\theta) = \int d^2s\, Q_n^m(s,\psi) e^{2\pi i \mathbf{s}\cdot \boldsymbol{p}} \tag{12.210}$$

where

$$Q_n^m(s,\psi) = \sqrt{n+1}\frac{J_{n+1}(2\pi s)}{\pi s} \begin{cases} (-1)^{(n-m)/2}(-i)^m\sqrt{2}\cos m\psi & m \neq 0 \\ (-1)^{(n-m)/2}(-i)^m\sqrt{2}\sin m\psi & m \neq 0 \\ (-1)^{n/2} & m = 0. \end{cases} \tag{12.211}$$

J_n is the nth order Bessel function, and s, ψ are the polar coordinates of normalized frequency space. As before, the m and n indices are often remapped to a single index j so that $W(p)Z_j(p,\theta)$ and $Q_j(s,\psi)$ are Fourier pairs. Sometimes it will be more convenient to express this Fourier relation in physical units. In that case,

$$W(r/R) Z_j(r/R,\theta) = R^2 \int d^2k\ Q_j(Rf,\psi) e^{2\pi i \mathbf{f}\cdot\mathbf{r}} \tag{12.212}$$

and

$$Q_j(Rf,\psi) = \frac{1}{R^2}\int d^2r W(r/R) Z_j(r/R,\theta) e^{-2\pi i \mathbf{f}\cdot\mathbf{r}}. \tag{12.213}$$

REFERENCES

[Ane12] Dragi Anevski. "Riemann-Stieltjes integrals." 2012.

[Fri66] D. L. Fried. "Optical Resolution Through a Randomly Inhomogeneous Medium for Very Long and Very Short Exposures." *Journal of the Optical Society of America (1917-1983)*, **56**(10):1372, October 1966.

[JG94] Erik M. Johansson and Donald T. Gavel. "Simulation of stellar speckle imaging." In James B. Breckinridge, editor, *Amplitude and Intensity Spatial Interferometry II*, volume 2200 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pp. 372–383, June 1994.

[Nol76] Robert J. Noll. "Zernike polynomials and atmospheric turbulence$*$." *J. Opt. Soc. Am.*, **66**(3):207–211, Mar 1976.

[OY59] A. M. Obukhov and A. M. Yaglom. "On the microstructure of atmospheric turbulence – A review of recent work in the U.S.S.R." *Quarterly Journal of the Royal Meteorological Society*, **85**(364):81–90, April 1959.

[Rod81] F. Roddier. "The effects of atmospheric turbulence in optical astronomy." *Progess in Optics*, **19**:281–376, January 1981.

[TSC61] V. I. Tatarski, R. A. Silverman, and Nicholas Chako. "Wave Propagation in a Turbulent Medium." *Physics Today*, **14**(12):46, January 1961.

[YN63] A. M. Yaglom and Gordon Newell. "An Introduction to the Theory of Stationary Random Functions." *Journal of Applied Mechanics*, **30**(3):479–479, 09 1963.

CHAPTER 13

Simulation of atmospheric turbulence

Contents

In this chapter, I review some of the techniques used to model time-evolving atmospheric turbulence, following the statistics derived in §12. As a starting point, I cover the standard Fourier-based method which can be used to simulate "frozen flow" turbulence, where the atmosphere is assumed to be composed of unchanging turbulent layers which translate across the telescope pupil at some wind speed, as well as non-frozen flow. Next I discuss limitations of the method, particularly in accurately representing low-spatial-frequency aberrations which often dominate astronomical wavefront error. Lastly, I review an alternate modal approaches which bypass this limitation.

13.1 Sampling correlated random variables

Suppose we want to draw a 0-mean random vector $\mathbf{z}$ with covariance matrix Σ:

$$\Sigma_{ij} = \langle z_i z_j \rangle. \tag{13.1}$$

To do so, we may perform a Cholesky decomposition of Σ:

$$\Sigma = AA^T. \tag{13.2}$$

To draw $\mathbf{z}$, we first draw a vector of uncorrelated Gaussian random variables with 0 mean and 1 variance, then take

$$\mathbf{z} = A\mathbf{x} \tag{13.3}$$

which has the proper covariance

$$\langle \mathbf{z}\mathbf{z}^T \rangle = \Sigma. \tag{13.4}$$

This method will be our starting point for simulating atmospheric turbulence.

13.2 Fourier method

13.2.1 Setup

First, we note that any phase field $\phi(\mathbf{r}_0 + \mathbf{r}, t_0 + t)$ which varies over position $\mathbf{r}$ and time t can be expanded about some point $\mathbf{r}_0$ and time t_0 in the mode basis of our choice, and that the expansion coefficients will change depending on our choice of $\mathbf{r}_0$ and t_0 as well as modal basis. We first choose to use the Fourier modes:

$$\psi_{\mathbf{f}}(\mathbf{r}) = e^{2\pi i \mathbf{f}\cdot\mathbf{r}}. \tag{13.5}$$

Denote the spectral representation of $\phi(\mathbf{r}_0+\mathbf{r}, t_0+t)$ in the above mode basis as $\hat{\phi}(\mathbf{f}, \mathbf{r}_0, t_0+t)$:

$$\phi(\mathbf{r}_0 + \mathbf{r}, t_0 + t) = \int \mathbf{df}\ \hat{\phi}(\mathbf{f}, \mathbf{r}_0, t_0 + t) e^{2\pi i \mathbf{f}\cdot\mathbf{r}}. \tag{13.6}$$

By the Wiener-Khinchin theorem, reviewed in §12, the autocorrelation of ϕ and its power spectral density Φ are Fourier pairs; equivalently, the autocorrelation of $\hat{\phi}$ *is* the power spectral density:

$$\left\langle \hat{\phi}(\mathbf{f}_1, \mathbf{r}_0, t_0)\hat{\phi}^*(\mathbf{f}_2, \mathbf{r}_0, t_0) \right\rangle = \delta(\mathbf{f}_1 - \mathbf{f}_2)\Phi(f_1). \tag{13.7}$$

For Kolmogorov turbulence, we would take $\Phi(f)$ from equations 12.54 or 12.56, but the method shown here is equally valid for other types of turbulence, e.g. von Kármán. If we suppose that each Fourier mode in our atmosphere-induced wavefront distortion has an associated decorrelation time $\tau(f)$, then the normalized temporal correlation of the Fourier modes also obeys

$$\frac{\left\langle \hat{\phi}(\mathbf{f}, \mathbf{r}_0, t_0)\hat{\phi}^*(\mathbf{f}, \mathbf{r}_0, t_0 + t) \right\rangle}{\left\langle \hat{\phi}(\mathbf{f}, \mathbf{r}_0, t_0)\hat{\phi}^*(\mathbf{f}, \mathbf{r}_0, t_0) \right\rangle} = e^{-t/\tau(f)}. \tag{13.8}$$

Finally, we note that the spatial covariance in Fourier space can be computed using the shift property of Fourier transforms. Specifically:

$$\left\langle \hat{\phi}(\mathbf{f}, \mathbf{r}_0, t_0)\hat{\phi}^*(\mathbf{f}, \mathbf{r}_0 + \mathbf{r}, t_0) \right\rangle = e^{2\pi i \mathbf{f}\cdot\mathbf{r}}\Phi(f). \tag{13.9}$$

13.2.2 Frozen flow

Neglecting for the moment temporal variation, we can draw a single random phase screen according to some power spectral density $\Phi(f)$ by noting that $\Phi(f)$ is essentially a diagonal covariance matrix, at least when discretized. Thus we can perform a random draw of the Fourier mode amplitudes, denoted $\hat{z}(\mathbf{f}) = \hat{z}(f_x, f_y)$ through

$$\hat{z}(f_x, f_y) = \sqrt{\Phi(f)}\, x(f_x, f_y) \tag{13.10}$$

where $x(f_x, f_y)$ is a field of uncorrelated Gaussian random variables with mean 0 and variance 1 covering the frequency space. From here, we may compute z, a realization of the phase screen, with a Fourier transform, as in 13.6. Due to the properties of discrete Fourier transforms, this phase screen will be periodic, and can be scrolled numerically to produce a frozen flow effect. However, this periodicity can be problematic for simulations of long-exposure observations.

13.2.3 Non-frozen flow

Atmospheric turbulence is not truly frozen: the phase structure of the turbulent layer will decorrelate over time, and modes at different spatial frequencies may decorrelate at different rates. From a more practical standpoint, temporal decorrelations also produce non-periodic turbulent phase screens, which is beneficial for long-exposure simulation. To model non-frozen flow, we will require not only the autocorrelation of the Fourier mode amplitudes but also the correlation between mode amplitudes for two consecutive draws of the phase screen. Suppose our phase screen translates at some velocity $\mathbf{v}$, representing the windspeed of the turbulent layer, and denote the time between two consecutive phase screen draws as Δt. Using our prior results, the latter covariance is

$$\langle \phi(\mathbf{f}, \mathbf{r}_0, t_0)\phi^*(\mathbf{f}, \mathbf{r}_0 + \mathbf{v}t, t_0 + \Delta t)\rangle = \Phi(f)e^{2\pi i\mathbf{f}\cdot\mathbf{v}\Delta t}e^{-\Delta t/\tau(f)}. \tag{13.11}$$

To simulate turbulence, we use the above to construct a single covariance matrix that correlates a random vector composed of two consecutive draws of the Fourier mode amplitudes, which we discretize and denote as the vectors $\hat{\mathbf{z}}^n$ and $\hat{\mathbf{z}}^{n+1}$, respectively; our new random vector is then

$$\begin{bmatrix} \hat{\mathbf{z}}^n \\ \hat{\mathbf{z}}^{n+1} \end{bmatrix}. \tag{13.12}$$

Define

$$\Sigma_{11} = \Sigma_{22} = \delta_{\mathbf{f}_1,\mathbf{f}_2}\Phi(f_1) \tag{13.13}$$

and

$$\Sigma_{12} = \Sigma_{21}^T = \Sigma_{21} = \delta_{\mathbf{f}_1,\mathbf{f}_2}e^{2\pi i\mathbf{f}_1\cdot\mathbf{v}\Delta t}e^{-\Delta t/\tau(f_1)}\Phi(f_1). \tag{13.14}$$

Here, $\delta_{\mathbf{f}_1,\mathbf{f}_2}$ is the Kronecker delta. Our covariance matrix is

$$\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}. \tag{13.15}$$

The above covariance is block diagonal, making a simple expression for the Cholesky decomposition possible. Let P and g be diagonal matrices such that

$$\begin{aligned} \Sigma_{11} &= \Sigma_{22} = P^2 \\ \Sigma_{21} &= \Sigma_{12} = P^2 g \\ &\Longrightarrow \\ P &= \sqrt{\Phi(f)} \\ g &= e^{2\pi i f v \Delta t} e^{-\Delta t/\tau(f)}. \end{aligned} \tag{13.16}$$

P encodes the power spectral density of the turbulence, while g encodes the translation speed and decorrelation time for each Fourier mode. We then apply a Cholesky decomposition to Σ:

$$\begin{aligned} \Sigma &= LL^T \\ L &= \begin{bmatrix} L_{11} & L_{12} \\ L_{21} & L_{22} \end{bmatrix}. \end{aligned} \tag{13.17}$$

It can be shown that

$$\begin{aligned} L_{11} &= P \\ L_{12} &= 0 \\ L_{21} &= Pg \\ L_{22} &= P\sqrt{1-g^2}. \end{aligned} \tag{13.18}$$

Then, to create a realization of $\hat{\mathbf{z}}^n$ and $\hat{\mathbf{z}}^{n+1}$ with the proper covariance, draw vectors of normal random variables $\mathbf{x}^n$ and $\mathbf{x}^{n+1}$, and correlate them using L:

$$\begin{bmatrix} \hat{\mathbf{z}}^n \\ \hat{\mathbf{z}}^{n+1} \end{bmatrix} = L \begin{bmatrix} \mathbf{x}^n \\ \mathbf{x}^{n+1} \end{bmatrix} = \begin{bmatrix} L_{11} & L_{12} \\ L_{21} & L_{22} \end{bmatrix} \begin{bmatrix} \mathbf{x}^n \\ \mathbf{x}^{n+1} \end{bmatrix}. \tag{13.19}$$

The above equation is equivalent to the the following two:

$$\begin{aligned} \hat{\mathbf{z}}^n &= L_{11}\mathbf{x}^n \\ \hat{\mathbf{z}}^{n+1} &= L_{21}\mathbf{x}^n + L_{22}\mathbf{x}^{n+1}. \end{aligned} \tag{13.20}$$

The equations can be combined to yield a single relation

$$\hat{\mathbf{z}}^{n+1} = L_{21}L_{11}^{-1}\hat{\mathbf{z}}^n + L_{22}\mathbf{x}^{n+1} \tag{13.21}$$

or

$$\hat{\mathbf{z}}^{n+1} = g\hat{\mathbf{z}}^n + P\sqrt{1-g^2}\mathbf{x}^{n+1}. \tag{13.22}$$

The above is an autoregressive process of order 1 which can be used to simulate an infinite sequence of phase screens, and is also given in [SPR15]. If $\tau(f) \to \infty$ at all f, then we simply have

$$\hat{\mathbf{z}}^{n+1} = e^{2\pi i \mathbf{f}\cdot\mathbf{v}\Delta t}\hat{\mathbf{z}}^n \tag{13.23}$$

which is the same as the "phase scrolling" method for frozen flow turbulence that we derived in §13.2.2.

13.3 Representation in Zernike basis

Fourier-based methods have a fundamental flaw in that they poorly sample the power spectral density at small spatial frequencies, due to the finite spatial extent of the simulation region.

This is problematic because wavefront errors in astronomical settings are often dominated by low-spatial-frequency modes (e.g. Kolmogorov turbulence, which peaks to a singularity at $k = 0$.) By switching to a modal basis which better represents low spatial frequencies, such as the Zernike modes, we may sidestep this issue.

We begin by expanding the phase $\phi(\mathbf{r}_0 + \mathbf{r}, t_0 + t)$ in terms of the Zernike modes Z_j. The Zernike mode coefficients a_j can be derived with Fourier's trick:

$$a_j(\mathbf{r}_0, t_0) = \frac{1}{R^2} \int \mathbf{dr}' \; W(r'/R) Z_j(\mathbf{r}'/R) \phi(\mathbf{r}_0 + \mathbf{r}', t_0). \tag{13.24}$$

Here, $W(p)$ denotes the transmission for a circular aperture, with $W(p) = 1$ for $|p| \leq 1$ and 0 otherwise, and R is the telescope aperture radius. In this representation, there are three correlations we need to deal with. For some displacement $\mathbf{r}$, there is the spatial correlation between the Zernike modes:

$$B_{jj}(\mathbf{r}, 0) = \langle a_j(\mathbf{r_0}, t_0) a_j(\mathbf{r_0} + \mathbf{r}, t_0) \rangle \,. \tag{13.25}$$

There is also a temporal correlation:

$$B_{jj}(0, t) = \langle a_j(\mathbf{r_0}, t_0) a_j(\mathbf{r_0}, t_0 + t) \rangle \,. \tag{13.26}$$

And finally, there is the cross-correlation between the Zernike modes themselves (the Noll matrix from §12.5.2):

$$B_{ij}(0, 0) = \langle a_i(\mathbf{r_0}, t_0) a_j(\mathbf{r_0}, t_0) \rangle \,. \tag{13.27}$$

All three correlations can be mixed to form a "spatio-temporal-modal" correlation

$$B_{ij}(\mathbf{r}, t) = \langle a_i(\mathbf{r_0}, t_0) a_j(\mathbf{r_0} + \mathbf{r}, t_0 + t) \rangle \,. \tag{13.28}$$

In the case of frozen flow, spatial-temporal correlation reduces to just spatial correlation.

$$B_{jj}(\mathbf{r}, t) = B_{jj}(\mathbf{r} + \mathbf{v}t, 0). \tag{13.29}$$

This means we only need the Noll matrix and the spatial correlation in order to simulate a translating Kolmogorov phase screen. The Noll matrix is known and can be calculated analytically. We will derive the spatial correlation ourselves. Substituting 13.24 into 13.25 yields

$$\begin{aligned}&\langle a_j(\mathbf{r_0}, t_0) a_j(\mathbf{r_0} + \mathbf{r}, t_0)\rangle = \\ &\frac{1}{R^4}\int \mathbf{dr}_1'\mathbf{dr}_2' W(r_1'/R) Z_j(\mathbf{r}_1'/R) W(r_2'/R) Z_j(\mathbf{r}_2'/R) \left\langle \phi(\mathbf{r}_0 + \mathbf{r}_1')\phi(\mathbf{r}_0 + \mathbf{r} + \mathbf{r}_2')\right\rangle .\end{aligned} \tag{13.30}$$

Introduce the Fourier transform of the Zernike modes, defined by

$$\begin{aligned}W(r/R) Z_j(r/R, \theta) &= R^2 \int \mathbf{df}\ Q_j(R\mathbf{f}) e^{2\pi i \mathbf{f}\cdot\mathbf{r}} \\ &= R^2 \int \mathbf{df}\ Q_j^*(R\mathbf{f}) e^{-2\pi i \mathbf{f}\cdot\mathbf{r}}.\end{aligned} \tag{13.31}$$

Note that the Q_j functions have closed-form expressions involving Bessel functions. The PSD of the turbulence, assumed to be isotropic, is

$$\Phi(f) = \left\langle \hat{\phi}\hat{\phi}^* \right\rangle . \tag{13.32}$$

Thus, we can convert 13.30 into Fourier space (making use of the shift theorem), yielding:

$$\begin{aligned}&\langle a_j(\mathbf{r_0}, t_0) a_j(\mathbf{r_0} + \mathbf{r}, t_0)\rangle \\ &= \int \mathbf{df}_1\mathbf{df}_2 Q_j^*(R\mathbf{f}_1) Q_j(R\mathbf{f}_2) e^{-2\pi i \mathbf{r}_0\cdot\mathbf{f}_1} e^{2\pi i(\mathbf{r}_0 + \mathbf{r})\cdot\mathbf{f}_2} \Phi(f_1)\delta(\mathbf{f}_1 - \mathbf{f}_2).\end{aligned} \tag{13.33}$$

Thus, the spatial correlation between the modes is

$$\langle a_j(\mathbf{r_0}, t_0) a_j(\mathbf{r_0} + \mathbf{r}, t_0) \rangle = \int \mathbf{df} |Q_j(R\mathbf{f})|^2 e^{2\pi i \mathbf{r}\cdot\mathbf{f}} \Phi(f). \tag{13.34}$$

As you might expect, the above formula says that to shift a phase screen, we need to apply a phase ramp in Fourier space and then transform to Zernike space using the Q_j functions.

The mixed correlation (spatial, temporal, and modal) is

$$B_{ij}(\mathbf{r}, 0) = \langle a_i(\mathbf{r_0}, t_0) a_j(\mathbf{r_0} + \mathbf{v}t, t_0 + t) \rangle = \int \mathbf{df} Q_i^*(R\mathbf{f}) Q_j(R\mathbf{f}) e^{2\pi i \mathbf{f}\cdot\mathbf{v}t} \Phi(f) e^{-t/\tau(f)} \tag{13.35}$$

assuming there is no temporal cross-correlation between modes. This matrix can be computed numerically using the desired power spectral density $\Phi(f)$ and used to simulate frozen-flow turbulence in the Zernike modal basis, using the same technique as in §13.2.

13.4 Representation in Karhunen-Loève basis

In the Karhunen-Loève (KL) expansion, a random process is represented by a sequence of statistically independent, simple random variables $\xi_j(t)$. These random variables are the mode coefficients of an orthogonal basis, which we use to reconstruct the random process. KL expansion is different than the Zernike expansion because the mode coefficients of the Zernike modes are not necessarily independent (and in fact are correlated via the Noll matrix for Kolmogorov turbulence.) We explicitly write the KL expansion of a stochastic field ϕ as

$$\phi(\mathbf{r}, t) = \sum_j \lambda_j \xi_j(t) \Lambda_j(\mathbf{r}) \tag{13.36}$$

where λ_j is some constant and Λ_j is the jth KL mode. Exploiting orthogonality the random mode coefficients ξ_j can be extracted via

$$\xi_i = \frac{1}{\lambda_i} \int_\Omega \phi(\mathbf{r}, t) \Lambda_i(\mathbf{r}) \mathbf{dr} \tag{13.37}$$

where Ω fixes integration to be over the telescope aperture. The KL modes are defined by enforcing statistical independence:

$$\langle \xi_i(t) \xi_j(t) \rangle = \delta_{ij} \tag{13.38}$$

We can us the above to compute the form of the KL modes. Define the spatial correlation of the phase field ϕ as B:

$$B(\mathbf{r}_1, \mathbf{r}_2) = \langle \phi(\mathbf{r}_1, t) \phi(\mathbf{r}_2, t) \rangle. \tag{13.39}$$

We can write B in terms of the KL modes:

$$\begin{aligned} B(\mathbf{r}_1, \mathbf{r}_2) &= \sum_i \sum_j \lambda_i \lambda_j \Lambda_i(\mathbf{r}_1) \Lambda_j(\mathbf{r}_2) \langle \xi_i(t) \xi_j(t) \rangle \\ &= \sum_j \lambda_j^2 \Lambda_j(\mathbf{r}_1) \Lambda_j(\mathbf{r}_2). \end{aligned} \tag{13.40}$$

Therefore, the following integral also holds:

$$\begin{aligned} \int_\Omega B(\mathbf{r}_1, \mathbf{r}_2) \Lambda_i(\mathbf{r}_1) \mathbf{dr}_1 &= \int_\Omega \sum_j \lambda_j^2 \Lambda_j(\mathbf{r}_1) \Lambda_j(\mathbf{r}_2) \Lambda_i(\mathbf{r}_1) \mathbf{dr}_1 \\ &= \lambda_i^2 \Lambda_i(\mathbf{r_2}) \end{aligned} \tag{13.41}$$

by orthonormality. The above is the “Fredholm integral equation,” which can be numerically solved to generate the KL modes.

From another perspective, we may represent the KL modes in another basis, such as the Zernike modes. In the context of Kolmogorov turbulence, moving from a Zernike to

KL representation is the same as *diagonalization* of the Noll matrix. Denote C the change-of-basis matrix that goes from the Zernike representation to the KL representation. To generate turbulence in this new basis, we can reuse the autoregressive method derived in prior sections. The covariance matrix in the KL basis is

$$\Sigma_{ij} = \begin{bmatrix} I & C^T B_{ji}(\mathbf{v}t, t)C \\ C^T B_{ij}(\mathbf{v}t, t)C & I \end{bmatrix} \tag{13.42}$$

where I is the identity matrix. Expressing the above covariance matrix using the Cholesky decomposition and reusing the same notation as before, I find

$$\begin{aligned} L_{11} &= I \\ L_{21} &= C^T BC \equiv X \\ L_{12} &= 0 \\ L_{22}L_{22}^T &= I - C^T BCC^T B^T C = I - C^T BB^T C \\ &= I - XX^T. \end{aligned} \tag{13.43}$$

To expand upon the last matrix, note that $L_{22}L_{22}^T$ is real and symmetric and thus admits an eigendecomposition of the form

$$L_{22}L_{22}^T = I - XX^T = UPU^T \tag{13.44}$$

where U is an orthogonal matrix and P is diagonal. Therefore

$$L_{22} = U\sqrt{P}. \tag{13.45}$$

As usual, the L matrices are used to generate consecutive draws of the vector of mode coefficients via the autoregressive model. From equation 13.21, the autoregressive method in the KL basis is:

$$\mathbf{a}^{n+1} = X\mathbf{a}^n + U\sqrt{P}\,\mathbf{x}^{n+1}. \tag{13.46}$$

[SPR15] Srikar Srinath, Lisa A. Poyneer, Alexander R. Rudy, and S. Mark Ammons. "Computationally efficient autoregressive method for generating phase screens with frozen flow and turbulence in optical simulations." *Optics Express*, **23**(26):33335, December 2015.

CHAPTER 14

Adaptive optics transfer functions

The purpose of this chapter is to derive the leaky integrator transfer function. The diagram of the AO closed loop is shown in Figure 14.2.1. Much of this analysis follows from [GL94], though the authors in that reference consider a simple integrator instead of a leaky one.

14.1 Open and closed-loop transfer functions

As a primer, I derive the generic form of the closed-loop transfer function, denoting input and output signals as $x(t)$ and $y(t)$, respectively; their representations in frequency space are $\hat{x}(f)$ and $\hat{y}(f)$. The layout is shown in Figure 14.0.1. If the loop were not closed the signal that would have been be added back to $\hat{x}$ would be $\hat{x}(f)g(f)h(f)$; therefore the open loop transfer function is just $g(f)h(f)$. To derive the closed loop transfer function, note the

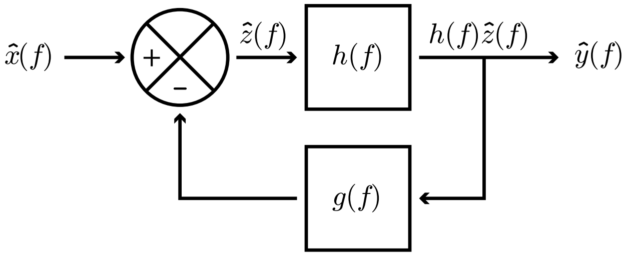

Figure 14.0.1: Generic block diagram for a closed loop system.

following:

$$\begin{aligned}
\hat{y}(f) &= h(f)\hat{z}(f) \\
\hat{z}(f) &= \hat{x}(f) - g(f)\hat{y}(f) \\
\hat{y}(f) &= h(f)\hat{x}(f) - h(f)g(f)\hat{y}(f) \\
\frac{\hat{y}(f)}{\hat{x}(f)} &= \frac{h(f)}{1+g(f)h(f)}.
\end{aligned} \tag{14.1}$$

The bottom line is the closed loop transfer function. The rejection transfer function is the above divided by the feed-forward transfer function $h(f)$, i.e. $1/(1+g(f)h(f))$. It allows us to compare the open- and closed-loop behavior of the system.

14.2 Overview of the AO loop

We will now apply the above analysis to the more complicated AO loop. For notational brevity, all quantities are assumed to be in frequency space, and I have dropped the circumflex symbols. Below, I give an overview of a typical AO loop, shown in Figure 14.2.1.

1. Aberrated wavefronts are represented as the vector $\mathbf{x}$, which is expressed in an infinite-dimensional modal basis.
2. This wavefront is corrected by a DM, becoming $\mathbf{y}$.
3. Next, the wavefront is acted upon by the wavefront sensor, for which we assume linearity; the interaction matrix in our infinite-dimensional basis is D_∞. Therefore, the signal produced by the sensor is $\mathbf{s} \equiv D_\infty \mathbf{y}$.
4. Because DMs don't have a infinite number of actuators, the matrix D_∞ is usually truncated to only include the DM's mirror modes (so now there are a finite, instead of infinite number of columns in the matrix). Denote this truncated interaction matrix as D, which has the associated pseudoinverse D^+ (the *control matrix*).

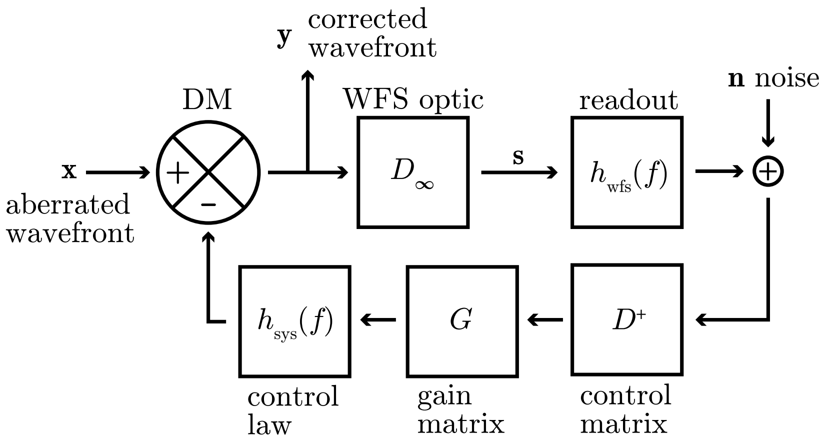

Figure 14.2.1: Overview of an AO loop, adapted from [GL94].

5. The wavefront sensor signal **s** is read out (i.e. integrated); the transfer function for this operation is denoted h_{wfs}.

6. The readout also introduces a noise component **n**.

7. The noisy signal is inverted back to a DM command using the control matrix D^+, then filtered using the diagonal gain matrix G.

8. Finally, the resulting signal is inserted into a control law with transfer function h_{sys}, which drives the DM correction.

I will now derive the transfer functions for this AO loop. Denote the input wavefront is **x**, and the corrected wavefront is **y**. From Figure 14.2.1 we note the following relation between the input and corrected wavefront:

$$\begin{aligned}\mathbf{y} &= \mathbf{x} - h_{\text{sys}}(f)GD^+h_{\text{wfs}}D_\infty\mathbf{y} + h_{\text{sys}}(f)GD^+\mathbf{n}\\ \mathbf{y} &= \mathbf{x} - h_{\text{ol}}(f)GD^+D_\infty\mathbf{y} + h_{\text{sys}}(f)GD^+\mathbf{n}\end{aligned} \tag{14.2}$$

where $h_{\text{ol}} \equiv h_{\text{wfs}}h_{\text{sys}}$ is the open loop transfer function. The transfer functions can be different for different modes. Next, we note that $D^+D_\infty\mathbf{u} = D^+D\mathbf{u}$ for $\mathbf{u}$ in the subspace of mirror modes; this relation does not hold otherwise. The corollary we will use is that $D^+ D_\infty D^+ D_\infty \mathbf{u} = D^+ D_\infty \mathbf{u}$ for *all* $\mathbf{u}$.

To deduce the transfer function, we must solve equation 14.2 for $\mathbf{y}$. To do so, we repeatedly substitute equation 14.2 into itself. I write the first two steps explicitly:

$$\begin{aligned}\mathbf{y} &= \mathbf{x} - h_{\mathrm{ol}}(f) G D^+ D_\infty \mathbf{y} + h_{\mathrm{sys}}(f) G D^+ \mathbf{n} \\ &= \mathbf{x} - h_{\mathrm{ol}} G D^+ D_\infty \mathbf{x} + h_{\mathrm{ol}}^2 G^2 D^+ D_\infty \mathbf{y} + h_{\mathrm{sys}} G D^+ \mathbf{n} - h_{\mathrm{ol}} h_{\mathrm{sys}} G^2 D^+ \mathbf{n} \\ &= \mathbf{x} - h_{\mathrm{ol}} G D^+ D_\infty \mathbf{x} + h_{\mathrm{ol}}^2 G^2 D^+ D_\infty \mathbf{x} - h_{\mathrm{ol}}^3 G^3 D^+ D_\infty \mathbf{y} \\ &+ h_{\mathrm{sys}} G D^+ \mathbf{n} - h_{\mathrm{ol}} h_{\mathrm{sys}} G^2 D^+ \mathbf{n} + h_{\mathrm{ol}}^2 f_{\mathrm{sys}} G^3 D^+ \mathbf{n}.\end{aligned} \tag{14.3}$$

If we continue, we get the series expansion

$$\begin{aligned}\mathbf{y} &= \mathbf{x} - h_{\mathrm{ol}} G \left[\sum_j (-h_{\mathrm{ol}}\, G)^j \right] D^+ D_\infty \mathbf{x} + h_{\mathrm{sys}} G \left[\sum_j (-h_{\mathrm{ol}} G)^j \right] D^+ \mathbf{n} \\ &= \mathbf{x} - \frac{h_{\mathrm{ol}} G}{1 + h_{\mathrm{ol}} G} D^+ D_\infty \mathbf{x} + \frac{h_{\mathrm{sys}} G}{1 + h_{\mathrm{ol}} G} D^+ \mathbf{n}.\end{aligned} \tag{14.4}$$

Note that the "ending" $G^j D^+ D_\infty \mathbf{y}$ term in the above summation goes to 0 as $j \to \infty$, since the gains are less than 1. Considering the ith mode, we can express equation 14.4 as

$$y_i = \frac{1}{1 + h_{\mathrm{ol}} g_i} x_i - \frac{h_{\mathrm{ol}} g_i}{1 + h_{\mathrm{ol}} g_i} \sum_{j > M} \left[D^+ D_\infty \right]_{ij} x_j + \frac{h_{\mathrm{sys}} g_i}{1 + h_{\mathrm{ol}} g_i} \sum_j D^+_{ij} n_j \tag{14.5}$$

where M is the dimension of the mirror control mode subspace (the first M columns of $D^+ D_\infty$ comprise an $M \times M$ identity matrix), and $g_i \equiv G_{ii}$ is the gain of mode i. Intuitively, equation 14.5 has three components: the first is the input signal filtered by the rejection transfer function, $h_{\mathrm{rej}} = 1/(1 + h_{\mathrm{ol}} g_i)$; the second is an aliasing term that is introduced by modes outside the mirror control space; the third corresponds to noise propagation through the loop. Next, I derive expressions for the transfer functions in the case of a leaky integrator.

The form of h_{wfs}: The action of detector readout is often approximated as convolution of the signal with a rectangular pulse starting at $t = 0$. Denote the rectangular function as $\mathrm{rect}(t/a)$ where a is the width of the pulse. This function is centered at 0. We will also work in terms of frames n, instead of time t, with $n \equiv t/T_e$. The corresponding non-dimensional frequency parameter is $s \equiv fT_e$. The relation between input and output signals is

$$y(n) = \left\{\mathrm{rect}(m - 0.5) \circledast x(m)\right\}(n) \tag{14.6}$$

where $\circledast$ denotes convolution and m is a dummy variable, so the transfer function (Fourier transforming both sides and making use of the convolution theorem) is

$$h_{\mathrm{wfs}} = \frac{\hat{y}(n)}{\hat{x}(n)} = \mathrm{sinc}(s)\, e^{-i\pi s}. \tag{14.7}$$

The sinc function is defined as

$$\mathrm{sinc}(s) = \frac{\sin(\pi s)}{\pi s}. \tag{14.8}$$

Changing $s \to fT_e$ gives the desired answer.

The form of h_{sys}: The assumed control law in [GL94] is a simple integrator; here we derive a similar result but for a leaky integrator. The particular leaky integrator we assume has the form

$$y(t) = l_i\, y(t - T_e) + x(t - \tau). \tag{14.9}$$

Here, $y(t)$ is the output (the leaky-integrated signal), l_i is the leakage for mode i, $x(t)$ is the input (the signal to be integrated), and τ is the end-to-end system latency. This latency is assumed to be dominated by hardware limitations. Note the delays: the leaky integrator uses information of the last value of y, computed one frame in the past, and also uses an input $x(t)$ that is delayed by τ due to hardware latency. For an alternate perspective, add τ to all arguments; this is equivalent to saying that it takes a time τ before the DM can react to the correction computed by the control law. Taking the Fourier transform of both sides gives

$$\begin{aligned} \hat{y}(f) &= l_i\, \hat{y}(f)\, e^{-2\pi i T_e f} + \hat{x}(f)\, e^{-2\pi i \tau f} \\ h_{\mathrm{sys}} &= \frac{\hat{y}(f)}{\hat{x}(f)} = \frac{e^{-2\pi i \tau f}}{1 - l_i\, e^{-2\pi i T_e f}}. \end{aligned} \tag{14.10}$$

h_{sys} for a simple integrator: [GL94] consider a simple integrator (i.e. convolution with a Heaviside step function). The Fourier transform of the Heaviside step function $H(n)$, in terms of frames n (not time t) is

$$\mathcal{F}\{H(n)\}(s) = \frac{1}{2\pi i s} + \frac{\delta(s)}{2}. \tag{14.11}$$

Note that the $\delta(s)$ contributes a singularity at the origin. Adding in the system latency delay of τ, the system transfer function for this control law at $f > 0$ is

$$h_{\mathrm{sys,I}} = \frac{e^{-2\pi i \tau f}}{2\pi i T_e f}. \tag{14.12}$$

The I subscript stands for "integrator." There is, however, a singularity at $f = 0$.

14.3 Rejection transfer function

Combining the previous analyses, the rejection transfer function for an AO loop with a leaky integrator control law is

$$\begin{aligned} h_{\mathrm{rej,LI}}(f) &= \frac{1}{1 + h_{\mathrm{sys}} h_{\mathrm{wfs}} g_i} \\ &= \left[1 + g_i \frac{\mathrm{sinc}(T_e f) e^{-\pi i f (T_e + 2\tau)}}{1 - l_i\, e^{-2\pi i T_e f}}\right]^{-1}. \end{aligned} \tag{14.13}$$

The LI subscript stands for "leaky integrator". The four parameters that completely determine a mode's rejection transfer function are the gain g_i, the leakage l_i, the detector integration time T_e, and the system latency τ. To compare with experiment, the modulus squared of the rejection transfer function should be equal to the ratio between closed and open loop PSDs.

The rejection transfer function in the case of a simple integrator control law is

$$h_{\text{rej,I}}(f) = \left[1 + g_i \frac{\operatorname{sinc}(T_e f) e^{-\pi i f (T_e + 2\tau)}}{2\pi i T_e f}\right]^{-1}. \tag{14.14}$$

Comparison: The behavior of a system with a simple integrator control law is markedly different than with a leaky integrator at small f:

$$\begin{aligned} \lim_{f \to 0} h_{\text{rej,I}} &= \frac{2\pi i f T_e}{g_i} \\ \lim_{f \to 0} h_{\text{rej,LI}} &= \frac{1 - l_i}{1 - l_i + g_i}. \end{aligned} \tag{14.15}$$

So the simple integrator transfer function keeps on decreasing as f approaches 0, while the leaky integrator's rejection transfer function bottoms out at a value set by the leak and gain. However, the benefit of the leaky integrator is that it explicitly removes the singularity at $f = 0$. Note that both rejection transfer functions approach 1 as $f \to \infty$.

REFERENCES

[GL94] E. Gendron and P. Lena. "Astronomical adaptive optics. I. Modal control optimization." *Astronomy and Astrophysics*, **291**(1):337–347, November 1994.

CHAPTER 15

Optimal linear phase retrieval

The purpose of this chapter is to derive the optimal linear control law when the noise in the wavefront sensor image is non-uniform. More generally, since the usual linear control law in the case of uniform image noise can be seen as a specific case of maximum likelihood estimation under the assumption of homoskedasticity (uniform noise), the goal of this document is to rederive maximum likelihood estimation (MLE) assuming heteroskedasticity (non-uniform noise).

15.1 Homoskedastic MLE and ordinary least squares

Suppose our system takes some input $\mathbf{x}$, applies a linear transformation A, and then yields some output signal $\mathbf{y}$, which is corrupted by noise $\mathbf{n}$. In the case of adaptive optics: $\mathbf{x}$, $\mathbf{y}$, A, and $\mathbf{n}$ are the phase aberration, WFS image, response matrix, and detector/photon noise, respectfully. Mathematically:

$$\begin{aligned} \mathbf{y} &= \mathbf{n} + A\mathbf{x} \\ y_i &= n_i + \sum_j A_{ij} x_j. \end{aligned} \tag{15.1}$$

I have written the above in both matrix-vector and indexed notation, for clarity. In homoskedastic MLE, we now assume that each measurement y_i is corrupted by a *uniform* Gaussian noise. In other words, each n_i is a Gaussian-distributed random variable with standard deviation σ:

$$p(n) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{n^2}{2\sigma^2}\right]. \tag{15.2}$$

We now consider a single datapoint, y_i. Combining the above two equations, the probability of obtaining this datapoint, given some input $\mathbf{x}$, is

$$p(y_i|\mathbf{x}) = p(n_i) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(y_i - \sum_j A_{ij}x_j)^2}{2\sigma^2}\right]. \tag{15.3}$$

The probability of obtaining *all* N datapoints is the joint probability:

$$p(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^{N} p(y_i|\mathbf{x}). \tag{15.4}$$

The goal of MLE is, as the name suggests, to find the value of $\mathbf{x}$ which maximizes the likelihood of the observed $\mathbf{y}$, using the above formula. For computational simplicity, we take the log of the above first, which does not change where the maximum occurs (because the log is a monotonically increasing function).

$$\begin{aligned}
\mathrm{argmax}_{\mathbf{x}}\ p(\mathbf{y}|\mathbf{x}) &= \mathrm{argmax}_{\mathbf{x}}\ \log\left[p(\mathbf{y}|\mathbf{x})\right] \\
\log\left[p(\mathbf{y}|\mathbf{x})\right] &= \log\frac{N}{\sqrt{2\pi\sigma^2}} - \frac{1}{2\sigma^2}\sum_{i=1}^{N}(y_i - \sum_j A_{ij}x_j)^2 \\
\mathrm{argmax}_{\mathbf{x}}\ p(\mathbf{y}|\mathbf{x}) &= \mathrm{argmin}_{\mathbf{x}}\ \sum_{i=1}^{N}(y_i - \sum_j A_{ij}x_j)^2.
\end{aligned} \tag{15.5}$$

Homoskedastic MLE reduces to least-squares minimization. Finally, to find the minimum, we differentiate the last line of the above and set it equal to 0:

$$\begin{aligned}
\frac{\partial}{\partial x_k}\sum_{i=1}^{N}(y_i - \sum_j A_{ij}x_j)^2 &= 0 \\
\sum_{i=1}^{N} A_{ik}(y_i - \sum_j A_{ij}x_j) &= 0 \\
A^T(\mathbf{y} - A\mathbf{x}) &= 0 \\
A^T(A\mathbf{x} - \mathbf{y}) &= 0 \\
A^TA\mathbf{x} &= A^T\mathbf{y}.
\end{aligned} \tag{15.6}$$

To solve the above, we thus take either the inverse (or more generally the pseudoinverse) of A^TA, giving

$$\mathbf{x} = (A^TA)^+A^T\mathbf{y}. \tag{15.7}$$

This particular value for $\mathbf{x}$ maximizes the likelihood of the observation $\mathbf{y}$, under the assumption of uniform noise (homoskedasticity), and is known as the "ordinary least-squares" (OLS) estimator.

Finally, we note the connection to A^+, the pseudoinverse of A (or the "control matrix" in AO). Using the properties of the pseudoinverse, we have

$$\begin{aligned}
(A^TA)^+A^T = A^+A^{T+}A^T = A^+(AA^+)^T &= A^+AA^+ = A^+ \\
\mathbf{x} &= A^+\mathbf{y}.
\end{aligned} \tag{15.8}$$

The above assumes A is real (as is the case for AO response matrices), but all the math can be extended to complex A, replacing transpose with conjugate transpose.

15.2 MLE for non-uniform but independent noise: weighted least squares

For non-uniform but independent noise, each random variable n_i associated with the noise in the measurement in y_i has an associated variance σ_i^2. Thus, we must repeat the derivation in equation 6, retaining the noise factor. The derivation is as follows. The maximum likelihood value of $\mathbf{x}$ solves

$$\begin{aligned} \frac{\partial}{\partial x_k}\sum_{i=1}^{N}\frac{(y_i-\sum_j A_{ij}x_j)^2}{2\sigma_i^2} &= 0 \\ \sum_{i=1}^{N}\frac{A_{ik}(y_i-\sum_j A_{ij}x_j)}{\sigma_i^2} &= 0. \end{aligned} \tag{15.9}$$

We define a "squared-weight" matrix W (the naming will become apparent later) as the diagonal matrix with

$$W_{ii}=\frac{1}{\sigma_i^2}. \tag{15.10}$$

Returning to matrix-vector notation, we have

$$\begin{aligned} A^TWA\mathbf{x} &= A^TW\mathbf{y} \\ \mathbf{x} &= (A^TWA)^{+}A^TW\mathbf{y}. \end{aligned} \tag{15.11}$$

This is the so-called "weighted-least-squares" (WLS) estimator. In terms of practical implementation in an AO system, we note that equation 15.11 can be recast in terms of ordinary least squares. Define

$$\begin{aligned} V &\equiv \sqrt{W} \\ A' &\equiv VA \\ \mathbf{y}' &\equiv V\mathbf{y}. \end{aligned} \tag{15.12}$$

Then:

$$\left(A'\right)^T A'\mathbf{x} = \left(A'\right)^T \mathbf{y}' \tag{15.13}$$

which has the same functional form as OLS (equation 15.8). Thus, we need to weight the response matrix and the wavefront sensor image inversely to measurement standard deviation; then we can proceed as usual. If the response matrix is measured empirically, only the wavefront sensor image needs to be weighted. Note that generally

$$(A')^+ = (VA)^+ \neq A^+V^+. \tag{15.14}$$

Aside: MLE for heteroskedastic noise with covariance Σ leads to generalized least squares.