

COSMIC SHEAR, OR ATMOSPHERE?
MODELING THE ATMOSPHERE FOR DARK ENERGY MEASUREMENTS
WITH THE VERA C. RUBIN OBSERVATORY

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF APPLIED PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Claire-Alice Hébert
September 2023

Abstract

The Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) will image tens of billions of galaxies over ten years—a 100 times increase over previous surveys—enabling the most precise measurements of dark energy properties to date. Realizing the full potential of this powerful dataset requires an unprecedented understanding of observational effects that, if uncorrected, can bias the science results. One such source is the blurring of images due to optical turbulence in the atmosphere, which dominates the point-spread function (PSF) for ground-based instruments. This atmospheric effect imprints spatially correlated noise on scales (and with amplitudes) similar to the cosmological signal we will study: the spatial correlation of positions and shapes of galaxies on the plane of the sky due to gravitational lensing of light by the dark matter in the Universe. High-fidelity simulated astronomical images are an important tool in developing and measuring the performance of image-processing algorithms that will be needed to accurately and precisely account for sources of correlated noise such as the atmospheric PSF.

This thesis presents a new simulation tool, PSF-WEATHER-STATION, that allows us to study and model the dependence of correlations in the atmospheric PSF on weather conditions at any observatory by leveraging data from weather forecasting models. We use this tool to simulate and study the correlations in the size and shape of the atmospheric PSF predicted for the Vera C. Rubin Observatory in Chile. We make quantitative predictions for two-point correlation functions (2PCF) that are used in analyses of cosmic shear. We observe a strong anisotropy in the two-dimensional 2PCF, which is expected based on observations in real images, and study the dependence of the orientation of the anisotropy on dominant wind directions at the observatory site.

We also explore the temporal behaviour of the atmospheric PSF using high-resolution speckle images of stars, recorded near Rubin Observatory, and simulate similar observations using PSF-WEATHER-STATION. We find differences in the time scales of correlation in the PSF between data and simulations, but observe a consistent dependence of PSF shape with image motion on short time scales.

Pour mon cher papa.

Acknowledgments

None of this would have happened without Pat. Since our first meeting, you have encouraged and guided my questions, slowly transforming my confusion and uncertainty into curiosity and confidence. You approach teaching, mentoring, and leadership with intention and empathy. You bring enthusiastic curiosity and sharp critical thinking to all questions and problems that come your way. Through guidance and example you have taught me to be a more precise and ethical scientist and a more earnest and optimistic human being. You will always have my deepest respect and gratitude.

I'm incredibly grateful for all past and present members of the Burchat group. We have shared offices and bonded over many frustrations: ambiguous definitions of ellipticity, installing stubborn software packages, and wrapping our heads around correlations of correlations. You all have become my friends as well as my colleagues. In particular, thank you Sowmya for your welcoming camaraderie and for introducing me to South Indian cuisine; Pierre-François for your energetic conversation and constant encouragement; Josh for patiently sharing your boundless knowledge; Sid for your sense of humor, generosity of spirit, and discerning taste in coffee; Shuang for introducing me to new snacks and your willingness to teach me spherical geometry. I interacted with many thoughtful and enthusiastic undergraduate students who have been a delight to work with: Ismael, Paxton, Jordan, Alex, Andrii, Mya, Peter, and Zory.

I want to thank the following people for their mentorship; you have engaged with and encouraged me in my work, taught me many things, and showed me, in a wide variety of ways, that I could be an academic as I am, without changing myself or my values. At KIPAC: Eric C., Adam M., Bruce M., and Aaron R.. In the DESC: Chihway C., Mike J., and Arun K.. At LANL CCS-6: Earl L. and Kary M.. At UCSB: Ram S..

The financial and programmatic support I received was a huge part of my journey in graduate school. I am indebted to the DOE Computational Science Graduate Fellowship (CSGF) program for giving me research freedom through financial support, for extending my knowledge through coursework in statistics and computing, for introducing me to a cohort of wonderful humans, and for two fantastic summers of non-thesis-related research at Los Alamos National Lab. The Stanford Enhancing Diversity in Graduate Education program showed me very early on that I wasn't

alone—Chris and Solomon, thank you. The Stanford Diversifying Academia, Recruiting Excellence fellowship program introduced me to scholarship and scholars that broadened my understanding of the academy and my place within it. I particularly thank Anika, who supported me during some especially difficult months of grieving. Stanford Women in Physics and the LGBTQIA+ crew in DESC (thank you, Sierra!) have been havens of belonging in a sea of otherness. Thank you to the Stanford Women in Physics and the Junior DESC Organization for trusting me with leadership and giving me the opportunity to grow into those roles.

My mental and physical health has been cared for over the years by incredibly empathetic and talented professionals. Thank you to Stephanie, Julie, and Cheryl for keeping me sane, to Stephanie and Abigail for fixing my hip, and to Maria and Lindsey for keeping me safe and optimistic through every step of my recovery.

I am eternally grateful to Shreya, the physics cycling crew, and the Canary Challenge for lighting the fire of my love for cycling, a sport which has since brought me countless hours of joy and adventure. Similarly, I want to thank the caretakers of the nature preserves in Santa Clara and San Mateo counties, the Stanford climbing wall, the staff at Alice’s Wednesday breakfasts, and everyone who has gone climbing, skiing, hiking, or cycling with me. I will treasure my memories of these places and adventures.

The friends I’ve made along the winding road of graduate school have made everything possible and worthwhile. I’m grateful for the years-long support of my Friendsgiving and cohort-dinner family; Abhijit, Brandon, Brett, Connie, Jon, Kate, Nicole, Ronen, Ruby, Stan, Tina, Will, Winston, Yihui, and Yuval. In particular, Yihui for sharing wisdom and good food, and Nicole for your kind heart and sense of adventure. Aditi, Ben, Cyndia, Jack, Jannicke, Kyle, Sean H., and Tony: each of you has brightened my days and filled my cup. Tori: thank you for your friendship and for always listening. The friends who bridge the divide of work-life have brought me joy and comfort both in the workplace and out: Theo, Melisa, Ji Won, and Emily. (Shout out to Verve Coffee, where I currently sit, for the lovely and innumerable writing sessions I’ve shared with Theo and Ji Won.) Thank you to Annie and Connor for countless hours on bikes in beautiful places together and for enthusiastic support of my endeavors, both academic and bike-related. I’m grateful for the years of laughter and adventure I shared with Ronen.

I am extraordinarily lucky to have close friendships that have spanned decades and distance. Gabi, you’ve been here since the beginning. Thank you, Nitin, and Jay for being my (Thanksgiving) family. Thank you Megan for your courage, for seeing me, and for always supporting me. Kelly: your optimism, empathy, spontaneity, breakfast burritos, and friendship have been invaluable.

Thank you to Sabina for exploring together, sharing good food and good conversation. Thank you for being there for me through these many challenging months.

To my home, Roo’s crew: you have been the light in the dark. Thank you to our dog/cat/goose, Roo, for being strange and wonderful and for letting us take care of you. Tina: thank you for

your loving acceptance and support. Thank you for all our conversations, big and small, and for sharing with me your political passions, your love of chili oil, and our Vinegar (Vingego) jokes. Kate has been my front-line support the entirety of grad school, in particular during shelter-in-place and during my acute grief. I wouldn't have made it to the end without you. Thank you for growing together, as QPPs and as dog co-parents, and for always knowing how to make me laugh.

Finally, I thank my family. Papa, for his curiosity, creativity, wit, and kindness. Maman, for everything. For holding us all together from the beginning, for thoughtful conversation, for chocolate cake in birthday packages. François and Susanna, for being both home and adventure; and for introducing the newest family member, petit Sebastian, whose infectious smile has cheered me up every day since his arrival. Marie, for showing me how to live a life authentic to myself. Bix, for your steadfastness and sense of humor. Anne, for daily friendship and unwavering belief in me. All together, for the laughter, trust, and good humor even in our darkest times. Un enorme merci à chacun et à tous.

Contents

Abstract	v
Acknowledgments	ix
1 Dark Energy, Cosmic Shear, and the Point Spread Function	1
1.1 Dark energy	1
1.1.1 Λ CDM	2
1.2 Weak gravitational lensing	6
1.2.1 Impact of cosmic shear	6
1.2.2 Lensed galaxy shapes	8
1.2.3 Shear correlation functions	10
1.2.4 Decomposition of shear correlation functions into E and B modes	13
1.3 Cosmic shear with Vera C. Rubin Observatory	21
1.4 Systematic effects for measurements of cosmic shear	24
1.4.1 Point spread functions	26
1.4.2 Impact of the PSF on weak lensing measurements	28
1.5 Motivation for this thesis	29
2 Imaging through the turbulent atmosphere	31
2.1 Introduction	31
2.2 The Physics of Atmospheric Turbulence	32
2.2.1 Kolmogorov turbulence	33
2.3 Turbulence distorted wavefronts	36
2.3.1 Von Kármán turbulence	41
2.3.2 Vertical structure of turbulence	42
2.3.3 Horizontal structure of turbulence	47
2.4 Outlook	49

3	Simulating realistic atmospheres: psf-weather-station	50
3.1	Motivation and context	50
3.1.1	How to simulate an atmospheric PSF	51
3.1.2	Earlier work on atmospheric PSF simulations	53
3.2	Inputs to PSF-WEATHER-STATION	54
3.2.1	Data products from General Circulation Models	54
3.2.2	Empirical model of optical turbulence profiles	60
3.2.3	Measurements at Cerro Pachón	64
3.3	Producing simulation input parameters	67
3.3.1	Optional additional correlations	70
3.3.2	Projecting direction from Earth to sky	71
3.4	Summary	73
3.4.1	Future work	73
4	Impact of wind on spatial correlations in atmospheric PSF	75
4.1	Motivation and context	75
4.2	Simulation study	75
4.3	Simulation results for isotropic correlation functions	79
4.4	Simulation results for anisotropic correlation functions	81
4.4.1	Large-scale separations of ~ 100 arcmin	84
4.4.2	Small-scale separations of ~ 1 arcmin	87
4.4.3	Summary of anisotropies in 2PCFs in simulations	90
4.5	Dependence of anisotropies on input wind parameters	92
4.6	Summary	99
5	Exploration of atmospheric wavefront data from DECam	101
5.1	Motivation	101
5.2	DECam wavefront sequences	101
5.3	Anisotropic two-point correlation function results	103
5.4	Discussion	109
6	Temporal evolution of the PSF at Cerro Pachon	111
6.1	Motivating questions	111
6.2	Fast PSF images from the Zorro speckle imager	112
6.2.1	Zorro observations and targets	114
6.2.2	Data selection and processing	114
6.3	Simulated observations	116
6.4	PSF parameter fitting	120

6.5	Time scales of PSF variation	126
6.5.1	Smoothing	126
6.5.2	Temporal correlations	130
6.6	Further studies: weather conditions and centroid motion	132
6.6.1	Dependence of PSF parameters on observing conditions	133
6.6.2	Impact of centroid motion	136
6.7	Wavelength dependence of PSF size	140
6.8	Summary and future work	143
7	Conclusions and future work	145
7.1	Summary	145
7.2	Future directions for this work	146
7.2.1	PSF modeling and interpolation	146
7.2.2	Impact of persistent on-sky correlations on co-added images	147
7.2.3	Exploring impact of weather patterns on the atmospheric PSF	147
7.2.4	Characterizing atmospheric turbulence using speckle images	147
7.2.5	More precise numerical modeling of the summit environment	148
7.3	Conclusion	148
A	Global Forecasting Models for psf-weather-station	149
B	Gallery of simulated PSF parameter correlations	151
C	Simulation tests	156
D	Exploration of adaptive optics telemetry	160
D.1	Wavefront data from adaptive optics instrumentation	160
D.2	Simulated wavefronts	161
D.3	Probing turbulence parameters	162
D.4	Summary	164

Chapter 1

Dark Energy, Cosmic Shear, and the Point Spread Function

1.1 Dark energy

Since the 1990s, cosmologists have gathered conclusive observational evidence that the Universe – i.e., space itself – is expanding at an increasing rate. Within our current cosmological model this can be described by an energy component, comprising $\approx 68\%$ of the energy density of the Universe. Aside from being sure it exists, we know almost nothing about the properties of this energy component, which is called “dark energy” to reflect our ignorance. Whereas the gravitational attraction of matter embedded in space slows the expansion of the Universe, dark energy increases the rate of expansion. Current theoretical attempts to explain dark energy often raise as many questions as they answer; for instance, the model currently supported by observation indicates that dark energy is consistent with the “cosmological constant” model, in which dark energy, described by a constant Λ , has an energy density that is constant in both time and space. However, theoretical arguments for this model suggest that if dark energy is described by a cosmological constant Λ , the density should be larger than the value we observe by a factor of approximately 10^{120} – “a truly enormous factor” (Albrecht et al., 2006). The nature of dark energy is therefore a fundamental problem in physics, which will require not only theoretical advances but accurate and precise observational measurements to constrain proposed models of dark energy and rigorously test Einstein’s Theory of General Relativity.

Despite these theoretical challenges, the current standard model of cosmology has so far produced predictions that are consistent with observations, given the current sensitivity (see next section). This standard cosmological model is termed “Lambda Cold Dark Matter” (Λ CDM), because it

assumes a cosmological-constant form of dark energy and that dark matter is “cold”, i.e., non-relativistic. Dark matter must be non-relativistic in order to form the large-scale structure we observe today. We review the theory of Λ CDM, as well as current observational evidence for the model, in Section 1.1.1. In Section 1.2, we describe weak gravitational lensing, a powerful probe of dark energy; in Section 1.3, we introduce Vera C. Rubin Observatory, an exciting instrument that is being commissioned on a mountain top in Chile and will enable the most precise measurements of dark energy properties to date; in Section 1.4, we discuss key observational challenges for weak lensing measurements. In Section 1.5, we outline how the work described in this thesis will address one of the dominant challenges.

1.1.1 Λ CDM

Our understanding of the Universe is that it started with a big bang 13.8 billion years ago¹, followed by a *very* short period ($\approx 10^{-33}$ second) of cosmic inflation – exponential expansion. The Universe has been expanding and cooling ever since; as it cooled, atoms were able to form, photons decoupled from matter, and eventually, structures (stars and then galaxies) started to form under the influence of gravity. In this section we will give an overview of relevant aspects of modern cosmology; for further reading, see Dodelson (2003); Ryden (2003); Weinberg et al. (2013).

We parameterize the expansion of space by the scale factor $a(t)$ that describes how the distance between two fixed points in space changes with time due to the expansion of space itself; $a(t)$ is typically set to 1 at $t = 0$, so that $a(t) < 1$ in the past. The rate at which $a(t)$ changes with time, $\dot{a}(t)$, describes the expansion rate of the Universe; the fractional rate of expansion is called the Hubble parameter:

$$H(t) \equiv \frac{\dot{a}(t)}{a(t)}. \quad (1.1)$$

The present-day value of $H(t)$ is called the Hubble constant H_0 ; the measured value is approximately 70 km/s/Mpc.²

It is common to use the cosmological redshift z to describe cosmological distances, rather than the scale factor, because z is an observable quantity. A photon of wavelength λ_{emit} , emitted by some distant object, will have a redshifted wavelength $\lambda_{\text{obs}} > \lambda_{\text{emit}}$ when observed, because the space through which the photon travels is expanding! The ratio of these wavelengths is given by the ratio of scale factors at the time of observation and emission. Redshift z is defined as $z = \frac{\lambda_{\text{obs}}}{\lambda_{\text{emit}}} - 1$, so it is related to the scale factor as,

$$\frac{1}{a(t_{\text{emit}})} = 1 + z, \quad (1.2)$$

where we have used that $a(t_{\text{obs}}) = 1$. Given this definition of redshift, $z = 0$ for light emitted at the

¹Most precise measurements are by NASA’s WMAP and ESA’s Planck spacecraft.

²There is some tension between the current measurements based on the cosmic microwave background radiation (the “early universe”), which are clustered around 67 km/s/Mpc, and the values measured by “late universe” probes, which cluster around 73 km/s/Mpc, each with uncertainties of about 1 km/s/Mpc.

current time and $z > 0$ in the past.

The Universe is filled with matter, radiation, and dark energy; we denote their relative present-day energy densities by Ω_m , Ω_r , and Ω_ϕ , respectively. These energy densities are, by definition, normalized by the “critical energy density” $\frac{3H_0^2}{8\pi G}$ such that for a flat (Euclidean) universe,

$$1 \equiv \Omega_m + \Omega_r + \Omega_\phi. \quad (1.3)$$

The current measured energy densities for each component are $\Omega_m = 32\%$ (27% dark matter and 5% ordinary “baryonic” matter) and $\Omega_\phi = 68\%$ (Collaboration et al., 2020); Ω_r is negligible. In other words, the energy density of the Universe is currently dominated by dark energy. However, when the Universe was younger, the energy density was first radiation dominated, and then matter dominated; as the Universe expands, these densities decrease (with radiation energy density decreasing faster due to the cosmological redshift) while the dark-energy density stays constant!

Current observations are consistent with a universe that is homogeneous and isotropic—i.e., there is no preferred location or direction—at scales larger than ≈ 100 Mpc. In General Relativity, a homogeneous and isotropic universe with the above components is governed by the Friedmann equation, which describes the redshift-dependence of the rate of change of the expansion of the Universe:

$$\frac{\dot{H}(z)}{H_0} = \Omega_m(1+z)^3 + \Omega_r(1+z)^4 + \Omega_\phi \frac{u_\phi(z)}{u_\phi(z=0)}, \quad (1.4)$$

where $u_\phi(z)$ is the (unscaled) energy density of dark energy. For a constant $u_\phi(z)$, denoted by $u_\Lambda(z)$,

$$\frac{u_\Lambda(z)}{u_\Lambda(z=0)} = (1+z)^{3(1+w)}. \quad (1.5)$$

More generally,

$$\frac{u_\phi(z)}{u_\phi(z=0)} = \exp\left(3 \int_0^z [1 + w(z')] \frac{dz'}{1+z'}\right), \quad (1.6)$$

where $w(z)$ is the equation of state parameter for dark energy: for a scalar field ϕ with pressure p , the equation of state is given by

$$w(z) \equiv \frac{p(z)}{u(z)}. \quad (1.7)$$

For dark energy described by a cosmological constant, the equation of state parameter is $w = -1$ for all time. Putting this together with Equation 1.5 and Equation 1.4, we find a dark energy component with no time dependence, as expected.

A model for $w(z)$ that includes evolution with redshift has two-parameters:

$$w(z) = w_0 + w_a(1 - a). \quad (1.8)$$

The equation of state parameter is initially $w_0 + w_a$ and, as t and a increase (z decreases), w reaches

the present value of w_0 . This simple model allows us to test deviations from a cosmological constant by searching for evidence of $w_a \neq 0$.

Equation 1.4 assumes a homogeneous and isotropic universe on large $\mathcal{O}(100 \text{ Mpc})$ scales. The existence of you, me, the Solar System, and even the Milky Way Galaxy are proof enough that our Universe does not satisfy these conditions on small scales. To adapt our theory to accommodate the presence of local inhomogeneities, we include a linear perturbation away from the average matter density $\bar{\rho}_m(t)$. These density fluctuations are described by

$$\delta_m(\mathbf{x}, t) \equiv \frac{\rho_m(\mathbf{x}, t) - \bar{\rho}_m(t)}{\bar{\rho}_m(t)}, \quad (1.9)$$

where \mathbf{x} is the co-moving coordinate.³

We can parameterize the time dependence in Equation 1.9 by a linear growth function $g(t)$, where $\delta_m(\mathbf{x}, t) \equiv \delta_m(\mathbf{x}, t_i) \frac{g(t)}{g(t_i)}$. Fluctuations in matter density will evolve gravitationally as

$$\ddot{g} + 2H(z)\dot{g} - \frac{3}{2}H_0^2(1+z)^3g = 0. \quad (1.10)$$

The growth of these density fluctuations $\delta_m(\mathbf{x})$ will be influenced by properties of dark energy and other cosmological parameters. Simple inflation models predict that the density fluctuations are described by a Gaussian random field—and observations so far support this hypothesis. In this case the statistical properties of $\delta_m(\mathbf{x})$ are fully described by the matter power spectrum $P_m(\ell)$ or its Fourier transform, the two-point correlation function $\xi(\theta)$ (more on this in Section 1.2). The power spectrum is defined by

$$\langle \tilde{\delta}_m(\ell) \tilde{\delta}_m(\ell') \rangle \equiv (2\pi)^3 P_m(\ell) \delta_D^3(\ell - \ell') \quad (1.11)$$

for Fourier mode ℓ ; δ_D^3 is the 3D Dirac delta function and $\tilde{\delta}_m(\ell)$ is the Fourier transform of the density field. Isotropy implies that $P(\ell)$ will be a function of only $k \equiv |\ell|$.

We can observe the distribution of matter in the nearby universe with multiple independent probes. Figure 1.1 shows measurements of the matter power spectrum with the cosmic microwave background by the Planck space telescope (blue, orange, and green points), with galaxy clustering and cosmic shear from the ground-based Sloan Digital Sky Survey (red points) and the Dark Energy Survey (yellow points), and with high redshift quasars from spectroscopic measurements by the Extended Baryon Oscillation Spectroscopic Survey (purple points). The curve corresponds to the best-fit prediction of the matter power spectrum based on only the Planck data – i.e., data at low values of k or large scales. Although the DES, SDSS, and eBOSS data represent very different methods and probe matter at different scales, they are consistent with the predicted curve fit to the Planck data.

³Comoving coordinates factor out the expansion of space, giving a distance that does not change in time due to the expansion of space.

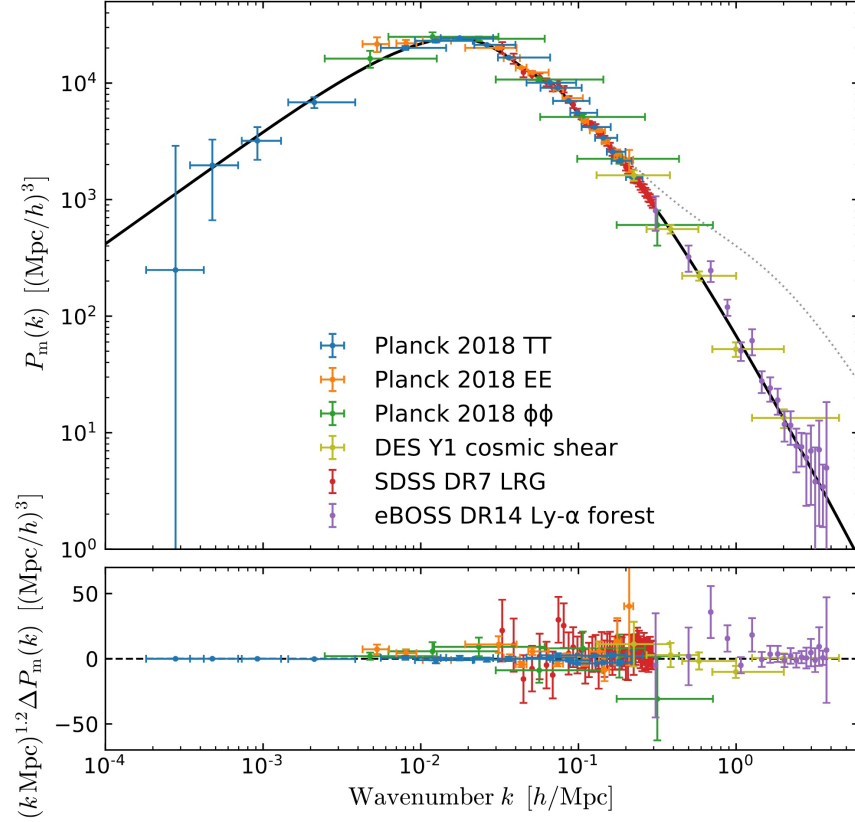


Figure 1.1: *Top*: Data points show inferences of the 3D linear matter power spectrum at $z = 0$ from various observational probes (Planck CMB data, largest scales; SDSS galaxy clustering, intermediate scales; SDSS Ly α clustering and DES cosmic shear data, smallest scales). See text for more details. The solid black curve is the theoretical expectation given the best-fit Planck 2018 Λ CDM model. The dotted line for reference shows the theoretical spectrum including non-linear effects. *Bottom*: deviation of the data from the Planck best fit Λ CDM 3D matter power spectrum. Figure from Chabanier et al. (2019).

We conclude with an important takeaway from this section: to further constrain proposed models of dark energy, we must make more precise – and accurate – measurements of w_0 and w_a . In the next section, we focus our review on a single, but powerful, cosmological probe of dark energy: weak gravitational lensing.

1.2 Weak gravitational lensing

Images of distant galaxies are distorted by weak gravitational lensing due to inhomogeneities at large scales in the mass distribution in the Universe, in the space between the galaxy and the observer. These distortions are called cosmic shear (see, e.g., Kilbinger (2015); Mandelbaum (2018); Weinberg et al. (2013) for excellent reviews). Spatial correlations in the distortions are sensitive to the properties and evolution of the matter density on large scales and the geometry of space. As the collaborations associated with current cosmological surveys (the Dark Energy Survey (DES⁴), Hyper Suprime-Cam (HSC⁵), and Kilo-Degree Survey (KiDS⁶)) complete their analyses of their entire surveys, cosmic shear is becoming one of the most precise probes of cosmology – in particular for measurements of the average matter density and the amplitude of matter fluctuations. In this section, we outline the relevant theory and review the basics of conducting measurements of cosmic shear. More extensive derivations of the relevant formalism summarized below can be found in Weinberg et al. (2013); Kilbinger (2015); Schneider et al. (2006).

1.2.1 Impact of cosmic shear

Suppose we observe a distant source, at co-moving distance χ . If the source is observed at an apparent angular position $\boldsymbol{\theta}$, what is the true angular position $\boldsymbol{\beta}$? The inhomogeneous mass distribution (large-scale structure) of the Universe sources a nonuniform gravitational potential Φ ; the gradient of this potential at each point along the path of a photon causes light to be deflected by some angle $d\boldsymbol{\alpha}$. The total deflection angle $\boldsymbol{\alpha}$ of a light ray is found by integrating $d\boldsymbol{\alpha} \propto \nabla\Phi$ along its path. Therefore,

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \boldsymbol{\alpha}, \quad (1.12)$$

– i.e., the deflection angle relates the true source angular position to the observed angular position.⁷ We can define a Jacobian matrix $A = \partial\boldsymbol{\beta}/\partial\boldsymbol{\theta}$ that, to first order,⁸ describes the mapping from $\boldsymbol{\theta}$ to

⁴<https://www.darkenergysurvey.org/>, Amon et al. (2022)

⁵<https://hsc.mtk.nao.ac.jp/ssp/>, Li et al. (2023)

⁶<http://kids.strw.leidenuniv.nl/>, Asgari et al. (2021)

⁷If there is more than one solution to Equation 1.12 we will have images at multiple locations on the sky—this occurs in strong lensing.

⁸We apply the Born approximation and then linearize the problem by assuming that the source is much smaller than the angular scales over which the matter density that produces the lensing varies.

β :

$$A_{ij} = \frac{\partial \beta_i}{\partial \theta_j} = \delta_{ij} - \frac{\partial^2}{\partial \theta_i \partial \theta_j} \psi(\theta, \chi), \quad (1.13)$$

where $\psi(\theta, \chi)$, called the *lensing potential*, is an integral along the line of sight over the gravitational potential:

$$\psi(\theta, \chi) = \frac{2}{c^2} \int_0^\chi d\chi' \frac{\chi - \chi'}{\chi \chi'} \Phi(\chi' \theta, \chi'). \quad (1.14)$$

The matrix in Equation 1.13 is symmetric, and is often parameterized in terms of scalar convergence κ and two-component shear (a “spin-two” quantity) γ :

$$A = \begin{pmatrix} 1 - \kappa - \gamma_1 & -\gamma_2 \\ -\gamma_2 & 1 - \kappa + \gamma_1 \end{pmatrix}. \quad (1.15)$$

Convergence changes the apparent size of a source, while shear is an anisotropic stretching of images (more on γ in Section 1.2.2).

We will now use the Poisson equation, which in comoving coordinates is

$$\nabla^2 \Phi = \frac{3H_0^2 \Omega_m}{2a} \delta_m, \quad (1.16)$$

to express κ in terms of matter fluctuations δ_m . From Equation 1.13 and Equation 1.15, convergence can also be written $\kappa = \nabla^2 \psi$. Together with Equation 1.14, this gives

$$\kappa(\theta) = \frac{3H_0^2 \Omega_m}{2c^2} \int_0^{\chi_{\text{lim}}} d\chi \frac{\chi q(\chi)}{a(\chi)} \delta_m(\chi \theta, \chi). \quad (1.17)$$

The function $q(\chi)$ is the *lens efficiency*, which, for a density fluctuation at some distance χ and given the observed distribution of source distances $n(\chi')$, is defined as

$$q(\chi) = \int_\chi^{\chi_{\text{lim}}} d\chi' n(\chi') \frac{\chi' - \chi}{\chi'}. \quad (1.18)$$

The integrals are taken to χ_{lim} , the co-moving distance to the furthest source in $n(\chi)$. We can think of $\kappa(\theta)$ as an effective surface mass density, as for each angle θ we have integrated δ_m along the line of sight.

The Limber equation states that a 2D projection (here, κ) of a 3D homogeneous and isotropic random field (δ_m) is also a homogeneous and isotropic random field. This means the power spectrum $P_\kappa(\ell)$ fully describes the field κ . The Fourier space analogy of Limber’s equation relates $P_\kappa(\ell)$ and $P_m(\ell)$:

$$P_\kappa(\ell) = \frac{9H_0^4 \Omega_m^2}{4c^4} \int_0^{\chi_{\text{lim}}} d\chi \frac{q^2(\chi)}{a^2(\chi)} P_m\left(\frac{\ell}{\chi}, \chi\right). \quad (1.19)$$

The power spectrum of the convergence due to gravitational lensing of sources from a mass distribution is related, through this integral, to the power spectrum of the mass distribution itself.

In practice, it is more feasible to measure the shear field $\gamma(\boldsymbol{\theta})$ from gravitational lensing than the convergence field $\kappa(\boldsymbol{\theta})$. This is primarily due to the unknown prior on the size of galaxies—how do we separate intrinsic variation in size from variation that is due to lensing? In contrast, we can assume that the distributions of galaxy shapes and orientations in the universe are random, with no preferred direction, by homogeneity and isotropy. The cosmic shear field $\gamma(\boldsymbol{\theta})$ is directly related in Fourier space to the convergence field by

$$\tilde{\gamma}(\boldsymbol{\ell}) = e^{2i\eta} \tilde{\kappa}(\boldsymbol{\ell}), \quad (1.20)$$

where η is the polar angle of vector $\boldsymbol{\ell}$. Therefore,

$$\langle \tilde{\gamma}(\boldsymbol{\ell}) \tilde{\gamma}^*(\boldsymbol{\ell}') \rangle = \langle \tilde{\kappa}(\boldsymbol{\ell}) \tilde{\kappa}^*(\boldsymbol{\ell}') \rangle = (2\pi)^2 P_\kappa(\ell) \delta_D^2(\boldsymbol{\ell} - \boldsymbol{\ell}'). \quad (1.21)$$

The power spectra of the shear and convergence fields are identical, so we can use Equation 1.19 to relate the shear field to the matter density field. Thus, we can constrain cosmological parameters by making measurements of cosmic shear from the weak gravitational lensing by the large scale structure of the Universe! It is common in observational cosmology to measure the real-space two-point correlation function of shear rather than the shear power spectrum $P_\gamma(\ell)$; in what follows we will focus on this real-space measure.

1.2.2 Lensed galaxy shapes

We have just reviewed how the shear field of background sources can theoretically be used as a cosmological probe. Here, we discuss how we can, in practice, carry out measurements of the shear field from observations of galaxy shapes.

As we saw in Equation 1.15, the impact of lensing on the image of an objects (referred to as the “source” object – e.g., the source galaxy) is described by convergence κ and shear γ . For example, lensing will turn the image of a circular source object with radius R into an ellipse with major and minor axes given by

$$a = \frac{R}{1 - \kappa - \gamma}, \quad b = \frac{R}{1 - \kappa + \gamma}. \quad (1.22)$$

Since galaxies are not intrinsically circular, it is necessary to quantify the effect of lensing on a source with non-zero ellipticity. Given the form of the Jacobian matrix A (Equation 1.15), we introduce “reduced shear” $g = \gamma/(1 - \kappa)$ to parameterize the lensing shear of galaxies, and write it as a complex number

$$g = g_1 + ig_2 = |g|e^{-2i\phi}. \quad (1.23)$$

Then A can be written as

$$A = (1 - \kappa) \begin{pmatrix} 1 - g_1 & -g_2 \\ -g_2 & 1 + g_1 \end{pmatrix}. \quad (1.24)$$

The complex reduced shear g is a spin-two (“spinor”) quantity. The magnitude g quantifies the impact of magnification and shear on the ellipticity of the object. The angle ϕ , defined over the range $(0, 180)^\circ$, corresponds to a polar angle 2ϕ for the spinor in the complex g plane. Therefore, a desired consequence of the factor of 2 in the exponential is that the ellipticity of the applied shear rotated by 180° in real space is identical to that of the unrotated shear.

To estimate shape (and size) from images of galaxies, which are in general not circular and not perfect ellipses, it is common to use the weighted second moments Q_{ij} of the galaxy intensity profile $I(\boldsymbol{\theta})$,

$$Q_{ij} = \frac{\int d^2\theta I(\boldsymbol{\theta}) W(\boldsymbol{\theta}) \theta_x \theta_y}{\int d^2\theta I(\boldsymbol{\theta}) W(\boldsymbol{\theta})}, \quad (1.25)$$

where $\boldsymbol{\theta}$ corresponds to angular position on the image and $W(\boldsymbol{\theta})$ is a weighting function. The square root of the variance σ^2 is a measure of galaxy size, where

$$\sigma^2 = Q_{xx} + Q_{yy}. \quad (1.26)$$

For galaxy shape, we define a complex ellipticity parameter e as

$$e = e_1 + ie_2 = \frac{Q_{xx} - Q_{yy} + 2iQ_{xy}}{Q_{xx} + Q_{yy}}. \quad (1.27)$$

The magnitude of e is given by $|e| = \frac{1-q^2}{1+q^2}$, where q is the ratio of the length of the minor to major axes of the ellipse described by the second-moment matrix. The complex ellipticity e and reduced shear g follow the same mathematics. We use the notation e for the imaged shape of a source, \hat{e} for the *estimate* of the shape of the source, and g for the true (reduced⁹) lensing shear that we wish to measure.

Figure 1.2 illustrates the interpretation of (e_1, e_2) ; each ellipse is drawn with the shape corresponding to the values of (e_1, e_2) at that point. If the original source galaxy were circular, then the measured values of (e_1, e_2) would correspond to the shear (g_1, g_2) . If $e_2 = 0$, then the orientation of the major axis of the ellipse is parallel to or perpendicular to the θ_x direction for positive or negative values of e_1 , respectively. If $e_1 = 0$, then the major axis of the ellipse lies parallel to or perpendicular to an axis rotated $+45^\circ$ with respect to the θ_x axis, for positive or negative values of e_2 , respectively. Non-zero e_1 and e_2 describe orientations in between the above special cases.

⁹We will drop the modifier “reduced” going forward and depend on the notation g to indicate reduced shear rather than shear γ .

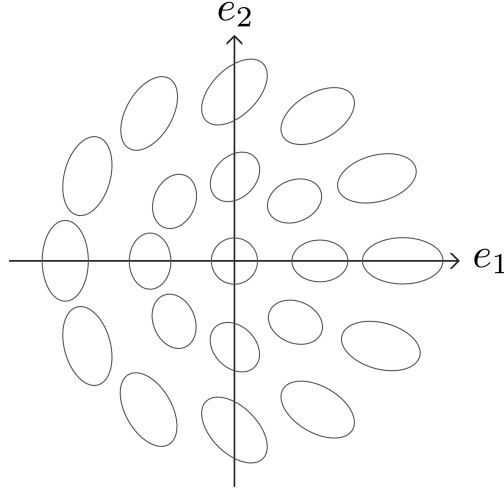


Figure 1.2: Interpretation of the real and imaginary components e_1 and e_2 of complex ellipticity e . Each ellipse is drawn with the shape corresponding to the values of (e_1, e_2) at that point. The axes could also be interpreted as tangential and cross components, e_t and e_\times , for a pair of galaxies as described in Section 1.2.3.

1.2.3 Shear correlation functions

In defining e and g above, we did not specify a coordinate system in which to measure ϕ . A natural coordinate system to use (in anticipation of calculating two-point correlation functions) is that defined by the separation vector between each pair of galaxies, rather than an arbitrary fixed coordinate system. Therefore, for each pair of galaxies, we define “tangential” and “cross” components of shear (g_t, g_\times) with respect to their separation vector. If angle α is the polar angle of the separation vector in a fixed coordinate system, then

$$g_t = -\operatorname{Re}(ge^{-i2\alpha}), \quad g_\times = -\operatorname{Im}(ge^{-i2\alpha}). \quad (1.28)$$

With this definition, g_t , for example, is positive or negative if the major axis of the shear ellipse lies parallel or perpendicular to the separation vector – similar to e_1 in Figure 1.2. Likewise, g_\times plays a role analogous to e_2 in Figure 1.2, but with respect to the separation vector. This coordinate system changes for every pair of galaxies but has a very important advantage in applications to gravitational lensing – namely, lensing from a spherically symmetric gravitational potential generates purely tangential shear distortions.

To build physical intuition for tangential and cross shear averaged over pairs of galaxies when the coordinate system for each pair is rotated (which will be necessary when measuring two-point correlation functions), we show examples of different orientations of ellipses with respect to the

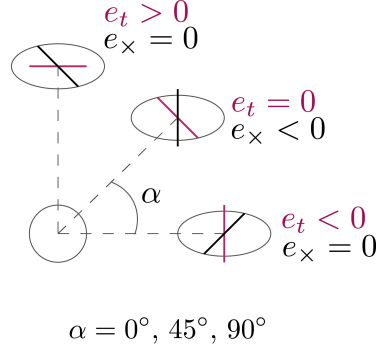


Figure 1.3: Sketch of tangential (red) and cross (black) components of ellipticity rotating as the polar angle α of the separation vector between two “galaxies” takes on values $\alpha = 0^\circ, 45^\circ, 90^\circ$. (The circular source in the lower left corner also has a rotating coordinate system of e_t and e_x , but it is not shown since the ellipticity magnitude is 0, leading to 0 tangential and cross components for the circular galaxy for all α .)

separation vector in Figure 1.3 for three pairs of “galaxies”, one of which is circular. In the figure, as α changes from 0 to 45° to 90° , the only nonzero component of e for this particular ellipse is first $e_t < 0$, then $e_x < 0$, and then $e_t > 0$.

Note that the examples in Figure 1.3 are just a handful of the many possible relative orientations; ellipses of different shapes and orientations will have different e_t and e_x values for a given separation vector.

As noted in Section 1.2, our goal is to measure real-space two-point correlation functions of shear. To this end, we are interested in measuring the average correlation between shapes and orientations (ellipticities) of pairs of galaxies, as a function of separation distance between the galaxies. We have four quantities that describe the ellipticities of a pair of sources (a, b) : $e_t^a, e_x^a, e_t^b, e_x^b$. We can combine these four numbers into the following three two-point correlation functions:

$$\langle e_t^a e_t^b \rangle(\theta), \langle e_t^a e_x^b \rangle(\theta), \langle e_x^a e_t^b \rangle(\theta), \langle e_x^a e_x^b \rangle(\theta), \quad (1.29)$$

where the angle brackets $\langle \dots \rangle$ denote the average over many pairs of galaxies (a, b) at different separations θ on the sky. In a parity-symmetric universe, for a given value of shear component e_t , the strength of e_x and $-e_x$ is the same, so the two-point correlation functions involving the cross correlation e_x vanish.

In Figure 1.4, we show example source pairs for special cases in which only e_t is nonzero or only e_x is nonzero. Gravitational lensing potentials can generate only the shear distortions shown in the upper left panel: $e_t e_t > 0$ and $e_x e_x = 0$; therefore, we expect only positive tangential correlations of galaxy pairs. However, many systematic effects can introduce nonzero $e_x e_x = 0$, as we will see later in this thesis. Therefore, these correlations provide useful signatures of potential systematic

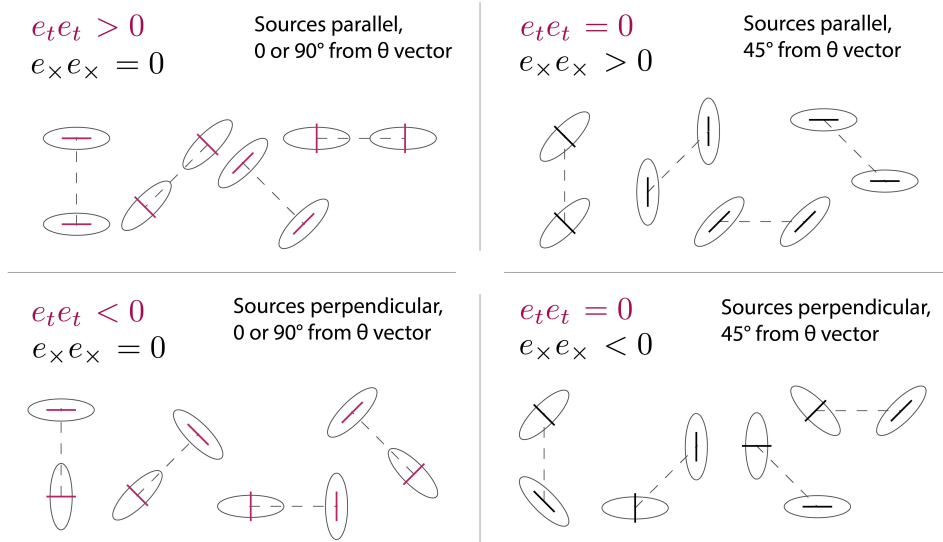


Figure 1.4: Examples of elliptical source pairs corresponding to four sets of two-point correlation values; only a few special cases are illustrated for each set. In the panels on the left, the major/minor axes of the ellipses are orthogonal to the separation vector between sources, so only the e_t components are nonzero. In the panels on the right, the major/minor axes of the ellipses are at 45° with respect to the separation vector, so only e_x is nonzero. The two-point correlations are 0 or positive in the upper panels and 0 or negative in the lower panels.

effects.

The shear two-point correlation function used in weak lensing has two components, denoted $\xi_{\pm}(\theta)$ (from now on, we will drop the superscripts a, b from the ordered pair of ellipticity components):

$$\xi_{\pm}(\theta) = \langle e_t e_t \pm e_x e_x \rangle(\theta) \quad (1.30)$$

These correlation functions can be related to Equation 1.21 to find

$$\xi_+(\theta) = \int_0^\infty \frac{d\ell \ell}{2\pi} J_0(\ell\theta) P_\kappa(\ell), \quad (1.31)$$

$$\xi_-(\theta) = \int_0^\infty \frac{d\ell \ell}{2\pi} J_4(\ell\theta) P_\kappa(\ell), \quad (1.32)$$

where $J_n(x)$ is the n -th order Bessel function of the first kind.

In a complete cosmological analysis, we generate a set of correlation functions ξ_{\pm}^{ij} that correlate galaxy shapes across redshift bins i and j . Binning the galaxy sample by redshift gives us a view of the evolution of the Universe: we see galaxies as they were in the younger Universe (when they emitted the light we now receive); therefore, high redshift galaxies have passed through a statistically different large-scale structure than galaxies at low redshift. The two-point correlation functions ξ_{\pm}^{ij}

(Troxel et al., 2018),

$$\xi_{\pm}^{ij} = \frac{\sum_{ab} w_a w_b (\hat{e}_t^{a,i} \hat{e}_t^{b,j} \pm \hat{e}_{\times}^{a,i} \hat{e}_{\times}^{b,j})}{\sum_{ab} w_a w_b}, \quad (1.33)$$

where w is a per-galaxy weights (usually given by measurement uncertainties), and the sums are over galaxy pairs (a, b) separated (within some range of angular separations $\Delta\theta$) by an angle θ on the sky. As an example, we show in Figure 1.5 the measurements of $\xi_{\pm}(\theta)$ across pairs of four redshift bins, from the DES Year 3 results¹⁰ (Amon et al., 2022). Note that the correlation functions plotted on the vertical axes are multiplied by the galaxy-pair separation θ . The redshift bins are labeled 1, 2, 3, and 4, in order of increasing redshift. The highest correlation values observed are in the panels labeled 4, 4 – i.e., the shapes of galaxies at the highest redshifts show the most evidence of correlation. Given that their light has passed through the highest volume of intervening matter, and that shapes are likely to be most highly correlated with objects whose light passes through approximately the same intervening matter, this is as expected.

1.2.4 Decomposition of shear correlation functions into E and B modes

As discussed in Section 1.2.1, two-point shape correlation functions $\xi_{\pm}(\theta)$ constitute observables of interest because of their dependence on cosmological parameters of interest (see Equation 1.19, Equation 1.31, and Equation 1.32). A related formulation of the same information is the decomposition of $\xi_{\pm}(\theta)$ into E - and B -modes of shear. Analogous to the vector fields in electricity and magnetism, we can think of the E -mode as the “gradient” and the B -mode as the “curl” of shear; see Figure 1.6 for an illustration of spatial patterns of shear distortions corresponding to pure E modes (left) and B modes (right).

From Figure 1.6, we see that for the pure E -mode shear patterns (images on the left of the figure), $\langle e_t e_t \rangle > 0$ and $\langle e_{\times} e_{\times} \rangle = 0$ for source pairs on opposite sides of the center of each image; in contrast, for pure B -mode shear patterns (images on the right of the figure), $\langle e_{\times} e_{\times} \rangle > 0$ and $\langle e_t e_t \rangle = 0$ for corresponding pairs. However, recall from Equation 1.30 that $\xi_{\pm} = \langle e_t e_t \rangle \pm \langle e_{\times} e_{\times} \rangle$. Therefore, to first approximation, one can think of ξ_+ and ξ_- as being proportional to the sum and difference of E and B modes, respectively.

To lowest order, gravitational lensing is a gradient field and thus produces only E -mode power Mandelbaum (2018). Cosmological B modes due to higher order lensing effects are expected to be much smaller than cosmological E modes – by a factor of at least 500, even at the smallest scales (Hilbert et al., 2009; Krause & Hirata, 2010). Nevertheless, there are efforts to make significant measurements of cosmological B modes; evidence of B modes could indicate a non-Gaussian primordial random field, which would have powerful implications for models of early-universe inflation (Meerburg et al., 2019). The large suppression of B modes renders them undetectable even with full-sky optical galaxy surveys; instead, the most sensitive searches for B modes are conducted with

¹⁰DES is the optical survey with the widest area cosmic-shear results as of this writing.

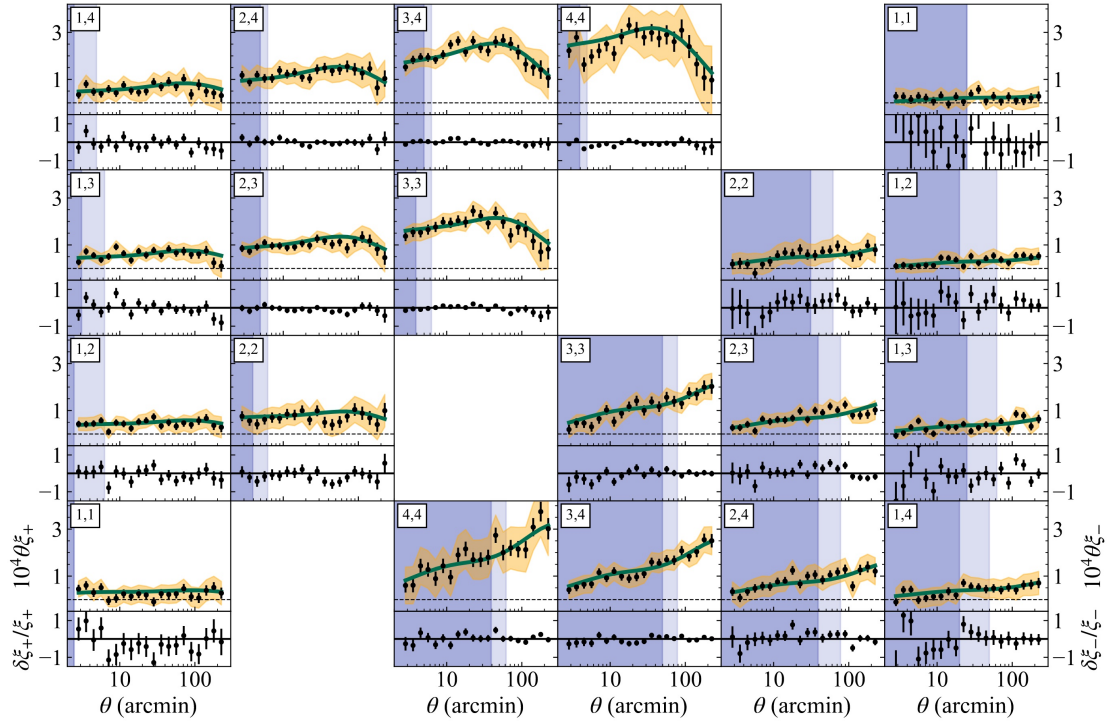


Figure 1.5: DES Y3 measurements (Amon et al., 2022) of cosmic shear two-point correlation functions: $\xi_+(\theta)$ (plots above and to the left of the diagonal) and $\xi_-(\theta)$ (below and to the right of the diagonal), scaled by the separation angle θ . The data points are in black, the yellow shaded region shows the DES Y1 uncertainties for comparison, and the green curves show the best fit Λ CDM theoretical predictions. The shaded blue regions indicate the ranges of θ that are excluded from the analysis (light and dark shading correspond to a theory-agnostic and a Λ CDM-optimized analysis). The lower region in each panel show the fractional difference between the measurements and the best-fit values.

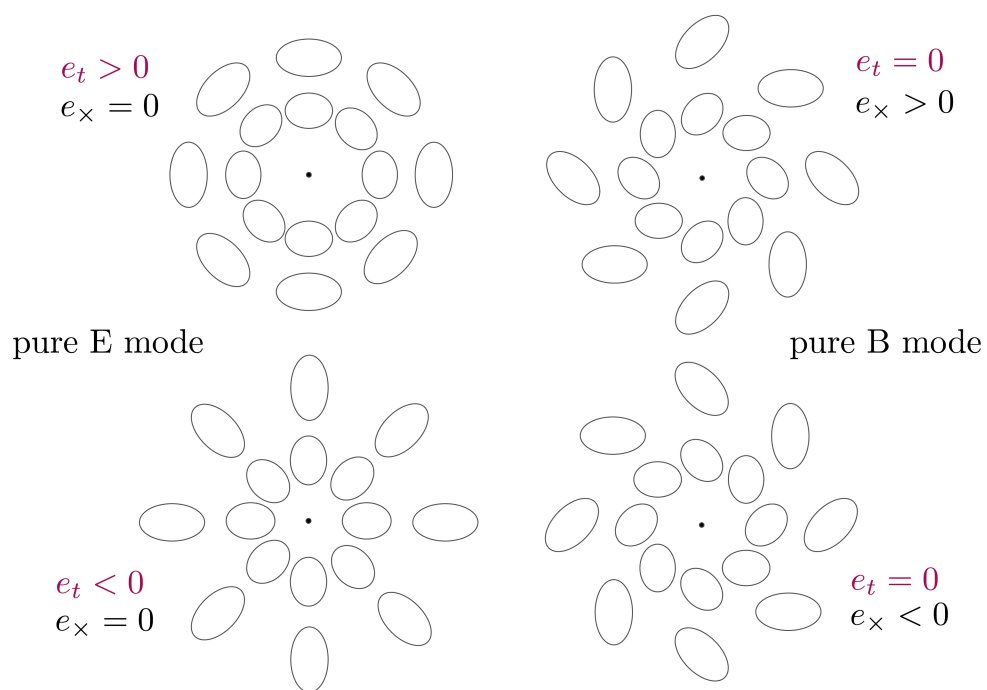


Figure 1.6: Illustration of pure E -mode (left) and B -mode (right) shear distortions. The upper two images correspond to 0 or positive tangential and cross shear components with respect to a line through the center of the group of shapes (indicated by the points) and the lower two images to 0 or negative shear components.

measurements of the cosmic microwave background radiation. Any cosmological signal we are able to detect with an optical survey is therefore expected to be purely E -mode, and any nonzero B -mode signal is very likely evidence of a non-cosmological effect – namely a systematic effect which we must understand and take into account.

As we saw in Equation 1.31 and Equation 1.32, the shear two-point correlation functions can be expressed in terms of the power spectrum $P_\kappa(\ell)$. Similarly, they can be expressed in terms of the E and B power spectra $P_E(\ell)$ and $P_B(\ell)$:

$$\begin{aligned}\xi_+(\theta) &= \int_0^\infty \frac{d\ell}{2\pi} J_0(\ell\theta) [P_E(\ell) + P_B(\ell)], \\ \xi_-(\theta) &= \int_0^\infty \frac{d\ell}{2\pi} J_4(\ell\theta) [P_E(\ell) - P_B(\ell)].\end{aligned}\tag{1.34}$$

Going from these expressions to the two-point correlation functions $\xi_\pm^{E/B}(\theta)$ is a nontrivial journey with Bessel function recurrence relations and integrations; the good news is that these measures of E - and B -mode correlations have been analytically derived in Crittenden et al. (2002) and Schneider et al. (2002). These initial results required knowledge of $\xi_-(\theta)$ and $\xi_+(\theta)$ to arbitrarily large and small scales, respectively. Therefore, for use in the context of cosmic shear surveys, which cover a finite range of angular separations, the data would need to be extrapolated beyond the measured range.

Rather than extrapolation, cosmic shear surveys generally relied on approximation techniques to calculate E and B mode correlations for systematic tests. In particular, the “Complete Orthogonal Sets of E and B mode-separating Integrals” (COSEBI) method, first presented in Schneider et al. (2010), has been widely used – notably in the analysis of cosmic shear survey results from both KiDS (Giblin et al., 2021) and DES (Gatti et al., 2021). COSEBIs are two-point statistics that separate well-defined E and B modes within a finite angular range. They are measured from $\xi_\pm(\theta)$ using oscillating filter functions; the number n of modes needed to fully capture cosmological information depends on the range of angular scales available, desired accuracy, angular bin size, etc.

In contrast to both of the above methods, more recent results published in Schneider et al. (2022) provide explicit equations for pure-mode correlation functions $\xi_\pm^{E/B}(\theta)$. The descriptor “pure-mode” refers to pure E and pure B modes that have been separated from “ambiguous” modes—denoted by $\xi_\pm^{\text{amb}}(\theta)$ —that cannot be uniquely ascribed to either E or B modes due to the finite interval of angular separations available. In equation form,

$$\xi_\pm(\theta) = \xi_\pm^E(\theta) \pm \xi_\pm^B(\theta) + \xi_\pm^{\text{amb}}(\theta).\tag{1.35}$$

The pure-mode results in Schneider et al. (2022) allow for a finite range of θ ($\theta_{\min} \leq \theta \leq \theta_{\max}$) over which ξ_\pm is measured, and can be used without extrapolating data to large/small scales. Another drawback of the COSEBI method has been the difficulty in tracing a detection in E or B mode to

a specific angular scale (Asgari et al., 2019); the pure-mode results, calculated directly from data, overcome this.

The full Schneider et al. (2022) results are summarized in the remainder of this section. We begin with the decomposition of $\xi_+(\theta)$:

$$\xi_+^E(\theta) = \frac{1}{2} \left[\xi_+(\theta) + \xi_-(\theta) + g(\theta) \right] - \frac{1}{2} \left[S_+(\theta) + S_-(\theta) \right], \quad (1.36)$$

$$\xi_+^B(\theta) = \frac{1}{2} \left[\xi_+(\theta) - \xi_-(\theta) - g(\theta) \right] - \frac{1}{2} \left[S_+(\theta) - S_-(\theta) \right], \quad (1.37)$$

where

$$g(\theta) = 4 \int_{\theta}^{\theta_{\max}} \frac{d\theta'}{\theta'} \xi_-(\theta') (1 - 3 \frac{\theta^2}{\theta'^2}), \quad (1.38)$$

$$S_+(\theta) = \int_{\theta_{\min}}^{\theta_{\max}} \frac{\theta' d\theta'}{\theta_{\text{avg}}^2} \xi_+(\theta') H_+(\theta, \theta'), \quad (1.39)$$

$$S_-(\theta) = \int_{\theta_{\min}}^{\theta_{\max}} \frac{d\theta'}{\theta'} \xi_-(\theta') H_-(\theta, \theta'), \quad (1.40)$$

with $\theta_{\text{avg}} = \frac{1}{2}(\theta_{\min} + \theta_{\max})$ and $H_{\pm}(\theta, \theta')$ defined below. The $S_{\pm}(\theta)$ functions are the terms that account for our limited knowledge of $\xi_{\pm}(\theta)$; as $\theta_{\min} \rightarrow 0$ and $\theta_{\max} \rightarrow \infty$, they will vanish. In the more general (and realistic) case, we must include these functions in the calculation of E and B modes from measured two-point correlation functions $\xi_{\pm}(\theta)$. H_{\pm} are polynomial functions in θ' , θ , and $R = \frac{\theta_{\max} - \theta_{\min}}{\theta_{\max} + \theta_{\min}}$:

$$H_+(\theta, \theta') = \frac{1}{8R^3} \left(4R^2 + 3 \left[\left(\frac{\theta}{\theta_{\text{avg}}} \right)^2 - 1 - R^2 \right] \left[\left(\frac{\theta'}{\theta} \right)^2 - 1 - R^2 \right] \right), \quad (1.41)$$

$$\begin{aligned} H_-(\theta, \theta') = & \frac{(1-R)^2}{8R^3} \left(3(1-R)^2 \left[(1+R)^4 - (1+4R+R^2) \left(\frac{\theta}{\theta_{\text{avg}}} \right)^2 \right] \left(\frac{\theta'}{\theta} \right)^{-2} \right. \\ & \left. + 3(1+R)^2 \left(\frac{\theta}{\theta_{\text{avg}}} \right)^2 - [3+6R+14R^2+6R^3+3R^4] \right). \end{aligned} \quad (1.42)$$

The pure-mode decomposition of $\xi_-(\theta)$ follows a similar pattern:

$$\xi_-^E(\theta) = \frac{1}{2} \left[\xi_+(\theta) + \xi_-(\theta) + f(\theta) \right] - \frac{1}{2} \left[V_+(\theta) + V_-(\theta) \right], \quad (1.43)$$

$$\xi_-^B(\theta) = \frac{1}{2} \left[\xi_+(\theta) - \xi_-(\theta) + f(\theta) \right] - \frac{1}{2} \left[V_+(\theta) - V_-(\theta) \right], \quad (1.44)$$

where

$$f(\theta) = 4 \int_{\theta_{\min}}^{\theta} \frac{\theta' d\theta'}{\theta^2} \xi_+(\theta') (1 - 3 \frac{\theta'^2}{\theta^2}), \quad (1.45)$$

and the $V_{\pm}(\theta)$ functions play the same role as $S_{\pm}(\theta)$ and are constructed similarly¹¹:

$$V_+(\theta) = \int_{\theta_{\min}}^{\theta_{\max}} \frac{\theta' d\theta'}{\theta_{\text{avg}}^2} \xi_+(\theta') K_+(\theta, \theta'), \quad (1.46)$$

$$V_-(\theta) = \int_{\theta_{\min}}^{\theta_{\max}} \frac{d\theta'}{\theta'} \xi_-(\theta') K_-(\theta, \theta'), \quad (1.47)$$

with the K_{\pm} polynomials defined in terms of H_{\pm} as

$$K_+(\theta, \theta') = \left(\frac{\theta_{\text{avg}}}{\theta} \right)^2 H_-(\theta', \theta), \quad (1.48)$$

$$K_-(\theta, \theta') = (1 - R^2)^2 \frac{\theta_{\text{avg}}^2}{\theta^2} H_+ \left([1 - R^2] \frac{\theta_{\text{avg}}^2}{\theta}, [1 - R^2] \frac{\theta_{\text{avg}}^2}{\theta'} \right). \quad (1.49)$$

By adding Equation 1.36 and Equation 1.37, Schneider et al. (2022) find that $\xi_+^{\text{amb}}(\theta) = S_+(\theta)$, and by subtracting Equation 1.43 from Equation 1.44, that $\xi_-^{\text{amb}}(\theta) = V_-(\theta)$. The function $\xi_+^{\text{amb}}(\theta)$ captures the knowledge, or lack thereof, at large θ , and $\xi_-^{\text{amb}}(\theta)$ at small θ . We are able to separate $\xi_{\pm}(\theta)$ into purely E and B modes in the limit of $\theta_{\max} \rightarrow \infty$ and $\theta_{\min} \rightarrow 0$, respectively, in which case $\xi_{\pm}^{\text{amb}} \rightarrow 0$. In these limits and in the absence of B modes, $\xi_{\pm}(\theta) = \xi_{\pm}^E(\theta)$.

Another limiting case to consider is that of a constant shear field of magnitude g_0 and phase ϕ_0 . In this case, $\xi_+(\theta) = \langle gg^* \rangle = g_0^2$ and $\xi_-(\theta) = \langle gge^{-4i\alpha} \rangle = g_0^2 e^{-4i\phi_0} \langle e^{-4i\alpha} \rangle = 0$. One can show, with some patience for algebra, that

$$\int_{\theta_{\min}}^{\theta_{\max}} \frac{d\theta' \theta'}{\theta_{\text{avg}}} H_+(\theta, \theta') = 1; \quad (1.50)$$

therefore for a constant $\xi_+(\theta)$, the function $S_+(\theta)$ has that same constant value (see Equation 1.39). In our example case, this means the ambiguous mode $S_+(\theta) = g_0^2$, and by Equation 1.36 and Equation 1.37, $\xi_+^E(\theta) = \xi_+^B(\theta) = 0$.

Results from Schneider et al. (2022) for the pure-mode decomposition applied to KiDS-1000 shear data are shown in Figure 1.7 for E modes (i.e., the cosmological results) and Figure 1.8 for B modes (i.e., measures of potential systematic effects). The upper right and lower left sets of panels in each figure show $\theta \xi_+^{E/B}(\theta)$ and $\theta \xi_-^{E/B}(\theta)$, respectively, with θ measured in arcminutes. Each panel corresponds to two-point statistics across a pair of redshift bins, numbered 1 through 5, with numbers increasing with higher redshift. The pure-mode $\xi_{\pm}^{E/B}(\theta)$ calculated from the data

¹¹Compared to the equations presented in Schneider et al. (2022), here I absorbed a factor of $(\theta'/\theta_{\text{avg}})^2$ into K_- for symmetry with the definitions of S_{\pm} .

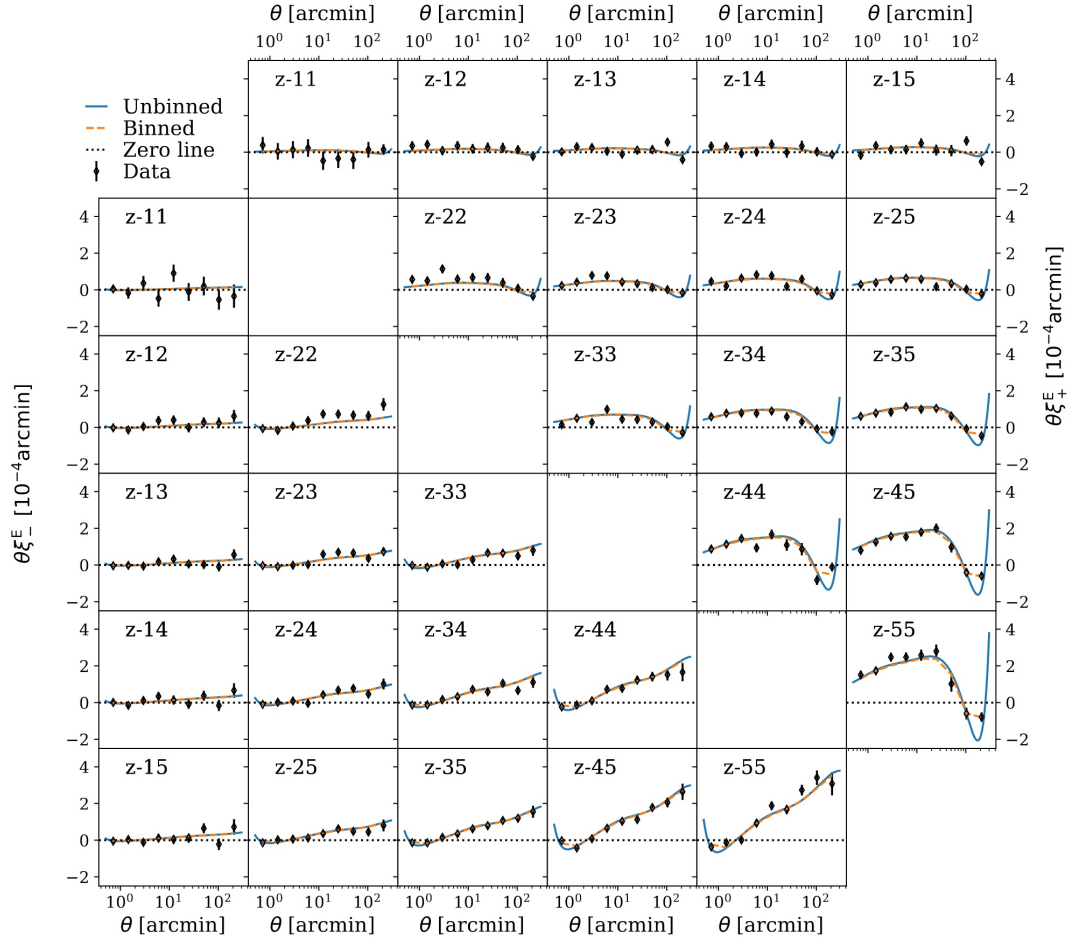
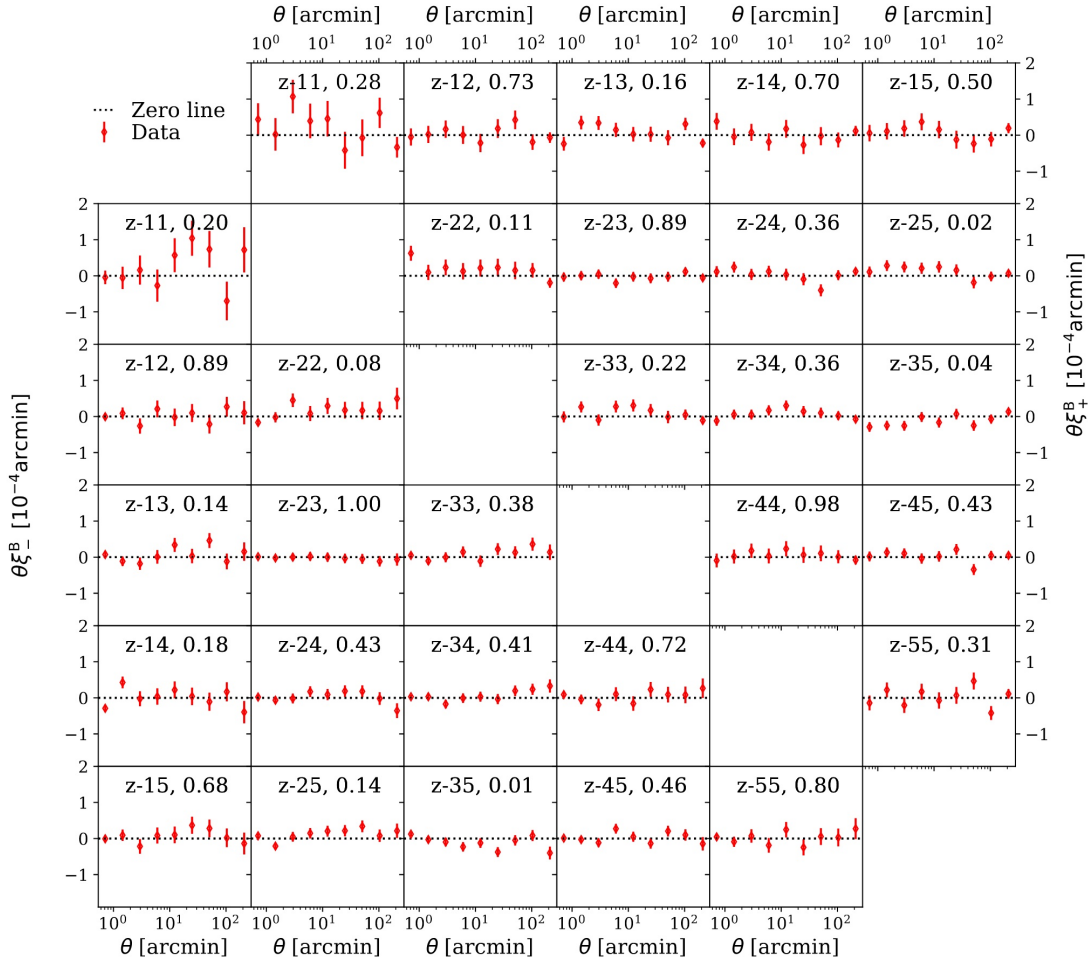


Figure 1.7: Measured pure-mode $\xi_{\pm}^E(\theta)$ for cosmic shear, from KiDS-1000 (Schneider et al., 2022).

Figure 1.8: Measured pure-mode $\xi_{\pm}^B(\theta)$ for cosmic shear, from KiDS-1000 (Schneider et al., 2022).

are shown as black points; the error bars correspond to the diagonal elements of the respective covariance matrices. The blue and orange curves in Figure 1.7 correspond to $\xi_{\pm}^E(\theta)$ calculated from the KiDS-1000 best-fit cosmology, which was performed using COSEBIs; see Asgari et al. (2021). Although the data points on the plot were not themselves directly used in the fit, the data and best fit curves are in good agreement. In Figure 1.8, the B -mode shear correlation functions are shown to be consistent with zero across angular scales and redshifts. The authors note, however, that the significance of B -modes depends nontrivially on the choice of data binning, and recommend the use of multiple binning schemes during checks for systematic effects.

1.3 Cosmic shear with Vera C. Rubin Observatory

The cosmic shear results described in Section 1.2.3 and Section 1.2.4 were produced by survey collaborations that have completed their data-taking and will release at most one more round of science results. Of the next generation of surveys and missions targeting cosmological questions, one of the most exciting for cosmic shear measurements is the ground-based Vera C. Rubin Observatory, which will be complemented by two space-based missions. During its first ten years, Rubin Observatory will conduct the Legacy Survey of Space and Time (LSST), imaging approximately 20 billion galaxies in six filters and reaching depths down to r -band magnitude ~ 26.9 . The LSST is expected to reach and then surpass by an order of magnitude the precision of existing surveys for cosmic shear.

The design of the observatory is optimized for a wide and fast survey of the sky; the 3.5-deg field of view and short visit times (30 s) allow for the whole sky to be imaged every three nights. Figure 1.9(a) highlights the compact optical system of Rubin Observatory, which allows for fast slewing between observations and a large light-collecting area with an 8.4-m outer-diameter annular primary mirror, which has an effective circular aperture diameter of 6.5 m. The three-mirror optical design of the telescope enables high-quality imaging over a wide field of view; with each additional mirror surface, optical aberrations at a higher order (e.g., astigmatism) can be corrected, thus improving imaging capabilities near the edges of the field of view. The projection of the LSST camera focal plane on an image of the sky shown in Figure 1.9(b) illustrates the wide field of view of the Rubin LSST Camera – 3.5 deg in diameter, approximately 50 times the area of the full moon. Each small square represents a 4k pixel by 4k pixel CCD, with a pixel size of 10 micron corresponding to a pixel scale of 0.20 arcseconds.

Rubin Observatory is in final stages of construction on Cerro Pachón in Chile; see Figure 1.10 for an aerial view of the dome and neighboring telescopes in the background. Commissioning efforts have started, the LSST camera is in the final stages of testing, and first-light for the survey is currently projected in late 2024.

Eight LSST science collaborations are already preparing for the arrival of the rich dataset from Rubin Observatory. Each collaboration focuses on a separate science interest: active galactic nuclei,

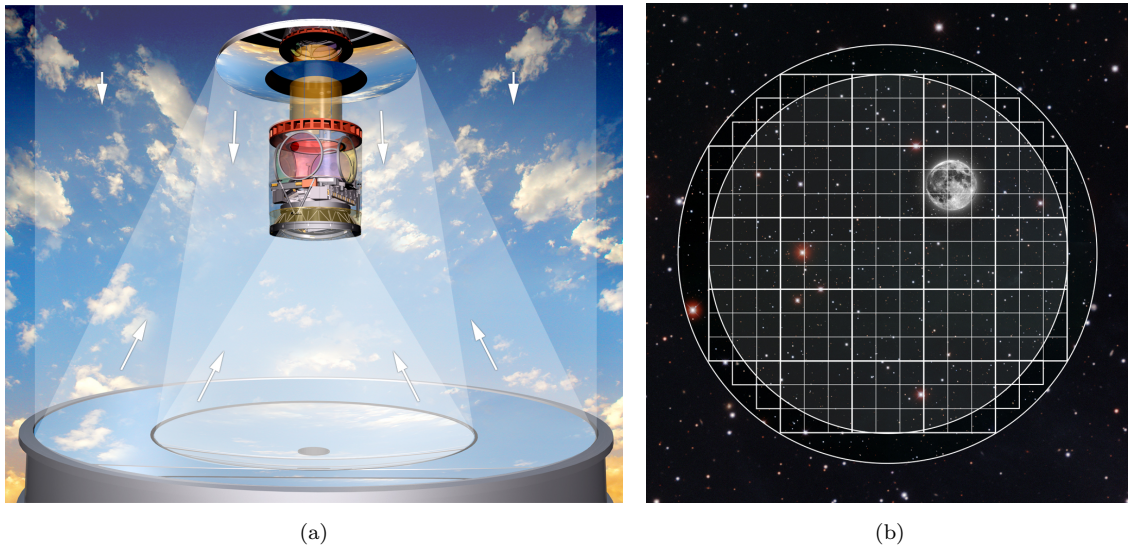


Figure 1.9: Schematics of (a) Rubin Observatory's three-mirror optical system and (b) the LSST Camera's focal plane projected on the sky. In (a), the direction of the arrows follow the path of photons through the optical system into the camera. The primary mirror (large annular region at bottom of the image) is 8.4 m in diameter. In (b), the focal plane is shown to scale on an image of the sky; the full moon subtends 0.52 deg within the 3.5-deg field of view (inner circle). Image credit: Rubin Observatory.



Figure 1.10: The Vera C. Rubin Observatory (foreground), as of April 2023. The Rubin Auxiliary Telescope (used for commissioning and measurements of atmospheric transmission) is up the road to the right of the Rubin dome. On the background ridge are the SOAR (left) and Gemini South (right) telescopes. Image credit: Rubin Observatory.

dark energy, galaxies, informatics and statistics, the solar system, stars in our Milky Way Galaxy and the local volume, strong gravitational lensing, and transient objects and variable stars.

The Dark Energy Science Collaboration (DESC) has been working since 2012 to prepare to make precise and accurate cosmological analyses with the Rubin LSST data. This work includes forecasting the precision we can achieve on measurements of cosmological parameters given the unprecedented volume of data: 100 times more than previous surveys.¹² In Figure 1.11, the Fisher forecasts for 68% confidence-level contours are shown for the ten-year survey, for the dark energy equation-of-state parameters w_o and w_a (see Equation 1.8). Contours for individual cosmological probes are plotted in different colors; galaxy clusters (light blue), strong lensing (SL; yellow), supernovae (SN; green), and the joint analysis of cosmic shear, galaxy-galaxy lensing, and galaxy clustering (3×2 pt; dark blue). The red contour shows the precision of existing surveys (“Stage III”) and the black contour is the combination of all above probes and pre-existing data. The power of the LSST data set is evident by the significant reduction in statistical uncertainty from the red to black contours.

To fully leverage the statistical power of the Rubin LSST dataset, it is crucial that any potential systematic effects – both systematic biases and uncertainties on those biases – are understood and constrained with similar accuracy and precision. This is a major focus of the ongoing work within the DESC, including the work described in this thesis.

1.4 Systematic effects for measurements of cosmic shear

Observational cosmologists have the thorny challenge of turning images of galaxies into constraints on fundamental parameters of a cosmological model describing our Universe. This task is often thought about in two steps: (1) turning images into “catalogs” of objects (in this case, galaxies) and their measured properties (e.g., position on the sky, size, shape, distance/redshift), and (2) using the catalogs to infer cosmological information via, in the case of cosmic shear, galaxy shape correlation functions (see Section 1.2.2). In the images-to-catalog process, we have the opportunity to mitigate many observational effects—for instance, correcting sensor anomalies during instrument signature removal—and we must ensure that our calibration and measurement algorithms do not introduce any significant biases. The catalog-to-cosmology step must account for any remaining systematic effects. In this thesis we focus on challenges in the images-to-catalog step; for detailed reviews of how cosmological information is inferred from weak-lensing shape catalogs, see Mandelbaum (2018); Weinberg et al. (2013); Kilbinger (2015); Troxel et al. (2018).

One of the crucial steps in making catalogs for weak lensing analyses is the estimation of lensing shear from the images of galaxies. Given a sample of galaxies distorted by a true shear g , we denote the estimated shear as $\langle \hat{e} \rangle$ (averaged over a sample of galaxies). For a perfectly accurate estimator, $\langle \hat{e} \rangle \equiv g$; however, for realistic estimators, the measurement may be biased. For example, “model

¹²The Dark Energy Survey imaged hundreds of millions of galaxies over 6 years of observation.

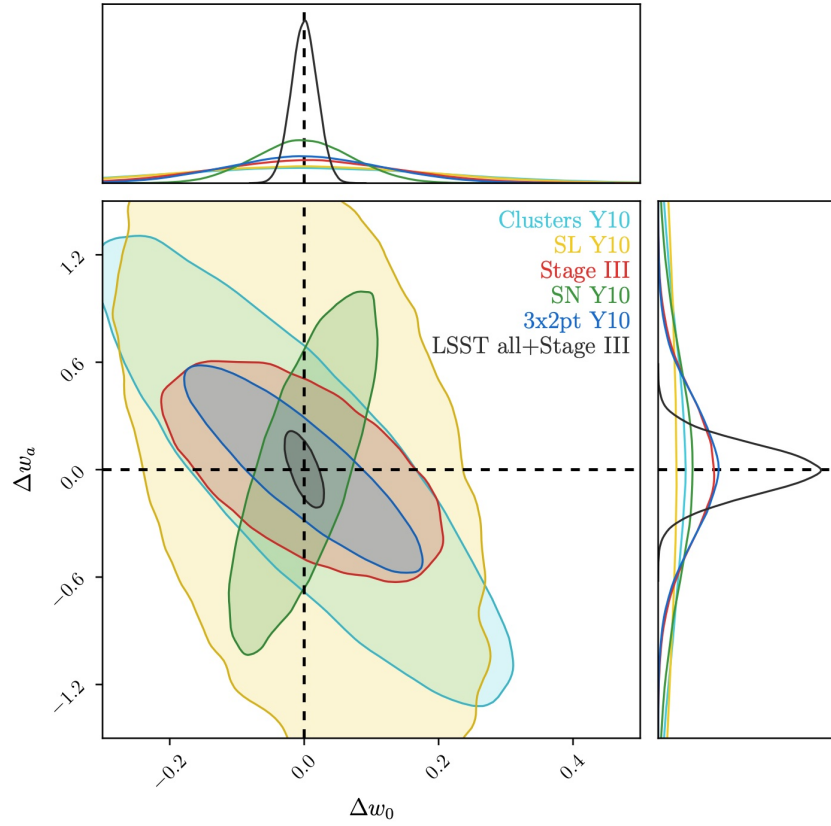


Figure 1.11: Forecast dark energy constraints for the full ten years of LSST data. Colored contours (other than red) show constraints from individual probes. The black contour is the joint precision of all probes when combined with existing “Stage III” surveys (red contour). Axes show Δw_o and Δw_a , the difference of dark energy equation-of-state parameters w_o, w_a from their fiducial values. Contours show 68% confidence level constraints. From the DESC Science Requirements Document (The LSST Dark Energy Science Collaboration et al., 2021a).

bias” may be introduced in the shear estimation step when galaxy shear is estimated via a maximum likelihood fit to a galaxy model; the choice of galaxy model, which is necessarily unrealistic (see the complex galaxy profile in Figure 1.12, for example), can introduce bias in the estimate of shear. These shear estimators can be calibrated with detailed simulations. More recently, self-calibration methods (“Metacalibration” and “Metadetection”), described in Huff & Mandelbaum (2017); Sheldon & Huff (2017); Sheldon et al. (2020, 2023), have been developed to measure the response of the galaxy to shear directly on data. All these methods require an accurate model of the PSF at the location of the galaxy.

The bias in a shear estimator is typically parameterized as follows:

$$\langle \hat{e} \rangle = (1 + m)g + c, \quad (1.51)$$

where m is called the multiplicative bias and c the additive bias (Heymans et al., 2006). The magnitude of the true cosmic shear signal g for any individual galaxy is 0.01; the requirement on mean multiplicative bias for the LSST DESC analysis of cosmic shear in the planned ten-year survey is $m < 3 \times 10^{-3}$. If m is greater than this value, the bias will be larger than the statistical precision of the survey. A requirement on additive bias is more challenging to quantify—see extensive discussion in Meyers & Burchat (2015a), where the authors determined that the variance of c could be as large as the statistical uncertainty if its value (for all additive systematic effects combined) exceeds 1.8×10^{-7} .

1.4.1 Point spread functions

In addition to shear, the imaged shape of the galaxy is also affected by the atmosphere, the optics, diffusion in the sensors, pixelization, noise, and other possible artefacts—as illustrated in Figure 1.12. The middle panel of Figure 1.12 shows the blurring of the galaxy image by the point spread function (PSF). The PSF is the transfer function of the atmosphere,¹³ the telescope optics, and the sensors, and causes distortions in the shapes and sizes of images of objects. These distortions from the PSF are one of the primary sources of systematic uncertainty for weak lensing if not modeled and corrected to very high accuracy. We discuss the PSF in this section and its impact on weak lensing measurements in Section 1.4.2.

The output of an imaging system is linear in intensity when illuminated with incoherent light—i.e., light waves whose phases vary in an uncorrelated way; see Liaudat et al. (2023) for a complete derivation. In this case, the transformed image intensity $I_{\text{image}}(\theta_x, \theta_y)$, as a function of image coordinates θ_x, θ_y , can be described as a convolution of the source image $S(\theta_x, \theta_y)$ by the PSF $I(\theta_x, \theta_y)$,

$$I_{\text{image}}(\theta_x, \theta_y) = S(\theta_x, \theta_y) \otimes I(\theta_x, \theta_y). \quad (1.52)$$

¹³For ground-based imaging—space-based missions bypass the atmosphere but sacrifice the large étendue (product of field of view and mirror size) enabled by optical complexity of a ground-based instrument like Rubin Observatory.

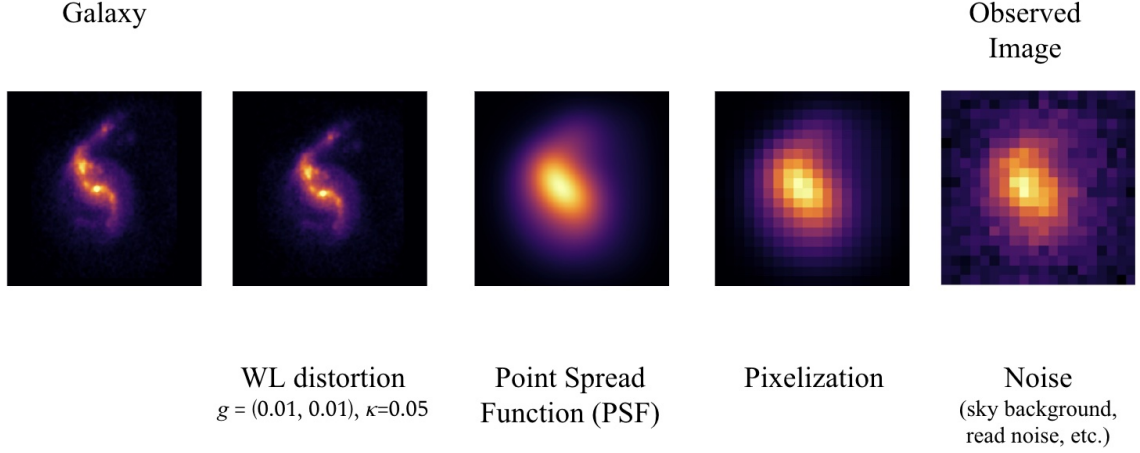


Figure 1.12: Illustration of weak gravitational lensing and subsequent observational effects impacting images of galaxies taken with Rubin Observatory. The image of a galaxy undergoes distortion due to cosmic shear and loss of resolution due to smearing by the PSF and pixelization in the detector. Image credit: Sowmya Kamath.

The intensity profile of the PSF can be described analytically under some reasonable assumptions¹⁴ (Liaudat et al., 2023; Goodman, 2005):

$$I(\theta_x, \theta_y) = \frac{a^2}{\lambda^2 z^2} \left| \int \int_{-\infty}^{\infty} P(u, v) e^{i \frac{2\pi}{\lambda} W(u, v)} e^{-i \frac{2\pi}{\lambda z} (u\theta_x + v\theta_y)} du dv \right|^2, \quad (1.53)$$

for constant amplitude a , wavelength λ , distance between telescope pupil and image plane z , pupil transfer function $P(u, v)$, and pupil coordinates u, v . $W(u, v)$ describes the optical path difference, at each pupil location u, v , between the aberrated incident wave and a perfect spherical wave.

The full width at half-maximum value (FWHM) of this intensity profile is often referred to as the PSF “size”, and is a measure of image quality. PSFs with large FWHM will produce very blurred images; on the other hand, one risks undersampling if the PSF size approaches the pixel size of a detector.

So far we have treated only the “local” spatial variation of the PSF – i.e., the variation of the intensity profile of the PSF. However, we must also consider the spatial variation of the PSF across the focal plane – i.e., how the PSF depends on the angle of incidence of the light rays or, equivalently, the spatial position on the focal plane. The PSF also depends on the spectral energy distribution of the source, as diffraction and refraction depend on wavelength, and both the atmosphere and sensors have a chromatic response (Meyers & Burchat, 2015b,a). The PSF changes temporally because both the state of the telescope and the atmosphere are dynamic. And, finally, the PSF depends on the

¹⁴Assumptions are that we consider scalar electric fields, the telescope aperture is large compared to the observing wavelength, and we are in regimes where the Fresnel and far-field Fraunhofer approximations apply.

length of an exposure, since atmospheric turbulence varies on time scales much shorter than the typical exposure times of tens to hundreds of seconds. So in general,

$$I = I(\theta_x, \theta_y, x, y, t, \lambda \Delta t_{\text{exp}}), \quad (1.54)$$

where we have used x, y to denote the PSF position on the focal plane.

As mentioned above, the PSF includes contributions from the atmosphere as well as the telescope optics and the sensors; because the component due to the atmosphere has very different behavior from the latter two (in terms of dependence on time, wavelength, and length of an exposure), the PSF is often described as the sum of two components:

$$I(\theta_x, \theta_y) = I_{\text{opt+sensor}}(\theta_x, \theta_y) + I_{\text{atm}}(\theta_x, \theta_y). \quad (1.55)$$

The list of effects contributing to each part of the PSF is extensive. For example, the optical PSF has contributions from diffraction, optical aberrations in the system of mirrors and lenses, mirror polishing errors, scattered light, etc. Sensors contribute through diffusion, chromatic effects, the so-called “brighter-fatter effect”¹⁵, charge-transfer inefficiencies, etc. See Liaudat et al. (2023) for a more exhaustive list as well as a detailed discussion of each contribution. The atmospheric PSF captures the dynamic effects from refraction in the turbulent atmosphere, which contributes effects at scales that are both small and large compared to the size of the field of view; the rest of this thesis will detail many aspects of the atmospheric PSF.

1.4.2 Impact of the PSF on weak lensing measurements

As one can see from Figure 1.12, the PSF has a very significant effect on the apparent size and shape of galaxies as seen from ground-based telescopes. In order to accurately estimate the shear induced by gravitational lensing, one must first correct for these subsequent distortions caused by the PSF. As we have seen, the PSF varies with time, so we must do this correction on each exposure in the survey. We must also account for the spatial variation across the field of view. This correction, for a single exposure image, is often accomplished via the following procedure:

1. From the objects detected in the image, select bright (but not saturated) stars. Stars are point sources with $S(\theta_x, \theta_y) = \delta(\theta_x, \theta_y)$; hence, their images are themselves images of the PSF.¹⁶
2. Measure the PSF from the image of each star. This can be done by fitting a parametric model to the intensity profiles, for example, or by interpolating the PSF from flux values in pixels to a standard grid and measuring the resulting intensity at each point. By whichever method

¹⁵Due to dynamical image distortions caused by accumulated charge during exposures in “thick” CCD sensors, such as those in the LSST Camera, bright stars appear slightly larger in size than faint stars.

¹⁶We exclude discussion of wavelength dependence here; see Meyers & Burchat (2015a) for details.

is chosen, the result of this step is, at each star location, a set of parameters p whose values describe the PSF profile.

3. Use an interpolation method (e.g., a fit to a polynomial or Gaussian Process interpolation) to infer the values of p at the locations of galaxies on the image, using the values measured at the locations of stars.
4. Reconstruct an image of the PSF at each galaxy location, from the PSF model used in (2) and the inferred parameters p from (3). These images can then be used in the shear estimation step to account for the effect of the PSF on each galaxy image.

This general method has been demonstrated previously in DES (Jarvis et al., 2020, 2016), HSC (Li et al., 2022), and KiDS (Giblin et al., 2021), and will be used for Rubin Observatory LSST. See also Mandelbaum (2018); Liaudat et al. (2023) for more details and reviews of various proposed algorithms.

The above modeling and interpolation algorithm provides a model PSF at each observed galaxy location. Errors in the size and shape of these PSF models will introduce multiplicative and additive shear biases (m and c in Equation 1.51), respectively, to the estimate of shear from the associated galaxy image. It is also possible that the errors in the PSF model are spatially correlated across the focal plane. The spatial two-point correlation functions of these errors determine the bias introduced to weak lensing observables – i.e. the galaxy shape correlation functions described in Section 1.2.2. Five correlation functions involving PSF errors (often referred to as the eponymous ρ statistics by Rowe (2010)) have been widely used to quantify the impact of second-moment PSF size and shape errors on measurements of cosmic shear; they depend on PSF shape residuals, PSF size, and PSF size residuals, and describe the additive terms in the galaxy shape correlation functions due to PSF modelling errors (Rowe, 2010; Jarvis et al., 2016).

1.5 Motivation for this thesis

For accurate (unbiased) measurements of cosmic shear, all other sources of correlated image distortions must be modeled and/or removed. One such source is the correlated blurring of images due to Kolmogorov turbulence in the atmosphere, which dominates the point spread function (PSF) for ground-based instruments such as DES, HSC, KiDS, and Rubin Observatory (e.g., Jarvis et al. (2016); Heymans et al. (2012)). Algorithms are under active development for more accurately modeling and interpolating the PSF across the focal plane (see, e.g., Jarvis et al. (2020); Bertin (2011)), correcting galaxy shapes for the PSF, and calibrating measures of cosmic shear from galaxy shapes (Sheldon & Huff, 2017; Huff & Mandelbaum, 2017; Sheldon et al., 2023; Gatti et al., 2021). A challenge looking ahead to the analysis of the LSST data from Rubin Observatory is to ensure these algorithms achieve low spatial correlations of PSF residuals, in line with the stringent requirements

on accuracy of shear estimation given the unprecedented precision afforded by the ten-year survey, and perform well on real data.

Both real astronomical images and high-fidelity simulated images are important tools for developing and optimizing these algorithms and measuring their performance. Simulations play a unique role because the input parameters we are trying to measure or infer, including shear, are known, and because data from existing surveys are not representative of the PSF at Rubin—Rubin has shorter exposure times, a larger field of view, and different local topography. In addition, physically motivated PSF correction algorithms that treat the optical and atmospheric components of the PSF separately require for method development images of each PSF component in isolation; this is only possible in simulations.

In this thesis, we explore the ways in which various types of data—measurements of weather and atmospheric turbulence (Chapters 2, 3, 4), wavefront data (Chapter 5, Appendix D), and short-exposure PSF images (Chapter 6)—can each be exploited to inform and validate our simulation tools, with the goal of producing simulated PSF images that are realistically complex and include relevant correlations with the environment. These state-of-the-art tools will play a role in using cosmic shear to measure cosmological parameters with both high precision and high *accuracy*.

Chapter 2

Imaging through the turbulent atmosphere

Big whirls have little whirls which feed on their velocity,
and little whirls have lesser whirls and so on till viscosity.
— Lewis Fry Richardson

2.1 Introduction

A significant source of noise for any ground-based astronomical imaging is the effect of the turbulent atmosphere. Turbulence in the atmosphere is stochastic and spatially correlated, and, as we know from watching stars appear to twinkle in the night sky, it is also time-varying on time scales as short as milliseconds. This sets it apart from other components of the PSF, since contributions to the PSF from telescope optics and sensor electronics evolve on much longer time scales. For astronomical applications we consider the atmosphere as a medium that refracts light, and are interested in understanding its accumulated effects on light rays coming from stars or galaxies and ending up in our telescope. As we have seen, we are in particular interested in the spatial correlations of noise for weak lensing measurements; we need to understand how individual objects are affected by this atmospheric refraction *and* how the effects vary over the focal plane.

The physics of atmospheric turbulence is complex and fluid-mechanical in nature, and as such is prohibitively expensive to simulate precisely even on modern day supercomputers. Nevertheless, it is possible to derive some statistical properties of turbulence, as we will do in the following chapter. Although some topics remain areas of active research, much of this formalism was first

described in the mid-twentieth century. This knowledge, combined with Fourier optics, helps us better understand the variations in the PSF over the focal plane of a telescope.

2.2 The Physics of Atmospheric Turbulence

A stationary atmosphere would have a vertical temperature gradient, and thus a uniform density at any given altitude. However, the Earth's atmosphere is in constant motion – what conclusions can we draw from this? The motion will mix pockets of air that are at different temperatures, leading to patches of over/under densities of air at any given altitude. These slight density changes cause slight index of refraction variations between adjacent air pockets. Due to the continual motion of these air pockets, the structure of temperature (and thus, refraction) in the atmosphere is variable and chaotic. This leads to two distinct effects, which impact imaging in different ways:

1. scintillation, which can lead to variation in the intensity of light entering the telescope, is not expected to be a significant effect for Rubin Observatory exposures due to averaging over the large aperture.
2. spatial variation of optical phase over the telescope aperture, which on very short exposures (for large telescopes) causes speckle patterns: the diffraction pattern of interfering wavefronts results in multiple images of the same star, each about the size of the telescope diffraction limit, distributed over an area the size of the seeing disc. For long exposures, these highly variable instantaneous speckle patterns blur together into a smooth point-spread function.

Here we are primarily interested in optical phase variations, and will accordingly focus on those derivations. Refer to Roddier (1981), Lawrence & Strohbehn (1970), Clifford (1978), and Léna et al. (1998) for much of the following formalism.

In addition to assuming a clear atmosphere (e.g., ignoring scattering and absorption due to aerosols or precipitation) we will assume that the atmosphere is in pressure equilibrium, meaning that we can treat turbulent fluctuations as isobaric processes; this is reasonable in the sense that pressure fluctuations are small and disperse very rapidly compared to the temperature fluctuations of interest for refractive index variation. In general, we can approximate the index of refraction in air by

$$n(\mathbf{x}, t) \simeq 1 + 0.00029 \frac{\rho(\mathbf{x}, t)}{\rho_{\text{air,STP}}} . \quad (2.1)$$

We can see from Equation 2.1 that deviations in the density of air $\rho(\mathbf{x}, t)$ will result in fluctuations in $n(\mathbf{x}, t)$; moreover, in a pressure equilibrium, temperature fluctuations will be the main driver of density changes. Taking these temperature fluctuations as isobaric, fluctuations in index of refraction

can be related to temperature and pressure, at optical wavelengths ($\lambda \approx 600 \text{ nm}$), by

$$\frac{dn}{dT} = \frac{2.9 \times 10^{-4}}{T} \frac{P/1000 \text{ mbar}}{T/273 \text{ K}}. \quad (2.2)$$

Turbulence in the atmosphere is sourced by both convection (heat in low layers of the atmosphere – e.g., from solar heating – leads to convective bubbles of warm air rising, mixing air of different temperatures) and wind shear (a fluid instability called the Kelvin-Helmholtz instability arises when the vertical wind-speed gradient $\frac{dv_x}{dz}$ between two adjacent layers is large). In the following subsections, we will explore theoretical models of turbulent flows and what they tell us about refractive index fluctuations and image quality.

2.2.1 Kolmogorov turbulence

The onset of turbulence is determined by the magnitude of the Reynolds number, $Re = uL/\nu$, which depends on only three quantities that characterize the flowing fluid: the flow speed u , the characteristic linear size L , and the kinematic viscosity of the fluid, ν . When viscous forces dominate, the Reynolds number is small and flow is laminar; when inertial forces and instabilities dominate, the Reynolds number is large and the flow is turbulent.

Turbulence typically dissipates quickly as kinetic energy is converted to internal energy via viscous shear stresses, unless there is an obstacle or injection of energy to sustain the turbulent flow. (Think, for example, of a cup of coffee; the fluid quickly settles after one stops stirring.) The presence of an obstacle is equivalent to a persistent energy source; in both cases fluid is displaced, creating a space devoid of fluid. Fluid then flows in to fill this void, forming a swirl on each side, followed by a short reverse current. This swirl and reverse flow in a turbulent fluid is called an *eddy*. In the context of astronomical imaging, we try to minimize the impact of obstacle-driven turbulence by building observatories on isolated mountain peaks and designing domes that disturb the natural flow of air as little as possible; however, we cannot control the large scale injection of energy into the atmosphere.

Whether driven by obstacles or an energy source, turbulence creates eddies of many different sizes, though most of the kinetic energy is contained in the largest structures – eddies of size $\sim L$. At these largest scales, the Reynolds number is much greater than 1 and instabilities in the large eddies cause them to fracture into smaller eddies, in effect transferring kinetic energy to smaller scales. This continues until the size of the eddies reaches the so-called inner scale of turbulence – some value $l_0 \ll L$, at which $Re \sim 1$ – where viscous forces take over and the energy is instead dissipated as heat. This flow of energy forming a hierarchy of eddies, from the large scales of energy input to the scale of viscous dissipation, is referred to as the turbulent energy “cascade”. This energy cascade is common to many types of turbulence that have in common a region of spatial scales intermediate between the scales of energy input and dissipation. For example, milk poured

into coffee creates turbulent flow first on scales on the order of the coffee cup's size but quickly cascades to smaller swirls as the liquids mix. As another example, Figure 2.1 shows turbulent eddies in a tank of water, visualized with the addition of a small amount of fluorescent dye and illuminated with a thin sheet of light. It is important to note that the chaotic nature of turbulence is possible only in a three dimensional system; this visualization technique is showing us a 2D slice through the three dimensional eddies.

Kolmogorov (1941) predicted a specific power law dependence for turbulent energy as a function of spatial frequency for scales within the intermediate range of spatial scales in this energy cascade. This range in which Kolmogorov's model is valid is referred to as the inertial range, and is bound on small spatial scales by the inner scale l_0 and at large scales by the energy input scale (the "outer scale") L_0 . In the case of the atmosphere, the inner and outer scales are typically on the order of millimeters and tens or hundreds of meters, respectively. Key to Kolmogorov's model, as we will see shortly, was the realization that between these inner and outer scales, and in a steady state, there are no direct forces on the fluid and viscous dissipation is insignificant. In this case, the energy input rate and the energy dissipation rate are equal. We denote this rate by ε , which has dimensions of energy per unit time per unit mass, or $\text{length}^2 \text{time}^{-3}$. The power spectrum as a function of spatial frequency κ is given by $E(\kappa)$, where the total energy per unit mass in a volume $d^3\kappa$ about a particular κ is $\int E(\kappa)d^3\kappa$, with κ a vector of spatial frequencies in x, y, z . Each component in κ has units of inverse length, and $\int E(\kappa)d^3\kappa$ has units of energy per unit mass, or $\text{length}^2 \text{time}^{-2}$, so the spectrum $E(\kappa)$ has dimensions of $\text{length}^{2+3} \text{time}^{-2}$.

The second realization that led Kolmogorov to his famous result is that in the inertial range – i.e., for any frequency $1/L_0 \leq \kappa \leq 1/l_0$ – the system has access to only the energy dissipation rate ε and the scale κ ; the energy spectrum $E(\kappa)$ can therefore depend on only these two variables. From here, simple dimensional analysis relating $E(\kappa) \propto \varepsilon^\alpha \kappa^\beta$ leads to Kolmogorov's power law spectrum for 3-dimensional turbulence, in terms of the vector of spatial frequencies κ :

$$E(\kappa) \propto \varepsilon^{2/3} |\kappa|^{-11/3}, \quad (2.3)$$

where taking the magnitude of κ indicates we assume isotropic turbulence. We recall that this model is, implicitly, valid only between spatial scales l_0 and L_0 ; modifications of the Kolmogorov model that extend the validity beyond the outer scale will be discussed in Section 2.3.1.

How does this cascade of kinetic energy, and Kolmogorov's power law describing it, relate to the quantities of interest for imaging through a turbulent atmosphere – i.e., fluctuations in temperature and index of refraction? In turbulent flow, the concentration of certain additives also follow Kolmogorov's law, if the additive is not reactive and does not affect the dynamics of the flow; the mixing of cold and warm air satisfy these conditions, so this result applies to temperature fluctuations, and by extension (via Equation 2.2) to variations in index of refraction (Obukhov, 1949; Yaglom, 1949; Corrsin, 1951). In other words, gas motion (which follows Equation 2.3) does not itself induce



Figure 2.1: Turbulent eddies visualized in a tank of water via the addition of fluorescent dye and illuminated with a thin light sheet (Sreenivasan, 1999).

refraction; rather, the motion moves air pockets of varying temperature within its turbulent flow, resulting in fluctuations in temperature (and therefore density and index of refraction), which follow the same statistics.

There is observational evidence for both wind and optical phase fluctuations following this general spectral form, within some range of length scales; Figure 2.2 shows two such measurements. The first panel shows measurements of power in wind fluctuations in the stratosphere (at ~ 19 km), recorded with LITOS (Leibniz Institute Turbulence Observations in the Stratosphere). This balloon-borne instrument records wind speed fluctuations as it travels vertically, sampling every ~ 2.5 mm with a device called a constant temperature anemometer, which operates by measuring the voltage fluctuations in a hot wire due to convective heat loss associated with ambient flow. The resulting power spectrum very closely matches the predicted Kolmogorov power law. Note that this experiment is measuring the turbulent spectrum in only one-dimension (vertical); the Kolmogorov prediction for the power law exponent for this case is $-\frac{5}{3}$, corresponding to the slope of the red dashed line in the figure. The solid red line combines a Kolmogorov power law in the inertial range, where $L > l_0$, and a viscous flow model $\propto |\kappa|^{-7}$ for $L < l_0$. The authors report an inner scale of $l_0 = 4.5$ cm, indicated with the vertical line. The second panel shows measured values of the power law exponent from optical phase data reconstructed from Shack-Hartmann wavefront sensors in the Gemini Planet Imager adaptive optics system. The optical phase power spectrum for each run was fit to a power law between spatial frequencies of 0.3 and 1 m^{-1} (this limited range reduces contamination from sensor aliasing and tip/tilt effects). The resulting exponents are plotted against values of ΔT_M , a proxy for dome turbulence, but show no clear dependence on this variable. These measurements, though noisy, are consistent with the Kolmogorov prediction of $-\frac{11}{3}$, corresponding to the horizontal line.

The results in Figure 2.2 are not unique; there is broad evidence to support the $|\kappa|^{-11/3}$ dependence (for intermediate values of κ) in Equation 2.3 throughout the literature on atmospheric turbulence. Observations of non-Kolmogorov turbulence do exist—certainly for scales $\kappa < 10^{-3}$ —and we will come back to these in Section 2.3.1. In the following, we review formalism from the literature to describe the implications of atmospheric turbulence to wide-field, ground-based imaging.

2.3 Turbulence distorted wavefronts

Generally, a wavefront describes the set of points in an electromagnetic field that have the same phase. In astronomy, we assume that wavefronts entering the atmosphere from far away astronomical sources are plane waves; however, by the time they reach the telescope, having traveled through a spatially and temporally varying turbulent atmosphere, they have incurred aberrations. We quantify these aberrations at the telescope pupil (defined to be at altitude $z = 0$) as the two-dimensional phase shift (or just “phase”) with respect to the original planar wavefront.

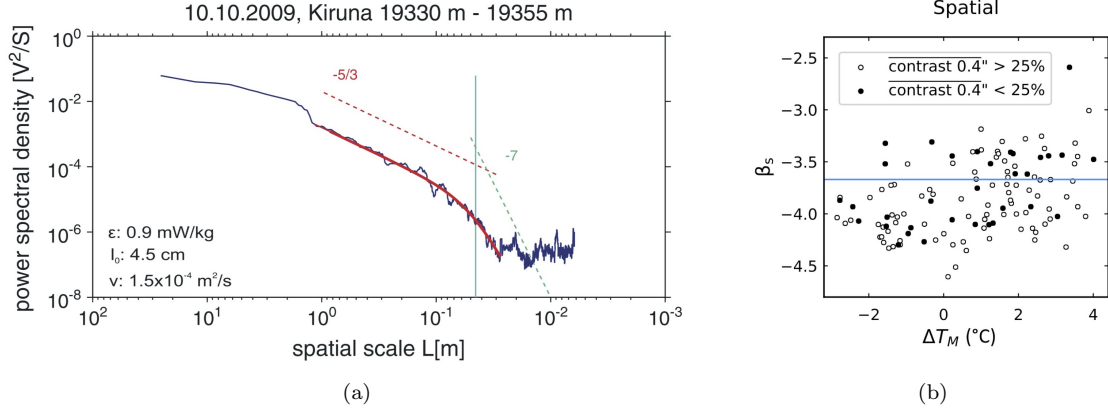


Figure 2.2: Atmospheric turbulence power spectrum measurements. a. Power in wind fluctuations in the stratosphere from Haack et al. (2014), measured by the balloon-borne instrument LITOS. The solid red line models a smooth transition from a Kolmogorov power law spectrum to the viscous subrange, with the turnover point (turquoise vertical line) given by the inner scale l_0 . b. 122 measurements of a power-law exponent from power spectra of optical phase variation from the Gemini Planet Imager (GPI) Shack-Hartmann wavefront sensor, as presented in Tallis et al. (2020). The values are plotted as a function of ΔT_M , which is a proxy for dome turbulence. The blue horizontal line indicates the expected Kolmogorov value of $-\frac{11}{3}$. e

More concretely, we define a wavefront with wavenumber $k = \frac{2\pi}{\lambda}$ propagating vertically downward toward a telescope aperture at altitude $z = 0$. We denote the phase of this wavefront by $\phi(\mathbf{x})$, where \mathbf{x} is a vector of horizontal coordinates defined over the telescope pupil. We will first consider a simple atmosphere with a single turbulent layer at height $z = h$ and vertical extent Δh . We assume that Δh is much larger than the correlation length scale of the refractive index – i.e., Δh contains many turbulent eddies – but small enough that diffraction effects within the layer are negligible. We will return to this *thin screen approximation* many times during this thesis.

The absolute phase of the wavefront when it reaches the telescope can be found by integrating the index of refraction over all altitudes $z > 0$. The PSF is insensitive to the mean phase, since only variations from this will cause interference in the telescope – this means we only need to integrate over the fluctuations in refractive index away from the mean. For a scalar field α (in this case, $\alpha = n$) we denote the fluctuations as

$$\theta_\alpha(\mathbf{x}, z, t) = \alpha(\mathbf{x}, z, t) - \langle \alpha \rangle. \quad (2.4)$$

In the case of a single turbulent layer, the bounds of the integral reduce to the extent of the layer, h to $h + \Delta h$:

$$\phi(\mathbf{x}) = k \int_h^{h+\Delta h} \theta_n(\mathbf{x}, z, t) dz. \quad (2.5)$$

Ultimately, we would like to find an analytic form for the autocorrelation function (defined precisely below) of the wavefront, as it will be a useful way to describe PSF quantities in relation to properties of turbulence that lead to variations in $n(\mathbf{x}, z, t)$.

Before proceeding further, we define some statistical properties that will be used in the derivations to follow. Given some variable such as temperature T , the fluctuation in T is described by $\theta_T \equiv T - \langle T \rangle$, and its autocorrelation, as a function of separation vector \mathbf{r} , is

$$B_T(\mathbf{r}) \equiv \langle \theta_T(\mathbf{x}) \theta_T(\mathbf{x} + \mathbf{r}) \rangle, \quad (2.6)$$

where the angle brackets denote averaging over positions \mathbf{x} . This autocorrelation can in principle be evaluated via a Fourier transform of the relevant power spectrum¹; however, for the Kolmogorov turbulence spectrum of Equation 2.3, this method is not viable since the integral $\int_{-\infty}^{\infty} |\boldsymbol{\kappa}|^{-11/3} e^{-i\boldsymbol{\kappa} \cdot \mathbf{r}} d\boldsymbol{\kappa}$ diverges. As shown by Tatarskii (1971), we can work around this by using the structure function $D_T(\mathbf{r})$, which has the advantage of a finite integral and is related to the autocorrelation:

$$D_T(\mathbf{r}) \equiv \langle |\theta_T(\mathbf{x} + \mathbf{r}) - \theta_T(\mathbf{x})|^2 \rangle \quad (2.7)$$

$$= 2(B_T(0) - B_T(\mathbf{r})). \quad (2.8)$$

Lastly, if $\phi(\mathbf{x})$ is Gaussian distributed, a signal $\psi(\mathbf{x}) = e^{-i\phi(\mathbf{x})}$ has the following useful property:

$$B_\psi(\mathbf{r}) = e^{-\frac{1}{2}D_\phi(\mathbf{r})}. \quad (2.9)$$

Since the phase $\phi(\mathbf{x})$ of a wavefront $\psi(\mathbf{x})$ distorted by the atmosphere is a collection of points sampled from a Gaussian random field, this property applies to the autocorrelation of the wavefront $B_\psi(\mathbf{r})$ and the structure function of the phase $D_\phi(\mathbf{r})$. This simplifies the derivation of wavefront autocorrelation to a derivation of *phase* autocorrelation $B_\phi(\mathbf{r})$: once we work out $B_\phi(\mathbf{r})$ all we have to do is apply Equation 2.8 and then Equation 2.9. As we know from Equation 2.5, $\phi(\mathbf{x})$ depends on the integral within the turbulent layer of the spatially varying θ_n , the deviation from the mean of the index of refraction, so that

$$B_\phi(\mathbf{r}) = \langle \phi(\mathbf{x}) \phi(\mathbf{x} + \mathbf{r}) \rangle \quad (2.10)$$

$$= k^2 \int_h^{h+\Delta h} \int_h^{h+\Delta h} \langle \theta_n(\mathbf{x}, z', t) \theta_n(\mathbf{x} + \mathbf{r}, z'', t) \rangle dz' dz'', \quad (2.11)$$

where we have pulled the average over horizontal position \mathbf{x} inside the z integrals. With the goal of recasting this expression in terms of the autocorrelation function for the refractive index, $B_n(\mathbf{r}, z) = \langle \theta_n(\mathbf{x}, z', t) \theta_n(\mathbf{x} + \mathbf{r}, z' + z, t) \rangle$, we change to a new variable $z = z'' - z'$; note that we can treat z'

¹Via the Wiener–Khinchin theorem.

as fixed while integrating over z . This gives

$$B_\phi(\mathbf{r}) = k^2 \int_h^{h+\Delta h} \int_{h-z'}^{h+\Delta h-z'} B_n(\mathbf{r}, z) dz' dz. \quad (2.12)$$

Under the thin screen approximation, $\theta_n(\mathbf{x}, z, t)$ oscillates many times within the vertical extent Δh . Because of this, the refractive index autocorrelation will be zero or negligible for vertical separations z on the order of Δh or greater, and we can equivalently write the integral over z as an integral from $-\infty$ to $+\infty$. Now that z' no longer appears in the limit of integration over z , the integral over z' reduces to Δh , and we find

$$B_\phi(\mathbf{r}) = k^2 \Delta h \int_{-\infty}^{\infty} B_n(\mathbf{r}, z) dz. \quad (2.13)$$

The phase autocorrelation depends linearly on the thickness of the turbulence layer, and on the integral of the refractive index autocorrelation.

Going back to Equation 2.8, we can use this result to solve for the wavefront phase structure function:

$$D_\phi(\mathbf{r}) = 2k^2 \Delta h \int_{-\infty}^{\infty} (B_n(0, z) - B_n(\mathbf{r}, z)) dz \quad (2.14)$$

$$= 2k^2 \Delta h \int_{-\infty}^{\infty} (B_n(0, 0) - B_n(\mathbf{r}, z) - B_n(0, 0) + B_n(0, z)) dz \quad (2.15)$$

$$= k^2 \Delta h \int_{-\infty}^{\infty} (D_n(\mathbf{r}, z) - D_n(0, z)) dz. \quad (2.16)$$

Up to this point we have not made any assumptions about the nature of the turbulence itself—only that the wavefront $\phi(\mathbf{x})$ is Gaussian distributed and about the geometry of the turbulent layer. To go further, we will need a functional form for $D_n(\mathbf{r}, z)$.

We will now summarize a derivation of the structure function for Kolmogorov turbulence; however, it should be noted that this description of a turbulent field is equivalent to Equation 2.3 ($E(\boldsymbol{\kappa}) \propto \varepsilon^{2/3} |\boldsymbol{\kappa}|^{-11/3}$), and can be rigorously derived by evaluating a Fourier integral for each term in Equation 2.8, based on the power spectral density from Equation 2.3. The general form for $D_n(\mathbf{r})$ can also be found with dimensional analysis arguments similar to those of Section 2.2.1, assuming locally homogeneous and isotropic incompressible flow. The structure function might depend, in that case, on the energy dissipation rate ε , the magnitude of the separation vector \mathbf{r} between two points in the velocity field, and the kinematic viscosity of the fluid ν . At scales $|\mathbf{r}| > l_0$ in the inertial range, however, ν should not appear in the equation since we assume negligible viscous dissipation. Dimensional arguments then lead us to the following structure function for temperature fluctuations:

$$D_T(\mathbf{r}) = C_T^2 |\mathbf{r}|^{2/3}, \quad (2.17)$$

where C_T^2 is called the temperature structure constant. As discussed previously, and shown in Equation 2.2, variations in temperature are directly related to variations in index of refraction, so the functional forms of the two structure functions are the same (as are the power spectra):

$$D_n(\mathbf{r}) = C_n^2 |\mathbf{r}|^{2/3}. \quad (2.18)$$

These structure constants for temperature and refractive index can also be related using Equation 2.2 as $C_n^2 = C_T^2 \frac{dn}{dT}$. We are now ready to resume our analysis of wavefront autocorrelation.

Using the above refractive index structure function, Equation 2.16 becomes

$$D_\phi(\mathbf{r}) = k^2 \Delta h \int_{-\infty}^{\infty} C_n^2 (\sqrt{|\mathbf{r}|^2 + z^2})^{2/3} - z^{2/3} dz. \quad (2.19)$$

After evaluating the above integral numerically and generalizing the result to include zenith angle ζ^2 , we apply Equation 2.9 and find the wavefront autocorrelation function to be

$$B_\psi(\mathbf{r}) = \exp \left[-\frac{1}{2} 2.914 k^2 |\mathbf{r}|^{5/3} C_n^2 \sec \zeta \Delta h \right]. \quad (2.20)$$

This result was the goal of this section; to facilitate interpretation we can go a step further and define a length scale r_0 , first proposed by Fried (1965) and thus referred to as the *Fried parameter*,

$$r_0 \equiv (0.423 k^2 C_n^2 \sec \zeta \Delta h)^{-3/5}. \quad (2.21)$$

As we will see in Section 2.3.2, this quantity has dimensions of length. With this definition, Equation 2.20 becomes

$$B_\psi(\mathbf{r}) = \exp \left[-3.44 \left(\frac{|\mathbf{r}|}{r_0} \right)^{5/3} \right]. \quad (2.22)$$

So we see that the autocorrelation of the wavefront falls off exponentially with separation distance relative to the Fried parameter. The constant factor in Equation 2.21 may seem arbitrary, but as defined the Fried parameter represents a very important length scale for imaging through a turbulent medium and can be interpreted as follows:

1. The phase variance within an aperture of size r_0 is approximately 1 rad^2 .
2. A long-exposure image through atmospheric conditions characterized by r_0 has the same angular resolution as an image taken with a telescope of diameter r_0 .

The Fried parameter, as we see in Equation 2.21, depends on zenith angle, turbulence structure constant, and wavelength (through the wavenumber k). More precisely, $r_0 \sim \lambda^{6/5}$, which means a long-exposure, seeing-limited PSF would have size that depends on wavelength as $\frac{\lambda}{r_0} \sim \lambda^{-1/5}$. The

²When observing at an angular distance ζ from zenith, the layer thickness Δh is increased by a factor of $(\cos \zeta)^{-1}$.

inverse wavelength dependence results in larger PSF sizes for images taken at shorter wavelengths. This is in contrast to diffraction-limited seeing which has a wavelength dependence of $\frac{\lambda}{D} \sim \lambda$.

The following subsections cover extensions and generalizations of these results as applied to model atmospheres with more complex vertical and/or horizontal turbulence³ structure.

2.3.1 Von Kármán turbulence

There is one more important theoretical model of turbulence to introduce here, one which addresses certain shortcomings of Kolmogorov’s power law. From Section 2.2.1 we are already aware that Kolmogorov’s power law allows for power at infinite spatial scales (small κ). While we can state, as we did upon its introduction, that the Kolmogorov power law model in Equation 2.3 is only valid for spatial scales in the inertial range (l_0, L_0) , a new model is needed in order to include a finite outer scale in calculations of a wavefront phase structure function. An important point to note is that there are two *distinct* outer scale parameters: the turbulence outer scale and the wavefront outer scale (Tokovinin et al., 1998). It is common to see references to an “outer scale” without specification in the literature; from the definitions below it is typically straightforward to distinguish between the two in context, but it is valuable nonetheless to clarify the distinction. The turbulence outer scale sets an upper bound on the turbulence strength at low frequencies, which would otherwise diverge as $\kappa \rightarrow 0$; it is the L_0 parameter we have referred to thus far in this thesis. The wavefront outer scale, on the other hand, represents spatial scales associated with the largest correlations in wavefront phase, and as such, it is the parameter more directly relevant for astronomical imaging contexts. It is defined with reference to the wavefront phase structure function $D_\phi(\mathbf{r})$, and thus is not itself a physical parameter of turbulence – however, Borgnino (1990) showed that it can be expressed as a weighted average of the turbulence outer scales along a propagation path. Following the notation from Tokovinin et al. (1998) and Tokovinin (2002), we denote the wavefront outer scale by \mathcal{L}_0 .

The wavefront phase structure function for Kolmogorov, found by combining Equation 2.9 and Equation 2.20, is

$$D_\phi(\mathbf{r}) = 6.88 \left(\frac{|\mathbf{r}|}{r_0} \right)^{5/3}, \quad (2.23)$$

which diverges for large $|\mathbf{r}|$. To define \mathcal{L}_0 more precisely, we recognize that at large scales \mathbf{r} the wavefront phase autocorrelation $B_\phi(\mathbf{r})$ should drop to zero, which means the structure function $D_\phi(\mathbf{r}) = 2(B_\phi(0) - B_\phi(\mathbf{r}))$ (Equation 2.8) saturates to twice the wavefront variance σ_ϕ^2 – the distance $|\mathbf{r}|$ at which this saturation occurs defines \mathcal{L}_0 :

$$6.88 \left(\frac{\mathcal{L}_0}{r_0} \right)^{5/3} = 2\sigma_\phi^2. \quad (2.24)$$

³In the rest of this thesis, we will use the term “turbulence” to refer specifically to the driver of optical phase variations – e.g., temperature gradients – and not the more general concept of fluid-mechanical turbulence.

This is an intuitive way to define the wavefront outer scale, but does not describe what happens in the transition region, i.e. the scales in between the inertial range (in which Equation 2.23 holds) and saturation (described by Equation 2.24). For this, we must turn to heuristic models; the one most commonly used in the literature was developed by von Kármán (1948) and describes the spatial power spectrum of wavefront phase fluctuations⁴ as

$$E_{\text{vK}}(\boldsymbol{\kappa}) = 0.0228 r_0^{-5/3} (|\boldsymbol{\kappa}|^2 + \mathcal{L}_0^{-2})^{-11/6}. \quad (2.25)$$

Turbulent phase fluctuations with this power spectral density recover Kolmogorov behavior for spatial frequencies $|\boldsymbol{\kappa}| \gg \mathcal{L}_0^{-1}$, while for $|\boldsymbol{\kappa}| \ll \mathcal{L}_0^{-1}$ the spectrum plateaus to a constant value. This change of behavior at low spatial frequencies from Kolmogorov spectrum turbulence has significant impact on the properties of PSFs affected by von Kármán turbulence. For example, Tokovinin (2002) presents an approximation, based on numerical results, for relating von Kármán and Kolmogorov seeing values as a function of $\frac{r_0}{\mathcal{L}_0}$:

$$\left(\frac{\epsilon_{\text{vK}}}{\epsilon_0} \right)^2 \approx 1 - 2.183 \left(\frac{r_0}{\mathcal{L}_0} \right)^{0.356}. \quad (2.26)$$

The approximation is quoted to have relative errors within $\pm 1\%$ for $\mathcal{L}_0/r_0 > 20$. As expected from removing power at large spatial scales, PSF sizes are comparatively *reduced* in von Kármán turbulence.

2.3.2 Vertical structure of turbulence

The results reviewed in the previous section were derived assuming a single thin layer of turbulence in the atmosphere at a single altitude. However, turbulence occurs at a range of altitudes and data suggest that the turbulent fraction – the fraction of a volume of air that is turbulent – starts to drop only above altitudes around 20 km (Haack et al., 2014). Realistically, then, we have to contend with a large volume of turbulent air; how do we modify our calculation of the wavefront autocorrelation to reflect this? To answer this, we first present the current literature on the variation of atmospheric turbulence with altitude relevant for astronomical imaging. This will motivate corresponding modifications to the calculations from Section 2.3.

In the results above, e.g., Equation 2.20, the so-called refractive index structure constant C_n^2 represents the strength of turbulence within the turbulent region. However, calling this a “constant” is a misnomer; turbulence, and therefore refractive index variation, is a spatially and temporally varying field, and thus depends on time, altitude, and horizontal coordinates! We are interested in the column of turbulent air that lies above the telescope during the course of an exposure. At any

⁴Analytical versions of phase structure function for this von Kármán model can also be found (Consortini & Ronchi, 1972; von Kármán, 1948) but are not particularly intuitive so we will not present them here.

given altitude, we treat variation of turbulence in time and horizontal coordinates as equivalent, under the assumption that patches of turbulence move across the telescope aperture as time elapses – e.g., due to wind – but the patches themselves do not evolve. This is the “Taylor frozen flow hypothesis” (Taylor, 1938), which assumes that the timescales of turbulent evolution are greater than the time it takes for patches to cross the telescope aperture. We will come back to horizontal and temporal variation of turbulence strength in the following section. In this section, we consider turbulence strength at some instance in time, in a column of air above a telescope.

At an instant in time, we are interested in the strength of turbulence as a *function* of altitude z : we will write this as $C_n^2(z)$. This is often referred to in the literature as the *optical turbulence profile* (OTP). In the derivation of wavefront autocorrelation in the previous section, the assumption of constant turbulence strength over the thin layer would have been equivalent to a top-hat function for the OTP: zero everywhere, and with constant value C_n^2 from altitude h to $h + \Delta h$. From the definition of C_n^2 in Equation 2.18, for which the structure function on the left-hand side is dimensionless, we see that the dimensions of the OTP are $\text{length}^{-2/3}$. The total turbulence J in some column of air is calculated by integrating the OTP over the vertical extent $h_0 \rightarrow h_1$:

$$J \equiv \int_{h_0}^{h_1} C_n^2(h) dh. \quad (2.27)$$

We refer to this quantity very literally as a “turbulence integral,” and it can be calculated over any subset, or the entirety, of the atmosphere. When the integration bounds are $0 \rightarrow \infty$, this quantifies the total turbulence strength in a column of air that a light wave from a distant star will traverse before entering the telescope. The turbulence integral has dimensions of $\text{length}^{1/3}$.

How can we measure OTPs and turbulence strength parameters? Image quality measurements have been an important tool for astronomers for many decades – e.g., as part of site characterization campaigns or as calibration instrumentation for existing observatories. As we near the era of extremely large telescopes, however, there has been active development in quantifying optical turbulence at all altitudes in an attempt to find the best observing sites for these powerful instruments; the performance of the adaptive optics systems planned for the European Extremely Large Telescope⁵, for example, depends on a highly sophisticated knowledge of OTPs. In Figure 2.2 we showed very high-resolution turbulence data taken with a microthermal balloon instrument. This type of data is expensive to gather, and the balloons travel slowly – meaning the result is not a snapshot of turbulence at *one* time as light rays would see it. Below we will focus instead on two (of many) more common measurement techniques in optical astronomy: Multi Aperture Scintillation System (MASS) and stellar SCIntillation Detection And Ranging (SCIDAR) (Vernin et al., 1990; Klueckers et al., 1997; Tokovinin & Kornilov, 2002; Tokovinin et al., 2003b). Both are scintillation techniques—scintillation is the fluctuation of light intensity entering a telescope pupil. Both MASS

⁵<https://www.eso.org/sci/facilities/eelt/>

and SCIDAR leverage the fact that turbulence at different altitudes produces peaks at different spatial frequencies in the power spectrum of scintillation.

We first focus on SCIDAR, which has the capability of measuring high-resolution C_n^2 profiles. This is done by imaging two target stars with a known angular separation and triangulating the correlations observed in their scintillation patterns. Traditionally this process is performed with a single camera, but Egner & Masciadri (2007); Shepherd et al. (2014) developed a generalized method of SCIDAR using two detectors that can thus leverage spatio-temporal cross-covariance of images as well as auto-covariance, improving the resolution to roughly 250 m) (the sensitivity is altitude dependent). Note that this instrument is sensitive to weak turbulence very near the instrument noise floor, since the small covariance peaks from weak turbulence move differently in the spatio-temporal cross-correlation than does noise. This method, called stereo-SCIDAR, has been used to carry out a monitoring campaign at Cerro Paranal (in northern Chile, home of the Paranal Observatory); examples of measurements from a night of observing are shown in Figure 2.3. The variation in turbulence strength C_n^2 is shown as variations in color as a function of altitude (vertical axis) and time (horizontal axis). In the right panel, the median of these data over time is shown as a single profile vs. altitude. An important distinction in the figure is that while regions in time that are shaded black for all altitudes correspond to times when no data are available, black sections within a profile denote *low* turbulence strength as indicated by the color scale at the bottom of the figure. We can learn two things from the depicted distribution of turbulence: first, profiles (a vertical slice through the data) have stratified regions of high and low turbulence, where the strength varies between the regions by 1 to 2 orders of magnitude. Second, there is temporal variation, over timescales on the order of a few hours, in the number, location, and strength of these turbulent regions. Both these points will be very relevant in Section 3.1.1 when we discuss simulations of atmospheric PSFs; the first point lends credibility to approximating the atmosphere as discrete layers of turbulence and the second motivates us to use OTP models that capture these variations.

Given the bands of persistent turbulence observed in the profiles in Figure 2.3, we are motivated to quantify the turbulence in vertical sections of the atmosphere with the turbulence integral defined in Equation 2.27. Since this integral is linear in altitude, calculating the integral over ranges of altitude $h_i \rightarrow h_{i+1}$, for n layers, $i = 1, \dots, n$, leads to a set of values

$$J_i = \int_{h_i}^{h_{i+1}} C_n^2(h) dh$$

that contribute to the total turbulence integral $J = \sum_i J_i$. This separation will be very useful for quantifying turbulence contributions from different regions in the atmosphere – for instance, it is common to distinguish between the *ground layer* (GL), which comprises the first ≈ 1 km of the atmosphere above ground, and the *free atmosphere* (FA), which is everything above the ground layer. This distinction separates turbulence in the ground layer, which is in general dominated by

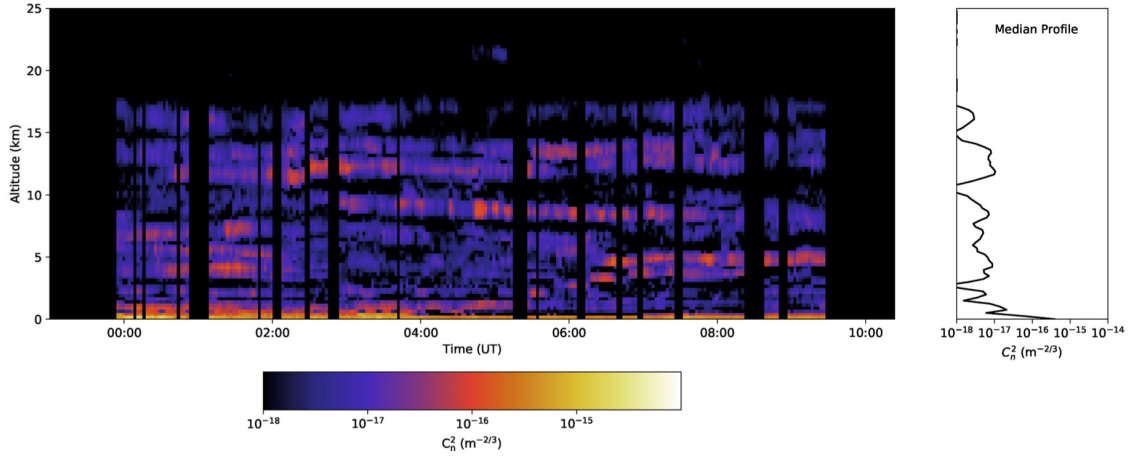


Figure 2.3: Stereo-SCIDAR measurements of optical turbulence profiles for the night beginning March 8, 2017, at Cerro Paranal, Chile. Black areas in the figure denote *different* aspects of the data: vertical black bars denote regions of missing data whereas black regions within a profile correspond to very low turbulence, according to the color bar. Figure adapted from Osborn et al. (2018).

turbulence around local structure and topology, from the “free” turbulence generated by convection and wind gradients. In this thesis, we will take the boundary between these atmospheric regions to be at 800 m above ground level.

Turbulence strengths for vertical sections of the atmosphere – i.e., the values of J_i – can be measured by a multi-aperture scintillation system (MASS), which, like SCIDAR, uses scintillation to quantify optical turbulence structure in the atmosphere. Details on the instrumentation and data processing can be found in Tokovinin et al. (2003a), and is also reviewed in Section 2.1 of Ogane et al. (2021). Single stars are imaged with a MASS and their intensity fluctuations due to scintillation are measured in four concentric annular apertures. The variance of the observed intensity in each annulus, and the covariances between annuli, can be related to the spatial power spectrum of intensity fluctuations, from which the turbulence strength of multiple layers can be estimated (see equations 6 through 12 in Ogane et al. (2021)). The number and location of layers estimated is set by the number and size of the concentric apertures – a typical MASS instrument measures 6 turbulence layers at 0.5, 1, 2, 4, 8, and 16 km above ground ⁶. The MASS instrument response functions are nearly triangular as a function of altitude, with overlap in intermediate regions such that turbulence at 6 km, for example, contributes equally to the layers centered at 4 km and 8 km.

An example of MASS turbulence strength measurements from a night of observation at Cerro Telolo is shown in panel (a) of Figure 2.4. The turbulence strength is represented by vertical bars

⁶The impact of turbulence on the scintillation power amplitude goes to zero as the altitude of turbulence goes to zero, so MASS measurements are not sensitive to turbulence near the ground

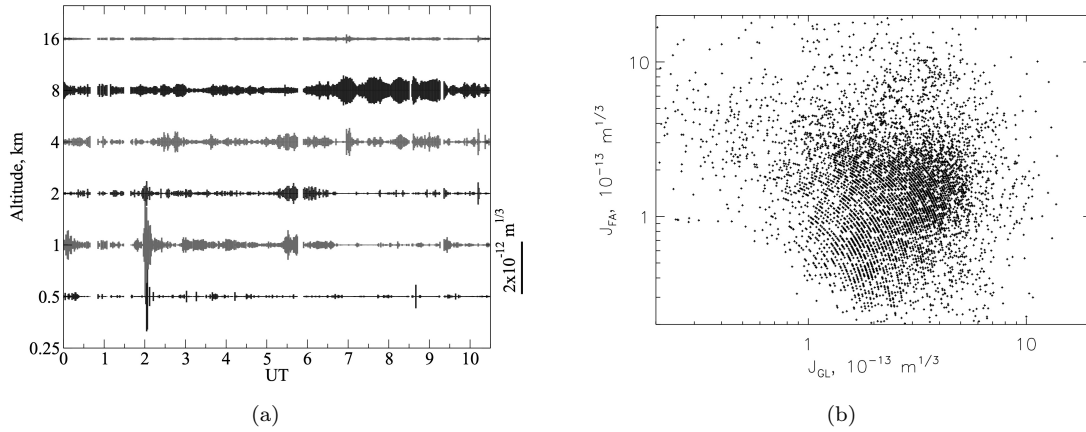


Figure 2.4: (a) MASS turbulence profiles for the night of July 25-26, 2002, at Cerro Telolo. Turbulence integrals J_i as a function of time (in hours) for each atmospheric layer are displayed as vertical bars, with the scale bar showing a reference intensity of $2 \times 10^{-12} \text{ m}^{1/3}$ (Tokovinin et al., 2003a). (b) Over 6000 values of turbulence integrals J_{GL} , J_{FA} for the ground layer (0-500 m) and J_{FA} free atmosphere (everything above 500 m) measured with MASS-DIMM and MASS, respectively, at Cerro Pachón. Data were collected in 2005 and are described in Tokovinin & Travouillon (2005).

with length proportional to J_i at each of the six MASS altitudes and displayed as a function of time; a MASS data accumulation is only a minute long so the temporal resolution is high. It is worthwhile noting that although the difference in altitude for each pair of adjacent layers, and thus the integration range for each J_i , is doubling with altitude, the altitude dependence of the integrated turbulence strength is in general not doubling and is highly variable even over the course of a single night.

As described above, MASS measurements are insensitive to turbulence near the ground, but provide complete information about the free atmosphere turbulence. To measure turbulence over a more complete range, a MASS is often combined with a differential image motion monitor (DIMM), an instrument that is sensitive to total seeing by measuring wavefront slope differences over two small apertures some distance apart (Sarazin & Roddier, 1990; Kornilov et al., 2007). Ground-layer turbulence integral values are thus calculated by subtracting the MASS free atmosphere values from coincident DIMM full-seeing measurements. Panel (b) in Figure 2.4 compares MASS measurements of free atmosphere turbulence strength (J_{FA} , vertical axis) with ground layer turbulence (J_{GL} , horizontal axis) obtained from the MASS-DIMM difference. At this particular site – the Gemini-South observatory at Cerro Pachón, Chile – the turbulence integrals in the ground layer and the free atmosphere are of comparable magnitude on average.

We are now ready to revisit some definitions from Section 2.3 with these new ideas in mind. For instance, instead of the simplified one-layer definition presented in Equation 2.21, the Fried parameter as defined in Fried (1965) includes the altitude variation of optical turbulence over the

full atmosphere:

$$r_0 = \left(2.914 k^2 \sec \zeta \int_0^\infty C_n^2(h) dh \right)^{-3/5}. \quad (2.28)$$

Typical values of C_n^2 at modern observatories range from 10^{-18} to $10^{-14} \text{ m}^{-2/3}$ (see Figure 2.3) and values of J range from 10^{-13} to $10^{-12} \text{ m}^{1/3}$ for both the GL and FA (see Figure 2.4). For an observation at zenith ($\sec \zeta = 1$) at $\lambda = 500 \text{ nm}$, typical values of the Fried parameter r_0 are 10 cm and 2.5 cm for turbulence integrals of $10^{-13} \text{ m}^{1/3}$ and $10^{-12} \text{ m}^{1/3}$, respectively.

How do these relate to parameters we are used to thinking about as observational cosmologists/astronomers? Image quality, commonly referred to as *seeing*, is usually quantified by the full-width at half-maximum value (FWHM) of the PSF in units of arcsec. The theoretical prediction for the FWHM of the long-exposure PSF for pure Kolmogorov turbulence is, in radians,

$$\epsilon_0 = 0.98 \frac{\lambda}{r_0}. \quad (2.29)$$

For the above range of Fried parameters (and assuming Kolmogorov turbulence), we predict seeing values of 1–4 arcsec. This theoretical seeing prediction overestimates the PSF size because pure Kolmogorov turbulence allows power at unphysically large spatial scales — as discussed in Section 2.3.1, the presence of an outer scale reduces the PSF size.

This size variation is a function of wavelength — though Equation 2.26 has no explicit λ dependence, r_0 does vary with λ as $\lambda^{-6/5}$ (see Equation 2.28). For Kolmogorov seeing, combining this with Equation 2.29 gives a wavelength dependence of PSF size of $\lambda^{-1/5}$. The power law exponent b for von Kármán turbulence has been estimated to be somewhere between -0.2 (pure Kolmogorov) and -0.5 (Xin et al., 2018; Carlsten et al., 2018). This exponent is also expected to vary with outer scale, where a smaller outer scale leads to a larger (in magnitude) value of b .

2.3.3 Horizontal structure of turbulence

In the previous section, we considered the effects on image quality of the variation in optical turbulence strength along the altitude axis above the telescope. Turbulence strength also varies along directions perpendicular to the vertical axis, so, when we extend beyond a single star to imaging an entire field of view, we must consider additional effects that we previously ignored. To explore these nuances, let us first explore in more detail the light that reaches us from a single star, lying infinitely far away (compared to the pupil diameter of the telescope) along the optical axis of the telescope. The atmosphere responsible for degrading the PSF of a star is the column of air along this optical axis, of diameter roughly equal to the telescope pupil diameter D . This is illustrated in panel (a) of Figure 2.5: the solid orange lines depict the limits of the volume of atmosphere through which the wavefront from the star travels (i.e., the section of the wavefront that will enter the telescope pupil) — since stars are effectively at infinity, the wavefronts are planar and the path

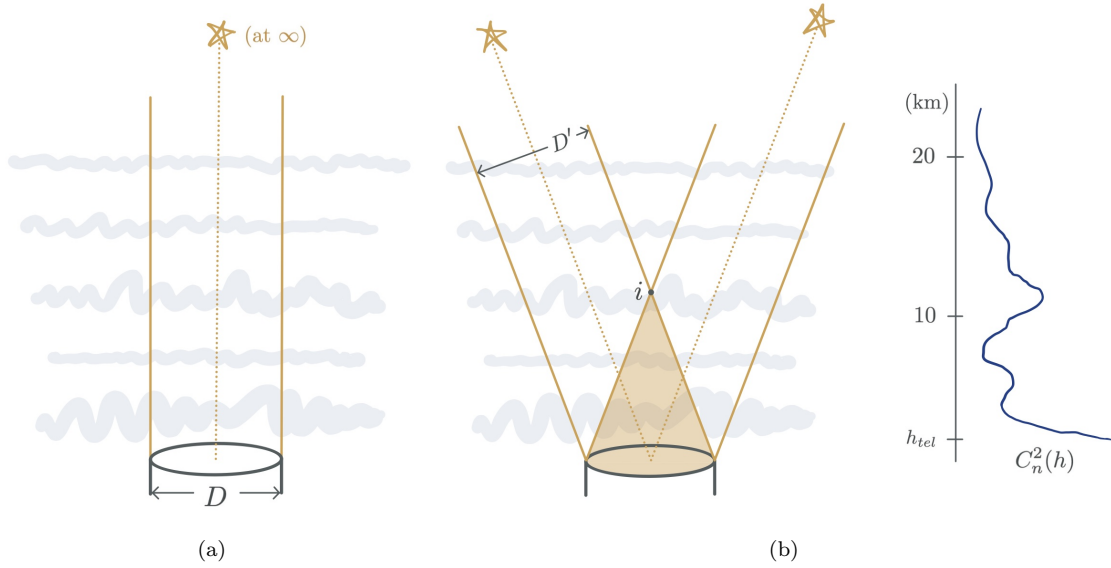


Figure 2.5: Schematic of imaging through a turbulent atmosphere, for (a) a single star and (b) two off-axis stars. The trajectory of the observed stellar wavefronts are shown in solid orange, with D, D' the diameters of the telescope pupil and atmospheric column, respectively. Turbulence is represented by thick wiggles that are roughly consistent with the cartoon $C_n^2(h)$ profile shown on the right – note that the horizontal and vertical scales in this schematic differ by a factor of $\sim 10^3$. The shaded region denotes the volume of atmosphere that is common to both wavefronts and i marks the highest point of overlap.

through the atmosphere is a column. If we assume that turbulence strength is relatively constant over a telescope diameter, we can use a single $C_n^2(h)$ profile to describe the variation with altitude of turbulence strength within the column of air for a single star.

However, the real imaging situation we face with Rubin Observatory, or any wide-field telescope, is more complex. Realistically, Rubin will image hundreds of stars in a single exposure, but for the sake of illustration, two will suffice; see panel (b) of Figure 2.5. Light from stars that are off the telescope axis still travels through a column of turbulent atmosphere before reaching the pupil – but the column associated with each star now has a diameter (D' , in the figure) that scales as the cosine of the star’s angle off-axis⁷. Two such columns will then pass through different turbulence – overlapping near the ground, as shown by the shaded region in Figure 2.5, but diverging at high altitudes.

In the upper atmosphere (altitude ~ 20 km), the two columns are separated by a distance of $\sim 20 \text{ km} \times 2 \tan(\frac{\theta_{\text{FoV}}}{2})$. For Rubin’s field of view of $\theta_{\text{FoV}} = 3.5^\circ$, this comes out to roughly 1 km, or two orders of magnitude larger than the diameter of the columns themselves. We assume, in the

⁷Note that the angle between the star and the telescope axis in the figure are not to scale because of the difference in the horizontal and vertical scales. For reference, the field of view for the Rubin telescope is 3.5° across.

case of a single star, that turbulence strength is relatively constant over distances $\sim D$ – but it is not clear whether this holds for distances $\sim 100 \times D$. So, in the context of wide field imaging, we understand references to quantities such as $C_n^2(h)$ and r_0 as being some kind of ensemble average over the large angular extent we image in each exposure. In fact, we have assumed this implicitly when including only a single profile in Figure 2.5.

Returning to the lower atmosphere, the region of overlapping turbulence for two stars elucidates how correlations in the PSF across the field of view arise. The overlap in turbulence sampled by each wavefront depends on the separation angle between the stars. For stars at opposite sides of the Rubin field of view, only a fraction of the ground layer turbulence (between the telescope and ~ 800 m) is in common since the highest point of overlap between the columns occurs at an altitude of $i = \frac{D}{2 \tan(3.5^\circ/2)} \sim 130$ m. Although the overlap in volume is small for these particular stars, two other factors also lead to correlations. First, the majority of stars have a separation angle much less than 3.5° , and thus a much higher overlap altitude i . Secondly, up until now we have treated the atmosphere as static, which it is absolutely not expected to be. Accounting for motion of the air will *increase* the effective overlap of turbulence in different columns, as wind tends to move turbulence horizontally, thus creating more correlation than one would expect from a static view. For example, with an exposure time of 30 s and a wind speed of 10 m/s, which is a typical speed a few kilometers above the telescope, the wind travels a distance ~ 3 times greater than the column separation at that altitude.

2.4 Outlook

Though we understand the phenomenological reasons for correlations in wavefronts, and thus PSFs, of stars at different locations on the field of view, there is limited literature that extends the formalism covered in this chapter to describing those correlations analytically. To make predictions of this kind, we turn to simulations, which will be covered in Section 3.1.1 of this thesis. The literature describing temporal variation of astronomical wavefronts is likewise limited, consisting mostly of narrow-field adaptive optics studies and measurements of the atmospheric “coherence time” τ_0 , the time taken for an r_0 sized patch of turbulence to drift a distance equal to its own size. This parameter is most relevant for setting wavefront sensor and deformable mirror sampling rates in order to achieve the best image quality on individual exposures. In the context of wide field observations for weak gravitational lensing, for which we have more stringent requirements on overall *bias* than exposure-by-exposure noise, a more relevant question is whether spatial correlations persist on time scales that could imprint a bias in shape measurements. We will explore this question in different ways throughout this thesis, using short-exposure data and high-fidelity simulations.

Chapter 3

Simulating realistic atmospheres: psf-weather-station

3.1 Motivation and context

With the goal of improving the fidelity of image simulations for ground-based instruments, we have produced a public software package¹ (PSF-WEATHER-STATION) that provides realistic site-specific environmental parameters with which to simulate atmospheric PSFs. An important scientific goal in simulating atmospheric PSFs at Rubin Observatory, or other wide-field instruments, is to predict how weather in the observatory environment impacts correlations in PSF parameters across the focal plane.

As we have learned in Chapter 2, the atmosphere is correlated via the von Kármán power spectrum described in Equation 2.25, and these spatial correlations translate into angular correlations in the size and shape of the atmospheric PSF in the associated exposure. Wind over the telescope plays an integral role in this process, as it moves correlated patches of turbulence through the atmospheric columns that impact the images of different objects, leading to correlations on angular scales larger than the patches. Although wind varies with altitude, if wind directions are consistent across altitudes, turbulence at all altitudes will move coherently, thereby imprinting a stronger correlation in the PSF than when wind directions at different altitudes are uncorrelated. Another relevant factor for PSF correlations is the altitude-dependence of the optical turbulence profile (OTP), which describes the contribution of each layer to the total turbulence strength. Interestingly, one of the drivers of atmospheric turbulence is wind itself – specifically, wind shear – so we expect that these two factors that influence spatial correlations in PSF parameters are not independent.

Therefore, it is important to include in the simulations altitude-dependent correlations among

¹<https://github.com/LSSTDESC/psf-weather-station>

wind speed, wind direction, and turbulence, as well as realistic temporal variations. The PSF-WEATHER-STATION package leverages local weather measurements and data products from global weather models to produce turbulence parameters and wind speeds and directions that are realistically correlated with altitude and in time. After reviewing the literature on simulated atmospheric PSFs in the rest of this section, we describe the data and models used within PSF-WEATHER-STATION in Section 3.2, and outline the range of package outputs in Section 3.3.

3.1.1 How to simulate an atmospheric PSF

Since it is computationally intractable to simulate atmospheric PSFs by calculating the trajectory of each photon through a turbulent (i.e., chaotic) 3D volume of atmosphere, approximate methods have been developed. As discussed in Section 2.3.2, measurements of optical turbulence with SCIntillation Detection And Ranging (SCIDAR) instruments show that the atmosphere is often stratified into regions of stronger turbulence separated in altitude by areas of relative calm—see in particular Figure 2.3 (Osborn et al., 2018; Osborn & Sarazin, 2018). Typically each extending only ~ 1 km vertically within the ~ 20 km column, these layers of stronger turbulence dominate the atmospheric contribution to the PSF. These observations motivate a simplified model of the atmosphere that consists of only 2-dimensional phase screens; the relative wavefront phase across such a plane encodes the impact of diffraction due to variations in index of refraction within a layer of atmosphere. A schematic of such an atmospheric simulation, with only two phase screens, is depicted in Figure 3.1.

The refractive index variations within each phase screen are a realization of von Kármán turbulence. We assume Taylor’s frozen flow hypothesis (Taylor, 1938), in which the time scales for changes in turbulence are longer than those for changes due to phase-screen drift from wind. Under this assumption, it is not necessary to evolve the turbulence structure during a simulated exposure; the variation in phase across the screen is “frozen” during an exposure, but the screen moves across the field of view to simulate the impact of wind. Each phase screen is assigned a “wind” speed and direction; for each time step Δt of the simulation, the phase screens are translated accordingly. The altitude dependence of wind speed and direction and of turbulence strength is modeled by including multiple screens, as illustrated in Figure 3.1. After each time step, the phase screen contributions within the atmospheric column (for each star in the field) are summed vertically. For each star, the sum is then Fourier transformed to focal plane coordinates and, after all time steps, added together to form the simulated image – i.e., the PSF. As a function of image coordinates θ_x, θ_y , the PSF has an intensity profile

$$I(\theta_x, \theta_y) \propto \sum_{\Delta t} \left| \mathcal{F} \left\{ P(u, v) e^{-\frac{2\pi i}{\lambda} W(u, v, \Delta t)} \right\} \right|^2, \quad (3.1)$$

where $P(u, v)$ is the aperture transfer function and $W(u, v, \Delta t)$ the wavefront, both in terms of pupil coordinates u, v . The sum over Δt represents the sum over all simulation time steps during an exposure.

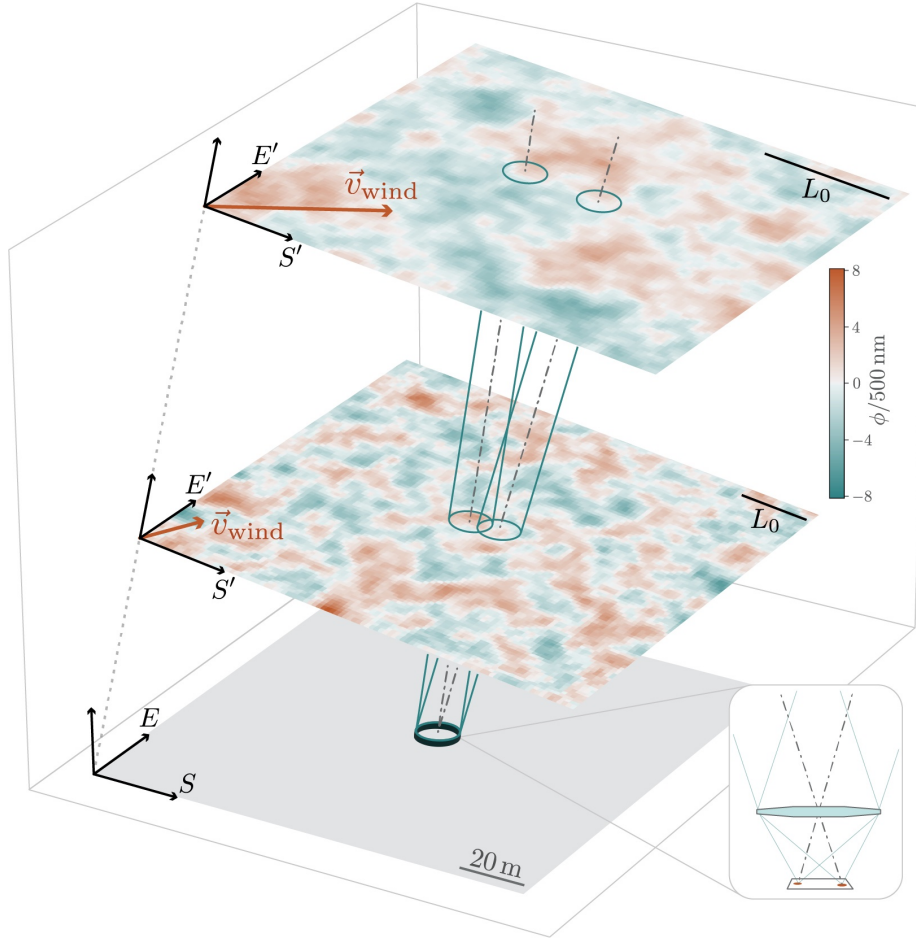


Figure 3.1: This schematic illustrates the simplified view of the atmosphere used for PSF simulations based on discrete phase screens. Lines of sight (dot-dashed grey) for two stars (whose images are located at extrema of the field of view) pass through two phase screens of von Kármán refractive index variations, each with different values for the outer scale L_0 . The phase offset incurred by light passing through each point on each screen is indicated by the color scale in units of wavelength. The columns (teal) associated with each line of sight show the path of starlight that will reach the telescope aperture (black), along with the relevant phase screen area. The wind vectors (orange arrows) show speed and direction of the wind in the plane of the screen. The primed coordinate systems are perpendicular to the telescope axis, and are related to the ground coordinate system via the altitude and azimuth of the pointing.

The phase screen is the building block of the simplified atmospheric model described above. Because turbulence integrals J_i add linearly, each phase screen contributes to the total turbulence with weight $w_i = J_i(\sum_i J_i)^{-1}$. Although one could use the turbulence integral J_i to generate the phase pattern for screen i , it is more natural to use the Fried parameter r_0 because r_0 determines the turbulence power spectrum amplitude. Given that $r_0 \propto J^{-3/5}$ (Equation 2.28), the contribution of the i th screen is $r_{0,i} = w_i^{-3/5} r_0$. By convention, r_0 is specified at $\lambda = 500$ nm and zenith angle $\zeta = 0$.

In summary, the input parameters for a simulation of PSFs across a single exposure are the outer scale L_0 , the atmospheric seeing parameterized by r_0 , the number of phase screens and their altitudes, the wind speed and direction for each screen, and the fractional contribution w_i of each screen to the total turbulence.

3.1.2 Earlier work on atmospheric PSF simulations

High-fidelity atmospheric PSF simulations have played an important role in astronomy – from the development of adaptive optics (Jolissaint, 2010; Madurowicz et al., 2018; Srinath et al., 2015) to the optimization of instrumentation and software for precision cosmology, including Rubin Observatory and the associated survey, LSST (Jee & Tyson, 2011; Chang et al., 2012; Peterson et al., 2015; The LSST Dark Energy Science Collaboration et al., 2021b). All the studies referenced here used the general approximation method of moving frozen-flow phase screens for simulating atmospheric PSFs as described above.

In preparation for analysis of data from Rubin Observatory, the LSST Dark Energy Science Collaboration (DESC) produced an extensive set of image simulations, called Data Challenge 2 (DC2) (The LSST Dark Energy Science Collaboration et al., 2021b). The atmospheric PSFs in these images were initialized as follows: wind speeds and directions were drawn randomly from uniform distributions (ranging 0 – 20 m/s and 0 – 360 deg respectively), and turbulence weights for each phase screen were drawn as small ($\sim 10\%$) uncorrelated Gaussian deviations from the set of fixed empirical turbulence parameters presented in Ellerbroek (2002). These relative weights for six discrete turbulence layers are based on both SCIDAR and balloon (thermosonde) measurements conducted for the 1998 Gemini site-testing campaign at Cerro Pachón in Chile (≈ 2 km from Rubin Observatory). The drawback of the Ellerbroek (2002) model is that while it may represent median conditions it contains none of the variability known to be present in the OTP at Cerro Pachón (Tokovinin & Travouillon, 2005), including variations in the altitude of turbulence. Therefore, the primary limitation of the turbulence inputs to the DC2 simulations is the unrealistic distributions over time (see Section 3.2); in addition, the wind parameters lack correlation across altitude.

A second image simulation suite, described in Peterson et al. (2015), draws wind speeds and directions, and turbulence strengths from historical data for a seven-layer atmospheric model. The wind parameters for each phase screen were drawn from Weibull distributions fit to data from the

National Oceanic and Atmospheric Administration National Centers for Environmental Prediction² (NOAA NCEP). The turbulence weights were drawn from lognormal distributions based on published quantiles of the turbulence integral at Cerro Pachón from Tokovinin & Travouillon (2005), which used the Gemini site-testing data (also used in Ellerbroek (2002)) augmented with more recent MASS-DIMM measurements. These quantiles are given for seven discrete altitudes, with evidence of independence between values at the ground versus in the higher atmosphere. Citing evidence in Tokovinin & Travouillon (2005) of high correlation among the high-altitude layers, however, Peterson et al. (2015) uses a single turbulence weight for the phase screens at higher (fixed) altitudes, as well as a single (different) value for the ground layer. While these parameters will more accurately capture any dominant weather conditions at the site than the simulations used in The LSST Dark Energy Science Collaboration et al. (2021b), Peterson et al. (2015) do not account for the correlations in altitude of a single parameter at any given time, nor the relationship between wind and turbulence strength in the free atmosphere.

In the rest of this chapter we present improvements on setting parameters for input to atmospheric PSF simulations, focusing on including important correlations (in time, or with altitude) among parameters. Our goal is to parametrize atmospheric conditions that, in an ensemble sense, are representative of a particular site and can be used to generate atmospheric PSFs across a wide field of view.

3.2 Inputs to psf-weather-station

PSF-WEATHER-STATION draws on several sources of data and physically motivated models to produce parameters.

In what follows, we separate the atmosphere into two regions based on the characteristic turbulence regime for those altitudes. The ground layer (GL) is typically defined as the region between ground level and 500 – 1000 m above the telescope, where complex topography and heat sources generate non-Kolmogorov eddies. The free atmosphere (FA) is defined as the region above the ground layer, where turbulence is generally well-described by Kolmogorov statistics. This separation into GL and FA plays an important role in many design choices for PSF-WEATHER-STATION.

3.2.1 Data products from General Circulation Models

The primary sources of data for PSF-WEATHER-STATION are data products from global weather forecasting organizations such as the European Centre for Medium-Range Weather Forecasts³ (ECMWF) and the National Oceanic and Atmospheric Administration National Centers for Environmental Prediction⁴ (NOAA NCEP). The global circulation models (GCM) used in weather forecasting cover the

²<https://www.ncei.noaa.gov/products/weather-climate-models/global-forecast>

³<https://www.ecmwf.int/en/forecasts/datasets/>

⁴https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php

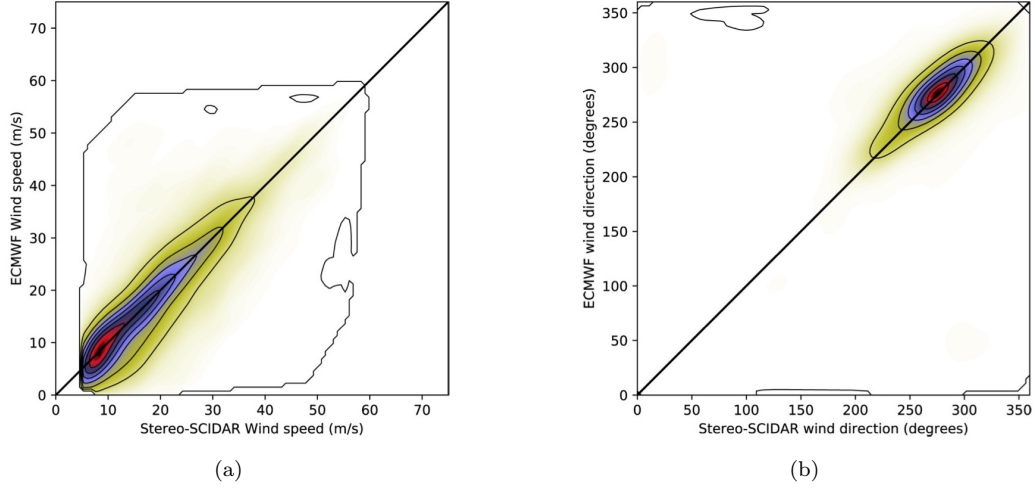


Figure 3.2: Comparisons of Stereo-SCIDAR recovered wind parameters with the forecast wind from ECMWF ERA5 at Cerro Paranal. (a) Wind speed – correlation coefficient: 0.82; average difference: 0.41 m/s; root-mean-square difference: 6.44 m/s. (b) Wind direction – correlation coefficient: 0.77; average difference: 1.59 deg; root-mean-square difference: 27.09 deg. Adapted from Osborn et al. (2018).

entire globe on a grid of 0.25-2.5 deg resolution and output predictions for dozens of environmental parameters at a number of different altitudes between 0 and ~ 80 km above sea level.

Although any of the available GCM data sets can be used in PSF-WEATHER-STATION, the results in the analyses described in this thesis are based on data from the ECMWF Reanalysis v5 (ERA5) catalog; these are hourly estimates of global weather conditions based on meteorological measurements assimilated into the GCM. ERA5 was chosen for its dense coverage in time – hourly – which is useful for sampling conditions through the night, and in altitude – 137 levels – which is important for capturing vertical wind gradients in the atmosphere. In Appendix A, we provide more details on the ECMWF and NOAA NCEP models and our reasons for choosing the ECMWF dataset.

The GCMs provide robust estimates of wind and temperature throughout the free atmosphere as a function of time. For instance, Figure 3.2 shows a comparison of (a) wind speed and (b) wind direction above Cerro Paranal (published in Osborn et al. (2018)) for the ERA5 data set and Stereo-SCIDAR measurements (see Section 2.3.2 for details). The datasets are highly correlated (correlation coefficients ~ 0.8) and the difference between corresponding temporal datapoints is low (see figure caption for numerical values), inspiring confidence in the fidelity of ERA5 outputs for use in astronomy.

The wind direction in Figure 3.2 is defined according to the meteorologic convention—namely, the direction *from* which the wind is blowing—and the angle is defined in degrees East of North. In this thesis we adopt the same convention.

In this thesis we will use ECMWF ERA5 data from the ERA5 (horizontal) grid point closest to Rubin Observatory (~ 1.9 km away).⁵ In Figure 3.3 we show the location of Cerro Pachón relative to the Chilean coast and the city of La Serena, as well as a zoomed-in view of the mountaintop. Rubin Observatory and Gemini South are visible from the satellite photo, and their locations are indicated by pins (Rubin Observatory is indicated by the LSST pin, Gemini South by the Cerro Pachón pin). The ERA5 grid point closest to Rubin is labeled by the yellow pin.

In Figure 3.4 we show the distribution of wind and temperature information from ERA5 for a 6-month period (May through October 2019) at the grid point discussed above. The ERA5 altitude profiles are sampled from this point at 00h, 02h, and 04h local time. We see clear persistent trends in weather conditions. The wind speeds start low near the ground (indicated by the orange dashed line) and increase with altitude up to ≈ 12 km (the jet stream altitude) before decreasing at higher altitudes. At all altitudes (above ≈ 5 km), there is a very dominant wind direction of ≈ 280 deg E of N (or ≈ 80 deg W of N), corresponding to winds coming from the ocean towards Cerro Pachón. Finally, temperature decreases steadily between the ground and ≈ 17 km (the troposphere) and then increases slightly with altitude in the stratosphere.

These are extremely useful data for describing the free atmosphere. However, interactions of the atmosphere near the ground are not accurately captured with GCMs because topographical features are not modeled at scales smaller than ≈ 1 km. Therefore, the accuracy of the GCM data (both initial conditions and predictions) is limited near the ground⁶. An additional complication for our particular case is the difference in the elevations of ground level at the ERA5 grid point (1.75 km) and Rubin Observatory (2.7 km)—note the sharp drop in terrain between Rubin and the grid point in Figure 3.3 (b).

These issues are illustrated in Figure 3.5, in which we compare measurements from the Gemini South weather tower to ERA5 outputs both at the grid point elevation and the altitude closest to the Gemini South elevation (≈ 2.74 km). As shown in the upper panels of Figure 3.5, the wind speed values for ERA5 at 2.76 km (purple histogram and points) are correlated with, and have a similar distribution to, the weather-tower measurements (orange histogram). The ERA5 values at 1.75 km, however, are peaked at much lower speeds. The lower panels of Figure 3.5 show that the wind direction for the lower ERA5 altitude (teal histogram and points) peaks just above 90 deg, while the weather-tower wind directions peak around 300 deg. With increasing altitude for ERA5 data, the difference with the weather-tower wind direction decreases; however, even for ERA5 data at ≈ 2.7 km, the wind direction peaks near North—approximately 60 deg away from the wind direction measured at the weather-tower. From Figure 3.4 we see the wind directions maintain this northerly skew until around 2 km above the elevation of the observatory.

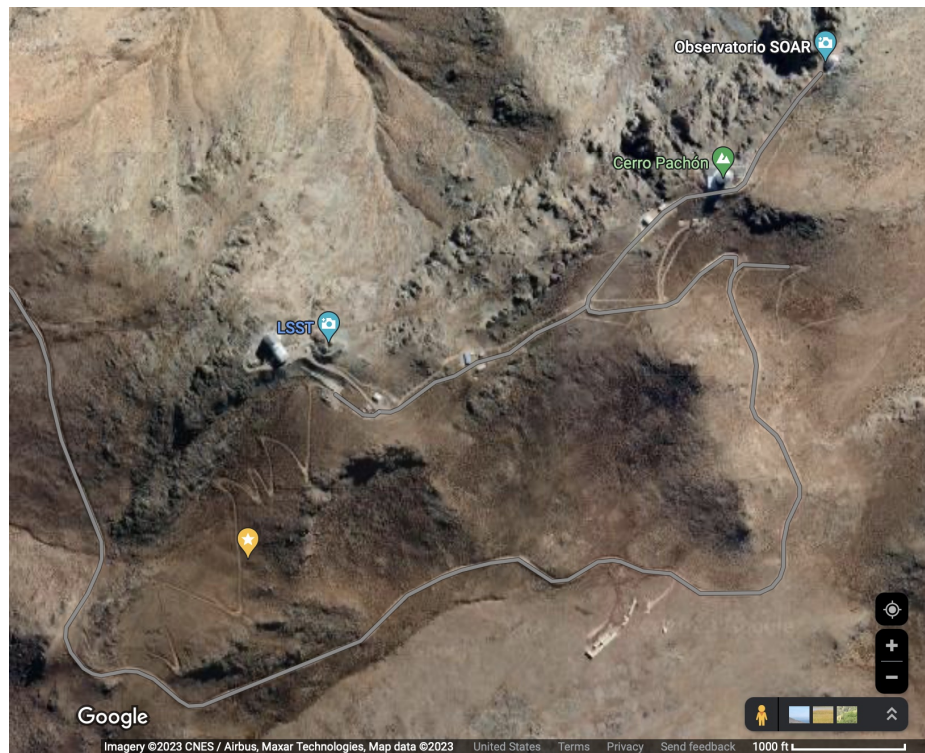
In PSF-WEATHER-STATION, we overcome this limitation in ECMWF data by using weather-tower

⁵Cerro Pachón is at longitude, latitude of -30.241 , -70.737 while the closest grid point is at -30.25 , -70.75 .

⁶In addition, because the grid point closest to Cerro Pachón is ≈ 1.9 km away, the topography, even if modeled with sufficient resolution, would not be representative of the observatory surroundings.



(a)



(b)

Figure 3.3: (a) Location of Cerro Pachón relative to the Chilean coast, and (b) location of Rubin Observatory (marked by the “LSST” pin), Gemini South (under the green Cerro Pachón pin), and the ECMWF ERA5 grid point closest to Rubin (yellow star pin). Image credit: Google Maps.

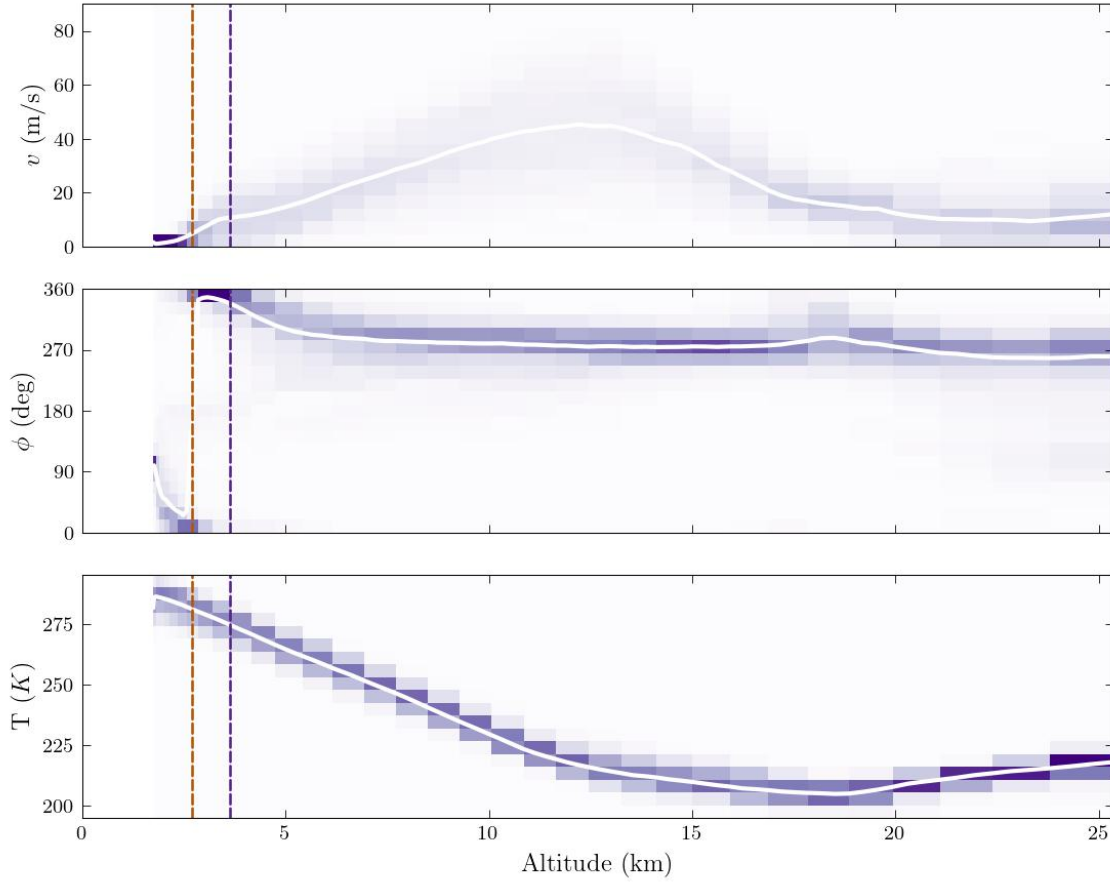


Figure 3.4: Distribution of ECMWF ERA5 values, over 6 months, of wind speed v (top), wind direction ϕ (middle; measured in degrees East of North), and temperature T (bottom) for a grid point ~ 1.9 km south-west of Cerro Pachón. The purple shading represents the density of data values. The median values, as a function of altitude above sea level, are shown by white solid curves. The vertical dashed lines correspond to the altitude of Cerro Pachón (orange) and the beginning of the free atmosphere (purple), as we have defined it at 800 m above ground. The bins in altitude are spaced such that each range contains ECMWF data at three altitudes, which are averaged.

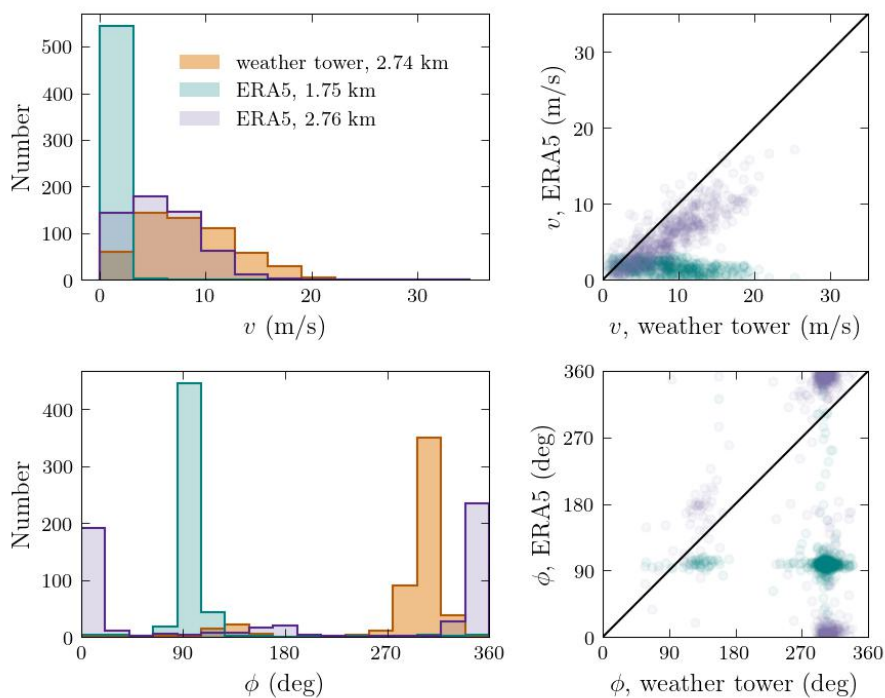


Figure 3.5: Distributions of wind speed v (top) and wind direction ϕ (bottom) as measured by the Gemini South observatory weather tower (2.74 km, orange shaded histograms) and predicted with ECMWF ERA5 at the grid point elevation (1.75 km, teal shaded histograms), and at the altitude closest to the observatory elevation (2.76 km, purple shaded histograms). The panels on the right show ERA5 values (teal and purple points corresponding to 1.75 km and 2.76 km, respectively) versus weather-tower measurements.

measurements rather than GCM data for the ground layer. Since weather-tower data are typically recorded every few minutes, their sampling times can be matched to the GCM data used for the free atmosphere (the matching is described in more detail in Section 3.2.3). The weather-tower measurements are optional but highly recommended inputs to PSF-WEATHER-STATION since the GL is of significant importance for the PSF, contributing (40 to 60)% of the turbulence at many observatories (Tokovinin & Travouillon, 2005; Tokovinin et al., 2005, 2003a).

3.2.2 Empirical model of optical turbulence profiles

There exist non-hydrostatic atmospheric models with sub-kilometer horizontal resolution that are successful in simulating optical turbulence around observatories (Masciadri et al., 2001, 2017); however, such models are computationally prohibitive when many realizations need to be simulated. On the other hand, useful parameterizations of optical turbulence as a function of environmental quantities – such as temperature, pressure, wind, and kinetic energy – can be adapted from this literature, as is done in (Osborn et al., 2018) (hereafter OS18). In particular, in OS18, wind shear is assumed to be the sole contributor to the kinetic energy term (i.e., wind shear is the source of turbulent mixing of air) and the temperature and pressure profiles from mesoscale simulations are replaced with GCM data. In exchange for coarser resolution and more limited accuracy, with minimal computational time the OS18 empirical model produces estimates of optical turbulence profile $C_n^2(h)$ as a function of altitude h , which, as shown in Figure 3.6 (adapted from OS18), are broadly consistent with stereo-SCIDAR measurements at Cerro Paranal.

With the above motivation, we now briefly summarize the OS18 model. Following Masciadri et al. (2001, 2017); Tatarskii (1971), OS18 finds the following expression for the OTP:

$$C_n^2(z) = k \left(8 \times 10^{-5} \frac{P(z)}{T(z)\theta(z)} \frac{\partial \theta(z)}{\partial z} \right)^2 L(z)^{4/3}, \quad (3.2)$$

for altitude z , coefficient k , temperature T , potential temperature θ , and a length scale $L(z)$. The potential temperature $\theta(z)$ is a quantity used in atmospheric science that describes the temperature a parcel of air would have if adiabatically brought to a pressure $P_0 = 1000$ mbar (standard pressure):

$$\theta(z) = T \left(\frac{P_0}{P} \right)^{R/c_p}, \quad (3.3)$$

where $R/c_p = 0.286$ is the ratio of gas constant to specific heat for air. $\theta(z)$ is of dynamical interest because it does not change in the absence of heating, cooling, evaporation, or condensation, whereas normal temperature is affected by, e.g., flow over obstacles.

The length scale $L(z)$ in Equation 3.2 is the scale of largest energy input to turbulent flow. It

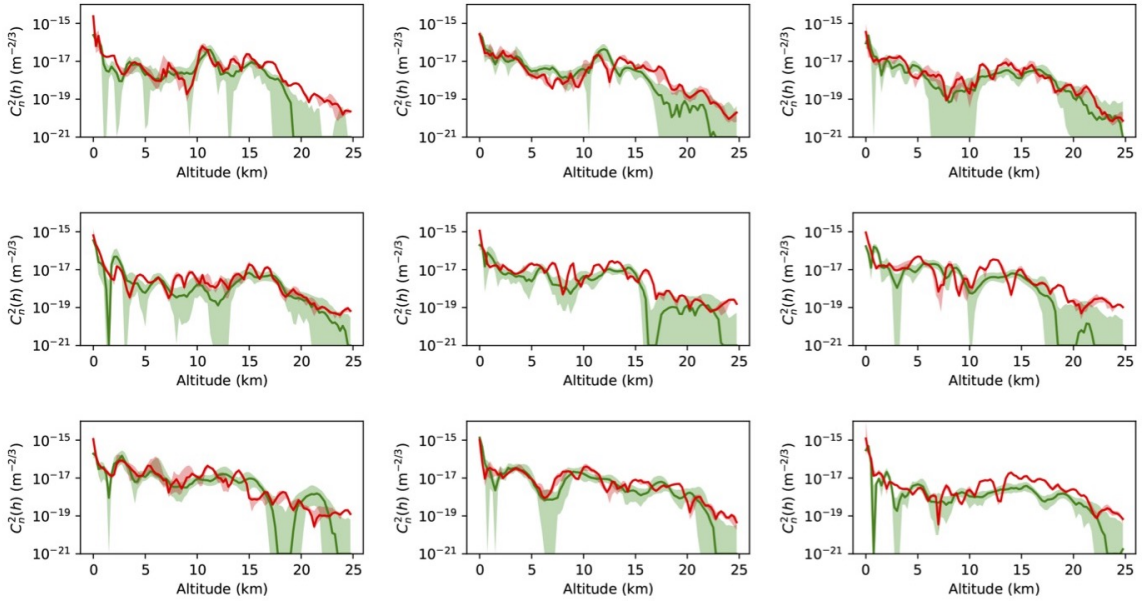


Figure 3.6: Example of vertical optical turbulence profiles at Cerro Paranal as measured by stereo-SCIDAR (green) and estimated with the OS18 model using ECMWF outputs (red). Altitude in these panels is measured relative to the elevation of the observatory. The profiles shown are the median for each night of observation, for observing times beginning on April 26 to 29 and July 22 to 26, 2016. The shaded regions correspond to the 25th to 50th percentile range. Adapted from OS18.

depends on turbulent kinetic energy E as

$$L(z) = \sqrt{\frac{2E\theta(z)}{g\partial\theta(z)/\partial z}}. \quad (3.4)$$

OS18 approximates this kinetic energy E as proportional⁷ to the vertical wind shear $\frac{\partial v}{\partial z}$:

$$E \propto \sqrt{\left(\frac{\partial u}{\partial z}\right)^2 + \left(\frac{\partial v}{\partial z}\right)^2}, \quad (3.5)$$

where u, v are the North and East components of wind velocity. OS18 do not motivate this choice explicitly, but one could expect the strength of wind shear to set the distance two blobs of turbulence in adjacent altitudes travel from each other horizontally over a given time, which is some measure of a maximum correlation length $L(z)$ for turbulence.

The OS18 final result (Equation 3.2) depends, intentionally, on quantities readily available from GCM data products: wind velocity (u, v) , pressure P , and temperature T . In Figure 3.7 we show the intermediate quantities of wind shear and potential temperature (top panel) along with resulting C_n^2 profiles for the FA. These curves are based on cubic spline interpolations of the ~ 80 data points in the range of altitudes included in the analysis: 800 m above the elevation of the observatory to 25 km above sea level⁸. By visual inspection of Figure 3.7 we see that the peaks in the wind shear profile correspond to peaks in the OTP, while the rising potential temperature curve suppresses the amplitude of these turbulence peaks with altitude.

We use the OS18 model to calculate the profiles C_n^2 for the wind and temperature data shown in Figure 3.4 (using also the corresponding pressures of each ERA5 altitude); these profiles are summarized in Figure 3.8. A consistent trend is visible in the median value (white solid line) as well as the distribution in purple: the OTP at Cerro Pachón generally decreases with altitude but has a secondary peak around 13 km (approximate altitude of the jet stream). This trend is qualitatively consistent with the results in Figure 3.6, despite the different locations.

We now explain why, in the above figures, we have not shown the C_n^2 estimates for the ground layer using OS18. One of the assumptions in Equation 3.2 is that the atmosphere is in hydrostatic equilibrium; however, in the GL, we typically do not expect this condition to hold. Additionally, turbulence in the GL likely has significant contributions from other sources than vertical wind shear, such as heat from the ground and air flow around complex topography. For these reasons, we expect the OS18 model to significantly under-predict GL turbulence.⁹ An additional issue to using OS18 for the GL turbulence is that we did not find good agreement between the ERA5 and measured values of wind near the ground (Figure 3.5 and Figure 3.4). Even if we trust the OS18 model could

⁷We note that although OS18 writes this as an equality, the units in this heuristic equation are mismatched.

⁸The interpolations produce smoother OTPs that point calculations, but do not introduce significant quantitative differences in integrated turbulence strength.

⁹We have observed this, but it is inconsistent with OS18 results shown in Figure 3.6

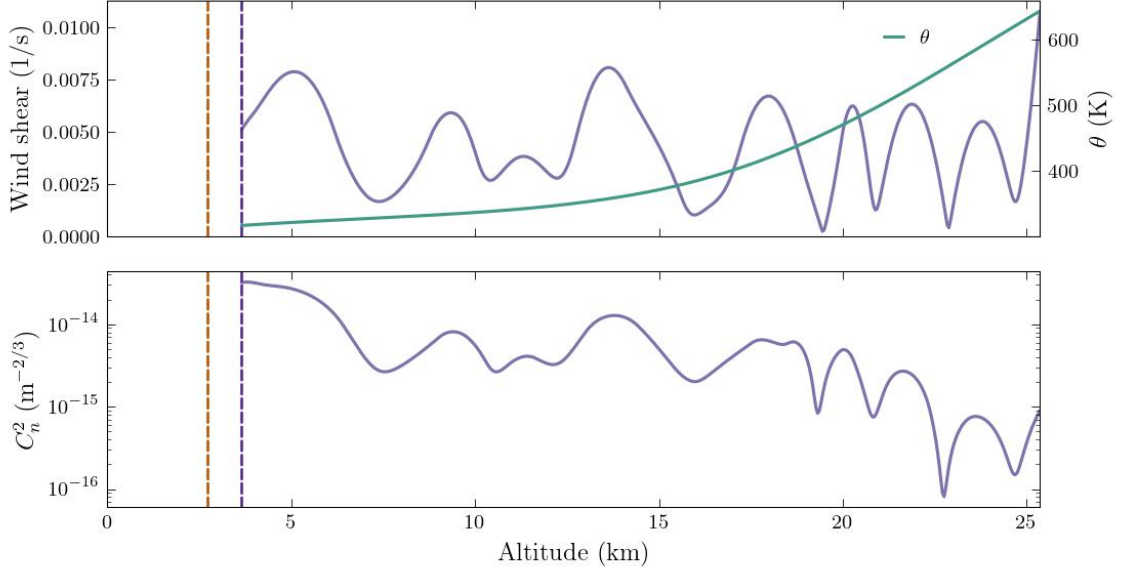


Figure 3.7: Altitude-dependent profiles in the free atmosphere for wind shear (purple curve) and potential temperature (green curve) in the top panel, and resulting optical turbulence profile C_n^2 (bottom panel), for the a ERA5 profile that is representative of the dominant weather conditions. The vertical dashed lines correspond to the observatory elevation (orange) and the beginning of the free atmosphere (purple) at 800 m above ground.

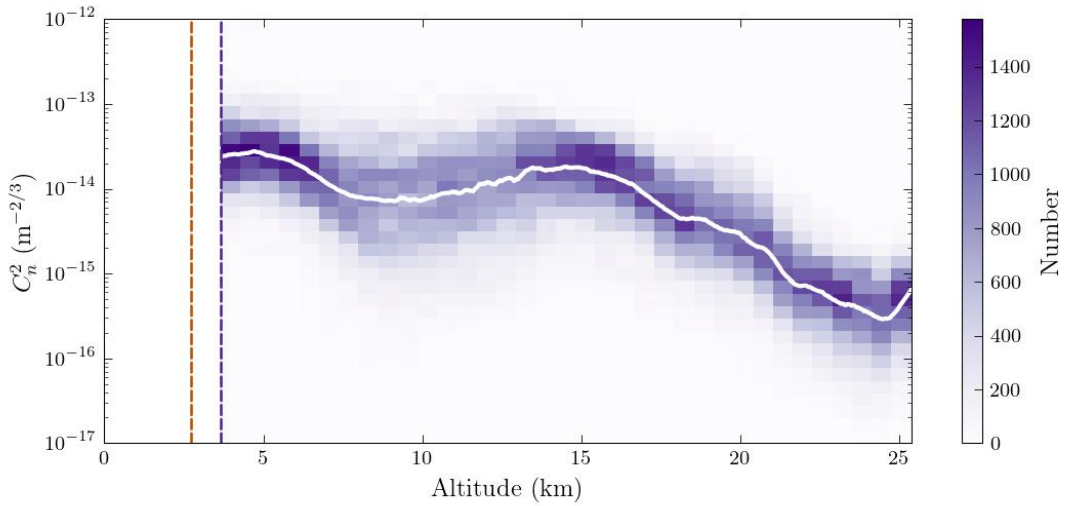


Figure 3.8: Distribution of $C_n^2(h)$ profiles calculated with the OS18 model for the 6 months of ECMWF ERA5 data shown in Figure 3.4 for Cerro Pachón. The white curve corresponds to the median value at each altitude. The vertical dashed lines correspond to the elevation of the observatory (orange) and the beginning of the free atmosphere (purple) at 800 m above ground.

sufficiently capture GL turbulence, we do not expect it to produce representative estimates from inaccurate values of wind. Therefore, we opt to use different methods to estimate GL turbulence. Because the GL does not span more than 1 km in altitude, it is reasonable to model it with a single phase screen; this avoids describing the potentially complex altitude-dependence of the OTP near the ground. Instead, we directly use the turbulence integral of the GL.

The OS18 model captures variations in turbulence strength with altitude, but not the absolute strength; instead it requires calibration for the total turbulence J parametrized with the coefficient k in Equation 3.2. The results shown in Figure 3.6 use a value of k calculated on half the SCIDAR measurements; in other words, the absolute amplitude of the red curves is set by the data, but the variation with altitude emerges from the ECMWF data. We achieve this through the inclusion of site-specific empirical distributions of altitude-specific seeing contributions, described in the next section.

3.2.3 Measurements at Cerro Pachón

In the previous sections we identified two areas where further improvement is needed. First, in Section 3.2.1, we saw that the data products from GCMs, using ECMWF ERA5 as an example, do not accurately capture information about wind velocity near the ground. Second, in Section 3.2.2, we noted that the OS18 model that PSF-WEATHER-STATION will use to calculate free atmosphere OTPs requires calibration for the total turbulence integral—so we require a turbulence integral value for the GL and for the FA. We address both these issues by including site-specific information in PSF-WEATHER-STATION.

As a source of wind velocity measurements that are more reliable than a GCM product, we optionally include in PSF-WEATHER-STATION local weather-tower measurements of wind velocity, as described in Section 3.2.1 and Figure 3.5. As examples, the grey points and histograms in Figure 3.9 show six months of wind direction and velocity measurements from the Gemini South weather tower at Cerro Pachón. The orange points represent data that are temporally matched by PSF-WEATHER-STATION to the ERA5 data from the same time period and are the average of all measurements made within 5 min of the ERA5 timestamps. The time series and projected histograms show consistency between the original and down-sampled data.

To determine the total turbulence integrals for the GL and the FA, we turn to the same literature used by Peterson et al. (2015). Measurements of the altitude dependence of atmospheric turbulence with MASS-DIMM instruments at a variety of sites show that turbulence contributions from the FA and the GL are *independent* (Tokovinin & Travouillon, 2005; Tokovinin et al., 2005; Chun et al., 2009). (For more detail on MASS-DIMM measurements, see Section 2.3.2 – in particular the discussion around Figure 2.4.) Motivated by this independence, the total turbulence in the GL and FA layers (J_{GL} and J_{FA}) are treated separately in PSF-WEATHER-STATION.

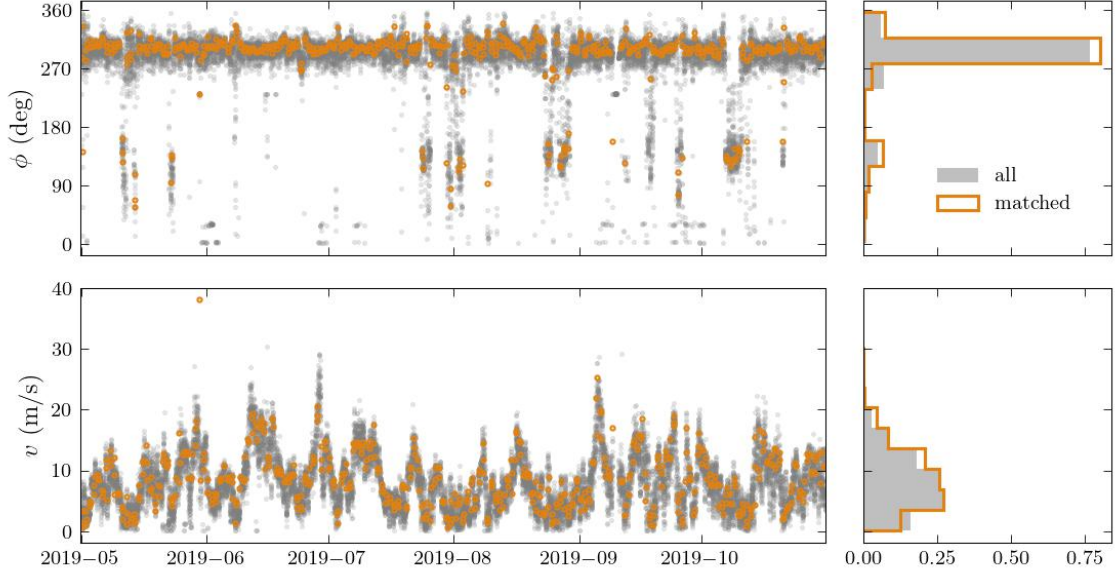


Figure 3.9: Weather-tower measurements of wind direction ϕ (top) and speed (v) (bottom) at Gemini South for the 6 month period of May to October 2019. All available data are shown in grey, and points in orange are the 5-minute average around the ECMWF timestamps. Panels on the right show the time series over the six months, with projected histograms on the right.

For each of FA and GL, we will rely on empirical distributions; these could be compiled from measurements made by any/all of MASS-DIMM, SCIDAR, SLODAR, thermosonde, etc. instruments (with a large enough number of nights to adequately capture variations). As seen in Figure 2.4 and in Tokovinin & Travouillon (2005); Tokovinin et al. (2005); Chun et al. (2009), turbulence integrals are well described by lognormal distributions; for both FA and GL, PSF-WEATHER-STATION takes as input the shape and scale parameters defining a log-normal distribution. It is straightforward to fit these parameters from quantiles (e.g., 25th, 50th, and 75th percentiles) of an empirical distribution, such as those published for Cerro Pachón (Tokovinin & Travouillon, 2005), Cerro Telolo (Tokovinin et al., 2003a; Els et al., 2009), and Mauna Kea (Tokovinin et al., 2005; Chun et al., 2009).

While the relative amount of turbulence contributed by each FA layer is calculated with GCM data and OS18 (detailed in the next section), PSF-WEATHER-STATION uses draws from the J_{FA} distribution to set the total contribution of the FA. Based on the quantiles of the MASS-DIMM distributions at Cerro Pachón presented in Tokovinin & Travouillon (2005), we illustrate in Figure 3.10 examples of PDFs and draws from J_{FA} and J_{GL} . This method of using random draws from empirical distributions does somewhat restrict the predictive capabilities of simulations run with PSF-WEATHER-STATION, as we do not expect to recover the seeing on individual nights; however, we expect to recover overall seeing statistics as well as spatial correlations of the PSFs.

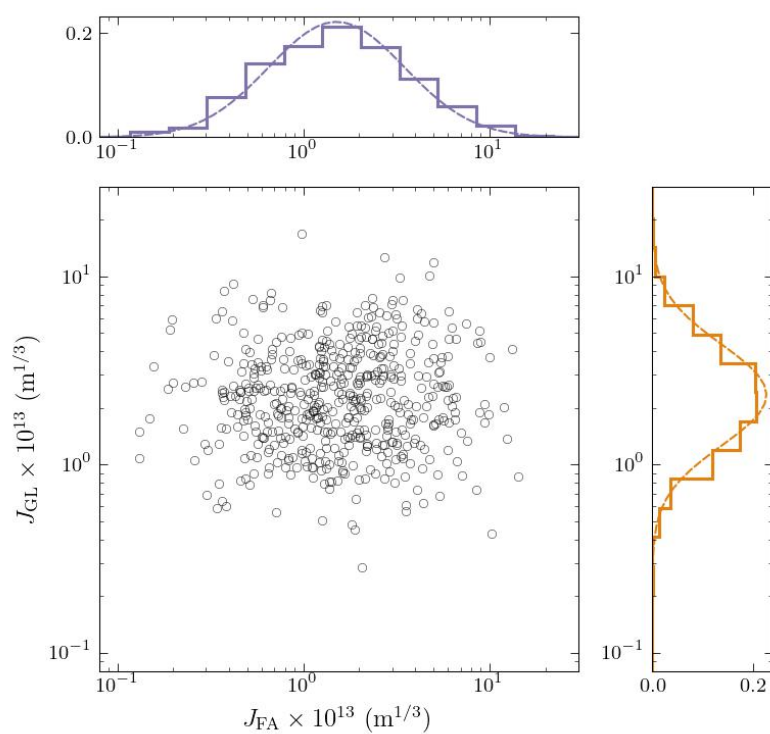


Figure 3.10: Draws from lognormal distributions, with parameters fit to published quantiles of GL and FA turbulence integral measurements from Tokovinin & Travouillon (2005). PDFs are shown as dashed lines and projected histograms of J_{GL} (orange, right) and J_{FA} (purple, top).

3.3 Producing simulation input parameters

The number of phase screens needed to replicate realistic atmospheric PSFs is not well studied in the literature; however, six or seven layers have typically been used (Jee & Tyson, 2011; Peterson et al., 2015; The LSST Dark Energy Science Collaboration et al., 2021b). At locations with fairly constant wind directions across all altitudes, it seems likely that this number suffices, since all turbulence is moving coherently. If wind speeds across altitudes are highly correlated as well, placing adjacent layers close together in altitude would be redundant. Beyond this reasoning, there is as yet no physically motivated model for distributing phase screens for simulating atmospheric PSFs. Therefore, the number of phase screens for which to provide parameters is not fixed in PSF-WEATHER-STATION; the default is to create one layer for the ground, and divide the free atmosphere into $N - 1$ equally spaced layers, for N desired layers. Within the equally spaced FA layers, we set the altitude of each phase screen to be the C_n^2 -weighted altitude. These choices offer some ability to model, as desired, the temporal variability in the number and altitude of turbulence layers observed in SCIDAR data (Osborn et al., 2018; Osborn & Sarazin, 2018).

In Figure 3.11 we illustrate the variability in layer altitude this model produces, for the six-months of Cerro Pachón data shown Figure 3.4, Figure 3.9, and Figure 3.10. The distribution of altitudes for each of the six layers is shown in the top panel, with the fixed altitudes of the six Ellerbroek (2002) layers and seven Tokovinin & Travouillon (2005) shown for reference. The altitude of the phase screen representing the ground layer is fixed in both cases, but the distributions of FA phase screen altitudes for PSF-WEATHER-STATION each cover roughly half the distance between layer boundaries. In the bottom panel, the PSF-WEATHER-STATION histogram is now weighted by the fraction of turbulence (normalized per set of six phase screens) contributed by each phase screen. We also show this weighted histogram for the models described in Peterson et al. (2015) and The LSST Dark Energy Science Collaboration et al. (2021b). Because the ground layer screen on average contributes $\approx 60\%$ of the turbulence, with the remaining 40% distributed between the five FA screens, we display this on a log scale.

We now outline the steps taken in PSF-WEATHER-STATION to go from the raw data, discussed in Section 3.2, to simulation-ready input parameters. We illustrate the process in Figure 3.12 with Cerro Pachón as an example site for the period May to October 2019. We use as data sources those illustrated in the previous section: wind speed and direction measurements from the weather tower at Gemini South, ECMWF ERA5 data products at the grid point closest to Rubin (-30.25 , -70.75), and J_{GL} and J_{FA} quantiles for Cerro Pachón from Tokovinin & Travouillon (2005). The plots in the top two rows of Figure 3.12 show the weather-tower (orange dots and histogram) and ERA5 GCM (purple curves and histogram) data; each GCM altitude profile has a corresponding co-temporal ground-layer data point. The frequency distributions on the right illustrate the range of wind speeds and the distinctly non-uniform distribution of wind directions, both near the ground and across most altitudes.

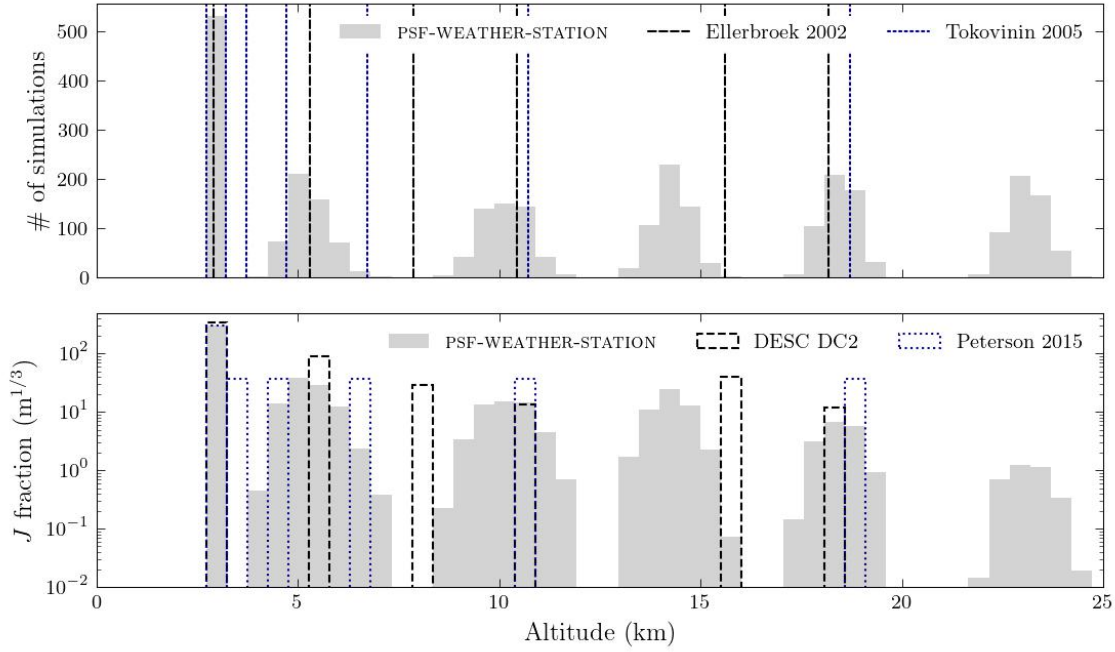


Figure 3.11: Altitude and turbulence distributions obtained using the 6 months of data described in this section and the turbulence integral values (from random draws) shown in Figure 3.10. Top panel shows distribution of altitudes (grey) for 6 phase screens, compared with the set of fixed altitudes used in Ellerbroek (2002) (black dashed lines) and Tokovinin & Travouillon (2005) (blue dotted lines). Bottom panel shows the distribution of altitudes for the same phase screens, weighted by the turbulence weights w (the fraction of total turbulence J contributed per screen, for each set of 6 phase screens). The analogous distribution for the The LSST Dark Energy Science Collaboration et al. (2021b) and Peterson et al. (2015) models are shown in the black dashed/blue dotted histograms (see Section 3.1.2 for more details).

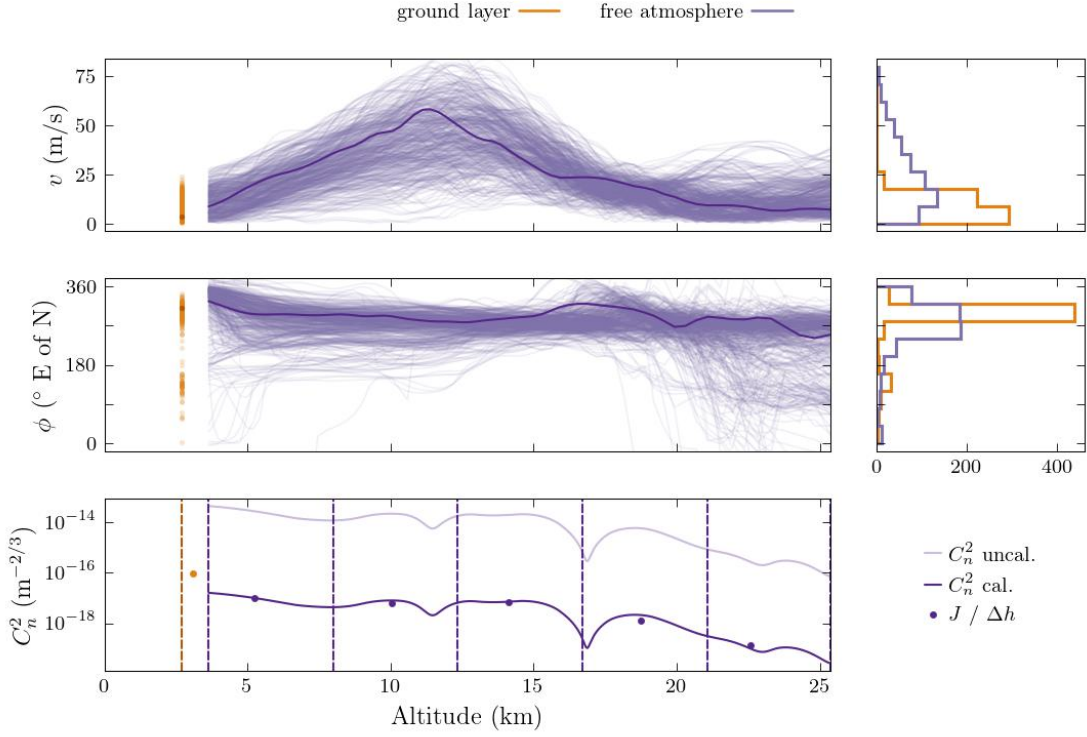


Figure 3.12: Six months of wind data at and above Cerro Pachón (05/01/2019 to 11/01/2019), processed with PSF-WEATHER-STATION. We plot wind speed (top) and meteorological direction (direction of wind *origin*; middle) as a function of altitude, and as a frequency distribution (right). Weather-tower measurements for the ground layer (50 m above the elevation of the observatory) are shown in orange; ECMWF ERA5 data for the free atmosphere are shown in purple. The heavy purple line in each panel corresponds to data from the same representative example time as in Figure 3.7. In the bottom panel, the uncalibrated $C_n^2(h)$ profile for the same time is shown in light purple; the calibrated profile (scaled by J_{FA}), is shown in dark purple. The dashed vertical lines depict the boundaries between the altitude bins used to calculate turbulence integrals J for each FA phase screen. The dots show the value of J divided by the corresponding range of altitudes for each of the phase screens, which are positioned at the C_n^2 -weighted position within each FA layer (purple dots) and at 200 m above the observatory for the ground layer.

The heavy purple curve in each plot in Figure 3.12 corresponds to a chosen example time that is representative of the dominant weather conditions. This representative example will serve to illustrate our process in this section. The OS18 output for this representative example is shown in the bottom panel of Figure 3.12; the light purple curve is the uncalibrated data and the dark curve is shifted by the calibration value chosen randomly from the J_{FA} distribution.

For the illustration in Figure 3.12, we calculate turbulence integrals for six layers of the atmosphere. In the bottom panel of the figure, the purple vertical dashed lines correspond to the boundaries between the six layers; the orange dashed line corresponds to the elevation of the observatory – here 2737 m at Gemini South. The dots correspond to the values of the turbulence integrals divided by the width of the layer, and are placed at 200 m above the elevation of the observatory (following The LSST Dark Energy Science Collaboration et al. (2021b)) and the C_n^2 -weighted position within each FA layer. The complete set of parameters p_{psfws} returned by PSF-WEATHER-STATION is $p_{\text{psfws}} = \{h, J(h), v(h), \phi(h)\}$, where $v(h)$ and $\phi(h)$ are the interpolated wind speed and direction at the height h of the phase screen. PSF-WEATHER-STATION does not currently provide estimates for \mathcal{L}_0 due to the lack of physically motivated models linking environmental parameters to the outer scale.

3.3.1 Optional additional correlations

Using random draws from turbulence distributions for ground turbulence integrals and calibration of C_n^2 for the FA is not an optimal solution since these draws are not temporally correlated with other environmental parameters. There are a range of studies in the literature correlating observed values of J_{GL} , or ground seeing fraction, with measured ground wind speeds. At Mauna Kea, there is strong evidence of correlation between J_{GL} and ground wind speed (Chun et al. (2009), data from 27k SLODAR profiles taken on 124 nights over 2 years). At Cerro Telolo, there is some evidence of correlation between ground seeing and ground wind speed coming from the south-west (secondary dominant wind direction), but no evidence of correlation with the dominant north-eastern wind (Tokovinin et al. (2003a), data from 17k MASS-DIMM profiles taken on 58 nights over 3 months). At Cerro Pachón, the evidence for correlation is mixed and seems to point to only very low-altitude turbulence being correlated with wind speed (Tokovinin et al. (2005), data from 500 SODAR profiles over 16 nights, and 7k MASS-DIMM profiles over 70 nights).

Given these variable observations of the degree of correlation between ground turbulence and wind speed, we include in PSF-WEATHER-STATION an option to correlate the random J_{GL} draws to values of wind speed $v(\text{GL})$ with a desired value of the correlation coefficient ρ (default value is 0). In more technical terms, we can choose samples from a joint distribution $p(J_{\text{GL}}, v(\text{GL}))$ rather than from each of the marginal distributions. This is achieved by (1) sorting both distributions (i.e. maximum correlation of joint distribution), (2) iterating through the samples of J and swapping each with a random pair (within some window in separation), randomizing the joint distribution

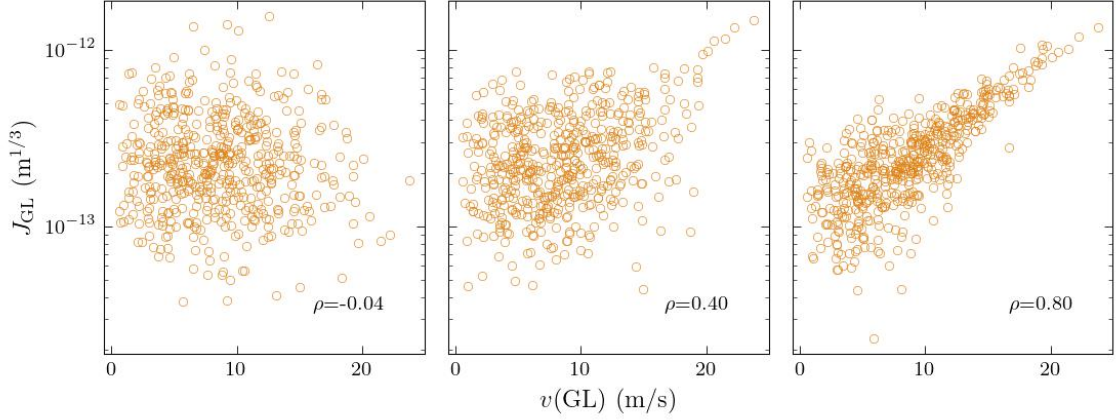


Figure 3.13: Example PSF-WEATHER-STATION output of ground turbulence J_{GL} and wind speed $v(\text{GL})$ for three different settings of ρ_{Jv} : 0 (left); 0.4 (center); 0.8 (right). The values of both v and J_{GL} are unchanged across panels, but permuted ordering produces a 2D distribution with varying amounts of correlation.

until the desired correlation coefficient is achieved. We illustrate this in Figure 3.13; the same set of J_{GL} and ground wind speed $v(\text{GL})$ are shown in each panel, with the input correlation coefficient increasing from left to right. The arrangement of points in the leftmost panel has spurious but negligible correlation ($\rho = 0.05$, for requested $\rho = 0$), and the other panels show perfect recovery of the desired level of correlation.

Although there is possible correlation of J_{GL} with ground wind *direction* (Tokovinin et al., 2003a), we have not yet implemented such an option in PSF-WEATHER-STATION. Likewise, because there is only limited empirical evidence of similar correlations between J_{FA} and FA wind speed (Tokovinin et al., 2003a) – and, in the OS18 model, turbulence already depends on wind shear – we do not include an option to introduce correlations between FA parameters.

3.3.2 Projecting direction from Earth to sky

All wind speeds and directions described and shown in the chapter thus far have been in the local North-East-Zenith (N-E-Z) coordinate system — i.e., the unprimed coordinates in Figure 3.1. In this section, we describe a change of coordinates to translate a wind velocity vector from the local N-E-Z coordinate system to its projection on the imaged plane of the sky. This is motivated in the context of PSF-WEATHER-STATION to provide seamless usage with the GalSim implementation of atmospheric PSF simulations: phase screens in GalSim are oriented perpendicular to the line of sight—see Figure 3.14—i.e., in the imaged plane of the sky (at zenith, this is identical to the local N-E-Z coordinates). Astronomical observations typically have a non-zero zenith angle; however, the wind is defined in local coordinates regardless of the altitude and azimuth of an observation’s



pointing.

An additional effect of this phase screen orientation is the *tilt* of the screen relative to the atmospheric layer it represents: there will be some distance between the edges of the two, causing turbulence to appear at a higher or lower altitude in the simulation. The maximum relevant value in a simulation is the distance to the atmospheric layer from the point on the phase screen corresponding to an observation at the edge of the field of view (indicated by teal arrow in Figure 3.14). For a angular field of view of 2β , a screen at height h , and a telescope diameter D , the distance along the

phase screen from the central line of sight to the point marked by the teal arrow is $\approx h \sin \beta + \frac{D}{2}$. The zenith angle ζ of the central line of sight gives the angle between the screen and layer; so the maximum relevant distance is $\approx (h \sin \beta + \frac{D}{2}) \sin \zeta$. For simulations of observations at Rubin Observatory, which has a diameter of $D \approx 8$ m, a field of view of 3.5 deg, and a maximum zenith angle around $\zeta \approx 45$ deg, the maximum distance is $\approx (3 + 0.02h)$ m — this will always be negligible compared to the height of the screen. For example, for the ground layer screen, which is typically placed at a height of 100 m or more, this distance is only ≈ 5 m.

We do, however, need to account for the difference in the line of sight distance when pointing away from zenith. As we can see in Figure 3.14, the distance between the telescope and the atmospheric layer increases by a factor of $\sec \zeta^{-1}$. There is also an increase in overall turbulence strength due to the increased airmass for an observation off zenith, which is given by the same factor of $\sec \zeta$. Both of these affects are accounted for in PSF-WEATHER-STATION for observations simulated away from zenith.

3.4 Summary

In this chapter we motivated the need for realism in atmospheric PSF simulations used in image simulations for studies of systematic effects for weak lensing and validation of mitigation techniques: local weather conditions are not random and uniform, and can in principle contribute correlations in the PSF that have previously been unaccounted for. We described PSF-WEATHER-STATION, a new software package that produces realistically correlated sets of input parameters needed to run atmospheric PSF simulations. The methods in PSF-WEATHER-STATION rely on two kinds of weather data: data products from global forecasting models, and parameters of log-normal distributions describing the distributions of turbulence strength at the ground and in the free atmosphere. A third kind of data, measurements of wind velocity from an on-site weather tower, are recommended but optional.

We used 6 months of data for Cerro Pachón, the location of Rubin Observatory, as an illustrative example. However, the PSF-WEATHER-STATION package is not specific to any observatory. The code is flexible and includes functionality to download necessary datasets from global weather forecasting models (see Appendix A for details) for any date and location. PSF-WEATHER-STATION has been code-reviewed and is fully documented, along with example notebooks, at <https://github.com/LSSTDESC/psf-weather-station/tree/master>.

3.4.1 Future work

There are additional functionalities that could be implemented in the PSF-WEATHER-STATION package. As mentioned above, there is some evidence of correlations between the wind direction and the strength of turbulence at the ground. This is an effect that depends highly on location, but

we could include a method to produce such correlations, with tunable parameters for the degree of correlation and the direction along which such a correlation is present.

A current limitation in frequency of samples one can draw from PSF-WEATHER-STATION is the frequency of GCM data product outputs (hourly, for ECMWF ERA5); it is currently not possible to draw parameters from PSF-WEATHER-STATION to simulate a series of observations with a higher frequency. One could expect the free-atmosphere weather to change more slowly than at the ground; therefore, because weather-tower measurements are sampled on the scale of minutes, a method could be constructed to draw more frequently from the ground parameters while using the same or interpolated free-atmosphere information. An implementation challenge for this method, given the current implementation in which random values of turbulence strength are drawn, would be to ensure some degree of continuity in that parameter.

In terms of parameters that PSF-WEATHER-STATION produces, there are a few useful additions that would require further research and modeling. The first is the outer scale \mathcal{L}_0 , for which we need a physically motivated model to link this parameter to local weather conditions. Another possible addition, which would broaden the use of PSF-WEATHER-STATION, is to use humidity and cloud cover values from the GCM to determine the density of precipitable water vapor along an atmospheric column—of use for calculation of the atmospheric transmission in optical wavelengths (Perrefort et al., 2019). This would be relevant for LSST since the stringent requirements on photometric precision requires careful calibration of effects such as atmospheric absorption.

Chapter 4

Impact of wind on spatial correlations in atmospheric PSF

4.1 Motivation and context

PSF-WEATHER-STATION uses multiple sources of telemetry and vetted models to generate sets of correlated input parameters for PSF simulations. In this chapter, we describe benchmark tests of these generated parameters. These tests aim to quantify how simulations run with PSF-WEATHER-STATION parameters compare to earlier generations of atmospheric PSF simulations with uncorrelated input parameters (Section 4.2). In particular, we examine the role of environmental input parameters, such as wind speeds and wind directions at different altitudes, on the atmospheric PSF statistics. This is motivated by the observation, in both ground layer measurements and GCM weather data, of persistent wind directions over the Rubin site (see Figure 3.4 and Figure 3.9); it is important to understand what, if any, effect this may have on PSF size and shape correlations.

4.2 Simulation study

This validation study compares atmospheric PSF simulations for Rubin Observatory generated with input parameters set in three different ways.

1. PSFWS: In the first case, input parameters are generated for Cerro Pachón using PSF-WEATHER-STATION with the data summarized in Figure 3.12. Six phase screens are used; the altitudes are allowed to vary according to the C_n^2 scheme described in Section 3.2.2.
2. BENCH: As a second case, we use as a benchmark the input atmospheric parameters used in the DESC Data Challenge 2 (DC2) image simulations (The LSST Dark Energy Science

Collaboration et al., 2021b). For each of six phase screens, the wind speed and direction are drawn from uniform distributions between 0-20 m/s and 0-360°, respectively. Small (~10%) Gaussian variations around the turbulence integrals from Ellerbroek (2002) are introduced, but the associated six altitudes remain fixed between simulations.

3. MATCH: As a third case we use the same values of input parameters as in BENCH – *except for the wind directions*, which are matched to the correlated wind directions used in PSFWS.¹ The motivation for this MATCH simulation is to identify whether differences between distributions of PSF parameters for the first two cases are mainly driven by the highly correlated wind directions in PSF-WEATHER-STATION.

In Figure 4.1 are shown the distribution of wind speed and directions corresponding to each of the three cases described above (PSFWS, BENCH, and MATCH). As expected, in the left panels the distributions of wind speed for BENCH and MATCH are consistent with each other; in the right panels, the distributions of wind directions for PSFWS and MATCH are identical. The distribution of FA wind speeds for PSFWS extends to values three times than BENCH. The PSFWS wind direction distributions in both GL and FA are very peaked, both at similar values of ≈ 270 deg, but the FA peak is broader. The distributions of input turbulence integrals are shown in Figure 3.11; the values for PSFWS are the distributions labeled PSF-WEATHER-STATION (solid grey), and BENCH and MATCH both use the distributions labeled “DESC DC2” (dash black).

For each of the three cases (PSFWS, BENCH, and MATCH), we simulate one 30-second, 3.5-deg² exposure (the expected exposure time and area for Rubin LSST) for each of the 531 time points in the six months of ERA5 and site telemetry data. Each triplet of simulations has the same outer scale (drawn from a truncated log normal distribution with median of 25 m and used for all phase screens), the same random seed for generating phase screens, and the same atmospheric seeing (drawn uniformly from 0.6 to 1.6 arcsec). The contribution of seeing from each phase screen varies according to the turbulence integrals used for that case (PSF-WEATHER-STATION outputs for PSFWS; randomized Ellerbroek (2002) for BENCH and MATCH). The distributions of values of simulation parameters that are common across the three simulation types are shown in Figure 4.2. Seeing FWHM and r_0 are functions of wavelength; the values shown correspond to $\lambda = 500$ nm but all simulations were drawn at the mean i -band wavelength $\lambda = 754.06$ nm. Given this observing wavelength and a value of L_0 , the FWHM and r_0 from the distributions in Figure 4.2 are converted to the von Kármán FWHM and r_0 for that exposure via Equation 2.26 and Equation 2.29.

All simulations described here are generated with the GalSim² software library (Rowe et al., 2015). We use the same GalSim implementation as described in The LSST Dark Energy Science Collaboration et al. (2021b), which follows ray-tracing methods developed in Jee & Tyson (2011) and

¹Although the altitudes of the FA screens are not exactly the same in BENCH and PSFWS, this has little impact because the wind directions are so strongly correlated across phase screens.

²<https://github.com/GalSim-developers/GalSim>

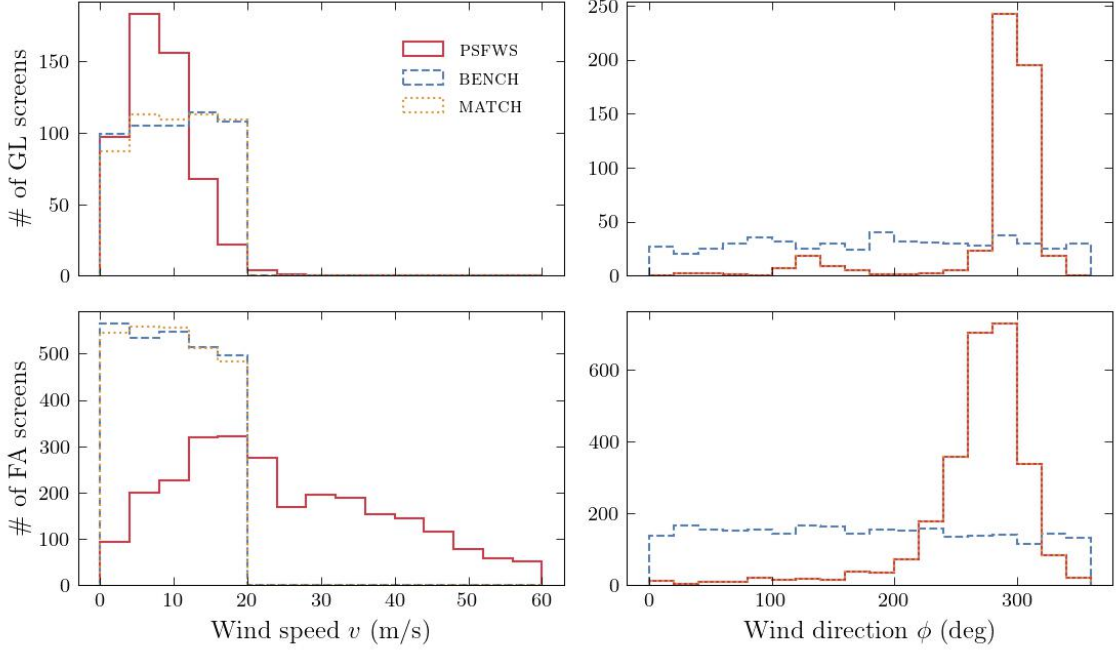


Figure 4.1: Summary of wind speed and direction input values to the three types of simulation input parameters: PSFWS (solid red), BENCH (dashed blue), MATCH (dotted yellow). Top panels are values for the ground layer (GL) screen, and bottom panels summarize values for all free-atmosphere (FA) layers.

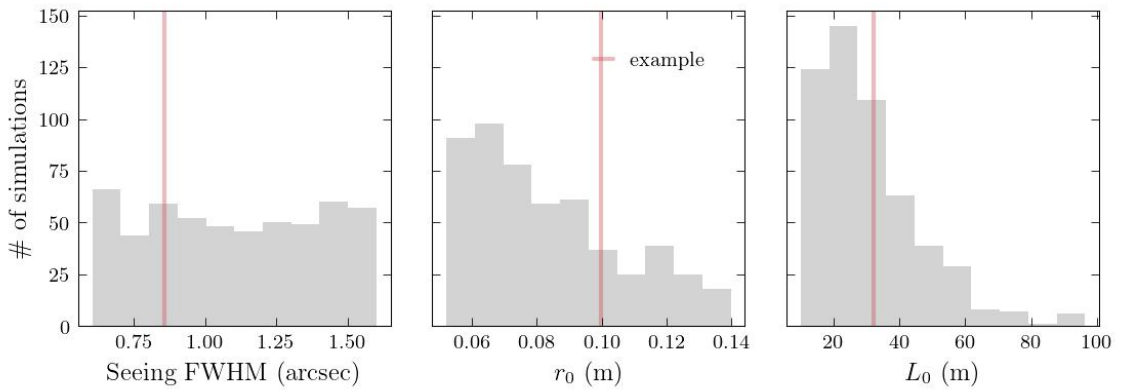


Figure 4.2: Summary of common input parameters to the three types of simulations. The histograms show, left to right, the input seeing FWHM in arcsec, the Fried parameters r_0 , and the outer scale L_0 . Vertical red lines show the value of each parameter used for the PSFWS representative example. The seeing and Fried parameter values are given at $\lambda = 500$ nm.

Peterson et al. (2015). The simulation includes, for each photon, a refractive kick proportional to local instantaneous phase screen gradients to treat large scale turbulence and a statistical “second kick” treatment of small scale turbulence. We do not enable any chromatic effects, the optics is treated as a perfect Airy function, and the sensor is simulated as a perfect photon collecting grid. No background light is included.

PSF images are generated at 50k random locations across the field of view with a pixel resolution of 0.2 arcsec. Each PSF is drawn with sufficient photons (10^6) such that Poisson fluctuations are not significant and then convolved with a Gaussian of 0.35 arcsec FWHM to account for the PSF contribution from optics and sensors. To avoid issues related to overlapping PSF images, each PSF is generated and measured individually on a 50×50 pixel grid and its second moments measured with the GalSim implementation of the HSM adaptive moments algorithm (Hirata & Seljak, 2003).

To compare only the effect of the difference in environmental parameters between PSFWS and BENCH, the PSFWS simulations are run with six phase screens, but their altitudes are allowed to vary according to the C_n^2 scheme described in Section 3.3. See Figure 3.10 and Figure 3.11 for the distributions of turbulence integrals and layer altitudes used in the simulation.

We ran an additional fourth type of simulation, which I call RANDOM, as a control set. In the RANDOM simulations, the PSF size and ellipticity distributions match those of the atmospheric PSF simulations; however, the size, ellipticity, and orientation of each PSF is chosen randomly to provide a spatially uncorrelated control sample. As with the other simulation types, we generated a set of 531 “exposures” with 10^4 PSFs in each exposure. These simulations will not generally be included in the discussion since they are intended only as a null test of certain calculations.

For each of the simulation inputs – PSFWS, MATCH, or BENCH – we expect the simulated turbulence to imprint structure in the distribution of PSF parameters across a given exposure. As an example of the PSFWS input, we display in Figure 4.3 the spatial distribution of the PSF size relative to the mean size (left panel) and the PSF shape ($|e|$ and orientation, right panel) of the simulated PSFs for a single exposure generated with one of the PSFWS parameter sets (the representative example described in Section 3.2.2 and depicted by the dark curves and dots in Figure 3.12). The PSF parameters are clearly correlated across the field of view and the correlations are not isotropic.

In the following sections, we quantify these correlations and their anisotropy for each type of simulation inputs. We compare ensemble statistics across exposures for each simulation type, with a focus on the spatial two-point correlation functions of PSF size and shape. All 2PCFs were computed using the open-source TreeCorr software package.³ Results for isotropic 2PCFs are presented in Section 4.3; we generalize to anisotropic 2PCFs in Section 4.4. In particular, we are interested the dependence on particular input weather parameters of these two-point correlation functions; this is discussed in Section 4.5.

³<https://rmjarvis.github.io/TreeCorr>

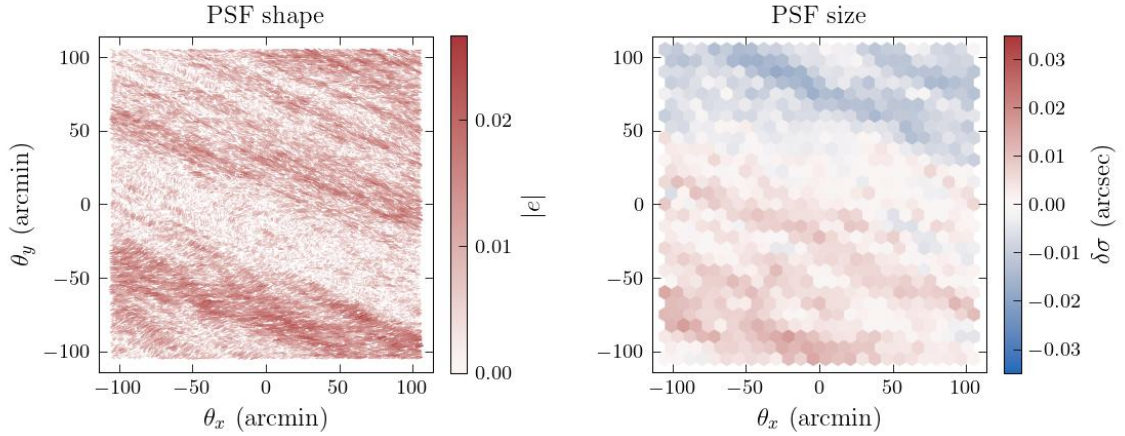


Figure 4.3: Spatial distributions of PSF ellipticity (right) and PSF size relative to mean size (left) across a simulated 3.5- deg² exposure generated with the PSF-WEATHER-STATION input parameters for the representative example depicted by the dark curves and dots in Figure 3.12. The orientation of each line in the left plot corresponds to the orientation of the major axis of the PSF shape, while the length and color contrast are proportional to the magnitude of the ellipticity e .

4.3 Simulation results for isotropic correlation functions

Two-point correlation functions of galaxy shapes were introduced and discussed in Section 1.2.2; we will use precisely the same formalism for the analysis of PSF shapes below. We are also interested in 2PCFs of PSF size; because using a PSF of the wrong size can lead to large bias in shear estimation, correctly measuring and interpolating the sizes over the field of view is a crucial part of the PSF estimation pipeline (Section 1.4.2).

The 2PCF of sizes is simpler than that of shapes, since size is a scalar rather than a spinor with two components. We define the size two-point correlation function as

$$C(\theta) = \langle \delta\sigma_a \delta\sigma_b \rangle(\theta), \quad (4.1)$$

where $\delta\sigma$ is the deviation in PSF size σ from the mean in that exposure, (a, b) denote a pair of PSFs, and the angle brackets indicate an average over all pairs at each separation θ .

This size correlation $C(\theta)$, measured on the PSFWS representative example, is shown in the left panel of Figure 4.4 as a function of separation θ in arcminutes. The right panel shows a summary of $C(\theta)$ for all simulation types and exposures. For each simulation type, the median value of $C(\theta)$ over the 531 simulations is plotted as a curve. In addition, the central 50 percentiles of values for PSFWS are indicated by the shaded region.

The RANDOM result in Figure 4.4 is consistent with zero for all separations, as expected. Compared to BENCH and MATCH, the PSFWS simulations show lower correlation at small separations,

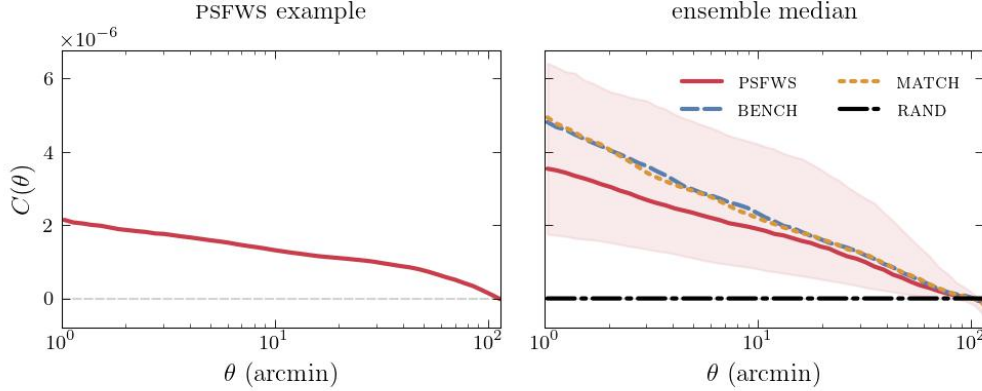


Figure 4.4: Isotropic two-point correlation functions of PSF size variation (Equation 4.1) for the representative PSFWS simulation (left) and the ensemble of simulations (right). The ensemble median correlations are shown as curves in the right figure for all simulation types: PSFWS (solid red), BENCH (dashed blue), MATCH (dotted yellow), and RANDOM (dot-dash black). The shaded area shows the 25th to 75th percentiles of PSFWS simulations (i.e., the central 50 percentiles).

which could be due to the higher free-atmosphere wind speeds in PSFWS: higher wind speeds indicate that more air passes over the telescope aperture during an exposure, lowering correlations overall by introducing more variation in optical turbulence. This effect is discussed further in Section 4.5, where we address why this difference in correlations is observed only at smaller separations.

We now return to the two-point correlation function of shape parameters presented in Section 1.2.2:

$$\xi_{\pm}(\theta) = \langle e_t e_t \pm e_{\times} e_{\times} \rangle(\theta). \quad (4.2)$$

In addition, we will use the decomposition of $\xi_{\pm}(\theta)$ into E - and B -modes from Schneider et al. (2022) described in Section 1.2.4:

$$\xi_{\pm}(\theta) = \xi_{\pm}^E(\theta) \pm \xi_{\pm}^B(\theta) + \xi_{\pm}^{\text{amb}}. \quad (4.3)$$

The ambiguous modes ξ_{\pm}^{amb} account for the modes that cannot be uniquely separated into E - or B -modes due to the finite range of angular separations for which we have information.

These eight 2PCFs ($\xi_{\pm}(\theta)$, $\xi_{\pm}^E(\theta)$, $\xi_{\pm}^B(\theta)$, $\xi_{\pm}^{\text{amb}}(\theta)$) for the representative PSFWS simulation are shown in Figure 4.5. The upper panels show $\xi_{+}(\theta)$ (left) and $\xi_{-}(\theta)$ (right), middle panels show the pure E - (red) and B - (black dotted) mode correlations, and lower panels show $\xi_{\pm}^{\text{amb}}(\theta)$. There are a few key differences between these results and the galaxy shape 2PCFs due to cosmic shear (see Figure 1.7 and Figure 1.8) Firstly, there is a strong B -mode signal in $\xi_{+}(\theta)$ (in fact, the E - and B -modes are practically indistinguishable), and $\xi_{-}(\theta)$ is approximately zero everywhere—these two statements are not independent since, as discussed in Section 1.2.4, $\xi_{-}(\theta)$ approximately represents

the difference between the E - and B -modes.

Secondly, the ambiguous mode $\xi_+^{\text{amb}}(\theta)$ has a nonzero constant value, as well as a quadratic increase at large separations that is expected because of lack of information about $\xi_+(\theta)$ at large θ . A constant value in ambiguous mode indicates a nonzero average shear, as described in Section 1.2.4—and also leads to a positive offset in $\xi_+(\theta)$. More specifically, the constant value of $\xi_+^{\text{amb}}(\theta)$ is equal to γ_0^2 , for average shear magnitude γ_0 . In the example shown in Figure 4.5, the constant value of ξ_+^{amb} corresponds to an average shear of approximately 2.5×10^{-3} . This relatively large average shear contributes to $\xi_+(\theta)$ being strictly positive in the range $\theta < 100$ arcmin shown in the top left panel in Figure 4.5, despite the negative correlations one expects at certain scales (and in particular orientations) due to stripe-like patterns in the PSF of an exposure (see Figure 4.3). In comparison, we *defined* the PSF size variation $\delta\sigma$ to have zero mean per exposure, and its 2PCF (shown in Figure 4.4) does cross zero at the largest separations.

In contrast to $\xi_+^{\text{amb}}(\theta)$, $\xi_-^{\text{amb}}(\theta)$ is zero at all separations, which indicates that at small angular scales, both 2PCFs of PSF shape can be completely separated into E - and B -modes.

Many of the above observations for a single simulation continue to hold for the ensemble of simulations. In Figure 4.6, we compare each of the eight PSF shape 2PCFs across the different simulation types, where the reported result for each simulation type is the median correlation value over all exposures, as in Figure 4.4. We have now separated the E - and B -modes into separate rows and the y axes of the right panels span only 10% of the corresponding panels on the left; otherwise the layout of the figure is the same as for Figure 4.5.

As noted previously, the PSFWS correlations have lower values at smaller separations, with the difference decreasing at larger separations. The E - and B -modes for each simulation type are nearly identical, and $\xi_-(\theta)$ is very small, here on the order of 10^{-6} . The constant value of ξ_+^{amb} for separations less than ≈ 30 arcmin indicates a nonzero average shear for the median case; however, the 25th percentile (lower border of shaded region) is very close to zero, so we expect roughly 25% of the data has a negligible average shear.

4.4 Simulation results for anisotropic correlation functions

In the previous section we discussed the isotropic two-point correlation functions of simulated PSFs. However, as noted previously, the PSF in each exposure is clearly not isotropic (see Figure 4.3, for example). In this section, we generalize the study of simulated exposures to anisotropic two-point correlation functions. Since the results from Section 4.3 showed $\xi_-(\theta)$ amplitudes negligible compared with $\xi_+(\theta)$, in what follows we include only results for the ξ_+ shape correlation functions.

In Figure 4.7 we illustrate two ways of representing anisotropies in the 2PCF. The upper panels show the PSF parameter variation over the focal plane for the PSFWS representative example. Overlain on both these panels is a pair of sources (black points) and two (equivalent) ways of specifying

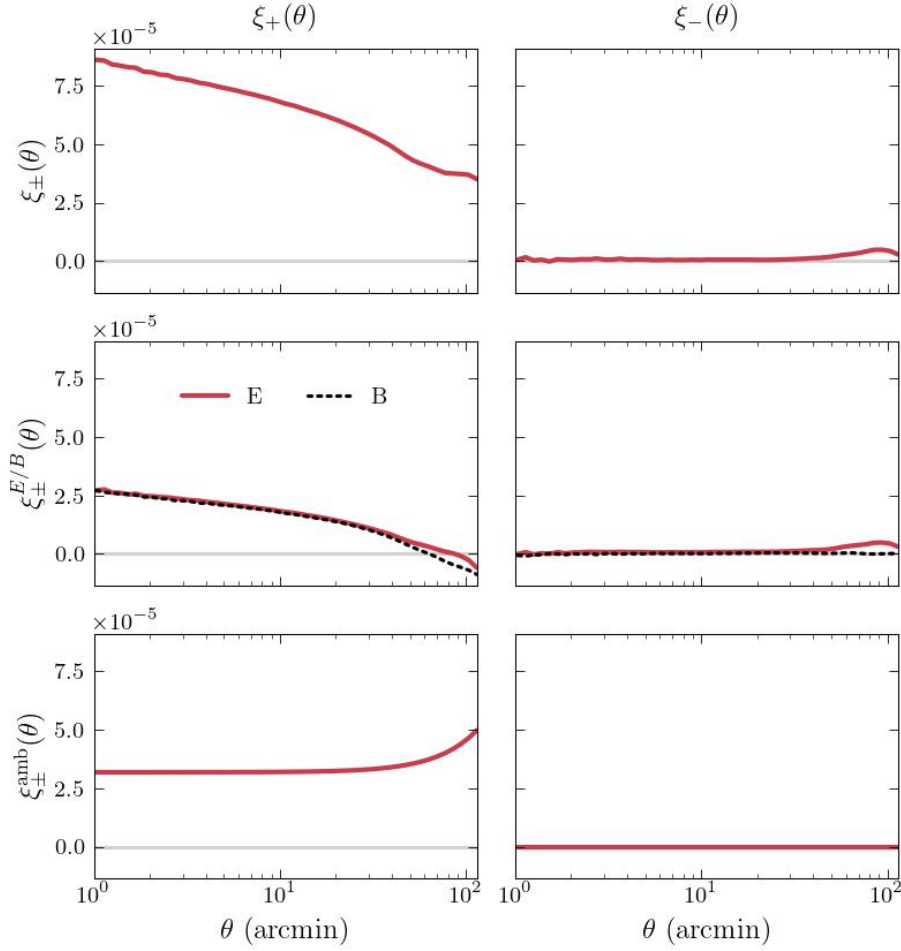


Figure 4.5: Isotropic two-point correlation functions of PSF shape for the PSFWS representative example. Within each column, the correlation $\xi(\theta)$ is shown in the uppermost panel, its E (red) and B (black) mode decomposition in the middle, and the ambiguous correlation in the lower panel. The left and right columns show results for $\xi_+(\theta)$ and $\xi_-(\theta)$, respectively.

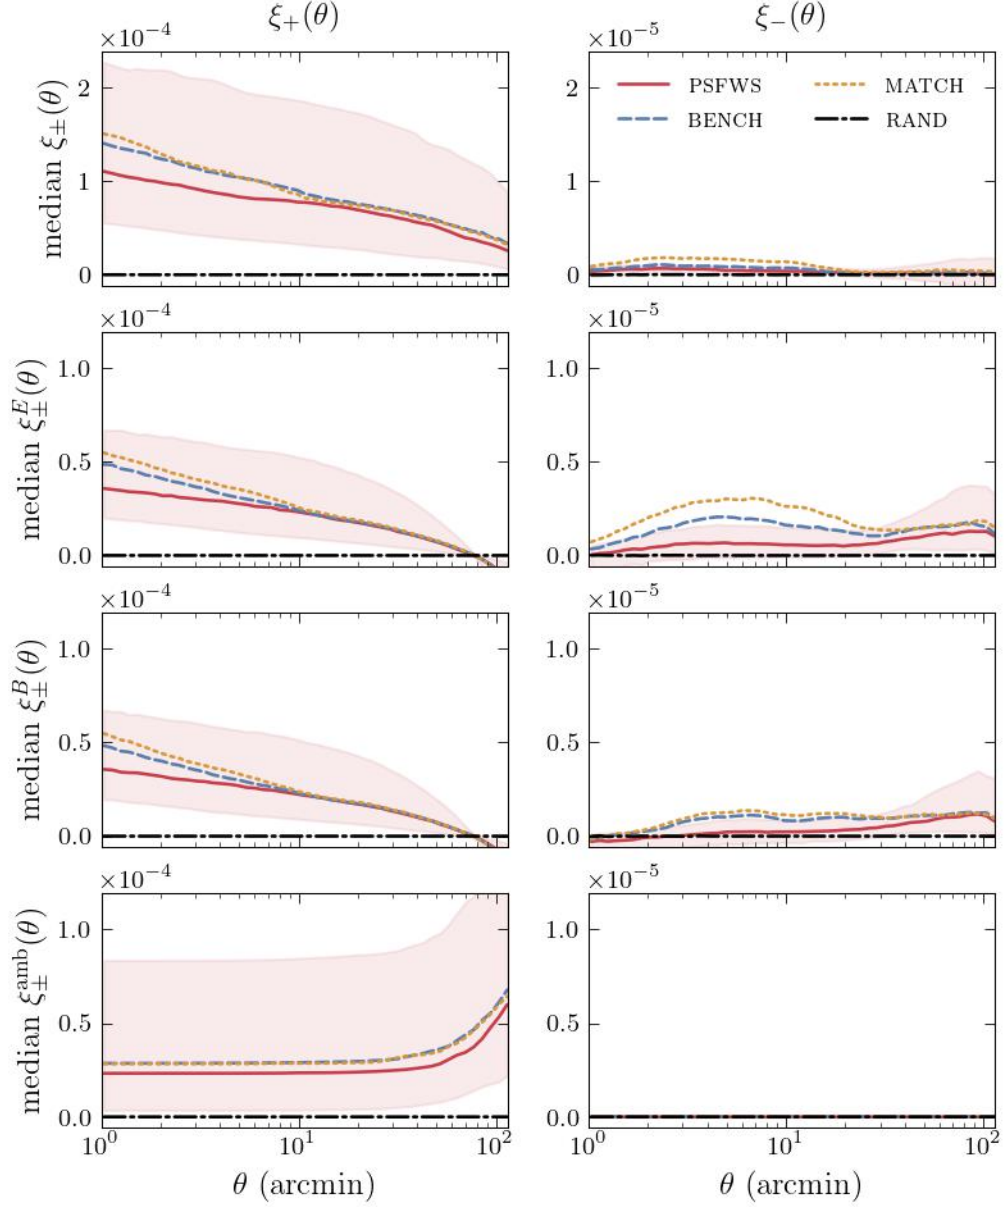


Figure 4.6: Isotropic two-point correlation functions of PSF shape for all simulation types. Within each column, the correlation $\xi(\theta)$ is shown in the uppermost panel, with its E , B , and ambiguous mode decomposition in the panels below. The left and right columns show the ensemble median values for $\xi_+(\theta)$ and $\xi_-(\theta)$, respectively. The red shaded region shows the 25th to 75th percentiles for the PSFWS results.

the vector separating them: the $\Delta\theta_x, \Delta\theta_y$ components (left) and the vector magnitude θ and angle α (right). Note that we measure α relative to the $+y$ -axis on the focal plane, to match the definition of the wind direction ϕ —all simulations are run at zenith, so the $+y$ -axis points North.

As in the isotropic case, we find the anisotropic 2PCF by averaging over all pairs of PSFs in an exposure—however, we now average over all pairs with a given separation *vector*. For instance, the panels in the second row in Figure 4.7 show the 2PCFs for PSF shape (left) and size (right) as a function of separation-vector components $\Delta\theta_x$ and $\Delta\theta_y$. We use the same symbols as in the previous section to denote the anisotropic 2PCFs for shape and size; however, rather than being functions of θ , they are now functions of two variables – e.g., $\xi_+(\Delta\theta_x, \Delta\theta_y)$. Since the 2PCF is agnostic under interchange of two PSFs within a pair, the function is symmetric under 180-degree rotations in the $\Delta\theta_x, \Delta\theta_y$ plane.

We are interested in quantifying the dominant orientations of anisotropy in these 2PCFs—for example, the orientations of the two distinct stripes of stronger correlation we see in both $\xi_+(\Delta\theta_x, \Delta\theta_y)$ and $C(\Delta\theta_x, \Delta\theta_y)$ in the panels in the second row in Figure 4.7. Therefore, we transform the 2PCFs into polar coordinates, separation θ and angle α (illustrated in the top right panel of Figure 4.7). The transformed functions $\xi_+(\theta, \alpha)$ and $C(\theta, \alpha)$ are shown in the panels in the bottom two rows for two ranges of θ : 12 to 120 arcmin, and 1.5 to 6 arcmin. The angle α on the horizontal axis spans only 180 deg because of the symmetry discussed above. This transformation of coordinates allows us to visually associate stronger correlations for a particular value of α with vertical stripes in polar plot.

The maximum separation θ shown in the third row of panels ($\theta = 120$ arcmin) is close to the largest scales probed in cosmic shear analyses, as shown in the figures for DES and KiDS-1000 in Section 1.2. (We cannot probe the largest scales used in cosmic shear analyses because of the limited size of the field of view of each simulated Rubin-like exposure.) The minimum separation θ shown in the bottom row of panels ($\theta = 1.5$ arcmin) corresponds to the smallest scales typically included in current analyses of cosmic shear. We separate these regions because we use higher resolution in separation θ to calculate the 2PCF at small values of θ than for large values of θ , since it would be unnecessarily computationally expensive to use high resolution for the larger separations.

In the remainder of this section we display 2PCFs in only polar coordinates. We first discuss the analysis of the anisotropic two-point correlations functions at “large-scale” separations of ~ 100 arcmin and then turn to “small-scale” separations of ~ 1 arcmin .

4.4.1 Large-scale separations of ~ 100 arcmin

In Figure 4.8, we display the ensemble median values of the 2PCFs for shape ($\xi_+(\theta, \alpha)$, left) and size ($C(\theta, \alpha)$, right) for separations θ in the range 12 to 120 arcmin, for the set of simulations of each type: PSFWS (top), MATCH (middle), and BENCH (bottom). (This is similar to the analysis described in Section 4.3 for the isotropic 2PCF.) For both the PSFWS and MATCH simulations, there is a very

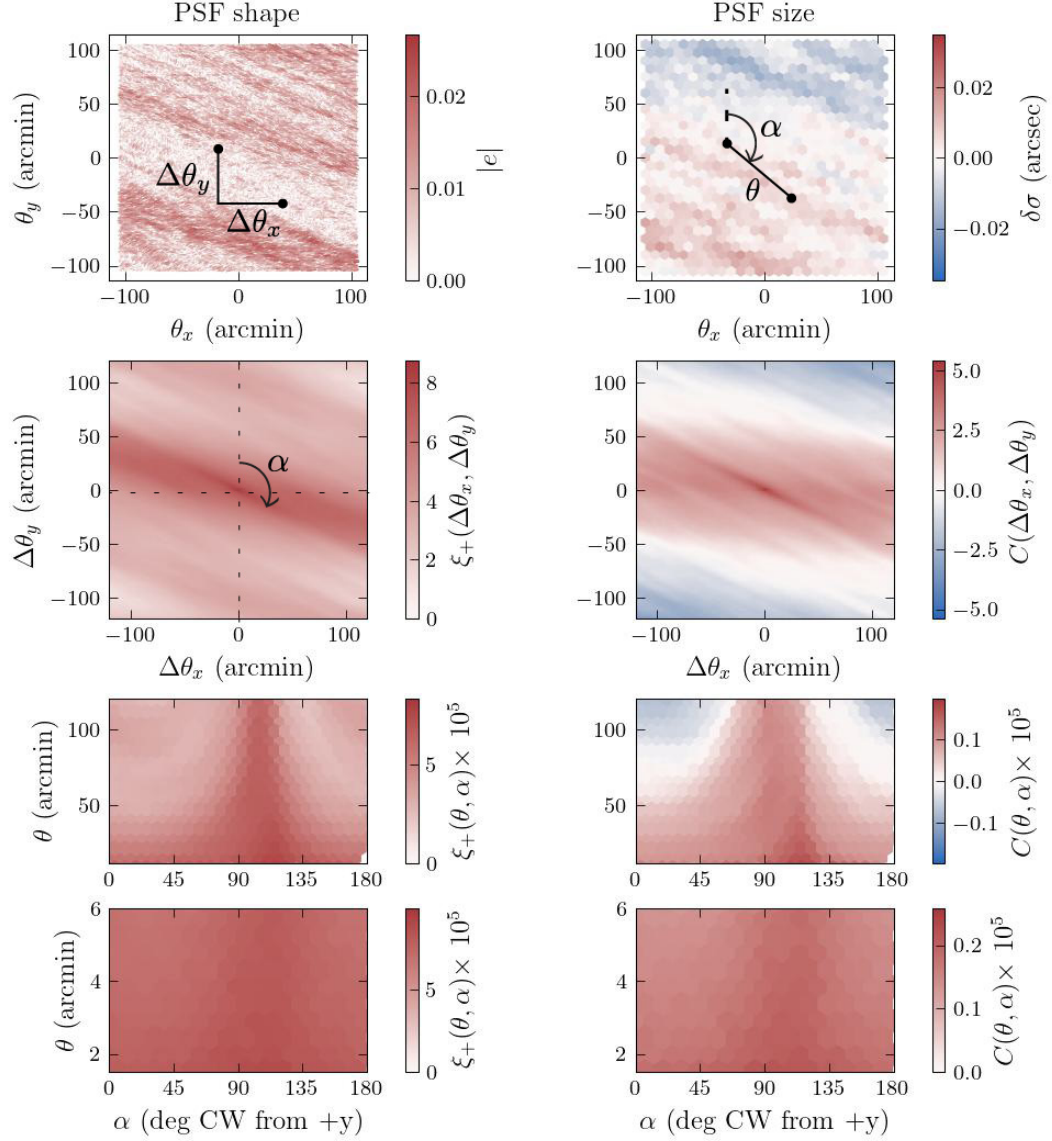


Figure 4.7: Illustration of anisotropic two-point correlation functions for the PSFWS representative example, for PSF shape (left column) and PSF size (right column). The top panels show simulated PSF parameters across the field of view overlain with diagrams of an example pair of points (black dots) and the separation between them, in Cartesian (left) and polar (right) coordinates. The panels in the second row show the dependence of the 2PCFs on the Cartesian coordinates $\Delta\theta_x$ and $\Delta\theta_y$. The panels in the bottom two rows display the same 2PCFs, but transformed into polar coordinates (separation θ versus angle α) for two ranges of θ : 12 to 120 arcmin, and 1.5 to 6 arcmin.

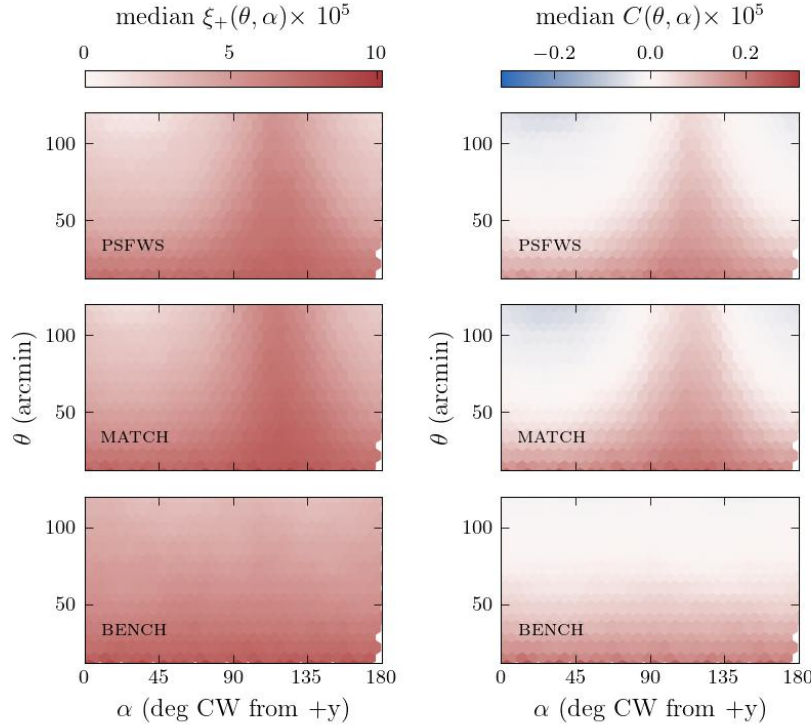


Figure 4.8: The median value of the anisotropic two-point correlation functions at separations θ in the range 12 to 120 arcmin, for PSF shape $\xi_+(\theta, \alpha)$ (left) and size $C(\theta, \alpha)$ (right), for the three simulation types: PSFWS (top), MATCH (middle), and BENCH (bottom).

apparent dependence of the shape and size 2PCFs on α ; in particular, both the PSF shape and size median values peak near the same value of α . In contrast, the median values of the 2PCFs for the BENCH simulations show no evidence of an α -dependence.

Given this visual evidence of dependence on α of the median 2PCF value, we quantify the anisotropy as illustrated in Figure 4.9 for the PSFWS representative example. We average the 2PCF over a range in separation θ to construct a one-dimensional 2PCF, as a function of α , from the two-dimensional version. We choose a θ range of 12 arcmin spanning $\theta = 108$ arcmin to $\theta = 120$ arcmin, indicated by the horizontal dashed lines in the top panels in Figure 4.9. The α range is separated into 30 bins in which we average the values of $\xi_+(\theta, \alpha)$ (or $C(\theta, \alpha)$) within the range in θ . The results of this calculation, denoted $\xi_+(\alpha)$ and $C(\alpha)$, are shown in the bottom panels. We find the bin in α at which this maximum occurs and label it α_{\max} . We repeat this procedure for the anisotropic 2PCFs for each of the simulated exposures, obtaining a set of one-dimensional functions $\xi_+(\alpha)$ each with an associated value of α_{\max} , and similarly for $C(\alpha)$.

We now determine the ensemble median values of $\xi_+(\alpha)$ and $C(\alpha)$ over each set of 531 simulations; these are shown in the left and right panels of Figure 4.10, for each of the three types of

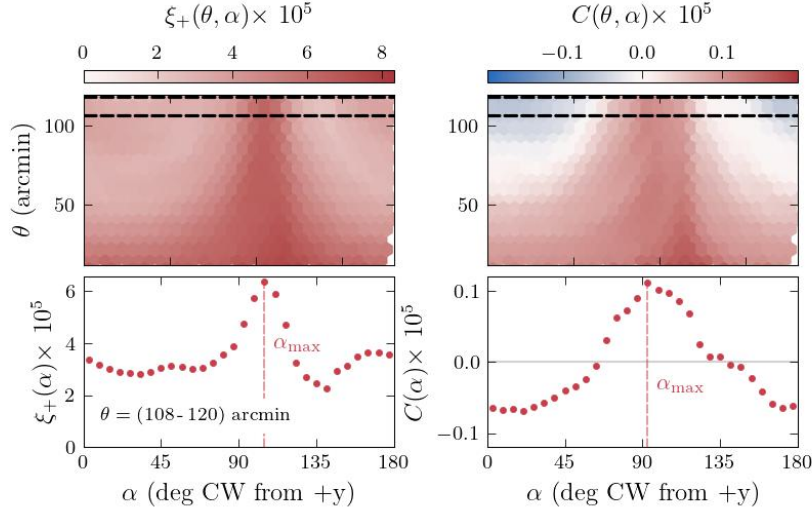


Figure 4.9: Illustration of method to find the orientation of maximum correlation at separations of ~ 100 arcmin for the anisotropic 2PCFs $\xi_+(\theta, \alpha)$ (left) and $C(\theta, \alpha)$ (right). The top panels show the same two-point correlation functions as the third row of panels in Figure 4.7. The points in the bottom panel correspond to the average value of the 2PCF for θ between 108 and 120 arcmin (the region between the horizontal lines in the top panels), in 6-deg bins in α . The vertical dashed lines indicate the angle α_{\max} at which $\xi_+(\alpha)$ and $C(\alpha)$ take on their maximum values.

simulations. The strong anisotropy for PSFWS and MATCH, but no evidence of α -dependence for BENCH, is consistent with Figure 4.8. Furthermore, the values of α at which the median values of the 2PCFs peak (≈ 120 deg) are consistent for both the PSF shape (left) and size (right), and for both PSFWS and MATCH.

In Figure 4.11 we plot the distributions of the 531 values of α_{\max} for PSF shape (left) and size (right) for each simulation type. The distributions of α_{\max} for both shape and size are strongly peaked around 120 deg for PSFWS and MATCH, while the BENCH histogram is consistent with a uniform distribution. The fact that the distributions for MATCH are much more consistent with PSFWS than BENCH indicates that the orientation of anisotropies in the 2PCF is dominated by the matched wind directions. We will explore this link to wind direction in the next section, but we first show the results of anisotropic correlations at separations of ~ 1 arcmin.

4.4.2 Small-scale separations of ~ 1 arcmin

We now repeat the analysis described above for the anisotropic correlations at separations of ~ 1 arcmin. In Figure 4.12, we display the ensemble median value of the 2PCF over the set of simulations of each type: PSFWS (top), MATCH (middle), and BENCH (bottom). For both PSF shape (left) and size (right) the amplitude of the median 2PCF value dominates the color scale, and any

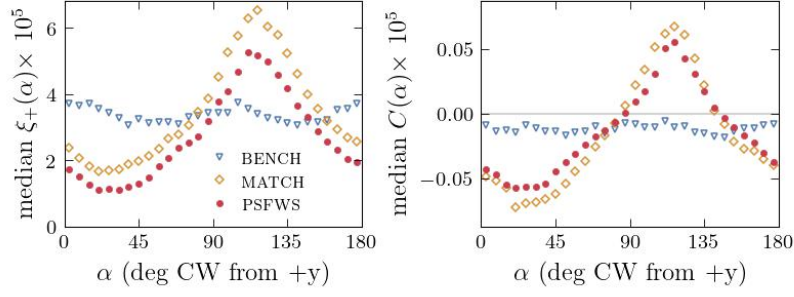


Figure 4.10: Ensemble median values of $\xi_+(\alpha)$ (left) and $C(\alpha)$ (right) averaged over separations θ between 108 and 120 arcmin and in 6-deg bins in α . Results are shown for three simulation types: BENCH (blue triangles), MATCH (gold diamond), and PSFWS (red dots).

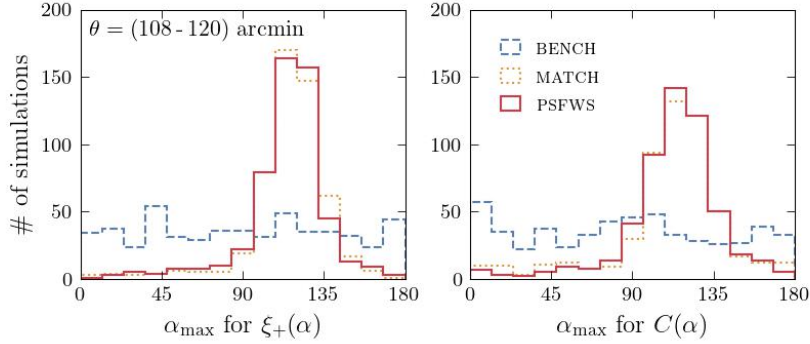


Figure 4.11: Histogram of α_{\max} values for $\xi_+(\alpha)$ (left) and $C(\alpha)$ (right) at large separations $\theta = (108 - 120)$ arcmin, for the three types of simulation: BENCH (blue dashed), MATCH (gold dotted), and PSFWS (red solid). The width of each α_{\max} histogram bin covers 12 deg – i.e., twice the width of the bins in α used for the two-point correlation functions $\xi_+(\alpha)$ and $C(\alpha)$.

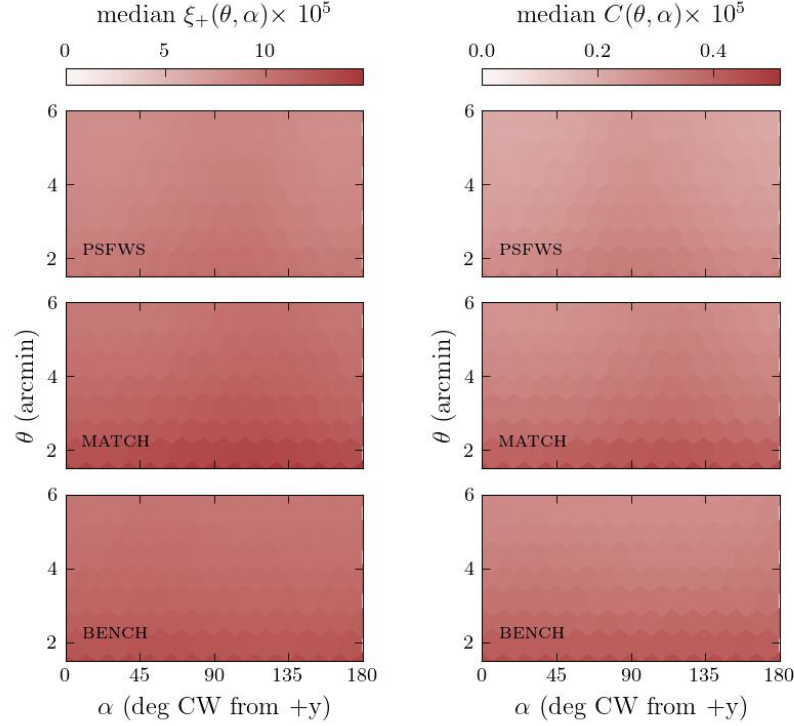


Figure 4.12: The median value of the anisotropic two-point correlation functions at separations of ~ 1 arcmin, for PSF shape $\xi_+(\theta, \alpha)$ (left) and size $C(\theta, \alpha)$ (right), for the three simulation types: PSFWS (top), MATCH (middle), and BENCH (bottom).

variation with angle α here is small in comparison. We will examine the anisotropy in the same way as above – by calculating a one-dimensional 2PCF in α for each simulated exposure and finding its maximum.

Figure 4.13 illustrates the method for the PSFWS representative example. For small separations, we average the 2PCF over a range of θ between 3 and 6 arcmin (the region between the dashed black horizontal lines in the top panels of the figure). Since there are fewer samples of the 2PCF at these separations than were available for separations of ~ 100 arcmin⁴, we use bins more coarsely sampled in α ; the resolution is now 10 deg per bin for a total of 18 bins.

The α dependence in this one-dimensional version of the 2PCF is more apparent than in the upper panels. For this individual example we can draw consistent conclusions between upper and lower panels: although the overall positive correlation dominates, we are able to see evidence of anisotropy with the maximum correlation at $\alpha \approx 80 - 85$ deg.

⁴Since these 2PCFs were calculated in Cartesian coordinates and transformed to polar coordinates, the density of information at small separations is low and limits the number of α bins we can use while keeping the number of points in each bin roughly uniform. The resolution could be increased for simulations run with a much higher density of PSFs.

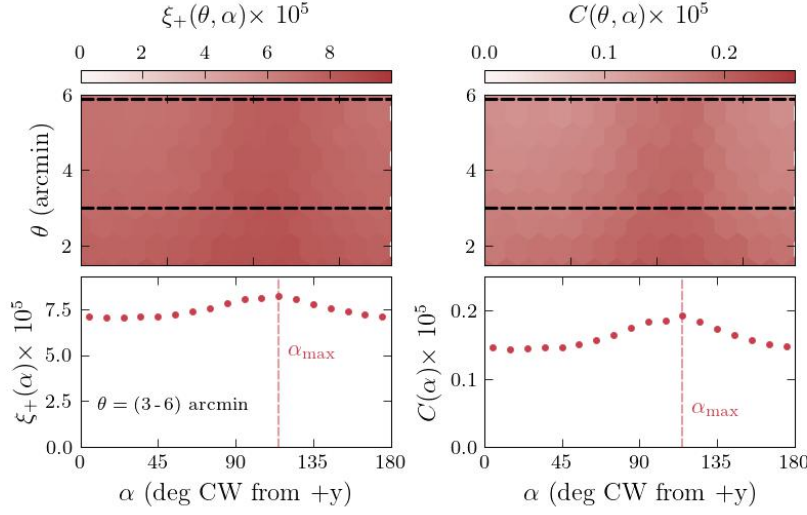


Figure 4.13: Illustration of method to find the orientation of maximum correlation at separations of ~ 1 arcmin for the anisotropic 2PCFs $\xi_+(\theta, \alpha)$ (left) and $C(\theta, \alpha)$ (right). The top panels show the same two-point correlation functions as the bottom panels in Figure 4.7. The points in the bottom panel correspond to the average value of the 2PCF for θ between 3 and 6 arcmin (the region between the horizontal lines in the top panels), in 10 deg bins in α . The vertical dashed lines indicate the angle α_{\max} at which $\xi_+(\alpha)$ and $C(\alpha)$ take on their maximum values.

In Figure 4.10, the results for the median values of the one-dimensional 2PCFs $\xi_+(\alpha)$ and $C(\alpha)$ for each set of 531 simulations are shown for each of the three simulation types. The median PSF *size* 2PCFs (right panel) have a non-uniform dependence on α for the PSFWS and MATCH simulations, while there is no evidence of such a dependence for the BENCH simulations. There is a weaker α dependence in the PSF *shape* correlations (left panel).

Finally, we summarize in Figure 4.15 the distribution of the 531 values of α_{\max} for each simulation type, for PSF shape (left) and size (right) for separations of ~ 1 arcmin. As we saw in Figure 4.11 for separations of ~ 100 arcmin, the distributions for PSFWS and MATCH are strongly peaked and very similar to each other, for both $\xi_+(\alpha)$ and $C(\alpha)$. However, the distributions for α_{\max} at separations of ~ 100 arcmin (Figure 4.11) are narrower, and peak at slightly higher values, than the distributions for α_{\max} at separations of ~ 1 arcmin, although direct comparisons are complicated by the different bin sizes in α_{\max} for the two separations.

4.4.3 Summary of anisotropies in 2PCFs in simulations

The results for both small- and large-scale separations point to a strong dependence of the orientation of the anisotropy in the 2PCFs on the profiles of the input wind parameters, which were matched between the PSFWS and MATCH simulations, but were set randomly in the BENCH simulation. In

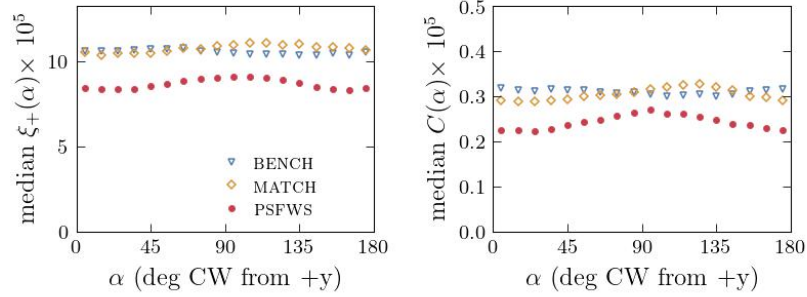


Figure 4.14: Ensemble median values of $\xi_+(\alpha)$ (left) and $C(\alpha)$ (right) at separations of ~ 1 arcmin , for $\theta = (3 - 6)$ arcmin and in 10 deg bins in α . Results are shown for three simulation types: BENCH (blue triangles), MATCH (gold diamond), and PSFWS (red dots).

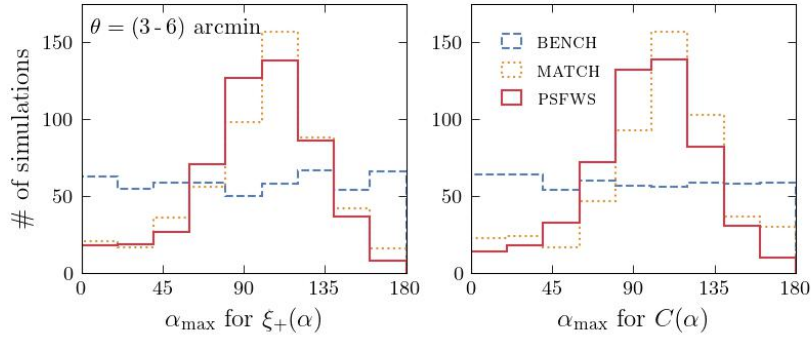


Figure 4.15: Histogram of α_{\max} values for $\xi_+(\alpha)$ (left) and $C(\alpha)$ (right) at small separations $\theta = (3 - 6)$ deg, for the three types of simulation: BENCH (blue dashed), MATCH (gold dotted), and PSFWS (red solid). The width of each α_{\max} histogram bin covers 20 deg – i.e., twice the width of the bins in α used for the two-point correlation functions $\xi_+(\alpha)$ and $C(\alpha)$.

the next section, we explore this dependence in more detail, using the α_{\max} values as a summary statistic of the orientation of anisotropy in the PSF of each exposure.

4.5 Dependence of anisotropies on input wind parameters

In this section, we probe the relationship between simulation input wind directions and the orientations α_{\max} along which the anisotropic 2PCFs for the output PSF parameters are maximum. As described in the previous section, we find α_{\max} values for the 2PCFs of both PSF size, $C(\alpha)$, and shape, $\xi_+(\alpha)$, at two scales – \sim few arcmin and ~ 100 arcmin – for 531 exposures simulated with each of 3 types of simulations. For wind direction, we use both the direction of the ground layer (GL), $\phi(\text{GL})$, and the direction of the free atmosphere (FA), $\phi(\text{FA})$, defined as the wind vector averaged over the five FA layers⁵.

In the scatter plots in Figure 4.16, we plot α_{\max} for the shape 2PCF $\xi_+(\alpha)$, versus $\phi(\text{GL})$ (left) and versus $\phi(\text{FA})$ (right), for separations of ~ 100 arcmin (top) and for separations of ~ 1 arcmin (bottom). The projected histograms in the top row illustrate how the distributions of GL and FA wind directions are uniform for BENCH simulations and peaked at the same dominant wind direction for PSFWS and MATCH simulations⁶. The projected histograms of α_{\max} on the far right are the same as the left panels of Figure 4.11 and Figure 4.15.

The Pearson correlation coefficient for each type of simulation is displayed on each scatter plot. A strong correlation between α_{\max} for separations of ~ 100 arcmin and $\phi(\text{GL})$ (top left scatter plot) exists for all three types of simulation, regardless of the distribution of ground wind directions. In contrast, for separations of ~ 1 arcmin, there is no evidence for a correlation between α_{\max} and $\phi(\text{GL})$.

In particular, the correlation is significant for simulations with realistic altitude-dependent FA wind directions and speeds (PSFWS, $\rho = 0.67 \pm 0.05$) but is roughly consistent with zero when directions and speeds are chosen randomly with altitude (BENCH, $\rho = 0.10 \pm 0.05$). When input FA wind directions (but not speeds) are matched to those of PSFWS, the correlation coefficient has an intermediate value (MATCH, $\rho = 0.28 \pm 0.07$). However, even though the correlation coefficients are statistically quite significant at separations of \sim few arcmin when FA wind directions are correlated across altitudes, the distributions of α_{\max} have higher variance than those at separations of ~ 100 arcmin.

To understand the FA results, we describe the expected angular scale of PSF correlations and a physical source of anisotropies.

⁵An equivalent definition is the speed-weighted average direction over the five layers.

⁶The distributions of $\phi(\text{GL})$ are identical for the PSFWS and MATCH simulations by definition of the “matched” simulation. The distributions of the altitude-averaged free-atmosphere wind direction $\phi(\text{FA})$ are slightly different for the PSFWS and MATCH simulations because the profiles of wind *speeds* are not matched, thereby impacting the average.

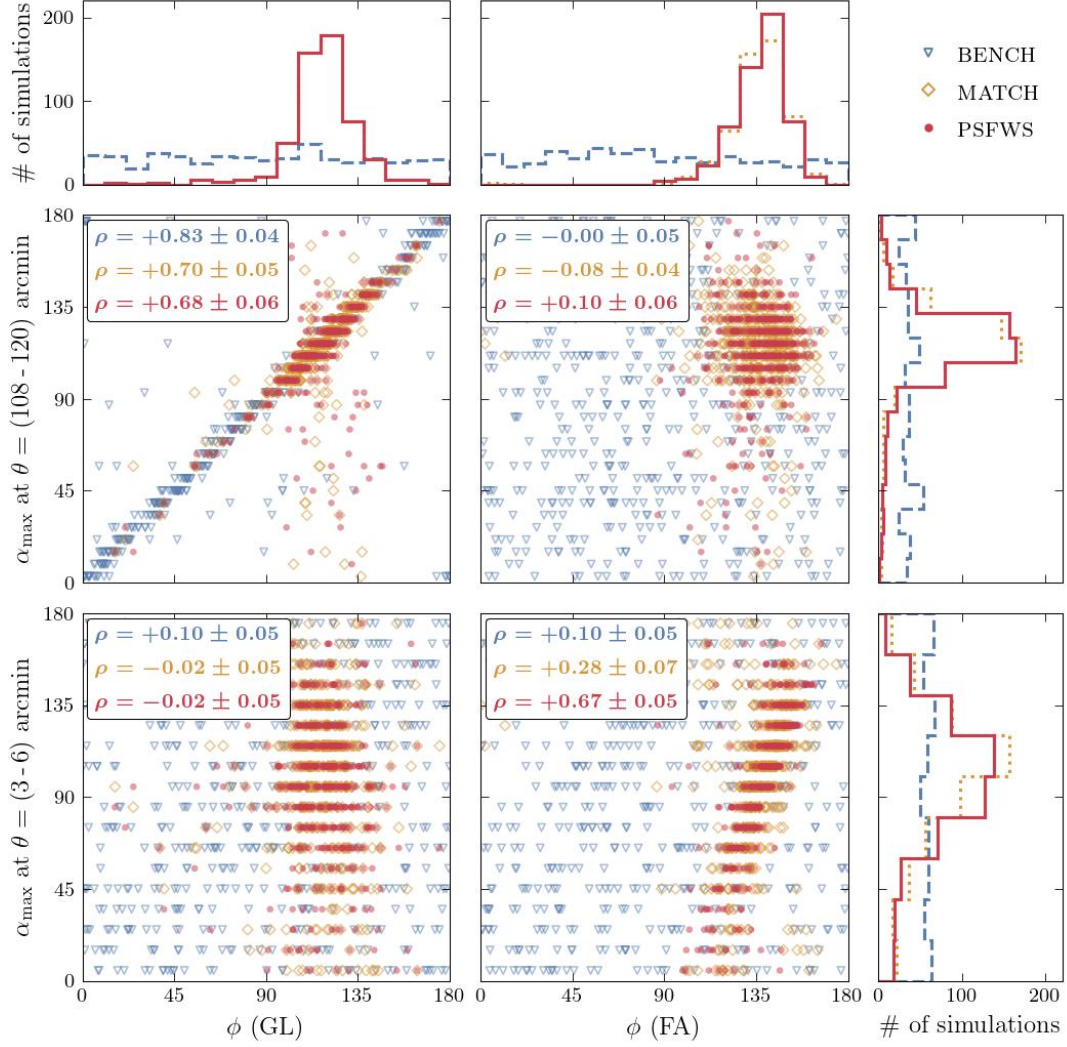


Figure 4.16: Orientation α_{\max} at which the shape 2PCF $\xi_+(\alpha)$ is maximum versus input wind direction ϕ for each of 531 simulated exposures, for the three types of simulations (BENCH, MATCH, and PSFWS), for separations of ~ 100 arcmin (middle row) and separations of ~ 1 arcmin (bottom row). The direction ϕ corresponds to the ground layer wind direction (GL, left column) and the average free-atmosphere wind direction (FA, middle column). The distributions of the four directions are shown in the projected histograms. Correlation coefficients for each type of simulation are reported for each scatter plot (BENCH in blue, MATCH in gold, PSFWS in red).

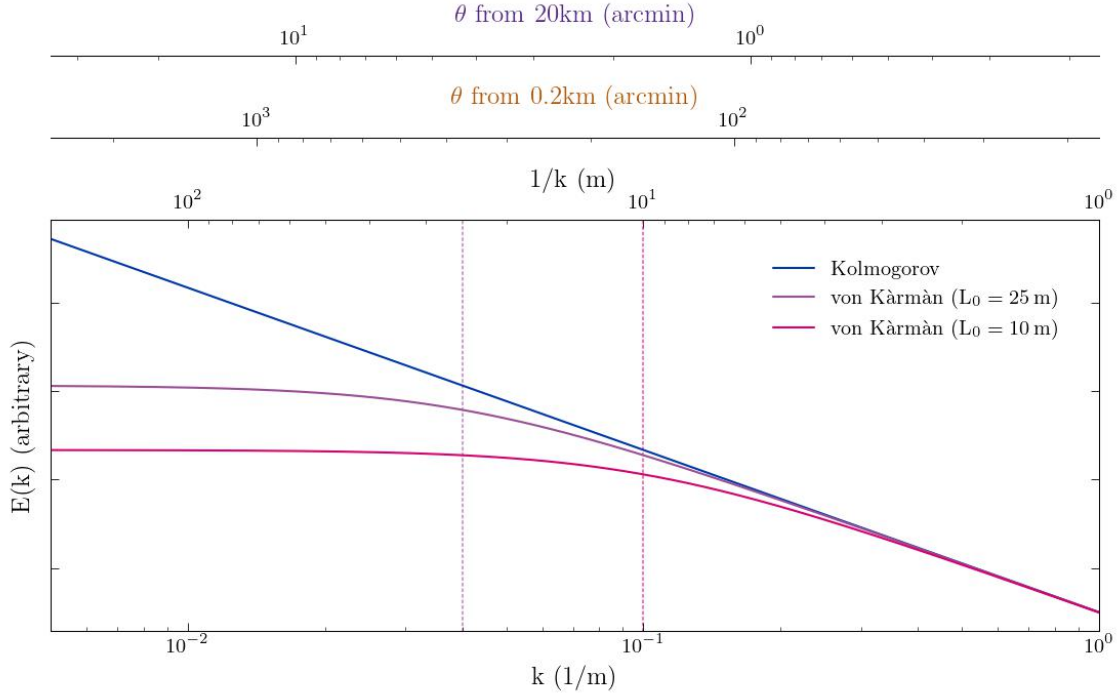


Figure 4.17: Log-log plot of Kolmogorov and von Kármán turbulence spectral energy density. The horizontal (bottom) axis is spatial frequency k ; the top axes show the spatial distance $1/k$ and the associated angular scales θ as seen from different heights (near ground at 200 m, orange; in the free atmosphere at 20 km, purple). von Kármán spectra with $L_0 = 25$ m (purple) and $L_0 = 10$ m (pink) are shown, with L_0 values indicated by horizontal dotted lines. The Kolmogorov spectrum ($L_0 = \infty$) is shown in blue.

Angular scales

Atmospheric features of a particular *physical* size that are located at lower altitudes have a larger *angular* size in the focal plane than features of the same physical size at higher altitudes. Therefore, at one instant in time an optical turbulence pattern near the ground results in spatial variations in the PSF over larger angular separations than the same turbulence pattern located near the top of the free atmosphere. Consider, for example, turbulence with a scale equal to the median outer scale L_0 . (The power spectrum flattens for spatial frequencies lower than $\approx 1/L_0$, as shown in Equation 2.25 and Figure 4.17) For $L_0 = 25$ m and a ground-layer screen at a height of 200 m, the angular scale of turbulence variations is $\arctan \frac{25}{200} \approx 400$ arcmin; for a phase screen at an altitude of 20 km, the angular scale is $\arctan \frac{25}{20000} \approx 4$ arcmin. This range of angular scales is approximately the same as the ranges of separation θ we use to calculate $\xi_+(\alpha)$ and $C(\alpha)$ in Section 4.4, and the associated values of α_{\max} .

Sources of PSF anisotropies and correlations with wind

While wind speeds and directions can change throughout the night, they are typically fairly stable during the course of an exposure. This is particularly true at Cerro Pachón, where the wind direction is persistently coming from the sea. During a 30-sec exposure then, we expect air at all altitudes to be moving coherently; different layers may be translating at different speeds, but in mostly the same direction. As a consequence, PSF images along the direction of the wind on the focal plane have “seen” much of the same optical turbulence, although in slightly different combinations due to wind speed variation with altitude. The shapes of these PSFs will thus be more correlated with each other than with those in a direction orthogonal to the wind.

The realistic inputs used to generate the PSFWS exposures include these correlated wind speeds and directions. In the case of BENCH simulations, there is no coherent motion of the turbulence in the FA because the wind speed and wind direction for each layer is chosen randomly. This results in a correlation coefficient ρ consistent with zero between FA wind direction and α_{\max} at separations of \sim few arcmin. The MATCH simulations have wind directions that are correlated across altitudes, but speeds are chosen randomly. The random speeds cause turbulence layers to move with greater differences in speed than the smoothly varying wind profiles in PSFWS, resulting in a suppressed correlation coefficient for MATCH simulations despite the wind directions being matched with PSFWS.

In Figure 4.18 we show the same set of histograms and scatter plots as in Figure 4.16, but now for the PSF size 2PCF $C(\alpha)$. The results are generally consistent with those for PSF shape. Although we still see a strong correlation between α_{\max} and $\phi(\text{GL})$ at separations of ~ 100 arcmin (upper left scatter plot), the values of the correlation coefficient are lower than they were for PSF shape, potentially due to noisier estimates of α_{\max} .

We can also explore the impact that wind speed has on the PSF. In Figure 4.19 we plot the variance of each PSF parameter (ellipticity components e_1 , e_2 , and size σ) across each simulated exposure versus the input ground-layer wind speed for that exposure. For each type of simulation input (PSFWS, BENCH, and MATCH), the variance of all three PSF parameters decreases with increasing ground-layer wind speed $v(\text{GL})$. This is expected because wind moves the phase screens across the field of view, washing out variations in PSF size or shape due to turbulence structure at all angular scales; therefore, the higher the wind speed, the more the variations are suppressed.

In Figure 4.20, correlations between variance in PSF size or shape parameters and free-atmosphere wind speed are found to be weaker. This is expected because, instantaneously, different PSFs across the focal plane are more likely to sample independent regions of the phase screen with increasing altitude (see illustration in Figure 3.1); the resulting variance in a PSF parameter is less likely to decrease due to motion of the screen, compared to the GL screen where different PSFs are more likely to sample overlapping regions of the phase screen.

These results from Figure 4.19 and Figure 4.20 help us explain the observation in Section 4.3

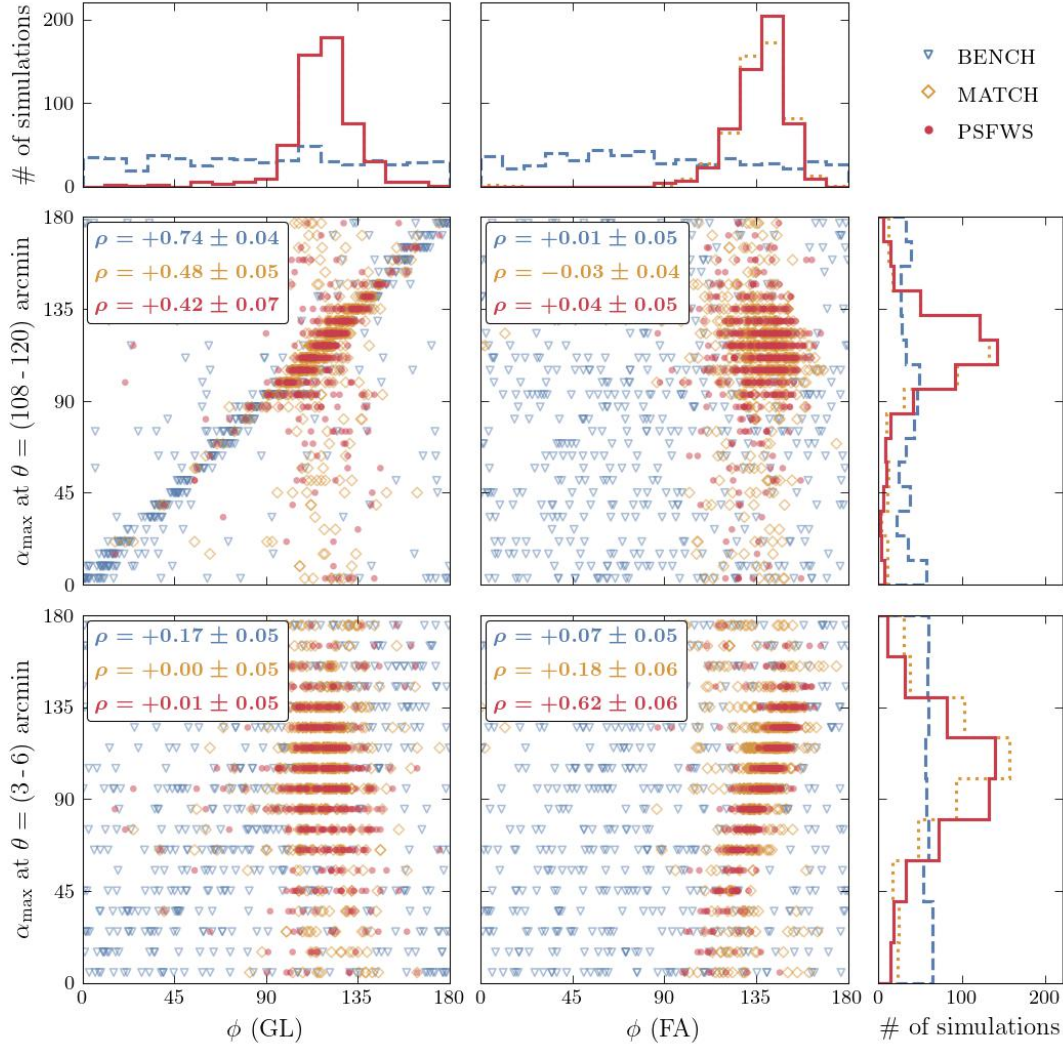


Figure 4.18: Orientation α_{\max} at which the size 2PCF $C(\alpha)$ is maximum versus input wind direction ϕ for each of 531 simulated exposures, for the three types of simulations (BENCH, MATCH, and PSFWS), for separations of ~ 100 arcmin (middle row) and separations of ~ 1 arcmin (bottom row). The direction ϕ corresponds to the ground layer wind direction (GL, left column) and the average free-atmosphere wind direction (FA, middle column). The distributions of the four directions are shown in the projected histograms. Correlation coefficients for each type of simulation are reported for each scatter plot (BENCH in blue, MATCH in gold, PSFWS in red).

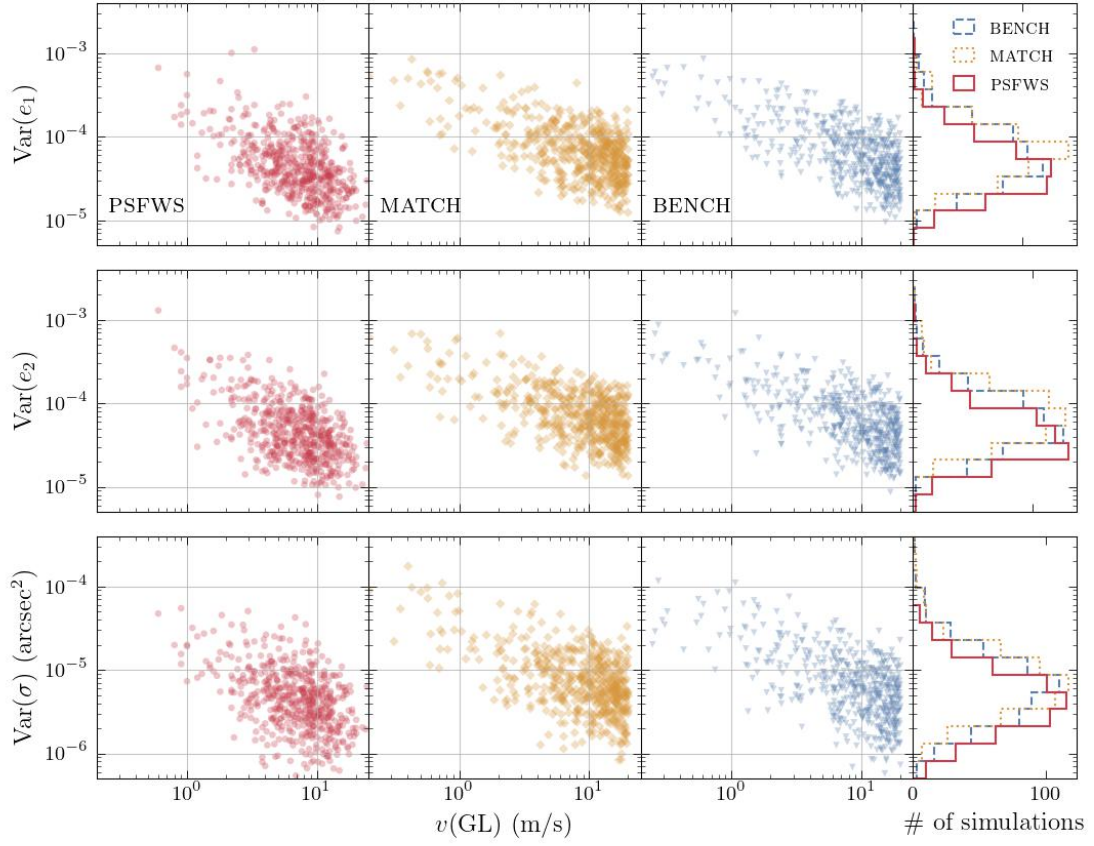


Figure 4.19: Variance of PSF parameters across each simulated exposure versus the input free-atmosphere (FA) wind speed for that simulation, for the three types of simulations (PSFWS, MATCH, and BENCH). PSF parameters are the PSF ellipticity components e_1 (top row) and e_2 (middle row), and PSF size σ (bottom row). The distributions of PSF parameter variance are shown in the projected histograms on the right.

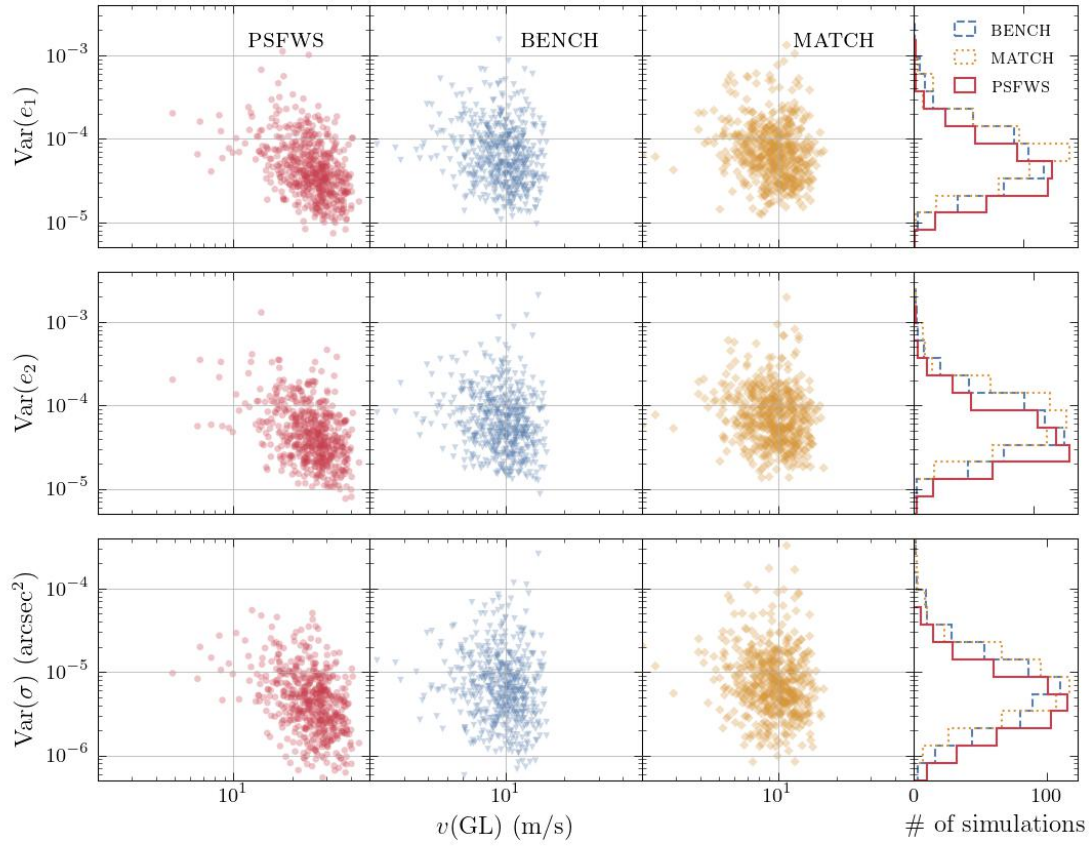


Figure 4.20: Variance of PSF parameters across each simulated exposure versus the input ground-layer (GL) wind speed for that simulation, for the three types of simulations (PSFWS, MATCH, and BENCH). PSF parameters are the PSF ellipticity components e_1 (top row) and e_2 (middle row), and PSF size σ (bottom row). The distributions of PSF parameter variance are shown in the projected histograms on the right.

that PSFWS has lower median values of isotropic two-point correlation functions at separations of ~ 1 arcmin than MATCH and BENCH, while at separations of ~ 100 arcmin the values converged. The PSFWS have higher average wind speed values in the FA than BENCH and MATCH; according to Figure 4.20, this indicates lower variance in the wavefront. As discussed above, the FA most affects small angular scales—so we expect that higher wind speeds in the FA lead to lower spatial correlation in the PSF. At the ground layer, the wind speed distribution of all simulation types are roughly consistent; so the correlation values are similar at separations of ~ 100 arcmin. We can draw these conclusions because the three simulation types were all generated with identical distributions of overall seeing and turbulence structure parameters. Therefore, the wind speed and direction, and relative distribution of turbulence strength with altitude, are the only drivers of differences in PSF parameter correlation anisotropy between simulation types.

4.6 Summary

In this chapter, we have explored both isotropic and anisotropic two-point correlation functions of PSF parameters in simulations, and related the input wind parameters in the simulations to summary statistics of 2PCF anisotropy.

The isotropic 2PCFs of atmospheric PSF shape correlations are found to have E and B modes of roughly equal magnitude, leading to negligible amplitudes in ξ_- . The simulated exposures generated with PSF-WEATHER-STATION parameters (PSFWS) have lower values of median (isotropic) 2PCF, for both PSF shape and size, than exposures generated with the input parameters used in the DESC Data Challenge 2 simulations (BENCH). This is especially true at small scales, because wind speeds in the free atmosphere (FA), which affect smaller angular scales, are much higher on average, and we find that higher wind speeds decrease overall correlations in the atmospheric PSF.

Comparisons of the anisotropic 2PCFs of PSF parameters, motivated, for example, by striation patterns observed in the maps of PSF size and shape across an exposure, also show differences among simulation types: dominant orientations in 2PCF anisotropy are seen over the ensemble of exposures for PSFWS and MATCH simulations (input parameters identical to BENCH except the wind directions, which are matched to PSFWS) but not in the BENCH simulations. At both separations of ~ 1 arcmin and separations of ~ 100 arcmin, there is evidence of correlations between the input wind directions and the orientation along which the 2PCFs are highest, which we denote α_{\max} . In particular, values of α_{\max} at separations of ~ 100 arcmin correlate strongly with the corresponding input ground layer wind direction for all simulation types. Values of α_{\max} at separations of ~ 1 arcmin, on the other hand, have some correlation with the average free-atmosphere wind direction in the PSFWS and MATCH cases (wind directions in these simulations are correlated across altitude), but not in the BENCH simulations. The fact that the anisotropies in 2PCFs at different angular scales depend on wind directions at different altitudes can be understood in terms of the different angles subtended

at different heights by turbulent structure of the same physical scale. High altitude layers imprint features on scales \sim few arcmin, while low altitude layers affect angular scales of ~ 100 arcmin.

The results of this simulation study raise further questions. Because of the dominant wind direction at Cerro Pachón, we expect to see a dominant orientation of the anisotropy in the 2PCF for PSF size and shape with respect to the ground coordinate system, across LSST exposures at Rubin Observatory. If these persistent anisotropies in the atmospheric PSF are observed in data, and persist through the coaddition process, it is crucial for weak lensing measurements that we ensure they can be adequately removed by our PSF modeling and correction algorithms. In the next chapter, we apply the analyses from this simulation study to observed PSFs in real data, and in Section 7, we discuss how wind directions map onto different patches of sky observed during the survey, and the potential impact of resulting PSF anisotropies on PSF modeling for weak lensing.

Chapter 5

Exploration of atmospheric wavefront data from DECam

5.1 Motivation

The study in Chapter 4 was conducted purely on simulated wide-field atmospheric PSFs. It is important and interesting to know whether the strong dependence on wind direction of anisotropy orientation in the spatial correlation of atmospheric PSF is reproduced in data. The challenge is that real images are affected by not only the atmospheric PSF but also include contributions from optics and sensors; these components can imprint additional, and different, spatial correlations in the PSF. In this chapter we attempt to measure spatial correlations in the atmospheric PSF at Cerro Tololo Inter-American Observatory (CTIO), an observatory 15 km northwest of Rubin and home to the Dark Energy Camera (DECam), the wide-field camera used for the Dark Energy Survey (DES). We use a set of wavefront measurements from DECam that were processed to remove optical contributions, following the method developed in Davis et al. (2016). We describe the wavefront sequences and data processing in Section 5.2, and discuss our analysis of two-point correlation functions in Section 5.3.

5.2 DECam wavefront sequences

Wavefront measurements can be made using out-of-focus images; these images are of the wavefront at the primary mirror rather than the (in-focus) point spread function. This image has the shape of the primary mirror—hence, the colloquial term “donut” image for the annular mirror in the DECam optical system. The optical contribution to these wavefront data can be estimated and then subtracted; the residual wavefront should then contain only the atmospheric signal. We summarize

the process of estimating and subtracting the optical contributions to the wavefront from data; the procedure is described in more detail in Davis et al. (2016). The wavefront image at focal plane position x, y is a function of pupil coordinates u, v , and can be decomposed into a linear combination of Zernike polynomials:¹

$$W(u, v)[x, y] = \sum_i a_i(x, y) Z_i(u, v). \quad (5.1)$$

This wavefront is fit to all the donut images in a single exposure; each coefficient a_i varies spatially. After this decomposition of the image, we remove the average value δ_i of each coefficient, as well as additional tilt terms for coma ($i = 5, i = 6$):

$$a'_i(x, y) = a_i(x, y) - \delta_i - (\theta_x y + \theta_y x)_{i=5,6}. \quad (5.2)$$

The atmospheric contribution to the wavefront is expected to vary much more quickly in time than the optical contribution. We use the approximation that, over a sequence of many out-of-focus exposures, the optical contribution to $a_i(x, y)$ is static and the atmospheric contribution is *dynamic and random*; thus, we recover the optical contribution to each coefficient by taking an average over a series of k exposures:

$$\bar{a}_i(x, y) = \langle a'_i(x, y) \rangle_{k \text{ exposures}}. \quad (5.3)$$

Subtracting these contributions from the coefficients for each exposure, we recover coefficients for the wavefront that are nominally due to only atmospheric effects:

$$a''_i(x, y) = a'_i(x, y) - \bar{a}_i(x, y). \quad (5.4)$$

I did not start with raw out-of-focus exposures; Aaron Roodman kindly supplied processed wavefront data in the form of the coefficients a_i , a'_i , δ_i , θ_x , θ_y , and $a''_i(x, y)$, for four sequences of 1000 30-s exposures. These data are binned over the focal plane into 18 bins per DECam CCD or a total of 1,116 bins over the 2.2-degree field of view. We reconstruct the atmospheric PSF for each bin by 1) calculating the wavefront at x, y with Equation 5.1 using the $a''_i(x, y)$ coefficients, and then 2) taking the Fourier transform of the wavefront to determine the PSF at x, y (see Equation 1.53). We set the focus term ($i = 4$) to zero during this process in order to get in-focus PSF images rather than donut images. From these images we measure PSF size and shape, using the HSM algorithm to calculate weighted second moments as in Chapter 4.

Results for two of the DECam exposures are shown in Figure 5.1. The PSF parameters plotted for one exposure in the top panels show striation patterns characteristic of the atmospheric PSF, whereas the exposure in the bottom panels shows no such pattern. The magnitude of the PSF parameters shown here are smaller than we might expect from simulations, with variations on the

¹Zernikes are a basis set of polynomials commonly used in optics to describe circular surfaces. We use the Noll convention here. Polynomials of increasing order i correspond to effects like focus, coma, astigmatism, etc.

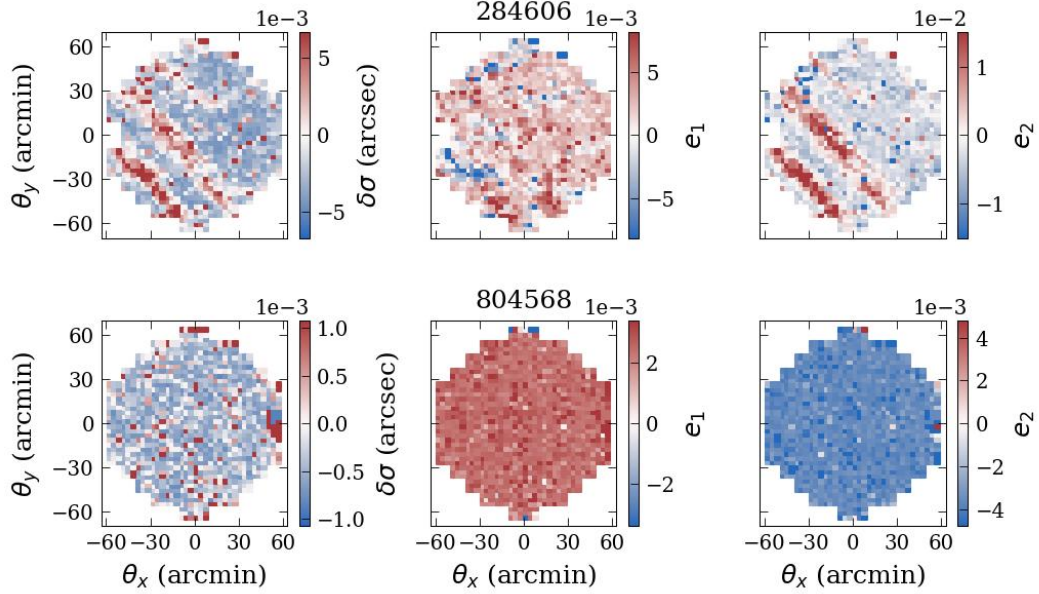


Figure 5.1: PSF moments, mapped across the field of view, for two exposures (284606 from 20140212 *i* band, top; 804568 from 20181218 *z* band, bottom); PSF size variation from mean size (left), PSF ellipticity component e_1 (middle) and e_2 (right). Moments are measured from reconstructed atmospheric PSFs using HSM.

order of 10^{-3} rather than 10^{-2} (see for example Figure 4.3, or find the average expected magnitudes given the correlation function amplitudes in Section 4.3 and Section 4.4). This could indicate that the optical-subtraction process described above has also removed some of the atmospheric contributions to the wavefront.

The four series of exposures that we analyze in this chapter were recorded in different bands (*i, g, r, z*) on three different nights. Measurements of the observing conditions—seeing and wind—over the course of observing are shown in Figure 5.2, plotted as a function of time from the start of the exposure series. Atmospheric seeing data from the DIMM at CTIO was not available for all the nights. Based on available data, the dominant wind direction at CTIO appears to be approximately from the North.

5.3 Anisotropic two-point correlation function results

Following the analysis in Section 4.4, we calculate anisotropic two-point correlation functions for PSF size, $C(\theta, \alpha)$, and shape, $\xi_+(\theta, \alpha)$ from the PSF moments of each DECam exposure. The binning of the data across the field of view allow only a coarse estimate of the 2PCF; we probe intermediate scales of ~ 10 arcmin, which is in the middle of the ~ 1 to 100 arcmin range analyzed in Section 4.4.

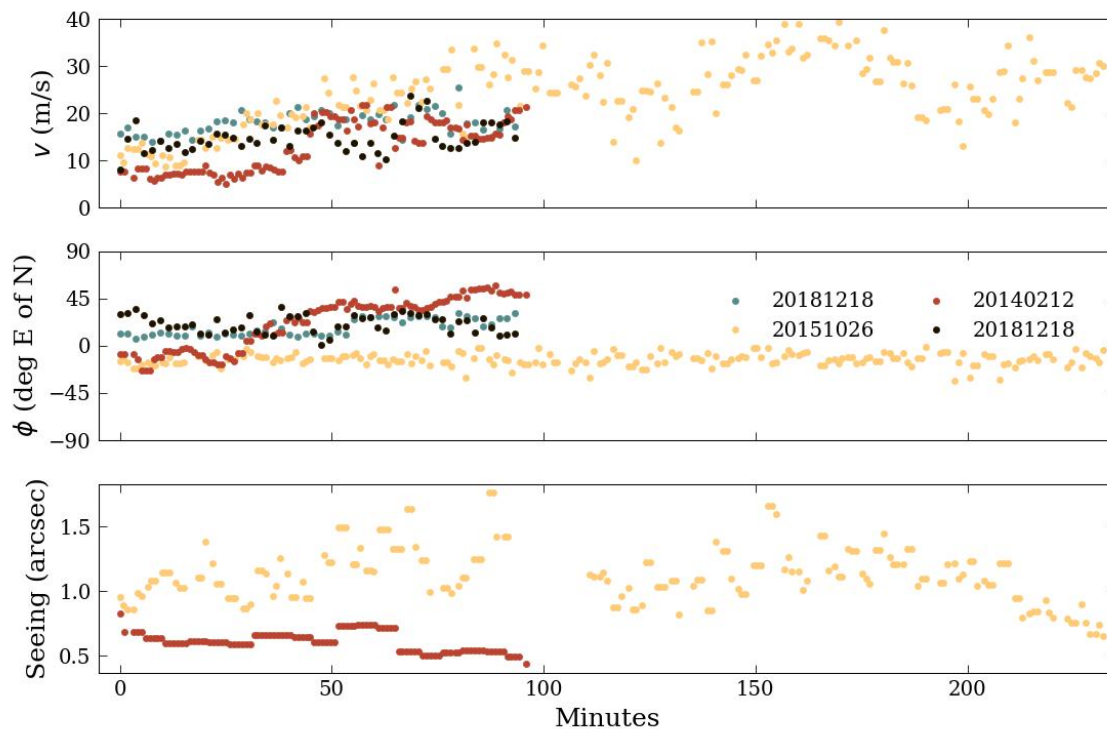


Figure 5.2: Telemetry values recorded at CTIO as a function of time since beginning of each series of exposures. The panels from top to bottom show wind speed v , wind direction ϕ , and seeing FWHM. Each series of exposures is identified by a number representing the date (yyyymmdd) and denoted by different color points: series 20140212 i band in red, 20151026 r band in orange, 20181218 g band in teal, and 20181026 z band in black.

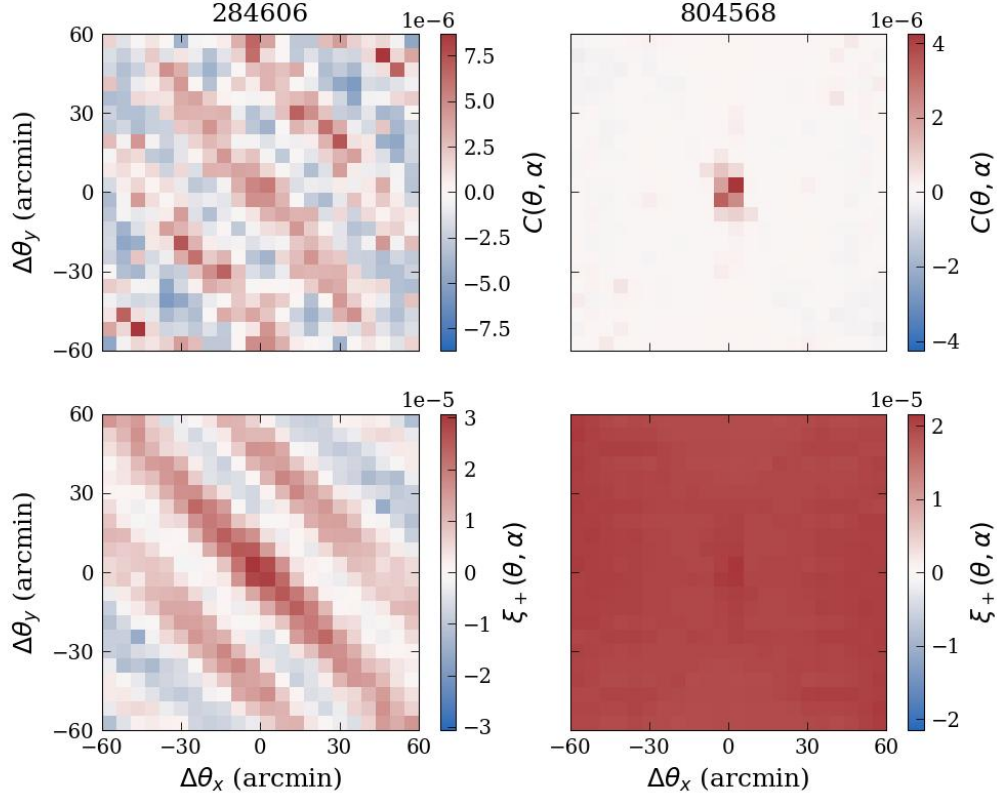


Figure 5.3: Two-point correlation functions of PSF size ($C(\theta, \alpha)$, top) and shape ($\xi_+(\theta, \alpha)$, bottom) for PSF moments shown in Figure 5.1 (284606 from 20140212 *i* band, left; 804568 from 20181218 *z* band, right).

The 2PCFs for the example exposures in Figure 5.1 are shown in Figure 5.3. The left panels of Figure 5.3 show the 2PCFs for the exposure in the top panels of Figure 5.1; we see the striation patterns in PSF parameters reflected in the structure of the 2PCF. The 2PCFs in the panels on the right (corresponding to the exposure shown in the bottom panels of Figure 5.1) have barely any structure, reflecting the fairly uniform values of the PSF parameters. As expected from the magnitude of variation in the PSF parameters shown in Figure 5.1, the correlation function amplitudes are smaller, especially for PSF size, than we see in simulated atmospheric PSFs.

These 2PCFs are fairly representative of the whole sample; a few 2PCFs have fairly obvious structure similar to those in the panel on the left, while others are mostly isotropic. The method of finding α_{\max} from Section 4.4 does not behave well on these images—many of the 2PCFs do not have a dominant orientation that can be identified, and others exhibit a strong correlation in stripes that

do not pass through the center of the 2PCF map. We conducted a simple hand scan of the 2PCF plots to identify those with visible anisotropies, and coarsely estimated their α_{\max} orientation. We matched any evident orientations of anisotropies in the images with the closest of ten orientations that were overlaid as lines on the images at evenly separated angles. Two participants independently scanned all images and their results were broadly consistent.

For those of 2PCF images identified as having visible anisotropy ($\sim 40\%$ of images), we denote the hand-estimated orientation of this anisotropy as $\hat{\alpha}_{\max}$. The results are summarized by the unfilled histograms in Figure 5.4 for the orientations of PSF size correlations (top) and PSF shape correlations (middle) along with the associated ground wind directions (shaded histograms). Scatter plots of orientations $\hat{\alpha}_{\max}$ versus the wind direction are in the bottom row. The estimated orientations $\hat{\alpha}_{\max}$ have been projected from the image plane to local ground coordinates according to the altitude and azimuth of each exposure. We see little evidence of correlation in any of the exposure series.

The 2PCFs of individual exposures are noisy and many have very low signal. Stacking the 2PCFs from multiple exposures (similar to the ensemble analysis from Section 4.4) might boost a signal of some common correlation if such a signal exists. Stacking the simulated 2PCFs in the previous chapter was straightforward since all exposures were simulated in a common coordinate system (local ground coordinates) since simulations were run at zenith. In data we need to account for the varying telescope pointings throughout the image series. The 20151026 *r*-band data set was an engineering test designed to measure the wavefront of the optical system as the telescope swept through a wide range of altitude angles; we do not pursue stacking the 2PCF plots for these images. The top two panels in Figure 5.5 show the altitude and azimuth angles for the remaining three image series. We estimate the change in the image coordinate system from these altitude and azimuth angle ranges by projecting a direction, fixed in ground coordinates, into image coordinates for each telescope pointing. These results are shown in the bottom panel of Figure 5.5. We determine that the $\approx 10\text{-}20^\circ$ change over the course of each exposure series is an acceptable amount of rotation for our stacking purposes, given this is roughly the size of one angle bin in Figure 5.4.

Figure 5.6 shows the results of stacking the 2PCFs of PSF shape (top) and PSF size (bottom) for each of the three image series. The solid black lines correspond to the ground wind direction, averaged over the course of each observation and projected to the coordinates of the average telescope pointing during observation. There is no consistent relationship between the angle of these lines and the orientation of the 2PCF anisotropies. The scale of the angular separations probed by these 2PCFs are in between those analyzed in Section 4.5, so it is not clear whether the ground wind direction or the free-atmosphere wind direction is the relevant point of comparison. We do not have measurements of the free atmosphere wind direction, but we can use the weather forecasting data as an estimate—we use a wind profile from the night of 2018/12/18 from the NOAA Global Forecasting System (GFS). The average direction from this profile is shown in all panels in Figure 5.6 as a dotted black line, projected in the same way as the ground wind direction. There is no consistent

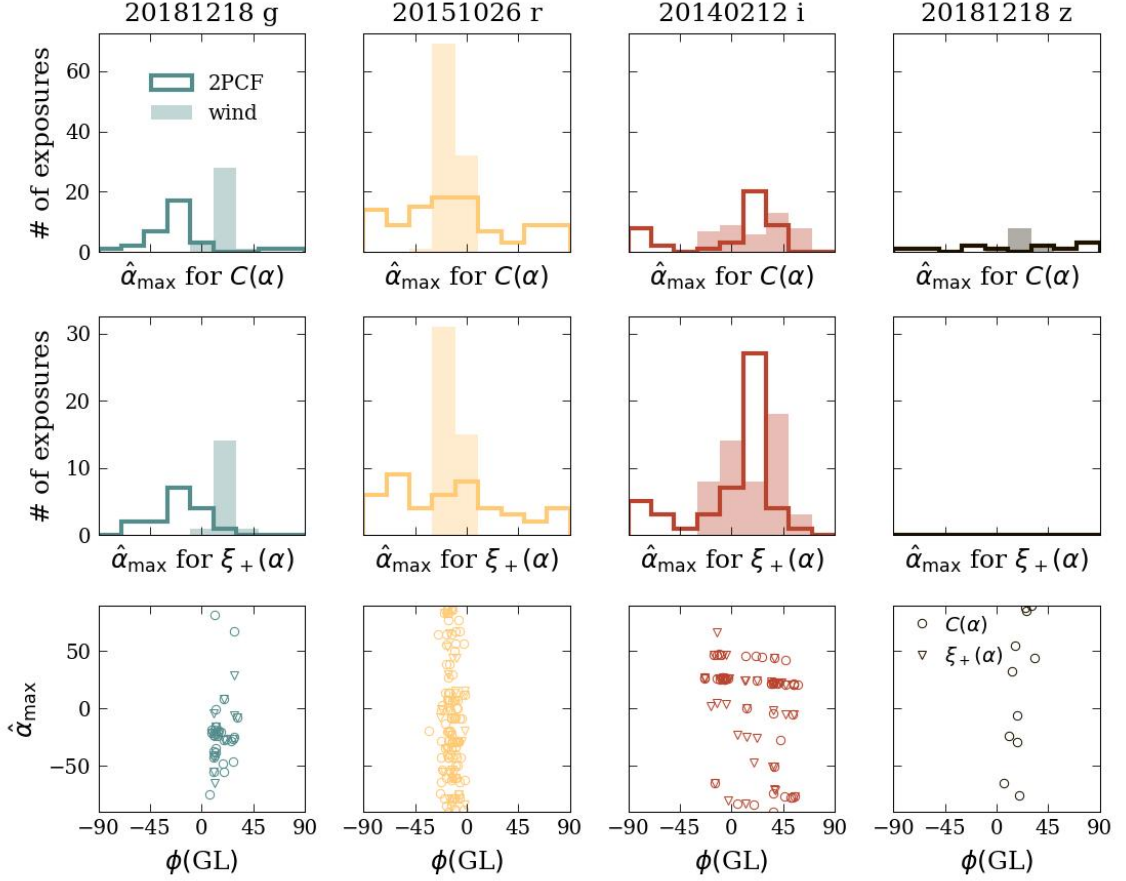


Figure 5.4: Results of a hand scan of anisotropic 2PCFs to estimate α_{\max} , the direction along which the correlation is maximum. The top and middle rows contain histograms of $\hat{\alpha}_{\max}$ for PSF size and shape, respectively. These are overlain with the distribution of wind directions corresponding to the exposure times of the data. Panels in the bottom row are scatter plots of $\hat{\alpha}_{\max}$ versus ground wind direction $\phi(\text{GL})$, for PSF size (circles) and shape (triangles). Each column of plots corresponds to one series of exposures.

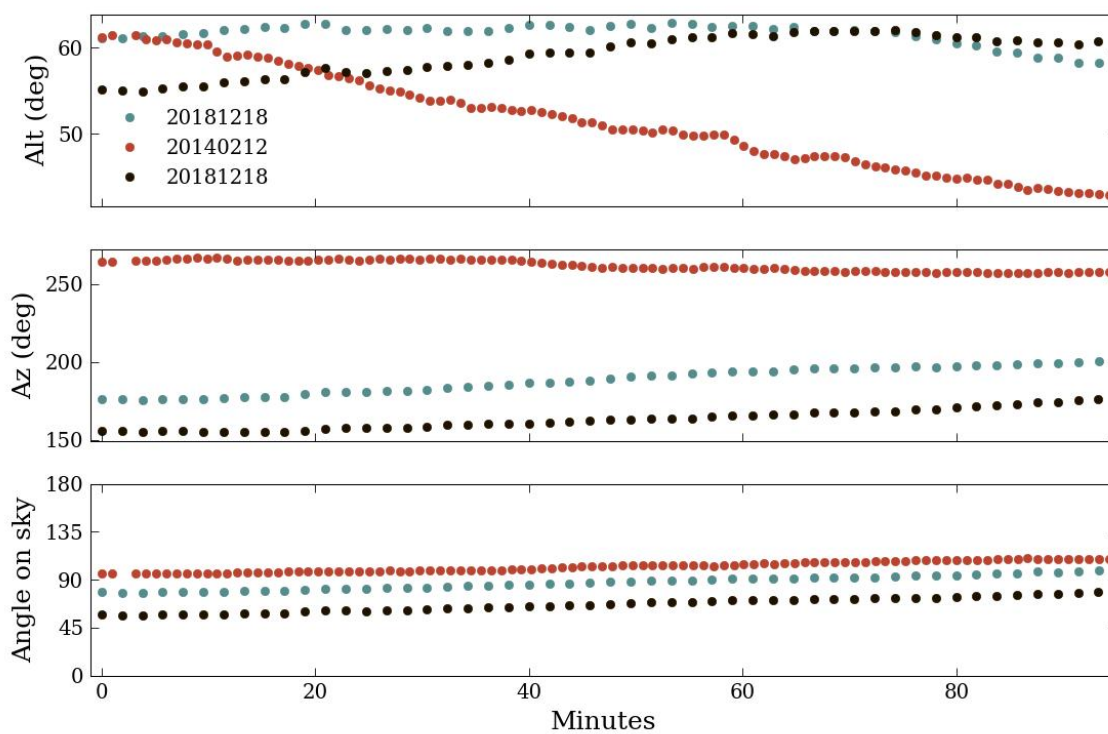


Figure 5.5: Telescope pointing at the time of each exposure and the resulting projected angle on the sky. Top and middle panels show the altitude and azimuth as a function of minutes since start of exposure series.

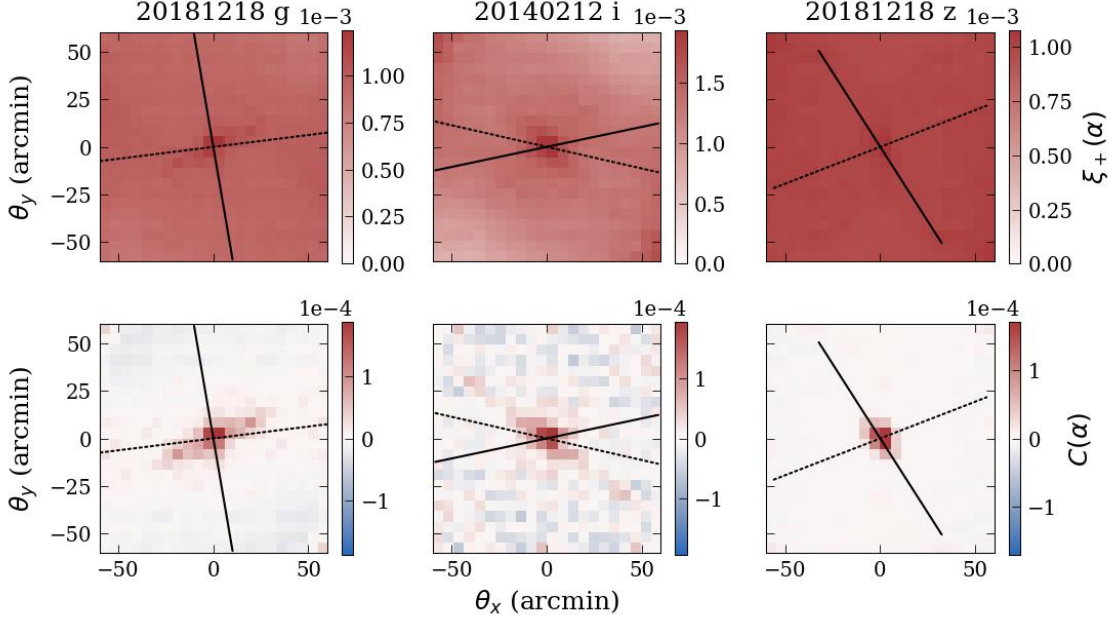


Figure 5.6: Stacked two-point correlation functions of PSF shape (top) and size (bottom) for three exposure series (20181218 g band, left; 20140212 i band, middle; 20181218 z band, right). Overlaid lines correspond to wind directions: solid black lines are the average ground layer wind direction, projected to the sky according to the average telescope pointing for each series; dotted lines for all panels is the 3AM wind direction on 2018/12/18 from the NOAA GFS Analysis, averaged across the free atmosphere and projected to the sky according to the average telescope pointing for each series.

relationship between this estimated free atmosphere wind and any anisotropies the stacked 2PCF images.

5.4 Discussion

The analysis of two-point correlation function anisotropy in DECam wavefront data showed no evidence of the dependence on wind direction found in Chapter 4. We discuss here the significant limitations of this analysis. As briefly discussed in Section 5.2, it is likely that the data processing, which aimed to remove static components of the wavefront, also removed wavefront contributions from the atmosphere that were common across the field of view. This would include effects on large scales from the ground layer, which were largely responsible for the persistent anisotropies found in simulation; thus it is perhaps not surprising that the ground wind direction was not predictive of anisotropies in the 2PCFs from the processed DECam data. On the small scales, we are limited by the binning of the data we received; each bin is ≈ 3 arcmin wide, so we are not sensitive to

correlations at the separation distances found to be related to the free atmosphere wind direction.

Given these complications in the data, it is not possible to draw conclusions about the presence, or lack thereof, of correlations between the two-point correlation functions of the PSF at CTIO with the wind directions at the time of observation.

Future work

It is possible that the atmospheric signal removed in the procedure described in Section 5.2 can be added back into the data, after static contributions have been removed. This signal would be in the δ_i terms – i.e., the average Zernike amplitude for each exposure. However, the out-of-focus images described here were recorded with the active optics system (necessarily) turned off, which means the δ_i terms may contain significant optical distortions to the wavefront as well as the atmospheric components. The optical wavefront estimated from the wavefront decomposition described in Section 5.2 can also be used to remove the optical PSF from in-focus (active optics on) images (Davis et al., 2016); an analysis of this type of data may be promising to explore.

Chapter 6

Temporal evolution of the PSF at Cerro Pachon

6.1 Motivating questions

The work described in the thesis so far has focused on exploring the spatial correlations in the PSF across a wide field of view—in particular the spatial correlations we might expect for Rubin Observatory. Temporal correlations in the PSF are also of relevance. The following two related questions motivate the analysis described in this chapter.

1. Rubin has shorter exposures than previous wide-field galaxy surveys—30 s for each visit.¹ The atmospheric PSF is very dynamic and featureful on short time scales due to the nature of turbulent structure in the atmosphere. The first step to interpolating and correcting for the PSF, before measuring shear from galaxy images, is to model the size and shape of the PSF itself. How quickly during an integrating exposure do these dynamic features “average out” in the atmospheric PSF to leave a smooth and stable intensity profile?
2. On what time scale does the PSF stay correlated with itself? In other words, how much temporal correlation (in contrast to spatial correlation across the focal plane) do we expect between the values of the PSF parameters in back-to-back Rubin exposures?

To explore these question we use short-exposure images from the Zorro speckle imager, an instrument mounted at Gemini South observatory (near Rubin, on Cerro Pachón). The Zorro speckle imager on Gemini South (and its identical companion, 'Alopeke, on Gemini North) is used primarily

¹It is still to be decided whether these 30-s images will be divided into two 15-s “snaps” or a single 30-s exposure.

for validation and characterization of exoplanet targets.² Speckle imaging enables direct observation of binary star systems with separations from 0.02 to 1.2 arcsec; via speckle interferometry, the individual stellar masses and binary orbital periods can be estimated from their color, separation, and position angles. Speckle instruments deliver images with high temporal and spatial resolutions; these features are necessary for the interferometric studies of exoplanet candidates, but also provide the opportunity for novel analyses of the turbulent atmosphere on short time scales.

We describe the Zorro instrument and data in more detail in Section 6.2. In Section 6.3, we describe how we simulate speckle observations using PSF-WEATHER-STATION parameters. The PSF analysis is described in Section 6.4. The main results are presented in Section 6.5, followed by further studies in Section 6.6. Conclusions are summarized in Section 6.8.

6.2 Fast PSF images from the Zorro speckle imager

In this section, we summarize the thorough description of the Zorro speckle imager and capabilities from Scott et al. (2021). The Gemini telescope provides an input beam to Zorro; the beam is collimated and split, through a dichroic element (break at 674 nm), into “red” and “blue” beams—see Figure 6.1. These beams each pass through a filter wheel and into a camera; the simultaneous operation of two optical channels is a key aspect of the instrument for exoplanet follow-up.

A speckle image from Zorro, with the settings for the data we will describe in this thesis, consists of a 256×256 pixel region, centered on the star (or binary stars) of interest. With the approximate pixel scale of $0.01 \text{ arcsec pixel}^{-1}$, these images are $\sim 2.5 \text{ arcsec}$ on a side. In addition to the high spatial resolution of the detectors, the Zorro instrument is capable of very short exposures due to the camera specifications. The red and blue beams are imaged by two identical Andor iXon Ultra 888 EX EMCCD cameras. These have electron-multiplying CCDs (EMCCD): before CCD readout, photoelectrons in these detectors are accelerated in a high-voltage region containing hundreds of electrodes. The impact ionization creates an avalanche of additional charge, and the signal can be amplified up to several thousand times. In electron-multiplying mode the detectors are sensitive to single photons and have read noise of < 1 electron (because the exposures are so short, the dark current is negligible). The CCDs are frame-transfer CCDs, operating shutterless with high-speed readout. In this way, Zorro can take series of 60 ms exposures with fast 20 MHz readout.

Transmission curves for different components of the instrument are shown in Figure 6.2; the observations described in this thesis use the narrow filters centered at 562 nm and 832 nm. Curves for the dichroic element, the Gemini telescope, and the EMCCDs are also shown.

²The exoplanet targets are from the NASA Kepler, K2, and Transiting Exoplanet Survey Satellite (TESS) missions, and those discovered by precision radial velocity and other measurements.

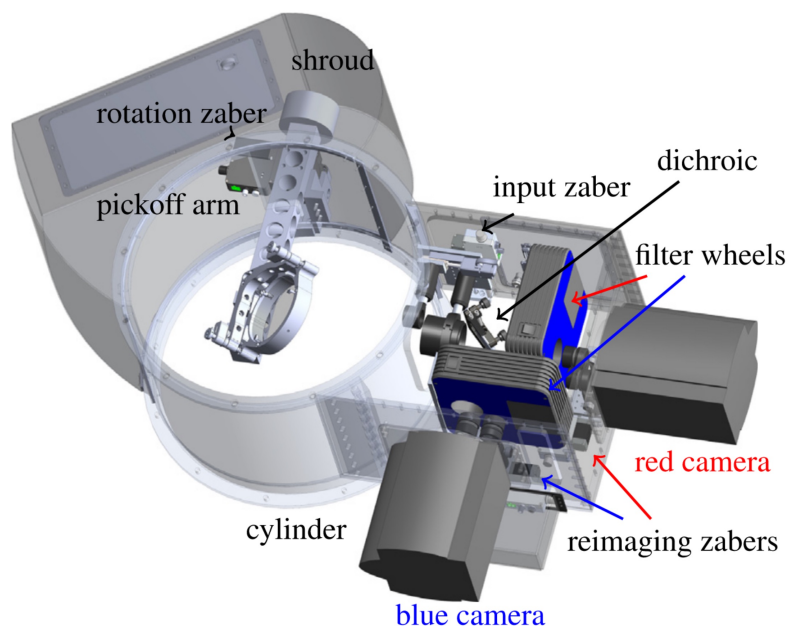


Figure 6.1: Instrument diagram showing components of the Zorro imager with main parts labeled. The dichroic element splits light into two beams, which pass through red and blue filters and are then imaged by the EMCCD cameras. Figure from Scott et al. (2021).

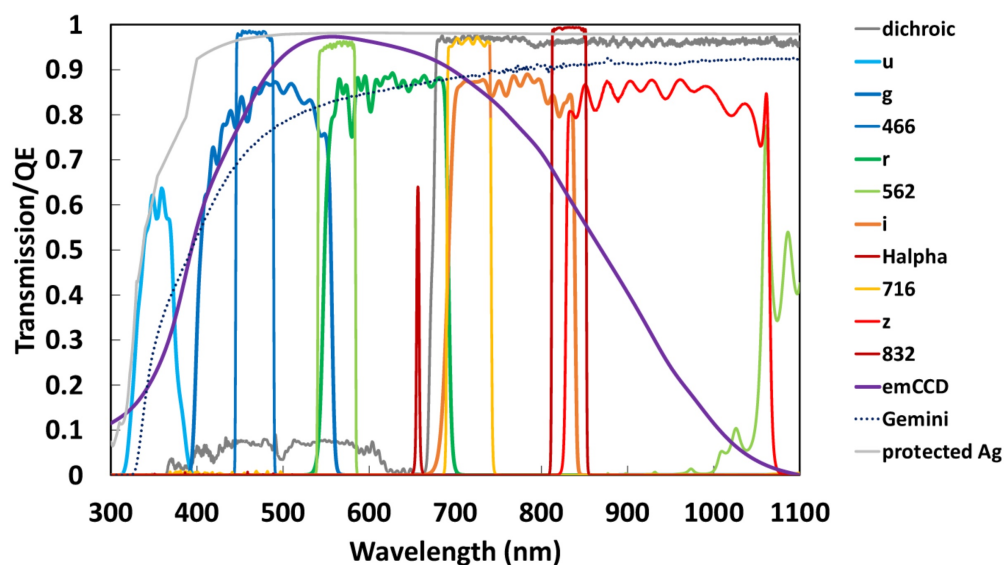


Figure 6.2: Transmission curves for various components of the Zorro instrument mounted on Gemini South. The dichroic element divides the incoming beam at 674 nm into a blue channel (dark grey curve) and a red channel. Each channel is instrumented with a filter wheel with two narrow-band filters for speckle observations; the data discussed here use the filters with transmission curves centered at 562 nm (light green) and 832 nm (dark red). Figure from Scott et al. (2021)

6.2.1 Zorro observations and targets

For the case of speckle interferometry, a single observation corresponds to data taken during one minute of observation: 1000 of these short 60 ms-exposure images. For bright targets and point-source calibration stars (which we refer to as “PSF standard stars”), one observation may be sufficient for the interferometry analysis; fainter targets (to be classified as either single or binary systems) require a set of multiple minute-long observations to reach the desired signal-to-noise ratio Howell & Furlan (2022). We will use *observation* to refer to a series of exposures, regardless of how many exposures (1000 or more) it is composed of. The data used here correspond to

- 117 observations of PSF standard stars each with an observation length of 60 s;
- 61 observations of fainter science targets each with an observation length of 300 s.

All data were taken between May 2019 and November 2019. Both observation types are useful in different ways: the PSF standard stars are bright, confirmed single sources, whereas the longer observations of science-target stars give us a longer sampling time of the atmosphere. We have simultaneous observations in both 562 nm and 832 nm filters so we have 2×117 and 2×80 observations of PSF standards and science targets, respectively; however, because these observations are sampling the same atmosphere, they are not independent observations of turbulence structure.

6.2.2 Data selection and processing

The 117 PSF standard-star and 61 faint science-target observations correspond to data that passed our selection criteria. Since we will eventually fit a model to the intensity profile of these PSFs images to estimate the PSF parameters (size and shape), we require that we have information about the wings of the PSF. Therefore, we select data based on the criteria that the PSF size (FWHM) plus the standard deviation of its centroid motion across the observation does not exceed half the image (128 pixels). In addition, we select for observations at airmass < 1.3 .

We describe two ways that we artificially integrate the exposures before extracting the PSF parameters from the images. In the first method, we stack exposures from a single observation to produce images of the PSF for a range of integration times. Figure 6.3 shows a selection of these integrated images from a randomly selected observation of a PSF standard star. Speckles, formed by refraction through turbulent structure in the atmosphere, are evident in the single exposure image in the top left panel. As the integration time increases the speckle structure is washed out and a smoother PSF profile forms.

The second method is demonstrated in Figure 6.4 and Figure 6.5. Rather than accumulating an integration over all exposures within an observation, we bin the 60 ms exposures into shorter consecutive time intervals and stack the images in each time interval, forming images for back-to-back exposure times less than the total for the observation. In Figure 6.4, for example, the same

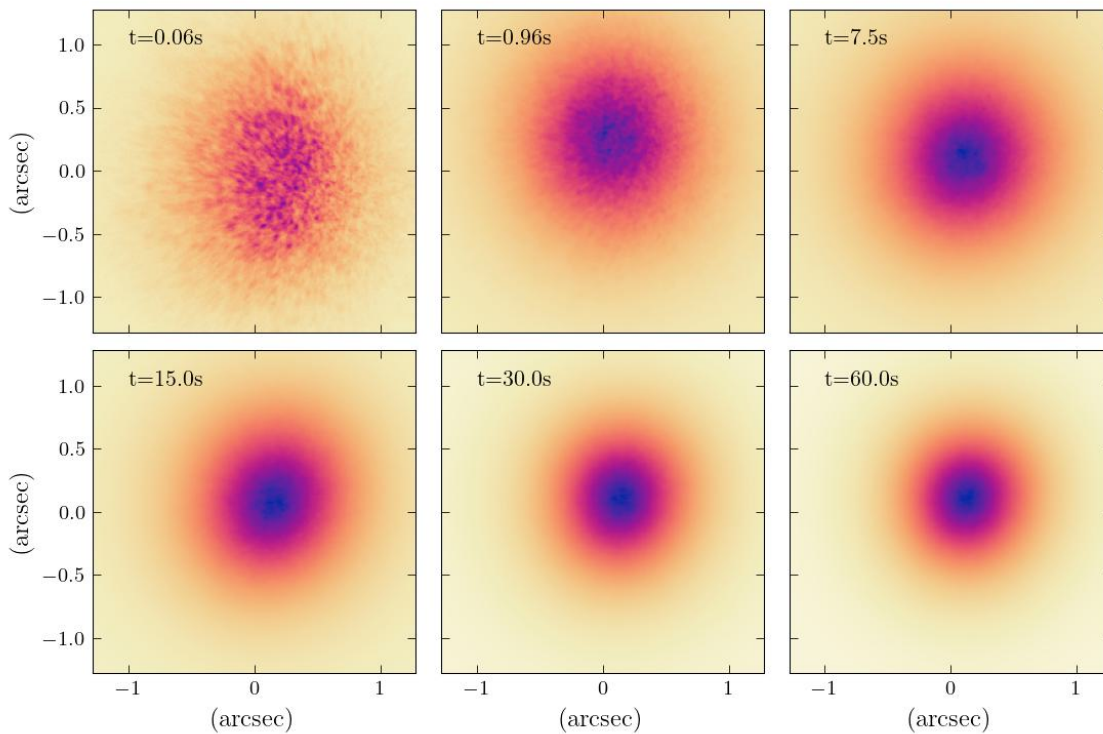


Figure 6.3: Example Zorro observation of a PSF standard star in the filter centered at 562 nm. Individual exposures have been stacked to produce a view of the PSF for integrated exposure times between 0.06 s and 60 s.

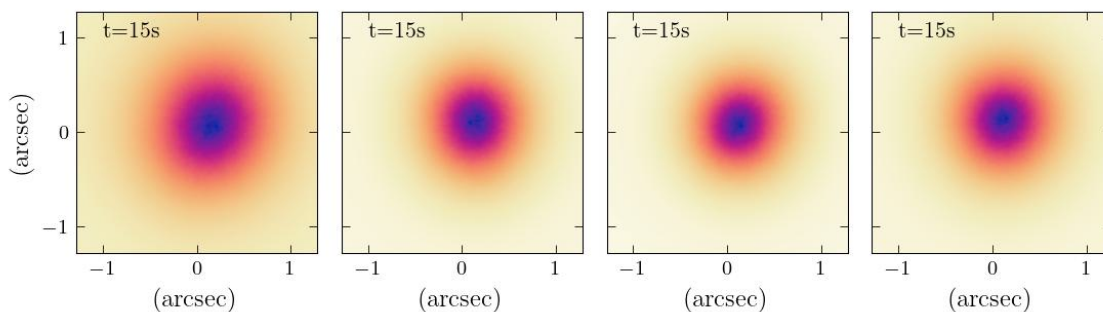


Figure 6.4: Example Zorro observation of a PSF standard star in the filter centered at 562 nm. Individual exposures have been stacked to produce PSF images corresponding to four contiguous 15-s exposures.

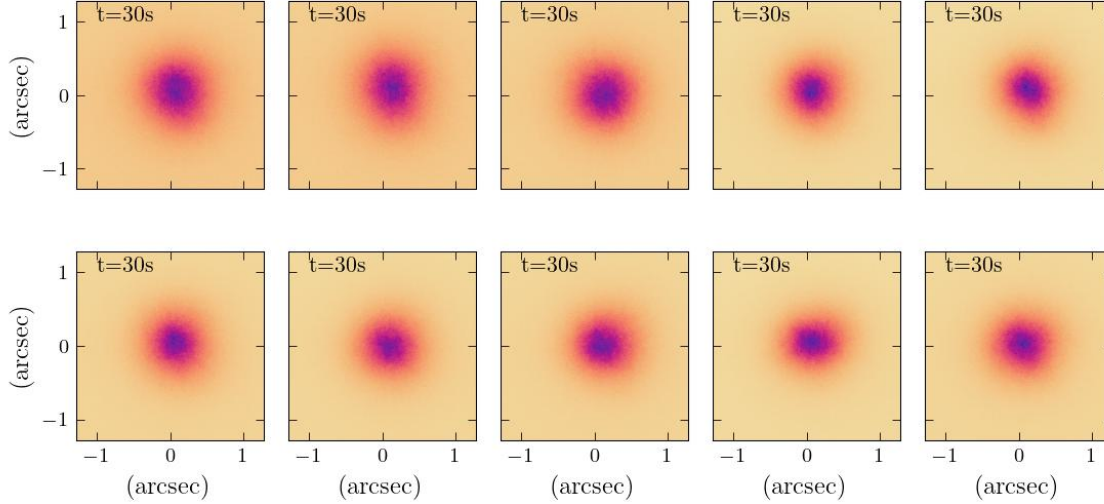


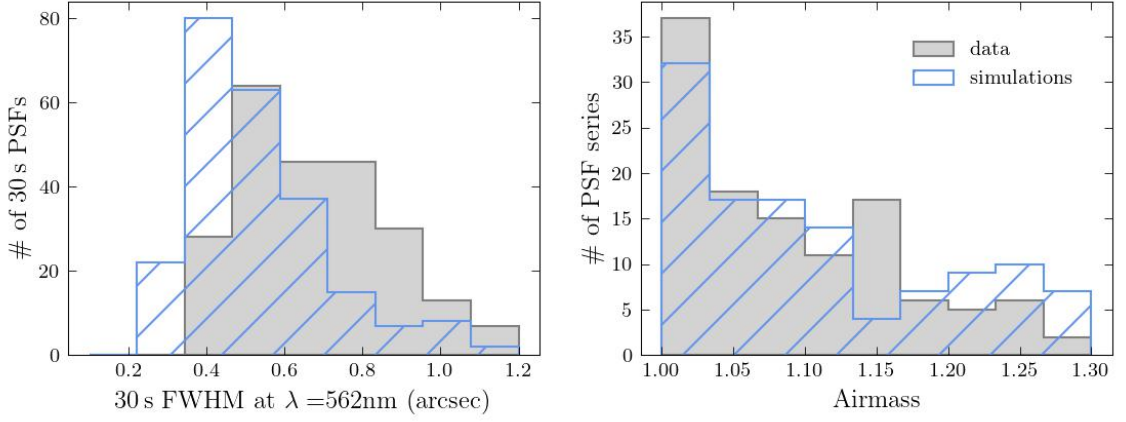
Figure 6.5: Example Zorro observations of a faint science target in the filter centered at 562 nm. Individual frames from a 300-s observation have been stacked to produce PSF images corresponding to ten contiguous 30-s exposures.

PSF standard observation as Figure 6.3 is binned into four 15-s integrated exposures. For each consecutive pair of 15-s images, we can see subtle variations in the shape of the PSF. Figure 6.5 shows the same procedure carried out on a randomly chosen observation of a faint science target, this time with each image corresponding to a 30-s integrated exposure.

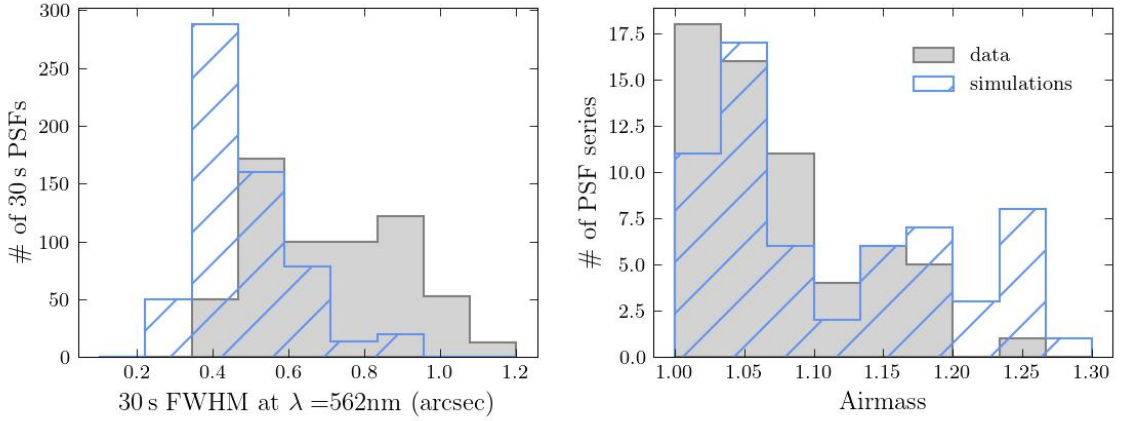
In Figure 6.6 we show the distributions of measured PSF FWHM and airmass for all the selected data (grey histograms), separated into PSF standard stars (top panels) and science targets (bottom panels). The PSF size is measured on each 30-s integrated exposure within the observation; i.e., twice for the PSF standard stars and ten times for the science targets. In the next section we will describe the simulated observations whose distributions are shown in blue hatched histograms.

6.3 Simulated observations

We simulate Zorro-like observations with the GalSim implementation of the simulation method described in Chapter 3, using PSF-WEATHER-STATION input parameters corresponding to the location of the Gemini South observatory and the time period between May 2019 and November 2019 over which the data were taken. The simulated images match the specifications of the Zorro images: observation series of 60-ms exposures with 0.01 arcsec pixel scale. The Gemini mirror diameter (8.1 m) and obscuration (13%) are used as inputs to the simulation. Unlike the PSF simulations described



(a) PSF standard stars



(b) Faint science targets

Figure 6.6: Distributions of PSF size (left) and airmass (right) for Zorro observations (grey shaded histograms) and simulated observations (hatched blue histograms) of PSF standard stars (a) and science targets (b). Values of airmass are selected only at the beginning of each observation, while PSF sizes are measured every 30 s.

in Chapter 4, for which we used the “photon-shooting” method (ray-tracing approximation) in generating the PSF images, the simulations here use the full Fourier-transform method; including the wave nature of light is necessary to produce the characteristic speckle patterns of short exposures. The Fourier method is slower and more computationally intensive than the photon-shooting method; however, as we are simulating single stars here, rather than the 50k stars per exposure in Chapter 4, the Fourier method is very computationally tractable.

The simulations were run *off* zenith, with airmass values chosen from the data distribution (grey histograms in Figure 6.6) and azimuth chosen randomly (in simulation the azimuth angle is only used in the transformation of wind parameters from local ground coordinates to pixel coordinates—see Section 3.3.2). The outer scale was, for each of six phase screens in the simulation, randomly chosen from a log normal distribution with median $L_0 = 25$ m truncated below 10 m and above 100 m. We adjust the total seeing from PSF-WEATHER-STATION outputs to match the measured distribution of atmospheric seeing values at Cerro Pachón from the LSST Science Book (LSST Science Collaborations & LSST Project, 2009), and use a Gaussian with FWHM= 0.35 arcsec to account for optical and sensor contributions to the PSF.

The simulations needed for this study are of significantly longer exposure times than those in Chapter 4: 60 s and 300 s compared to 30 s. This has the potential to be a computational problem depending on the phase screen setup of the simulation. As described in Chapter 3, the simulations use phase screens (2D arrays) that are translated over the telescope aperture, according to some wind speed, during the course of the exposure. The phase screens typically “loop” – i.e., they use periodic boundary conditions (PBC); otherwise, as described in Appendix C, the storage demands of the large arrays can become prohibitive. We explore the effect of the PBC method in Appendix C and determined that this method does not introduce bias in the parameters of simulated PSFs when compared to those from simulations run with phase screens large enough to avoid PBC (no looping). This enables us to simulate observations of 300 s, which would otherwise be too memory-intensive.

The simulated observations are processed in the same way as the Zorro data, including the selection based on size and centroid motion. We ran ~ 150 simulations of both standard-star and science-target observations; we randomly sampled from those that passed the selection step to match the number of datasets.

The distributions of PSF size and airmass for the simulated observations are shown in Figure 6.6 as blue hatched histograms. The simulated PSFs are smaller on average than those in the Zorro data. This could be due to a real optical contribution in the Zorro images that is larger than the optical component we assumed in the simulation – a Gaussian with FWHM of 0.35 arcsec. Another possibility is that there is significantly more image motion in the data than the simulated images—this is described in more detail in Section 6.6.2.

Figure 6.7, Figure 6.8, and Figure 6.9 show simulated data processed in the same ways as the Zorro data in Figure 6.3, Figure 6.4, and Figure 6.5. The most stark difference between the sets of

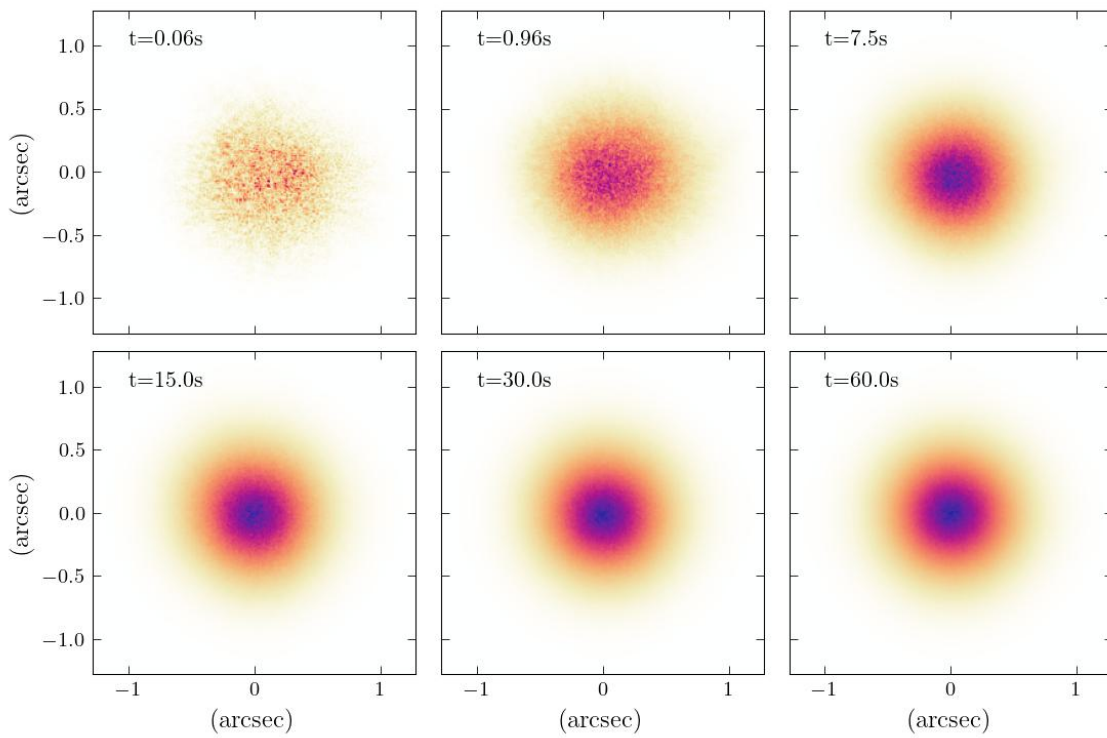


Figure 6.7: Example simulation of a series of speckle images for a wavelength of 562 nm. Individual exposures have been stacked to produce a view of the PSF for integrated exposure times between 0.06 s and 60 s.

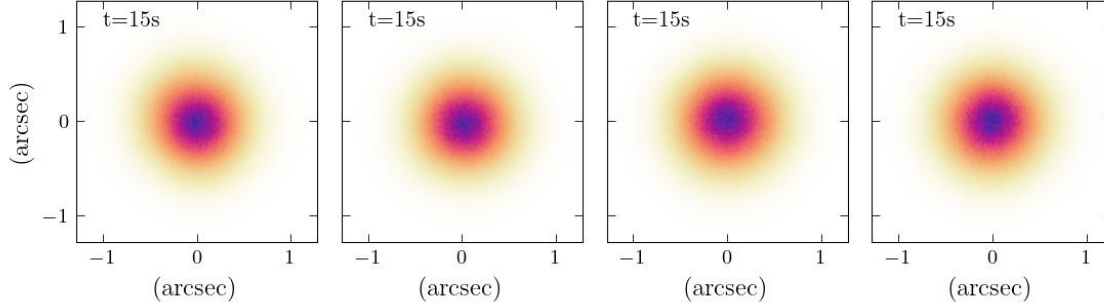


Figure 6.8: Example simulation of a speckle series at $\lambda = 562 \text{ nm}$; individual exposures have been stacked to produce PSF images corresponding to four contiguous 15-s exposures.

figures is the background level—in all cases, the white end of the color scale in the images is fixed at zero. This difference is due to the fact that the simulated observations are generated without a background level. The simulated PSF images in Figure 6.7 retain features from the speckles for longer integrated exposure time than the data—the 15-s image still has some visible patterns within the central core. The 30-s PSF images in Figure 6.9, on the other hand, show much less variability over the 300 s than the data in Figure 6.5.

6.4 PSF parameter fitting

In the previous section, we made some qualitative observations on the behavior of the PSF by inspecting the series of PSF images produced by binning or gradually integrating individual exposures. To draw more quantitative conclusions, we extract information about the PSF parameters—we again use PSF size and shape. In this section, we describe the parameter fitting method and explore its performance and results on both data and simulation.

There are two general ways to estimate the size and shape parameters of the PSF: second moments and profile fitting. In Chapter 4, for example, we used weighted second moments. In this case, we will use profile fitting, which is motivated by the potential to use the high-resolution profiles from the Zorro data to compare the performance of various models of PSF profiles—e.g. Kolmogorov and von Kármán profiles, which are the analytic shapes predicted for an infinite exposure through Kolmogorov or von Kármán turbulence, respectively. We performed the fitting with the LMFIT PYTHON package,³ using GalSim implementations of Gaussian, Kolmogorv, and von Kármán profiles. The fit parameters are the PSF size (FWHM), shear parameters e_1 and e_2 , a background level (1-2 orders of magnitude lower than the signal), total flux, and centroid location x, y on the image. We note that the analyses in this chapter use a slightly different definition of ellipticity than the previous

³<https://lmfit.github.io/lmfit-py/>

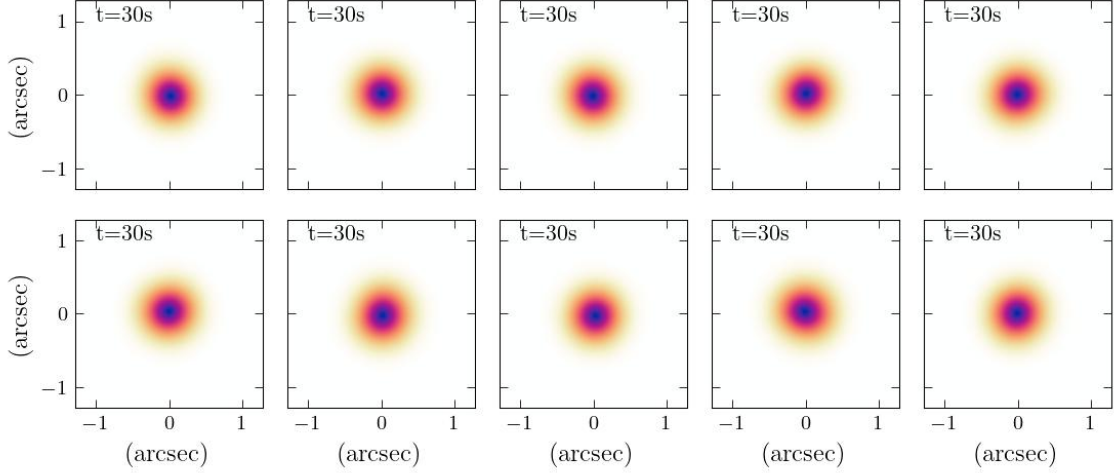


Figure 6.9: Example simulation of a series of speckle observations of a fainter star at $\lambda = 562$ nm. Individual frames from a 300-s observation have been stacked to produce PSF images corresponding to ten contiguous 30-s exposures.

chapters—the magnitude of ellipticity here is defined as $|e| = \frac{1-q}{1+q}$, where q is the ratio of the length of the minor to major axes of the ellipse described by the second-moment matrix (see Equation 1.27). The von Kármán profile has an additional fit parameter, the outer scale L_0 .

We ran preliminary tests with simulated images of a simple Kolmogorov profile. We generated 200 of these images, each with a different realization of Poisson noise, and ran our fitting method on each. The distributions of the best-fit values for the PSF parameters and the background level are shown in Figure 6.10; the vertical line on each plot indicates the input value for the simulations. The PSF shape parameters, as well as flux and centroid x, y (not shown), are robustly recovered, but we see a systematic underestimation of the background level and an overestimation of PSF FWHM. These parameters are degenerate to some degree—their values in the fit are negatively correlated with a correlation coefficient of $\rho \approx -0.4$. Nevertheless, as the bias in estimated FWHM values is only ≈ 0.01 arcsec (compared to the true FWHM of 0.60 arcsec), we deem this fitting method to be accurate enough for our purposes in this chapter.

We ran a similar fitting test with von Kármán images and a von Kármán profile fit; in this case, the L_0 and background parameters are highly degenerate, with a correlation coefficient of $\rho \approx +0.75$. This indicates that L_0 most affects the wings of the PSF. Therefore, the PSF wings are confounded with background level, and we cannot reliably recover a L_0 estimate from these images; we focus on Gaussian and Kolmogorov ($L_0 = \infty$) profile fitting in what follows in this chapter.

We compare the fits to Gaussian and Kolmogorov profiles on the Zorro data and simulations. The Kolmogorov profile is defined in Fourier space, for spatial frequency k in radians⁻¹, as the

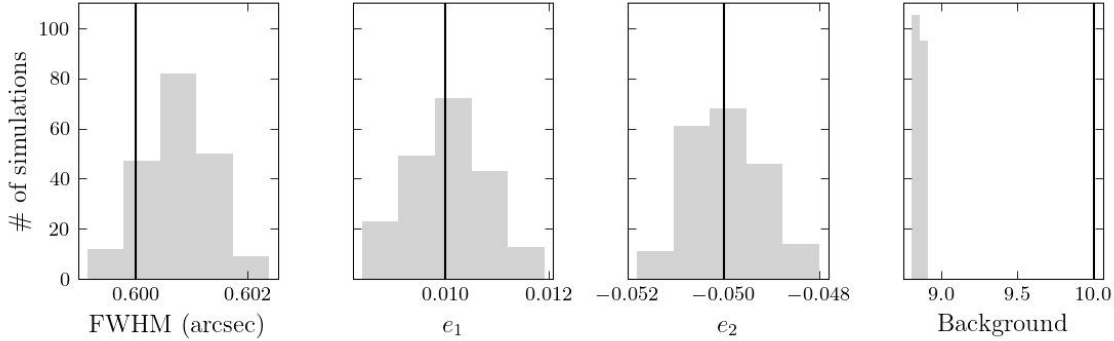


Figure 6.10: Fit parameters recovered from test simulations of Kolmogorov PSFs with different random noise realizations. Black vertical lines indicate true values input to the simulations.

transfer function

$$T(k) \propto e^{-D(k)/2}, \quad (6.1)$$

where $D(k)$ is the wavefront phase structure function (see Section 2.3 and Equation 2.23, or Roddier (1981)):

$$D(k) = 6.88 \left(\frac{\lambda k}{2\pi r_0} \right)^{5/3}. \quad (6.2)$$

In Figure 6.11, we show images for 1-s exposures for a PSF standard star for Zorro data (left top image) and simulations (left bottom image), and the corresponding fits to Gaussian (middle column) and Kolmogorov (right column) profiles. The bottom row in each panel shows the “pull” of the model and data, defined as the difference between the best-fit model and the data (model value minus data value), divided by the square root of the data, where the data is the number of detected photons.

For the fits to data in the top panel of Figure 6.11, the Kolmogorov seems to better capture the bright core and heavier wings in the data than the Gaussian model, whose pull image shows rings of alternating under- and over-estimation of flux with increasing radius. The simulation fit results (bottom panel) are more consistent for the Gaussian and Kolmogorov models but the Gaussian model appears to perform marginally better. Away from the core, we observe higher pull in the simulation than the data; this is likely in part due to smaller pixel values, thus smaller values in the pull denominator, at those locations. In both cases, the pull distributions exhibit the small-scale speckle patterns in the images, which the smooth profiles do not model.

In Figure 6.12 we show the results of fits to 60-s exposures of PSF standard stars. The speckles in the data have been nearly washed out at this exposure time and the observations drawn from Figure 6.11 are more obvious in the pull distributions. The Kolmogorov fit to the data performs better than the Gaussian model, while the Gaussian model better describes the simulated image. One dimensional profiles through the center of each image are shown in Figure 6.13 for the 1-s and 60-s

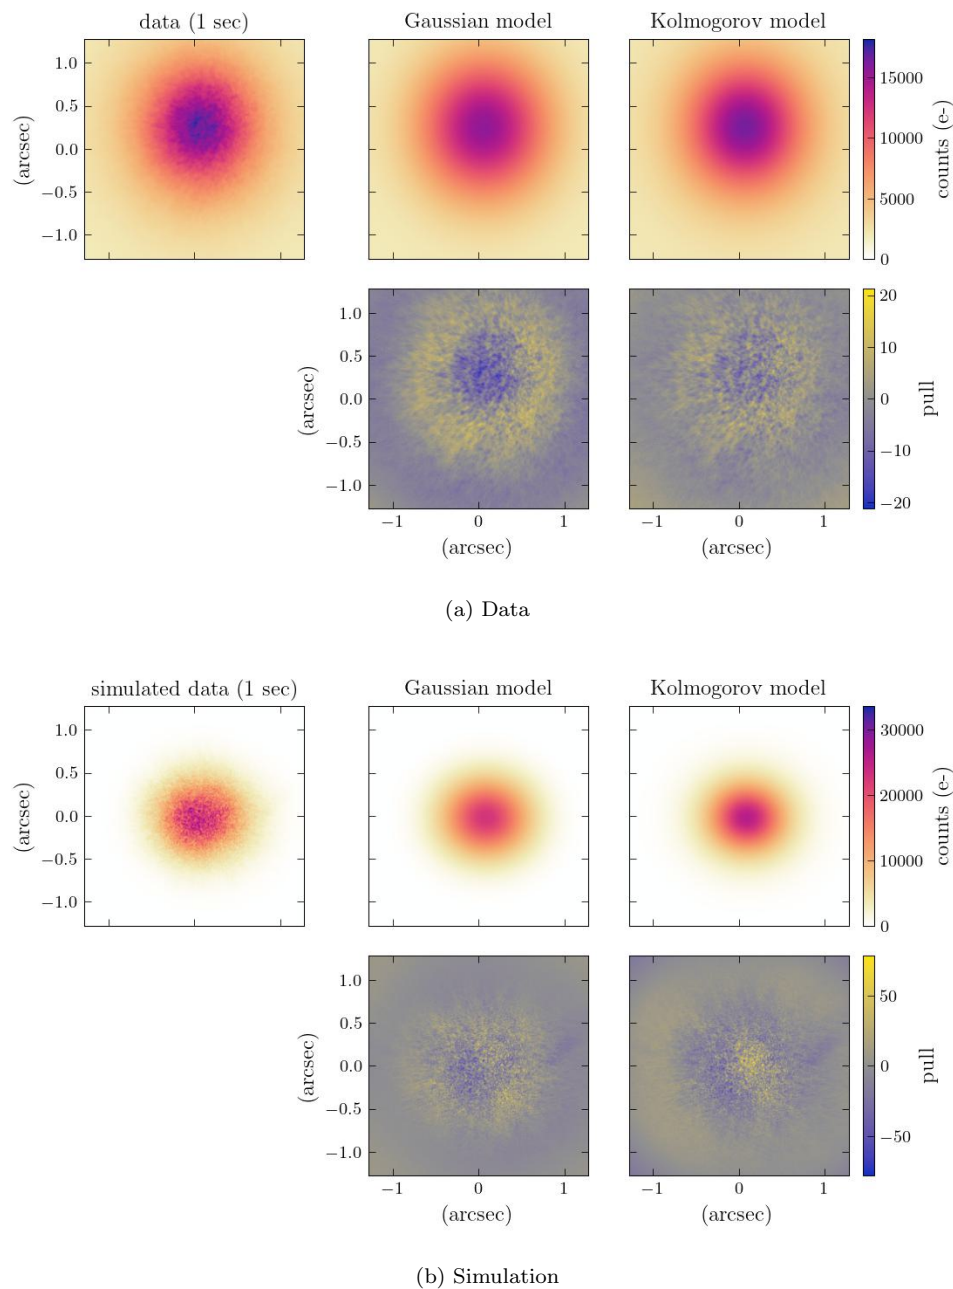
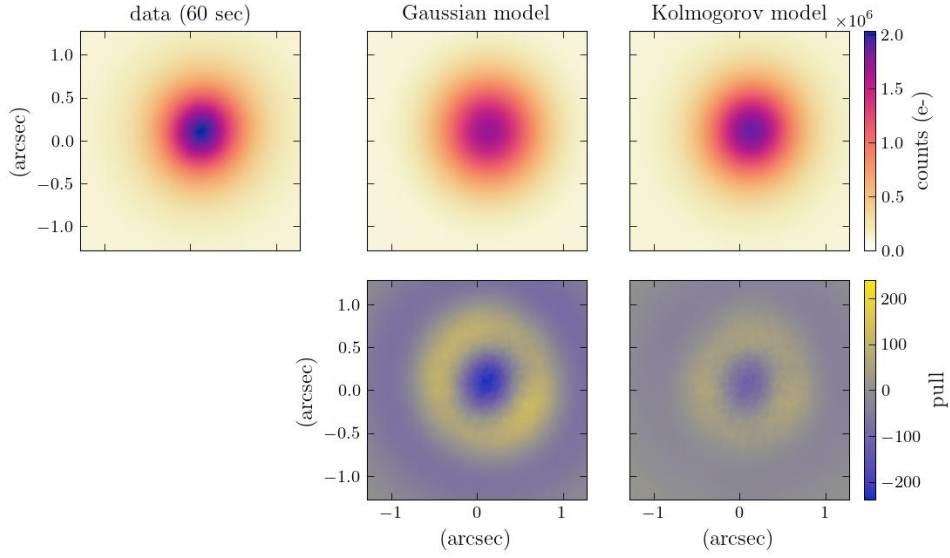
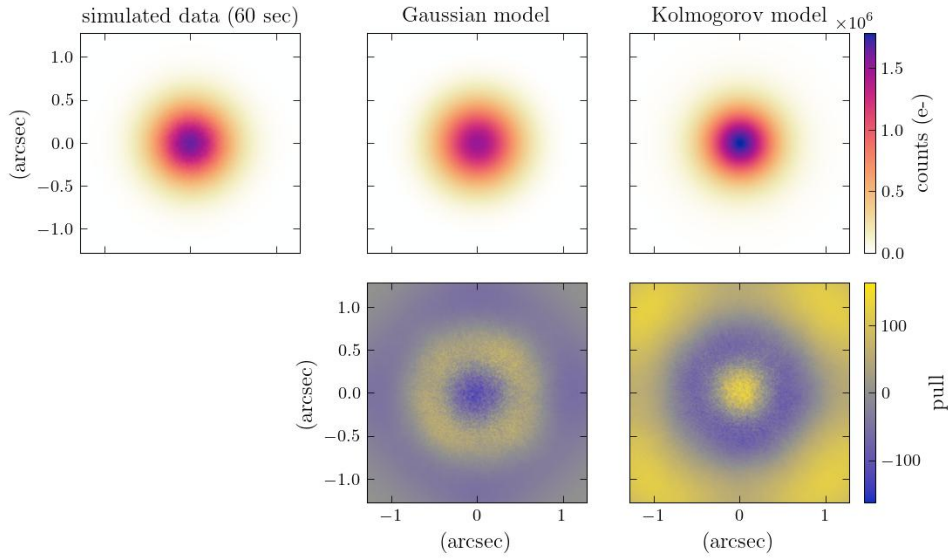


Figure 6.11: PSF profiles for a 1-s exposure (at $\lambda = 562 \text{ nm}$) for (a) Zorro data and (b) simulation. From left to right in the top panels are images of the PSF, the best-fit Gaussian model, and the best-fit Kolmogorov model. The bottom plots show the “pull” for each fit, defined as the difference between the best-fit model and the data, divided by the square root of the data, where the data is the number of detected photons.

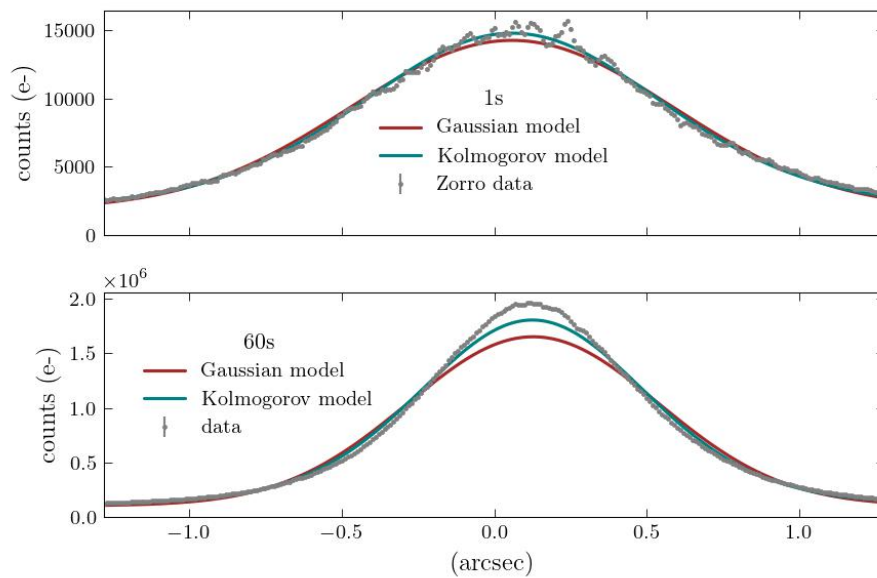


(a) Data

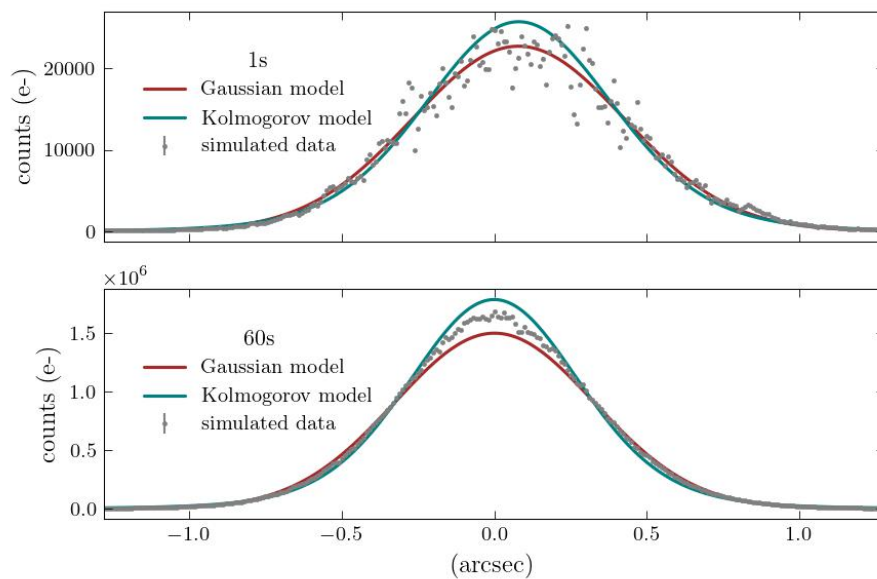


(b) Simulation

Figure 6.12: PSF profiles for a 60-s exposure (at $\lambda = 562$ nm) for (a) Zorro data and (b) simulation. From left to right in the top panels are images of the PSF, the best-fit Gaussian model, and the best-fit Kolmogorov model. The bottom plots show the pull distribution for each model.



(a) Data



(b) Simulation

Figure 6.13: One-dimensional slice through PSF profiles (grey points) and best-fit models (Gaussian, red; Kolmogorov, teal) for (a) Zorro data and (b) simulation. Top and bottom panels for each of data and simulation show the PSF profiles for 1-s and 60-s exposures, respectively.

images for data (top panels) and for simulation (bottom panels). These show more clearly the width and height of model profiles compared to the measured or simulated profile. For a 60-s exposure, the Gaussian model in red cannot capture both the wings and the central peak accurately—to match the width of the peak the Gaussian would necessarily underestimate the flux in the tails—and so compromises on both in the fit.

For the rest of this chapter we use PSF parameter measured on images using the Kolmogorov profile model (for both data and simulation), as it more accurately describes the data and is more physically motivated than the Gaussian model.

6.5 Time scales of PSF variation

We outlined questions of interest for Rubin Observatory regarding the temporal behavior of the atmospheric PSF in Section 6.1. With the PSF parameters measured from both Zorro and simulated images of varying exposure times, as described in the previous sections, we are now equipped to explore how 1) the atmospheric PSF evolves with exposure time (Section 6.5.1) and 2) the atmospheric PSF is temporally correlated with itself (Section 6.5.2).

6.5.1 Smoothing

We will first explore the evolution of the PSF during an integrating exposure using measurements of the PSF on a series of images such as those shown in Figure 6.3 and Figure 6.7. Because the short exposure images are significantly noisier for observations of faint science targets, we include only the PSF standard-star observations in this analysis. Apriori we expect the atmospheric PSF ellipticity to decrease with integration time since the shape of an instantaneous PSF depends on the structure of turbulence that the wavefront passed through; because the structure is rapidly and randomly changing, the integrated PSF should get rounder. However, we do not expect the size of the instantaneous PSF to vary significantly over the course of an exposure because the size is set not by the specific structure of turbulence but by the spatial scales of variation of the turbulence field (r_0 and L_0). One effect that could increase the PSF size with integration is image motion, caused by evolution of the wavefront “tilt”, which leads to motion of the centroid of the PSF.

Scatter plots of the PSF ellipticity components e_1 and e_2 for each PSF standard-star observation are plotted in Figure 6.14 for Zorro data (top) and simulations (bottom), for integrated exposure times of 0.06, 1.1, 14, 37, and 60 s. The cloud of points in each scatter plot shrinks in size with increasing exposure time, as we expect. However, we find that the spread in e_1 and e_2 , and therefore the magnitude of the ellipticity, is significantly larger in Zorro data than in simulations across all exposure times; in other words, the simulated PSFs are much *rounder*, on average.

We investigate the ellipticity dropoff with exposure time, and the difference between data and simulation, further in Figure 6.15. The PSF ellipticity magnitude e is now plotted as a function

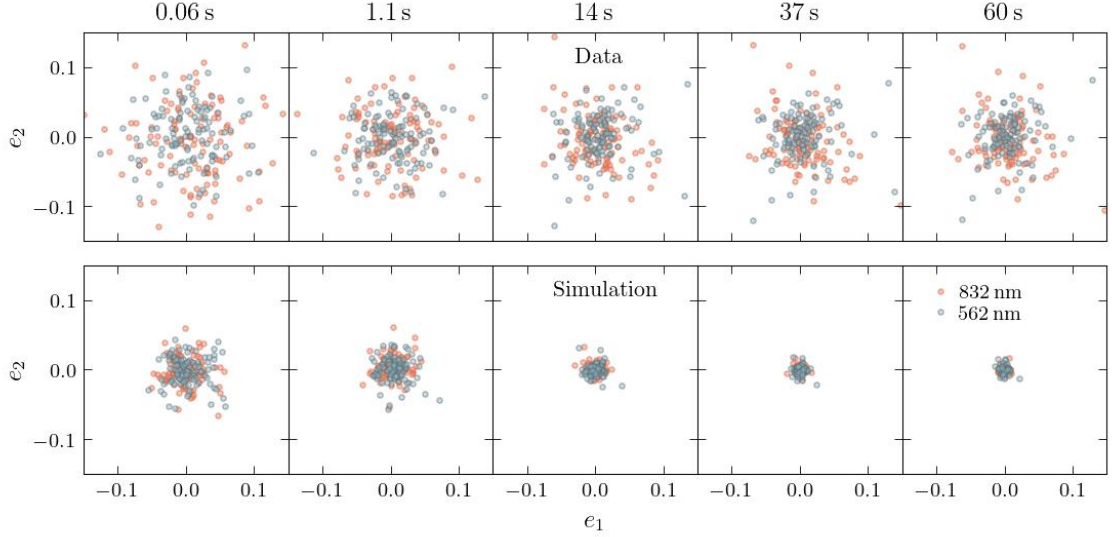


Figure 6.14: Scatter plots of PSF shape parameters e_1 and e_2 for observations (top) and simulations (bottom) of PSF standard stars, with increasing exposure times from 0.06 s to 60 s, from left to right. Observation wavelengths are $\lambda = 562$ nm (blue points) and $\lambda = 832$ nm (red points).

of integrated exposure time, where the data points correspond to the median ellipticity across all observations at approximately logarithmically increasing exposure times. The shaded regions show the 25th to 75th percentiles of the measurements. We observe that the spread in ellipticity magnitude is much larger in Zorro data (top panel) than in the simulations (bottom panel). The Zorro measurements in the redder filter show more elliptical PSFs on average than those in the bluer filter; this is consistent with smaller PSF sizes expected at longer wavelengths.

Each curve on the plot is a power-law fit to the median ellipticity e as a function of exposure time t_{exp} , of the form

$$e(t_{\text{exp}}) = b t_{\text{exp}}^{\alpha}. \quad (6.3)$$

The value of the exponent from the fit is $\alpha \approx -0.1$ for the Zorro data and $\alpha \approx -0.2$ for the simulations—i.e., we find that the PSF ellipticity decreases fractionally more quickly in simulations than in Zorro data, in addition to having lower values at all exposure times.

Figure 6.16 shows the evolution of PSF size with increasing integrated exposure time. As we saw in Figure 6.6, the median PSF size (as well as the spread in sizes) is larger for Zorro data than in the simulations for all exposure times. The Zorro data show a slight increase in median PSF size over the course of the exposure, while in simulation the median size is approximately constant. This difference could be due to image motion being relatively larger in the data; we discuss the impact of image motion in Section 6.6.2. The relative size of the PSF in the different filters follows the expected trend in both data and simulation—i.e., the median size is smaller for redder wavelengths.

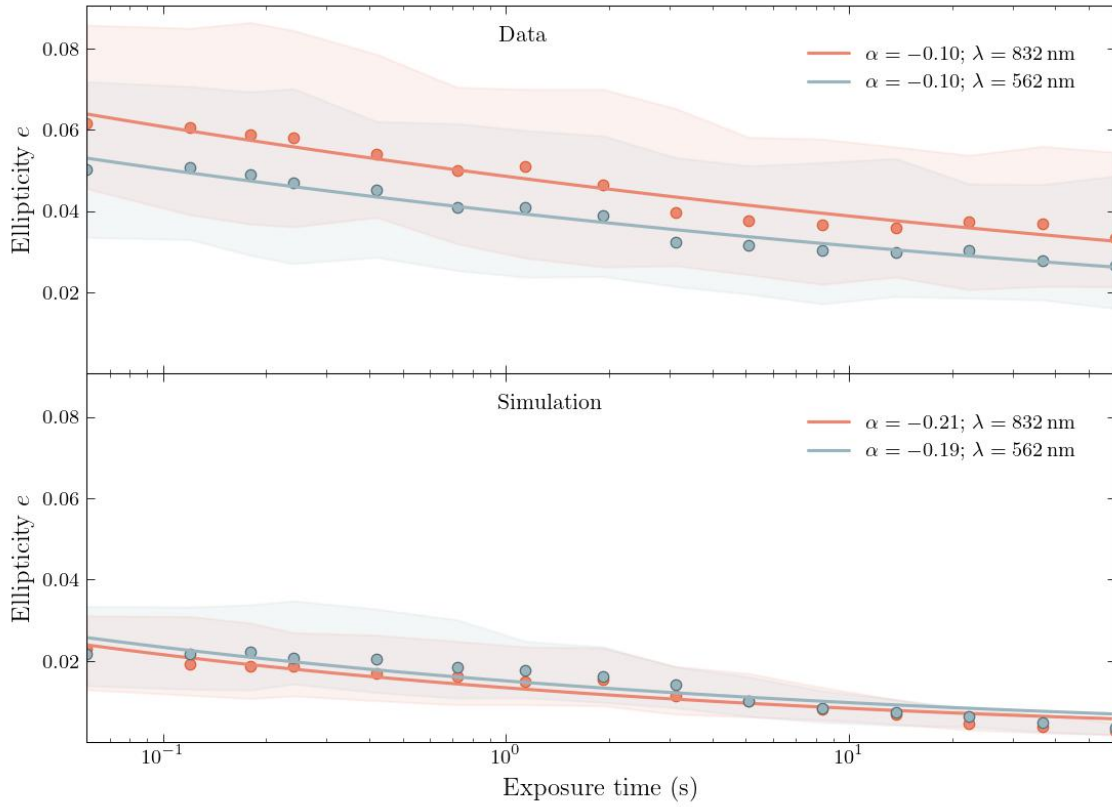


Figure 6.15: PSF ellipticity magnitude e as a function of increasing exposure time (from 0.06 s to 60 s) in Zorro data (top) and simulations (bottom) for PSF standard-star observations. Dots represent the median ellipticity and shaded regions span the 25th to 75th percentiles of the ellipticity values. For each observation wavelength ($\lambda = 562$ nm, blue; $\lambda = 832$ nm, red), the curves show the best-fit power law; the corresponding values of the best-fit exponent α are shown in the figure labels.

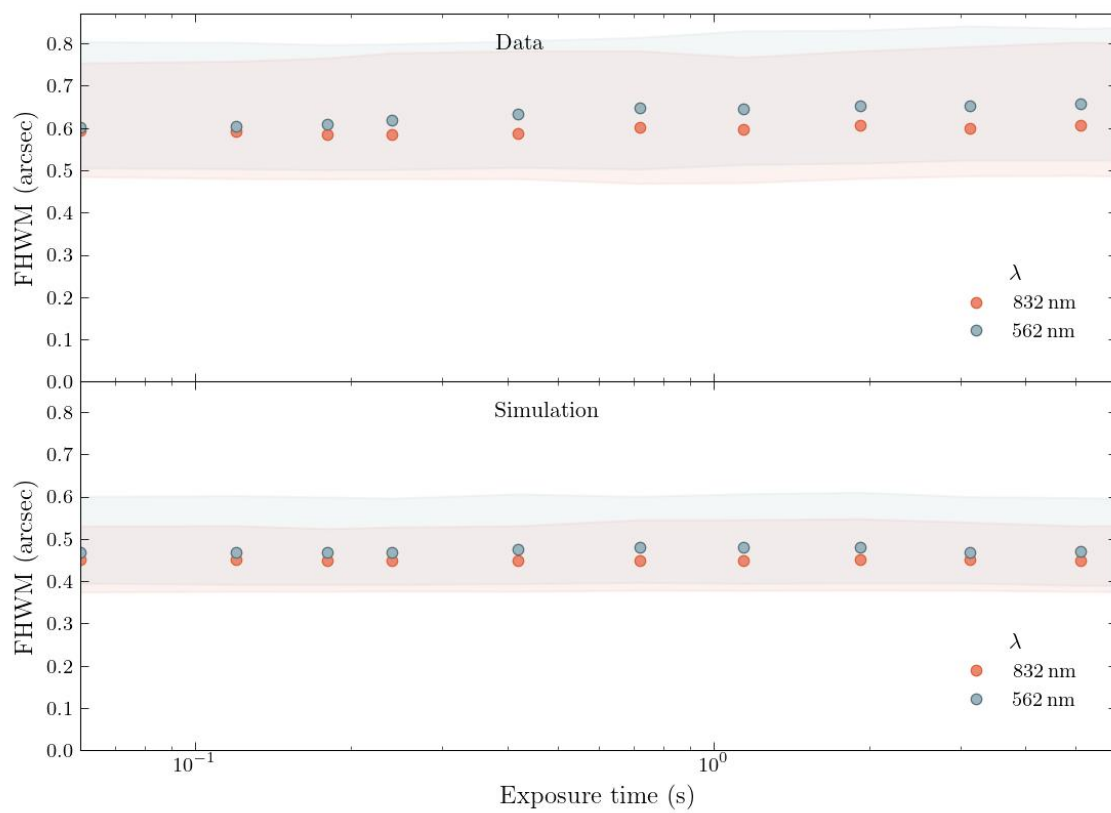


Figure 6.16: PSF size (FWHM) as a function of increasing exposure time (from 0.06 s to 60 s) in Zorro data (top) and simulations (bottom) for PSF standard-star observations. Dots represent the median FWHM and shaded regions span the 25th to 75th percentiles of the FWHM values. Observation wavelengths are $\lambda = 562$ nm (blue points) and $\lambda = 832$ nm (red points).

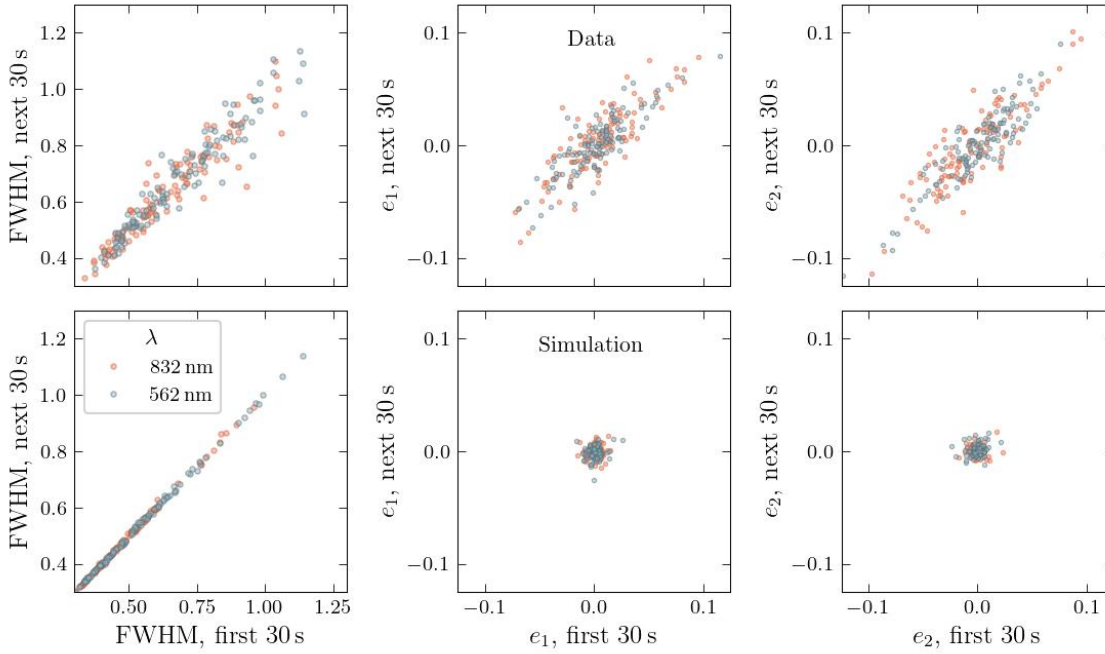


Figure 6.17: PSF size (left) and shape parameters e_1 (middle) and e_2 (right) for Zorro data (top) and simulations (bottom) for PSF standard-star observations. In each scatter plot, each point corresponds to the values of PSF parameters for two consecutive 30-s exposures with the PSF parameter measured from the first 30-s exposure plotted on the horizontal axis, and the parameter from the following 30-s exposure on the vertical axis.

6.5.2 Temporal correlations

We now address the second question introduced in Section 6.1: how correlated are parameters of the PSF from one exposure to the next? Following the same reasoning described in the previous section, we expect the PSF size to be very temporally correlated and the PSF shape to be mostly random from one exposure to the next.

We focus on 30-s exposures in this section, to match the Rubin exposure time. We begin with the PSF parameters measured from back-to-back 30-s exposures from the PSF standard-star observations. The correlations between parameters in successive exposure are shown in Figure 6.17. As expected the PSF sizes (left panels) show strong correlation—tighter in simulation than data, consistent with the results in Figure 6.16. The scatter in this relation for data increases for larger PSFs. PSF shape parameters e_1 and e_2 show a strong correlation between successive exposures for the Zorro data, while there is no evidence of correlation in simulations, where the spread in shape parameters is much smaller. Note that even for Zorro data, the evidence for correlation would be much weaker if we considered for the first 30 s only values of $|e_i|$ that are comparable to the values observed in simulations.

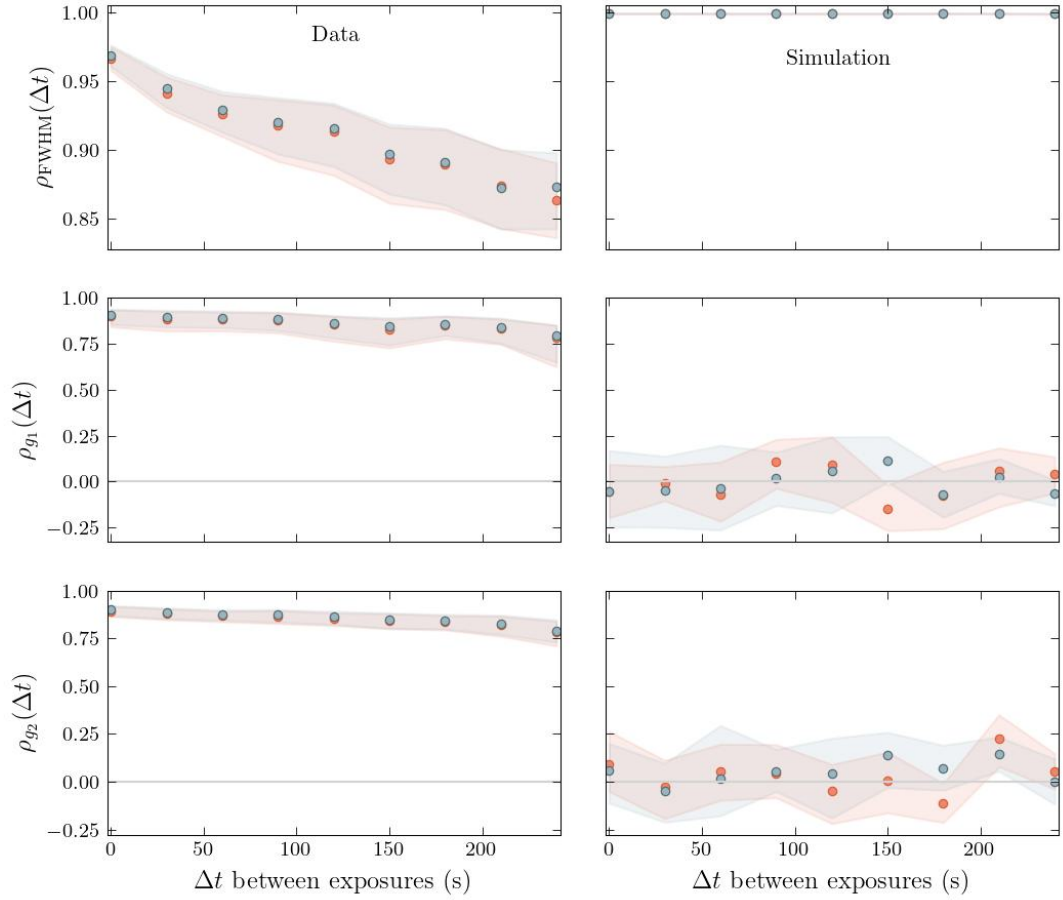


Figure 6.18: Correlation coefficient for PSF parameters (size, top; e_1 , middle; e_2 , bottom) as a function of time Δt between the beginning of one 30-s exposure and the beginning of another 30-s exposure in the same 300-s observation, for Zorro data (left) and for simulations (right) of science targets. Pearson correlation coefficients for each parameter are evaluated across pairs of 30-s exposures within each of 61 300-s observations. Points show median correlation and shaded regions span the 25th to 75th percentiles (estimated with a jackknife method).

We explore the unexpected correlations in PSF ellipticity further using the 61 faint science-target observations. We divide the 300-s total observation time into 10 consecutive 30-s PSF images. We calculate Pearson correlation coefficients between each pair of 30-s PSF images, with the jackknife method.⁴ The results are shown in Figure 6.18, where the correlation coefficients have been plotted as a function of the separation Δt between the beginning of one 30-s exposure and the beginning of a later 30-s exposure. The points are the median correlation coefficient over all jackknife samples for all pairs of PSFs separated by the value of Δt , and the shaded regions span the 25th to 75th percentiles of those same samples.

The simulation results (right panels in Figure 6.18) are consistent with our prior observations (based on Figure 6.15, Figure 6.16, and Figure 6.17): the small variations in PSF size, which we observed to be highly correlated between consecutive 30-s exposures, remain highly correlated between non-consecutive 30-s exposures; the (small) variations in PSF ellipticity component show no significant correlation between any pairs of 30-s exposures.

The data results (left panels in Figure 6.18) show very different trends from the simulation. The PSF size correlation decreases throughout the 300 s, with a final correlation coefficient of $\rho \approx 0.87$ —still highly correlated, but significantly less than in simulation. A possible explanation for the different behavior of the PSF in simulation versus data might be that the turbulent phase in simulation phase screens are all drawn with precisely the same r_0 value; thus there is no cause for evolution in PSF size. However, in reality r_0 is a local spatial average and, therefore, wavefronts at different points in time during an exposure can sample turbulence with varying r_0 . While the temporal correlation coefficients for the shape parameters in simulation are consistent with zero, the shape parameters in data are strongly correlated in time; although they decrease with increasing Δt , after 270 s the median correlation coefficient is still $> +0.75$.

This result—persistent temporal correlations in PSF ellipticity—is not consistent with our expectations or simulations. Before drawing conclusions from these high correlation values, a thorough investigation of potential systematic effects would be necessary. Though our goal in this work is to understand aspects of the *atmospheric* PSF, other effects are unavoidable. The Zorro data were recorded with an optical instrument, which contributes an optical PSF. Other potential sources of noise or bias are present in a real telescope—for example, telescope “shake” due to wind, deformations in the optical system due to gravity, dome seeing, etc. We identify and explore some of these potential contributions to the PSF in the next section.

6.6 Further studies: weather conditions and centroid motion

In this section we extend the analysis of the Zorro data, and simulated observations, beyond the temporal behavior of the PSF. We first relate measured PSF parameters to observing conditions (the

⁴We repeatedly calculate the correlation coefficient over a subsample of the data, leaving just one datapoint out of the sample each time. The jackknife is similar to, but predates, the bootstrap method.

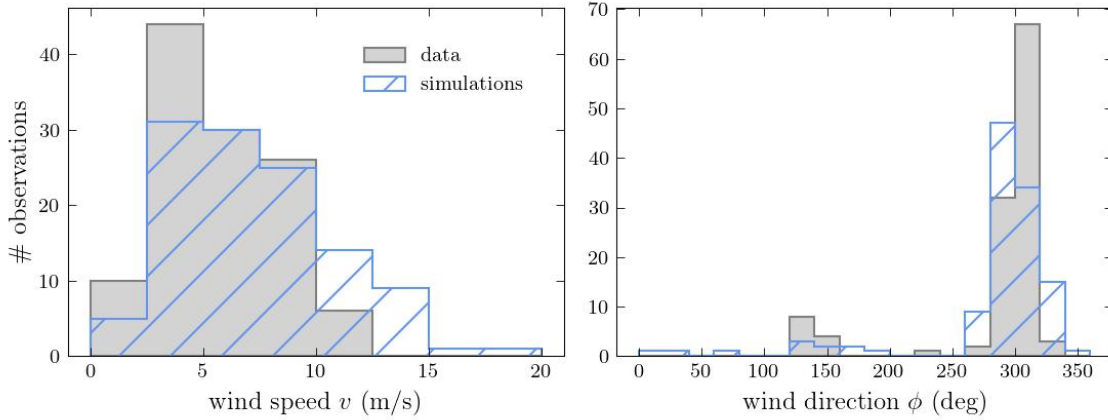


Figure 6.19: Distribution of wind speeds v (left) and directions ϕ (right) for data measured by the Gemini South weather tower within 5 minutes of associated observation (grey shaded histograms) and for simulations generated with input wind parameters from PSF-WEATHER-STATION (blue hatched histograms).

wind speed and direction at the time of observation/simulation, and airmass) and then thoroughly investigate the impact of image motion on PSF parameters.

6.6.1 Dependence of PSF parameters on observing conditions

We saw in Chapter 4 that the spatial correlation of the atmospheric PSF depends on both wind speed and direction. Do parameters of individual PSFs also have such a dependence? We explore this question using PSF parameters measured from 30-s observations and simulations of PSF standard stars. The wind speed and direction at Gemini South are recorded by the weather tower every five minutes; therefore, we are able to match each Zorro observation with a measurement of wind speed and direction taken within 2.5 min of the images.

The wind speed and direction distributions are shown in Figure 6.19 for both the Gemini South data (grey shaded histogram) and simulation input (blue hatched histogram). The distributions are mostly consistent—by design, since PSF-WEATHER-STATION drew from the Gemini South weather tower records for these simulations. The distributions of wind speed are slightly different, with the simulation reaching higher speeds than the observations: Gemini South doesn’t operate with winds much higher than 10 m/s due to wind shaking the dome. The vast majority of observations (and simulations) were taken when the wind direction was ≈ 300 degrees (wind blowing from the ocean), and a few were taken with wind direction ≈ 120 degrees (wind blowing from the mountains to the ocean). In data, we expect warmer, more turbulent air when the wind is blowing from the mountains compared to from the sea. The simulation, on the other hand, does not capture any potential differences in turbulence strength and air temperature due to wind origin. Therefore, if

the wind origin has an impact on any PSF parameters, we expect to see it only in the data.

We compare the PSF size, ellipticity magnitude, and image motion with the wind speed and direction in Figure 6.20. For this plot only, we plot PSF size adjusted for airmass (PSF size scales as airmass raised to a power of 0.6). We characterize the overall amount of image motion within an exposure using the standard deviation of the PSF centroid position, where the position is measured on stacks of five 60-ms exposures for a more robust estimate. Since we are interested in any potential dependences on wind direction, we color-code the points according to whether the wind direction is onshore ($\phi > 240$ degrees; teal) or offshore ($\phi < 240$; orange). The correlation coefficients on each panel use both teal and orange data and are calculated with a jackknife method.

In all panels of Figure 6.20, wind speeds tend to be lower when the wind is offshore; in fact, offshore winds do not exceed ≈ 7.5 m/s. This trend is interesting to note; however, the wind direction does not appear to have a significant effect on the PSF parameters shown in this figure.

The PSF size in data shows a positive but weak correlation with wind speed ($\rho = +0.4 \pm 0.1$); no significant correlation is observed in simulation. This could be explained by wind shake in the telescope, a phenomenon not present in the simulation; however, we do not observe a significant correlation between image motion and wind speed (bottom left) in the data, which we would expect if this were the case. The PSF shape in data has a correlation coefficient consistent with zero, while in simulation there is a weak negative correlation, which is expected: because higher-speed winds blow turbulence more quickly across the aperture, the resulting PSFs have seen less correlated turbulence leading to smaller values of ellipticity. Image motion in the data does not show a clear dependence on wind speed, while in simulation there is a very weak and not very significant negative correlation—similarly to the case of PSF shape, the wavefront may have smaller overall “tilt” when it has seen turbulent phases that are less correlated.

The size and shape of the PSF may also depend on the airmass of the observation; at larger airmass, light travels through a larger volume of air, which can increase the size of the PSF (as discussed in Chapter 2). The PSF shape can also be affected because light is refracted by the atmosphere slightly more, or less, depending on wavelength; this effect, called atmospheric chromatic dispersion, smears the PSF slightly along the zenith direction. However, because the Zorro filters are rather narrow, this effect is likely too small to measure on the ~ 200 PSFs analyzed here, given the many other sources of variation in PSF shape. We compare the PSF shape and size to airmass in Figure 6.21; PSF parameters are measured on the first 30-s image for both types of target observations for Zorro data and for simulation. We do not observe a clear dependence of the values of either PSF parameter on airmass. Though we expect PSF size to depend on airmass, it is not the only determining factor; a large variation in atmospheric seeing could obscure the airmass dependence. The lack of dependence of PSF shape on airmass is somewhat reassuring, in that the we do not see evidence of distortions related to telescope pointing.

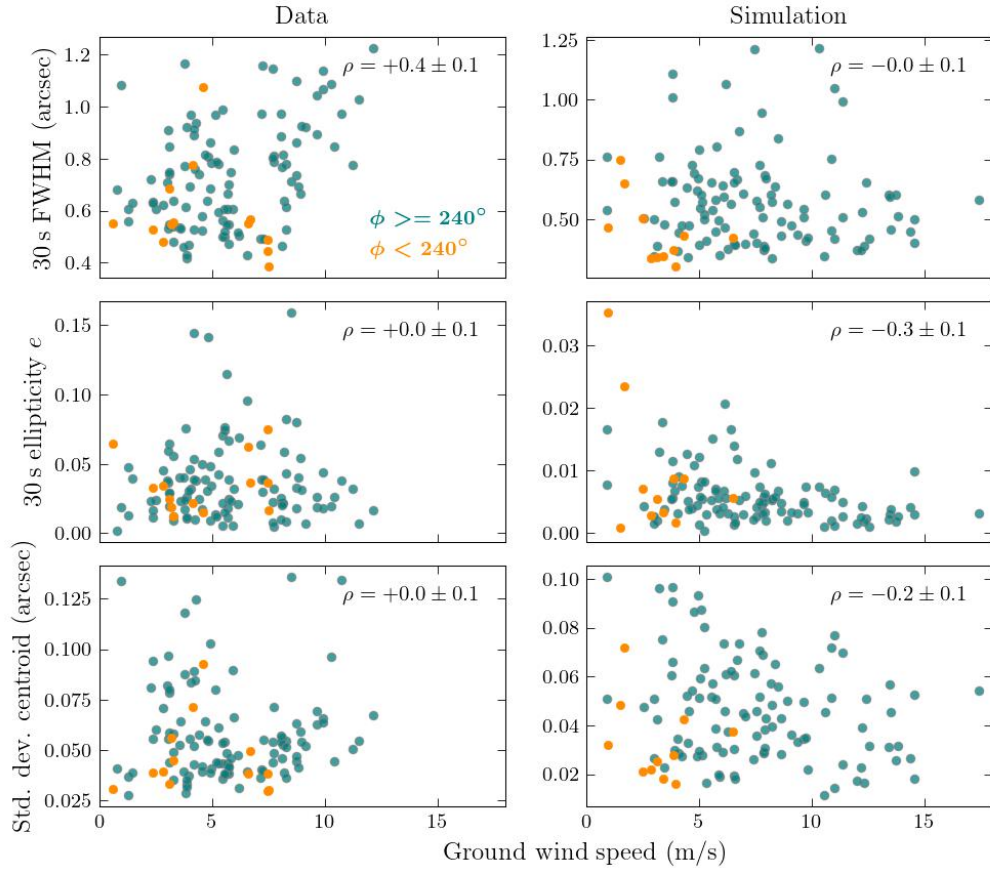


Figure 6.20: PSF parameters (size, top; ellipticity e , middle; standard deviation of centroid position, bottom) versus wind speed at the time of observation for Zorro data (left) and for simulations (right). The color of each point indicates the wind direction ϕ : teal points correspond to dominant on-shore wind directions (see peak around 300° in Figure 6.19); orange points correspond to off-shore winds (peak around 120° in Figure 6.19).

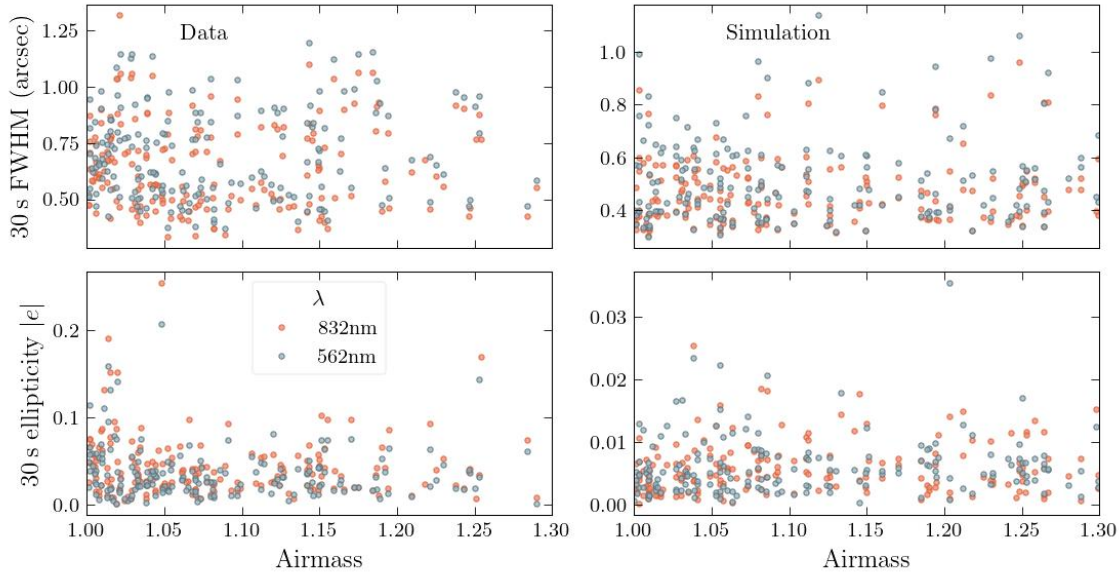


Figure 6.21: PSF size and shape parameters from (left) all Zorro observations (both PSF standard stars and faint science targets) and (right) simulations. PSF size (FWHM; top row) and shape (e ; bottom row) are plotted against airmass. The PSF parameters are measured from the first 30 s of the observation in $\lambda = 562\text{ nm}$ (blue points) and $\lambda = 832\text{ nm}$ (red points).

6.6.2 Impact of centroid motion

We now further investigate the impact that image motion has on measured PSF parameters, both as an exploration and as a search for potential systematic effects in the data. Image motion for each observation is evaluated as described in the previous section: the PSF position is measured on stacked images of 0.3 s total integration time, so that the position is measured 100 times during a 30-s exposure and 200 times during a 60-s exposure. In this section we use only observations of PSF standard stars.

Statistics of these positions, as measured on 60-s observations from data and simulation, are shown in Figure 6.22: the mean radial distance from the center of the image (left) and the standard deviation of the positions (right). The standard deviation of the positions across each observation quantifies the amount of image motion. The simulations exhibit shorter mean distances from the center of the image than the data, which is not surprising given possible imperfections in telescope pointing in data but not in simulations. Simulations show reduced image motion compared to data; the peak in simulation values is roughly half that of the data. None of the panels show evidence of wavelength dependent differences in position or motion; the mean positions in the blue and red filters in the data are 0.15 arcsec and 0.16 arcsec, respectively, and 0.06 arcsec and 0.05 arcsec in simulation.

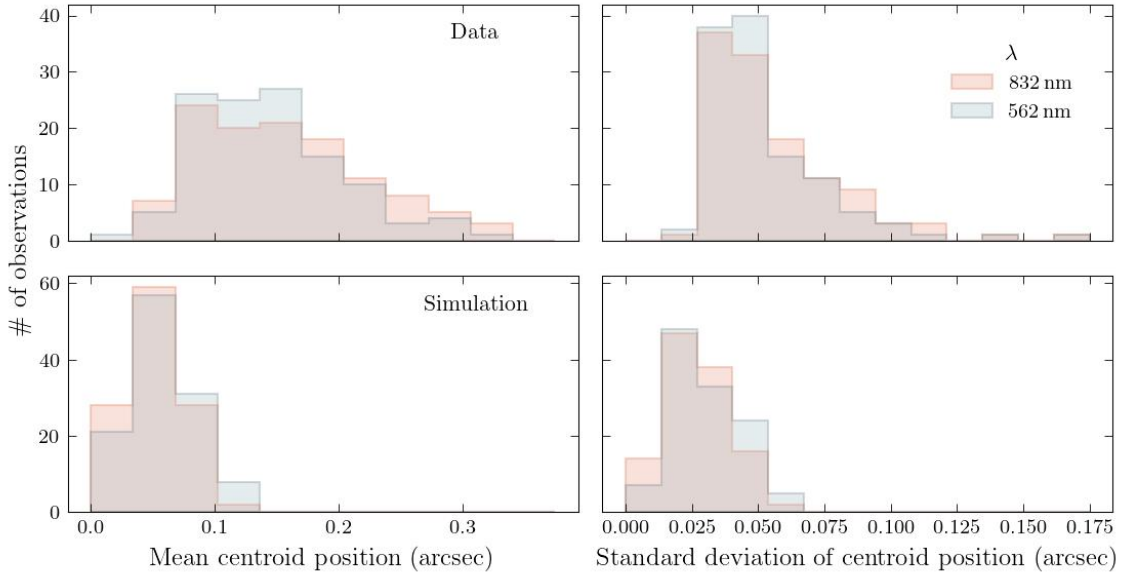
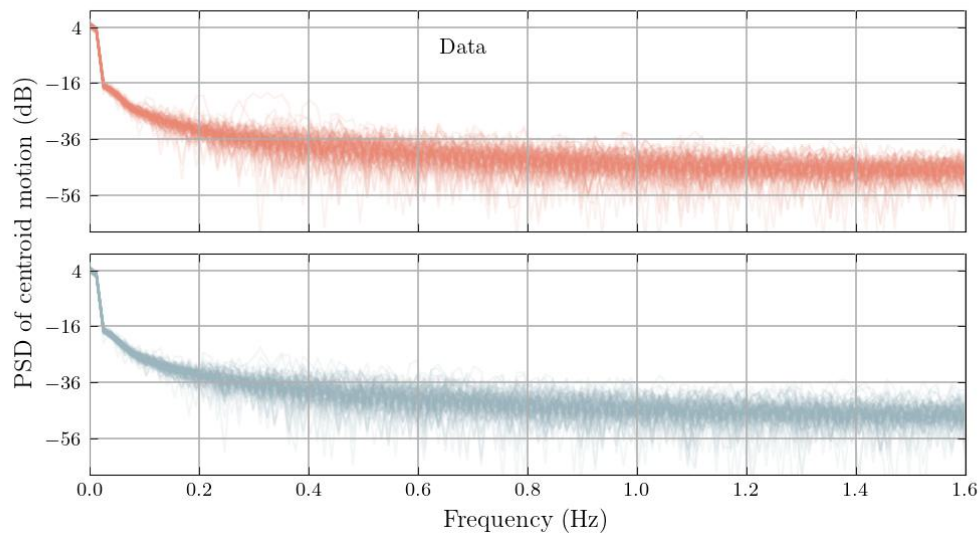


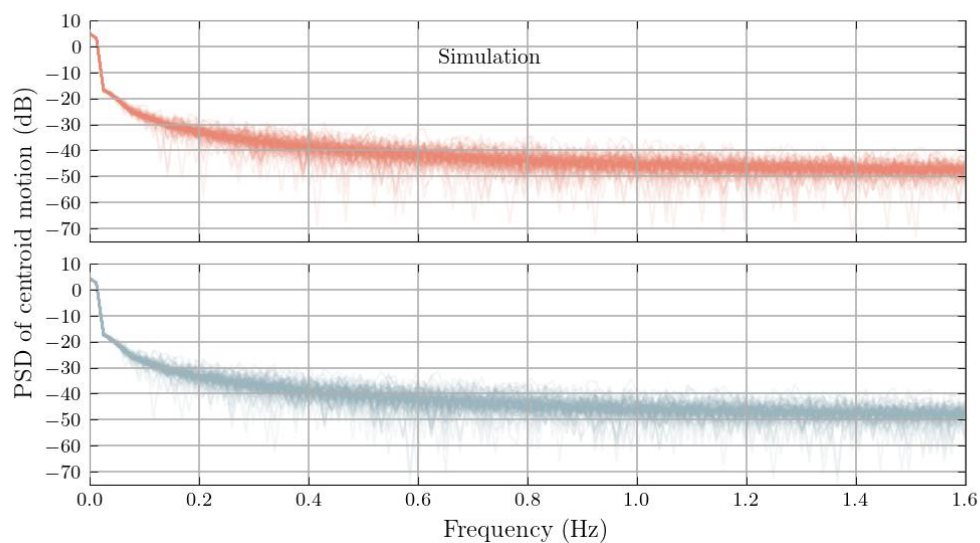
Figure 6.22: Distributions of mean (left histograms) and standard deviation (right histograms) of the radial distance of the PSF centroid from the center of the image across each of 117 60-s observations, for all Zorro observations of PSF standard stars (top) and simulated PSFs (bottom). Observation wavelengths are $\lambda = 562$ nm (blue shaded histograms) and $\lambda = 832$ nm (red shaded histograms).

It is not just the standard deviation of centroid position that is of interest, but also the frequency of the motion—are there particular time scales on which the PSF is moving? A frequency analysis could reveal any periodic motion, such as telescope shake, that is present in the data. Power spectral densities (PSD) of image motion, for each of the 60-s observations, are shown in Figure 6.23 for (a) data and (b) simulation. Curves for observations in the blue (centered at 562 nm) and red (centered at 832 nm) filters are plotted in blue (bottom panels) and red (top panels), respectively. We do not see peaks in the power spectrum at any particular frequency other than at very low frequencies of ≈ 0.025 Hz (period of ≈ 40 s). This peak is visible in both simulation and data, so we do not attribute it to effects not included in the simulations – e.g. telescope shake.

For the remainder of this section we relate PSF parameters to image motion; for relevance to Rubin exposure times, we use only 30-s, rather than 60-s, exposures. In Figure 6.24 we show scatter plots of measured PSF size (FWHM) and ellipticity magnitude e versus the standard deviation of centroid position within each 30-s exposure, for data (left) and simulation (right). There is evidence that PSF size is correlated with image motion in both data and simulation; in both cases, an increase in PSF size from 0.5 arcsec vs 1 arcsec is associated with a 0.05 arcsec increase in standard deviation of position—this is, very roughly, a 10% effect. The PSF ellipticity e , on the other hand, decreases as image motion increases. There are two reasons we might expect this. As image motion increases the size of the PSF, its ellipticity will decrease because convolution with an effective PSF due to motion



(a) Data



(b) Simulation

Figure 6.23: Power spectral density (PSD) for the PSF centroid motion over 60-s observations for (a) Zorro data and (b) simulations for $\lambda = 562$ nm (blue curve) and $\lambda = 832$ nm (red curve).

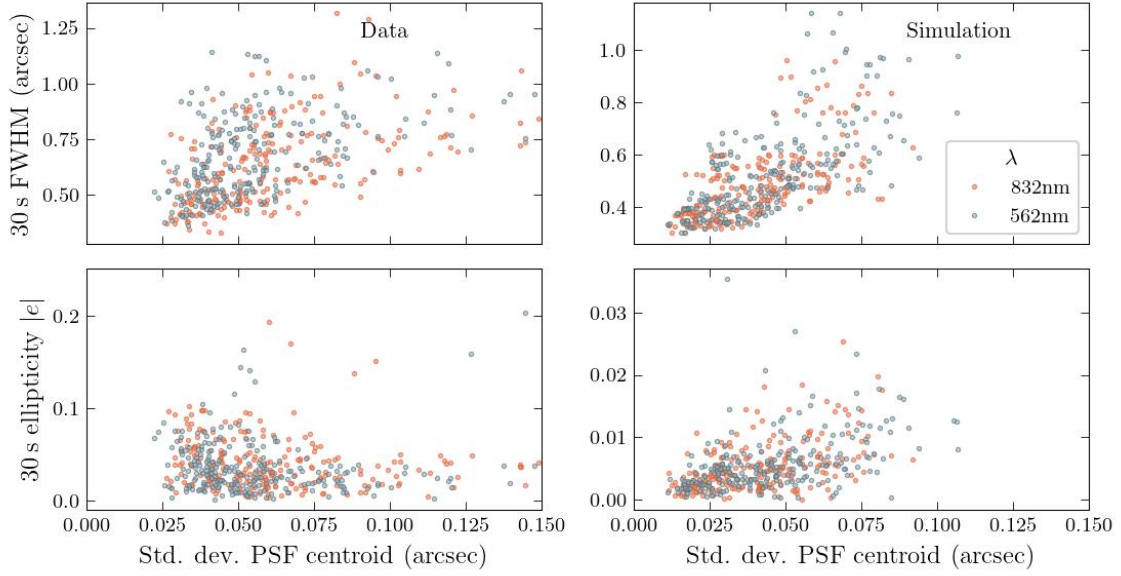


Figure 6.24: PSF size and shape parameters from (left) Zorro observations of PSF standard stars and (right) simulations. PSF size (FWHM; top row) and shape (e ; bottom row) are plotted against standard deviation of the PSF centroid motion. The PSF parameters are measured from both 30-s exposures in each 60-s observation in $\lambda = 562\text{ nm}$ (blue points) and $\lambda = 832\text{ nm}$ (red points).

makes the image rounder. The PSF shape in simulation, however, has the opposite dependence on centroid motion compared to data. However, as we have already seen, the ellipticity magnitudes are very different in data and simulation, with the average value of e being about an order of magnitude smaller in simulation than in data. Since the simulated PSFs are already fairly round, more centroid motion possibly serves to imprint more ellipticity rather than smoothing it out.

Up to this point we have treated image motion as a scalar quantity—how much is the position moving—without considering the image coordinate directions along which the PSF centroid position moves. Rather than the ellipticity magnitude, for this comparison we use the shape parameters e_1 , describing PSF shape along image coordinates x and y , and e_2 , describing shape along the 45-degree line. If image motion were the only driver of ellipticity, we expect that e_1 and e_2 are set by centroid motion along the relevant axes. For example, if the PSF centroid moves mostly along the x axis we expect a PSF with $e_1 > 0$ and $e_2 = 0$. Or, if the PSF moves only on the $x = y$ line, we expect $e_2 > 0$ and $e_1 = 0$. To quantify the image motion along image coordinates x and y , we calculate the second central moments of PSF centroid position: variance in the x -direction (σ_{xx}^2) and the y -direction (σ_{yy}^2), as well as the covariance between x and y , σ_{xy}^2 .

Scatter plots of PSF shape parameters versus centroid second moments in data are shown in Figure 6.25, with correlation coefficients estimated using a jackknife resampling method. We capture excess motion along either x or y by the difference of σ_{xx}^2 and σ_{yy}^2 . There is evidence of modest

statistically significant positive correlation between e_1 and $\sigma_{xx}^2 - \sigma_{yy}^2$, and between e_2 and σ_{xy}^2 . *Although increased image motion in data decreases PSF ellipticity overall, motion of the PSF centroid along specific image directions acts to increase the relevant ellipticity component.*

The scatterplots for simulation are shown in Figure 6.26—we note that the axis limits are different from those in the figure for data. The correlation results are similar for simulation and data, although the correlation coefficients between e_1 and $\sigma_{xx}^2 - \sigma_{yy}^2$, and between e_2 and σ_{xy}^2 , are larger and more significant in simulation than in the data.

6.7 Wavelength dependence of PSF size

The size of the PSF is a wavelength-dependent quantity, as briefly discussed at the end of Section 2.3.2. The dependence on wavelength λ is described by a power law, $\text{FWHM} \propto \lambda^b$. For pure Kolmogorov turbulence, $b = -0.2$; as the outer scale decreases, the magnitude of b is predicted to increase, with measured values as low as -0.4 . The Zorro data provide an opportunity to measure this power-law exponent, since the simultaneous measurements in two filters mitigates potential contamination from time-dependent effects, allowing a direct measurement of b . For a PSF FWHM measured in filters centered at wavelengths λ_r and λ_b ,

$$b = -\frac{\log(\text{FWHM}_b) - \log(\text{FWHM}_r)}{\log(\lambda_b) - \log(\lambda_r)}. \quad (6.4)$$

We calculate this value for each pair of simultaneous 30-s PSFs from observations of PSF standard stars. Results are plotted as a function of FWHM_b in the left panel of Figure 6.27; we use the PSF size as a proxy to probe the dependence of b on L_0 . The projected histograms on the right show distributions of the calculated values of b in data (grey shaded histogram) and simulation (hatched blue histogram). The median measured values are $b = -0.22$ from data (grey line) and $b = -0.18$ in simulation (blue line). The Kolmogorov prediction is marked with the orange dashed line.

This measurement is very sensitive to systematic changes in PSF size. For example, the light from each filter is measured on a different camera; therefore, small changes in the pixel scale for each camera can skew the median measured b . We have used the pixel scale as calibrated during the first half of 2019 as described in Scott et al. (2021); however, it is likely to drift during the 6-month span that our data covers (May through October)—indeed, the calibration values listed for the latter half of 2019 show a 3% decrease in pixel scale in the red filter and a 0.7% increase in the blue filter. Our estimated magnitude of b would increase if this effect were taken into account (see Equation 6.4).

Image motion can lead to an error in the calculated value of b by increasing the PSF size for both wavelengths, as we observed in Section 6.6.2, and thus reducing the apparent dependence of PSF size on wavelength. If image motion is introduced by a non-atmospheric source, this effect would also systematically decrease the calculated magnitude of b . We do not have an estimate of

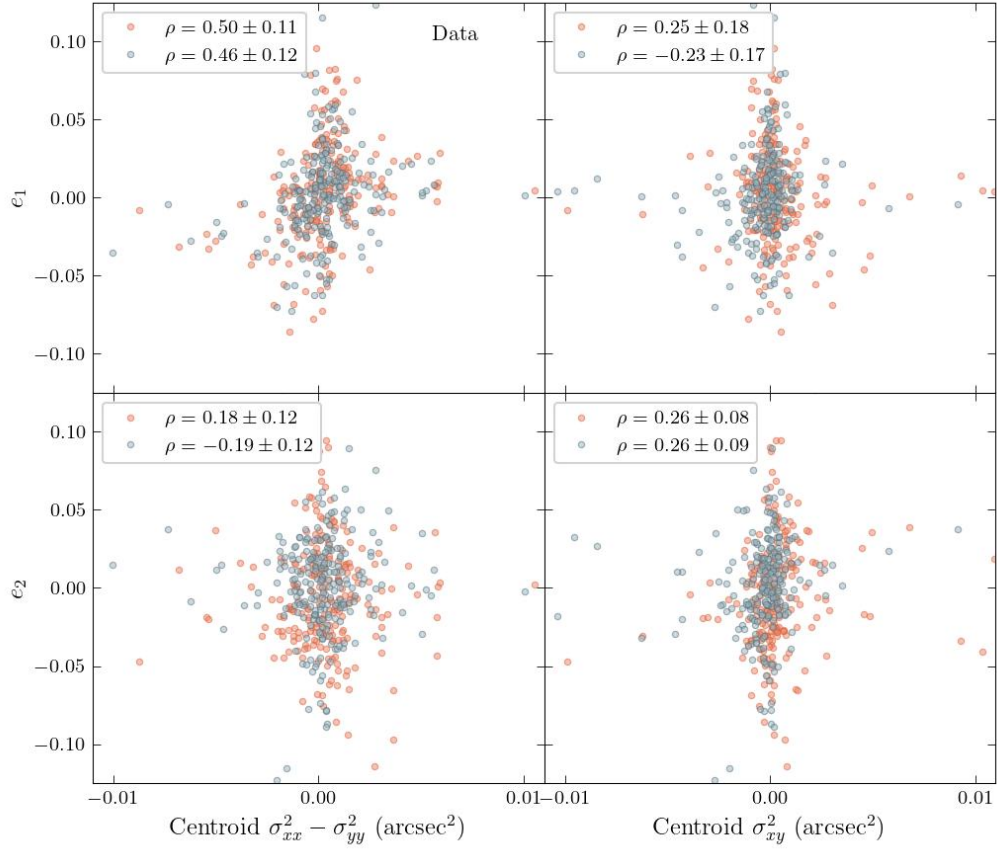


Figure 6.25: PSF shape parameters e_1 (top) and e_2 (bottom) versus second central moments of the per-exposure PSF centroid position distributions for Zorro observations of PSF standard stars. Centroid $\sigma_{xx}^2 - \sigma_{yy}^2$ (left) and σ_{xy}^2 (right) are calculated from x and y PSF positions measured on 200 PSF images for 0.3-s exposures. PSF shape parameters are measured from the 60-s integrated exposures. Pearson correlation coefficients are reported for observations at $\lambda = 562 \text{ nm}$ (blue points) and $\lambda = 832 \text{ nm}$ (red points) for each panel.

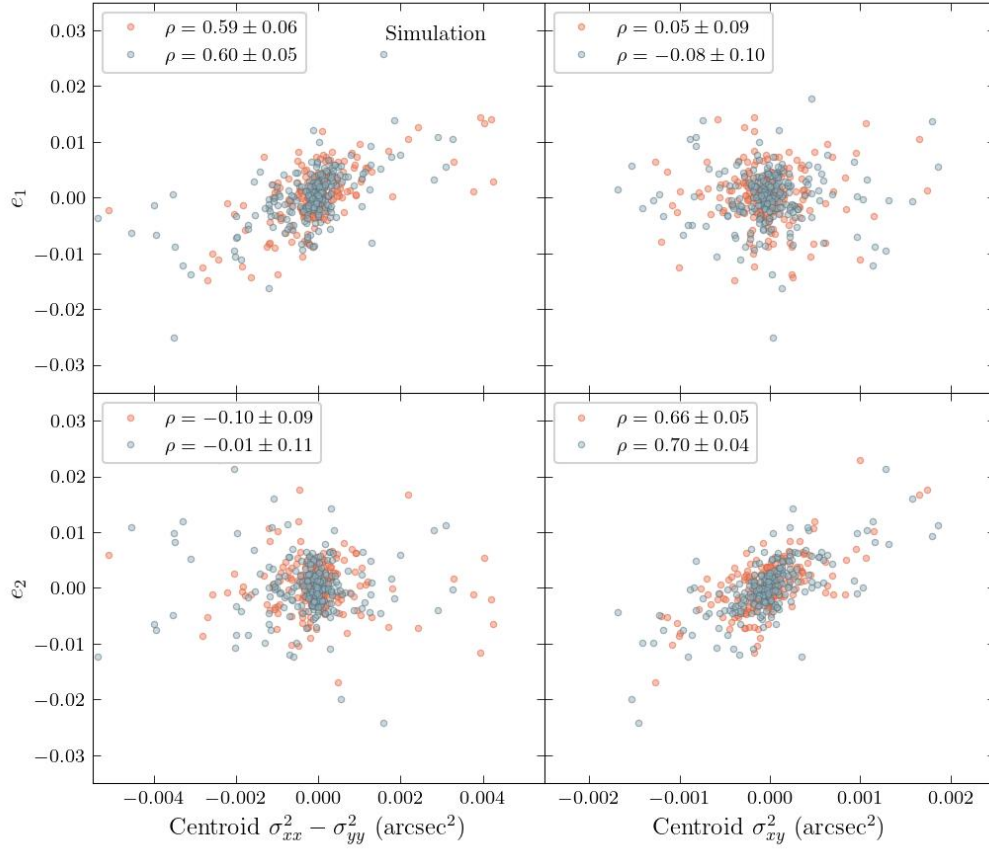


Figure 6.26: PSF shape parameters e_1 (top) and e_2 (bottom) versus second moments of the per-exposure PSF centroid position distributions for simulated observations of PSF standard stars. Centroid $\sigma_{xx}^2 - \sigma_{yy}^2$ (left) and σ_{xy}^2 (right) are calculated from x and y PSF positions measured on 200 PSF images for 0.3-s exposures. PSF shape parameters are measured from the 60-s integrated exposures. Pearson correlation coefficients are reported for observations at $\lambda = 562 \text{ nm}$ (blue points) and $\lambda = 832 \text{ nm}$ (red points) for each panel.

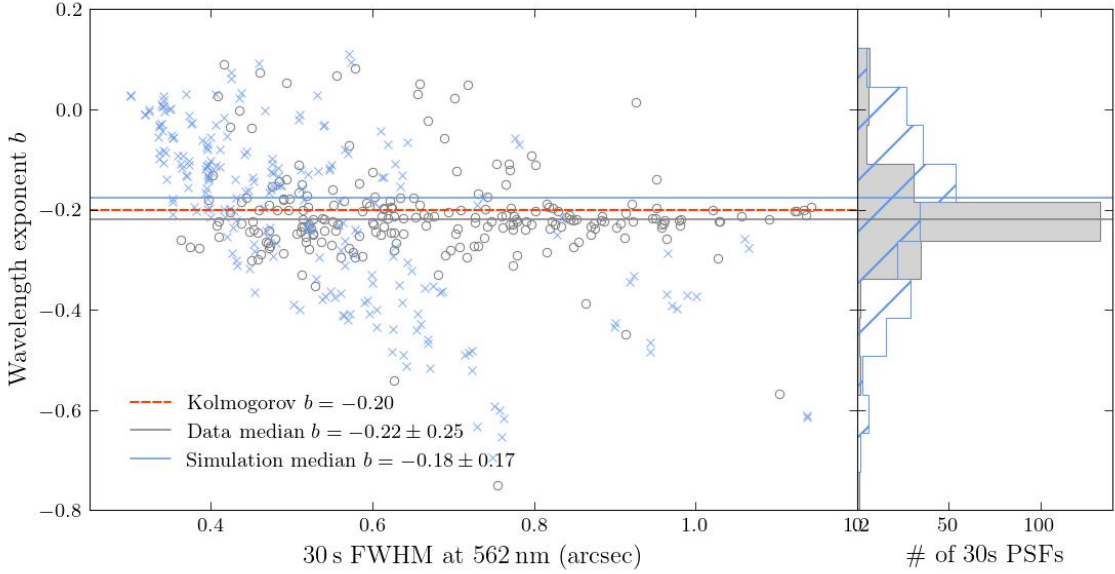


Figure 6.27: Measured values of the power law exponent b for the wavelength dependence λ^b of PSF size (FWHM), as a function of PSF size at $\lambda = 562$ nm. Each pair of simultaneous observations in narrow filters centered at $\lambda = 562$ nm and $\lambda = 832$ nm are used to estimate a value of b . Estimated values of b from Zorro observations of PSF standard stars (grey circles and shaded histogram) and simulations (blue crosses and hatched histogram) are shown in the scatter plots and projected distributions. The median value of the exponent in data is $b = -0.22$ (grey line) and in simulation is $b = -0.18$ (blue line), while the prediction from pure Kolmogorov turbulence (infinite outer scale) is $b = -0.20$ (orange dashed line). See text for a discussion of potential systematic uncertainties in these values of b .

the systematic uncertainties on our estimates of b , since the impact of both of the above effects on the value is difficult to quantify.

6.8 Summary and future work

In this exploratory study we found high temporal correlation in PSF size across both data and simulation, and high temporal correlation in PSF shape in data but not in simulation. Analysis of image motion over the course of an exposure showed that the PSF size and shape at the end of the exposure both depend on the motion. In particular, image motion contributed about 10% to the PSF size in both data and simulation. Increased magnitude of image motion decreased the PSF ellipticity magnitude in data, but motion of the PSF centroid along specific directions increases the relevant ellipticity component, in both data and simulation.

Future work

We found no evidence of particular sources of image motion in a frequency analysis; however, the data showed much more motion than the simulated observations. To mitigate contamination from contributions to this motion from telescope mis-centering, for example, we could stack the images by aligning the PSF centroid of each image. Comparison with the results presented here may provide insights into some of the discrepancies between data and simulation we observed.

Chapter 7

Conclusions and future work

7.1 Summary

In this thesis, I have described my work on modelling and characterizing key aspects of the atmospheric point spread function: its anisotropic spatial correlations across a wide field of view during a single exposure, its dependence on atmospheric parameters (wind speed, wind direction, turbulence strength – all as a function of altitude), and its temporal evolution, including temporal correlations. These studies address the goal of making accurate and precise measurements of cosmic shear with large-field surveys such as the Legacy Survey of Space and Time (LSST) at Rubin Observatory. A challenge for making unbiased measurements of two-point statistics of galaxy shapes is the development of PSF modeling and interpolation methods that accurately capture the anisotropic two-point correlations for PSF size and shape on different scales across the field of view. Robust and high-fidelity simulated images provide an important tool for developing and testing improved methods for PSF correction that perform to the desired accuracy for Rubin.

In Chapter 3 we presented a new software tool (PSF-WEATHER-STATION) for producing high-fidelity correlated weather and turbulence parameters as input to atmospheric PSF simulations, enabling realistic images of point spread functions across a wide field of view. The PSF-WEATHER-STATION software package, configurable to any observatory, is public at <https://github.com/LSSTDESC/psf-weather-station> and includes installation instructions, documentation, and tutorial notebooks.

In Chapter 4 we used realistic, altitude-dependent weather and turbulence input provided by the PSF-WEATHER-STATION package to simulate the atmospheric PSF across the Rubin Observatory field of view and characterize the predicted two-point correlation functions (2PCFs) we can expect for LSST. We explored the impact of the realistic weather conditions input to the simulations on summary statistics of the two-point correlation functions, and found:

1. The variance in PSF size and shape across a single exposure decreases as wind speed increases.
2. The values of the 2PCFs for PSF shape in a single exposure are of the same order as (or larger than) the expected 2PCFs for cosmic shear over the range of angular separations used in cosmic shear analyses: \sim few arcmin to \sim 100 arcmin.
3. There exist dominant wind directions at Rubin Observatory, which in turn lead to dominant orientations of anisotropies in the 2PCF that is relevant for studies of cosmic shear (ξ_+). At scales \sim 100 arcmin, these anisotropies are due to strong correlations with ground-layer wind direction; at scales \sim few arcmin, they are due to correlations with free-atmosphere wind direction. These results can be understood in terms of the different angles subtended at different heights by turbulent structure of the same physical scale.

Large scale correlations are more easily modeled than small scales, so the realistic treatment of correlations in the upper atmosphere is important for development of new algorithms.

In Chapter 5 we repeat the 2PCF analysis on several exposures from the camera used for the Dark Energy Survey (DECam), and find that the wavefront method for removing optical PSF effects likely also removes part of the static atmospheric signal that leads to persistent correlations.

Finally, in Chapter 6 we explore time-dependent aspects of the PSF using novel speckle-image datasets recorded at the Gemini South Observatory (near Rubin Observatory) and compare these images with simulations generated with PSF-WEATHER-STATION parameters as input. We find inconsistencies in temporal correlations in PSF shape between data and simulation. We demonstrate that image motion plays an important role in generating PSF ellipticity.

7.2 Future directions for this work

7.2.1 PSF modeling and interpolation

We are now using PSF-WEATHER-STATION to generate simulations of LSST-like images of PSFs. We are exploring different modeling and interpolation algorithms for the PSF profiles (polynomial fits, Gaussian process interpolation), and measuring residual correlations. The next step for the LSST Dark Energy Science Collaboration will be to propagate residual correlations to cosmology. This work will allow us to set requirements on the accuracy we must achieve when using stars to model and interpolate the PSF across the focal plane and will inform further development of state-of-the-art algorithms.

The discussion in this thesis has focused solely on the second moments of the PSF—size and ellipticity—but accounting for higher-order moments of the PSF during PSF modeling and interpolation will be necessary to achieve the stringent requirements on cosmic shear for future surveys such as Rubin LSST (Zhang et al., 2021, 2023).

7.2.2 Impact of persistent on-sky correlations on co-added images

Because of the dominant wind direction at Cerro Pachón, we expect to see a dominant orientation of the anisotropy in the 2PCF for PSF size and shape with respect to the ground coordinate system, across LSST exposures at Rubin Observatory. The mapping of wind direction on the ground onto sky coordinates for each exposure depends on the pointing of the telescope. Therefore, the degree to which the dominant orientation in the ground coordinate system will wash out when exposures for a single field on the sky are combined depends on the observing strategy for the survey. Using the observing strategy for the 300-square-degree DESC DC2 simulation, we find that the dominant wind direction will translate to a persistent on-sky anisotropy; further study is needed to understand the implications for the full LSST survey. PSF simulations produced with PSF-WEATHER-STATION input can be used to study this question for a particular survey strategy.

7.2.3 Exploring impact of weather patterns on the atmospheric PSF

We have a tool, PSF-WEATHER-STATION, to link weather at the observatory with the atmospheric PSF. There are weather patterns on large time scales that will impact the weather at Rubin Observatory to some extent, such as el Niño and la Niña. These refer to the large scale warming, or cooling, of surface water in the Pacific, resulting in periods of increased, or decreased, temperature and rainfall along the Western coast of South America, lasting from a few to tens of months. These periods are predictable months in advance; understanding their impact on the PSF at Rubin could potentially be used to inform the observing strategy of Rubin LSST. Based on meteorological records of these phenomena, we can use PSF-WEATHER-STATION to draw on archival weather forecasting data products during these periods and predict the resulting changes in the PSF.

7.2.4 Characterizing atmospheric turbulence using speckle images

In our analysis of the speckle data from the Zorro instrument we have utilized the high temporal resolution; the high spatial resolution of the images remains unexplored. Fully resolved speckles from refraction through atmospheric turbulence, like those imaged with Zorro, contain information about the structure of that turbulence. Parameters of turbulence (e.g. r_0 , L_0) could potentially be measured from a Fourier analysis of the speckle patterns.

Zorro has been regularly taking data since going online in 2019; those data are available on the Gemini archive. Once we understand and quantify sources of systematic errors in the temporal and/or spatial analysis of the speckle data, we could leverage the large quantity of data to conduct statistically significant analyses.

7.2.5 More precise numerical modeling of the summit environment

The atmospheric PSF simulations described in this thesis have higher fidelity than previous simulations developed for ground-based, wide-field observatories; however, many approximations in the treatment of turbulence remain. The approximations we have used are not valid near the ground where turbulence is complex and non-Kolmogorov due to topographical effects and sources of heat. Precise numerical simulations of turbulence in the telescope environment *are* possible using either Reynolds-averaged numerical simulations and large-eddy simulations. These simulations are extremely expensive both in human and computer time – it is unrealistic to run these models for anything close to nightly.

Instead, one can run a grid of these expensive simulations, sampling from the relevant environmental parameters of temperature, wind direction, and wind speed, as well as variable dome configurations: whether baffles are in place, and the strength of the inside fans. A Gaussian process emulator to interpolate the 5-dimensional surface can be trained on these simulation results for fast queries. This is still an expensive process: before simulations can be run, the simulation mesh of the local topography must be set up and validated and dependence on boundary conditions understood. An understanding of the accuracy desired from the numerical simulations—for example, the amount of acceptable error in modeled turbulence given an accuracy requirement on the resulting PSF—is an important prerequisite for an undertaking of this computational complexity.

7.3 Conclusion

We have many challenges on the road to fully leveraging the statistical promise of the Rubin LSST data for high-precision measurements of cosmic shear. Vera C. Rubin Observatory is in final stages of construction; commissioning of the telescope and camera is underway and soon the system will collect its first “on sky” photons. The commissioning period is an exciting opportunity to explore, test, and improve the capabilities of our instruments and algorithms before the survey begins in earnest!

Appendix A

Global Forecasting Models for psf-weather-station

As summarized in Section 3.2, multiple organizations around the world produce high-quality weather models and forecasts; we focus here on those from the European Centre for Medium-Range Weather Forecasts (ECMWF) and the National Oceanic and Atmospheric Administration National Centers for Environmental Prediction (NOAA NCEP). Data products from these global models can be very useful for studies of the atmosphere for astronomical applications. Here we summarize the types of data available and considerations for their use in atmospheric PSF simulations.

Both ECMWF and NCEP make available two types of data: (1) *analysis* products are the best estimate of the state of the atmosphere, produced by combining a numerical weather prediction model with a variety of observations through a process called data assimilation; (2) *forecast* products are the numerical predictions (based on initial analysis products) for some time into the future. Analysis and forecast data are available in real-time (of use for weather forecasting) and as *reanalysis* data products: state-of-the-art data assimilation and numerical modeling methods applied to archival data (highly relevant for long-term climate monitoring)¹.

The 5th generation ECMWF reanalysis (ERA5) catalog covers the time period from 1940 to the present and is extensively documented² (Hersbach et al., 2020). ERA5 analyses are available hourly, with forecasts initialized at 00h and 18h UTC. At the time of writing, all ECMWF archival data (including ERA5) and subsets of real-time forecasts are available publicly under creative commons³.

All NOAA NCEP data are available publicly; real-time forecasts are from the Global Forecast System (GFS), and several reanalysis efforts are available for specific time periods⁴. The NCEP

¹All available atmospheric reanalysis datasets are summarized at <https://reanalyses.org/atmosphere/comparison-table>

²<https://confluence.ecmwf.int/display/CKB/ERA5>

³<https://www.ecmwf.int/en/forecasts/accessing-forecasts/licences-available>

⁴<https://psl.noaa.gov/data/gridded/data.ncep.reanalysis.html>

analysis products are available every 6h.

Data products are output on a uniform grid over the Earth’s surface, with resolution 31 km $\approx 0.28^\circ$ for ERA5 and NCEP GFS, and 2.5° for NCEP reanalysis.

Data products are available at heights corresponding to specific levels of pressure rather than specific altitudes. Since higher spatial resolution in the vertical direction allows for more accurate capture of important wind gradients in the atmosphere, we use the output type with densest vertical coverage – called *model levels*. ECMWF uses 137 model levels. NCEP uses 127 for the period since February 3, 2021, and 64 prior to that time. Model levels follow terrain at the Earth’s surface, and the conversion to altitude uses temperature and specific humidity⁵. This conversion has been implemented in PSF-WEATHER-STATION following ECMWF documentation and example code.

We choose to use ECMWF data products because the temporal and spatial resolution of available data is higher than for NCEP. In addition, ECMWF documentation is more detailed and accessible.

⁵<https://confluence.ecmwf.int/display/CKB/ERA5%3A+compute+pressure+and+geopotential+on+model+levels%2C+geopotential+height+and+geometric+height>

Appendix B

Gallery of simulated PSF parameter correlations

In this appendix, we provide a more extensive gallery of outputs from the simulations, based on PS-FWS inputs, described in Chapter 4. Similar to Figure 4.7 in the main text, Figure B.1, Figure B.2, Figure B.3, and Figure B.4 together show PSF shapes and their spatial two-point correlation functions from 32 randomly sampled simulated exposures. Each simulation output is summarized in a trio of panels: “whisker” plots showing PSF ellipticity—the orientation of each line corresponds to the orientation of the major axis of the PSF shape, while the length and color contrast are proportional to the magnitude of the ellipticity e —across the focal plane (top), and the two-point correlation function of PSF shape, $\xi_+(\theta, \alpha)$, at ~ 100 - arcmin separations (middle) and ~ 5 - arcmin separations (bottom). The color scales are common across each pair of 2PCF images.

The wind speeds and directions are indicated by the white vertical lines: the value of the angle corresponds to the wind direction, and the length of the line is proportional to the wind speed. The ground wind is shown on the ~ 100 arcmin separation image; the line length is relative to a wind speed of 25 m/s for the full height of the plot. The speed-weighted average direction of the free atmosphere wind is shown on the ~ 5 arcmin panels; the line length is relative to a wind speed of 65 m/s.

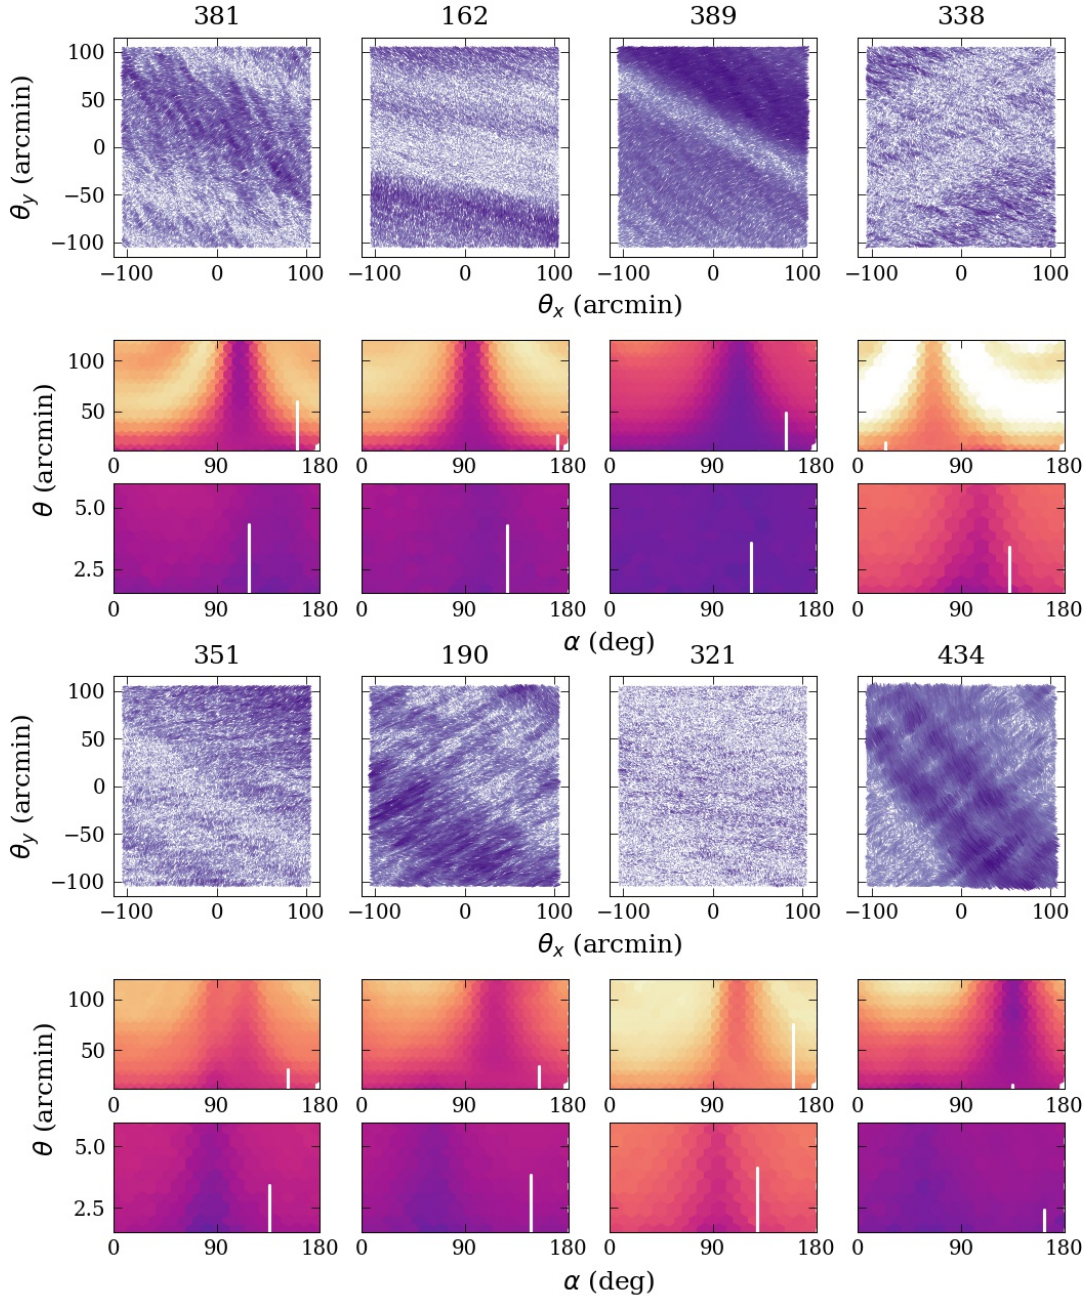


Figure B.1: Summary of simulated PSF shapes for eight PSFWS simulations. See text for details.

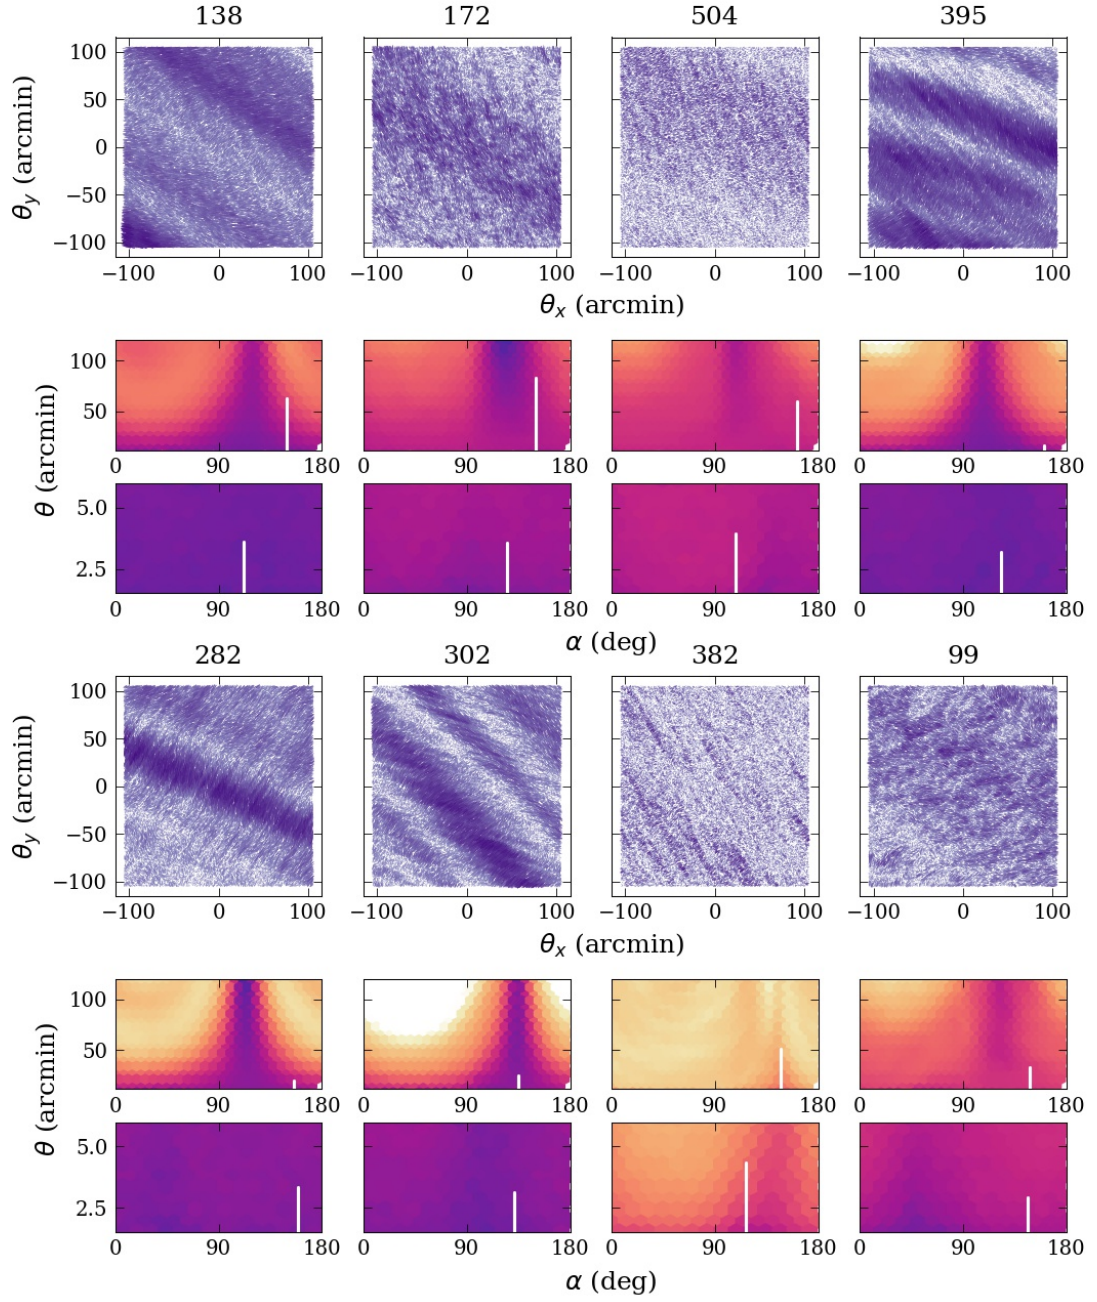


Figure B.2: Summary of simulated PSF shapes for eight PSFWS simulations. See text for details.

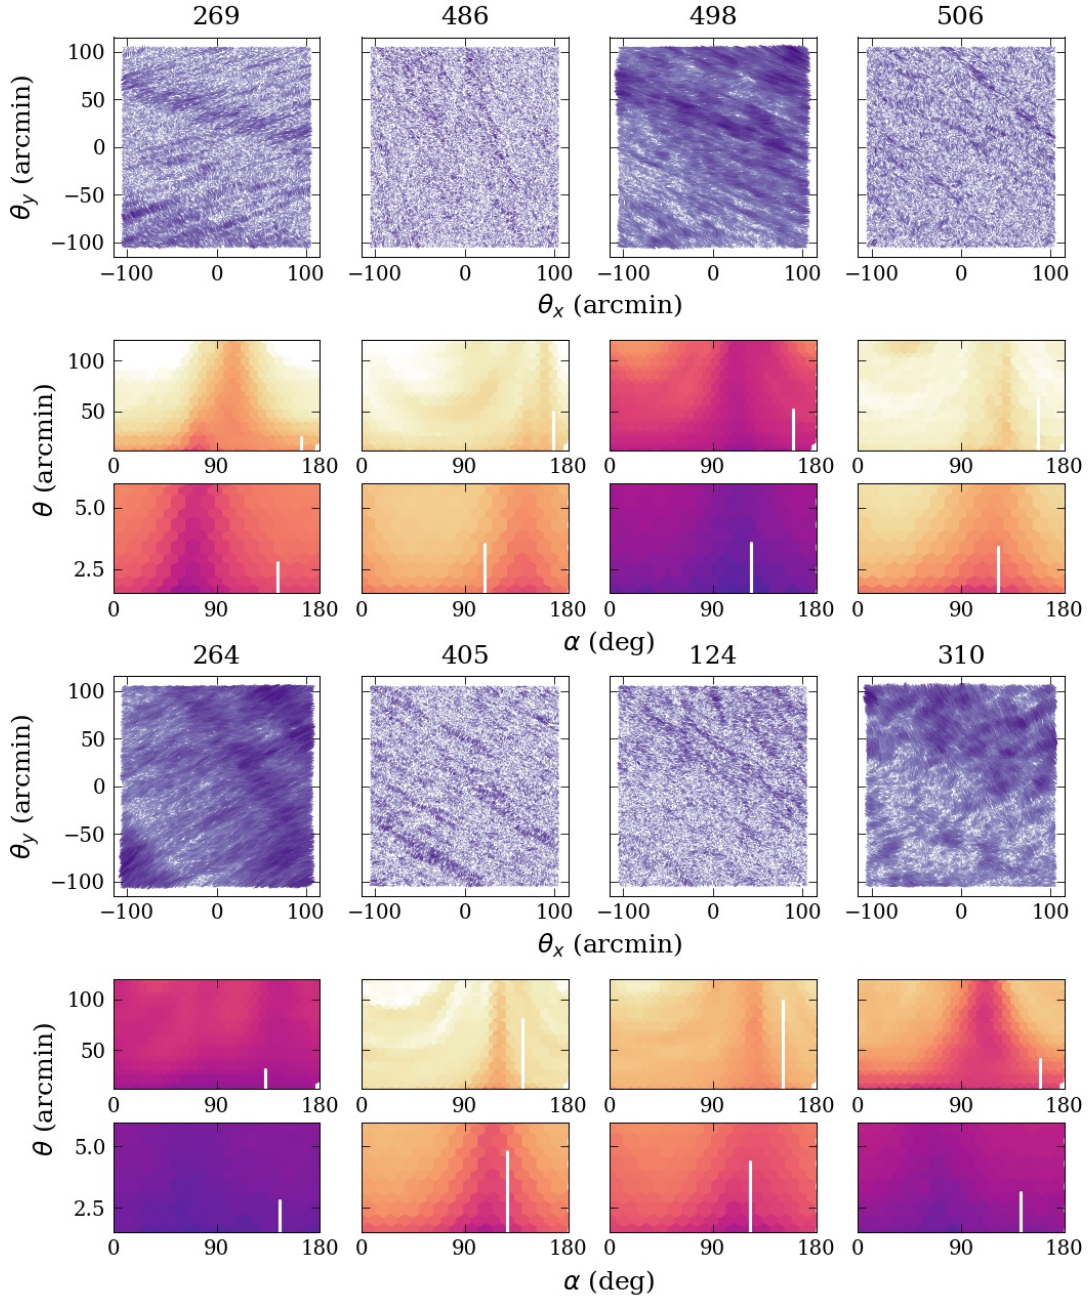


Figure B.3: Summary of simulated PSF shapes for eight PSFWS simulations. See text for details.

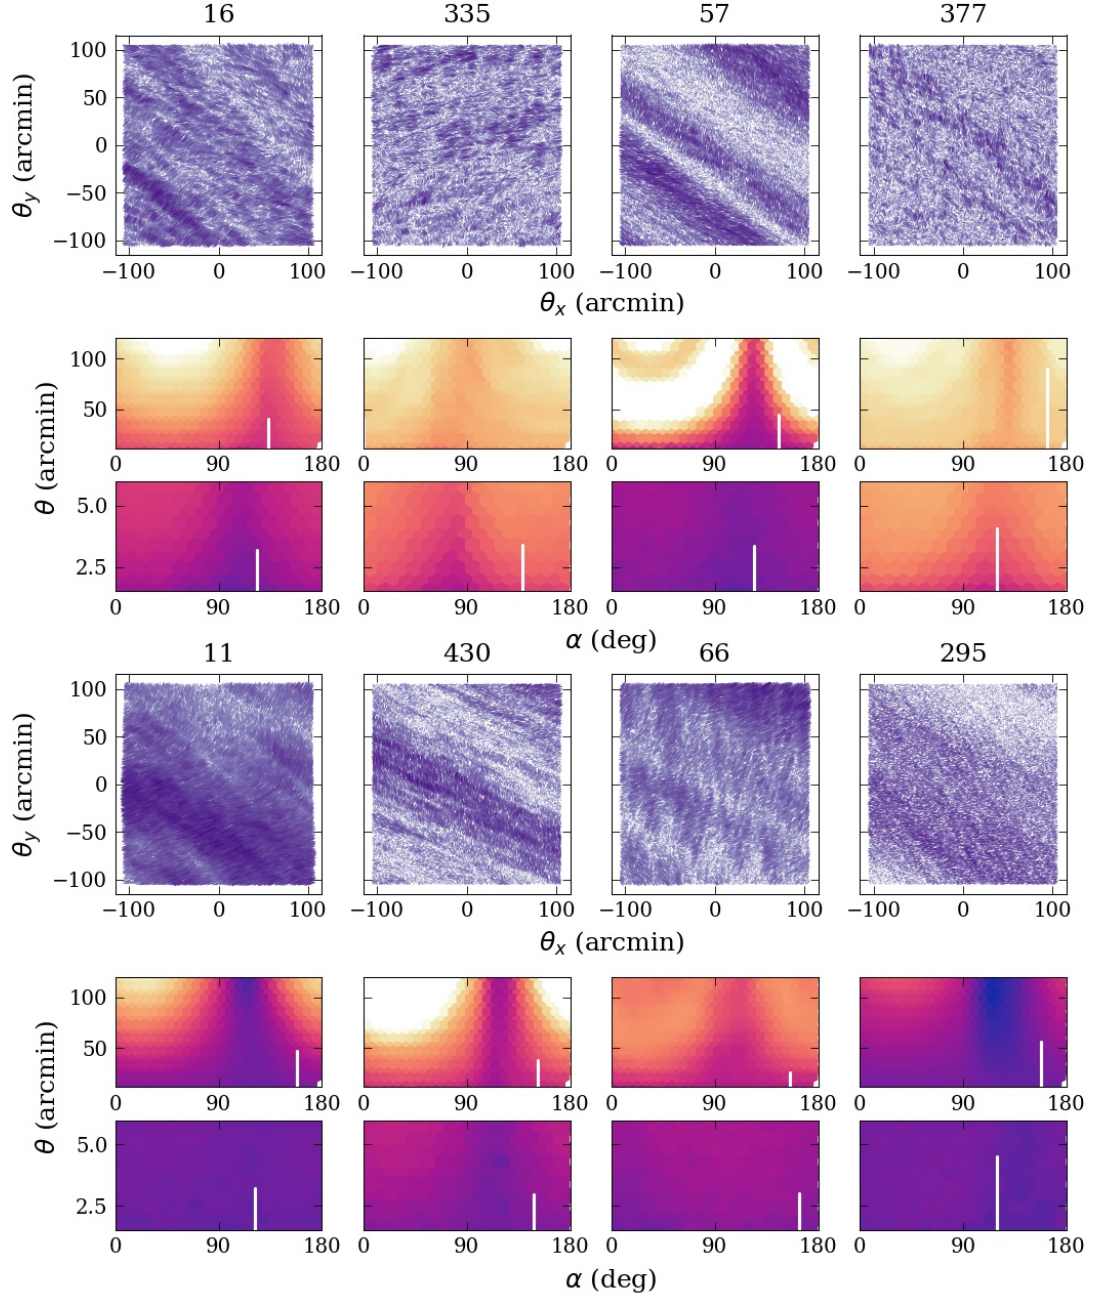


Figure B.4: Summary of simulated PSF shapes for eight PSFWS simulations. See text for details.

Appendix C

Simulation tests

The atmospheric simulation methods described in Section 3.1.1 become intractable in the long exposure regime: for an exposure time T , one must generate phase screens that are of size $\sim T \times v_{\max}$, where v_{\max} is the maximum speed for any of the simulated altitudes; typical values from PSF-WEATHER-STATION are $v_{\max} \approx 50$ m/s. In GalSim, a square array of optical turbulence strength values—i.e., phase shifts—sampled from the von Kármán power spectrum is stored in memory. The spatial resolution of these phase screens is given by the *screen_scale* parameter, whose value is roughly that of the Fried parameter, $0.1 - 0.2$ m. For a nominal Rubin exposure time of $T = 30$ s, we thus need to store a set of phase screen arrays of dimension N^2 , with $N \approx 1500 \text{ m} / 0.1 \text{ m} \approx 10^4$. This memory use scales as T^2 (since $N \sim T$). Therefore, going from an LSST exposure to a DES exposure, for example, results in a factor of ≈ 3 longer exposure time or a factor of 9 increase in memory requirement. This is just doable on computing clusters with large memory allocations; however, longer exposure times such as the 300 s simulations described in Chapter 6, or more modest memory availability, would present significant hurdles.

In this appendix, we describe two different approximation methods that reduce memory requirements. We ran simulations to quantify the effect of each approximation on the PSF statistics relative to the standard simulation method, which we denote as the “full” simulations.

The first approximation method uses GalSim-native periodic boundary conditions for the phase screens. The phase screens are smaller than $T \times v_{\max}$; when the simulation runs out of phase screen during the exposure, it simply loops back to the beginning – wrapping the phase screen. We refer to this method as a “wrap” simulation. This method allows a significant reduction in memory requirements, since one can use phase screens that are multiple times smaller than $T \times v_{\max}$; however, the wrapping of the phase screen could introduce extraneous spatial correlations due to repeated ray tracing through the same atmosphere – or at least similar atmosphere, depending on the exact details of the wrapping and boundary condition implementation.

The second approximation uses the linearity in exposure time to reduce memory: given identical

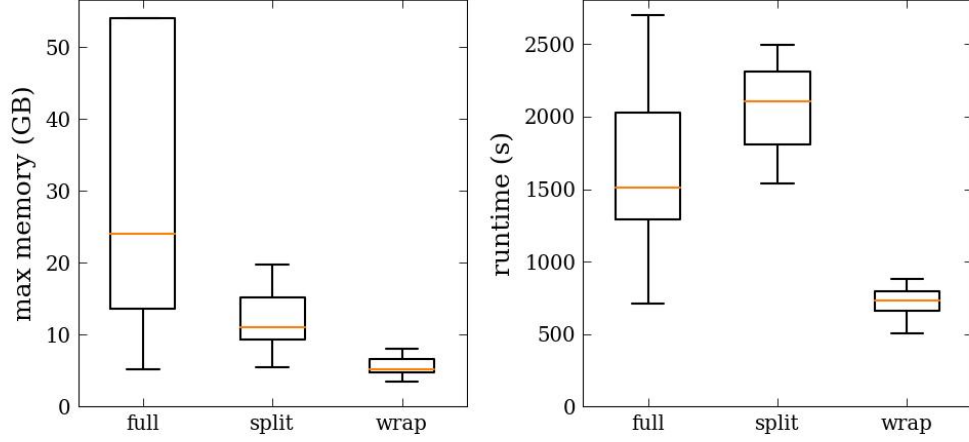


Figure C.1: Maximum memory usage in GB (left) and CPU runtime in seconds (right) for the three sets of 50 test simulations (full, split, and wrap). For each set in each panel, the median value is shown in red, the box spans a range that includes the two middle quartiles of values (i.e., the central 50%), and the error bars span all the data. Two runtime outliers (with runtimes ≈ 10000 s) are not shown for the full simulation.

phase screens, simulating several short back-to-back exposures and stacking the resulting PSFs should give the same result as one continuous exposure. Our implementation of the back-to-back exposure method is approximate: the phase screen in the first short simulation is a cutout (i.e., a subsection) of the single, full phase screens, and the screens for the remaining exposures are randomly generated from the same power spectrum—it is nontrivial to do otherwise given the variety of wind directions, pixelization, and the delayed instantiation of the phase screen arrays in GalSim. This method might result in a *less* correlated PSF than the wrap method, and relative to the full simulation, but would decrease memory use linearly in the number of separate exposures (which we call splits) used. We will refer to this method as: “split”.

We simulated 50 exposures of 10,000 PSFs on an LSST-size focal plan (3.5 deg) for each of the three methods (full, split, wrap). Other than the approximations described above, each of the 50 exposures was set up identically for the three cases. The total exposure time for each simulation was 30 s; this was divided into three 10-s exposures in the split condition. The boundary conditions in the wrap method were set such that the phase screen “wrapped” a maximum of three times during the 30-s exposure; i.e., the maximum phase screen size was $10\text{ s} \times v_{\text{max}}$.

The simulations were run on 12 CPUs at the Stanford research computing cluster, Sherlock¹. The memory and computational costs for the three simulation setups are summarized in Figure C.1, with memory usage in the left panel and runtime in the right panel. A simple calculation predicts

¹<https://www.sherlock.stanford.edu>

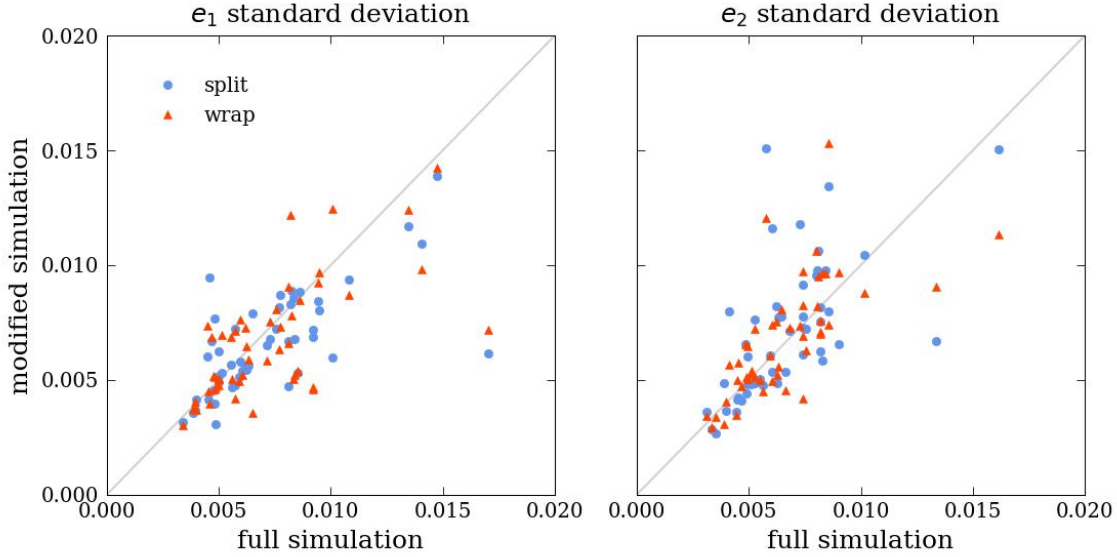


Figure C.2: Standard deviation of each component of PSF ellipticity across the focal plane for each of the modified simulations (vertical axis) versus the result for the full simulation (horizontal axis). Blue dots: split simulations, red triangles: wrapped simulations.

that the split scenario will need ~ 3 times as much memory as the wrapped simulation, and the full scenario ~ 9 times more. The results in Figure C.1 are roughly consistent with this estimate, but the medians (in red) show a shallower dependence with factors of ~ 2 and ~ 5 more memory than the wrapping method, for split and full simulations respectively. Computational cost comparison aside, we must also verify whether the relevant PSF statistics for the modified simulations are consistent with those for the full simulation.

We illustrate two metrics of spatial correlation of PSF parameters in Figure C.2. We plot the standard deviation across the field of view of the PSF ellipticity parameters e_1 (left) and e_2 (right) for each exposure. The results of the modified simulations are plotted against the same quantity from the full simulation. Both approximations show some random scatter but correlations between the full and modified results exists. The second metric of spatial correlation we consider is the dominant orientation of the spatial two-point correlation functions of PSF parameters, as described in detail in Section 4.4. In Figure C.3 we show the difference, between the modified and full simulations, in correlation orientations in e_1, e_2 , and PSF size σ . The results for most simulations for both approximation methods lie in the peak near 0 difference in angle, and the tails of the residual distributions are roughly symmetric.

Taken together, these results show that modifications to the atmospheric PSF simulation method that optimize computational efficiency (with either the split or wrapping scheme) introduce some scatter, but no significant bias in the distributions of PSF spatial correlations, both in orientation

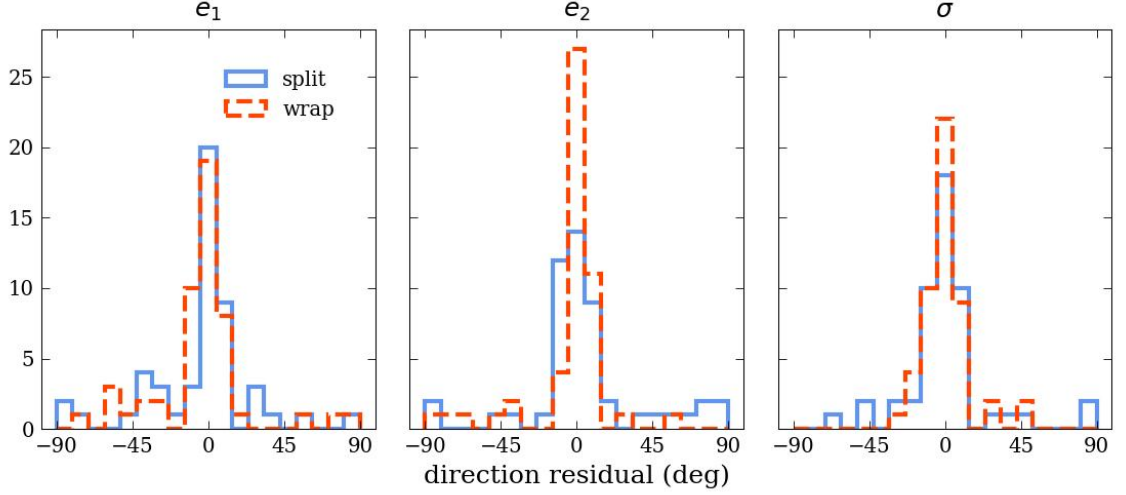


Figure C.3: Summary of 2PCF directions, as described in Section 4.1. Each histogram shows the difference in 2PCF direction, in degrees, between the modified and full simulations. Blue: split, red: wrap.

and magnitude. Therefore, we deem it reasonable to use either of these methods. For the simulations described in this thesis we will use the wrapping method due to its superior computational performance, although we use a larger maximum phase screen size than in the tests described in this section.

Appendix D

Exploration of adaptive optics telemetry

As we have seen throughout this thesis, it is important to model atmospheric contributions to the PSF accurately; current methods rely on simulations of atmospheric turbulence. In this appendix we describe an early exploratory project published in Hébert et al. (2018). Here we use data from the Gemini Planet Imager (GPI) to validate these atmospheric turbulence simulations and explore the simulation input parameter space. The GPI instrument uses adaptive optics (AO) to successfully image exoplanets at Gemini South, less than 2 km from the LSST site on the Cerro Pachón ridge in Chile. The signals recorded by the GPI AO system can be used to reconstruct wavefronts at the pupil, which can then be compared to the simulation-based predictions.

D.1 Wavefront data from adaptive optics instrumentation

The GPI adaptive optics system consists of a Shack-Hartmann wavefront sensor and two micro-electricalmechanical system (MEMS) deformable mirrors (Macintosh et al., 2007). Over a 60 s GPI exposure, the wavefront sensor uses 2×2 quad-cell centroiding to measure the slope of the wavefront on a 48×48 grid of subapertures (18 cm sampling). These measurements and corrections are performed at 1 kHz speed in the AO control system.

The phase of the wavefront is reconstructed in Fourier modes. High frequency and low frequency modes are separated and corrected by the tweeter and woofer deformable mirrors, respectively (Poyneer & Dillon, 2008). By saving the commands sent to these mirrors, one can later reconstruct the phase of the wavefront at the pupil (Snyder et al., 2016).

The reconstructed wavefront telemetry from GPI can provide information about the atmospheric conditions and parameters at Cerro Pachón. In particular, by comparing simulated wavefronts to

these data, we can validate certain aspects of atmospheric simulations (e.g., in GalSim). In addition, the data can serve to inform the choice of input simulation parameters. To explore the effect of various parameters, we calculate the temporal power spectrum of the wavefront, as well as the variance of the phase across the pupil. The datasets used here are 20 s or 30 s long.

D.2 Simulated wavefronts

The results discussed here are from simulations of three layers of atmosphere, each of which uses Von Kármán turbulence (von Kármán, 1948) with an outer scale of $L_0 = 25$ m. Unless specified, the simulations used Taylor’s frozen-flow hypothesis (Taylor, 1938). The parameters for the turbulence and wind for each layer are given in Table 1. The altitude and Fried parameter r_0 for each layer are taken from median seeing atmospheric measurements at Cerro Pachón (Snyder et al., 2016). The wind speed for each layer is drawn randomly from a uniform distribution between 0 and the maximum wind speed, which increases with altitude as shown in the table.

For comparison with the AO telemetry data, simulations were generated for continuous integrated exposures of 20 or 30 second, depending on the data of interest. The simulation time step used was 1 ms, to match the cadence of the GPI AO system. To test dependence of wavefront variation on turbulence strength, a second set of simulations was generated with Fried parameters a factor of $2^{3/5}$ smaller than those listed in Table D.1 (i.e., with turbulence strength C_n^2 a factor of 2 smaller). For a third set of simulations, we used the original values of r_0 shown in the table, but modified the “frozen-flow” hypothesis by introducing atmospheric boiling (Srinath et al., 2015): every 10 time steps (i.e., every 0.01 s), a small fraction $1 - \alpha$ of the turbulence amplitude in the phase screens is replaced by a random realization of Van Kármán turbulence. The fraction used here is $\alpha = 0.99$, so that 1% of the phase is replaced. In each of the three sets of simulations, wind directions were drawn randomly for each turbulence layer. The simulations were run with different numbers seeding the random generation of wind parameters/turbulence screens.

Table D.1: Simulation parameters for the three turbulence layers. The wind speed for each layer is drawn randomly from a uniform distribution between 0 and the maximum wind speed.

Layer altitude [km]	max wind speed [m/s]	r_0 [cm]
0	10	19
0.5	20	43
7.6	30	36

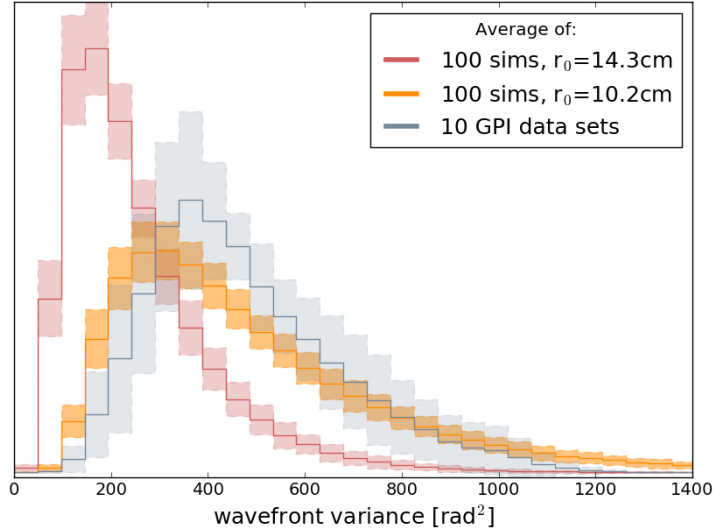


Figure D.1: Histograms of the wavefront variance across the telescope aperture computed for each kHz sample for GPI data and two simulations with different Fried parameters (20,000 variances corresponding to 20s total time). The solid bold histogram corresponds to the average variance in each bin for 10 data sets or 100 simulations while the shaded region indicates the $\pm 1\sigma$ spread.

D.3 Probing turbulence parameters

The variance of the reconstructed wavefronts across the telescope pupil provides one measure of the strength of turbulence. For GPI AO wavefront data and for GalSim simulations, we compute the variance at each point in time and then histogram the variance for a set of time samples. The peak of such a histogram is a measure of the average turbulence strength, and the spread indicates how much the strength varies over the total time – 20s for this analysis.

Three variance histograms are plotted in Figure D.1: for reconstructed GPI wavefront data, and for two simulations with turbulence strengths that are different by a factor of two, as described above. The solid bold histogram of each color corresponds to the average variance in the bin for 10 datasets (gray) or 100 simulations (orange and red), and the shaded regions indicate the $\pm 1\sigma$ (standard deviation) spread in the values of the variance in each bin. We see that in the simulations, the distribution of wavefront variance depends strongly on the value of r_0 , and that the simulation with stronger turbulence ($r_0 \sim 10.2$ cm) is a much better match to these datasets than ($r_0 \sim 14.3$ cm). Rather than quote the r_0 values for each of the layers in the simulation, we report the effective Fried parameter. For a set of atmospheric layers indexed by i , the effective Fried parameter is given by: $(\sum_i r_{0,i}^{-5/3})^{-3/5}$.

While r_0 is an important parameter in determining the turbulence strength, we see that in both the simulations (with fixed values of r_0) and the data the variance has a wide spread about the peak value.

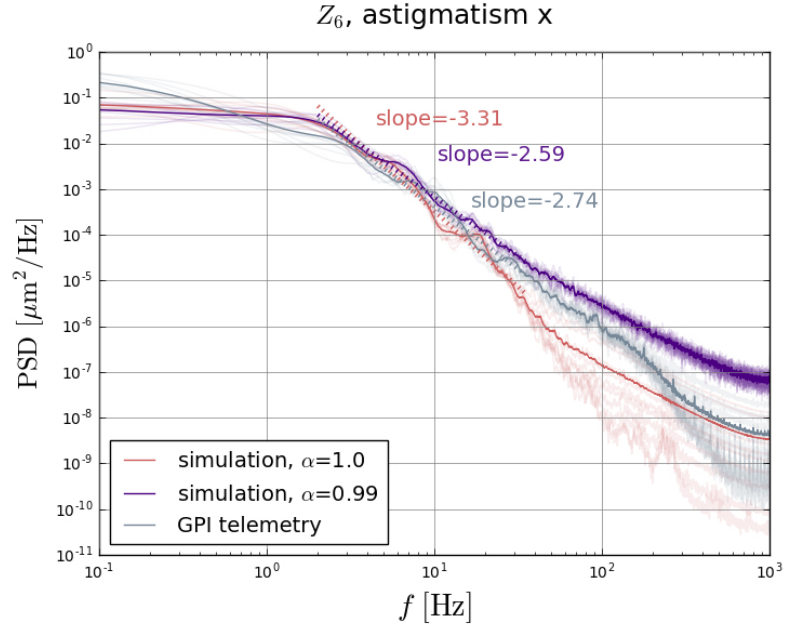


Figure D.2: Temporal power spectral density for Zernike coefficient Z_6 for GPI data (gray) and two sets of simulations (boiling off in red; boiling on in purple). The slopes are extracted by fitting a power law to the approximately linear region of the curves between 2 Hz and 40 Hz. Each pale line is the PSD from one simulation/dataset, and the average of the ten is shown in bold.

The amount of power at each temporal frequency – called the power spectral density (PSD) – is another way of characterizing turbulence (Snyder et al., 2016). In particular, we can compare the slopes of the PSDs for data and simulations with the power law slope expected for a Kolmogorov turbulence model. As outlined in Ref 78, we decompose the wavefronts from reconstruction of GPI telemetry and from simulations into Zernike functions, a complete basis set of orthogonal functions describing optical aberrations. The PSDs are then computed for each Zernike coefficient. As an example, we plot in Figure D.2 the PSD for the coefficient corresponding to Z_6 , as smoothed periodograms. The dark and pale curves correspond to the mean and spread, respectively, for ten realizations of two simulations – boiling off ($\alpha = 1.0$) and boiling on ($\alpha = 0.99$) – and ten datasets.

The expected slope of the power spectrum for spatial Kolmogorov turbulence is $-\frac{11}{3} \sim -3.67$. We expect the temporal power spectrum to have this same value due to the correspondence between spatial and temporal evolution of turbulence screens. We compare the Kolmogorov slope with those calculated from the PSDs in Figure D.2, for data and for the two simulations. We fit the slopes from the middle region of the spectrum (2 to 40 Hz) where the curves are approximately linear. In this region, the simulation with no atmospheric boiling is the closest to Kolmogorov, but its slope is less steep than the $-\frac{11}{3}$ prediction – perhaps due the presence of an outer scale, or the finite nature of the 20-second interval. Either of these factors could suppress power at low frequencies, reducing

the overall slope. The GPI data slope has an even weaker dependence on frequency than this first simulation. However, we find that adding atmospheric boiling to the simulation results in better agreement with the data. Varying the boiling parameter α could lead to a closer match, as could varying other parameters that influence the PSD slope – e.g. the outer scale.

D.4 Summary

A significant concern in the comparison of data and simulation is the lack of knowledge about correct input parameters for the atmospheric simulations – e.g., wind speed, wind direction, variation with altitude, etc. GPI telemetry from a wider range of seeing conditions could potentially guide this choice, especially when combined with data from seeing monitors and wind information from the National Oceanic and Atmospheric Administration (NOAA) databases. In particular, these data could inform a realistic range of simulation inputs to improve the study of PSF parameter behavior. In general, a greater range of parameters in both observational datasets and simulations will enable a more systematic study of the simulation input parameters and the asymptotic PSF behavior. We have addressed the concluding thoughts from this early exploratory study in the rest of this thesis.

Bibliography

- Albrecht, A., Bernstein, G., Cahn, R., et al. 2006, Report of the Dark Energy Task Force, Tech. Rep. FERMILAB-FN-0793-A, 897600, doi: 10.2172/897600
- Amon, A., Gruen, D., Troxel, M. A., et al. 2022, Physical Review D, 105, 023514, doi: 10.1103/PhysRevD.105.023514
- Asgari, M., Heymans, C., Hildebrandt, H., et al. 2019, Astronomy & Astrophysics, 624, A134, doi: 10.1051/0004-6361/201834379
- Asgari, M., Lin, C.-A., Joachimi, B., et al. 2021, Astronomy & Astrophysics, 645, A104, doi: 10.1051/0004-6361/202039070
- Bertin, E. 2011, in ASP Conference Proceedings, Vol. 442 (Boston, Massachusetts: Astronomical Society of the Pacific), 435. <https://ui.adsabs.harvard.edu/abs/2011ASPC..442..435B>
- Borgnino, J. 1990, Applied Optics, 29, 1863
- Carlsten, S. G., Strauss, M. A., Lupton, R. H., Meyers, J. E., & Miyazaki, S. 2018, Monthly Notices of the Royal Astronomical Society, 479, 1491, doi: 10.1093/mnras/sty1636
- Chabanier, S., Millea, M., & Palanque-Delabrouille, N. 2019, Monthly Notices of the Royal Astronomical Society, 489, 2247, doi: 10.1093/mnras/stz2310
- Chang, C., Marshall, P. J., Jernigan, J. G., et al. 2012, Monthly Notices of the Royal Astronomical Society, 427, 2572, doi: 10.1111/j.1365-2966.2012.22134.x
- Chun, M., Wilson, R., Avila, R., et al. 2009, Monthly Notices of the Royal Astronomical Society, 394, 1121, doi: 10.1111/j.1365-2966.2008.14346.x
- Clifford, S. F. 1978, in Laser Beam Propagation in the Atmosphere, ed. J. W. Strohbehn, Vol. 25 (Springer Berlin Heidelberg), 9–43, doi: 10.1007/3540088121_16
- Collaboration, P., Aghanim, N., Akrami, Y., et al. 2020, Astronomy & Astrophysics, 641, A6, doi: 10.1051/0004-6361/201833910

- Consortini, A., & Ronchi, L. 1972, *Applied Optics*, 11, 1205, doi: 10.1364/AO.11.001205
- Corrsin, S. 1951, *Journal of Applied Physics*, 22, 469, doi: 10.1063/1.1699986
- Crittenden, R. G., Natarajan, P., Pen, U., & Theuns, T. 2002, *The Astrophysical Journal*, 568, 20, doi: 10.1086/338838
- Davis, C. P., Rodriguez, J., & Roodman, A. 2016, in *SPIE Astronomical Telescopes + Instrumentation*, Edinburgh, United Kingdom, 990668, doi: 10.1117/12.2233366
- Dodelson, S. 2003, *Modern Cosmology* (Elsevier Science). <https://books.google.com/books?id=3oPRxdXJexcC>
- Egner, S. E., & Masciadri, E. 2007, *Publications of the Astronomical Society of the Pacific*, 119, 1441, doi: 10.1086/524850
- Ellerbroek, B. L. 2002, *Journal of the Optical Society of America A*, 19, 1803, doi: 10.1364/JOSA.19.001803
- Els, S. G., Schoeck, M., Bustos, E., et al. 2009, *Publications of the Astronomical Society of the Pacific*, 121, 922, doi: 10.1086/605450
- Fried, D. L. 1965, *Journal of the Optical Society of America*, 55, 1427, doi: 10.1364/JOSA.55.001427
- Gatti, M., Sheldon, E., Amon, A., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 504, 4312, doi: 10.1093/mnras/stab918
- Giblin, B., Heymans, C., Asgari, M., et al. 2021, *Astronomy & Astrophysics*, 645, A105, doi: 10.1051/0004-6361/202038850
- Goodman, J. W. 2005, *Introduction to Fourier Optics*, 3rd ed. (Englewood, CO: Roberts and Company)
- Haack, A., Gerding, M., & Lübken, F.-J. 2014, *Journal of Geophysical Research: Atmospheres*, 119, 10,605, doi: 10.1002/2013JD021008
- Hersbach, H., Bell, B., Berrisford, P., et al. 2020, *Quarterly Journal of the Royal Meteorological Society*, 146, 1999, doi: 10.1002/qj.3803
- Heymans, C., Rowe, B., Hoekstra, H., et al. 2012, *Monthly Notices of the Royal Astronomical Society*, no, doi: 10.1111/j.1365-2966.2011.20312.x
- Heymans, C., Van Waerbeke, L., Bacon, D., et al. 2006, *Monthly Notices of the Royal Astronomical Society*, 368, 1323, doi: 10.1111/j.1365-2966.2006.10198.x

- Hilbert, S., Hartlap, J., White, S. D. M., & Schneider, P. 2009, *Astronomy & Astrophysics*, 499, 31, doi: 10.1051/0004-6361/200811054
- Hirata, C., & Seljak, U. 2003, *Monthly Notices of the Royal Astronomical Society*, 343, 459, doi: 10.1046/j.1365-8711.2003.06683.x
- Howell, S. B., & Furlan, E. 2022, *Frontiers in Astronomy and Space Sciences*, 9, 871163, doi: 10.3389/fspas.2022.871163
- Huff, E. M., & Mandelbaum, R. 2017, 15, doi: <https://doi.org/10.48550/arXiv.1702.02600>
- Hébert, C.-A., Macintosh, B., & Burchat, P. 2018, in *Ground-based and Airborne Telescopes VII*, ed. R. Gilmozzi, H. K. Marshall, & J. Spyromilio (Austin, United States: SPIE), 198, doi: 10.1117/12.2314311
- Jarvis, M., Sheldon, E., Zuntz, J., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 460, 2245, doi: 10.1093/mnras/stw990
- Jarvis, M., Bernstein, G. M., Amon, A., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 501, 1282, doi: 10.1093/mnras/staa3679
- Jee, M. J., & Tyson, J. A. 2011, *Publications of the Astronomical Society of the Pacific*, 123, 596, doi: 10.1086/660137
- Jolissaint, L. 2010, *Journal of the European Optical Society: Rapid Publications*, 5, 10055, doi: 10.2971/jeos.2010.10055
- Kilbinger, M. 2015, *Reports on Progress in Physics*, 78, 086901, doi: 10.1088/0034-4885/78/8/086901
- Klueckers, V. A., Wooder, N. J., Adcock, M. A., Nicholls, T. W., & Dainty, C. 1997, in *Optical Science, Engineering and Instrumentation*, ed. L. R. Bissonnette & C. Dainty, San Diego, CA, 234–240, doi: 10.1117/12.283886
- Kolmogorov, A. 1941, *Doklady Akademii Nauk S.S.S.R.*, 30, 301
- Kornilov, V., Tokovinin, A., Shatsky, N., et al. 2007, *Monthly Notices of the Royal Astronomical Society*, 382, 1268, doi: 10.1111/j.1365-2966.2007.12467.x
- Krause, E., & Hirata, C. M. 2010, *Astronomy & Astrophysics*, 523, A28, doi: 10.1051/0004-6361/200913524
- Lawrence, R., & Strohbehn, J. 1970, *Proceedings of the IEEE*, 58, 1523, doi: 10.1109/PROC.1970.7977

- Li, X., Miyatake, H., Luo, W., et al. 2022, *Publications of the Astronomical Society of Japan*, 74, 421, doi: 10.1093/pasj/psac006
- Li, X., Zhang, T., Sugiyama, S., et al. 2023, *Hyper Suprime-Cam Year 3 Results: Cosmology from Cosmic Shear Two-point Correlation Functions*, arXiv. <http://arxiv.org/abs/2304.00702>
- Liaudat, T. I., Starck, J.-L., & Kilbinger, M. 2023. <http://arxiv.org/abs/2306.07996>
- LSST Science Collaborations, & LSST Project. 2009, *LSST Science Book 2.0*, LSST Science Collaborations and LSST Project. arXiv:0912.0201
- Léna, P., Lebrun, F., & Mignard, F. 1998, *Observational Astrophysics*, 2nd edn. (Berlin, Heidelberg: Springer)
- Macintosh, B., Graham, J., Palmer, D., et al. 2007, *Comptes Rendus Physique*, 8, 365, doi: 10.1016/j.crhy.2007.04.007
- Madurowicz, A., Macintosh, B. A., Ruffio, J.-B., et al. 2018, in *Adaptive Optics Systems VI*, ed. D. Schmidt, L. Schreiber, & L. M. Close (Austin, United States: SPIE), 230, doi: 10.1117/12.2314267
- Mandelbaum, R. 2018, *Annual Review of Astronomy and Astrophysics*, 56, 393, doi: 10.1146/annurev-astro-081817-051928
- Masciadri, E., Lascaux, F., Turchi, A., & Fini, L. 2017, *Monthly Notices of the Royal Astronomical Society*, 466, 520, doi: 10.1093/mnras/stw3111
- Masciadri, E., Vernin, J., & Bougeault, P. 2001, *Astronomy & Astrophysics*, 365, 699, doi: 10.1051/0004-6361:20000050
- Meerburg, P. D., Green, D., Flauger, R., et al. 2019, *Bulletin of the American Astronomical Society*, 51, 107, doi: https://ui.adsabs.harvard.edu/link_gateway/2019BAAS...51c.107M/doi:10.48550/arXiv.1903.04409
- Meyers, J. E., & Burchat, P. R. 2015a, *The Astrophysical Journal*, 807, 182, doi: 10.1088/0004-637X/807/2/182
- . 2015b. <http://arxiv.org/abs/1505.02307>
- Obukhov, A. M. 1949, *Im. Akd. Nauk SSSR, Ser. Geogra. Geofz.*, 13, 58
- Ogane, H., Akiyama, M., Oya, S., & Ono, Y. 2021, *Monthly Notices of the Royal Astronomical Society*, 503, 5778, doi: 10.1093/mnras/stab105
- Osborn, J., & Sarazin, M. 2018, *Monthly Notices of the Royal Astronomical Society*, 480, 1278, doi: 10.1093/mnras/sty1898

- Osborn, J., Wilson, R. W., Sarazin, M., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 478, 825, doi: 10.1093/mnras/sty1070
- Perrefort, D., Wood-Vasey, W. M., Bostroem, K. A., et al. 2019, *Publications of the Astronomical Society of the Pacific*, 131, 025002, doi: 10.1088/1538-3873/aaecb7
- Peterson, J. R., Jernigan, J. G., Kahn, S. M., et al. 2015, *The Astrophysical Journal Supplement Series*, 218, 14, doi: 10.1088/0067-0049/218/1/14
- Poyneer, L. A., & Dillon, D. 2008, in *MOEMS-MEMS Micro and Nanofabrication*, ed. S. S. Olivier, T. G. Bifano, & J. A. Kubby, San Jose, CA, 68880H, doi: 10.1117/12.772754
- Roddier, F. 1981, in *Progress in Optics*, Vol. 19 (Elsevier), 281–376, doi: 10.1016/S0079-6638(08)70204-X
- Rowe, B. 2010, *Monthly Notices of the Royal Astronomical Society*, doi: 10.1111/j.1365-2966.2010.16277.x
- Rowe, B., Jarvis, M., Mandelbaum, R., et al. 2015, arXiv:1407.7676 [astro-ph]. <http://arxiv.org/abs/1407.7676>
- Ryden, B. 2003, *Introduction to Cosmology* (Addison-Wesley). <https://books.google.com/books?id=z27vAAAAAAAJ>
- Sarazin, M., & Roddier, F. 1990, *Astronomy and Astrophysics*, 227, 294. <https://ui.adsabs.harvard.edu/abs/1990A&A...227..294S>
- Schneider, P., Eifler, T., & Krause, E. 2010, *Astronomy and Astrophysics*, 520, A116, doi: 10.1051/0004-6361/201014235
- Schneider, P., Kochanek, C. S., & Wambsganss, J. 2006, *Gravitational lensing: strong, weak, and micro*, Saas-Fee advanced course / Swiss Society for Astrophysics and Astronomy No. 33 (Berlin ; New York: Springer)
- Schneider, P., Van Waerbeke, L., & Mellier, Y. 2002, *Astronomy & Astrophysics*, 389, 729, doi: 10.1051/0004-6361:20020626
- Schneider, P., Asgari, M., Jozani, Y. N., et al. 2022, *Astronomy & Astrophysics*, 664, A77, doi: 10.1051/0004-6361/202142479
- Scott, N. J., Howell, S. B., Gnilka, C. L., et al. 2021, *Frontiers in Astronomy and Space Sciences*, 8, 716560, doi: 10.3389/fspas.2021.716560

- Sheldon, E. S., Becker, M. R., Jarvis, M., Armstrong, R., & LSST Dark Energy Science Collaboration. 2023, *The Open Journal of Astrophysics*, 6, 10.21105/astro.2303.03947, doi: 10.21105/astro.2303.03947
- Sheldon, E. S., Becker, M. R., MacCrann, N., & Jarvis, M. 2020, *The Astrophysical Journal*, 902, 138, doi: 10.3847/1538-4357/abb595
- Sheldon, E. S., & Huff, E. M. 2017, *The Astrophysical Journal*, 841, 24, doi: 10.3847/1538-4357/aa704b
- Shepherd, H. W., Osborn, J., Wilson, R. W., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 437, 3568, doi: 10.1093/mnras/stt2150
- Snyder, A., Srinath, S., Macintosh, B., & Roodman, A. 2016, in *SPIE Astronomical Telescopes + Instrumentation*, ed. H. J. Hall, R. Gilmozzi, & H. K. Marshall, Edinburgh, United Kingdom, 990642, doi: 10.1117/12.2234362
- Sreenivasan, K. R. 1999, *Reviews of Modern Physics*, 71, S383, doi: 10.1103/RevModPhys.71.S383
- Srinath, S., Poyneer, L. A., Rudy, A. R., & Ammons, S. M. 2015, in *Imaging and Applied Optics 2015* (Arlington, Virginia: OSA), JW4F.5, doi: 10.1364/AOMS.2015.JW4F.5
- Tallis, M., Bailey, V. P., Macintosh, B., et al. 2020, *Journal of Astronomical Telescopes, Instruments, and Systems*, 6, 1, doi: 10.1117/1.JATIS.6.1.015002
- Tatarskii, V. I. 1971, *The effects of the turbulent atmosphere on wave propagation*
- Taylor, G. I. 1938, *Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences*, 164, 476, doi: 10.1098/rspa.1938.0032
- The LSST Dark Energy Science Collaboration, Mandelbaum, R., Eifler, T., et al. 2021a, *The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document*, arXiv. <http://arxiv.org/abs/1809.01669>
- The LSST Dark Energy Science Collaboration, Abolfathi, B., Alonso, D., et al. 2021b, *The Astrophysical Journal Supplement Series*, 253, 31, doi: 10.3847/1538-4365/abd62c
- Tokovinin, A. 2002, *Publications of the Astronomical Society of the Pacific*, 114, 1156, doi: 10.1086/342683
- Tokovinin, A., Baumont, S., & Vasquez, J. 2003a, *Monthly Notices of the Royal Astronomical Society*, 340, 52, doi: 10.1046/j.1365-8711.2003.06231.x
- Tokovinin, A., & Kornilov, V. 2002, 266, 104. <https://ui.adsabs.harvard.edu/abs/2002ASPC..266..104T>

- Tokovinin, A., Kornilov, V., Shatsky, N., & Voziakova, O. 2003b, *Monthly Notices of the Royal Astronomical Society*, 343, 891, doi: 10.1046/j.1365-8711.2003.06731.x
- Tokovinin, A., & Travouillon, T. 2005, *Monthly Notices of the Royal Astronomical Society*, 365, 1235, doi: 10.1111/j.1365-2966.2005.09813.x
- Tokovinin, A., Vernin, J., Ziad, A., & Chum, M. 2005, *Publications of the Astronomical Society of the Pacific*, 117, 395, doi: 10.1086/428930
- Tokovinin, A. A., Ziad, A., Martin, F., et al. 1998, in *Adaptive Optical System Technologies*, ed. D. Bonaccini & R. K. Tyson, Vol. 3353 (Kona, HI: SPIE), 1155 – 1162, doi: 10.1117/12.321672
- Troxel, M., MacCrann, N., Zuntz, J., et al. 2018, *Physical Review D*, 98, 043528, doi: 10.1103/PhysRevD.98.043528
- Vernin, J., Crochet, M., Azouit, M., & Ghebrebrhan, O. 1990, *Radio Science*, 25, 953, doi: 10.1029/RS025i005p00953
- von Kármán, T. 1948, *Proceedings of the National Academy of Sciences*, 34, 530, doi: 10.1073/pnas.34.11.530
- Weinberg, D. H., Mortonson, M. J., Eisenstein, D. J., et al. 2013, *Physics Reports*, 530, 87, doi: 10.1016/j.physrep.2013.05.001
- Xin, B., Ivezić, Z., Lupton, R. H., et al. 2018, *The Astronomical Journal*, 156, 222, doi: 10.3847/1538-3881/aae316
- Yaglom, A. M. 1949, *Doklady Akademii Nauk S.S.S.R.*, 69, 743
- Zhang, T., Mandelbaum, R., The LSST Dark Energy Science Collaboration, et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 510, 1978, doi: 10.1093/mnras/stab3584
- Zhang, T., Almoubayyed, H., Mandelbaum, R., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 520, 2328, doi: 10.1093/mnras/stac3350