Title
Predictive Dynamic Digital Holography

Author
Sulaiman, Sennan David

Publication Date
2017

UNIVERSITY OF CALIFORNIA
Los Angeles

Predictive Dynamic Digital Holography

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Aerospace Engineering

by

Sennan David Sulaiman

2017

ABSTRACT OF THE DISSERTATION

Predictive Dynamic Digital Holography

by

Sennan David Sulaiman
Doctor of Philosophy in Aerospace Engineering
University of California, Los Angeles, 2017
Professor James S. Gibson, Chair

Digital holography has received recent attention for many imaging and sensing applications, including imaging through turbulent and turbid media, adaptive optics, three-dimensional projective display technology and optical tweezing. It holds several advantages over classical methods for wavefront sensing and adaptive-optics correction, chief among these being significantly fewer and simpler optical components. A significant obstacle for digital holography in real-time applications, such as wavefront sensing for high-energy laser systems and high-speed imaging for target tracking, is the fact that digital holography is computationally intensive; it requires iterative virtual wavefront propagation and hill-climbing algorithms to optimize sharpness criteria. This research demonstrates real-time methods for digital holography based on approaches for optimal and adaptive identification, prediction, and control of optical wavefronts. The methods presented integrate minimum-variance wavefront prediction into dynamic digital holography schemes to accelerate the wavefront correction and image sharpening algorithms. Further gains in computational efficiency are demonstrated in this work with a variant of localized sharpening in conjunction with predictive dynamic digital holography for real-time applications. This “subspace correction” method optimizes sharpness of local regions in a detector plane by parallel independent wavefront correction on reduced-dimension subspaces of the complex field in a spectral plane.

The dissertation of Sennan David Sulaiman is approved.

Pei-Yu Chiou

Lieven Vandenberghe

Tsu-Chin Tsao

James S. Gibson, Committee Chair

University of California, Los Angeles

2017

To my parents, David and Fatima, whose love and support have carried me through all my years, both the highs and lows. I would not be where I am today without them.

To my advisor, Professor Steve Gibson, who bestowed me with tremendous wisdom and knowledge. I will always cherish the academic opportunities you've given me.

To my friend and mentor, Chuck Rathfelder, whose intellectual discussions and advice helped me realize my own potential. I always enjoy our conversations.

And finally, to my friends and colleagues I have known throughout my UCLA career,
especially Molly Lower, Anthony Jung, Monika Ujkic, Khai Tran, Nikesh Patel,
and all my sailing friends. You have given me unforgettable memories that I will treasure.

Not all those who wander are lost.
- J.R.R. Tolkien

TABLE OF CONTENTS

1 Introduction **1**

1.1 Background Review of Digital Holography 2

1.2 Introduction of Wavefront Correction and Prediction 3

1.3 Introduction of Local Sharpening 5

2 Simulation Setup and Relevant Digital Holography Theory **7**

2.1 Simulation Setup 7

2.2 Digital Holography Recording 11

2.3 Numerical Reconstruction in Digital Holography 15

3 The Aero-Optical Wavefronts **17**

4 Sharpness Algorithm and Wavefront Modeling **22**

4.1 Sharpness Metrics 22

4.2 Modal Representation of Aero-Optical Wavefronts 25

4.3 Computational Considerations 28

5 Predictive Dynamic Digital Holography **32**

5.1 Subspace Identification 32

5.2 Predictive Dynamic Digital Holography 34

    5.2.1 Initialization of the Sharpening Algorithm 34

    5.2.2 The Prediction Filter 36

    5.2.3 Identification of the Prediction Filter 36

    5.2.4 Other Considerations 38

    5.2.5 The Overall Predictive Digital Holography Scheme 39

6 Wavefront Prediction Filter Analysis **42**

6.1 The Ideal Prediction Filter: P_∞ 43

6.2 Finite k_{ID} Prediction Filters: P_4 and P_8 46

7 Wavefront Prediction Simulation and Results **52**

7.1 Results 52

7.1.1 Sharpening Initializations M_0 and M_1 vs. P_∞ 53

7.1.2 Sharpening Initializations P_4 and P_8 vs. P_∞, M_0 and M_1 59

7.1.3 Local Sharpening vs. Global Sharpening 62

7.1.4 Strehl Ratios of P_8 vs. P_∞ and M_0 64

7.2 Predictive Dynamic Digital Holography in the Far-Field 67

8 Local Sharpening and Subspace Correction **68**

8.1 Single Patch Sharpening 68

8.2 Block Sharpening 70

8.3 Results 73

9 Conclusion **81**

References **84**

LIST OF FIGURES

1.1 Predictive dynamic digital holography overview 4

1.2 Region-of-interest (local) versus global sharpening 6

2.1 Off-axis detector plane recording geometry 9

2.2 The digital-holography simulation model that enables wavefront estimation and correction using image sharpening. 10

2.3 Sample hologram recorded on detector plane 12

2.4 Spatially separated holographic terms 14

3.1 Sample aero-optical wavefront 20

3.2 Temporal correlation index for aero-optical wavefronts 20

3.3 Temporal correlation index for simulation generated wavefronts 21

3.4 Strehl ratios for aero-optical wavefront sequence 21

4.1 Sample image situation for sharpness metrics with a high β 22

4.2 Diffracted-limited USAF bar chart reconstructed on image plane 23

4.3 Cumulative power plot of first 50 Zernike modes 27

4.4 Representation of first 15 Zernike modes 28

4.5 Phase error power of benchmark metric 29

6.1 Power spectral densities of selected modes for identified P_∞ prediction filter . . 44

6.2 Relative system norm errors of $P_{k_{\text{ID}}}$ relative to P_∞ 45

6.3 Bode plots of selected modes for certain k_{ID} identified prediction filters 47

6.4 Power spectral densities of selected modes for identified P_4 and P_8 prediction filters 49

6.5 Fractional power of modal sequences for the P_∞ prediction filter 50

6.6 Fractional power of modal sequences for the P_4 and P_8 finite prediction filters . 50

6.7 Normalized temporal power of modal prediction error sequence for P_4, P_8, and P_∞ prediction filters . . . 51

7.1 Pupil plane phase error power for M_0, M_1, and P_∞ filters after k sharpening iterations for S_1 sharpness metric . . . 54

7.2 Histogram of sharpness iterations required for wavefront convergence for M_0, M_1, and P_∞ filters . . . 55

7.3 Array of region of interest reconstructed example images after k sharpening iterations using M_0 and P_∞ filters . . . 57

7.4 Pupil plane phase error power for M_0, M_1, and P_∞ filters after k sharpening iterations for S_2 sharpness metric . . . 58

7.5 Pupil plane phase error power for M_0, P_4, P_8, and P_∞ filters after k sharpening iterations for S_1 sharpness metric . . . 60

7.6 Histogram of sharpness iterations required for wavefront convergence for P_4, P_8, and P_∞ filters . . . 61

7.7 Reconstructed region-of-interest irradiance error power for M_0, M_1, P_4, P_8, and P_∞ filters after k sharpening iterations for S_1 sharpness metric . . . 63

7.8 Strehl ratio for corrected wavefront sequence using M_0 filter after $k_f = 10$ sharpness iterations . . . 65

7.9 Strehl ratio for corrected wavefront sequence using P_8 filter after $k_f = 10$ sharpness iterations . . . 66

7.10 Strehl ratio for corrected wavefront sequence using P_∞ filter after $k_f = 10$ sharpness iterations . . . 66

7.11 Fraunhofer regime wavefront correction. Time-series average pupil plane phase error power for M_0, P_8, and P_∞ methods versus number of sharpness iterations . 67

8.1 Scheme for initial subspace wavefront generation pertaining to one ROI of the image . . . 69

8.2 Region of interest area A_1 (left) and A_2 and subspace block partitions 74

8.3 Pupil plane phase error power for M_0 and P_8 filters using global and subspace wavefront correction after k sharpening iterations for S_1 sharpness metric 75

8.4 Histogram of the number of sharpness iterations required for wavefront correction convergence for the M_0 and P_8 methods using global and subspace techniques . 77

8.5 Time-series average irradiance RMS error for M_0 and P_8 methods versus number of sharpness iterations using global and local sharpening and subspace correction over a Region of Interest A_1 78

8.6 Time-series average irradiance RMS error for M_0 and P_8 methods versus number of sharpness iterations using global and local sharpening and subspace correction over Region of Interest A_2 79

LIST OF TABLES

2.1 Parameters for digital holography simulation 8

3.1 Notre Dame AAOL experimental data details 17

4.1 Normalized cumulative power for $N = 15$ 27

6.1 Parameters for identified models 43

ACKNOWLEDGMENTS

I would like to express my gratitude to Dr. Mark Spencer for his guidance and expertise in the learning stages of the topic of digital holography and his collaboration during this research. The results presented in this work are based on aero-optical data obtained by Professor Eric Jumper and the University of Notre Dame aero-optical research group, for which I am grateful.

VITA

2012	B.S. (Aerospace Engineering), UCLA.
2014	M.S. (Aerospace Engineering), UCLA.
2013–2015	Teaching Assistant, Mechanical and Aerospace Department, UCLA.
2014–2017	Graduate Research Assistant, Mechanical and Aerospace Department, UCLA.
2013–present	Member of Technical Staff, The Aerospace Corporation

PUBLICATIONS

S. Sulaiman and J.S. Gibson, “Predictive dynamic digital holography,” [Submitted for review to J. Opt. Soc. Am. A] (2017)

S. Sulaiman and J.S. Gibson, “Predictive digital holography with local image sharpening and subspace wavefront correction,” [Submitted for review to J. Opt. Soc. Am. A] (2017)

S. Sulaiman and J.S. Gibson, “Local sharpening and subspace wavefront correction with predictive dynamic digital holography, Proc. SPIE Optics and Photonics, San Diego, California, August 2017

S. Sulaiman and J.S. Gibson, “Predictive digital holography for wavefront correction,” Proc. SPIE Optics and Photonics, San Diego, California, August 2016.

CHAPTER 1

Introduction

Holography is a technology that has been recognized since 1948 [1], providing a means to capture phase information of a target object that was otherwise unavailable through a two-step process of recording and reconstructing wavefronts. However, its application has stagnated due to various limiting factors over the decades. Initially, these factors were composed entirely around physical constraints, such as a source of coherent illumination. Eventually, once the digitization of holography matured, it involved the low spatial resolution of electronic sensors. Most recently, with the advent of inexpensive, high-resolution CCD and electronic sensors, the burden of limitation has shifted to the vast swath of image processing methods and computational costs of digital holography. This dissertation proposes contributions to alleviate the computational burden experienced by digital holography, with the aim of making real-time application a more practical endeavor. Two areas will be discussed in detail, both centered on the reconstruction portion of holography. First, a system identification and prediction filter approach is introduced during wavefront reconstruction to expedite the process. Second, a local correction scheme that reduces the effective pixel size needed to numerically propagate the wavefronts and holograms is discussed.

In this chapter, a literature review of digital holography is given, followed by an introduction to wavefront prediction for optical wavefronts. In Chapter 2, the simulation model and theory for digital holography recording and reconstruction is presented. Chapter 3 discusses the aero-optical wavefront data [2, 3] used in this dissertation, while Chapter 4 explains the image sharpness algorithm and the model representation of the atmospheric wavefronts. Chapter 5 details the subspace identification algorithm and prediction filters used in predictive wavefront correction. Analysis of these filters is provided in Chapter 6. A simulation developed to integrate these various parts is culminated in results shown in Chapter 7. A study of local sharpening, subspace correction, and their results is given in Chapter 8. The dissertation concludes with a summary of this research in Chapter 9.

1.1 Background Review of Digital Holography

Optical holography was proposed in 1948 [1] by Denis Gabor. He demonstrated a two-step process of recording information of the full complex field by interfering a coherent reference wave and diffracted waves from an object onto a photographic plate. The image could be realized by illuminating this holographic recording with the same reference wave in the reconstruction step. In Gabor's original work, the reference beam is superimposed in general alignment with the diffracted waves onto the recording plate, coined "on-axis" holography. This method injected the holograms with the presence of the DC and twin image terms produced by the interference, leading to a corrupted reconstruction. Since a suitable source of coherent light was not available at the time, interest in holography had to wait until the invention of the laser in 1960.

In 1962, Leith and Upatnieks modified the Gabor hologram setup by off-setting the reference wave during interference in "off-axis" holography [4]. This breakthrough avoided the corrupting zero-order and twin image terms by spatially separating the holographic terms. After this discovery, holography experienced proliferation of activity, primarily due to the capability of measuring physical metrics by changes in the phase field. This includes areas of study in vibrational modes [5] and refractive index measurement [6].

Goodman and Lawrence were able to digitally record holograms and electronically reconstruct images in 1967 [7]. This published work was the first time digital holography was achieved, eliminating the need for photographic material during recording and the coherent reference source for reconstructing images. Instead, electronic sensors and computers replaced the recording and reconstruction procedures. However, digital holography's growth was stunted by the inconvenient and expensive digital procedures at the time. Meanwhile, in 1974, Muller and Buffington developed an image processing technique known as image sharpening [8]. While not directly related to digital holography, this important image processing development orchestrated advances in phase-retrieval [9,10] and wavefront correction algorithms [11,12].

In 1994, Schnars and Jüptner were able to record and numerically reconstruct a Fresnel hologram using a CCD array for the first time [13]. Naturally, this advancement expanded holography into new areas of study, including data storage [14], numerical auto-focusing [15,16], and image processing [17,18]. One particular field of study that greatly benefited from the advent of digital holography is microscopy, allowing researchers to study cells from various perspectives without intrusion [19]. In particular, the capability of a reconstructed object image to be numerically focused at any distance allowed several target objects to be focused and analyzed individually within the same captured image [15,16,19].

1.2 Introduction of Wavefront Correction and Prediction

Wavefront correction is not exclusive to digital holography. This is a central area of effort for the field of classical adaptive optics. Digital holography holds several advantages over classical wavefront sensing usually used in adaptive optics applications. Chief among these is requiring fewer optical components in operation. For example, a Shack-Hartmann sensor can be replaced by numerical wavefront estimation. Furthermore, a lens is not required to focus the image on a focal plane array; instead, the scattered light from a digital holography setup can be recorded directly on the detector medium. While the advent of holography through digital means allows for a broad realm of applied image processing methods, the computational cost involved in wavefront estimation and correction prevents it from seeing more widespread use in real-time applications, such as target tracking. Earlier work has shown that real-time digital holography can achieve 24 frames per second, by conjoining several high-speed parallel processors [20]. While this work demonstrates the potential of exploiting parallel processing for digital holography, it does not directly address the crux of the issue regarding the number of extremely expensive operations needed during reconstruction. Additionally, the development of that work did not consider wavefront correction during

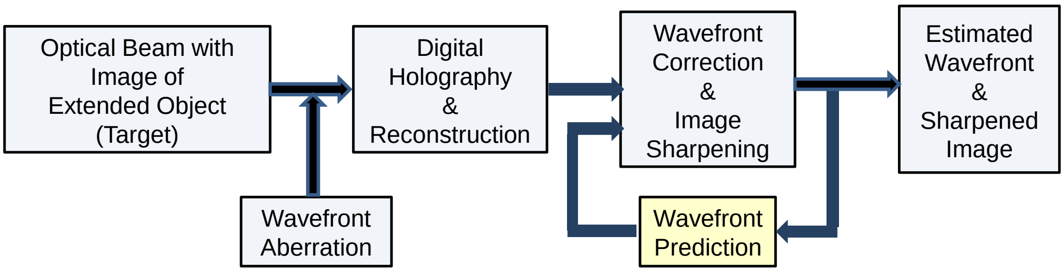

Figure 1.1: Predictive dynamic digital holography overview. Wavefront prediction creates a closed loop.

reconstruction, a fact that adds enormous computational complexity to the reconstruction process. Undoubtedly, this would lower the reported frame-per-second capability of digital holography. This dissertation builds on earlier work involving wavefront prediction for adaptive optics and applies it to digital holography in order to short-circuit the computationally intensive algorithms for sharpness optimization and wavefront correction.

A linear time-invariant (LTI) state-space prediction filter, identified by subspace system identification, for modeling, prediction, and correction of optical wavefronts has been used previously [21–24]. However, in those applications, sample wavefront sequences used to identify the prediction filters were measured directly by a Shack-Hartmann sensor. For digital holography, such wavefront sensors are unnecessary. Only a sequence of CCD intensity patterns are measured and a prediction filter is identified solely with that knowledge. Figure 1.1 highlights this important aspect of the scheme by the wavefront prediction feedback loop. This feedback interaction between wavefront prediction and image sharpening has no corresponding feature in the previous applications of wavefront prediction. The previously proposed methods for wavefront prediction involved a wavefront sensor that measures the wavefront sequence, and a prediction filter (either linear time-invariant or adaptive) uses the measured wavefronts to generate predicted wavefronts; the prediction does not affect the wavefront measurement. In this dissertation, only images recorded on the detector plane are measured directly. The wavefronts are estimated by digital holography and image sharpening, and the prediction accelerates the sharpening and improves the estimates of the wavefronts.

Even though this work focuses on an LTI prediction filter for dynamic wavefront correction in a digital holography setup, it is not unreasonable to think that a fully adaptive filter can be used in lieu of an LTI filter. While such filters are not presented in this dissertation, the results in Chapter 7 demonstrate the greatly positive implications and feasibility for implementing a fully adaptive predictive algorithm during real-time digital holography dynamic wavefront estimation.

1.3 Introduction of Local Sharpening

Often in application, only a particular region of an extended object is required to be imaged. A region-of-interest (ROI) can readily be sharpened by isolating the region on the image plane, usually at the expense of degradation in areas outside the region. Since the target area is only a subset of the full pixel array, the sharpness optimization does not necessarily minimize the wavefront estimation error. Therefore, the corrected phase field produces less sharpening outside the ROI and is usually a poor estimate of the true wavefront when compared to global sharpening. This fact may limit applications interested in wavefront sensing. Figure 1.2 demonstrates these effects of local sharpening. The top row illustrates the full image and the ROI obtained by global sharpening. The bottom two images show the full image and the ROI obtained by local sharpening. Clearly, the local area experiences improvement in quality. However, the wavefront estimate used to reconstruct that ROI is not a good estimate of the full wavefront and indeed degrades the remaining pixels of the full image.

Earlier works have studied the effects of ROI sharpening [11, 25, 26]. The increasing attention to localized image sharpening is due to its demonstrated potential in improving image correction distorted by anisoplanatic wavefronts [27]. Anisoplanatism is the phenomena where the propagation paths of beams are not spatially invariant, as in the case of atmospheric turbulence with various wavefronts that are weakly correlated over a field-of-view. By partitioning a recorded image, one could presumably sharpen and splice several near-isoplanatic patches together to reproduce a higher quality image than one corrected over the entire field-of-view. Despite calculating the sharpness cost metric over a subset of pixels on the image plane, local sharpening must still propagate the full hologram pixel-grid to the image plane for the sharpness algorithm to administer its optimization. Therefore, for each iteration of the sharpness optimization, the method of localized sharpening makes inefficient use of the measured hologram information when interest is relegated to only a small region-of-interest.

While part of this dissertation aims to dramatically reduce the number of extremely expensive sharpness iterations required in wavefront estimation through prediction, the other part focuses on reducing the computational cost of the remaining few iterations. This is done by taking a subspace of the hologram that characterizes the ROI on the image plane, effectively reducing the size of the necessary propagations and substantially improving operational efficiency. This “subspace correction” method additionally permits the potential to perform wavefront estimation, a feature that is infeasible with earlier local sharpening methods.

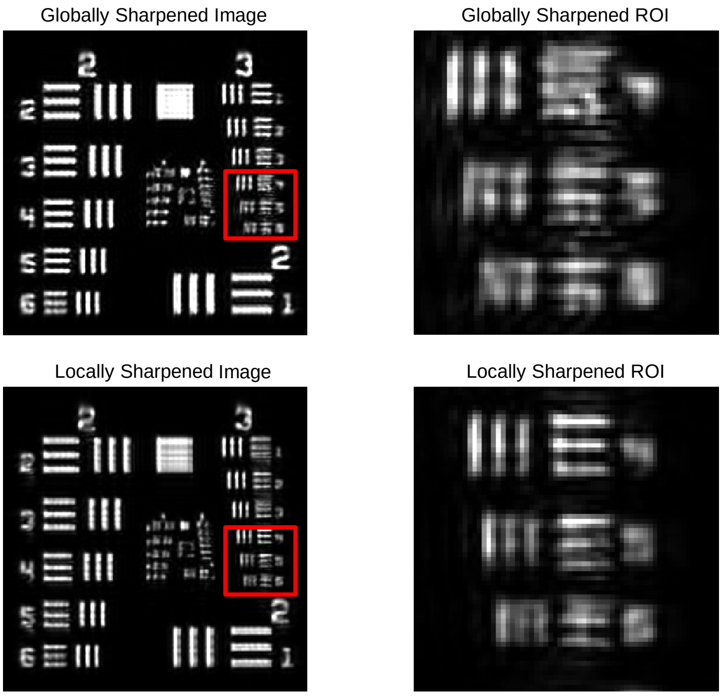

Figure 1.2: Top row: image sharpened globally. Bottom row: image sharpened locally. Both rows: complete image on left, region of interest (ROI) on right.

CHAPTER 2

Simulation Setup and Relevant Digital Holography Theory

Unlike conventional photography, holography allows the capture of the phase of complex wavefields. This information permits the reconstruction of different perspectives of an object. Digital holography can be conveniently described in two distinct procedures: recording and numerical reconstruction.[1]

2.1 Simulation Setup

For reference, Table 2.1 summarizes the optical parameters for this simulation. The distance, z, and off-axis angle, θ, are chosen to avoid aliasing issues, while remaining in the near-field regime. Executing predictive dynamic digital holography in the near-field regime is not a requirement; in fact, some results will later be shown to demonstrate its viability in the far-field as well.

A typical digital holography recording geometry is shown in Figure 2.1 and outlines the associated signal-processing details. This off-axis image plane recording geometry [28] is extremely common among wave propagation simulations. First, it is assumed that the object is uniformly illuminated with the on-axis monochromatic light from a master-oscillator laser, which has a wavelength of λ. The reflected light from the object is coherently imaged onto a detector. The detector is located in an image plane, and so this work will henceforth refer to the image plane as the detector plane. For all intents and purposes, the aero-

---

[1] Dr. Mark Spencer's guidance and collaboration were essential to the theoretical and computational development of the digital holography model here.

Table 2.1: Parameters for digital holography simulation

Parameter Name [Symbol]	Value [units]
Object Image Size [M]	256×256 [pixels]
Aperture Diameter [D]	1 [cm]
Pixel Pitch [L]	10 [micron]
Grid Spacing [$\frac{L}{M}$]	0.039 [micron]
Wavelength [λ]	600 [nm]
Propagation Distance [z]	4.3 [cm]
Off-Axis Angle [θ]	0.3 [deg]
Simulation Run-Time	2000 [frames]

optical aberrations used in this dissertation are assumed to exist in the exit-pupil plane of the coherent-imaging system, which has a circular aperture of diameter d. To determine the amplitude-spread function of the coherent imaging system [29], the simulated exit-pupil plane is zero-padded from a 256×256 pixel grid to a 512×512 pixel grid and Fourier transformed, ultimately generating the coherent-imaging systems impulse response. The object is also zero-padded from a 256×256 pixel grid to a 512×512 pixel grid and then convolved with the amplitude-spread function to simulate the effects of the isoplanatic phase errors, $\psi(\zeta, \eta)$, commonly found when imaging in the presence of aero-optical aberrations. This convolution provides the complex field corrupted by wavefront disturbances.

The corrupted complex field is ultimately interfered with the tilted-planar monochromatic light from a local oscillator. Next, the resulting digital hologram is Fourier transformed to go to the spectral plane, where the appropriate complex pupil function is band-pass filtered. The resulting complex pupil function is then zero-padded from a 256×256 pixel grid to a 512×512 pixel grid, inverse Fourier transformed, and magnitude squared to obtain the reconstructed image. This last step captures the reconstructed image intensity, $I(\zeta, \eta)$, and enables image sharpening, which is the subject of Chapter 4.

In Figure 2.2, a more complicated recording geometry is used for the simulation in this dissertation. As such, the subsequent theory is given with this model in mind. The primary

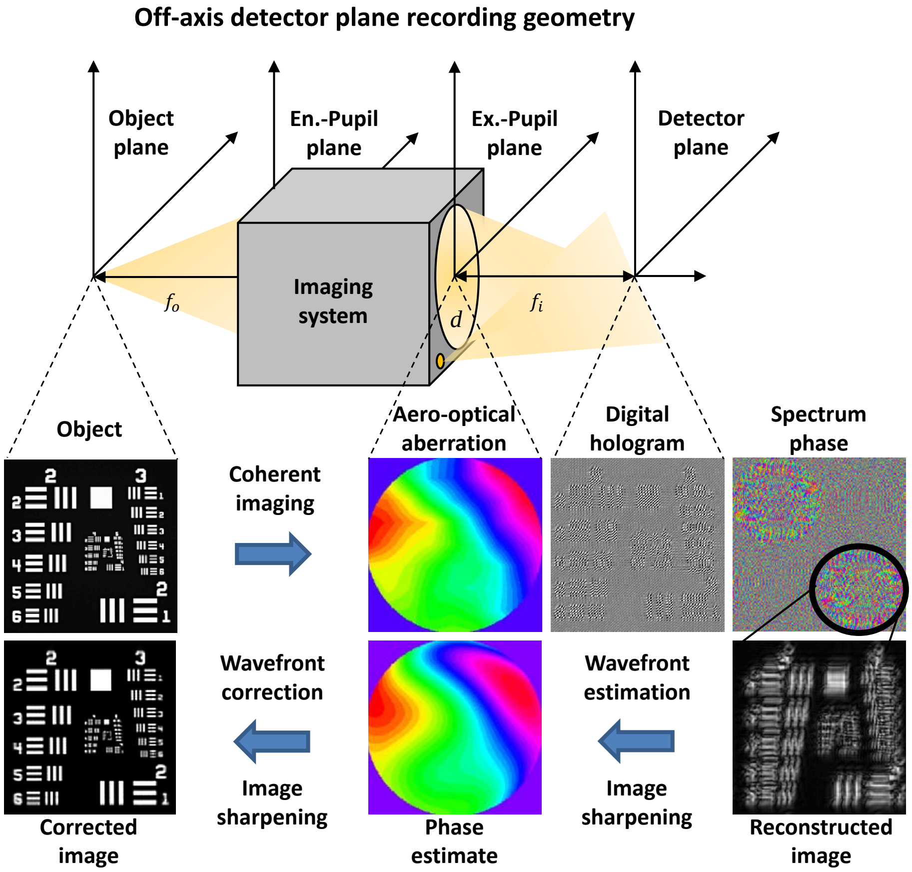

Figure 2.1: Off-axis detector plane recording geometry.

difference is that the object is placed further back to some initial plane and then propagated some distance z to the object plane. This defocused object, rather than the original object, is coherently imaged onto the detector plane. The purpose of imaging a defocused object is to represent phase error commonly found in real applications, such as imaging in the presence of aero-optical aberrations. The defocused object is imaged onto the detector plane as before in Figure 2.1, however, during image reconstruction, the process includes back-propagation to the initial plane in addition to the inverse Fourier transformation. The pupil plane phase estimation is not altered by this change in the optical setup, as the defocus is attributed to wave propagation rather than the pupil plane aberrations.

Before moving on, it is important to note that the simulation described above does not include the effects of speckle. With that said, the goal of this dissertation is to introduce the combination of wavefront prediction and system identification with digital holography

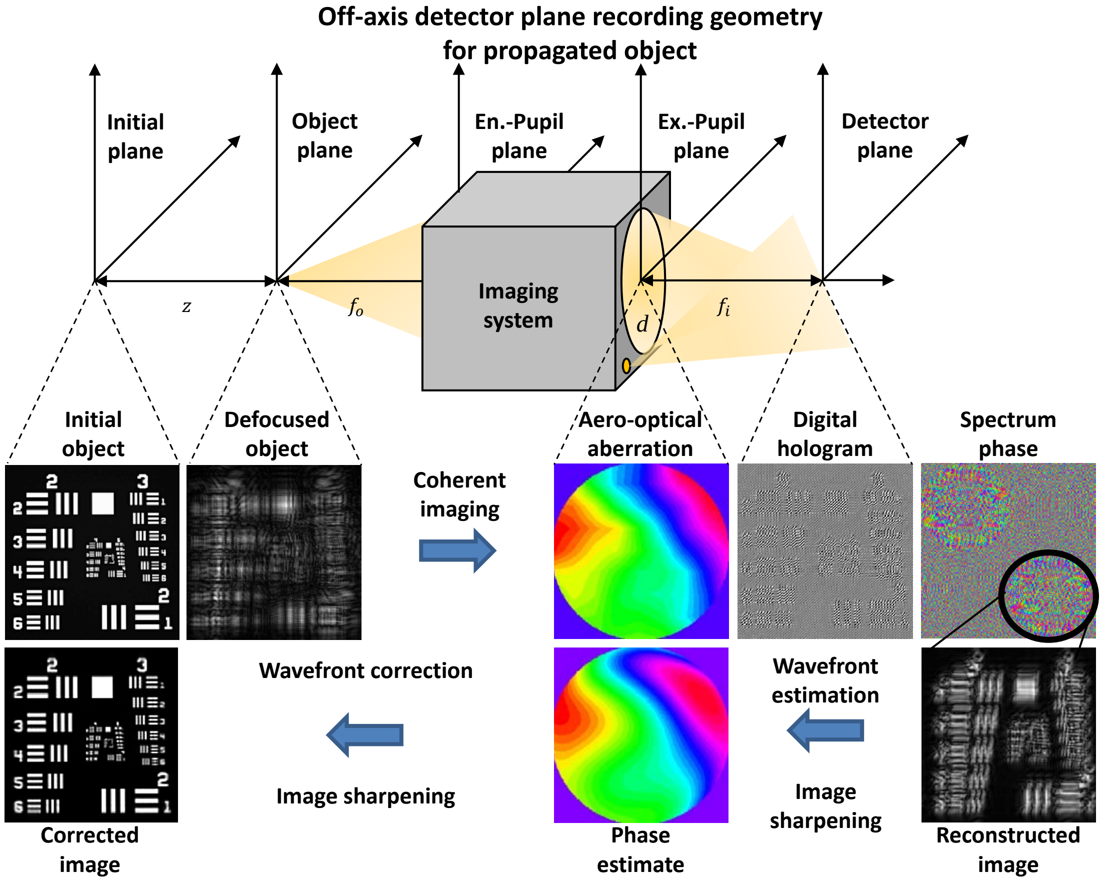

Figure 2.2: The digital-holography simulation model that enables wavefront estimation and correction using image sharpening.

and image sharpening, and to present an initial, although detailed, simulation study of the interplay between the methods from the two different fields. The effects of speckle will be an important topic of future investigations; however, including them here would require even more simulation results in an already rather elaborate initial simulation study of the proposed methods, making the overall simulation results less clear, and therefore, less impactful.

From the point of view of the methods introduced in the following chapters, speckle amounts to noise. The minimum-variance prediction filters used in this dissertation and the subspace system identification methods for determining the filters typically yield optimal performance in the presence of noise, although sufficiently high levels of noise (including speckle) can be expected to degrade performance, as is true for methods for signal and image processing.

2.2 Digital Holography Recording

Recording the amplitude and phase of a coherently illuminated object is achieved through interferometry. The local oscillator (reference beam) complex field, $r(\zeta, \eta)$ and object complex field, $u(\zeta, \eta, z)$ can be represented initially as,

$$r(\zeta, \eta) = |r(\zeta, \eta)| \exp[j\phi_r(\zeta, \eta)], \tag{2.1}$$

$$u(\zeta, \eta, 0) = |(u(\zeta, \eta, 0)| \exp[j\phi_u(\zeta, \eta, 0)], \tag{2.2}$$

where ζ and η are spatial coordinates, ϕ_r is the phase of the LO, and ϕ_u is the phase of the object field. Generally, a plane wave LO beam is desired, thus, ϕ_r will be 0. However, off-axis recording will introduce an offset angle tilting the plane waves, so the general form of Eq. 2.1 is ultimately useful. After flood illuminating the object, the diffracted waves travel some distance z to an object plane, where it interferes with the local oscillator (see Fig. 2.2).

Two common recording propagators are: the Fresnel diffraction method and the angular spectrum method [30, 31]. As the name suggests, the former uses the Fresnel diffraction integral to propagate the diffracted complex field to the object plane. Employing the paraxial approximation, it can be written as,

$$u(\zeta, \eta, z) = -\frac{j\tilde{k}}{2\pi z} \exp(j\tilde{k}z) \exp[\frac{j\tilde{k}}{2z}(\zeta^2 + \eta^2)]\mathcal{F}\Big\{u(\zeta, \eta, 0) \exp[-\frac{j\tilde{k}}{2z}(\zeta_0^2 + \eta_0^2)]\Big\}, \tag{2.3}$$

where $\tilde{k}$ is the angular wavenumber, $\tilde{k} = \frac{2\pi}{\lambda}$, and λ is the wavelength of propagating light. The $\mathcal{F}$ operator symbolizes the two-dimensional Fourier transform,

$$\mathcal{F}[\cdot] = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} (\cdot) \exp\left(-\frac{j\tilde{k}}{z}(\zeta\zeta_0 + \eta\eta_0)\right) \mathrm{d}\zeta \mathrm{d}\eta. \tag{2.4}$$

Angular spectrum uses two 2D-Fourier transforms to propagate plane wave spectral components and reconstruct the field at a plane some distance away. That is, the object complex field is first brought into the spectral domain before propagating by the free space transfer function, and then brought back to the spatial domain. For this reason, the angular spectrum method is also known as the spectral method. Using an angular spectrum propagator, the diffracted object field in Eq. 2.2 arrives at the object plane at a distance z by,

$$u(\zeta,\eta,z) = \mathcal{F}^{-1}\Big\{\mathcal{F}\{u(\zeta,\eta,0)\}\exp\Big[-jz\tilde{k}\lambda\sqrt{\frac{1}{\lambda^2}-f_\zeta^2-f_\eta^2}\Big]\Big\}, \tag{2.5}$$

where f_ζ and f_η are spatial frequency coordinates. For scalar wave propagation in the paraxial regime, the two propagators are equivalent. Therefore, the choice of propagator is often contingent on the physical properties of the optical system. For the simulation work presented in this dissertation, the angular spectrum method is used for holographic recording, while either the Fresnel diffraction integral or angular spectrum method are used during reconstruction.

The recorded intensity of the incident light on the CCD (i.e., detector plane) is the local oscillator superimposed with the coherently imaged object, written as the intensity of the sum of Eq. 2.1 and Eq. 2.5, $h(\zeta,\eta) = |u(\zeta,\eta,z) + r(\zeta,\eta)|^2$, which expanded yields,

$$h(\zeta,\eta) = |r(\zeta,\eta)|^2 + |u(\zeta,\eta,z)|^2 + u^*(\zeta,\eta,z)r(\zeta,\eta) + u(\zeta,\eta,z)r^*(\zeta,\eta). \tag{2.6}$$

As a reminder, the propagated object in Eq. 2.6 is coherently imaged onto the detector plane, represented by a convolution with a circular pupil function, $\mathcal{A}(\zeta,\eta)$,

$$u(\zeta,\eta,z) = |(u(\zeta,\eta,z)|\exp[j\phi_u(\zeta,\eta,z)] \star \mathcal{F}\{\mathcal{A}(\zeta,\eta))\}. \tag{2.7}$$

Equation 2.6 fully describes the hologram recorded on the CCD. Despite only measuring the intensity of the superimposed beams, the fringe pattern encodes both phase and amplitude

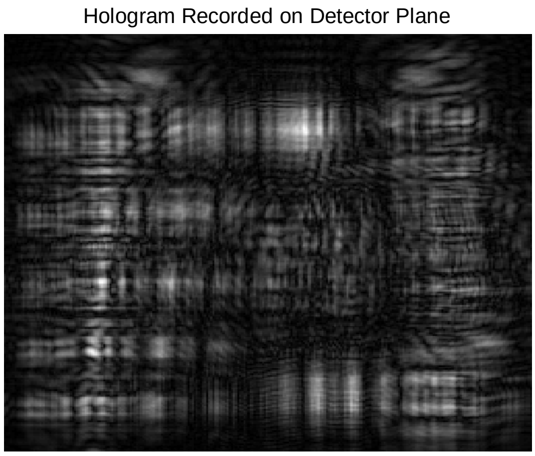

Figure 2.3: Sample hologram recorded on detector plane.

information that can be used to numerically reconstruct the image. An example of a recorded hologram fringe pattern is shown in Figure 2.3. Inverse Fourier transforming the hologram in Eq. 2.6 to the hologram's spectral plane leads to,

$$H(f_\zeta, f_\eta) = |R(f_\zeta, f_\eta)|^2 + |U(f_\zeta, f_\eta)|^2 + U(f_\zeta, f_\eta)^* \star R(f_\zeta, f_\eta) + U(f_\zeta, f_\eta) \star R(f_\zeta, f_\eta)^*, \quad (2.8)$$

where $\star$ denotes convolution. The four terms in Eq. 2.8 are, respectively: the autocorrelation of the LO and object complex fields constituting a DC term, a distorted holographic twin-image term, and the virtual image term (object complex pupil function).

When the diffracted object and reference beams propagate along the same axis, as with in-line holography, the DC and twin-image terms are superimposed with the object complex pupil function in the center of the detector plane, contaminating the image reconstruction. In order to avoid this amalgamation of parts, one of two methods of holographic recording are often utilized: phase-shift interferometry and off-axis holography. Phase-shift interferometry maintains the base in-line holographic recording setup with the addition of polarizing plates [32]. These plates allow for the capture of four phase shift permutations of recorded interference patterns (holograms) on the CCD. This collection of patterns are used for phase field retrieval of the object complex pupil function and avoid processing out the DC and out-of-focus twin-image terms. However, this procedure requires a static target to capture the four interferometry holograms. For the work investigated in this dissertation, a dynamic wavefront sequence in real-time is involved. Therefore, the second method of digital holographic recording, off-axis holography, is used instead.

In order to avoid overlapping of the various holographic constituents, the terms are spatially separated by introducing the local oscillator to an offset angle from the hologram axis, as in Figure 2.2. Since we have full authority of the local oscillator, the phase of the LO in Eq. 2.1 becomes a function of the offset angle, $\phi_r(\zeta, \eta, \theta) = \exp(j\tilde{k}\zeta \sin\theta)$. These separated terms are reconstructed and displayed in Fig. 2.4, to plainly distinguish between them. Clearly, the zeroth-order term is suppressing the image content of the hologram in the left capture. In the right capture, the virtual image term appears in the top-left and the out-of-focus holographic twin-image term in the bottom-left. To ensure that the frequency

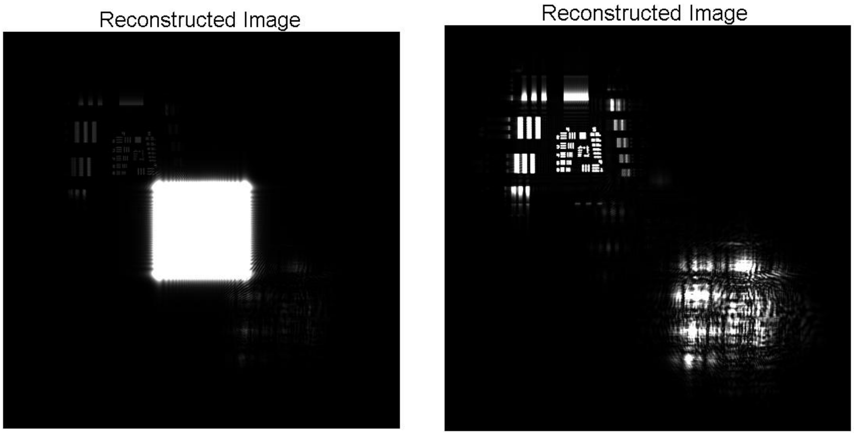

Figure 2.4: Left: Unfiltered, Right: Filtered. Reconstructed image, DC term (middle), and twin-image term.

resolution of the spectral domain of the detector plane is large enough for the spatial content of the object signal to maintain its spatial resolution and be distinctly identifiable, the detector plane is zero-padded. This prevents aliasing and pixel truncation issues that may occur from a lower frequency resolution spectral domain. Unfortunately, zero-padding does not prevent aliasing occurring from under-sampling. Under-sampling can occur from an offset angle that is not large enough or from the object propagated at an insufficient distance from the imaging system. Increasing the offset angle in off-axis holography can solve aliasing from under-sampling, but requires the imaging system be placed at a further distance from the object. Consequently, this limits the available bandwidth of a CCD during off-axis recording.

Once the hologram has been successfully recorded on the detector plane, a plethora of options are available to realize the image. Historically, realizing an image through optical holography required a coherent light with the same wavelength used in recording to illuminate the fringe pattern on the photographic plate [1, 4]. With the shift to a digital paradigm, reconstructing a hologram has taken a numerical approach.

2.3 Numerical Reconstruction in Digital Holography

Numerical reconstruction begins with suppressing and filtering undesired terms from the hologram spectral plane. The left capture of Fig. 2.4 illustrates the initial situation on an $\ell \times m$ grid of pixels. Degradation of image reconstruction in the left capture is largely due to the DC term [33]; the effect can be mitigated by,

$$\bar{h}(\zeta, \eta) = h(\zeta, \eta) - \frac{1}{(\ell + m)} \sum_{(\zeta, \eta)} h(\zeta, \eta), \tag{2.9}$$

to produce the situation in the right capture. For the simulation developed in this work, the hologram was subsequently tilted by the same phase as the local oscillator, $\phi_r(\zeta, \eta)$. This allows the object complex pupil function, $U(f_\zeta, f_\eta)$, to be easily cropped with a spatial filter, by centering it in the spectral plane of the hologram. A smooth spatial filtering profile, such as Gaussian or Tukey windows [34], is used to extract the appropriate complex pupil function while preserving its frequency content, particularly near the edges.

Once spectral domain filtering is accomplished, the image can be reconstructed by first inverse Fourier transforming the object complex pupil function on the hologram spectral domain, $U(f_\zeta, f_\eta)$, to the detector plane and then back-propagating to the initial plane with the Fresnel diffraction integral (or angular spectrum method). Once again, the spectral domain is zero-padded before the transformation in order to preserve the spatial resolution of the image on the detector plane. Using inverse Fresnel diffraction, the full back-propagation can be written as,

$$u(\zeta, \eta, 0) = \frac{j\tilde{k}}{2\pi z} \exp(j\tilde{k}z) \exp[\frac{j\tilde{k}}{2z}(\zeta^2 + \eta^2)] \mathcal{F}^{-1}\Big\{\mathcal{F}\{U(f_\zeta, f_\eta)\} \exp\big[j\frac{\pi}{\lambda z}(\zeta^2 + \eta^2)\big]\Big\}. \tag{2.10}$$

For an optical system with an ideal point-spread function, Eq. 2.10 is sufficient to reconstruct the image as it was recorded. Introducing isoplanatic wavefront aberration, $\psi(\zeta, \eta)$, as in the presence of aero-optical effects, to the diffracted object beam will cause errors in the complex field. While amplitude and phase disturbances both affect the integrity of the reconstructed image, the reconstruction is much more robust to amplitude errors. Therefore, the focus of this dissertation will be on phase degradation as it is the primary source of image error [31]. To represent the complex field on the object plane passing through the wavefront, the object is convolved with the point-spread function at the detector plane. The corrupted complex field of the imaged object then becomes,

$$u_\varepsilon(\zeta,\eta,z) = |(u(\zeta,\eta,z)| \exp[j\phi_u(\zeta,\eta,z)] \star \mathcal{F}\{\mathcal{A}(\zeta,\eta)\exp(j\psi(\zeta,\eta))\}, \tag{2.11}$$

where the expression in the 2-D Fourier Transform is the complex pupil function injected with the isoplanatic wavefront aberrations, $\psi(\zeta,\eta)$.

Figure 2.2 demonstrates an overview of the diffracted object wavefield passing through phase aberration in the digital holography simulation for the research presented here. After reconstructing the corrupted image (Eq. 2.11) from the spectral plane of the hologram by numerical propagation through Eq. 2.10, the resultant image is originally corrupted. This initially degraded reconstruction is demonstrated in Figure 2.2. If the wavefront aberration were known perfectly, the phase error at the pupil plane could simply be subtracted out from the corrupted image in Eq. 2.11 to reproduce the diffracted-limited object. Unfortunately, the wavefront fields are generally not known *a priori*, particularly for real-time applications. Instead, this phase field is estimated, $\hat{\psi}(\zeta,\eta)$, through wavefront sensors or numerical methods, and corrected on the spectral plane of the hologram to generate an estimate of the image, given as,

$$\begin{aligned}\hat{u}(\zeta,\eta,0) = \frac{j\tilde{k}}{2\pi z}\exp(j\tilde{k}z)\exp[\frac{j\tilde{k}}{2z}(\zeta^2+\eta^2)] \\ \times \mathcal{F}^{-1}\Big\{\mathcal{F}\{\{U_\varepsilon(f_\zeta,f_\eta)\exp[-\hat{\psi}(f_\zeta,f_\eta)]\}\}\exp\big[j\frac{\pi}{\lambda z}(\zeta^2+\eta^2)\big]\Big\}.\end{aligned} \tag{2.12}$$

To take advantage of digital holography's extremely minimalist nature with regards to necessary electro-optical equipment, a numerical algorithm is used for phase retrieval rather than expensive, bulky, and space-constraining sensors. In this dissertation, the sharpness optimization algorithm is utilized to determine an estimation of the pupil plane phase errors (i.e., wavefronts) used for correction in the corrupted image, by maximizing sharpness criteria. Before introducing this algorithm, this dissertation will analyze the aero-optical wavefronts present in the next chapter.

CHAPTER 3

The Aero-Optical Wavefronts

This section details the aero-optical wavefront sequence used for predictive dynamic wavefront correction in digital holography. Notre Dame's Airborne Aero-Optics Laboratory (AAOL) obtained experimental data from aero-optical wavefronts produced by turbulence over a flat-windowed turret [2, 3, 35]. Flight tests involved two planes in constant formation transmitting continuous-wave lasers between them. A sequence of 8000 wavefront aberrations were recorded at a sampling rate of 16 kHz in these experiments. Details of the flight tests are given in Table 3.1.

Figure 3.1 (left) shows a sample aero-optical wavefront frame. Data was not recorded on the outside regions of the annular aperture, nor on the obscuration in the center. To avoid the obscuration at the center of the wavefront's 30×29 pixel aperture, the simulation pupil plane phase errors are generated by taking a section of 9×9 pixels and interpolating the data into a grid of 256×256 pixels. Three such 9×9 pixel patches were superimposed at each wavefront frame to invent the simulation pupil plane phase errors, while maintaining

Table 3.1: Notre Dame AAOL experimental data details

Turret Azimuthal Angle	119 deg
Turret Elevation Angle	57 deg
Freestream Mach	0.36
Altitude	4570 m
Target Distance	50 m
Aperture Size	10.1 cm
Sampling Rate	16 kHz

temporal correlation of the dynamic wavefront series. The phase errors were then magnified by 6, producing a wavefront like in Fig. 3.1 (right). The final dynamic sequence comprises RMS wavefront errors in the range of $[-7, 8]$ microns, where the simulation wavefront sample in Fig. 3.1 is a typical aberration field in the sequence. The magnitude of atmospheric turbulence is greatly above that of the flight tests to put the sharpness algorithm and prediction filters under rigorous testing. Finally, as explained in Chapter 4, the piston and tip-tilt Zernike modes (i.e., the first three modes) used as part of modeling the aero-optical wavefronts are adjusted for separately from image sharpening. Therefore, the pupil plane phase field shown is with those modes removed.

A requirement for wavefront correction through predictive dynamic digital holography is that the wavefront sequence be temporally correlated. This requirement arises due to the use of a subspace system identification algorithm to identify an LTI filter (or implementing fully adaptive filters). One might expect that wavefronts from atmospheric turbulence captured at a 16 kHz sampling rate would exhibit such correlation. Temporal correlation for the original aero-optical wavefront sequence and for the modified wavefront sequence are shown in Figures 3.2 and 3.3, respectively. These figures support that temporal correlation exists in the wavefront sequences.

Phase errors, defined on a circular pupil plane (Fig. 3.1) are zero-padded into a 512×512 pixel grid and Fourier transformed to generate the optical system's point spread function, or impulse response of the optical system. Ideal optical systems that are diffracted-limited beget a point spread function that resembles the Airy disk. Imperfect systems corrupted with wavefront errors have broader point spread functions and loss of peak intensity. Characterizing the quality of an optical system under atmospheric turbulence, relative to a theoretically perfect scenario, can be done using the Strehl ratio metric [36]. Lord Rayleigh determined that point sources that can be resolved within $\frac{\lambda}{4}$, or a Strehl ratio of 0.82, were considered practically diffracted-limited [37]. However, this benchmark is at best a minimum; high-end optical systems may desire greater performance for their needs. Generally speaking, this criterion for point-spread functions marks the point that diffraction, rather than atmospheric aberration, comprises the majority of image degradation. That is to say, it does not grant a reprieve from achieving further atmospheric wavefront correction. Many scientific and industrial systems require Strehl ratios up to 0.90, beyond the Rayleigh limit of resolving point sources [38]. Various methods for calculating the Strehl ratio exist, though the industry-standard is the "Marechal approximation",

$$SR \approx 1 - (2\pi\sigma)^2, \tag{3.1}$$

where σ is the root-mean-square of the wavefront over the aperture [39]. This approximation is simply the first two terms of the exponential expansion, also known as the "extended Marechal approximation" [38, 40],

$$SR \approx \exp(-\sigma^2). \tag{3.2}$$

In Figure 3.4, the Strehl ratio using Eq. 3.2 is calculated and plotted for the modified wavefront sequence used for the pupil plane phase errors in this dissertation. With an average uncorrected Strehl of 0.275, the plot makes it evident that the digital holography optical system is severely degraded and requires extensive wavefront sensing and correction to achieve at least the Rayleigh criterion. Without predictive methods, this effort would be computationally (and possibly physically) expensive. Before discussing predictive dynamic digital holography for wavefront estimation, the next chapter will explain the sharpening optimization algorithm and model representations for this aero-optical wavefront data.

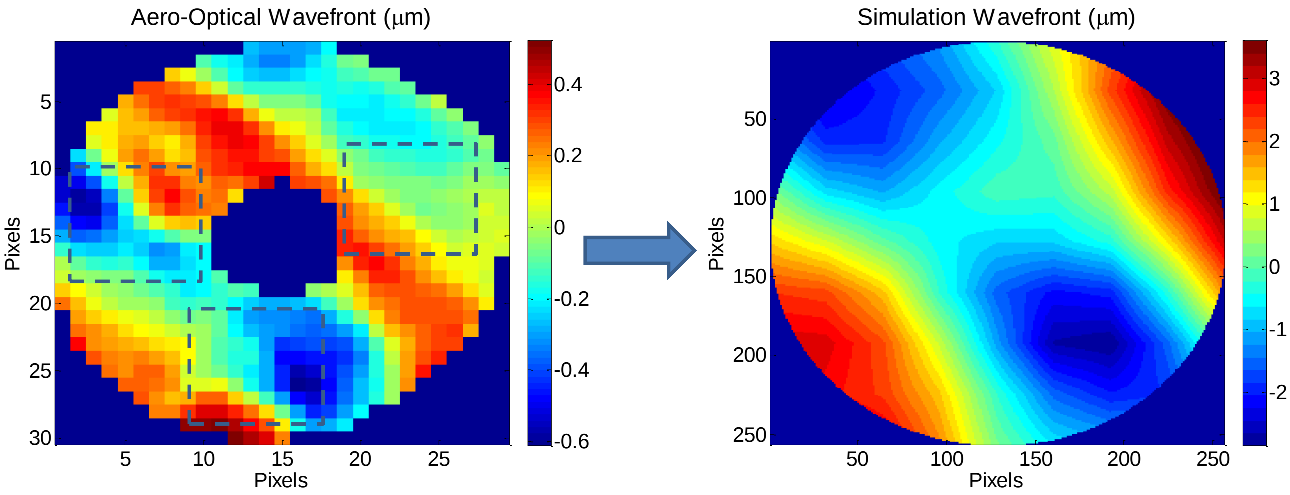

Figure 3.1: Sample aero-optical and constructed simulation wavefront. Wavefront errors are expressed in microns (μm). Tilt and piston modes removed.

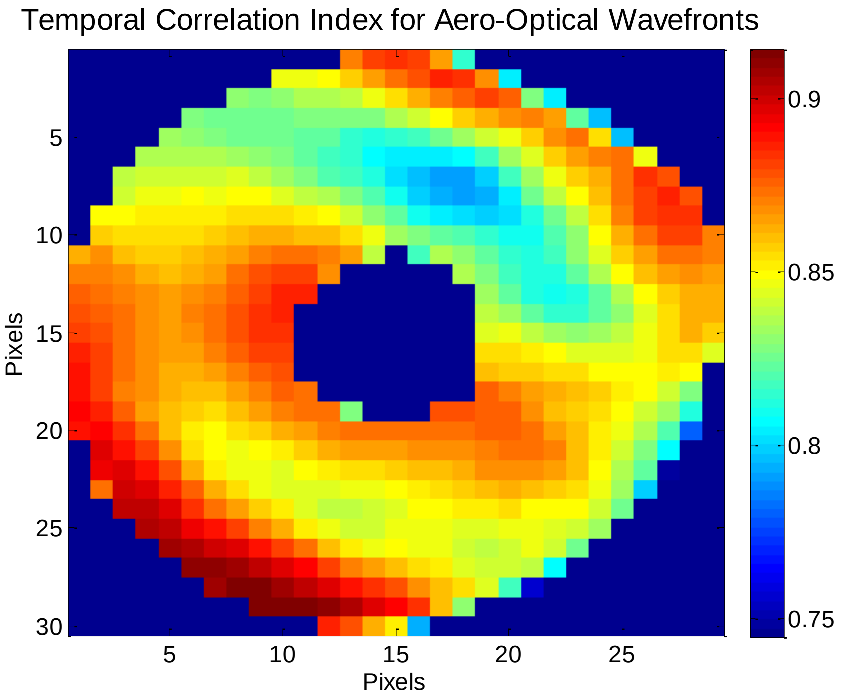

Figure 3.2: Temporal correlations for aero-optical wavefronts over full sequence.

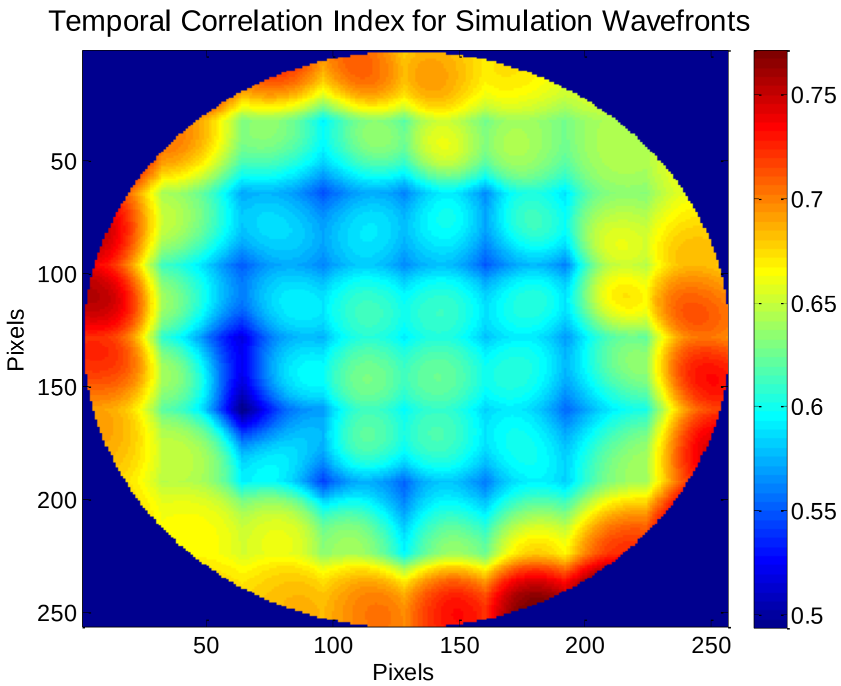

Figure 3.3: Temporal correlations for simulation constructed wavefronts over full sequence. Piston and tilt modes removed.

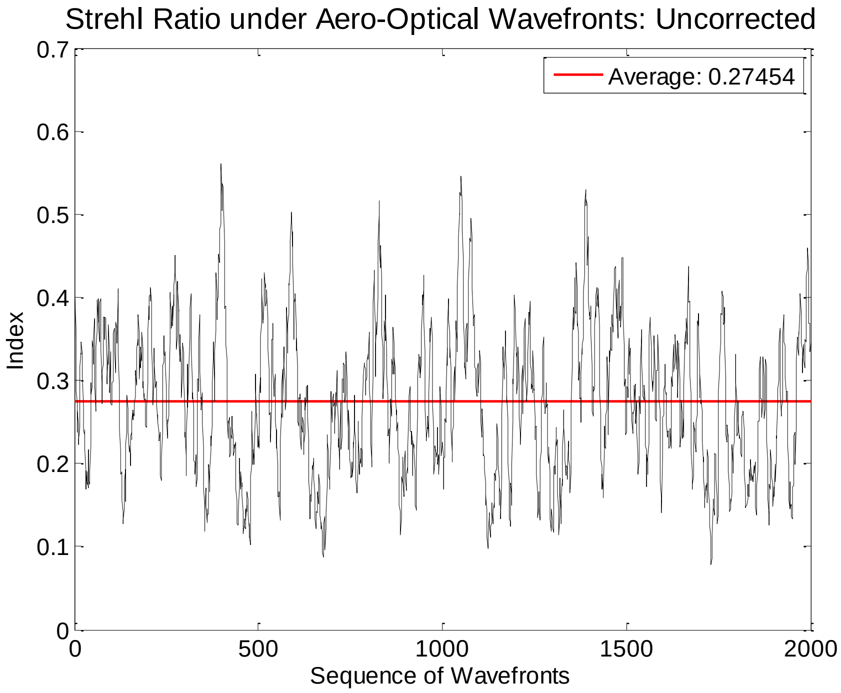

Figure 3.4: Strehl Ratios calculated with simulation-modified pupil plane phase errors. Piston and tilt modes not included in wavefront sequence.

CHAPTER 4

Sharpness Algorithm and Wavefront Modeling

4.1 Sharpness Metrics

The sharpness algorithm [8] performs non-linear, non-convex optimization to maximize a criterion metric measuring the reconstructed image intensity. Several possible sharpness metrics are listed in Ref. [8], where a typical choice seen in application [26, 41] is often,

$$S = \sum_{(\zeta,\eta)} I(\zeta,\eta)^{\beta}. \tag{4.1}$$

The choice of which to metric to use, or what value of β to assign in the case of Eq. 4.1, depends on the nature of the illuminated object. Some metrics, such as those in Eq. 4.1 with a $\beta \geq 2$, concentrate on making bright pixels brighter, while others, such as $\beta < 2$, focus on making dark pixels darker [9]. Unsurprisingly, a significant consideration when

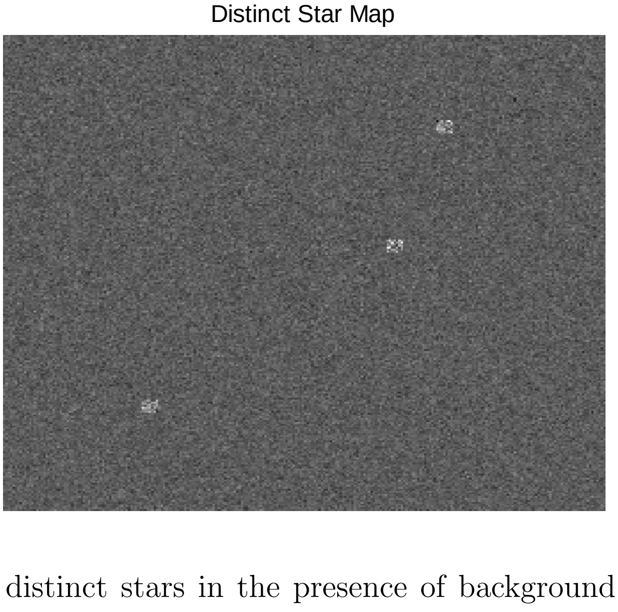

Figure 4.1: Image of distinct stars in the presence of background noise, a situation well suited for sharpness metrics with a high β ($\beta \geq 2$).

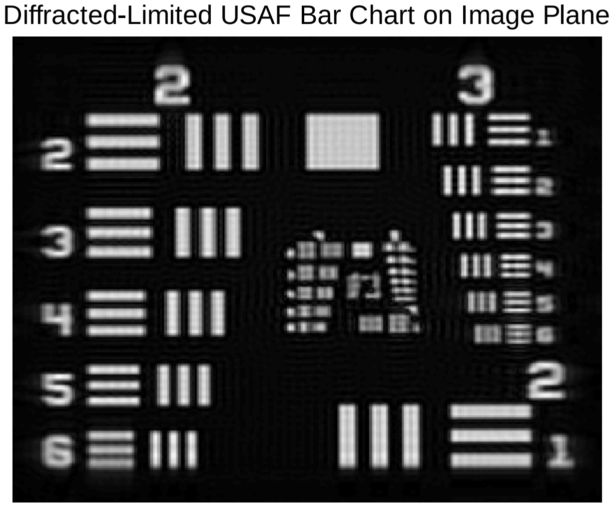

Diffracted-Limited USAF Bar Chart on Image Plane

Figure 4.2: Diffracted-limited USAF bar chart reconstructed on the detector plane - the ideal scenario.

choosing the sharpness index is the expected amount of targeted "bright points" in the image. Attempting to sharpen a few bright stars, for example as in Fig. 4.1, would ideally use a high β ($\beta \geq 2$) value for best performance and minimal computation. The work done by Fienup and Miller [9] shows the solution space for this case is smooth for images with fewer, distinct bright targets. In this dissertation, a U.S. Air Force bar chart (Fig. 4.2) is the target image during wavefront estimation. It consists of several locales of bright pixels parsed intermittently among very low-intensity dark regions. With this in mind, two metrics are considered in this work,

$$S_1 = \sum_{(\zeta,\eta)} I(\zeta,\eta), \tag{4.2}$$

$$S_2 = \sum_{(\zeta,\eta)} I(\zeta,\eta)\log I(\zeta,\eta), \tag{4.3}$$

where I is the estimated image intensity after reconstruction, i.e.,

$$I(\zeta,\eta) = |\hat{u}(\zeta,\eta,0)|^2. \tag{4.4}$$

Since Eq. 4.4 dictates that only the intensity of the estimated image is necessary for sharpness optimization, the phase terms in front of the integral in Eq. 2.12 can be dropped such that,

$$\hat{u}(\zeta,\eta,0) = \mathcal{F}^{-1}\Big\{\mathcal{F}^{-1}\{\{U_\varepsilon(f_\zeta,f_\eta)\exp[-\hat{\psi}(f_\zeta,f_\eta)]\}\}\exp\Big[j\frac{\pi}{\lambda z}(\zeta^2+\eta^2)\Big]\Big\}. \tag{4.5}$$

The Shannon Entropy index, S_2, imposes similar sharpness optimization behavior as indices of $\beta < 2$ in Eq. 4.1 [9]. Specifically, it behaves similarly to the case of $\beta = 1$ (Eq. 4.6). While S_2 inherently makes dark pixels darker, S_1 is indiscriminate in its intensity distribution over the grid of pixels. In other words, it is impossible to distinguish between adding intensity to a dark or already bright pixel once the intensities are summed to determine the sharpness cost index. In light of this, S_1, in its current form, is practically meaningless for intensity optimization. To combat this, the sharpness index S_1 is maximized over a subset $\mathcal{D}$ of the brightest pixels, such that,

$$S_1 = \sum_{(\zeta,\eta)\in\mathcal{D}} I(\zeta,\eta), \tag{4.6}$$

so that sharpness optimization focuses on increasing the intensity of the brightest target pixels. In this dissertation, set $\mathcal{D}$ consists of only the 25% of pixels with greatest intensities. This maintains the smooth solution space inherent for low β values for images with dark regions, while targeting the desired features of the image.

Assuming the unaberrated image maintains the highest sharpness index, the optimization necessarily minimizes the wavefront estimation error, $\psi - \hat{\psi}$. Muller's development of sharpness metrics shows this is true only for isoplanatic wavefronts [8]. As a result, maximizing the sharpness theoretically yields the optimal wavefront estimate, for the time interval over which the wavefront is constant. This minimizing phase error estimate is subsequently used to reconstruct the optimal image through Eq. 4.5. In the case that the incoming wavefront is not spatially invariant, or anisoplanatic, other works have shown that local sharpening is capable of addressing image correction [27, 42]. Local sharpening involves isolating a subset of intensities from the image plane and determining its sharpness cost index. This technique of image sharpening targets only a subset of the full pixel array, and so the optimization routine does not necessarily minimize the wavefront estimation error. The resultant image is not corrected with the optimal wavefront estimate in this case, making local sharpening ill-suited for wavefront estimation. Additionally, it performs inefficient computations by using the full hologram spectrum to sharpen a subset of image pixels. Recent work has demonstrated a more efficient approach to local sharpening to potentially address the anisoplanatic issue, called subspace correction [12]. This technique will be investigated further in Chapter 8.

4.2 Modal Representation of Aero-Optical Wavefronts

Optimizing sharpness metrics produces phase estimates of the aero-optical wavefronts that can be parameterized to a Zernike polynomial expansion [29, 43–45], written as,

$$\hat{\psi}(t,k)(\zeta,\eta) = \sum_{n} a_n(t,k) Z_n(\zeta,\eta), \tag{4.7}$$

where $\hat{\psi}(t,k)(\zeta,\eta)$ is the pupil plane phase estimate (i.e., wavefront estimate) at time step t and sharpening iteration k, with spatial coordinates ζ and η. On the right side of Eq. 4.7, $Z_n(\zeta,\eta)$ are Zernike basis functions and $a_n(t,k)$ are the Zernike coefficients at time step t and k sharpening iterations, where n indicates the Zernike mode. The wavefront estimate is used in Eq. 4.5 to reconstruct an image estimate, and the intensity of this image is used to determine the sharpness cost index from either Eq. 4.6 or Eq. 4.3. It should be noted that Zernike expansion is not a requirement for representing the pupil plane phase estimate. Another consideration includes Gaussian basis functions. Zernikes were chosen here for their simplicity in defining an orthonormal set of basis functions over a circular aperture. This is the case for the pupil plane phase errors designed in this study.

Maximizing sharpness corrects phase disturbances due to defocus, astigmatism, and higher aberration terms; however, the first three Zernike modes of piston, tip, and tilt are simple translations of the phase and thus do not contribute to optimizing image sharpness. Consequentially, the solution space for the sharpness cost function involves an infinite number of solutions that satisfy the maximization. In practice, these modes are determined through techniques involving centroid tracking used regularly in adaptive optics applications [46, 47]. For the purposes of wavefront estimation and correction in this dissertation, these modes are assumed to have been determined and corrected; therefore, they are not included in the actual sharpness optimization. Instead, these low-order modal effects are corrected in the pupil plane phase estimate by least-squares projection.

For each positive integer N, the least-squares projections of the true pupil plane wavefronts onto the subspace spanned by the first N Zernike modes can be written as,

$$\tilde{\psi}_N(t)(\zeta,\eta) = \sum_{n=1}^{N} \tilde{a}_n(t) Z_n(\zeta,\eta), \qquad t = 1, 2, ..., \tag{4.8}$$

where $Z_n(\zeta, \eta)$ are the Zernike modes and $\tilde{a}_n(t)$ are the Zernike coefficients at time-step t determined by the least-squares fit,

$$\tilde{a}(t) = \underset{\tilde{a}_n}{\operatorname{argmin}} \sum_{(\zeta,\eta)} (\hat{\psi}_N(t)(\zeta, \eta) - \psi(t)(\zeta, \eta))^2 \quad t = 1, 2, ..., \tag{4.9}$$

for each time-step in the dynamic wavefront sequence. Since a least-squares projection represents the best estimator for the pupil plane phase estimate, the resulting fit, $\tilde{\psi}_N(t)$, are the benchmark wavefronts and the corresponding coefficients, $\tilde{a}(t)$, are the benchmark Zernike coefficients. These projections are used as a benchmark metric to quantify the theoretical limit of image sharpening within the scope of the simulation presented in this work.

Parameterizing the aero-optical wavefronts by polynomial expansion requires an analysis of the modal powers contained within the dynamic wavefront sequence. The normalized cumulative power contained within the first 50 modes is calculated in order to determine a sufficient number of Zernike modes needed to capture key spatial frequency content features of the sequence. For a time interval $T = [t_1, t_1 + 1, t_1 + 2, ..., t_2] = [t_1 : t_2]$, the normalized cumulative power in the first N Zernike modes is defined by,

$$CP(N, T) = \frac{\sum_{t \in T} \|\tilde{\psi}_N(t)\|_F^2}{\sum_{t \in T} \|W(t)\|_F^2}, \tag{4.10}$$

where $\|\cdot\|_F$ is the Frobenius norm of a wavefront and $W(t)$ are the true pupil plane wavefronts. The normalized cumulative power, $CP(N, T)$, includes the power in the first three Zernike modes (i.e., the piston and tip-tilt modes). The normalized cumulative power with these piston and tip-tilt modes removed is defined by,

$$\overline{CP}(N, T) = \frac{\sum_{t \in T} \|\tilde{\psi}_N(t) - \tilde{\psi}_3(t)\|_F^2}{\sum_{t \in T} \|W(t) - \tilde{\psi}_3(t)\|_F^2}. \tag{4.11}$$

Figure 4.3 show the results of the normalized cumulative power from Eq. 4.10 as a function of N for $1 \leq N \leq 50$ for the first half of the wavefront sequence. These frames mark the first half of the wavefront sequence that will be used in Chapter 5 for system identification. In Figure 4.3, the piston, tip, and tilt modes constitute the first three lower-order Zernike modes in red. These modes are not actively sharpened through the optimization algorithm

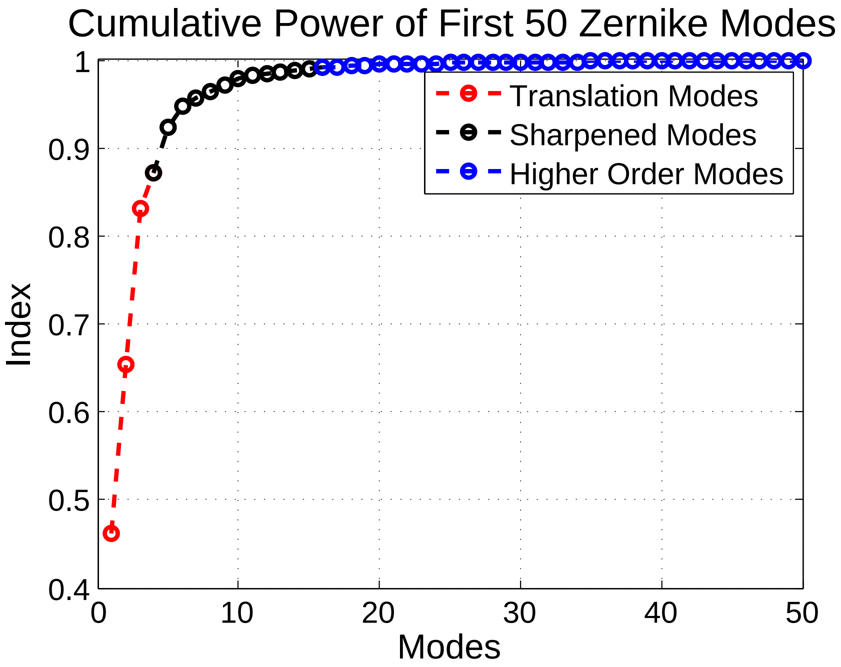

Figure 4.3: Normalized cumulative power, relative to the total wavefront power in the first 4000 wavefront frames, of First 50 Zernike Modes. Mode 15 = 99.06%.

Table 4.1: Normalized cumulative power for $N = 15$

	$CP(15,T)$	$\overline{CP}(15,T)$
$T = [1:4000]$	0.9906	0.9470
$T = [4001:8000]$	0.9903	0.9482
$T = [4001:6000]$	0.9896	0.9465

and are corrected exactly with the first three benchmark Zernike coefficients. Illustrated in black are those first 15 Zernike modes that are actively sharpened. The fraction of the power contained in these first 15 modes is 99.06%. The power in the first $N = 15$ modes for other segments, with and without the translational modes, of the data are shown in Table 4.1.

With virtually all of the spatial content captured by a polynomial expansion of 15 Zernike modes for the frames in Figure 4.3, the remaining modes are higher-order terms that are not used to model the pupil plane phase estimates. The first 15 Zernike modes represent the translational modes, defocus, astigmatism, spherical aberration, and higher orders of these degrading optical factors. Figure 4.4 provides a chart of these modes.

Figure 4.5 shows the pupil plane phase error power of the benchmark wavefronts relative to the true wavefront sequence for 2000 frames, with piston and tilt modes removed from the sequence. The 2000 frames used are the subsequent frames after the first 4000 wavefronts (i.e., first half) of the sequence, and are ultimately the operational run-time length for the simulation results in Chapter 7. Figure 4.5 demonstrates the fidelity of the theoretical performance limit achieved by the projected wavefronts. This benchmark wavefront sequence contains 94.65% of the fractional power in the first 15 modes. Also shown is an average case of a sample wavefront and fitted wavefront pair. Most of the lower-order spatial frequency content, as well as some higher-order characteristics, is captured by the projection on a Zernike polynomial expansion. However, the spatial frequency agreement is not so fine-tuned as to encroach on over-sharpening, an issue that can arise from image sharpening optimization. This, among other considerations, is studied next.

4.3 Computational Considerations

A quasi-Newton or conjugate gradient algorithm can be used to optimize the sharpness objective over the Zernike coefficients. The most expensive computing operation for each function evaluation in the sharpness optimization involves two 2-D Fourier transforms. For

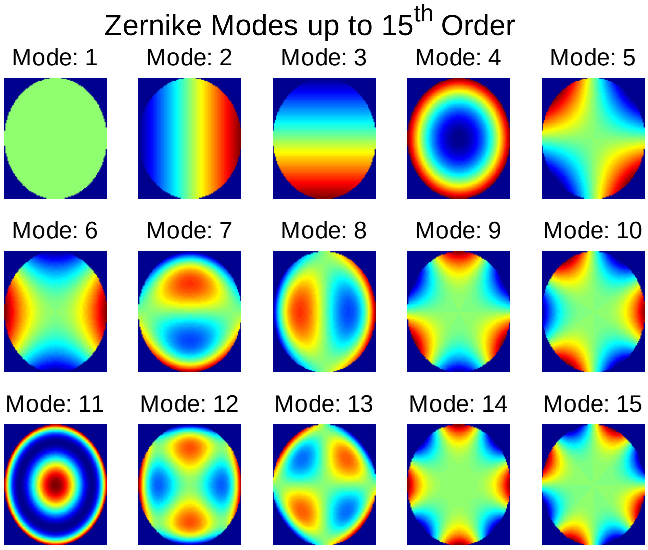

Figure 4.4: Representation of the first 15 Zernike modes over circular pupil.

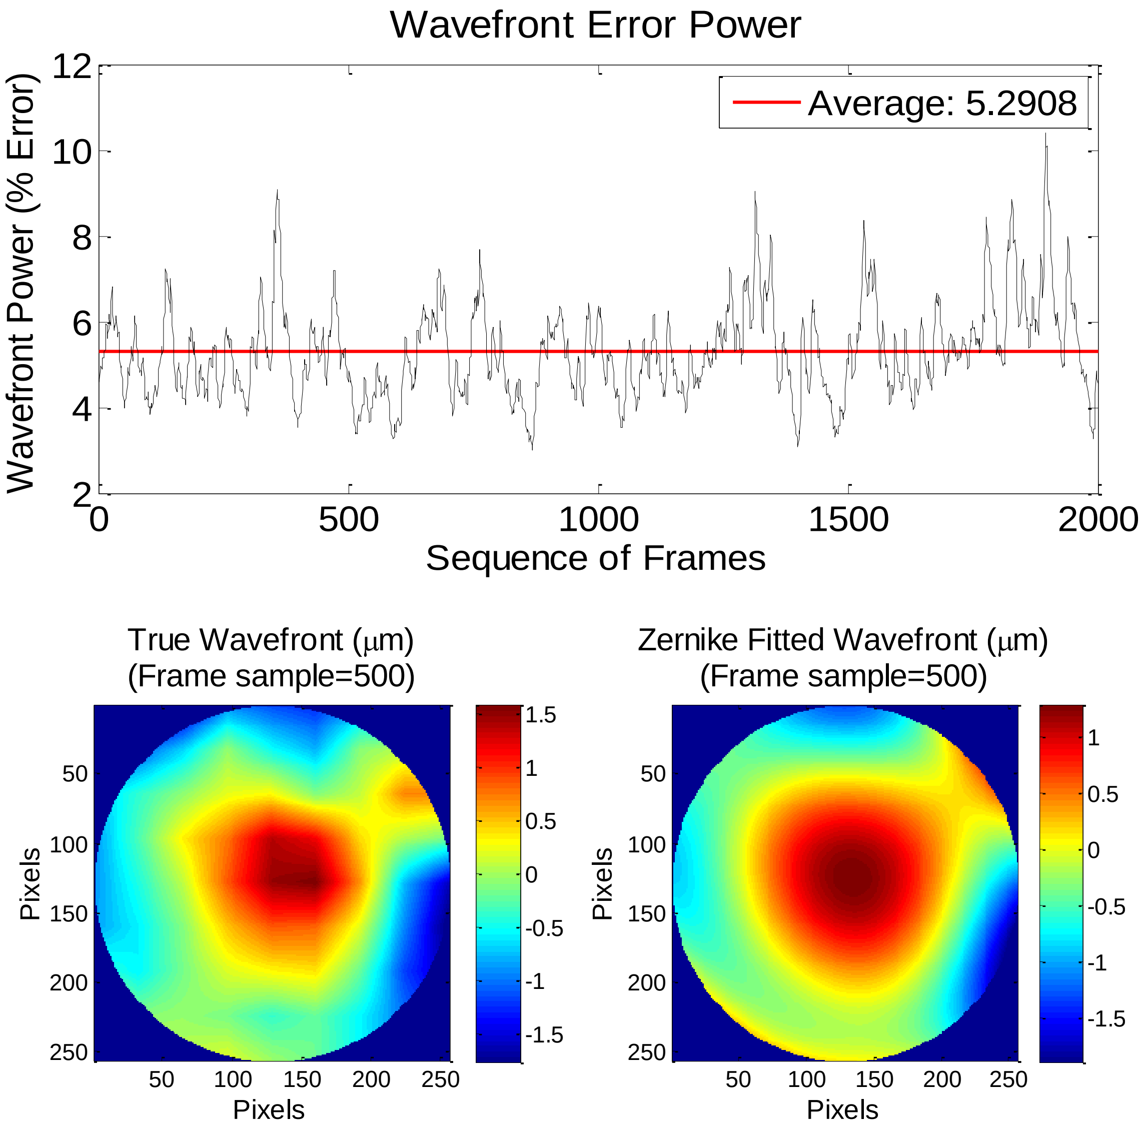

Figure 4.5: Top plot: time-series of pupil plane phase error power of true wavefronts relative to benchmark metric. Bottom images: sample of Zernike-fitted wavefront correction. Tilt and piston modes removed from wavefront sequence.

problems involving large pixel arrays or an exhaustive set of Zernike modes, this can quickly become overbearing for a gradient-based solver; thus, an analytic gradient should be computed to improve the computational performance of this optimization. An analytic gradient for the algorithm is evaluated by,

$$\begin{aligned}\frac{\partial S_1}{\partial a_n} = -2\sum_{(\zeta,\eta)} \mathrm{Im}\Bigg\{ & U_\varepsilon(\zeta,\eta)\exp[-\hat{\psi}(\zeta,\eta)]Z_n(\zeta,\eta)\ \otimes \\ & \mathcal{F}^{-1}\Big\{\exp[-j\frac{\pi}{\lambda z}(\zeta^2+\eta^2)]\mathcal{F}^{-1}\{|\hat{u}(\zeta,\eta,0)|^2\hat{u}(\zeta,\eta,0)\}\Big\}^*\Bigg\},\end{aligned} \tag{4.12}$$

for the S_1 metric, where $\otimes$ denotes the Kronecker product. The gradient vector is of N length corresponding to the number of Zernike coefficients. This analytic gradient improves reconstruction and wavefront correction performance by dramatically reducing the necessary function evaluations per iteration. Unfortunately, each sharpness iteration still involves four 2-D Fourier transforms for Fresnel propagation and gradient vector calculation (as in the model in Fig. 2.2). From here, the computational burden will have to be alleviated by targeting the number of sharpness iterations required. Predictive dynamic digital holography will explore this avenue in Chapter 5.

Theoretically, including higher Zernike modes in the wavefront parameterization will improve the correction fit, such that $n \to \infty$ replicates the true wavefront. Practically, there are a couple considerations that discourage using a considerable polynomial modal order. First, this quickly introduces increasing computational complexity on the sharpness algorithm. As Figure 4.3 demonstrates, for the scenario studied here, this trade-off for improved wavefront realization has greatly diminished returns beyond 10 Zernike modes. Second, extending the Zernike coefficients such that it is able to capture precise spatial structure can compromise the sharpening optimization by a phenomena known as over-sharpening [25, 26]. Over-sharpening is the result of the sharpening algorithm producing wavefront corrections that drive the final image intensities into star-like structures – that is, when the structure of the USAF bar chart in Fig. 4.2 is destroyed and steers toward the star chart in Fig. 4.1. Obviously, a point-source final image will result only if the wavefront correction approaches a state where it is the conjugate of the object complex pupil function.

Generally, the propagated object phase has very fine spatial structure. As long as the resolution of spatial features of the Zernike estimate are not on the order of these object phase fields, this cause of over-sharpening can be avoided. From studying the behavior of sharpness metrics, over-sharpening behaves similarly to metrics that use a high-valued β parameter. Indeed, this is another potential cause to over-sharpening and further reason to ensure the integrity of the sharpness metric used in particular imaging applications.

CHAPTER 5

Predictive Dynamic Digital Holography

5.1 Subspace Identification

Subspace system identification [48,49] has been used to identify state-space models for wavefront sequences in the past [22–24, 50]. However, while these past works implemented subspace identification in the context of adaptive optics with a wavefront sensor, subspace identification will be pioneered in this dissertation for wavefront estimation through digital holography. The subspace identification algorithm developed in [21], based on a recursive least-squares lattice filter, is used to identify linear time-invariant (LTI) prediction filters for aero-optical wavefront sequences in this dissertation. This algorithm is described in this section.

The LTI discrete-time state-space system is characterized as follows:

$$\begin{aligned} x(t+1) &= Ax(t) + Ke(t) \\ y(t) &= Cx(t) + e(t), \end{aligned} \qquad t = 0, 1, 2, \ ... \tag{5.1}$$

where the $A, K,$ and C matrices are models to be identified, $x(t)$ and $y(t)$ are the state and output vectors, respectively, and $e(t)$ is the input vector (usually temporally white noise). Consider that the one-step prediction of the output at time t is,

$$\hat{y}(t) = Cx(t), \tag{5.2}$$

then the innovations sequence, $e(t) = y(t) - \hat{y}(t)$, can be used to rewrite Eq. 5.1 into its closed-loop form,

$$\begin{aligned} x(t+1) &= [A - KC]x(t) + Ky(t) \\ \hat{y}(t+1) &= C[A - KC]x(t) + CKy(t), \end{aligned} \tag{5.3}$$

where the closed-loop system, $A - KC$, is stable. For a discreet-time state-space system, the stability criterion is achieved when the magnitudes of all eigenvalues are less than unity. This LTI state-space model is the prediction filter, where the output $\hat{y}(t+1)$ is the one-step prediction of $y(t+1)$. To identify the A, K, and C matrices, as well as the covariance matrix of the input, Q_{ee}, the subspace algorithm in [21] computes the Hankel matrix from the QR factorization of the input-output sequence,

$$[U_f \quad U_p \quad Y_f] = [Q_1 \quad Q_2 \quad Q_3] \begin{bmatrix} R_{11} & R_{12} & R_{13} \\ 0 & R_{22} & R_{23} \\ 0 & 0 & R_{33} \end{bmatrix}, \tag{5.4}$$

where U_f, U_p, and Y_f are stacked input and output column vectors, respectively. The size of these input-output sequences is determined by a designed forward/backward sliding window size, $n_{\max}$. Subscripts p and f denote the past and forward finite-number of data in the input-output sliding window. The following least-squares operation yields the Hankel matrix,

$$H = (R_{22} \backslash R_{23}), \tag{5.5}$$

which is subsequently used in the singular value decomposition:

$$H R_{22}^T = U \Sigma V^T. \tag{5.6}$$

The identified system order, i, is designed based on modal reduction of Σ. Using the first i columns of U and V, U_i and V_i, the finite-interval observability and controllability matrices for some positive integers p and q are computed as,

$$\mathcal{O}(p) = U_i \Sigma_i^{1/2} \qquad \mathcal{C}(q) = \Sigma_i^{1/2} V_i^T R^{-T}. \tag{5.7}$$

Finally, the system matrices, A, K, and C are identified as,

$$A = \mathcal{O}(p-1)^\dagger \mathcal{O}(p) \tag{5.8}$$

$$K = \text{first column block of } \mathcal{C}(q) \tag{5.9}$$

$$C = \text{top row block of } \mathcal{O}(p), \tag{5.10}$$

where † represents the Moore-Penrose inverse.

5.2 Predictive Dynamic Digital Holography

5.2.1 Initialization of the Sharpening Algorithm

The optimization routine for sharpness maximization is quite computationally intensive. Even with analytic gradients reducing the number of function evaluations in the optimization routine, reconstruction and gradient-vector calculation using Fresnel propagation involves four 2-D FFTs for each sharpness maximization iteration. Reducing the number of sharpness iterations required for near-optimal wavefront correction becomes of paramount importance in real-time operation. The results presented in Chapter 7 show that the number of sharpening iterations required to reach near-optimal wavefront correction can be reduced significantly by wavefront prediction. Specifically, the initial Zernike coefficients used for sharpness optimization at each time step (i.e., frame) are determined by prediction, based on the time-history of previous optimal coefficients.

Chapter 7 will compare wavefront estimation results obtained with prediction-based sharpening initialization to results obtained with two sharpening initialization methods with no prediction. The two non-prediction methods of initialization used are defined as follows:

$$\begin{aligned} &\underline{\text{Sharpening Initialization } \mathrm{M}_0} \text{ (No Prediction)} \\ a_n(t,0) &= 0, \qquad t = 1, 2, \ldots \\ \hat{\psi}(t,0)(\zeta,\eta) &= 0 \end{aligned} \tag{5.11}$$

$$\begin{aligned} &\underline{\text{Sharpening Initialization } \mathrm{M}_1} \text{ (No Prediction)} \\ a_n(t,0) &= a_n(t-1,k_f), \qquad t = 1, 2, \ldots \\ \hat{\psi}(t,0)(\zeta,\eta) &= \textstyle\sum_n^N a_n(t,0) Z_n(\zeta,\eta) \\ k_f &= \text{final sharpening iteration at time } t-1 \end{aligned} \tag{5.12}$$

Since the M_0 initialization simply sets the initial Zernike coefficients to zero to begin sharpening at each time step, it does not take advantage of temporal correlation in the wavefront sequence to reduce the number of sharpening iterations. Initialization M_1 attempts to exploit temporal correlation in the simplest way by beginning sharpening at each time step t with the final Zernike coefficients from sharpening at time-step $t-1$. For wavefront sequences that share similar modal spatial frequency content across each time-step, Eq. 5.12 can be appropriate. However, issues arise when shifts in the estimated Zernike modes from one time-step to another may cause the sharpness to deviate further away from the near-optimum converged value of the current time-step, or even break down for a large enough variation. Chapter 7 will demonstrate such occurrences arising.

One can argue reasonably that initialization M_1 is a kind of prediction, but it uses no prediction filter or other prediction algorithm, so it does not seem appropriate in the context of this dissertation to refer to initialization M_1 as using prediction. This dissertation proposes the following class of predictive sharpening initializations:

$$\begin{array}{rl} & \underline{\text{Sharpening Initialization } \mathrm{P}_{k_{\mathrm{ID}}}} \text{ (Prediction)} \\ a(t,0) & = \hat{a}(t), \qquad t = 1, 2, \ldots . \\ \hat{\psi}(t,0)(\zeta,\eta) & = \sum_{n=1}^{N} a_n(t,0) Z_n(\zeta,\eta) \end{array} \tag{5.13}$$

Here, the vector $\hat{a}(t)$ is a one-step prediction of the final sharpened Zernike vector $a(t, k_f)$, which is obtained after k_f sharpening iterations at time-step t. The prediction filter in Section 5.2.2 generates $\hat{a}(t)$ from the Zernike vectors $a(\tau, k_f)$ $(\tau = t-1, t-2, \ldots)$. Unlike the previous two non-prediction methods, this filter exploits the temporal and spatial correlations of the wavefront sequence.

The k_{ID} in the notation $\mathrm{P}_{k_{\mathrm{ID}}}$ refers to the number of sharpening iterations used to construct a sample (i.e., training) input wavefront sequence for identification of the prediction filter, as discussed in Section 5.2.3. Generally, $k_{\mathrm{ID}} \neq k_f$. In fact, k_f could vary with time without changing the predictive initialization of the LTI filter. Even though k_f could vary with time, in applications k_f most likely will be set to a constant value determined by real-time computational capacity. For this dissertation, simulation results presented subsequently compare performance obtained with different values of constant k_f. In principle, k_{ID} is not required to be constant, as would be the situation for a fully adaptive filter. Results in Chapter 7 will investigate various constant k_{ID} LTI filters and the strong implications for a feasible fully adaptive filter.

5.2.2 The Prediction Filter

The vector $\hat{a}(t)$ in Sharpness Initialization $\mathrm{P}_{k_{\mathrm{ID}}}$ is generated by a prediction filter that has the state-space form

$$\begin{aligned} x(t+1) &= Ax(t) + K[a(t,k_f) - \hat{a}(t)] \\ \hat{a}(t) &= Cx(t) \end{aligned} \qquad t = 1, 2, \dots, \tag{5.14}$$

or

$$\begin{aligned} x(t+1) &= [A - KC]x(t) + Ka(t,k_f) \\ \hat{a}(t+1) &= C[A - KC]x(t) + CKa(t,k_f) \end{aligned} \qquad t = 1, 2, \dots . \tag{5.15}$$

The matrices $A, K,$ and C are identified using subspace system identification, from a sample wavefront sequence as described in the next section. The vector $x(t)$ is the internal filter state, which usually has no direct physical interpretation.

From Eq. 5.14, the prediction filter can be seen to have the structure of a Kalman predictor, and the subspace system identification method produces matrices $A, K,$ and C that, at least theoretically, minimize the prediction-error covariance averaged over time. However, whereas classical Kalman filter design requires *a priori* knowledge of process dynamics and noise covariances, none of those are required by the system identification scheme here. The subspace system identification scheme requires only a sample data sequence to determine the matrices for the optimal (i.e., minimum-variance) filter. In other words, the identified model captures the statistics of the dynamic wavefront sequence and its corresponding prediction filter delivers minimum-variance prediction error.

5.2.3 Identification of the Prediction Filter

Identification of the matrices $A, K,$ and C in the prediction filter begins with a sample sequence of reconstructed images in the presence of phase aberrations (i.e., wavefronts). These phase aberrations are assumed to have approximately the same spatial and temporal statistics as the wavefront sequence that will be present subsequently when the identified prediction filter is utilized in operation. Digital holography and image sharpening are applied to each image in the sample sequence, yielding a sample sequence of wavefront estimates. This image sharpening for the wavefront training sequence uses initialization M_0 and k_{ID} sharpening iterations for each sample wavefront. The k_{ID} here is the subscript in the notation $\mathrm{P}_{k_{\mathrm{ID}}}$. In general, k_{ID} is different from and independent of k_f, the number of sharpening iterations performed at each time step during image sharpening. Generally, k_{ID} may be time-varying in the construction of a sample wavefront sequence for identification of the prediction filter, but for simulations results presented in this dissertation, k_{ID} is constant during the construction of each sample sequence.

After the sample sequence of wavefront estimates is constructed, the wavefronts are vectorized to form the sample vector sequence $y(t)$. Then a subspace system identification algorithm [21], described in Section 5.1, uses the vectorized sample sequence of wavefronts to identify the $A, K,$ and C matrices in the state-space model,

$$\begin{aligned} x(t+1) &= Ax(t) + Ke(t) \\ y(t) &= Cx(t) + e(t). \end{aligned} \tag{5.16}$$

In this model, $y(t)$ is the sequence of Zernike vectors that parameterize the sample sequence of wavefront estimates, $x(t)$ is an internal state vector and $e(t)$ is a (theoretically) white sequence. The identified $A, K,$ and C matrices then are used in the prediction filter in Eq. 5.14.

Prediction filters of the form used here, identified by subspace system identification, have been used previously for modeling, prediction, and correction of optical wavefronts [22–24, 50]. A significant difference between those previous applications and the problem here is the following: In the previous applications, the sample wavefront sequence used to identify the prediction filter was measured directly by a Shack-Hartmann sensor; here, only the sample detector plane intensity patterns are measured, and the sample wavefronts are estimated by digital holography and image sharpening, as described above. The implications and advantages to removing bulky, expensive electro-optical equipment such as a wavefront sensor should be immediately recognized.

The notation $\mathrm{P}_{k_{\mathrm{ID}}}$ for prediction-based sharpening initialization includes k_{ID} because the prediction filter changes with the value of k_{ID}. Two different values of k_{ID} usually produce different estimated phase profiles from the same sample image, unless the iterative sharpening essentially converges for both values of k_{ID}. Hence, different values of k_{ID} can lead to different sequences $y(t)$ and therefore different $A, K,$ and C matrices for the prediction filter, even though the same sequence of sample images is used to generate the sequence of sample wavefronts. It should be noted that, while the input $a(t, k_f)$ and output $\hat{a}(t+1)$ for the filter change with k_f, the filter itself does not depend on k_f.

5.2.4 Other Considerations

One practical consideration when implementing the prediction filter in digital holography numerical reconstruction is the issue of outliers in the dynamic wavefront sequence. Some wavefronts cause exceedingly corrupted images that the prediction filter (indeed, any sharpness initialization method discussed) still produces poor quality reconstructed images. For the prediction filters, these outliers may potentially pollute the time-history of wavefront estimates with statistics that are not representative of the typical wavefront. In the spirit of real-time application, the primary focus in this work is in numerical reconstruction with very low values of sharpness iterations, k_f. In that sense, these particularly degrading wavefronts are practically unusable, albeit for high, finite-valued k_f, though they may still be sufficiently correctable. It is not difficult to remove these outlier wavefronts in an offline application, given knowledge of the wavefront sequence. Unfortunately, in online operations, it is safe to assume that such knowledge is not available. With that in mind, after reconstructing an image estimate, a normalized relative error norm of the current image and a time-history of image estimates is calculated, i.e.,

$$G(t) = \frac{\|w_\tau(\zeta,\eta) - \hat{u}(\zeta,\eta,0)\|_F}{\|w_\tau(\zeta,\eta)\|_F}, \tag{5.17}$$

where $w_\tau(\zeta,\eta)$ represents the time-history of reconstructed image estimates, $\hat{u}_\tau(\zeta,\eta,0)$ for $\tau = t-1, t-2, \ldots$. If this error norm yields a value below the desired error tolerance, the image estimate is saved in the time-history by means of an exponential smoother, written as:

$$w_{\tau+1}(\zeta,\eta) = \alpha\hat{u}(\zeta,\eta,0) + (1-\alpha)w_\tau(\zeta,\eta), \tag{5.18}$$

where α is an arbitrary weighting index such that $0 \leq \alpha \leq 1$. Values of α closer to zero put less weight on the current image estimate, while values of α closer to 1 put less weight on the exponential time-history of image estimates. If the error norm yields a value above the desired error tolerance, the prediction filter ignores the phase error estimate for this outlier and uses the previous time-step phase error estimate. This real-time filter effectively behaves as a high-pass filter would in post-processing. Using a first-order exponential smoothing model is not the only solution to outlier filtering, and the reader is encouraged to pursue more elegant solutions. For the purposes of demonstrating effectiveness of wavefront prediction in digital holography in this work, the solution presented here is sufficient.

It should be noted that outlier wavefronts are a different problem from an overall change in the wavefront statistics trend during application. While the former was addressed above, the latter can be solved by implementing a fully adaptive filter prediction algorithm rather than an LTI filter as discussed in this chapter. Even though this is not directly investigated in this work, the results of small k_{ID} in the $\text{P}_{k_{\text{ID}}}$ filters reported in Chapter 7 heavily imply that such algorithms can be implemented to the same effect concluded in this study.

5.2.5 The Overall Predictive Digital Holography Scheme

Algorithm 1 summarizes the overall scheme for predictive digital holography, wavefront estiamtion, and image sharpening. The optical beam reflected from a flood-illuminated object is distorted by a sequence of wavefront aberrations before reaching the detector plane. Predictive digital holography and image sharpening process the sequence of detector plane intensity patterns as described in this research to reconstruct wavefront estimates and sharpened images. Sharpening Initialization $\text{P}_{k_{\text{ID}}}$ uses the predicted vector of modal coefficients generated by the prediction filter in the form in Eq. 5.15 to reduce the number of sharpening iterations required to obtain the wavefront estimates. Although the wavefront estimates are represented in this work as linear combinations of Zernike modes, other modes could be used without changing initialization $\text{P}_{k_{\text{ID}}}$, the form of the prediction filter, or the method identifying the prediction filter.

An important aspect of this scheme is the wavefront prediction feedback loop shown in Figure 1.1, and detailed by Algorithm 1. This feedback interaction between wavefront prediction and image sharpening has no corresponding feature in previous applications of wavefront prediction. In previously proposed methods for wavefront prediction, a wavefront sensor measures the wavefront sequence, and a prediction filter (either linear time-invariant or adaptive) uses the measured wavefronts to generate predicted wavefronts; the prediction does not affect the wavefront measurement. Here, only the detector plane images are measured directly, by a CCD or similar recording sensor. The wavefronts are constructed by digital holography and image sharpening, and the prediction accelerates the sharpening and improves the estimates of the wavefronts.

Algorithm 1: Pseudo-Code for Predictive Dynamic Digital Holography

```
Initialize: $x_0 = 0$, $\hat{a}_0 = 0$;
for $t \leftarrow 1$ to $T$ do
    Propagate object to object plane with Eq. 2.5 (or Eq. 2.3);
    Coherently image object and determine hologram with Eq. 2.6 and Eq. 2.11;
    Transform hologram to spectral domain;
    Determine one-step dynamics, $x(t+1)$, and one-step Zernike prediction, $\hat{a}(t+1)$,
      through Eq. 5.15;
    for $k \leftarrow 1$ to $k_f$ do
        if $(S_{opt} - S) < \epsilon_1$ then
            Realize wavefront estimate from estimated Zernike modes as in Eq. 4.7;
            Correct wavefront with wavefront estimate and reconstruct on image
              plane with Eq. 4.5;
            Calculate sharpness index, $S$ as in Eq. 4.6 or Eq. 4.3;
            Calculate analytic gradient, as in Eq. 4.12;
        else
            Break;
        end
    end
    if $G(t) < \epsilon_2$ then
        Update outlier filter using exponential smoothing, as in Eq. 5.18;
    end
end
```

CHAPTER 6

Wavefront Prediction Filter Analysis

This chapter will analyze and investigate various prediction filters that are identified using the subspace identification methods described in Chapter 5. First, an ideal prediction filter is introduced as P_∞. For practical real-time applications, finite k_{ID} are identified as P_4 and P_8. Before running the simulation, an LTI prediction filter of the form Eq. 5.15 is identified offline using a sample sequence of the first 4000 of the 8000 wavefront estimates. Even though the filter is identified offline in this paper, it is not a requirement, and Chapter 7 will demonstrate the viability for an LTI (or adaptive) filter being identified online. Table 6.1 below summarizes the design parameters used to identify these state-space models. These design parameters were chosen based on an optimization routine minimizing the normalized temporal power of the modal prediction error, $\bar{e}_n^2(t)$, that is,

$$\begin{aligned} \underset{i,n_{\max}}{\operatorname{argmin}} \quad & \bar{e}_n^2(t) = \frac{\sum_{t\in T} \|e_n(t)\|^2}{\sum_{t\in T} \|y_n(t)\|^2} \\ \text{subject to} \quad & (A - KC) \text{ stable} \end{aligned} \tag{6.1}$$

where $e_n(t)$ is the prediction error sequence and $y_n(t)$ the wavefront estimate sample sequence for the n^{th} modal coefficients. Specifically, $y_n(t)$ is the estimated Zernike coefficient vector sample sequence parameterizing the aforementioned wavefront sequence. A constraint is applied to ensure the closed-loop system for the identified state-space filter, Eq. 5.15, is stable, such that the magnitude of the eigenvalues of the closed-loop filter are less than unity. All of these filters are modeled as minimum-variance prediction filters, where, theoretically, they should yield white prediction error sequences. Thus, the quality for how well the identified state-space prediction filters match the statistics of the sample wavefront sequence input data is best determined by the whiteness of the associated prediction error sequences.

Table 6.1: Parameters for identified models

Filter	Order, i	Input-Output Size, $n_{\max}$
P_∞	34	5
P_4	56	13
P_8	59	13

6.1 The Ideal Prediction Filter: P_∞

For certain predictive sharpening results presented in this dissertation, the label P_∞ is used. In this case Sharpening Initialization $P_{k_{\mathrm{ID}}}$ (Eq. 5.13) is used, but the sample wavefront sequence fed into identification of the prediction filter is not constructed by wavefront estimates invoked by digital holography or the image sharpening algorithm. Rather, the benchmark wavefronts from Eq. 4.8 is used as the sample sequence $y(t)$ for identifying the prediction filter. When this benchmark prediction filter is used in simulations of predictive digital holography and image sharpening, the sharpening initialization is the following:

Sharpening Initialization P_∞ (Benchmark Prediction)
Same as Sharpening Initialization $P_{k_{\mathrm{ID}}}$ except that
the benchmark prediction filter generates $\hat{a}(t)$.

The benchmark prediction filter is useful in addressing the important question of how the performance of predictive digital holography and image sharpening depends on the number of sharpening iterations used to generate the sample sequence of wavefront estimates for identification of a prediction filter. For each of the true phase aberrations, this construction yields the Zernike coefficient vector for the closest approximation in the subspace spanned by the first N Zernikes. In this manner, it is assumed that *a priori* knowledge of at least the sample wavefront input sequence is known. The subspace identification algorithm then has the best possible knowledge to generate a state-space model of the prediction filter for the statistics of the aero-optical input data.

Of course, the true phase aberrations are not available in real applications, so that the P_∞ case here would not be an option, but the results for this case provide a theoretical limit

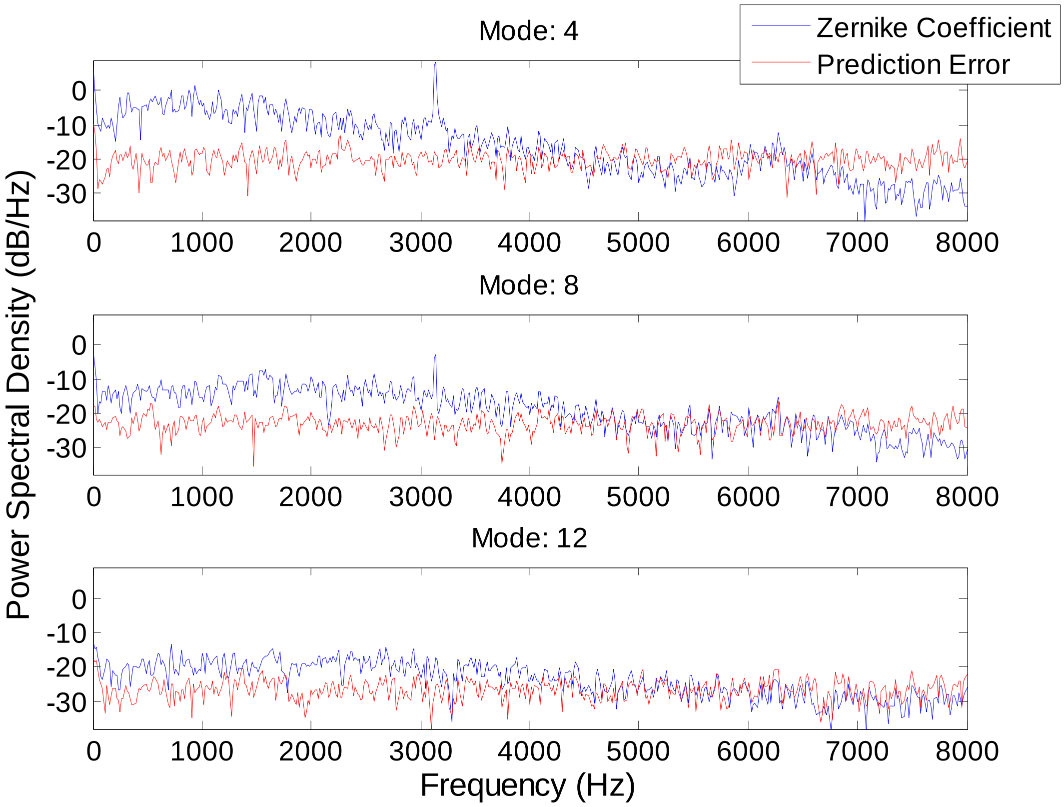

Figure 6.1: Identified P_∞ filter state-space model. Power spectral densities of selected modes and their prediction errors.

for evaluating the performance of predictive image sharpening with realistically identified prediction filters (i.e., with finite k_{ID}). In Figure 6.1, the power spectral densities for selected Zernike modes are shown for an identified P_∞ prediction filter using 15 Zernike modes. Clearly, the bandwidths of the higher modes become increasingly broadband; thus, these modes consequently less identifiable. However, there are still enough identifiable features to capture the statistics at these higher modes, particularly at lower frequencies where the bandwidth is not so broadband. Additionally, while not perfectly white, it is important to observe that the prediction error sequences are nearly white across the spectrum.

It is not unreasonable to assume that when the sample wavefronts are constructed by image sharpening for finite k_{ID}, the constructed wavefronts should approach the projections of the true wavefronts onto the first N Zernikes as k_{ID} increases. That the performance of corrected wavefront aberrations for finite k_{ID} asymptotically approach the performance behavior of the benchmark wavefronts is supported by results presented in Chapter 7. Furthermore, one would expect the characteristics of the finite prediction filters identified by wavefront estimates conceived by finite k_{ID} sharpening iterations to converge to the ideal filter as $k_{\mathrm{ID}} \to \infty$. Figure 6.2 shows the normalized system norm error of various identified state-space models relative to the P_∞ system. It is evident that the relative system error norm begins to asymptotically converge, although not to zero as one might theoretically conclude. Upon closer analysis in the following section, the identified models have particularly good agreement at the lower frequencies, even for the higher modes. Figure 6.1 also indicates that the lower modes have some bands at the higher frequencies that contain useful information. However, the high frequency bands for higher modes are difficult to identify. Increasing the k_{ID} values in the estimated wavefront sample sequence quickly saturates the amount of information that can be identified for lower frequencies and does not alleviate the broadband issue for higher frequencies. From Figure 6.2, one might be quick to discount the P_4 prediction filter; however, it is worth pointing out that the vast majority of the power of the wavefront sequence is contained in the lower modes, where identification does extremely

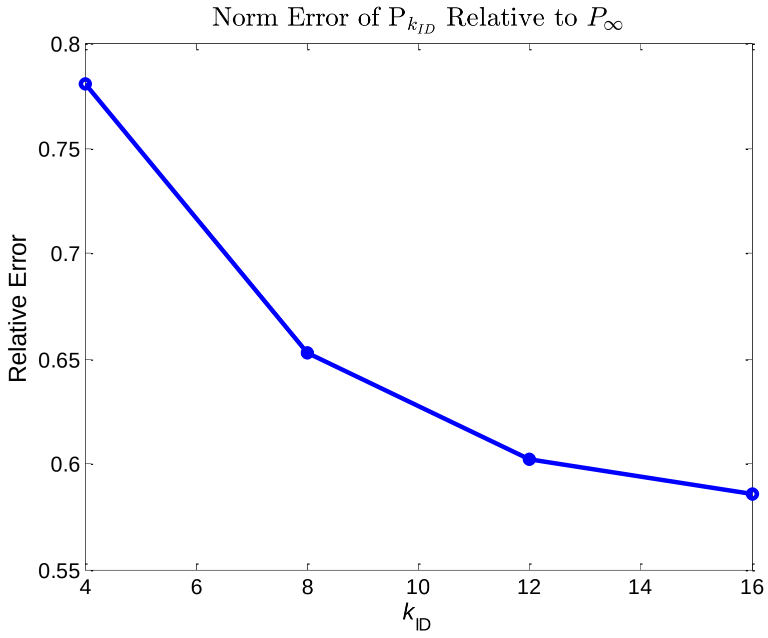

Figure 6.2: Normalized relative system norm errors of identified state-space models for select k_{ID} relative to P_∞.

well, even for the P_4 filter. Chapter 7 will demonstrate the viability of the P_4 filter. First, a closer analysis of the finite prediction filters is provided next.

6.2 Finite k_{ID} Prediction Filters: P_4 and P_8

Since the true wavefronts - and hence the $\tilde{a}(t)$ sequence and P_∞ filter - are not known during operation *a priori*, two additional filters, P_4 and P_8, are identified to address the possibility of generating an LTI prediction filter from limited estimated wavefront knowledge. These prediction filters are identified with finite k_{ID} and thus the sequence $y(t)$ is formed by image sharpening with fixed sharpness iterations of $k_{\mathrm{ID}} = 4$ and $k_{\mathrm{ID}} = 8$, respectively, over the sample sequence. Even in the context of severe wavefront aberration, small k_{ID} were chosen well below the number of iterations required for optimization convergence to analyze the capability of predictive dynamic digital holography in real-time with an LTI or adaptive filter. As in the case with the ideal filter, the sample sequence used to identify these finite prediction filters were the first 4000 of the 8000 modified wavefront sequence described in Chapter 3.

In Figure 6.3, the bode plots of some of the modes for the identified $P_{k_{\mathrm{ID}}}$ 15 state MIMO systems are shown. Each listed mode represents its diagonal entry of the MIMO system, where the majority of the frequency response power for the respective modes are contained. The norm system errors from Figure 6.2 suggests that there is a large discrepancy of low k_{ID} filter approximations relative to the benchmark prediction filter. However, the bode plots demonstrate that the responses of the finite prediction filters agree very well for lower frequencies, even for the higher order modes. Significant departures from the benchmark prediction filter response are located primarily near the Nyquist frequency of the filters. Figure 6.4 provides power spectral density plots of the two finite filters for some of the modal sequences. It demonstratively indicates a drop in power for the higher frequencies for all modes, as was the case for the benchmark prediction filter. Furthermore, the disparagement in the higher modes of the finite filter approximations to $P_{k_{\mathrm{ID}}}$ in Figure 6.3 is of less consequence as the mode number increases. This is due to the fractional power of the wavefront estimates contained in higher order modes rapidly decreasing, as shown in Figure 6.5 and Figure 6.6. Calculating the fractional power in the prediction error sequence as a function of modal order seen in those figures is of a similar form to that of Eq. 4.10,

$$EP(n,T) = \frac{\sum_{t\in T} \|e_n(t)\|^2}{\sum_{t\in T} \|W(t)\|_F^2}. \tag{6.2}$$

Therefore, despite the disagreement in the high frequency and high modal order responses, the finite prediction filters are capable of sufficiently approximating the performance of the ideal infinity filter, as results in Chapter 7 affirm.

Whether infinite or finite, the prediction filters modeled from subspace identification are

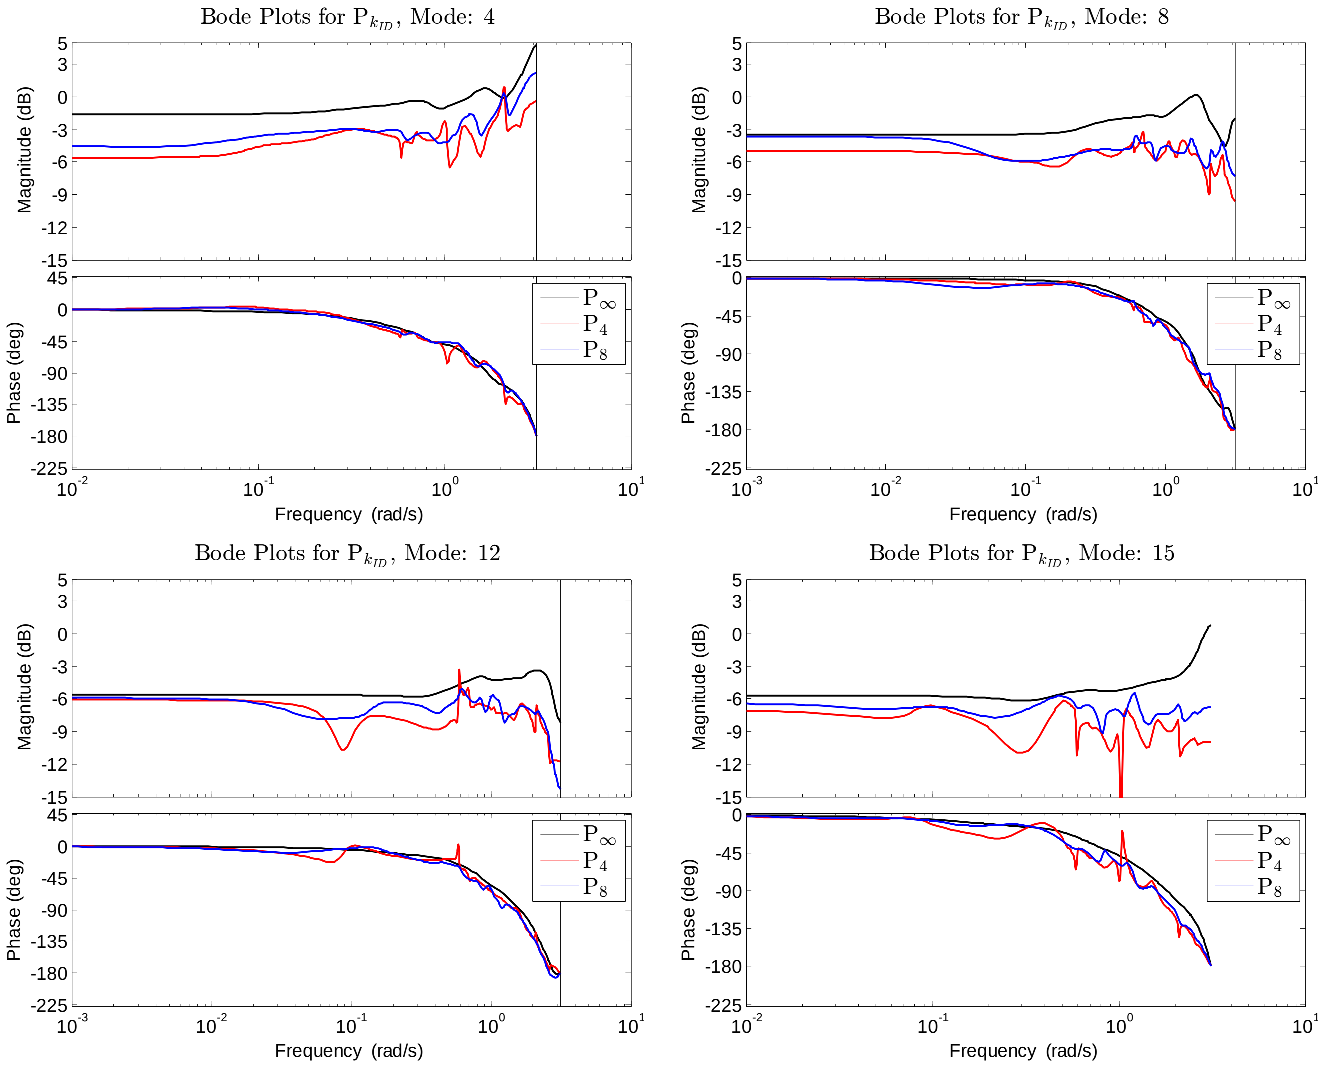

Figure 6.3: Bode plots of various modes for P_4, P_8, and P_∞ prediction filters.

qualified best by how broadband - or white - the bandwidth of the prediction error sequence is. This error sequence is the difference of the estimated (either sharpened or benchmark) wavefront sequence and the predicted wavefront sequence,

$$e(t) = y(t) - \hat{y}(t). \tag{6.3}$$

The closer the error is to theoretical white noise, the higher fidelity in the captured statistics of the sample wavefront sequence. Similar to the analysis of the P_∞ filter, the power spectral density plots of the finite filters show that this is largely the case, albeit imperfect. Finally, the power spectral density plots in Figures 6.1 and 6.4 suggest that it is difficult to identify the state-space models for the higher order modes due to the increasing bandwidth. This is of course the case, as interpreted by the calculated normalized temporal power of the prediction error, $\bar{e}_n^2(t)$ from Eq. 6.1, in Figure 6.7. However, the prediction error of these filters is highest in higher order modes that contain relatively low power, as demonstrated by Figures 6.5 and 6.6.

The following chapter integrates all aspects of this dissertation up to this point, culminating in an in-depth look into the simulation results.

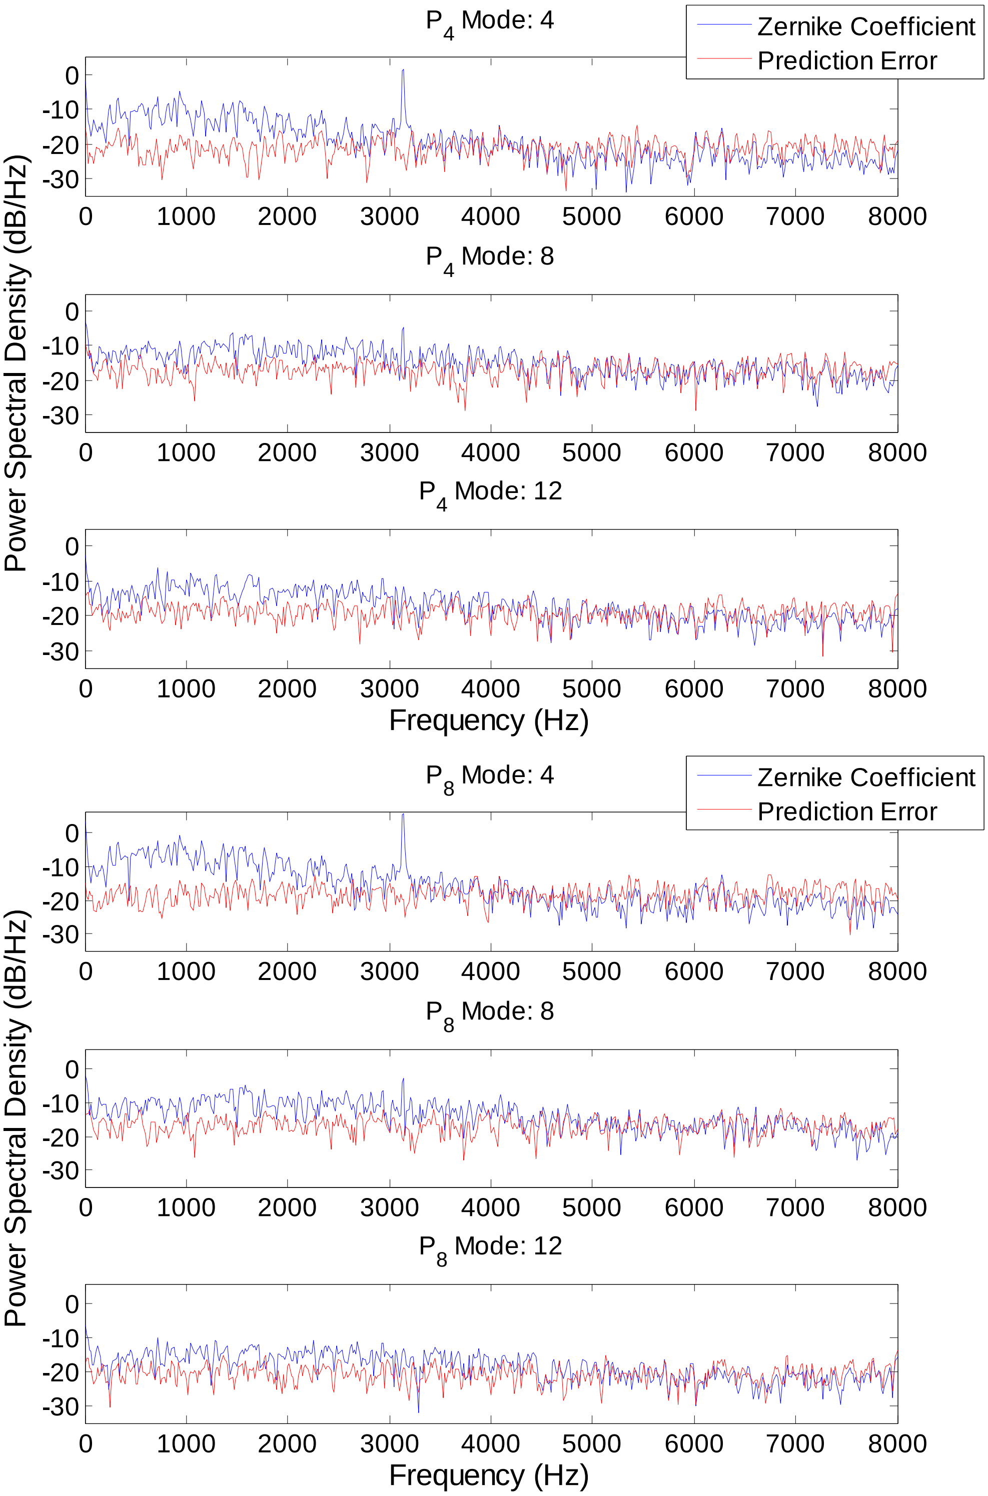

Figure 6.4: Power spectral densities of selected modes and their prediction errors. Identified P_4 (top) and P_8 (bottom) filter state-space models.

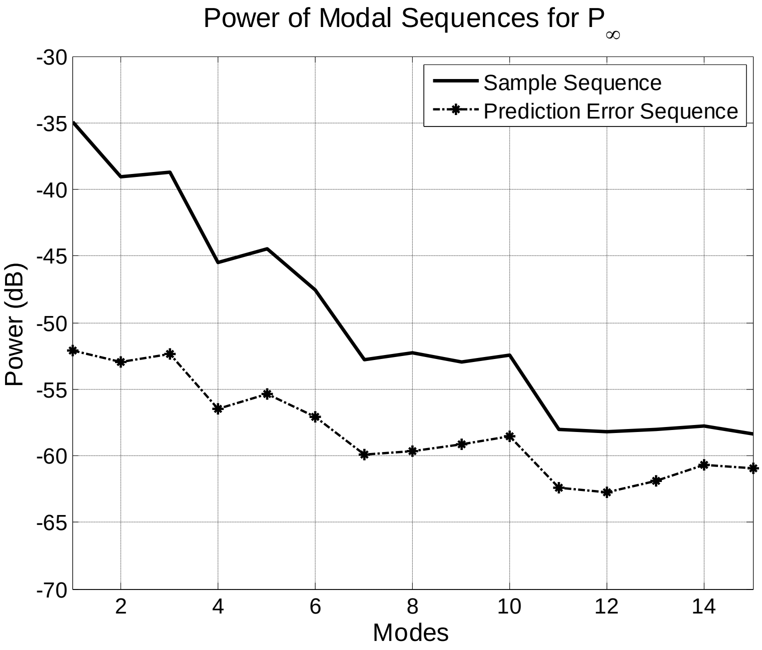

Figure 6.5: Fractional power of modal sequences for the P_∞ prediction filter; sample estimated wavefront Zernike parameterization, $y_n(t)$ (Eq. 4.10), and corresponding prediction error, $e_n(t)$ (Eq. 6.2).

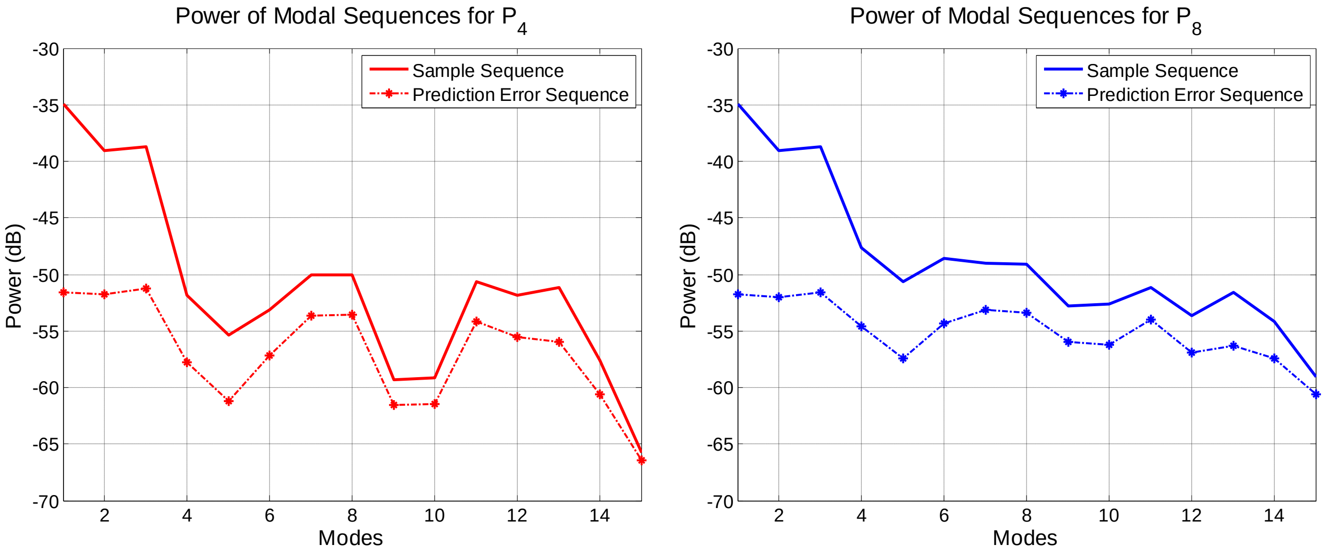

Figure 6.6: Fractional power of modal sequences for the P_4 (left) and P_8 (right) finite prediction filters; sample estimated wavefront Zernike parameterization, $y_n(t)$ (Eq. 4.10), and corresponding prediction error, $e_n(t)$ (Eq. 6.2).

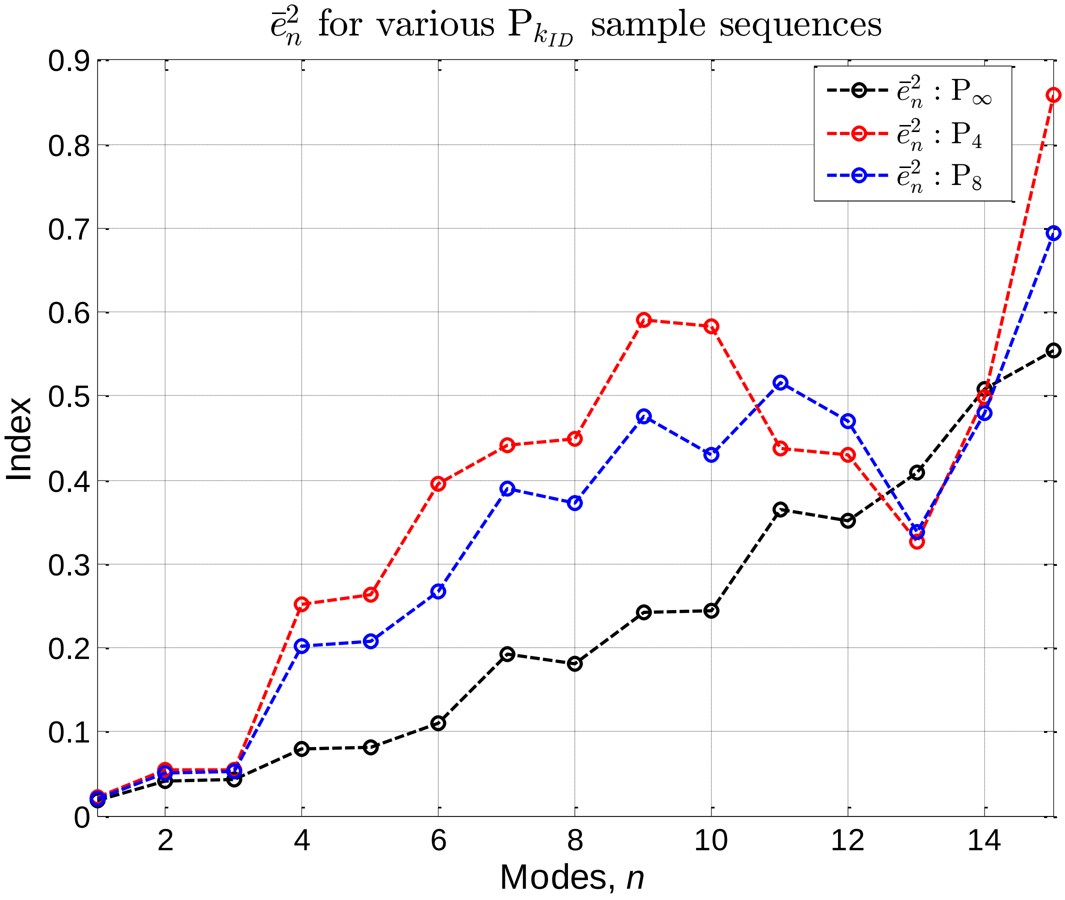

Figure 6.7: Normalized temporal power of modal prediction error sequence for P_4, P_8, and P_∞ prediction filters.

CHAPTER 7

Wavefront Prediction Simulation and Results

This chapter assembles digital holography, image sharpening, and predictive wavefront correction into a simulation. Algorithm 1 describes the overall scheme and the results it yields are presented here.

7.1 Results

For all of the results presented in this section, the first 4000 frames from the sequence of disturbance wavefronts were used to identify prediction filters. The next 2000 disturbance frames then were used to generate the performance evaluation results shown in the subsequent figures. Thus, the prediction filters were identified from one set of disturbance frames and then evaluated on a different set of frames. It is important, of course, that the wavefronts used for identification and those used for evaluation have similar statistics.

In the simulations reported here, Zernike modes 4–15 are estimated and corrected by digital holography and image sharpening. Modes 1–3 (piston and tip-tilt), which cannot be estimated by image sharpening, are corrected with the first three benchmark Zernike coefficients. The phase errors reported are the errors between the wavefront estimates and the benchmark wavefronts defined in Eq. 4.8, i.e., the error in the phase estimates for modes 4–15, since modes 1–3 are corrected exactly.

Several subsequent figures show pupil plane phase error power. For each frame, phase error power is the mean-square phase error over the frame. In most plots, the phase error values shown are time averages over the 2000 evaluation frames. The phase error values are given as % Error because the phase error power is normalized by the time-averaged power in the benchmark modes 4–15.

7.1.1 Sharpening Initializations M_0 and M_1 vs. P_∞

Figures 7.1–7.4 compare results obtained with the non-predictive sharpening initializations M_0 and M_1 and results obtained with the benchmark sharpening initialization P_∞ (which uses the benchmark prediction filter). The sharpness metric S_1 was used for Figs. 7.1–7.3, while the sharpness metric S_2 was used for Fig. 7.4.

In Fig. 7.1, the top plot presents the time-averaged pupil plane phase error power as a function of the number of sharpening iterations (i.e., as a function of k_f). The bottom plots in Fig. 7.1 show how the phase error power varies with time over the 2000 evaluation frames. For P_∞ (benchmark prediction) and 10 sharpening iterations, the time-averaged phase error power is 14.4%. Both the non-prediction methods, M_0 and M_1, perform nearly indistinguishably from each other. The primary reason for M_1 not performing better than M_0, and indeed at times under-performing, is that initializing time-step $t+1$ with the final Zernike coefficient vector at time t does not necessarily start the optimization at time-step $t+1$ in the correct direction. This increases the number of iterations needed to correct the wavefront error. Even after $k_f = 10$ sharpening iterations, M_0 and M_1 have not reached the level of correctness as the prediction method, P_∞ at $k_f = 4$ iterations. In fact, the sharpness initializations M_0 and M_1 have yet to reach a converging state by $k_f = 10$ iterations, exceeding 5% error of the near-optimal wavefront power.

Observation of the time series in the left bottom plot in Fig. 7.1 shows that, with initialization M_0 and $k_f = 4$ sharpening iterations, a non-trivial number of frames have large aberrations, with errors surpassing 100% error. Meanwhile, the same time-series under P_∞ correction is robust to these damaging pupil plane phase fields. Special attention should be brought to the early stages of the optimization, where P_∞ starts wavefront correction within only 15% power of near-optimal correction. Consequently, the prediction filter arrives at the final-value much sooner than its counterparts, in particular, at $k_f = 6$ sharpness iterations when the phase error power is within 3% of its final value. This suggests that the prediction filter is capable of reducing the number of sharpness iterations needed to reach near-optimal wavefront correction, by exploiting temporal correlation in the wavefront sequence to inject

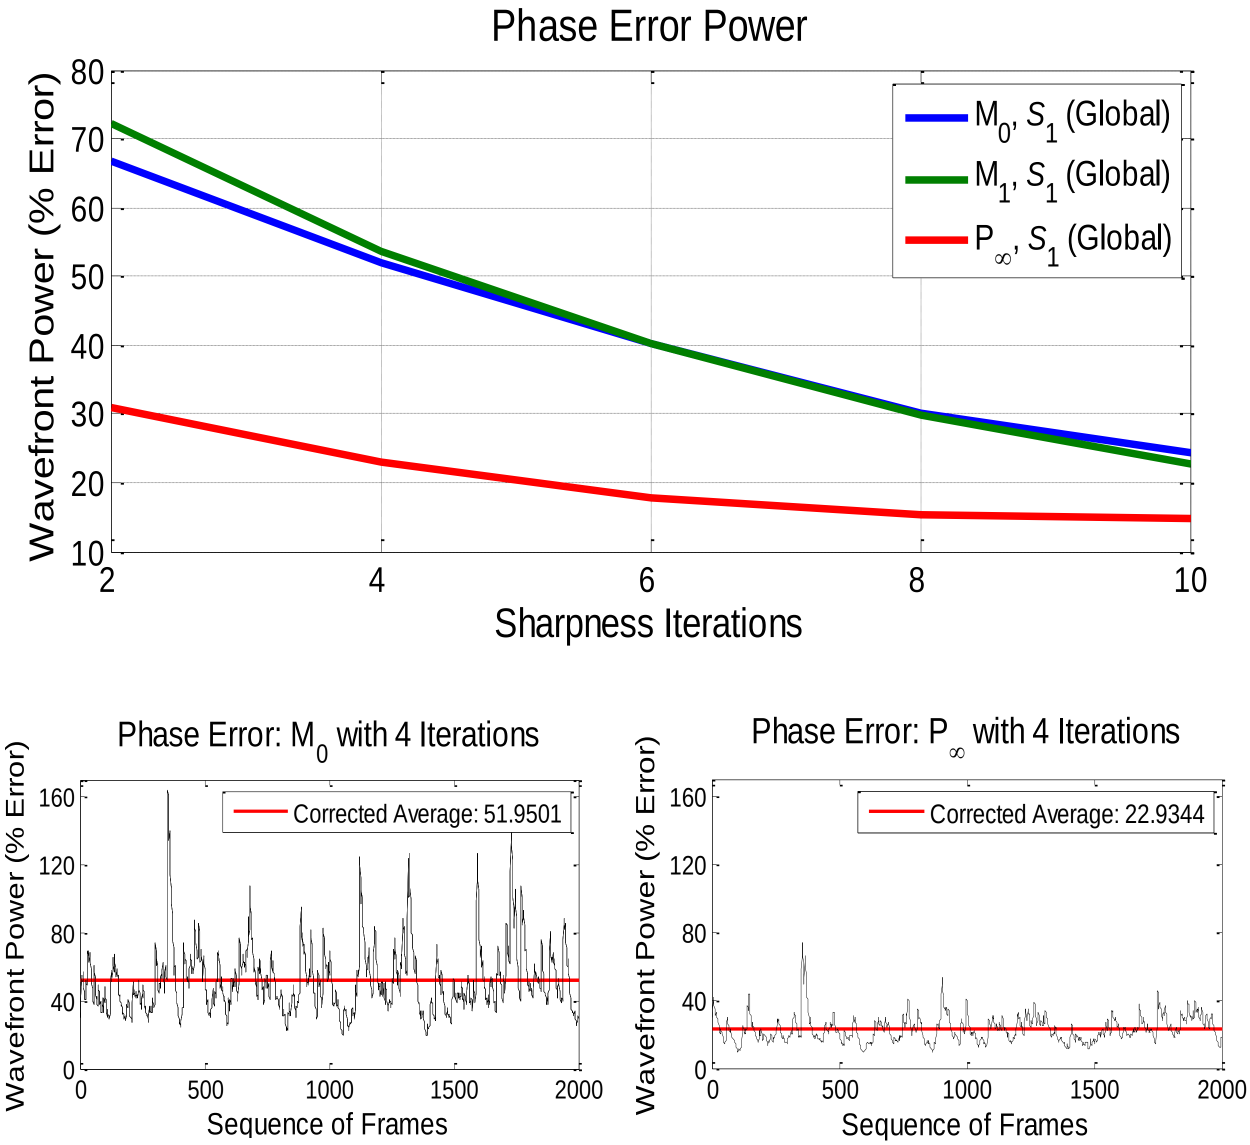

Figure 7.1: Top plot: Time-series average pupil plane phase error power for M_0, M_1, and P_∞ methods versus number of sharpness iterations using the S_1 metric. Bottom plots: Time-series of phase error power after four sharpness iterations for (left) M_0 and (right) P_∞.

a significant increase in computational efficiency.

Figure 7.2 shows a histogram of the number of sharpness iterations per frame required to achieve within 0.01% convergence of the final-value pupil plane power using the M_0, M_1, and P_∞ methods. As expected, M_0 and M_1 have similar results, with the higher number of iterations needed for M_1 convergence suggesting that the approach tends to start the optimization routine in adverse directions, relative to the zero-initialization of M_0, more often than not. Keeping the perspective of an application operating in real-time in mind, non-prediction algorithms arrived at near-optimal correction by $k_f = 10$ iterations in less than 1% of frames for the entire sequence, compared to 60% of the wavefronts utilizing a P_∞ algorithm. The required number of iterations exceeded 30 for over 57% of the frames

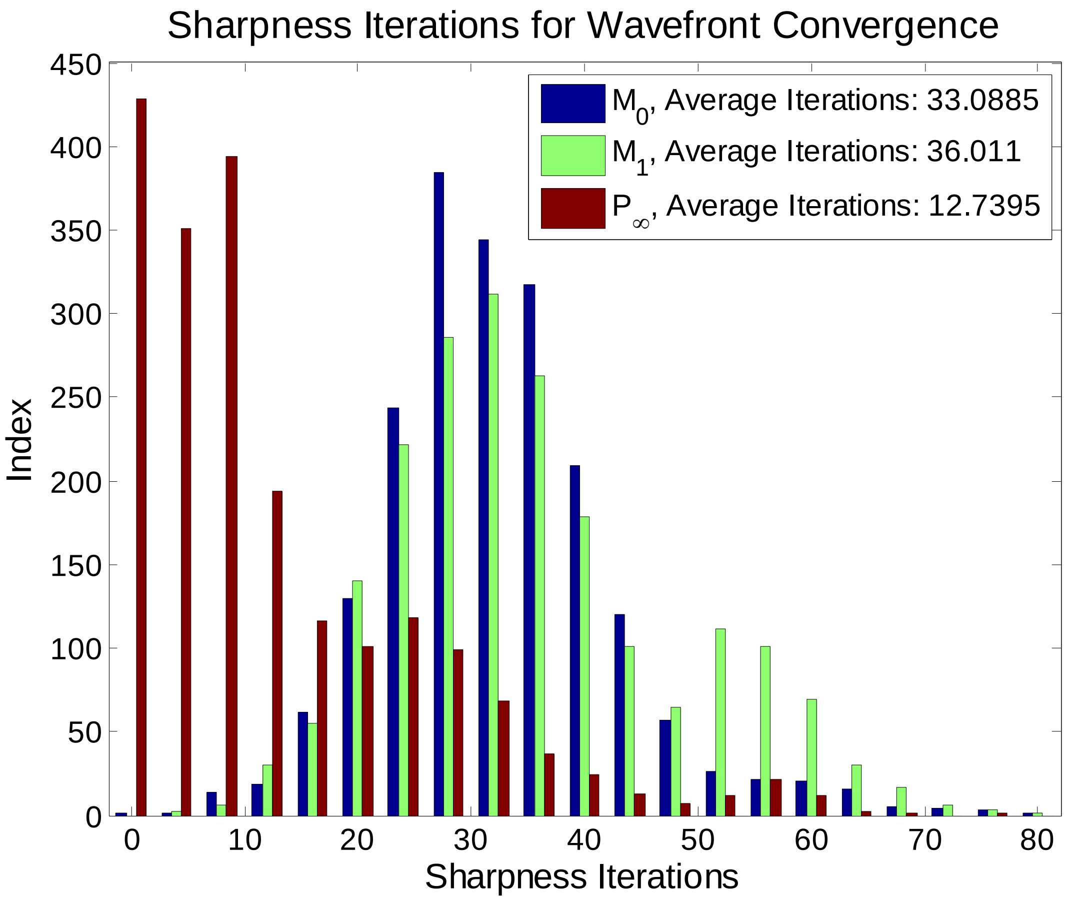

Figure 7.2: Histogram of the number of sharpness iterations required for wavefront correction convergence for the M_0, M_1, and P_∞ methods.

under a non-prediction approach, whereas over 90% of frames have already converged by that point using a prediction approach. This is unsurprising considering the average wavefront power for only P_∞ seems to be converging by $k_f = 8$ sharpness iterations in Fig. 7.1. A particularly interesting result in Figure 7.2 is that 21% of the frames have immediately achieved sufficient wavefront correction after the one-step predict ahead with P_∞, requiring no sharpness iterations to achieve its near-optimal value.

Figure 7.3 shows sample irradiance images of a 60 × 60 pixels ROI after k sharpness iterations and $k_f = 8$ using the M_0 and P_∞ initializations for global sharpening and P_∞ under local sharpening. The average uncorrected RMS image irradiance relative error for the 2000 frame wavefront sequence is 46.3%, thus the sample frame shown is a particularly corrupted case-study, as indicated by $k = 0$ under M_0. The figure demonstrates the prediction filter produces a clearer image and, as expected, local sharpening further improves this particular region of the full image. Recall that unlike the prediction filters, M_0 does not use time-history wavefront correction information. In the prediction cases, the $k = 0$ frame is already near or better than the image quality at $k_f = 8$ for M_0. This amount of initial performance is what allows predictive holographic reconstruction to circumvent much of the computational burden. By $k = 2$ iterations, the generated images from the prediction filters are better than those after $k_f = 8$ iterations for M_0. This snapshot of the dynamic wavefront correction sequence is indicative of the performance gap of the different methods for fixed k iterations on average.

In Figure 7.4, time-series average pupil plane phase error power versus sharpness iterations is reported for the Shannon entropy metric S_2, for sharpening initializations M_0, M_1, and P_∞. Procedures M_0 and P_∞ generate results fairly similar overall to sharpness under the S_1 metric, with comparable convergence behaviors; though, S_2 does not perform as well, particularly in the early stages of correction where there is noticeably worse estimated wavefront sensing. It is worth pointing out the detail that the P_∞ prediction filter is largely insensitive to the change in sharpness metric in these results. For dynamic wavefront sequences with turbulent spatial frequencies from one time-step to another, as in this simulation, M_1 does poorly. At $k_f = 8$ iterations, the M_1 sharpness initialization deviates

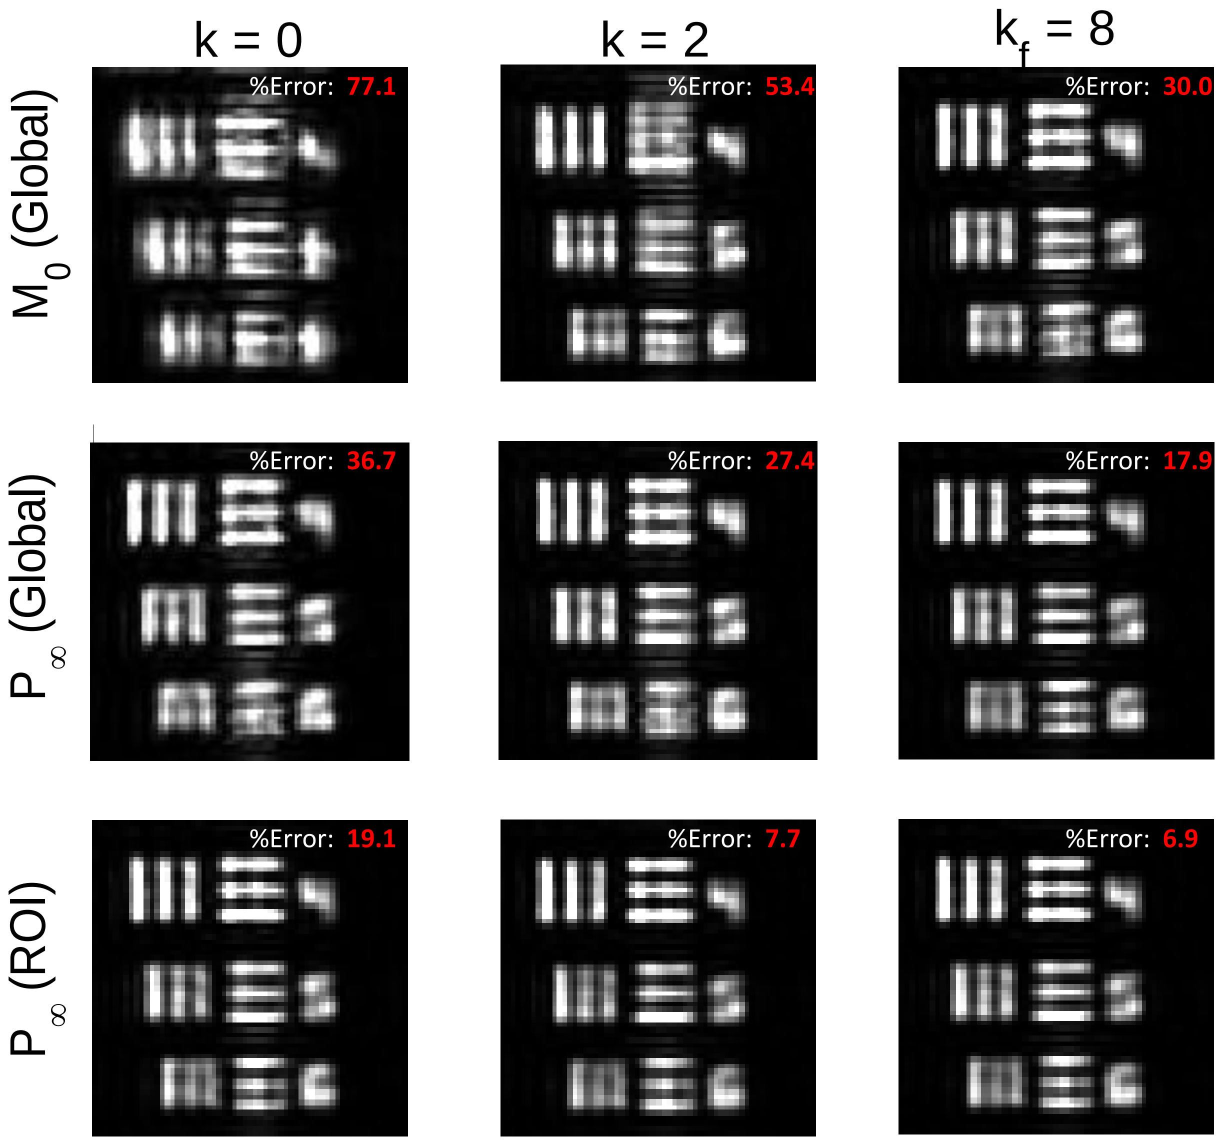

Figure 7.3: Specific irradiance images of a ROI. Each column is a sample image after k sharpness iterations and $k_f = 8$ using (top row) global sharpening for M_0, (middle row) global sharpening for P_∞, and (bottom row) local sharpening for P_∞.

from a nondecreasing shape. This occurs for several reasons: First, it should be noted that while sharpness maximization necessarily minimizes pupil plane wavefront error, as in the consistent M_0 algorithm, M_1 may nullify this property if it begins the sharpness routine in a sufficiently incorrect direction. That is, sharpening a frame at time-step t after $k_f = 6$ iterations initializes the $t + 1$ time-step far enough from its optimal value that $k_f = 8$ iterations are unable to mitigate the poor initialization choice in that number of fixed iterations. This occurs for a sufficient amount of frames to skew the average from its nominal nondecreasing shape. Normal behavior resumes for high enough k_f sharpness iterations. Second, the S_2 metric bears some responsibility, as it includes the intensity at every pixel of the image rather than only 25% of the brightest pixels, as in the S_1 metric. The consequence of this is that image errors from severe wavefront aberrations will be considered at every pixel, including

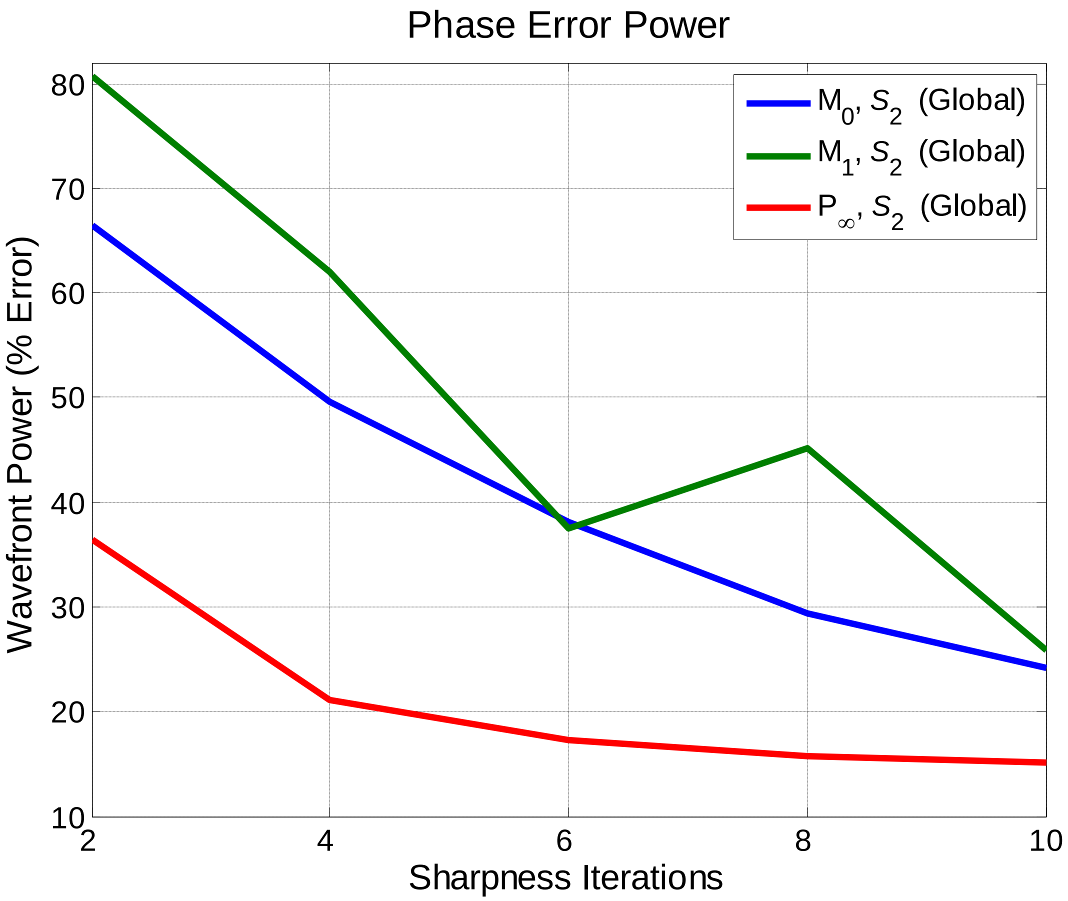

Figure 7.4: Time-series average pupil plane phase error power for M_0, M_1, and P_∞ methods versus number of sharpness iterations using the S_2 metric.

dark pixels that receive dispersed light, which may lead to sharp gradient changes. This is especially detrimental to M_1 as it is the most sensitive to significant gradient changes. Choosing a sharpness metric depends on knowing the nature of the target image, with some metrics performing better than others for particular images, as Fig. 7.4 demonstrates.

7.1.2 Sharpening Initializations P_4 and P_8 vs. P_∞, M_0 and M_1

In Fig. 7.5, the top plot shows the time-averaged pupil-plane phase error power as a function of the number of sharpening iterations for three prediction filters: P_4, P_8 and P_∞. As in previous phase error power plots, each data point in Fig. 7.5 is an average of the phase error power over the 2000 evaluation frames, with the phase correction for each frame determined with k_f sharpness iterations. The time series of frame-by-frame phase error power are shown for P_4 and P_8 with four sharpening iterations (i.e., $k_f = 4$) for each evaluation frame.

As discussed in Section 5.2.3, the prediction filters P_4 and P_8 are identified from the sample sequence of 4000 wavefront estimates reconstructed from detector plane holograms by image sharpening with $k_{\mathrm{ID}} = 4$ and $k_{\mathrm{ID}} = 8$, respectively. Thus, P_4 and P_8 are determined by a method that can be used in real applications.

The key take-aways from these results are described as follows. First, even though the P_4 prediction filter is identified with only $k_{\mathrm{ID}} = 4$ sharpness iterations, it still outperforms the non-prediction sharpness initialization M_0, improving wavefront estimation for a low number of sharpness iterations. However, P_4 is over 5% of the near-optimal corrected wavefront power by $k_f = 10$ sharpness iterations, indicating it has not reached a converging state as quickly as one might desire.

Additionally, the sample time-series indicates that P_4 exhibits several frames that were poorly corrected, suggesting a higher k_{ID} may be required. This leads to a second important point: Obviously, there is no benefit to a predictive approach if the number of finite k_{ID} iterations required to approximate P_∞ is large enough to be computationally disadvantageous. The prediction filter, P_8, in Fig. 7.5 shows that using wavefront estimates from filters using $k_{\mathrm{ID}} = 8$ iterations for identification generates a filter that performs nearly as well as the

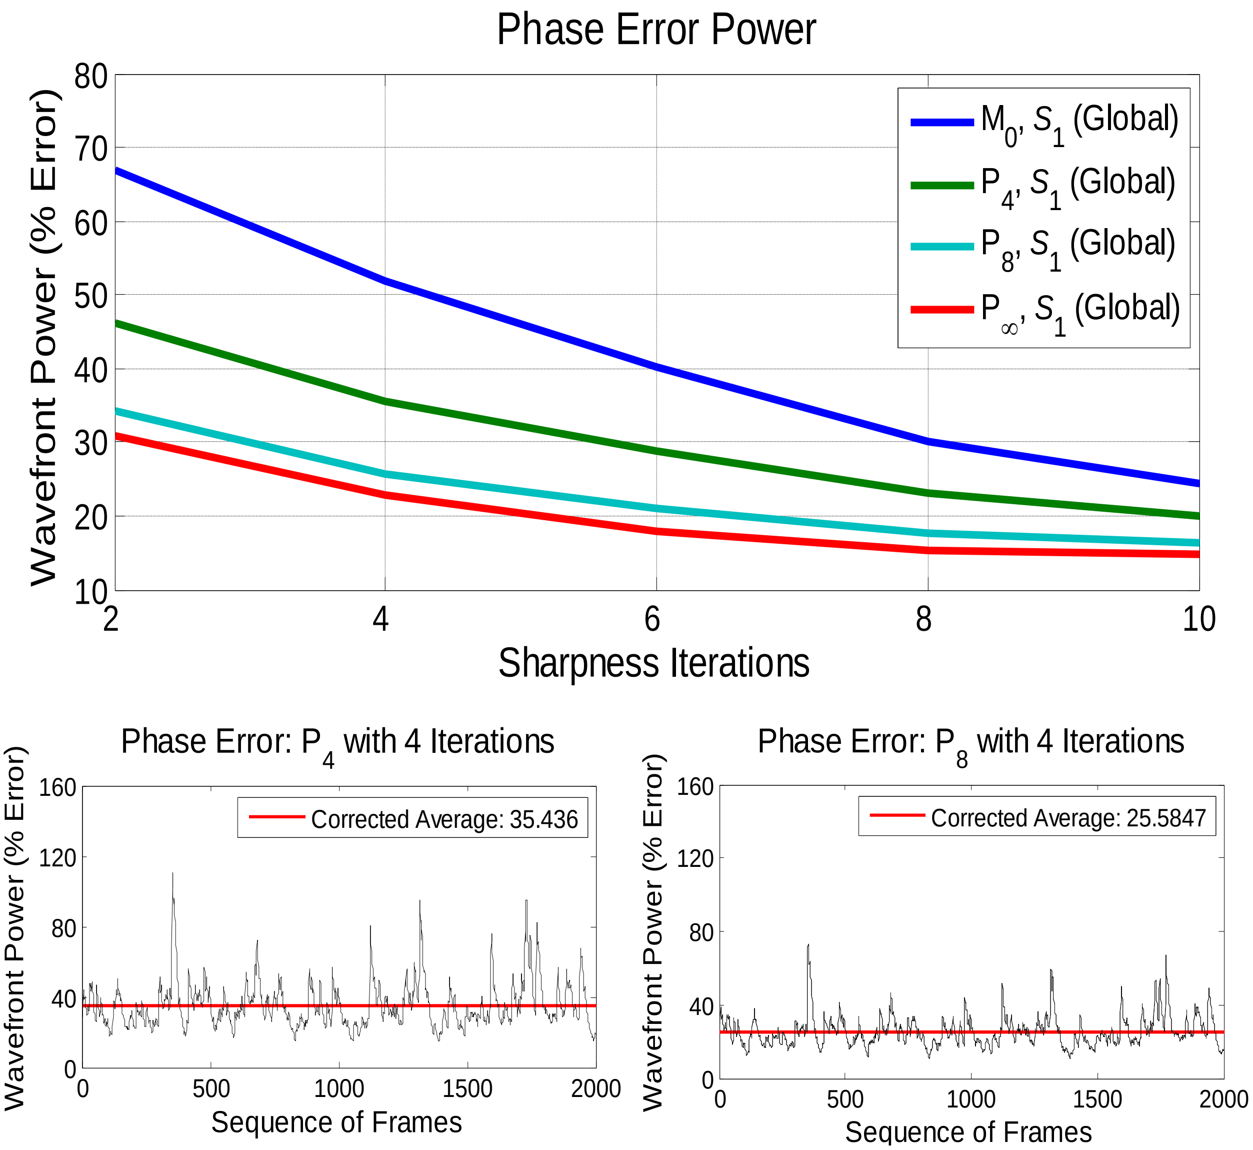

Figure 7.5: Top plot: Time-series average pupil plane phase error power for M_0, P_4, P_8, and P_∞ methods versus number of sharpness iterations using the S_1 metric. Bottom plots: Time-series of phase error power after four sharpness iterations for (left) P_4 and (right) P_8.

ideal case of having complete knowledge of the wavefronts (i.e., P_∞). This has important positive implications for future work involving the integration of a fully adaptive predictive algorithm during online digital holography wavefront sensing.

Figure 7.6 displays a histogram of the number of sharpness iterations per frame required to achieve within 0.01% of the final-value pupil-plane power using the prediction filters P_4, P_8, and P_∞. Compared to the non-prediction methods from Fig. 7.2, the average number of iterations is significantly improved even with low k_{ID} values. When examining the results from Fig. 7.5 and Fig. 7.6, it is important to remember the wavefront aberration sequence is an extremely scaled up degraded version of real-world aero-optical data. The significance of this is that P_8, and indeed even lower k_{ID}, may be highly accurate approximations of

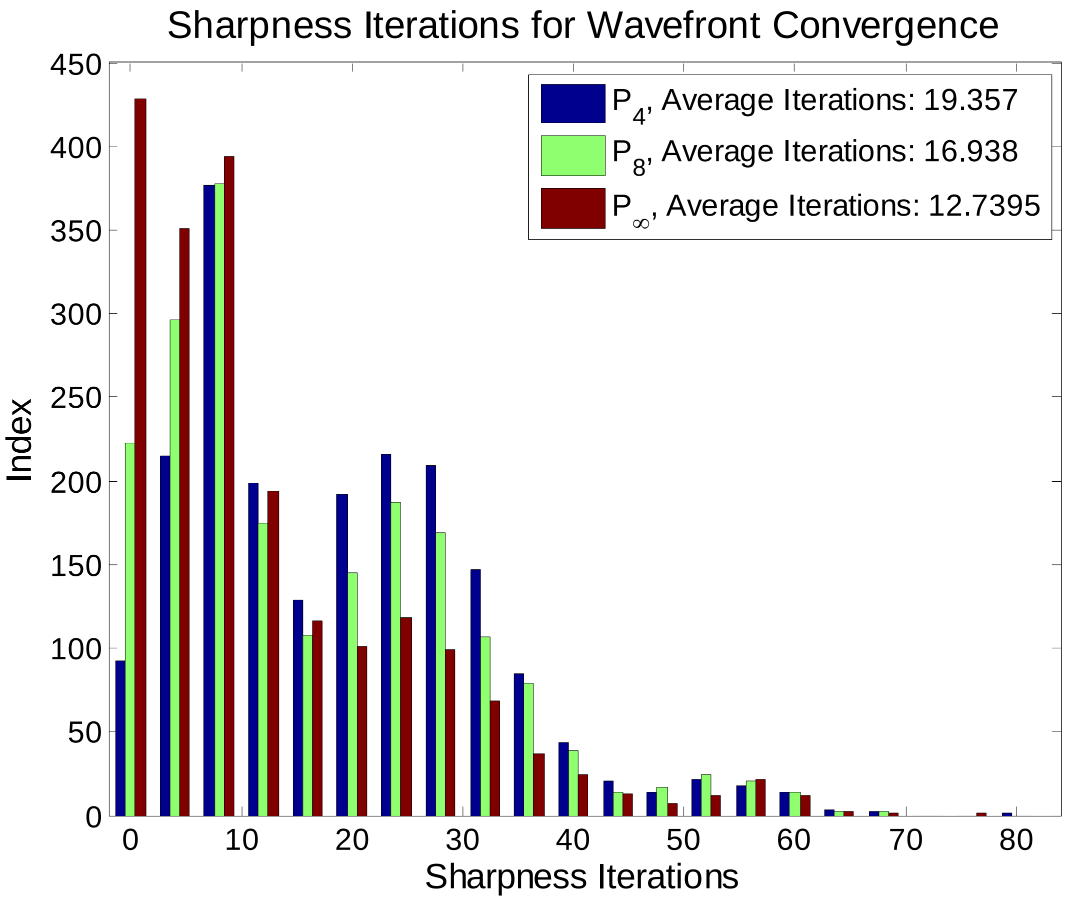

Figure 7.6: Histogram of the number of sharpness iterations required for wavefront correction convergence for the P_4, P_8, and P_∞ methods.

the P_∞ theoretical limit. For P_8, about 45% of the frames in the wavefront sequence reach near-optimality before 10 iterations, with 12% immediately converging after the propagation step in the prediction filter. Even though P_4 is roughly 46% slower than the ideal scenario, it is a 74% improvement over M_0. Clearly then, identifying the LTI prediction filter does not need to be run for an extended number of k_{ID} to experience computational improvement. Regarding a fully adaptive predictive algorithm, the implication is that digital holography can be executed by running a prediction filter that is identified in real-time with remarkably lower k_{ID} than the time it takes for M_0 to converge.

7.1.3 Local Sharpening vs. Global Sharpening

Figure 7.7 evaluates the effectiveness of the methods in this paper for sharpening the 60×60 pixel ROI shown in Figs. 1.2 and 7.3. Since sharpening a relatively small ROI usually does not minimize the wavefront error, performance evaluation for ROI sharpening is based on ROI image error. This image error is defined as the error between a final sharpened ROI image and the ROI image obtained by correcting the wavefront disturbance with the benchmark wavefronts. For each of the 2000 evaluation frames, the ROI image error is the pixel-by-pixel difference between the ROI irradiance of the sharpened image and the ROI irradiance with the benchmark wavefront correction. The measure of ROI image error plotted in Fig. 7.7 is the RMS value of the ROI image error computed over the 60×60 set of ROI pixels and the 2000 evaluation frames, normalized by the RMS value of the irradiance for the ROI pixels in the 2000 frames with the benchmark wavefront correction.

Several insights can be taken from these results. First, as expected, local sharpening with the prediction filter P_8 improves image quality over global sharpening for the ROI. As in the phase plots in Fig. 7.5, the image irradiance for P_8 using global sharpening enters a converging state by $k_f = 8$ sharpness iterations. Local sharpening meanwhile begins its optimization near the image reconstructed by the minimizing global wavefront correction and quickly converges to the minimizing wavefront correction associated with the ROI in as quickly as $k_f = 4$ iterations. Second, an interesting result occurs for local sharpening with M_0 and M_1

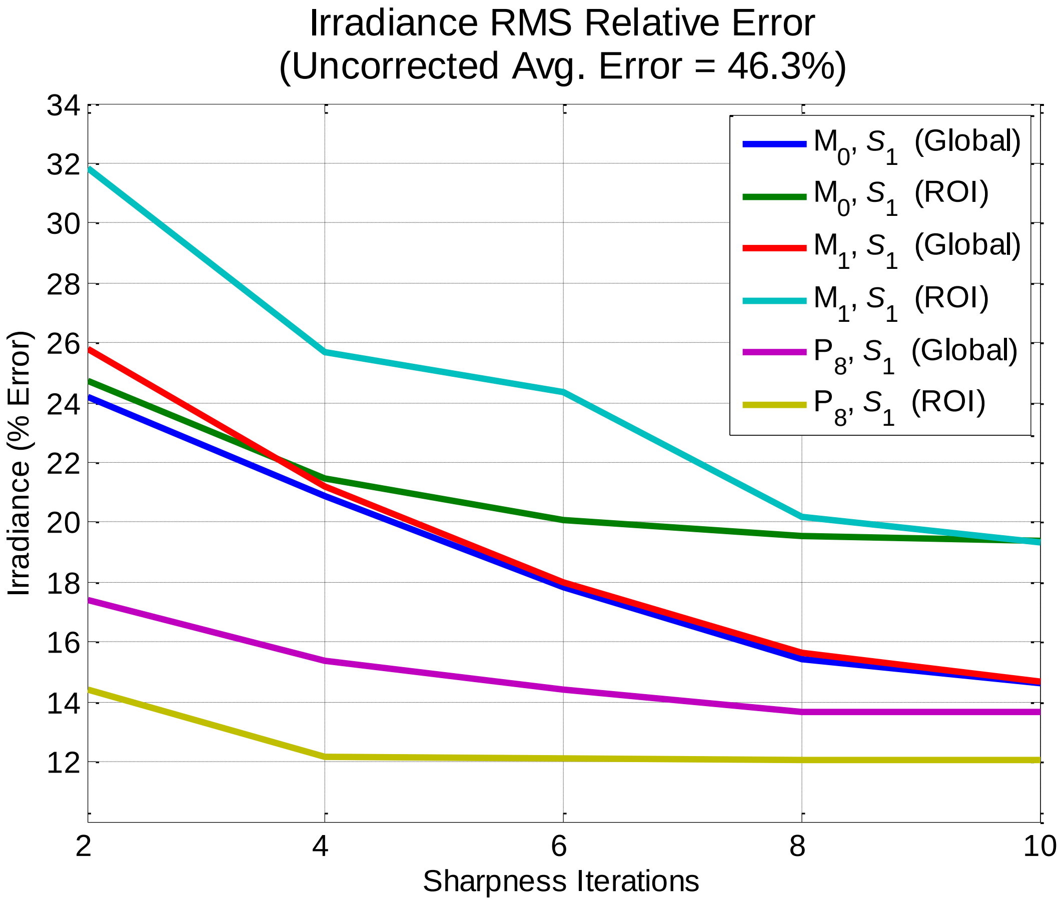

Figure 7.7: Time-series average irradiance RMS error for M_0, M_1, and P_8 methods versus number of sharpness iterations using global and local sharpening over an ROI.

in that they appear to be generating images of the ROI with worse quality than a global sharpening approach. This seemingly contradictory data is revealed to be because of 15% of the wavefront frames with severe aberration having irradiance RMS errors exceeding 100%, even after some fixed k_f iterative corrections. The term "outliers" is purposefully refrained from being used as virtually all of these frames are eventually corrected appropriately by the global variants of the non-prediction schemes, let alone any of the prediction variants. As such, this issue is evidently confined to ROI sharpening for non-prediction methods.

Optimizing an ROI has unique artifacts that are not a concern for a global image approach, including stray light from outside the region, loss of light due to dispersion, and the appearance of localized tilt for the subset of wavefront pixels pertaining to the ROI. These artifacts are greatly amplified by the increased severity of the wavefront disturbance sequence used in the simulation. The worst of these frames constitute the culprits that skew the time-series average value. Since the artifacts are inherent to the phase errors in a global sharpening approach, minimizing the wavefront error will correct them at no additional cost—hence the smooth, continuously decreasing behavior of the global M_0 and M_1 curves in spite of these worst frames. What is especially interesting is that unlike the non-prediction counterparts, the local sharpening prediction algorithm seems robust to these damaging frames, as both the time-series and their corresponding averages displayed in Fig. 7.7 show those frames are corrected without issue, to the point where it leads to the best image quality results. Similar robust behavior was hinted at in the time-series in Fig. 7.1.

Despite the artifacts present in the ROI with severely aberrated frames, the prediction filter mitigates much of the artifact existence that the non-prediction methods fail to correct. It does so by initializing the Zernike coefficient vector based on the temporal and spatial correlation of the wavefront from the past history, which involves more typical wavefront disturbances. When these frames are removed universally from all time-series, the local sharpening for M_0 and M_1 meet expected behavior for ROI sharpening, reconstructing images with better clarity than corresponding global variants. Albeit, those results are still outperformed by even the global sharpening version of P_8.

7.1.4 Strehl Ratios of $\mathbf{P}_8$ vs. $\mathbf{P}_\infty$ and $\mathbf{M}_0$

The Strehl ratios for the corrected wavefront sequence under the M_0, P_8, and P_∞ filter approaches are given in Figures 7.8, 7.9, and 7.10. From Figure 3.4, the 2000 frame wavefront sequence is initially far below the Rayleigh limit of 0.82, hovering around 0.275. After sharpness optimization with $k_f = 10$ iterations, the M_0 method is still considerably below the Rayleigh limit and would be considered unacceptable even for most amateur optics applications. On the other hand, both prediction filters have virtually reached the Rayleigh limit, with several individual frames corrected well beyond the diffraction criterion. Just over 60% and 70% of corrected wavefronts for the P_8 and P_∞ filters, respectively, are above

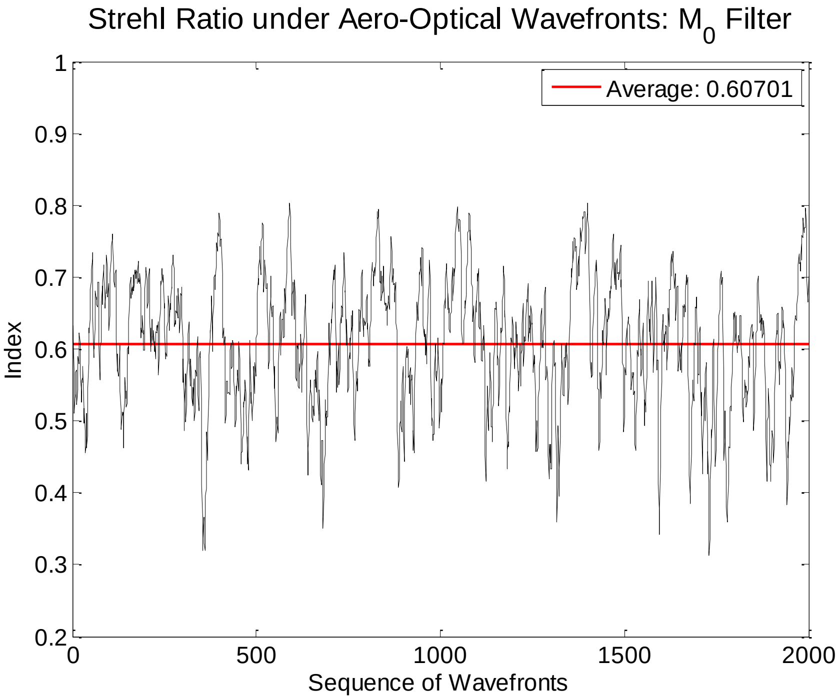

Figure 7.8: Strehl ratio for corrected wavefront sequence using M_0 filter after $k_f = 10$ sharpness iterations.

this criterion. Recall that the Rayleigh criterion does not symbolize the pinnacle of wavefront correction; it merely indicates the point at which diffraction overtakes atmospheric turbulence as the primary source of phase error. This is why a great deal of frames under the predictive filters are corrected well beyond the Rayleigh criterion.

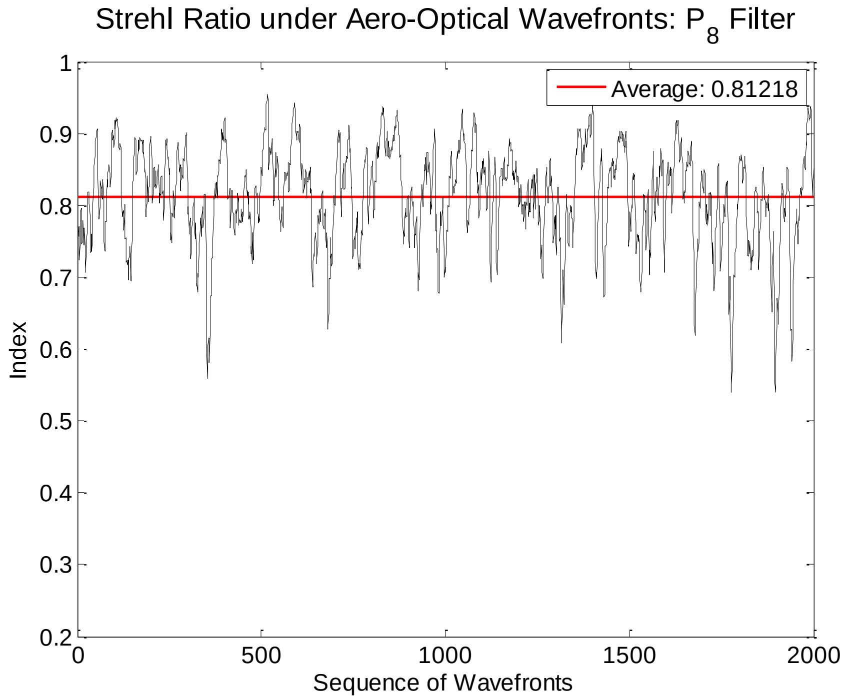

Figure 7.9: Strehl ratio for corrected wavefront sequence using P_8 filter after $k_f = 10$ sharpness iterations.

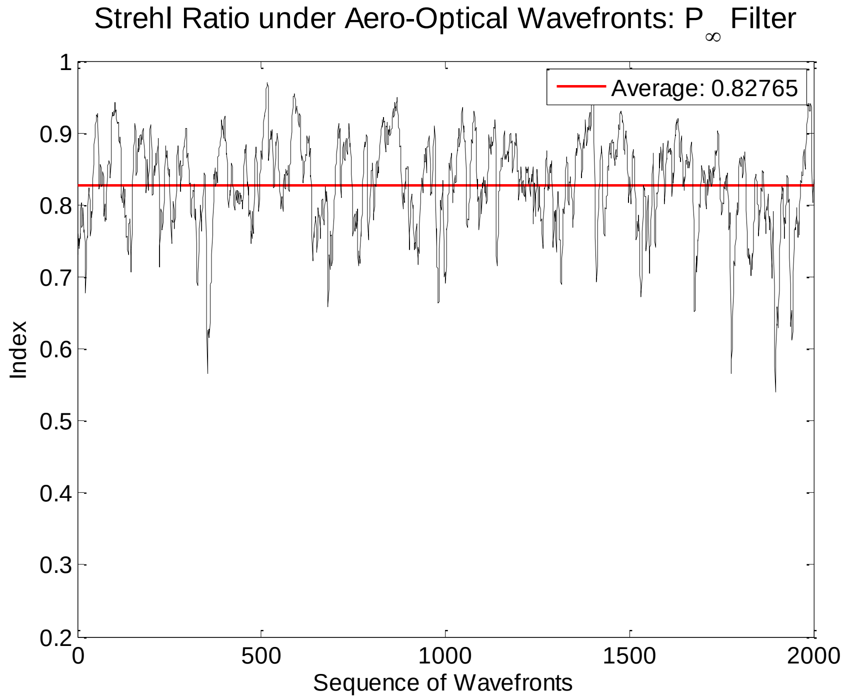

Figure 7.10: Strehl ratio for corrected wavefront sequence using P_∞ filter after $k_f = 10$ sharpness iterations.

7.2 Predictive Dynamic Digital Holography in the Far-Field

When the propagation distance onto the object plane is far enough such that the incoming complex field is planar (or closely planar), then the propagation is in the far-field regime. In this case, the Fresnel diffraction integral simplifies to Fraunhofer diffraction, written as,

$$u(\zeta, \eta, z) = \frac{\exp(j\tilde{k}z)}{j\lambda z} \exp\big(j\frac{\tilde{k}}{2z}(\zeta^2 + \eta^2)\big)\mathcal{F}\{u(\zeta, \eta, 0)\}. \tag{7.1}$$

Quickly, one notices that Eq. 7.1 is simply proportional to the 2-D Fourier transform. Figure 7.11 shows results of wavefront correction in the far-field regime under Fraunhofer diffraction. The plot represents time-series average pupil plane phase error power for M_0, P_8, and P_∞. Similar behavior for the various filters exist as they did for wavefront estimation in the near-field regime. This demonstrates predictive wavefront estimation for digital holography is not restricted to a particular diffraction region.

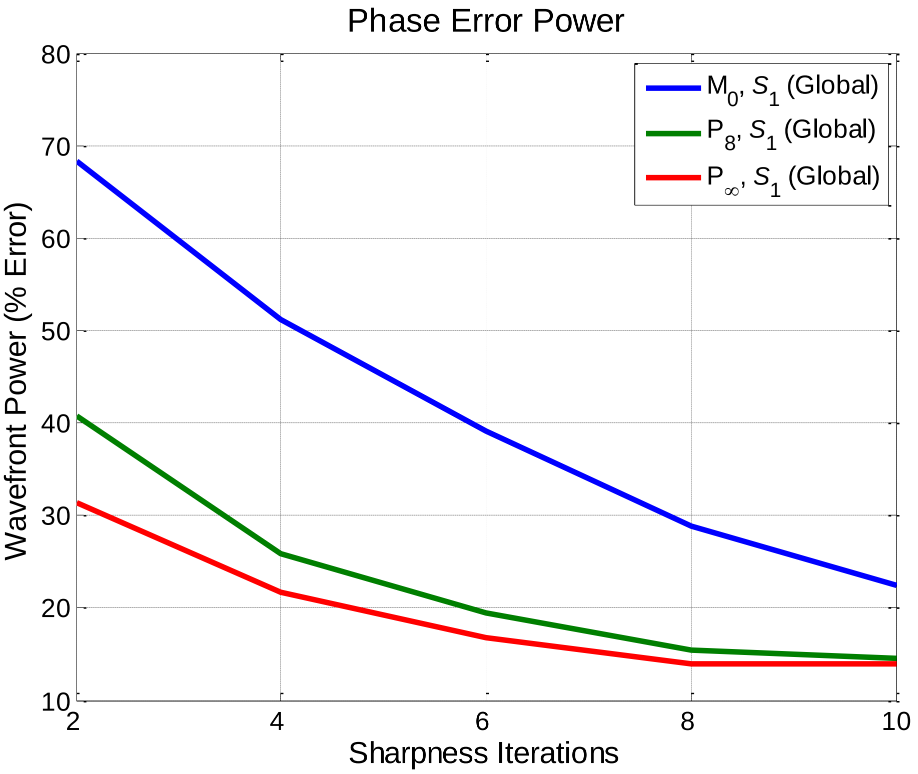

Figure 7.11: Fraunhofer regime wavefront correction. Time-series average pupil plane phase error power for M_0, P_8, and P_∞ methods versus number of sharpness iterations using the S_2 metric.

CHAPTER 8

Local Sharpening and Subspace Correction

Often in application, only a particular region of an extended object is required to be imaged. An ROI can readily be sharpened by isolating the region of the reconstructed image, usually at the expense of degradation in areas outside the region (Fig.1.2). The resulting wavefront estimate is therefore not one that minimizes the pupil plane aberration over the entire image, but rather for that particular area, a fact that may limit applications interested in wavefront sensing. Earlier works have studied the effects of ROI sharpening [11, 25, 26]. Particularly, localized sharpening has demonstrated potential image improvement in anisoplanatic images over the ROI [27], increasing interest in localized image sharpening. By partitioning an extended object over its image, one could presumably sharpen and splice several near-isoplanatic patches together to reproduce a higher quality image than one corrected over the entire field-of-view. However, for each iteration of the sharpness optimization, the method of localized sharpening still involves correcting the wavefront through full-size pixel grid propagation of the spectral plane of the hologram during image reconstruction. A more efficient technique of localized sharpening aimed at real-time application is introduced by way of subspace correction.

8.1 Single Patch Sharpening

Subspace wavefront correction tackles the linear combination of pixels on the spectral plane of the hologram that span the subspace which characterize an ROI on the image. Figure 8.1 demonstrates this concept over one 60×60 pixel grid ROI. Recall that the wavefront is applied at the circular pupil plane of the interferometry simulation. The spectral plane of the hologram retains this physical feature, as seen by the original 256×256 pixel array. Since light from every pixel on the pupil plane contributes to the image, a specific ROI is isolated from the resulting image and propagated back to the spectral plane to obtain the subset of pixels on the spectral plane that span the ROI. Notice that the circular pupil structure of the spectral plane for the "localized" wavefront is maintained. This wavefront is not localized in the spatial sense; rather, there exists a wavefront 60×60 pixel array subset constituting a linear combination of pixels on the pupil plane representing the ROI. Initially, it is necessary to perform propagation of the full spectral plane to generate the isolated area and back-propagate to the subspace wavefront. However, once this reduced-size subspace wavefront

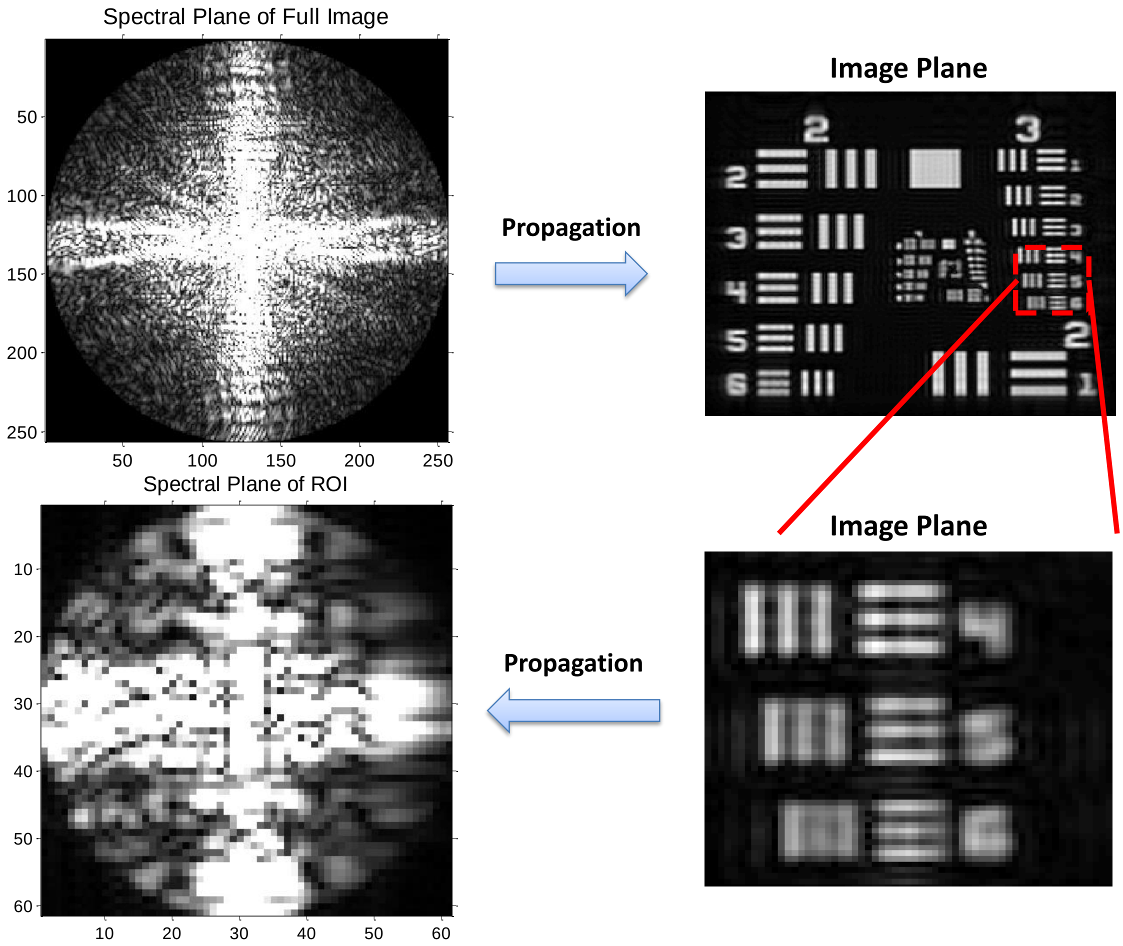

Figure 8.1: Scheme for initial subspace wavefront generation pertaining to one ROI of the image.

is captured, it is all the information needed for subsequent sharpness optimization over that ROI. In other words, each iteration during optimization for the example in Fig. 8.1 will only need to propagate the 60×60 grid rather than the 256×256 grid, drastically reducing the computation cost needed for the sharpening algorithm.

An issue unique to local sharpening not experienced by its global counterpart is the fact that artifacts attenuate the quality of the ROI, including stray light from outside the region, dispersed light from within the region, and the presence of localized phase tilt in the region. Since the initial subspace wavefront is captured only at the start of optimization, some information is unavoidably lost when sharpening over the initial subspace wavefront. Compared to local sharpening, which has access to the global wavefront to potentially mitigate these artifacts after each iteration, subspace correction is restricted to only the information initially captured by the subspace wavefront. Simulation results have shown that this loss effect causes subspace sharpened local regions to perform within 1% RMS image error as those with global sharpening, but more noticeably behind those with local sharpening. The trade-off is, of course, a significant improvement in algorithm run-time. One solution to this loss of quality during subspace correction is through mosaicking several block partitions of the global image, each independently sharpened in parallel through subspace correction. This preserves the advantage of computation efficiency and enables the capability of wavefront estimation, a feature lost with conventional local sharpening. With knowledge of the wavefront estimate, local sharpening artifacts can be circumvented. The results presented in this chapter incorporated such a mosaicking procedure, discussed in the next section.

8.2 Block Sharpening

It is possible to reconstruct the wavefront estimate for the entire image through a partition and splicing technique. The idea is for several block partitions of the image to be sharpened by parallel and independent subspace wavefront corrections. While the corrected wavefront for each individual block does not correspond to the near-optimal wavefront estimate for the global image, splicing these blocks together reconstructs a full field-of-view sharpened image.

With knowledge of the final and initial complex fields of the image, the wavefront estimate can be back-solved from Eq. 2.12, and projected onto the Zernikes by Eq. 4.7. Algorithm 2 provides a summary of this parallel subspace wavefront correction procedure using predictive methods. Special attention should be brought to lines 8 and 9 of Algorithm 2, where a block ROI is initially partitioned from the reconstructed full-sized image and back-propagated using inverse Fresnel (or another appropriate propagator) propagation,

$$U_\varepsilon(f_\zeta, f_\eta, z) = \mathcal{F}\Big\{\mathcal{F}\{\{\hat{u}(\zeta, \eta, 0)\}\} \exp\big[-j\frac{\pi}{\lambda z}(\zeta^2 + \eta^2)\big]\Big\}. \tag{8.1}$$

This reduced-size pixel block is then fed into the image sharpening algorithm, where the iterative procedure involving reconstruction in line 12 of Algorithm 2 is done much more efficiently than previously with the full-sized image (i.e., Algorithm 1 and local sharpening). Once the sharpening optimization concludes, the reconstructed image estimate for each l block is mosaicked together back onto the full field-of-view image plane. Wavefront estimates pertaining to each l block represent the near-optimal wavefront for each l subspace. The l blocks are corrected through independent and parallel processors so that the run-time for wavefront estimation remains on the order of one reduced-size ℓ block. In order to obtain the wavefront estimate for the full image, the estimated object complex pupil function, $\hat{U}(f_\zeta, f_\eta, z)$, is obtained by back-propagating using Eq. 8.1. Knowing both the final, corrected and initial, corrupted complex fields on the pupil plane, the pupil plane phase estimate is back-solved by,

$$\hat{\psi}(t,k)(\zeta,\eta) = -\frac{1}{j}\ln\left(\frac{\hat{U}(f_\zeta, f_\eta, z)}{U_\varepsilon(f_\zeta, f_\eta, z)}\right). \tag{8.2}$$

Finally, this wavefront estimate is parameterized onto Zernike basis functions, represented by Eq. 4.7, by projecting the wavefront onto a Zernike polynomial of order N. The achievement of this subspace algorithm is the capability of wavefront estimation while retaining an efficient method for local sharpening. Since each reduced-size pixel block is independently sharpened from one another through subspace correction, wavefront estimation and image sharpening over the entire image can be done much faster than either conventional global or local sharpening. The next section reflects on results using Algorithm 2.

Algorithm 2: Pseudo-Code for Predictive Subspace Wavefront Correction

```
Initialize: $x_0^{(1)} = x_0^{(2)} = \cdots = x_0^{(l)} = 0$, $\hat{a}_0^{(1)} = \hat{a}_0^{(2)} = \cdots = \hat{a}_0^{(l)} = 0$;
for $t \leftarrow 1$ to $T$ do
    Propagate object to object plane with Eq. 2.5 (or Eq. 2.3);
    Coherently image object and determine hologram with Eq. 2.6 and Eq. 2.11;
    Transform hologram to spectral domain;
    for processor $l$ in parallel do
        Determine $x^{(l)}(t+1)$ and $\hat{a}^{(l)}(t+1)$, by Eq. 5.15;
        Reconstruct global image with Eq. 4.5;
        Back-propagate ROI partition $l$ to capture subspace wavefront with Eq. 8.1;
        for $k \leftarrow 1$ to $k_f$ do
            if $(S_{opt} - S) < \epsilon_1$ then
                Realize subspace wavefront as in Eq. 4.7;
                Correct subspace wavefront and reconstruct image with Eq. 4.5;
                Calculate sharpness index, $S$ as in Eq. 4.6 or Eq. 4.3;
                Calculate analytic gradient, as in Eq. 4.12;
            else
                Break;
            end
        end
        if $G(t) < \epsilon_2$ then
            Update outlier filter using exponential smoothing, as in Eq. 5.18;
        end
        Splice block partition $l$ image for global field-of-view image;
    end
    Back-propagate ROI to pupil plane and back-solve phase estimate with Eq. 8.2;
    Realize global field-of-view corrected wavefront estimate as in Eq. 4.7;
end
```

8.3 Results

The results presented here are produced from the methods described in Section 8.2, where the same dynamic sequence of 8000 aero-optical wavefronts used for the simulations in Chapter 7 are used here. For the simulation in this section, the 256 × 256 pixel USAF bar chart is partitioned into four 128 × 128 pixel blocks. At each wavefront frame in the wavefront sequence, the subspace wavefront for each block is projected onto an initial set of 15 Zernike coefficients as in the form seen in Eq. 5.11. This method of sharpness algorithm initialization, M_0, is typical of a polynomial-expanded wavefront representation. The piston, tip, and tilt Zernike modes were not sharpened through optimization, but instead determined through least-squares projection for each time-step t in the dynamic wavefront sequence (Eq. 4.8). These modes were not included during the back-solve for the wavefront estimate in order to avoid phase wrapping issues. Incidentally, since they simply represent translation of the image, they do not influence the sharpening metric, and can be reintroduced at anytime during post-processing.

In an effort to push digital holography for real-time applications, a sharpness initialization using prediction, as in Eq. 5.13, is also considered here. For the results in this section, an LTI prediction filter with $k_{\mathrm{ID}} = 8$ (i.e., P_8) is identified from a sample subspace wavefront sequence for each block partition. These LTI filters run independent and parallel from each other, determining the one-step prediction of the final sharpened Zernike coefficient vector for their corresponding block partitions. The first 4000 wavefront frames are reserved to generate the sample input sequence for identifying the P_8 prediction filter, while the subsequent 2000 frames constitute the time-series run-time data for the results in this simulation.

Three different approaches to image sharpening are studied: global correction, conventional local sharpening, and subspace correction. For each approach, the results are measured against the metric of wavefront estimates and reconstructed images generated by projecting the Zernike coefficient vectors directly onto the true wavefront sequence. Figure 8.2 illustrates the subspace block partitions used for wavefront estimation and the two ROI areas highlighted are analyzed by the aforementioned sharpening techniques under investigation.

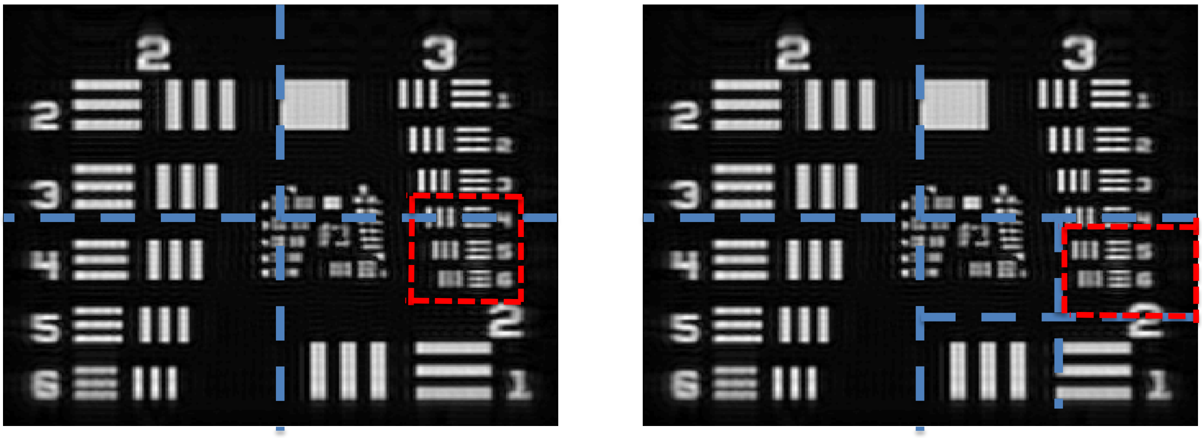

Figure 8.2: ROI area A_1 (left) and A_2 (right); dashed red and subspace block partitions; dashed blue

Figure 8.3 presents a pupil plane phase error power time-series average versus sharpness iterations for the metric defined in Eq. 4.6. Each data point is an average of the time-series of phase error power after k_f sharpness iterations for a 2000 wavefront sequence. Particular examples of these time-series are shown in Fig. 8.3. The subspace wavefront estimates were determined by block partition splicing, back-solving, and projection onto 15 Zernike coefficients, while the global wavefront estimates were determined simply by the resulting Zernike coefficient vector after k_f sharpening iterations.

Mosaicking subspace blocks enables estimation of the wavefront otherwise inaccessible through local sharpening. For this reason, plots of wavefront estimation using conventional local sharpening are not shown. The resulting plots demonstrate that, for the case of M_0, subspace correction has a noticeable improvement in wavefront estimation over global correction. One explanation for this is that the subspace wavefront represents a set of pixels spanning the wavefront that uniquely characterizes each partition block. Thus, sharpening over each partition produces the wavefront that uniquely minimizes the phase error for the corresponding ROI.

For conventional local sharpening, the resulting wavefront similarly minimizes the phase error for the particular local region. Unfortunately, this wavefront correction is done on a global scale with light from linearly dependent pixels that affect the full image, typically

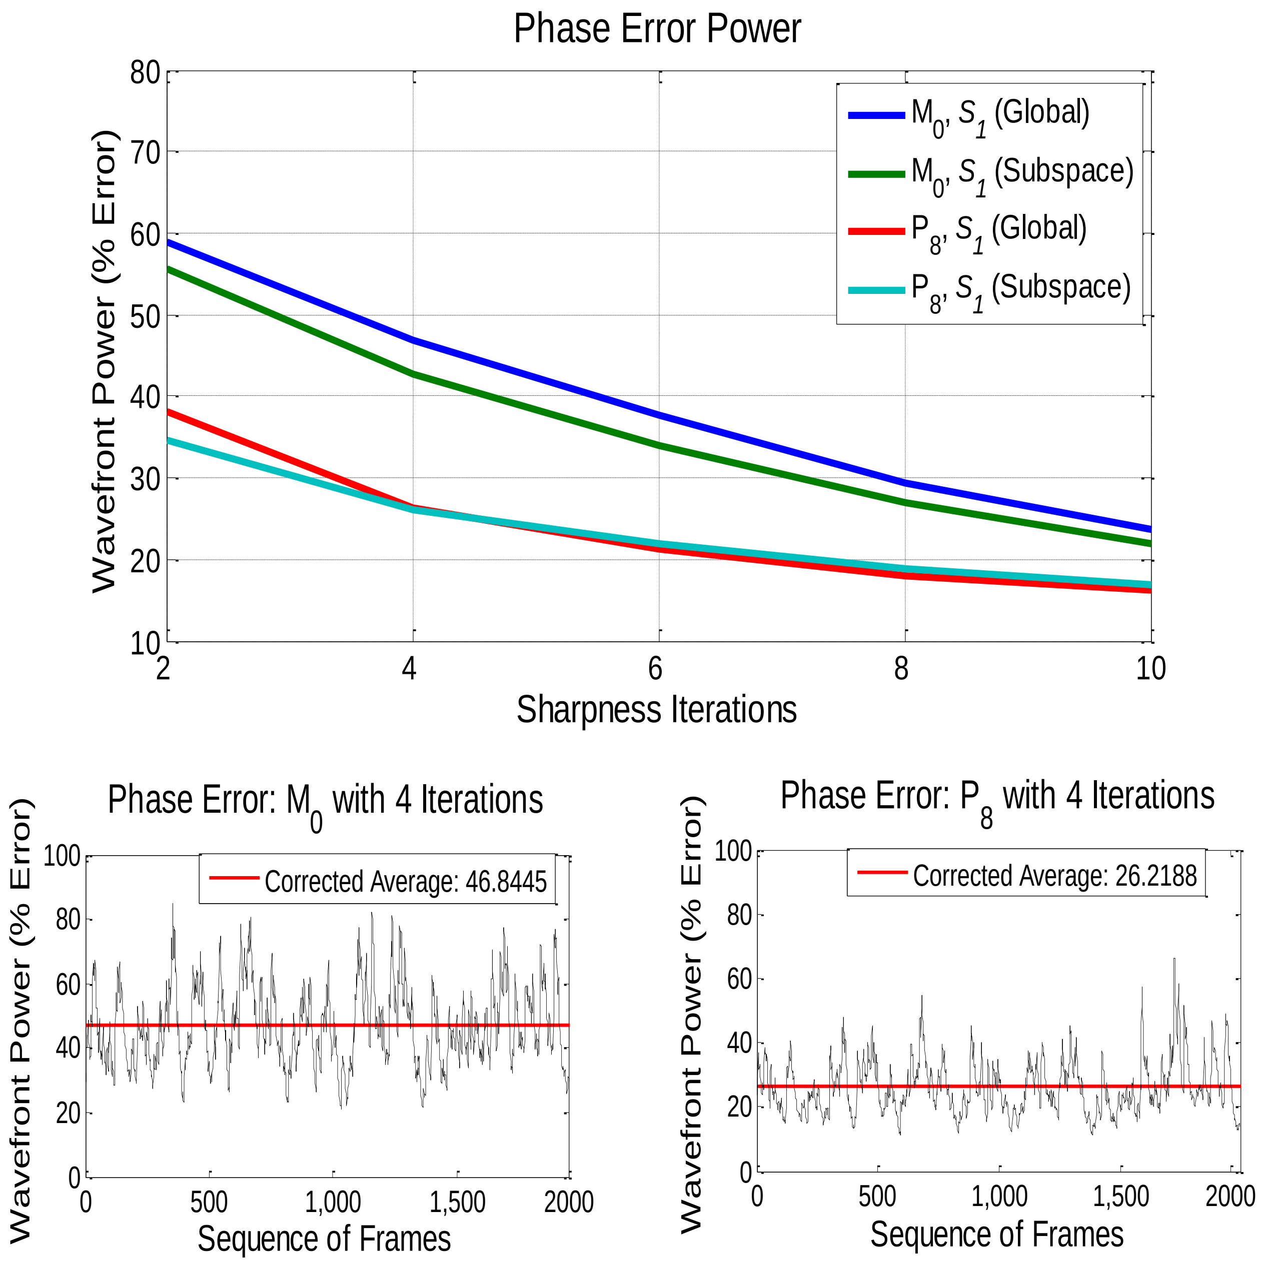

Figure 8.3: Top plot: Time-series average pupil plane phase error power for M_0, and P_8 methods versus number of sharpness iterations using the S_1 metric. Bottom plots: Time-series of phase error power after four sharpness iterations for: (left) M_0 global correction and (right) P_8 subspace correction.

leading to degradation of the global image outside the local region. This is why local sharpening is ill-suited for wavefront estimation and correcting, as in Figure 8.3. Mosaicking the partition areas together through subspace correction creates an aggregate image where the wavefront estimate is consequently not penalized for any part of the image; the subspace wavefront the ROI is corrected over is unique to only that ROI and independent of the rest of the full pupil plane complex field.

Finally, in Figure 8.3, the prediction filter, P_8, using a subspace correction approach is capable of outperforming the non-prediction methods. The prediction method approach is intended to reduce the number of sharpness iterations and improve convergence rate. It is paramount to remember, however, that the primary goal of subspace correction is to speed up the digital holography sharpness algorithm by efficient wave propagation of the remaining iterations without sacrificing the quality of the wavefront estimation. Figure 8.3 indicates that such a sacrifice does not occur. It maintains a very comparable behavior to the global correction approach done with P_8, with the only recognizable difference being at $k_f = 2$ sharpness iterations. This close behavior is most likely due to the fact that, through prediction, a significant percentage of frames in the wavefront sequence has already essentially converged by $k_f = 4$ iterations, leaving little room for improvement with regards to wavefront estimation through subspace correction. The M_0 method requires a greater number of iterations for convergence, leading to more distinctive differences for that case.

Figure 8.4 shows a histogram of the number of sharpness iterations per frame required to achieve within 0.01% convergence of the final-value pupil plane power using the M_0 and P_∞ methods for the global sharpening and subspace correction techniques. As expected from Fig. 8.3, the global and subspace variants of M_0 have similar performance metrics. Both variants of the P_8 filter performed comparably as well, with 45% of frames converging prior to 10 sharpening iterations. The ramification of Fig. 8.4 is that the subspace approach speeds up sharpening of each iteration without losing performance in the way of wavefront estimation. Combined with the prediction filter retaining its ability to converge in half the time of non-prediction methods, the implication lends itself to viable real-time digital holography operations.

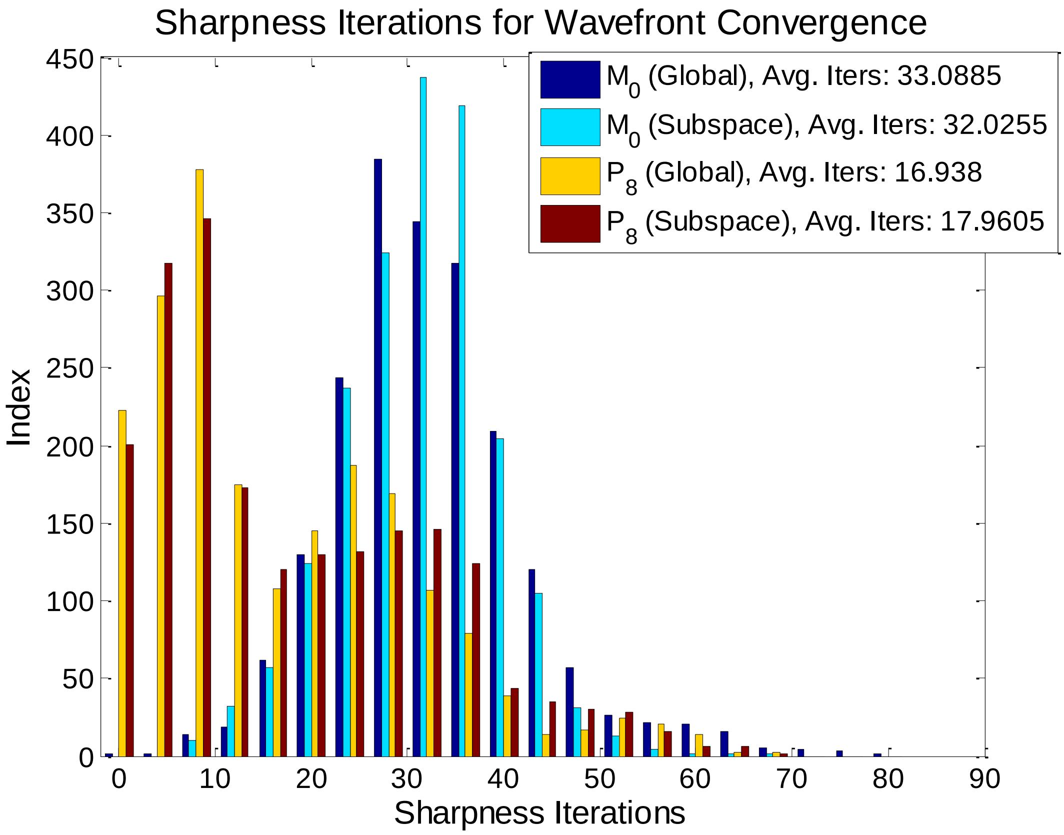

Figure 8.4: Histogram of the number of sharpness iterations required for wavefront correction convergence for the M_0 and P_8 methods using global and subspace techniques.

Figure 8.5 describes the time-series average irradiance RMS error versus sharpness iterations over ROI area A_1 in Fig. 8.2, after global, local, and subspace correction. One possible point of concern was the potential loss of information due to artifacts that may cause a degradation in image quality after subspace correction. A method described by Section 8.2 is used to bypass this issue. After reconstructing the full image from the projected Zernike coefficient vector determined by back-solving Eq. 2.12, an ROI can be readily windowed as in global correction. Ideally, one would inscribe the ROI entirely as part of a partitioned block for subspace correction. For the simulation presented here, the ROI in Fig. 8.1 is partly contained by two of the four 128 × 128 pixel blocks.

First, it can be observed that the subspace method for both non-prediction, M_0, and prediction, P_8, filters performs better than global correction. This is owed to the block procedure that provides knowledge of the full wavefront estimate that would otherwise cause

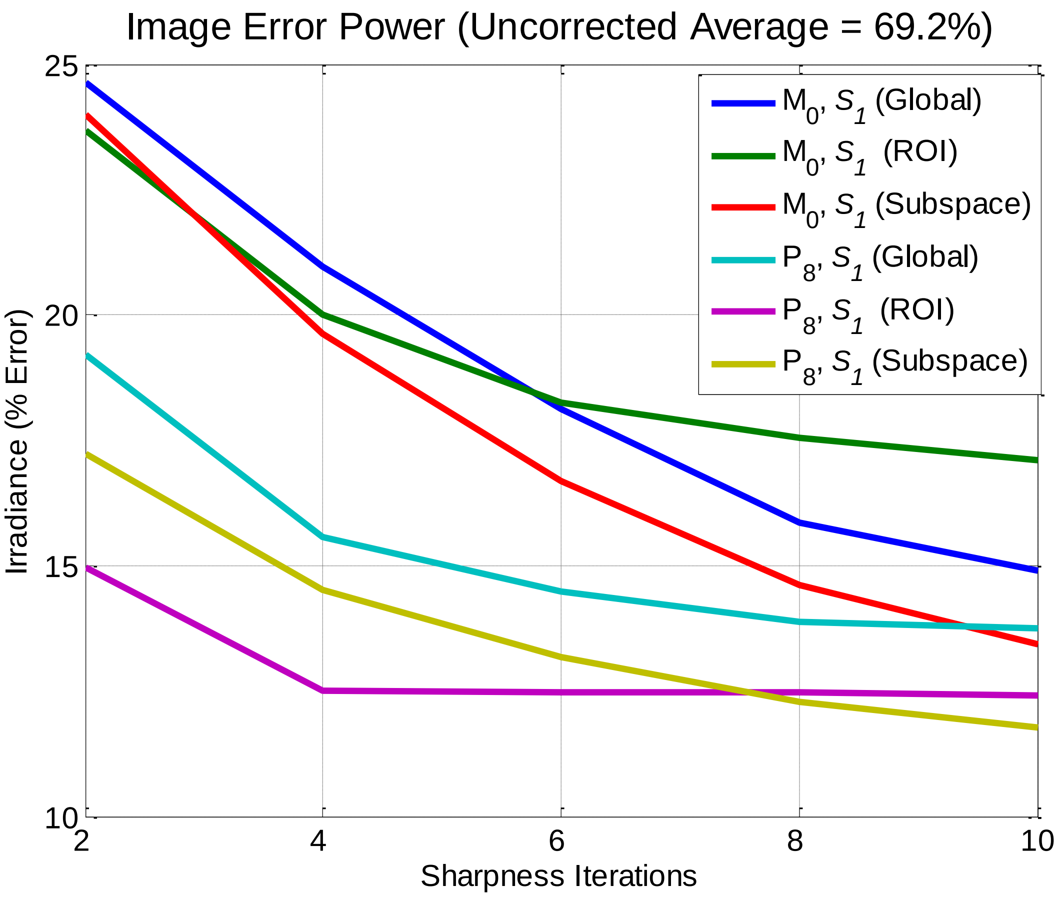

Figure 8.5: Time-series average irradiance RMS error for M_0 and P_8 methods versus number of sharpness iterations over A_1 using global and local sharpening and subspace correction.

some loss in quality due to localized artifacts. In the case of M_0 with local sharpening, these artifacts are prevalent in enough outlier frames that the time-average appears worse than global correction. About 7% of the 2000 frame sequence experiences such ample deterioration. Removing these frames would restore the expected behavior of the M_0 local method relative to the global method. The prediction filter for local sharpening does not experience as many of these outlier wavefronts in the dynamic sequence because the prediction filter is capable of beginning the sharpening at near-optimal convergence, frame-by-frame. Finally, because the simulation performed here did not include a partitioned block exclusively for the ROI, it is not expected that the irradiance of the image be corrected better than that with local sharpening, which calculates the metric cost function directly over only the ROI. Indeed, this outcome is reflected in Figure 8.5 by the P_8 curves for local and subspace

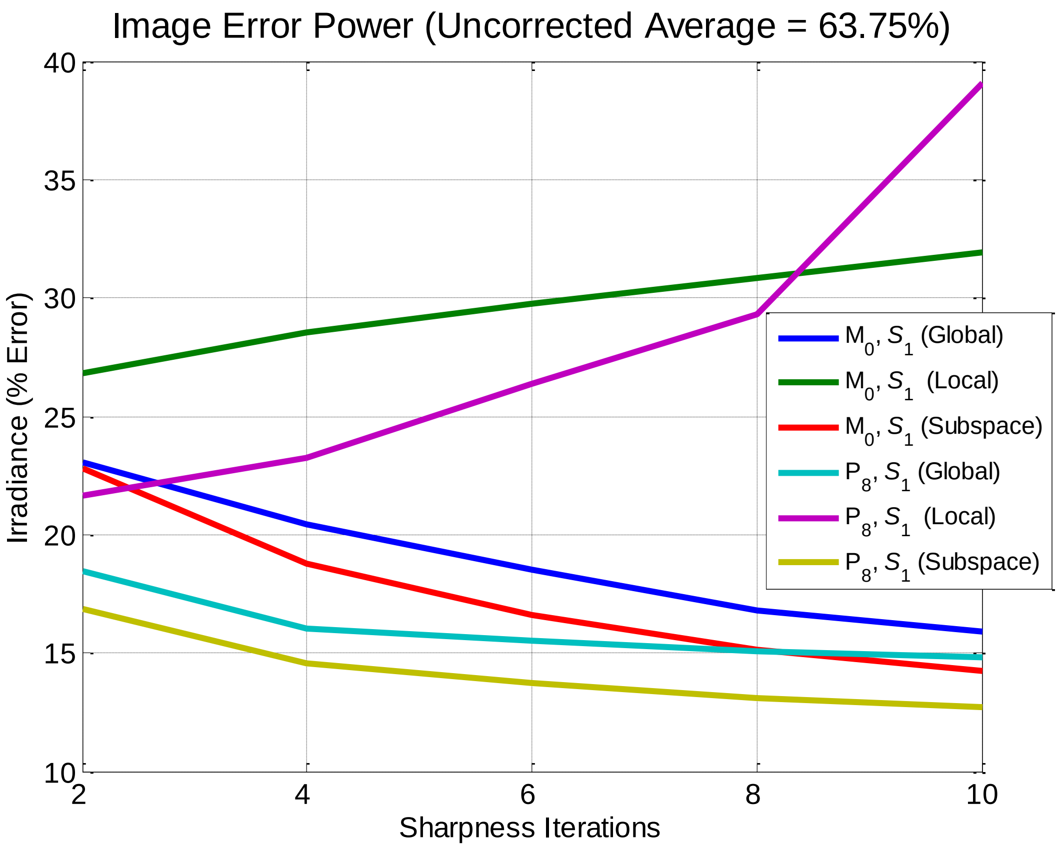

Figure 8.6: Time-series average irradiance RMS error for M_0 and P_8 methods versus number of sharpness iterations over A_2 using global and local sharpening and subspace correction.

correction.

Figure 8.6 shows the time-series average irradiance RMS error versus sharpness iterations over ROI area A_2 in Fig. 8.2, after global, local, and subspace correction. While the global and subspace approaches for M_0 and P_8 sharpening algorithms perform as expected, the local sharpening variants diverge immediately with poor results. Inspecting the two ROI, A_1 and A_2, the former consists of the bright pixels of interest naturally padded by darker pixels around its perimeter. Artifacts present in the local sharpening method are mitigated due to this natural padding, providing the performance demonstrated in Fig. 8.3. For A_2, significant portions of the image are located around the border of the ROI. The light from these frames is likely to exhibit the dispersed and stray light artifacts due to the ineffective window size for this region, despite both A_1 and A_2 being equal sizes [27]. Hence, both prediction and non-prediction filters for local sharpening are contingent on the particular ROI areas and may be highly susceptible to artifacts tied to this method. Conversely, the subspace method is insensitive to the choice of ROI, as the analysis of the two ROI areas seem to suggest.

CHAPTER 9

Conclusion

Digital holography has demonstrated its capability to be a lightweight and cost-effective solution for wavefront estimation and correction in the presence of atmospheric turbulence. However, the Fourier and Fresnel propagations required during sharpness optimization prevent it from experiencing widespread real-time implementation. The goal of this research was to solve this computational complexity by employing methods based on approaches for optimal and adaptive identification, prediction, and control of optical wavefronts. The end result is a novel method for a combined identified LTI prediction filter and sharpness algorithmic approach to dynamic wavefront correction that both significantly improves initial pupil plane phase error estimates and convergence rates, dubbed predictive dynamic digital holography. Results presented in this dissertation provide substantial evidence of the feasibility for implementing an online adaptive prediction filter to aid in the push for digital holography in real-time operations.

The techniques presented in this dissertation also aim to exploit the processing power of contemporary computers and systems to efficiently perform sharpness optimization for digital holography wavefront estimation and image reconstruction. In particular, subspace correction is a new sharpness optimization method that boasts a significant increase in speed for computation run-time compared to its global and local correction counterparts. The results presented in this work show image sharpening with subspace correction outperforming global sharpening and rivaling conventional local correction for an ROI. Furthermore, subspace correction is capable of wavefront estimation, a feature not found through local sharpening. This enables both full image reconstruction and wavefront correction abilities. Results have shown that these reconstructed wavefronts are as good, if not better, than those determined by global sharpening. Combined with predictive dynamic digital holography, subspace correction provides an avenue for real-time applications to address the algorithmic complexity that prevents it from practical operations.

In Chapters 2 and 4, the theory and application of digital holography recording, reconstruction, and sharpening algorithms are described. Various propagators and techniques for image reconstruction are investigated to examine the robustness and versatility of wavefront prediction in the scope of digital holography. The primary diffraction propagators used are Fresnel and the angular spectrum method. Meticulous attention is focused on avoiding issues of both aliasing during holographic recording and reconstruction, and the concept of over-sharpening during image sharpening optimization. The sequence of atmospheric wavefronts, obtained by the University of Notre Dame's Airborne Aero-Optics Laboratory, is analyzed and modified in Chapter 3 and modeled using Zernike modes, detailed in Chapter 4. The Strehl ratios of the modified wavefront sequence used for the simulation are shown to be highly disruptive with an average of 0.27.

Chapter 5 reviews subspace identification based on a recursive least-squares lattice filter. The heart of predictive dynamic digital holography is introduced in this chapter: an identified state-space LTI filter is integrated with the sharpness optimization algorithm to improve the performance of numerical reconstruction for digital holography. These prediction filters, $\mathrm{P}_{k_{\mathrm{ID}}}$, are built alongside the non-prediction filters, M_0 and M_1, to ascertain the performance and wavefront correction differences. Past research work studied wavefront prediction in the context of adaptive optics with the aid of a Shack-Hartmann sensor. However, the research presented here is unique because wavefront prediction is done in the context of digital holography, specifically without any wavefront sensor to supplement the process.

Chapter 6 analyzes the wavefront prediction filter in more detail. A benchmark prediction filter, P_∞, is formulated from the true wavefronts projected directly onto a Zernike polynomial expansion, behaving as the theoretical limit of prediction that can be achieved. This benchmark is compared directly to feasible, finite prediction filters P_4 and P_8. This chapter shows the finite prediction filters have close approximations to the behavior of the benchmark prediction filter, P_∞. The power spectra presented capture similar statistics of the power for the various Zernike modes for both the ideal and finite prediction filters. Finally, the prediction error sequence is shown to have relatively low power in the less-identifiable higher modes, verifying the prediction filter for operation.

Results for non-prediction and prediction wavefront correction methods are presented in Chapter 7. Wavefront estimation, image irradiance, and Strehl ratios are among the metrics investigated. The results show that the finite prediction filters, even P_4, are significant improvements to the non-prediction filters by all metrics studied. The finite filter P_8 is demonstrated to be a particularly close approximation in all regards to the benchmark prediction filter. One implication for the achievements of the finite LTI filters is the validation for a time-varying k_{ID} or adaptive filter. Overall, the results dictate the ability of wavefront prediction to solve the computational burden that is native to numerical image reconstruction in digital holography, with better wavefront estimation than conventional methods.

A different approach to local sharpening, labeled subspace correction, is introduced in Chapter 8. Local sharpening has been studied in past works and holds a special role in the area of correcting for anisoplanatism in wavefronts. However, it is limited to only image irradiance correction and not wavefront estimation. Additionally, despite operating over a particular area of an image, local sharpening still propagates over the full pixel grid and corrects over the full wavefront. Both inefficiencies with local sharpening are addressed by subspace correction. The results in this chapter demonstrate equivalent image correction to local sharpening in a significantly improved numerical procedure. Not only does subspace correction, unlike local sharpening, enable the ability to wavefront estimate and correct, it does so better than global correction methods.

Predictive wavefront techniques developed in this dissertation are coupled with subspace correction to ultimately provide the most efficient version of predictive dynamic digital holography and its numerical reconstruction algorithm. This algorithm is intended to provide the foundation to solving the complexity issue for digital holography in real-time applications.

REFERENCES

[1] D. Gabor, “A new microscopic principle,” Nature **161**, 777–778 (1948).

[2] C. Porter and et. al., “Flight measurements of aero-optical distortions from a flat-windowed turret on the airborne aero-optics laboratory (aaol),” in “AIAA Plasmadynamics and Lasers Conf., AIAA, Honolulu, Hawaii, 2011. Proceedings of the 2011,” , vol. 42 (2011), vol. 42.

[3] S. Gordeyev and et. al., “Aero-optical environment around a cylindrical turret with a flat window,” AIAA J. **49**, 308–315 (2011).

[4] E. Leith and J. Upatnieks, “Reconstructed wavefronts and communication theory,” J. Opt. Soc. Am. **52**, 1123–1130 (1962).

[5] R. L. Powell and K. A. Stetson, “Interferometric vibration analysis by wavefront reconstruction,” J. Opt. Soc. Am. **55**, 1593–1598 (1965).

[6] D. W. Sweeney and C. M. Vest, “Reconstruction of three-dimensional refractive index fields from multidirectional interferometric data,” Appl. Opt. **12**, 2649–2664 (1973).

[7] J. Goodman and R. Lawrence, “Digital image formation from electronically detected holograms,” Appl. Phys. Lett. **11**, 77–79 (1967).

[8] R. A. Muller and A. Buffington, “Real-time correction of atmospherically degraded telescope images through image sharpening,” J. Opt. Soc. Am. **64**, 1200–1210 (1974).

[9] J. Fienup and J. Miller, “Aberration correction by maximizing generalized sharpness metrics,” J. Opt. Soc. Am. **20**, 609–620 (2003).

[10] S. T. Thurman and J. R. Fienup, “Phase-error correction in digital holography,” J. Opt. Soc. Am. A **25**, 983–994 (2008).

[11] S. Sulaiman and J. S. Gibson, “Predictive dynamic digital holography,” [Submitted for review to J. Opt. Soc. Am. A] (2017).

[12] S. Sulaiman and J. S. Gibson, “Predictive dynamic digital holograph with local image sharpening and subspace wavefront correction,” [Submitted for review to J. Opt. Soc. Am. A] (2017).

[13] U. Schnars and W. Jüptner, “Direct recording of holograms by a ccd target and numerical reconstruction,” Appl. Opt. **33**, 179–181 (1994).

[14] L. Hesselink, “Fundamental issues related to digital holographic data storage,” in “LEOS ’99. IEEE Lasers and Electro-Optics Society 1999 12th Annual Meeting,” (1999).

[15] C. J. Mann, L. Yu, and M. K. Kim, “Movies of cellular and sub-cellular motion by digital holographic microscopy,” BioMedical Engineering OnLine **5** (2006).

[16] S. Grilli, P. Ferraro, S. D. Nicola, A. Finizio, G. Pierattini, and R. Meucci, "Whole optical wavefields reconstruction by digital holography," Opt. Express **9**, 294–302 (2001).

[17] G. Liu and P. Scott, "Phase retrieval and twin-image elimination for in-line fresnel holograms," J. Opt. Soc. Am. **4**, 159–165 (1987).

[18] J. R. Fienup, "Synthetic-aperture radar autofocus by maximizing sharpness," Opt. Lett. **25**, 221–223 (2000).

[19] P. Marquet and et. al., "Digital holographic microscopy: a noninvasive contrast imaging technique allowing quantitative visualization of living cells with subwavelength axial accuracy," Opt. Lett. **30**, 468–470 (2005).

[20] T. Shimobaba, Y. Sato, J. Miura, M. Takenouchi, and T. Ito, "Real-time digital holographic microscopy using the graphic processing unit," Opt. Express **16**, 11776–11781 (2008).

[21] N. Y. N. Chen and J. S. Gibson, "Subspace system identification using a multichannel lattice filter," in "American Control Conference, 2004. Proceedings of the 2004," , vol. 1 (2004), vol. 1, pp. 855–860.

[22] J. Tesch and J. Gibson, "Optimal and adaptive control of aero-optical wavefronts for adaptive optics," J. Opt. Soc. Am. **29**, 1625–1638 (2012).

[23] J. Tesch, J. Gibson, and M. Verhaegen, "Receding-horizon adaptive control of aero-optical wavefronts," Optical Engineering **52**, 071406–071406 (2013).

[24] J. Tesch, T. Truong, R. Burruss, and J. Gibson, "On-sky demonstration of optimal control for adaptive optics at palomar observatory," Opt. Lett. **40**, 1575–1578 (2015).

[25] A. E. Tippie, A. Kumar, and J. R. Fienup, "High-resolution synthetic-aperture digital holography with digital phase and pupil correction," Opt. Express **19**, 12027–12038 (2011).

[26] J. C. Marron, R. L. Kendrick, N. Seldomridge, T. D. Grow, and T. A. Höft, "Atmospheric turbulence correction using digital holographic detection: experimental results," Opt. Express **17**, 11638–11651 (2009).

[27] A. E. Tippie and J. Fienup, "Region-of-interest sharpness correction," in "Frontiers in Optics 2011/Laser Science XXVII," (Optical Society of America, 2011), p. FTuF2.

[28] M. F. Spencer and et al., "Deep-turbulence wavefront sensing using digital-holographic detection in the off-axis image plane recording geometry," Optical Engineering **56** (2016).

[29] J. Schmidt, *Numerical Simulation of Optical Wave Propagation with Examples in MATLAB* (SPIE Press Book, 2010).

[30] T. Poon and J. Liu, *Introduction to Modern Digital Holography* (Cambridge University Press, 2014).

[31] J. Goodman, *Introduction to Fourier Optics* (MaGraw-Hill, 1996), 2nd ed.

[32] M. Kim, *Digital Holographic Microscopy: Principles, Techniques, and Applications* (Springer, 2011).

[33] T. Kreis and W. Juptner, "Suppression of the dc term in digital holography," Optical Engineering **36**, 2357–2360 (1997).

[34] E. Cuche, P. Marquet, and C. Depeursinge, "Spatial filtering for zero-order and twin-image elimination in digital off-axis holography," Appl. Opt. **39**, 4070–4075 (2000).

[35] S. Gordeyev and E. Jumper, "Fluid dynamics and aero-optics of turrets," Progress in Aerospace Sciences **46**, 388–400 (2010).

[36] K. Strehl, "Aplanatische und fehlerhafte abbildung im fernrohr," Zeitschrift fr Instrumentenkunde pp. 362–370 (1895).

[37] L. Rayleigh, "Investigations in optics, with special reference to the spectroscope," Philosophical Magazine **8**, 261–274 (1879).

[38] J. W. Hardy, *Adaptive Optics for Astronomical Telescopes* (Oxford University Press, 1998), 1st ed.

[39] A. Marechal, "Etude des effets combins de la diffraction et des aberrations gomtriques sur limage dun point lumineu," Rev. Opt. **2**, 257–277 (1947).

[40] V. Mahajan, "Strehl ratio for primary aberrations in terms of their aberration variance," J. Opt. Soc. Am **73**, 860–861 (1983).

[41] R. G. Paxman and J. C. Marron, "Aberration correction of speckled imagery with an image-sharpness criterion," in "Proc.SPIE Statistical Optics," , vol. 0976 (1988), vol. 0976, pp. 0976 – 0976 – 11.

[42] S. T. Thurman and J. R. Fienup, "Correction of anisoplanatic phase errors in digital holography," J. Opt. Soc. Am. A **25**, 995–999 (2008).

[43] E. Acosta, S. Bará, M. A. Rama, and S. Ríos, "Determination of phase mode components in terms of local wave-front slopes: an analytical approach," Opt. Lett. **20**, 1083–1085 (1995).

[44] D. L. Fried, "Statistics of a geometric representation of wavefront distortion," J. Opt. Soc. Am. **55**, 1427–1435 (1965).

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

[46] P. Wizinowich and et. al., "Near-infrared wavefront sensing," in "Proc.SPIE Adaptive Optics Systems V," , vol. 9909 (2016), vol. 9909, pp. 9909 – 9909 – 13.

[47] R. Tyson, *Principles of Adaptive Optics* (CRC Press Book, 2010), 3rd ed.

[48] T. Katayama, *Subspace Methods for System Identification* (Springer, 2005).

[49] P. van Overschee and B. D. Moor, *Suibspace Identification for Linear Systems* (Springer, 1996).

[50] A. Faghihi, J. Tesch, and J. Gibson, "Identified state-space prediction model for aero-optical wavefronts," Optical Engineering **52**, 071419–071419 (2013).