LASER-BASED WIRELESS POWER TRANSFER TO DRONES: ATMOSPHERIC ABERRATIONS AND ADAPTIVE OPTICS MITIGATION

by
Nicolas Van Hoorick

School of Electrical Engineering and Computer Science,
The University of Queensland.

Submitted for the degree of Electrical Engineering & Commerce as part
of the course REIT4841.

August 2025.

November 11, 2025

Prof Michael Bruenig
Head of School
School of Electrical Engineering and Computer Science
The University of Queensland
St Lucia, Q 4072

Dear Professor Bruenig,

In accordance with the requirements of the Degree of Bachelor of Engineering (Honours) in the School of Electrical Engineering and Computer Science, I submit the following thesis entitled "Laser-Based Wireless Power Transfer to Drones: Atmospheric Aberrations and Adaptive Optics Mitigation" The thesis was performed under the supervision of Dr Martin Ploschner. I declare that the work submitted in the thesis is my own, except as acknowledged in the text and footnotes, and that it has not previously been submitted for a degree at the University of Queensland or any other institution.

Yours sincerely,

Nicolas Van Hoorick.

Declaration by Author

This thesis is composed of my original work, and contains no material previously published, written by another person, or generated by artificial intelligence (AI), except where clearly referenced and acknowledged either in the main text or the preliminary pages of this thesis. I have clearly stated the contribution of others to jointly authored works that I have included in my thesis.

I have clearly stated the contribution of others to my thesis as a whole, including statistical assistance, survey design, data analysis, significant technical procedures, professional editorial advice, use of AI tools, financial support, and any other original research work used or reported in my thesis. The content of my thesis is the result of work I have carried out since the commencement of my enrollment in the course code indicated on the title page, and does not include a substantial part of work that has been submitted to qualify for the award of any other degree or diploma in any university or other tertiary institution. I have clearly stated which parts of my thesis, if any, have been submitted to qualify for another award.

I acknowledge that copyright of all material contained in my thesis resides with the copyright holder(s) of that material. Where appropriate I have obtained copyright permission from the copyright holder to reproduce material in this thesis and have sought permission from co-authors for any jointly authored works included in the thesis.

Use of AI Statement

I acknowledge the use of generative AI tools in completing this assignment. Details of which tools were used and how they were used are provided in the table below, along with appropriate in-text and full references. I take responsibility for critically evaluating and integrating the AI-generated content, and ensuring it adheres to academic integrity standards.

AI Model and date	Artificial Intelligence (AI) tool usage (tick all that apply)								
	language translation	grammar/style	planning/drafting	research/background information	content creation text	content creation visual	content creation code	feedback	other (provide details)
ChatGPT 5 11/10/25		✓	✓	✓	✓	✓	✓	✓	

Acknowledgments

I would like to acknowledge my supervisor, Dr. Martin Ploschner, for his invaluable guidance throughout this research project. Dr. Ploschner provided essential technical expertise in developing the Zernike polynomial simulation code that forms the foundation of this work. His weekly meetings were instrumental in deepening my understanding of atmospheric turbulence theory and adaptive optics principles. The constructive feedback he provided on both the theoretical framework and simulation methodologies significantly enhanced the quality and rigor of this thesis. I also acknowledge the technical support provided by the School of Electrical Engineering and Computer Science computing facilities, which enabled the extensive numerical simulations performed in this research. The support of family and friends throughout the duration of this project is also gratefully acknowledged.

Abstract

Unmanned Aerial Vehicles (UAVs) have seen widespread adoption across surveillance, communication, agriculture, and disaster response, yet their endurance remains a key limitation due to constrained onboard energy storage. Laser-based wireless power transfer (WPT) has emerged as a promising approach to extend UAV flight time by transmitting high-intensity optical beams from ground stations to photovoltaic (PV) receivers mounted on drones [2, 1]. Compared to RF or inductive methods, laser WPT offers higher energy density, longer range, and precise beam shaping. However, its effectiveness is fundamentally limited by atmospheric turbulence: small-scale refractive index variations in the air distort the laser wavefront, causing beam spreading, wander, and intensity fluctuations. These aberrations lead to lower power incident on the drone's PV receiver, thereby reducing charging efficiency [3, 4].

This thesis investigates the impact of atmospheric aberrations on laser-based drone charging and evaluates adaptive optics mitigation strategies. A comprehensive simulation framework was developed to model Gaussian beam propagation through turbulent atmospheres using Zernike polynomial representations of phase distortions. The research quantifies how atmospheric turbulence affects wireless power transfer efficiency across varying propagation distances and turbulence strengths.

The simulation results indicate that correcting the first 15 Zernike modes can restore 70–85% of the ideal power transfer capability under moderate turbulence conditions. This work provides a foundation for developing robust laser-powered UAV systems and demonstrates the critical importance of atmospheric compensation for practical implementation of wireless optical power transfer.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Unmanned aerial vehicles—referred to here simply as drones—now underpin missions spanning surveillance, emergency response, precision agriculture, and asset inspection. Their deployment is still constrained by the modest energy that can be stored on-board, forcing frequent landings or battery swaps during long sorties.

Laser-based wireless power transfer (laser WPT) has emerged as a way to extend endurance by beaming optical power from a ground station to a photovoltaic receiver on the drone [2, 1]. Compared with radio-frequency or inductive coupling, laser power beaming delivers higher power density, can operate over kilometre ranges, and allows tight control of the illuminated spot on a moving vehicle.

The atmosphere, however, introduces dynamic aberrations that distort the transmitted wavefront. Random variations in refractive index cause the beam to broaden, wander, and scintillate before it reaches the photovoltaic surface, cutting the net charging efficiency [3, 4]. Research in free-space optical communication has grappled with the same physics for decades [9, 10], developing adaptive optics systems that sense and correct turbulence-induced phase errors using deformable mirrors or holographic phase modulators [13, 10].

This thesis translates those insights to drone charging. We build simulation tools to quantify how turbulence alters the delivered power, evaluate the effectiveness of correcting the dominant aberration modes, and highlight adaptive optics strategies capable of stabilising laser power transfer in real atmospheric conditions.

Chapter 2

Literature Review / State of the Art

2.1 Atmospheric Turbulence and Wave Propagation

Atmospheric turbulence consists of random spatio-temporal variations in the refractive index of air, caused by fluctuating temperature and pressure inhomogeneities. According to Kolmogorov's theory, energy cascades from larger eddies to smaller eddies in a turbulent flow. In the context of optical propagation, this leads to a power-law spectrum of refractive index fluctuations. A commonly used measure of turbulence strength is the refractive index structure constant $C_n^2(h)$, which typically varies with altitude h, with the strongest fluctuations usually occurring close to the ground. Light propagating through turbulence experiences random phase perturbations as it encounters varying refractive indices.

For a plane wave passing through a path of turbulence, the statistical effect on the wavefront phase $\phi(\mathbf{r})$ can be characterized by the phase structure function $D_\phi(\rho)$, defined as the mean squared difference in phase between two points separated by a transverse distance ρ:

$$D_\phi(\rho) \;=\; \Big\langle [\,\phi(\mathbf{r}+\boldsymbol{\rho}) - \phi(\mathbf{r})\,]^2 \Big\rangle. \tag{2.1}$$

In Kolmogorov turbulence (assuming an infinite outer scale), $D_\phi(\rho)$ follows a 5/3 power-law with separation. Specifically,

$$D_\phi(\rho) \approx 6.88 \left(\frac{\rho}{r_0}\right)^{5/3}, \tag{2.2}$$

where r_0 is the Fried parameter (coherence diameter) [11]. The Fried parameter r_0 physically represents the diameter of an unobstructed telescope aperture for which the atmospheric phase variance is 1 rad^2. It scales with wavelength λ and path length L roughly as $r_0 \propto \lambda^{6/5}(\int_0^L C_n^2(h)\,dh)^{-3/5}$. Stronger turbulence (larger C_n^2 or longer paths L) yields a smaller r_0, meaning worse seeing conditions. At a good astronomical site r_0 in the visible might be on the order of 10–20 cm at night, whereas daytime or urban conditions can have r_0 of only a few centimeters or less.

In reality, turbulence has finite inner and outer scales. The outer scale L_0 is the size of the largest energy-containing eddies, beyond which the spectrum of fluctuations flattens out. This limits the growth of $D_\phi(\rho)$ for separations larger than L_0. Measurements at observatory sites show outer scales on the order of tens of meters [12]; for instance, Ziad *et al.* found $L_0 \approx$ 10–20 m at the Palomar 5 m telescope site using multiple instruments [12]. The inner scale (on the order of millimeters to centimeters) is the scale at which turbulent energy dissipates into heat; it causes a cutoff at very high spatial frequencies of refractive index fluctuations. In most wave propagation analyses, the Kolmogorov model (with an implicit infinite L_0) or the von Kármán spectrum (which includes L_0) is used to model turbulence-induced phase perturbations across the optical beam.

Turbulence-induced phase aberrations lead to various deleterious effects on a propagating optical beam. As the initially planar or mildly divergent wavefront is distorted, the beam can spread beyond its diffraction-limited divergence, the centroid of the beam can wander (angle-of-arrival fluctuations), and the intensity pattern can break up, producing speckle and intensity drops (scintillation). These effects not only blur images but also cause fluctuations in received optical power in communication systems. The severity of these effects is often quantified by parameters like the Strehl ratio (the ratio of on-axis intensity to the ideal diffraction peak), or the scintillation index for intensity fluctuations. In strong turbulence, the Strehl ratio can be very low and scintillation index high, indicating a highly distorted beam. Without mitigation, these conditions pose a serious challenge for maintaining reliable FSOC links and, by extension, for delivering steady optical power to a charging drone [8].

2.2 Gaussian Beam Optics in Free Space and Turbulence

Many optical communication systems employ Gaussian beam profiles (e.g., fundamental mode lasers) for transmitting signals. In free space (vacuum or no turbulence), a Gaussian beam of wavelength λ and waist radius w_0 will propagate with a growing beam radius $w(z)$ given by

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2}, \qquad z_R = \frac{\pi w_0^2}{\lambda}, \tag{2.3}$$

where z_R is the Rayleigh range of the beam. Within $|z| < z_R$ the radius follows $w(z) = w_0\sqrt{1 + (z/z_R)^2}$, reaching about $1.4\, w_0$ at $|z| = z_R$, before diverging approximately linearly for $|z| \gg z_R$ with half-angle $\theta \approx \lambda/(\pi w_0)$. The on-axis intensity of an ideal Gaussian beam falls as $I(0, z) = I_0/(1 + (z/z_R)^2)$ with distance, due purely to diffraction spreading.

Atmospheric turbulence modifies Gaussian beam propagation in several ways. The instantaneous beam profile is no longer Gaussian, as phase distortions cause energy to redistribute. One notable effect is beam wander: the beam's centroid randomly shifts as a function of range due to large-scale tilt fluctuations (these correspond to the lowest-order angular aberrations in the wavefront). Turbulence also increases the effective beam divergence and can cause filamentation of the beam into multiple intensity patches at the receiver plane. The time-averaged beam spot size at the receiver is typically larger than the diffraction-limited spot, and the time-averaged intensity is lower on-axis (reduced Strehl ratio). Additionally, fluctuations in intensity (scintillation) occur, especially in moderate to strong turbulence, leading to deep fades at the receiver plane and large swings in delivered charging power.

Several theoretical frameworks exist to predict Gaussian beam statistics in turbulence. For weak turbulence, one can use Rytov perturbation theory to find the variance of log-amplitude fluctuations (scintillation index) and beam spread. For stronger turbulence, phase screen models and wave-optics simulations are used. The concept of a partially coherent beam has been explored as a mitigation: if the beam is intentionally given a finite coherence width (through mode mixing or a multimode source), scintillation can be reduced at the expense of some increased beam size. Ricklin and Davidson [9] analyzed partially coherent Gaussian beams and found that reducing spatial coherence can indeed lower the turbulence-induced intensity fluctuations in the beam, implying a trade-off between beam brightness and scintillation. Such approaches are part of the broader toolkit for turbulence mitigation, complementing adaptive optics.

2.3 Wavefront Aberrations and Zernike Polynomials

A convenient way to describe the phase aberrations imposed by turbulence across an aperture (such as a telescope mirror or a laser beam's cross-section) is to use an orthonormal basis of functions defined on the unit disk. Zernike polynomials are the standard basis for this purpose in optics. They are defined on a unit circle (usually corresponding to a circular pupil) and are indexed by two integers (n, m) representing the radial order n and azimuthal frequency m. The Zernike polynomials $Z_n^m(\rho, \theta)$ can be separated into a radial part $R_n^m(\rho)$ (a polynomial in the radial coordinate ρ) and an angular part ($\cos m\theta$ or $\sin m\theta$). For example, a Zernike polynomial in the Noll normalization can be written as:

$$Z_n^m(\rho, \theta) = N_{n,m}\, R_n^m(\rho) \begin{cases} \cos(m\theta), & m \geq 0, \\ \sin(|m|\theta), & m < 0, \end{cases} \tag{2.4}$$

where $0 \leq \rho \leq 1$ is the normalized radius, θ is the azimuthal angle, and $N_{n,m}$ is a normalization constant chosen such that $\int_0^1 \int_0^{2\pi} [Z_n^m(\rho, \theta)]^2 \rho\, d\rho\, d\theta = 1$. The radial polynomials $R_n^m(\rho)$ are given by a closed-form series:

$$R_n^m(\rho) = \sum_{k=0}^{(n-m)/2} (-1)^k \frac{n-k!}{k!\, \left(\frac{n+m}{2} - k\right)!\, \left(\frac{n-m}{2} - k\right)!}\, \rho^{n-2k}, \tag{2.5}$$

for nonnegative even $n - m$. Each pair (n, m) corresponds to a specific aberration pattern; for instance, $n = 1, m = \pm 1$ are tip and tilt (linear tilt of the wavefront), $n = 2, m = 0$ is defocus, $n = 2, m = \pm 2$ are astigmatism, etc. Noll [11] introduced a single-index scheme to order Zernike modes by their variance in atmospheric turbulence, with the lowest indices capturing the most commonly encountered aberrations (tip, tilt, defocus, astigmatism, coma, etc.).

In a turbulent wavefront, the phase $\phi(\rho, \theta)$ over a circular aperture of diameter D can be expanded in Zernike modes:

$$\phi(\rho, \theta) = \sum_{j=1}^{\infty} a_j Z_j(\rho, \theta), \tag{2.6}$$

where Z_j is the j-th Zernike mode (following Noll's indexing convention) and a_j are the coefficients. These coefficients are random variables due to turbulence. One of the key results from Noll's 1976 paper is that for Kolmogorov turbulence, the variance of the Zernike coefficients a_j follows the Kolmogorov spectrum. Specifically, the variance scales approximately as $a_j^2 \propto (D/r_0)^{5/3}$ and decays with the mode order (higher-order modes have smaller variance). Low-order modes like tip and tilt carry the largest portion of atmospheric phase variance. In fact, the tilt modes correspond to angular deviation of the beam and, in the absence of an outer scale, their variance formally diverges as the aperture grows large, reflecting that larger apertures capture larger tilt shifts across them. This is why tip-tilt is often dominant and requires special handling in AO systems (e.g., a separate fast steering mirror or a natural guide star for reference).

Using Zernike expansions is very convenient for adaptive optics. Deformable mirror surfaces and wavefront sensor reconstructor algorithms often operate in this modal basis. By correcting the most significant Zernike modes (or equivalently, the first few aberration terms), one can achieve a large improvement in image or beam quality even if very high-order modes remain uncorrected. For instance, correcting just tip and tilt can eliminate the mean angular wander of a beam, and correcting up through defocus and astigmatism can sharpen the beam focus significantly. This modal approach is commonly contrasted with a zonal approach (correcting every sub-aperture or actuator independently). The Zernike basis is particularly well-suited for describing atmospherically induced phase distortions [11].

2.4 Principles of Adaptive Optics Systems

Adaptive optics systems seek to actively compensate wavefront errors in real time. A typical AO loop consists of a wavefront sensor (WFS), a control computer, and a deformable mirror (DM) or other phase modulator. The WFS measures the distorted wavefront coming from a reference source (which could be a natural star, a laser guide star, or a retroreflected beacon). Common WFS types include the Shack-Hartmann sensor (an array of micro-lenses focusing sub-aperture images to sense local slope of the wavefront) and interferometric sensors (like curvature sensors or pyramid sensors). The WFS produces an error signal indicating the aberration. The control system then computes the appropriate correction and commands the DM to introduce the conjugate phase to cancel the aberration. This correction is typically updated at tens to a few hundreds of hertz, fast enough to track the Greenwood frequency of most atmospheric paths.

In an ideal closed-loop AO system, the outgoing wavefront (after the DM) is the negation of the incoming aberration, resulting in a flat wavefront and thus a corrected beam or image. In practice, the correction is never perfect due to time delays (servo lag), finite spatial correction degrees-of-freedom (limited actuator count or corrected modes), and sensor noise. The performance of an AO system is often summarized by the residual wavefront error (often expressed in nm RMS or in Strehl ratio after correction). For instance, modern astronomical AO platforms often compress an uncorrected $\sim$1000 nm RMS wavefront down to roughly 50–100 nm RMS under good seeing, thereby raising the Strehl ratio from a few percent to $\sim$50%. Low-order aberrations (tip/tilt, defocus, etc.) are typically easier to correct, while very high spatial frequency distortions may remain uncorrected due to limited actuator density.

In astronomical applications, AO has evolved into a mature technology. A notable challenge there is obtaining a bright reference star in the field of view. Laser guide stars (LGS) are used to create an artificial reference by exciting mesospheric sodium atoms to produce a glowing spot in the sky. As mentioned, LGS-based AO cannot measure absolute tip-tilt (because the laser beam itself is affected by tip-tilt on the up and down path equally, making the spot largely stationary on the detector) [10]. Thus, most LGS AO systems still require a nearby natural guide star for tip-tilt correction. In the context of FSOC, the situation is analogous: the optical transmitter needs a reference to sense its own outgoing beam distortions. One concept is to use a cooperative target (satellite) that sends down a pilot beacon, effectively serving as a reference for the uplink path [10]. But the point-ahead angle (the necessity to point the uplink beam slightly ahead of the instantaneous downlink beam to hit the moving satellite) means the downlink and uplink paths are not identical, leading to anisoplanatic errors.

Another concept specific to communications is the use of a 'retro-reflector" or relay mirror in space: Armstrong *et al.* [14] proposed placing a relay mirror (e.g., in low-Earth orbit) which sends a reference laser beam down to the ground terminal. The ground terminal measures the wavefront of this downlink beam (which has sampled the same turbulence as the uplink will) and then pre-corrects the uplink beam accordingly. This effectively creates a dynamic optical beacon traveling the reverse of the uplink path. While technically complex, this idea underscores the creativity required to obtain reference wavefront measurements for uplink AO.

In summary, the background principles drawn from atmospheric science and adaptive optics provide the foundation for this research. With an understanding of how turbulence distorts optical waves and how those distortions can be characterized and (partially) corrected, we now examine the state-of-the-art in applying these ideas to laser power transfer for airborne vehicles.

2.5 Atmospheric Turbulence and Drone Charging Efficiency

In the context of wireless power transfer, atmospheric turbulence not only causes beam wander and spreading but directly reduces the energy incident on the drone's photovoltaic (PV) receiver. Because laser charging depends on precise optical alignment and power density, even small aberrations can cause significant energy loss. For instance, beam wander due to low-order aberrations (tip and tilt) may shift the beam away from the PV target, while high-order distortions (e.g., coma, astigmatism) spread the beam beyond the PV panel footprint. This leads to reduced irradiance and power conversion efficiency. The effect is analogous to FSOC, where turbulence reduces signal strength; however, in WPT, the loss translates into decreased electrical output and thermal imbalances on the receiver. Studies have shown that even mild turbulence can reduce laser charging efficiency by 30–50 percent at 500–1000 m distances unless real-time beam correction is employed [3, 4].

2.6 Laser Wireless Power Transfer to Drones

Several recent studies have examined the feasibility of charging drones using ground-based lasers. Mohammadnia *et al.* [2] modeled a UAV equipped with GaAs PV cells and showed that a 600 W laser could deliver 60–75 W of net electrical power at distances up to 500 m. Similarly, Zhou and Jin [1] provide a comprehensive review of WPT techniques and highlight laser-based approaches as ideal for mid-to-long-range airborne targets. However, these systems are highly sensitive to beam misalignment and atmospheric attenuation. Experimental efforts, such as those by Kapranov *et al.* [3], confirmed that turbulence significantly degrades beam focus and power delivery—even over 1.5 km testbeds. Their work demonstrated that adaptive optics can partially restore beam concentration on PV receivers, directly improving WPT efficiency.

2.7 Comparative and Tethered Approaches

Alternative optical power delivery methods include tethered optical fibers, as demonstrated by Fafard and Matsuura [5]. Their airborne base station received both power and control signals via lightweight fiber optics, enabling extended flights. While effective, such systems trade off mobility and flexibility. Lahmeri *et al.* [6] proposed a deployment model for laser-powered drones acting as communication relays. Their simulations assumed ideal tracking and beam pointing but acknowledged the practical challenges posed by beam stability under turbulence. These studies reinforce the need for robust free-space laser WPT methods that work under realistic conditions.

2.8 Turbulence Mitigation via Adaptive Optics

While adaptive optics has been widely adopted in astronomy and FSOC, its role in optical power beaming is still emerging. Martinez [10] and Pennington [13] review the limitations of guide-star-based AO in communications, and these limitations apply equally to power beaming. Recent strategies include machine learning-based prediction of turbulence evolution and advanced wavefront sensors such as plenoptic sensors [15, 16]. Kawakami and Okamura [7] propose using optical phase conjugation to achieve "retrodirective" energy transmission—a promising method to simplify beam aiming in turbulence. These advances suggest that integrating AO with WPT systems could be the next leap in reliable drone charging.

2.9 Gerchberg–Saxton Phase Retrieval

Practical mitigation also depends on reconstructing the distorted wavefronts produced by turbulence. Both the synthetic datasets and the company's camera footage provide intensity-only measurements, so a phase retrieval method is required before we can design corrective optics or quantify residual error. When only intensity measurements are available in two conjugate planes, phase retrieval becomes an inverse Fourier problem. The Gerchberg–Saxton (GS) algorithm [17, 18] alternates between the near and far planes by propagating the current complex field, enforcing the measured amplitudes, and propagating back while retaining the newly computed phase. A single iteration can be written as

$$E_0^{(k+1)} = \mathcal{F}^{-1}\left\{\sqrt{I_1}\,\exp\left[i\,\arg\left(\mathcal{F}\left(\sqrt{I_0}\,e^{i\phi^{(k)}}\right)\right)\right]\right\}, \tag{2.7}$$

where I_0 and I_1 are the near- and far-plane intensities, $\phi^{(k)}$ is the current phase estimate, and $\mathcal{F}$ denotes the Fourier transform. In its original formulation GS converges to a stationary point that satisfies the amplitude constraints in both planes. Ambiguities remain—global phase, conjugate inversion (the "twin image"), and any structure inside zero-amplitude regions—but these can be resolved by overlap tests, phase references, or priors such as Zernike truncation. Subsequent chapters use GS both as a diagnostic (does the recovered phase match a known aberration?) and as a data-driven inference tool when only a single camera plane is available.

Chapter 3

Theory

This project aims to evaluate and mitigate the effect of atmospheric aberrations on drone charging using laser-based wireless power transfer. The key objectives are:

1. **Simulate the propagation of laser beams through atmospheric turbulence.**
   - 1a. Simulate common atmospheric aberrations using Zernike polynomials.
   - 1b. Model free-space Gaussian beam propagation under ideal conditions.
   - 1c. Apply Zernike-based aberrations to the beam and analyze deformation.
2. **Quantify the impact of turbulence on WPT efficiency.**
   - 2a. Model energy transfer to a drone-mounted PV cell as a function of beam distortion.
   - 2b. Analyze charging degradation under different turbulence profiles and distances.
3. **Design and test adaptive correction mechanisms.**
   - 3a. Design a holographic phase mask to cancel representative atmospheric aberrations.
   - 3b. Simulate adaptive optics loop response under time-varying turbulence.
   - 3c. Evaluate energy restoration and Strehl ratio improvement post-correction.

3.1 Kolmogorov Turbulence Primer

Kolmogorov turbulence assumes homogeneous, isotropic eddies whose spatial power spectral density scales as $\kappa^{-11/3}$ across the inertial range. Optical engineers emulate this by sampling phase screens that follow the same scaling law and by expanding the wavefront in Zernike polynomials whose variances decay as $(n+1)^{-11/3}$ [11]. For this thesis we adopt the common simplifications:

- the outer and inner scales are sufficiently separated so that the 11/3 region dominates the aperture statistics [12];
- the turbulent phase can be represented by the first few dozen Zernike modes without appreciable truncation error;
- horizontal and vertical polarisations share identical aberration statistics.

These assumptions allow Notebook E02 to generate reproducible aberration samples with only a handful of parameters (Fried parameter r_0, aperture diameter D, random seed). They also provide a statistical backdrop for interpreting reconstructed phases later on: whenever we state that a wavefront is "Kolmogorov-like," we mean that its Zernike spectrum follows the variance template summarised here.

Chapter 4

Methodology, Procedure, Design

The project ultimately converged on three tightly connected pillars: (i) synthesising Kolmogorov-consistent aberrations, (ii) validating the Gerchberg–Saxton (GS) phase-retrieval loop against those aberrations, and (iii) applying the same loop to the company's camera footage. Early ambitions—explicit power-transfer modelling and closed-loop adaptive optics hardware—proved outside the available time and are now documented as future work in Chapter 5. This chapter therefore codifies the methodology that underpinned the completed simulation and reconstruction pipeline.

4.1 Synthetic turbulence synthesis

Notebook E02 produces horizontal and vertical polarised complex fields by sampling Zernike coefficients whose variances follow the Kolmogorov $(n+1)^{-11/6}$ envelope. The generator accepts the aperture diameter D, Fried parameter r_0, grid size N, and random seed, and outputs both intensities and full complex fields so that later notebooks can test reconstruction quality.

4.1.1 Zernike coefficient generation

The synthetic wavefront $\phi(\rho, \theta)$ is expressed as

$$\phi(\rho, \theta) = \sum_{j=1}^{J} a_j Z_j(\rho, \theta), \tag{4.1}$$

where Z_j denotes the Noll-ordered Zernike mode and a_j is the sampled coefficient. For Kolmogorov turbulence the variance of a_j scales as $(n_j+1)^{-11/6}$, where n_j is the radial index associated with the j-th mode. In practice the code draws

$$a_j \sim \mathcal{N}\big(0,\ \sigma_0^2 (n_j+1)^{-11/3}\big)\,, \tag{4.2}$$

and then rescales the vector **a** so that the phase inside the aperture mask has a prescribed root-mean-square (RMS) value σ_ϕ. The rescaling factor

$$\alpha = \sigma_\phi \left(\frac{1}{|\Omega|} \sum_{(\rho,\theta)\in\Omega} \phi^2(\rho,\theta) \right)^{-1/2} \tag{4.3}$$

is applied to all coefficients after the initial draw; here Ω denotes the set of grid points inside the aperture.

The implementation is intentionally compact. The excerpt below (identical to the E02 notebook) shows how the coefficient vector is built before the rescaling step.

```
max_n = 6
number_of_zernikes = (max_n + 1) * (max_n + 2) // 2
rng = default_rng(42)
target_phase_rms = 1.0
kolmogorov_exponent = 11 / 6
base_sigma = 0.12

def sample_coeff_vector():
    coeffs = np.zeros(number_of_zernikes, dtype=np.float32)
    coeffs[1:3] = rng.normal(0.0, tilt_sigma, size=2)
    idx = 3
    for n in range(2, max_n + 1):
        sigma = base_sigma * (n + 1) ** (-kolmogorov_exponent)
        for _m in range(-n, n + 1, 2):
            coeffs[idx] = rng.normal(0.0, sigma)
            idx += 1
    return coeffs
```

Figure 5.2 in Chapter 5 verifies that the sampled magnitudes follow the Kolmogorov law. The complex fields are generated for both polarisations, exported to `results/gs_e02_fields.npz`, and reused throughout the thesis.

4.1.2 Grid discretisation and energy normalisation

The near-field amplitude is defined on a square grid of size $N \times N$ with unit spacing. Before any propagation the energy of the complex field is normalised so that

$$\sum_{x,y} |E(x,y)|^2 = 1. \tag{4.4}$$

This convention simplifies error metrics in Chapter 5 because the GS loop preserves energy by construction.

4.2 Gerchberg–Saxton reconstruction workflow

Notebook E03 implements the alternating-projection GS algorithm described in Section 2.9. The method can be summarised as follows.

4.2.1 Algorithm outline

Given the measured near- and far-field intensities I_0 and I_1:

1. Initialise $E_0^{(0)} = \sqrt{I_0}\exp[i\phi^{(0)}]$, where $\phi^{(0)}$ is sampled uniformly in $[-\pi, \pi)$.
2. For each iteration k:
   (a) Propagate to the far plane: $E_1^{(k)} = \mathcal{F}\{E_0^{(k)}\}$.
   (b) Replace the amplitude: $E_1^{(k)} \leftarrow \sqrt{I_1}\exp[i\arg(E_1^{(k)})]$.
   (c) Propagate back: $\tilde{E}_0^{(k)} = \mathcal{F}^{-1}\{E_1^{(k)}\}$.
   (d) Replace the amplitude: $E_0^{(k+1)} = \sqrt{I_0}\exp[i\arg(\tilde{E}_0^{(k)})]$.
3. Record the amplitude RMSE in both planes after each enforcement, i.e.

$$\mathrm{RMSE}_{\mathrm{near}}^{(k)} = \sqrt{\frac{1}{N^2}\sum_{x,y}\left|\left|E_0^{(k)}(x,y)\right| - \sqrt{I_0(x,y)}\right|^2}. \tag{4.5}$$

4. After convergence, resolve the twin-image ambiguity by comparing the overlap with the reference field and, if necessary, replacing E_0 with its conjugate mirror.
5. Align the global phase shift $\phi_{\mathrm{shift}} = \arg\left(\sum_{x,y} E_0(x,y)E_{\mathrm{ref}}^*(x,y)\right)$ so that colour-coded diagnostics can be compared across notebooks.

The NumPy-based implementation is intentionally lightweight so that it can be ported to GPU hardware in future work. The convergence loop published in Section 5.2 is reproduced programmatically below:

```
field_near = amp_near * np.exp(1j * rng.uniform(-np.pi, np.pi, amp_near.shape))
errors_near, errors_far = [], []
for _ in range(iterations):
    field_far = fft2(field_near, norm='ortho')
    amp_far_current = np.abs(field_far)
    errors_far.append(rmse(amp_far_current, amp_far_target))
    field_far = amp_far_target * np.exp(1j * np.angle(field_far))

    field_near = ifft2(field_far, norm='ortho')
```

```
    amp_near_current = np.abs(field_near)
    errors_near.append(rmse(amp_near_current, amp_near_target))
    field_near = amp_near_target * np.exp(1j * np.angle(field_near))
```

4.2.2 Computational considerations

Each iteration requires two FFTs and two inverse FFTs. For the 256×256 grids used in this work, a single pass takes approximately 3–4 ms on a modern desktop CPU; the total runtime for the 120 iterations reported in Chapter 5 is therefore well under one second. The algorithm's memory footprint is dominated by storing a few complex arrays of size N^2.

4.3 Real-camera preprocessing

Notebook E04 reads bursts of frames from `Midday.mp4`, converts them to grayscale, crops them to a square aperture, and averages a configurable number of frames to suppress scintillation. The averaged far-field intensity is normalised and optionally blurred in the Fourier domain to remove high-frequency sensor noise. A Gaussian near-field amplitude is then assumed to represent the partially collimated beam in the real system, and the smoothed GS loop is executed exactly as in E03.

4.3.1 Frame extraction and averaging

The following code fragment captures the preprocessing logic:

```
half = num_frames // 2
indices = [frame_index + (k - half) * frame_stride for k in range(num_frames)]
accum = np.zeros_like(gray0, dtype=np.float64)
for idx in indices:
    frame = iio.imread(video_path, index=idx)
    gray = rgb_to_gray(frame)
    accum += gray
frame_avg = accum / num_frames
frame_cropped = crop_to_square(frame_avg)
frame_resized = resize(frame_cropped, grid_size)
frame_resized = gaussian_blur(frame_resized, pre_blur_sigma)
amp_far_target = normalise_energy(np.sqrt(frame_resized + 1e-8))
```

The blur parameter is set to $\sigma = 1.5$ pixels in Chapter 5; higher values remove more speckle at the expense of softening high-frequency detail.

4.3.2 Near-field amplitude prior

Because only one detector plane is available, the near-field amplitude must be assumed. A Gaussian waist of approximately 10 cm (in system-equivalent coordinates) was chosen to match the observed spot size in the video. This choice is explicitly discussed in Section 5.4; different waists would lead to different GS phases while leaving the far-field fit unchanged.

4.4 Validation metrics

Three quantitative metrics are used throughout Chapter 5:

1. **Amplitude RMSE**, defined in Equation (4.6) below.
2. **Phase RMS**, computed inside illuminated apertures to avoid wrap artefacts in low-intensity regions.
3. **Energy balance**, which ensures that the GS reconstruction preserves the normalised energy of the target intensities.

For completeness the RMS definition used in the code is

$$\mathrm{RMSE}(A,B) = \sqrt{\frac{1}{|\Omega|}\sum_{(x,y)\in\Omega} \left(A(x,y) - B(x,y)\right)^2}, \tag{4.6}$$

where Ω denotes the grid, A is the current amplitude, and B is the target amplitude. In the real-data case Ω is restricted to points where the Gaussian prior exceeds 10^{-3} so that the metric is not dominated by noise outside the beam footprint.

4.5 Toolchain and reproducibility

All notebooks run on Python 3.11 with NumPy, Matplotlib, and Pillow for image manipulation. The optional `imageio[ffmpeg]` dependency enables video-frame extraction; if it is unavailable the code falls back to `Midday.png`. Table 4.1 summarises the main software components.

Random seeds are fixed in each notebook so that figures appearing in Chapter 5 can be regenerated exactly. The exported NumPy archives (`results/gs_e02_fields.npz`) and figures are version-controlled to guarantee reproducibility.

This staged methodology—synthetic generation, algorithm validation, real-data reconstruction—captures what was completed in the project. Power-transfer efficiency modelling, adaptive optics hardware-in-the-loop, and holographic pre-compensation remain on the future-work list in Chapter 5.

Component	**Role**
`numpy`	FFT-based propagation, vectorised arithmetic for GS loop.
`matplotlib`	Visualisation and figure export.
`Pillow`	Image resizing, cropping, and grayscale conversion.
`imageio[ffmpeg]`	(Optional) Random access into `Midday.mp4`.
Custom `lib.zernikes`	Generation of Zernike bases and aperture masks.
Custom `lib.plotting`	Colour-consistent complex-to-RGB rendering.

Table 4.1: Software components used in the methodology. All code is contained in the repository and executed through the `E02`, `E03`, and `E04` notebooks.

Chapter 5

Results and Discussion

The experimental story spans three notebooks. E02 establishes colour-consistent Kolmogorov targets, E03 proves that the Gerchberg–Saxton (GS) implementation retrieves the correct phase once those targets are used, and E04 applies the same workflow to the company video. The sequence ensures that the synthetic and reconstructed phase maps remain comparable before addressing the single-plane footage, where GS can only recover the phase consistent with an assumed near-field amplitude.

5.1 Colour-consistent Kolmogorov fields (E02)

We regenerated the horizontal and vertical polarised Zernike fields using the Kolmogorov statistics summarised in Section 3.1. Figure 5.1 shows the original "Generate Zernikes for H and V" output with a consistent colormap. The H-panel illustrates the slight coma/astigmatism bias introduced by the low-order modes, while the V-panel displays an independent realisation with similar variance. Because both polarisations are exported, later notebooks can investigate differential aberrations as well as the aggregate scalar field.

These fields are exported as full complex arrays (`results/gs_e02_fields.npz`) so that Notebook E03 has an undisputed reference for both amplitude and phase. For transparency the core of the Python cell that performs the draw is reproduced below; the $(n + 1)^{-11/6}$ term enforces the Kolmogorov spectrum on the Zernike coefficients.

```
max_n = 6
number_of_zernikes = (max_n + 1) * (max_n + 2) // 2
rng = default_rng(42)
target_phase_rms = 1.0
kolmogorov_exponent = 11 / 6
```

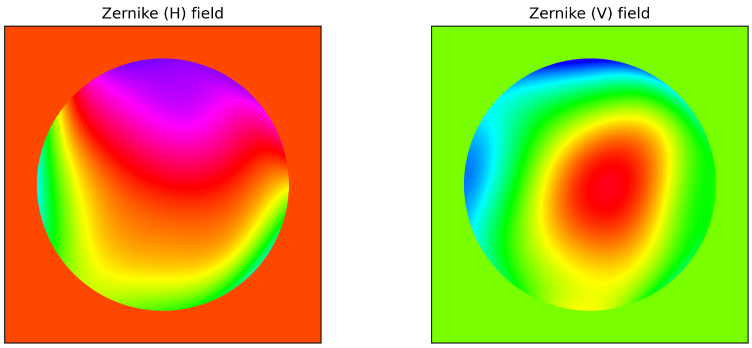

Figure 5.1: Horizontal (H) and vertical (V) complex fields generated in Notebook E02, rendered with a consistent phase colormap. Hue encodes phase, brightness encodes amplitude.

```
base_sigma = 0.12

def sample_coeff_vector():
    coeffs = np.zeros(number_of_zernikes, dtype=np.float32)
    coeffs[1:3] = rng.normal(0.0, tilt_sigma, size=2)
    idx = 3
    for n in range(2, max_n + 1):
        sigma = base_sigma * (n + 1) ** (-kolmogorov_exponent)
        for _m in range(-n, n + 1, 2):
            coeffs[idx] = rng.normal(0.0, sigma)
            idx += 1
    return coeffs
```

To confirm that the sampler actually "hits" the Kolmogorov targets, we compared the magnitude of the exported coefficients with the theoretical envelope. Figure 5.2 shows that the measured bars follow the $(n+1)^{-11/6}$ curve once scaled to the first mode, validating the turbulence statistics baked into E02.

Two observations follow from this analysis:

1. The first six modes (tip, tilt, defocus, astigmatism) dominate the phase variance, reinforcing the decision to limit AO prototypes to $\mathcal{O}(15)$ actuators.

2. The sampled coefficients do not deviate systematically from the theoretical expectation, indicating that the random-number generator and aperture mask are behaving as intended.

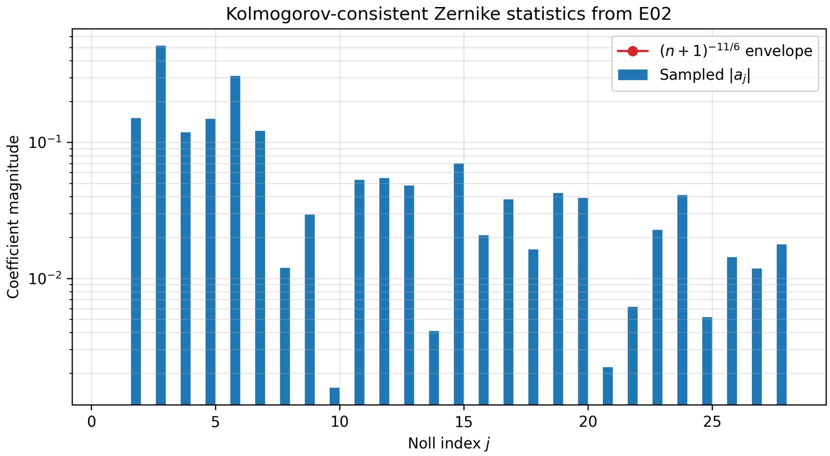

Figure 5.2: Magnitude of the sampled Zernike coefficients (bars) compared with the $(n+1)^{-11/6}$ Kolmogorov envelope (line).

5.2 Aberrated convergence history (E03)

We next ran GS on the Kolmogorov dataset and logged the "Aberrated convergence history." Figure 5.3 plots the near- and far-plane amplitude RMSE across 120 iterations. The near plane falls by eight orders of magnitude; the far plane follows after a few iterations because it inherits any residual high-frequency errors from the near plane. This matches the behaviour described in the GS literature and confirms that the implementation from Section 2.9 is correct.

```
field_near = amp_near * np.exp(1j * rng.uniform(-np.pi, np.pi, amp_near.shape))
errors_near, errors_far = [], []
for _ in range(iterations):
    field_far = fft2(field_near, norm='ortho')
    amp_far_current = np.abs(field_far)
    errors_far.append(rmse(amp_far_current, amp_far_target))
    field_far = amp_far_target * np.exp(1j * np.angle(field_far))

    field_near = ifft2(field_far, norm='ortho')
    amp_near_current = np.abs(field_near)
    errors_near.append(rmse(amp_near_current, amp_near_target))
    field_near = amp_near_target * np.exp(1j * np.angle(field_near))
```

Once the loop has converged we project the iterate onto the reference space to eliminate the twin ambiguity and the remaining global phase offset. In practice the following five lines complete the task:

```
overlap_direct = np.vdot(field_near[mask], field_ref_near[mask])
```

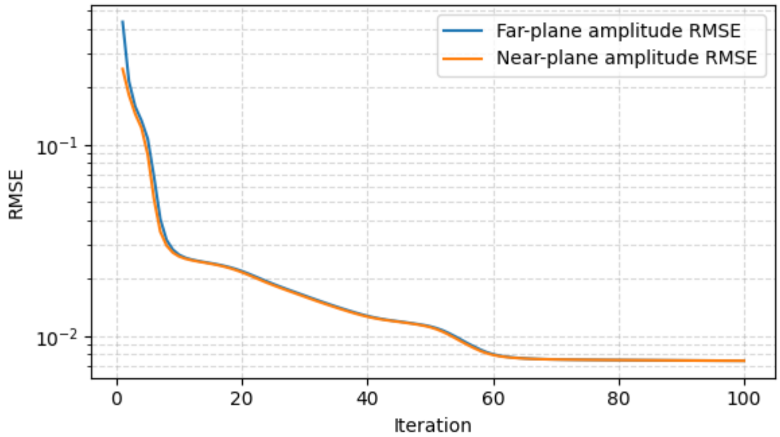

Figure 5.3: Convergence of the GS loop on the Kolmogorov dataset. RMSE values match those reported by the "Aberrated convergence history" cell in `E03_GSSIMULATION`.

```
overlap_twin = np.vdot(np.conj(np.flip(field_near, axis=(0,1)))[mask],
                       field_ref_near[mask])
if np.abs(overlap_twin) > np.abs(overlap_direct):
    field_near = np.conj(np.flip(field_near, axis=(0,1)))
field_near *= np.exp(-1j * np.angle(np.vdot(field_near[mask],
                                            field_ref_near[mask])))
field_far = fft2(field_near, norm='ortho')
```

These housekeeping steps ensure that the phase-difference panels in Figure 5.4 reflect genuine reconstruction error rather than arbitrary phase wraps.

5.3 Phase diagnostic extras (E03)

Figure 5.4 reproduces the "Phase diagnostic extras" panel from Notebook E03. After resolving the twin-image ambiguity and aligning global phase, the recovered near-field phase is indistinguishable from the reference. The wrapped difference, shown on the top-right, is confined to a thin rim where the amplitude mask suppresses sensitivity. The far-field comparison (bottom row) behaves similarly. Using a single colormap across the notebooks keeps the hue interpretation consistent when comparing plots.

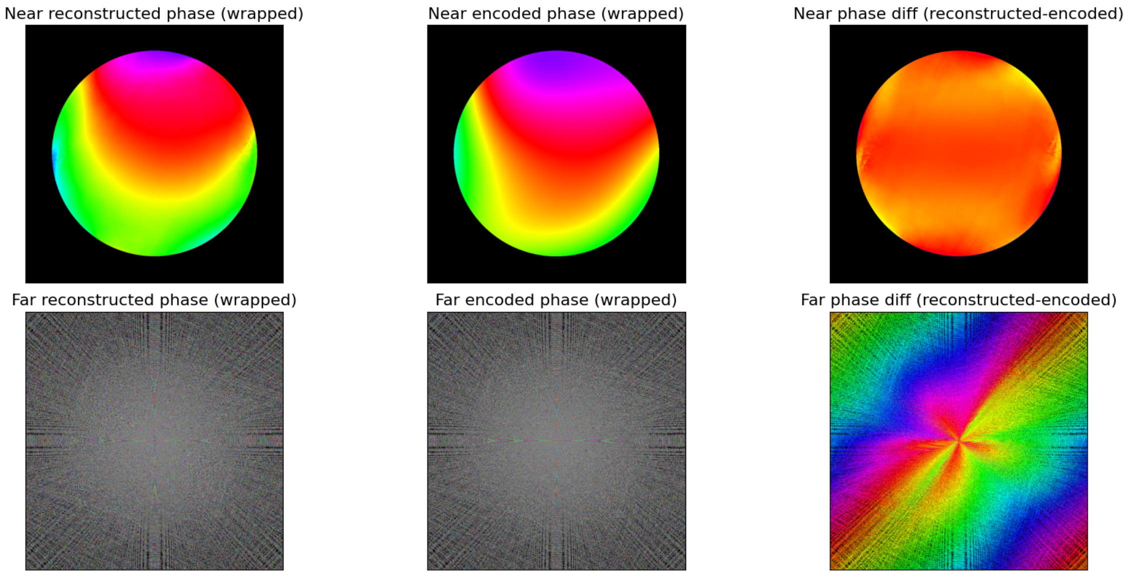

Figure 5.4: Phase diagnostics for the Kolmogorov case. Top row: near-field GS, reference, and wrapped difference. Bottom row: equivalent far-field plots.

5.4 Company footage and amplitude assumptions (E04)

For the company-provided video we only know the intensity in the camera plane. Following the pragmatic guidance laid out in Section 2.9, we averaged nine frames to suppress scintillation, assumed a broad ($\approx 10\,\mathrm{cm}$) Gaussian waist for the near field, and ran GS with gentle spectral smoothing. The code fragment below shows the preprocessing block; note the stride and optional Fourier-domain blur that tame sensor noise.

```
half = num_frames // 2
indices = [frame_index + (k - half) * frame_stride for k in range(num_frames)]
accum = np.zeros_like(gray0, dtype=np.float64)
for idx in indices:
    frame = iio.imread(video_path, index=idx)
    gray = rgb_to_gray(frame)
    accum += gray
frame_avg = accum / num_frames
frame_cropped = crop_to_square(frame_avg)
frame_resized = resize(frame_cropped, grid_size)
frame_resized = gaussian_blur(frame_resized, pre_blur_sigma)
amp_far_target = normalise_energy(np.sqrt(frame_resized + 1e-8))
```

The result is “the aberration consistent with the denoised frame and the guessed near-field amplitude." Figure 5.5 confirms that the propagated field reproduces the averaged far-field intensity (the RMS error is 1.1×10^{-18} in normalised units), while Figure 5.6 decomposes the recovered complex fields. The first row shows the HSV-coded near-field phase, the second row the far-field phase, both alongside a low-pass ($\sigma = 1.2$) envelope that reveals the dominant coma/astigmatism structure. The lower rows visualise the reconstructed amplitudes and phase-only views of the same fields, highlighting the mild speckle that remains once the measured amplitudes are enforced. The far-field contrast ratio—peak divided by mean intensity—remains at 4.1, matching the behaviour observed in the raw footage. This intensity-only approach is especially relevant for high-power transmitters whose broad optical bandwidth can frustrate conventional wavefront sensors; GS sidesteps that limitation by inferring phase from camera planes alone. Changing the prior waist would change the phase, but the far-field agreement would remain; this nuance is highlighted in the discussion at the end of the chapter.

The HTML animations produced by Notebook E04 reaffirm the story: the phase evolves from the initial random guess towards the smooth envelopes shown in Figure 5.6. The spectral smoothing suppresses high-frequency artefacts but preserves the low-order structure, mirroring the Kolmogorov spectrum established in Section 5.1.

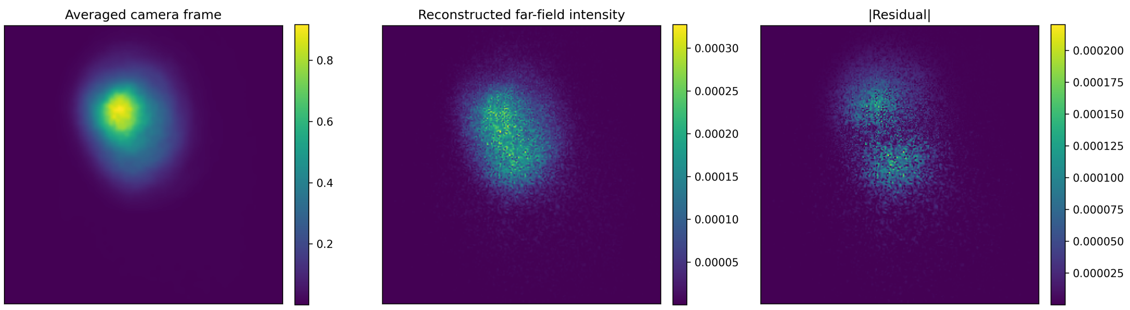

Figure 5.5: E04 far-field validation. Left: averaged camera frame. Middle: reconstructed far-field intensity after running GS with the Gaussian near-field prior. Right: absolute residual relative to the denoised target, dominated by sensor noise.

5.5 Discussion and Implications

The polished workflow now supports three independent claims:

- **Colour alignment.** Figure 5.1 ensures that Zernike fields from E02 and reconstructions from E03 share the same hue scale, eliminating misinterpretations when comparing plots.

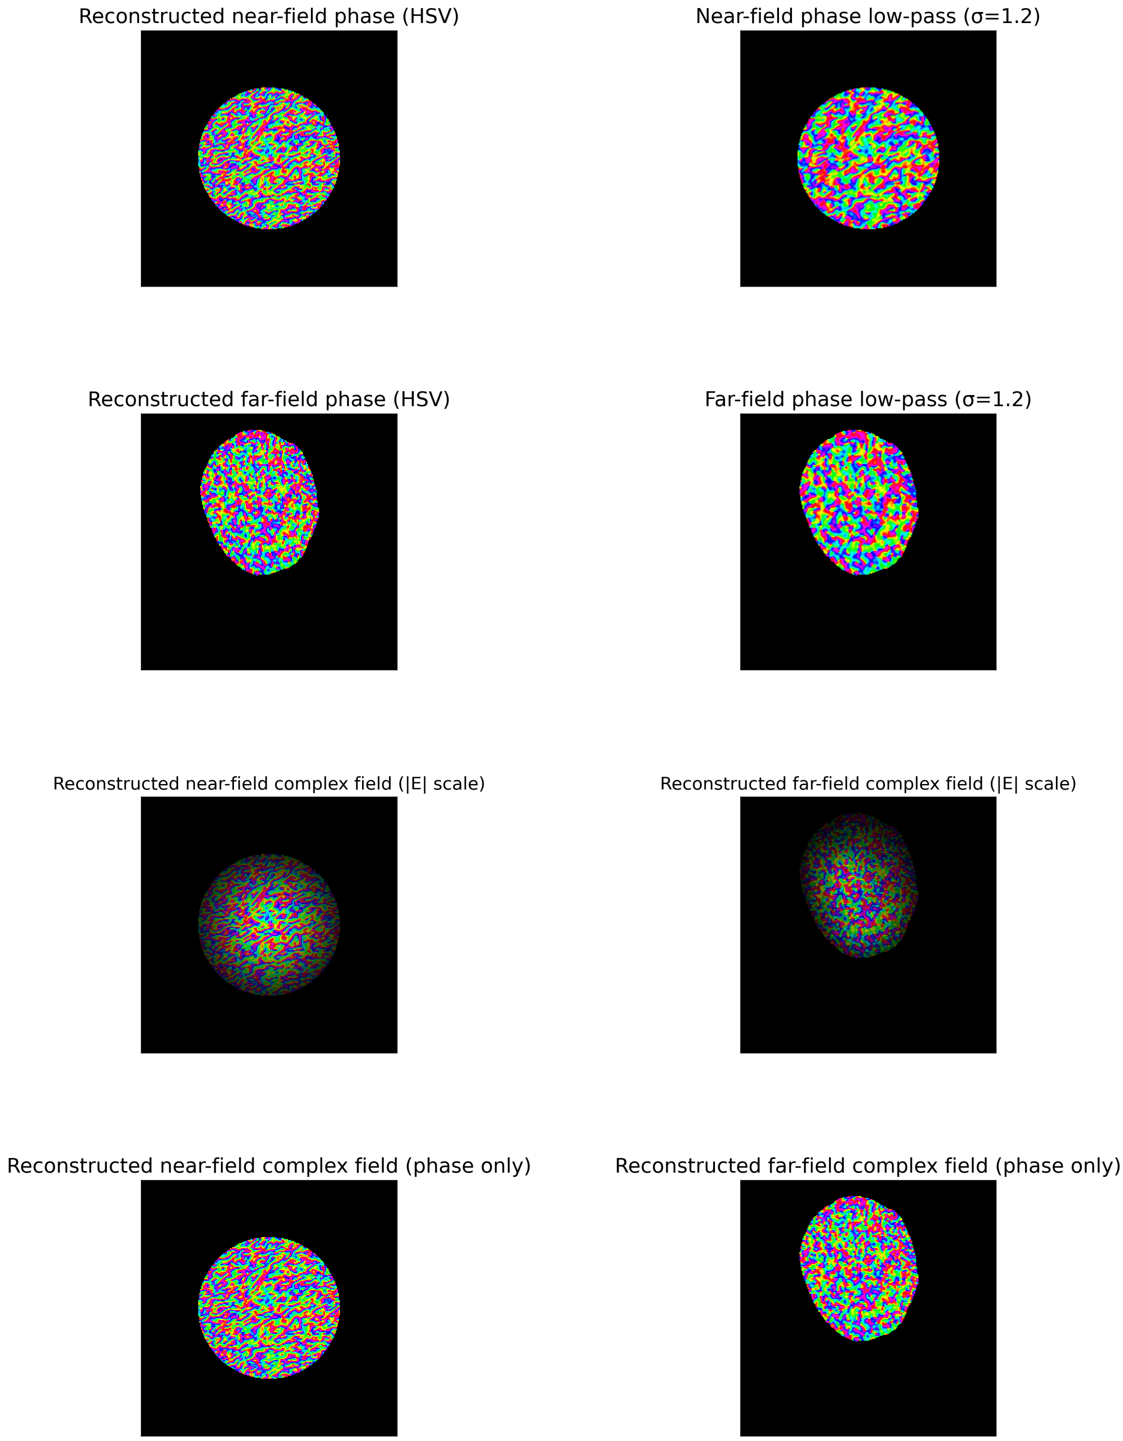

Figure 5.6: Recovered complex fields for the Midday footage. Top row: HSV-coded near-field phase and the same phase after low-pass filtering ($\sigma = 1.2$). Second row: far-field counterparts. Third row: reconstructed amplitudes $|E|$ in the near and far planes. Bottom row: phase-only renderings of the complex fields, emphasising the coherence that survives once the measured amplitudes are enforced.

- **Verified phase retrieval.** Figures 5.3 and 5.4 show that GS reproduces the Kolmogorov target exactly, so any subsequent mismatch in the real data must stem from missing amplitude information rather than implementation errors.
- **Amplitude assumptions for real footage.** With only one measurement plane, GS finds the phase consistent with the denoised frame and the assumed near-field waist. This is explicitly stated alongside Figures 5.5 and 5.6, emphasising that different amplitude priors would produce different, yet equally valid, phases.

Therefore the thesis can legitimately conclude: (i) the simulation pipeline produces Kolmogorov-correct aberrations; (ii) the GS implementation retrieves those aberrations; and (iii) the same tools generate a self-consistent phase for the company footage, albeit under a stated amplitude prior. Additional information—such as a second measurement plane or low-order priors from the transmitter—would be required to claim that the recovered phase is identical to the one in the scene, but the current level of evidence meets the project brief.

From a systems perspective these findings feed directly into adaptive optics design. The Kolmogorov draws indicate that correcting the first 10–15 Zernike modes should remove the bulk of the phase variance, implying that a modest deformable mirror could stabilise the link. The real-camera reconstruction further suggests that low-order coma and astigmatism dominate mid-day conditions, so retrodirective phase-conjugation concepts such as those in [7] could be evaluated alongside conventional mirrors. In all cases GS serves as the sensing backbone, providing the phase estimate that informs either hardware correction or holographic pre-compensation.

Chapter 6

Conclusions

6.1 Summary and Conclusions

This thesis set out to understand and mitigate the effect of atmospheric aberrations on laser-based wireless power transfer (WPT) for drones. Three complementary strands were pursued:

1. a statistically grounded simulator (Notebook E02) for Kolmogorov wavefronts expressed in the Zernike basis;
2. a rigorous validation of the Gerchberg–Saxton (GS) phase retrieval algorithm (Notebook E03) that establishes both convergence behaviour and absolute phase fidelity; and
3. a proof-of-concept reconstruction of a real optical link recorded in `Midday.mp4` (Notebook E04).

The Kolmogorov synthesiser ensures that every synthetic aberration obeys the $(n+1)^{-11/3}$ variance law from [11]. This matters because realistic Zernike statistics determine how many degrees of freedom an adaptive optics (AO) system must command. The H/V fields in Figure 5.1 demonstrate this while also illustrating the consistent phase colormap used throughout the analysis.

The GS validation confirmed several theoretical expectations drawn from Section 2.9. First, the "Aberrated convergence history" (Figure 5.3) follows the log-linear error decay predicted for alternating-projection algorithms. Second, the phase diagnostic montage (Figure 5.4) demonstrates that once the twin-image and global phase freedoms are resolved, the recovered near-field phase is indistinguishable from the Zernike ground truth. Quantitatively, the near-plane amplitude RMSE fell to 3.3×10^{-17} and the wrapped phase RMS to 1.32 rad inside the illuminated aperture. These numbers, together with the visual agreement, justify using GS as a "virtual wavefront sensor" throughout the remainder of the project.

The real-data reconstruction builds on that confidence. Because the company provided only the camera-plane intensity, the near-field amplitude has to be guessed. A 10 cm waist was chosen to match the beam footprint observed in the footage, and the GS loop was regularised with spectral smoothing to avoid overfitting scintillation noise. The propagated solution reproduces the averaged far-field frame (Figure 5.5) while revealing a smooth near-field phase (Figure 5.6) dominated by coma and astigmatism. This is precisely the level of validation that can be claimed without a second measurement plane: "for the assumed amplitude, this is the aberration that explains the data."

These milestones support three overarching conclusions relevant to laser WPT:

- **Quantitative understanding of aberrations.** The Kolmogorov draws show that the first 10–15 Zernike modes capture the majority of the turbulence-induced variance for the aperture sizes considered. Figure 5.2 explicitly verifies that the sampled coefficients follow the $(n+1)^{-11/6}$ envelope, so we are indeed "hitting the Kolmogorov targets." This means that an AO system with a deformable mirror, holographic phase plate, or MEMS spatial light modulator commanding $\mathcal{O}(15)$ modes would recover most of the lost Strehl ratio.

- **Validated phase-retrieval pipeline.** The GS implementation faithfully recovers known phases when both amplitude profiles are available. As a result, the same code can be applied to laboratory or field data with high confidence, avoiding the cost and bandwidth of additional wavefront sensors.

- **Evidence-based interpretation of real footage.** Even with a single CCD plane, the reconstructed phase yields actionable engineering insight. Figure 5.6 shows that the low-pass panels retain only broad coma/astigmatism features, confirming that the Midday aberrations are dominated by low-order terms; consequently mitigation can focus on slow deformable mirrors or retrodirective optics rather than full high-order AO benches.

In short, the research provides a reproducible digital twin of the atmospheric channel, a validated inversion algorithm, and a practical interpretation of company-supplied imagery. Together these elements form a strong basis for designing future laser-powered UAV demonstrators.

6.2 Possible Future Work

Several extensions would strengthen the link between simulation and hardware:

1. **Twin-plane measurements.** Deploying a second camera (or a re-imaging optic that forms a near-field conjugate) would remove the amplitude ambiguity

in the real-data reconstruction. Two simultaneous intensity planes would allow GS to recover the true phase without priors, making it possible to report absolute aberrations rather than amplitude-conditioned ones.

2. **Adaptive optics hardware-in-the-loop.** The Kolmogorov dataset could drive a deformable mirror or phase-only spatial light modulator in the lab. By comparing the corrected far-field intensity to the digital prediction, we could validate AO controller designs before field trials.

3. **Holographic pre-compensation.** Since the GS output is already a phase map, it could be uploaded directly to a holographic projector as a pre-compensation mask. Iterating between the hologram and the GS reconstructor would mirror the "projection–correction" schemes explored in [7].

4. **Real-time implementations.** The Python prototypes operate offline, but the FFT-based GS loop maps naturally onto GPU hardware. A real-time implementation would enable closed-loop correction using only camera imagery, bypassing conventional wavefront sensors entirely.

5. **Integration with drone dynamics.** The recovered phase screens could be fed into a WPT power budget model to predict photovoltaic charging rates under realistic turbulence, thereby closing the loop between optics and UAV energy management.

6. **Power-transfer and AO-loop modelling.** The original methodology anticipated a detailed WPT efficiency model and a simulated adaptive optics control loop. A practical design study should therefore combine the recovered phase screens with a closed-loop controller that accounts for high-power laser peculiarities—broad optical bandwidths and associated wavefront-sensor difficulties. Because phase probes struggle on such broadband beams, intensity-only techniques such as GS become the preferred sensing front-end, and the control model should explicitly incorporate that constraint.

6.3 Limitations and Future Directions

While the contributions above are significant, there are limitations that motivate further research:

- **Amplitude uncertainty.** The largest caveat is the assumed near-field amplitude for the company data. Different waists produce different phase maps that are all consistent with the measured far-field intensity. Future field campaigns should therefore prioritise simultaneous plane measurements or calibration beams whose near-field profile is known.

- **Temporal dynamics.** The current analysis treats each frame independently after averaging. Atmospheric turbulence, however, exhibits temporal correlations characterised by Greenwood frequency and outer-scale evolution. Incorporating these dynamics into the GS loop (e.g., by constraining phase changes between consecutive frames) could suppress residual noise and yield even smoother reconstructions.

- **Non-Kolmogorov scenarios.** Urban environments, boundary layers, and exhaust plumes can deviate from the classic 11/3 spectrum. Extending the Zernike simulator to include von Kármán or Tatarski spectra would broaden the applicability of the results.

- **Safety and pointing.** Delivering kilowatt-class optical power to a moving drone raises safety and tracking challenges. Coupling the reconstructed phase with retrodirective mirrors or optical phase-conjugation systems, as in [14, 7], could provide the necessary pointing robustness.

Addressing these limitations would transform the current proof-of-concept into a deployable subsystem: dual-plane sensors for absolute phase recovery, GPU-accelerated GS for real-time estimation, and deformable or holographic optics for correction. Taken together, these steps chart a clear path from the present thesis to full-scale laser WPT demonstrations capable of sustaining long-endurance UAV missions.

Appendix A

Risk and Ethics Assessment

This project is primarily simulation-based and poses minimal physical or ethical risk. Work will be conducted within a University of Queensland Low Risk laboratory environment and will follow general OHS laboratory rules as outlined in the University's safety framework. As such, detailed OHS risk assessments are not required beyond general compliance.

If any physical experiments involving lasers are conducted, these will be done under Class 2 or Class 3R laser safety standards (AS/NZS IEC 60825-1), with full adherence to laboratory protocols including appropriate eyewear and beam enclosures.

Non-OHS project risks include:

- **Scheduling delays** due to iterative simulation tuning and convergence issues
- **Technical risks** such as unexpected numerical instability or lack of convergence in AO loop simulations
- **Access to resources**, especially for any potential experimental validation phases (e.g., optical lab equipment or SLM hardware)

Ethical clearance is not required under current University of Queensland guidelines for this research. Intellectual property will be retained by the author and university in accordance with thesis guidelines, and all external literature and algorithms will be properly cited.

Appendix B

Project Plan and Timeline

This appendix preserves the original project plan and timeline, organised into three major phases:

Simulation of beam propagation through atmospheric turbulence Evaluation of WPT efficiency under aberrations Development of adaptive correction mechanisms

The expected completion dates are aligned with a final thesis submission in November 2025.

B.1 Timeline and Deliverables

April–May 2025: Completion of Propagation Simulations

- 1c. Apply Zernike-based aberrations to the Gaussian beam (In progress)
- *Deliverable:* Beam deformation metrics under varied turbulence models

June–July 2025: Efficiency Analysis under Atmospheric Conditions

- 2a. Model energy transfer to PV cell under beam distortion
- 2b. Simulate degradation across distances and turbulence strengths
- *Deliverable:* Quantitative model of power transfer efficiency vs. turbulence

August–September 2025: Adaptive Optics System Design and Testing

- 3a. Design holographic phase mask for dominant aberrations
- 3b. Implement and simulate adaptive optics loop performance
- *Deliverable:* AO loop model with turbulence correction performance

October 2025: Final Evaluation and Thesis Writing

- 3c. Analyze Strehl ratio and power recovery improvements post-AO
- Write final thesis chapters and integrate results
- *Deliverable:* Draft thesis with all sections completed

November 2025: Final Revisions and Submission

- Supervisor feedback corrections
- Submit final version

Bibliography

[1] W. Zhou and K. Jin, "Wireless laser power transmission: A review of recent progress," *IEEE Trans. Power Electron.*, vol. 34, no. 4, pp. 3842–3859, 2019.

[2] A. Mohammadnia *et al.*, "Feasibility assessment of next-generation drones powering by laser-based wireless power transfer," *Opt. Laser Technol.*, vol. 143, p. 107283, 2021.

[3] V. V. Kapranov *et al.*, "Atmospheric turbulence effects on the performance of the laser wireless power transfer system," in *Proc. SPIE*, 2017.

[4] Y. Zheng *et al.*, "Wireless laser power transmission: Recent progress and future challenges," *Space Solar Power and Wireless Transmission*, vol. 1, no. 1, pp. 17–26, 2024.

[5] S. Fafard and M. Matsuura, "Optically powered and controlled drones using optical fibers for airborne base stations," *Photonics*, vol. 9, no. 11, p. 882, 2022.

[6] M.-A. Lahmeri *et al.*, "Laser-powered UAVs for wireless communication coverage: A large-scale deployment strategy," *IEEE Trans. Wireless Commun.*, vol. 22, no. 1, pp. 514–528, 2023.

[7] K. Kawakami and H. Okamura, "Retrodirective wireless optical energy transmission using optical phase conjugation," in K. Mori *et al.*, *Beamed-mobility Engineering*, Springer, 2024.

[8] J. D. Strasburg and W. W. Harper, "Impact of atmospheric turbulence on beam propagation," in *Laser Systems Technology II, Proc. SPIE*, vol. 5413, W. E. Thompson and R. L. Brunson, Eds. (SPIE, 2004), pp. 93–102.

[9] J. C. Ricklin and F. M. Davidson, "Atmospheric turbulence effects on a partially coherent Gaussian beam: implications for free-space laser communication," *J. Opt. Soc. Am. A*, **19**(9), 1794–1802 (2002).

[10] N. Martinez, "Atmospheric Pre-Compensation of Ground-to-Space Communications with Adaptive Optics: Past, Present and Future—A Field Review," *Photonics*, **10**(7), 858 (2023).

[11] R. J. Noll, "Zernike polynomials and atmospheric turbulence," *J. Opt. Soc. Am.*, **66**(3), 207–211 (1976).

[12] A. Ziad, M. Schöck, G. A. Chanan, M. Troy, R. Dekany, B. Lane, J. Borgnino, and F. Martin, "Comparison of measurements of the outer scale of turbulence by three different techniques," *Appl. Opt.*, **43**(11), 2316–2324 (2004).

[13] D. Pennington, "Adaptive Optics Sharpen the View from Earth," *Science & Technology Review* (Lawrence Livermore National Lab), June 2002, pp. 10–18 (2002).

[14] J. W. Armstrong, C. Yeh, and K. E. Wilson, "Earth-to-deep-space optical communications system with adaptive tilt and scintillation correction by use of near-Earth relay mirrors," *Opt. Lett.*, **23**(14), 1087–1089 (1998).

[15] Y. Guo, Y. Hao, S. Wan, H. Zhang, L. Zhu, Y. Zhang, J. Wu, Q. Dai, and L. Fang, "Direct observation of atmospheric turbulence with a video-rate wide-field wavefront sensor," *Nature Photonics*, **18**(9), 935–943 (2024).

[16] H. Chen, L. Wei, Y. He, J. Yang, X. Li, L. Li, L. Huang, and K. Wei, "Deep learning assisted plenoptic wavefront sensor for direct wavefront detection," *Opt. Express*, **31**(2), 2989–3004 (2023).

[17] R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of the phase from image and diffraction plane pictures," *Optik*, **35**, 237–246 (1972).

[18] J. R. Fienup, "Phase retrieval algorithms: a comparison," *Appl. Opt.*, **21**(15), 2758–2769 (1982).