INTERFEROMETER FOR MEASURING DYNAMIC CORNEAL TOPOGRAPHY

by

Jason Daniel Micali

_____

A Dissertation Submitted to the Faculty of the

COLLEGE OF OPTICAL SCIENCES

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2015

THE UNIVERSITY OF ARIZONA

GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Jason Daniel Micali entitled “Interferometer for Measuring Dynamic Corneal Topography” and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy

_____ Date: 07/20/2015

John E. Greivenkamp

_____ Date: 07/20/2015

Jim Schwiegerling

_____ Date: 07/20/2015

James C. Wyant

Final approval and acceptance of this dissertation is contingent upon the candidate’s submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

_____ Date: 07/20/2015

Dissertation Director: John E. Greivenkamp

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at the University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that an accurate acknowledgment of the source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the head of the major department or the Dean of the Graduate College when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

SIGNED: Jason D. Micali

ACKNOWLEDGEMENTS

I would like to thank my advisor, John Greivenkamp, for his support and guidance during my time at the College of Optical Sciences. I would also like to thank the other members of my committee, Jim Schwiegerling and Jim Wyant, for their counsel. I am also thankful to the funding sponsor of this project, for which this research would not be possible.

A special thanks to my friends and lab mates who have provided support in a myriad of ways; this includes everything from discussing research to being on the receiving end of a prank. I am also indebted to my friends and coworkers outside of the college who provided me with the encouragement and support that I needed to pursue this endeavor.

Most importantly, I would like to thank my family, who has always supported me in my education and career. I would not be where I am today if it was not for them. And to my wife Briana, whose love, patience, and sacrifice in moving to the Arizona desert allowed this to be possible.

And finally, I would be remiss if I did not thank El Güero Canelo and Caffé Lucé, who literally provided me with the strength to complete this research.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

The cornea is the anterior most surface of the eye and plays a critical role in vision. A thin fluid layer, the tear film, coats the outer surface of the cornea and serves to protect, nourish, and lubricate the cornea. At the same time, the tear film is responsible for creating a smooth continuous surface where the majority of refraction takes place in the eye. A significant component of vision quality is determined by the shape of the cornea and stability of the tear film. It is desirable to possess an instrument that can measure the corneal shape and tear film surface with the same accuracy and resolution that is currently performed on common optical elements.

A dual interferometer system for measuring the dynamic corneal topography is designed, built, and verified. The completed system is validated by testing on human subjects. The system consists of two co-aligned polarization splitting Twyman-Green interferometers designed to measure phase instantaneously. The primary interferometer measures the surface of the tear film while the secondary interferometer simultaneously tracks the absolute position of the cornea. Eye motion, ocular variation, and a dynamic tear film surface will result in a non-null configuration of the surface with respect to the interferometer system. A non-null test results in significant interferometer induced errors that add to the measured phase. New algorithms are developed to recover the absolute surface topography of the tear film and corneal surface from the simultaneous interferometer measurements. The results are high-resolution and high-accuracy surface topography measurements of the *in vivo* cornea that are captured at standard camera frame rates.

This dissertation will cover the development and construction of an interferometer system for measuring the dynamic corneal topography of the human eye. The discussion starts with the completion of an interferometer for measuring the tear film. The tear film interferometer is part of an ongoing research project that has spanned multiple dissertations. For this research, the instrument was tested on human subjects and resulted in refinements to the interferometer design. The final configuration of the tear film

interferometer and results from human subjects testing are presented. Feedback from this instrument was used to support the development and construction of the interferometric corneal topographer system. A calibration is performed on the instrument, and then verified against simulated eye surfaces. Finally, the instrument is validated by testing on human subjects. The result is an interferometer system that can non-invasively measure the dynamic corneal topography with greater accuracy and resolution than existing technologies.

1 INTRODUCTION AND BACKGROUND

The cornea is the foremost optical element of the human eye and plays a critical role in producing images on the retina. It is the largest avascular tissue in the human body (i.e. does not contain any blood vessels), which is necessary for it to remain optically transparent. Oxygen and nutrients are instead supplied by an external mechanism: the tear film. It is a thin, multi-layer of fluid that is 4 - 10 microns thick and composed of three layers: A mucin, aqueous, and lipid layer (Frank J Holly, 1980; Mishima, 1965). The tear film has additional roles of protecting, cleaning, and lubricating the cornea. Figure 1.1 shows a diagram of the human eye and structures of the cornea and tear film. The tear film mucin layer coats the surface of the cornea, creating a hydrophilic layer that holds the tear film to the cornea. The majority of the tear film thickness comes from the central aqueous layer, a solution that contains salts, glucose, enzymes, and proteins. The outermost layer is the lipid layer, an oil surface approximately 100 nm thick that protects the tear film by inhibiting evaporation of the aqueous layer (F J Holly & Lemp, 1977; McDonald, 1968).

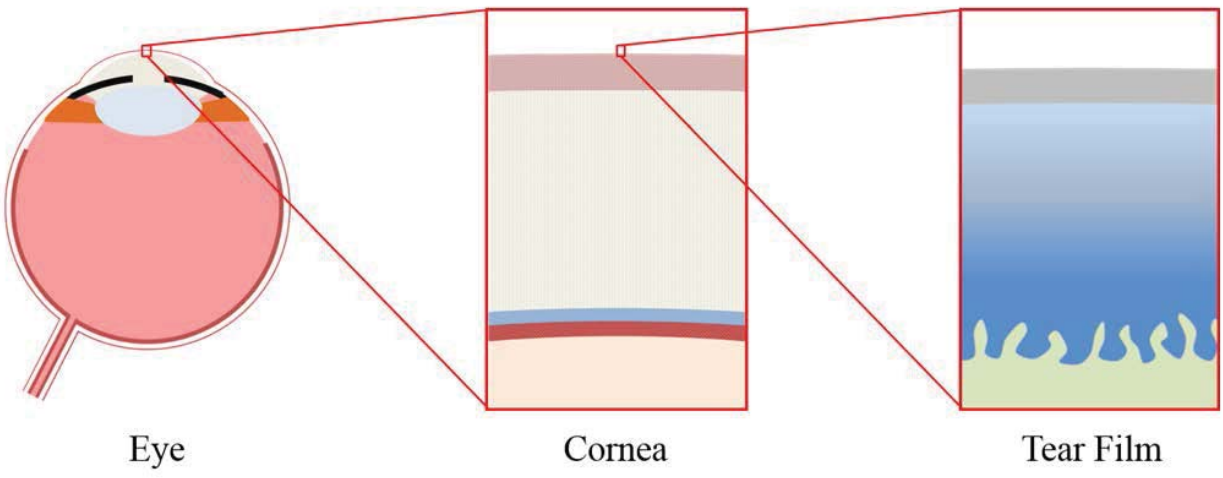

Figure 1.1: Structure of the human cornea and tear film.

The tear film is the first optical interface with the air, resulting in the largest change in refractive index of the eye. Combined with the underlying curvature of the cornea, the tear film is responsible for approximately 70% of the eye's total refractive power. A significant portion of vision quality is determined by the shape of the cornea and stability of the tear film. However, the nature of the tear film does not guarantee a smooth nor stable optical surface. Blinking is required to refresh the tear film to replenish nutrients, remove foreign substances, and reform an optically smooth surface. In a normal eye, a blink initially destabilizes the tear film then rapidly smooths itself over the surface of the cornea within a few seconds due to surface tension of the fluid (Benedetto, Clinch, & Laibson, 1984). The tear film will begin to thin due to evaporation after 5 seconds and may begin to breakup after 15 seconds (Németh et al., 2002). Breakup in the tear film typically begins with pits or canyons forming on the surface which continue to grow in size until the recurrence of another blink. Irregularities in the tear film can result in decreased visual performance (Kasprzak & Licznerski, 1999), physical discomfort (Frank J Holly, 1980), and uncomfortable contact lens wear (Guillon, 1986).

The shape of the cornea can be measured with an ophthalmic instrument known as a corneal topographer. Corneal topographers rely on the first surface reflection from the cornea, which is the tear film interface. Therefore, the overall topography of the cornea is determined by the base shape of the cornea and the structure of the tear film. Most corneal topographers are either unable to resolve the tear film or are at the limit of being able to resolve it. There exist other topographer systems that are capable of higher resolution measurements, but have other drawbacks that limit their capabilities. For example, some systems require the installation of a diffuse substance into the tear film, which alters the tear film and cannot therefore be considered a true measurement of the tear film. Other systems are capable of resolving the tear film structure, but provide no information on the topography of the cornea. These systems will be discussed in more detail in Section 1.6. What is presented in this dissertation is a new corneal topographer system that can make non-invasive measurements of the corneal topography and tear film with greater spatial and height resolution than any known system.

The development, verification, and results for a new interferometric corneal topographer instrument will be discussed in the following chapters of this dissertation. Background is presented on the variation in corneal topography and eye motion that adds a unique challenge to classical interferometric metrology. This is followed by a description of soft contact lenses and their interaction with the tear film. Systems for measuring corneal topography and the tear film are reviewed. Additional context for this research is provided through a brief history of soft contact lens metrology systems that have been developed in this research group. This chapter concludes with a discussion of the goals and requirements for the system that is presented in this dissertation.

1.1 TOPOGRAPHY OF THE HUMAN EYE

One of the challenges in the design of an interferometer to measure the topography of the human cornea is that the shape of the cornea will vary from person to person. The average corneal shape is best approximated by an asphere with a base radius of 7.8 mm and a -0.25 conic constant, which is an oblate ellipsoid (Daily & Coe, 1962; P M Kiely, Smith, & Carney, 1982; Stenstrom, 1948). The range of radii that were measured in these studies varied from 6.7 to 9.4 mm. The distributions were normally distributed with a standard deviation of 0.3 mm about a mean of 7.8 mm. The conic constant, or asphericity, was also measured in a range from -0.76 to +0.47. The asphericity measurements were normally distributed with a standard deviation of 0.27 about the mean value of -0.25.

A more accurate description of the cornea includes meridional variations of the base radius, which results in a toroidal shape of the cornea. The toroidal shape of the cornea is more commonly referred to as corneal astigmatism. Mean variation between meridional radii was measured to be 0.1 – 0.2 mm, where the variation was defined as the difference between the maximum radius and minimum radius (P M Kiely et al., 1982; Patricia M Kiely, Smith, & Carney, 1984). Approximately 74% of the study population had a larger radius close to the horizontal (i.e. within 45º).

Additional studies would indicate that a toroidal description of the cornea is incomplete (Guirao & Artal, 2000; Iskander, 2009; J. T. Schwiegerling & Greivenkamp, 1997). Surface deviation from a toroid can exceed a few microns or more due to the complex surface structure of the cornea. Tear film dynamics, especially the reforming and thinning, should be expected to add an additional range of surface variation that is comparable to the thickness of the tear film (i.e. 4 – 10 microns). The complex and dynamic surface topography of the cornea presents a significant challenge to making an interferometric measurement.

1.2 EYE MOTION

The advantage gained by using interferometry to acquire high resolution measurements becomes a disadvantage when it comes to the increased sensitivity to misalignment. This is generally not an issue for traditional interferometry, where stable mounting and fine resolution adjustment mechanisms can be used to align and maintain the test surface relative to the interferometer. These methods could be extended to the application of interferometry on the eye. However, a non-invasive measurement of the *in vivo* (i.e. living) eye would not allow for an equivalent degree of securing the subject's head and/or eye. Therefore, it must be accepted that a range of random eye motion will be present during measurements that will result in a misaligned surface with respect to the interferometer.

Eye motion can be classified into two domains: voluntary and involuntary. The most basic voluntary eye motion is gaze, where a coordination of eye and head motion allows a person to fixate on and track an object. This is corrected by having the subject fixate on a stationary target. Head motion is further minimized by placing the subject's head in a restraining fixture. A common ophthalmic fixture used with ophthalmic instruments, such as a corneal topographer, is comprised of a chin-cup and forehead rest (Figure 1.2). One study showed that head motion in a typical ophthalmic headrest was in the range of ±100 μm and strongly correlated to the subject's breathing and heartbeat (Kasprzak & Iskander, 2010).

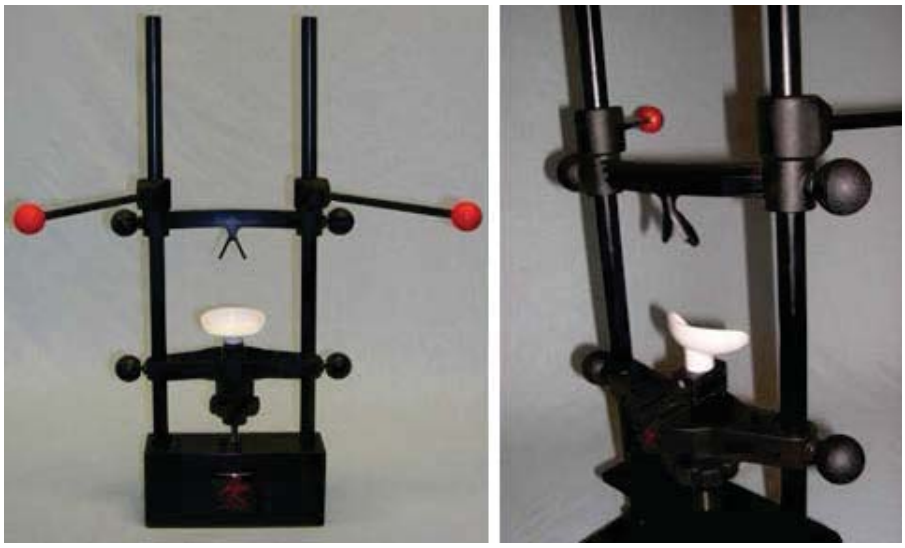

Figure 1.2: Commercial ophthalmic head fixture.
Image source: (Primeau, 2011)

Minimization of gaze is accomplished by directing the subject to lock their visual gaze on a stationary target. Despite the ability of a subject to fixate, the eye will make small, random, involuntary motions (Barlow, 1952; Ratliff & Riggs, 1950). Involuntary eye motion is decomposed into three types of motion: Saccades, drifts, and physiological nystagmus or tremors. Saccades are quick, large flicks of the eye that occur one to three times per second with an average magnitude of 6 arc-minutes (Steinman, Haddad, Skavenski, & Wyman, 1973). Drifts are slow, cyclic motions that occur at a rate of a few arc-minutes/second (Nachmias, 1961). Physiological nystagmus is a small, random, high frequency (50 - 100 Hz) motion with a magnitude slightly greater than two minutes of arc (Steinman et al., 1973). In general, random eye motion is within 3 - 7 arc-minutes of its mean fixation position. The eye's center of rotation is approximately 14.8 mm from the surface of the cornea (Fry & Hill, 1962), which results in 3 - 7 arc-minutes of involuntary eye motion being equivalent to 13 - 30 μm of surface translation.

Previous studies have shown that color, luminance, contrast, or image quality of the fixation target have no effect on fixation except in the case where any of these factors render the target barely visible (Boyce, 1967; Steinman, 1965). However, it was found that differences in fixation target sizes resulted in changes in saccades and drift (Rattle, 1969; Steinman, 1965). External distractions, such as motion in the

periphery or auditory stimuli, were also found to increase the amount of microsaccades (Engbert & Kliegl, 2003; Rolfs, Engbert, & Kliegl, 2005). A non-invasive measurement of the cornea and tear film must therefore accept that a certain amount of motion will be present. If eye motion is left unaccounted for, it will significantly degrade or inhibit the ability of the instrument to measure the cornea and tear film.

1.3 SOFT CONTACT LENSES

A soft contact lens is a thin (70 - 200 μm) hydrogel material that is placed directly on the surface of the eye. A hydrogel is a hydrophilic polymer that is highly absorbent, with the water content of a fully hydrated hydrogel lens ranging anywhere from 24% to 74% (Tighe, 2013). A recent development in the last decade has been to introduce silicone into the hydrogel, which increases oxygen permeability allowing for longer continuous wear. By varying water content, silicone, and other parameters of a contact lens, a manufacturer can adjust the optical characteristics, comfort, and a number of other factors.

When worn, the soft contact lens sits on a cushion of tear film that is approximately 2 μm thick (Nichols & King-Smith, 2003). A pre-lens tear film forms on the anterior surface of the contact lens that has similar function to the tear film on the bare eye. Vision correction is an obvious requirement for the design of a contact lens; however, an equally important requirement is comfort. Stability of the tear film can greatly influence the comfort of the contact lens. The hydrophobic property of the silicone that is added to the hydrogel lenses contributes to an increase in tear film break-up and can reduce vision quality. Yet at the same time, the hydrophobic properties of silicone are desired to increase oxygen permeability for improved comfort. Therefore, it is desirable to have an instrument that is capable of resolving the interaction of the tear film surface with the contact lens to better study and characterize contact lens materials.

1.4 APPLICATIONS OF CORNEAL TOPOGRAPHY

Keratoconus is a degenerative eye disorder that causes a structural weakening of the cornea. The reduced structural integrity of the cornea results in thinning and an outward bulging of the cornea (Figure 1.3). Symptoms may not exist early in the stages of keratoconus, but can later develop into reduced vision quality and eventually a significant loss in vision (Rabinowitz, 1998). Detection generally occurs after vision quality has been significantly impaired, whereas corneal topographers can be used to provide a quantitative and repeatable method for early detection of keratoconus (Maguire & Bourne, 1989).

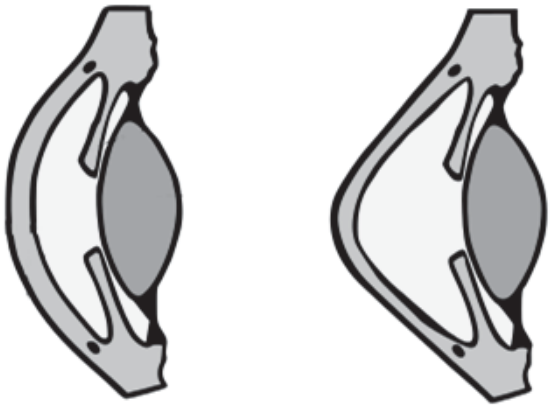

Figure 1.3: The effect of keratoconus on the cornea.
(Left) A normal cornea. (Right) Cornea with Keratoconus.
Image source: (Wikimedia Commons, 2009b)

One of the most common eye diseases is keratoconjunctivitis sicca, more commonly known as dry eye syndrome. Dry eye is a disorder of the tear film characterized by tear deficiency or excessive evaporation, resulting in discomfort and even damage to the corneal surface (Lemp, 1995). It affects 4 - 6 % of the adult population and as much as 15 % in the elderly (> 65 years of age) (Schaumberg, Dana, Buring, & Sullivan, 2009; Schaumberg, Sullivan, Buring, & Dana, 2003; Schein, Muñoz, Tielsch, Bandee-Rohe, & West, 1997). However, dry eye syndrome is difficult to diagnose due to the reliance on a subject's self-reporting of their symptoms and limited clinical measures (Schein, Tielsch, Muñoz, Bandeen-Roche, & West, 1997). Corneal topographers have been demonstrated to provide a quantitative method for diagnosing dry eye (Kojima et al., 2004; Kopf et al., 2008; Lakhbir S Mengher, Bron, Tonge, & Gilbert,

1985). A higher resolution corneal topographer has the potential to significantly improve the detection and characterization of subjects with dry eye syndrome.

Another application of corneal topography is in the field of refractive surgery. Degraded vision quality is commonly the result of an incorrect amount of power within the eye. Vision can be corrected by corrective eyewear, contact lenses, or refractive surgery. Refractive surgery is one of the most invasive methods, which involves modifying the overall power of the eye by reshaping the cornea. A common refractive surgery method is LASIK (Laser-Assisted in situ Keratomileusis), where a laser is used to ablate a layer of the cornea to form a new effective anterior surface with the desired refractive properties for proper vision. Topographic maps of the cornea are used to determine the amount of reshaping that is necessary to improve vision. Improved performance may be gained by high resolution characterization of the dynamic tear film to isolate the tear film structure from the base corneal shape. Pre- and post-characterization of the cornea is also useful in understanding the healing process of the eye (Jaycock, Lobo, Ibrahim, Tyrer, & Marshall, 2005).

A final example for the use of a high resolution corneal topographer is also the motivation for this instrument: Soft contact lens metrology. The value of the contact lens market is placed at about $2.5 billion, with 39.2 million contact lens wearers in the United States and an international market that is estimated to be about 3 times larger (Nichols, 2015). The majority of contact lens wear is vision correction or therapeutic, with approximately 5% cosmetic (e.g. colored lenses). When worn, the contact lens sits on the cornea between a cushion of tear film and an outer-layer of tear film (Nichols & King-Smith, 2003). The role of the pre-lens tear film is nearly identical to the tear film on the bare eye. Vision quality and comfort are strongly influenced by the interaction of the tear film with the contact lens material. The ability to measure the corneal topography and dynamics of the tear film structure with high resolution allows for a more complete characterization of a contact lens in its natural environment. Feedback from the instrument that will be presented in this dissertation can be used to develop and refine contact lens materials and designs that could improve contact lens comfort and vision correction.

1.5 A BRIEF HISTORY OF SOFT CONTACT LENS METROLOGY

For a little over a decade, the focus of this research group has been on the application of interferometry to soft contact lens metrology (Greivenkamp et al., 2014). This has led to the development of a number of instruments for characterizing the optical properties a soft contact lens. One of the first systems was a Mach-Zehnder interferometer for measuring the transmitted wavefront of the contact lens (Williby, Smith, Brumfield, & Greivenkamp, 2003). The next system measured the index of refraction of the hydrogel materials (Goodwin, 2007). And a final system measured the surface profile of both sides of the contact lens (Heideman, 2014). The combination of measurements from these three systems allows for a complete characterization of the optical properties of the contact lens.

However, these systems are limited to testing the soft contact lenses submersed in a saline solution, which is not necessarily representative of their intended use. As discussed in Section 1.3, the contact lens sits between layers of the tear film. Exposure to the environment, interaction with the tear film, variation in corneal topography, and external forces such as blinking, results in a dynamic environment that adds an additional layer of complexity to the design of a contact lens. The next phase in soft contact lens metrology is the ability to measure the contact lenses *in vivo* (i.e. measured directly on the living eye).

The development of an interferometer for measuring the *in vivo* tear film and corneal topography has been accomplished in three phases. The first phase demonstrated an interferometer for measuring the dynamics of a fluid layer on a surface (Primeau, Greivenkamp, & Sullivan, 2011; Primeau, 2011). This was an *in vitro* measurement (i.e. outside of the body) of a simulated tear film layer on a contact lens surface placed on an artificial eye surface. This system has the advantage of being able to characterize the dynamics of a fluid layer surface on a contact lens without the need for a human subject. It also demonstrated that an interferometric measurement the *in vivo* tear film was possible. This led to the second phase, the development of the Tear Film Interferometer (TFI) for *in vivo* measurements of the tear film surface on a contact lens (Primeau & Greivenkamp, 2012b; Primeau, 2011). The TFI was built, but remained untested while waiting for approval to test on human subjects. The first part of this dissertation is a continuation of

the work of Primeau. Once the instrument was approved for testing, multiple human subjects were tested. Feedback from human subjects testing was used to improve the design of the TFI. The successful demonstration of the TFI has led to the third phase, an interferometer for measuring the dynamic corneal topography of *in vivo* eye.

1.6 CORNEAL TOPOGRAPHERS AND TEAR FILM MEASUREMENTS

A review of corneal topographers and tear film measurement methods is necessary to determine the requirements and goals for a next generation corneal topographer. This is not an exhaustive review of corneal topographers, nor does it provide any history on the subject of corneal topography. What is presented is a summary of instruments or methods that are currently in use, represents the current state of the art, or provides necessary background on this subject.

1.6.1 STAINING

A common method for evaluating the tear film involves an eye care professional applying a fluorescent fluid to the eye to aide in the inspection (Figure 1.4). The degree of fluorescence gives a qualitative measure of the tear film thickness and corneal health. However, this method is qualitative and invasive, making it unreliable for proper tear film characterization (S. Patel, Murray, McKenzie, Shearer, & McGrath, 1985). Despite these limitations, it is a simple and inexpensive diagnostic tool for detecting abrasions, ulcers, and infections on the cornea.

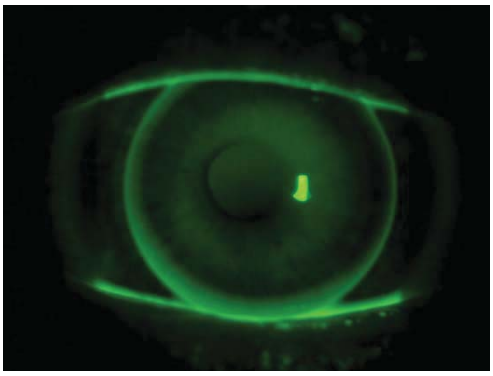

Figure 1.4: Staining of the tear film to evaluate a contact lens fit.
Image source: (Keratoconus NZ, 2015)

1.6.2 REFLECTION-BASED

Reflection based corneal topographers use the specular reflection from the surface of the cornea. An illuminated target is placed in front of the cornea and imaged (Figure 1.5). Comparing the known target to the measured image can be used to determine the surface topography. Evaluation of these devices reveal that surface height accuracies are on the order of 5 μm RMS (>100 μm PV) (Priest & Munger, 1998). The spatial resolution of these systems is 100 - 200 μm (C. J. Roberts, 1994).

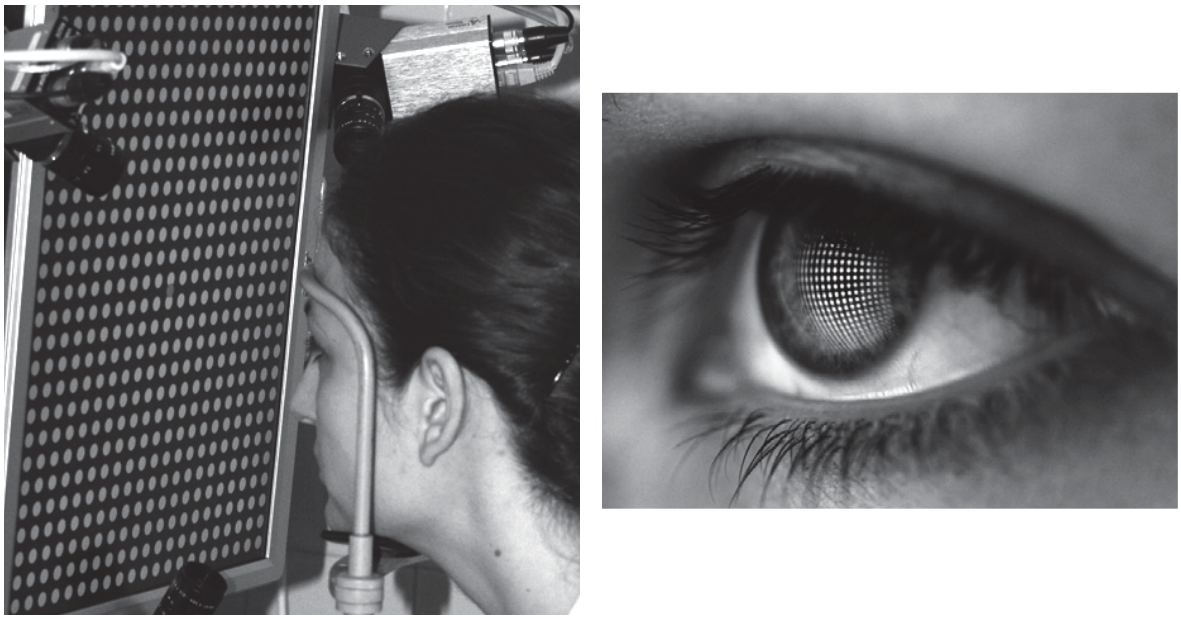

Figure 1.5: Image from a reflection-based corneal topographer.
(Left) Target placed in front of human subject. (Right) Image from the specular reflection of the cornea.
Left and right image source: (Fazekas, Soumelidis, Bodis-Szomoru, & Schipp, 2008)

An arbitrary target can be used; however, it is more common for the target to be structured, such as an array of dots, circles, or squares. The most common pattern is a series of concentric circles that alternate in contrast (Figure 1.6). Systems that use a series of concentric circles are generally referred to as Placido disc systems. Modern computerized Placido disc topographers place a camera at the center and behind the concentric circular targets. The camera captures images of the targets that are reflected off of the cornea and a computer routine can analyze the image to determine the shape of the cornea.

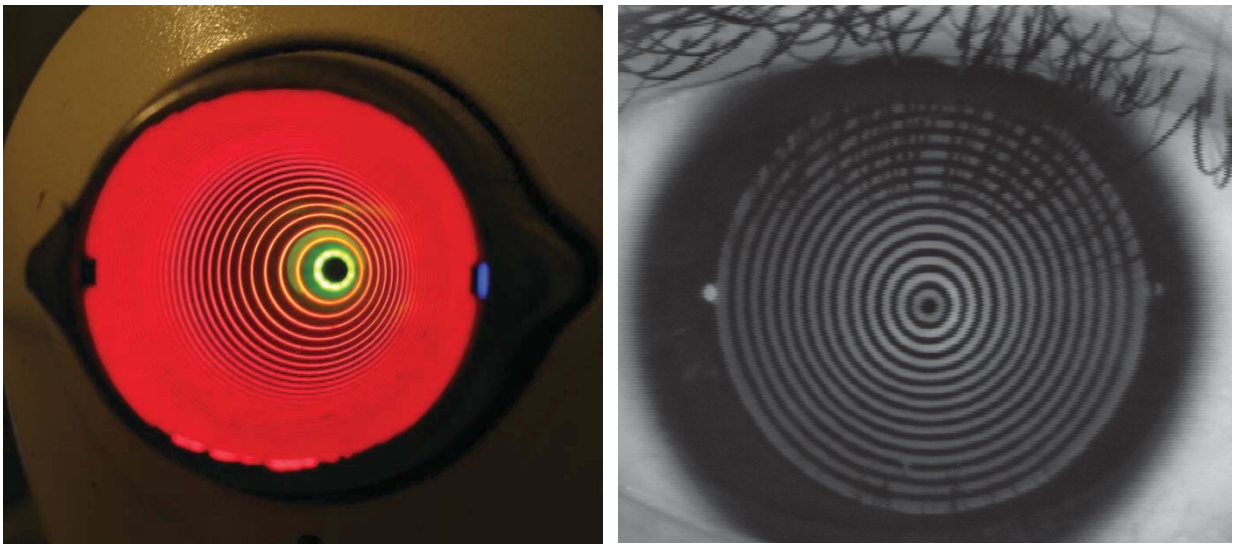

Figure 1.6: Placido disc corneal topographer.
(Left) Image looking into a Placido disc topographer. (Right) Image from a Placido disc topographer showing the reflected mires. Left image source: (Wikimedia Commons, 2009a)

Reconstruction of the corneal topography from the measured image requires four pieces of information: The cardinal points of the optical system, coordinates of target features with respect to the optical system, the corresponding target feature coordinates at the image plane, and the distance from the device to the corneal vertex. The cardinal points and target features are determined from the design of the topographer. A calibration of the instrument can be performed by measuring a known surface to improve the accuracy of these pieces of information. The distance measurement can be accomplished by one of several methods. One of the methods used to determine the distance is by means of an optical trip-line to constrain measurements to a fixed distance. An optical trip-line uses a small source at one end of the instrument and a detector at the other end. The illumination source can be seen in the left image in Figure 1.6 as a faint blue spot on the right side of the targets and the detector is opposite on the left side. The source also appears in the right image on the left side as a white dot. The trip-line works by detecting a transition in the incident power and capturing and image when a certain threshold power is measured. A calibration process will include an accurate determination of the trip-line distance that is necessary for the reconstruction methods. The final piece of information is determined from the captured image of the targets. A software routine can analyze the captured image to determine the corresponding location of the

target features. The target features that are used for Placido disc systems are the transitions between the bright and dark areas of the circle which can be found by an edge detection method.

The topography of the cornea can be recovered from the four pieces of information just described by any number of methods. One method determines the slope of the cornea by calculating the displacement between measured target features (Doss, Hutson, Rowsey, & Brown, 1981). The topography can be reconstructed by integrating across these slope calculations. Other methods use an iterative method to trace rays through the optical system and off of the reconstructed corneal model. The corneal model is perturbed to minimize errors from the ray trace and measured data (Halstead, Barsky, Klein, & Mandell, 1995; J. Schwiegerling, Greivenkamp, & Miller, 1995).

The commercial corneal topography system that will be used in this paper will be a Placido disc system. It is the Keratron Piccolo, manufactured by Optikon (Roma, Italy). The system uses an optical trip-line to determine the distance from the subject's cornea vertex to the device, which is referred to by the manufacturer as an Eye Positioning Control System (EPCS). The manufacturer also claims that the EPCS has the ability to correct for decentration of the cornea. The method by which the topography is constructed is unknown, but it is assumed that the method uses the slope/integration method. The system uses 14 concentric circles for the target, which are designed to be equally spaced when imaged off of a sphere with a radius of curvature of 7.849 mm. The system measures 7,168 points over an 11 mm diameter on a normal eye, which corresponds to an approximate spatial resolution of 116 μm. The system also claims to resolve to ±0.01D (power) or 1 μm (height), but makes no claims of accuracy. These claims will be verified and compared to the corneal topography system described in this dissertation later in Chapter 5.

1.6.3 RASTER-STEREOGRAPHY

Raster-stereographic corneal topography systems use structured illumination that is projected onto the surface of the cornea. These systems differ from reflection-based systems by treating the cornea as a diffuse surface. Surface topography is determined from the measured distortions of the projected image. However, the cornea has very little diffuse reflectance and so a foreign substance has to be added to the tear film to increase scattering. The introduction of a foreign substance alters the behavior of the tear film and invalidates any measurements related to it (L S Mengher, Pandher, & Bron, 1986). Surface height accuracies for these systems has been shown to be 3 μm RMS (> 100 μm PV), with 200 μm spatial resolution (Belin, Cambier, Nabors, & Ratliff, 1995; Priest & Munger, 1998; Tang, Collins, Carney, & Davis, 2000).

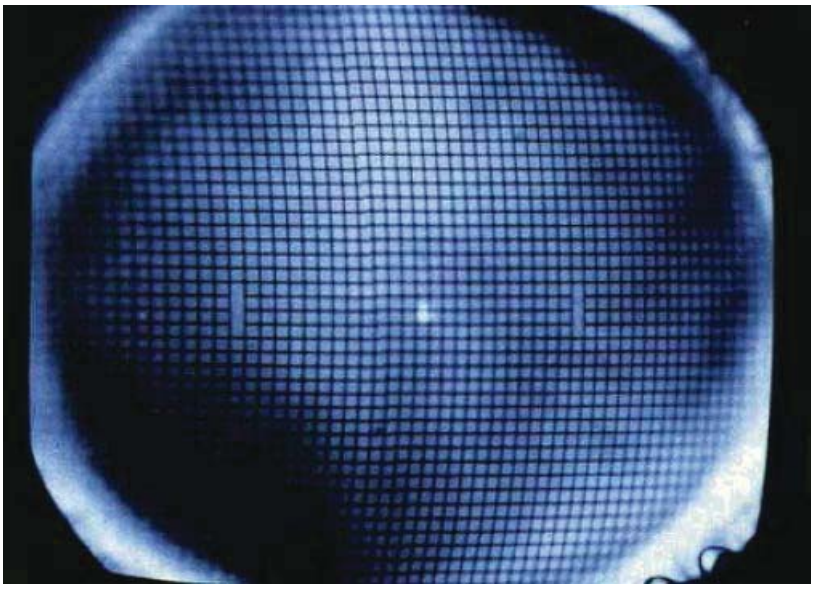

Figure 1.7: Image from the PAR CTS, a raster-stereographic corneal topographer.
Image source: (Belin et al., 1995)

1.6.4 MOIRÉ-DEFLECTOMETRY

Moiré-deflectometry is another form of projection based topography. The difference from the raster-stereographic method is that the structured illumination involves pairs of gratings that are super-imposed on the cornea to produce moiré-interference patterns. This method also requires the installation of a

foreign agent into the tear film to increase diffuse scattering and therefore has the same issues as the raster stereographic system. Surface height accuracies for these systems has been shown to be on the order of 0.5 μm RMS (3.5 μm PV), with 30 μm spatial resolution (Chang et al., 2014; Franciscus H M Jongsma, De Brabander, Hendrikse, & Stultiens, 1998).

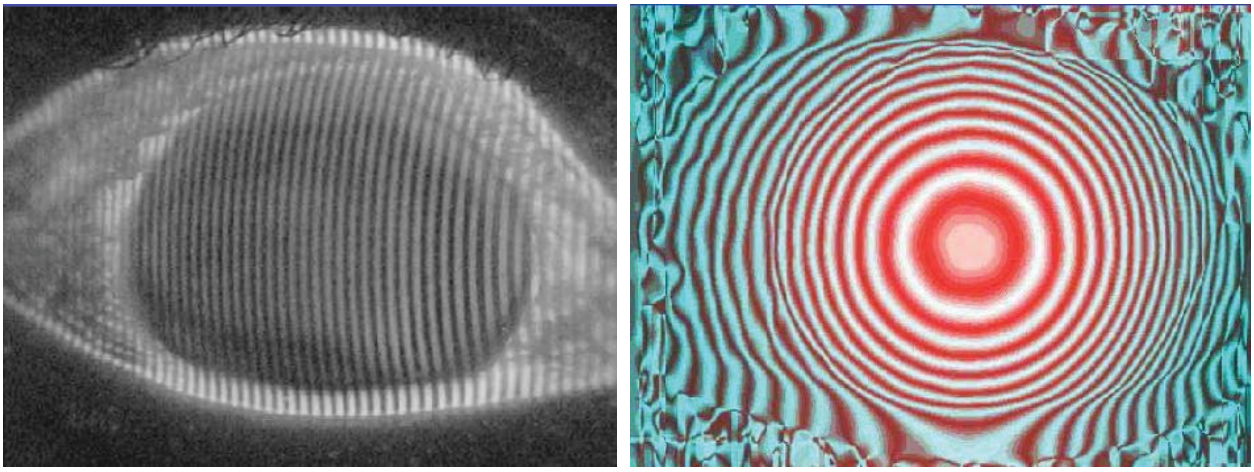

Figure 1.8: Images from Maastricht Shape Topographer, a Moiré-deflectometry topographer. (Left) One of the two sets of projected fringes. (Right) Processed surface height contours. Image source: (F H M Jongsma, Brabander, & Hendrikse, 1997)

1.6.5 SLIT SCANNING

Slit scanning devices can also be classified as a form of projection based topography. A thin slit is projected onto the surface of the cornea and from the measured distortions of the image the topography can be reconstructed (Figure 1.9). The projected slit is imaged from off-axis and is able to measure scattered light from the anterior and posterior surfaces of the cornea, as well as the anterior surface of the lens. By scanning the slit across the cornea, topography of all three surfaces can be reconstructed. Accuracy of these systems has been shown to be 0.2 μm RMS (11 μm PV), with 200 μm spatial resolution (Cairns, McGhee, Collins, Owens, & Gamble, 2002). However, these systems can require up to 30 seconds to acquire a full measurement. The reported surface height accuracies were derived from a stationary surfaces, which means that on eye performance could be significantly reduced by eye motion and tear film dynamics.

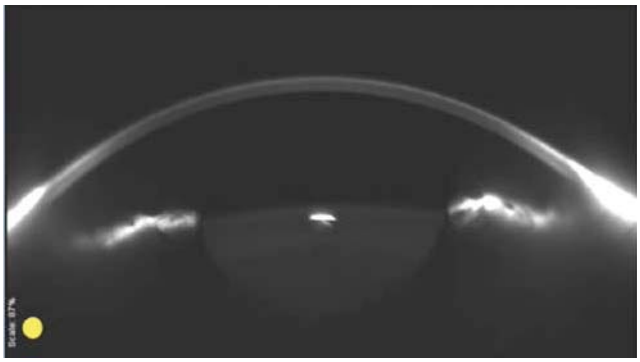

Figure 1.9: Image from GALILEI Dual Scheimpflug Analyzer, a slit scanning topographer.
Image source: (C. J. Roberts & Züger, 2006)

1.6.6 CONFOCAL MICROSCOPY

Confocal microscopy works by imaging a small volume element of a scattering medium; an aperture placed at the image conjugate filters out all light except from the small volume. Light from on-axis that is out of plane with the object is defocused enough that only a small fraction of the light gets through the aperture. By scanning the small volume object across the object, a three dimensional structure of the object can be built up. When used to characterize the surface topography of the human cornea, surface height accuracies have been shown to be on the order of $4 - 25$ μm with spatial resolutions of $0.5 - 2.0$ μm (D. V Patel & McGhee, 2007). While confocal microscopy has the capability of forming high resolution three-dimensional images, the scanning method requires lengthy scan times for larger surface areas and may suffer similar issues as the previous system.

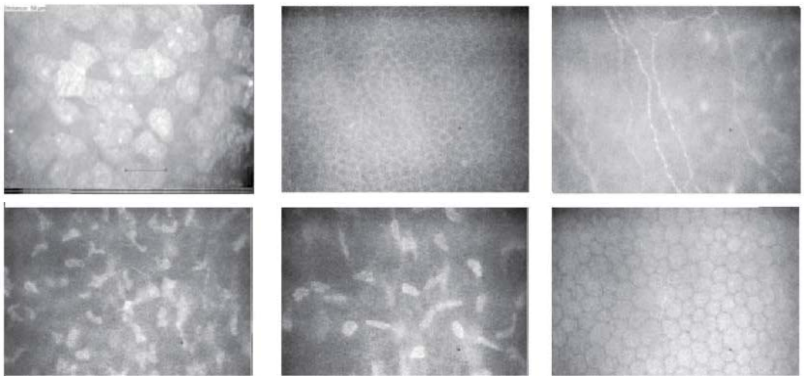

Figure 1.10: Image from a confocal microscopy system.
The images are of the various layers within the cornea.
Image source: (Tavakoli, Hossain, & Malik, 2008)

1.6.7 Optical Coherence Tomography

Optical Coherence Tomography (OCT) has similarities to confocal microscopy but instead uses a low coherence light source and adjustable reference path to scan axially through the medium. The resolution of OCT systems is comparable to confocal microscopy, but also suffers from the same scanning issues (Drexler et al., 2001).

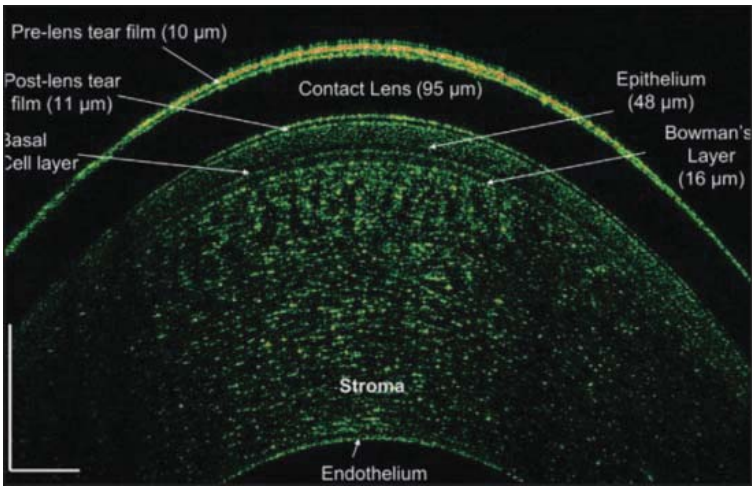

Figure 1.11: Image from an OCT system.
The image shows a cross-sectional view of the cornea with a contact lens.
Image source: (J. Wang et al., 2011)

1.6.8 Holography

Holographic methods require the creation of a reference hologram that has been recorded from either the cornea or a reference surface. Most systems use two wavelength holography to reduce the sensitivity of the system in order to resolve fringes (Lam, Gaskill, & Wyant, 1984). Surface height accuracies of these systems claimed to be as low as 0.25 µm, although it is estimated that at best they could resolve 0.50 µm (Baker, 1990; Friedlander, Mulet, Buzard, & Granet, 1991). However, it appears as though holography has fallen out of use beyond these experimental demonstrations. It is hypothesized that the instrument was too complicated for practical use.

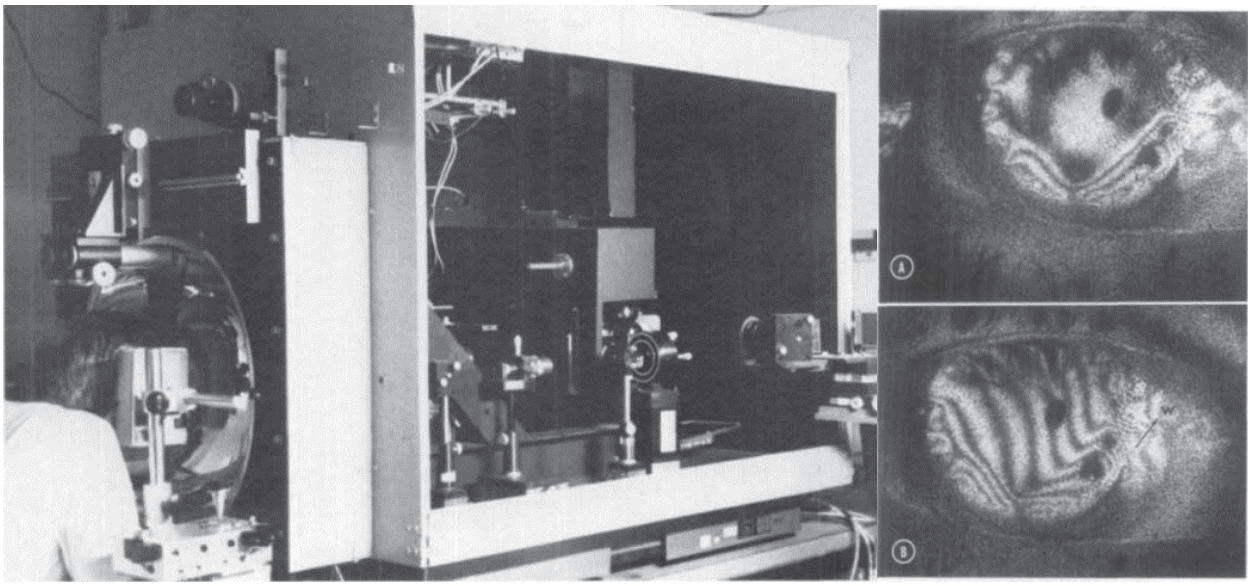

Figure 1.12: Holographic imaging of the cornea.
The instrument is shown on the left with a subject sitting in front of it.
Image source: (Calkins, Hochheimer, & Stark, 1981)

1.6.9 INTERFEROMETRY

Interferometry is a promising technique for measuring the surface topography of the cornea with high resolution. Phase shifting interferometry is capable of providing surface height resolutions < 0.001 μm on optical surfaces (Bruning et al., 1974). Interferometry works by measuring the irradiance of two superimposed wavefronts. The measured irradiance is a cyclical function of the difference in relative phase between the two wavefronts, which is commonly referred to as interference fringes. If at least one of the wavefronts has interacted with the optical element of interest, the relative phase information can be related back to a property of the optical element. For example, if one of the wavefronts has been reflected off of the surface of the optical element, the recovered phase information can be used to determine the shape of the optical surface. This type of measurement can be done with a Twyman-Green interferometer configuration as shown in Figure 1.13 on the left. A Twyman-Green configuration splits a collimated wavefront into two paths with a beamsplitter. One path is reflected (upper path as shown in Figure 1.13) and the other path is transmitted (right path, Figure 1.13). One path is designated as the reference path, which contains a high-quality mirror that reflects the wavefront without introducing errors into the

wavefront. The other path is designated as the test path, where a mirror of unknown quality can be placed. The reflected wavefronts from the two paths are recombined at the beamsplitter and the interference fringes can be observed (lower path, Figure 1.13). The phase that is recovered from the measured interference fringes corresponds directly to the surface figure error.

Although the example in Figure 1.13 shows planar surfaces being tested, the system can be easily extended to test non-planar surfaces. In the test path, a series of optics or an optical assembly can be placed to reshape the collimated wavefront into an arbitrary shape to match the desired shape of the test surface. The test path optical assembly that is used to generate the desired wavefront shape is more commonly referred to as an objective, null optic, converger, or diverger. It has been common to use a spherical wavefront, as most non-planar optical surfaces tend to be spherical. However, the demand for aspheric and freeform surface shapes has been increasing in recent years, requiring more complicated and customized objectives. In the case of the corneal topography system that is described in this dissertation, a Twyman-Green interferometer will be used with a customized objective designed to match the aspheric shape of the human cornea.

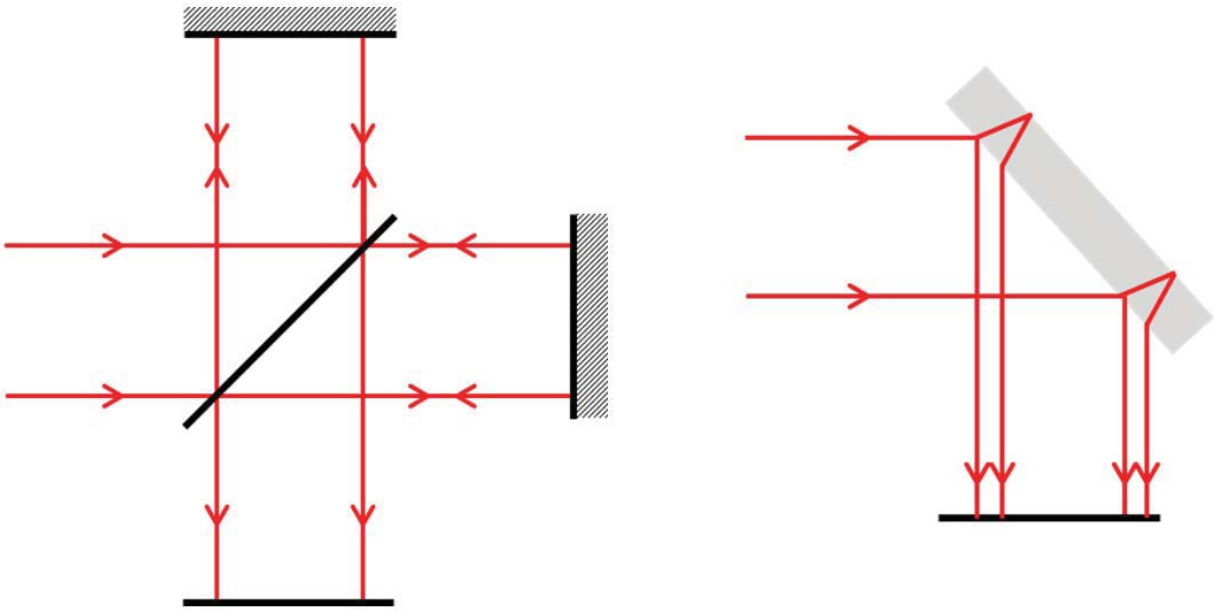

Figure 1.13: Interferometer configurations.
(Left) Twyman-Green. (Right) Lateral shearing.

Another interferometer configuration that is discussed in this section is a lateral shearing interferometer (LSI). The LSI interferes the test wavefront with a shifted copy of itself. The right image in Figure 1.13 shows a common LSI configuration. The test wavefront is incident on a plate surface. The front surface reflects a portion of the wavefront, and the remaining wavefront transmits and reflects off of the back of the surface. The two reflected wavefronts are combined and result in interference fringes. The angle of the plate and thickness produces an offset, or shear, of the secondary wavefront. The secondary wavefront should be a copy of the first wavefront, assuming that the plate has not introduced additional wavefront error. The shear between the two wavefronts results in a set of interference fringes that relates directly to the average slope of the wavefront along the direction of the shear. By increasing the thickness of the plate, the shear can be increased and the average of the wavefront slope will be increased. However, the increased amount of shear will result in an effective reduction in the LSI to resolve height differences. It is for this reason that LSI has been used to measure the cornea, which was to reduce the sensitivity to eye motion by increasing the amount of shear (Section 1.2). Dubra, et al, built a LSI system that measured a 3.4 mm diameter region of the cornea (Dubra, Paterson, & Dainty, 2005). Although they showed that the system could resolve surface height features < 0.01 μm, when the dynamics of the human eye were factored in, surface height errors were shown to exceed 1 μm.

The University of Wrocław has conducted a number of studies using interferometry to characterize the *in vivo* tear film. An initial attempt using a Twyman-Green to directly measure the surface of the tear film was overly sensitive to eye motion and difficult to align (Kasprzak, Kowalik, & Jarański, 1995). Measurements were captured by a framegrabber at 25 FPS over a 4.5 mm diameter area on the eye. Surface topography was calculated through fringe tracing and ordering. Fringe densities were generally too high to evaluate, so analysis was limited to 256 x 256 pixels covering a 2 mm diameter area on the tear film. However, topography could not be determined with complete certainty, due to the fact that fringe tracing results in lower spatial resolution, reduced height accuracy, and ambiguity in the sign of the

topography. No claims are made for surface height accuracy. Spatial resolution is estimated to be 50 - 150 μm. This system appears to have been discontinued in favor of lateral shearing interferometry.

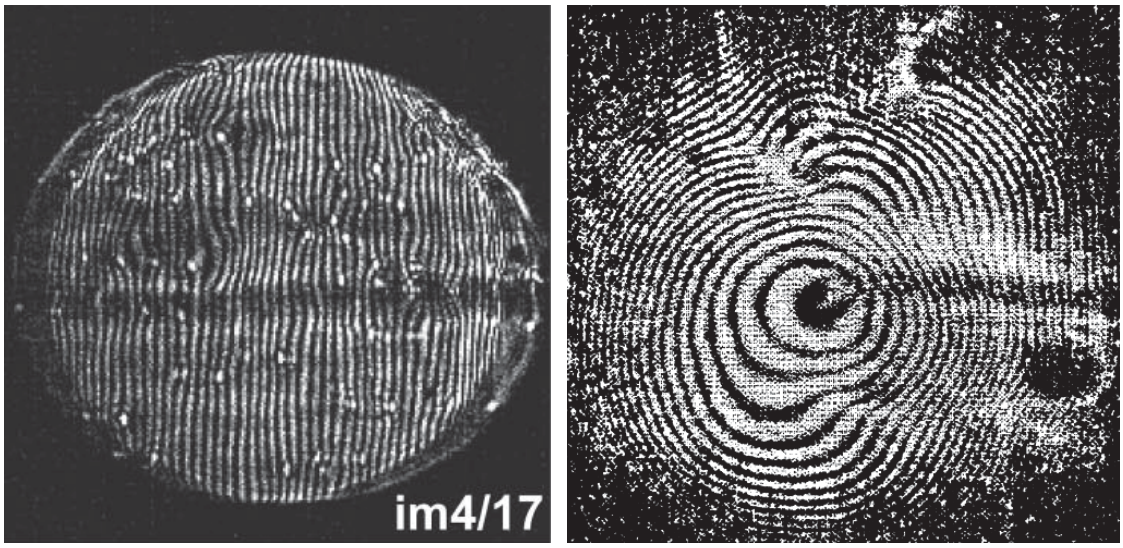

Figure 1.14: Image of the tear film with interferometry.
(Left) Lateral shearing interferometer. Image source: (Dubra et al., 2005)
(Right) Tywman-Green interferometer. Image source: (Kasprzak et al., 1995)

1.6.10 SUMMARY

A performance summary of the different systems that have been discussed is presented in Table 1.1.

Specification	Placido	Raster	Moiré	Slit	Confocal	OCT	Holo.	Interfer.
Coverage [Ø mm]	8-12	> 20	> 20	7	< 1	15	12-20	3-5
Height Accuracy [RMS μm]	5-100	3-6	0.5-5	0.2-2	4-25	2-4	0.5	1.0
Spatial [μm]	100-200	200	30	200	0.5	0.5	25	50
Dynamics [FPS]	25	30	30	< 1	< 1	< 1	?	20-50
Measures topography	✓	✓	✓	✓				
Measures tear film							✓	✓
Non-invasive	✓			✓	✓	✓	✓	✓

Table 1.1: Performance summary of known cornea and tear film measurement systems.

1.7 REQUIREMENTS FOR AN INTERFEROMETRIC CORNEAL TOPOGRAPHER

The top level requirements for the interferometric corneal topographer are:

- Measure a significant portion of the population
- Measure the optically important area on the cornea
- Measure the corneal topography and tear film structure simultaneously
- Measure the dynamics of the cornea and tear film
- Non-invasive measurement
- Measurement resolution at least an order of magnitude, or better, than currently known systems

In addition to the above requirements, it was desired to make use of as many components as possible from the existing Tear Film Interferometer system (Primeau, 2011). This is a budget constraint that places a number of constraints on the design that will influence the performance of the system. The specifications for the system are given in Table 1.2 and discussed in the following sections.

Specification	
Corneal coverage	6 mm
Height accuracy Low frequency Mid/High frequency	 5.0 μm 0.05 μm
Spatial resolution	6 μm
Dynamics FPS Duration	 30 Hz > 60 s

Table 1.2: Interferometric corneal topographer specifications.

1.7.1 CORNEAL COVERAGE

The system will cover a 6 mm diameter area on the nominal human eye. The coverage is fixed by the converging optics that will be reused from the TFI system (Section 2.1). An increase in coverage would require a new converger design, which is the second most expensive optical component in the system. In addition to the converger, a new sensor would have to be purchased to account for the increase in coverage, which is the most expensive component in the system. That is because the amount of eye motion will remain constant regardless of the test area on the cornea, yet the sensitivity to motion will

increase proportional to the increase in test area. The dynamic range of the interferometer will have to be extended to accommodate the increase in sensitivity, which is accomplished by increasing the sampling at the detector (i.e. a larger format sensor).

1.7.2 SURFACE HEIGHT ACCURACY

The summary of corneal topographer capabilities (Table 1.1) showed that the best performing system that could potentially meet our requirements was the reflection-based (e.g. Placido disc) system. Systems such as the raster-stereographic or moiré-deflectometry systems can be excluded because they are non-passive, requiring a foreign substance to be added to the tear film. The remaining systems are excluded because they do not provide any shape information. Therefore, the goal of our corneal topographer should be to exceed the capabilities of the Placido disc system for absolute surface shape information.

However, there is an issue with defining a single surface height accuracy requirement. An absolute surface height accuracy requirement can obscure relative or high spatial-frequency surface height information. For example, previous interferometric measurements of the tear film showed sub-micron surface height accuracies, but contained no information about the topography of the cornea. Whereas a corneal topographers can have relatively good height accuracy, but when a best-fit eye shape is subtracted there is no information about the tear film structure. A better definition for a height accuracy requirement needs at least two domains: low and mid-to-high spatial-frequency height accuracies. A low spatial-frequency height accuracy requirement captures the ability of the system to measure the topography; the base radius of curvature, asphericity (i.e. aspheric conic constant), and corneal astigmatism. Residual surface features are captured by the mid/high spatial-frequency requirement, which covers the ability of the system to resolve the tear film structure.

1.7.3 SPATIAL RESOLUTION

The system will resolve a 6 x 6 μm area on the nominal human eye. Re-use of the converger and 1 MP sensor system as discussed in Section 1.7.1 determines this requirement.

1.7.4 DYNAMICS

The system will capture surface topographic measurements at a rate of 30 frames per second (FPS). This requirement is fixed by the re-use of the sensor system. This framerate will be sufficient to resolve most of the notable dynamic features in the tear film based on previous testing with the *in vitro* and *in vivo* systems (Micali, Greivenkamp, & Primeau, 2015; Primeau & Greivenkamp, 2012a). Little emphasis is placed on capture duration in literature. Therefore, the requirement of 60 seconds was based on experience with the TFI, which seemed to be a reasonable time to capture the tear film dynamics.

2 TEAR FILM INTERFEROMETER

The Tear Film Interferometer (TFI) is an instantaneous phase-shifting Twyman-Green interferometer for measurement of the *in vivo* tear film. As previously discussed in Section 1.5, the TFI is the second phase in the development of a system for measuring a dynamic fluid surface, and the first implementation of an interferometer for *in vivo* measurements of the tear film. Detailed discussion on the development of the TFI can be found in the references (Primeau & Greivenkamp, 2012b; Primeau, 2011). However, a lengthy approval process for human subjects testing left the TFI in an incomplete state.

This research picks up where the work of Primeau concluded. Human subjects testing was eventually approved and the TFI was validated against human subjects. Feedback from human subjects testing has led to numerous improvements in the design of the TFI system and has been used to support the development of the interferometric corneal topographer system that will be discussed in the remaining chapters of this dissertation. This chapter discusses the final configuration of the TFI, a number of lessons learned while testing, and presents results from human subjects testing.

2.1 INSTRUMENT DESCRIPTION

The TFI is a polarization-splitting Twyman-Green interferometer designed to instantaneously measure the wavefront reflected off of the anterior surface of the tear film. A diagram of the TFI system is shown in Figure 2.1 and a photograph of the completed system is shown in Figure 2.2. The laser source is a near-infrared solid-state laser (λ = 785nm) and is modulated by an Acousto-Optical Modulator (AOM) that is synchronized to the camera's electronic shutter (refer to Section 2.3). A continuously variable neutral density filter (CVND) is used to adjust the laser power level to an eye safe level at the output of the interferometer. Following the laser shutter, a power meter continuously monitors the laser source to ensure eye safe laser levels are maintained at all time. The coherence length of the laser is approximately 1 m, requiring optical path matching of the reference and test arms to maximize fringe visibility at the detector (refer to Sections 2.5 and 3.2.1).

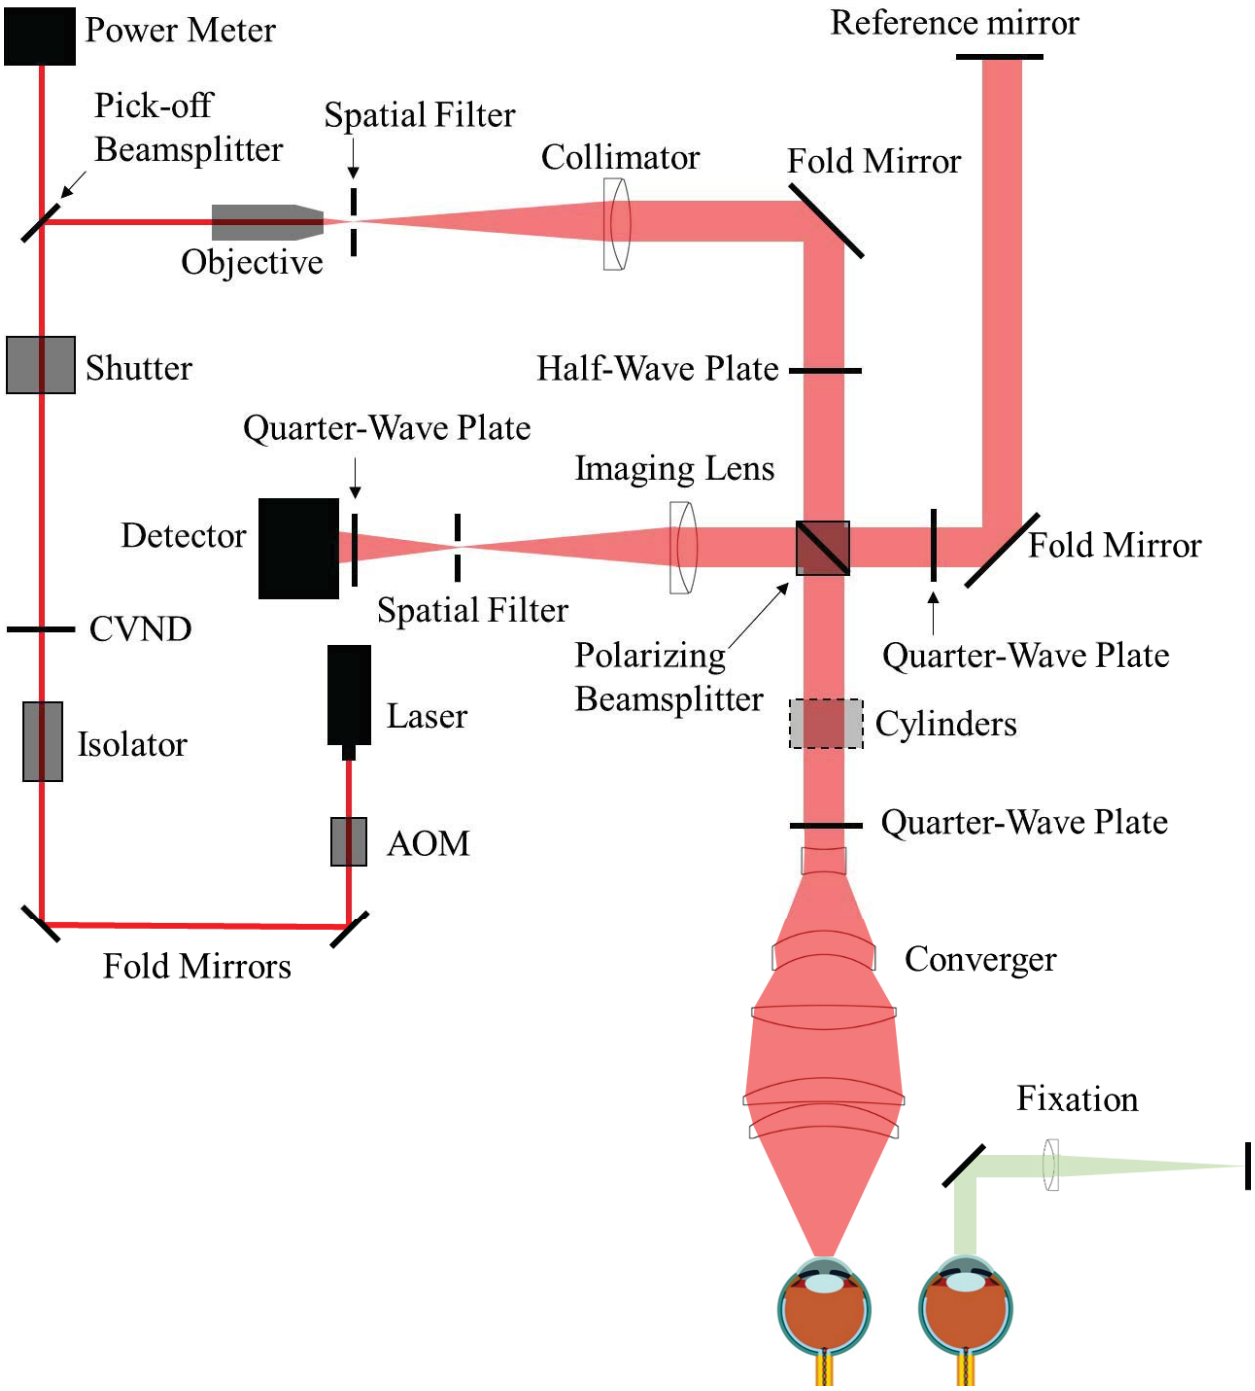

Figure 2.1: Layout of the Tear Film Interferometer.

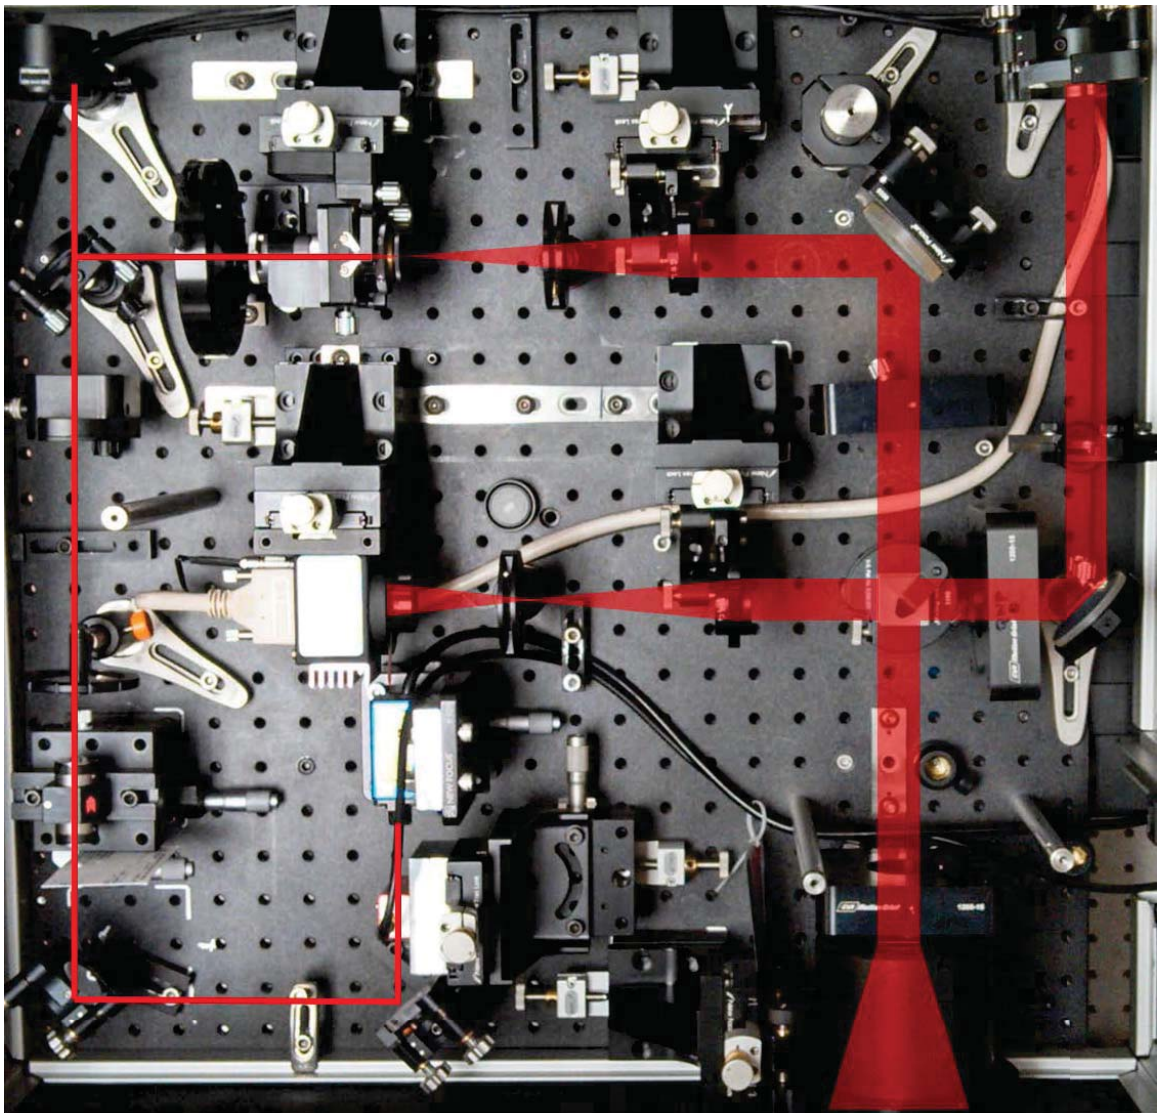

Figure 2.2: Image of the TFI hardware with optical path overlaid in red.

The test path contains the converger; a custom built lens assembly (Photon Gear, Inc., Ontario, NY) designed to generate a wavefront to match the surface shape of the human cornea (Figure 2.3). The converger is designed to null against an asphere having a base radius of 7.8 mm and a -0.25 conic constant (Section 1.1). The lens operates at f/1.3, covering a 6 mm diameter zone on the average human eye. A Stokes lens, or a pair of crossed cylinders, can be placed in the test arm to compensate corneal astigmatism. These are typically left out to minimize stray reflections resulting from one of the four surfaces, but may be necessary for subjects with larger amounts of corneal astigmatism. The reflected

wavefront from the tear film surface is interfered with the reference wavefront and detected by a 1MP Pixelated Camera Kit (4D Technology Corporation, Tucson, AZ). The camera uses a pixelated phase-mask aligned to the detector array that acquires four phase shifted interferograms in a single camera frame (Millerd et al., 2004). Phase is recovered from an individual frame by a spatial convolution method that is discussed in detail in Section 4.1.1.

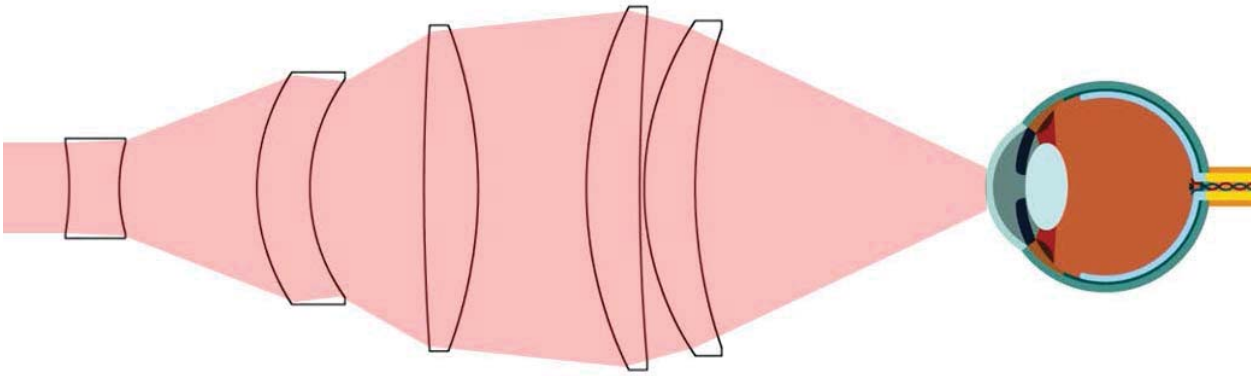

Figure 2.3: Converger lens with eye model.

A test subject will sit in front of the TFI and place their head in an ophthalmic headrest. The headrest consists of a chin cup, horizontal crossbar with a nose pad to lean their forehead against, and a pair of guiding rods to support the head at the temples. A view looking into the TFI with the headrest is shown in Figure 2.4. The TFI is mounted on top of a pair of linear translation stages that provide fine alignment of the device to the subject. A vertical stage is placed under the headrest for a third degree of fine alignment of the test subject to the interferometer. Eye motion is minimized by using off-eye fixation that is placed external to the system (green path, Figure 2.1). An illuminated target is placed at the focus of a 250 mm focal length lens to present a distant image to the subject. The fixation path is folded approximately 50 mm from the subject to fit the fixation target and optics within the space between the subject and the interferometer. Once the subject fixates on the target, the TFI is aligned to the non-dominant eye and measurements are made.

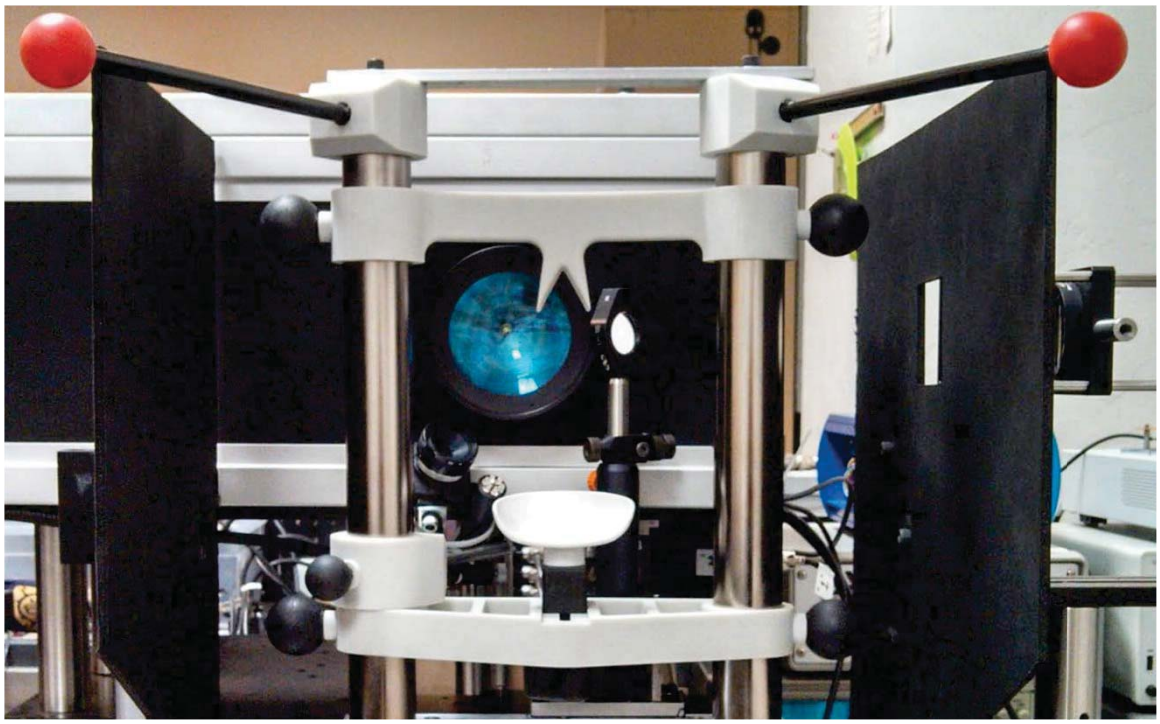

Figure 2.4: View looking into TFI system.
Configuration shown will test on subject's left eye with the fixation set for the right eye.

2.2 STABILITY

Significant eye and head motion will be present in measurements made with the TFI (Section 1.2). This motion will add significant wavefront error to the measurements which can potentially exceed the dynamic range of the instrument, resulting in loss of data and difficulty in aligning the instrument. The functional dynamic range for the 1 MP pixelated camera was measured by the manufacturer to be approximately 170 waves/radius (corresponding to a 22.2 mrad slope error measured at the tear film surface) (Millerd et al., 2004). Dynamic range can be impacted by vignetting, which is discussed in more detail in Section 2.4.

It was predicted that for the duration of collection (i.e. multiple sequential frame captures) wavefront slope introduced by eye motion, head motion, and varying ocular geometry, would be 105 waves/radius (13.7 mrad slope error) (Primeau & Greivenkamp, 2012b). It should be noted that it was assumed that all motion was centered about a mean position to which the TFI could be aligned and that surface slope

errors from the tear film structure could be ignored. From the predicted motion and dynamic range expectation of the sensor, it was estimated that at least 50% of the captured data captured should be useable. Usable was defined in this instance to be where slope errors did not exceed 100 waves/radius.

However, the first few tests with human subjects revealed significant issues with the TFI. Alignment of the instrument to test subjects was extremely difficult, sometimes taking minutes. Once aligned, only a few seconds of data could be captured with less than 30% of the data being usable. The issue was that motion far exceeded all previous predictions and was severely limiting the functionality of the TFI. The motion could be attributed to two sources: Subject stability in the headrest and their inability to maintain sufficient fixation. Although both of these motions were accounted for in the initial design, the sensitivity of the design to these motions was underestimated. This required an iterative approach to redesigning both the headrest and fixation assemblies, which are discussed in the following sections.

After the redesign was implemented, a trained operator could align the TFI to a subject in less than a half minute. Additionally, the quantity of valid data from a collection increased from less than 30% to over 90%. Furthermore, there has been no determinable limit to the duration of a data collection, with acquisition times up to 120 seconds having been repeatedly demonstrated. Duration is ultimately limited by subject fatigue and their ability to comfortably remain still in the headrest while maintaining fixation.

2.2.1 HEAD STABILIZATION

The TFI head stabilization fixture (shown in Figure 2.4) is a modified reproduction of a commercial head positioning system from Arrington Research (Scottsdale, AZ). The commercial system was found to have insufficient stability for this application, so the modifications involve strengthening critical components. The current fixture now consists of a pair of vertically mounted 1 ft. tall, 1.5 in. diameter, dampened steel posts, securely mounted to a base plate. The chin cup and forehead rest were rapid prototyped out of glass-filled nylon and tripled in thickness from the original parts. An aluminum rod is also used to connect the two posts along the top to further constrain lateral motion. The pair of 0.5 in. black guiding

rods from the original device has been re-used; they are not critical to the over-all stability of the fixture, but are still necessary to maintain subject stability. These guiding rods do not physically constrain the subject's head, but act as a physiological reminder for the subject to hold their head still. At any point during the test the subject can back out of the headrest with little effort. Finally, the layout was adjusted to place the subject's head to rest as close to the base of the fixture as possible, reducing the significant moment arm that previously existed. A force gauge verified that the stiffness of the new fixture has been increased by over 160 times the previous setup.

The head stabilization fixture is mounted to a vertical stage to allow for fine alignment of the TFI to the subject. The ideal configuration would be to place a vertical stage under the interferometer, such that all three degrees of adjustment would be with respect to the subject. During the design of the TFI it was decided to re-use stages from a previous setup as a cost savings. The vertical stage cannot support the full weight of the interferometer, but can support the weight of the head fixture and a human head. Unfortunately, the vertical stage is the weakest support in the overall head fixture setup. However, the overall stiffness of the setup, including the head stabilization fixture and stage, is still a 15 times improvement in stiffness over the initial design.

2.2.2 FIXATION

Although it was discussed in Section 1.2 that eye motion was a substantial fraction of motion, it was assumed that fixation occurred under a number of ideal conditions. It was evident that in addition to head stabilization, improvements in fixation could be made. The first change was to the fixation target and size, as previous studies have shown that differences in fixation target sizes resulted in changes in saccades and drift (Rattle, 1969; Steinman, 1965). More recent studies have concluded that the optimal fixation target shape to minimize random eye movement is a set of cross-hairs surrounded by a circular target with a dot in the center (Thaler, Schütz, Goodale, & Gegenfurtner, 2013). External distractions, such as motion in the peripheral vision or auditory stimuli, were also found to increase the amount of microsaccades (Engbert & Kliegl, 2003; Rolfs et al., 2005). There is an important distinction between the TFI fixation

and the systems described in literature: The TFI uses monocular off-eye fixation, whereas the systems described in the literature used monocular on-eye fixation.

In the original implementation, the fixation target was an image of the Old Main building on the University of Arizona campus (Figure 2.5, first from left). The experiments of Boyce and Steinman, et al, were replicated on the TFI to verify the optimal choice of fixation targets (Boyce, 1967; Thaler et al., 2013). The experiment was modified to account for the off-eye fixation used on the TFI, such that eye tracking system would monitor the eye under test while the subject performed monocular fixation with the opposite eye. A commercial eye tracker (Arrington Research, Scottsdale, AZ) is attached to the side of the converger system and images the eye at a 45º angle. Subjects were placed in front of the TFI in a normal test configuration, with the exception that the TFI laser was disabled. Subjects were then instructed to fixate on a series of different targets (Figure 2.5) for up to 2 minutes. Additionally, subjects were tested with different permutations of illumination schemes: Red, green, blue, and white illumination; low or high illumination; and constant or modulated illumination. A total of 6 subjects were tested with the eye tracking system.

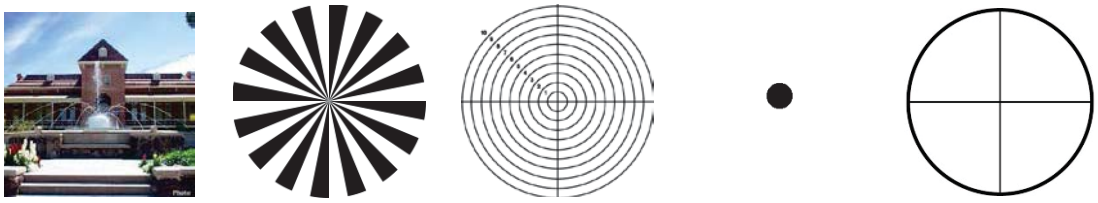

Figure 2.5: Fixation targets tested on the TFI.

The results were in good agreement with the previous tests of Boyce and Steinman; the largest contributor to eye motion was the target shape. Target performance is ordered in Figure 2.5, with the worst starting on the left (image of the Old Main building) to the best performing on the right (cross-hairs with circle). The improvement gained by switching from the image of the building to the crosshairs was a decrease in random eye motion by 50%. The distributions shown in Figure 2.6 are from a single test subject, but are

representative of every subject tested. Although more difficult to quantify due to the limited data, it appears as though saccades were also reduced by 50% by choice of the cross-hair target.

An interesting aside, although Steinman reports no variation for illumination color, he did not test the color blue (Steinman, 1965). The current test results showed an almost 100% increase in eye motion for targets that were illuminated with a solid blue light. Otherwise, all other color schemes were in agreement with Steinman, producing negligible variation in eye motion. When given a choice, most subjects chose a constant green illumination. Therefore the implemented illumination system uses a fixed, constant green illumination scheme with the cross-hair target.

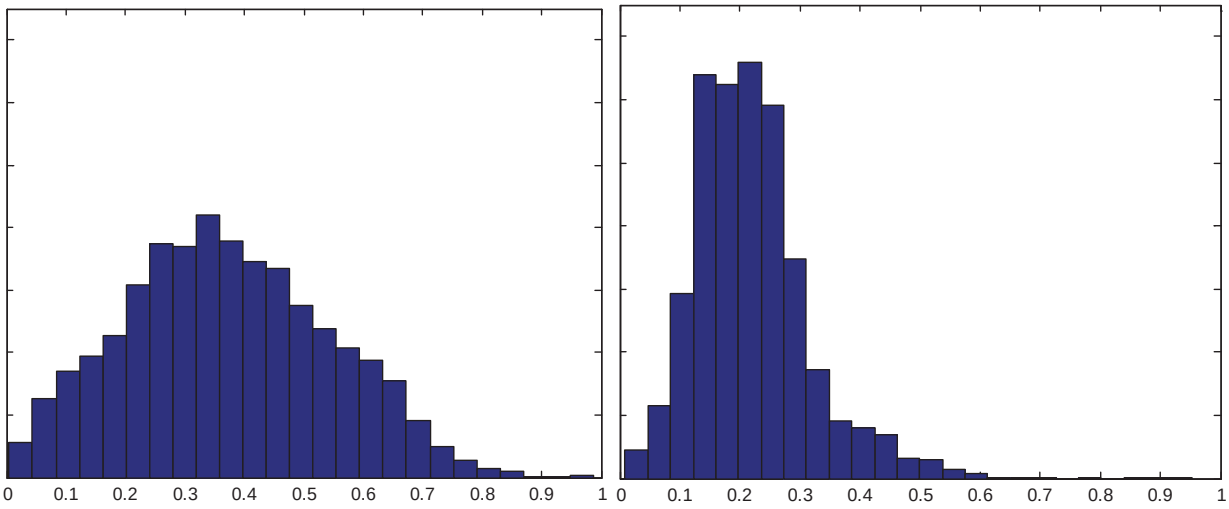

Figure 2.6: Eye motion during target fixation.
Histogram shows radial deviation from mean position. (Left) Fixation on image of building. (Right) Fixation on crosshairs with circle. Gaze direction was not calibrated for absolute positions; therefore the horizontal axes have no units.

The next correction to the TFI was to minimize external distractions, which have been found to increase the amount of microsaccades (Engbert & Kliegl, 2003; Rolfs et al., 2005). This was accomplished by adding baffling around the subject's head to minimize peripheral distractions. Instrumentation displays, such as the power meter, were easily visible by the subject when sitting in front of the instrumentation. The baffles were cut from matte black foam board and can be seen in Figure 2.4 to the left and right of the headrest. The front panel of the TFI enclosure was originally a specular black plastic panel that a subject

could see directly in their line of sight during a test. This panel was replaced by a diffuse black panel. Although no empirical data from the TFI has been collected on the use of baffles, it is expected to provide a benefit based on the results from previous studies.

2.3 SOURCE MODULATION

Source modulation was not implemented in the original design of the interferometer. It was thought that the 785 nm wavelength laser source was sufficiently below the human visual response that modulation would be unnecessary. However, every test subject has reported being able to see the 785 nm source. Subjects report that the source appears as a large, faint, dark red spot that fills their vision and is filled with “smaller blobs” that were either static or moving across their field of vision. It is speculated that the static blobs were a combination of dust particles within the interferometer and floaters in the subject’s eye. Other structures were reported to move downward, which would be consistent with the observed tear film dynamics: Upon blinking, the tear film would be pulled upwards by surface tension and the upper lid. The perceived downward motion of the structures is a result of an inverted image due to the source coming to focus inside the eye. Because the laser source subtends 45 degrees of the subject’s field, resulting in significant peripheral distraction, visual perception of the source must be minimized to guarantee stable fixation during testing.

The first attempt at correction was to reduce the laser power and increase camera gain. However, even at the lowest usable levels, approximately 50 μW or 1.8 μW/mm^2 at the surface of the cornea, the source was still perceptible to most subjects. Therefore it was decided to modulate the laser source. The laser was modulated by placing an acousto-optic modulator (AOM) directly after the laser source with the modulation synchronized to the camera’s electronic shutter. The nominal camera frame rate operates at 33 ms with camera integration times limited to 100 μs, resulting in less than 0.5% duty cycle of laser source illumination. Once the laser source was modulated, subjects reported no longer being able to see the laser source.

Modulating the laser source provides an added benefit: a reduction in detector blooming and a subsequent increase in fringe modulation. An example of the added blooming suppression that resulted from modulating the source in use is shown in Figure 2.7. The detector used in the 4D pixelated camera (Truesense Imaging KAI-1010) specifies blooming suppression to be greater than 100X, which is sufficient for nominal operation. However, the large fringe densities that are measured at the detector, which are a result of eye motion, ocular variation, misalignment, and tear film dynamics, are significantly reduced in modulation by detector sampling (Gaskill, 1978). By modulating the source and reducing blooming, fringe modulation can be improved by as much as 10%.

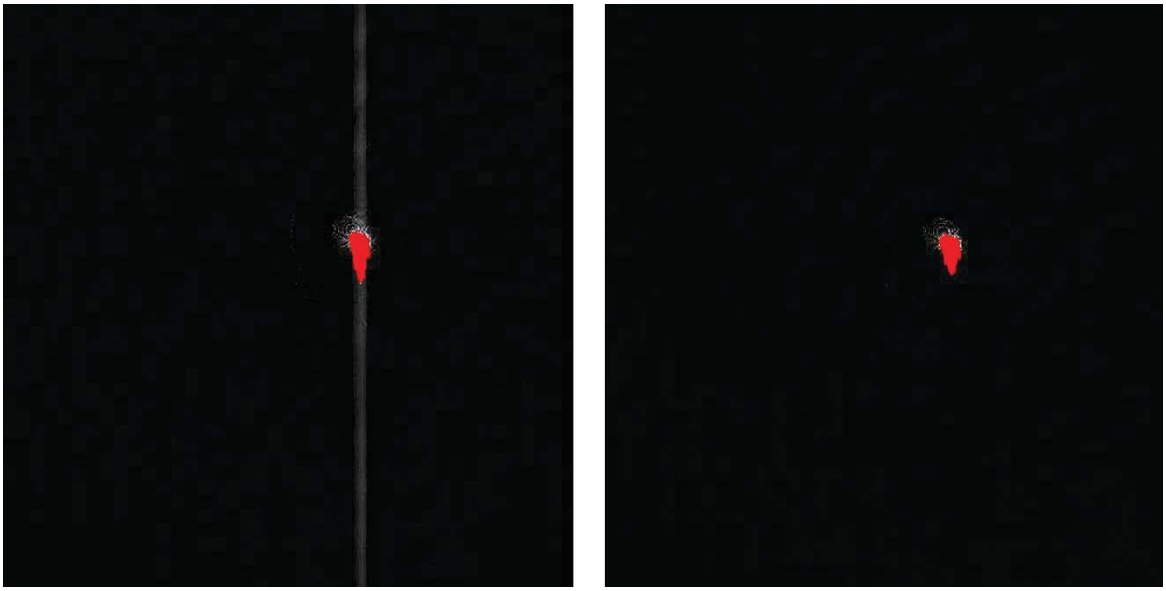

Figure 2.7: Detector blooming without source modulation (left) and with modulation (right).

2.4 VIGNETTING

Vignetting occurs when an aperture other than the system aperture blocks part or all of a beam bundle. In an interferometer, the field of view can be considered to be small or non-existent and apertures are sized accordingly. This is rarely an issue with traditional interferometry, where parts are measured in a null or near-null configuration and little consideration needs to be made for “off-axis” beams. However, measurements made with the TFI differ from traditional interferometry in that eye motion and ocular variation will result in most measurements being non-null. These non-null measurements will result in

vignetting and a loss of information if left unaccounted for. In general, vignetting will limit the dynamic range of the interferometer.

The unfolded test path of the TFI is shown in Figure 2.8, with the detector on the left and the corneal surface on the right. Ray bundles are launched from three normalized field positions at the detector surface, 0, 0.7, and 1.0 field points. The detector area is 9.6 x 9.6 mm, so the 1.0 field position corresponds to a 4.8 mm off-axis field point. The aperture stop is not used to define the ray bundles, as commonly done in a classical optical system design. Instead, the effective chief-ray for each field position is determined by the reference path: Rays are traced from the source, through the reference path, to the detector. A null measurement would require that the test rays are co-incident with the reference rays, so the test-path chief-rays are defined as coincident to the reference rays. The maximum ray angle deviation from each ray is defined by the functional dynamic range of the camera sensor; 170 waves/radius, which equates to 29.4 mrad at the detector. Applying the allowable ray angle deviations to the nominal chief ray angles produces the test ray bundles, which are then launched from the detector surface towards the cornea surface model.

The results of the ray trace are shown in Figure 2.8, with the red, green, and blue rays representing the 0.0, 0.7, and 1.0 field points, respectively. Three areas are blown-up to show locations of possible vignetting. The first is the imaging lens (bottom-left) which shows that the 1.0 field points (blue) will begin to vignette for large deviations of the eye surface. There may be additional vignetting from the mechanical components used to mount the imaging lens which are not shown in the model. The next surface is the polarizing beam-splitter cube (bottom-right). Although it would appear that the 1.0 field point is potentially vignetting, this assumes a perfect alignment of the polarizing beam-splitter cube in the interferometer. The system has no sensitivity to translational misalignment of the polarizing beam-splitter and it can be expected that vignetting is occurring here as well. Additionally, the clear aperture of the polarizing beam-splitter is specified to 90% which may introduce additional complications.

Finally, the largest potential contributor to vignetting is the converger (top). It should be noted that the exact clear apertures for the lenses in the converger are unknown. The model uses clear apertures provided by the manufacturer's lens prescription. The manufacturer of the converger assembly was unable to provide any additional information when asked. However, if it was assumed that the modelled clear apertures are representative of the as-built system, then significant vignetting will occur at the edge of field for non-null alignments of the test surface.

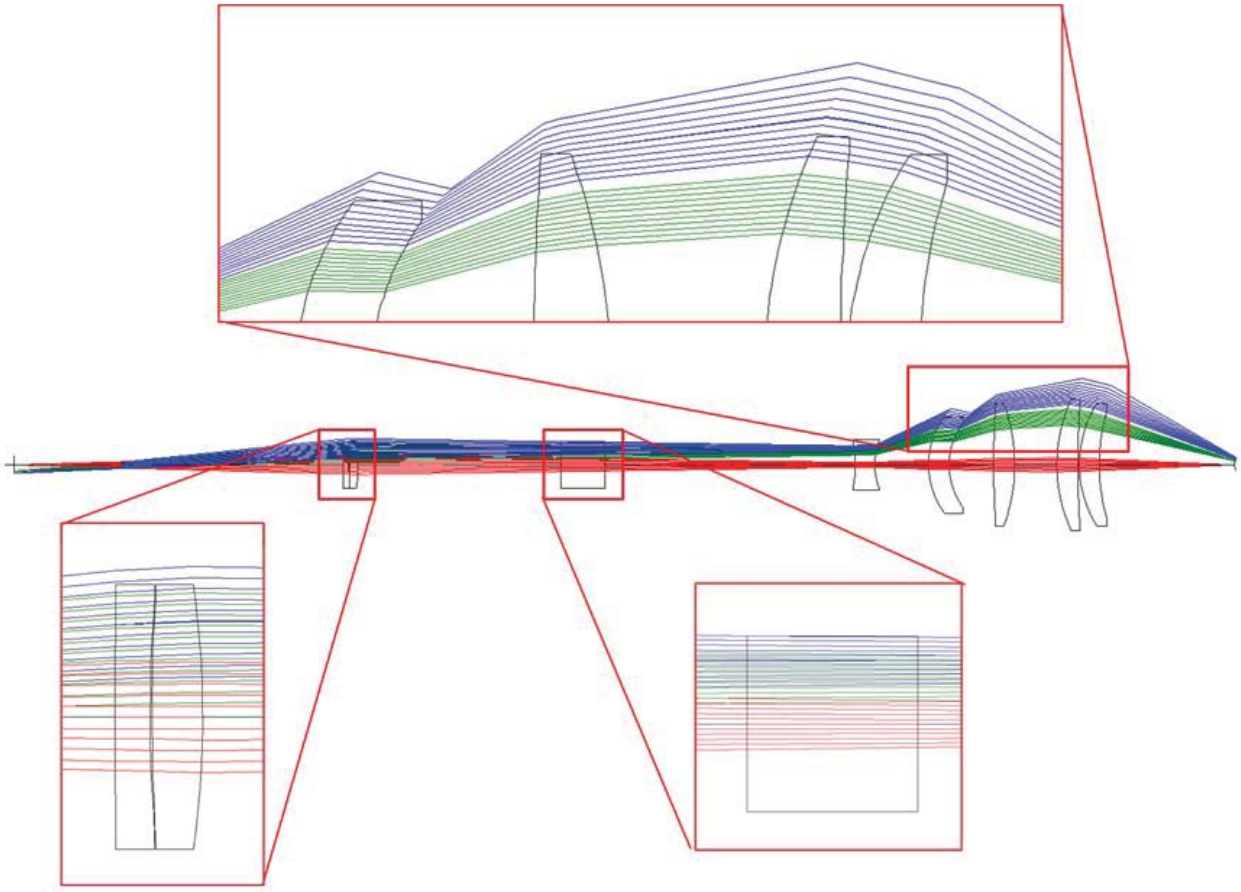

Figure 2.8: Potential vignetting sources in the TFI.

Vignetting is not a significant issue for this system, but should be carefully considered. The above analysis shows that roughly 10-15% of a measurement by area may be lost for extreme motions of the eye. This analysis is supported by measurements that have been made with the TFI: The edges of data are clipped and their location randomly varies frame-to-frame as a function of the random eye motion.

Measurements made with a test surface are shown in Figure 2.9. The image shown on the left was made near null and fills the full circular aperture of the display. The same test surface was laterally translated by 75 μm and defocused by 100 μm, resulting in the measurement shown on the right. The top portion of the measurement is clipped due to vignetting. When analyzing data, vignetted regions of data could be ignored, but it alters the definition of "usable" data. For instance, the data may still be useful if 80% of the data was present. Vignetting will be addressed in the new corneal topographer system and attempts will be made to minimize it (Section 3.2).

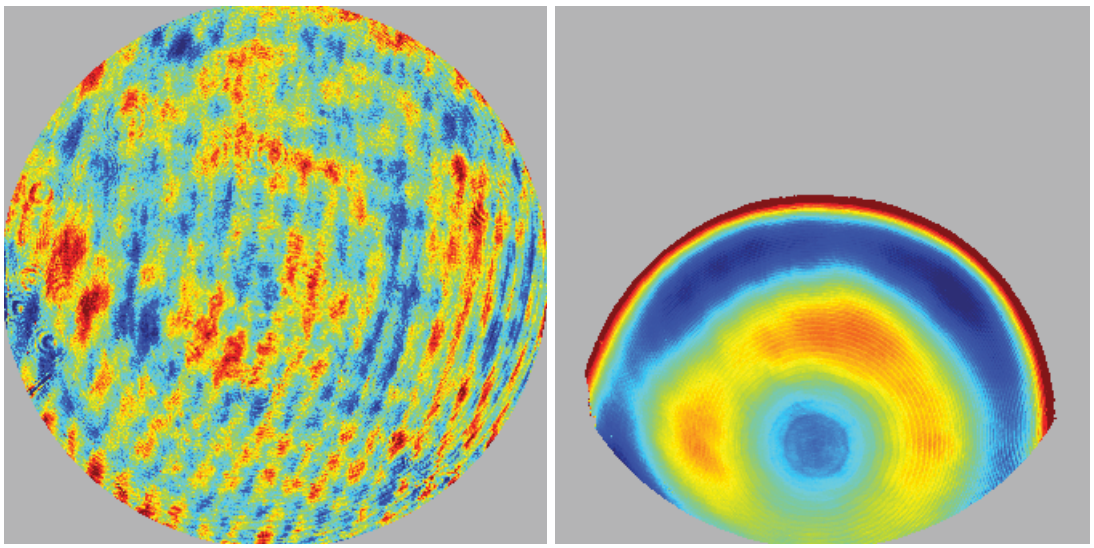

Figure 2.9: Example of vignetting occurring in a TFI measurement.
(Left) Null measurement. (Right) Test part is laterally decentered and defocused.

2.5 SOURCE COHERENCE

Early measurements made with the TFI would occasionally result in large drop-outs of data. These drop-outs were a result of a loss of fringe modulation, or visibility. Modulation loss was traced back to the laser source (Innovative Photonics I0785SH0100B-TH-L), which was found to be relatively unstable. The source is an external cavity laser (ECL) and is extremely sensitive to temperature drifts and mechanical instability (Rudder, 2006). The laser is temperature controlled by use of a thermoelectric cooler (TEC) and the TFI is not a source of mechanical shock or vibration. It is suspected that the set-point of the TEC would drift over a period of a few days to weeks, depending on laser usage and lab conditions. This small

drift was sufficient enough to destroy the coherence of the laser source. A more detailed discussion on source coherence and examples of reduced modulation can be found in Section 3.2.1.

The first correction was to re-adjust the TEC set-point anytime a drop in fringe modulation was noticed (refer to Section 3.2.1). Once the set-point was adjusted, the coherence length could be measured. One method to evaluate the coherence length is to measure the fringe modulation as a function of the optical path difference. A typical coherence length measurement is shown in Figure 2.10. The converger assembly was removed for this measurement to avoid issues with dispersion. A plane mirror was used in the test path and translated along the collimated interferometer beam. At each translation point, the modulation and path difference was measured. Six measurements were made for the example in Figure 2.10 and a Gaussian curve was fit to the data (red curve).

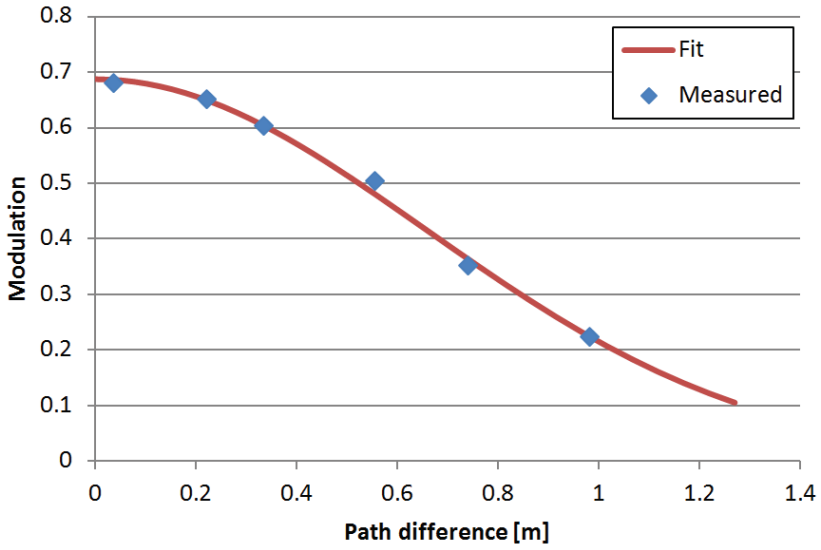

Figure 2.10: Measured coherence length of TFI laser source.

The Gaussian function used for the fit is given by the form:

$$V = Ae^{-\frac{l^2}{2\sigma^2}} \qquad (2.1)$$

For a Gaussian curve, the coherence length (ℓ_c) is defined at the $1/e^2$ point. The coherence length for this measurement was $l_C = 0.77\ m$.

The coherence as a function of path difference is directly relatable to the spectral function of the source through a Fourier transform and therefore we can assume that the laser spectrum is Gaussian (Goodman, 1985). This would result in a spectral bandwidth given by:

$$\Delta\nu = \frac{c}{l_c}\sqrt{\frac{2}{\pi}} = 310\ MHz \qquad (2.2)$$

And a spectral linewidth of:

$$\Delta\lambda = \frac{\lambda^2}{\Delta\lambda}\sqrt{\frac{2}{\pi}} = 0.64\ pm \qquad (2.3)$$

The vendor specifications of the linewidth is $< 100\ pm$, therefore the laser appears to be operating within specification. However, as an added precaution, the TFI was modified to path match the test and reference arms as best as possible. Although re-adjusting the TEC set-point and verifying modulation is sufficient to get the TFI operational, it is not considered optimal and will cause problems in the future. This issue is addressed in the design of the corneal topographer system (Section 3.2.1).

2.6 DATA ANALYSIS

One of the unforeseen issues with interferometric measurements of the tear film was data processing and analysis. The *in vitro* system successfully utilized a commercial software package (4Sight, 4D Corporation, Tucson, AZ) to capture and process interferometric data. An initial attempt was made to reapply the same methodology that was used on the *in vitro* system to the TFI, however 4Sight was unable to successfully process most of the interferometric data collected on the tear film. One example is shown on the left in Figure 2.11, where numerous artifacts are present in the processed wavefront. The

difficulty results from the high fringe densities and low signal-to-noise that result from eye motion, ocular variation, and tear film structure. The solution was to develop custom software routines to process measurements that had been made on TFI. As a comparison, the same data processed with the custom software is shown on the right in Figure 2.11. The same custom software will be used on the corneal topographer system and is discussed in more detail in Section 4.2.

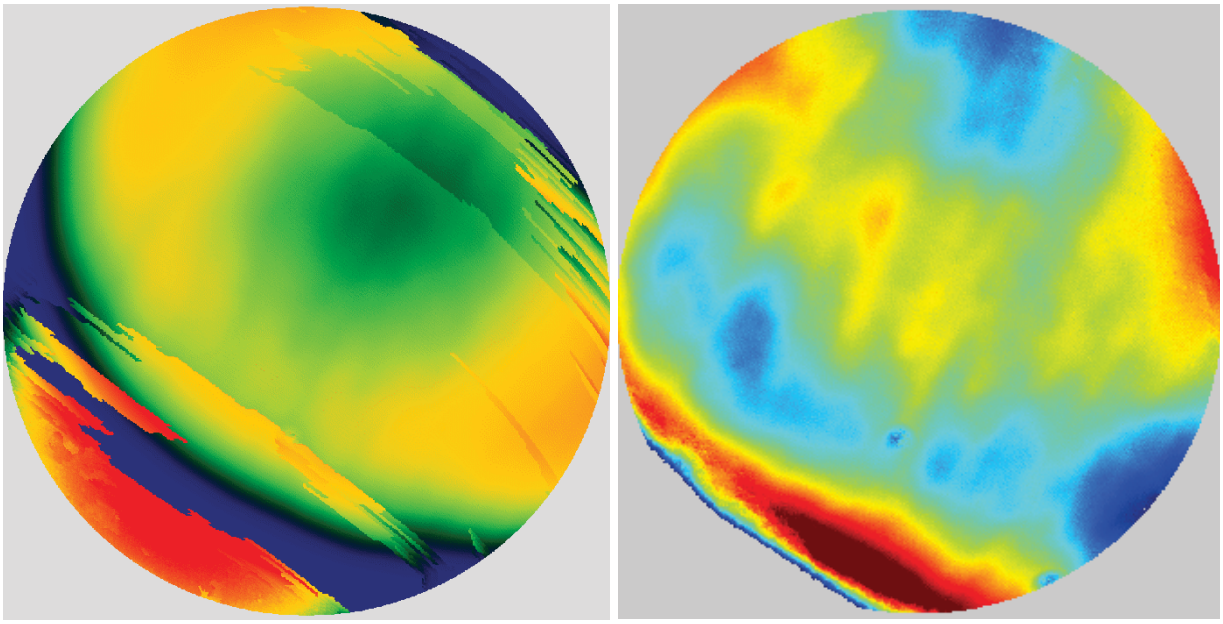

Figure 2.11: Example of phase unwrapping breakdown with TFI data.
(Left) Data processed with commercial software package. (Right) Processed with custom software.

2.7 RESULTS

The TFI was approved for human subjects testing by the University of Arizona Institutional Review Board and adheres to the tenants of the Declaration of Helsinki. Informed consent is obtained from a subject prior to testing. Testing was performed on healthy adults, aged 18 to 60 years. Subjects were not allowed to participate if they had any known eye disease or have undergone refractive surgery. Subjects outside of this age range or with an eye disease may introduce potential complications that inhibit the TFI system's ability to collect data.

Nine human subjects were examined using the TFI. A trained operator can align the TFI to a subject in under a minute. Acquisition times up to 120 seconds have been repeatedly demonstrated with no appreciable change in data quality. Subjects will sit an average of 5 minutes in the TFI before requiring a small break; the TFI will have to be realigned to the subject upon resumption of testing. Test protocol limits the total time with a subject to 60 minutes. However, subject fatigue becomes noticeable after 20-30 minutes at which point the test is ended. In a typical data set, approximately 90% of the frames captured will be useable. The other 10% of data is lost due to transient events such as large saccades or drifts that move the cornea out of interferometer capture range, saccadic motion that is faster than the camera exposure time, blinking (i.e. eyelid and eyelashes obstruct the cornea), loss of fixation that occurs around a blink, and rapid axial motion of the eye during a blink (Doane, 1980). The reconstructed surface topography has an average lateral resolution of 6 μm, surface height resolution better than λ/30 or 25 nm, and 35 ms temporal resolution of dynamic surface features. The following two sections show tear film measurements that were made on bare eyes and eyes with contact lenses.

2.7.1 BARE EYE RESULTS

This section presents results from human subjects testing on the bare eye. Movie sequences for some of the data that is presented in this section can be found in the publication: "Dynamic measurement of the corneal tear film with a Twyman-Green interferometer" (Micali et al., 2015). What is presented here instead is a subset of these movies, or sequential frames of measurements that demonstrate the spatial, height, and temporal capabilities of the TFI.

The first sequence of images in Figure 2.12 shows the measured interferogram on the left and the final processed surface information on the right. An approximately 6 mm diameter region on the cornea is displayed. The color-height scale is set to a peak-to-valley (PV) scale of 0.75 μm (±0.375 μm); red is away from the surface of the cornea, blue is towards the retina. The topography of this subject required that the first 16 Zernike polynomial terms to be removed to provide an optimal display range. For TFI data, only high spatial-frequency variation of the tear film surface was examined. Removal of the 16

Zernike terms allows the high spatial-frequency variation to be seen. The top pair of images in Figure 2.12 were taken a few seconds after a blink has occurred and the tear film has had time to stabilize. The smaller circular bubble features that are approximately 200 μm in diameter are mucus globules. The second set of images were taken approximately 1 second after a partial blink had occurred. A partial blink results in the valley formation that appears in the lower third of the image. The topography of the partial blink bears strong resemblance to tear film thickness measurements that have been reported in literature (Braun, King-Smith, Begley, Li, & Gewecke, 2015). The final set of images were taken approximately 1 second after a second partial blink and results in the valley formation near the center of the image.

Additional examples are shown in Figure 2.13 through Figure 2.17 and demonstrate various phenomena. The interpretation and significance of these examples are beyond the scope of this dissertation, but are presented to demonstrate the capabilities of this instrument. Figure 2.13 demonstrates high spatial-frequency structure that was exhibited by a subject. Figure 2.14 is from another subject that exhibited low spatial-frequency drying structure. Figure 2.15 shows a subject that had a large number of mucin globules present in their tear film. The evolution of the tear film in Figure 2.15 demonstrates that the gradient in surface tension across the tear film results in an upward drift of the tear film surface (King-Smith, Fink, Hill, Koelling, & Tiffany, 2004). Figure 2.16 and Figure 2.17 are examples from two different subjects of the tear film drying over a period of 8 and 45 seconds, respectively.

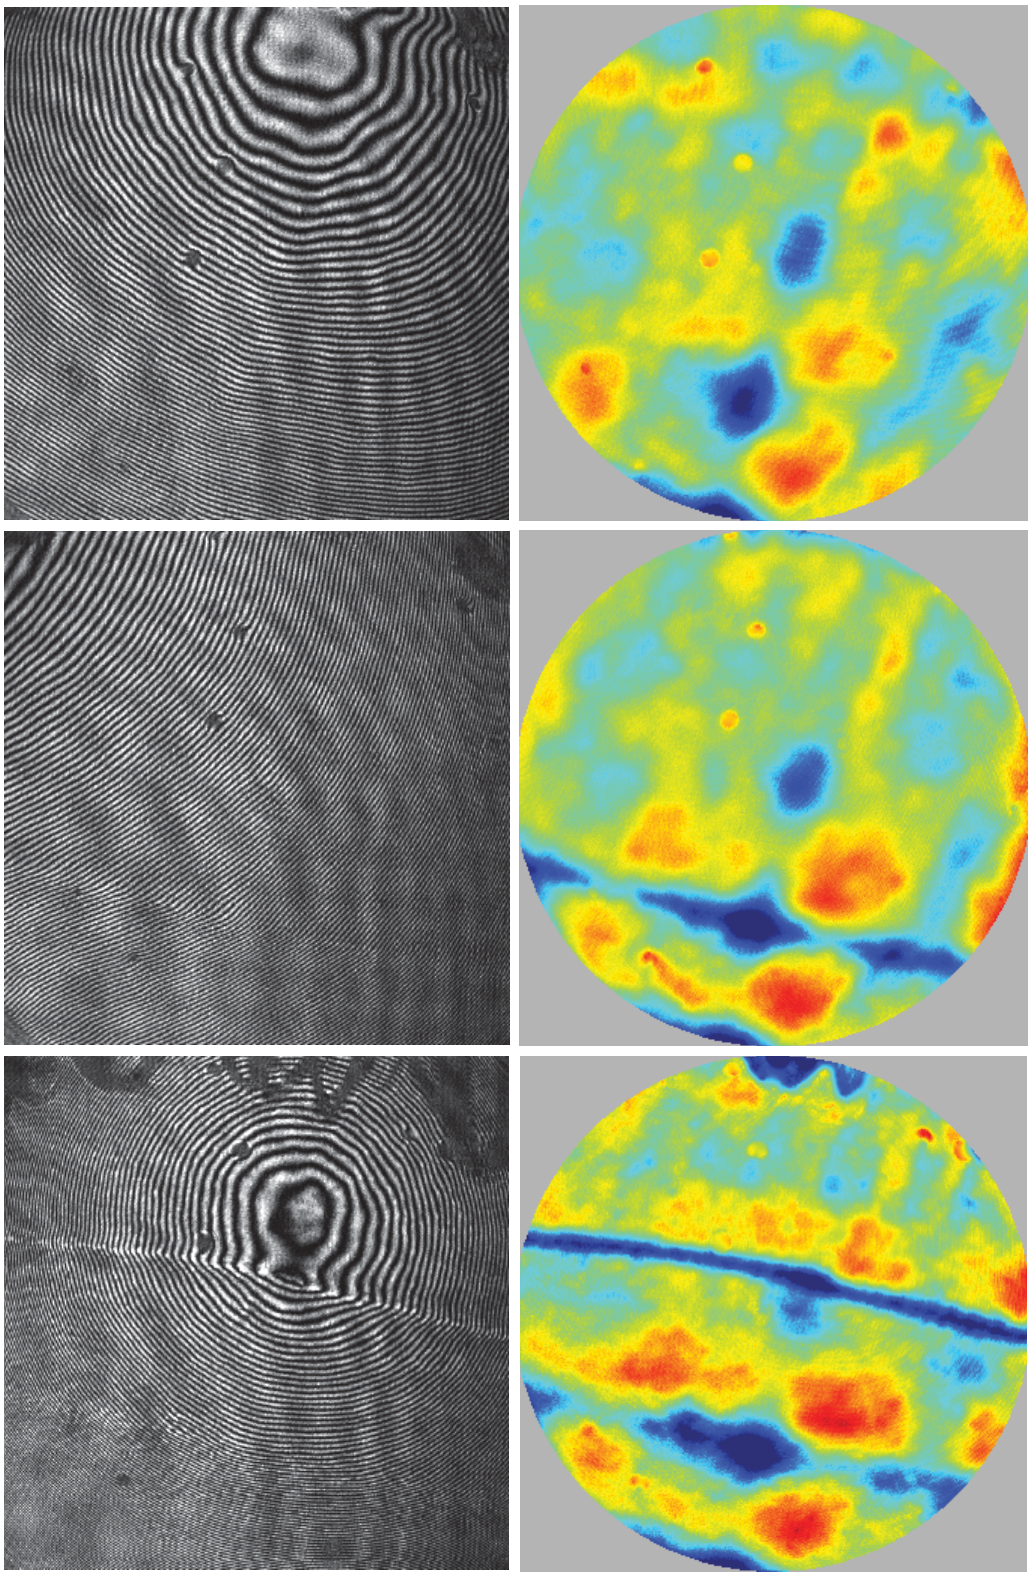

Figure 2.12: Comparison of fringes (left) to surface data (right) with a partial blink structure. 6 mm diameter displayed. ±0.375 µm color height scale; red is away from surface of cornea. First 16 Zernike terms subtracted. First still is with stable tear film (top). Second still is within 1 s after a partial blink (middle) and third still is within 1 s after a following partial blink (bottom).

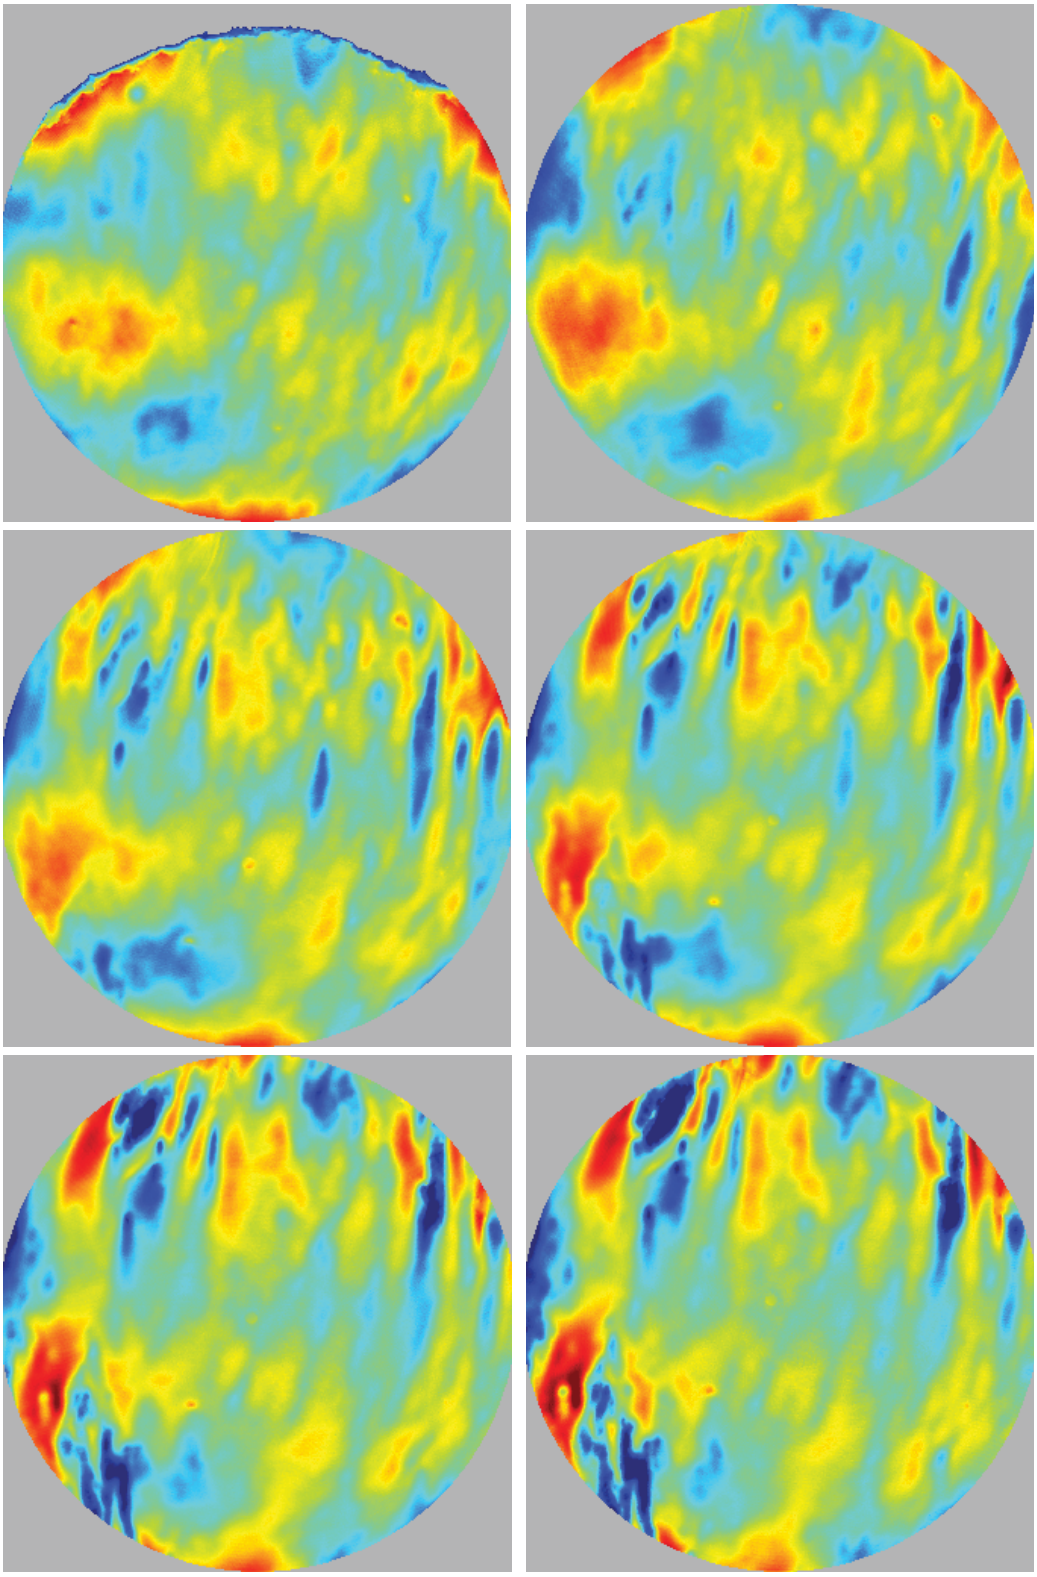

Figure 2.13: Evolution of tear film on subject exhibiting high spatial-frequency drying. 6 mm diameter displayed. ±0.75 μm color height scale; red is away from surface of cornea. First 16 Zernike terms subtracted. First still occurs ~100 ms after a blink. Each still is separated in time by 1 s.

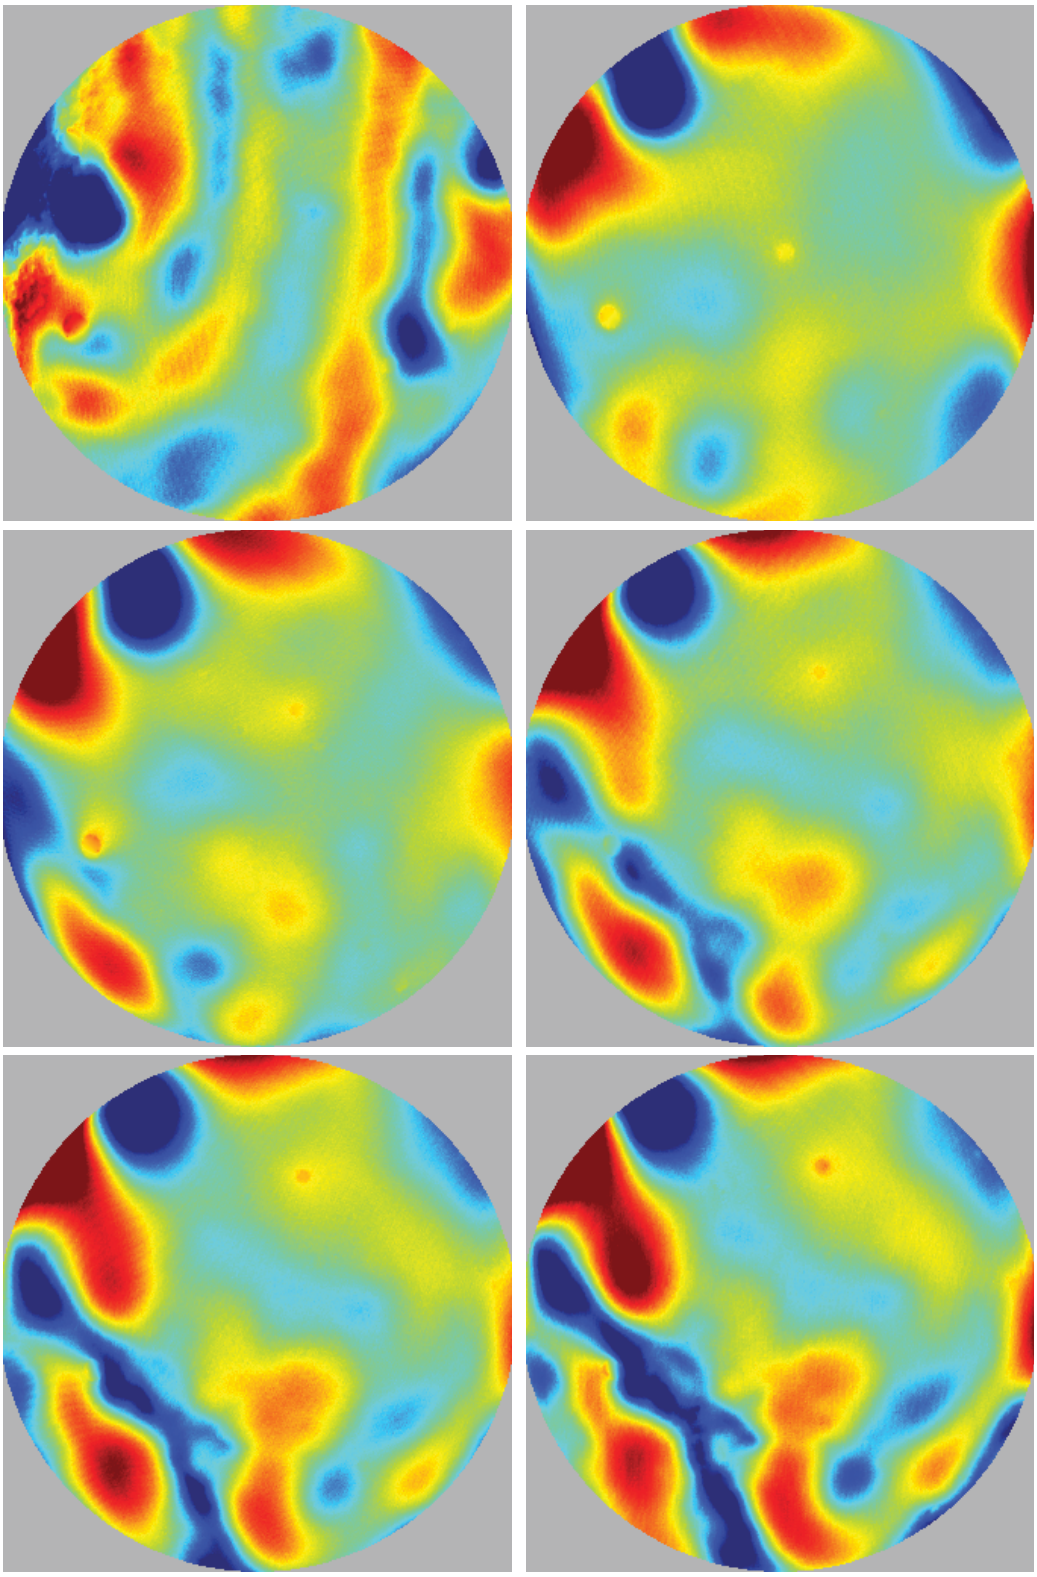

Figure 2.14: Evolution of tear film on subject exhibiting low spatial-frequency drying. 6 mm diameter displayed. ±0.75 μm color height scale; red is away from surface of cornea. First 16 Zernike terms subtracted. First still occurs ~100 ms after a blink. Each still is separated by 2.25 s.

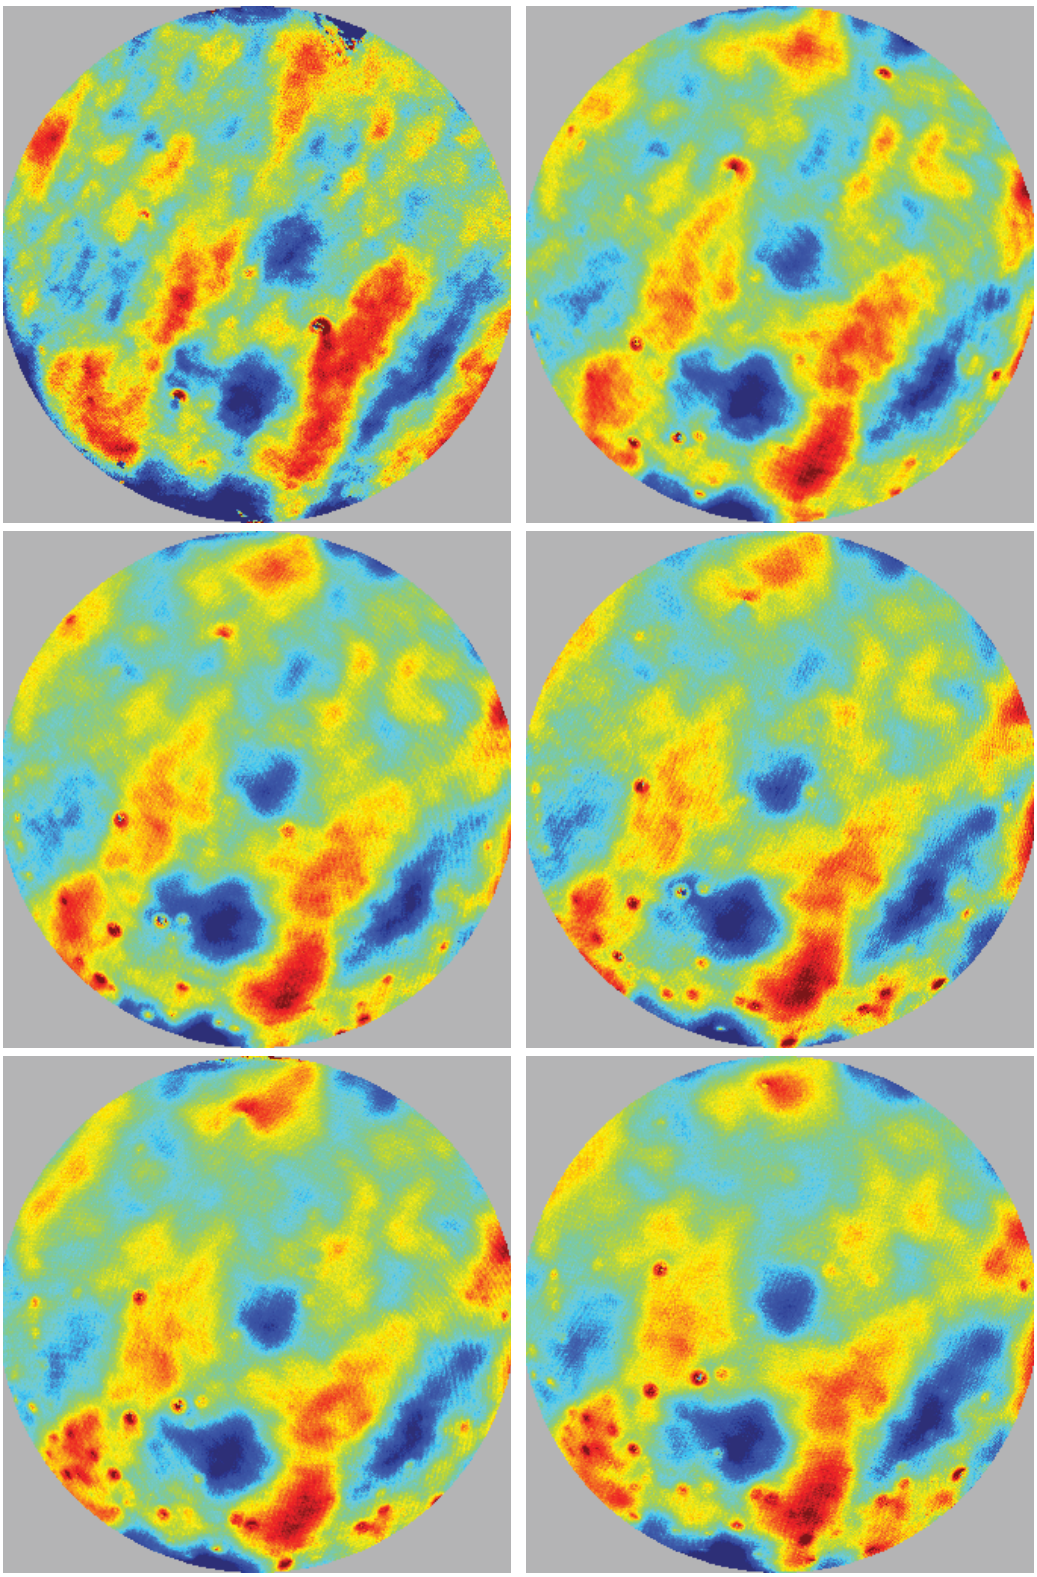

Figure 2.15: Mucin globules pulled upwards by surface tension.
6 mm diameter displayed. ±0.25 μm color height scale; red is away from surface of cornea. First 16 Zernike terms subtracted. First still occurs ~100 ms after a blink. Each still is separated in time by 1 s.

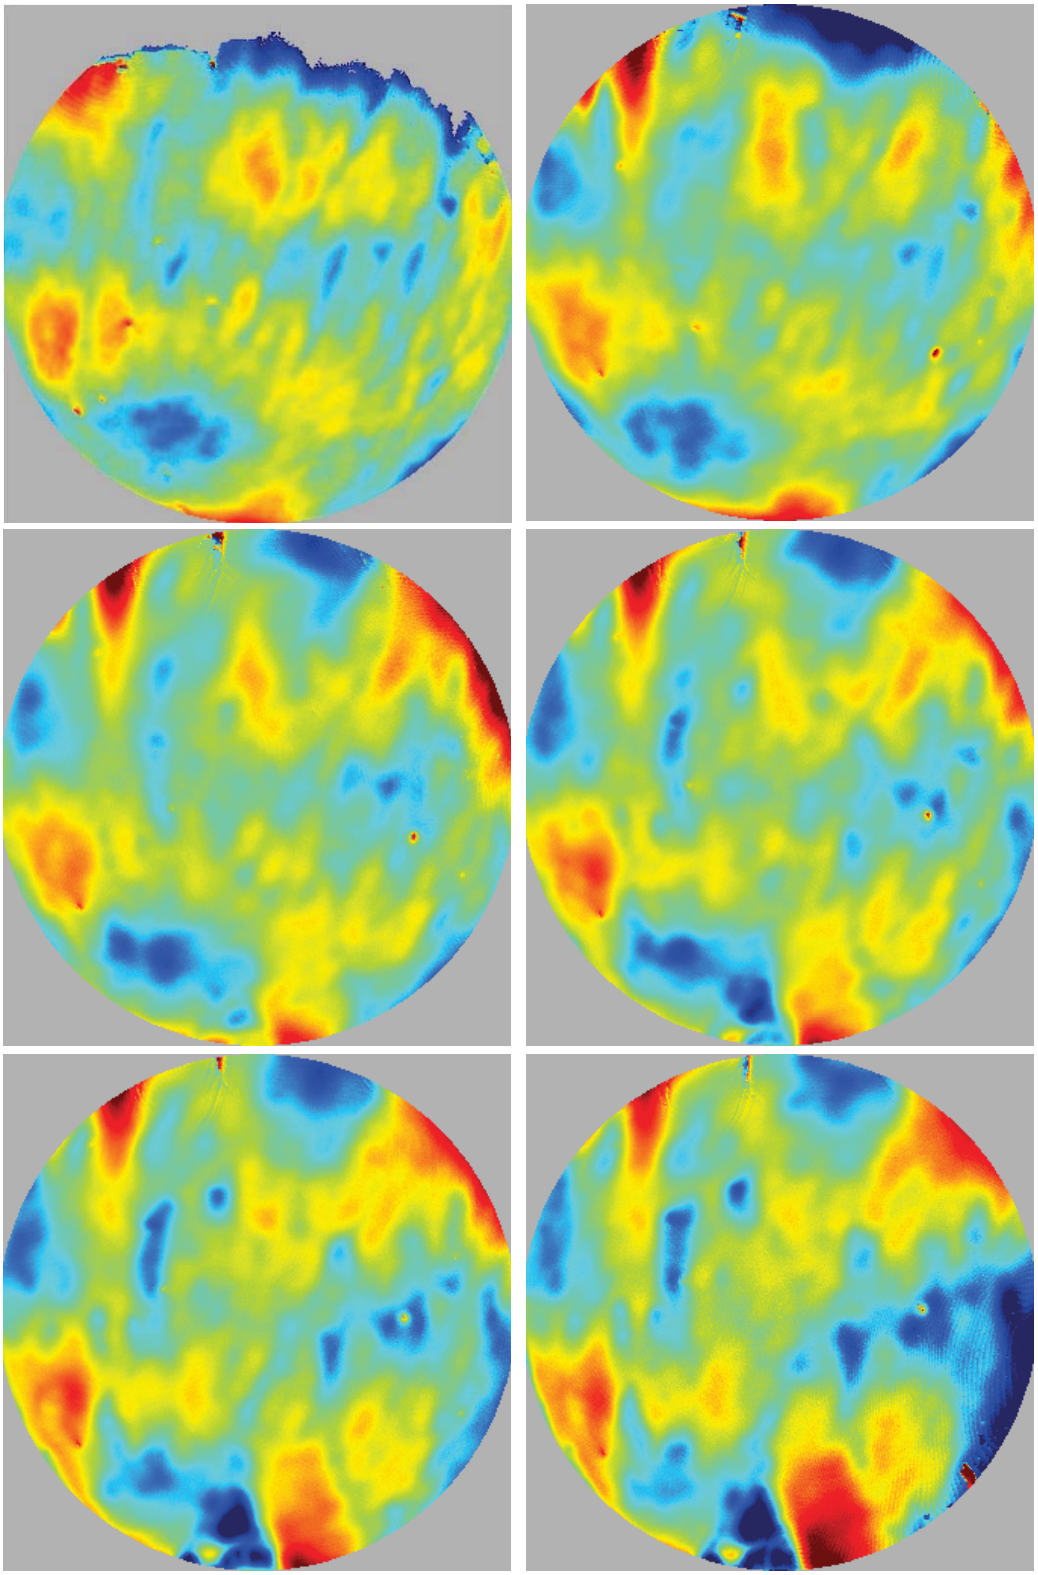

Figure 2.16: Tear film surface evolution after a blink.
6 mm diameter displayed. ±0.75 μm color height scale; red is away from surface of cornea. First 16 Zernike terms subtracted. First still occurs ~100 ms after a blink. Each still is separated in time by 1.5 s.

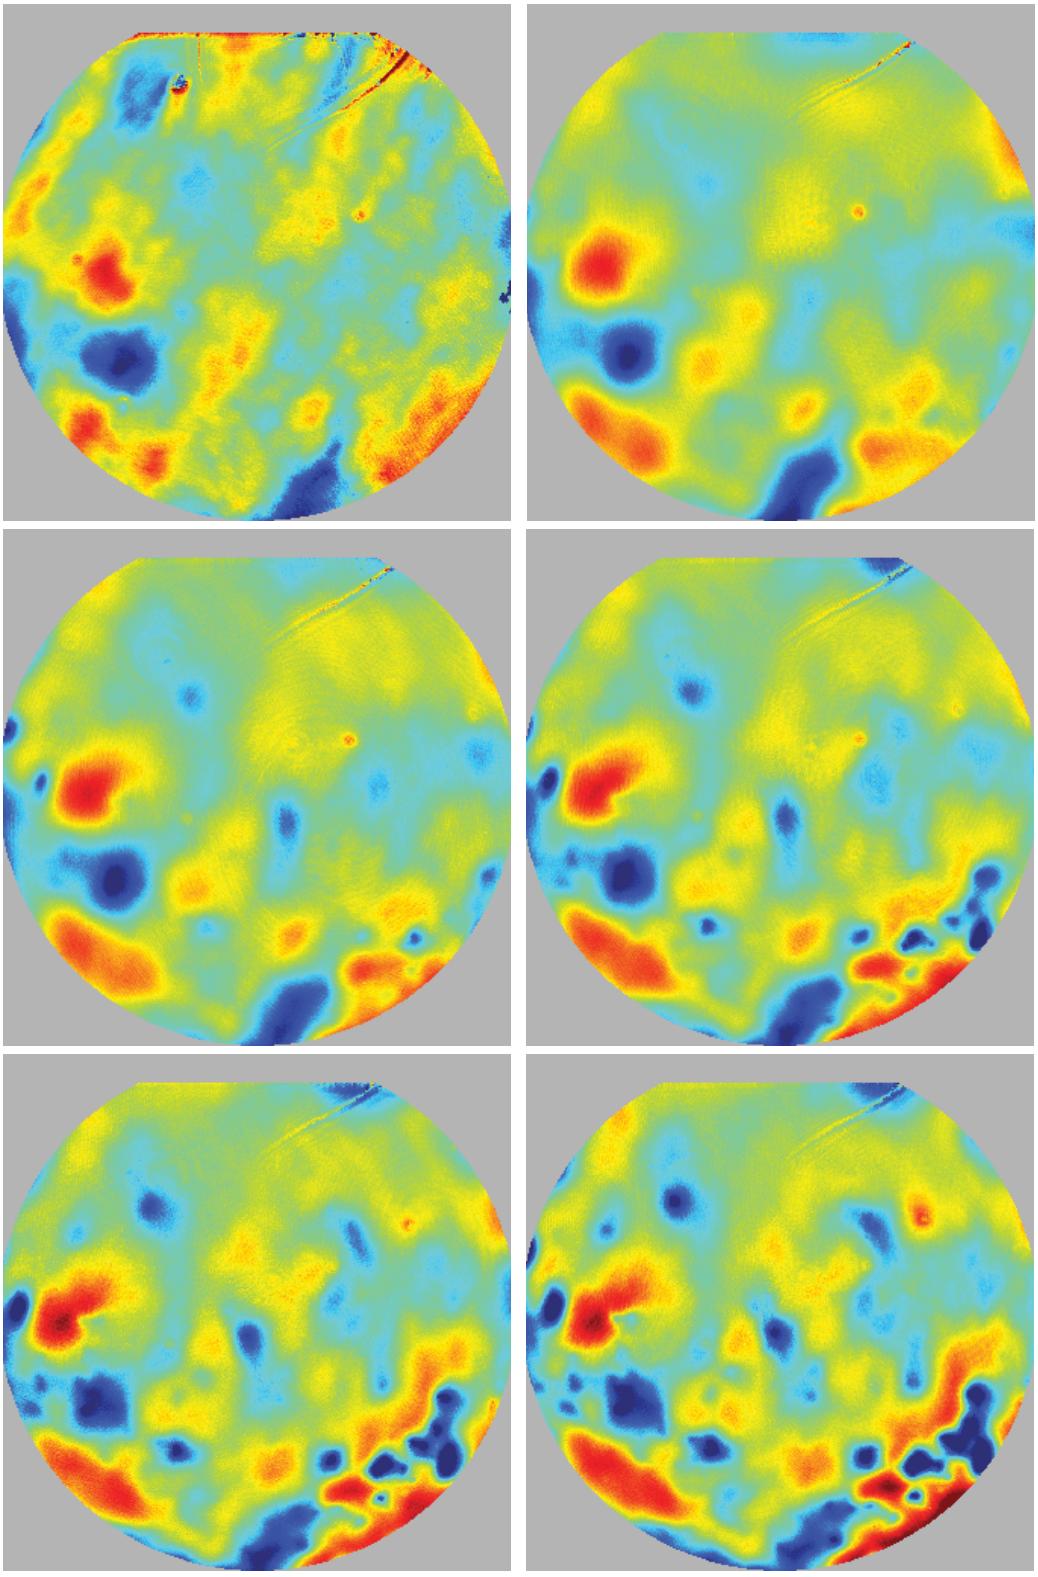

Figure 2.17: Tear film surface evolution after a blink.
6 mm diameter displayed. ±0.375 µm color height scale; red is away from surface of cornea. First 16 Zernike terms subtracted. First still occurs ~100 ms after a blink. Each still is separated in time by 9 s.

2.7.2 CONTACT LENS

This section presents results from human subjects wearing soft contact lenses. Similar to the previous section, a subset of sequential image frames are presented. Testing subjects with contact lenses presented a slight increase in testing difficulty. The first issue is that blinking displaces the contact lens by a small amount. The displacement is large enough to result in a loss of data for approximately a full second until the lens has resettled into the capture range of the TFI. Figure 2.18 shows an example of a contact lens settling after a blink. The red arrow to the right of each image follows an artifact in the tear film that allows the location of the contact lens to be tracked between frames. Contact lens motion can also be inferred by the display boundaries (i.e. the lower half of the data appears over a period of time). The second issue is that the surface topography of the tear film that results from the underlying contact lens results in deviations from the designed aspheric wavefront generated by the converger optics. This has the result of producing larger fringe densities at the edges and loss of data. Therefore, most of the results shown in this section have been reduced to an approximate 5 mm display diameter. A final issue with testing on contact lenses is that the interaction of the tear film with the contact lens is significantly more complex than the tear film on the bare eye, resulting in recorded interferograms that are significantly more difficult to process (refer to Section 2.6).

Figure 2.19 through Figure 2.22 contain examples of how the tear film structure evolves in the presence of contact lenses. From these examples it can be seen that the behavior of the tear film on a contact lens is drastically different from the tear film on the bare eye. The primary change in tear film behavior appears to be an increase in breakup and surface stability, which may be due to the presence of the hydrophobic silicone in the contact lens (Section 1.3). The tear film also appears to behave differently between various contact lenses, which may be explained by material and/or subject differences. These examples demonstrate how the TFI could be used to develop and characterize contact lens materials.

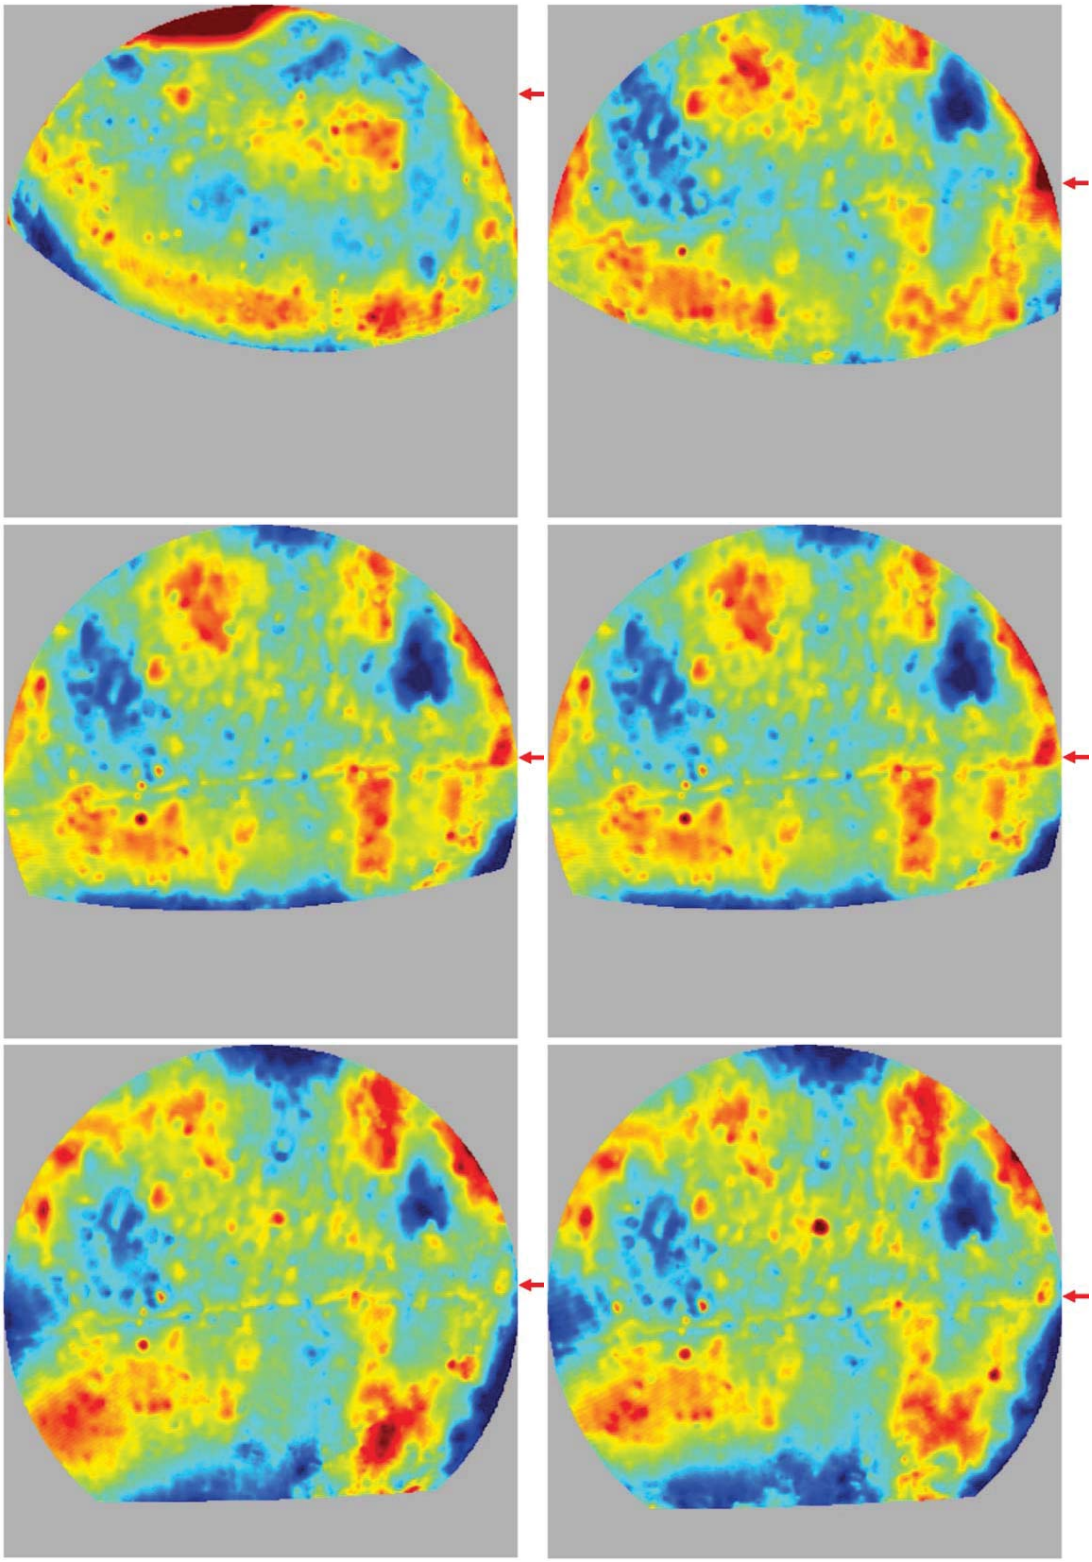

Figure 2.18: Biofinity toric (CooperVision). Contact lens settling after a blink. 5 mm diameter displayed. ±0.75 µm color height scale; red is away from surface of cornea. First 16 Zernike terms subtracted. First still occurs ~1 s after a blink. Each still is separated in time by 200 ms.

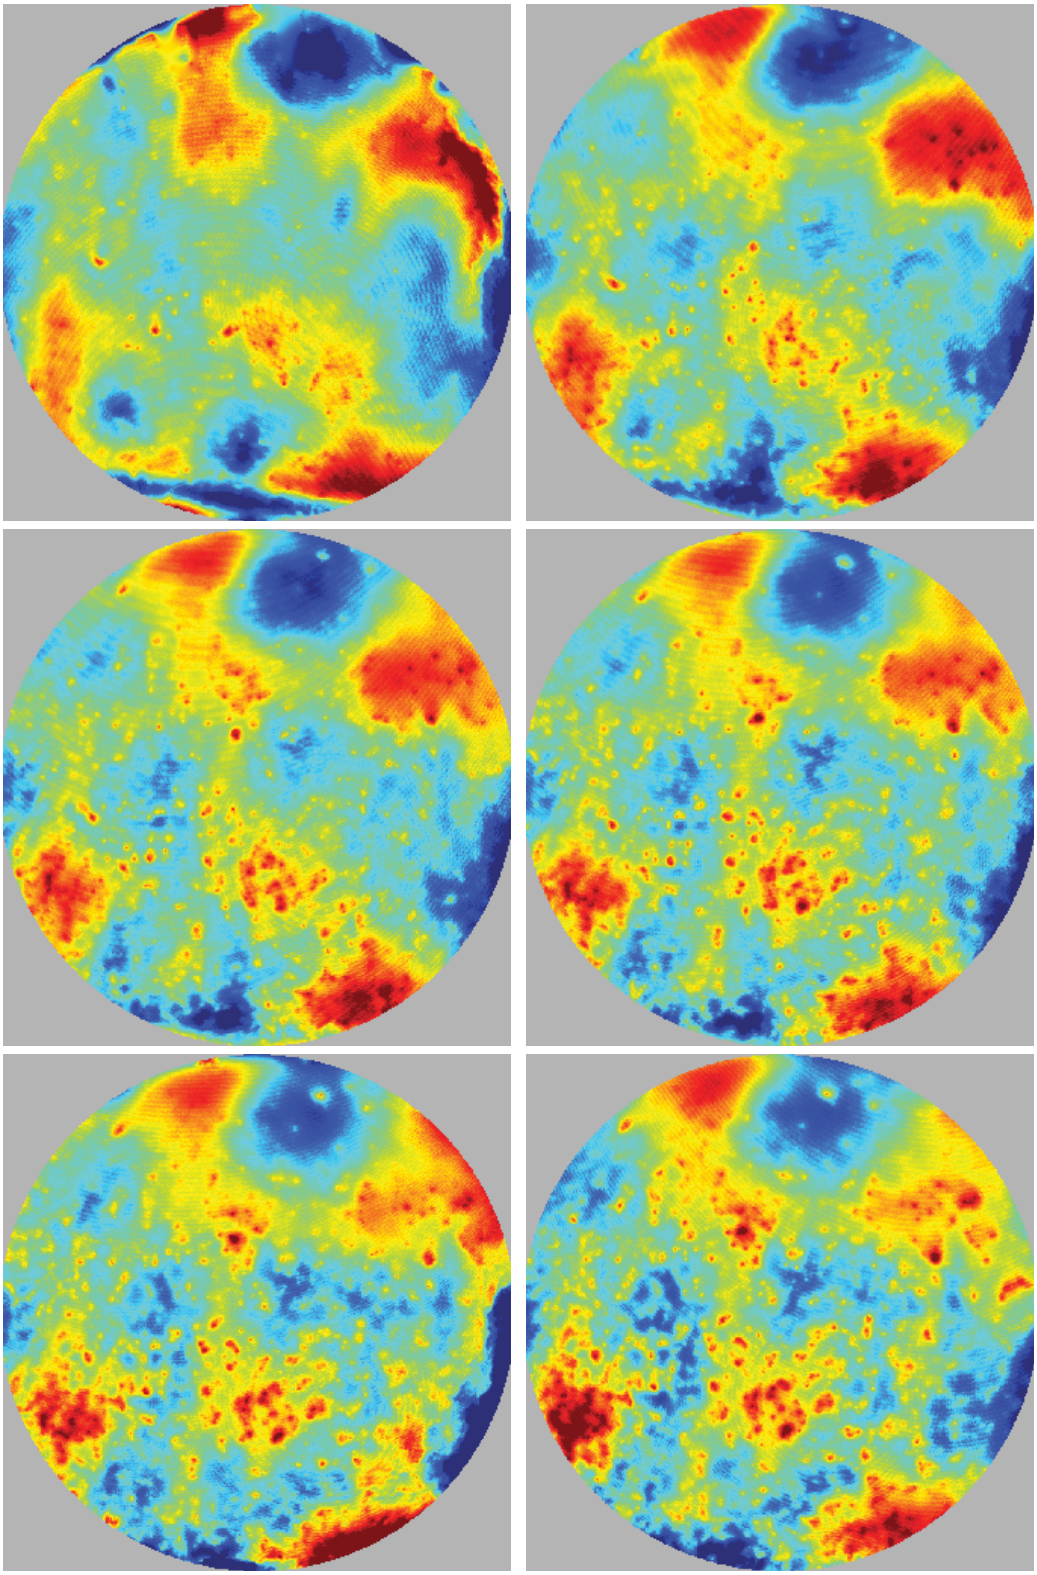

Figure 2.19: AquaClear (CooperVision). Tear film surface evolution after a blink. 5 mm diameter displayed. ±0.375 µm color height scale; red is away from surface of cornea. First 16 Zernike terms subtracted. First still occurs ~1 s after a blink. Each still is separated in time by 200 ms.

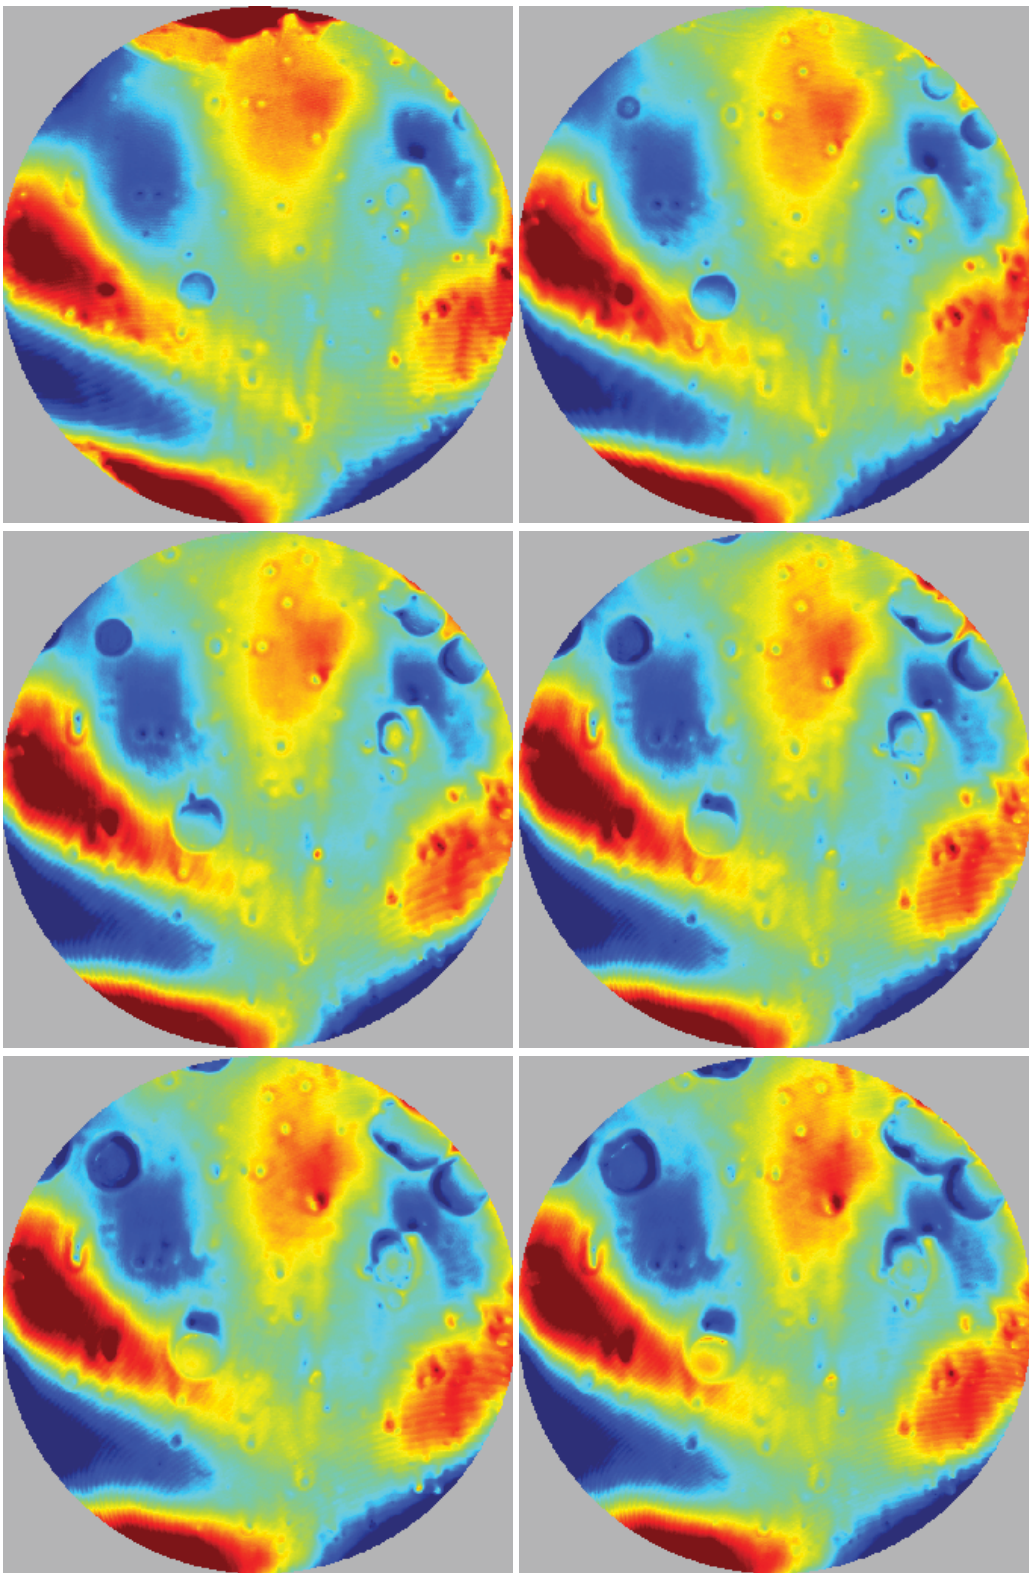

Figure 2.20: Biofinity (CooperVision). Tear film surface evolution after a blink. 5 mm diameter displayed. ±0.75 µm color height scale; red is away from surface of cornea. First 16 Zernike terms subtracted. First still occurs ~1 s after a blink. Each still is separated in time by 200 ms.

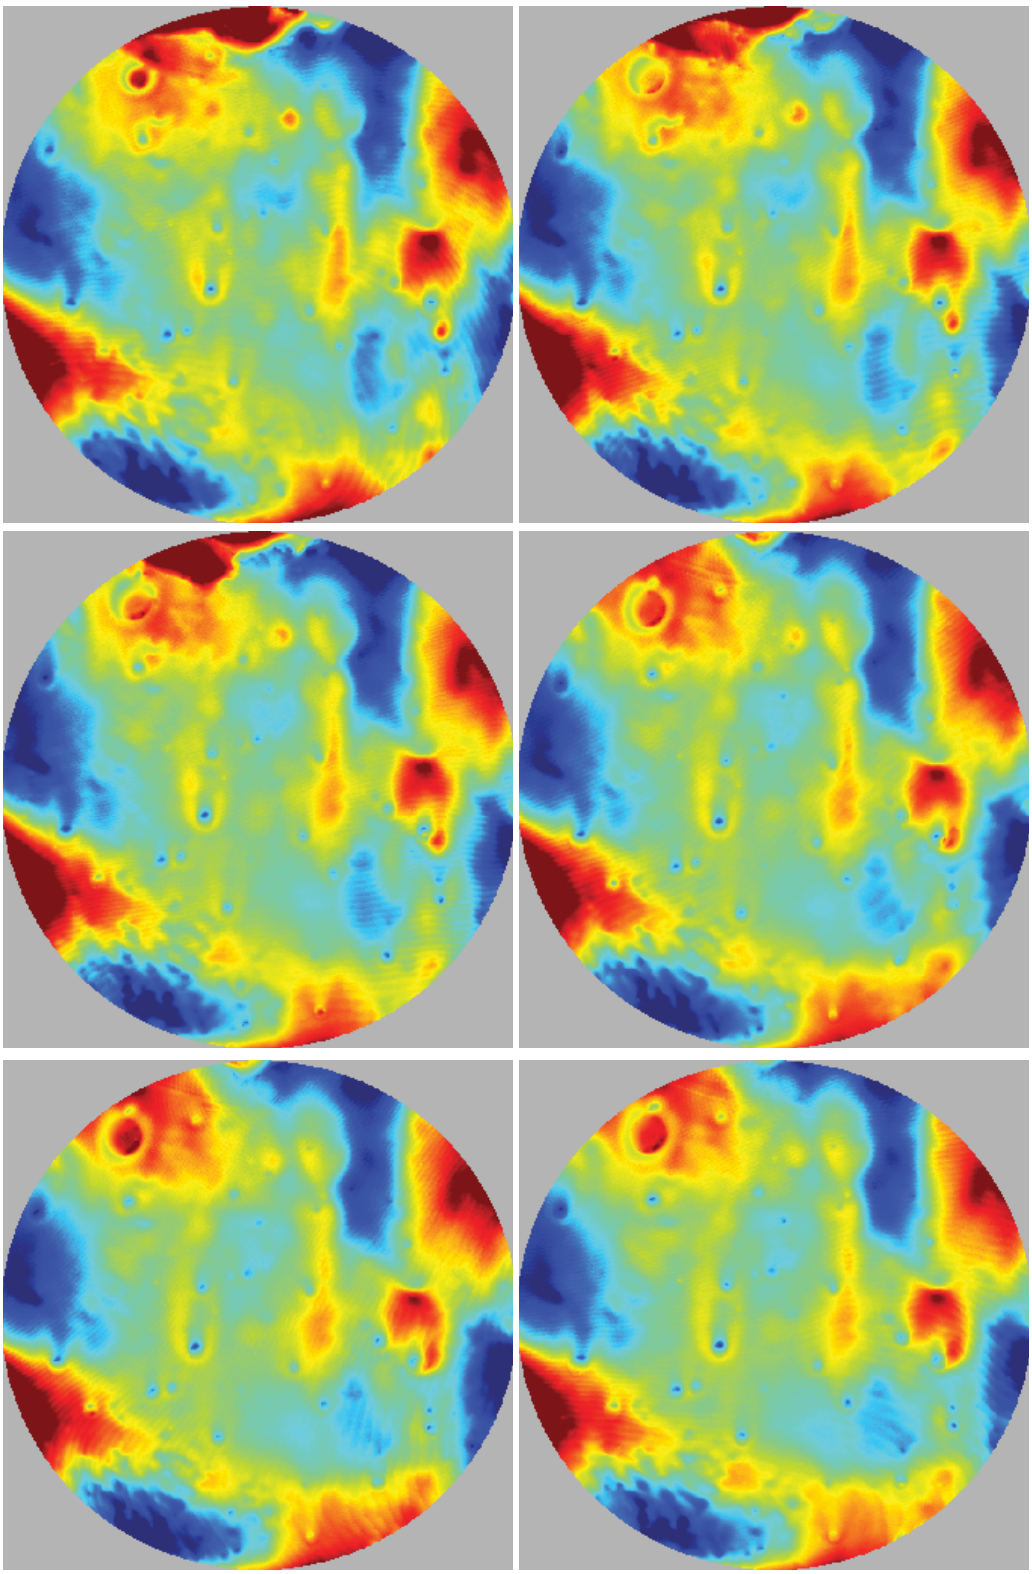

Figure 2.21: Biofinity (CooperVision). Tear film surface evolution after a blink. 5 mm diameter displayed. ±0.5 µm color height scale; red is away from surface of cornea. First 16 Zernike terms subtracted. First still occurs ~1 s after a blink. Each still is separated in time by 70 ms.

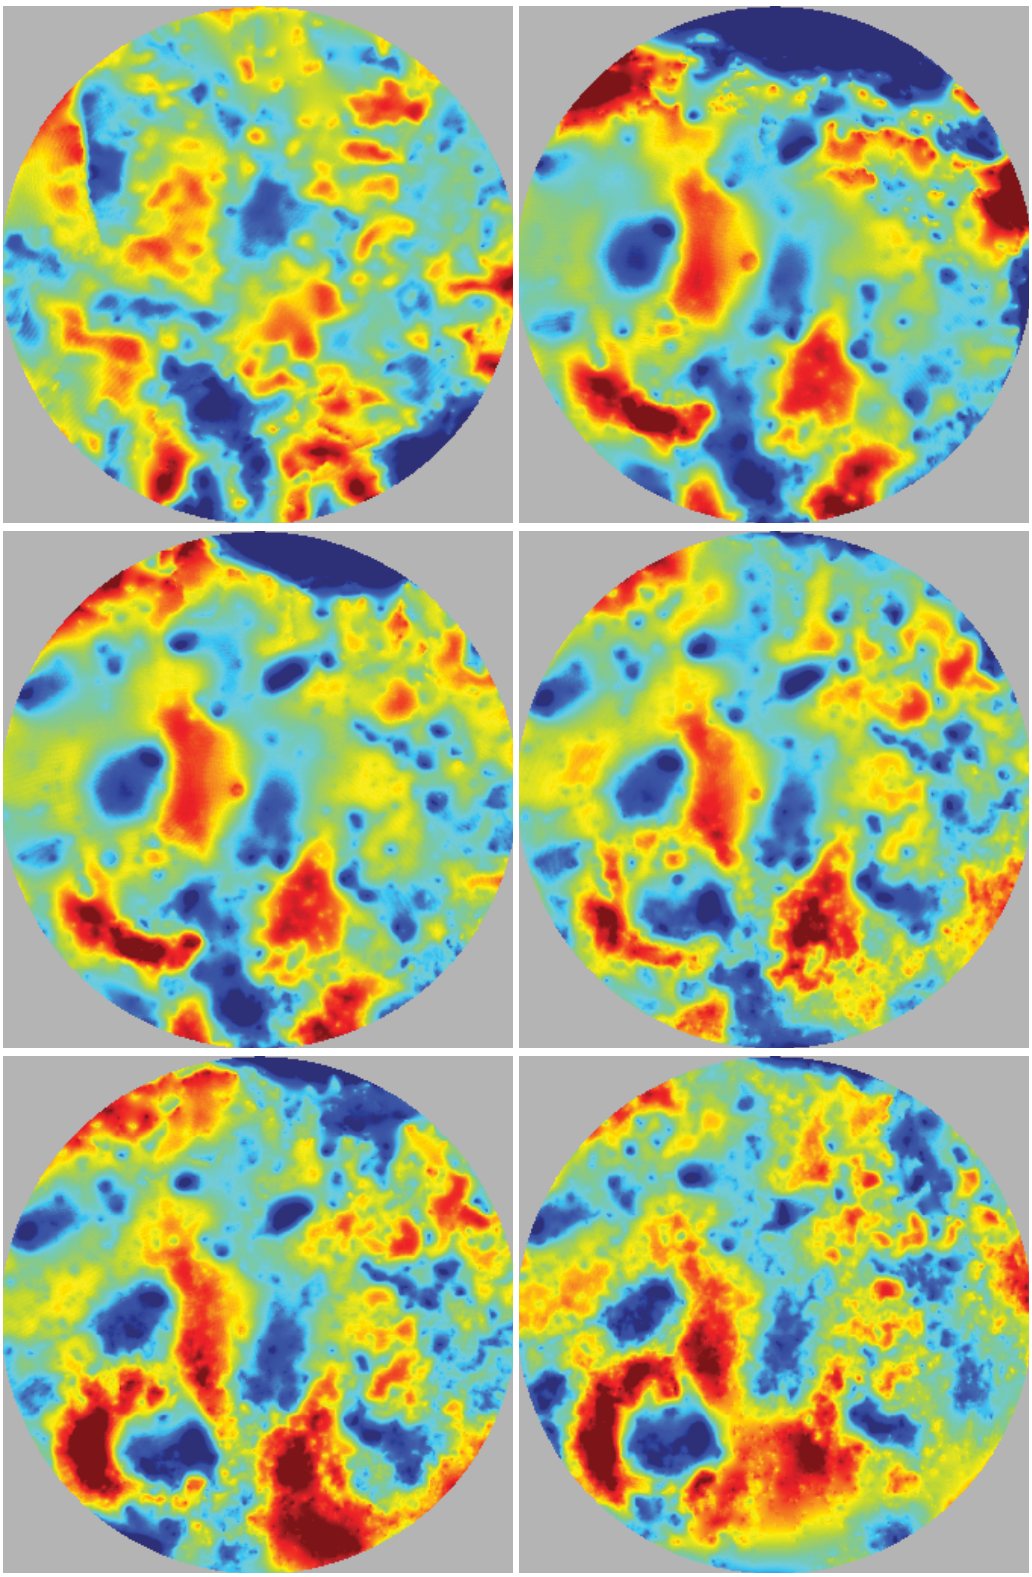

Figure 2.22: Oasys (Acuvue). Tear film surface evolution after a blink.
4 mm diameter displayed. ±1.0 µm color height scale; red is away from surface of cornea. First 16 Zernike terms subtracted. First still occurs ~1 s after a blink. Each still is separated in time by 1.6 s.

2.8 CONCLUSIONS

The Tear Film Interferometer has been demonstrated as a viable and practical technology for the non-invasive characterization of the dynamic topography of the *in vivo* corneal tear film. Several human subjects have been examined using this system, with a demonstrated surface height resolution of 25 nm and spatial resolution of 6 μm. Acquisition times of up to 120 seconds have been obtained. The combination of height resolution, spatial resolution, and measurement area has been improved over existing *in vivo* technologies.

The TFI has been tested on a small subset of people to verify its capabilities. The true value of the TFI is that it is a basic science instrument developed to investigate the human tear film and its dynamic behavior. It is expected that further testing with the TFI will yield a better understanding of the tear film and support the development of new contact lens materials. The results from this phase of interferometer development will support the development of the next interferometer for measuring the dynamic corneal topography. This system will extend the capabilities of the TFI even further.

3 SURFACE PROFILING INTERFEROMETER FOR NON-INVASIVE ASSESSMENT OF LOCAL TOPOGRAPHY

This chapter presents an interferometer for measuring dynamic corneal topography, or the Surface Profiling Interferometer for Non-invasive Assessment of Local TOPography (SPINALTOP). SPINALTOP is a dual-interferometer system that can instantaneously measure the dynamic topography of the human cornea and dynamic tear film surface. Sequential measurements allow for high resolution observation of the corneal topography and tear film dynamics.

3.1 SYSTEM OVERVIEW

SPINALTOP consists of a pair of polarization-splitting Twyman-Green interferometers and a fixation assembly. All three sub-systems are co-aligned and combined at the input to the converger assembly (Figure 3.1). A schematic of the full system is shown in Figure 3.2 and Figure 3.3 and the final as-built system is shown in Figure 3.4. Beam paths for the three sub-systems have been overlaid on the images.

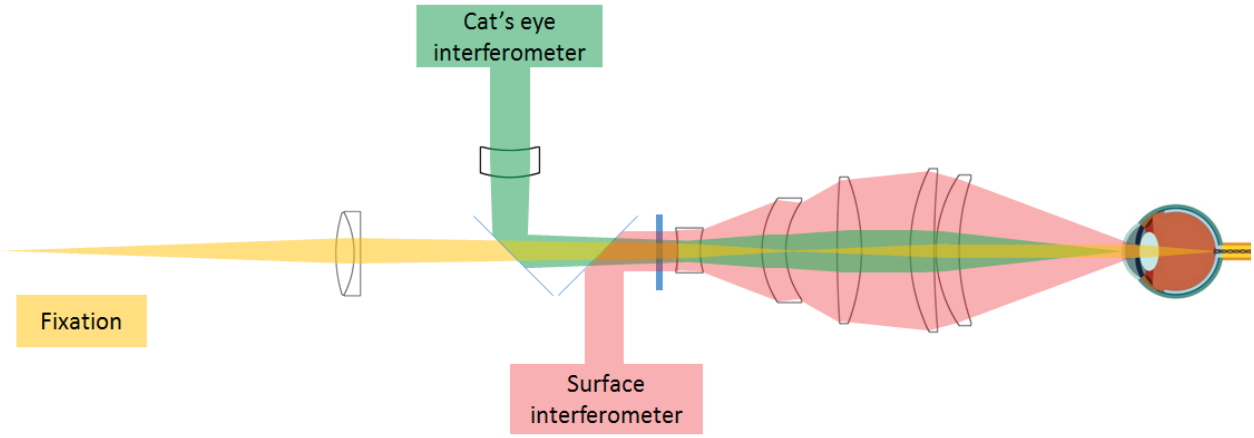

Figure 3.1: This is SPINALTOP.

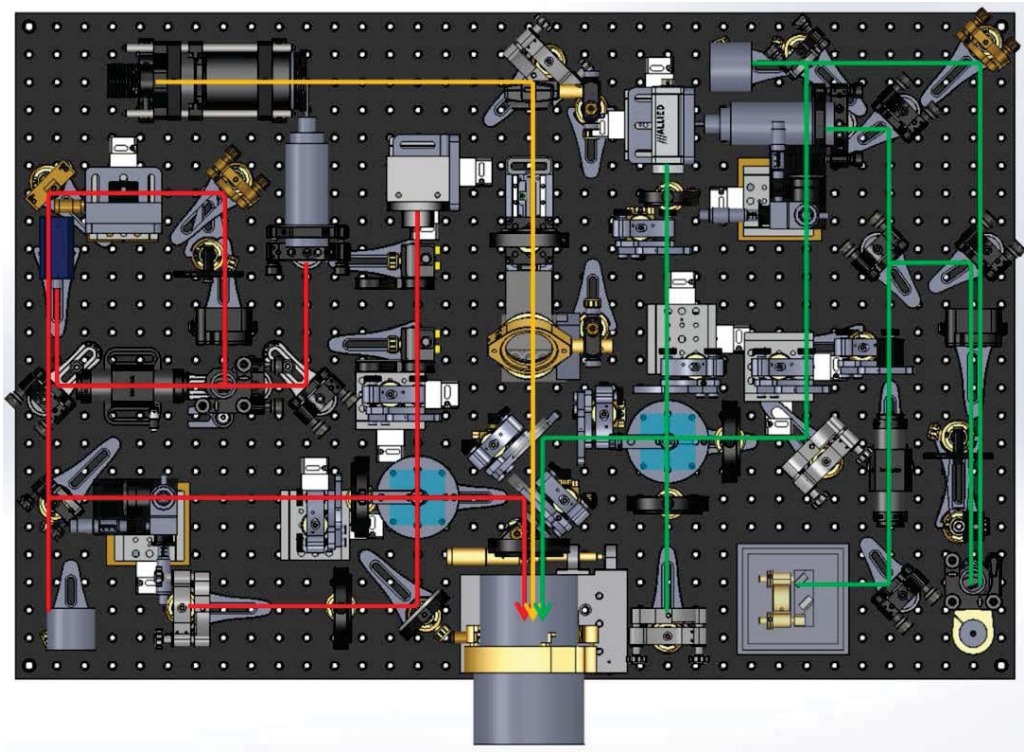

Figure 3.2: Schematic of SPINALTOP system (top view).
The surface interferometer, cat's eye interferometer, and fixation assembly paths are overlaid in red, green, and yellow, respectively.

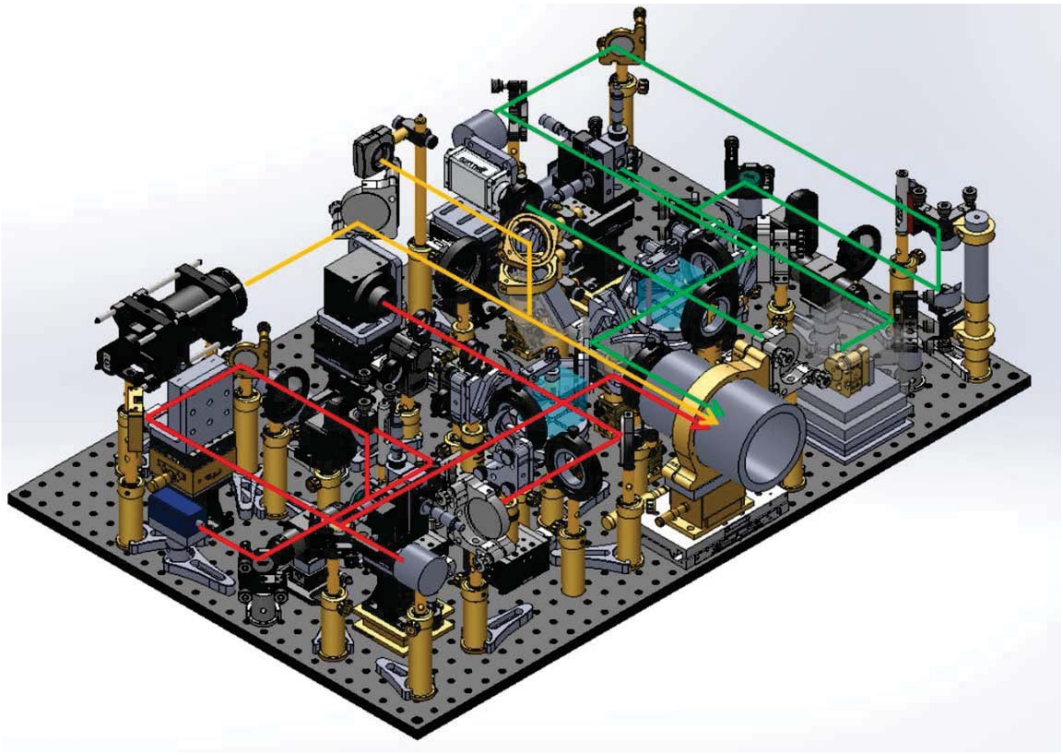

Figure 3.3: Schematic of SPINALTOP system (isometric view).

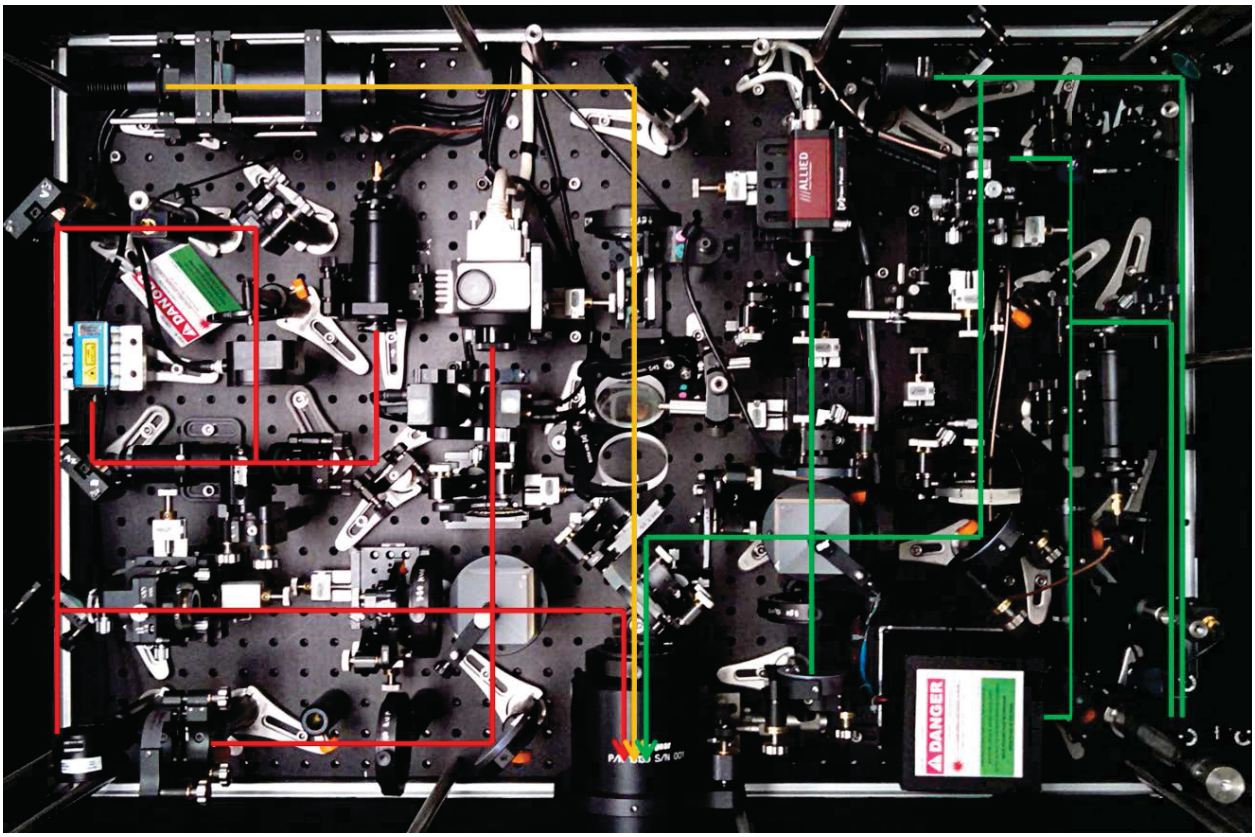

Figure 3.4: Top view of the SPINALTOP hardware.

The first sub-system in SPINALTOP is the surface interferometer, which is an adaptation of the TFI system (Chapter 2). Modifications have been made to the surface interferometer based on feedback from testing with the TFI and is discussed in Section 3.2. The second sub-system is the cat's eye interferometer. It is designed to produce a simultaneous cat's eye measurement along with the surface interferometer measurement (refer to Figure 3.1). The cat's eye measurement allows for the recovery of the absolute radius of curvature, but will be used to recover more information about the test surface. A detailed description is given in Section 3.3. The combination of measurements from the two interferometer sub-systems is necessary to reconstruct the topography of the cornea. The method by which the topography is recovered is discussed in Chapter 5. The addition of another laser source in the cat's eye interferometer that will illuminate the eye requires a re-evaluation of laser safety for the system (Section 3.6).

The final sub-system in SPINALTOP is the fixation assembly which provides in-line fixation with the instrument (Section 3.4). The TFI system used external fixation that was off-eye, meaning that the eye used for fixation was the opposite of the eye being tested. This did not have an impact on a subject's ability to maintain fixation, but did limit present several limitations.

The SPINALTOP system is built on a 2' x 3' aluminum optical breadboard and mounted to a pair of motorized linear stages (NEAT XYL-1515-SM, New England Affiliated Technologies). The headrest is mounted to a motorized vertical stage (NEAT ZE-100-SM, New England Affiliated Technologies). The three stages provide fine alignment of the subject to the interferometer, which is the same setup that was used on the TFI. An external view of the SPINALTOP system is shown in Figure 3.5. From a subject's point of view, the operation and external layout of SPINALTOP is no different than the TFI, with the exception of fixation. A subject will sit down, place their head in the head rest, and the operator will align SPINALTOP to the subject and capture data. The difference from the TFI system is that SPINALTOP will be capturing simultaneous surface and cat's eye data that will be used to reconstruct high resolution absolute topographic measurements of the cornea. A number of surface parameters can be extracted from the topography, such as the radius of curvature, asphericity, and corneal astigmatism. These are parameters that could not be recovered from TFI measurements. TFI measurements required that at least 12 Zernike had to be subtracted from the measurement.

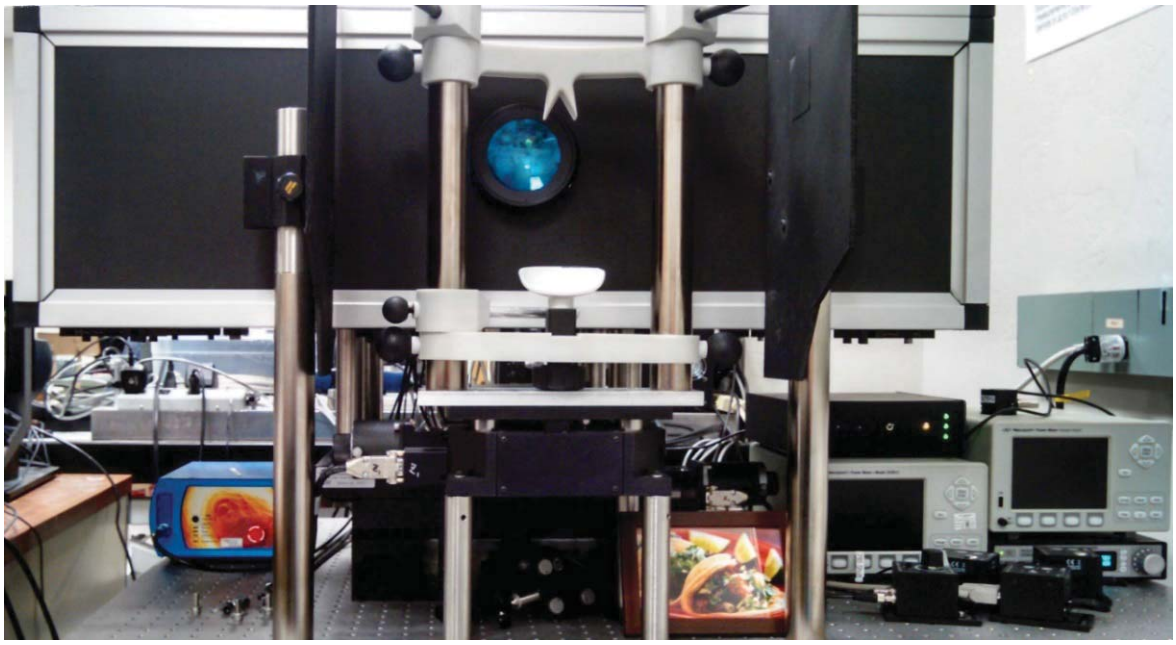

Figure 3.5: External view of the SPINALTOP hardware.

3.2 SURFACE INTERFEROMETER

The surface interferometer is functionally equivalent to the TFI with a few modifications, so the discussion will be kept brief. A more in depth discussion about the surface interferometer can be found in the previous chapter or the references (Micali et al., 2015; Primeau & Greivenkamp, 2012b; Primeau, 2011). The surface interferometer is a polarization-splitting Twyman-Green interferometer that operates at a near-infrared (NIR) wavelength of 785 nm. The laser source is continually monitored to ensure eye safe operation of the laser system. The test path contains the same custom built converger assembly designed to the mean shape of the human cornea. Reference and test path wavefronts are interfered and measured with a pixelated camera that provides an instantaneous measurement of the phase. The individual components of the surface interferometer are discussed in more detail in the following sections with emphasis on changes that have been made from the TFI system.

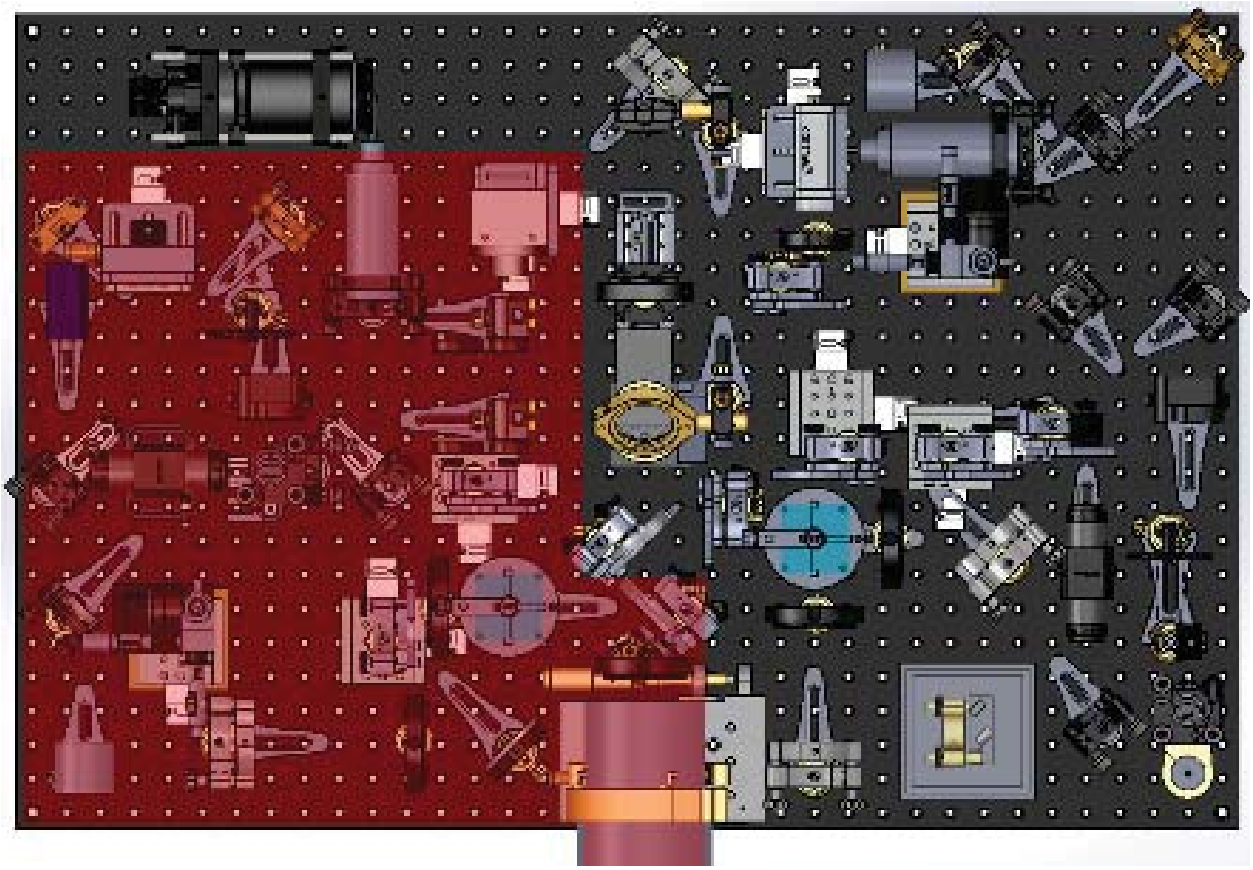

Figure 3.6: Hardware layout of the surface interferometer (highlighted in red).

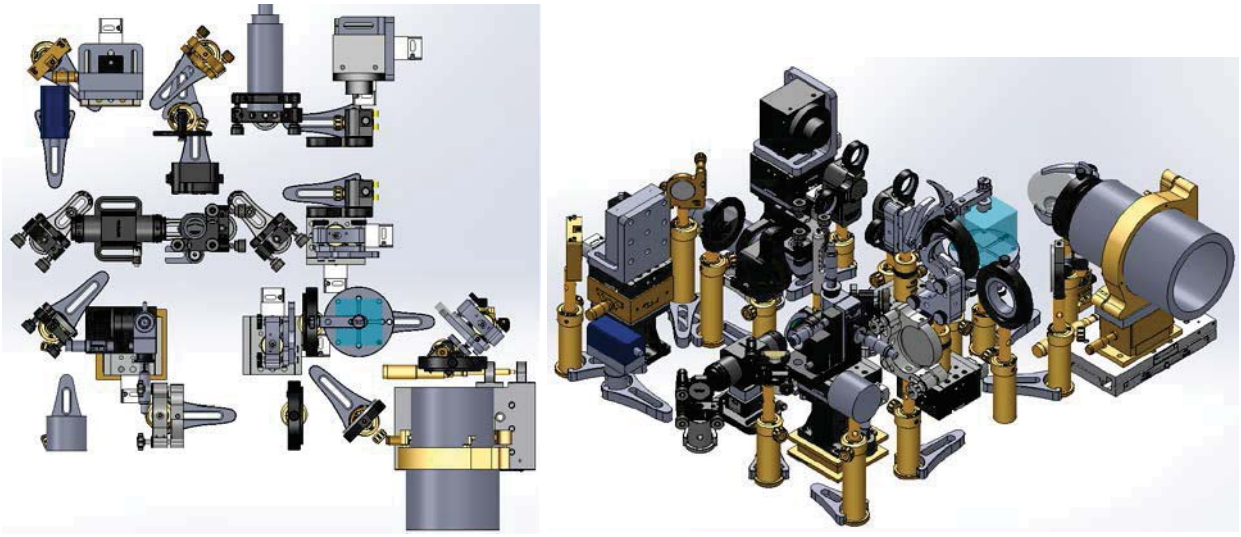

Figure 3.7: Isolated surface interferometer hardware layout.
(Left) Top view. (Right) Isometric view.

3.2.1 SOURCE ASSEMBLY

The surface interferometer source assembly is shown in Figure 3.8 with key components labelled. One of the requirements for SPINALTOP was to re-use as much of the previous TFI hardware as possible, which includes the problematic laser source (Section 2.5). The laser source is a 785 nm solid state laser (Innovative Photonics I0785SH0100B-TH-L) with an average output power of 100 mW. After a fold, the laser is passed through a free-space optical isolator (Thorlabs IOT-5-780-VLP) with a specified isolation of 55 dB. The isolator was upgraded from the TFI isolator (Thorlabs IO-5-780-VLP) which had a specified isolation of 33 dB. The diode laser is extremely sensitive to feedback and the previous level of isolation was insufficient. Optical feedback results in a loss of stability in the laser and a loss of fringe modulation. Feedback can also permanently damage the laser.

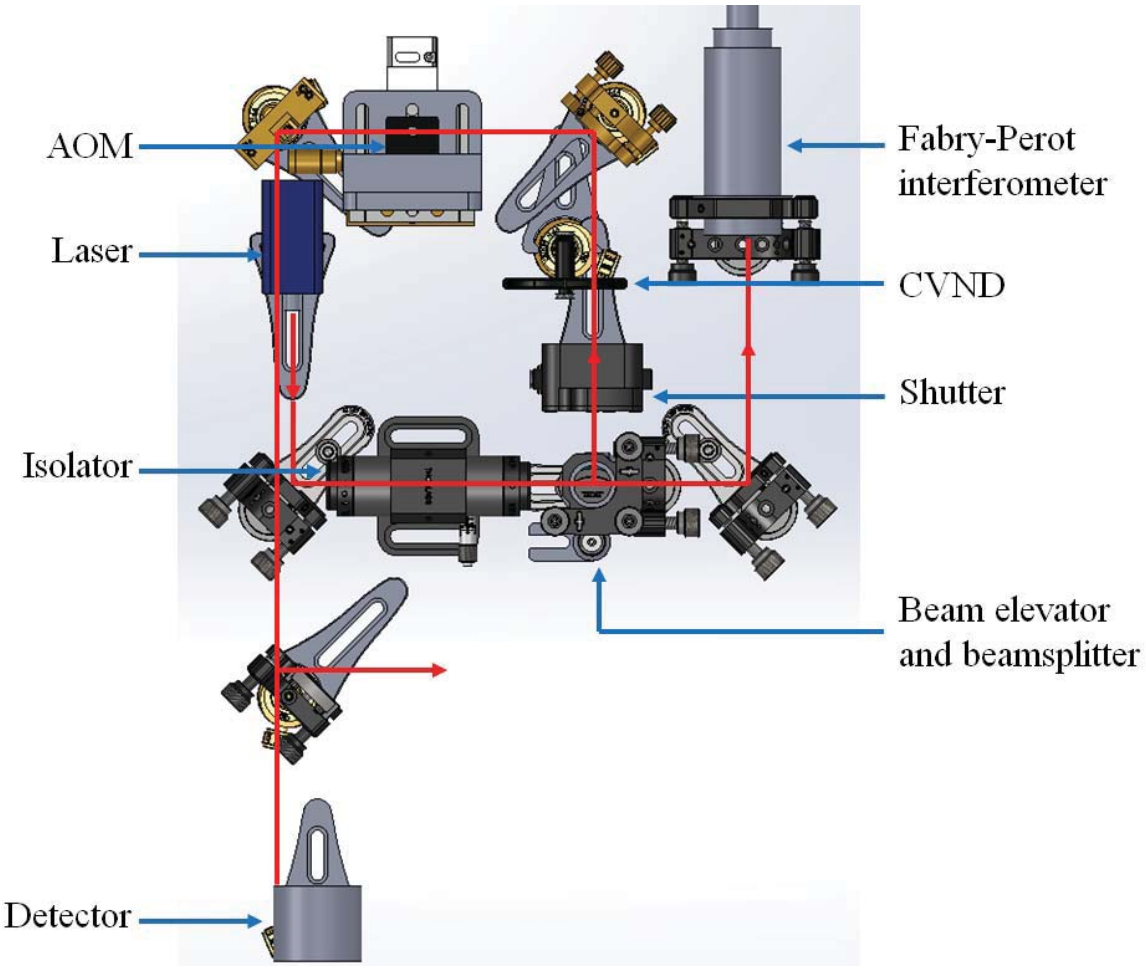

Figure 3.8: Surface interferometer source assembly.

Following the optical isolator, a pick-off beamsplitter transmits 10% of the laser to a fold mirror and feeds into a scanning Fabry-Pérot interferometer (Thorlabs SA200-5B). The Fabry-Pérot interferometer is a high resolution spectrum analyzer that provides continuous feedback on the laser source spectrum. The Free Spectral Range (FSR) is 1.5 GHz and has a Finesse ($\mathcal{F}$) of approximately 250. The resolution ($\Delta\nu$) can be determined from the relationship:

$$\Delta\nu = \frac{FSR}{\mathcal{F}} = 7.5\ MHz \qquad (3.1)$$

At λ =785 nm the resolution is also given by:

$$\Delta\lambda = \frac{\Delta\nu \cdot \lambda^2}{c} = 0.015\ pm \qquad (3.2)$$

The FSR of the Fabry-Pérot interferometer (1.5 GHz) is equivalent to 3.1 pm. The linewidth of the laser source based on measurements made in Section 2.5 was 0.64 pm, therefore the Fabry-Pérot interferometer should be sufficient to resolve the linewidth of the laser source without exceeding the measurement range.

The Fabry-Pérot interferometer consists of two high reflectivity mirrors that form a confocal cavity (Hercher, 1968). Piezoelectric transducers (PZT) allow the cavity length (mirror separation) to be varied, which provides a method to select the transmitted wavelength. A voltage ramp applied across the PZT provides a linear scan through the source spectrum. The transmitted light is captured by a photodiode detector and measured on an oscilloscope that is triggered off of the PZT voltage ramp. Figure 3.9 shows a measurement of the 785 nm laser source that has been properly tuned and allowed to stabilize. The PZT drive voltage appears as a linear ramp in the background. The three peaks are periodic repetitions of the source spectrum that result from scanning the Fabry-Pérot cavity through multiple modes of the laser wavelength.

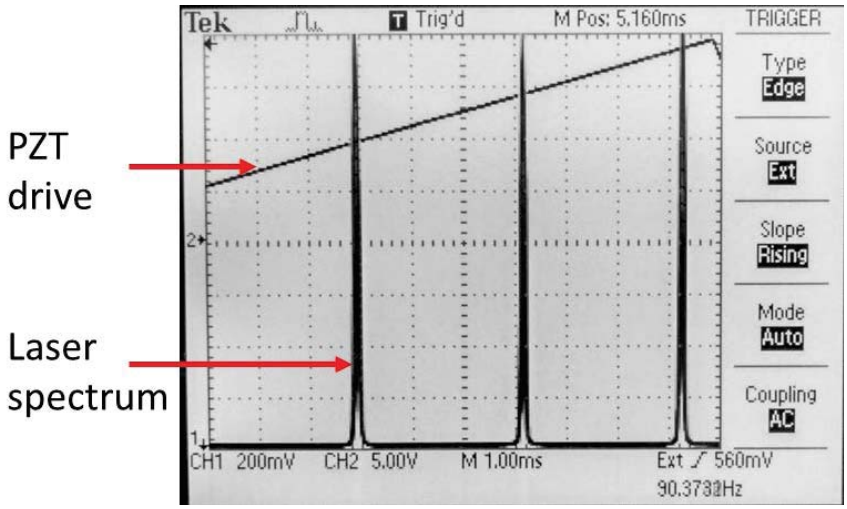

Figure 3.9: Surface interferometer laser source spectrum measurement with a scanning Fabry-Pérot interferometer.

Knowledge of the interferometer FSR and measured peak separation allows for a calibration of the oscilloscope display. The peak-to-peak separation from the measurement shown in Figure 3.9 is 3.25 ms. The resulting calibration is then 1.5 GHz / 3.25 ms = 0.4625 GHz/ms. Note that the unit of Hz is being applied to the optical source spectrum and the unit of ms is being applied to the oscilloscope time measurement. The laser source FWHM is then measured on the oscilloscope to be approximately 100 μs, or:

$$\Delta\nu = 0.4625\frac{GHz}{ms}\cdot 100\ \mu s = 46.2\ MHz \qquad (3.3)$$

The coherence length for this source, assuming a Gaussian spectrum is:

$$\ell_c = \frac{c}{\Delta\nu}\sqrt{\frac{2}{\pi}} = 5.2\ m \qquad (3.4)$$

The ability to fine tune the laser source and achieve a coherence length of 5.2 m over the previously measured 0.7 m demonstrates the utility of having feedback from the Fabry-Pérot interferometer. Another significant improvement is that coherence length does not have to be measured indirectly through the use

of scanning the test or reference arms. This method is instantaneous and allows for quick and easy adjustment of the laser without having to tear apart or move components within the interferometer.

Typical operation of the surface interferometer requires that the laser source is left on for at least 30 minutes to stabilize. Once the laser has stabilized, the spectrum is verified with the Fabry-Pérot interferometer. If the output is unstable or appears poor, the TEC set-point on the laser controller is adjusted (Figure 3.10). The spectrum shown in Figure 3.9 represents a stable and well-tuned laser. Figure 3.11 shows two degraded configurations of the laser. The spectrum on the left represents a semi-stable state which could provide sufficient fringe modulation if used. The spectrum on the right represents an even more degraded case where reduced fringe modulation will occur unless the optical path lengths in the two arms of the interferometer are matched. Figure 3.12 shows measurements made on a glass reference surface with the laser in two different states. The measurement on the left was made with the laser in a stable state (Figure 3.9) and the measurement on the right was made when the laser was in a degraded state (Figure 3.11, left).

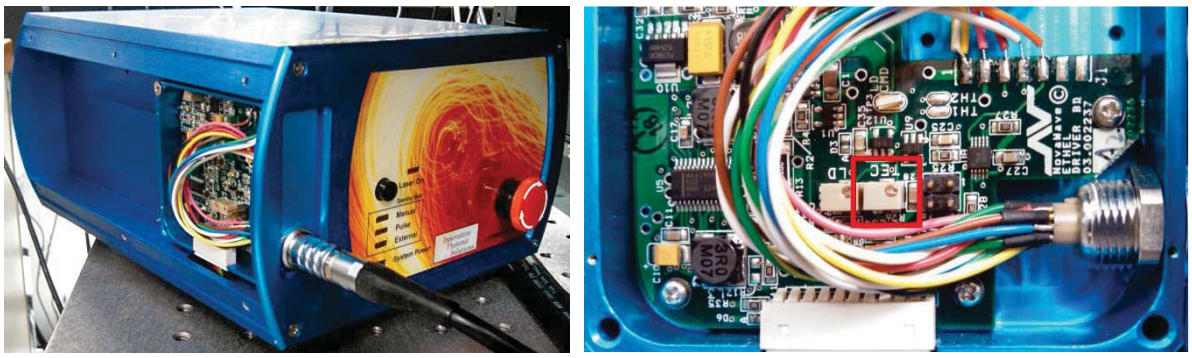

Figure 3.10: Laser control box (left) showing TEC set-point pot (right, red box).

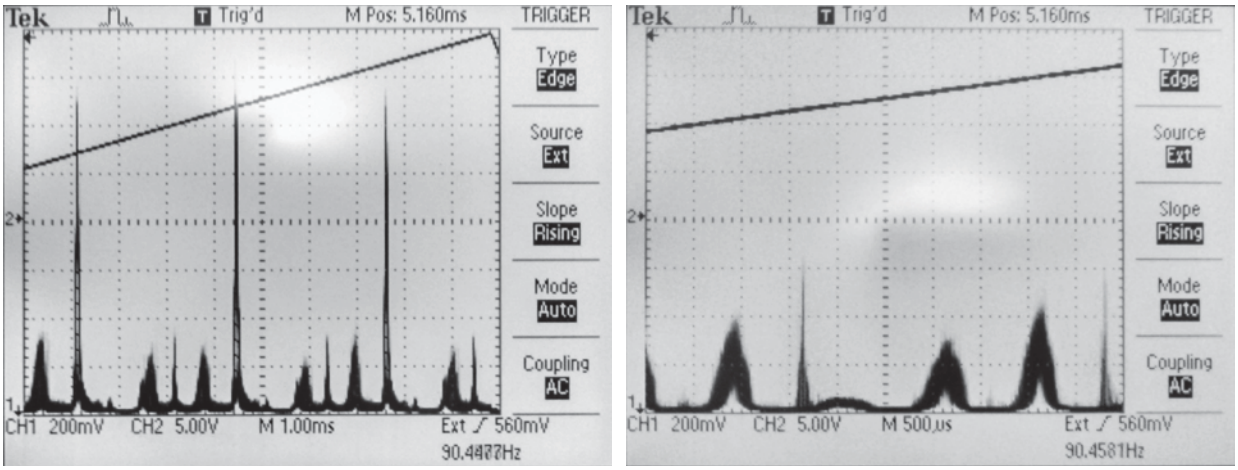

Figure 3.11: 785 nm laser spectrum for two degraded cases.

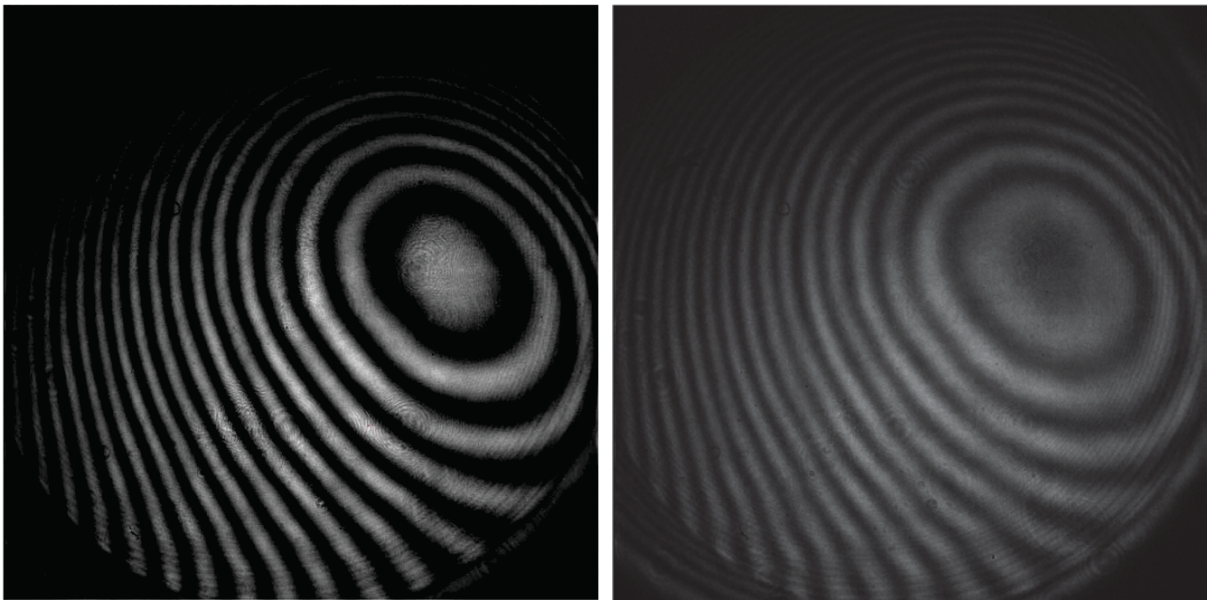

Figure 3.12: Fringe visibility example for two laser states.
(Left) Stable and well-tuned laser. (Right) Degraded tuning of the laser.

Returning back to the discussion of the source assembly (Figure 3.8); after the pick-off beamsplitter, the remaining 90% of light is reflected through a beam elevator. The beam elevator raises the effective beam plane to approximately 75 mm above the starting beam plane, or 160 mm above the optical breadboard. After the beam elevator, the light passes through the shutter and continuously variable neutral density filter (CVND). The shutter is part of the safety-interlock system (Section 3.6). The CVND controls the laser power that will reach the eye and is adjusted to eye safe levels.

The next element in the system is an Acousto-Optical Modulator (AOM). An AOM consists of a PZT that introduces high frequency sound waves into a transparent crystal. The mechanical strain introduced by the sound waves changes the materials refractive index, which is known as the acousto-optic effect (Korpel, 1981). The sound waves introduce a periodic modulation in the refractive index in which incident light will experience Bragg diffraction. By aligning the interferometer to the 1st diffracted order, the AOM is used as a shutter by modulating the PZT drive. The output of the 1st diffracted order is aligned to a pick-off mirror that folds the beam to the rest of the interferometer. When the drive voltage is removed from the PZT, the AOM no longer acts as a diffraction grating and transmits all light through the 0th order. Light in the 0th order does not reach the pick-off mirror and does not reach the rest of the interferometer.

Operation of the AOM as a Bragg device requires that the incident laser beam should enter the device at the Bragg angle:

$$\theta_B = \frac{\lambda}{2\Lambda} \qquad (3.5)$$

Where λ is the wavelength of light, and Λ is the sound wavelength, which is analogous to a grating line spacing. The acousto-optical material used is Tellurium Dioxide (TeO_2) which has a longitudinal sound velocity of $V = 4260\, m/sec$. The drive frequency for the PZT is $F = 80\, MHz$ which gives a sound wavelength:

$$\Lambda = \frac{V}{F} = \frac{4260\, m/sec}{80\, MHz} = 53.25\, \mu m \qquad (3.6)$$

Operation at $\lambda = 785\, nm$ results in a Bragg angle of $\theta_B = 7.4\, mrad$. The first diffracted order operates at twice the Bragg angle, so the angular separation between the on and off (i.e. deflected and non-deflected) beams will be $2\theta_B = 14.7\, mrad$. The angular deviation is important, because it requires enough path length from the deflection point and pick-off mirror to isolate the two beams. The nominal laser beam width requires a minimum of 3 mm separation, which results in a minimum path length of 200 mm.

Finally, the AOM drive voltage is synchronized to the electronic shutter (i.e. integration time) of the surface interferometer's camera (Section 3.5). The AOM response time is limited by the sound velocity, which for this device is approximately 100 ns and is sufficiently fast for this use. Typical exposure times are approximately 33 μsec with the camera frame rate of 30 FPS (33 ms), resulting in a 0.1% duty cycle. This effectively reduces the perceived irradiance of the laser source by human subjects by a factor of 1000.

After the AOM, a pick-off mirror is used to split the beam(s) between the laser power monitor and the interferometer. In place of an amplitude beamsplitter that was previously used on the TFI, the pick-off mirror has been replaced by a D-mirror (Figure 3.13). The advantage of the D-mirror in this configuration is that the AOM transmits two beams separated by a small angle. When the AOM is disabled, the non-deflected beam will pass over the D-mirror and is incident on the laser monitoring detector. For the 0.1% of the time that the AOM is enabled, the deflected beam will be incident on the D-mirror and folded towards the rest of the interferometer.

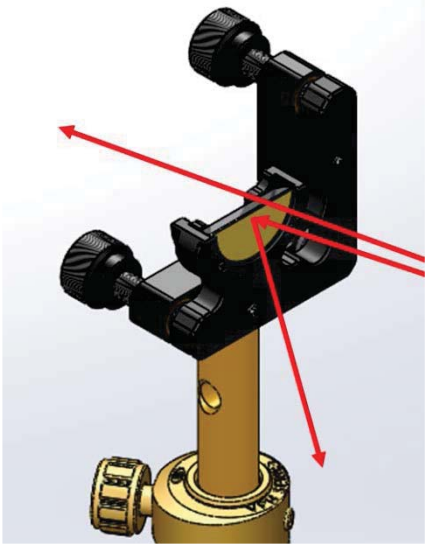

Figure 3.13: D-mirror configuration used to pick-off laser after AOM.

3.2.2 COLLIMATION ASSEMBLY

The collimation assembly is functionally the same as the TFI. After the D-mirror, the laser passes into a 40X microscope objective (Newport M-40X) and is spatially filtered with a 5 µm pinhole (Newport 910PH-5). After the spatial filter, the divergent light is collimated by a 150 mm focal length aplanat (CVI LAP-150.0-30.0-780). Following the collimating lens is a zero-order half-wave plate (HWP) (CVI QWPO-780-12-2-R15). Rotation of the HWP controls the orientation of the linear polarized light that enters the Polarizing Beam-Splitter (PBS), which allows for control over the relative intensity transmitted between the test and reference arms. The PBS has been upgraded from the previous TFI with a larger 2" cube to avoid issues of vignetting (Section 2.4) (Edmund Optics 48-866).

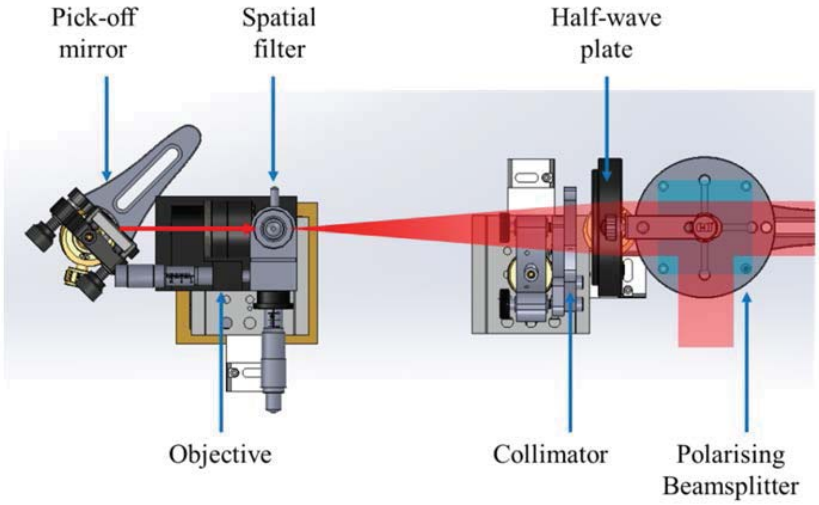

Figure 3.14: Surface interferometer collimation assembly.

3.2.3 REFERENCE PATH

The reference path has not been changed from the TFI other than the layout to support the new system. The path length is still extended to match the test arm despite the fact that this may no longer be necessary. The Fabry-Pérot interferometer (Section 3.2.1) allows for optimization of the source coherence length and has been shown to be sufficient enough to not require path matching. SPINALTOP was designed before the Fabry-Pérot interferometer was purchased, so it was assumed that the best coherence length for the 785 nm source was 0.7 m (Section 2.5).

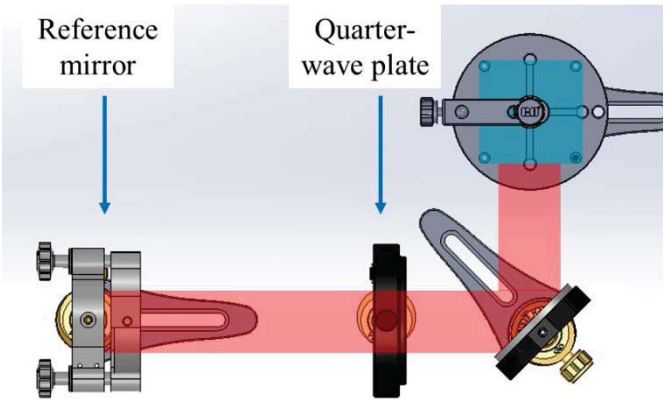

Figure 3.15: Surface interferometer reference path.

3.2.4 TEST PATH

The test path is shown in Figure 3.16. After exiting the polarizing beamsplitter, the light is reflected from a dichroic beamsplitter. The dichroic beamsplitter was custom manufactured by Spectral Products (Putnam, CT). The dichroic uses a 2.0” diameter, 0.375” thick BK7 window with a 30 arc-minute wedge. The coating was designed to reflect at 785 nm for the surface interferometer and to transmit at 850 nm (cat’s eye interferometer) and visible (400-700 nm, fixation assembly). A spectral measurement on a witness sample from the coating run shows 98.4% reflectance for 785 nm, 91.3% transmittance for 850 nm, and an 80% average transmission at visible (Figure 3.17).

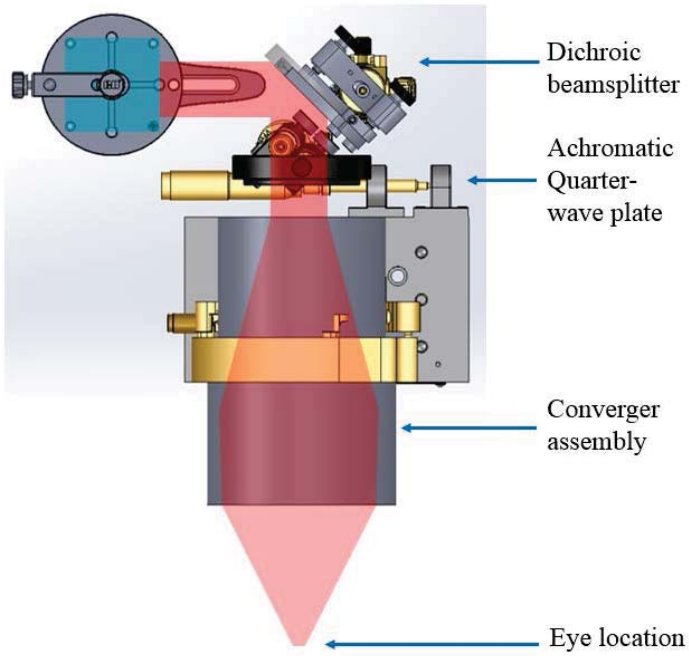

Figure 3.16: Surface interferometer test path.

Following the dichroic beamsplitter, the light passes through an achromatic quarter-waveplate (OptiSource LLC, ACWP-650-1100-38.10-R20). The achromatic waveplate was designed to accommodate both the 785 nm and 850 nm sources, and has no effect on the incoherent illumination used for the fixation assembly. The clear aperture of the quarter-wave plate was increased to 38.0 mm from 25.0 mm to minimize any potential vignetting that was previously discussed in Section 2.4.

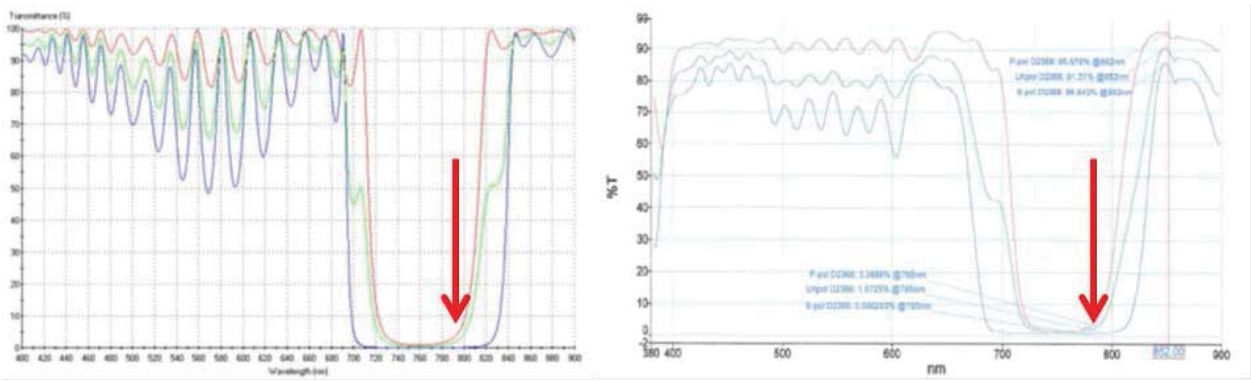

Figure 3.17: Spectral transmittance for the surface dichroic beamsplitter.
(Left) Designed. (Right) Measured. The red arrow is at 785 nm.

Directly after the wave-plate is the converger assembly, which reshapes the collimated beam into a converging wavefront designed to match the mean shape of the human cornea: A convex surface with a 7.8 mm base radius and a -0.25 conic constant. The requirement to re-use as much of the previous TFI hardware places a number of constraints on the new system design. The first is the issue of vignetting; if vignetting is occurring within the converger assembly, there is nothing that can be done about it. A second issue is that it limits flexibility of the design for the cat's eye interferometer and fixation assembly. A final issue is that the limited access to components within the converger assembly prevents proper characterization of the converger, which is required for the surface reconstruction methods that will be discussed in Chapter 4.

3.2.5 IMAGING PATH

Wavefronts from the test and references arms are reflected back through their respective paths and re-combine in the imaging arm. The re-combined wavefronts first pass through the imaging lens, which is designed to image the test surface (i.e. nominally shaped cornea) onto the detector. Typically, the imaging path is designed to minimize retrace errors (Huang, 1993; Z Malacara & Malacara, 1995). But, as will be discussed in Chapter 4, these errors will become unavoidable. Therefore, the imaging system was designed with the goal of simplifying the interferometer design, which in turn simplifies the reconstruction processing. The imaging lens is a 100 mm focal length, plano-convex lens (Newport, KPX187AR.16). The clear aperture has been selected to be 50 mm to minimize potential vignetting issues (Section 2.4). The unfolded imaging path layout is shown in Figure 3.19.

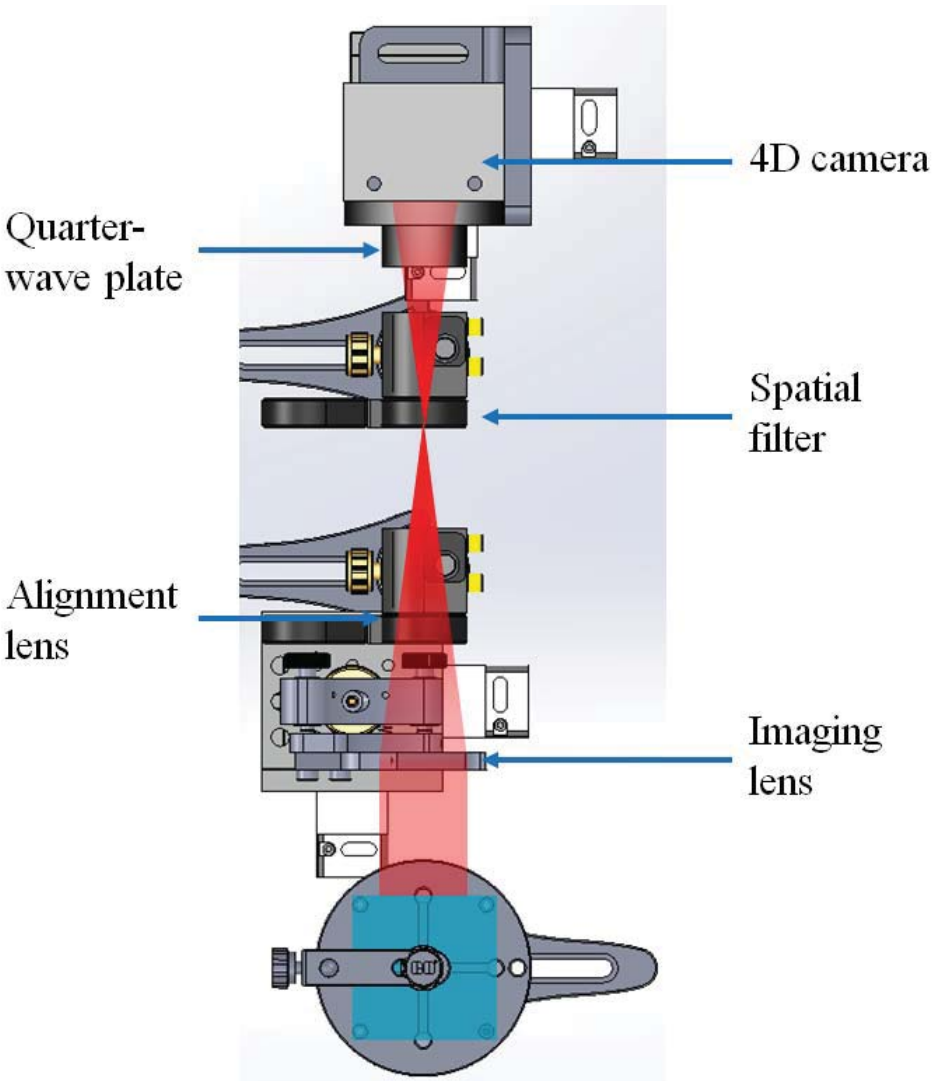

Figure 3.18: Surface interferometer imaging path layout.

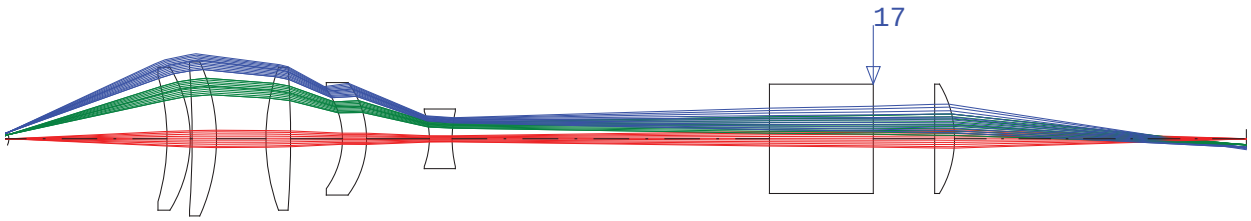

Figure 3.19: Unfolded surface interferometer imaging path.

Following the imaging lens, a pair of motorized flip-in mounts (Thorlabs MFF101) is used for alignment. The flip-in mount furthest from the imaging lens is at the focus of the imaging lens where an iris is placed. The iris is the effective aperture stop for the imaging system, or a low-pass filter for surface slope measurements. The aperture stop is also necessary to minimize stray light. The aperture is stopped down to limit rays that would generate fringes that exceed the Nyquist-sampling of the detector. This is accomplished by aligning the interferometer to a test surface and decentering the surface until aliased tilt fringes are produced on the detector. The iris is then stopped down until the aliased fringes are blocked.

Directly after the imaging lens and before the aperture stop is an alignment lens. The alignment lens is a -100 mm focal length bi-concave lens (Newport KBC037AR.16). Along with the lens is an absorptive neutral density filter with an optical density of 6.0 (Thorlabs NE60B-B). During normal operation, the alignment lens is stowed and the aperture stop is in the beam path. In alignment mode, the aperture stop is moved out of position and the alignment lens flips-in. This effectively shifts the focus of the spot to the surface of the detector. On the detector two spots should appear: A fixed focus spot from the reference path and another spot from the test path. The neutral density filter reduces the amount of energy that is being focused on the detector; without the filter, the focused beams would saturate and/or potentially damage the detector.

The final element in the imaging path is the 4D camera system (4D Technologies, Tucson AZ). More discussion detailing the sensor and synchronization with the AOM can be found in Section 3.5.1. Mounted to the face of the camera is a quarter-wave plate. The two linear polarization states of the test and reference arms are converted to circular polarization. The pixelated mask of the 4D sensor records the interference fringes and allows for the instantaneous phase to be recovered (refer to Section 4.1.1). The resulting measurement contains all of the surface feature information of the cornea and tear film. However, retrace error, distortion, and ambiguous phase recovery requires additional information to recover the true shape of the cornea. The cat’s eye measurement that is discussed in the next section provides this additional information and the algorithms discussed in Chapter 4 recover the true shape.

3.3 CAT'S EYE INTERFEROMETER

The cat's eye measurement is necessary to recover the absolute corneal topography from the surface measurement. The requirement to re-use the converger assembly placed a number of restrictions on the design of a system for simultaneously acquiring the cat's eye measurement. The result of this restriction was a second, nearly duplicate, interferometer built and co-aligned to the surface interferometer to provide the necessary cat's eye measurement.

The cat's eye interferometer sub-system is a polarization-splitting Twyman-Green interferometer that operates at the NIR wavelength of 850 nm. Laser safety is handled the same way as it has been for the TFI and surface interferometers, which is reviewed in Section 3.6. The cat's eye interferometer is folded in-line with the surface interferometer with the addition of secondary optics to shift the focus of the 850 nm beam directly on to the surface of the cornea, producing a cat's eye measurement. Phase is measured instantaneously through a spatial carrier method (Section 4.1.2). The spatial carrier method is a cheaper alternative to the pixelated phase camera system used in the surface interferometer, but at the cost of reduced resolution. The following sections will discuss the layout and operation of the cat's eye interferometer.

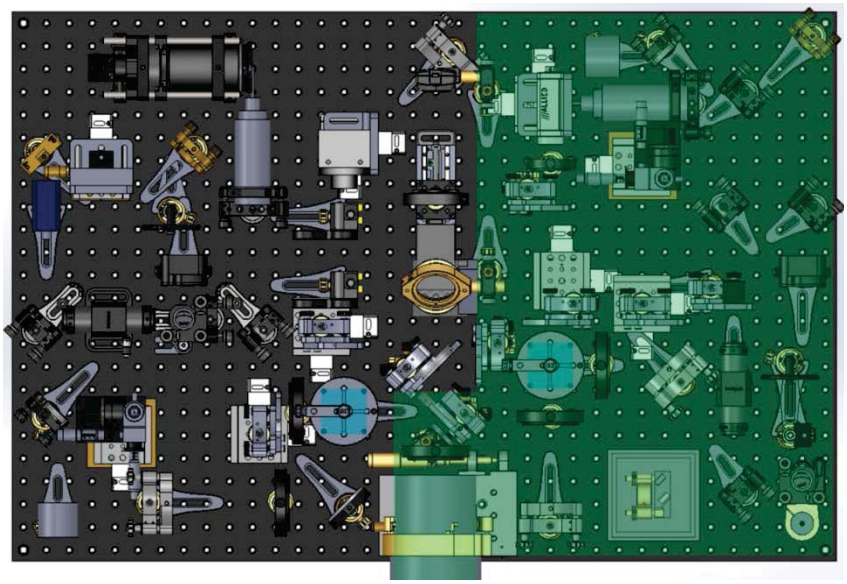

Figure 3.20: Hardware layout of the cat's eye interferometer (highlighted in green).

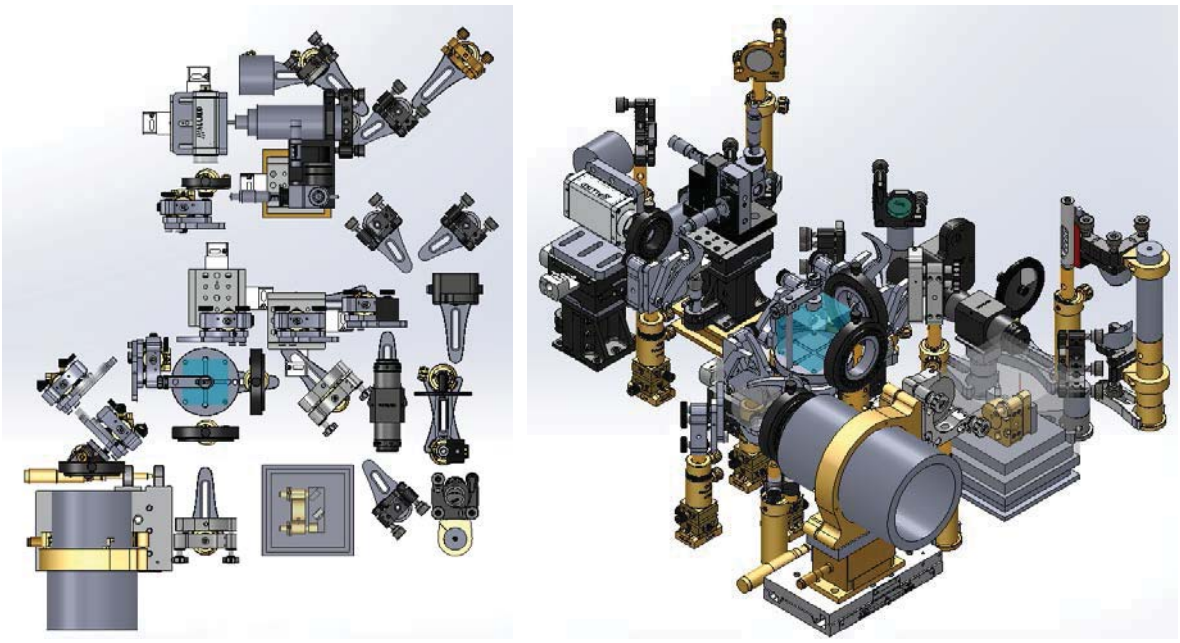

Figure 3.21: Isolated cat's eye interferometer hardware layout.
(Left) Top view. (Right) Isometric view.

3.3.1 SOURCE ASSEMBLY

The source assembly for the cat's eye interferometer is similar to the surface interferometer. The difference in the layout that is shown in Figure 3.22 is a result of breadboard space constraints and a slightly larger space requirement for the custom laser assembly. One of the lessons learned from the TFI system was to invest in a better laser source; however, the budget was limited for this system. It was shown in Section 3.2.1 that with active monitoring, a semi-stable laser source could be made to work with the interferometer. The solution was to build the laser source in-house and duplicate the in-line monitoring of the surface interferometer source to ensure stability and an optimized coherence length.

The concept for building a custom laser source was inspired by the Atomic Physics Group (a research group within the College of Optical Sciences), where it is common to build narrow bandwidth laser sources that can be built for less than $1000 (Wieman & Hollberg, 1991). The custom laser that is used for this system is an external-cavity diode laser (ECDL) in a Littrow configuration (Figure 3.23) (Fleming & Mooradian, 1981). The ECDL consists of a semiconductor diode followed by a grating that sets up an

external cavity. The semiconductor device used for the laser was a single mode laser diode operating at 850 nm (QPhotonics QLD-850-150S). The bandwidth of the laser diode is $\Delta\lambda = 1\,nm$, equivalent to $\ell_c = 0.5\,\mu m$. The output face of the diode is a small rectangular region, which results in a large asymmetric divergent output. The diode used for this laser has an 8º divergence along one axis and 18º in the orthogonal axis. The output has to be collimated before the external cavity grating, so a 6.24 mm focal length, 0.40 NA aspheric collimating lens is used (Thorlabs LT110P-B). The collimated light is elliptical with approximate dimensions of 0.9 x 2.0 mm.

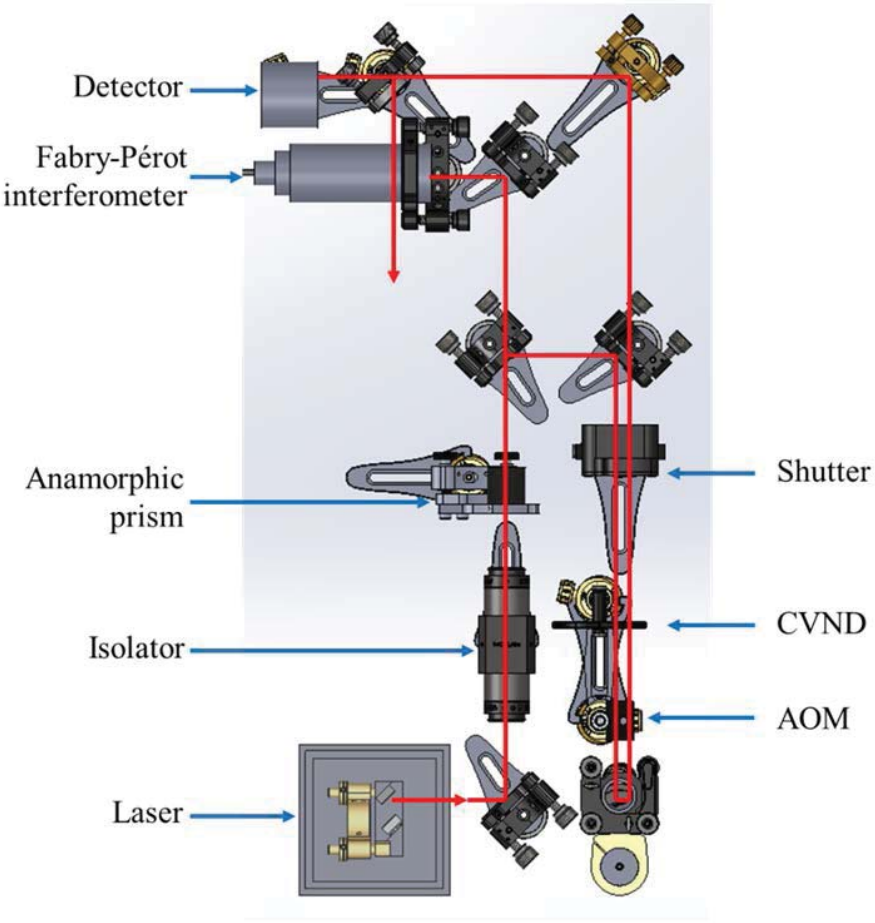

Figure 3.22: Cat's eye interferometer source assembly.

The elliptical beam is incident on a ruled reflective diffraction grating with 1800 Grooves/mm (Thorlabs GR13-1850). The grating is aligned relative to the diode to maximize the resolving power. In other words, the long-axis of the elliptical beam is projected along the grating to cover as many line-pairs as possible.

This will ensure that the optical feedback set up by the grating can result in a narrow linewidth. The grating is aligned in the Littrow configuration, where the first-order diffracted beam is fed back into the diode. The optical feedback is what sets up the cavity and allows for fine selection of wavelength. However, it is important to note that the selected grating is optimized for use at 500 nm. The efficiency for the grating at 850 nm is 10% in this configuration, which has been purposely chosen to protect the diode from oversaturation and possible damage.

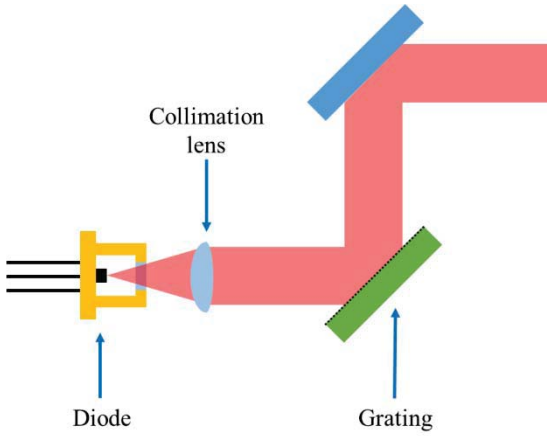

Figure 3.23: External cavity diode laser in a Littrow configuration.

Alignment and stability of the grating is critical to the overall operational stability of the laser. Therefore, consideration has to be made for mounting and alignment of the grating. First, the angle of the grating can be solved from the grating equation:

$$\sin(\theta_i) - \sin(\theta_o) = \frac{m\lambda}{d} \qquad (3.7)$$

Where the incident and output angles are given by θ_i and θ_o, respectively, the diffracted order m, wavelength λ, and groove spacing d. The Littrow configuration results in $\theta = \theta_i = -\theta_o$ and $m = 1$, allowing the angle to be solved for:

$$\theta = \sin^{-1}\left(\frac{\lambda}{2d}\right) = \sin^{-1}\left(\frac{850\,nm}{2\,(1/1800\,mm)}\right) = 49.9° \qquad (3.8)$$

The grating can be mounted at the nominal angle followed by an additional mirror to simplify alignment. By mounting the mirror and grating on the same mount, the grating/mirror combination acts as a beam steering device which maintains the nominal output angle of the laser. Although beam translation will occur when the grating/mirror assembly is adjusted, the expected translation is small enough that alignment down-path in the interferometer can be maintained if an adjustment has to be made to the laser source assembly.

The grating/mirror assembly is mounted on a tip/tilt stage directly in front of the laser diode. The tip/tilt stage is a repurposed 1” mirror mount. The diode can be mounted into the mirror mount and the grating/mirror assembly mounted to the base (Figure 3.24). The tip/tilt adjusters on the mirror mount are used to fine-align the grating to the diode. Alignment of the laser is accomplished by monitoring power output and adjusting the tip/tilt to optimize output power. The process starts by setting the drive current of the laser diode to just below its lasing threshold and watching for a spike in power as the tip/tilt is adjusted. The spike in power is referred to as a flash point and indicates that optical feedback has occurred and the diode started lasing. This can be repeated by continually dropping the diode output power below lasing levels and searching for the flash point until an optimal alignment is found. A final tuning is performed when the laser is fed into the Fabry-Pérot interferometer.

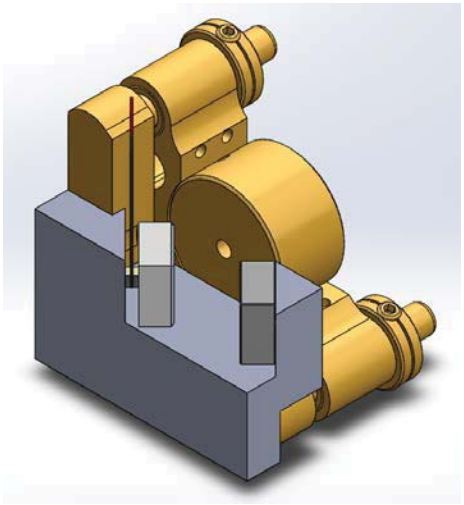

Figure 3.24: ECDL mirror and grating setup.

Figure 3.25 is a Fabry-Pérot interferometer measurement of the laser source after tuning. The oscilloscope measurements are calibrated (left) using the method described in Section 3.2.1 and the linewidth is determined from the measurement (right). The measured linewidth of the source shown is 29 MHz which results in coherence length of 8.3 m. This is sufficient for this application and so further optimization on the source is no longer necessary. If desired, the laser cavity can be fine-tuned to reduce the linewidth, extending the coherence length of the laser source.

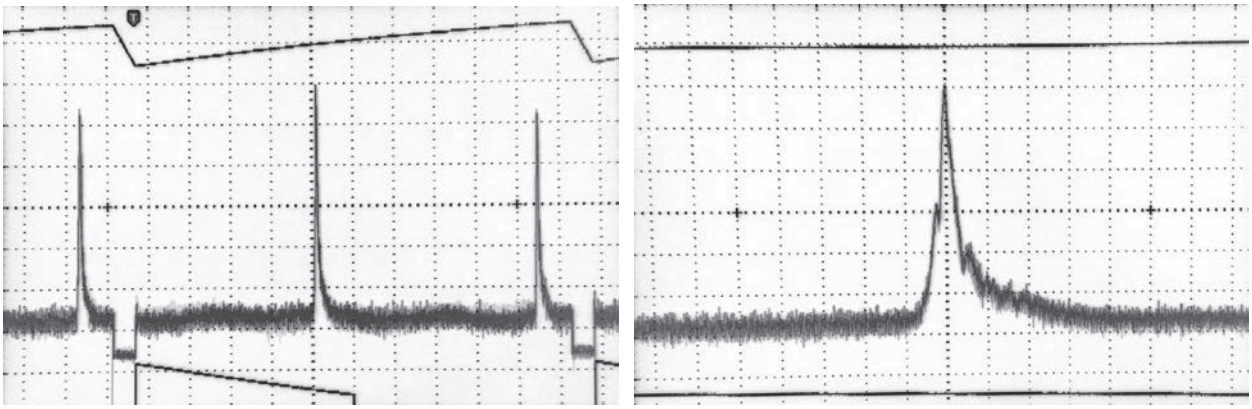

Figure 3.25: Fabry-Pérot interferometer measurements of the cat's eye 850 nm laser.
(Left) Oscilloscope set to view multiple orders for calibration. (Right) Spectrum of tuned ECDL.

The laser assembly is mounted to a 100 x 100 x 25 mm aluminum base plate (Figure 3.26). A thermoelectric cooler (TEC) (Laird CP14,71,10,L1,W4.5) is placed between the baseplate and a large aluminum heat sink. A 10 kΩ thermistor (Arroyo Instruments, 1600-10K) is mounted to the diode collimation tube for monitoring. The TEC is controlled by a temperature controller (Arroyo Instruments, 5240) and can maintain thermal stability to within 4 mK over an hour, or 10 mK over 24 hours. The temperature controller set point is kept at 25.0 C (298.15 K).

The final operation required for this laser is to reshape the beam. The elliptical output has an aspect ratio of 2:1, so a 2X anamorphic prism pair (Thorlabs PS875-B) is used to convert the beam to a 2 mm diameter circular beam. However, the laser is sensitive to optical feedback, requiring an isolator between the anamorphic prism and laser. The isolator is able to support a maximum beam diameter of 4.7 mm.

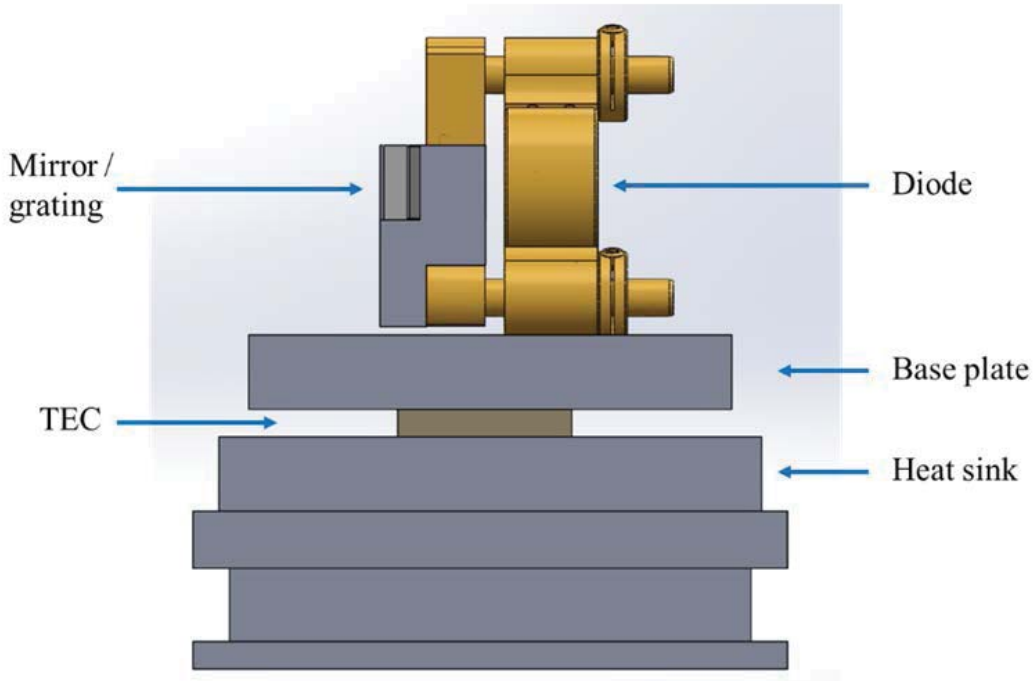

Figure 3.26: Side-view of the cat's eye laser source assembly.

Following the anamorphic prism, the system is functionally identical to the surface source assembly. A pickoff beamsplitter sends 10% of the light to a Fabry-Pérot interferometer to monitor the laser. The remaining is 90% of light passes is reflected through the shutter, CVND, AOM, and is finally split between to the laser power detector and on to the collimation assembly.

3.3.2 COLLIMATION ASSEMBLY

The collimation for the cat's eye interferometer is similar to the surface interferometer. The same 40X objective and 5 μm pinhole are used for focusing and spatially filtering the laser source. The collimation lens is a 200 mm focal length achromatic doublet designed for the NIR (Thorlabs AC508-200-B). A fold mirror following the objective is required due to limited space on the breadboard. The folded beam passes through a custom zero-order half-wave plate designed for 850 nm (OptiSource LLC, QWPO-852-15-2-R20). Following the waveplate is a 2" polarizing beamsplitter cube (Edmund Optics, 48-868).

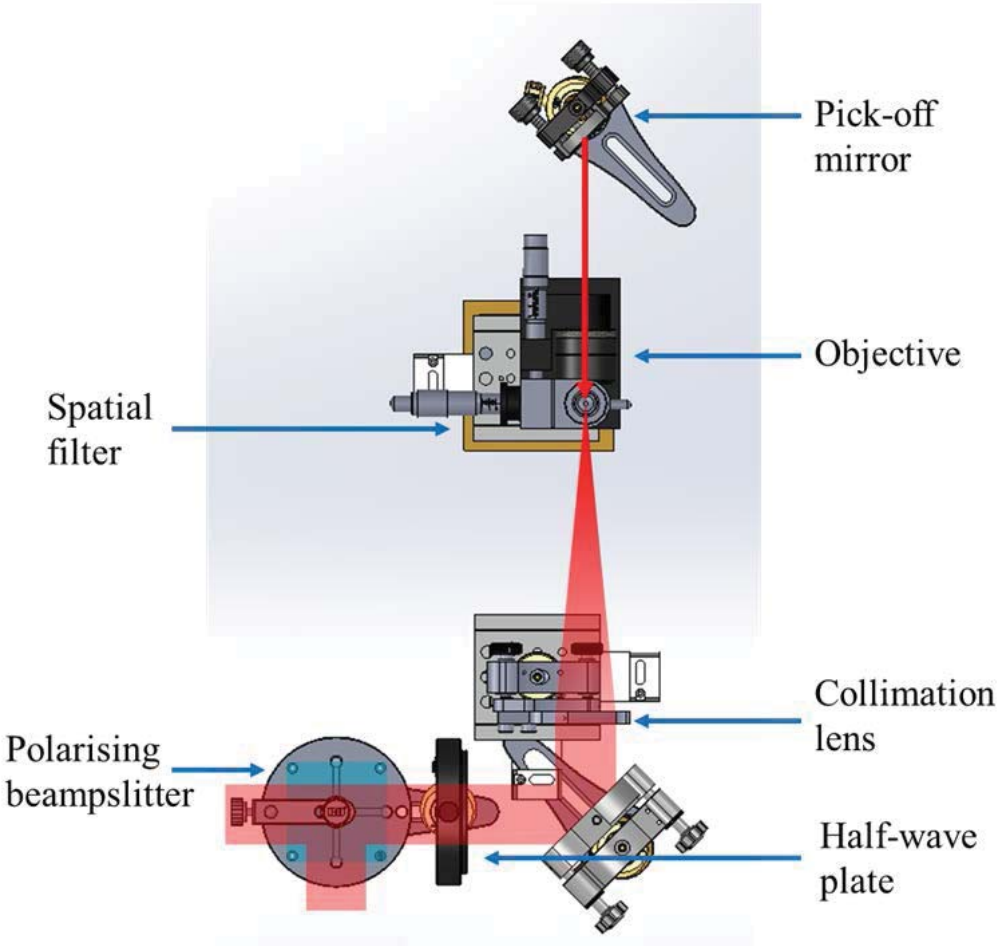

Figure 3.27: Cat's eye collimation assembly.

3.3.3 REFERENCE PATH

The reference path contains a custom zero-order quarter-wave plate designed for 850 nm (OptiSource LLC, QWPO-852-15-4-R20) and a 2" protected silver reference mirror (Edmund Optics, 43-416-577). The ability to tune the laser and ensure a longer coherence length meant that it was no longer required to path match the test and reference arms, so little effort was expended in trying to extend the reference arm as was previously done on the surface interferometer and TFI.

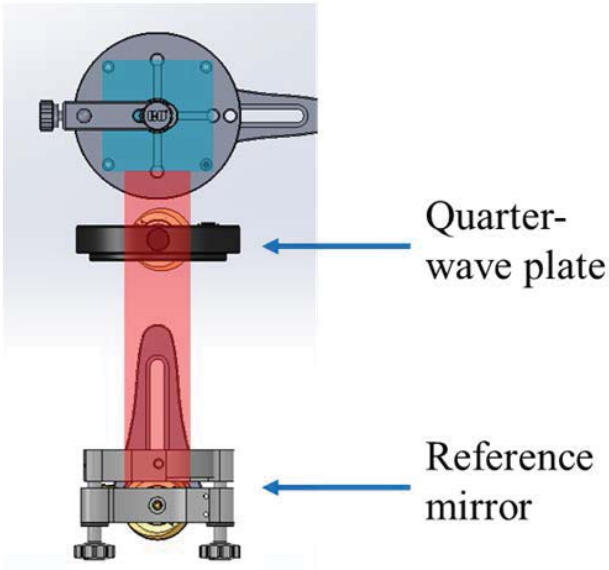

Figure 3.28: Cat's eye reference path.

3.3.4 TEST PATH

The test path is shown in Figure 3.29. The first element in the test path following the polarizing beam splitter is the cat null lens. The cat null is designed to shift the focus of the 850 nm laser source to the surface of the cornea when a nominally shaped cornea is placed at a confocal position relative to the surface interferometer. The system is aligned so that a nominally shaped cornea can be nulled simultaneously on the surface interferometer in a confocal configuration and nulled on the cat's eye interferometer in a cat's eye configuration. A precision glass surface will be used to represent the nominally shaped cornea when aligning and testing the two interferometer systems (more detail in Section 4.3.2.3). As previously discussed, the purpose of the cat's eye interferometer is to provide a supplementary measurement that is necessary to reconstruct the surface topography of the cornea. The cat's eye measurement is used to determine the axial location of the cornea. Surface reconstruction accuracy is a directly related to the resolution of the cat's eye measurement. Therefore, the cat null must provide a sufficiently fast f/# measurement that produces enough sensitivity to track axial motion.

A cat's eye measurement can be easily obtained by introducing defocus in the test arm of the cat's eye interferometer. However, significant spherical aberration will be present for three primary reasons. First,

the converger was designed with an amount of spherical aberration to accommodate the aspheric test surface. Second, the cat null focus is different from the designed focus which introduces additional spherical aberration. Finally, the use of an off-design wavelength will introduce a small amount of spherochromatism. It was found that a meniscus lens could be used for the cat null to compensate for most of the induced spherical aberration.

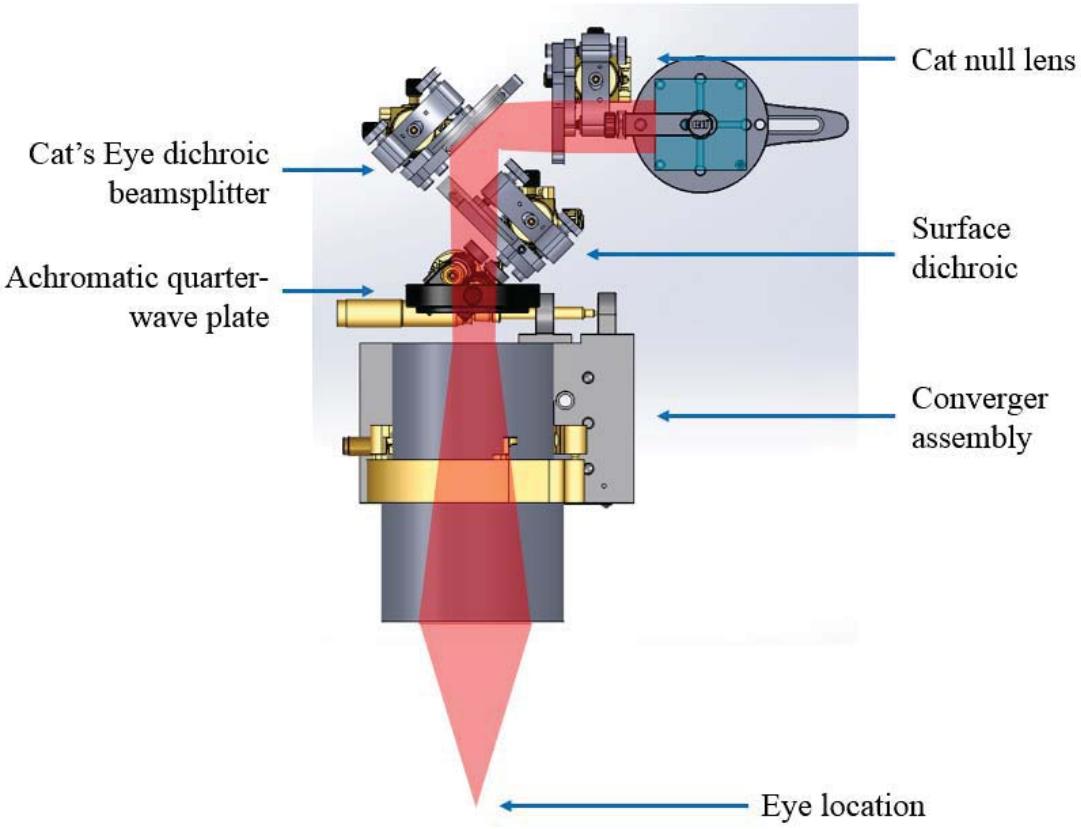

Figure 3.29: Cat's eye test path.

However, a complication arises from the hardware re-use requirement. The cat's eye sub-system has been treated as an add-on to a pre-existing system. Ideally, the converger would be designed with simultaneous wavelength operation in mind (Section 6.2.1). Folding in a secondary interferometer system required the use of a dichroic beamsplitter to combine the two systems, which added an additional 75 mm of path length to the cat's eye interferometer. This path length increase, as will be shown, nearly halves the working f/#. The interferometer could be laid out such that the cat's eye interferometer feeds into the

converger first, but this was avoided because it would require the surface interferometer to transmit through the dichroic beamsplitter and the surface interferometer is given priority. Another solution would be to include a beam expander as part of the cat null, but this increases the complexity of the design and would vignette on the dichroic beamsplitters. The fixation could be placed in-line with the cat's eye interferometer, but because of the previous issues with fixation stability (Section 2.2.2), priority was given to the fixation assembly (Section 3.4). The final decision was to keep the cat's eye interferometer as simple as possible for the purpose of this setup and thus a cat null was chosen.

The cat null was designed by adding a negative optical element 175 mm from the back vertex of the first converger optical element. The distance was determined by how close the optical element could be fit into the system, which includes the pair of dichroic beamsplitters. The optimized solution resulted in a 320 mm focal length meniscus lens. However, wavefront quality in the cat's eye interferometer is not as strict a requirement as it is in the surface interferometer, so an inexpensive off-the-shelf solution was sought. The closest lens was a 300 mm focal length meniscus lens from Thorlabs (LE1985-B). Using the off-the-shelf lens instead of the custom lens, the system was re-optimized for distance only. This changed the lens spacing from 175 mm to 192 mm, which resulted in a slower working f/# and further reduced the sensitivity of the cat's eye system. The final configuration is shown in Figure 3.30. The residual error for this design is less than $\lambda/20$ PV.

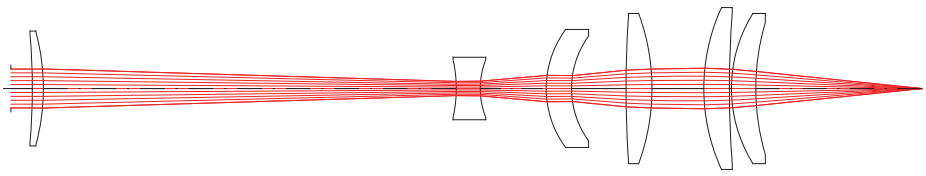

Figure 3.30: Cat null lens optical layout.

It might appear that an important oversight is the absence of the surface dichroic beamsplitter from the cat null design. The presence of a thick wedge at a 45º in a focusing beam will introduce astigmatism. However, the purpose of the cat's eye system is not to measure the surface quality of any particular

feature but to resolve axial changes in distance. Although astigmatism will be present in measurements, the calibration process can account for this (Section 4.3.3). The additional axial displacement due to the windows can be compensated by focusing the cat null.

The defocused beam emerging from the cat null is reflected by the second dichroic beamsplitter in the system, another custom fabricated beamsplitter from Spectral Products. The beamsplitter reflects approximately 99.4% at 850 and transmits 90% of visible light (Figure 3.31). The beam then transmits through the first dichroic beamsplitter, the achromatic waveplate, and finally the converger. The beam will come to focus at the surface of the nominally shaped cornea that is aligned to the surface interferometer.

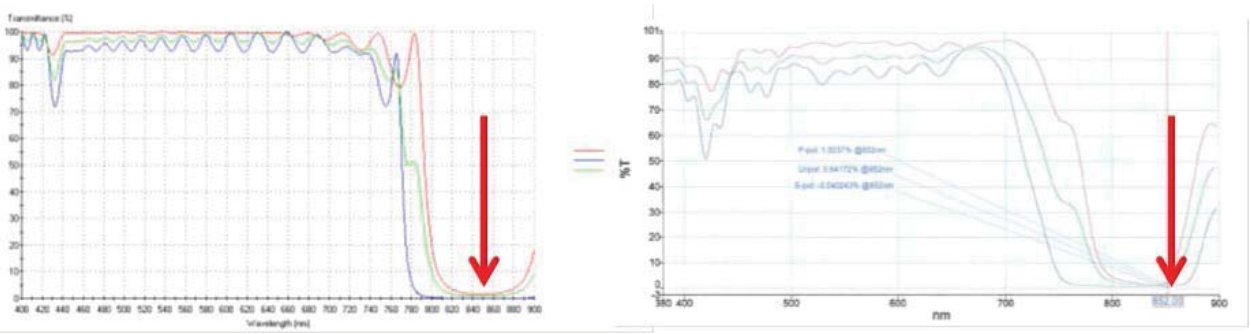

Figure 3.31: Spectral transmittance for the cat's eye dichroic beamsplitter.
(Left) Designed. (Right) Measured. The red arrow is at 850 nm.

3.3.5 IMAGING PATH

Test and reference wavefronts are combined in the imaging path. A hardware layout is shown in Figure 3.33 and the optical layout is shown in Figure 3.32. The first element in the imaging path is the imaging lens, although this is a slight misnomer. The cat's eye interferometer is not making a surface measurement, so there is no object conjugate to be imaged. In this configuration, the focus of the imaging lens would be conjugate to the surface of the cornea. Instead, the imaging lens is setup to provide a point focus for both the test and reference wavefronts. Near the focus of the imaging lens is a Wollaston prism, which generates a large amount of tilt between the test and reference wavefronts. The large tilt term acts

as a spatial carrier for the phase difference between the two wavefronts and can be recovered from a single snapshot (Section 4.1.2).

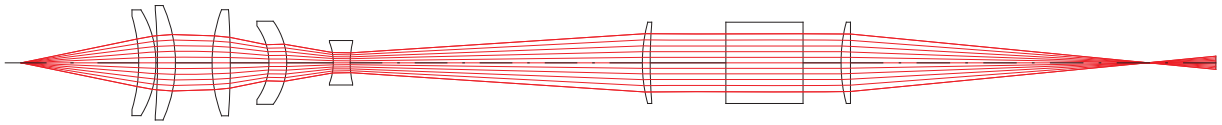

Figure 3.32: Optical layout for the cat's eye imaging path.

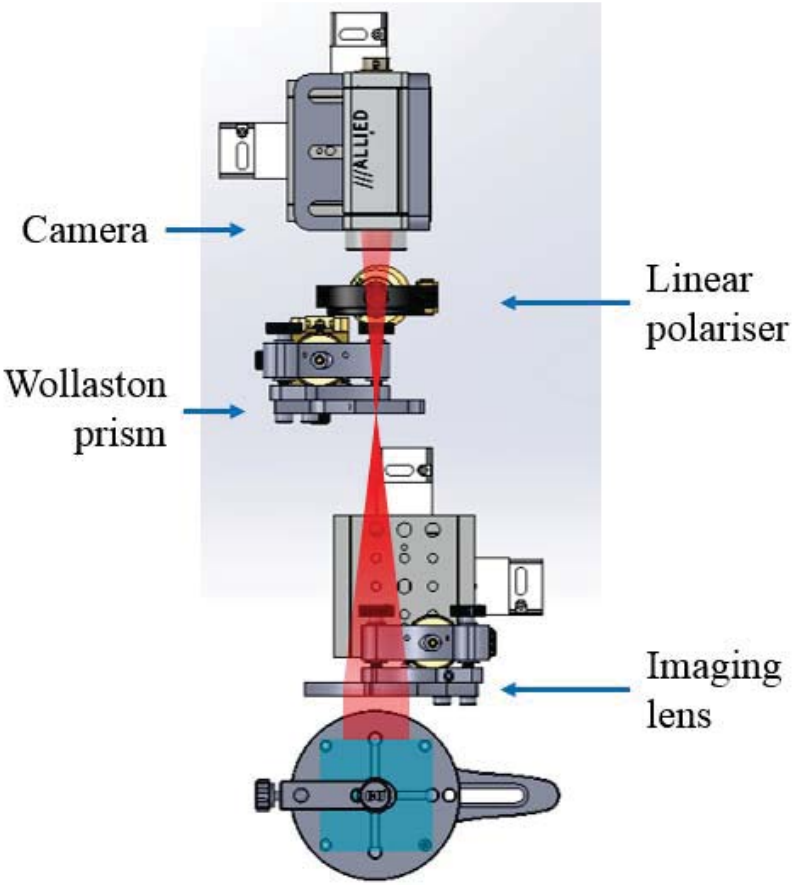

Figure 3.33: Cat's eye imaging path.

A Wollaston prism is an optical element that splits unpolarized light into two orthogonally polarized outputs (Figure 3.34). It consists of two birefringent right angle prisms cemented together such that the fast axes are perpendicular. With proper alignment of the Wollaston prism, the orthogonal linear polarization states of the test and reference wavefronts will offset generating tilt in the observation plane.

Translation of the prism along the optical axis provides a method for adjusting the amount of tilt between the two emerging wavefronts.

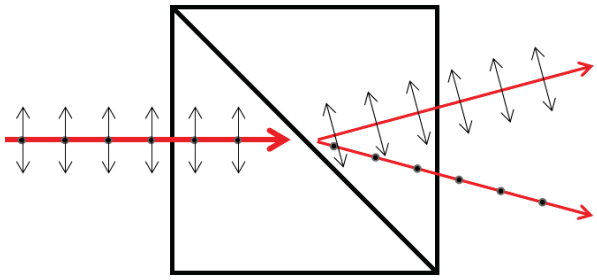

Figure 3.34: Wollaston prism.

The orthogonal linear polarization of the sheared wavefronts will not produce interference. A linear polarizer is placed between the prism and detector, oriented 45º relative to both polarization states. While observing interference fringes, the Wollaston prism can be adjusted along the optical axis to introduce enough tilt for the spatial carrier method (refer to Section 4.1.2).

3.4 FIXATION ASSEMBLY

The TFI placed the fixation assembly external to the interferometer. The system was setup such that the subject's non-dominant eye would be tested with the interferometer, while the dominant eye would fixate on the target. This was sufficient for initial testing of an *in vivo* interferometer system, but had limitations that are corrected in this new system. Most ophthalmic instruments, such as corneal topographers, have in-line fixation which allows for easy translation between eyes and ensures that the eye is nominally aligned to the optical axis of the instrument.

An in-line fixation is a more elegant solution compared to the external fixation, which required two separate alignments – one to align the fixation assembly to the subject's eye, followed by a separate alignment of the instrument to the subject's other eye. Another advantage for moving the fixation inline is that it allows for accommodation to the subject's visual errors. The external fixation was limited to only a few diopters of correction, whereas an adjustable re-imaging lens internal to the system can allow for

±10 diopters of correction. In addition to the fixation system, an attempt was made to place a pupil imaging system in line with the system to gather additional information on eye motion and behavior.

3.4.1 FIXATION

The fixation system is designed to present a collimated, or distant, object for the subject to fixate on (Figure 3.36). External fixation with the TFI was accomplished by placing the target at the focus of a lens. With in-line fixation, the converger lens now has to present the collimated object. The back focus of the converger is approximately 17 mm inside of the assembly, which will require the fixation optics to image the target into the converger. The reimaging optics will also have to compensate for additional aberrations introduced by the converger optics and the pair of dichroic beamsplitters in the beam path that are rotated 45º along the same axis.

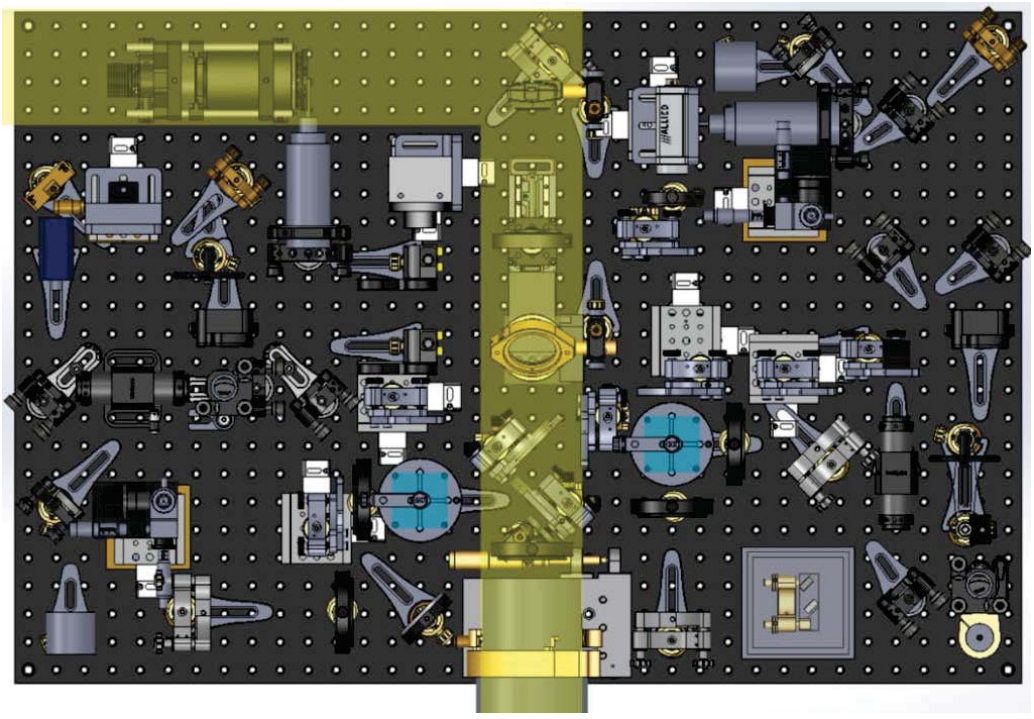

Figure 3.35: Hardware layout of the fixation assembly (highlighted in yellow).

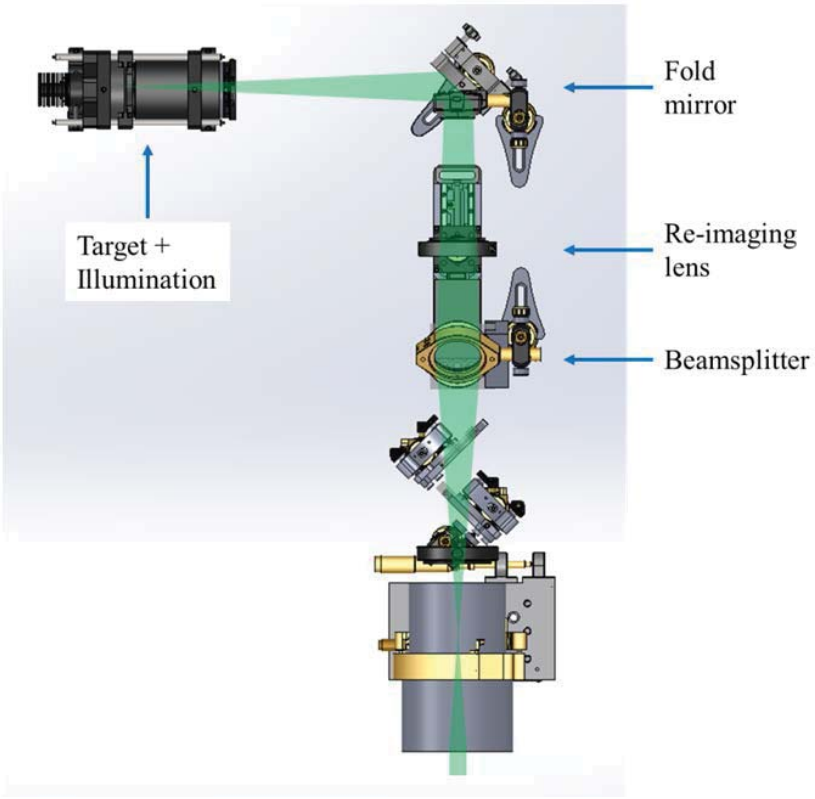

Figure 3.36: Fixation hardware layout.

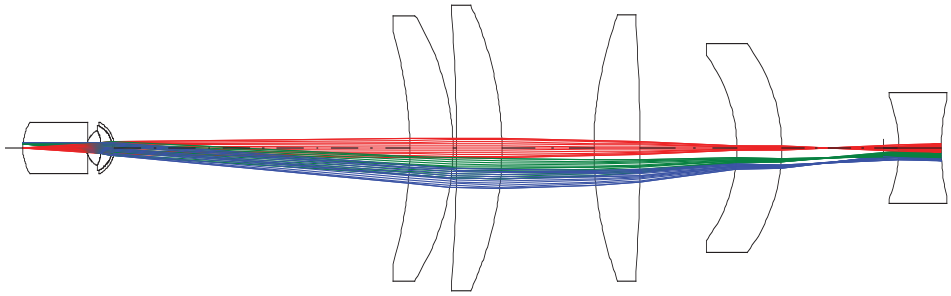

Figure 3.37: Location of back focus of the converger optics.

The parameters for the fixation assembly are derived with the intention to maintain fixation. The first requirement is to cover the fovea with some additional margin. The fovea is a 1.5 mm wide area of the retina consisting mostly of tightly packed cones that provides high resolution vision. Fixation involves centering the target on the fovea. Extending the designed fixation field of view to 3 mm area on the retina will ensure that fixation zone is controlled, thereby minimizing potential distracting artifacts that might appear in the extra-foveal zone. The aperture for the system will be set by the subject's pupil and will be

designed for a light adapted environment of 3 mm. Testing is typically done in a brightly lit room; additionally, the fixation illumination will be kept at sufficiently bright levels to force the subject's pupil to stop down. This has two benefits; the first is that it will help minimize aberrations that will still be present. The second is that by designing the system to match to a smaller pupil of the eye, fixation will constrain the eye location and gaze direction to a smaller region. This minimizes the uncertainty in locating the placement of the eye with respect to the interferometer when reconstructing the corneal topography.

Another requirement is to maintain a level of magnification that would support current existing fixation targets in the lab. The targets have line widths on the order of 50 μm and features (e.g. circular bullseye pattern) on the order of 1 mm. If the magnification is too large, the outer concentric circle could fall out of the fixations field of view, and if the magnification is too small the line widths will no longer be resolvable. The final requirement is that the system should be able to accommodate at least +/-5 D of correction. When testing, subjects will have to remove eyewear or contact lenses (for bare eye tests).

The system is designed in reverse to simplify the ray tracing process. The object is set at the retina with a 1.5 mm half field of view. The Arizona eye model is used to simulate eye performance, which includes accommodation and refractive errors (J. T. Schwiegerling, 2004). The aperture stop is set at the pupil with a 3 mm diameter. Based on results from the TFI, green was a preferable color for target illumination, therefore the design is set for 546 nm (the e-line, a mercury transition line). Choosing a monochromatic source reduces the complexity of the fixation optics. The eye is placed such that the cornea is at the nominal confocal position, 77.4 mm away from the vertex of the last lens of the converger. An initial ray trace without a fixation system shows where the intermediate image is located, between the first two lenses of the converger (Figure 3.37). The resulting magnification is -2.14.

The path length from the converger to the reimaging lens is nominally set at 250 nm. This accounts for the distance that is covered by the achromatic waveplate and dichroic beamsplitters (130 mm), additional

compensating windows (90 mm), and margin for design variation and focus adjustment (30 mm). The two dichroic beamsplitters are rotated at 45º along the same axis and will introduce astigmatism in the fixation path. A pair of equal thickness windows are placed in the fixation path and are rotated by an equal amount along the orthogonal axis. The compensating windows generate an equal amount of astigmatism that is rotated by 90º relative to the astigmatism generated by the dichroic beamsplitters. The addition of two orthogonal astigmatic terms results in an overall power term. The power error can be corrected by adjusting the axial position of the imaging lens.

The optical path length from the intermediate image to the initial lens is approximately 300 mm. The initial design starts with a 150 mm focal length NBK7 biconvex lens to setup a 1:1 magnification system. Optimization on the radii does not yield very much improvement, resulting in at PV of 11.3 waves of error over all fields. Splitting the lens and optimizing yields a slightly meniscus lens followed by a biconvex lens with a PV error of 9.1 waves. The initial 1:1 magnification requiring an approximately 150 mm lens is unnecessarily constraining the performance of the system, so the distances are relaxed slightly. The total path length from the intermediate image to the reimaging lens can be increased to 330 mm while maintaining space constraints within the SPINALTOP enclosure. The error for the relaxed system is 2.6 waves, PV. The magnification from the intermediate image to the target is now 1:1.5, or from the eye to target is 1:2.1. This is an acceptable level of magnification of the system.

The next step in the design process is to replace the custom lenses for commercial off-the-shelf lenses. The second lens from the split doublet has radius of curvature of 212 and 7100 mm, which is nearly a plano-convex lens with a 400 mm focal length. A close match is a 400 mm plano-convex lens from Thorlabs (LA1725-A). Replacing the second lens with the Thorlabs part and optimizing the system again yields a first lens that has radius of curvature of 480 and 230. A reasonably close lens is a 300 mm bi-convex lens from Thorlabs (LB1917-A) that has radius of curvature of 308 mm. The image location is optimized for the final configuration, and the resulting PV error is still 2.6 waves. The final configuration is shown in Figure 3.38 and residual errors in Figure 3.39. The residual design error will not be an issue.

Previous studies on fixation have shown that image quality of the fixation target have no effect on fixation except in the case where any of these factors render the target barely visible (Boyce, 1967; Steinman, 1965).

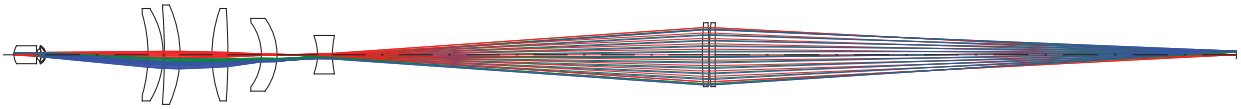

Figure 3.38: Fixation optical path layout.

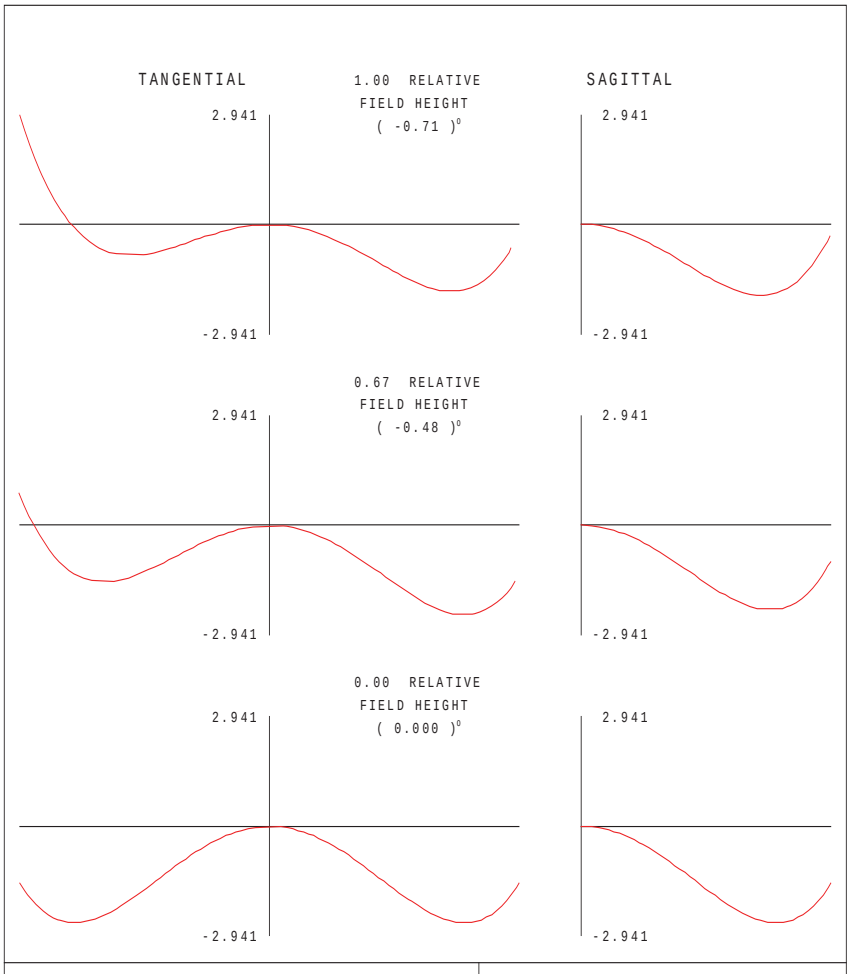

Figure 3.39: Residual OPD errors for fixation imaging path.

To verify that the fixation lens can accommodate for variation in eye power, the system is reversed in the lens design software (Figure 3.40). The target to reimaging lens distance is 394 mm, and reimaging lens to converger is 280 mm, resulting in a reimaging track length of 674 mm. The track length is fixed while allowing the reimaging lens position to vary. The eye model can be updated with an amount of accommodation and the reimaging lens position is optimized in the model. A +10 D accommodation applied to the eye requires a 9 mm translation of the reimaging lens from the nominal position. A -10 D accommodation requires 7 mm translation in the opposite direction. Adjustment of the fixation lens is accomplished by mounting the lens to a 50 mm motorized translation stage (Thorlabs, MTS50-Z8E).

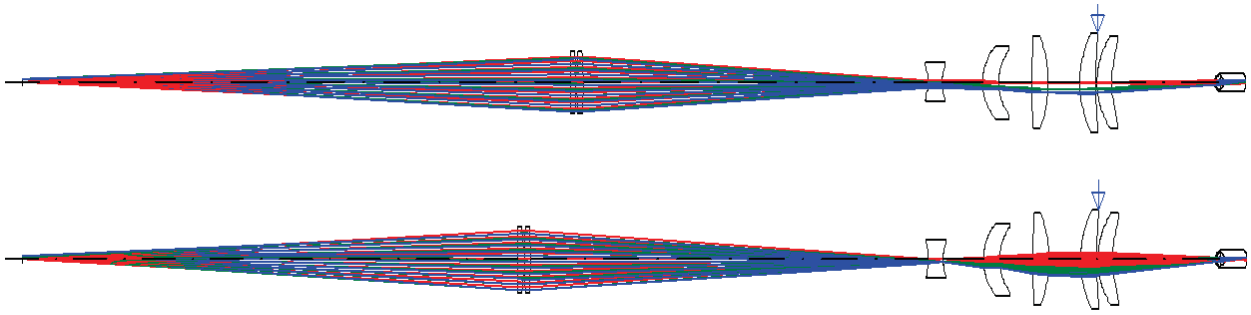

Figure 3.40: Accommodation of the fixation lens to account for variation in eye power.

The entrance pupil (i.e. conjugate to the pupil of the eye) is 115 mm from the target. A baffle tube is built around the illumination and target up to the entrance pupil location, where an iris is placed (Figure 3.41). The baffle is necessary to minimize stray light, which is a concern if it reaches the subject. Stray features in the subject's field can result in an increase in saccadic motion (Section 1.2). The iris can also act as the aperture stop for the system, bypassing the need to rely on the subject's pupillary response. The iris is set to match the nominal design of a 3 mm eye pupil diameter.

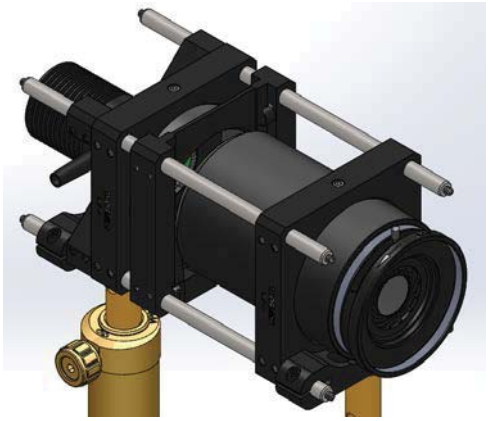

Figure 3.41: Baffle system built around the fixation illumination and target.

3.4.2 PUPIL TRACKING

Placement of an in-line fixation system opens access to direct on-eye imaging. An attempt was made to implement a system that could image the human subject's pupil (Figure 3.42). The TFI contained an external pupil camera that viewed the eye obliquely, which was placed below the converger assembly. A pupil imaging system could be potentially improved by placing it in line with the rest of the system. The fixation path contains two additional windows for astigmatism compensation. It is easy to replace one of the windows with a beamsplitter to pick off residual light scattered from the pupil.

A few problems arise when attempting to image the eye pupil through the converger. First, the pupil is near the focus of the converger assembly. The converger is designed with effectively zero field of view and the portion of the pupil that is of interest is off-axis. Secondly, the converger is designed for an aspheric wavefront at a NIR wavelength. The focal length of the converger and expected pupil size will result in a divergent beam from the rear (i.e. collimated input) of the converger. By the time the beam is accessible, the total beam footprint will be greater than 50 mm in diameter. Instead of attempting to capture the full 50 mm diameter, only the top half of the eye pupil will be captured to reduce the burden on the design of on the pupil imaging system.

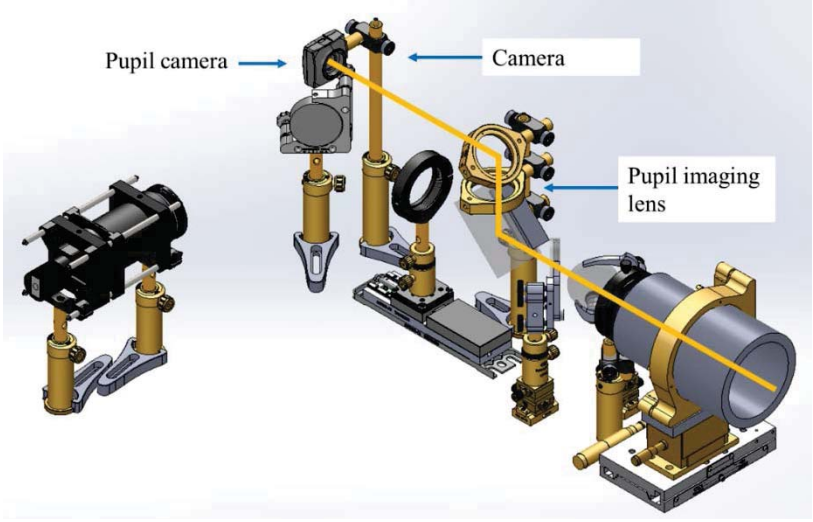

Figure 3.42: Pupil imaging pick-off path in the fixation assembly.

A ray trace showing an off-axis system capturing half of the pupil is shown in Figure 3.43. The nominal object location is placed 1.5 mm off-axis to match the expected location of the pupil. A full field of view is defined as ±1.0 mm about the nominal object. The remaining optics in the eye are built up around the object using the Arizona eye model, such that the object (i.e. pupil) is immersed in an aqueous solution. The aperture stop is placed at the first element of the converger assembly that is closest to the eye. The rays are traced 250 mm through the system to the closest accessible plane in the system. A number of issues can be observed at this point. The first is that the beams are severely aberrated and will require significant correction to be able to image the pupil. The second issue is that every field point will vignette on the last two optical elements of the converger. Intermediate optics, the achromatic waveplate, and dichroic beamsplitters, will result in additional vignetting. These are temporarily ignored in the design of the pupil imager.

The image plane is defined by a CMOS camera that was available in the lab (Thorlabs, DCC1545M). The camera uses a ½" sensor format that has an active area of 4.8 x 6.4 mm. The back focal distance of the lens is allowed to vary. A biconvex singlet is used as a starting point and optimized. The lenses are split and further optimized. This results is shown in Figure 3.43 (top), but very little correction was be made. Continued optimization and adjustment of the design started to converge towards the design shown in

Figure 3.43 (bottom). However, the image quality was still not optimal. More lenses could be added, but it does not appear that reasonable image quality can be achieved with a simple and cheap imaging system.

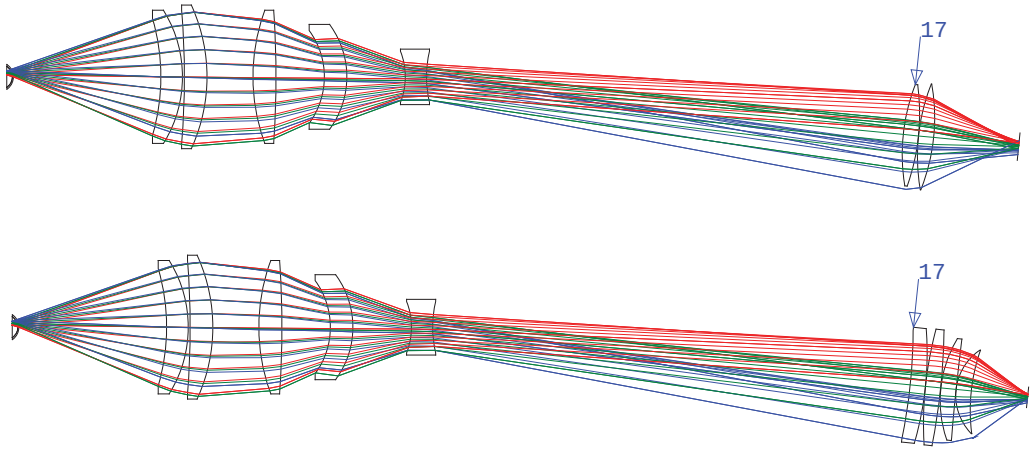

Figure 3.43: Optical layout for a pupil imaging system.
(Top) Initial design. (Bottom) Optimized design.

A cost effective solution for a pupil imaging system does not appear viable. Future converger designs may be able to account for a pupil imager, but the current converger assembly does not appear to support this. Despite this, an attempt was made to image through the converger using available optical components in the lab. The best solution that was found was only able to resolve faint shadows that provided little to no information on the location of the pupil. For this reason, the pupil imaging system was abandoned on SPINALTOP.

3.5 CAMERAS

The cat's eye interferometer provides a supplementary measurement to the surface interferometer measurement that allows for the absolute topography of the cornea to be reconstructed. This requires that the cat's eye measurement is made simultaneously with the surface measurement. This is accomplished by synchronizing the two camera systems. This section will discuss the two camera systems and the method used for synchronization.

3.5.1 SURFACE CAMERA

The 4D pixelated sensor consists of pixelated phase-mask that is aligned to the detector array of a commercial camera system. The digital camera is an Illunis MMV-1010 (Minnetonka, MN) based on the Kodak KAI-1010 interline transfer CCD (Figure 3.44). The KAI-1010 has a resolution of 1008 x 1018 with 9.0 μm square pixels and can operate up to 30 FPS. Analogue-to-Digital Converters (ADC) output 10-bit digital frame data. The MMV-1010 communicates through a Camera Link interface to a BitFlow Neon-CLB framegrabber card (Woburn, MA). One of the framegrabber features includes programmable region of interest (ROI) sub-windowing, camera control signals, and synchronization inputs. The framegrabber also contains an on board function generator that will be used to generate the camera control signals and synchronized to the cat's eye camera.

Figure 3.44: Camera system used by the 4D pixelated sensor.
(Left) MMV-1010. (Right) KAI-1010.

3.5.2 CAT'S EYE CAMERA

The cat's eye camera is an AVT Manta G419B-NIR, a Gigabit Ethernet camera (Figure 3.45). It is based on the CMOSIS CMV4000, a 2048 x 2048 CMOS sensor with 5.4 μm square pixels, 12-bit output, and operates at 28.6 FPS. The camera is NIR enhanced with quantum efficiency greater than 30% at 850 nm (Figure 3.46). The camera has a programmable ROI which is necessary to increase the frame rate to match the surface camera. For SPINALTOP operation, the ROI is decreased to 1974 x 1974 to ensure that

the frame rate can be synchronized to 30 FPS. A 12-pin Hirose power and I/O connect provides a method for external synchronization and triggering to control the camera capture rates.

Figure 3.45: Cat's eye camera, AVT Manta G419B-NIR.

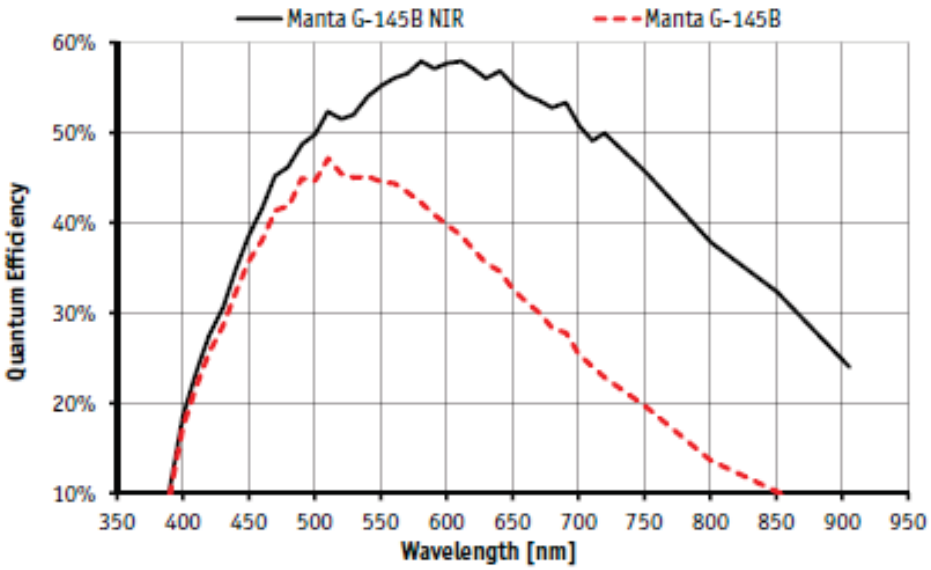

Figure 3.46: Quantum efficiency curve for the cat's eye camera.

3.5.3 SYNCHRONIZATION

Synchronization is accomplished through the BitFlow framegrabber Camera Control Tables (CTABs). The CTABs are a set of programmable waveform generators that can be used to control I/O signals. The waveform will be used to control the surface camera acquisition rates and duration, triggering cat's eye acquisition, and modulating their respective AOMs. The waveforms are binary, consisting of a High or Low output. The waveform is directed to a user selected general purpose input/output (GPIO) port, with

the levels set by a jumper on the framegrabber board. The GPIO levels are set to +5 V, which is required for operation with the cat's eye camera.

The initial waveform was a 30 Hz square wave with a 0.1% duty cycle. This would drive both the surface and cat's eye cameras to operate at 30 frames per second. The 0.1% duty cycle would drive the surface camera's electronic shutter, or integration time, to 33.3 μsec. The cat's eye camera was set to trigger off of the rising edge of the waveform to initialize an acquisition, while software commands could be used to set the integration times. Both AOMs were modulated with the same waveform.

Acquisition worked as expected, both the surface and cat's eye cameras would simultaneously acquire data at 30 FPS. However, the spectral leakage from the dichroic beamsplitter was sufficient enough that both of the laser sources were visible to the other interferometers. Although the leakage is only a percent or two, the two systems produced nearly focused beams on the other systems' camera. The focused energy was comparable to the nominal signals and would disrupt data at the center of the field of view. One advantage of the leakage was that it could be used to aide in the co-alignment of the two systems. Figure 3.47 is an example of the spectral leakage present on a surface measurement. The image on the left shows a red spot saturating the center of the detector, even with the camera gain and integration times at their lowest. The image on the right is with the 850 source disabled and no leakage is present.

The solution was to stagger the exposure times between the two systems. A diagram of the staggered output is shown in Figure 3.48. The CTABs were programmed to generate a second waveform that was triggered off of the first waveform. The first waveform is the initial 30 Hz square wave with a 0.1% duty cycle. The second waveform is duplicate 30 Hz square wave with a 0.1% duty cycle that triggers off of the falling edge of the first waveform. The first waveform is passed to the surface camera to drive the frame rate and acquisition, which is also directed to a GPIO to drive the surface AOM. The second waveform is passed to another GPIO to drive the cat's eye AOM and to trigger the cat's eye camera acquisition. Triggering the second waveform from the falling edge of the first waveform allows the

integration times of the surface camera to be modified without having to change any software or hardware parameters. Figure 3.49 is a capture of the generated waveforms displayed on an oscilloscope. The upper signal triggers the surface camera and the lower signal triggers the cat's eye camera. Surface interferometer integration times generally do not exceed 100 μsec. Therefore, the cat's eye collections are nearly simultaneous with surface collections, but without spectral leakage.

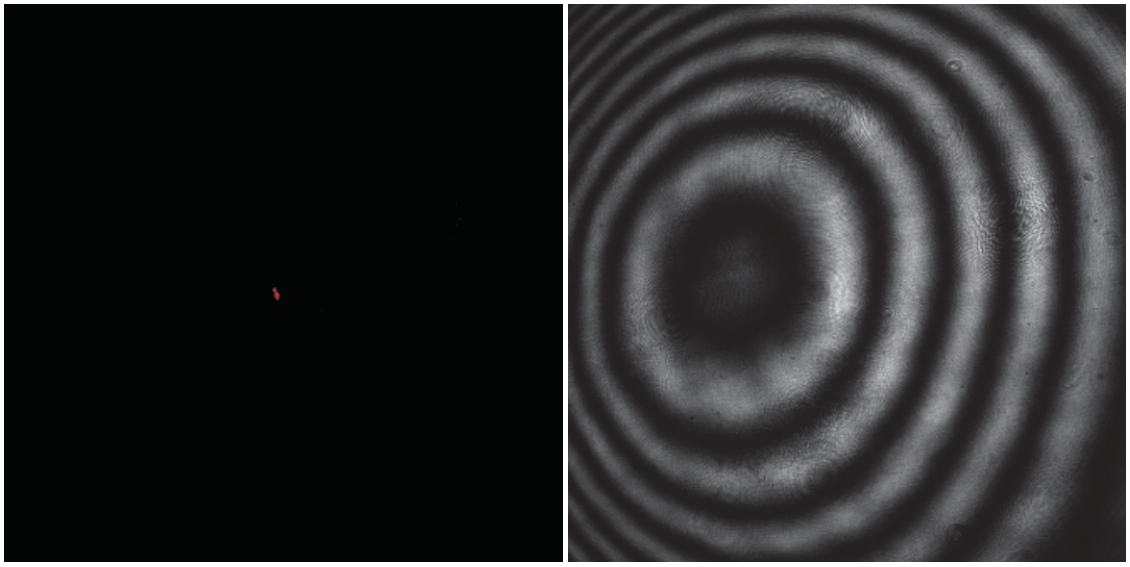

Figure 3.47: Spectral leakage from the cat's eye system into the surface interferometer. (Left) Leakage from the 850 nm cat's eye source is focused at the center of the camera, saturating the display (red pixels). (Right) Nominal operation of the surface interferometer.

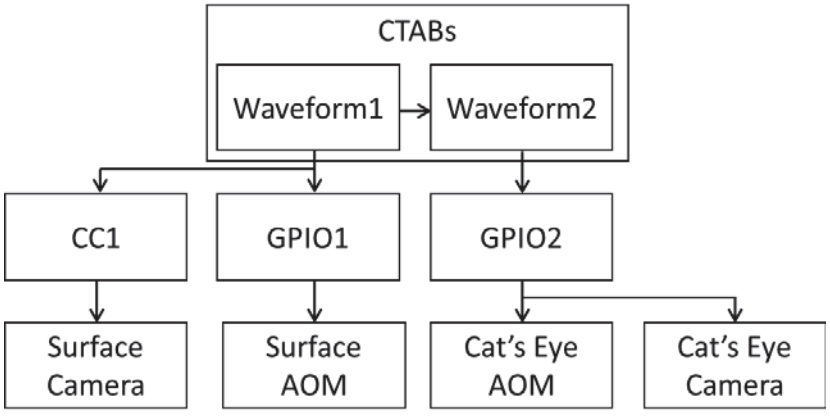

Figure 3.48: Timing diagram for camera synchronization.

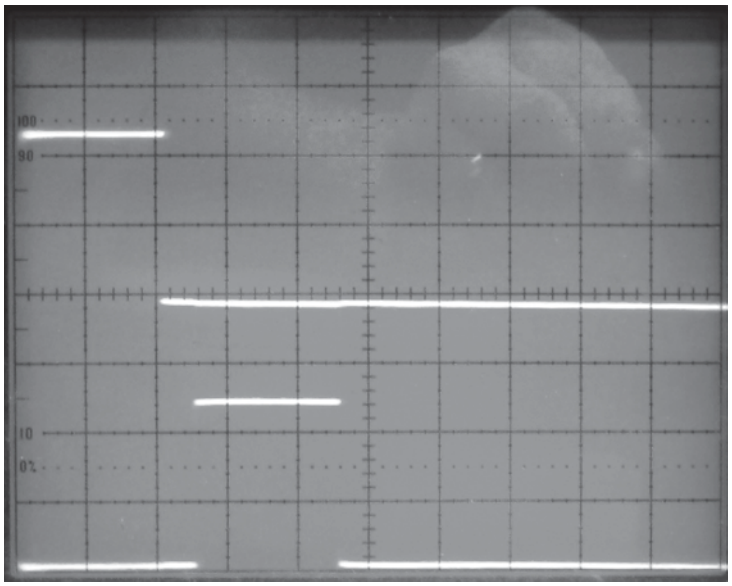

Figure 3.49: Oscilloscope measurement verifying timing offsets.

3.6 LASER SAFETY

The changes made from the TFI system to SPINALTOP can be considered significant with the addition of a secondary laser source. Detailed analysis previously established that the TFI was eye safe (Primeau, Goldstein, & Greivenkamp, 2012; Sliney, 2010). The previous analysis provides the necessary framework for demonstrating eye safety on SPINALTOP. A critical difference between the systems that was not previously addressed was consideration for a multi-wavelength source. The ANSI recommendation for multi-wavelength laser emission is to derive the maximum permissible exposure (MPE) for each wavelength and analyze the exposure limits for the cumulative values (ANSI, 2007). This section repeats prior analysis for the cat's eye interferometer, reviews the surface interferometer, and addresses changes that have been made to the fixation assembly.

3.6.1 CAT'S EYE INTERFEROMETER

The cat's eye interferometer module operates with the same geometry and in the same spectral domain as the surface interferometer (near infrared, NIR, or IR-A). This allows for reapplication of the laser safety analysis of Primeau, et al, and Sliney (Primeau et al., 2012; Sliney, 2010). The following sections mirrors their analysis with changes made to reflect differences in the new system.

3.6.1.1 RETINAL EXPOSURE LIMITS

The maximum permissible exposure (MPE) for retinal exposure is derived as a function of source geometry and wavelength as follows:

$$MPE = 1.8 \cdot C_A \cdot C_E \cdot t^{0.75} \cdot 10^{-3} \left[\frac{\text{J}}{\text{cm}^2}\right] \quad (3.9)$$

where C_A is a spectral correction factor, C_E is an extended source correction, and t is the exposure duration. The spectral correction factor for IR-A lasers is given by:

$$C_A = 10^{0.02(\lambda - 700)} \quad (3.10)$$

where λ is the wavelength in nanometers. The cat's eye interferometer uses an 850 nanometer source, resulting in $C_A = 1.995$. The extended source correction is given by the equation:

$$C_E = \frac{\alpha^2}{\alpha_{max}\alpha_{min}} \quad \text{for } \alpha > \alpha_{max} \quad (3.11)$$

where α is the visual angle in milliradians. The minimum visual angle α_{min} is defined as 1.5 mrad and the maximum α_{max} is 100 mrad. The visual angle is defined by the source geometry:

$$\alpha = 2\tan^{-1}\left(\frac{d_r}{2f_e}\right) \quad (3.12)$$

where d_r is beam diameter on the retina and f_e is the focal length of the eye. The focal length of the eye is defined as $f_e = 17$ mm. The cat's eye interferometer and reshaping optics system generates an f/5 cone with the focus at the tear film, resulting in a beam diameter on the retina of $d_r = 3.36$ mm (Figure 3.50). This results in a visual angle of $\alpha = 200$ mrad, which is considered an extended source. The extended source correction factor is given by $C_E = 266.7$.

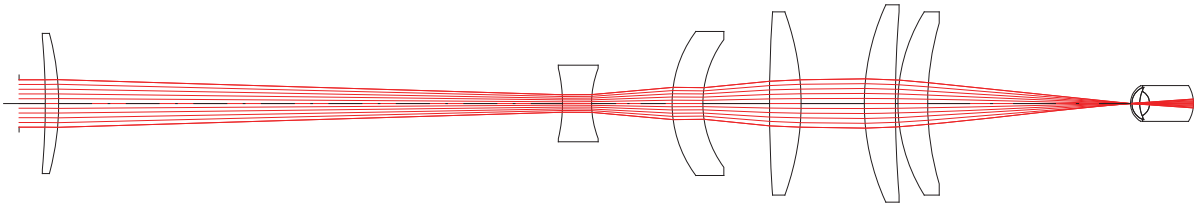

Figure 3.50: Geometry of the cat's eye laser illumination area on the retina.

The final term for MPE is exposure duration t, which is given by $t = 100$ s for all durations greater than 100 s. The MPE for retinal exposure for the cat's eye interferometer is then given by:

$$MPE = 1.8 \cdot 1.995 \cdot 266.7 \cdot 100^{0.75} \cdot 10^{-3} = 30.3 \left[\frac{\text{J}}{\text{cm}^2}\right] = 0.303 \left[\frac{\text{W}}{\text{cm}^2}\right] \text{ for } t = 100\text{s} \quad (3.13)$$

The MPE is an irradiance that is averaged over the 0.7 cm limiting aperture of the pupil as defined by ANSI. An Accessible Emission Limit (AEL) is related to the MPE by:

$$AEL = MPE \cdot AREA_{limiting} \quad (3.14)$$

where the limiting area is calculated as:

$$AREA_{limiting} = \pi \left(\frac{0.7 \text{ cm}}{2}\right)^2 = 0.385 \text{ [cm}^2\text{]} \quad (3.15)$$

Therefore, the AEL for the cat's eye system is:

$$AEL = 0.303 \left[\frac{\text{W}}{\text{cm}^2}\right] \cdot 0.385 \text{ [cm}^2\text{]} = 0.117 \text{ [W]} \quad (3.16)$$

The cat's eye interferometer output power is limited to 1 mW and therefore has a safety factor of 117X. For comparison, the surface interferometer operates with a safety factor of 1350X.

3.6.1.2 RETINAL EXPOSURE LIMITS WITH EYE MOTION, SYSTEM ALIGNMENT, AND OCULAR VARIATION

This section repeats the analysis of Primeau, et al. (2012), where the AEL for retinal exposure is calculated for a range of ocular variations and system alignment. The results for the cat's eye interferometer are shown in Figure 3.51. The worst case ocular variation and alignment configuration for the cat's eye interferometer results in an AEL of 14 mW. The 1 mW operation of the cat's eye interferometer yields a 14X safety margin. For comparison, the worst case AEL including variations for the surface interferometer is 577 mW, or an operating safety margin of 577X. The difference between the cat's eye and surface interferometers has to do with the numerical aperture of the beam coming from the converger; that cat's eye beam operates at a lower NA.

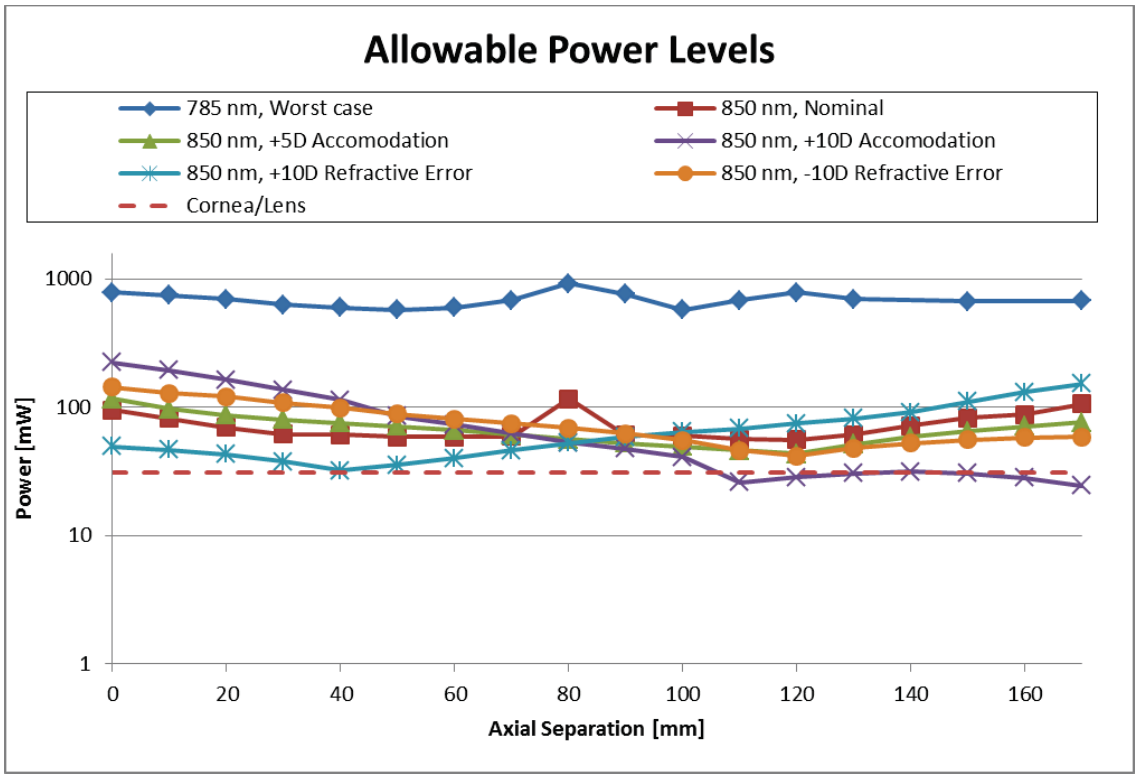

Figure 3.51: Accessible emission limits for retinal exposure, factoring in eye motion, system alignment, and ocular variation.

3.6.1.3 THERMAL LIMITS ON THE RETINA

The upper limit for retinal exposure to large sources as given by the ISO 15004-2, para. 5.4.1.5, expressed as a retinal irradiance limit is:

$$E_{VIS-IR} = 600 \left[\frac{\text{mW}}{\text{cm}^2}\right] \quad (3.17)$$

The illuminated area over the retina for the cat's eye interferometer module can be calculated from the previous illuminated diameter $d_r = 3.36$ mm, resulting in an area of 0.091 cm^2. The limiting power for the cat's eye interferometer is 1 mW resulting in a retinal exposure of 11.0 $\text{mW} \cdot \text{cm}^{-2}$. The thermal limit for the cat's eye interferometer has a 600 / 11 = 54X safety factor. For comparison, the surface interferometer operates with an 845X safety factor.

3.6.1.4 CORNEA AND CRYSTALLINE LENS EXPOSURE LIMITS

As determined in previous analysis, the cornea and crystalline lens have exposure limits of 31.4 mW. The cat's eye and surface interferometers therefore both have a 31X safety margin for cornea and crystalline lens exposure limits.

3.6.1.5 SOURCE MODULATION

The function of the AOM is to modulate, or chop, the laser source. This does not affect the power levels of the beam that are exposed to the eye, but does limit the exposure times. However, laser safety will be determined by the worst case scenario, which is when the AOM is not modulating the source. Instead, the analysis of the previous sections will be revisited to provide a typical safety level.

Source modulation with the AOM is synchronized to the camera shutter rates at 30 FPS, or 33.3 ms. Duty cycles are not expected to exceed 1%, or an active time of 0.3 ms every 33.3 ms. The maximum allowable duration for a single test is 1 hour, or 3600 seconds, per day. Applying the duty cycle of source modulation reduces the total possible exposure time to $t = 36\,ms$, which was previously limited to the

fixed value of $t = 100\ ms$. The reduced exposure time increases the AEL to a value of 18 mW, from the previous 14 mW. Therefore, it is expected that the typical operation of the cat's eye interferometer will operate with at least 18X safety factor.

3.6.1.6 CAT'S EYE LASER SAFETY SUMMARY

The worst case analysis for the cat's eye interferometer shows that the limiting exposure results from ocular variation and system alignment. The value for the limiting exposure is 14 mW. The cat's eye interferometer will operate with power levels less than 1 mW and is therefore eye safe. Typical operation with source modulation should provide an additional 30% increase in safety levels, but is not used to determine the overall laser safety. SPINALTOP laser safety is determined by the addition of both cat's eye and surface interferometers, and therefore the surface interferometer must be revisited before we can declare the full system as eye safe.

3.6.2 SURFACE INTERFEROMETER

The surface interferometer is functionally identical to the TFI, except for the addition of the AOM. Therefore, the AEL for the surface interferometer is the same as the TFI, which was determined by cornea and crystalline lens exposure of 31.4 mW. Cornea and crystalline lens exposure levels are fixed quantities, not dependent on temporal effects; therefore source modulation in the surface interferometer has no effect on the overall laser exposure limits.

3.6.3 FIXATION

The fixation module uses one white LED that diffusely illuminates a fixation target. Modifications to the TFI will place the module internal to the system and co-aligned with the two interferometers modules. The original analysis determined that the white LED operates far below any exposure level and requires no further assessment. The internal placement of the module does not change operating parameters of the fixation target and illumination and therefore does not require any further assessment. Furthermore, the addition of optics in line with the fixation will further reduce illumination levels. The dichroic

beamsplitters will cause a loss of 72% and the converger assembly has 10 optical surfaces that have a laser-line V-coat at 785 nm and will cause additional loss.

3.6.4 LASER EXPOSURE SUMMARY

The limiting exposure for the 785 nm surface interferometer was previously determined to be 31.4 mW from the cornea and crystalline lens. The same analysis performed above for the cat's eye interferometer determined that the limiting exposure is 14 mW from retinal exposure in a worst case analysis of ocular variation and system alignment. Independently, the surface and cat's eye interferometer systems will operate with 31X and 14X safety margins, respectively. Worst case analysis considers the case where SPINALTOP operation simultaneously illuminates the cornea with both laser systems, therefore the cumulative safety factor is:

$$\left(\frac{1}{31}+\frac{1}{14}\right)^{-1} = 10 \qquad (3.18)$$

SPINALTOP will operate with a safety factor of at least 10X and is therefore considered eye safe.

Typical operation of SPINALTOP will modulate both the surface and cat's eye laser sources. The surface interferometer exposure does not change with modulation, however the cat's eye interferometer safety exposure limit will increase to an 18X safety level. Additionally, the two lasers are staggered in their operation and will not be exposed to the eye at the same time (Section 3.5.3). Therefore, typical operation is limited by the cat's eye system at an 18X safety level.

3.6.5 INTERLOCK

Nominal operation of SPINALTOP is expected to be eye safe; however, the use of an interlock system ensures that system operation remains eye safe. The interlock system consists of an internal power monitor that will trip if the laser output levels of SPINALTOP exceed a threshold, which is set to 1 mW for each laser source as discussed in the previous section. The internal power monitor consists of a silicon

photodetector (Newport 918D-SL-OD2R) that is designed to operate in the 400 - 1100 nm range. The photodetectors are placed directly after the pick-off mirrors in the surface interferometer (Section 3.2.1) and cat's eye interferometer (Section 3.3.1). The photodetectors are biased and read out by an optical power meter (Newport 2936-R).

The detector output voltages are calibrated to the output of the SPINALTOP system, which effectively monitors system output levels to ensure that it operates at eye safe levels at all times. The calibration process determines the mapping between the output power of the system to the internal laser power measurement. The output of the system is measured with a third photodetector that is kept external to the system. The same type of photodetector and power meter that are used for internal monitoring are used for external measurements. The external photodetector and power meter are calibrated to NIST traceable standards. The calibration process is performed on each sub-system independently. A range of power levels are evaluated by adjusting the CVND (Sections 3.2.1 and 3.3.1). The internal and external measurements are simultaneously measured and recorded. An example set of data is shown in Figure 3.52. The output power measurements (in milliwatts) at the nominal eye location are plotted against the internal measurement levels (in millivolts). A linear mapping is determined from the measurement set that allows the corresponding output power to be determined from an internal measurement without requiring the external power meter. The internal threshold can be determined from the calibration curve for the interlock system. The threshold for the surface and cat's eye interferometer are both set to 1 mW as discussed in the previous section. For the set of data shown in Figure 3.52, the calibration mapping determined that the internal threshold would be set to 58.5 mV. In other words, if the internal power monitor measured a signal greater than 58.5 mV, this would mean that the laser output of the system is greater than 1 mW, which would trip the interlock system. Typically the interlock threshold is set to a 0.9 mW and the sub-system is set to operate somewhere between 0.6 – 0.8 mW.

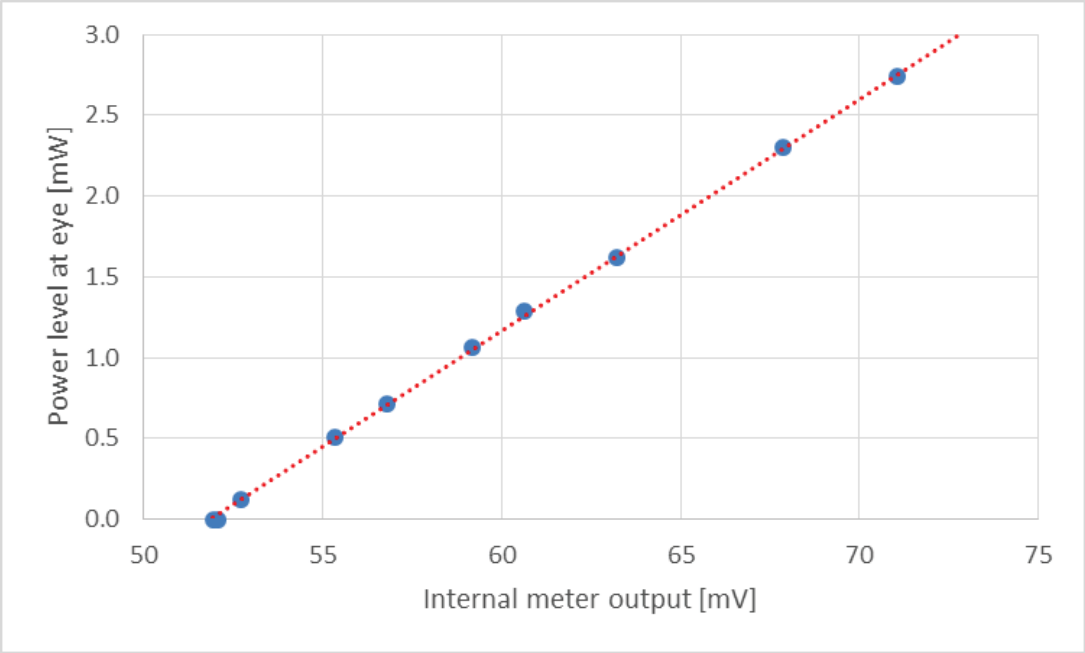

Figure 3.52: Calibration mapping for internal laser power meters.

The electronic interlock system is shown in Figure 3.53, which has been significantly modified from the TFI interlock system. The two major drivers for the change in the interlock system are the use of source modulation and the addition of a second laser source. The modulated source is exposed to the eye for approximately 100 μs. In the TFI implementation of the interlock, the digital acquisition switch that was used to monitor laser power could not detect pulses of that duration. Two changes were made in SPINALTOP to account for source modulation. The first was to use the D-mirror as a pick-off method (Sections 3.2.1 and 3.3.1). By passing the non-deflected beam to the laser monitor, the monitoring system was measuring a signal that was present 99.9% of the time. The second change was to use a set of high-speed comparators (Maxim, MAX942) to monitor laser power levels. The comparators have an 80 ns response time that is well below any modulation that would be used in the system and fast enough to detect any potential spikes in power.

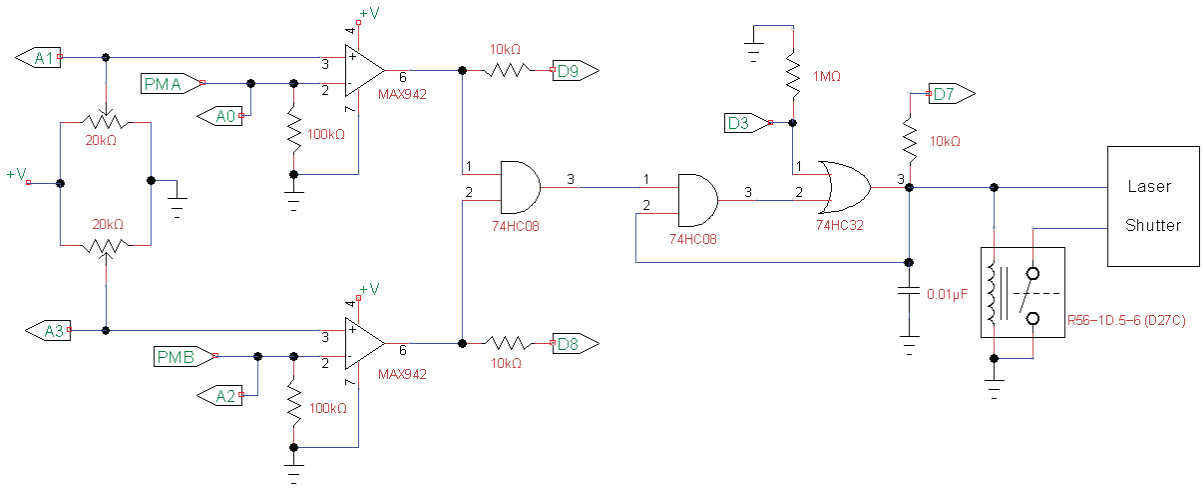

Figure 3.53: Interlock schematic for SPINALTOP.

The outputs of the internal photodetectors (PMA and PMB, Figure 3.53) are compared to a threshold voltage with the high-speed comparators. The threshold voltage is set by a potentiometer against a reference voltage source. The threshold voltage is determined by the calibration process as described at the beginning of this section. While the set-point voltage is greater than the photodetector voltage (i.e. laser output levels are less than the threshold), the comparator will maintain a high output (+5V). The two outputs from the set of comparators are fed into a logic AND gate. The output of the AND gate is only high if both comparator outputs are high (i.e. both lasers are operating below the threshold limits). By itself, this is not sufficient. A power spike would temporarily disable the output of the AND gate and then immediately restore itself. This is accounted for by passing the output of the initial AND gate into a second AND gate that has its output in a feedback loop to itself. That way, when a power spike occurs, the initial AND gate drops followed by the second AND gate. The feedback from the second AND gate causes it to now remain low, regardless of the primary AND gate. This mechanism provides the tripping action.

The AND gate will require an external signal to re-engage the laser interlock system (i.e. a reset). A logic OR gate follows the second AND gate to provide the reset or an override. It can force the feedback line to high and allow the system to resume operation if the primary AND gate is at high. The final output, which

is common with the feedback line, is passed to a reed relay (NTE Electronics Inc., R56-1D.5-6). The reed relay is used to open and close the interlock circuits on the laser and shutter circuits. When the laser systems are operating within proper eye safe levels, the system can be reset, and the reed relay will close allowing the shutter to be re-opened and the system to resume lasing.

A number of inputs and outputs were added to the interlock system to allow for remote monitoring and control of the interlock. Both analogue and digital outputs (A0-3, D2-9 in the schematic, Figure 3.53) are passed to an Arduino (Arduino Uno). The Arduino is an open source microcontroller board based on the ATmega328 (Atmel) (Figure 3.54). The Arduino can be programmed to monitor voltage levels and power status which is then forwarded through a serial interface (RS-232), which can in turn can be read out with the SPINALTOP software interface. Additional status LEDs are placed on the outside of the interlock for immediate feedback (Figure 3.55).

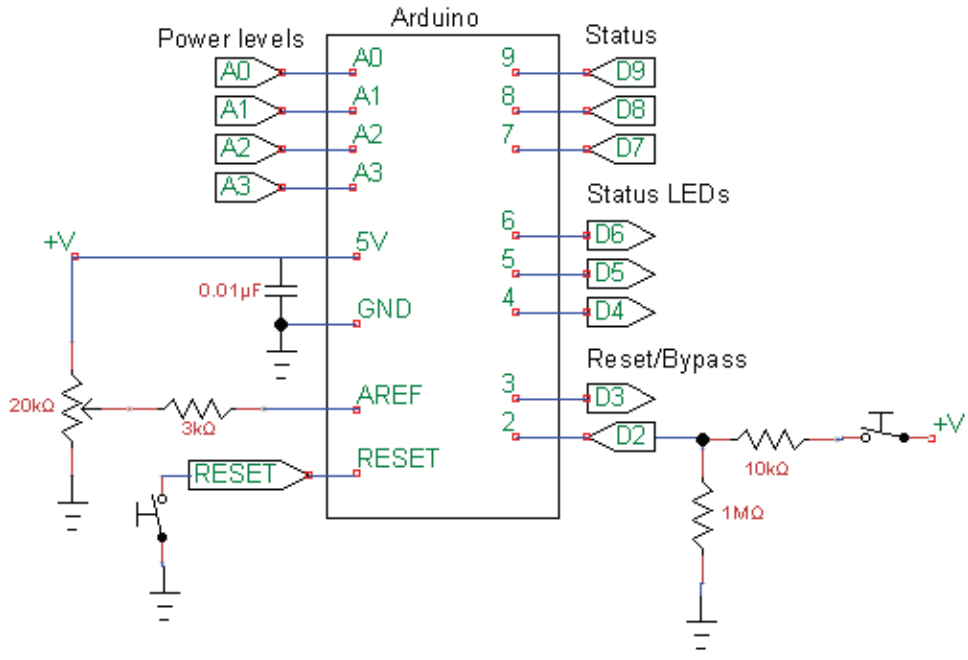

Figure 3.54: Schematic for Arduino interface to interlock.

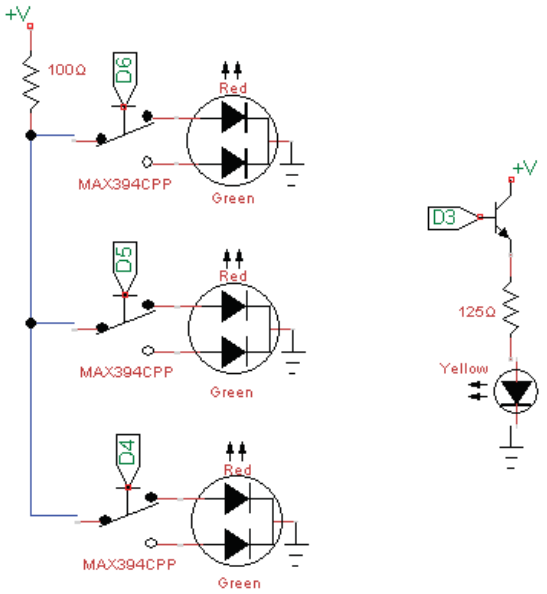

Figure 3.55: Schematic for interlock status LEDs.

3.7 SOFTWARE INTERFACE

The SPINALTOP system is controlled and operated through a custom graphic user interface (GUI) that has been developed in C# using the .NET Windows Forms library classes. The interface is shown in Figure 3.56. The software is responsible for controlling every critical component within the SPINALTOP system, which includes the surface camera, cat's eye camera, and three translation stages. The software can also interface with interlock system (Section 3.6.5), alignment system (Section 3.2.5), and fixation (Section 3.4). One of the major advantages of having the custom software is the ability to perform simultaneous data capture and display. The TFI system required the user to align the system to the subject and acquire a burst collect. During the burst, the live display would freeze until data collection had completed. During this collection time, it was unknown if the subject had drifted out of alignment or if the alignment was even optimal (i.e. not at the center of saccadic motion). The current software is able to perform a continuous burst (i.e. unlimited data capture) and provide a live display of the fringes for feedback which provides the operator with flexibility to terminate the burst or the ability to perform fine corrections of the system to the subject without interrupting data collection.

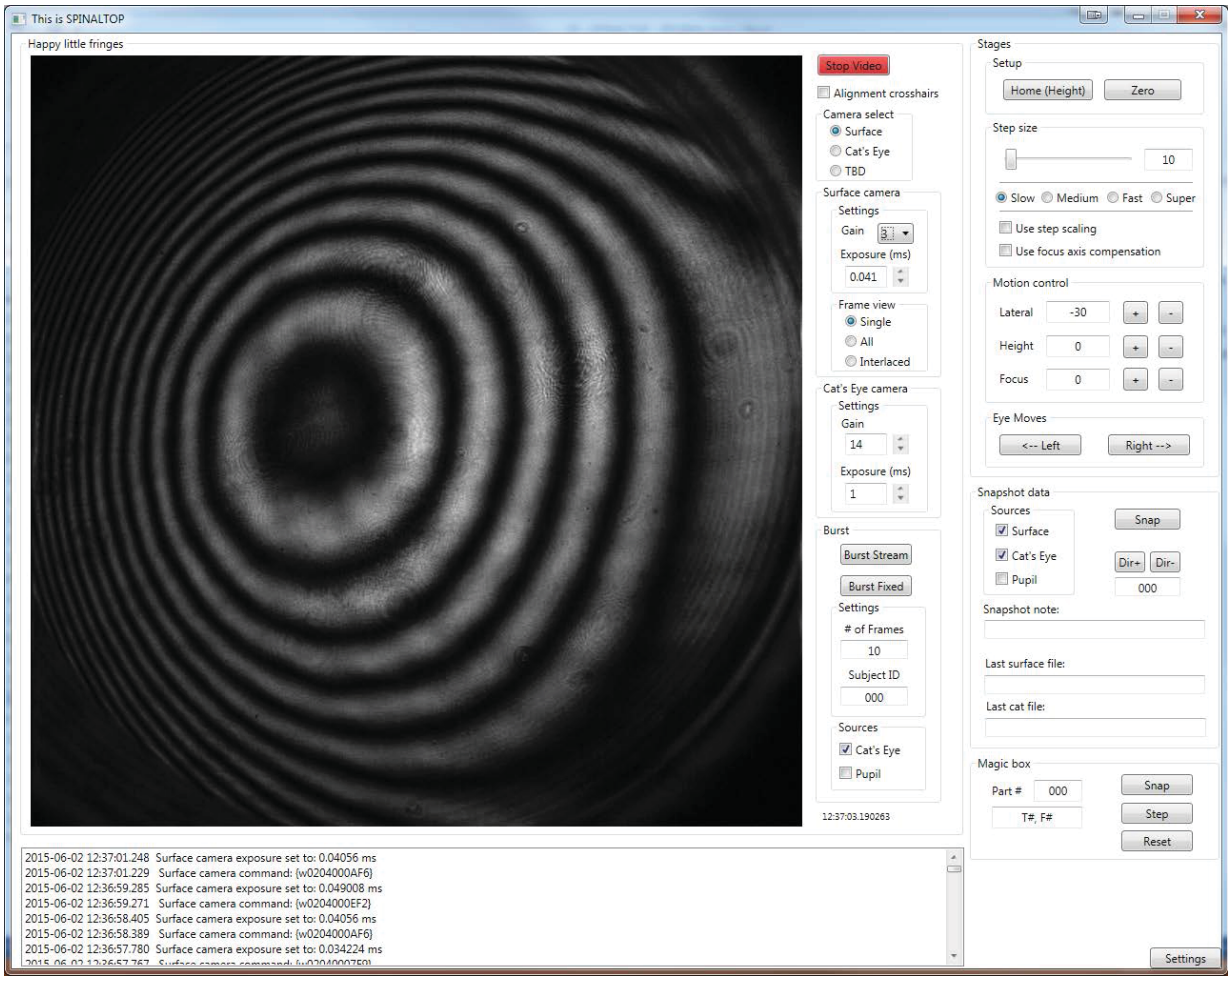

Figure 3.56: SPINALTOP software interface.

4 ALGORITHMS AND ANALYSIS

Reconstruction of the corneal topography from the simultaneous surface and cat's eye measurements is not a trivial process. The process to recover the topography and tear film structure that is discussed in this chapter is outlined in Figure 4.1. The development of the algorithms began with the TFI (Section 2.6). Large eye motion, ocular variation, and tear film dynamics generate complex, high-density fringe measurements that exceeded the capabilities of commercial software to recover the phase. Custom phase unwrapping software had to be developed to recover any surface information (Section 4.2). The same issues that were present with the TFI for phase unwrapping are also present in SPINALTOP.

What sets SPINALTOP apart from the TFI is the ability to reconstruct the absolute surface shape of the cornea and overlying tear film. Although interferometric characterization of optical surfaces can yield sub-nanometer height resolution (Bruning et al., 1974), the adaption of an interferometer to test surfaces other than planar or spherical shapes, such as the cornea, presents significant difficulties. Interferometric testing can be broken down into two domains: Null and non-null testing. A null test measures an optical component against a reference and the measured optical path differences correspond to surface figure errors. Conic surfaces with an axis of revolution can be tested in a null configuration using their conjugate foci, but are limited to a very small subset of shapes (D. Malacara, Creath, Schmit, & Wyant, 2007). More complex surfaces can be tested in a null configuration by use of compensating optics, but can be prohibitively expensive due to manufacturing and verification. Computer generated holograms can be used to reduce the complexity of the compensating optics, but are still costly due to manufacturing and verification (Burge, 1995). The cost of compensating optics may be amortized for a mass produced optic, but its use will be limited to a specific surface shape and configuration. For large or expensive components, the cost of the compensating optics must be included in the project budget.

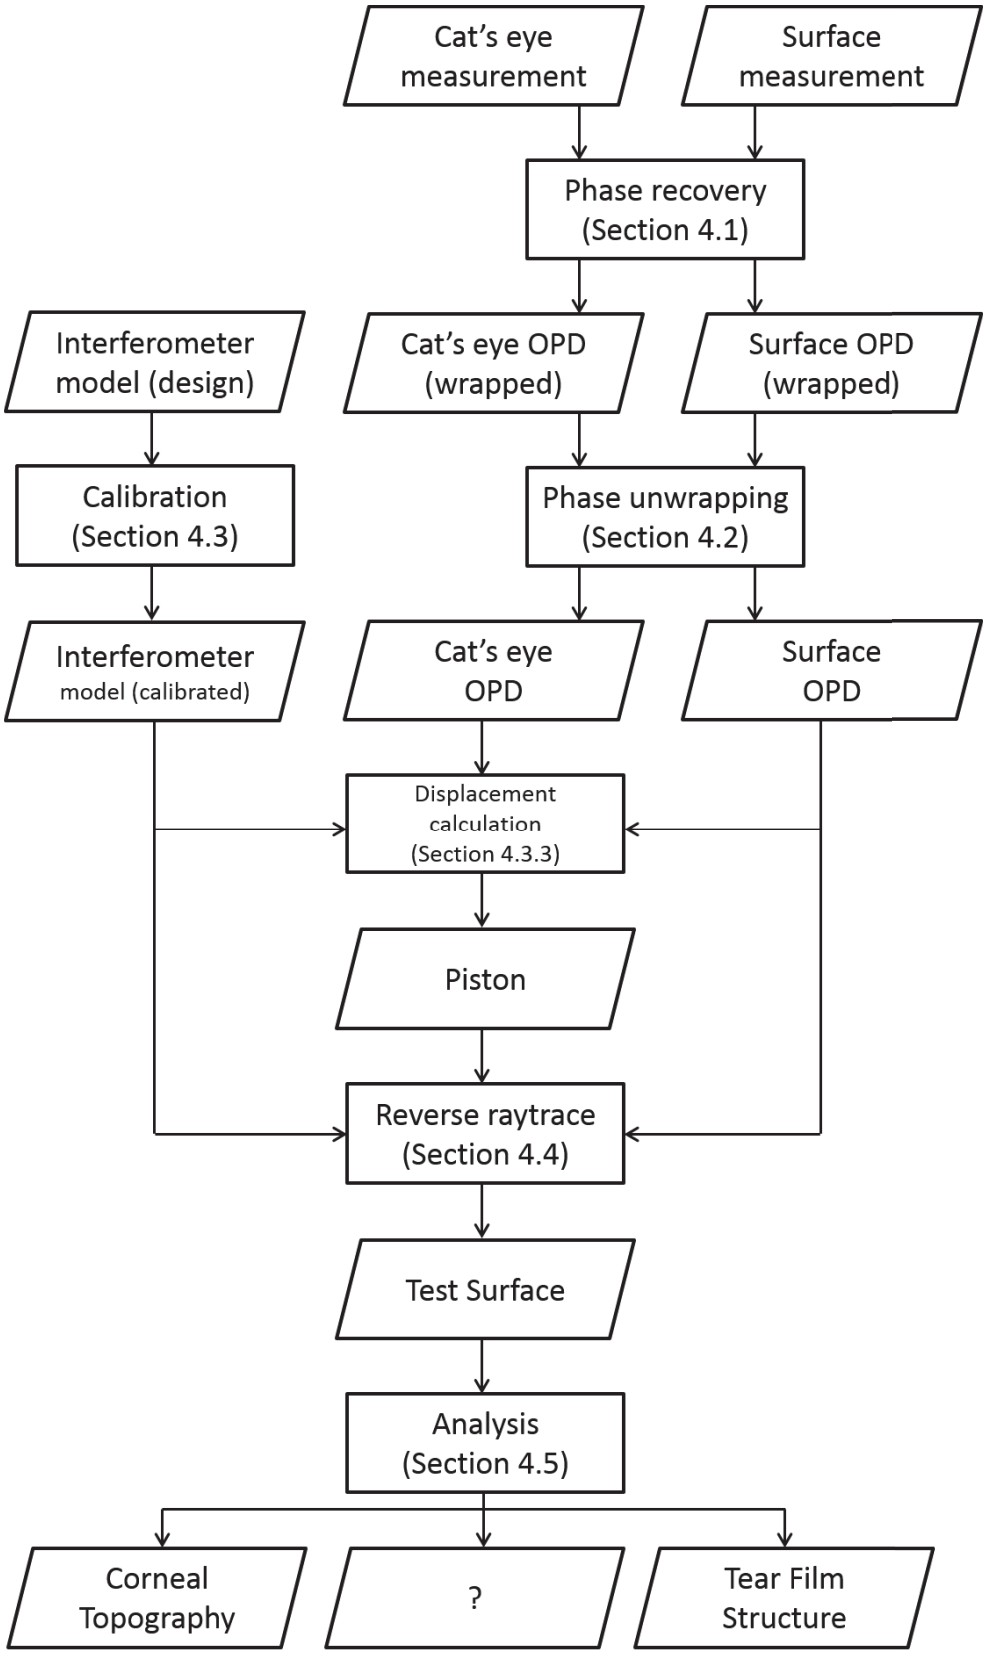

Figure 4.1: Outline of operations to recover surface topography.

However, the range of ocular surfaces shapes, eye motion, and tear film dynamics, does not allow for the TFI or SPINALTOP to test in a null configuration. These systems test in what is referred to as a non-null configuration. A non-null configuration exists when the test surface shape deviates from the reference surface or wavefront by a significant amount, generating a large number of fringes at the observation plane, even at an optimum alignment. One issue with a non-null test configuration is dynamic range – fringes must be resolvable to measure the test surface. The standard sampling requirement is to limit fringe frequency in the detector plane to less than two pixels per fringe (Gaskill, 1978). Large format detector arrays are the easiest way to extend the dynamic range of an interferometer system. Sub-Nyquist interferometry can extend the range of sensors beyond the standard sampling limits, but requires a specialized detector to undersample the fringes (Gappinger & Greivenkamp, 2004a; Greivenkamp, 1987). Shearing interferometry extends dynamic range by measuring surface slopes instead of surface height (Honda, Huang, Tsujiuchi, & Wyant, 1987; Wyant, 1987). Long-wavelength interferometry is another method to extend the dynamic range at the cost of height resolution (Kwon, Wyant, & Hayslett, 1980). Multiple-wavelength interferometry simulates long-wavelength interferometry in the visible domain, overcoming material, source, and detector limitations as compared to long-wavelength interferometry, but will have dispersion issues and still results in reduced height resolution (Cheng & Wyant, 1984; Creath, Cheng, & Wyant, 1985; Murphy, Brown, & Moore, 1999). Dynamic range limitations can also be avoided by the use of scanning or sub-aperture stitching techniques (Küchel, 2006; Murphy et al., 2006). All of these techniques deal with the problem of capturing the interferometric fringe data.

The larger issue with non-null interferometry is system induced error, where measured optical path differences no longer directly correspond to surface figure errors. These errors are more commonly referred to as retrace errors, originating from the geometric perspective that rays generated by the test surface no longer trace the same path as null rays through the interferometer optics (Selberg, 1991). Considering that the observed wavefront at the detector plane is conjugate to the test surface by means of the interferometer optics, retrace errors can be described in terms of imaging aberrations (Huang, 1993;

Jóźwicki, 1991; Murphy, Brown, & Moore, 2000a). Examples in literature have shown that these errors can exceed 10% of the total surface departure (Geary, Yoo, & Si, 1992; Greivenkamp & Gappinger, 2004; Lowman & Greivenkamp, 1994a). For the range of expected ocular variation (Section 1.2), surface departures will exceed 100 μm; not including eye motion or tear film variation. The amount of retrace error introduced by the non-null configuration of SPINALTOP will be significant if not accounted for.

There are a number of methods to reduce or eliminate retrace error, but they are generally impractical or limited in their capability. Retrace error can be minimized through the design of the imaging system, but is impractical for a generic test and potentially more complicated than the design of a null compensator (Z Malacara & Malacara, 1995). Characterization of the interferometer in terms of its imaging aberrations lends itself to an inverse problem, but will be limited in resolution by the polynomial representation of the wavefront error and is unable to allow for full characterization of interferometer errors (Huang, 1993; Murphy et al., 2000a; Murphy, Brown, & Moore, 2000b). Numerically derived mappings can simplify the process, but have not yet been shown to provide adequate compensation for retrace errors (Kurita, Saito, & Kato, 1987). Empirically mapped interferometer errors allow for quick corrections, but are limited to low spatial frequency features that deviate from a null configuration (C J Evans & Bryan, 1993; Chris J Evans, 1994).

The usual method for eliminating retrace errors is by reverse ray tracing and it is often considered to be the best method. Reverse ray tracing uses a model of the interferometer and the measured optical path differences to simulate the reverse propagation of the test wavefront from the detector to the test surface. The reverse ray trace operation effectively backs out the interferometer induced errors and is limited only to the accuracy of the interferometer model. One of the earliest documented uses of reverse ray tracing on a surface tested in reflection was inconclusive, as the test surface that was used in the experimental setup was not independently verified (Honda et al., 1987). Even if the test surface had been well characterized, it is suspected that the results would have still contained errors due to interferometer model inaccuracies and an incomplete description of the mapping of optical path differences to surface figure error. Kurita

later addresses these issues and successfully demonstrates a reverse ray trace correction on aspheres, but the method is only valid for a rotationally symmetric asphere tested with a spherical reference wavefront (Kurita, 1989). This method is also limited by not providing a way to account for interferometer hardware errors. It is also suspected that reverse ray tracing is used in commercial interferometers used for testing aspheric surfaces. Reverse ray tracing in a transmission test has been successfully demonstrated (Williby et al., 2003). Surface reconstruction of optical elements tested in transmission have also been discussed, but require *a priori* information of the optical elements (Seong & Greivenkamp, 2008).

The primary limitation of reverse ray tracing is that the achieved measurement accuracy is limited by the accuracy to which the interferometer model is known. Reverse optimization is a technique used to determine the configuration of an optical system by acquiring multiple measurements of the system on a known surface and adjusting parameters in the model to minimize differences between real and simulated measurements. Known changes to the system are made for each measurement. An improved interferometer model results (Gappinger & Greivenkamp, 2004a; Hwan J Jeong & Lawrence, 1988; Powell, 2000). The method was extended to solve for an unknown test surface by using it in place of the known surface and including it as a variable in the reverse optimization process (Gappinger & Greivenkamp, 2004a, 2004b; Greivenkamp & Gappinger, 2004; Lowman & Greivenkamp, 1985). The issue with using reverse optimization to solve for the test surface is that the process is intensive, requiring multiple iterations of ray traces to converge on a solution for the test surface. Spatial resolution of the reconstructed test surface is limited by the number of rays that are traced and is proportional to the amount of time required to solve for the surface. As a result, reverse optimization is better suited and better thought of as a system calibration step. The calibrated interferometer model can then be used to solve for the unknown test surface by reverse ray tracing methods.

What is discussed in this chapter is the full gamut of processing required to reconstruct the absolute shape of the corneal surface. The first section discusses standard phase recovery techniques for the two sub-systems. The recovered phase will need to be unwrapped, which is discussed in the following section.

These two sections are common to both the TFI and SPINALTOP. However, the remaining sections will focus on algorithms that are unique to SPINALTOP. The third section focuses on the interferometer calibration by means of reverse optimization. Reverse ray tracing can be applied to the calibrated interferometer model to recover the absolute surface topography of the cornea. There is limited information that exists in literature that discusses the problem of reverse ray tracing as it applies to this system. Therefore, an entire section is devoted to reverse ray tracing and developing the methods that are necessary to reconstruct the surface. A final section is included that discusses how to handle the data that comes out of the reconstruction process.

4.1 PHASE RECOVERY

The surface and cat's eye interferometers employ two different techniques to measure the phase from a single snapshot of data. The surface interferometer uses a spatially phase-shifted method by means of a micropolarizer array (Section 3.2.5) and the cat's eye interferometer uses a spatial carrier method by means of the Wollaston prism (Section 3.3.5). Recovering the phase from a measurement is a required first step before any other processing can occur. This section discusses the methods used for each system.

4.1.1 SPATIAL PHASE-SHIFTED METHOD

The sensor that is used in the surface interferometer is a 4D pixelated phase camera (4D Technology Corporation, Tucson, AZ). A pixelated phase-mask aligned to the detector array acquires four phase shifted interferograms in a single camera frame (Millerd et al., 2004). The phase mask consists of a replicated array of linear polarizers oriented in four different orientations in a 2 x 2 super-pixel (Figure 4.2). Kothiyal and Delisle (1985) showed that by having the test and reference in orthogonal circular polarization states, the irradiance of the interfering beams directly following a linear polarizer is given by:

$$I(x,y) = \frac{1}{2}\left(I_R + I_S + 2\sqrt{I_R I_S}\cos(\phi(x,y) + 2\alpha)\right) \qquad (4.1)$$

Where I_R and I_T correspond to the reference and test irradiances, respectively. The phase is given by ϕ and the orientation of the linear polarizer is given by α. The four polarizers in a super-pixel are rotated in increments of 45º to produce four-unique phase shifted irradiances that are measured by the detector. The replicated array of super-pixels produces a set of four unique phase-shifted interferograms that are interlaced in the final detector output. This is the spatial-phase shifted method. The raw output from a measurement made with this sensor is shown in Figure 4.3. A de-interlaced image can be extracted (e.g. only pixels with the 0º linear polarizer) and the result is shown in Figure 4.4.

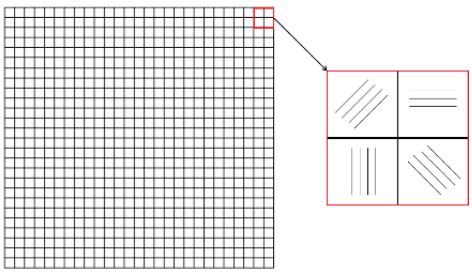

Figure 4.2: Super pixel on the micropolarizer grid array.

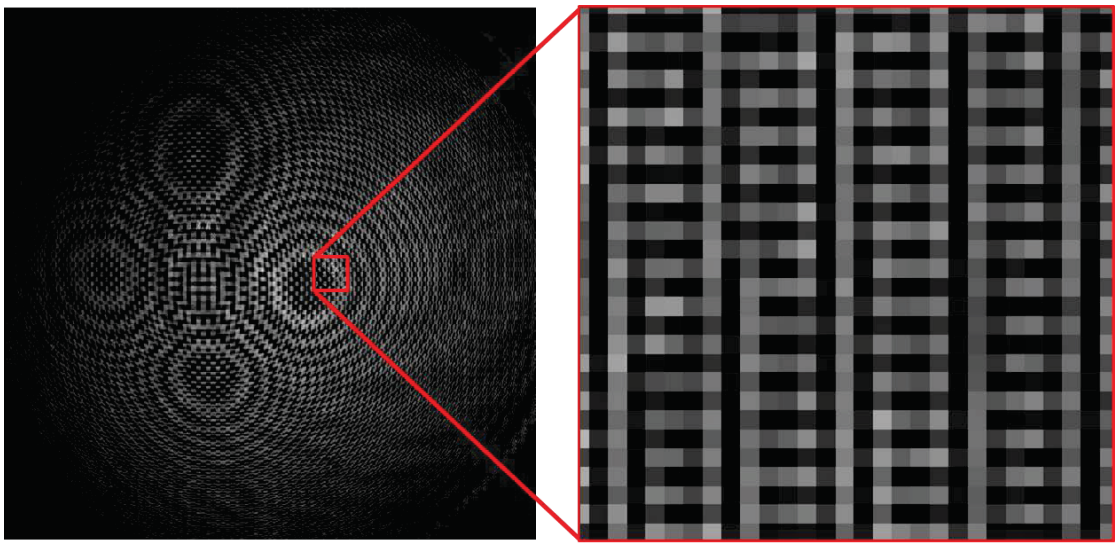

Figure 4.3: Raw spatial phase-shifted image containing the four interlaced phase-shifted interferograms.

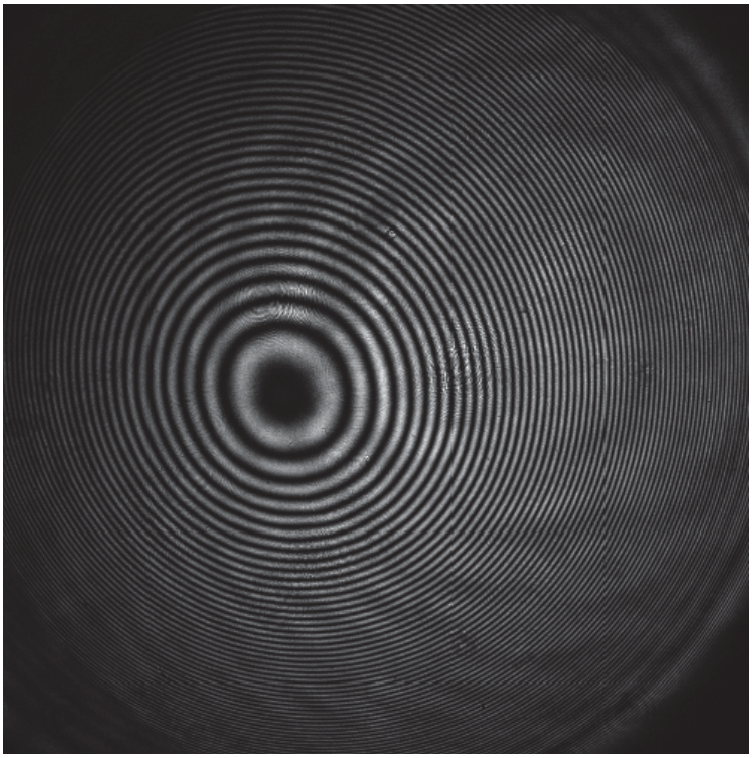

Figure 4.4: A de-interlaced interferogram.

Four individual frames can be de-interlaced from the image and the phase can be recovered by conventional N-bucket algorithms (Zacarias Malacara & Servín, 2010). However, spatial convolution methods are more simple and effective for recovering the phase (Millerd et al., 2004). The spatial convolution method that is used in this dissertation to recover phase from the pixelated phase mask is a circular carrier 9-step method (Kimbrough, 2006). The phase is recovered from a 3 x 3 neighborhood of pixels and has the advantage that the recovered phase value is co-incident with the center pixel (Figure 4.5). The left image in Figure 4.5 shows the relative phase shift between the individual irradiance values and the indexing convention is shown on the right. The recovered phase ψ from the 9-step method is given by:

$$\psi_5 = \text{atan2}(2[I_2 + I_8 - I_4 - I_6], -I_1 - I_3 + 4I_5 - I_7 - I_9) \qquad (4.2)$$

The atan2 function is a commonly used function in computer programming to recover the quadrant of the computed angle, which cannot be recovered from the standard arctangent operation. In other words, the

result of the atan2 function is an angle that is in the range of $(-\pi, \pi]$, whereas the arctangent results in an angle in the range of $\left(-\frac{\pi}{2}, \frac{\pi}{2}\right]$. The atan2 function can be expressed as:

$$\operatorname{atan2}(y,x) = \begin{cases} \operatorname{atan}\left(\frac{y}{x}\right), & x > 0 \\ \operatorname{atan}\left(\frac{y}{x}\right) + \pi, & y \geq 0, x < 0 \\ \operatorname{atan}\left(\frac{y}{x}\right) - \pi, & y < 0, x < 0 \\ +\frac{\pi}{2}, & y > 0, x = 0 \\ -\frac{\pi}{2}, & y < 0, x = 0 \\ \text{undefined}, & y = 0, x = 0 \end{cases} \qquad (4.3)$$

Phase can be recovered for the entire array of data with this method, with exception of the boundaries where phase information will be missing. Although there are alternative methods to extract the phase at the boundaries, it is easier to ignore the boundary data. The captured image is a 1008 x 1008 array and the recovered phase array is 1006 x 1006.

π	$-\frac{\pi}{2}$	π
$\frac{\pi}{2}$	0	$\frac{\pi}{2}$
π	$-\frac{\pi}{2}$	π

1	2	3
4	5	6
7	8	9

Figure 4.5: 3 x 3 neighborhood of pixels.
(Left) Phase shift value. (Right) Pixel index. Recovered phase will be coincident with pixel 5.

A consequence of measuring irradiance (Eq. 4.1) is that the recovered phase, through the use of Eq. 4.2, is a 2π modulo operation of the true phase (ϕ):

$$\psi = \phi - 2\pi \left\lfloor \frac{\phi}{2\pi} \right\rfloor \qquad (4.4)$$

where the ⌊ ⌋ is a rounding operation. The recovered phase is generally referred to as a wrapped phase and the process of recovering the true phase from the wrapped phase is known as phase unwrapping. The details of phase unwrapping will be discussed in Section 4.2. The wrapped phase that has been recovered from the example set of data (Figure 4.3) is shown in Figure 4.6.

The DC irradiance (I_{DC}) can be recovered from the interlaced data using the 3 x 3 pixel neighborhood with a weighted average, where the weighting is determined by the prevalence of a phase term:

$$I_{DC} = \frac{1}{16}(4I_5 + 2(I_2 + I_4 + I_6 + I_8) + I_1 + I_3 + I_7 + I_9) \qquad (4.5)$$

The modulation (γ) can be recovered by the same method:

$$\gamma = \frac{\sqrt{2(I_2 + I_8 - I_4 - I_6)^2 + (-I_1 - I_3 + 4I_5 - I_7 - I_9)^2}}{8 \cdot I_{DC}} \qquad (4.6)$$

Modulation and the DC irradiance will be used to support phase unwrapping. Figure 4.7 shows the recovered DC irradiance (left) and modulation (right) for the example set of data that has been used in this section.

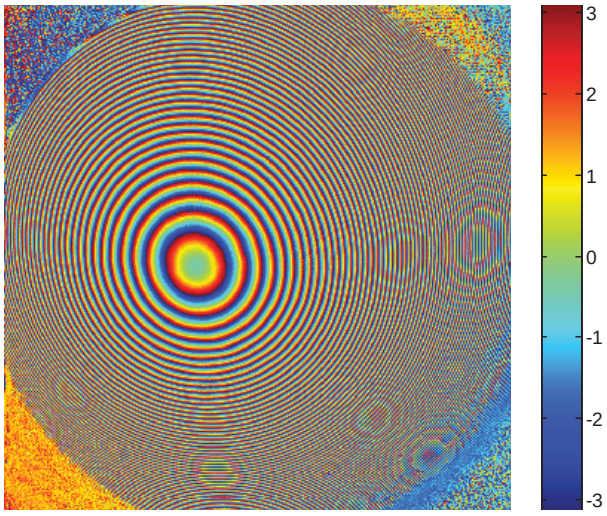

Figure 4.6: Recovered wrapped phase.
Note: Aliasing that appears in the image is a result of the compressed image display and is not in the actual measured data.

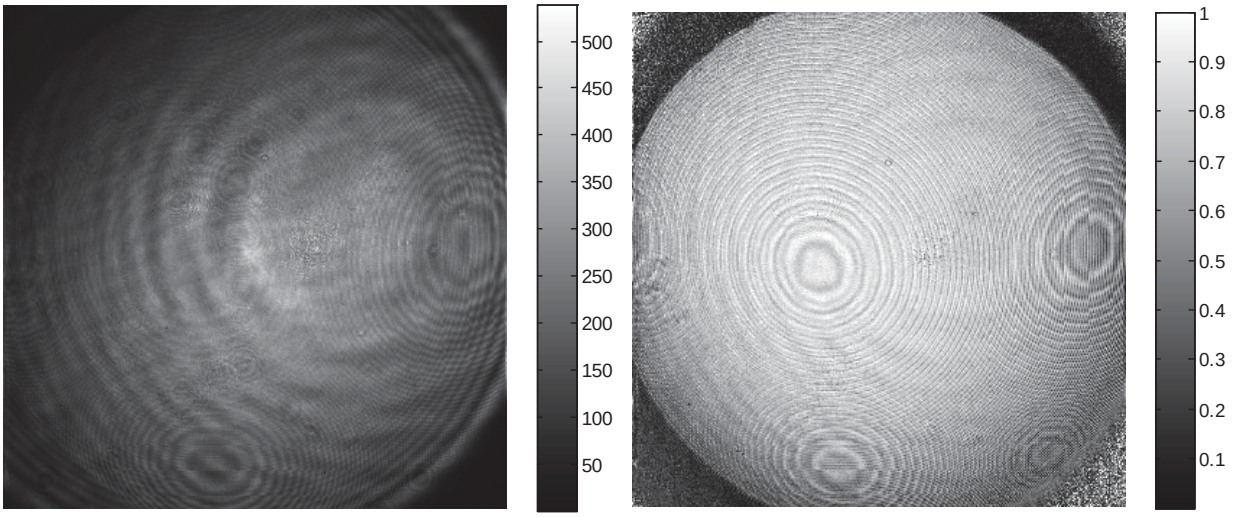

Figure 4.7: Recovered information for the spatial phase-shifted method.
(Left) DC Irradiance. (Right) Modulation.

4.1.2 SPATIAL CARRIER METHOD

The cat's eye interferometer captures phase in a single snapshot through a spatial carrier (Takeda, Ina, & Kobayashi, 1982; Takeda, 1990). The spatial carrier method is accomplished by introducing a large tilt term into the interferogram. The cat's eye system uses a Wollaston prism to offset the test and reference wavefronts to introduce the tilt (Section 3.3.5). The measured irradiance $I(x,y)$ for an interference pattern can be described by:

$$I(x,y) = I_1(x,y) + I_2(x,y)\cos\big(\phi(x,y)\big) \qquad (4.7)$$

where I_1 and I_2 relate to the test and reference beam irradiances, and ϕ represents the phase. The relationship between I_1 and I_2 and the test and reference irradiances can be derived from Eq. 4.1. Tilt along the x-axis introduces a linear phase ramp that results in the measured irradiance:

$$I(x,y) = I_1(x,y) + I_2(x,y)\cos\big(2\pi f_0 x + \phi(x,y)\big) \qquad (4.8)$$

The spatial carrier can be described in terms of a frequency f_0, which is directly proportional to the amount of tilt introduced along the x-axis. An example interferogram with a spatial carrier is shown in Figure 4.8.

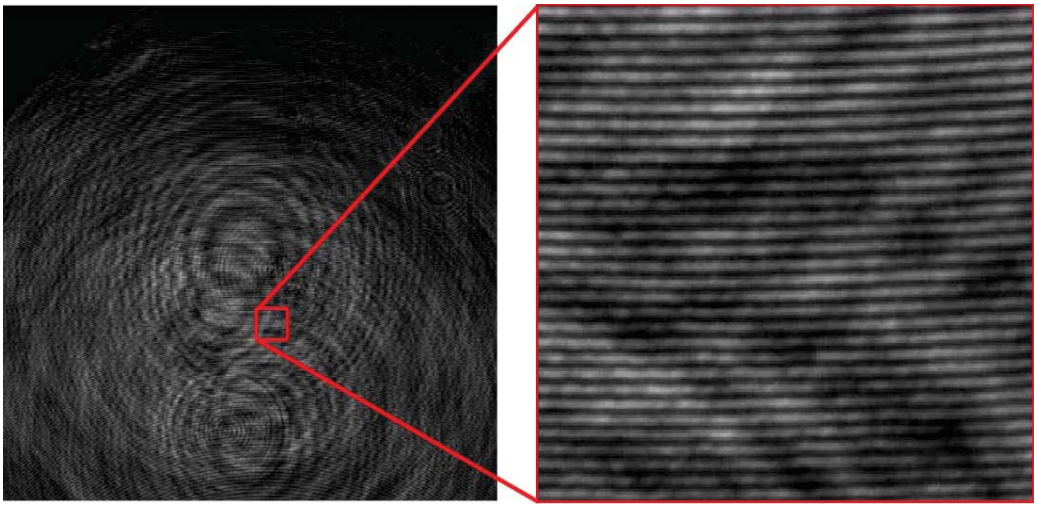

Figure 4.8: Measurement with the spatial carrier method.

The cosine function in the irradiance term can be described in terms of complex exponentials:

$$I(x,y) = I_1(x,y) + I_C(x,y)e^{2\pi i f_0 x} + I_C^*(x,y)e^{-2\pi i f_0 x} \quad (4.9)$$

where the complex irradiance term I_C is defined as:

$$I_C(x,y) = \frac{1}{2} I_2(x,y)e^{i\phi(x,y)} \quad (4.10)$$

In this form, it can be seen that the complex irradiance that carries the phase is riding a carrier term along the x-direction $\left(e^{2\pi i f_0 x}\right)$. This allows the complex irradiance term to be isolated in frequency space. A Fourier transform along the x-direction on the irradiance results in:

$$\mathcal{I}(f_x,y) = \mathcal{I}_1(f_x,y) + \mathcal{I}_C(f_x - f_0,y) + \mathcal{I}_C^*(f + f_0,y) \quad (4.11)$$

where the script letters denote the Fourier transformed data $\mathcal{I} = \mathcal{F}\{I\}$, and f_x is the spatial frequency along the x direction. If the variation in phase is small in comparison to the carrier frequency, the spectrum that carries the phase ($\mathcal{I}_C$) will be isolated from the amplitude variation of $\mathcal{I}_1$. Figure 4.9 illustrates the separation of the phase spectrum in Fourier space. The left image shows an example where the carrier frequency f_0 is not large enough to separate the term carrying the phase $\mathcal{I}_C$ from the central term $\mathcal{I}_1$, whereas the image on the right has a large enough carrier frequency to isolate the terms. The phase carrying spectrum $\mathcal{I}_C$ can be isolated leaving only $\mathcal{I}_C(f_x - f_0, y)$ which is then translated by f_0 to center it (Figure 4.10). Figure 4.11 shows the Fourier transformed set of data for the example used in this section. For this particular example, the tilt has been introduced along the y-direction and the red box outlines the window where the phase carrying spectrum will be extracted.

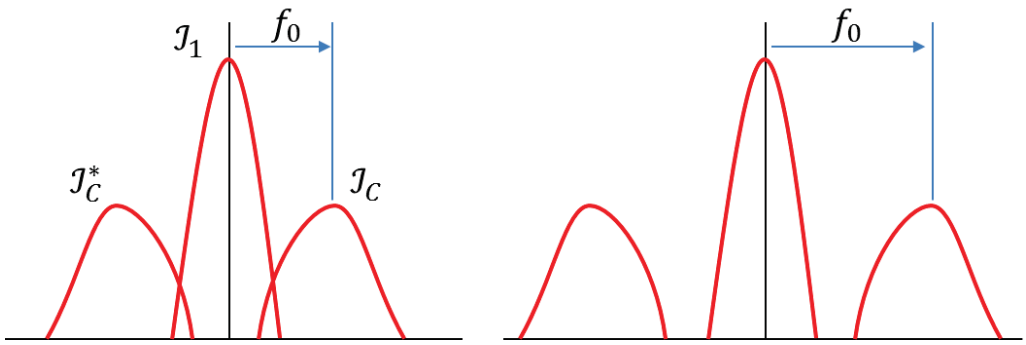

Figure 4.9: One-dimensional Fourier transform of the irradiance with a spatial carrier. (Left) The carrier frequency f_0 is not large enough to isolate the phase carrying data from the central term. (Right) A correct carrier frequency isolates the phase carrying data.

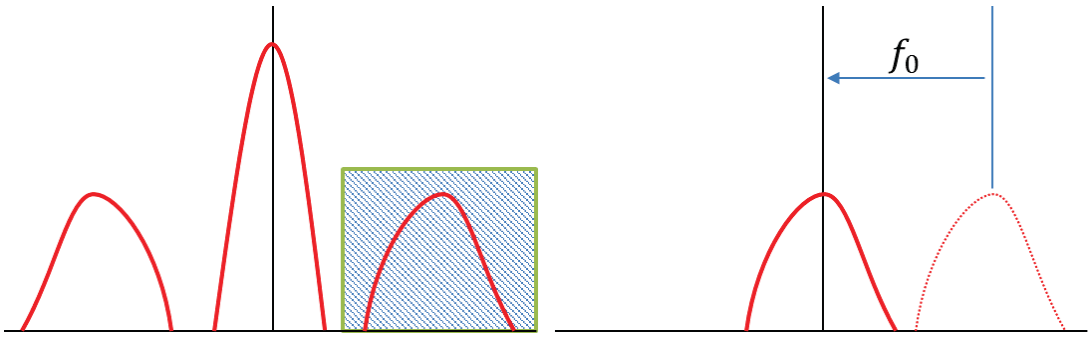

Figure 4.10: Isolation of phase carrying signal. (Left) Phase carrying data is windowed. (Right) Windowed data is centered on-axis.

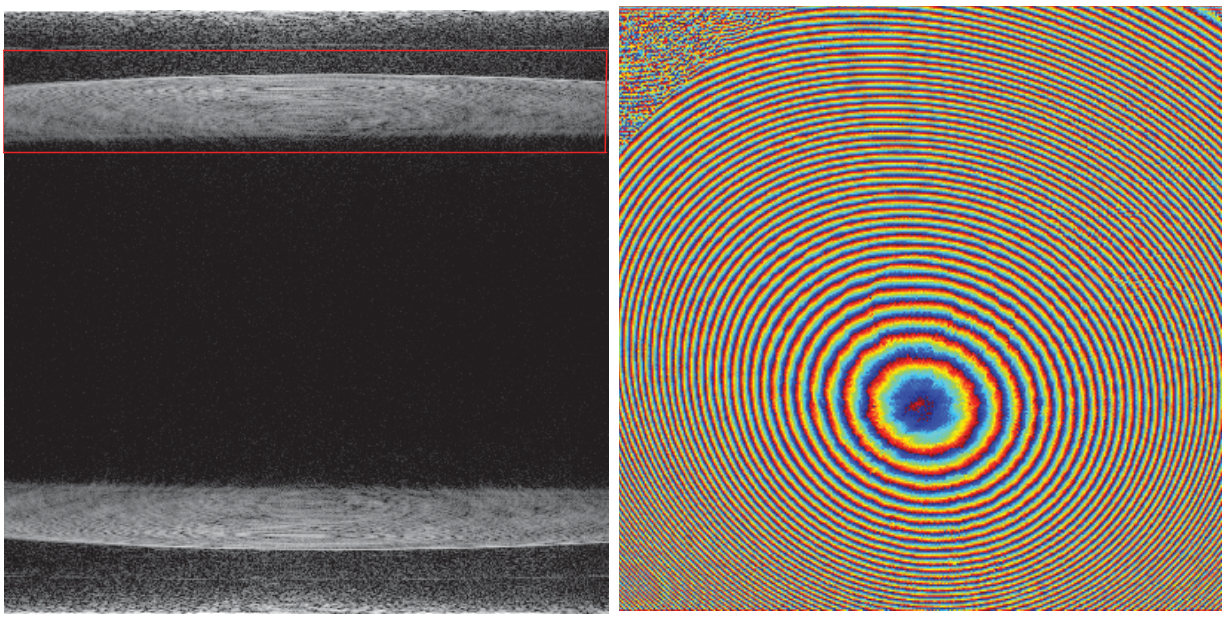

Figure 4.11: Fourier transformed data (left) and recovered phase (right).

The filtered irradiance $(\mathcal{I}_{F,C})$ spectrum can then be described as:

$$\mathcal{I}_{F,C}(f_x, y) = \mathcal{I}_C(f_x - f_0, y) \quad (4.12)$$

The inverse Fourier transform of the filtered irradiance $\mathcal{I}_{F,C}$ results in an isolated irradiance term $(I_{F,C})$:

$$I_{F,C}(x, y) = I_C(x, y) = \frac{1}{2} I_2(x, y) e^{i\phi(x,y)} \quad (4.13)$$

The phase can be isolated by taking a logarithm of both sides:

$$\log\left(I_{F,C}(x, y)\right) = \log\left(\frac{1}{2} I_2(x, y) e^{i\phi}\right) = \log\left(\frac{1}{2} I_2(x, y)\right) + i\phi \quad (4.14)$$

The phase is now isolated to the imaginary component of the complex function. The recovered phase (ψ) can be described in terms of the filtered irradiance as:

$$\psi = \text{imag}\left\{\log\left(I_{F,C}(x, y)\right)\right\} \quad (4.15)$$

Where the recovered phase from the spatial carrier method is also unresolved by a modulo of 2π as a result of measuring the irradiance. Phase unwrapping is also required for the spatial carrier method to recover the true phase (Section 4.2). The recovered phase from the example is shown in Figure 4.11 (right).

The DC irradiance and modulation terms can also be recovered from the spatial carried data. First, the irradiance (Eq. 4.8) is rearranged:

$$I(x,y) = I_1(x,y)\left[1 + \frac{I_2(x,y)}{I_1(x,y)}\cos\big(2\pi f_0 x + \phi(x,y)\big)\right] \qquad (4.16)$$

The DC irradiance term is then $I_1 = I_{DC}$. It can be filtered out from the Fourier transformed data (Eq. 4.11) by windowing out the central peak:

$$\mathcal{I}_{F,DC}(f_x, y) = \mathcal{I}_1(f_x, y) \qquad (4.17)$$

And then applying an inverse Fourier transform:

$$I_{F,DC}(x,y) = I_1(x,y) = I_{DC}(x,y) \qquad (4.18)$$

Modulation is given by the relationship:

$$\gamma = \frac{I_2(x,y)}{I_1(x,y)} \qquad (4.19)$$

The second irradiance term I_2 is isolated from the initial set of filtered data (Eq. 4.13) by extracting the real component:

$$A(x,y) = \text{real}\left\{\log\left(I_{F,C}(x,y)\right)\right\} = \log\left(\frac{1}{2}I_2(x,y)\right) \qquad (4.20)$$

And solving for I_2:

$$I_2(x,y) = 2e^{A(x,y)} \quad (4.21)$$

The resulting irradiance and modulation for the example in this section are shown in Figure 4.12.

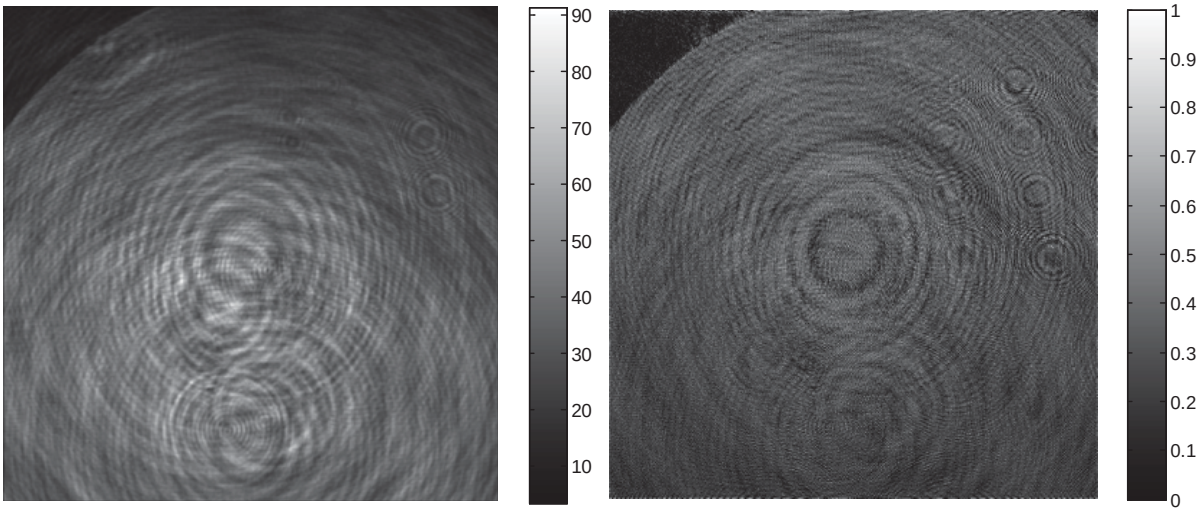

Figure 4.12: Recovered information for the spatial carrier method.
(Left) Modulation. (Right) DC Irradiance.

4.2 PHASE UNWRAPPING

The measured irradiance for interference patterns are a cyclical function of the measured phase:

$$I(x,y) = I_1(x,y) + I_2(x,y)\cos\big(\phi(x,y)\big) \quad (4.22)$$

Therefore, the recovered phase will be a modulo 2π of the true phase:

$$\psi = \mathcal{W}[\phi] = \phi - 2\pi\left\lfloor\frac{\phi}{2\pi}\right\rfloor \quad (4.23)$$

Where $\mathcal{W}$ is the wrapping operator, which removes a modulus number of 2π steps from the data. The result is that the recovered phase data is mapped between $(-\pi,\pi]$. When the phase exceeds $\pm\pi$, it is wrapped around and results in discontinuities in the recovered phase. An example wrapped function is

shown in Figure 4.13 for a quadratic phase (e.g. defocus). The wrapped phase for the image is shown on the left, and a cross-sectional plot along the center of the data is shown on the right. The process of removing 2π discontinuities and recovering the surface is known as phase unwrapping and is a widely researched topic (Ghiglia & Pritt, 1998).

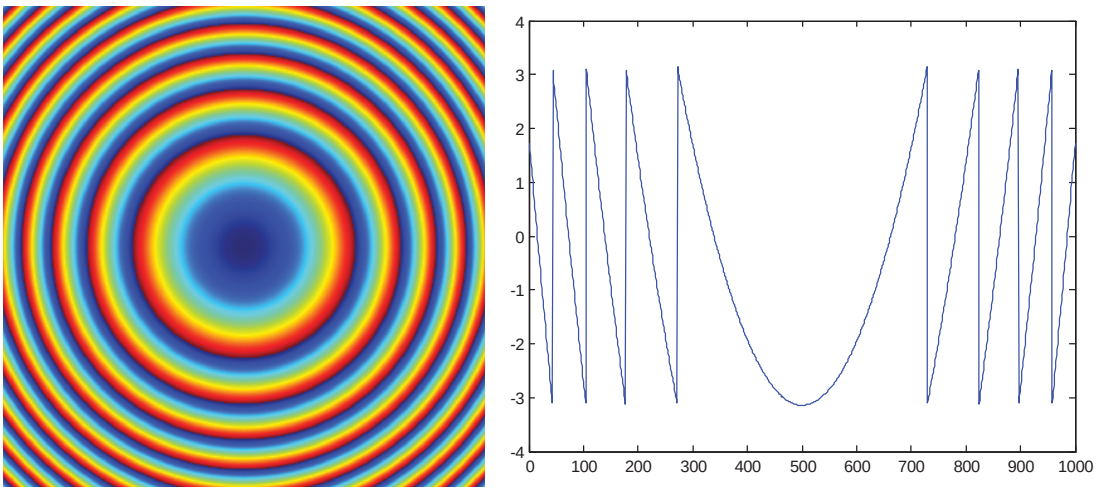

Figure 4.13: Example of wrapped phase.

Conventional interferometric testing typically results in data with high signal-to-noise ratios and minimal defects in the image. This reduces the burden on the phase unwrapping process, which tend to be sensitive to noise and defects in the phase surface. However, the tear film surface is unlike most optical surfaces. First, the structure of the tear film contains a large amount of mid- to high-spatial frequency content. This becomes an issue when structures exceed the dynamic range of the instrument, generally where tear film break-up occurs, or when mucus globules or other particles appear in the tear film surface. Eye disease, such as dry eye syndrome, may increase the prevalence of mid- to high-spatial frequency structures. These structures have to be avoided or minimally processed during the phase unwrapping routine. Second, since data is processed on a frame-by-frame basis, eye motion will induce a significant amount of tilt and defocus, further increasing fringe density in a given frame of data. Finally, ocular variation will contribute additional low order surface features, typically astigmatism and spherical aberration. The interaction of all these effects will generate dense fringe patterns on the detector that are reduced in

modulation by detector sampling. The reduced modulation results in a reduced signal-to-noise ratio, increasing the difficulty of a successful phase unwrapping. An example of a breakdown during phase unwrapping is shown in Figure 4.14 on the left. The discontinuities result from a failure of the phase unwrapping algorithm to properly resolve 2π discontinuities. Path following algorithms, where pixels are processed from a neighborhood queue, tend to propagate these errors across the rest of the image. A successful phase unwrapping of the same data is shown in Figure 4.14 on the right.

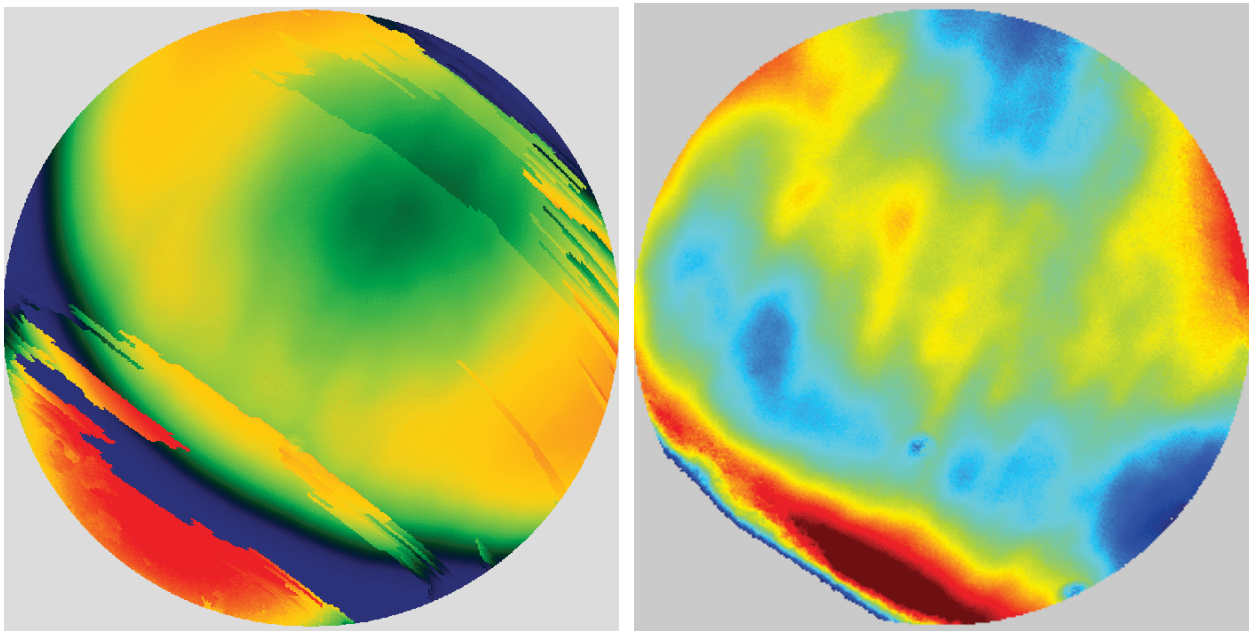

Figure 4.14: Example of a phase unwrapping breakdown.
(Left) Breakdown occurring in commercial software. (Right) Same phase data processed with a more robust algorithm. Note: The surface on the left was processed with a different application than on the right, resulting in different color maps.

The example shown in Figure 4.14 is representative of most of the data that has been captured on the TFI and SPINALTOP. In fact, a significant portion of the data will be even more complicated, especially when measuring the tear film interaction with a contact lens. Because of this, it was necessary to develop customized software routines to process the type of data that was being recorded with the TFI system and what was to be expected on SPINALTOP. This section covers the entire process of recovering the unwrapped phase, which is outlined in Figure 4.15.

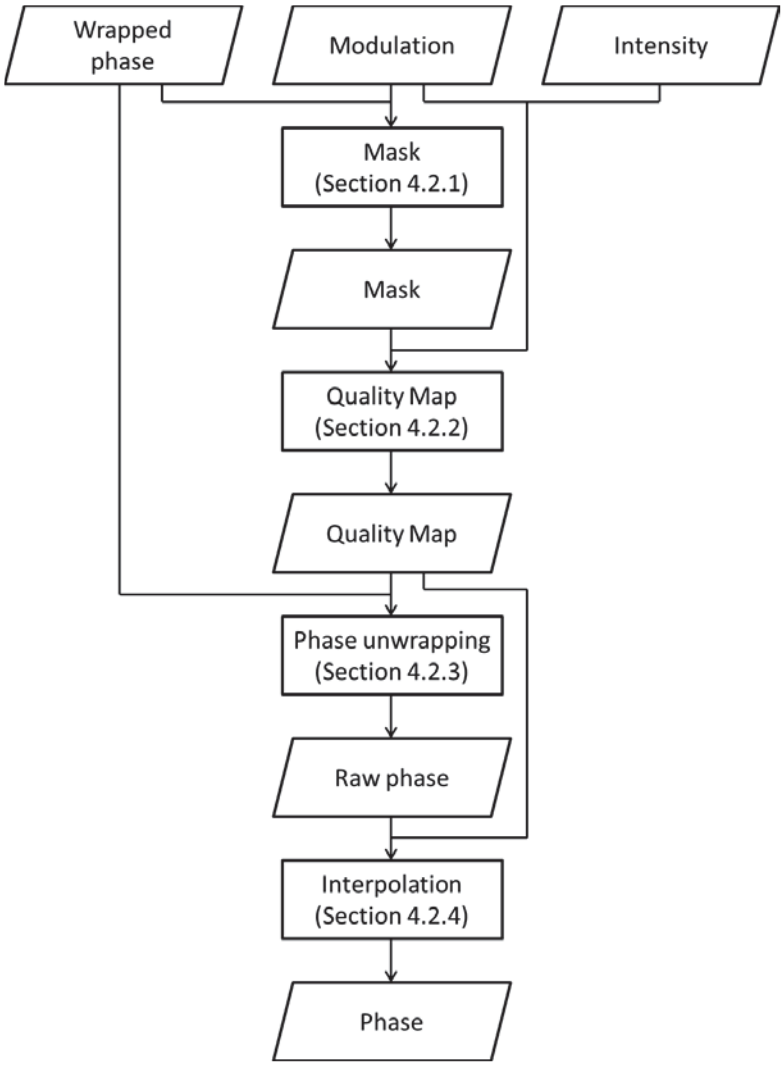

Figure 4.15: Outline of the phase unwrapping process.

4.2.1 MASKS

A set of masks are required for the phase unwrapping process. A mask contains a listing of pixels and their status. A status can denote whether a pixel is good or bad, or if a pixel has been phase unwrapped, to name a few examples. Masks are commonly formatted to match the dimensions of the image or data such that a one-to-one correspondence can be established between the two arrays; hence the name mask. This section defines a number of masks which will be used to support the phase unwrapping routines (Section 4.2.3) and interpolation routine (Section 4.2.4).

4.2.1.1 MODULATION

The most basic mask is a modulation mask (M_{mod}). It is generated by applying a threshold (γ_t) to the modulation (γ) that was recovered from the phase measurement (Section 4.1):

$$M_{mod}(i,j) = \begin{cases} 0, & \gamma(i,j) < \gamma_t \\ 1, & \gamma(i,j) \geq \gamma_t \end{cases} \quad (4.24)$$

where i,j are used to represent pixel indices. An initial guess for a threshold would be just above the background level modulation. Background modulation can be estimated by blocking the test or reference path and recording the modulation. This may not provide a complete characterization of the background, as it could block stray reflections from surfaces in either the test or reference path that may contribute to a modulation signal. The surface interferometer typically has an average background modulation level $\gamma_t = 15\%$. The modulation mask for an example set of data with this threshold is shown in Figure 4.16.

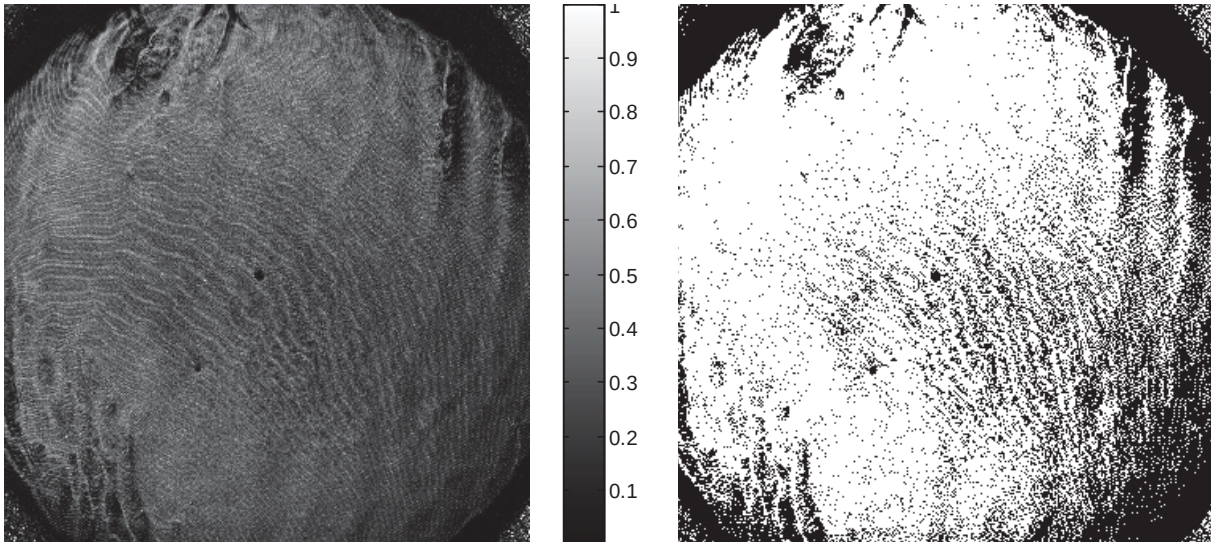

Figure 4.16: Modulation mask example.
(Left) Modulation map. (Right) Mask resulting from a 15% threshold.

However, it was found that using the background noise as a threshold was insufficient for a large percentage of tear film measurements. Eye motion results in a large amount of tilt that generates dense fringes across the field that contributes to an overall constant reduction in modulation. The recovered

modulation for such a measurement could be equal to or less than the background modulation. This would result in whole frames of valid data being marked as “bad” or a significant portion of data being masked out. The example shown in Figure 4.16 has a large number of pixels masked that are potentially usable.

A solution to the false-positives that resulted from the masking in Eq. 4.24 was to update the definition of the mask to:

$$M_{mod}(i,j) = \begin{cases} 0, & \gamma(i,j) < \gamma_{t1} \\ 1, & \gamma_{t1} \leq \gamma(i,j) < \gamma_{t2} \\ 2, & \gamma(i,j) \geq \gamma_{t2} \end{cases} \quad (4.25)$$

Where a 0 denotes a “bad” pixel, a 1 denotes a “questionable” pixel, and a 2 denotes a “good” pixel. Two different thresholds are used that meet the requirement:

$$0 \leq \gamma_{t1} \leq \gamma_{t2} \leq 1 \quad (4.26)$$

A tiered mask provides a level of confidence on the “goodness” of data, which can be used to adjust the behavior of the algorithms. For example, the phase unwrapping routine can process “good” pixels without any special treatment, while avoiding “questionable” and “bad” pixels. A second iteration of the phase unwrapping routine can then be run where the “good” pixels are already marked as completed, “questionable” pixels are processed, and the “bad” pixels are still avoided. This allows good, yet questionable, pixels to be processed while at the same time minimizing the impact from bad data that manages to slips through the masking process.

4.2.1.2 RESIDUES

Ghiglia, Mastin, and Romero originally described inconsistencies in the recovered phase that inhibited phase unwrapping (Ghiglia, Mastin, & Romero, 1987). Goldstein, Zebker, and Werner later described these phase inconsistencies as phase residues, analogous to residues encountered in complex analysis (Goldstein, Zebker, & Werner, 1988). Residues complicate the phase unwrapping process by presenting a

set of data that is inconsistent with an implicit assumption that the phase surface that is being unwrapped is smooth and continuous. Therefore, residues are best handled by masking them out of any processing routines.

In complex analysis, the closed path integral of a complex function is given by:

$$\oint f(z)dz = 2\pi i \cdot Res(f, \alpha) \qquad (4.27)$$

Where $Res(f, \alpha)$ represents the number of residues, or singularities, over the complex function $f(z)$. Figure 4.17 shows a surface that contains a single residue at the center of the image that appears as a singularity.

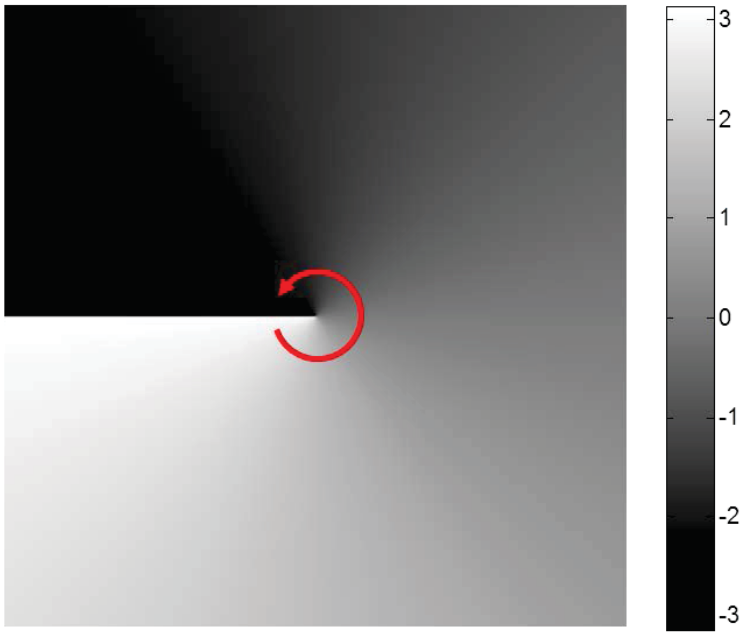

Figure 4.17: Residue occurring around a singularity.

The analogous application to phase is given by the closed path integral of the gradient:

$$\mathcal{R} = \oint \nabla\phi(x, y) \cdot dr \qquad (4.28)$$

Where one or more residues exist for $\mathcal{R} \neq 0$. The implementation of a discrete residue calculation for the measured phase is given by numerically integrating counter-clockwise over the closed path integral of a 2 x 2 neighborhood of pixels:

$$\begin{aligned}\mathcal{R}(i,j) = [\phi(i,j) + \phi(i-1,j)] + [\phi(i-1,j) + \phi(i-1,j-1)] \\ +[\phi(i-1,j-1) + \phi(i,j-1)] + [\phi(i,j-1) + \phi(i,j)]\end{aligned} \quad (4.29)$$

Numerical errors and approximations from the discretized phase measurement will result in non-zero values for all values of $\mathcal{R}$, and therefore a threshold is also applied determining the location of residuals. If a residue is detected, it would be indeterminate which of the four pixels used in the calculation resulted in the discontinuity. Therefore, all four pixels will be masked out if a residue is detected. The residue mask (M_{res}) is initially allocated assuming all phase points are valid:

$$M_{res}(i,j) = 1 \; \forall \; i,j \in \phi \quad (4.30)$$

A residue is defined by exceeding a threshold ($\mathcal{R}_t$) and the 2 x 2 neighborhood is removed from the mask:

$$\begin{aligned}\text{if } abs\big(\mathcal{R}(i,j)\big) > \mathcal{R}_t \; \rightarrow M_{res}(i,j) = M_{res}(i-1,j) = M_{res}(i,j-1) \\ = M_{res}(i-1,j-1) = 0\end{aligned} \quad (4.31)$$

Trial and error processing of tear film measurements has shown that $\mathcal{R}_t = 0.01$ is sufficient to capture most residues with minimum false-positives. Residues are treated as “bad” and are ignored until the data is replaced by interpolation (Section 4.2.3). Figure 4.18 shows the calculated residue mask for the same example set of data from Figure 4.16.

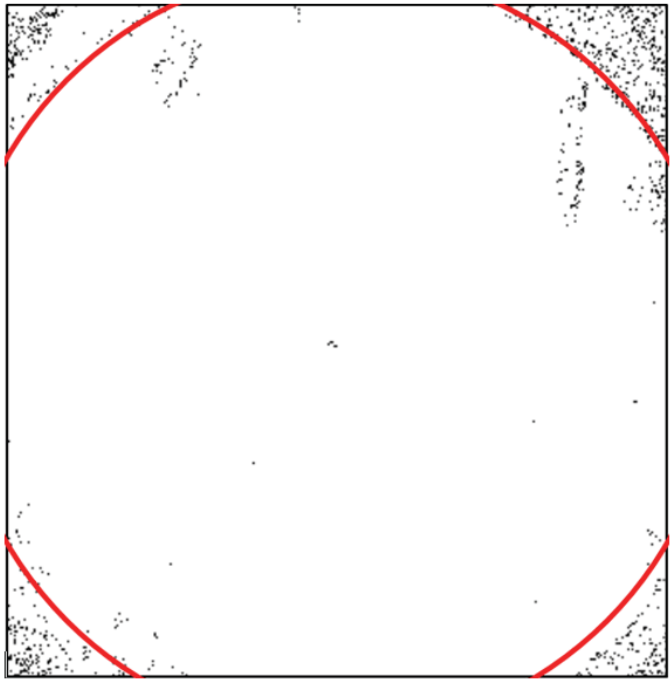

Figure 4.18: Example residue mask.
The red lines have been added to distinguish between signal and background data, which is located in the four corners.

4.2.1.3 PHASE DERIVATIVE VARIANCE

The phase derivative variance metric was developed for Interferometric Synthetic Aperture Radar (InSAR) applications where modulation data is not available (Ghiglia & Pritt, 1998; Pritt, 1996). This has been found to be a useful metric to supplement modulation when it fails with the tear film measurements. The phase derivative variance works by comparing the instantaneous slope at a given pixel to the local neighborhood of pixel slopes. The phase slope ($\widetilde{\Delta}$) is defined as forward difference of the measured phase:

$$\widetilde{\Delta}_x(i,j) = \phi(i+1,j) - \phi(i,j) \qquad (4.32)$$

$$\widetilde{\Delta}_y(i,j) = \phi(i,j+1) - \phi(i,j) \qquad (4.33)$$

The recovered phase will contain 2π phase discontinuities and will need to be removed to accurately represent the phase slope (Δ):

$$\Delta_x(i,j) = \begin{cases} \widetilde{\Delta}_x(i,j) - 2\pi, & \widetilde{\Delta}_x(i,j) > \pi \\ 2\pi - \widetilde{\Delta}_x(i,j), & \widetilde{\Delta}_x(i,j) < -\pi \\ \widetilde{\Delta}_x(i,j), & -\pi \leq \widetilde{\Delta}_x(i,j) \leq \pi \end{cases} \quad (4.34)$$

$$\Delta_y(i,j) = \begin{cases} \widetilde{\Delta}_y(i,j) - 2\pi, & \widetilde{\Delta}_y(i,j) > \pi \\ 2\pi - \widetilde{\Delta}_y(i,j), & \widetilde{\Delta}_y(i,j) < -\pi \\ \widetilde{\Delta}_y(i,j), & -\pi \leq \widetilde{\Delta}_y(i,j) \leq \pi \end{cases} \quad (4.35)$$

The $k\ x\ k$ neighborhood mean of phase slopes ($\overline{\Delta}$), where $k = 2N + 1$, is given by:

$$\overline{\Delta_x}(i,j) = \frac{1}{(2N+1)^2} \sum_{m=i-N}^{i+N} \sum_{n=j-N}^{j+N} \Delta_x(m,n) \quad (4.36)$$

$$\overline{\Delta_y}(i,j) = \frac{1}{(2N+1)^2} \sum_{m=i-N}^{i+N} \sum_{n=j-N}^{j+N} \Delta_y(m,n) \quad (4.37)$$

The directional phase derivative variance (σ^2) is given by the expression:

$$\sigma_x^2(i,j) = \frac{1}{(2N+1)^2} \sum_{m=i-N}^{i+N} \sum_{n=j-N}^{j+N} \left(\Delta_x(m,n) - \overline{\Delta_x}(i,j)\right)^2 \quad (4.38)$$

$$\sigma_y^2(i,j) = \frac{1}{(2N+1)^2} \sum_{m=i-N}^{i+N} \sum_{n=j-N}^{j+N} \left(\Delta_y(m,n) - \overline{\Delta_y}(i,j)\right)^2 \quad (4.39)$$

Where the phase derivative variance can then be found by the expression:

$$\sigma_{pdv}^2(i,j) = \sigma_x^2(i,j) + \sigma_y^2(i,j) \quad (4.40)$$

A mask (M_{pdv}) can be derived from the phase derivative variance by applying a threshold (σ_t^2):

$$M_{pdv}(i,j) = \begin{cases} 0, & \sigma_{pdv}^2(i,j) \geq \sigma_t^2 \\ 1, & \sigma_{pdv}^2(i,j) < \sigma_t^2 \end{cases} \quad (4.41)$$

Figure 4.19 shows an example data set with the calculated gradient image (left) and resulting mask (right).

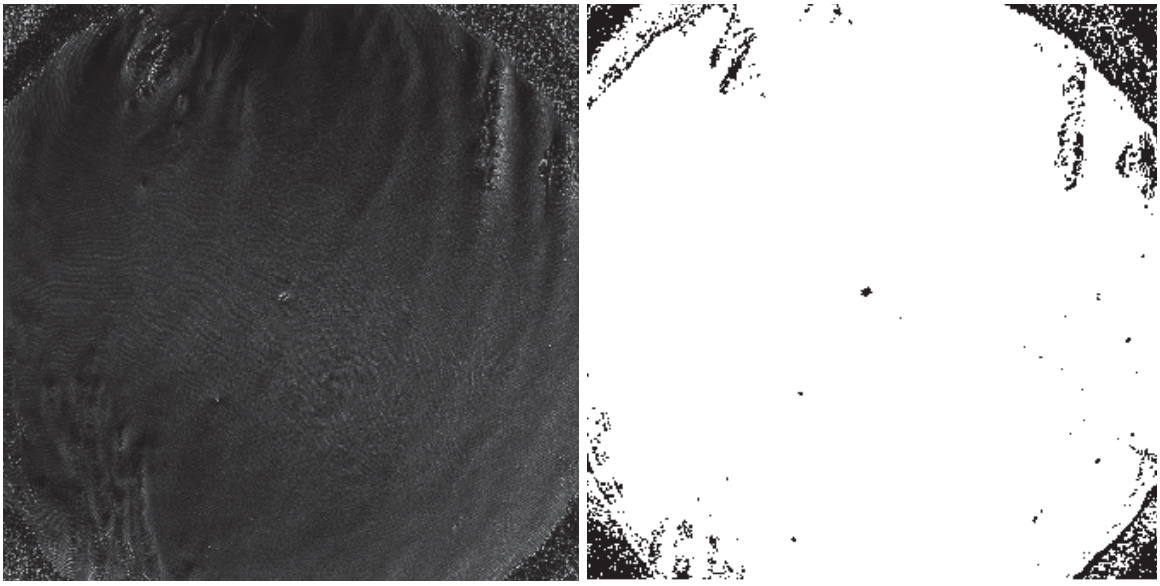

Figure 4.19: Example phase derivative variance mask.
(Left) Gradient of the phase. (Right) Resulting mask for a given threshold.

4.2.1.4 IRRADIANCE

The irradiance mask (M_{irr}) is constructed in the same way as the modulation mask (Section 4.2.1.1). A threshold (I_t) determines what pixels are masked:

$$M_{irr}(i,j) = \begin{cases} 0, & I(i,j) < I_t \\ 1, & I(i,j) \geq I_t \end{cases} \quad (4.42)$$

During the development of the previous masks, it was noted that pixels in the far corners would still pass through the mask. This can be seen in the examples in Figure 4.16, Figure 4.18, and Figure 4.19. The corner areas obviously contain invalid or bad data because those areas are vignetted by apertures in the

system, so no information exists in there. A fixed mask (next section) can be used to isolate these areas; however, the fact that not all pixels in these regions are captured demonstrates a limitation in the previous masks. It should be expected that dead zones will appear in the measured tear film data. These dead-zones are typically the result of mucin globules that create surfaces that are too steep to resolve, areas of tear film breakup where no signal is returned, or eyelashes that appear in the data. Figure 4.20 shows an example of an irradiance measurement (left) and a calculated mask (right). The corners are easily caught with minimal hits in the center of the image.

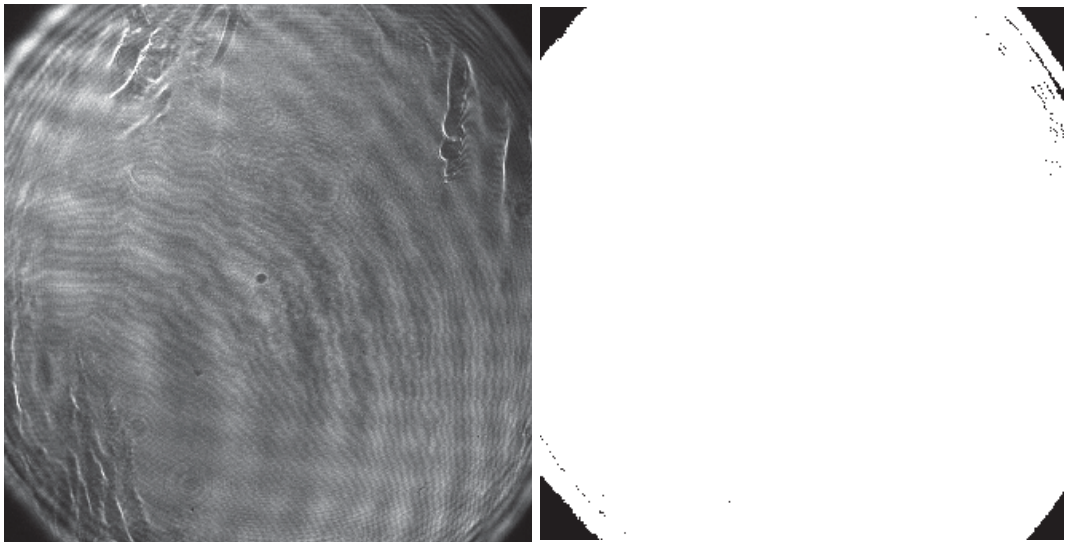

Figure 4.20: Example of an irradiance mask.
(Left) Irradiance measurement. (Right) Irradiance mask.

4.2.1.5 USER DEFINED MASK

A user defined mask is a mask that has been predetermined by the operator (i.e. does not require an algorithmic solution). This type of mask is most commonly used when the apertures of a system or surface under test is known. In the previous section, it was pointed out that it is known that the corners of the data do not contain signal because of the clear apertures in the system limiting the data that makes it to the detector. Additionally, dead or noisy pixels can be added to the user defined mask to ensure that these pixels are removed from any processing. An example of a user defined mask that would supplement the

masks shown in the previous sections is shown in Figure 4.21. The user defined mask will be denoted as: M_{User}

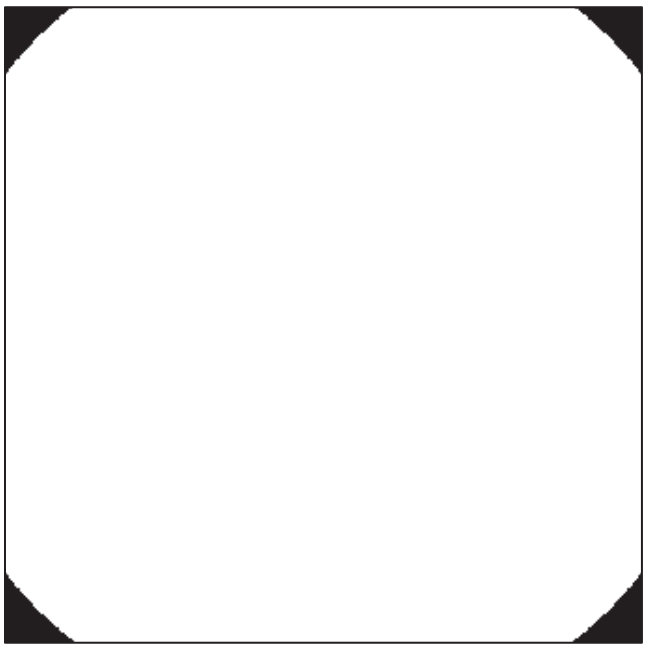

Figure 4.21: Example of a user defined mask.

4.2.1.6 FINAL MASK

The four masks from the previous sections can be combined to create a system mask to denote the location of all "bad" pixels by a simple disjunction operation:

$$M = (M_{mod} = 0) \cup M_{res} \cup M_{pdv} \cup M_{Irr} \cup M_{User} \quad (4.43)$$

4.2.2 QUALITY MAP

A quality map is effectively a mask that contains a continuous range of values associated with each pixel in the phase image. The values are weightings that describe the quality of a pixel. Pixels that are weighted higher than other pixels may be considered more trustworthy. Quality maps are used to guide or weight critical operations, such as phase unwrapping. A number of phase unwrapping algorithms are sensitive to

noise, so it is more preferable to operate on low-noise/high-quality pixels first before attempting to process lower-quality pixels.

Phase shifting interferometry provides a straightforward method for recovering the modulation of the measured phase (Section 4.1), which has been shown to be a useful quality map (Y. Xu & Ai, 2003). However, using modulation as a quality map has the same issues as the modulation mask (i.e. sensitivity to high slope measurements, Section 4.2.1.1). The phase variance (Section 4.2.1.3) is used to supplement modulation as a quality map. A combination of modulation and phase variance was more successful at recovering tear film measurements than any one metric individually. Further improvements were found by including irradiance data and reweighting the quality map using the modulation and phase variance masks. The final quality map ($\mathcal{Q}$) is defined as:

$$\mathcal{Q}(i,j) = \left[\left(\mathcal{M}(i,j) + \mathcal{P}(i,j)\right) \cdot \mathcal{I}(i,j)\right]^{\mathcal{W}(i,j)} \quad (4.44)$$

Where $\mathcal{M}$ is modulation, $\mathcal{P}$ is the slope variance error, $\mathcal{I}$ is the irradiance, and $\mathcal{W}$ is a weighting function that will be defined in the next paragraph. The slope variance weighting is inverted from the calculated slope variance to give higher weights to small errors. However, the inversion can quickly blow up where slope errors are small and is defined as such:

$$\mathcal{P} = \begin{cases} \left(\sigma_{pdv}^2\right)^{-1}, & \left|\sigma_{pdv}^2\right| > 0.01 \\ 100, & \left|\sigma_{pdv}^2\right| \leq 0.01 \end{cases} \quad (4.45)$$

The irradiance is combined directly without modification, as this was found to work sufficiently well, although a more robust algorithm could be developed by re-weighting the irradiance. Finally, the whole value is reweighted by $\mathcal{W}$ which is defined by the masks:

$$\mathcal{W} = \begin{cases} w, & (M_{mod} = 2) \cup M_{pdv} = 1 \\ 1, & (M_{mod} = 1) \cup M_{pdv} = 1 \\ 0.5, & \text{otherwise} \end{cases} \quad (4.46)$$

Where the weighting factor w was determined empirically as $w = 3.33$. Figure 4.22 is a comparison of the modulation (left) to the quality map (right). The quality map is more uniform over areas where it is apparent that data exists, and is able to selectively weight areas in the high slope regions that would typically cause issues. Other areas where breakup and mucin globules in the tear film are also weighted lower with the quality map, which do not show up as well in the modulation.

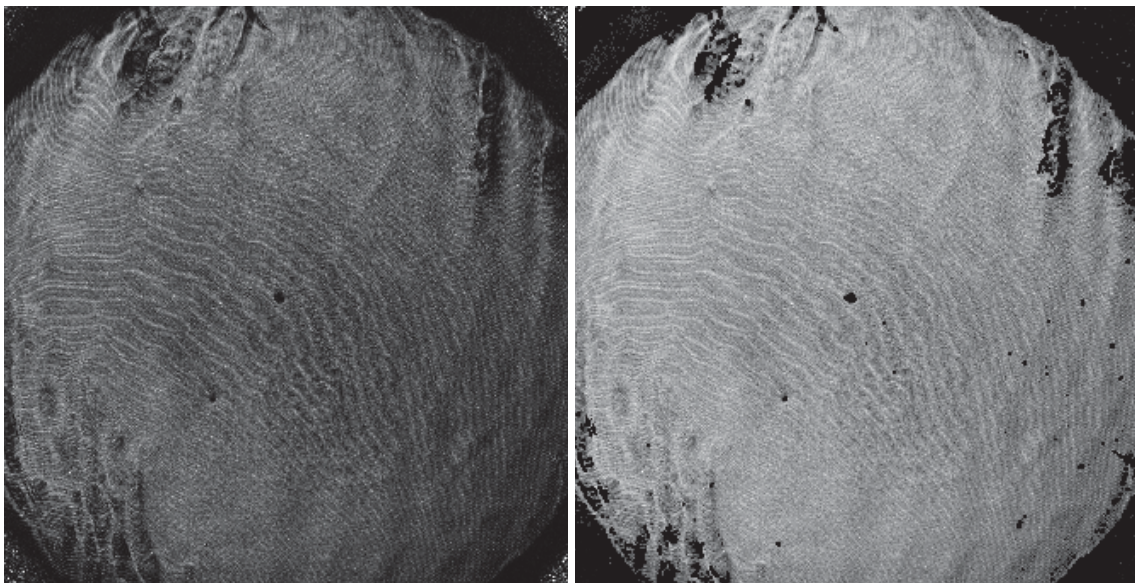

Figure 4.22: Comparison of modulation (left) to the equivalent quality map (right).

4.2.3 PHASE UNWRAPPING METHODS

The process of recovering the original phase from the wrapped phase is called phase unwrapping. Phase unwrapping is an extensive topic and has been widely researched (Ghiglia & Pritt, 1998; Judge & Bryanston-Cross, 1994; D. Malacara, 2007; Zacarias Malacara & Servín, 2010). Therefore, it is not the

goal of this section to exhaustively research all of the various algorithms, but to describe the algorithms that were found to be successful in unwrapping the type of data that has being measured with the TFI and SPINALTOP. In particular, two algorithms were found to work with these datasets: A path-following method and a minimum-norm method. Both of these methods will be discussed in detail in the following two sections.

4.2.3.1 PATH-FOLLOWING METHODS

Phase unwrapping algorithms can be broken up into two domains: Local and global. Local phase unwrapping routines rely on a small neighborhood of data to determine the phase unwrapping correction. These methods are also more generally known as path-following methods, the reason for which will become clear in the following sections. Before describing the quality-guided method, it will be necessary to describe the series of methods that build up to it. In addition to providing background for this method, it will demonstrate why more complicated techniques are required to process TFI and SPINALTOP data.

4.2.3.1.1 RASTER METHOD

An early discussion of phase unwrapping was by Itoh (Itoh, 1982). Itoh's method was described for a 1-dimensional signal, but can be easily extended to 2-dimensions. The method requires that the phase slope does not exceed π over a measurement, which is consistent with the Nyquist sampling requirement to unambiguously resolve interferometric data (Gaskill, 1978). A slope change between two adjacent measurement values, or pixels, that has a magnitude greater than π can therefore be attributed to a wrapping discontinuity. Phase unwrapping is performed by detecting the discontinuity and adjusting the data by an integer amount of 2π until the discontinuity is removed.

The 1-D phase difference between neighboring pixels is given by the forward difference method:

$$\Delta(i) = \psi(i+1) - \psi(i) \qquad (4.47)$$

And the unwrapped phase is recovered:

$$\tilde{\phi}(i) = \psi(i) + 2\pi \cdot \left\lfloor \frac{\Delta(i)}{2\pi} \right\rceil \qquad (4.48)$$

where the $\lfloor \quad \rceil$ represents a rounding operation. If the magnitude of the phase difference exceeds π, the rounding operation will carry it to an integer value that then allows for a 2π correction to the phase. The unwrapping process then proceeds to the next value at $\psi(i+1)$, where the process is repeated. An example of this process is shown in Figure 4.23. The first image (top-left) shows an arbitrary curve that has been wrapped, with a discontinuity appearing between third and fourth data points. The second image (top-right) shows that the 4th data point would be detected as wrapped, and corrected by a 2π offset. The process is repeated in the next two images and results in the final unwrapped curve.

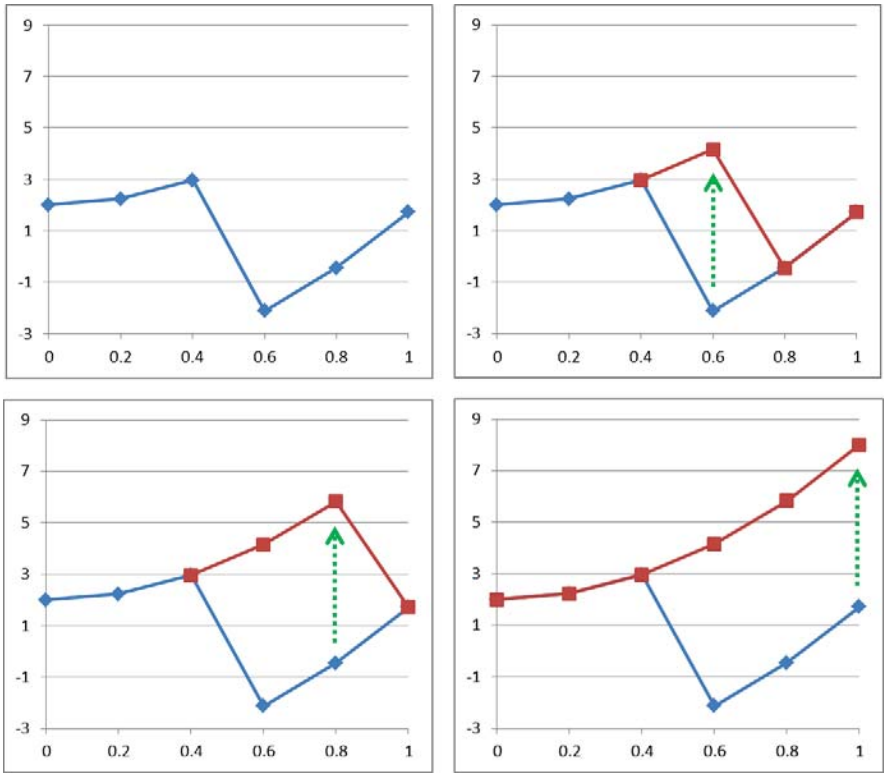

Figure 4.23: One dimensional phase unwrapping example.

Itoh pointed out that phase unwrapping could become corrupted by noise. Noise complicates the detection or phase discontinuities, resulting in an incorrect 2π adjustment being applied or bypassed. Because of the

reliance on previous phase unwrapped data points, errors introduced by noise may be propagated across the image. An example of noise corrupting the unwrapping process is shown in Figure 4.24. The red plot along the middle is a wrapped function. The uppermost blue plot is the proper unwrapped function. Noise is added to the wrapped function, unwrapped, and the result is plotted below the previous plot for display purposes. This is repeated three more times, with increasing noise, until the noise dominates the measurement and the result is unusable. The second to last blue plot is an example of what is most commonly seen in the unwrapped data, where a few data points will be incorrectly offset due to noise. The offset propagates across a small area until another noise error occurrence, which may or may not push it back into a proper offset. The phase unwrapping breakdown example shown in Figure 4.14 (left) is a result this issue, but applied across two dimensions.

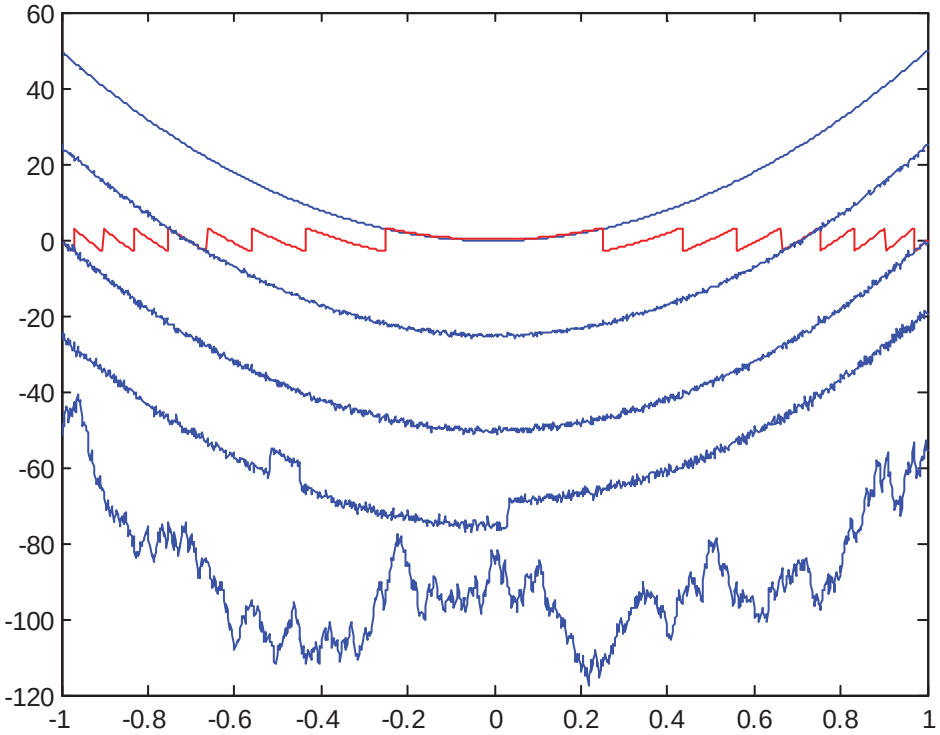

Figure 4.24: Example of noise corrupting the phase unwrapping process.

Before addressing how noise is handled, phase unwrapping in 2-dimensions needs to be discussed. The most straightforward extension of Itoh's method to 2-dimensions is by a Raster scan method. In a 2-dimensional image, a corner of the phase image can be selected as the starting point. The entire row can

be phase unwrapped using the 1-dimensional method. When the phase unwrapping routine reaches the end of the row, the pixel in the orthogonal direction is unwrapped for a single step. The process can be repeated along the new row, in reverse, until the edge is reached. The process can be repeated until every pixel has been unwrapped. This process is illustrated with Figure 4.25.

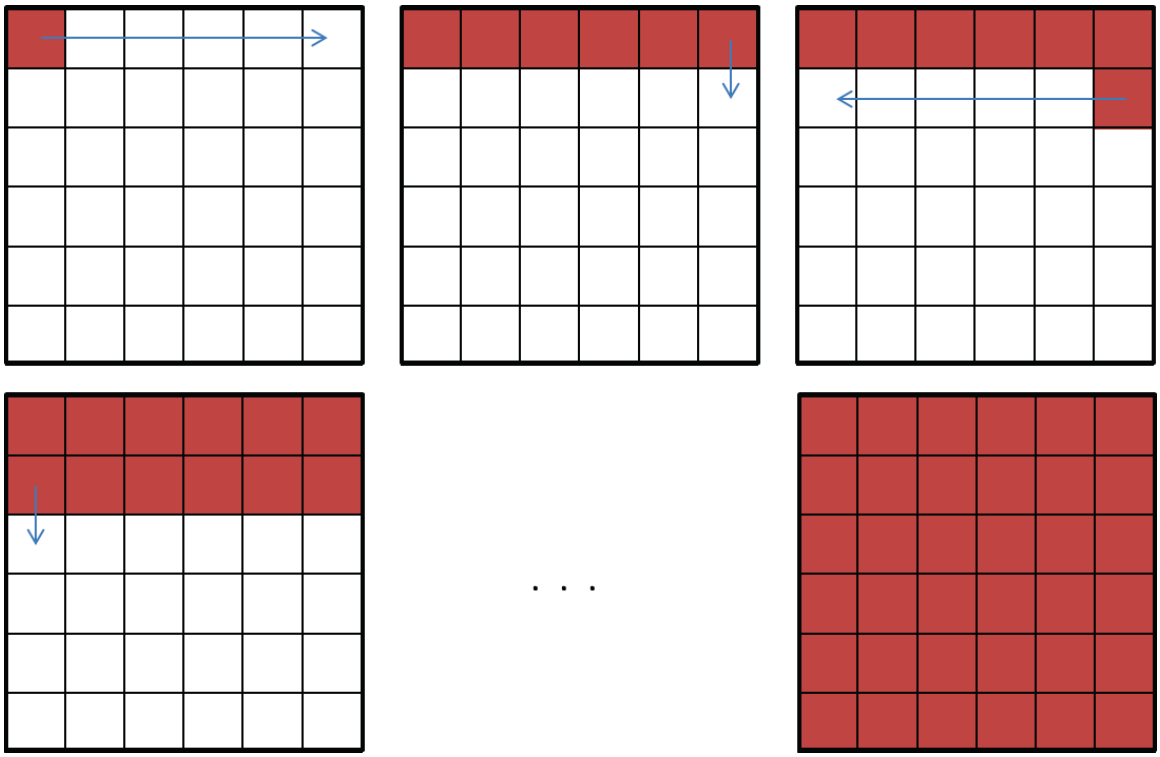

Figure 4.25: Raster method for phase unwrapping.

The main disadvantage of the Raster method is noise, as demonstrated with Figure 4.24. A second disadvantage is that the routine blindly unwraps the entire image, which does not allow for a method to skip or bypass pixels. An entire chapter was devoted to detection of "bad" pixels (chapter 4.2.1), so it is obvious that the next section discusses a method for how to utilize these masks in the phase unwrapping operation.

4.2.3.1.2 FLOOD-FILL METHOD

A flood-fill operation may be recognizable to most people as the bucket tool in a computer graphics editor. The user selects a color to fill an arbitrarily shaped region, and then selects a pixel within that

region, and the application proceeds to fill the area with that color. An example of this operation is shown in Figure 4.26, where the upper circle is filled with the color red. Only the region enclosed by the black lines is filled.

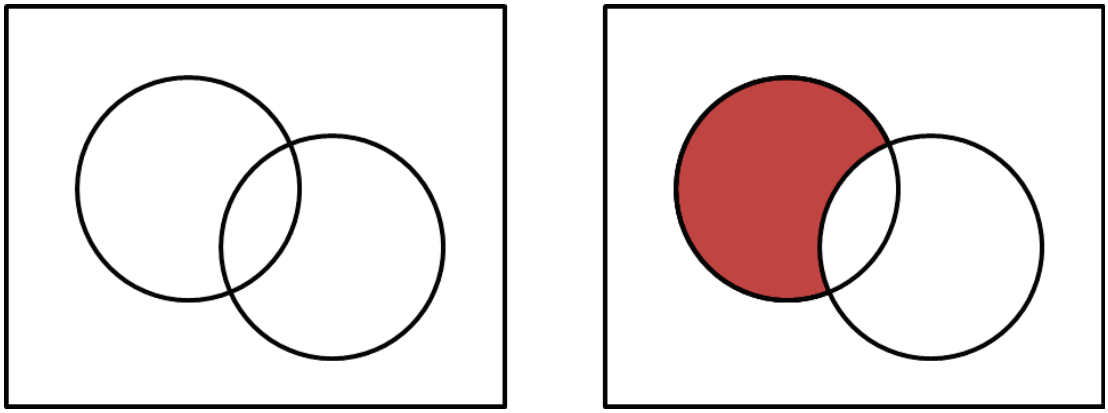

Figure 4.26: Flood-fill operation in a computer graphics program.

The flood-fill phase unwrapping algorithm is an extension of this method. A pixel is chosen as a starting point, which is known as the seed. From this pixel, a queue is built up of neighboring pixels that are available. The queue will typically exclude pixels that have already been unwrapped or that have been denoted as bad by a mask (Section 4.2.1). Once the queue has been updated, the next pixel is processed. In the case of a graphics editor, this pixel will have the value updated by the user requested color. For phase unwrapping, the pixel is unwrapped by Itoh's method using a neighboring pixel. The process appends the new neighborhood of available pixels to the queue, and repeats until every pixel has been unwrapped. An illustration of the flood-fill phase unwrapping method with a masked pixel array is shown in Figure 4.27, where the black pixels represent masked or "bad" pixels. The seed pixel is started near the center, shown in red, and is propagates outwards until the entire image is processed.

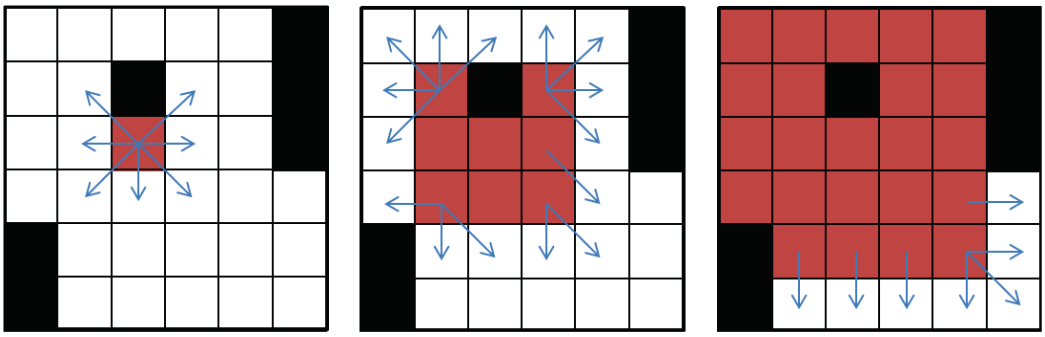

Figure 4.27: Flood-fill example with masks.

The Flood-fill method still has significant drawbacks as a phase unwrapping algorithm for the type of data acquired with the TFI and SPINALTOP systems. Despite the fact that this method has the capability to avoid “bad” pixels, the definition of bad is ambiguous and susceptible to many false-positives. Attempting to mask out every “bad” pixel may result in masking out a significant portion of valid data. This effect can be seen with Figure 4.28, where the flood-fill algorithm is applied to the image on the left with no mask. The image has relatively few defects other than the one large defect that splits the lower third of the image from the rest. The image on the right shows what happens as the mask thresholds are increased until the effects of the discontinuity near the center are minimized. A significant portion of valid data has been thrown away in an attempt to mask out a small portion of bad data. This example demonstrates that a more robust phase unwrapping algorithm is still required.

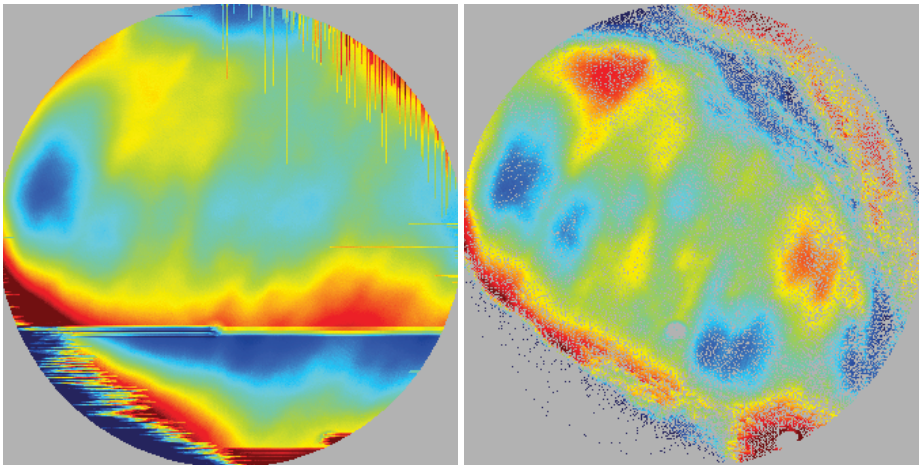

Figure 4.28: Flood-fill phase unwrapping algorithm example.
(Left) No mask. (Right) Mask with thresholds increased until the defect is removed.

4.2.3.1.3 QUALITY-GUIDED METHOD

The next step in the evolution of path-following methods is the quality-guided phase unwrapping algorithm (Bone, 1991; Lim, Xu, & Huang, 1995; Y. Xu & Ai, 2003). The quality-guided method is a flood-fill algorithm that prioritizes the pixel queue through the use of a quality metric. After a seed pixel is selected, a neighborhood queue can be constructed. Each position in the queue now contains an indexed pixel and an associated value that ranks the quality of that pixel. The quality metric that will be used for this process is the same quality map that was described in Section 4.2.2. When one or more pixels are added to the queue, the queue is sorted by quality, with the pixel that contains the highest quality metric placed at the top of the queue. The process repeats until all pixels in the image have been processed.

The quality-guided algorithm is a superior method to the previous discussed algorithms. The flood fill algorithm is only capable of avoiding pixels that are known to be bad. In most cases, it is difficult to determine the bad pixels. The example in Figure 4.28 demonstrates this issue, where a large portion of potentially valid data has been thrown away to accommodate the shortcomings of the phase unwrapping algorithm. The next logical step involves not only avoiding known bad pixels, but to prioritize the path of unwrapping based on the quality of pixels. This relies on the assumption that high-quality pixels (e.g. high signal to noise, good modulation, low-slope variation) will successfully phase unwrap and minimize potential phase unwrapping errors. By processing the highest-quality pixels first and working down to the lower-quality pixels, the algorithm avoids problematic regions and returns to those areas when they are all that remains to be processed. Figure 4.29 illustrates how a quality-guided phase unwrapping algorithm can differ from a standard flood-fill algorithm (Figure 4.27). Figure 4.30 demonstrates the quality-guided method on the same data set as Figure 4.28. Minimal thresholding was required for the masks and therefore most of the data was preserved in the final unwrapping solution.

The advantage of the quality-guided algorithm is that the problematic regions can be unwrapped without throwing away data; even if bad pixels make it through the filters, they will minimally effect the unwrapping operation. However, the quality-guided method is not the final solution. The method is still

dependent on a single pixel to determine the phase unwrapping correction. The quality map is still susceptible to noise and is not perfect. The final evolution of this algorithm is discussed in the next section.

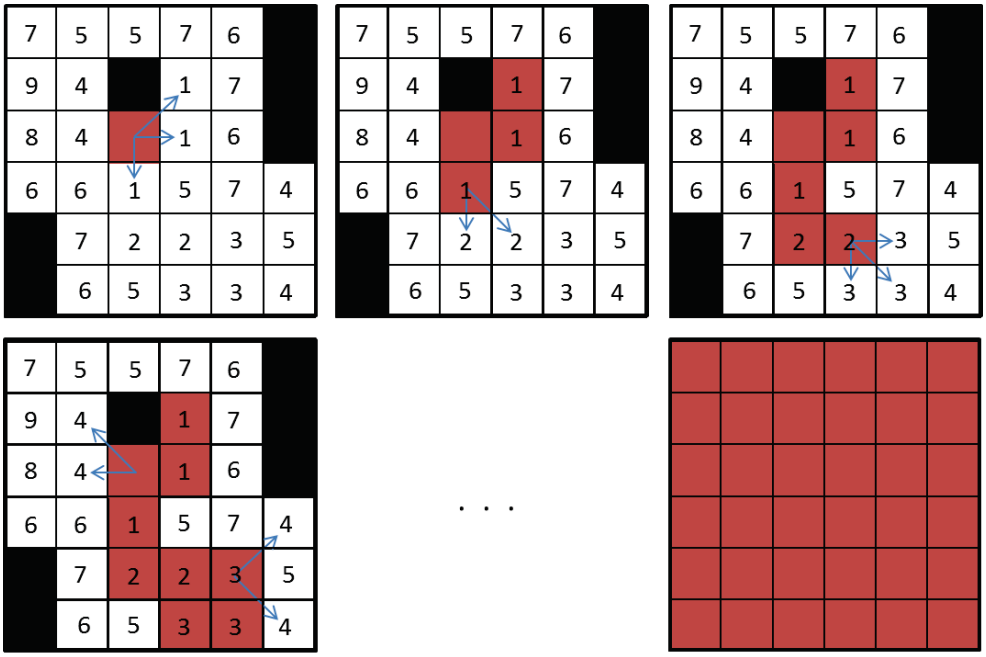

Figure 4.29: Quality-guided method.

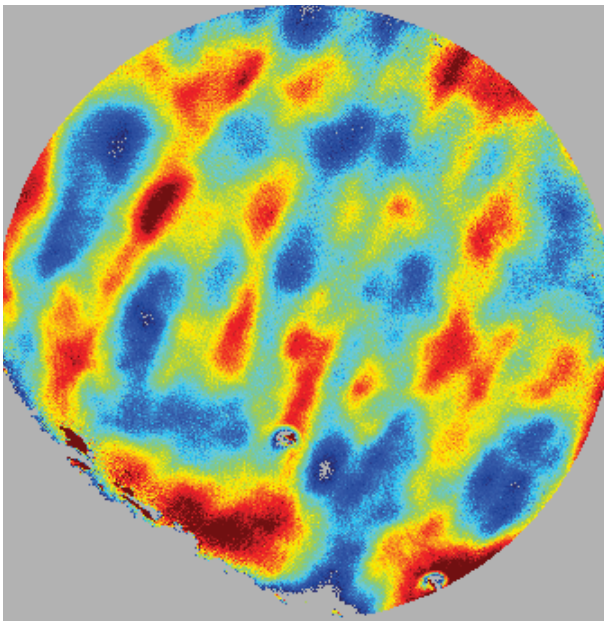

Figure 4.30: Quality-guided phase unwrapping result.

4.2.3.1.4 PATH-FOLLOWING PHASE UNWRAPPING METHOD

Xu and Cumming proposed a method to extend the phase unwrapping process to a local neighborhood to provide a consistency check on the phase correction offset (W. Xu & Cumming, 1999). The path-following methods up to this point have used a single neighboring pixel to derive the unwrapping offset. This final method determines the offset from a neighborhood of pixels to provide a better estimate of the unwrapping offset. The method for phase unwrapping is described in full in this section, as it is the method that is primarily used in processing TFI and SPINALTOP data. The steps are as follows:

1. **Seeding**

   Path-following methods require a seed pixel that is assumed to be "good," from which all neighboring pixels are unwrapped. The phase unwrapping process grows outward from the seed until the entire phase image is unwrapped. Generally, the seed pixel can be placed at the center of the phase image, but is not recommended for tear film measurements where there is no guaranteed good data (e.g. due to a breakup zone or mucin globule that happens to be at the center). Another method to select the seed may be to select the highest rated pixel from the quality map; however, this may occur near a region of bad data or happens to be a noisy result. The method that is used combines both of these methods in an attempt to find a seed pixel that is in a region of good data and close to the center. The process for this is as follows:

   First, the quality map is reweighted using the distance from the center of the phase image. The phase image center is defined as the origin and the edges are ± 1. A seed-quality mask ($\mathcal{S}$) can then be defined as:

$$\mathcal{S}(i,j) = \frac{\mathcal{Q}(i,j)}{1 + r(i,j)} \qquad (4.49)$$

   Where $r(i,j)$ is the radial distance from the center of the phase image for the pixel at (i,j), and $\mathcal{Q}$ is the quality map. The denominator contains an offset of 1 to prevent arbitrarily large values

from appearing at the center and skewing the results. Additionally, the center pixel would be undefined.

Next, a windowed sum is applied to the seed mask to determine the region with the highest quality:

$$\check{S}(i,j) = \frac{1}{N^2} \sum_{m=i-N/2}^{i+N/2} \sum_{n=j-N/2}^{j+N/2} \mathcal{S}(i,j) \qquad (4.50)$$

The seed point (i_s, j_s) is then selected where the windowed mask is at a maximum:

$$\{i_s, j_s \in \check{S} : \check{S}(i,j) = \max \check{S}\} \qquad (4.51)$$

The seed point is now considered the active pixel (i_a, j_a).

Figure 4.31 shows a quality map (left) and the resulting seed mask (right). This particular example has a drop-out at the center of the image, due to a mucin globule in the tear film. The seed mask not only avoids the drop-out, but reduces the weighting in the area surrounding the drop-out. The seed point that is picked for this image lies slightly to the left of the drop-out.

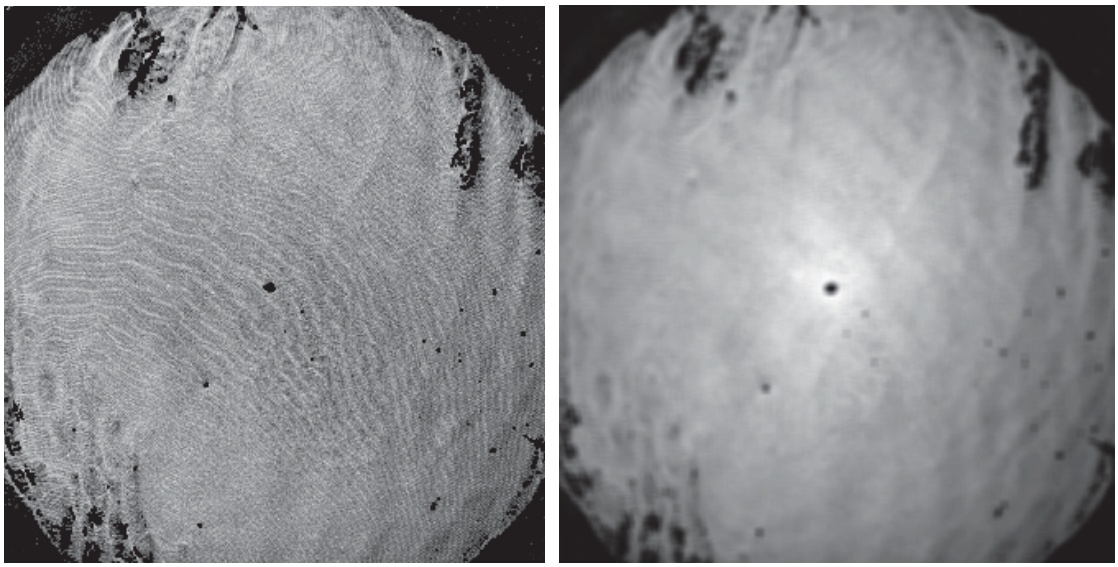

Figure 4.31: Quality map (left) vs. a seed map (right).

2. Update queue

All available pixels in the neighborhood of the active pixel will be added to the queue. An available pixel is:

- not masked (i.e. not bad)
- not been unwrapped
- not already queued
- within phase image boundaries

Only the four immediate pixels are inspected for the queue:

$$(i_a, j_{a-1}), (i_{a-1}, j), (i_a, j_{a+1}), (i_{a+1}, j_a)$$

The quality metric for each queued pixel is also included.

3. Unwrap next pixel

The next pixel to be unwrapped is selected from the queue and is determined by the pixel associated with the largest quality metric. The pixel is removed from the queue and is now the active pixel.

4. Consistency check

An $N\,x\,N$ neighborhood of pixels around the active pixel are checked for availability to support the phase unwrapping operation. An available pixel in this case is a pixel that has been previously unwrapped and is sequentially adjacent to the active pixel. Diagonal pixels can also be used for this operation. A typical value for the neighborhood is a value of N equal to 3 or 5. For noisy data, a value of $N = 5$ is recommended. Figure 4.32 shows an example for a $5\,x\,5$ neighborhood that will be used to demonstrate the process. The active pixel is in green, pixels that have been unwrapped are in red, and white pixels are unprocessed. The blue arrows show the integration path that will be used to determine the phase unwrapping offset value.

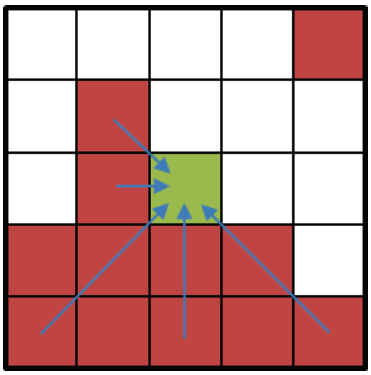

Figure 4.32: 5 x 5 neighborhood consistency check.

The consistency check begins by building up a prediction for the active pixel. The prediction is based on an extrapolation from the neighborhood of available pixels. For a single adjacent pixel at position (i_n, j_n), the i^{th} predictor is the value of the adjacent pixel:

$$\tilde{\phi}_p^i = \tilde{\phi}(i_n, j_n) \qquad (4.52)$$

For more than one adjacent pixel along a continuous line, an extrapolation method must be used. In the case of $N = 5$, the largest available number of adjacent pixels would be 2. Therefore a linear predictor would be used on the adjacent pixels $(i_n, j_n), (i_{\tilde{n}}, j_{\tilde{n}})$:

$$\tilde{\phi}_p^i = 2\tilde{\phi}(i_n, j_n) - (i_{\tilde{n}}, j_{\tilde{n}}) \qquad (4.53)$$

Figure 4.32 shows that two of the neighboring pixels are single predictors. The other three neighbors will require linear predictors. The composite prediction is formed by a weighted sum of all of the predictors:

$$\tilde{\phi}_p = \frac{\sum_i w^i \tilde{\phi}_p^i}{\sum_i w^i} \qquad (4.54)$$

Where the weight w is determined by the number of adjacent pixels used for the prediction, the quality map, or any combination of the two.

5. **Unwrap**

   If the consistency check is used, the predictor is given by the previous step ϕ_p. When the consistency check is skipped, the predictor is the single pixel value used in the normal quality-guided method. The phase correction is determined the same way as it has been previously done, by determining the offset between the phase at the active pixel and the predictor:

$$\Delta = \tilde{\phi}_p - \psi(i_a, j_a) \qquad (4.55)$$

   And the unwrapped phase can be solved for:

$$\tilde{\phi}(i_a, j_a) = \psi(i_a, j_a) + 2\pi \cdot \left\lfloor \frac{\Delta}{2\pi} \right\rfloor \qquad (4.56)$$

6. **Repeat**

   The process repeats step 2 through 5 until the entire phase image is unwrapped. Disjointed regions are not considered as this has been rarely seen to occur in the measurements.

4.2.3.2 MINIMUM-NORM METHODS

The path-following methods described in the previous section were shown to be highly dependent on the quality of data. Through the proper selection of filters, masks, and a robust algorithm, most tear film measurements can be processed. However, the path-following method described in the previous section was found to be limited to tear film measurements of the bare eye. Most contact lens measurements resulted in extremely complex tear film structures that the quality-guided methods could not unwrap without significant phase unwrapping errors. Fringes measured from a subject wearing a contact lens are shown in Figure 4.33 (left) with the resulting unwrapped image using the quality-guided method on the right. No amount of mask thresholding is able to adequately recover the image. The solution was to switch to a minimum-norm phase unwrapping method. This section will present the mathematical

formalism of a minimum-norm method for phase unwrapping and review the method that is used on more complex surface measurements of the tear film.

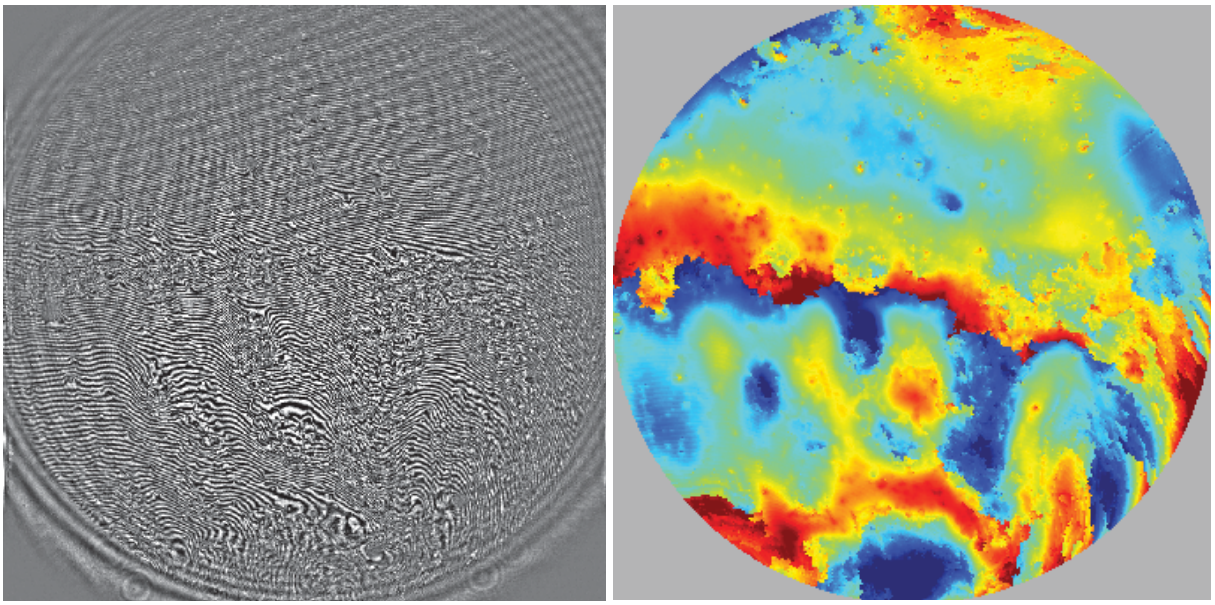

Figure 4.33: Contact lens measurement example.
(Left) Fringes. (Right) Unwrapped phase using the quality guided method.

4.2.3.2.1 BACKGROUND

A minimum-norm method seeks to minimize the L^p-norm error between the gradients of the solution phase (i.e. unwrapped phase) and the wrapped phase. The L^p-norm error for a function f in the space $\mathcal{X}$ is:

$$\|f\|_p = \left(\int_{\mathcal{X}} |f|^p \right)^{1/p} \qquad (4.57)$$

The function f is defined by the difference between the gradients of the unwrapped phase ($\tilde{\phi}$) and wrapped phase (ψ):

$$f = \left|\tilde{\phi}_x - \psi_x\right|^p + \left|\tilde{\phi}_y - \psi_y\right|^p \qquad (4.58)$$

Where the partial derivatives are of the form:

$$\psi_x = \frac{\partial \psi}{\partial x} \tag{4.59}$$

The solution $\tilde{\phi}$ that minimizes the L^p-norm $\|f\|_p$ can be found by the calculus of variations, which shows that $\|f\|_p$will be at a local minima or maxima where $|f|^p$ is stationary. The stationary condition can found by using the Euler-Lagrange equation which results in:

$$\frac{\partial}{\partial x}\left(\frac{\partial f}{\partial \tilde{\phi}_x}\right) + \frac{\partial}{\partial y}\left(\frac{\partial f}{\partial \tilde{\phi}_y}\right) = 0 \tag{4.60}$$

Using the identity:

$$\frac{\partial}{\partial z}|z|^p = \frac{\partial}{\partial z}(z^2)^{p/2} = pz|z|^{p-2} \tag{4.61}$$

the partial derivatives of f (Eq. 4.58) with respect to $\tilde{\phi}_x$ and $\tilde{\phi}_y$ can be rewritten as:

$$\frac{\partial f}{\partial \tilde{\phi}_x} = p\left(\tilde{\phi}_x - \psi_x\right)|\phi_x - \psi_x|^{p-2} \tag{4.62}$$

$$\frac{\partial f}{\partial \tilde{\phi}_y} = p\left(\tilde{\phi}_y - \psi_y\right)\left|\phi_y - \psi_y\right|^{p-2} \tag{4.63}$$

Substituting Eqq. 4.62 and 4.63 into the Euler-Lagrange equation (Eq. 4.60) yields:

$$\frac{\partial}{\partial x}\left(\tilde{\phi}_x - \psi_x\right)|\phi_x - \psi_x|^{p-2} + \frac{\partial}{\partial y}p\left(\tilde{\phi}_y - \psi_y\right)\left|\phi_y - \psi_y\right|^{p-2} = 0 \tag{4.64}$$

Two weighting functions are defined as:

$$W_x = g(x,y) = |\phi_x - \psi_x|^{p-2} \quad (4.65)$$

$$W_y = h(x,y) = \left|\phi_y - \psi_y\right|^{p-2} \quad (4.66)$$

Which reduces Eq. 4.64 in the form:

$$\frac{\partial}{\partial x} W_x\left(\tilde{\phi}_x - \psi_x\right) + \frac{\partial}{\partial y} W_y\left(\tilde{\phi}_y - \psi_y\right) = 0 \quad (4.67)$$

The nonlinear partial differential equation (PDE) in Eq. 4.67 is the most general form for a minimum-norm method to solve for the unwrapped phase. By selecting a value for p and defining weights, the PDE can be simplified into a manageable problem. The next few sections work through a progression of techniques to solve for the unwrapped phase by starting with this equation.

4.2.3.2.2 FFT METHOD

The first minimum-norm phase unwrapping method that is explored is the most efficient phase-unwrapping algorithm, the FFT method. Setting $p = 2$ results in an L^2-norm, which is known as the Euclidean norm or least-squares error. The weighting functions (Eqq. 4.65 and 4.66) become $W_x = W_y = 1$ and the nonlinear PDE reduces to:

$$\tilde{\phi}_{xx} + \tilde{\phi}_{yy} = \psi_{xx} + \psi_{yy} \quad (4.68)$$

Equation 4.68 is now conveniently in the form of Poisson's equation:

$$\nabla^2 \tilde{\phi} = \rho \quad (4.69)$$

where ∇^2 is the Laplacian operator:

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \quad (4.70)$$

And the function ρ is defined as the Laplacian of the known wrapped phase:

$$\rho = \psi_{xx} + \psi_{yy} = \nabla^2 \psi \quad (4.71)$$

An explicit solution to Poisson's equation can be found by expressing the equation in terms of a Fourier transform, where the Fourier transform is defined as:

$$F(k) = \mathcal{F}\{f(x)\} = \frac{1}{\sqrt{2\pi}} \int f(k) e^{ikx} dk \quad (4.72)$$

Using the Fourier transform relation:

$$\mathcal{F}\left\{\frac{d^n}{dx^n} f(x)\right\} = (ik)^n F(k) \quad (4.73)$$

Poisson's equation can be expressed in terms of the Fourier transform pairs. The left side of equation is:

$$\mathcal{F}\{\nabla^2 \tilde{\phi}\} = -k^2 \widetilde{\Phi} \quad (4.74)$$

And the right side of the equation is:

$$\mathcal{F}\{\rho\} = \mathrm{P} \quad (4.75)$$

The full expression for Poisson's equation in terms of the Fourier transform pairs is:

$$-k^2 \widetilde{\Phi} = P \quad (4.76)$$

The term that corresponds to the unwrapped phase can be isolated:

$$\widetilde{\Phi} = -\frac{\mathrm{P}}{k^2} \qquad (4.77)$$

And the unwrapped phase can be recovered by an inverse Fourier transform:

$$\tilde{\phi} = -\frac{1}{\sqrt{2\pi}} \int \frac{\mathrm{P}}{k^2} e^{ikx} dk \qquad (4.78)$$

The solution for the unwrapped phase using Poisson's equation can be easily adapted to the discrete mapping of wrapped phase. The function ρ is calculated by taking the numerical derivative of the wrapped phase ψ and the Fourier transform is accomplished through a Fast-Fourier Transform (FFT). The FFT method for phase unwrapping is effectively an unweighted least squares solution to the unwrapped phase.

The main advantage to this method is that it is extremely quick, requiring a small fraction of the amount of time as the path-following method. Another advantage is that the least-squares solution operates over the entire phase image, and not localized regions, which can significantly minimize the impact of noise. Path-following methods can propagate noise generated artifacts across the entire unwrapped phase image, whereas the least-squares solution will distribute the effects of noise through a local region and maintain continuity of the whole image. The final advantage of this method is that it is incredibly simple to program. The FFT method is applied to the set of data previously shown in Figure 4.33 and the results are shown in Figure 4.34.

However, there are a number of disadvantages to this method that will be corrected for in the following sections. The first issue is that the least-squares solution works on the entire 2-D image. The FFT method does not use masks or weights and is therefore susceptible to noise. Despite the fact that noise is averaged out through the least-squares method, a well-known issue with least-squares fitting is that an outlier can skew the fit. An example is shown in Figure 4.35 with a monotonically increasing set of data that has had

random noise added to it. A single value is marked as an outlier and a fitting is performed. The green line is a fit that includes the outlier and the red line is a fit that excludes the outlier. The same effect will occur with a least-squares phase unwrapping method, where noise and drop-outs are known to exist and will affect the unwrapped data.

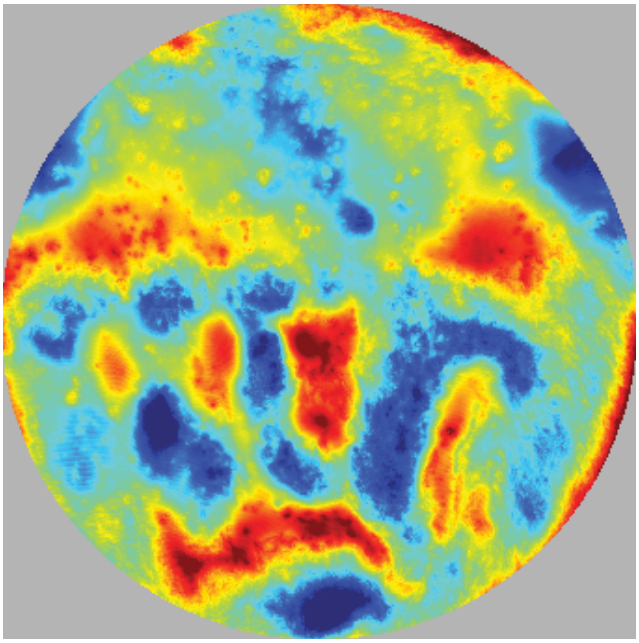

Figure 4.34: Contact lens measurement phase unwrapped with FFT method.

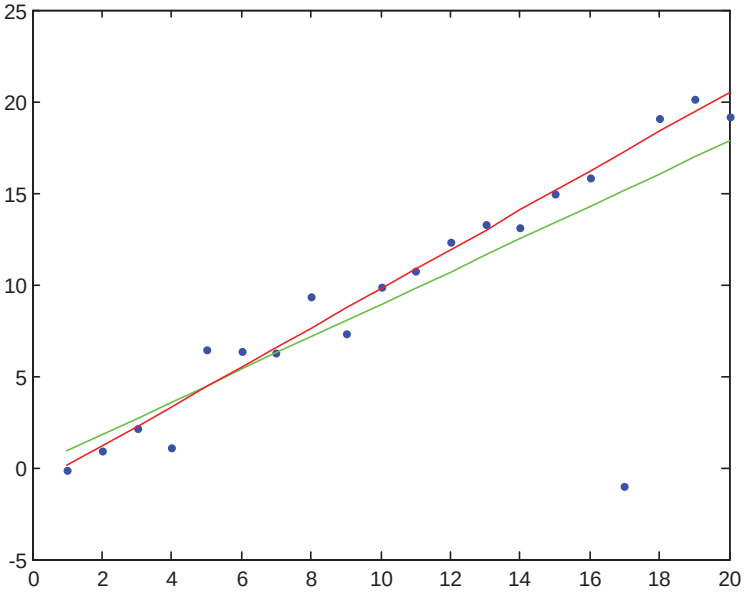

Figure 4.35: Least-squares fitting with an outlier.
Green line is a fit with the outlier. Red line is a fit that ignores the outlier.

A final issue with the FFT method is that the unwrapped phase is not congruent to the wrapped phase: That is to say, that if the unwrapped phase were wrapped, it would not be equal to the original wrapped phase. This is more noticeable with the outlier problem demonstrated with Figure 4.35. With the path-following methods, the changes to the value of phase for a pixel are integer steps of 2π, which maintains the original phase information. Whereas, a minimum-norm method seeks to minimize the error between the gradients of the wrapped and unwrapped phase and does not guarantee that this relationship is maintained. In other words, minimum-norm methods may introduce additional surface deviations that are not representative of the true phase surface. However, a straightforward solution exists. Ghiglia and Pritt defined a congruence operation that seeks to maintain the gradients by comparing the wrapped phase to a re-wrapped function of the unwrapped phase (Ghiglia & Pritt, 1998):

$$\tilde{\tilde{\phi}} = \tilde{\phi} + \mathcal{W}\{\psi - \tilde{\phi}\} \qquad (4.79)$$

When the congruence operation is applied to the FFT data in Figure 4.34, the resulting phase surface is shown in Figure 4.36. The unwrapping defects appear similar to the propagated phase unwrapping defects that show up in path following methods. However, these defects result from the skewed surface fit as a result of noise. One option is to ignore the congruence operation and assume that the best-fit phase unwrapping solution still represents the true phase. Otherwise, a more robust minimum-norm method will be needed.

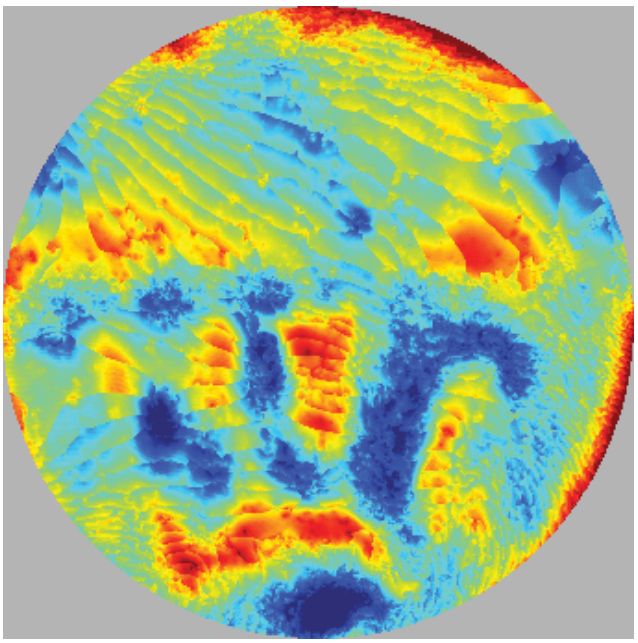

Figure 4.36: Congruence operation applied to unwrapped phase using the FFT-method.

4.2.3.2.3 PICARD ITERATION METHOD

As demonstrated in the previous section (Figure 4.35), a viable application of a minimum-norm method will require a means to mask and/or weight pixels. The FFT method is an unweighted least-squares solution and cannot directly incorporate a weighting matrix. This section discusses a method that uses Picard iteration to solve for the weighted least-squares solution to Poisson's equation (Ghiglia & Romero, 1994). The weighted and unweighted least-squares solutions can be described with linear algebra. The unweighted least-squares solution can be derived from a set of linear equations:

$$\mathbf{Ax} = \mathbf{b} \qquad (4.80)$$

Where **A** is an M x N array, **x** is an N x 1 vector, and **b** is an M x 1 vector. The least-squares solution is derived from the set of normal equations:

$$\mathbf{A}^T\mathbf{Ax} = \mathbf{A}^T\mathbf{b} \qquad (4.81)$$

And the algebraic solution:

$$\mathbf{x} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{b} \quad (4.82)$$

This can be applied to Poisson's equation by first rewriting it in the form:

$$\mathbf{P}\tilde{\boldsymbol{\phi}} = \boldsymbol{\rho} \quad (4.83)$$

Where the matrix $\mathbf{P} = \mathbf{A}^T\mathbf{A}\mathbf{x}$ performs a discrete Laplacian operation on the vector of unwrapped phase $\tilde{\boldsymbol{\phi}}$ and the vector $\boldsymbol{\rho} = \mathbf{A}^T\mathbf{b}$ contains the values of the discrete Laplacian operation on the wrapped phase (Eq. 4.71). The matrix form of Poisson's equation (Eq. 4.83) will have to be related to the set of normal equations (Eq. 4.81) to derive the weighted least-squares solution.

An array of weights $\mathbf{W}$ can be applied to the unweighted linear equation Eq. 4.80, which yields:

$$\mathbf{WAx} = \mathbf{Wb} \quad (4.84)$$

The resulting set of normal equations are given by:

$$\mathbf{A}^T\mathbf{W}^T\mathbf{WAx} = \mathbf{A}^T\mathbf{W}^T\mathbf{Wb} \quad (4.85)$$

The matrix $\mathbf{Q}$ is defined as a weighted discrete Laplacian operation on the vector of unwrapped phases:

$$\mathbf{Q} = \mathbf{A}^T\mathbf{W}^T\mathbf{WA} \quad (4.86)$$

And the vector $\mathbf{c}$ is defined as a container for the modified and weighted set of Laplacian operated wrapped phases:

$$\mathbf{c} = \mathbf{A}^T\mathbf{W}^T\mathbf{Wb} \quad (4.87)$$

The normal equations for the weighted-least squares problem (Eq. 4.85) can be reduced to the form:

$$\mathbf{Q}\tilde{\boldsymbol{\phi}} = \mathbf{c} \qquad (4.88)$$

A solution to the weighted least-squares problem uses a fixed-point iteration method, which is sometimes called a Picard iteration method (Press, 2007). The Picard iteration is first applied by splitting the matrix **Q** into two terms:

$$\mathbf{Q} = \mathbf{P} + \mathbf{D} \qquad (4.89)$$

Where **P** is the Laplacian operation (Eq. 4.83) and **D** is the difference matrix. Equation 4.88 can be rewritten into the form:

$$(\mathbf{P} + \mathbf{D})\tilde{\boldsymbol{\phi}} = \mathbf{c} \qquad (4.90)$$

And re-ordered into the final form:

$$\mathbf{P}\tilde{\boldsymbol{\phi}} = \mathbf{c} - \mathbf{D}\tilde{\boldsymbol{\phi}} \qquad (4.91)$$

The Picard iteration method applied to this set of equations results in an iterative solution:

$$\mathbf{P}\tilde{\boldsymbol{\phi}}_{k+1} = \mathbf{c} - \mathbf{D}\tilde{\boldsymbol{\phi}}_{k} = \boldsymbol{\rho}_{k} \qquad (4.92)$$

Where the index *k* represents the iteration number.

A summary of the Picard method for solving for the unwrapped phase follows:

1. Calculate **c**

   The weighted phase Laplacian (Eq. 4.87) can be written in terms of numerical derivatives and discrete weight values:

$$c(i,j) = W(i,j)\Delta_x(i,j) - W(i-1,j)\Delta_x(i-1,j) + W(i,j)\Delta_y(i,j) - W(i,j-1)\Delta_y(i,j-1) \quad (4.93)$$

   Where W is a matrix of user-defined weights (quality map, section 4.2.2), and the discrete phase derivatives Δ_x,Δ_y which are defined by Eqq. 4.34 and 4.35.

2. Initialize the solution vector

   The iterative method requires an initial guess for the unwrapped phase vector $\widetilde{\boldsymbol{\Phi}}$. The most straightforward solution is to fill the vector with zeros $\widetilde{\boldsymbol{\Phi}}_{k=0} = 0$. However, a better solution is to use the FFT method (section 4.2.3.2.2) to provide an initial guess for $\widetilde{\boldsymbol{\Phi}}_{k=0}$.

3. Solve for $\boldsymbol{\rho}_k$

   The Picard iteration (Eq. 4.92) gives the solution for the current value of $\boldsymbol{\rho}_k$:

$$\boldsymbol{\rho}_k = \mathbf{c} - \mathbf{D}\widetilde{\boldsymbol{\Phi}}_k = \mathbf{c} - (\boldsymbol{Q} - \boldsymbol{P})\widetilde{\boldsymbol{\Phi}}_k \quad (4.94)$$

4. Solve for $\widetilde{\boldsymbol{\Phi}}_{k+1}$

   The current estimate for the unwrapped phase $\left(\widetilde{\boldsymbol{\Phi}}_{k+1}\right)$ can be determined from Eq. 4.92, which lends itself to an unweighted least-squares method. The FFT method (section 4.2.3.2.2) can be used for this step as well.

5. Iterate

   Steps 3 and 4 can be repeated until either the desired number of iterations is reached or the residual errors have converged to a user defined threshold.

6. Congruence

   The final iteration can be processed with the congruence operation (Eq. 4.79) to ensure that the correct unwrapped phase gradients are maintained.

The Picard method is straightforward and easily implemented, however the convergence rate is extremely slow. Figure 4.37 is a comparison of the FFT method (left) to the Picard method (right). The weighted-least squares method provides a phase unwrapping method for difficult sets. However, the Picard method is not the final method that is used for phase unwrapping. The purpose of this section was to provide enough background to discuss a more efficient algorithm for solving the weighted least-square equation.

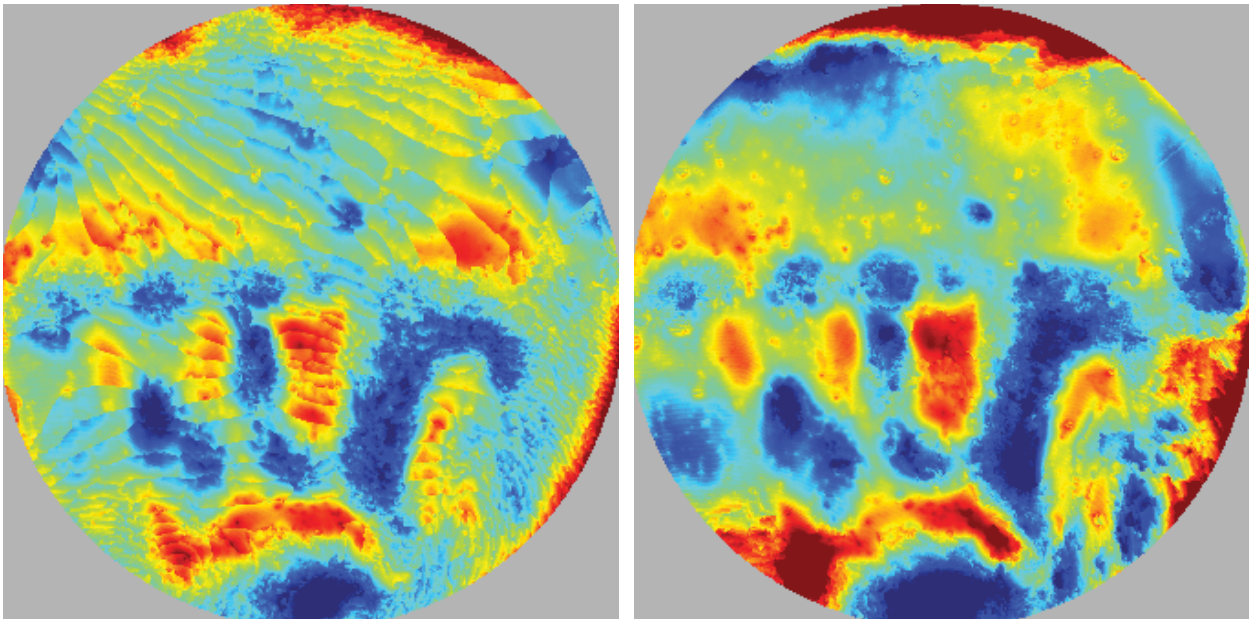

Figure 4.37: Comparison of FFT-method (left) to the Picard iteration method (right). The congruence operation has been applied to both sets of data.

4.2.3.2.4 PRECONDITIONED GRADIENT METHOD

The conjugate gradient method is an efficient algorithm for solving a set of liner equations. The algorithm is similar to the Picard method, but derives the iterative corrections as a sequence of vectors that are conjugate to the residuals. The details of the conjugate method are beyond the scope of this dissertation, but more information can be found in the references (Barrett et al., 1994; Gollub & Van Loan, 1990; Hestenes & Stiefel, 1952).

The conjugate gradient method is guaranteed to converge in N iterations for an $N\,x\,N$ problem, but the rate of convergence can be increased by preconditioning. Preconditioning is effectively a step by which one of the underlying matrices in the process is transformed into a form that is more efficient for the algorithm. A major advantage of the Preconditioned Conjugate Gradient (PCG) method is that it can incorporate the previous FFT method into the solution. The PCG method is described as follows:

1. Initialize the residuals $\mathbf{r}_0$

   The residuals are initialized with the weighted phase Laplacians given by Eq. 4.93, $\mathbf{r}_0 = \mathbf{c}$

2. Initialize the solution vector $\widetilde{\boldsymbol{\phi}}_0$

   This is the same operation that was performed as the Picard method, where the FFT method is used to provide an initial guess for $\widetilde{\boldsymbol{\phi}}_{k=0}$.

3. Preconditioner $\mathbf{z}_k$

   The preconditioner $\mathbf{z}_k$ is given by the equation $\mathbf{P}\mathbf{z}_k = \mathbf{r}_k$, which can be solved by the prior FFT method.

4. Conjugate vector $\mathbf{p}$

   If $k = 0$, the initial conjugate vector is set equal to the preconditioner $\mathbf{p}_1 = \mathbf{z}_0$

Otherwise, the conjugate vector is solved for by its conjugation constraint against the residuals which results in a form that appears similar to a Gram-Schmidt ortho-normalization:

$$\beta_{k+1} = \frac{\mathbf{r}_k^T \mathbf{z}_k}{\mathbf{r}_{k-1}\mathbf{z}_{k-1}} \quad (4.95)$$

Which is applied directly to the conjugate:

$$\mathbf{p}_{k+1} = \mathbf{z}_k + \beta_{k+1}\mathbf{p}_k \quad (4.96)$$

5. Update solution $\widetilde{\boldsymbol{\phi}}_{k+1}$

The next iteration of the unwrapped phase is given by:

$$\alpha_{k+1} = \frac{\mathbf{r}_k^T \mathbf{z}_k}{\mathbf{p}_{k+1}^T \mathbf{Q}\mathbf{p}_{k+1}} \quad (4.97)$$

$$\widetilde{\boldsymbol{\phi}}_{k+1} = \widetilde{\boldsymbol{\phi}}_k + \alpha_{k+1}\mathbf{p}_{k+1} \quad (4.98)$$

6. Update residual $\mathbf{r}_{k+1}$

The residual vector contains the current solution error $\mathbf{r} = \mathbf{c} - \mathbf{Q}\widetilde{\boldsymbol{\phi}}$ and is efficiently solved from the form:

$$\mathbf{r}_{k+1} = \mathbf{r}_k - \alpha_{k+1}\mathbf{Q}\mathbf{p}_{k+1} \quad (4.99)$$

7. Iterate

Steps 3 through 6 are repeated until a maximum number of iterations are reached, or until the norm of the residuals converge to a desired threshold.

8. Congruence

The final iteration can be processed with the congruence operation (Eq. 4.79).

The PCG algorithm is a robust weighted least-squares method for recovering the unwrapped phase from complex and noisy tear film measurements. However, the method does come at the cost of increased processing time and memory usage. The PCG method will require an average of 4X more processing time than the quality-guided method and is therefore limited to use on more complicated sets of data. The PCG method is a more efficient version of the Picard method, so a comparison between the two methods is unnecessary. Instead, Figure 4.38 shows a comparison of the PCG method (right) to the quality-guided method (left) to demonstrate the capabilities of the PCG method. The largest phase unwrapping discontinuity that occurs in the quality-guided method propagates vertically along the center of the image. A number of additional discontinuities exist near the top and bottom-right edges of the data. The PCG method is able to process most of these areas with little failure.

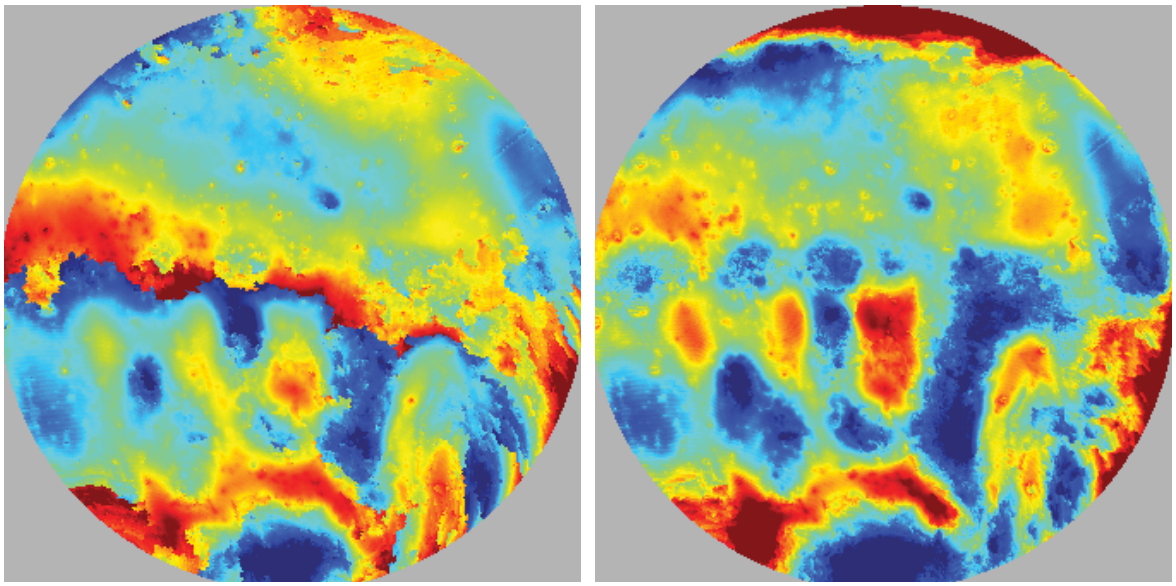

Figure 4.38: Comparison of quality-guided method (left) to the PCG method (right).

4.2.4 INTERPOLATION

Phase unwrapped images will contain missing data, resulting from masked out data (e.g. due to high-noise or high-slope). Path-following methods avoid masked off pixels and so the pixel is never unwrapped. Minimum-norm methods will generate a solution at a masked pixel, which act as a weighted neighborhood interpolant. Therefore, only phase unwrapped images that were processed by a path-

following method will require interpolation. The weighted neighborhood interpolant can be extended to use as a correction for masked pixels. The process is as follows:

1. For each masked pixel in the unwrapped image:
    a. Determine the neighborhood of valid pixels.

       The neighborhood is defined as an $N\ x\ N$ group of pixels centered about the current pixel. A valid pixel is a pixel that is not on the mask list.

    b. Threshold neighborhood

       Interpolation is only performed if a threshold number of neighbors exist, (i.e. valid number of pixels $\geq$ neighborhood threshold). If this criterion is not met, the pixel is not interpolated and the process skips to step e.

    c. Interpolation estimate

       A number of methods can be used to estimate the masked pixel value from the neighborhood of valid pixels. The method that is used has been adapted from the consistency check in the path-following phase unwrapping routine (Section 4.2.3.1.4). A weighted extrapolation value is determined from a sequential set of pixels in-line with the masked pixel (Figure 4.39). The $N\ x\ N$ square neighborhood of pixels can have a maximum of 8 extrapolation values. The average of the extrapolation values is used to fill the masked pixel.

    d. Update mask

       The interpolated pixel is removed from the mask and the process repeats.

    e. Repeat

       Steps a through d are repeated for every pixel on the mask list.

2. Repeat

   After one iteration, the mask may have been updated. In the case that the mask did not change (i.e. no pixels were interpolated) the operation is terminated. Otherwise, Step 1 is repeated for a user defined number of iterations.

The iterative processing within the interpolation routine is designed to minimize the impact on tear-film produced artifacts. Artifacts in the tear film due to break-up or mucin globules generally exceed a few pixels in area (e.g. 10 x 10 or more). If the neighborhood threshold is set high enough, only the bordering pixels on the border of these features will be selected for interpolation. The iterative routine will slowly fill these areas in and should therefore be limited in duration. Actual errors in the unwrapped image are generally small, sometimes only a pixel or two in size and are patched on the first iteration. Moderate sized errors will only be filled in partially due to the neighborhood threshold. However, the iterative method should fix these errors. Therefore, some adjustment is required to determine neighborhood size and threshold as well as number of iterations to clean up the image while at the same time attempting to maintain features. Typical numbers that are used requires a minimum of 3 neighbors and no more than 5 iterations to interpolate out bad data, but still maintain areas of known structure (i.e. tear break-up sites). An example of the interpolation routine is shown in Figure 4.40.

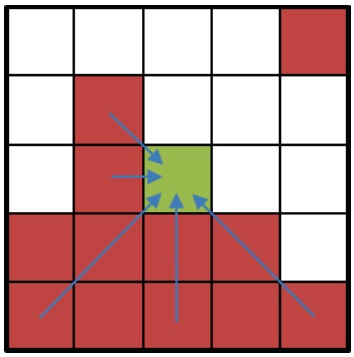

Figure 4.39: 5 x 5 neighborhood of pixels used for interpolation.

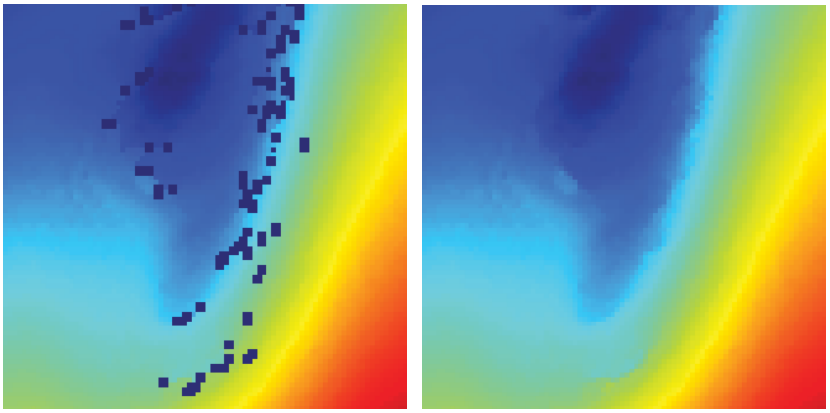

Figure 4.40: Interpolation example.
(Left) Drop-outs in image due to mask (dark blue areas). (Right) Interpolated correction applied.

4.3 CALIBRATION

Instrument calibration involves measuring a known standard to determine the corrections necessary for the instrument to produce the correct measurements. The International Bureau of Weights and Measures defines calibration as an "operation that, under specified conditions, in a first step, establishes a relation between the quantity values with measurement uncertainties provided by measurement standards and corresponding indications with associated measurement uncertainties and, in a second step, uses this information to establish a relation for obtaining a measurement result from an indication" (JCGM, 2008). The accuracy with which the corneal surface and tear film structure can be reconstructed through reverse ray tracing is highly dependent on the interferometer model for SPINALTOP. The calibration procedure determines the state of the interferometer by use of measurement standards. The calibration process is multi-staged process that operates on both of the interferometer sub-systems. The discussion is split between the two interferometer sub-systems in the following sections, although the calibration process is technically a full system calibration.

4.3.1 RAY TRACING

Ray tracing is a significant part of the reconstruction process, including calibration. The calibration methods described in the following sections will require the use of software to simulate optical path lengths through the interferometer model to generate optical path difference maps. Previous studies in

literature have demonstrated that commercial optical design software can be used for reverse ray tracing and reverse optimization (Gappinger & Greivenkamp, 2004a; Lowman & Greivenkamp, 1985). Additional studies have extended commercial software to completely describe an interferometer (B D Stone, 2000; Bryan D. Stone & Thompson, 2009).

Despite previous success using commercial software for interferometer modelling, all of the analysis discussed in this dissertation, including ray tracing, is done in MATLAB (MathWorks, 2014). MATLAB was found to have significant advantages over commercial ray tracing software. Propagation of large ray sets (e.g. millions of rays) could be optimized in MATLAB through vectorized computations, multi-core CPUs, and GPU processing. MATLAB also offers direct integration with image processing and optimization routines.

Although ray tracing is relatively simple (it's all triangles), it is necessary to first verify the ray tracing software. Verification is accomplished by tracing a set of rays in a commercial software program and comparing those set of rays to the same set of rays traced in MATLAB. The commercial software package that is used is Code V (Synopsys, Mountain View, CA). Every lens variation and perturbation that will be used in MATLAB is tested in Code V first. This includes aspheric elements, non-rotationally symmetric surfaces, tilts and decenters, Zernike surfaces, and surface deformation maps. For each change in the model, on- and off-axis and skew rays are chosen to best test the ray tracing software. The outputs from Code V include surface intersection coordinates, direction cosines, and optical path lengths. The same operations are performed in MATLAB and the outputs are compared.

The results are that the MATLAB ray tracing code is in good agreement with Code V, with no error exceeding 1e-10 units. For surface intersections, ray heights, sags, and optical path lengths, this would correspond to 1e-10 mm, or 0.1 pm. The reason for this level of error is due to the round-off error of the display; that is, the Code V output only contained 10 digits after the decimal. One of the lenses used for testing the software is shown in Figure 4.41, Figure 4.42, and Figure 4.43. The lens is an f/2 Double-

Gauss lens used with monochromatic light at 546.1 nm. Figure 4.41 shows a lens drawing from Code V and Figure 4.42 shows the same lens drawn with the MATLAB software. Figure 4.43 shows another test case where one lens has been decentered and another lens has been tilted (left) and another test case where a random Zernike polynomial is mapped onto the surface of a lens (right).

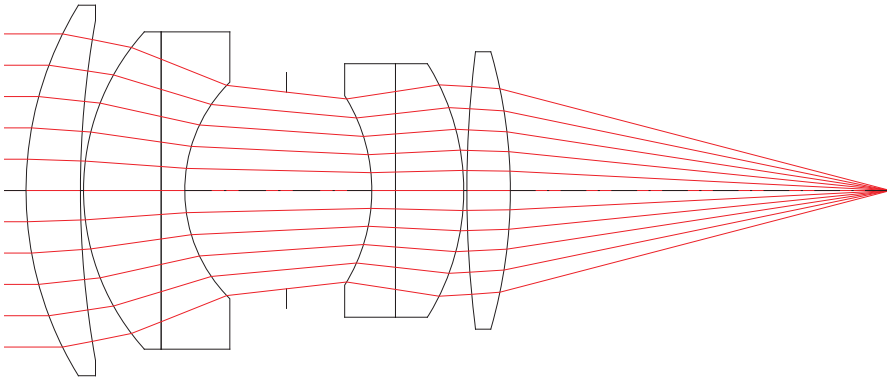

Figure 4.41: Code V lens drawing of an f/2 Double-Gauss lens.

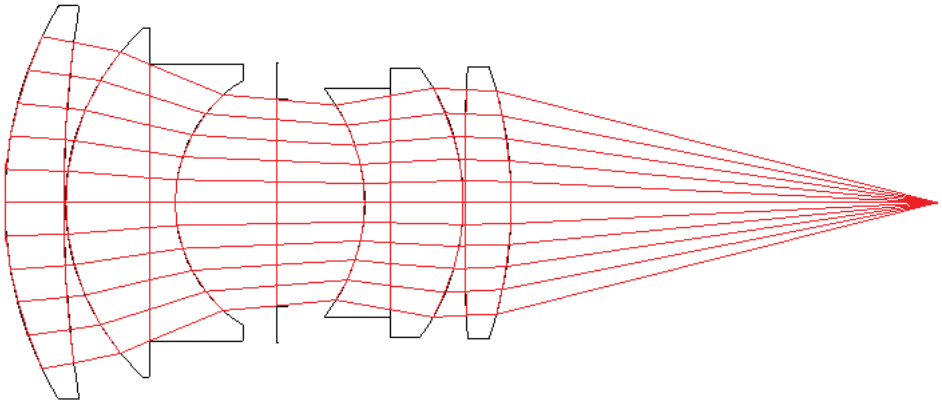

Figure 4.42: MATLAB lens drawing of an f/2 Double-Gauss lens.

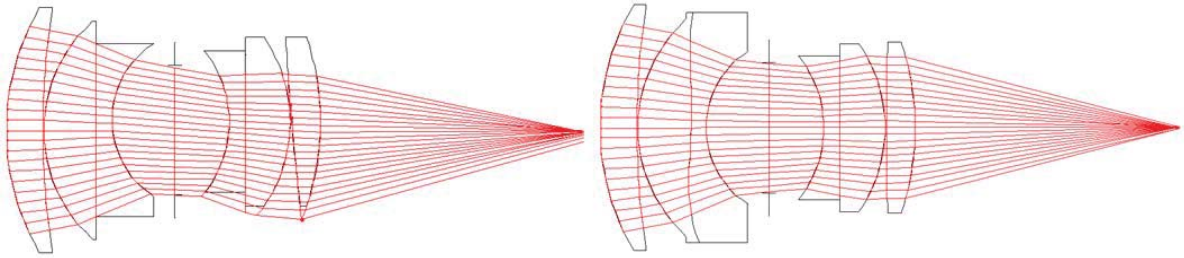

Figure 4.43: Double-Gauss with lens tilts and decenters (left) and Zernike surfaces (right). The lens drawings were done in MATLAB.

4.3.2 SURFACE INTERFEROMETER CALIBRATION

The surface interferometer is the primary instrument by which surface topography is determined. Recovery of surface topography will be achieved through the use of reverse ray tracing, which corrects for retrace errors introduced by the non-null test configuration of the system. The accuracy of the reverse ray trace, and subsequently the accuracy of the reconstructed surface topography, is highly dependent on the accuracy of the interferometer model. Therefore, calibration of the surface interferometer is the process by which the state of the interferometer is determined.

Reverse optimization is the method by which an accurate model of the interferometer can be determined. The concept of reverse optimization as a calibration step has its roots in the alignment of optical assemblies. One of the earliest examples in literature is the alignment of a three mirror off-axis telescope (H J Jeong, Lawrence, & Nahm, 1987). The telescope was built and nominally aligned, followed by a series of end-to-end Hartmann measurements made across multiple field positions. The state of the telescope is determined by comparing the measured data to data simulated by a ray tracing program. An optimization routine perturbs the placement of the optical elements in the system to minimize the difference between the measured and simulated data. The method is referred to as reverse optimization because the optimization routine starts with a perfect model of the system and adjusts the elements in the model to match measured data resulting in an imperfect model.

In the above example, reverse optimization was used to provide feedback to the operators to perform fine adjustments on individual mirror placements to improve system performance. The same methods can be applied to an interferometer, where measurements against known standards can be used to determine the state of the interferometer and the system alignment can then be fine-tuned. However, the alignment sensitives of individual optical elements required for accurate reverse ray tracing may be impossible to achieve. Additionally, not all parameters are adjustable (e.g. refractive index variations, part manufacturing errors, etc…). Reverse optimization is instead used to determine the state of the interferometer to provide an accurate model description of the interferometer for use with reverse ray

tracing (Gappinger & Greivenkamp, 2004a; Lowman & Greivenkamp, 1985). The reverse optimization method that is used to calibrate the surface interferometer will be discussed in this section. The method has been derived from the works of Gappinger et al, Lowman et al, and Hwan et al (2004a; 1985; 1987).

4.3.2.1 COMPONENT CHARACTERIZATION

Before construction of the SPINALTOP system, every critical optical component was characterized, with the exception of the converger assembly. Although it is desired to characterize the converger assembly, it would require a complete disassembly and is considered an unnecessary risk. Unfortunately, the manufacturer of the converger does not have information on any of the five optical elements within the converger. The converger is one of the most critical components in the system and will negatively impact the calibration accuracy due to the uncertainty of the components within it. However, a calibration will be performed with reasonable accuracy to demonstrate the capabilities of SPINALTOP, with knowledge that this is an area that can be improved in the future.

Surface characterization of optical components was performed on a WYKO 6000, a commercial digital phase-shifting Fizeau interferometer (Figure 4.44). A plano magnification target was used to calibrate the plate scale of the interferometer for plano surface measurements (Figure 4.45). The magnification target was custom manufactured by Apollo Optical Systems and is a plano aluminum surface with a series of concentric grooves cut into the face. The magnification target is from a previous research project in the group, but no documentation could be found for the target. However, the dimensions of the magnification target are easily determined. The calibration surface is a 2” diameter reflective aluminum mirror with 10 equally spaced grooves cut into the face. Groove width and spacing between grooves were measured to be 0.04878”. Parity is verified by placing identifying markers in the interferometer path (Figure 4.46). The surface measurements can then be mapped into the interferometer model with proper scale and parity. A typical measurement is shown in Figure 4.47.

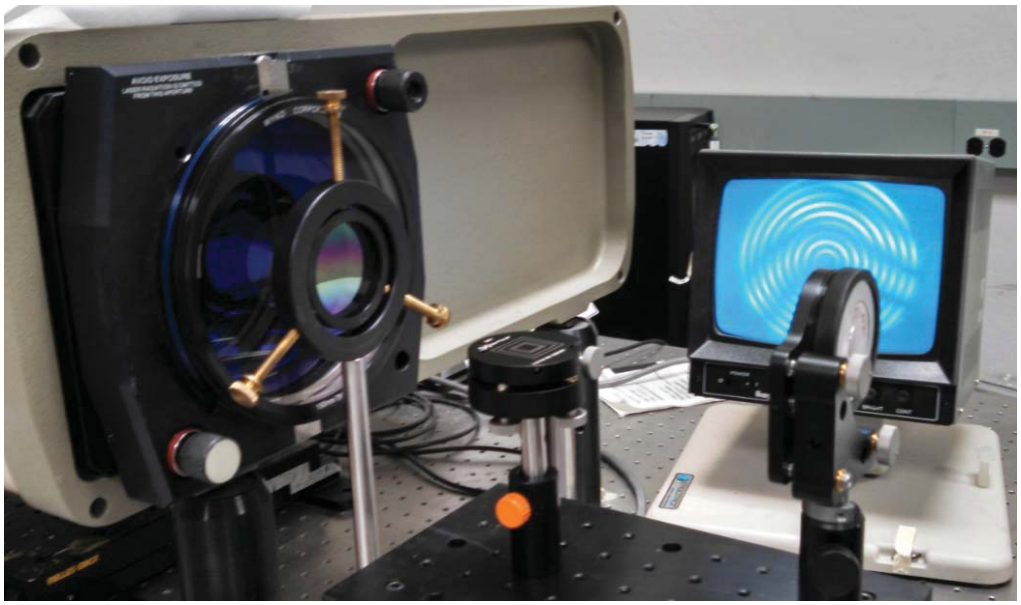

Figure 4.44: Characterization of components on the WYKO 6000.

Figure 4.45: Magnification calibration target.
(Left) Image of the target. (Right) Interference pattern recorded from the face of the target.

Figure 4.46: Verification of parity.
(Left) Temporary fiducials added over target. (Right) Interferogram of target with fiducials.

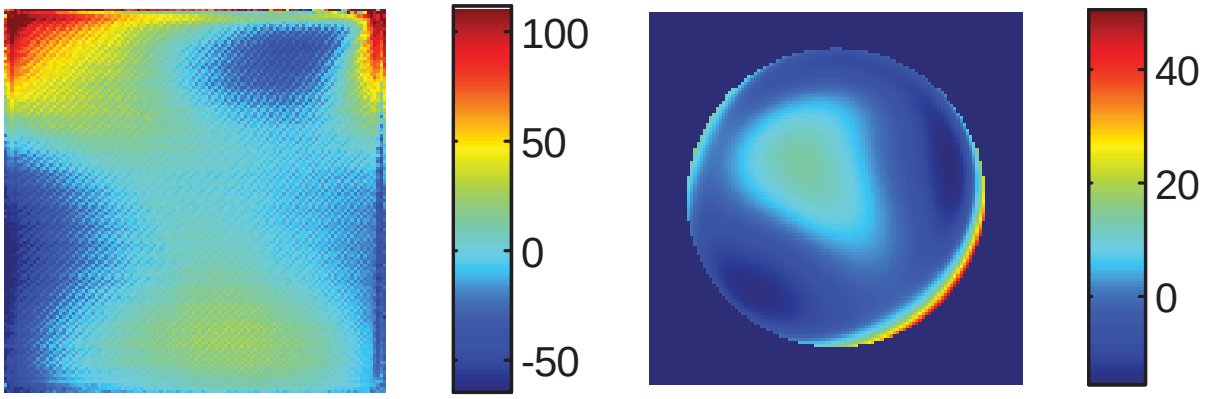

Figure 4.47: Example component measurements.
(Left) One of the faces of the polarizing beamsplitter. (Right) Surface of a dichroic beamsplitter.

Only one curved surface in the interferometer needs to be measured, and that is the non-plano face of the imaging lens. The lenses in the converger assembly would be characterized as well, but they are currently inaccessible. The collimation lens could also be characterized, but adds an unnecessary complication. Instead, the model uses a planar surface for the source and the rays are launched normal to the surface. Small surface perturbations can be added to the surface to correct for collimation errors. Errors from the collimation lens are relatively small in comparison to other errors that will be picked up in the system. This can be seen in the output from a shear plate that is placed after the collimation lens (Figure 4.48).

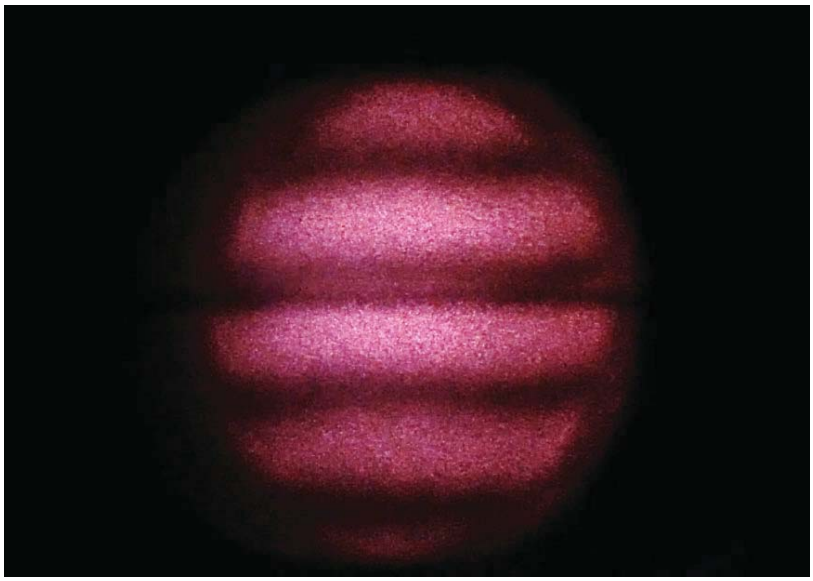

Figure 4.48: Shear plate measurement of the collimated 785 nm source.

4.3.2.2 INITIAL INTERFEROMETER MODEL

An initial model of the interferometer is constructed from the hardware layout (Figure 4.49, left). The layout model provides a list of optical elements and their nominal locations for the interferometer model (Figure 4.49, right). The interferometer elements are updated with direct measurements of components in the system. Non-critical distances are measured with a pair of calipers. The interferometer is constructed and aligned using the breadboard plane and thread-hole matrix to define the optical axis, which nominally places all the elements in orthogonal planes (e.g. all fold mirrors are 45º). Component measurements from the previous section are mapped onto the model surfaces.

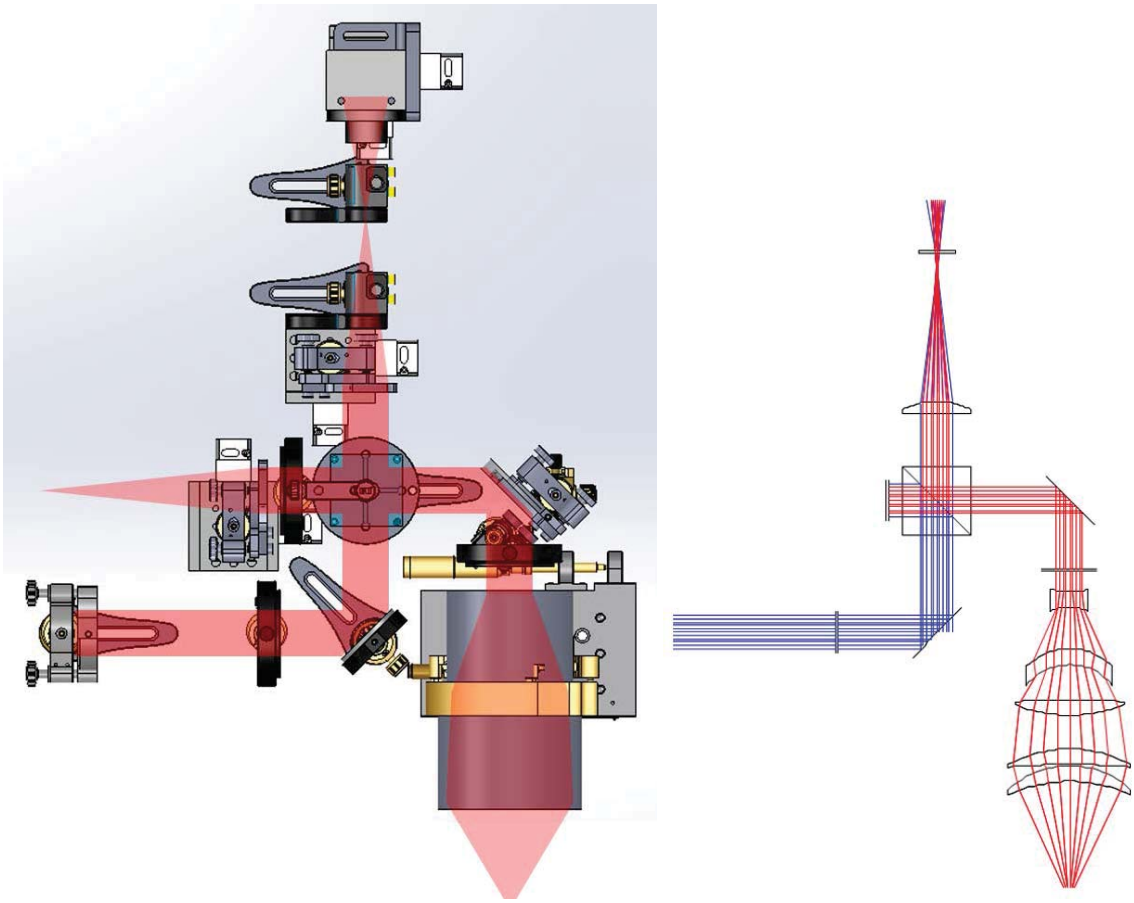

Figure 4.49: Interferometer optical model (right) derived from the hardware layout (left).

4.3.2.3 REFERENCE CHARACTERIZATION

A final prerequisite for the interferometer calibration is verifying the reference surfaces. The first surface is a reference eye model, or the "golden surface" (Figure 4.50, left). The golden surface is an aspheric glass part designed to match the nominal shape of the cornea. The surface has a 7.8 mm radius of curvature and a -0.25 conic constant. The surface was manufactured by QED (Rochester, NY) and has a figure error of $\lambda/20$ PV and $\lambda/100$ RMS (λ =632.8 nm) over a 7 mm clear aperture (Figure 4.50, right). The golden surface was used to align the converger optics assembly and to align and verify both interferometer systems in SPINALTOP. The golden surface is the known standard for a series of measurements on the surface interferometer to provide feedback from the reverse optimization.

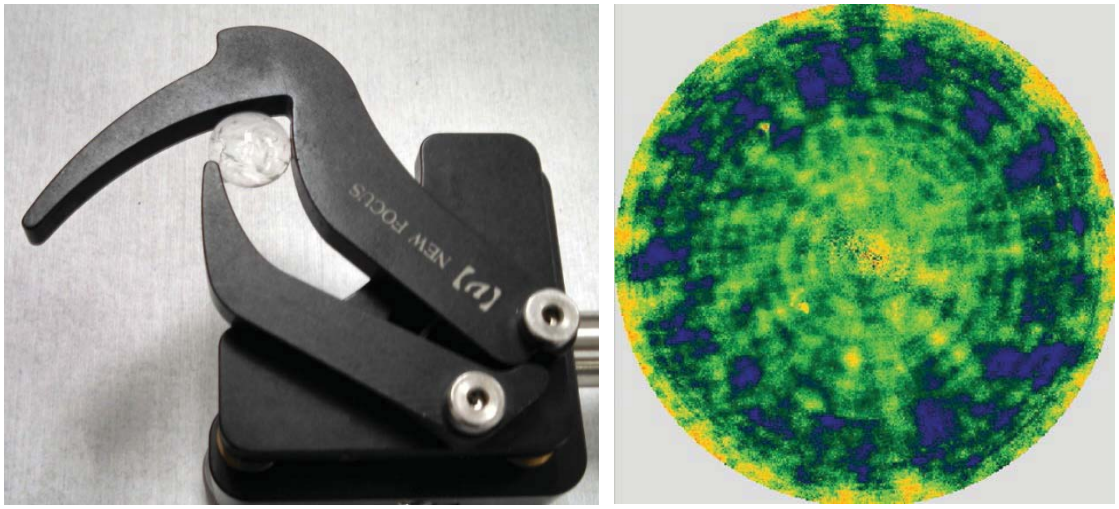

Figure 4.50: Golden surface used for alignment and verification.
(Left) Golden surface in optical mount. Note: The back of the surface has been covered with Vaseline to minimize back surface reflection. (Right) Residual errors as reported by the manufacturer. The display area is approximately 6.5 x 6.5 mm and the color height scale is set to approximately ±0.15 nm

A secondary reference surface is used with the reverse optimization of the surface interferometer. The surface is a distortion target and provides feedback on the alignment of the imaging system. The base figure of the distortion target is the same as the golden surface, a 7.8 mm asphere with a conic constant of -0.25. The aspheric surface is diamond-turned out of brass and was manufactured by the project sponsor. Circular markings approximately 70 μm in diameter are cut into the surface in a 0.5 mm spacing grid of

points (Figure 4.51). Three alignment fiducials are etched into the surface: A crosshair at the center, a single notch 2 mm above the crosshair (at 12 o'clock position), and a double notch 2 mm to the right of the crosshair (at 3 o'clock position). The fiducials are approximately 150 μm in dimension and are used as a centering and parity check for the imaging system.

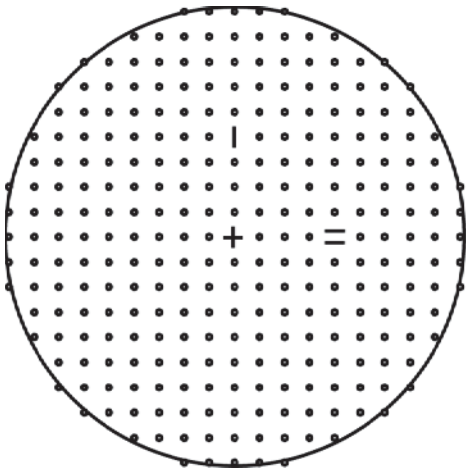

Figure 4.51: Distortion target layout.

The surface figure error of the distortion target is not as critical as knowledge of the alignment targets on the surface. The distortion target is placed under a microscope (Olympus, BH2-UMA) with a 5X infinity corrected microscope objective. The target is mounted to a set of translation stages and adjusted with manual micrometers that have 0.0001" (2.5 μm) resolution (Figure 4.52). A camera system is attached to the microscope to improve alignment feedback and provide additional magnification. The camera software has an option to overlay a crosshair on the image, allowing for repeatable centration on target features (Figure 4.53). The location of alignment features within a 6 mm diameter area on the target were measured and the results are recorded in Table 4.1.

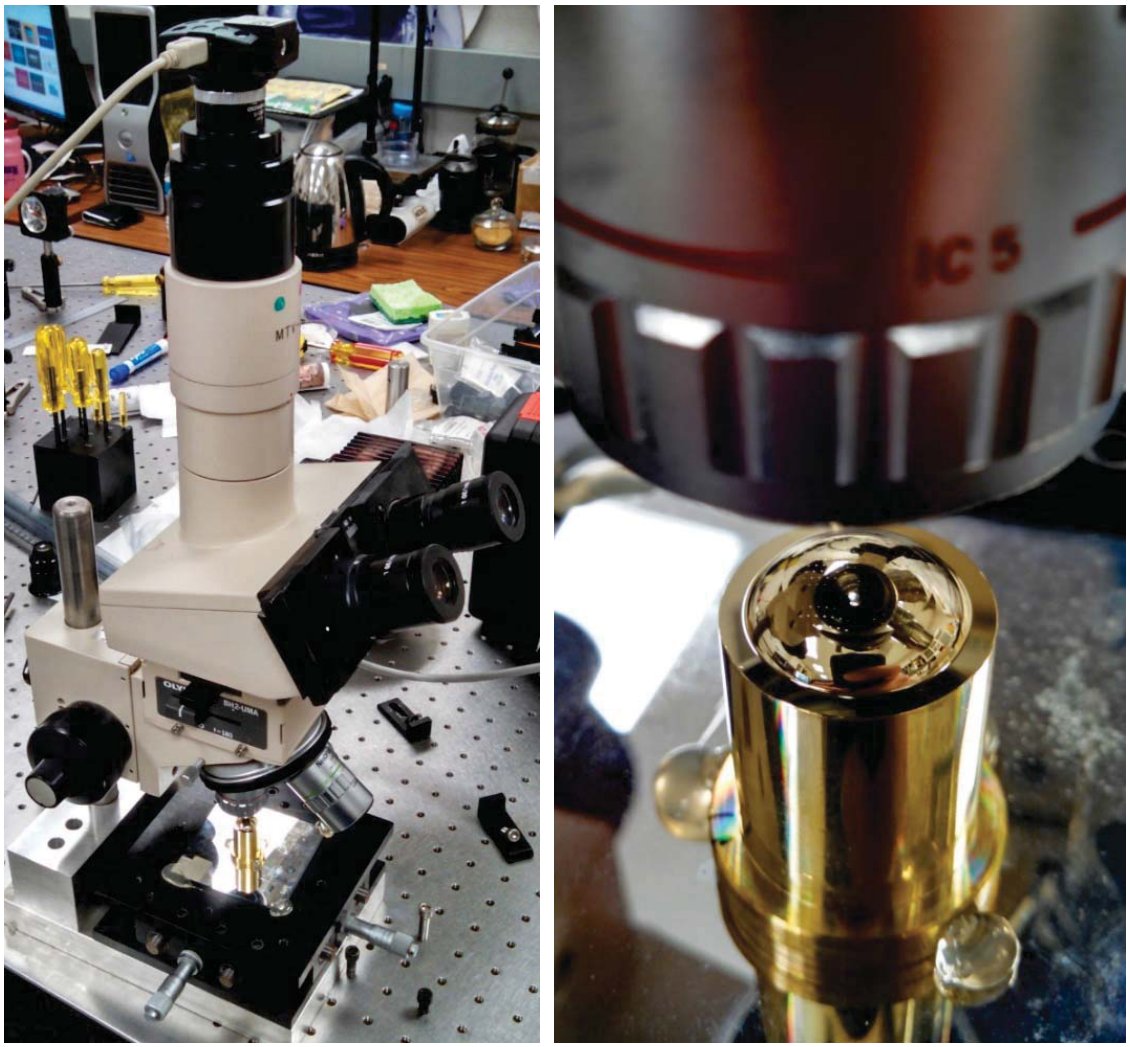

Figure 4.52: Verification of distortion target features under a microscope.

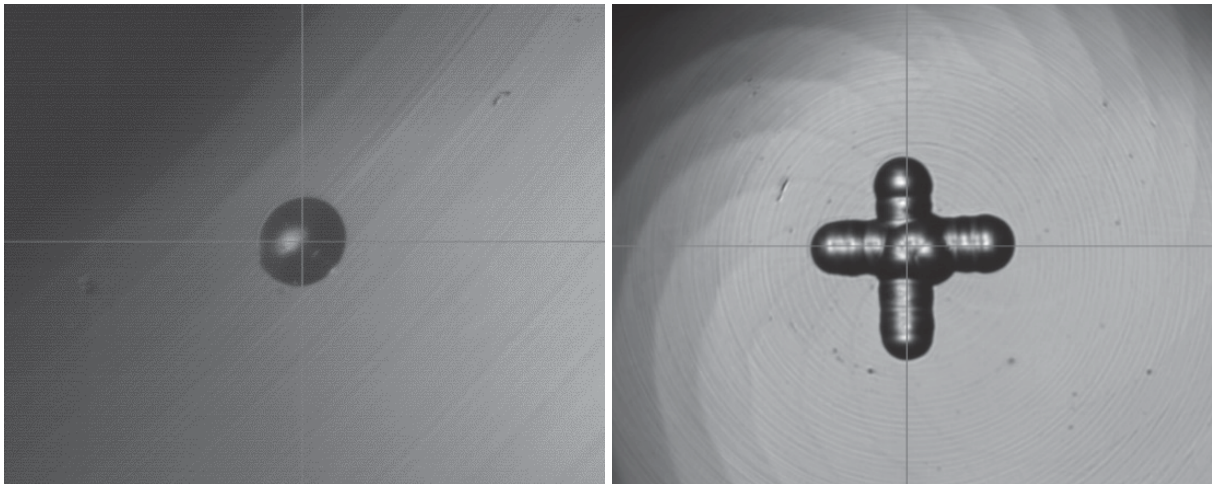

Figure 4.53: Image of features on the distortion target through the microscope.
(Left) One of the circular targets. (Right) Center cross-hair target.

The distortion target is a reflective test surface similar to the golden surface. However, the high reflectivity of the brass surface, compared to the bare glass of the golden surface or tear film on the human eye, will require the relative path intensities of the interferometer to be re-tuned to optimize visibility. Relative path intensities are tuned by adjusting the half-wave plate directly after the collimation lens (Section 3.2.2). The distortion target is aligned (i.e. nulled) and a frame of data is captured (Figure 4.54). A map of the modulation (Section 4.1.1) provides a high contrast image of the distortion targets (Figure 4.55). A software routine was developed to automatically locate the centroid of every target within the central 90% of the clear aperture. The nominal resolution of the interferometer imaging system is 6 μm, however the software routine can determine sub-pixel centroids.

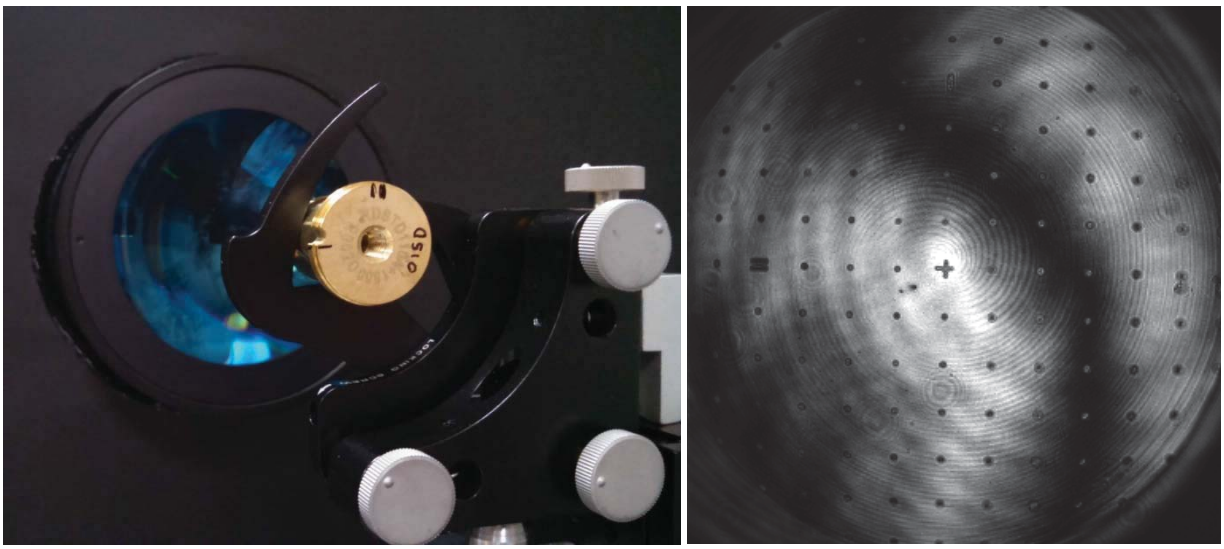

Figure 4.54: Measurement of the distortion target with the surface interferometer.
(Left) The distortion target is clocked and aligned to the interferometer. (Right) Interferogram from the distortion target.

Meas. #	X [inch]	Y [inch]		Meas. #	X [inch]	Y [inch]		Meas. #	X [inch]	Y [inch]
1	0.1480	0.0916		37	0.0647	-0.1218		73	-0.0611	0.1227
2	0.1502	0.0627		38	0.0638	-0.1506		74	-0.0615	0.0932
3	0.1516	0.0329		39	0.0298	0.1520		75	-0.0615	0.0627
4	0.1526	0.0030		40	0.0305	0.1240		76	-0.0614	0.0317
5	0.1528	-0.0271		41	0.0310	0.0946		77	-0.0611	0.0007
6	0.1526	-0.0575		42	0.0315	0.0642		78	-0.0607	-0.0310
7	0.1514	-0.0867		43	0.0319	0.0334		79	-0.0606	-0.0624
8	0.1186	0.1211		44	0.0326	0.0027		80	-0.0602	-0.0938
9	0.1207	0.0931		45	0.0335	-0.0287		81	-0.0595	-0.1243
10	0.1226	0.0637		46	0.0340	-0.0604		82	-0.0584	-0.1533
11	0.1236	0.0335		47	0.0340	-0.0921		83	-0.0892	0.1485
12	0.1240	0.0027		48	0.0340	-0.1227		84	-0.0910	0.1214
13	0.1247	-0.0278		49	0.0336	-0.1523		85	-0.0919	0.0924
14	0.1250	-0.0586		50	-0.0003	0.1519		86	-0.0924	0.0623
15	0.1240	-0.0888		51	0.0000	0.1237		87	-0.0924	0.0312
16	0.1228	-0.1180		52	0.0001	0.0942		88	-0.0923	0.0000
17	0.0883	0.1496		53	0.0004	0.0637		89	-0.0919	-0.0314
18	0.0903	0.1228		54	0.0006	0.0328		90	-0.0919	-0.0629
19	0.0917	0.0940		55	0.0019	0.0014		91	-0.0910	-0.0938
20	0.0929	0.0643		56	0.0026	-0.0295		92	-0.0900	-0.1238
21	0.0937	0.0336		57	0.0027	-0.0611		93	-0.0879	-0.1517
22	0.0945	0.0029		58	0.0028	-0.0926		94	-0.1198	0.1197
23	0.0951	-0.0282		59	0.0028	-0.1236		95	-0.1211	0.0911
24	0.0951	-0.0591		60	0.0030	-0.1532		96	-0.1225	0.0613
25	0.0953	-0.0903		61	-0.0303	0.1516		97	-0.1228	0.0307
26	0.0943	-0.1202		62	-0.0305	0.1235		98	-0.1228	-0.0003
27	0.0929	-0.1483		63	-0.0307	0.0939		99	-0.1229	-0.0317
28	0.0595	0.1513		64	-0.0306	0.0632		100	-0.1221	-0.0625
29	0.0607	0.1237		65	-0.0302	0.0321		101	-0.1211	-0.0932
30	0.0617	0.0945		66	-0.0296	0.0007		102	-0.1191	-0.1221
31	0.0624	0.0642		67	-0.0291	-0.0305		103	-0.1488	0.0890
32	0.0631	0.0337		68	-0.0288	-0.0620		104	-0.1508	0.0602
33	0.0636	0.0027		69	-0.0285	-0.0933		105	-0.1517	0.0300
34	0.0644	-0.0284		70	-0.0283	-0.1243		106	-0.1516	-0.0007
35	0.0649	-0.0599		71	-0.0278	-0.1536		107	-0.1518	-0.0316
36	0.0648	-0.0911		72	-0.0603	0.1506		108	-0.1505	-0.0615
								109	-0.1488	-0.0914

Table 4.1: Measured locations of the distortion target features.

The alignment fiducials help determine the ordering of the measured target locations which can be matched to the measured distortion feature locations. Because the distortion feature locations are relative they will have to be adjusted to account for alignment variations of the target with respect to the interferometer. The adjustments are for rotation and offset only, as these operations preserve the relative spacing. The amount of rotation and offset are recovered through a minimization routine. The routine starts by mapping the distortion target features into detector space using the derived magnification from the measurements between the two spaces. The residual error is determined by the norm of positional errors between the mapped and measured points. The optimization routine then seeks to minimize the residual error by applying an offset and rotation correction to the reference values. The final mapping from the alignment target to the detector can then be used as a fixed variable in the reverse optimization routine. A Levenberg–Marquardt algorithm (LMA) is used to optimize the parameters, which is discussed in the next section.

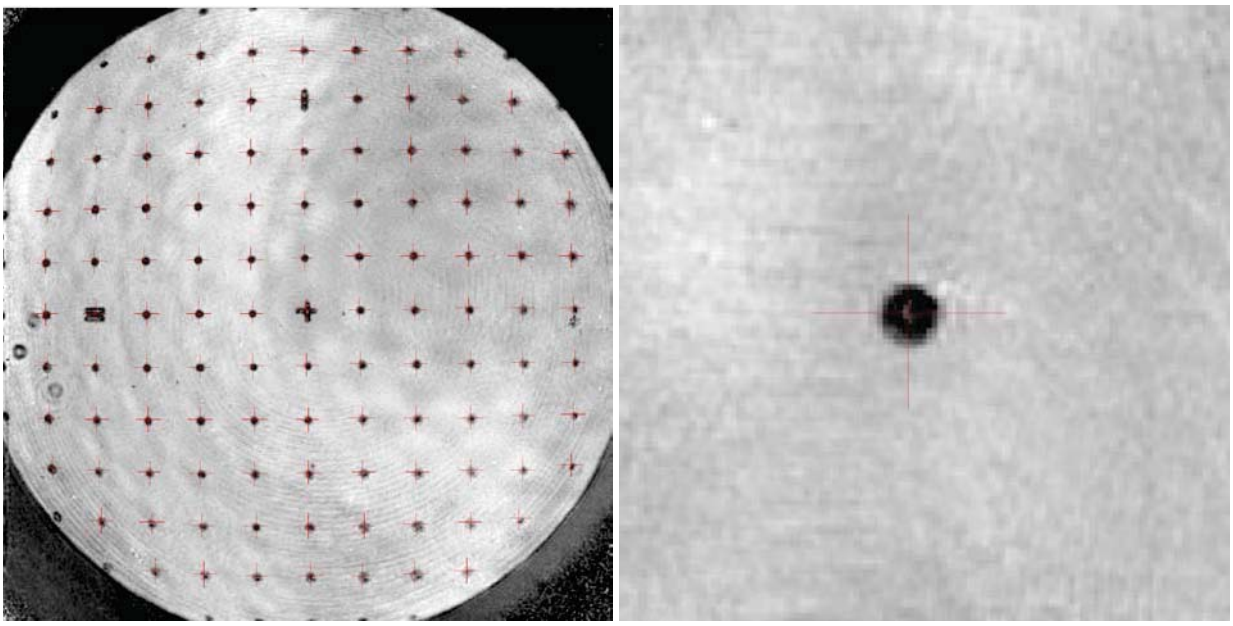

Figure 4.55: Modulation map of targets with centroid locations marked in red.

4.3.2.4 REVERSE OPTIMIZATION

Reverse optimization is the final step in the calibration of the surface interferometer. The optimization routine seeks to minimize the error between real and simulated measurements by adjusting parameters in the interferometer model. The solution for an interferometer model that minimizes the error best represents the state of the interferometer and is therefore suitable for use in the reverse ray trace operations to recover the test surface. An outline of the merit function that the optimization routine seeks to minimize is shown in Figure 4.56. The merit function definition and optimization routines are derived from the works of Gappinger et al, Lowman et al, and Hwan et al (2004a; 1985; 1987). However, a number of significant changes have been made since the prior work that warrants a discussion. This includes a review of the two optimization methods that are use, followed by details on the merit function.

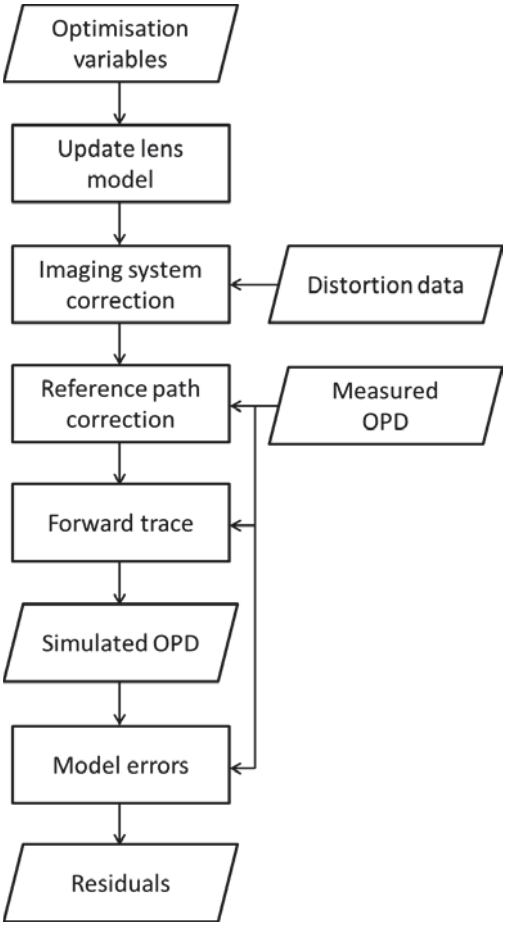

Figure 4.56: Outline of the merit function used for reverse optimization.

4.3.2.4.1 OPTIMIZATION

An optimization problem consists of minimizing or maximizing the output of a function through a systematic determination of inputs. For the case of reverse optimization, the goal is to minimize the error between simulated data generated with the interferometer model and real measured data. The function is then a method to simulate interferometer measurements that provides output in the form of optical path differences (OPD). The inputs to the model set the state of the interferometer model, which includes optical element spacing, element decenter/tilt/spacing/thickness, and lens radii corrections. The residual errors are determined from the differences between the real and simulated optical path difference measurements. An optimization routine seeks to minimize the residual errors by systematically adjusting the input parameters to the interferometer model. The topic of optimization is beyond the scope of this dissertation, so the discussion is limited to a brief review of the two methods that are utilized for reverse optimization.

This first optimization routine is the Levenberg–Marquardt algorithm (LMA), which is also known as the damped least-squares method (Levenberg, 1944; Marquardt, 1963). The LMA is primarily used by a number of sub-routines within the reverse optimization merit function, in addition to a number of other software routines used with SPINALTOP (e.g. distortion target alignment, section 4.3.2.3). LMA is a robust method for solving non-linear least squares problems of the form:

$$\min_{\mathbf{p}} \|f(x;\mathbf{p})\|_2^2 = \min_{\mathbf{p}} \sum_i f_i^2(x;\mathbf{p}) \qquad (4.100)$$

Where the function $f(x;\mathbf{p})$ is generally called an objective function, or in optical design and optimization it is referred to as a merit function. The optimization parameters are captured by $\mathbf{p}$ and independent variables x. The merit function that will be used by the LMA routine is defined as:

$$f_i(x;\mathbf{p}) = \frac{1}{w_i}\left[\phi_i(x) - \tilde{\phi}_i(x;\mathbf{p})\right] \qquad (4.101)$$

where w_i is a weighting factor applied to the difference between ϕ and $\tilde{\phi}$, the target and model functions, respectively. For most applications of the LMA that are used within this project, the target function ϕ will be the measured OPDs and the model function $\tilde{\phi}$ will be the simulated OPDs. Therefore, the discussion will be in terms of OPDs, but can be easily extended to any other problem (e.g. distortion target coordinates in place of OPD, Section 4.3.2.3). The simulated OPDs are a function of model parameters $\mathbf{p}$, which are the corrections to optical elements in the interferometer model. The indexing variable i indicates that the error function is made up of multiple observations, or measurements.

The sum of squares term in Eq. 4.50 can be described in matrix form as a function of the model parameters:

$$\begin{aligned}\chi^2(\mathbf{p}) = \sum_i f_i^2(x;\mathbf{p}) &= \left(\boldsymbol{\phi} - \widehat{\boldsymbol{\phi}}(\mathbf{p})\right)^T \mathbf{W}\left(\boldsymbol{\phi} - \widehat{\boldsymbol{\phi}}(\mathbf{p})\right) \\ &= \boldsymbol{\phi}^T\mathbf{W}\boldsymbol{\phi} - 2\mathbf{y}^T\mathbf{W}\tilde{\boldsymbol{\phi}} + \tilde{\boldsymbol{\phi}}^T\mathbf{W}\tilde{\boldsymbol{\phi}}\end{aligned} \qquad (4.102)$$

where $\mathbf{W}$ is the weighting matrix, a diagonal matrix with $W_{ii} = 1/w_i^2$. The minimum of the optimization problem must occur at a critical point of the χ^2 function, or where the following is true:

$$\frac{\partial}{\partial \mathbf{p}}\chi^2(\mathbf{p}) = 0 \qquad (4.103)$$

There are numerous methods to solve for Eq. 4.103 where an analytic or straightforward solution does not exist. LMA is one of these methods, which is an adaption of two other commonly used methods: The Gradient descent and Gauss-Newton method.

The Gradient descent method is also known as the method of steepest descent. The algorithm works by first estimating the function gradient for the current parameter estimate. The negative gradient determines the direction of steepest descent, or where the function should minimize the quickest if the parameters are updated along this direction. The gradient of the sum of squares term is:

$$\frac{\partial}{\partial \mathbf{p}}\chi^2(\mathbf{p}) = -\left(\boldsymbol{\phi} - \hat{\boldsymbol{\phi}}(\mathbf{p})\right)^{\mathrm{T}} \mathbf{W}\left(\frac{\partial}{\partial \mathbf{p}}\hat{\boldsymbol{\phi}}(\mathbf{p})\right) = \left(\boldsymbol{\phi} - \hat{\boldsymbol{\phi}}(\mathbf{p})\right)^{\mathrm{T}} \mathbf{WJ} \qquad (4.104)$$

Where **J** is the Jacobian matrix, which describes the local sensitivity of the function $\hat{\boldsymbol{\phi}}$ to variations in the parameter **p**. The individual elements of the Jacobian matrix are given by:

$$J_{ij} = \frac{\partial \tilde{\phi}_i(x, p_j)}{\partial p_j} \qquad (4.105)$$

The negative of the gradient, or steepest descent, provides the direction of the optimization step:

$$\mathbf{h} = \gamma \mathbf{J}^{\mathrm{T}} \mathbf{W}\left(\boldsymbol{\phi} - \hat{\boldsymbol{\phi}}(\mathbf{p})\right) \qquad (4.106)$$

where **h** is generally referred to as the perturbation value, and the value γ determines the size of the step used for the steepest descent. A number of methods exist to determine the step size, but is currently unnecessary as this is just the precursor to the LMA.

The other method that comprises the LMA is the Gauss-Newton method, which is a modification of Newton's method. The Gauss-Newton method starts with a first order Taylor series expansion on the perturbed model $\left(\widehat{\boldsymbol{\phi}}\right)$ to estimate the change for a given perturbation ($\mathbf{h}$):

$$\widehat{\boldsymbol{\phi}}(\mathbf{p}+\mathbf{h}) \approx \widehat{\boldsymbol{\phi}}(\mathbf{p}) + \left(\frac{\partial}{\partial \mathbf{p}}\widehat{\boldsymbol{\phi}}(\mathbf{p})\right)\mathbf{h} = \widehat{\boldsymbol{\phi}}(\mathbf{p}) + \mathbf{J}\mathbf{h} \qquad (4.107)$$

Substitution of the approximation back into the sum of squares (Eq. 4.102) yields:

$$\chi^2(\mathbf{p}+\mathbf{h}) \approx \boldsymbol{\phi}^{\mathrm{T}}\mathbf{W}\boldsymbol{\phi} + \widehat{\boldsymbol{\phi}}^{\mathrm{T}}\mathbf{W}\widehat{\boldsymbol{\phi}} - 2\boldsymbol{\phi}^{\mathrm{T}}\mathbf{W}\widehat{\boldsymbol{\phi}} - 2\left(\boldsymbol{\phi}-\widehat{\boldsymbol{\phi}}\right)^{\mathrm{T}}\mathbf{W}\mathbf{J}\mathbf{h} + \mathbf{h}^{\mathrm{T}}\mathbf{J}^{\mathrm{T}}\mathbf{W}\mathbf{J}\mathbf{h} \qquad (4.108)$$

The goal of the Gauss-Newton method is to minimize the sum of squares error through the perturbation variable:

$$\frac{\partial}{\partial \mathbf{h}}\chi^2(\mathbf{p}+\mathbf{h}) \approx -2\left(\boldsymbol{\phi}-\widehat{\boldsymbol{\phi}}\right)^{\mathrm{T}}\mathbf{W}\mathbf{J} + 2\mathbf{h}^{\mathrm{T}}\mathbf{J}^{\mathrm{T}}\mathbf{W}\mathbf{J} \qquad (4.109)$$

Solving for the minimum (Eq. 4.103) results in the set of normal equations:

$$\left(\mathbf{J}^{\mathrm{T}}\mathbf{W}\mathbf{J}\right)\mathbf{h} = \mathbf{J}^{\mathrm{T}}\mathbf{W}\left(\boldsymbol{\phi}-\widehat{\boldsymbol{\phi}}\right) \qquad (4.110)$$

The perturbation value can be solved resulting in a form similar to Newton's method.

The set of normal equations that describe the Steepest descent (Eq. 4.106) and the Gauss-Newton (Eq. 4.110) are combined to form a set of normal equations that describes the Levenberg-Marquardt method:

$$\left(\mathbf{J}^{\mathrm{T}}\mathbf{W}\mathbf{J} + \lambda\,\mathrm{diag}\left(\mathbf{J}^{\mathrm{T}}\mathbf{W}\mathbf{J}\right)\right)\mathbf{h} = \mathbf{J}^{\mathrm{T}}\mathbf{W}\left(\boldsymbol{\phi}-\widehat{\boldsymbol{\phi}}\right) \qquad (4.111)$$

Where λ is a dampening vector applied to the diagonals of $\mathbf{J}^{\mathrm{T}}\mathbf{W}\mathbf{J}$. Large values of λ will result in a Steepest descent method update to the perturbation value $\mathbf{h}$, and smaller values of λ will result in a Gauss-Newton update. The value for λ is typically initialized to a large value and decreases as the function

approaches a minimum. Although a number of methods exist to update λ, the method described by Marquardt is sufficient for the uses described in this dissertation (Marquardt, 1963). The solution to the set of normal equations can be found by any number of means (Press, 2007).

The LMA is a fast and robust method to solve for a local minimum. The merit function has a number of sub-routines that uses the LMA that will be discussed in the next section. The starting solutions for the sub-functions are typically close to the global minimum and will quickly converge on the solution. However, LMA does not work on the main merit function, where it tends to get trapped near the initial solution and fails to solve for a global minimum. It is most likely due to the fact that the main merit function is too noisy and non-linear for a proper gradient to be determined.

The reverse optimization merit function requires a global search method that does not depend on gradient information. The method that is used is the pattern search method (PSM) (Hooke & Jeeves, 1961). The PSM is from a class of optimization routines known as direct-search methods. In basic terms, the PSM searches a neighborhood of parameters for a minimum value. If a minimum value is found, the parameter is updated to the new location and the process repeats. If a minimum value is not found in the neighborhood, the neighborhood size is decreased and the entire process repeats until convergence on a minimum. The routine is extremely simple at the cost of being extremely inefficient. For the merit function that will be used, the PSM can take >100 times longer than the LMA method. This is not an issue, as the reverse optimization procedure only needs to be run once per calibration cycle and can be left to run overnight on a computer.

4.3.2.4.2 MERIT FUNCTION

The optimization methods described in the previous section seek to minimize the sum of squares of a function's output by varying a set of input parameters. This section will discuss the parameters, the function, and the output. A flow chart showing the various stages of the merit function is shown in Figure 4.56. The first step is to define the model parameters, or optimization variables. These are the set of

parameters that describe the state of the interferometer: Lens element decenter, tilts, spacing, thickness, and surface curvature. It is unnecessary to apply variables to every surface in the interferometer. Too many variables will slow the optimization process down and decrease the likelihood of a solution being found.

A sensitivity analysis of the interferometer model can be used to determine which variables should be used for the reverse optimization. The analysis can determine which optical elements are the most sensitive to differences between the interferometer model and the real system. Differences between the model and interferometer will result in an inaccurate reconstruction of the test surface through the reverse ray trace operation (an example is shown in Section 4.4.4.3). A sensitivity analysis can be performed as follows. First, one of the parameters in the interferometer model is perturbed by a small amount. A forward ray trace with the perturbed model produces a set of optical path differences. The optical path differences are then fed into the reverse ray trace operation to reconstruct the test surface (Section 4.4). The difference between the simulated and reconstructed test surface determines the parameters sensitivity. Repeating this operation for every optical element in the interferometer can determine which parameters are the most sensitive to errors and would benefit the most from an optimization.

There is one exception to the list of optimized parameters, which is the imaging system. The optical path differences that are used as feedback for the optimization routines are insufficient to fully characterize the imaging system. In other words, the OPD residuals cannot be used to accurately determine the focus of the imaging system. The imaging system parameters are better determined through secondary information, such as tracking the pupil. However, this requires a detailed knowledge of the limiting apertures inside the system, which are mostly unknown due to the uncharacterized converger. Instead, the distortion target (Section 4.3.2.3) is used to determine the imaging system parameters. The LMA method is extremely efficient in determining these parameters; therefore, the corrections are done in a sub-routine within the main merit function to minimize overhead on the less efficient PSM.

In a step that is similar to the imaging system correction, the reference path can be corrected by a sub-routine optimization. The only parameter in the reference path that needs to be optimized is the tip and tilt of the reference mirror surface. The error can be quickly determined by analyzing the cat's eye position measurement of the golden surface. At a cat's eye position, the retro-reflected test wavefront is inverted. The inversion in a double-pass configuration will cancel out the odd-order aberrations that are generated in the test path. Therefore, residual odd-order aberrations are mostly a result of errors in the reference path, the most noticeable of which is tilt. The tip and tilt error introduced by the reference mirror can be quickly fixed through an LMA routine.

The next step is to generate the set of simulated optical path differences with the current iteration of the interferometer model. The simulation is performed by adjusting the location of the golden surface in the model to match the location that was used for capturing the real OPD measurements, and then performing a forward ray trace to generate the simulated OPD measurements. However, this requires accurate location knowledge of the golden surface during the actual measurements. Although this information could have been determined through proper characterization of the test setup and the interferometer optics, it relies on the converger assembly which is not characterized. Instead, stage location coordinates are recorded during adjustment of the golden surface, which provides relative positioning of the golden surface in a separate coordinate frame. What remains to be determined is a transformation of the golden surface coordinate frame to the interferometer coordinate frame.

The coordinate frame transformation between the golden surface and interferometer is done through a third sub-routine within the merit function. Six parameters are required for the coordinate transformation (translation and rotation) which are determined by minimizing the differences between the measured and simulated OPDs. The LMA method is used for this sub-routine optimization. At first glance, this may appear as though the reverse optimization routine is being bypassed by this sub-routine. The main concern being that interferometer model errors could be masked, or compensated, by adjustments of the golden

surface. However, a sufficient number of perturbations of the golden surface are measured to minimize the occurrence of this happening.

The final step in the merit function is to determine the residuals: the errors between the measured and simulated OPDs. The coordinate frame transformation from the previous step can provide the residuals. Despite the fact that the sub-routine optimizes the alignment of the golden surfaces to minimize OPD residuals, the residuals will still be non-zero for an imperfect interferometer model. The residuals from the sub-routine are then forwarded to the output of the merit function. Additionally, the quality map (section 4.2.2) can be used as the weighting factor w_i (Eq. 4.101). The final residual error is determined and passed to the PSM optimization routine. The results of the calibration for the surface interferometer are presented in the next chapter.

4.3.3 CAT'S EYE INTERFEROMETER CALIBRATION

The purpose of the cat's eye interferometer is to track the axial position of the corneal surface. The terminology that will be used throughout this dissertation will describe the position of the cornea as a displacement relative to some known location in the interferometer. Knowledge of the corneal displacement is a necessary parameter to reconstruct the absolute shape, as will be discussed in Section 4.4. The cat's eye measurement is made from a retro-reflection off of the test surface when it is placed at the focus of the interferometer, assuming an accessible focus exists. If the test surface is translated along the optical axis of the interferometer from the cat's eye to confocal position, or vice versa, the translated distance corresponds to the radius of the surface (Figure 4.57, left). A typical radius measurement will require both measurement locations to be nulled, although deviations from null could be corrected through knowledge of the interferometer configuration.

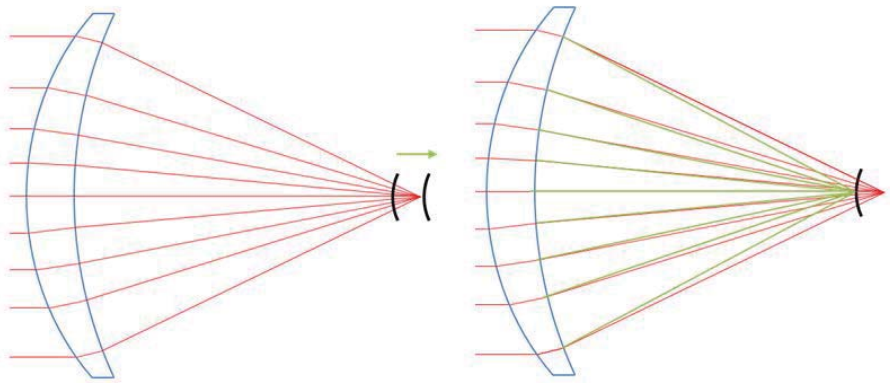

Figure 4.57: Cat's eye measurements.
(Left) Translation from confocal to cat's eye. (Right) Instantaneous cat's eye and surface measurement.

The cat's eye interferometer works off of the principle of a cat's eye measurement by measuring the retro-reflection that results from a beam focused on the corneal surface. The cat's eye interferometer is co-aligned to the surface interferometer, but comes to focus slightly before the surface interferometer focus (Figure 4.57, right). A secondary lens is used in the test path of the cat's eye interferometer to offset the location of focus (Section 3.3.4). The cat null is aligned such that when the golden surface is placed at the null of the surface interferometer, the cat's eye interferometer also produces a null fringe. A surface that deviates from the nominal shape of the cornea will not be able to simultaneously produce a null fringe on both interferometers. The measured differences in defocus from both interferometric measurements could be used to determine the displacement of the test surface relative to a nominal location.

An approximation for the displacement could be derived from the amount of defocus measured in the cat's eye measurement by the relationship:

$$\delta z \approx 4(f/\#)^2 \mathrm{Z}_3 \qquad (4.112)$$

Where δz is the amount of translation along the optical axis, $Z_3 = 2\rho^2 - 1$ is the Zernike fringe polynomial term for defocus, and $f/\#$ is the working f-number of the cat's eye interferometer. In practice, determination of the displacement of the test surface will not be as straightforward. Ocular variation and eye motion will result in large displacements of the corneal surface from the cat's eye focus. The convex

and irregular shape of the eye, coupled with these large motions, will result in a distortion of the cat's eye wavefront that will skew the displacement calculation. The result is that calculated displacement will have to be a function of measurements from both interferometer sub-systems:

$$\delta z = f\left(\phi_{surf}, \phi_{cat}\right) \tag{4.113}$$

Where ϕ_{surf} and ϕ_{cat} are the measured optical path differences from the surface and cat's eye interferometer sub-systems, respectively. The measured optical path differences can be described by the coefficients of a Zernike polynomial fit. The Zernike coefficients can in turn be used to describe the displacement through a polynomial fit. The relationship that was empirically derived for the SPINALTOP system is:

$$\begin{aligned}\delta z = {} & a_1 + a_2 C_3 + a_3 C_3^2 + a_4 S_3 + a_5 S_3^3 + a_6 S_3^4 \\ & + a_7 S_1 + a_8 S_2 + a_9 S_6 + a_{10} S_7\end{aligned} \tag{4.114}$$

where a_i are the polynomial coefficients for the displacement function and C_j and S_j represent the j^{th} Zernike polynomial coefficients for the cat's eye and surface OPD measurements, respectively.

The calibration process consists of measuring the system against a series of reference surfaces of various shapes and determining the polynomial coefficients from the captured set of measurements. This process will be dependent on a calibrated surface interferometer, because the absolute location of the reference surfaces will have to be established by means of the LMA optimization sub-routine that was described in Section 4.3.2.4.2 for aligning the golden surface to the interferometer model. Results for a cat's eye interferometer calibration are presented in the next chapter.

4.4 REVERSE RAY TRACING

The process of reverse ray tracing has been described in the literature as a means to correct for retrace errors introduced by a non-null interferometric test (Honda et al., 1987; Kurita, 1989; Seong &

Greivenkamp, 2008; Williby et al., 2003). However, there has been limited discussion with regards to reflection-based non-null interferometric tests. The lack of available information was discovered during the development of the reverse ray tracing software that was to be used with SPINALTOP. Most discussions of reverse ray tracing were related to transmission testing. The few discussions applying reverse ray tracing to reflection-based testing were incomplete. This necessitated the development of a method for reverse ray tracing a non-null interferometric measurement for a reflection-based interferometer. The details of this method are presented in this section.

4.4.1 REQUIREMENTS

4.4.1.1 INTERFEROMETER CALIBRATION

The accuracy with which a test surface can be reconstructed through reverse ray tracing is highly dependent on the interferometer model. The details of a calibration process have been discussed in Section 4.3.

4.4.1.2 SOURCE CONFIGURATION

The method described in this section assumes that the source is spatially and temporally coherent. Laser interferometers are typically spatially filtered by placing a pinhole at the focus of the objective before the collimating lens (Hariharan, 2003). A spatial filter at the source constrains every ray in the system to have originated from a common point source. The measured optical path difference at each detector element in the observation plane of the interferometer can then be accurately modelled as the difference between the optical path length for one reference ray and one test ray. Rays are traced for a single wavelength which assumes the system is temporally coherent, or is illuminated by a monochromatic light source.

4.4.1.3 SURFACE MEASUREMENT

Measurements of optical surfaces in a non-null configuration have additional and perhaps non-obvious requirements that are detailed in the reference (Sullivan & Greivenkamp, 2007), but are briefly summarized:

1. The interferometer must be able to resolve interference fringes. A number of methods to extend the dynamic range of an interferometer were discussed at the beginning of this chapter.
2. The range of test surfaces must be accounted for in the interferometer design to prevent vignetting. The authors use the term beam blocking instead of vignetting to emphasize that if a ray is vignetted (in a spatially coherent interferometer), all information is lost for the test surface location from which the ray would have intersected.
3. Caustics must be avoided in the imaging path. The relative ray order at the test part and detector must be maintained so that ray correspondence can be obtained. Caustics will result in ambiguities in the measured optical path difference which will generate false surface features on the reconstructed test surface.

By adhering to these requirements, the recorded interference fringes can be used to determine the full set of optical path differences that correspond to the test surface. The optical path differences will be used to determine the set of optical path lengths from the detector to the test surface, which are required for reverse ray tracing.

4.4.1.4 DISPLACEMENT MEASUREMENT

Phase recovery from interferogram analysis will result in a constant phase ambiguity that is common to all measured path differences. The absolute phase on a given fringe is unknown; the measured phase is a modulo 2π of the true phase, or a modulo of λ (the wavelength) of the true optical path length. This ambiguity in the measured optical path differences will result in indeterminate optical path lengths that are needed for the reverse ray trace operation. For example, a null measurement made with a planar or spherical surface will result in a wavefront that is invariant to propagation. The measured optical path differences can be directly related to surface figure errors without requiring knowledge of the total optical path lengths. In contrast, a complex surface shape will result in a wavefront that evolves as it propagates. Therefore, the accuracy of the reverse ray tracing method will be dependent on knowledge of the absolute optical path lengths.

The method described in this section will require knowledge of the test surface location relative to a specific location on the interferometer to recover the true optical path lengths. The specific location is on a surface that is defined with respect to the interferometer and will be referred to as an analysis surface. The details of the analysis surface are discussed in Section 4.4.2.1. Because it is expected that a range of optical surfaces will be tested in a variety of alignment configurations, this secondary measurement will be described as a displacement measurement. The displacement measurement determines the distance between the analysis surface and the test surface.

Any number of methods can be used to measure the displacement of the test surface from the analysis surface. For example, a null optic with an accessible focus (e.g. a transmission sphere) can be used to measure the test surface at both cat's eye and confocal positions. The distance that the test surface was translated from cat's eye to confocal is measured. The calibrated interferometer model can provide the locations of both the cat's eye and analysis surfaces, which can then be used to determine the displacement of the test surface with respect to the analysis surface. Other methods for determining test surface displacement can involve a low coherence source and reference mirror tracking (F H M Jongsma et al., 1997), point source focus tracking (Rottenkolber & Podbielska, 1996), or spot displacement measurements (Heynen & Kahn, 1988). Some ophthalmic applications use a light-based trip line or a shadow to detect when the eye (i.e. the test surface) crosses a known position in space (Tripoli, Cohen, Holmgren, & Coggins, 1995). In a fixed optical system, a magnification target can be used to determine the conjugate distances (Williby et al., 2003). In some instances, *a priori* information of a surface feature (e.g. radius) or an estimate for the test surface location relative to the interferometer can be used with acceptable levels of reconstruction accuracy. The method that is utilized by SPINALTOP to track the displacement of the test surface is by means of the cat's eye interferometer sub-system (Section 4.3.3).

4.4.2 DEFINITIONS

4.4.2.1 ANALYSIS SURFACE

The analysis surface is a virtual construct in test space that represents a wavefront originating from the source. It is a surface of constant optical path length from the source, and its shape encodes the interferometer aberrations associated with this propagation. The analysis surface acts as a basis from which the test surface will be reconstructed. It will only need to be generated once and can be reused for multiple measurements, unless the interferometer is modified or re-calibrated.

The method that is used to generate the analysis surface starts by tracing a single on-axis ray from the source (i.e. at the spatial filter) to a nominal test surface using the calibrated interferometer model. A surface will be required for this operation as interferometer misalignments may result in a non-axial intercept with the surface. The nominal test surface could be a model of an expected test surface or the mean shape of all surfaces that will be tested with the interferometer configuration. Next, a cone of rays is traced from the source to test space, where each ray is traced a total optical path length equal to that of the initial ray. The point of termination for every ray in the test space defines the analysis surface (Figure 4.58).

The density of rays used to generate the analysis surface is dependent upon the desired accuracy from the reconstruction method. For example, if significant mid- to high-spatial frequency structure is generated by optics between the source and test part, then a large number of rays will be need to be traced. Otherwise, these structures will appear as features printed onto the final reconstructed test surface.

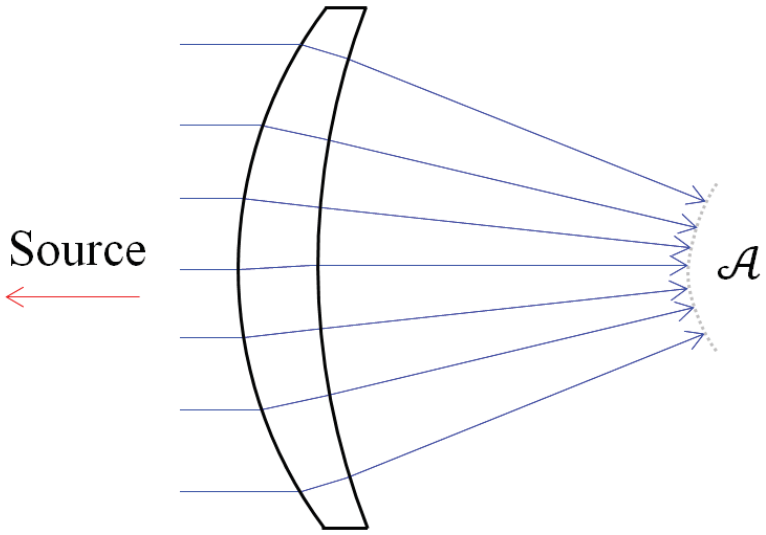

Figure 4.58: Creation of the analysis surface.
Rays are propagated by an equal optical path length from the source (from the left) and the point of termination for each ray defines the analysis surface (represented by $\mathcal{A}$).

The analysis surface can be saved as a surface in the model by a number of means. One suggested method is to use a Zernike polynomial representation of the surface, which acts as low-pass filter to minimize noise and stray light effects (Lowman & Greivenkamp, 1994b). Commercial ray tracing software allows for custom surface deformation models to be mapped over optical elements, which may require the point cloud to be interpolated into a Cartesian coordinate system. Otherwise, the raw set of points corresponding to the terminated ray locations and the direction cosines can be saved and operated on directly during reverse ray trace operations.

4.4.2.2 TEST PATH DOMAINS

As discussed in Section 4.4.1.4, the total optical path length through the test path will be recovered by means of the analysis surface. The operation requires that the set of test optical path lengths ($\mathbf{OPL}_{test}$) be decomposed into three domains as illustrated with a Twyman-Green interferometer in Figure 4.59 and expressed as:

$$\mathbf{OPL}_{test} = \mathbf{OPL}_1 + \mathbf{OPL}_2 + \mathbf{OPL}_3 \qquad (4.115)$$

A variable that represents a set, or collection, will be denoted by bold and non-italicized characters (e.g. **OPL**) to emphasize that the reconstruction methods discussed in this section are performed on a discrete sets of data. Constant or scalar values will be represented by italic characters (e.g. OPL).

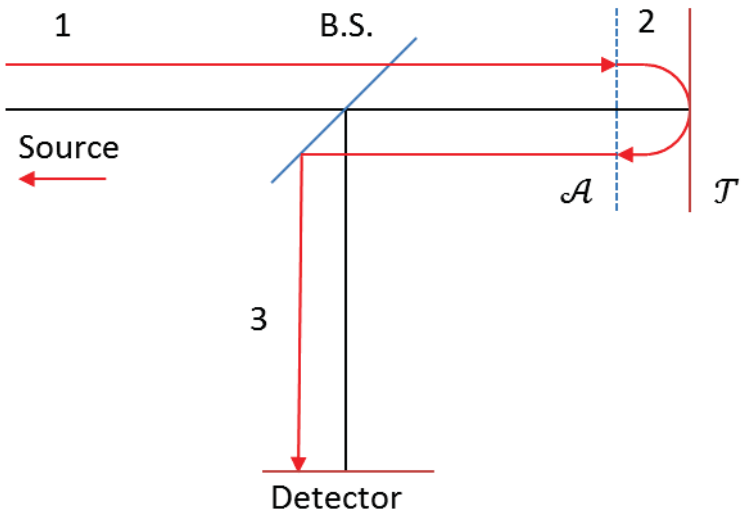

Figure 4.59: Decomposition of total optical path lengths through the interferometer. B.S. is a beamsplitter. $\mathcal{A}$ is the analysis surface, and $\mathcal{T}$ is the test surface.

The first domain is captured by the analysis surface, which includes propagation from the source into test space and terminating at the analysis surface. The optical path length for every ray in the first domain ($\mathbf{OPL}_1$) are all equal and are instead represented as a constant optical path length:

$$\mathbf{OPL}_1 = OPL_1 \qquad (4.116)$$

The second domain contains what will be referred to as the residual optical path lengths ($\mathbf{OPL}_2$), which capture the optical path length that a ray requires to propagate from the analysis surface to the test surface and back to the analysis surface after reflection. The third and final domain contains the set of optical path lengths ($\mathbf{OPL}_3$) of the reflected ray from the analysis surface to the detector surface.

The reconstruction method that is described in this section will use reverse ray tracing to propagate rays from the detector surface based upon the measured test wavefront to the analysis surface (domain 3). The residual optical path lengths (domain 2) can be determined using the analysis surface (domain 1) and reversed rays (domain 3). The final operation is to determine the fraction of the residual optical path length that would complete the reverse ray trace propagation from the analysis surface to the test surface.

To reiterate the discussion from Section 4.4.1.4, a displacement measurement is required to track the distance or offset of the test surface relative to the analysis surface. In practice, most measurements will be made with respect to some aspect of the interferometer and not the analysis surface. The calibrated interferometer model can be used to transform the measurement into the coordinate system of the analysis surface to determine the displacement of the test surface with respect to the analysis surface.

4.4.2.3 REFERENCE OPTICAL PATH LENGTHS

The reference optical path lengths are generated once and are reused for multiple measurements, unless the interferometer is modified or re-calibrated. The process uses a forward ray trace to propagate a set of rays from the source through the reference path, until they terminate at the detector surface. These path lengths will encode the aberrations associated with the reference arm of the interferometer model.

An iterative method can be used to solve for the set of rays that intersect the center of every detector element. Although this requires additional computation, it is a one-time operation and avoids operations such as interpolation to remap the data to match the detector pixel locations.

4.4.2.4 TEST OPTICAL PATH LENGTHS

The set of optical path differences measured at the detector is defined as the difference in optical path lengths for the set of rays originating from the reference and test paths:

$$\mathbf{OPD} = \mathbf{OPL}_{test} - \mathbf{OPL}_{ref} \quad (4.117)$$

where $\mathbf{OPD}$ is the set of physical optical path differences at the detector, $\mathbf{OPL}_{test}$ is the set of physical optical path lengths from the test path, and $\mathbf{OPL}_{ref}$ is the set of physical optical path lengths from the reference path. An estimate for the set of test optical path lengths ($\widetilde{\mathbf{OPL}}_{test}$) can be derived from the measured optical path difference $\left(\widetilde{\mathbf{OPD}}\right)$ and the modelled reference optical path lengths $\left(\mathbf{OPL}_{ref}\right)$:

$$\widetilde{\mathbf{OPL}}_{test} = \widetilde{\mathbf{OPD}} + \mathbf{OPL}_{ref} \quad (4.118)$$

It is currently assumed that the interferometer model is an accurate representation of the hardware and therefore the modelled and physical reference optical path lengths are equal. Inaccuracies in the interferometer model will invalidate this assumption and will be discussed in section 4.4.4.3.

There is an important distinction between the physical and measured optical path differences that was previously discussed in section 4.4.1.4. The optical path differences that are determined from interferometric measurements are indeterminate by a constant optical path length (C) such that:

$$\widetilde{\mathbf{OPD}} = \mathbf{OPD} + C \quad (4.119)$$

The constant C is a direct consequence of the modulo 2π or λ ambiguity. Ideally, the set of test optical path lengths could be recovered if the constant C was known:

$$\mathbf{OPL}_{test} = \left(\widetilde{\mathbf{OPD}} - C\right) + \mathbf{OPL}_{ref} \quad (4.120)$$

4.4.3 RECONSTRUCTION ALGORITHM

The test surface reconstruction algorithm starts by determining the reverse ray trace parameters using the derived test optical path lengths $\left(\widetilde{\mathbf{OPL}}_{test}\right)$ as found through eq. 4.118. The rays are traced in reverse from the detector surface to the analysis surface ($\mathbf{OPL}_3$, domain 3, refer to Figure 4.59). The residual optical path lengths ($\mathbf{OPL}_2$) can be determined from the analysis surface geometry ($\mathbf{OPL}_3$), reverse ray optical path lengths, and independently obtained knowledge of the test surface displacement. Finally, the optical path lengths from the analysis surface to the test surface are determined from the residual optical path lengths.

4.4.3.1 REVERSE RAY TRACE

A ray trace requires two inputs: A starting position and a direction of propagation. Ray positions can be arbitrary, but optimal spatial resolution can be achieved by using the location of each detector element. The estimated direction cosines $\left(\tilde{L}, \tilde{M}, \tilde{N}\right)$ of the rays corresponding to the estimated test wavefront at the detector are given by a normalized vector that is orthogonal to the wavefront:

$$\tilde{\mathbf{L}}_{test} = -\Delta_x \widetilde{\boldsymbol{OPL}}_{test} \quad (4.121)$$

$$\widetilde{\mathbf{M}}_{test} = -\Delta_y \widetilde{\mathbf{OPL}}_{test} \quad (4.122)$$

$$\widetilde{\mathbf{N}}_{test} = \sqrt{1 - \tilde{\mathbf{L}}_{test}^2 - \widetilde{\mathbf{M}}_{test}^2} \quad (4.123)$$

Where Δ_x and Δ_y represent discrete numerical derivatives along the x and y directions, respectively. The signs will be dependent on the software's sign conventions and method for reverse ray tracing.

The derived rays are reverse traced through the interferometer until they intersect the analysis surface. The reverse rays will have propagated by the optical path length in domain 3 ($\mathbf{OPL}_3$). An estimate for the residual optical path lengths ($\widetilde{\mathbf{OPL}}_2$) can then be determined:

$$\widetilde{\mathbf{OPL}}_2 = \widetilde{\mathbf{OPL}}_{test} - \mathbf{OPL}_3 - OPL_1 \qquad (4.124)$$

A consequence of solving for the residual optical path lengths by using the estimated optical path lengths results in an ambiguity in the estimated residual optical path lengths ($\widetilde{\mathbf{OPL}}_2$). The ambiguity in the estimated test optical path length (which is a constant) can be combined with OPL_1 into a single constant C', which would provide a solution for the residual optical path length:

$$\mathbf{OPL}_2 = \widetilde{\mathbf{OPL}}_{test} - \mathbf{OPL}_3 + C' \qquad (4.125)$$

The definition for the estimated residual optical path lengths can therefore be simplified to a form:

$$\widetilde{\mathbf{OPL}}_2 = \widetilde{\mathbf{OPL}}_{test} - \mathbf{OPL}_3 \qquad (4.126)$$

Such that:

$$\mathbf{OPL}_2 = \widetilde{\mathbf{OPL}}_2 + C' \qquad (4.127)$$

The goal is to then determine the value for C', which is discussed in the following section.

A previous study demonstrated the use of a fitted polynomial function to directly map the optical path differences into test space as an alternative to the reverse ray trace operation (Kurita et al., 1987). However, they reported errors of up to 5% in surface reconstruction with their method, which would be insufficient for most non-null applications. Reconstruction errors from a polynomial mapping function

could be substantially improved by extending the mapped parameters and continuing the reconstruction that is described in this paper. The extended mapping parameters would require a position and direction cosine in detector space as an input to determine the analysis surface position, direction cosine, and residual path lengths that would be equivalent to the reverse ray trace operation that has just been described.

4.4.3.2 RESIDUAL PATH CORRECTION

The next step in the reconstruction algorithm is to solve for the constant offset C' (Eq. 4.127), which is necessary to recover the residual optical path lengths. The offset C' is common to all residual optical path lengths. Therefore, if the true residual optical path length is known for at least one of the estimated residual optical path lengths, then C' can be solved for. This can be written as:

$$C' = OPL_{2,s} - \widetilde{OPL}_{2,s} \qquad (4.128)$$

Where $OPL_{2,s}$ is the true residual optical path length for a specific ray and $\widetilde{OPL}_{2,s}$ is the estimated residual optical path length for the same specific ray. The geometry of the analysis surface (Section 4.4.2.1) and displacement measurement (Section 4.4.1.4) can be used to determine C'.

Figure 4.60 shows the double-pass geometry for a forward ray trace through the analysis surface ($\mathcal{A}$) to the test surface ($\mathcal{T}$), which is then reflected back to the analysis surface. The specific ray that is used to solve for C' must be coincident with the test surface at the location that is associated with the displacement measurement. This is a result of the fact that the test surface shape is unknown, and therefore the method of measurement that was used to determine the displacement can only be isolated to a small region or point on the test surface. For example, if the displacement distance was determined by translating the test surface along the optical axis between the confocal and cat's eye positions, the specific ray would be coincident with the location of the null optic focus on the test surface that was used for the cat's eye measurement. The displacement is denoted by d in Figure 4.60 and must be converted to units

of optical path length. The displacement axis (DA) is the axis along which the displacement was determined, which in the case of a cat's eye to confocal measurement is ideally coincident with the optical axis. The variable $l_{s,i}$ is the incoming optical path length that results from a forward ray trace from the analysis surface to the test surface for the specific ray. Likewise, $l_{s,o}$ is the outgoing optical path length of the specific ray reflected from the test surface and propagated back to analysis surface. The sum of the incoming and outgoing optical path lengths must equal the residual optical path length:

$$OPL_{2,s} = l_{s,i} + l_{s,o} \qquad (4.129)$$

The geometry shown in Figure 4.60 is not a requirement for the surface reconstruction algorithm and is only used as an example. The test surface can intersect or lie on the other side of the analysis surface. In the case that the test surface comes before the analysis surface, the residual optical path length will simply become negative.

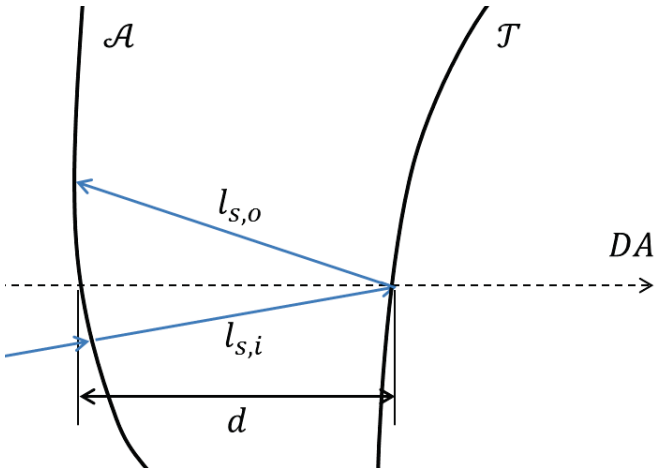

Figure 4.60: Geometry of the test and analysis surface used to determine the residual optical path correction factor.

The next step is to recover C' by locating a reverse ray that is coincident to the specific outgoing ray ($l_{s,o}$ in Figure 4.60). But, the outgoing ray is dependent on deflection from an unknown test surface. Therefore, a closed form solution for C' cannot be determined with the information currently available. However, the problem lends itself to an iterative solution: An initial estimate for a test surface can be

constructed to solve for C', which in turn allows the reconstruction algorithm to solve for an updated estimate to the test surface. The process iterates until a desired convergence on the test surface is achieved.

Two approximations are used to solve for an initial value of C'. First, it is assumed that the curvature of the analysis surface is small over the region where the incoming and outgoing rays intercept. This allows the analysis surface in the vicinity of the ray intersection to be approximated by a plane. The second approximation is that a reverse ray can be used as the specific ray that intersects the displacement measurement location at the analysis surface, and not a forward ray at the test surface intersection. The magnitude of the error for C' will be comparable to the amount of surface sag that was ignored by these approximations, which in general is a fraction of the overall residual optical path length. After the first iteration, the test surface estimate and analysis surface sag can be used to provide an improved estimate of the residual path correction term.

These approximations will result in the geometry shown in Figure 4.61. The test surface can be represented as a plane $(\mathcal{T}')$ since the current problem setup only requires knowledge of a single point on the test surface. The angle α is the angle of incidence that the reverse ray makes with the analysis surface approximation $(\mathcal{A}')$. The angle β is the angle of incidence that the reverse ray makes with the displacement axis (DA) along which the displacement distance (d) was derived.

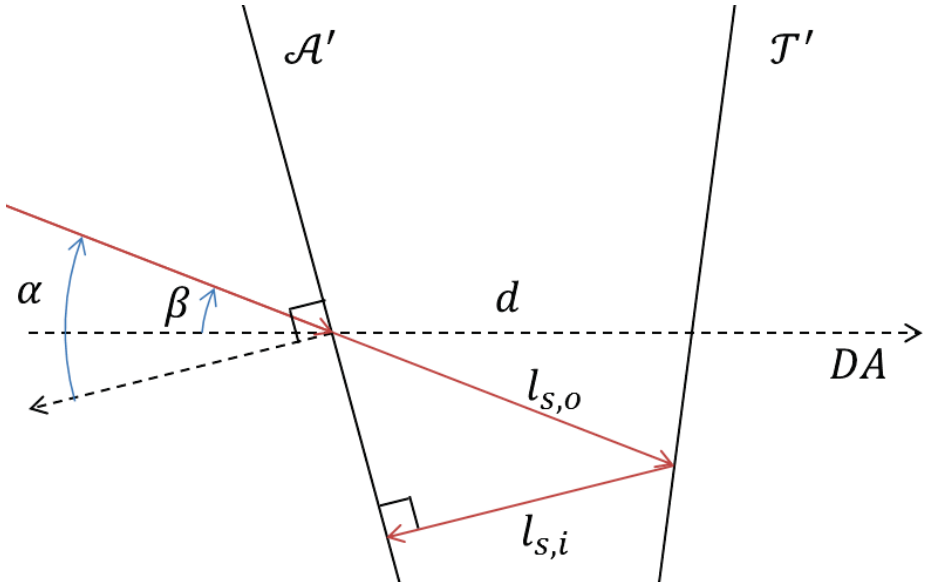

Figure 4.61: Geometry of the test and analysis surfaces
with approximations to determine the residual optical path correction term.

The solution to this geometry now requires a specific reverse ray that simultaneously intersects the analysis surface and displacement axis. The ray can be found by iteratively solving for a reverse ray that meets these requirements or by interpolating across a group of existing reversed rays in the neighborhood of the intersection to estimate the properties of the specific ray.

By definition of the analysis surface, the reversed ray that is reflected from the test surface must be normally incident to the analysis surface. That is because the analysis surface represents the wavefront originating from the source, and a reversed ray that has reflected from the test surface must now be coincident with a ray that is normal to that wavefront. This results in a reflected angle from the test surface that is twice the angle that is formed between the approximated analysis and test surfaces. The solution for the residual optical path length can then be derived with trigonometric relationships:

$$OPL_{2,s} = l_{s,i} + l_{s,o} = \frac{d\cos\left(\frac{\alpha}{2} - \beta\right)}{\cos\left(\frac{\alpha}{2}\right)}(1 + \cos\alpha) \qquad (4.130)$$

A standard interferometer configuration would result in the analysis surface being normal to the optical axis such that $\alpha = \beta$. Aberrations in the interferometer would result in the case where $\alpha \neq \beta$ but a reasonable approximation could be made that $\alpha \approx \beta$ and the remaining optical path distance could be estimated by:

$$OPL_{2,s} \approx d(1 + cos\,\alpha) \quad (4.131)$$

And for a null measurement where $\alpha \approx 0$, this would reduce to the form:

$$OPL_{2,s} = 2d \quad (4.132)$$

which is the expected case for a null measurement in double-pass, where displacement along the optical axis results in twice the amount of optical path length.

Finally, the residual optical path length for the specific ray Eq. 4.130 can be used to solve for the residual path correction term Eq. 4.128. The correct amount of residual optical path lengths for the set of reverse rays can be determined with Eq. 4.127.

4.4.3.3 TEST SURFACE RECONSTRUCTION

The test surface is reconstructed by propagating the set of reversed rays by a fraction of the residual optical path length Eq. 4.127. The residual optical path length (OPL_2) for every reverse ray can be split into two terms:

$$OPL_2 = l_i + l_o \quad (4.133)$$

Where l_i is the incoming optical path length, and l_o is the outgoing optical path length, and defined in the same manner as eq. 4.129.. The geometry that will be used for this derivation is shown in Figure 4.62. From this geometry, it can now be seen that the objective is to recover the value of the outgoing optical path length l_o, which will allow the reverse ray to be propagated from the analysis surface to the test

surface. It should be noted that the geometry shown in Figure 4.62 is an example configuration and the math that follows will still be valid if the test surface intersects or precedes the analysis surface.

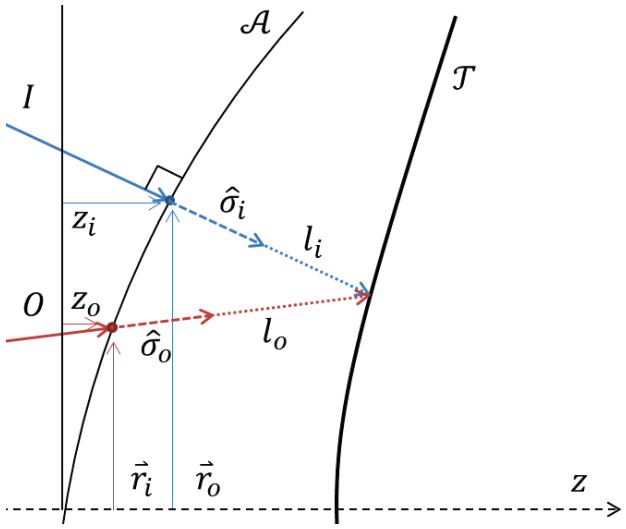

Figure 4.62: Geometry of the test and analysis surfaces for the test surface reconstruction. l_i and l_o are the path lengths from the analysis to test surface.

The variables in Figure 4.62 are defined relative to the coordinate system of the analysis surface ($\mathcal{A}$) and are defined as follows:

$\vec{r}$ is the 2-dimensional radial coordinate of a point on the analysis surface

z is the analysis surface sag at $\vec{r}$, such that $z = \mathcal{A}(\vec{r})$

$\vec{R}$ is the 3-dimensional radial coordinate of a point on the analysis surface, $\vec{R} = (\vec{r}, z) = (x, y, z)$

$\hat{\sigma}$ is the direction cosine of the propagating ray, $\hat{\sigma} = (L, M, N)$

Subscript o denotes an outgoing ray, i denotes an incoming ray

A solution requires the existence of a forward ray that shares a common intersection on the test surface with a reversed ray:

$$\vec{R}_o + \hat{\sigma}_o l_o = \vec{R}_i + \hat{\sigma}_i l_i \qquad (4.134)$$

The incoming optical path length defined by Eq. 4.133 can be used to solve for l_i:

$$\vec{R}_o + \hat{\sigma}_o l_o = \vec{R}_i + \hat{\sigma}_i (l - l_o) \quad (4.135)$$

A dot product with the direction cosine for the incoming ray is applied to both sides of the equation:

$$\hat{\sigma}_i \cdot \vec{R}_o + \hat{\sigma}_i \cdot \hat{\sigma}_o l_o = \hat{\sigma}_i \cdot \vec{R}_i + \hat{\sigma}_i \cdot \hat{\sigma}_i (l - l_o) \quad (4.136)$$

Using the two vector identities:

$$\hat{\sigma}_i \cdot \hat{\sigma}_i = 1 \quad (4.137)$$

$$\vec{A} \cdot \vec{B} = \left|\vec{A}\right|\left|\vec{B}\right| \cos\theta \quad (4.138)$$

The cosine of the angle between the incoming and outgoing ray angles (θ) can be described as:

$$\hat{\sigma}_i \cdot \hat{\sigma}_o = \cos\theta \quad (4.139)$$

Such that Eq. 4.136 can be reduced to:

$$\hat{\sigma}_i \cdot \vec{R}_o + l_o \cos\theta = \hat{\sigma}_i \cdot \vec{R}_i + l - l_o \quad (4.140)$$

Finally, solving for l_o:

$$l_o = \frac{l + \hat{\sigma}_i \cdot \left(\vec{R}_i - \vec{R}_o\right)}{(1 + \cos\theta)} = \frac{l}{(1 + \cos\theta)} + \frac{\hat{\sigma}_i \cdot \left(\vec{R}_i - \vec{R}_o\right)}{(1 + \cos\theta)} \quad (4.141)$$

The solution for the final optical path length l_o can be determined for every reverse ray at the analysis surface to solve for the test surface.

The final path length l_o can be split into two terms: An obliquity-corrected double-pass scale factor and a sag correction. The first term is the standard double-pass scale factor that includes a correction for the additional accumulated path length that is a result of the ray deflection at the test surface. The deflected

ray will also accumulate additional path length errors over the sag of the analysis surface, which is captured by the second term. Measurements made at null will result in incoming and outgoing rays that are nearly coincident with each other, such that $\theta \approx 0°$ and $\vec{R}_i \approx \vec{R}_o$. The resulting null correction distance then reduces to the expected double-pass scale factor $l_0 = l/2$.

The solution is still dependent on the unknown test surface, so one last approximation will be made to provide an initial estimate for the test surface. The obliquity-corrected double-pass scale factor is the dominate correction term for l_o. The sag correction term is dependent on the sag distance and slope error of the test surface, which is at least an order of magnitude less than the double-pass term. This allows for an initial guess at l_o by assuming that the second term is negligible:

$$\hat{\sigma}_i \cdot \left(\vec{R}_i - \vec{R}_o\right) \approx 0 \qquad (4.142)$$

Therefore, an initial guess for the final optical path lengths can be given by:

$$l_o \approx \frac{d}{(1 + \cos\theta')} \qquad (4.143)$$

where θ' is the angle formed between the known outgoing ray (i.e. a reverse ray) and an estimate for the incoming ray. The angle can be solved by the identity:

$$\cos\theta' = \hat{\sigma}_i' \cdot \hat{\sigma}_o \qquad (4.144)$$

where $\hat{\sigma}_i'$ can be estimated by using the surface normal at $\vec{R}_i \approx \vec{R}_o$. The surface normal at location $\vec{R}_o$ is written as:

$$\hat{\sigma}_i' = \mathcal{N}_{\mathcal{A}}\left(\vec{R}_o\right) \qquad (4.145)$$

where $\mathcal{N}_{\mathcal{A}}$ is the operation that returns an orthonormal vector to the analysis surface.

The set of reverse rays that are at the analysis surface can now be propagated by l_o as determined by Eq. 4.143. The point of termination of each reversed ray can be used to describe the test surface. However, the test surface was reconstructed through a series of operations that required approximations which may not result in a sufficiently accurate reconstruction. Under certain circumstances, such as a near-null measurement, the initial test surface reconstruction may be sufficient. When more accuracy is required, the initial test surface estimate can be used as an input to all of the parameters where approximations were used. This provides an iterative method for solving for the test surface shape and is summarized in the next section.

4.4.3.4 RECONSTRUCTION ALGORITHM

The complete algorithm for reconstructing the test surface from a non-null measurement is outlined in this section.

1. Estimate the test path optical path lengths for the set of rays originating at the detector:

$$\widetilde{\mathbf{OPL}}_{test} = \widetilde{\mathbf{OPD}} + \mathbf{OPL}_{ref} \quad (4.146)$$

2. Determine the direction cosines for the set of test rays at the detector:

$$\tilde{\mathbf{L}}_{test} = -\Delta_x \widetilde{\boldsymbol{OPL}}_{test} \quad (4.147)$$

$$\tilde{\mathbf{M}}_{test} = -\Delta_y \widetilde{\mathbf{OPL}}_{test} \quad (4.148)$$

$$\tilde{\mathbf{N}}_{test} = \sqrt{1 - \tilde{\mathbf{L}}_{test}^2 - \tilde{\mathbf{M}}_{test}^2} \quad (4.149)$$

3. The parameters from steps 1 and 2 can be used to reverse ray trace a set of rays from the detector to the analysis surface.

4. Determine the estimated residual optical path lengths at the analysis surface:

$$\widetilde{\mathbf{OPL}}_2 = \widetilde{\mathbf{OPL}}_{test} - \mathbf{OPL}_3 \qquad (4.150)$$

5. Estimate the correct optical path length for the specific ray associated with the displacement measurement:

$$OPL_{2,s} = l_{s,i} + l_{s,o} = \frac{d\,cos\left(\frac{\alpha}{2} - \beta\right)}{cos\left(\frac{\alpha}{2}\right)}(1 + cos\,\alpha) \qquad (4.151)$$

6. Update the set of residual optical path lengths from Step 4 using knowledge of the correct residual optical path length from Step 5:

$$\mathbf{OPL}_2 = \widetilde{\mathbf{OPL}}_2 + C' = \widetilde{\mathbf{OPL}}_2 + \left(OPL_{2,s} - \widetilde{OPL}_{2,s}\right) \qquad (4.152)$$

7. Provide an initial estimate for the set of final optical path lengths from the analysis surface to the test surface:

$$\mathbf{l}_o \approx \frac{\mathbf{l}}{(1 + \cos\boldsymbol{\theta}')} \qquad (4.153)$$

8. Continue propagating the set of reverse rays from the analysis surface by $\mathbf{l}_o$ to acquire an estimate for the test surface.

$$\vec{\mathbf{R}}_{\mathcal{T}} = \vec{\mathbf{R}}_o + \hat{\boldsymbol{\sigma}}_o \mathbf{l}_o \qquad (4.154)$$

9. Update the residual path correction term (l_c) using the current estimate for the test surface. The specific ray associated with the surface displacement measurement (l_d) can now be solved for directly (refer to Figure 4.60). The reverse ray that is coincident with the forward ray shown in Figure 4.60 can be determined by iteratively solving for the reverse ray that intersects the test surface at the displacement measurement location or by interpolating across a group of existing

reversed rays that are in the neighborhood of the intersection. Once found, the residual optical path lengths can be updated with Eq. 4.152.

10. Update all of the analysis surface parameters:

    Incoming ray propagation distance:

$$\tilde{\mathbf{l}}_i = \mathbf{l} - \mathbf{l}_o \tag{4.155}$$

    Analysis surface intersection:

$$\vec{\mathbf{R}}'_i = \vec{\mathbf{R}}_o + \hat{\boldsymbol{\sigma}}_o \mathbf{l}_o - \hat{\boldsymbol{\sigma}}'_i \tilde{\mathbf{l}}_i \tag{4.156}$$

    Analysis surface sag correction:

$$\vec{\mathbf{R}}'_i = \left(\mathbf{R}'_{i,x}, \mathbf{R}'_{i,y}, \mathcal{A}\left(\mathbf{R}'_{i,x}, \mathbf{R}'_{i,y}\right)\right) \tag{4.157}$$

    Incoming ray direction cosine:

$$\hat{\boldsymbol{\sigma}}'_i = \mathcal{N}_{\mathcal{A}}\left(\vec{\mathbf{R}}'_i\right) \tag{4.158}$$

    Cosine correction:

$$\cos\boldsymbol{\theta}' = \hat{\boldsymbol{\sigma}}'_i \cdot \hat{\boldsymbol{\sigma}}_o \tag{4.159}$$

11. Determine the propagation distance to test surface:

$$\mathbf{l}_o = \frac{\mathbf{l} + \hat{\boldsymbol{\sigma}}_i \cdot \left(\vec{\mathbf{R}}'_i - \vec{\mathbf{R}}_o\right)}{\left(1 + \cos\boldsymbol{\theta}'\right)} \tag{4.160}$$

12. Repeat Steps 8 through 11 until a desired convergence on $\mathbf{l}_o$.

13. The final test surface is then given by the set of points where the reversed rays terminate:

$$\vec{\mathbf{R}}_{\mathcal{T}} = \vec{\mathbf{R}}_o + \hat{\boldsymbol{\sigma}}_o \mathbf{l}_o \qquad (4.161)$$

4.4.4 VERIFICATION AND EXAMPLES

The method described in this section has been verified with a number of interferometer systems and test surfaces. One of these tests will be presented here. The effects of noise in measurements are considered in a second example. A final example considers inaccuracies in the interferometer model.

4.4.4.1 VERIFICATION

The interferometer model that is used for this example is the unfolded TFI system (Section 2.1) and is shown in Figure 4.63. The model consists of a point source on the left and is collimated by an air spaced doublet. Following the collimation lens is a glass beamsplitter cube. The reference path is folded upward and the test path continues the right. After reflection, both test and reference wavefronts are recombined after the beamsplitter in the imaging path. The imaging path is the lower path of the interferometer and contains an imaging lens and terminates at the detector surface.

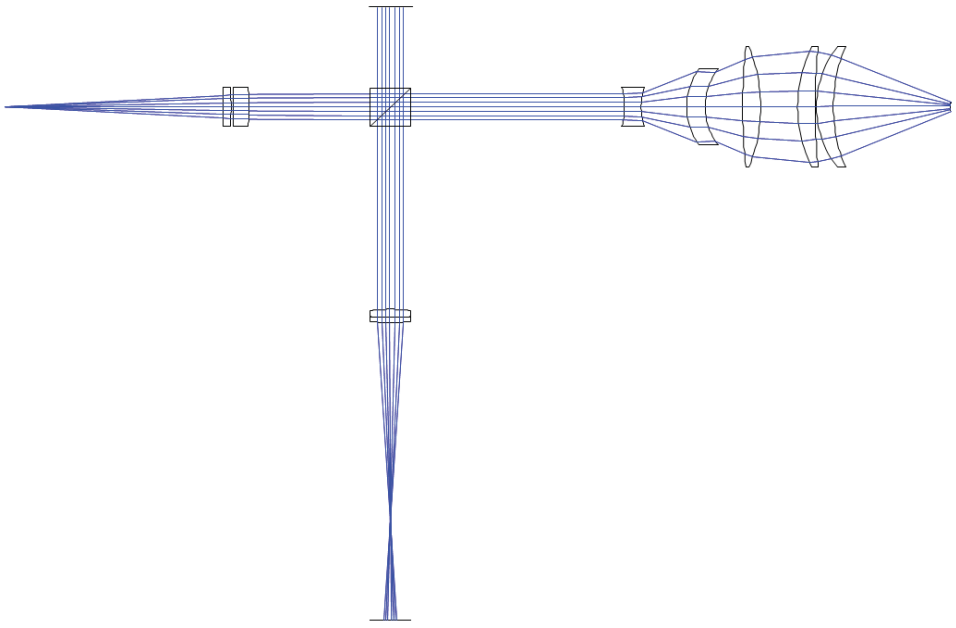

Figure 4.63: Unfolded TFI model used to demonstrate reconstruction method.

The simulated surfaces used on this system are designed to mimic natural eye motion, ocular variation, and tear film dynamics. Eye motion is simulated by applying random decenters to the base aspheric shape with magnitudes of 100 microns or more in all three directions. Ocular variation is simulated by a summation of random Zernike polynomials, using terms 4–15 with random coefficient magnitudes up to 20 microns (Zernike Fringe ordering, see Reference (Wyant & Creath, 1992)). The dynamics of the tear film result in complex surface structures that contain large amounts of mid- to high-spatial frequency information. The variation of the tear film surface is expected to be a few microns, no more than the 4 - 10 μm thickness of the tear film itself. The tear film structure was found to be best simulated by summing a random number of the following terms:

- High order Zernike polynomials (Z16–Z101) with random coefficients.
- A random number of localized Gaussian functions with random position, width, and amplitude.
- A random number of localized sinc functions with random position, width, and amplitude.

The Zernike polynomial functions provide most of the mid-spatial frequency structure. The large number of polynomials used requires that the coefficient magnitudes are kept less than a micron. The localized functions generate most of the high-spatial frequency information and also simulate structures such as pits and breakup that are seen in tear film measurements.

The randomly generated surface that will be used for the examples in this section is shown in Figure 4.64. The surface on the left is the full surface error that will be added onto the base conic shape of the eye. The random surface does not contain tilt or power, which is instead introduced by a lateral decenter of the eye model by 100 microns in both directions and an axial decenter by 175 microns. The surface on the right in is the same surface with the first 35 Zernike terms removed to show the mid- to high-spatial frequency structure.

An interferometer measurement is simulated by forward tracing both the reference and test paths and taking the difference to acquire the optical path differences (OPD) at the detector. Piston is removed from

the OPD measurement to simulate the ambiguity that would result in a real measurement. The displacement value (l_d) can be extracted from the interferometer model, although in practice this would be accomplished by measuring the translation from a cat's eye to confocal configuration. The left image in Figure 4.65 shows the simulated optical path difference that would be measured at the detector. The image on the right shows the optical path difference with the first 35 Zernike terms removed. The images show a parity change that comes from the imaging system. Grey values represent areas where no information exists due to vignetting by elements in the optical path.

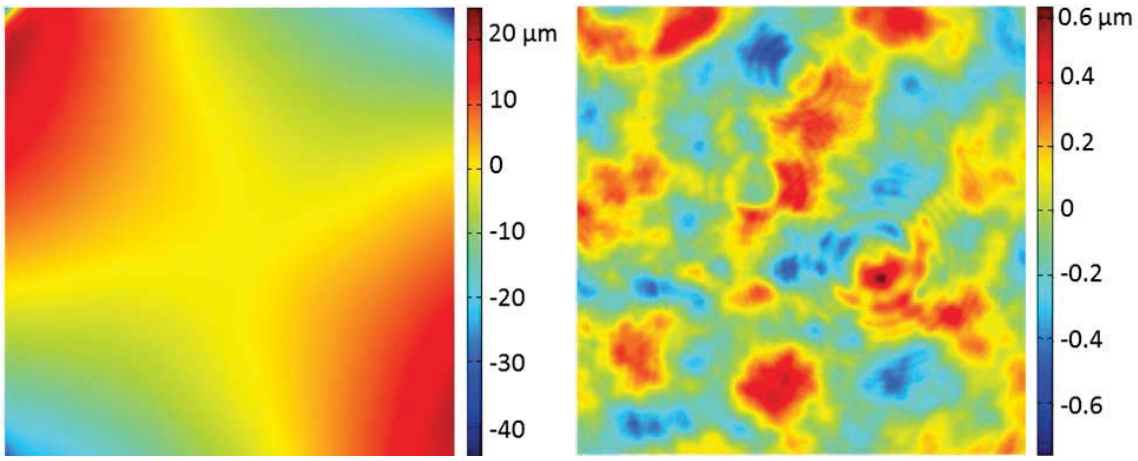

Figure 4.64: Random generated surface that is mapped onto the simulated test surface. The area shown is 6 x 6 mm. (Left) Full surface error (PV 70 µm). (Right) First 35 Zernike terms removed (PV 1.4 µm).

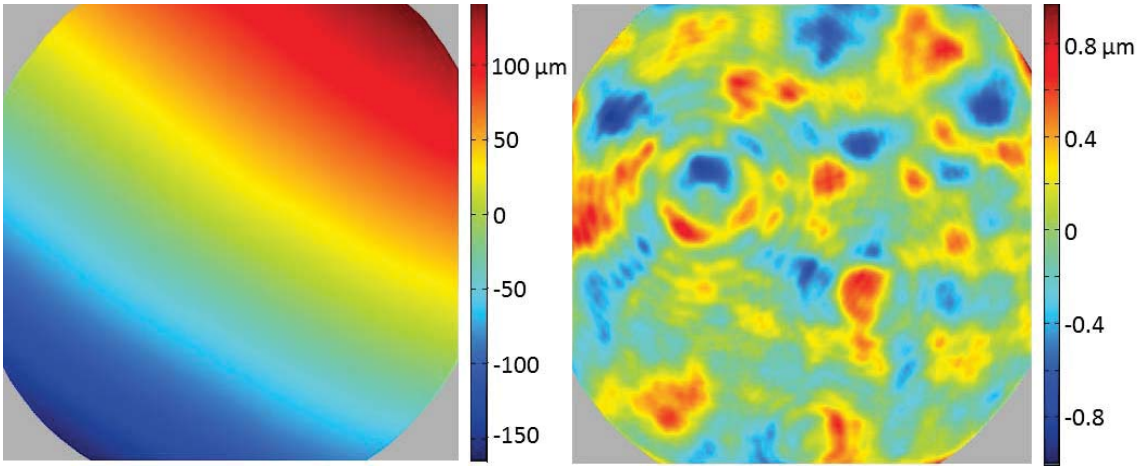

Figure 4.65: The simulated optical path difference measured at the detector. Detector area is 9 x 9 mm. (Left) Full optical path difference (PV 305 µm). (Right) OPD with first 35 Zernike terms removed (PV 1.9 µm).

The simulated optical path differences can now be used to recover the test surface by means of the reconstruction algorithm as discussed in this paper. The left image in Figure 4.66 is the surface sag of the reconstructed test surface. The right image is the reconstruction error, the difference between the reconstructed and known surface. The final reconstruction error is 0.05 nm Peak-to-Valley, or 0.01 nm RMS. Most of the error is attributed to numerical errors that could be potentially reduced by more robust operations. The grey areas in these images also represent areas where no information exists, which is a result of the vignetted OPD measurements.

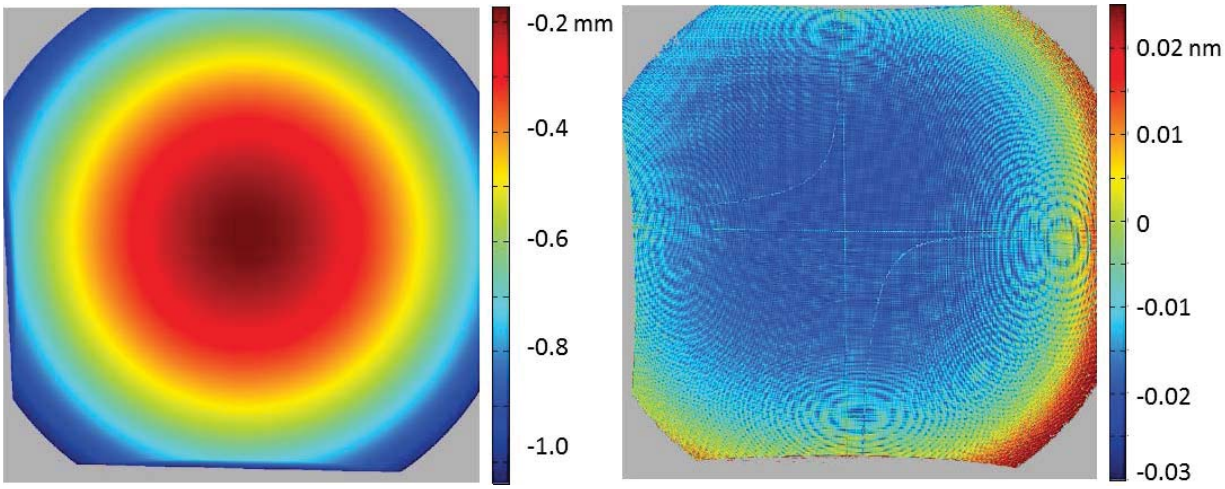

Figure 4.66: Simulated test surface reconstruction.
(Left) Surface sag. (Right) Reconstruction error (PV 0.052 nm, RMS 0.01 nm).

To demonstrate the need for reverse ray tracing, the system performance was simulated in the absence of this correction process. A typical double-pass null interferometer test assumes that the measured optical path differences are twice the measured surface departure from an ideal surface. Applying this methodology to the simulated optical path differences and correcting for parity results in the error shown in Figure 4.67 on the left. The result is 6 μm PV, or 0.9 μm RMS of error, after tilt and power have been removed. A plate scale, or distortion correction, can be applied to the optical path differences in an attempt to minimize effects from distortion in the imaging system. However, this results in the error shown on the right in Figure 4.67, which produces slightly more errors than the previous attempt.

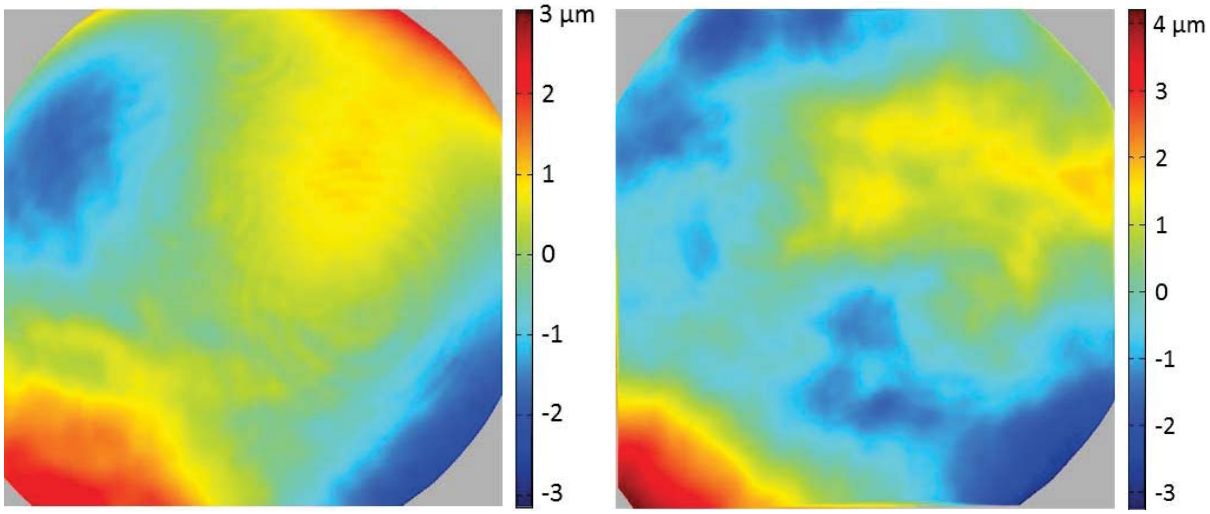

Figure 4.67: Retrace errors resulting from simulated surface figure.
(Left) Error when assuming that surface features are OPD/2 (PV 6 µm, 0.9 µm RMS).
(Right) Error when also applying a distortion correction (PV 7.1 µm, 0.9 µm RMS).

4.4.4.2 MEASUREMENT NOISE

The results of the previous section demonstrated the reconstruction method under ideal conditions. As a practical example, the method will be tested against a measurement that contains noise. A normally distributed random value having a standard deviation of 40 nm is applied to the simulated optical path differences on a pixel by pixel basis. The resulting error from applying the reconstruction method to the noisy optical path differences is shown on the left in Figure 4.68. The 40 nm of noise (RMS) on the measured optical path differences resulted in 67 nm of noise (RMS) in the reconstructed surface. The noise appears to map directly to the reconstructed test surface, however a more robust set of operations could improve the results. As an example, a 5 x 5 mean filter is applied to the simulated noisy optical path differences measurement before passing it through the reconstruction algorithm. The result is shown on the right in Figure 4.68, which is significantly improved over the noisy data with errors on the order of 9 nm PV, and 1.9 nm RMS. Note that the display scale for the image on the left in Figure 4.68 is approximately 50X larger than on the right.

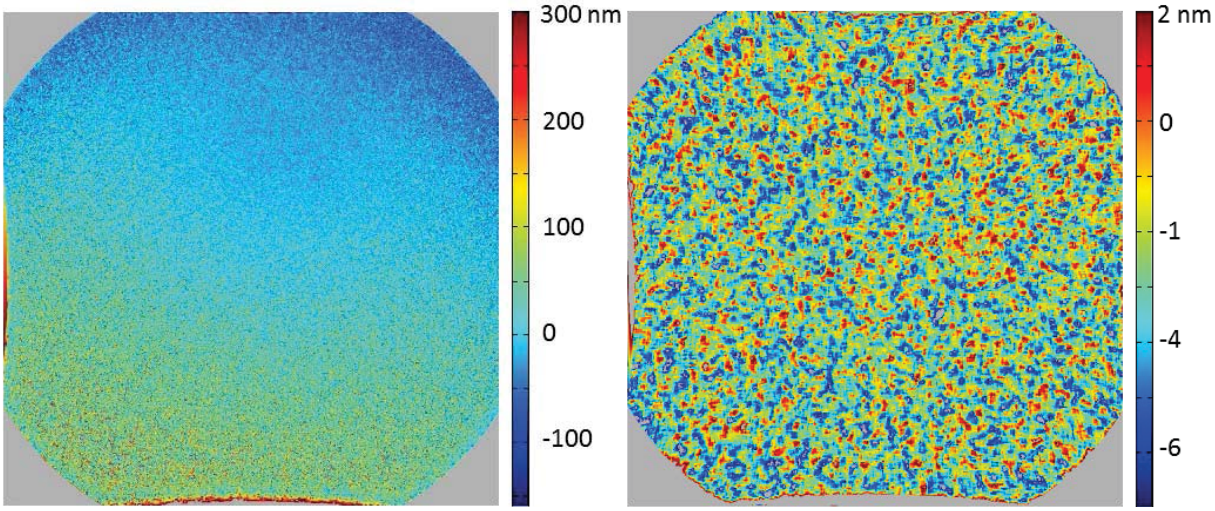

Figure 4.68: Reconstruction errors with measurement noise.
(Left) Reconstruction error with 40 nm (1σ) noise added to the simulated OPD measurement (460 nm PV, 67 nm RMS). (Right) Error when a 5 x 5 windowed mean filter is applied to noisy OPD (9 nm PV, 1.9 nm RMS). Note: The display scale on the left is approximately 50X larger than on the right.

4.4.4.3 INTERFEROMETER MODEL ACCURACY

A final consideration when applying the reconstruction method that is described in this section is the required accuracy of the interferometer model. The method effectively backs out interferometer induced errors by simulating the propagation of the measured test wavefront in reverse from the detector to the test surface. However, this method relies on an accurate model of the interferometer. Differences between the model and the true state of the hardware will result in a differential error that is introduced into the simulated wavefront. For example, a decentered spherical optical surface that is not captured in the interferometer model will result in a linear phase error that is mapped onto the reconstructed test surface. A large enough decenter between the model and hardware would start to result in a noticeable amount of coma. Additionally, more errors will be accumulated as the wavefront is propagated through the rest of the interferometer model.

A model induced error can be demonstrated with the same set of simulated optical path differences from the previous two examples. In the previous examples, the same model was used to generate the set of wavefronts and to reverse ray trace. A separate model can be used for the reverse ray trace operation to

simulate the effects from uncertainties in the interferometer model. This method will only be valid as a demonstration for perturbations in the test path unless the reference optical path lengths are recalculated (Section 4.4.2.3).

As an example, the second to last lens (second closest to test surface) is laterally decentered by 25 μm in the interferometer model, which now differs from the original model used to generate the simulated optical path differences. The resulting error of a surface reconstruction with the perturbed interferometer model is shown in Figure 4.69. The left image is the difference between the reconstructed and actual test surfaces. The dominant error is tilt, which is expected for a laterally decentered lens in this system. A lateral decenter of a lens will result in a lateral shift in the location of focus which would be indistinguishable from a lateral translation of the test surface. It could therefore be assumed that tilt, and similarly power, are a result of errors in locating the reconstructed test surface and could be removed. The resulting residual error with tilt and power removed is shown on the right in Figure 4.69. The residual error now consists of 1.3 μm of mostly coma, which could be interpreted as a residual from a decentered aspheric surface.

However, the process of removing terms invalidates the method of absolute surface reconstruction. It would be indeterminable from this example set of data whether the residual coma was an interferometer induced artifact due to misalignment (i.e. retrace error), surface figure error, or any combination thereof. Furthermore, this example only demonstrated a single surface decenter. The null lens used for this example contains five optical elements, with individual variations in decenter, tip/tilt, radius of curvature, thickness, spacing, index homogeneity, and surface figure errors. Additionally, there are optics in the source, reference, and imaging paths that must be considered. The cumulative effects from every individual element can become quite large if unaccounted for. Therefore, a model of the interferometer must be built to some level of accuracy in order to reconstruct an arbitrary test surface.

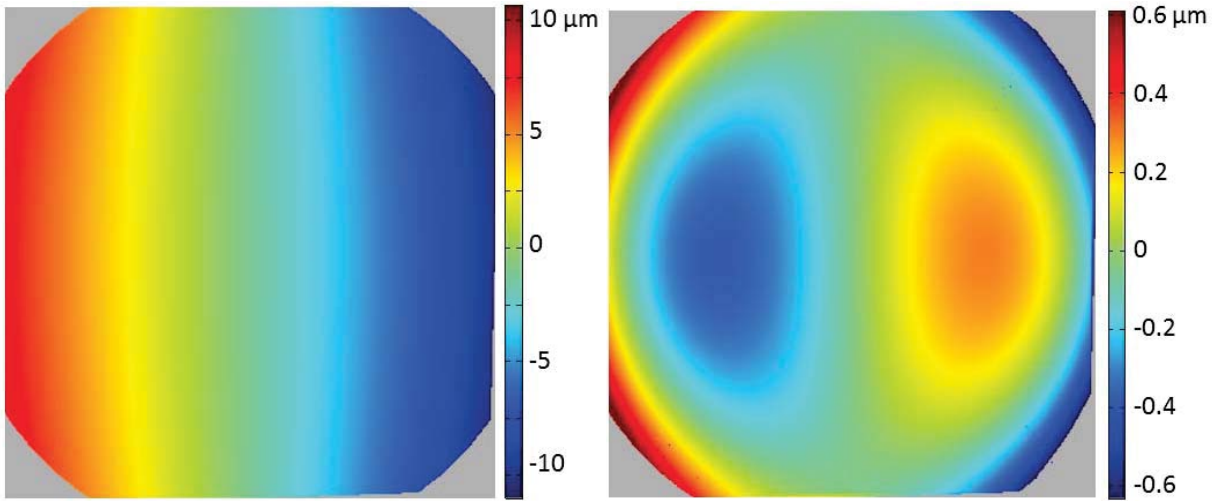

Figure 4.69: Reconstruction errors due to interferometer model errors.
(Left) Reconstruction error with a decenter of a single lens in the null optics (25 µm PV). (Right) Same error with tilt and power removed (1.3 µm PV, 0.25 nm RMS).

4.4.5 CONCLUSION

A method for reconstructing surfaces measured in a non-null configuration on a reflection-based interferometer is presented. The method is capable of reconstructing surfaces with high accuracy, increased spatial resolution, and is less computationally demanding compared to previous methods in literature. Accuracies are demonstrated to 0.05 nm PV and 0.01 nm RMS for a well characterized interferometer and low noise measurement. Residual errors are indicative of numerical errors and could be improved by more robust calculations. Measurement noise will have an impact on reconstruction, but can be minimized with appropriate filtering techniques. The accuracy of the surface reconstruction method is ultimately limited by the accuracy of the interferometer model, which can be improved through calibration (Section 4.3). This method is necessary to measure the absolute shape of the cornea due to the expected range of ocular surface variation, eye motion, and tear film structure.

4.5 ANALYSIS

The reverse ray trace operations from the previous section results in a 3-dimensional point cloud that describes the surface topography of the cornea and tear film (see example Figure 4.76). Sequential measurements provide information on the temporal dynamics. This information can be used for

quantitative analysis on various parameters of the eye, cornea, and tear film. The development of methods to extract information from these measurements as they relate to phenomenology of the eye is beyond the scope of this dissertation. Instead, what is presented in this section are a basic set of methods for handling the results from the previous section. These operations account for measurement noise and errors in the data, surface fitting to extract basic information, and data conversion for display and qualitative assessment. These methods may be necessary precursors for any future algorithm development to study the cornea and tear film.

4.5.1 FILTERING

Measurement noise, extreme surface features in the tear film, interferometer errors (e.g. dust, dead pixels), stray light, and vignetting are examples of conditions that will introduce errors into the reconstructed test surface. These errors generally result in large discontinuities or noisy areas which can corrupt analysis. Therefore, it is necessary to develop a set of metrics to isolate good or usable data from bad data. The first and most straightforward metric is a slope variance, which is similar to the phase derivative variance of Section 4.2.1.3. A second method is developed to eliminate rays that may be caustic. Finally, a clustering algorithm is used to eliminate outliers.

4.5.1.1 SLOPE FILTER

The interferometer system is capable of resolving a small range of surface slopes. Noise and other errors in the measurements may result in reverse traced rays that can exceed the capabilities of the interferometer. A slope filter can detect these rays and remove them. As discussed in Section 2.2, the 1 MP pixelated camera can resolve up to 170 waves/radius at the detector. At the test surface, this would result in ray slopes of approximately 22.2 mrad for a nominal eye surface. The slope filter would look for rays that exceed 22.2 mrad from a null ray. Information from the reverse ray trace operation can be easily adapted to this problem. One of the calculated reverse ray trace parameters was the angle between the

incoming forward ray and reverse ray (Eq. 4.139). A threshold can be applied to this angle to eliminate bad rays with high slopes.

4.5.1.2 CAUSTIC FILTER

Caustics will result in a loss in correspondence mapping of rays from the detector plane to the test surface. The interferometer should be designed to avoid caustics at the image plane (Section 4.4.1.3). However, caustics may still appear at the reconstructed test surface due to noise or other defects in the reconstructed surface. A simplified example of a caustic ray is shown in Figure 4.70. Rays originate from a surface on the left, propagate through an optical system, and intersect the surface on the right. A parity change may occur, which is shown in the figure as a reverse ordering of points. However, a caustic ray is formed on the right that results in points 4 and 5 losing their relative correspondence to the object plane on the left. In other words, the relative ray ordering of rays from the left (1,2,3,4,5) no longer match those on the right (1,2,3,5,4). The physical phenomenon that is occurring at a caustic is that a portion of the wavefront has folded over itself at some point in the propagation between the two surfaces. The wavefront that is observed will contain ambiguities in the regions where the wavefront has folded and surface information corresponding to that wavefront information will be lost or corrupted. In the reverse ray trace process, an equivalent phenomenon will result in corrupted surface information. A method is developed in this section to detect the presence of caustics and remove them.

Detection of caustics on a 2-dimensional surface is more complicated than the example shown in Figure 4.70. The view shown in Figure 4.70 is a 1-dimensional cross-section of an optical system that gives the appearance of a caustic with rays 4 and 5. The system could be rotated along the optical axis producing a different cross-sectional view that does not result in a caustic of rays 4 and 5. This is in fact what commonly occurs with the measurements of the tear film structure. The point cloud from the reconstructed surface in Section 4.4.4.1 is shown Figure 4.71. The points are projected into the x-y plane (i.e. sag ignored). The full point cloud over a 6 x 6 mm area is shown on the left, and a 50 x 50 μm area from the center of the data set is shown on the right. The reverse rays start from a Cartesian grid

originating at the detector surface, however the high spatial-frequency structure in the tear film results in the irregular distribution of points shown at the test surface. There is no clearly definable correspondence mapping from the detector grid to the test surface.

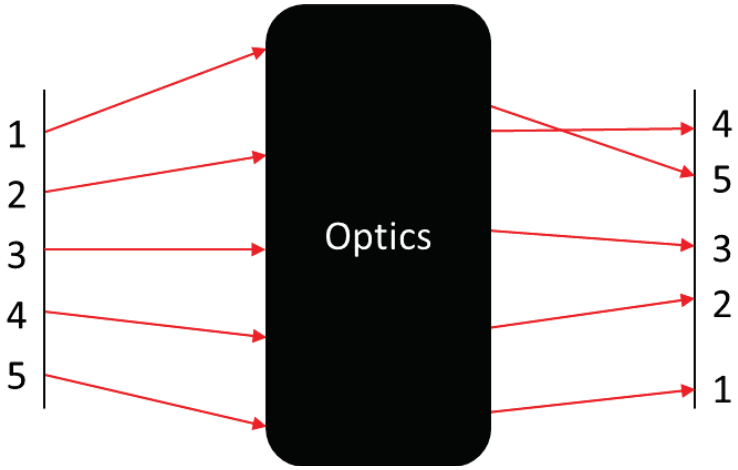

Figure 4.70: Simple caustic ray example.

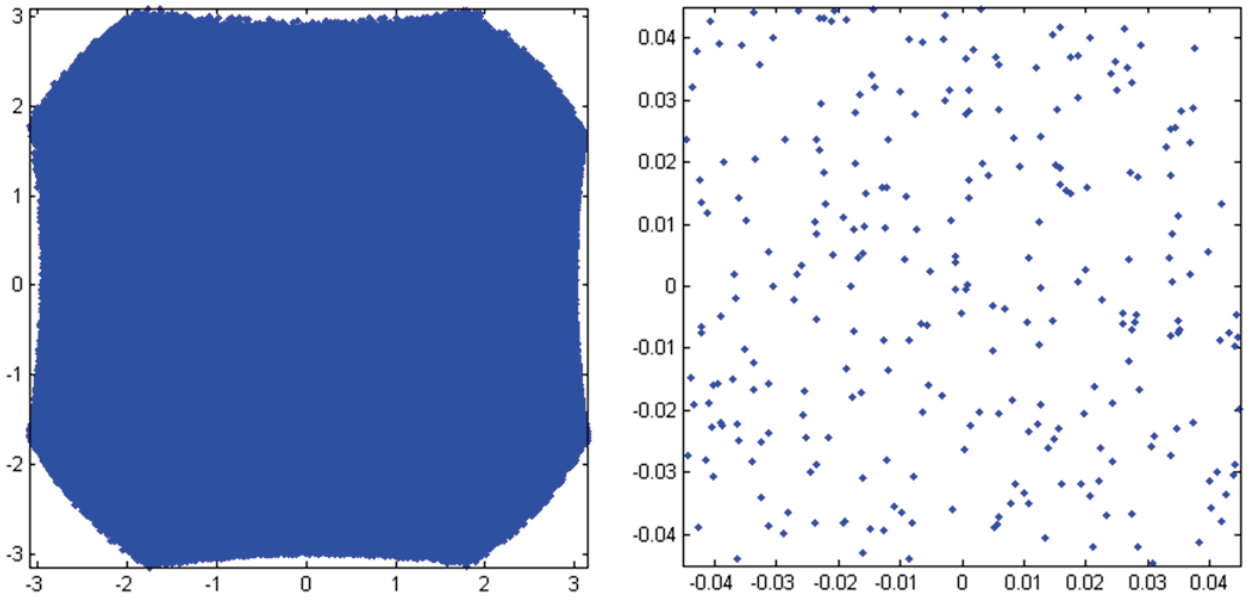

Figure 4.71: Reconstructed point cloud.
(Left) Full 6 x 6 mm area. (Right) Center 50 x 50 µm area.

The method that will be used to determine ray correspondence and detect caustics starts with the example shown in Figure 4.72. The image on the left represents a well-ordered ray set that originates from the

detector surface. Each ray is indexed so that it can be tracked through the optical system. The image on the right represents how the rays could appear at the test surface after having been propagated through the optical system. The top and right sides of the images are labelled with the ray indices to provide cross-sectional sets of data to track ray correspondence, similar to the example in Figure 4.70. This method of ordering would indicate that many of the rays are caustic. For example, the first row of the nominal ray set contains rays 1,2, and 3, whereas the concept of a row after propagation no longer appears exists. Are rays 1, 5, and 2 considered to be the corresponding top row? Or do rays 1, 2, and 3 form a distorted row across the top of the propagated ray set? In some cases it would appear that rays have formed caustics by crossing into other "rows", such as ray 5 appearing near the top, or ray 5 appearing in a column near rays 1, 4, and 7. This example shows that this form of ray ordering to determine caustics in a 2-dimensional case is insufficient and a new method will have to be devised.

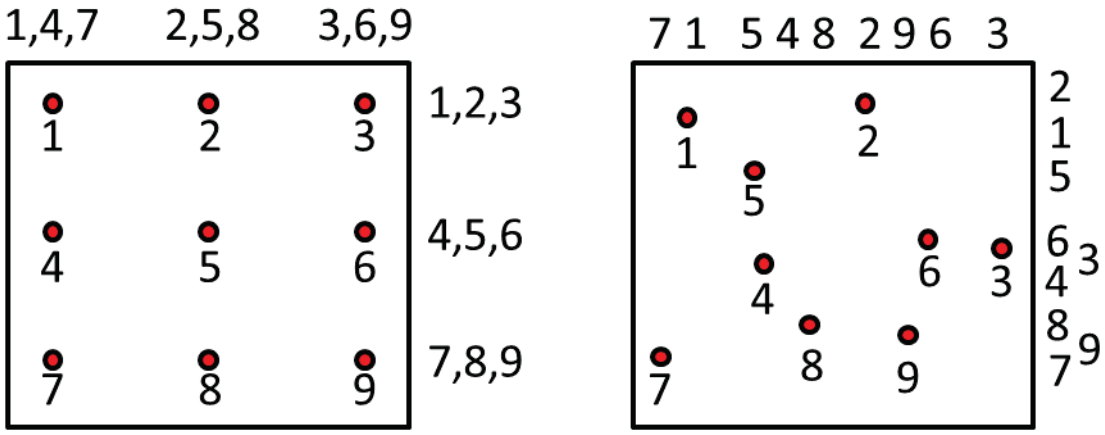

Figure 4.72: 2-dimensional caustic example.

The method that is used to detect caustics works by extending the concept of a point correspondence in 1-dimension to an area correspondence in 2-dimensions. The set of points on a surface can be used to partition the surface into triangulated sub-divisions, a process known as triangulation. Triangulation generates a set of unique non-overlapping triangles that can be used to describe the surface. A Delaunay triangulation minimizes the number of small angles of the triangulated sub-divisions (Lee & Schachter,

1980). The area correspondence method that will be used to detect caustics works by verifying that the ordering of triangulated divisions of the propagated wavefront is maintained.

A Delaunay triangulation of the Cartesian grid of points for the set of initialized reversed rays at the detector plane is trivial. An example is shown on the left in Figure 4.73. Each triangulated division now represents a subset of the wavefront that will be propagated through the interferometer optics and tracked for correspondence. Every triangulated sub-division is defined by a set of three points, or rays, so that the wavefront sub-division is effectively propagated through the optics by the propagation of the set of three rays. At the plane of the test surface, the surface divisions can be reconstructed using the previously defined point relations. The example set of propagated rays from Figure 4.72 (right) are now shown in Figure 4.73 (center) with the triangulated divisions added. The ordering of the set of points shown in the center image are the same as the points shown on the right in Figure 4.72. The surface sub-divisions are not labelled because it can be observed that the relative surface divisions appear to have maintained correspondence, and therefore no caustics exist. To demonstrate a loss of correspondence, ray 4 (refer Figure 4.72) is perturbed towards the right until it becomes caustic (Figure 4.73, right, yellow dot). The caustic ray results in overlapping surface divisions, which is consistent with the definition that a caustic occurs where the wavefront folds over on itself.

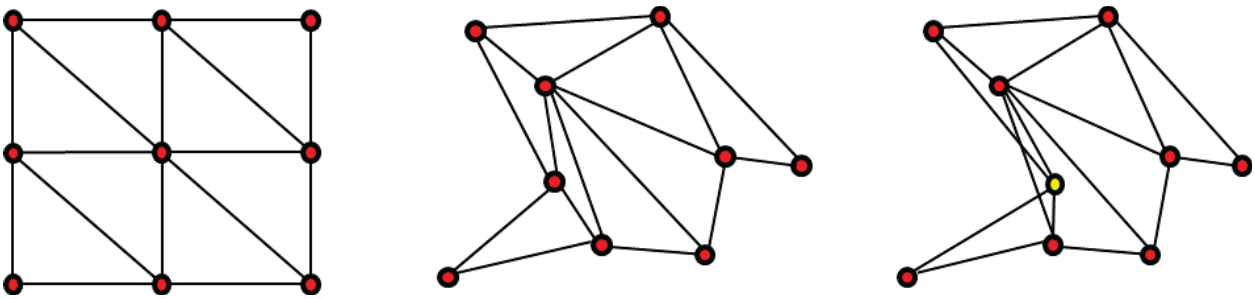

Figure 4.73: Triangulation of rays before and after propagation.
(Left) Rays at a detector surface. (Middle) Rays after propagation. (Right) Rays after propagation with a caustic ray shown in yellow.

Although the caustic ray can be detected by inspection of the propagated sub-divisions as demonstrated in Figure 4.73, the method will have to be described in terms of an algorithm for automated software detection. The algorithm is relatively straightforward. The triangulated sub-divisions are determined at the detector surface. The reverse ray trace operation is performed and results in a set of points at the test surface. The triangulated sub-divisions can be re-established in test spacing using the mapping previously defined at the detector. The filter routine iterates through each of the triangulated divisions at the test plane and searches for any points that exist within the division. If a ray, or point, is found within the triangulated division, it is a caustic ray and is filtered from the data set.

4.5.1.3 CLUSTER FILTER

The cluster filter is used to remove outliers from the set of reconstructed data. Most outliers are formed near regions of high-noise or break-up areas. In these regions, most of the data may have been previously filtered out by previous operations. A few rays may still get through, in which case an outlier detection method is used to remove them. The cluster filter works by generating a region of contiguous points within the main cluster and removing points that are outside of this region. The left image in Figure 4.74 is an example of a set of rays that have terminated in test space. The points are binned into a 2-dimensional grid (Figure 4.74, middle). A pixel is denoted as active if at least one ray intercepts a bin. The result is shown on the right in Figure 4.74. Visual inspection of this example would determine that an outlier exists at the top right of the image. An algorithmic description is still required for software.

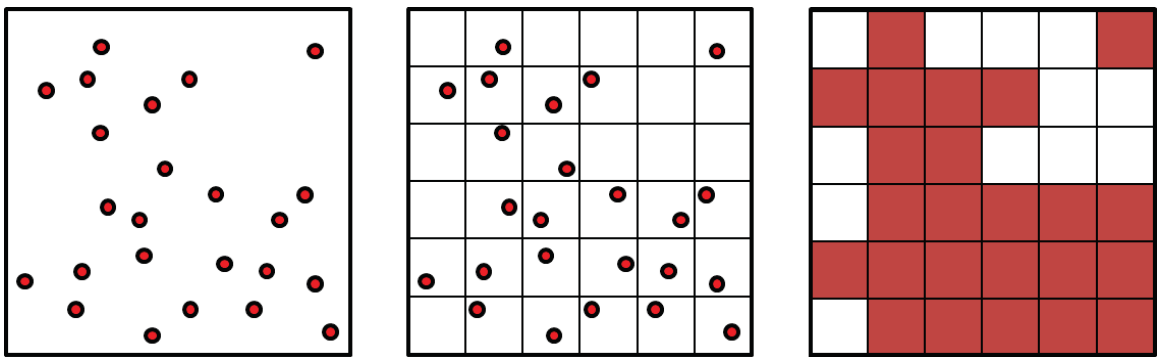

Figure 4.74: Cluster filter example.
(Left) Set of points at test surface. (Center) Binning of points. (Right) Binned image.

The method for determining the outlier can be accomplished through a region growing method (Sonka, Hlavac, & Boyle, 1993). Region growing is a common process for image segmentation, which is used to isolate a cluster of data from another. Region growing with the 2-dimensional binned image is in fact a form of the flood-fill algorithm that has been previously described (Section 4.2.3.1.2). Region growing starts with a seed pixel. The seed should start in a region of known good data, which generally occurs near the center of the main cluster. A region growing method identifies active pixels that are in the neighborhood of the seed. Every pixel that is in the neighborhood that is available (i.e. has been binned) is now marked as an active pixel and added to the queue. Next, a pixel is removed from the top of the queue, neighbors are added to the queue, and the process repeats until the queue is empty. What remains is a list of pixels that form a common region, or a cluster. Pixels that are not in this list have been isolated from the cluster and are classified as outliers. Rays that are associated with the outliers are filtered from the data set. The example from Figure 4.74 is processed in Figure 4.75. A seed pixel is chosen (green, left) and the region growing method fills the region (right). The outlier is still marked red and the ray(s) that are associated with this pixel would be removed.

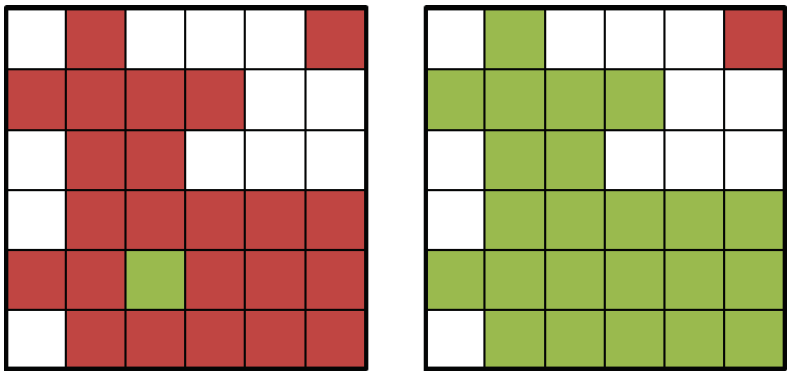

Figure 4.75: Region growing method to isolate outliers.
(Left) Seed pixel shown in green. (Right) Finished region.

4.5.2 BICONIC SURFACE FITTING

The filtered point cloud contains all of the information related to the topography of the cornea and the overlying tear film. The next step is to evaluate the topography of the cornea. The overall shape of the

cornea is best represented by a generalized biconic surface (P M Kiely et al., 1982). The sag for a generalized biconic surface is:

$$Z = \frac{r^2 \cos^2(\theta - \theta_z)/R_x + r^2 \sin^2(\theta - \theta_z)/R_y}{1 + \sqrt{1 - (1 + K_x) r^2 \cos^2(\theta - \theta_z)/R_x^2 - (1 + K_y) r^2 \sin^2(\theta - \theta_z)/R_y^2}} \quad (4.162)$$

Where the radius R_x and conic constant K_x is along the θ_z meridian, and the radius R_y and conic constant K_y are measured along the orthogonal meridian. Eye motion also has to be considered in the surface description before a fitting to the data can be performed. Lateral motion can be captured in the radial term:

$$r = (X - X_0)^2 + (Y - Y_0)^2 \quad (4.163)$$

Where the center of the cornea is located at (X_0, Y_0). Longitudinal displacement of the eye can be accounted for by directly applying an offset to the sag:

$$Z' = Z + Z_0 \quad (4.164)$$

Finally, tip and tilt of the eye can be accounted for by applying coordinate rotations along the X and Y directions:

$$R_x(\theta_x) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\theta_x & -\sin\theta_x \\ 0 & \sin\theta_x & \cos\theta_x \end{bmatrix} \quad (4.165)$$

$$R_y(\theta_y) = \begin{bmatrix} \cos\theta_y & 0 & \sin\theta_y \\ 0 & 1 & 0 \\ -\sin\theta_y & 0 & \cos\theta_y \end{bmatrix} \quad (4.166)$$

Fitting the point cloud to the generalized biconic should result in the following parameters that describe the topography of the eye: R_x, R_y, K_x, K_y, X_0, Y_0, Z_0, θ_x, θ_y, θ_z. The eye model surface description is

non-linear, so a least-squares solution is not easily derived. The LMA algorithm (Section 4.3.2.4.1) is used to solve for the parameters.

The point cloud from the example reverse ray trace problem in Section 4.4.4.1 is shown in Figure 4.76. The figure demonstrates how the mean topography of the cornea will dominate the dynamic range of the display. However, once the biconic fit is performed, the mean shape can be removed as shown in Figure 4.77. The underlying structure of the tear film can now be observed, however the point cloud is no longer an optimal medium to visualize the data. The next section will discuss a remapping function to map the set of points into a more usable format.

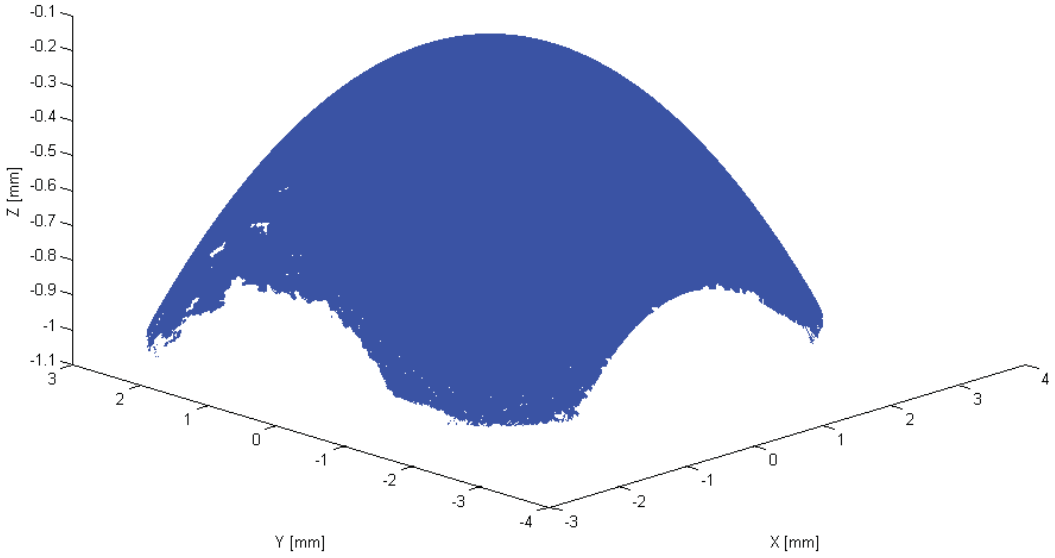

Figure 4.76: Point cloud from reconstructed topography.

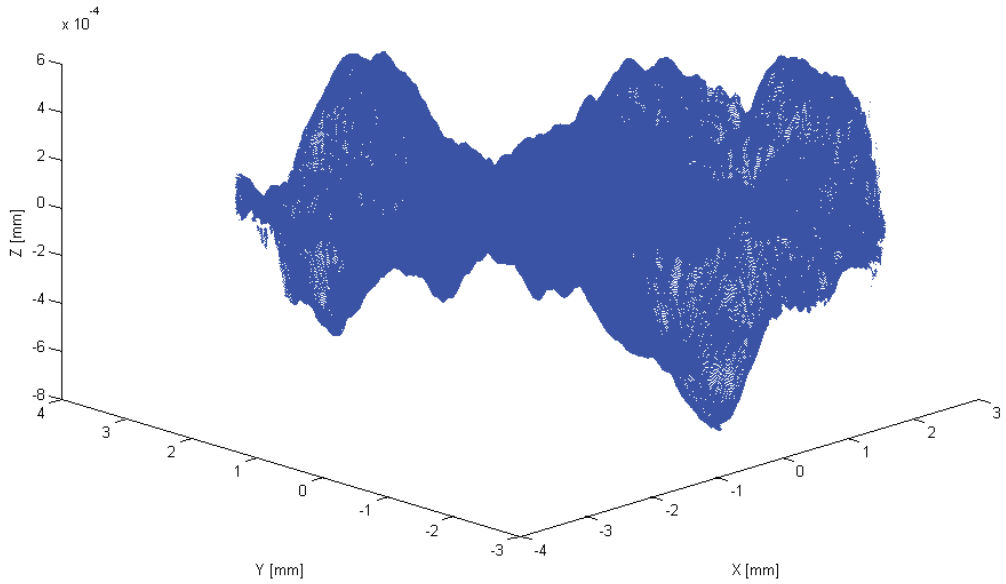

Figure 4.77: Best-fit generalized conic removed from point cloud.

4.5.3 COORDINATE REMAPPING

The reconstructed point cloud contains all of the information describing the topography of the cornea and tear film structure, but the irregular spacing of data makes it difficult to analyze or visualize as demonstrated in the previous section. Other than the generalized biconic surface fitting, the remaining operations will be best performed on a regularly spaced grid of data. This will require an operation to interpolate the sag information (i.e. Z-coordinates) into a standard X-Y Cartesian grid. The method described in this section takes advantage of the pre-existing Delaunay triangulation from Section 4.5.1.2 and uses a simple linear interpolation method to map the data into a regularly spaced grid.

The example set of irregularly spaced points from Section 4.5.1.2 is shown in Figure 4.78 on the left. The points are projected into the X-Y plane and have an associated sag value (Z) associated with them. The Delaunay triangulation results in a set of triangulated divisions as shown by the middle image of Figure 4.78. The desired mapping coordinates are shown on the right as a set of uniformly spaced black points. An interpolated value for the remapping coordinates can be determined by a simple interpolation of the surrounding vertices of any triangulated subset.

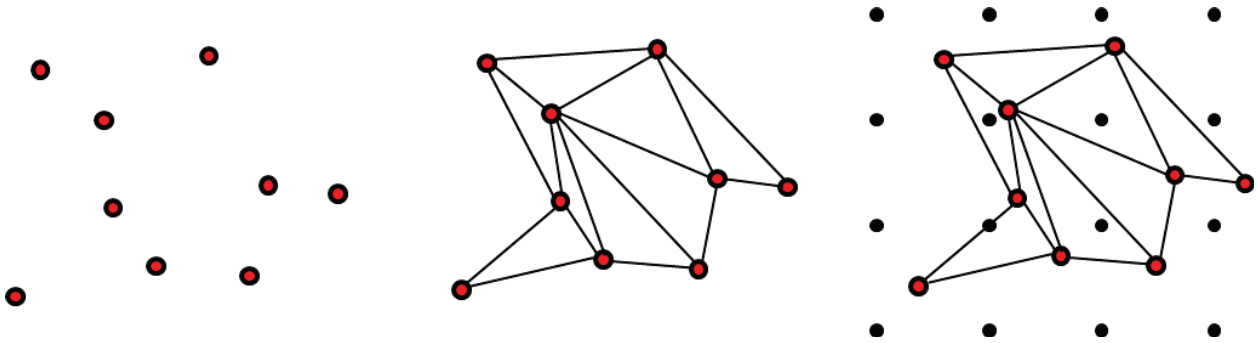

Figure 4.78: Coordinate interpolation from triangulated set of data.
(Left) Rays in test space. (Middle) Triangulated surfaces. (Right) Interpolation grid mapped over data.

A simple method for interpolation is to take the barycentric coordinates of the remapped triangulated sub-divisions and to derive a weighted average of the three sag values associated with the vertices (Renka, 1984). Barycentric coordinates correspond to the relative areas of the sub-divisions within a simplex. The simplex in this case is the triangulated sub-division, and the barycentric sub-divisions are formed by connecting the desired interpolant to the vertices of the triangle. This is illustrated with Figure 4.79. First, the desired interpolation point is selected; this is shown as the dark black circle in Figure 4.79 (left). The triangulated sub-division that bounds the desired interpolation point is isolated (Figure 4.79, middle). The coordinates of the three vertices are given by r_1, r_2, and r_3 and the desired interpolation coordinate by r'. The barycentric coordinates are defined as the normalized areas b_1, b_2, and b_3 of the barycentric sub-divisions shown on the right.

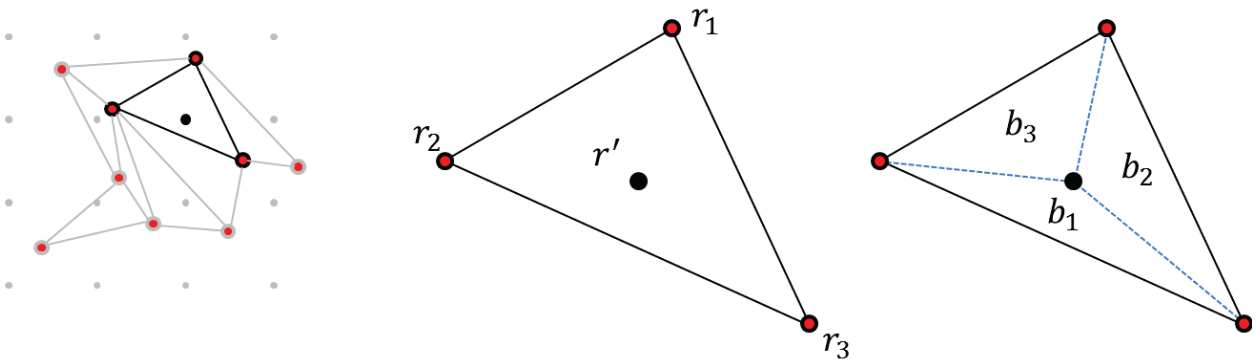

Figure 4.79: Interpolation of a point in a triangulated surface sub-division.
(Left) Starting point cloud and triangulated divisions. (Middle) Isolated triangulated sub-division. showing the location for the interpolated point. (Right) Barycentric coordinates.

The barycentric coordinates must satisfy the normalization condition:

$$\sum_{i=1}^{3} b_i = 1 \qquad (4.167)$$

The area of a triangle (A) is given by vectors $\vec{u}$ and $\vec{v}$ that originate from a vertex is:

$$A = \frac{1}{2}|\vec{u} \times \vec{v}| \qquad (4.168)$$

where $\times$ is the cross-product operation. Allowing for an arbitrary rotation of the triangle, a more generalized form for the area as a function of the three vertices is:

$$A = \frac{1}{2}|(\vec{r}_2 - \vec{r}_1) \times (\vec{r}_1 - \vec{r}_3)| \qquad (4.169)$$

The barycentric coordinate can then be determined from the area of each sub-triangle relative to the total area of the triangle:

$$b_i = \frac{A_i}{A} = \frac{\left|(\vec{r}_j - \vec{r}') \times (\vec{r}_k - \vec{r}')\right|}{|(\vec{r}_2 - \vec{r}_1) \times (\vec{r}_1 - \vec{r}_3)|} \qquad (4.170)$$

Where the subscripts j and k denote the opposite vertices from the barycentric coordinate. The interpolated sag value z' at the remapped coordinate r' is given by a weighted average of the sag values at each vertex, where the barycentric coordinates are the weighting values:

$$z' = \sum_{i=1}^{3} b_i z_i \qquad (4.171)$$

The interpolation method described in this section is applied to the set of data from Figure 4.77 and results in Figure 4.80.

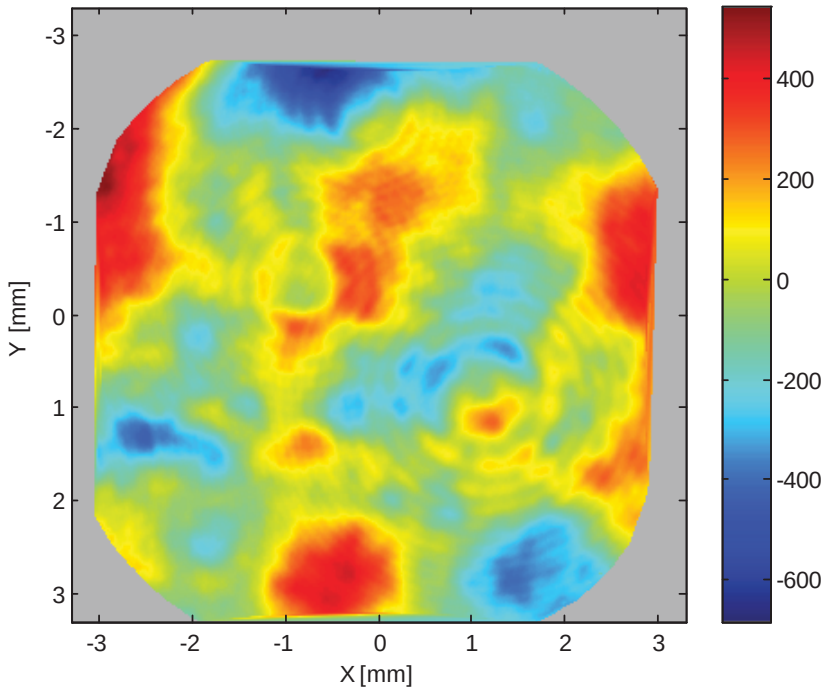

Figure 4.80: Interpolated surface height of the biconic-removed point cloud. Surface height units are in nm.

4.6 RECONSTRUCTION AND ANALYSIS SUMMARY

This chapter describes the process for reconstructing high resolution corneal topography maps from the set of interferometric measurements collected by SPINALTOP. It was demonstrated that raw interferometric data is insufficient for high resolution measurement of the corneal topography. Reverse ray tracing is required to reconstruct the corneal topography, which first requires a calibrated interferometer model. A reverse ray trace operation will result in a set of data that represents a high resolution mapping of the corneal topography. The shape of the cornea can be estimated by fitting a generalized biconic surface to the data. Surface fit residuals contain information on the tear film structure and are interpolated into a uniform grid. Corneal topography and tear film structure can be analyzed on a frame by frame basis to determine the dynamic characteristics. These methods will be applied to the results shown in the next chapter that presents from measurements on test surfaces and human subjects.

5 SPINALTOP RESULTS

This chapter presents the results from the SPINALTOP system. The results are divided into four sections. The first section covers the results from a calibration of the system. The calibrated system is verified against a set of reference surfaces and presented in the following section. The results are compared against measurements from a commercial corneal topographer system. Results from human subjects testing are presented in the next section. The final section discusses the difference between the SPINALTOP and TFI systems.

5.1 CALIBRATION

The details of the calibration process were previously discussed in Section 4.3. Calibration is effectively performed at a system-level; however, the results are divided between the surface and cat's eye interferometer sub-systems to highlight the overall importance of that sub-system to the to the calibration step. Reverse optimization is a subset of the calibration process that primarily operates on the surface interferometer (Section 4.3.2), which distinguishes it from the displacement calculation that is primarily a function of the cat's eye interferometer (Section 4.3.3).

5.1.1 SURFACE INTERFEROMETER CALIBRATION

5.1.1.1 SETUP

A set of eight measurements of the golden surface (Section 4.3.2.3) are used in the reverse optimization process. The first measurement is at null and defines the origin for the measurement set. The null OPD measurement is shown at the center of Figure 5.1. Six additional measurements are made by translating the golden surface by ±75 μm along each axis of the interferometer. The resulting lateral OPD measurements are shown in Figure 5.1 surrounding the null measurement, which correspond to the relative part translation (e.g. above the null measurement is the +Y measurement). The two axial translations (i.e. ±Z) are shown on the left and center in Figure 5.2. A final measurement is made at cat's eye, approximately 7.8 mm from the confocal position (Figure 5.2, right).

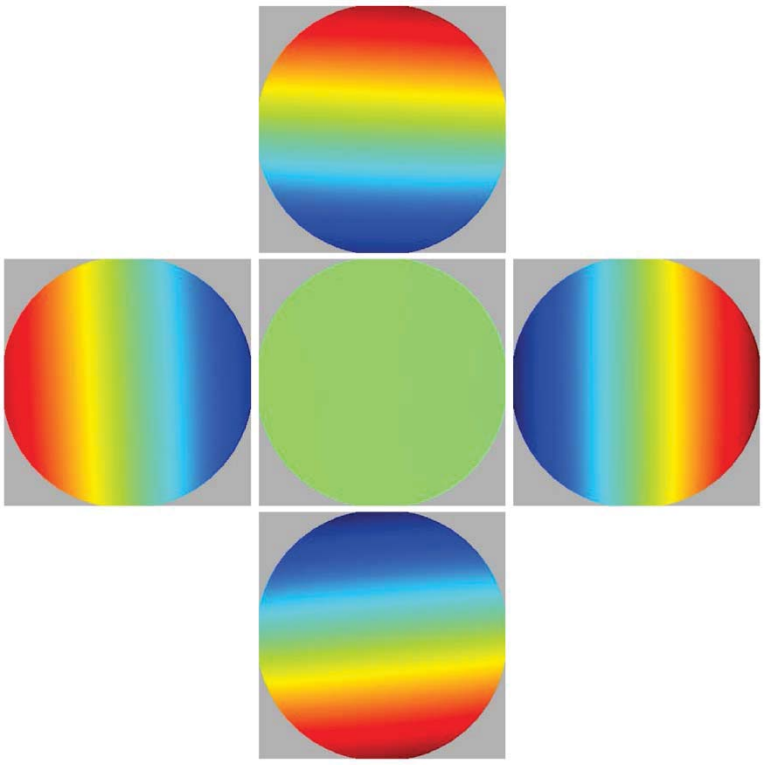

Figure 5.1: OPD measurements of golden surface used with reverse optimization. Center OPD is at null. Surrounding OPD measurements result from translating the golden surface by 75 μm away from null. ±75 μm color height scale.

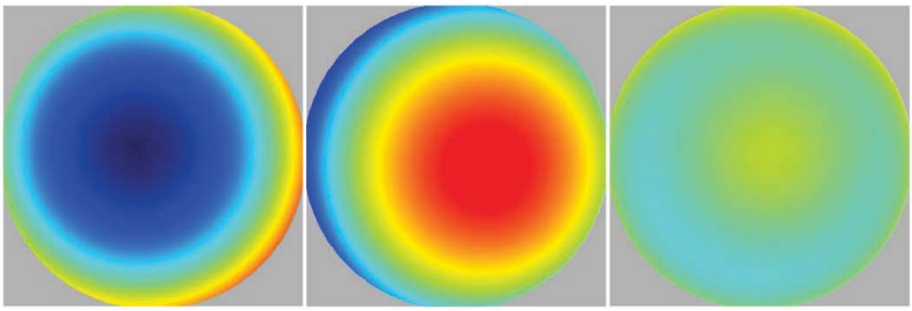

Figure 5.2: Additional OPD measurements of golden surface used with reverse optimization. Left and center OPDs are at ±75 μm defocused from null. Right OPD is a cat's eye measurement. ±12.5 μm color height scale.

The set of eight OPD measurements are passed to the reverse optimization merit function as targets. The optimization routine minimizes differences between modelled and measured OPDs by adjusting parameters within the model. The interferometer has been aligned such that the optical axis is maintained along the holes and plane of the breadboard. This sets the nominal angle between folds to 90 degrees.

Angular deviations for most surfaces have minimal impact on the reverse optimization process and therefore every optical element in the reference path, beamsplitter, test path dichroic, is assumed perfect in its alignment. Surface figure errors for a number of surfaces are assumed perfect, which has been verified through external measurements (Section 4.3.2.1). For example, the fold mirrors were verified to have surface figure errors less than λ/20 PV (λ=632.8nm) and contribute a small amount of error compared to the remaining residual errors that will be presented in this section. By making these assumptions, the reverse optimization and ray trace routines are greatly sped up. The remaining surfaces are mapped into the interferometer model by surface map deformations or Zernike polynomial representations.

The most sensitive optic in the interferometer is the converger assembly. An unfortunate consequence of reusing the converger assembly means that it cannot be fully characterized without dismantling the assembly. Because of the significant risk that would be involved in dismantling one of the most expensive components in the system, it has been left alone. During the reverse optimization process, it was initially assumed that the converger assembly was perfect. Alignment of the converger assembly relative to the interferometer is still optimized, but limited by the converger. Some of the variables within the converger were opened up to the reverse optimization routine, which showed improvement.

5.1.1.2 RESULTS

The initial iteration of the reverse optimization process provides feedback on the baseline interferometer model. It is useful for demonstrating the limitations of an un-calibrated system. It should also be noted that the use of a single surface for optimization does constitute a comprehensive evaluation of the system, which is why a separate section is devoted to system verification (Section 5.2). However, the residual errors that are presented in this section provide a valuable insight in the capabilities of the system. The residuals from a baseline system are shown in Figure 5.3, which are defined as the difference between the measured and simulated optical path differences. The order of residuals in the figure starting from the top-left and progressing to the bottom-right are: Null, +X, +Y, -X, -Y, +Z, -Z, and cat's eye. The P-V display

scale is 200 μm and χ^2 error of 123 μm (Eq. 4.102). The dominant error in all of the residuals appears mostly as defocus, which is indicative of a focusing error in the imaging system. It was discovered that a parameter in the imaging path was incorrectly entered (i.e. user error) and was fixed. This is not a fair example of an uncalibrated system, but has been left in as an example of how important a starting model can be.

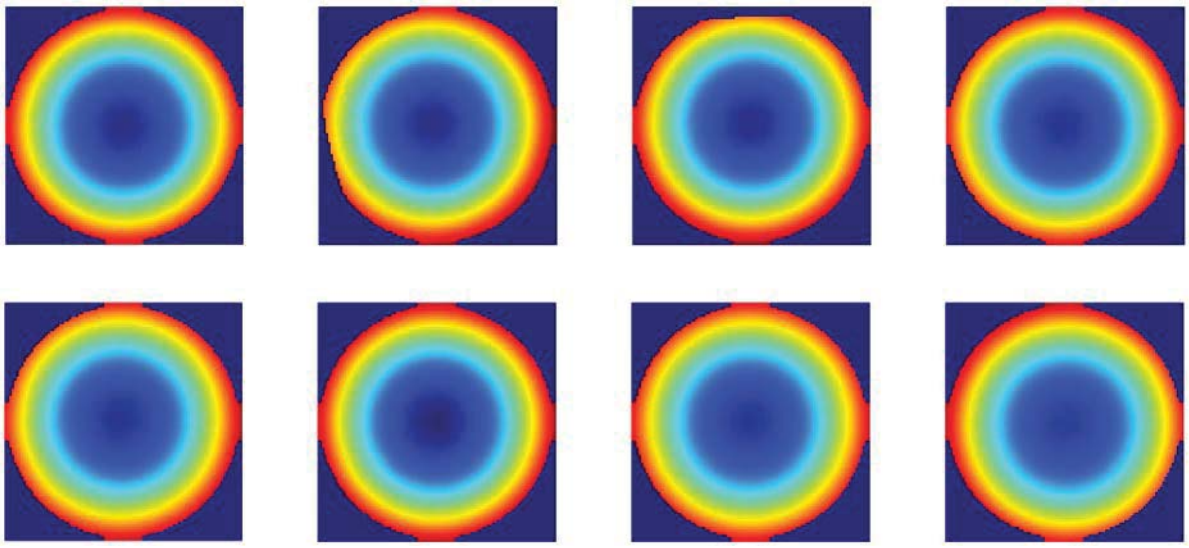

Figure 5.3: Baseline residuals for reverse optimization.
Residual order starting from top-left: Null, +X, +Y, -X, -Y, +Z, -Z, and cat's eye residuals. ±100 μm color height scale

The reverse optimization process was re-run once the model error was corrected. A more appropriate baseline set of residuals is shown in Figure 5.4. The errors are more reasonable and matched expectations for an un-optimized yet well aligned system. The display scale is set to 650 nm PV (±325 nm), with the cat's eye measurement saturating the color display scale. Because the reference mirror tilt has not been corrected yet in the model, the cat's eye measurement contained over 5 μm of error. Had the display been scaled to the cat's eye measurement, features in the other residuals would have been obscured. The χ^2 error for the baseline is 1248 nm. The residual amount of astigmatism present in the null measurement is characteristic of a converger tilt error. The remaining residuals were determined to come from a combination of the converger tilt and imaging system.

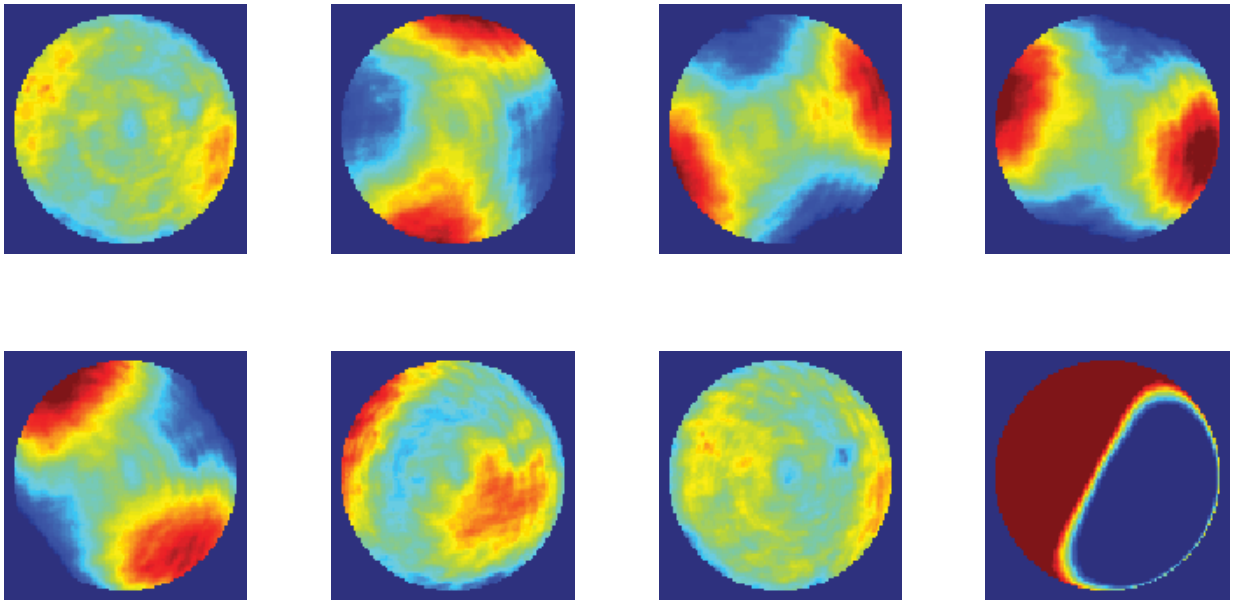

Figure 5.4: Corrected baseline residuals for reverse optimization.
±325 nm color height scale, χ^2 = 1248 nm. The cat's eye measurement exceeds 5 μm of error

The residual errors from an optimization run are shown in Figure 5.5. The display scale is set to 500 nm and the χ^2 error is 249 nm, with the cat's eye error still exceeding 700 nm of error. The residual errors are most apparent with the lateral translations of the golden surface, which is an issue. Random eye motion will result in measurements that are mostly dominated by tilt. These results indicate that a half-micron or more of random surface figure error will be uncorrected for during the retrace operation.

The residual errors are most likely due to the uncharacterized converger assembly. A total of five optical elements are contained in the converger assembly, which are the most sensitive optical elements in the interferometer system. Although a complete characterization of the converger is not possible at this time, parameters within the converger can be opened up to the optimization routine. This includes lens decenters, lens thickness, element spacing, and radius of curvature. Surface features, index homogeneity, and surface irregularity will have to be ignored; however, the dominant residual errors are low-spatial frequency structures which are most likely induced by lens decenter, thickness, and spacing errors.

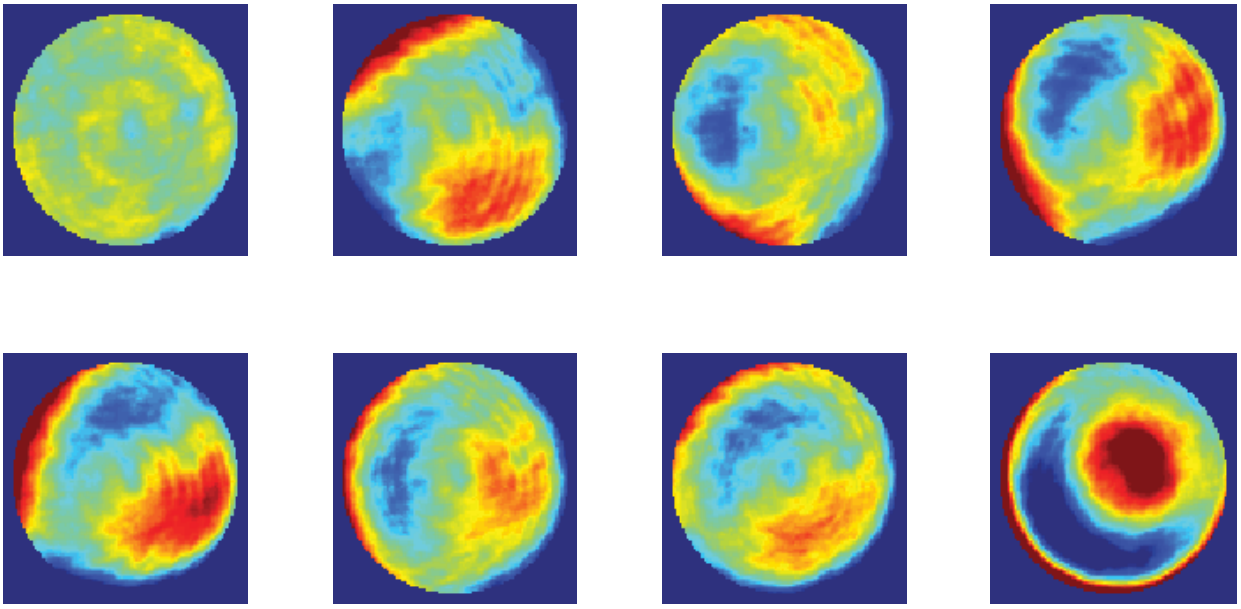

Figure 5.5: Reverse optimization residuals for initial run.
±250 nm color height scale, χ^2 = 249 nm. The cat's eye measurement exceeds 700 nm of error.

The time it takes for an optimization routine to converge on a solution is exponentially proportional to the number of variables. The pattern search method that is used for reverse optimization is slow, requiring a few hours to converge on a solution. The number of variables within the converger assembly is greater than the all of the currently free variables within the interferometer system. Opening up the full converger to optimization would increase the processing time to a few days or more. Therefore, only a few variables are opened up for optimization. The most sensitive optical element in the converger assembly is the second optical element (fourth lens from the eye). Because of the sensitivity of this element, the manufacturers placed adjusters on the fourth lens to act as a compensator. The converger assembly meets the requirements for a standard null interferometric test (i.e. on-axis) but non-null measurements (i.e. off-axis) performance is not guaranteed as can be seen by the initial results of the reverse optimization. In order to minimize the processing overhead for the optimization routine, decenter, thickness, and spacing for only these two lenses are opened up.

The OPD residuals that result from an optimization with the added converger variables are shown in Figure 5.6. The display scale is set to 160 nm and the χ^2 error is 98 nm. RMS figure errors for the

individual residuals are less than 35 nm. The second lens in the converger (the most sensitive) required 28 μm of decenter correction and 23 μm of thickness correction. The compensating element (fourth lens in the converger) required 64 μm of decenter correction and did not require a thickness correction. However, the radius on the last surface of the fourth lens (closest to the eye) required a 200 μm adjustment. The radius of curvature for this surface is 430 mm, resulting in an error of 0.046%. For reference, a typical high precision tolerance on a radius of curvature is 0.05% and center thickness is 25 μm. Therefore, the adjustments required for the elements in the converger assembly can be seen as reasonable. The imaging lens required a 200 μm radius of curvature correction and 100 μm thickness corrections, which are also consistent with commercial (i.e. mass produced) manufacturing tolerances. Finally, the imaging lens and camera system required decenter and axially translations that ranged from 0.5 – 1.0 mm. The magnitude of these values is not surprising considering how the system is assembled and aligned.

Many of the residuals have been improved through reverse optimization, but some still contain a small amount of low-spatial frequency errors. These errors are most likely due to the remaining un-optimized variables within the converger assembly. However, now that the interferometer model is converging on a solution, the mid- to high-spatial frequency errors begin to dominate the residuals. Some of these features come from within the converger, which can be verified by rotating the converger assembly and observing that they rotate with the converger. The remaining high-spatial frequency structure is due to stray light, diffraction, fringe print-through, and additional optical components in the interferometer that are left out of the model. A particular stray light feature, or spurious reflection, that prints through on all of the residuals can be seen in a background image (Figure 5.7) as a slightly off-axis, low power term. The vertical fringes and pattern of five high-powered fringes in a cross disappear when the aperture stop is placed in the system. Because most of the dominant residual errors are high frequency and the remaining low frequency terms are understood to be from the converger, the calibration process for the surface interferometer is considered complete.

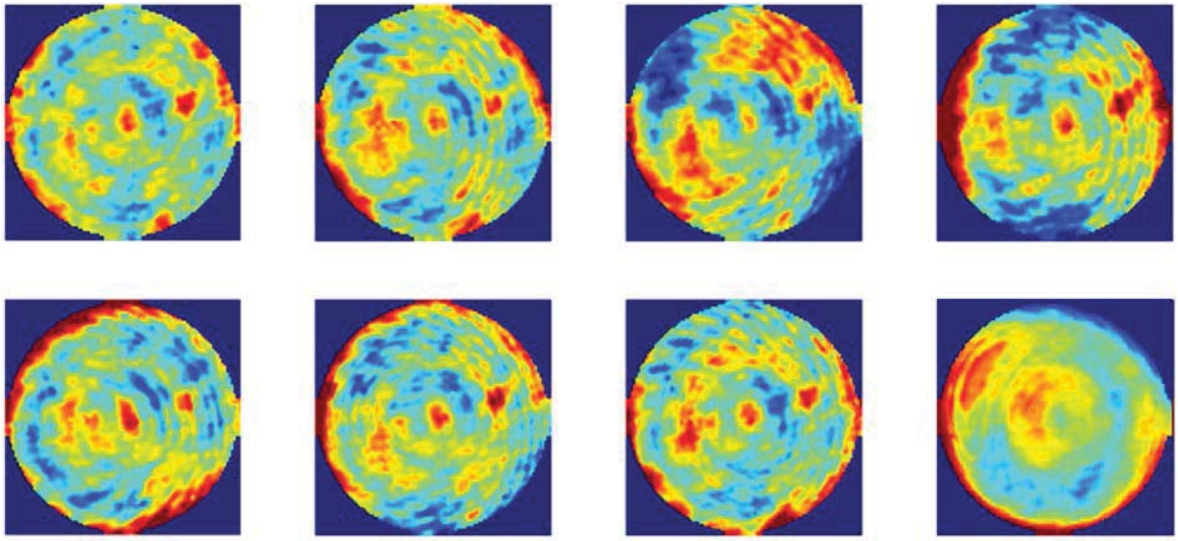

Figure 5.6: Reverse optimization residuals for final run.
±80 nm color height scale, $\chi^2 = 98$ nm.

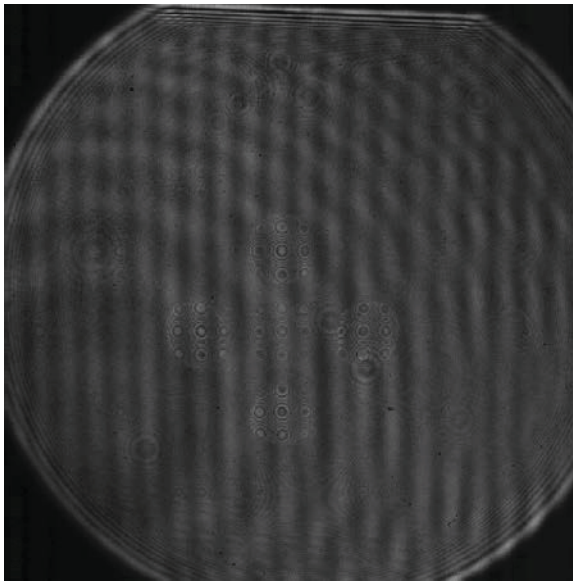

Figure 5.7: Residual stray light in the surface path.

5.1.2 CAT'S EYE CALIBRATION

Calibration of the cat's eye interferometer requires a set of surfaces with known surface shapes that have varying curvatures that can be used to calibrate the interferometer (Section 4.3.3). The set of surfaces that will be used for this calibration are shown in Figure 5.8. They are a set of diamond turned brass surfaces with varying curvatures. Each surface is aspheric with a -0.25 conic constant to ensure that the parts can

be reasonably nulled on the interferometer. The steepest radius of curvature is 7.0 mm, with each surface incrementing by 0.25 mm to a final surface with a 9.0 mm radius. The curvatures were selected to sufficiently probe the range of expected surface curvatures that would be tested on SPINALTOP (i.e. expected variation in corneal topography, Section 1.1). The radius of curvature for each surface was verified to within 10 μm on a Zygo Verifire with a 4" f/0.75 transmission sphere and a translation slide with a linear encoder.

Figure 5.8: Cat's eye curvature reference surfaces.

Each one of the curvature reference surfaces are aligned to the surface interferometer in a confocal position and nulled. The surface is measured with the SPINALTOP system; i.e. simultaneously with both the surface and cat's eye interferometers. The surface is translated by 75 μm along the X and Y axis and measured for a total of five measurements, including the null measurement. The surface is translated by 75 μm along the Z axis, and the five translated sets of measurements are repeated. This process is repeated for a total of 25 points, 5 X/Y points at each of the 5 Z locations.

The absolute position of each test surface is determined by a two-step process. The first is tracking the relative stage positions, which are accurate to a few microns. The stage coordinates, surface information, and surface measurements are passed to a routine that optimizes the location of the surface. This routine is the same routine that was discussed in Section 4.3.2.4.2 that was used to fine tune the absolute positioning of the reference surfaces relative to the interferometer. The optimization routine for this step requires a

calibrated interferometer model that was discussed in the previous section. Zernike polynomials are fit to the cat's eye and surface measurements and the coefficients are saved. The set of coefficients and displacement values are used to solve for the displacement polynomial coefficients (Section 4.3.3). The displacement polynomial can then be used for the reverse ray trace operation as described in Section 4.4.1.4.

A set of Zernike coefficients that were measured during a calibration are shown in Figure 5.9. The left image contains the Zernike defocus coefficient from the cat's eye interferometer as a function of displacement of the reference surface. The linear relationship is expected, based on Eq. 4.112. The slope of the line is approximately 10.9 wv/mm = 0.00927. Using Eq. 4.112, the f/# is estimated to be f/5.2, which is reasonably close to the expected working f/# of the cat's eye interferometer. The vertical spread in the measurements was expected as discussed in Section 4.3.3 (i.e. Eq. 4.112 is insufficient to characterize axial displacement).

The second set of data shown on the right in Figure 5.9 is a plot of the Zernike defocus term for the surface interferometer measurement as a function of displacement of the reference surface. This data shows a completely different relationship between the Zernike coefficients and the axial displacement. However, this data provides additional information which shows the dependence on the surface shape. The nine surfaces can be isolated in the data by the 9 vertical clusters of data. Each vertical cluster of data corresponding to a surface consists of five data sets. Each of these five data sets contains five points. These five by five sets of data correspond to the axial and lateral translations of the reference surface. This data provides additional unique information that will be used to support the surface displacement calculation.

The combination of information from these two examples and the additional Zernike coefficients that are calculated during the calibration measurements can be used to determine the displacement polynomial to track the axial displacement of the test surface (Eq. 4.114). The residual error from a polynomial fitting

with the calibration data is shown in Figure 5.10. The set of 9 surfaces with 25 measurements each provides a total of 225 data unique data points. The error is determined by comparing the polynomial derived displacement of each surface to the known location. The polynomial residual error is 1.5 μm RMS and 8.4 μm PV error.

Although it is difficult to relate the error in determining the axial displacement of the test surface to the absolute surface reconstruction due to the non-linearity of the reverse ray tracing process, an analogous result does exist. The cat's eye interferometer measurement is equivalent to a standard cat's eye measurement that is performed by axially translating a surface from confocal to cat's eye. This information is generally used to determine the radius of curvature of the test surface. Therefore, it could be said that the cat's eye interferometer can resolve the radius of curvature to 1.5 μm RMS. However, the reverse ray tracing process that will be used to reconstruct the absolute surface topography of the cornea applies to any arbitrary surface shape, not just spheres. Additionally, the methods described in this dissertation do not require an axially translation of the surface (e.g. due to lateral eye motion) nor does it require that the cat's eye and surface measurements are null.

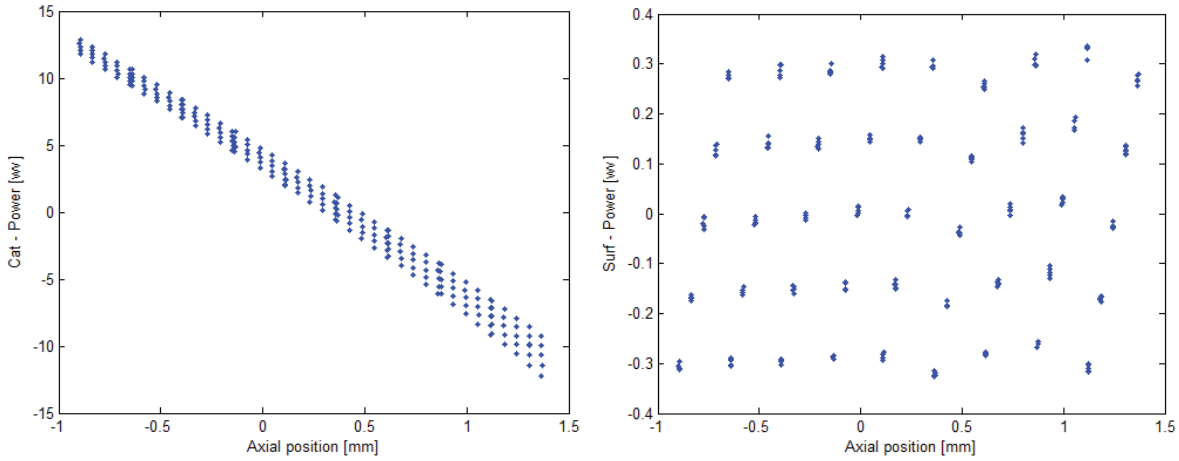

Figure 5.9: Zernike polynomial coefficients vs. reference surface displacement.
(Left) Cat's eye Zernike defocus coefficient. (Right) Surface Zernike defocus coefficient.

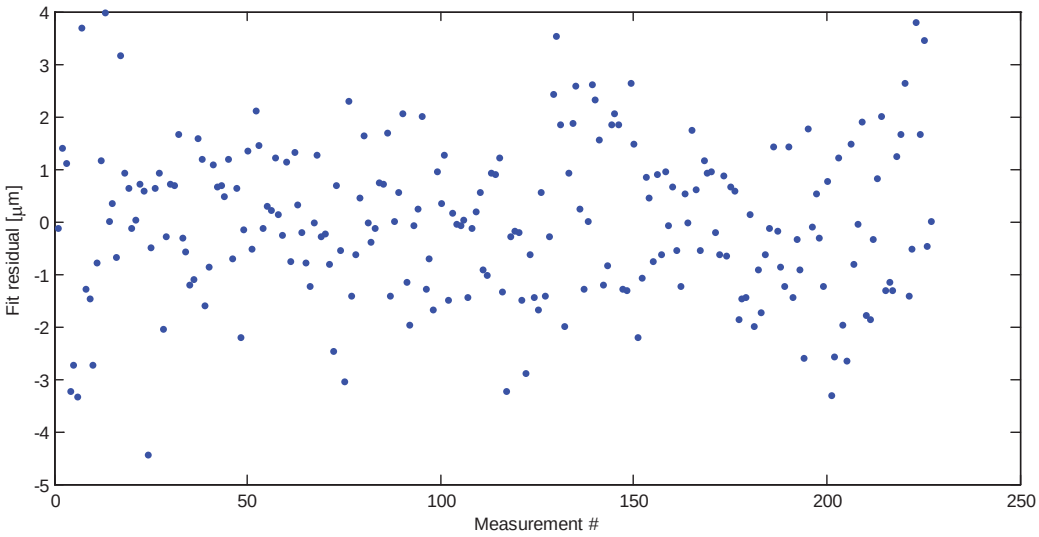

Figure 5.10: Displacement polynomial fit residuals.
Residual error is 1.5 µm RMS and 8.4 µm PV.

5.2 VERIFICATION

5.2.1 SETUP

The SPINALTOP system can be verified by measuring a set of known surfaces and comparing them to the reconstructed surface results. The verification surfaces should simulate the expected range of test surfaces that will be measured by the system. The primary error source will be retrace errors, which are induced through low-spatial frequency errors in the interferometer (e.g. lens decenters). Mid- to high-spatial frequency errors will also be of interest, so the surface finish will need to be good. Misalignment of a test surface with respect to the interferometer can introduce a range of low to mid spatial-frequency errors, but a variety of surface shapes would be more appropriate. The curvature reference surfaces for the cat's eye interferometer system are re-used for verification. Because these surfaces are a part of the system calibration, an independent set of surfaces is also used. These surfaces are diamond turned brass surfaces of various aspheric and toroidal shapes that differ from the cat's eye calibration surfaces. The set of verifications surfaces are shown in Figure 5.11 with the parameters for all of the brass surfaces listed in Table 5.1.

Figure 5.11: SPINALTOP verification surfaces.

Part ID	Rx	Ry	Kx	Ky	Usage
001	-7.00	-7.00	-0.25	-0.25	Cat's eye calibration
002	-7.25	-7.25	-0.25	-0.25	
003	-7.50	-7.50	-0.25	-0.25	
004	-7.75	-7.75	-0.25	-0.25	
005	-8.00	-8.00	-0.25	-0.25	
006	-8.25	-8.25	-0.25	-0.25	
007	-8.50	-8.50	-0.25	-0.25	
008	-8.75	-8.75	-0.25	-0.25	
009	-9.00	-9.00	-0.25	-0.25	
010	-7.00	-7.00	-0.25	-0.25	SPINALTOP verification
011	-8.00	-8.00	-0.50	-0.50	
012	-8.00	-8.00	0.00	0.00	
013	-8.00	-8.00	0.50	0.50	
014	-8.00	-8.00	1.00	1.00	
101	-7.00	-7.15	-0.25	-0.25	
102	-8.00	-8.20	-0.25	-0.25	
103	-8.80	-9.00	-0.25	-0.25	
104	-8.00	-8.10	-0.25	0.25	
105	-8.00	-8.10	-0.25	0.50	
015C	-7.80	-7.80	-0.25	-0.25	Distortion calibration
015D	-7.80	-7.80	-0.25	-0.25	

Table 5.1: Brass surface parameters.

The surfaces were measured on the WYKO 6000 with a 4” f/0.75 transmission sphere to verify that surface errors were minimal and that the basic shape was correct. Part 010 was found to be incorrectly manufactured; it was originally intended to be an asphere with a -8.0 mm radius of curvature and -1.00 conic constant. However, instead it appears to be a duplicate of part 001. Two measurements of the brass surfaces are shown in Figure 5.12, Part 005 is on the left and Part 104 is on the right. The color-height scale is set to 50 nm (±25 nm). The display is approximately 5 x 5 mm, which was limited by the interferometer to resolve fringes for the aspheric surfaces. The first 25 Zernike (Fringe) polynomial terms have been removed from the measurements. Each surface shows a small manufacturing defect near the center. Additional tooling marks can be seen in all of the parts, such as grooves or “rings” formed by the diamond turning process. The absolute shape was not verifiable with equipment that was available at the time. However, the manufacturer of the surfaces manufactured a pair of spherical surfaces before and after manufacturing the verification surfaces. These surfaces were used to verify that the machine was operating within tolerance.

It was later discovered that surface reconstruction errors were slightly higher for the non-rotationally symmetric verification surfaces. Inspection of the residual errors revealed the presence low-spatial frequency features that could not be explained by an incomplete calibration (Section 5.2.2.4). This led to a review of the WYKO measurements to see if the features were visible on the toric verification surfaces that may have been obscured by the removal of the first 25 Zernike terms. Instead, a subset of terms were removed: Tilt (Z1-2), power (Z3), primary and secondary spherical (Z8, Z15), and primary and secondary astigmatism (Z4-5, Z11-12). The results are shown in Figure 5.13, Figure 5.14, and Figure 5.15. The residual features that appear in the WYKO measurements appear to correlate with the SPINALTOP reconstruction residuals that will be presented in Section 5.2.2.4. Additionally, part 102 (Figure 5.13, right) contains a circular artifact that is located approximately 1.5 to 2.0 mm from the center, which is the most likely case for a manufacturing error. This artifact results in a step discontinuity that separates the inner 3 to 4 mm diameter area to the outer area. It cannot be determined from this data if

the step is larger than half a wavelength due to the 2π ambiguity that results from the recovery of the phase (refer to Section 4.2). These residual surface errors will negatively impact the results from the verification of the SPINALTOP system.

Ideally, these features could be removed from the final results or included as part of the surface description. However, this would require a more involved characterization of the surfaces that is not currently available at the time. For example, the surface area that was covered by the WYKO was approximately 5 x 5 mm. These surfaces would need to be characterized out to 7 x 7 mm for SPINALTOP. Additionally, absolute surface shape information would be required which cannot be accurately determined on the currently available systems in the lab (e.g. radius of curvature, conic constant). The errors created by these surfaces will be noticeable, but can be tolerated for the current instrument demonstration. Future iterations of the SPINALTOP system that improves the accuracy and resolution of the system will require a more accurate characterization of these verification surfaces.

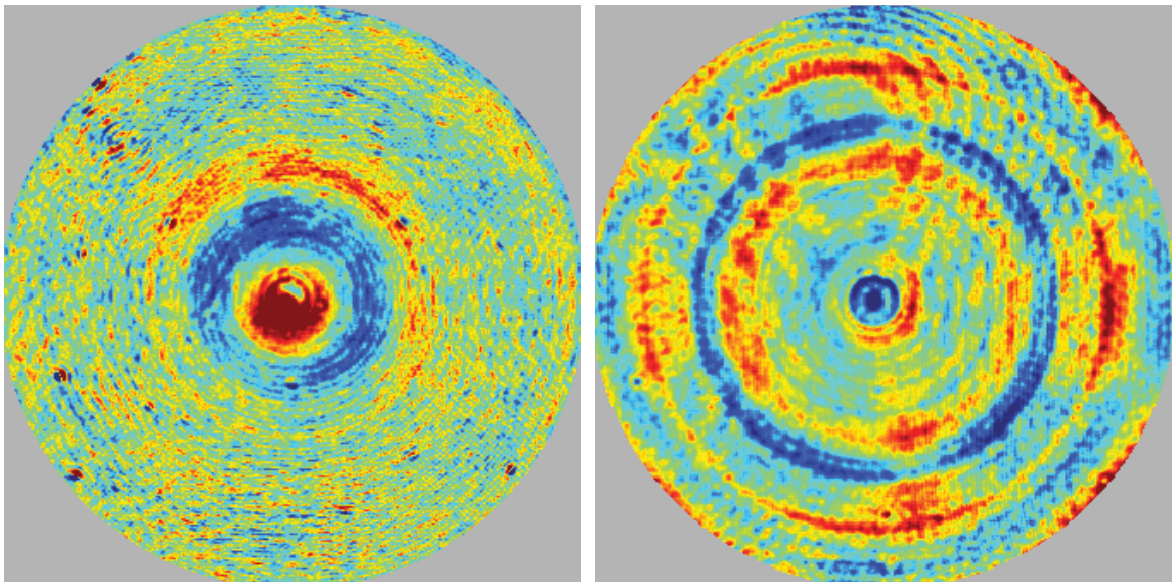

Figure 5.12: Brass surface residual errors as measured on the WYKO 6000.
(Left) Part 005. (Right) Part 104. ±25 nm color height scale. 24 Zernike terms removed. Displayed area is approximately 5 x 5 mm.

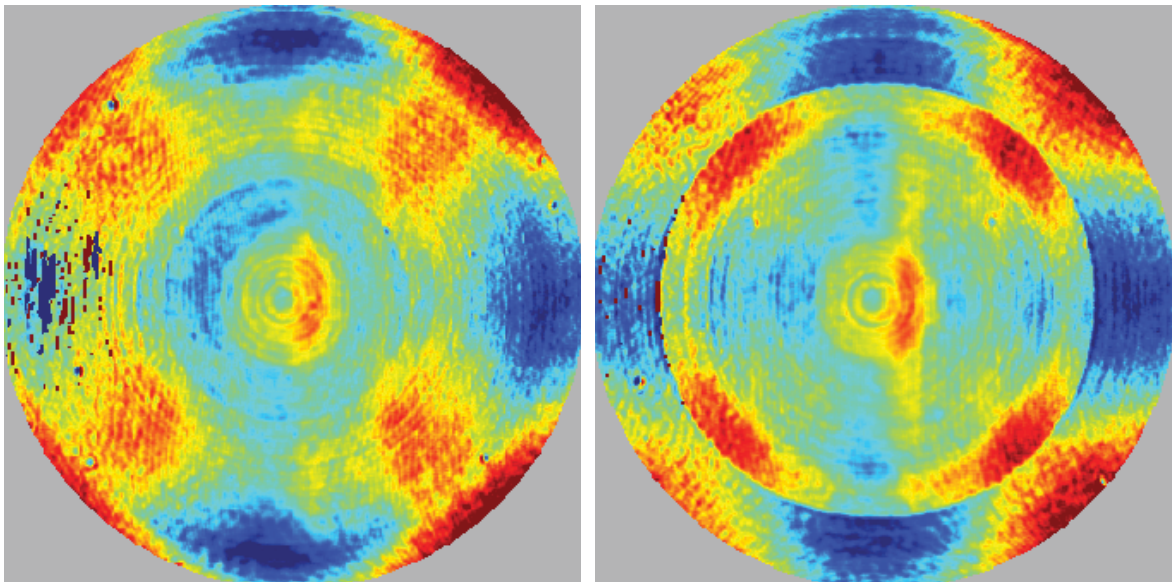

Figure 5.13: Brass surface manufacturing errors, measured on the WYKO 6000.
(Left) Part 101. (Right) Part 102 – note the additional discontinuity with a radius of approximately 2 mm.
±75 nm color height scale. 10 Zernike terms removed. Displayed area is approximately 5 x 5 mm.

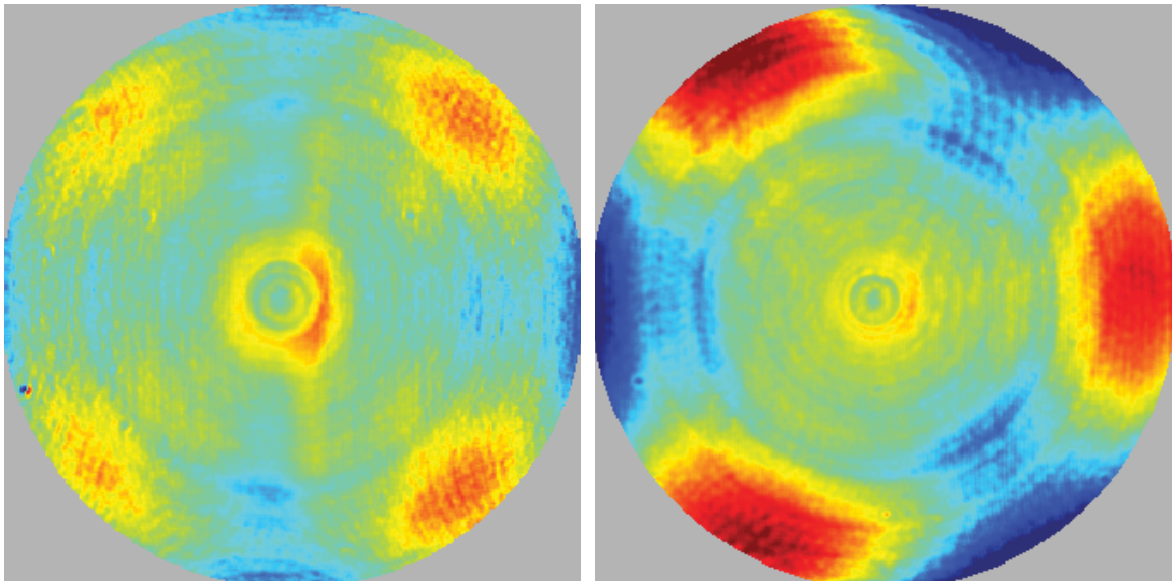

Figure 5.14: Brass surface manufacturing errors, measured on the WYKO 6000.
(Left) Part 103. (Right) Part 104. ±120 nm color height scale. 10 Zernike terms removed. Displayed area is approximately 5 x 5 mm.

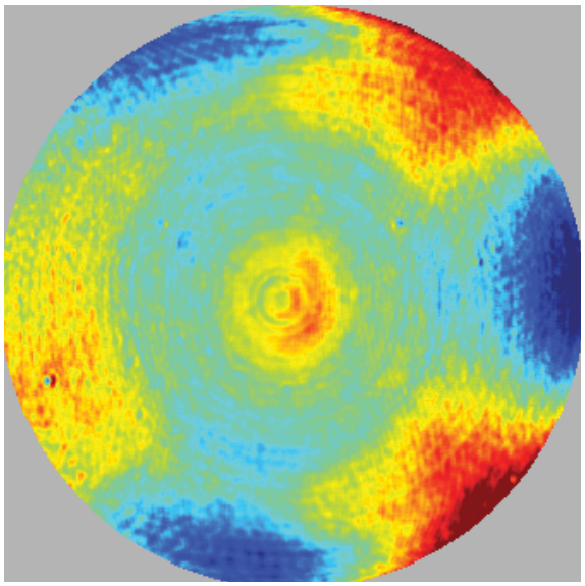

Figure 5.15: Brass surface manufacturing errors, measured on the WYKO 6000. Part 105. ±100 nm color height scale. 10 Zernike terms removed. Displayed area is approximately 5 x 5 mm.

The brass surfaces are tested on SPINALTOP in the same configuration as the cat's eye calibration (Section 5.1.2): A set of 25 measurements that span a range of lateral and axial translations. This will test the system's ability to measure a range of surface shapes that includes large random motions that will occur with human subjects testing when eye motion is present.

In addition to measuring the set of brass surfaces on SPINALTOP, each surface was measured with a commercial corneal topographer. The topographer is a Keratron Piccolo (Optikon, Roma, Italy) and uses a Placido disc to recover topography (Section 1.6.2). The surfaces were translated and measured with the Keratron Piccolo in the same configuration as the SPINALTOP measurements, with the exception of axial motion. The Keratron Piccolo uses a light trip-line to detect when the surface passes a specific location on the instrument and captures the data. A measurement of an axially displaced surface would violate the intended operating conditions of the instrument and is therefore not be included in this analysis.

5.2.2 RESULTS

Results from measurements on the brass surfaces for SPINALTOP and the Keratron Piccolo are presented in this section. The overall result is that SPINALTOP performance is nearly an order of magnitude better than the Keratron Piccolo instrument. A summary of the calculated accuracies is listed in Table 5.2 in the format of Bias ± 1σ. A few outliers have been excluded from the summary, which will be discussed in the following sections.

Accuracy	**SPINALTOP**	**Keratron Piccolo**
Radius [µm]	3.1 ± 4.2	-25.1 ± 50.8
Conic	0.0001 ± 0.0091	0.0413 ± 0.0653
Power [D]	-0.0166 ± 0.0224	0.1344 ± 0.2673
Astigmatism [D]	0.0010 ± 0.0086	-0.0732 ± 0.4597
Residual fit RMS [nm]	25.2 (<112.3)	305.5 (<886.8)
Residual fit PV_{98} [nm]	129.1 (<607.9)	1503.7 (<4339.1)
Spatial resolution [µm]	6	>100

Table 5.2: Summary of SPINALTOP and Keratron Piccolo performance.

5.2.2.1 BOX AND WHISKER PLOTS

Measurement results are presented graphically in this section using box and whisker plots. Two different forms of these plots are used. Figure 5.16 is a box and whisker plot showing the error in the estimation of the radius of curvature for the cat's eye reference surfaces (Parts 001 – 009). The raw set of topographic data is taken from a system and fit with a generalized biconic surface (Eq. 4.162). The fitted parameters (e.g. radius) are then compared to the expected parameters. The box height represents the root-mean-square (RMS) error in the estimation of the parameter and is centered about the mean error, or bias, of the measurement. The whiskers represent the maximum and minimum estimation errors for that parameter. The same set of box and whisker plots are used to show the error for the conic constant (Figure 5.17) and corneal power (Figure 5.18), where corneal power is defined as:

$$\Phi = \frac{n_k - 1}{R} \qquad (5.1)$$

where $n_k = 1.3375$ which is the keratometric index of refraction; an effective index used when describing the total power of the cornea resulting from an anterior surface measurement and accounting for the negative power introduced by the posterior surface.

The biconic fit residuals are described by absolute errors and require a modified box and whisker plot. Figure 5.19 shows the peak-to-valley error on the 98th percentile of data (stylized as PV_{98}), which is commonly used to minimize the impact of outliers. The box height represents the mean value of the PV_{98} errors and the whisker end represents the maximum PV_{98} error for that set of surface measurements. The same plot is used for the RMS residual errors shown in Figure 5.20.

Five plots are presented for the cat's eye curvature reference surfaces: Radius (Figure 5.21), conic constant (Figure 5.22), corneal power (Figure 5.23), biconic fit residual RMS (Figure 5.24), and a biconic fit residual PV_{98} (Figure 5.25). The same four plots are repeated for the SPINALTOP verification surfaces. A final plot is included at the end, which is the error in estimating the corneal astigmatism

(Figure 5.26). Corneal astigmatism is defined as the difference in corneal power along the meridional planes.

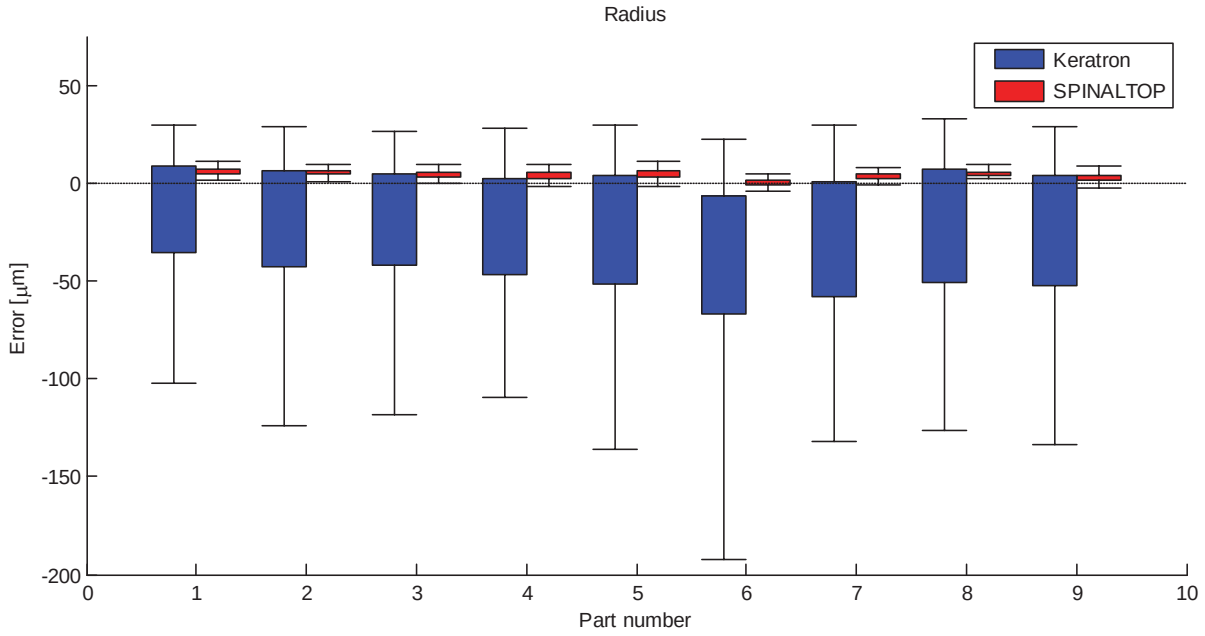

Figure 5.16: Radius estimation errors with cat's eye curvature reference surfaces. Errors are in microns.

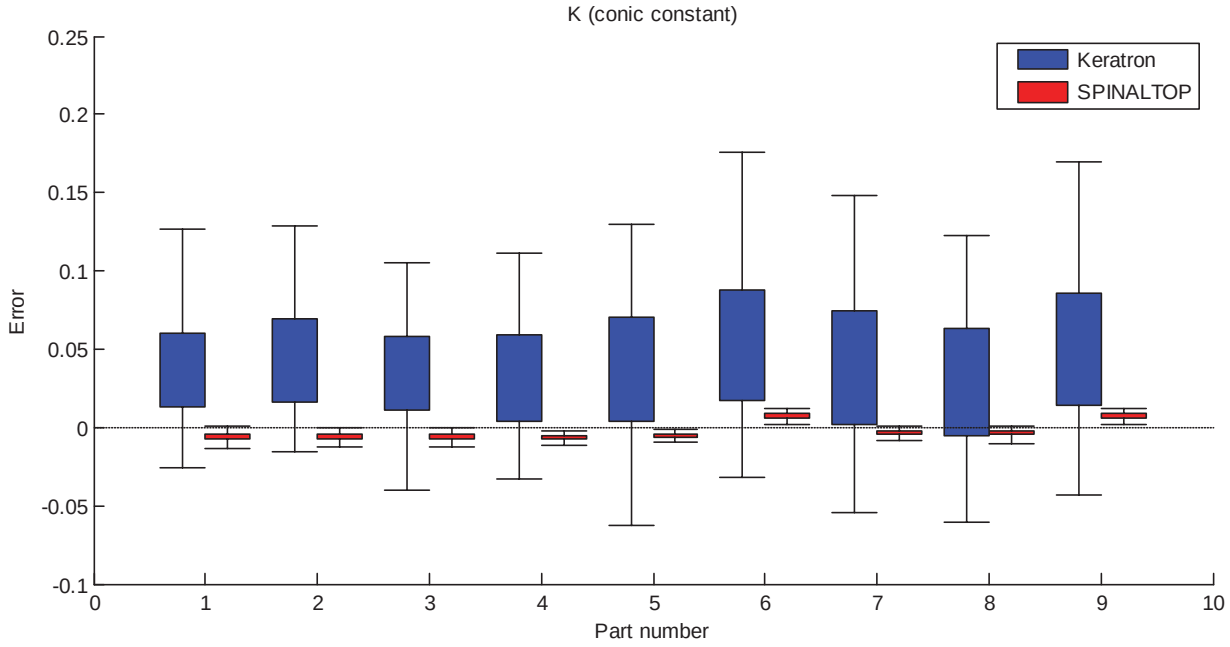

Figure 5.17: Conic constant estimation errors with cat's eye curvature reference surfaces. Errors are unitless.

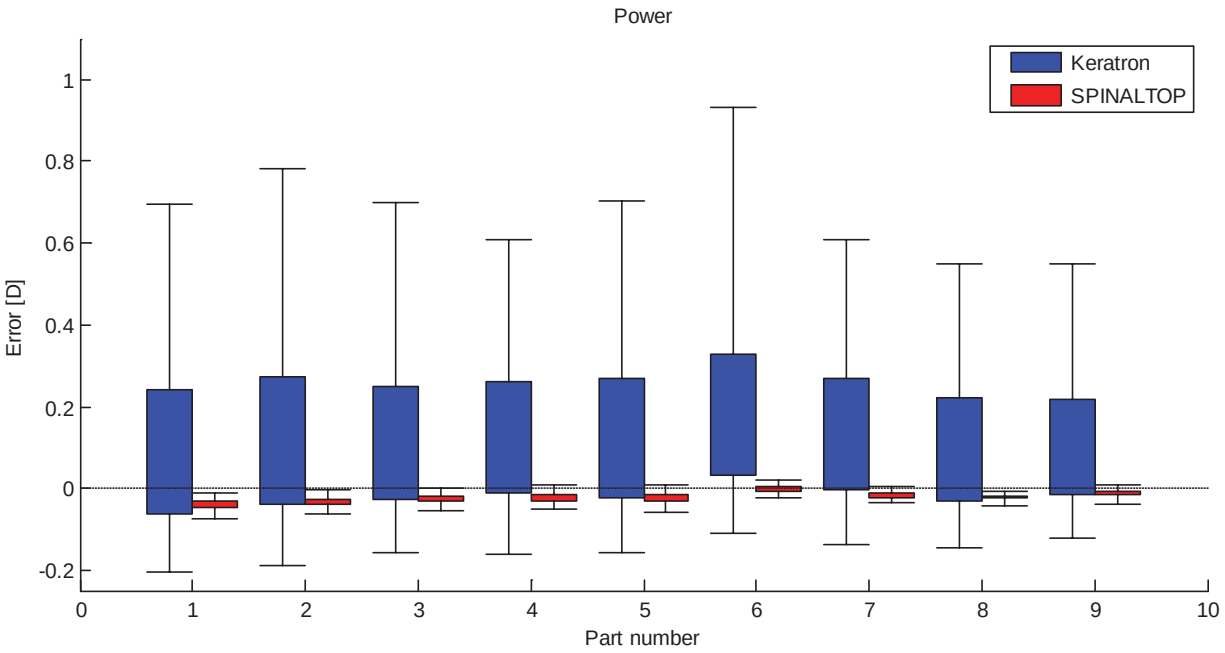

Figure 5.18: Corneal power estimation error with cat's eye curvature reference surfaces. Errors are in diopters.

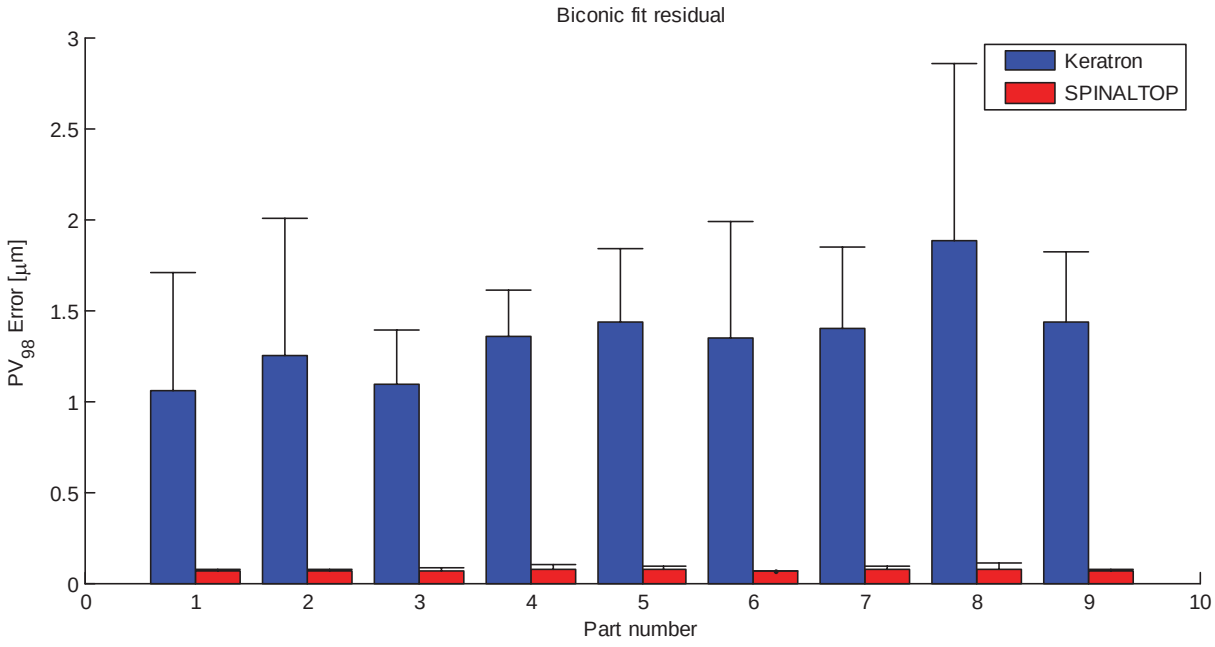

Figure 5.19: Biconic fit, PV_{98} residual errors with cat's eye curvature reference surfaces. Errors are in microns.

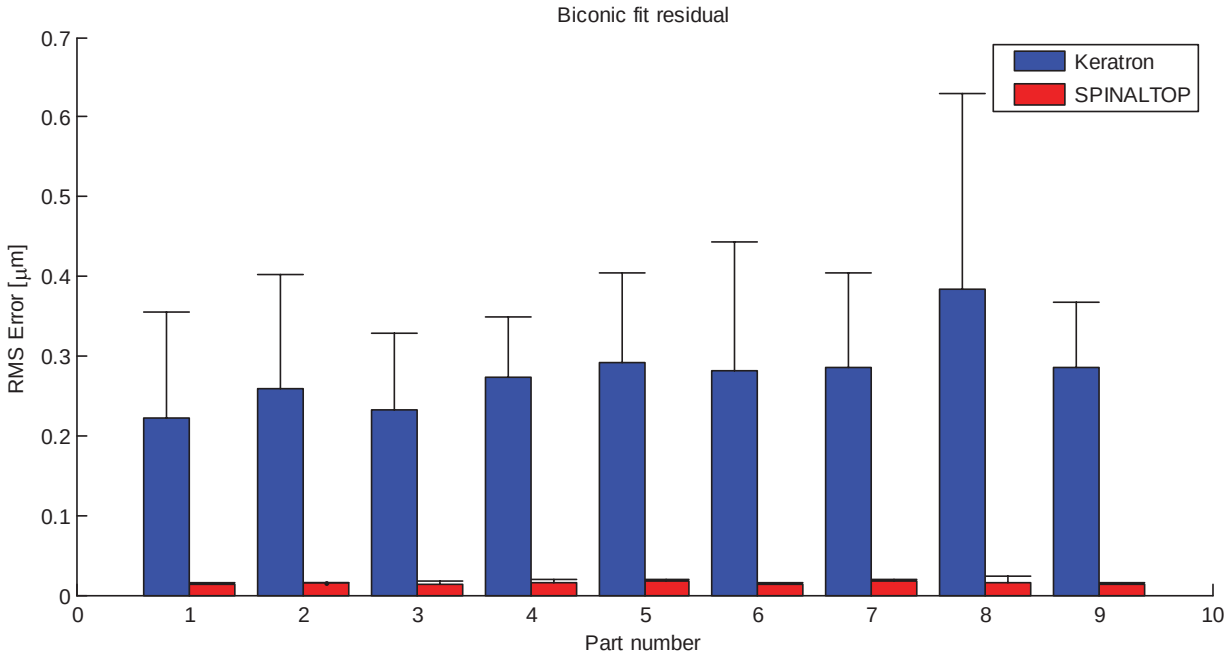

Figure 5.20: Biconic fit, RMS residual errors with cat's eye curvature reference surfaces. Errors are in microns.

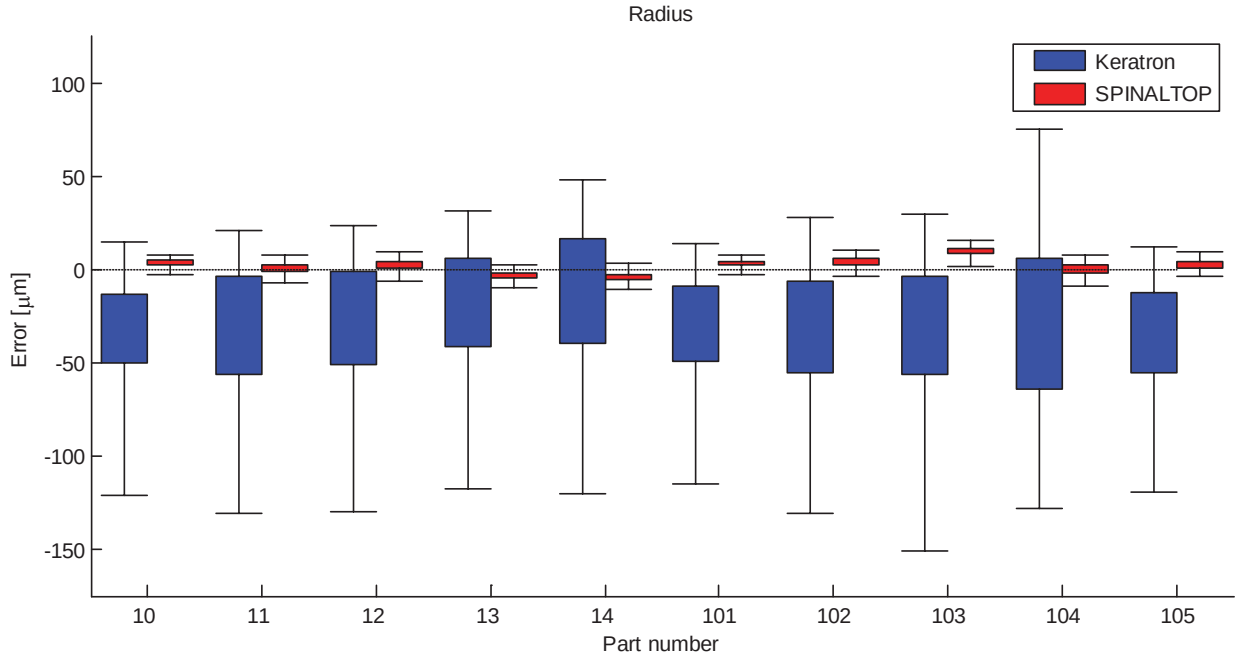

Figure 5.21: Radius estimation errors with SPINALTOP verification surfaces. Errors are in microns.

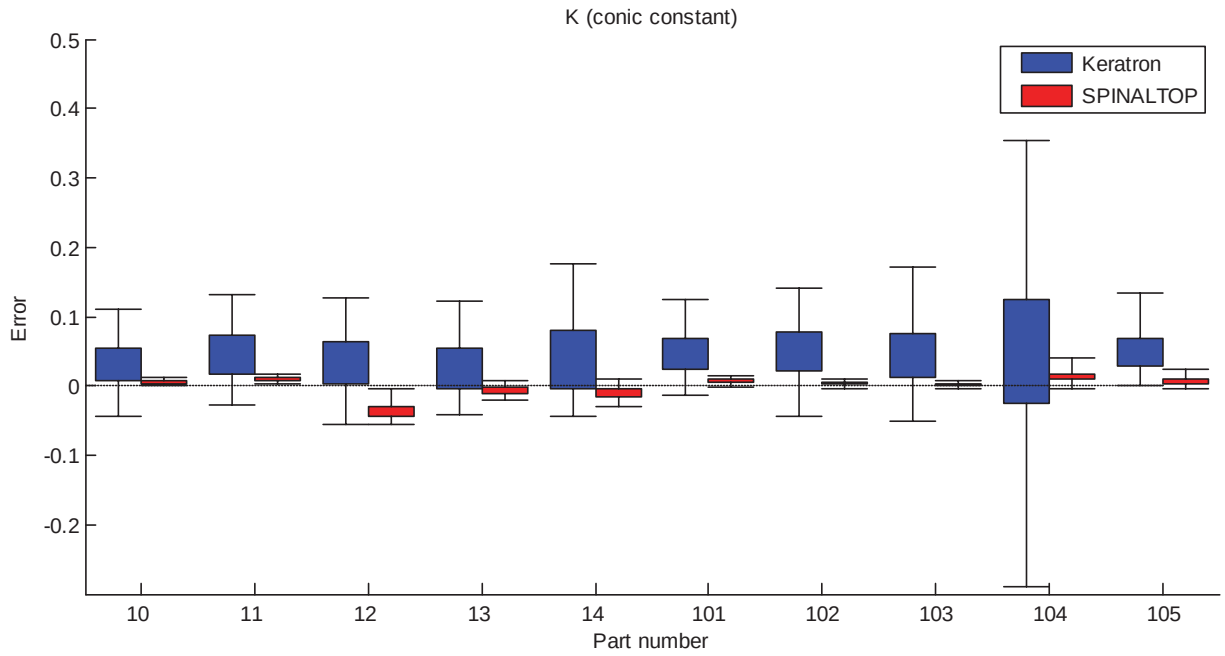

Figure 5.22: Conic constant errors with SPINALTOP verification surfaces. Errors are unitless.

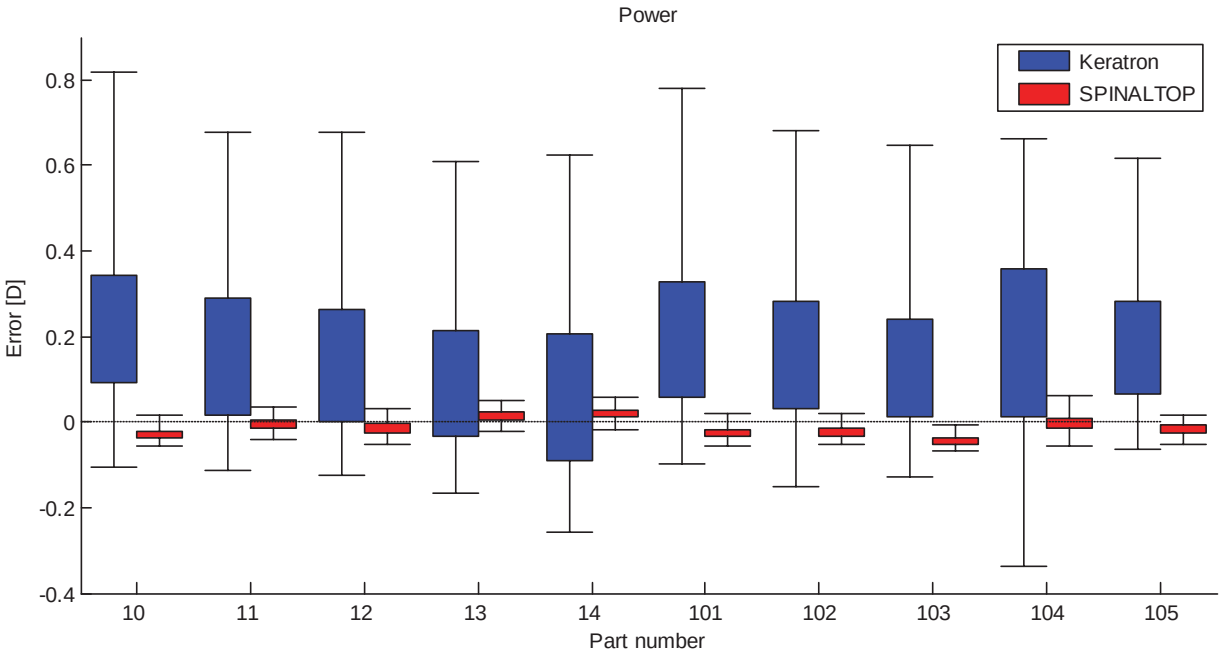

Figure 5.23: Corneal power estimation error with SPINALTOP verification surfaces. Errors are in diopters.

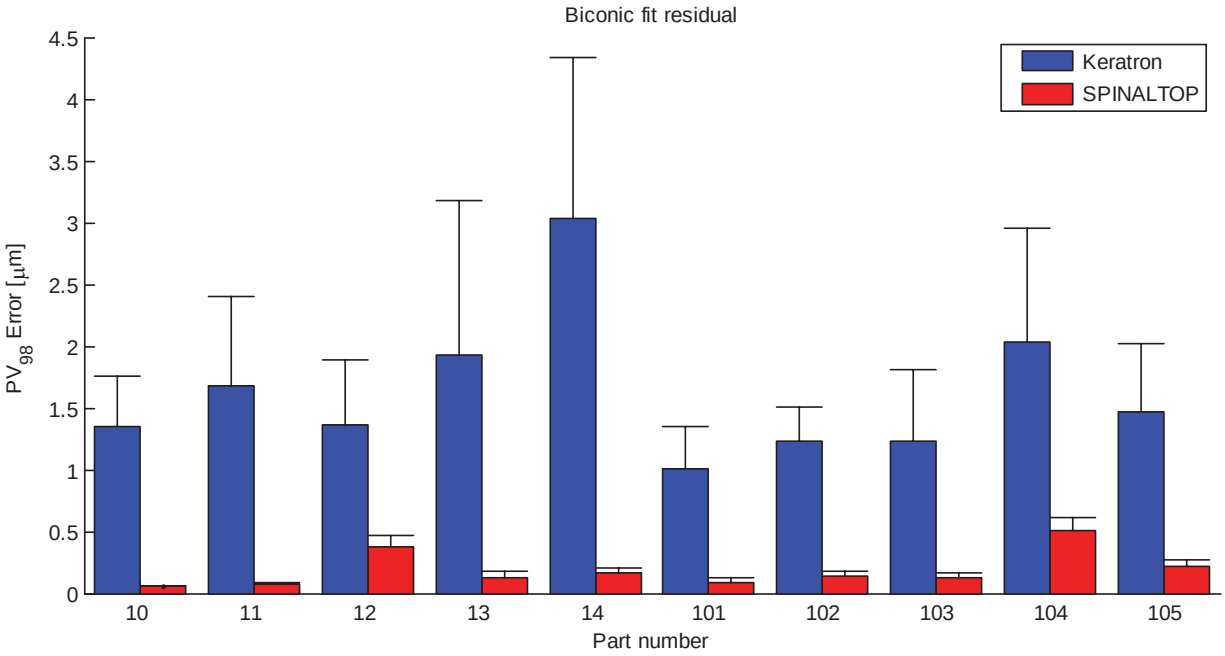

Figure 5.24: Biconic fit, PV_{98} residual errors with SPINALTOP verification surfaces. Errors are in microns.

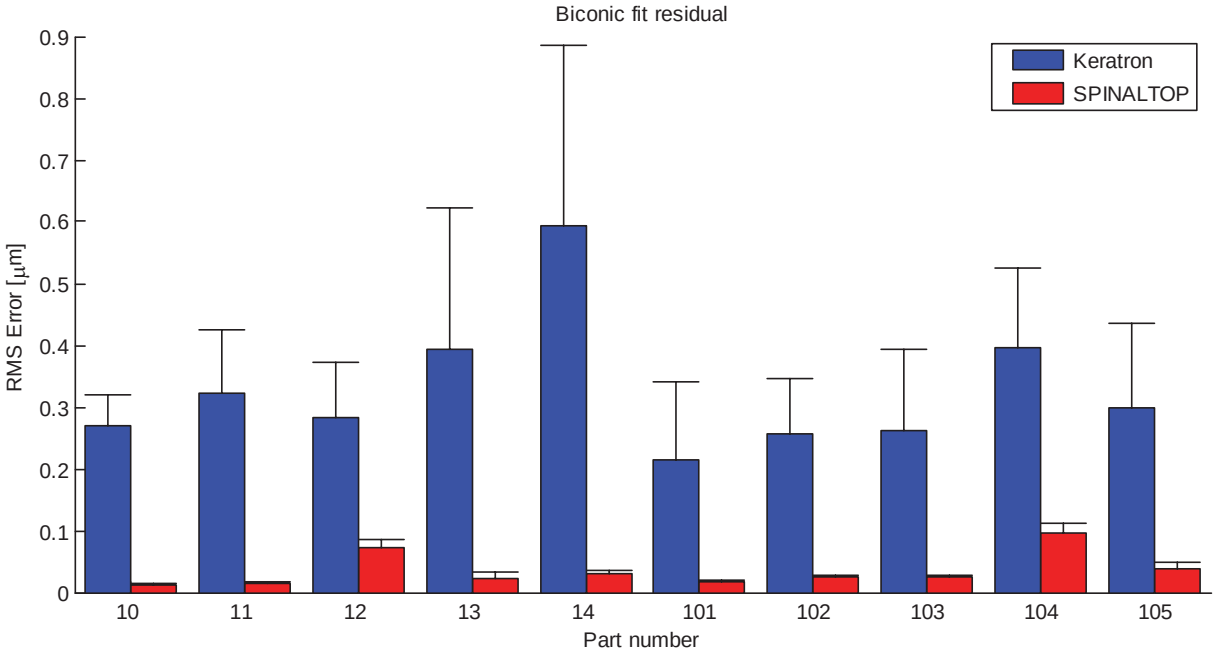

Figure 5.25: Biconic fit, RMS residual errors with SPINALTOP verification surfaces. Errors are in microns.

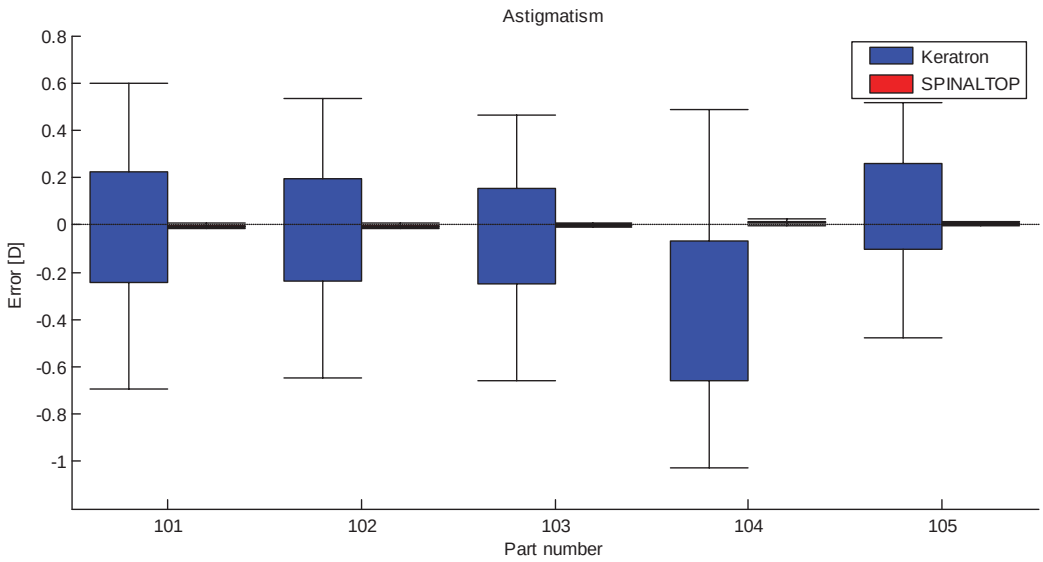

Figure 5.26: Corneal astigmatism estimation errors with biconic reference surfaces. Errors are in diopters.

5.2.2.2 TABULATED RESULTS

This section contains the tabulated results that correspond to the box and whisker plots from the previous section. Table 5.3 contains a listing of the errors for the estimation of the radius of curvature for both the cat's eye curvature surfaces and the SPINALTOP verification surfaces. The bias, standard deviation, root-mean-square (RMS), and maximum error are included. Table 5.4 contains a similar listing of the errors in estimating the conic constant errors, and Table 5.5 contains the errors in estimating the corneal power. Table 5.6 lists the biconic fit residuals, which includes RMS and PV_{98} errors. Finally, Table 5.7 contains the listing of the errors in estimating the amount of corneal astigmatism.

Part Number	Bias [μm]		Stdev [μm]		RMS [μm]		Max [μm]	
	SPINALTOP	Keratron	SPINALTOP	Keratron	SPINALTOP	Keratron	SPINALTOP	Keratron
1	5.6	-13.2	2.1	44.1	6.0	44.7	10.9	102.5
2	5.3	-18.6	2.1	49.2	5.7	51.1	9.6	123.9
3	4.1	-18.6	2.5	46.2	4.8	48.5	9.3	118.3
4	3.9	-22.5	2.6	49.0	4.6	52.5	9.3	109.7
5	4.5	-23.8	3.1	55.6	5.4	58.9	11.0	135.9
6	0.4	-36.7	2.1	60.6	2.1	69.2	4.5	192.6
7	3.5	-28.7	2.5	59.3	4.3	64.2	7.7	132.2
8	4.8	-22.0	1.5	57.6	5.0	59.9	9.3	126.3
9	2.4	-24.4	2.3	56.6	3.3	60.0	9.0	133.6
10	4.2	-31.8	2.2	37.1	4.7	48.0	8.1	120.9
11	0.6	-29.7	3.6	52.7	3.6	59.0	7.9	130.5
12	2.7	-25.5	4.2	50.0	4.9	54.7	9.7	130.1
13	-3.0	-17.4	3.1	47.1	4.3	48.8	9.3	117.6
14	-3.8	-11.3	3.2	56.3	5.0	55.7	10.7	120.5
101	3.6	-29.0	2.3	40.3	4.3	48.6	7.9	114.9
102	4.7	-30.6	3.6	48.6	5.9	56.2	10.7	131.0
103	10.2	-29.9	3.2	53.2	10.7	59.5	15.6	150.7
104	0.5	-29.0	3.8	70.2	3.8	73.9	8.7	127.7
105	3.0	-33.7	3.3	42.5	4.4	53.2	9.5	118.9
All	3.1	-25.1	4.2	50.8	5.2	56.6	15.6	192.6

Table 5.3: Radius of curvature errors for cat's eye curvature reference and SPINALTOP verification surfaces.

Part Number	Bias [µm]		Stdev [µm]		RMS [µm]		Max [µm]	
	SPINALTOP	Keratron	SPINALTOP	Keratron	SPINALTOP	Keratron	SPINALTOP	Keratron
1	-0.0062	0.0369	0.0034	0.0468	0.0070	0.0584	0.0135	0.1270
2	-0.0062	0.0424	0.0030	0.0535	0.0069	0.0669	0.0120	0.1288
3	-0.0062	0.0345	0.0030	0.0477	0.0069	0.0576	0.0129	0.1054
4	-0.0063	0.0317	0.0026	0.0550	0.0068	0.0620	0.0119	0.1109
5	-0.0052	0.0373	0.0025	0.0658	0.0058	0.0738	0.0095	0.1301
6	0.0073	0.0525	0.0028	0.0699	0.0078	0.0857	0.0123	0.1763
7	-0.0030	0.0382	0.0025	0.0718	0.0039	0.0793	0.0081	0.1479
8	-0.0033	0.0289	0.0027	0.0690	0.0043	0.0728	0.0101	0.1224
9	0.0072	0.0498	0.0026	0.0719	0.0077	0.0856	0.0122	0.1694
10	0.0054	0.0312	0.0033	0.0483	0.0064	0.0562	0.0125	0.1107
11	0.0102	0.0448	0.0032	0.0565	0.0106	0.0707	0.0175	0.1314
12	-0.0119	0.0582	0.0142	0.0604	0.0184	0.0825	0.0295	0.1523
13	-0.0065	0.0254	0.0081	0.0601	0.0103	0.0635	0.0195	0.1218
14	-0.0098	0.0376	0.0097	0.0826	0.0137	0.0884	0.0299	0.1753
101	0.0070	0.0459	0.0036	0.0430	0.0079	0.0620	0.0137	0.1254
102	0.0041	0.0496	0.0030	0.0541	0.0051	0.0722	0.0109	0.1410
103	0.0028	0.0438	0.0025	0.0610	0.0037	0.0735	0.0085	0.1703
104	0.0134	0.0493	0.0080	0.1507	0.0156	0.1534	0.0406	0.3545
105	0.0061	0.0481	0.0068	0.0410	0.0091	0.0624	0.0231	0.1341
All	0.0001	0.0413	0.0091	0.0653	0.0091	0.0772	0.0406	0.3545

Table 5.4: Conic constant errors for cat's eye curvature reference and SPINALTOP verification surfaces.

Part Number	Bias [D]		Stdev [D]		RMS [D]		Max [D]	
	SPINALTOP	Keratron	SPINALTOP	Keratron	SPINALTOP	Keratron	SPINALTOP	Keratron
1	-0.0387	0.0891	0.0148	0.3015	0.0414	0.3052	0.0753	0.6956
2	-0.0339	0.1168	0.0133	0.3129	0.0364	0.3247	0.0616	0.7822
3	-0.0249	0.1100	0.0151	0.2749	0.0290	0.2880	0.0559	0.6988
4	-0.0217	0.1245	0.0145	0.2730	0.0261	0.2922	0.0525	0.6076
5	-0.0238	0.1234	0.0162	0.2899	0.0287	0.3067	0.0582	0.7049
6	-0.0022	0.1790	0.0102	0.2955	0.0103	0.3375	0.0223	0.9331
7	-0.0166	0.1319	0.0115	0.2742	0.0201	0.2965	0.0358	0.6083
8	-0.0211	0.0952	0.0065	0.2515	0.0220	0.2614	0.0410	0.5487
9	-0.0102	0.1000	0.0095	0.2335	0.0139	0.2472	0.0373	0.5487
10	-0.0286	0.2171	0.0155	0.2519	0.0325	0.3265	0.0561	0.8185
11	-0.0029	0.1544	0.0190	0.2745	0.0190	0.3074	0.0416	0.6769
12	-0.0142	0.1327	0.0222	0.2614	0.0262	0.2858	0.0515	0.6768
13	0.0142	0.0901	0.0174	0.2466	0.0223	0.2552	0.0491	0.6110
14	0.0202	0.0573	0.0168	0.2949	0.0261	0.2912	0.0566	0.6260
101	-0.0243	0.1931	0.0155	0.2711	0.0288	0.3259	0.0547	0.7786
102	-0.0242	0.1559	0.0185	0.2507	0.0303	0.2885	0.0540	0.6797
103	-0.0436	0.1263	0.0135	0.2266	0.0456	0.2531	0.0682	0.6457
104	-0.0023	0.1861	0.0229	0.3458	0.0228	0.3824	0.0610	0.6630
105	-0.0155	0.1734	0.0174	0.2190	0.0231	0.2740	0.0503	0.6176
All	-0.0166	0.1344	0.0224	0.2673	0.0279	0.2988	0.0753	0.9331

Table 5.5: Diopter errors for cat's eye curvature reference and SPINALTOP verification surfaces.

Part Number	RMS mean [nm]		RMS Max [nm]		PV_{98} mean [nm]		PV_{98} max [nm]	
	SPINALTOP	Keratron	SPINALTOP	Keratron	SPINALTOP	Keratron	SPINALTOP	Keratron
1	13.8	221.6	15.6	354.8	63.0	1059.7	71.5	1702.0
2	14.6	259.0	16.1	401.5	67.1	1248.0	73.2	2000.6
3	14.3	231.9	17.4	328.4	66.7	1088.7	83.1	1389.4
4	16.0	274.3	20.4	349.6	75.0	1359.1	100.7	1608.4
5	16.5	292.0	18.7	403.7	77.3	1431.7	89.9	1841.7
6	13.2	282.1	14.7	444.0	63.7	1344.8	68.9	1989.2
7	16.6	285.3	18.7	404.9	77.8	1397.0	89.2	1847.1
8	16.1	383.4	23.9	630.0	77.2	1877.4	106.7	2857.2
9	12.5	286.4	15.3	367.2	62.4	1433.5	75.3	1822.0
10	12.2	269.1	14.2	320.3	57.7	1351.0	65.3	1760.0
11	13.6	323.3	15.9	425.6	66.6	1674.0	82.8	2406.6
12	71.6	281.9	85.3	372.1	380.0	1366.3	468.8	1881.6
13	23.6	394.6	32.0	622.1	123.2	1932.3	171.3	3174.8
14	31.0	592.8	36.0	886.8	165.0	3039.2	201.1	4339.1
101	17.8	214.2	20.7	340.3	91.7	1003.3	126.3	1350.6
102	24.5	255.8	28.4	345.0	131.6	1229.7	171.9	1511.1
103	24.3	261.0	27.8	394.2	128.8	1235.4	164.3	1808.4
104	97.2	397.1	112.3	525.5	509.3	2031.2	607.9	2951.0
105	37.9	298.9	48.2	436.1	214.3	1468.0	270.5	2015.3
All	25.2	305.5	112.3	886.8	129.1	1503.7	607.9	4339.1

Table 5.6: Biconic fit residual errors for cat's eye curvature reference and SPINALTOP verification surfaces.

Part Number	Bias [D]		Stdev [D]		RMS [D]		Max [D]	
	SPINALTOP	Keratron	SPINALTOP	Keratron	SPINALTOP	Keratron	SPINALTOP	Keratron
101	-0.0057	-0.0107	0.0053	0.4685	0.0077	0.4384	0.0152	0.6953
102	-0.0049	-0.0216	0.0053	0.4309	0.0072	0.4036	0.0150	0.6476
103	-0.0021	-0.0479	0.0047	0.4009	0.0051	0.3780	0.0119	0.6609
104	0.0118	-0.3639	0.0069	0.5899	0.0135	0.6610	0.0261	1.0274
105	0.0056	0.0782	0.0048	0.3639	0.0073	0.3493	0.0161	0.5137
All	0.0010	-0.0732	0.0086	0.4597	0.0087	0.4598	0.0261	1.0274

Table 5.7: Astigmatism estimation errors for toric verification surfaces.

5.2.2.3 CAT'S EYE CURVATURE REFERENCE SURFACE DISCUSSION

Biconic fit residuals for two of the cat's eye curvature reference surfaces are shown in Figure 5.27. The display area is 7 x 7 mm and the color-height scale has been set to 60 nm. Both surfaces were defocused by 150 μm and translated by 75 μm. The left image is from Part 001 which has a radius of 7 mm. The diameter of the reconstructed surface area is a function of the converger f/# and radius of curvature of the part (R) and could be estimated:

$$D_{test} = \frac{R}{f/\#} = \frac{7}{1.3} = 5.4\ mm \qquad (5.2)$$

The right image is from part 009 which has a radius of 9 mm, which results in a test diameter of 6.9 mm. Figure 5.27 demonstrates the advantage of SPINALTOP over the TFI, where the absolute surface height and shape can be determined on SPINALTOP. If these measurements had been made on the TFI, the absolute scale would be unknown and both surfaces would have been incorrectly scaled to a 6 mm diameter area.

The biconic fit residuals for the same two cat's eye curvature reference surfaces (Parts 001 and 009) that were measured on the Keratron Piccolo are shown in Figure 5.28. The display area is 7 x 7 mm and the color-height scale is 2.6 μm, which is 43X larger than the color-height scale used for the SPINALTOP measurements. The difference in height resolution is best exemplified in Figure 5.29 with a side-by-side comparison of the biconic fit residuals for Part 009 measured on SPINALTOP (left) and the Keratron Piccolo (right) that have been set to the same color-height scale of 1.8 μm. To further illustrate the difference in performance between the two systems, the best set of data (i.e. minimum residual error) was chosen for the Keratron Piccolo measurement and the worst set of data was chosen for the SPINALTOP measurement.

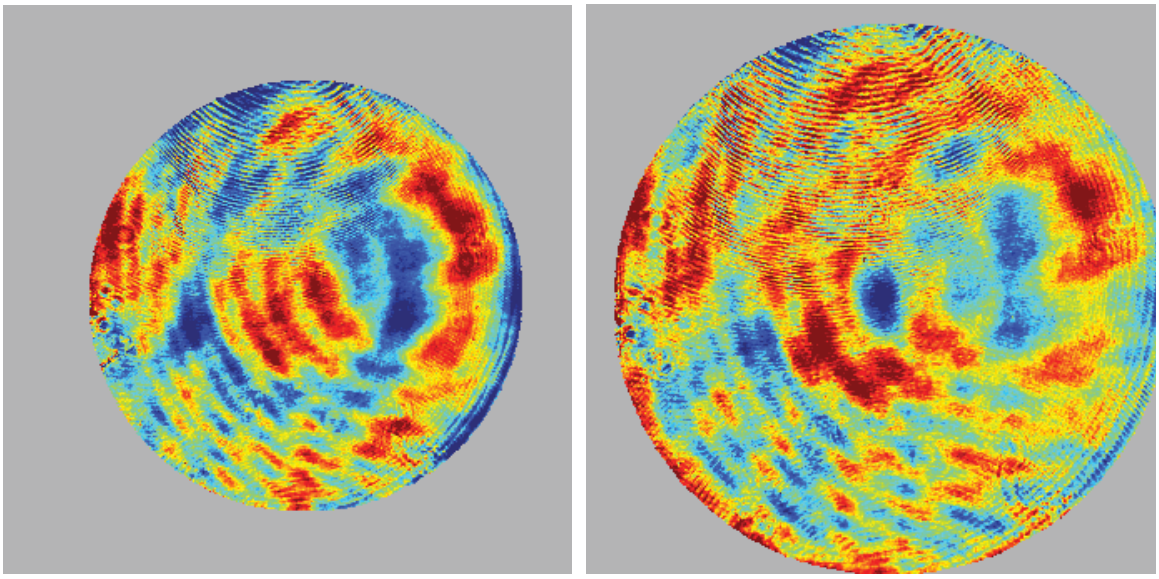

Figure 5.27: Biconic fit residuals for reference surfaces measured on SPINALTOP.
(Left) Part 001, (Right) Part 009. Both surfaces have been defocused by 150 µm and translated by 75 µm.
Display area is 7 x 7 mm. 30 nm color height scale.

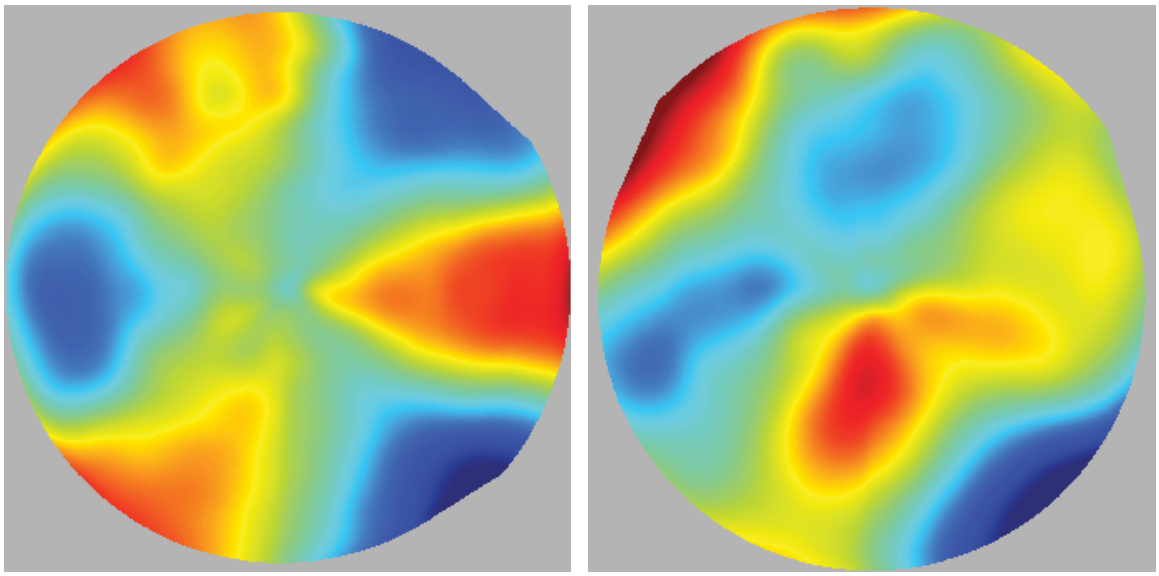

Figure 5.28: Biconic fit residuals for cat's eye surfaces measured on the Keratron Piccolo.
(Left) Part 001, (Right) Part 009. Surfaces have had a small amount of decenter. Display area is 7 x 7 mm.
±1.3 µm color height scale.

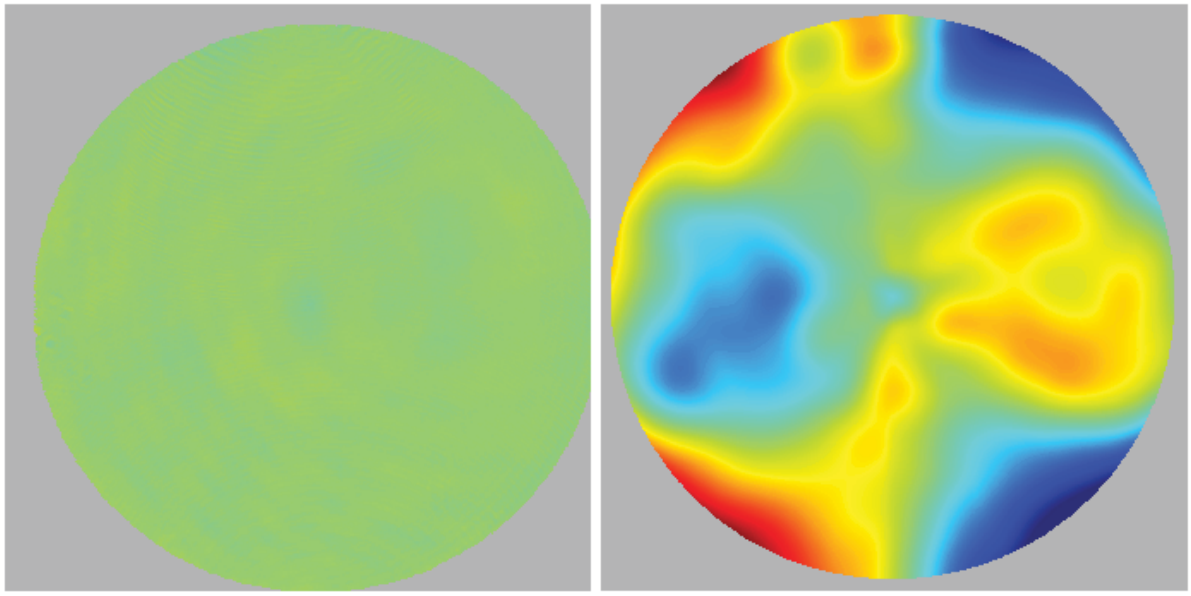

Figure 5.29: Side-by-side comparison of Part 009 (cat's eye curvature reference surface) biconic fit residuals from SPINALTOP (Left) and the Keratron Piccolo (Right). Display area is 7 x 7 mm. ±0.9 µm color height scale. SPINALTOP measurement was worst from the set. Keratron measurement is best from the set.

5.2.2.4 SPINALTOP VERIFICATION SURFACE DISCUSSION

Results for the verification surface are similar to the cat's eye reference surfaces, with some exceptions that will be discussed in this section. The biconic fit residuals for verification surfaces 011 and 014 are shown in Figure 5.30. The display area is 6.5 x 6.5 mm and the color-height scale has been set to 80 nm (±40 nm). Both surfaces were defocused by 150 µm and translated by 75 µm. Part 011 (left) has a radius of 8 mm and a conic constant of -0.50. Part 014 (right) has a radius of 8 mm and a conic constant of +1.0. A trend is present in the data for parts 011 to 014, where the amount of residual spherical aberration in the residuals increases as a function of the conic constant. Part 011 (left) has the least amount of residual spherical aberration and part 014 (right) appears to have the most. Figure 5.31 is a comparison of the biconic fit residuals for part 014 that has been measured with SPINALTOP (left) and the Keratron Piccolo (right). The color-height scale has been set to 1.3 µm (±0.65 µm) and a display area is 6 x 6 mm.

It was initially assumed that the spherical aberration that showed up in the residuals was a result of an incomplete characterization of the converger assembly; a radius, thickness, or spacing parameter may need correction. However, the potential manufacturing errors of the non-axially symmetric parts (parts 101 - 105) that were discussed in Section 5.2.1 raises questions about the verification surfaces. If the toroidal surfaces have residual manufacturing errors, than it may be possible that the other reference and verification surfaces have errors, such as the residual spherical aberration described above. However, with the toric (i.e. non-rotationally symmetric) surfaces, the manufacturing residuals can be easily isolated due to their non-rotational symmetry. Figure 5.32 is a comparison of the residual errors for part 103 that have been measured on SPINALTOP (left) and the WYKO (right). The SPINALTOP data area has been truncated to 5 x 5 mm to match the WYKO display area. The residuals appear to correlate between the two measurements; although, uncorrected distortion in the WYKO measurements may be skewing the display (e.g. the WYKO display area may cover more than 5 x 5 mm on the part). Based on these measurements, it is highly likely that the increase in reconstruction errors for the SPINALTOP system is due to the brass surfaces. A more complete characterization of these surfaces may improve the accuracy results of the SPINALTOP system.

A partial characterization of one of the verification surfaces is provided to demonstrate the utility of a full characterization. A complete characterization is not possible with the equipment available, but the WYKO measurement contains enough information for a partial characterization. The surface residual for part 105 that is shown in Figure 5.15 from the WYKO measurement is believed to be primarily due to manufacturing error. The surface residual is a result of subtracting 10 Zernike polynomial terms from the raw surface measurement that best represent the base surface figure shape and low spatial-frequency structure that are due to retrace errors (refer to Section 5.2.1). The dominant residual errors are best represented by trefoil (Z9/10) and pentafoil (Z16/17), so the manufacturing surface residual is represented by a Zernike polynomial with these four terms. The WYKO measurement covers a 5 mm diameter area on the part, so the polynomial fit can be used to extrapolate out to a 6 mm diameter area. The

manufacturing surface residual can then be added to the generalized biconic surface that is being subtracted from the reconstructed SPINALTOP measurement (Section 4.5.2). The results are shown in Figure 5.33; the generalized biconic residual for the uncorrected surface is shown on the left and the generalized biconic residuals for the corrected surface are shown on the right. The display area is 6 x 6 mm for both measurements, the color height scale for the uncorrected measurement is ±100 nm and for the corrected measurement it is ±30 nm. This example demonstrates that most of the residual errors for the SPINALTOP verification surfaces may be due to the uncharacterized verification surfaces and not a result of the system.

As a final demonstration, Figure 5.34 shows a comparison of part 104 (the worst performing part) measured on SPINALTOP (left) and on the Keratron Piccolo (right). The display area is 6.5 x 6.5 mm and the color-height scale has been set to 2.5 μm (±1.25 μm). Despite the residual manufacturing errors in this surface, they are barely noticeable in comparison to the measurement made on the Keratron Piccolo system.

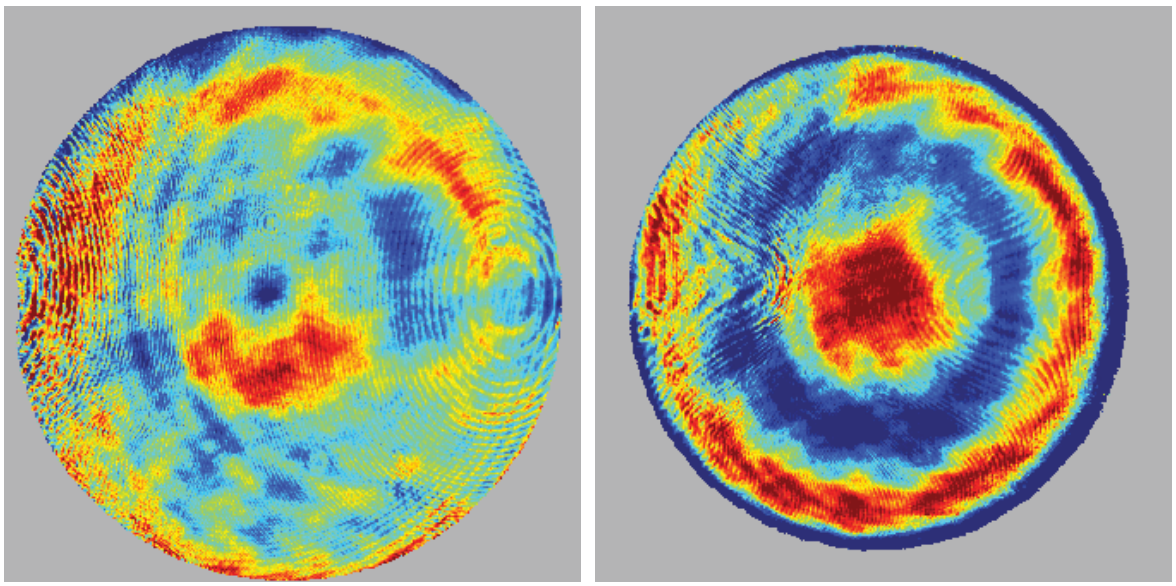

Figure 5.30: Biconic fit residuals for verification surfaces measured on SPINALTOP. (Left) Part 011, (Right) Part 014. Both surfaces have been defocused by 150 μm and translated by 75 μm. Display area is 6 x 6 mm. ±40 nm color height scale.

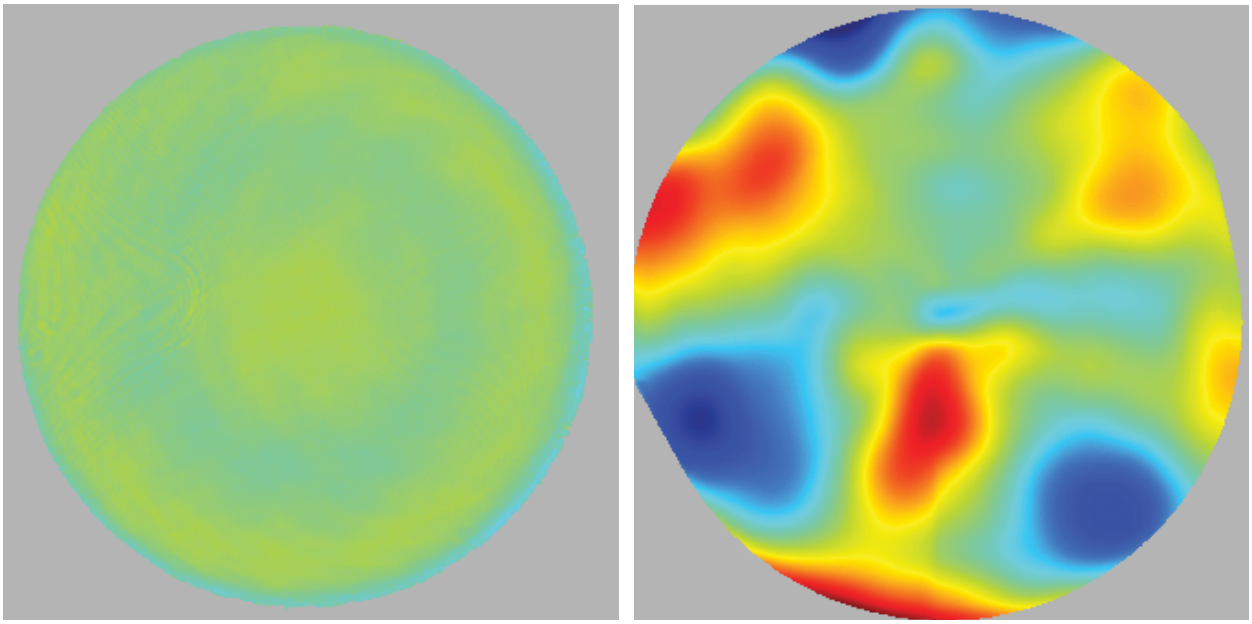

Figure 5.31: Side-by-side comparison of Part 014 (SPINALTOP verification surface) biconic fit residuals from SPINALTOP (Left) and the Keratron Piccolo (Right). Display area is 6 x 6 mm. ±0.65 µm color height scale.

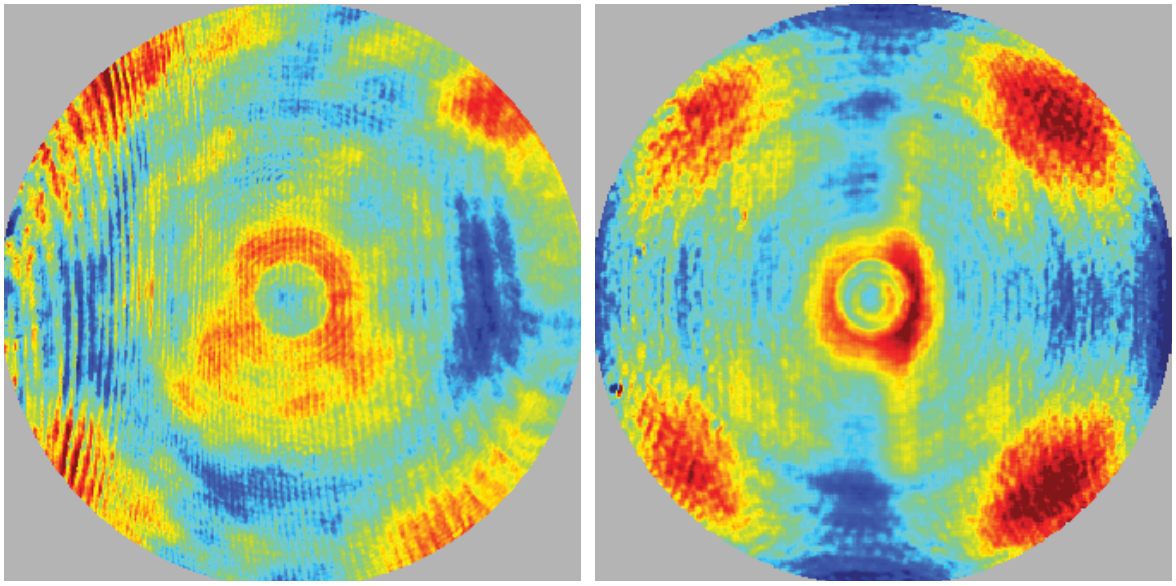

Figure 5.32: Part 103 surface residual comparisons. (Left) SPINALTOP measurement (Right) WYKO measurement. ±60 nm color height scale. Displayed area is approximately 5 x 5 mm. SPINALTOP residual from best fit generalized biconic. WYKO measurement has had 10 Zernike terms removed.

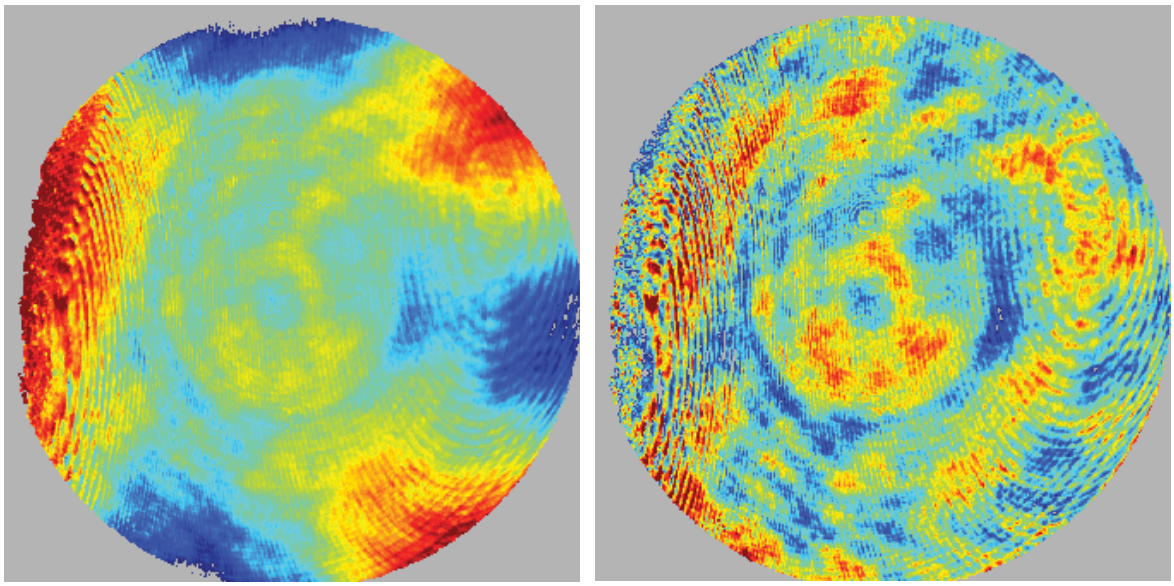

Figure 5.33: Surface correction applied to SPINALTOP measurement of Part 105.
(Left) Uncorrected measurement. ±100 nm color height scale.
(Right) Corrected measurement. ±30 nm color height scale.
Displayed area is 6 x 6 mm. Residual from best fit generalized biconic.

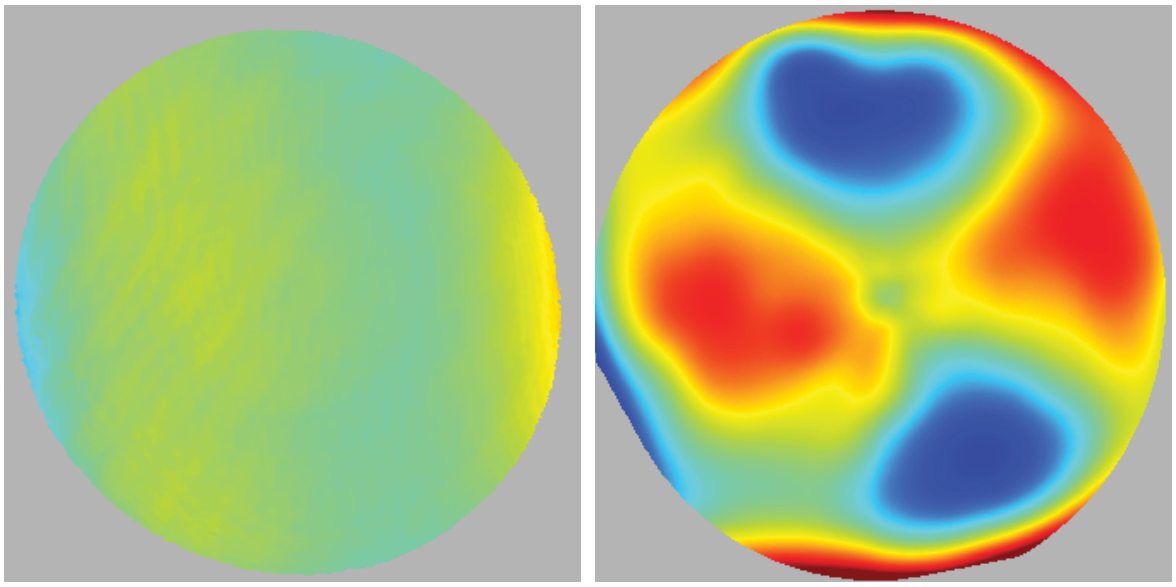

Figure 5.34: Side-by-side comparison of Part 104 (SPINALTOP verification surface) biconic fit residuals from SPINALTOP (Left) and the Keratron Piccolo (Right).
Display area is 6 x 6 mm. ±1.0 μm color height scale.

5.2.2.5 SPATIAL RESOLUTION

The spatial resolution of the SPINALTOP system is nearly two orders of magnitude greater than the Keratron Piccolo. Measurements of the reference and verification surfaces that were presented in the previous sections qualitatively demonstrate that this is the case. This section provides a more quantitative measure of the spatial resolution of the two systems. Figure 5.35 shows a side-by-side image of the distortion target captured by SPINALTOP (left) and the Keratron Piccolo (right). The Placido disc system operates by detecting transition edges between the white and black circles and using these to detect slope changes on the surface of the part (Section 1.6.2). From Keratron Piccolo data shown in Figure 5.35, it can be seen that 9 rings, or 18 transitions, over a 3 mm radius occur. This allows for 18 rings / 3 mm = 167 μm spacing resolution of the Keratron Piccolo. The dimensions of the standard target fiducials (circles) are 70 μm and the alignment fiducials are 150 μm, which means that the Keratron should not be able to resolve the standard fiducials and may barely resolve the alignment fiducial. This is verified with the processed data shown in Figure 5.36. The best fit generalized biconic surface has been removed from the data to enhance the display to show the mid to high spatial-frequency information. The left image shows a 6 x 6 mm area over the surface of the part and the right image shows the central 1 x 1 mm area. The distortion features are spaced by 0.5 mm and are approximately 70 μm in diameter. However, the Keratron cannot resolve any of these features. Nor can it be determined if the Keratron resolved any of the 150 μm alignment fiducials. Figure 5.37 shows a 1.5 x 1.5 mm area of the raw image that demonstrates the sampling issue with the Keratron and why it cannot resolve anything below of 167 μm.

The processed image from SPINALTOP for the distortion target is shown in Figure 5.38. The left image shows a 6 x 6 mm area over the surface of the part and the right image shows the central 1 x 1 mm area. In both images, the distortion features are easily visible. Figure 5.39 is zoomed in over one of the alignment fiducials. Approximately 11 pixels over the 70 μm area can be resolved, which is equivalent to a 6.4 μm spatial resolution. The properties of the imaging system and detector predicted that this would nominally be 6.0 μm, which is in good agreement with the measurement. Distortion in the imaging

system will result in variations of spatial resolution across the image, in particular at the edge of the field of view. However, the nominal spatial resolution of the SPINALTOP system has been demonstrated to far exceed that of the Keraton Piccolo.

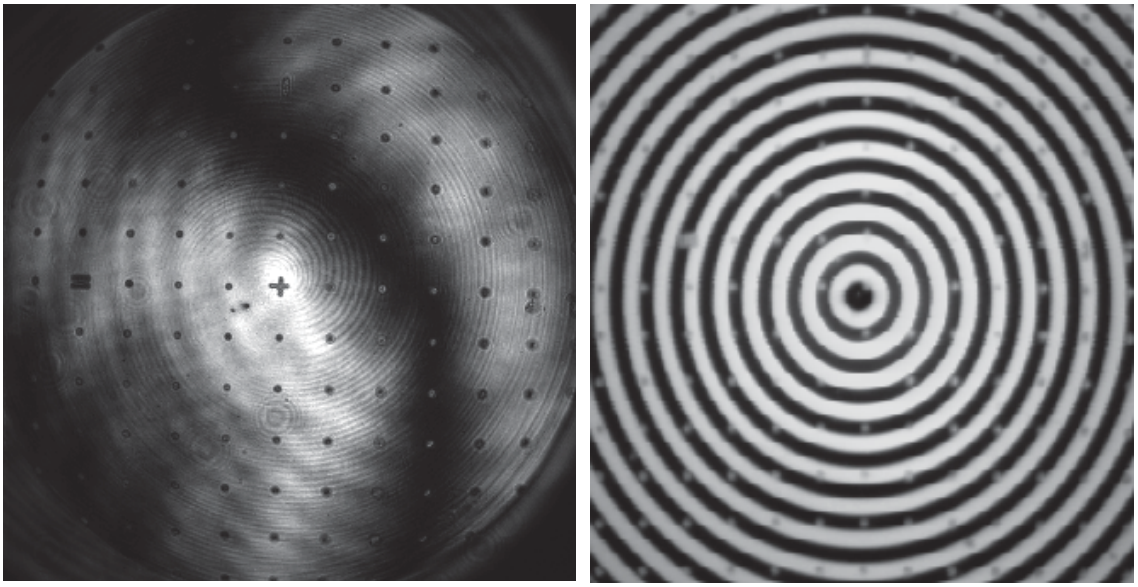

Figure 5.35: Images of part 015D (distortion target).
(Left) SPINALTOP. (Right) Keratron Piccolo.
Display area is 6 x 6 mm.

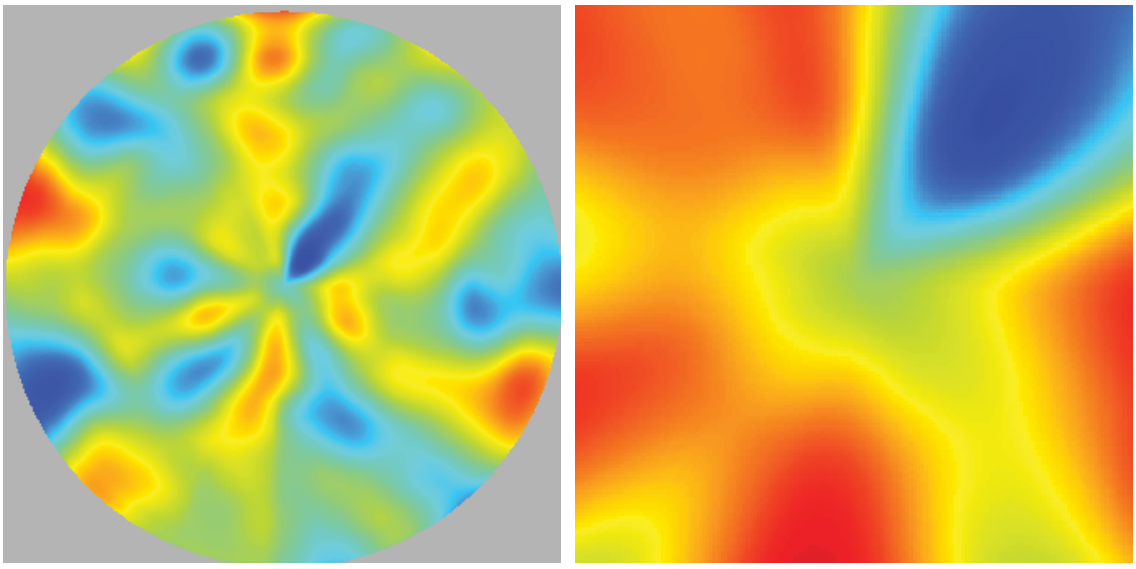

Figure 5.36: Distortion target surface height map with biconic removed (Keratron).
(Left) Display area is 6 x 6 mm. ±450 nm color height scale; red is away from surface.
(Right) Central area of target zoomed in, display area is 1 x 1 mm. ±300 nm color height scale.

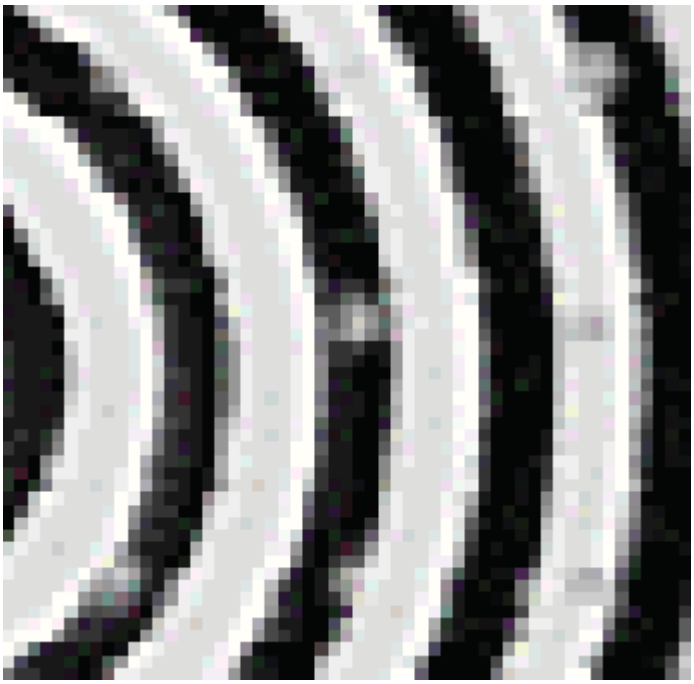

Figure 5.37: Distortion target features on Keratron raw measurement.
Display area is approximately 1.5 x 1.5 mm.

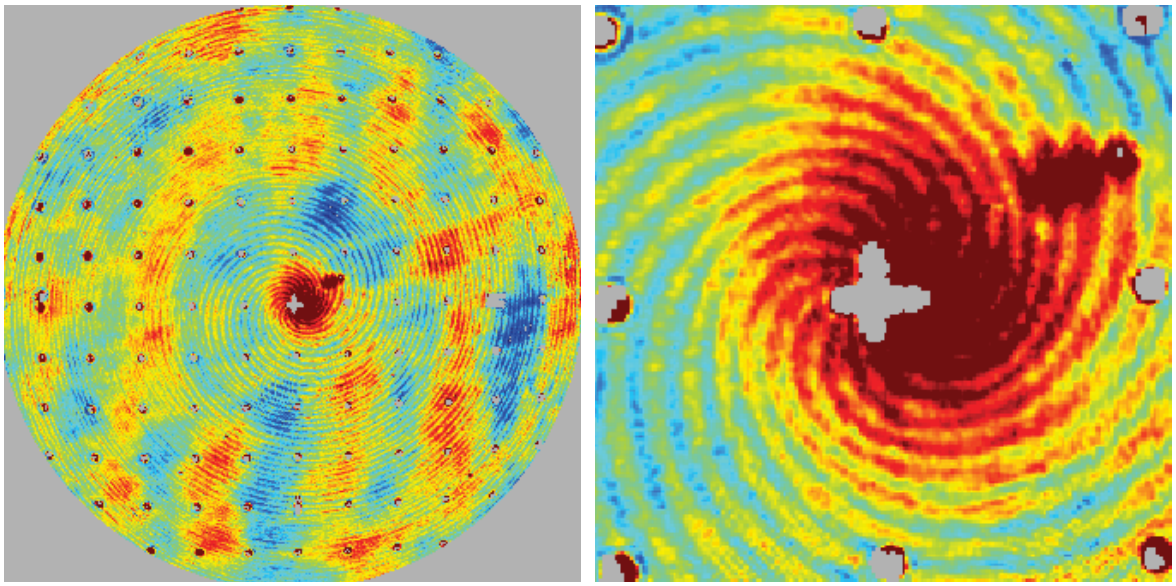

Figure 5.38: Distortion target surface height map with biconic removed (SPINALTOP).
(Left) Display area is 6 x 6 mm. ±90 nm color height scale; red is away from surface.
(Right) Central area of target zoomed in, display area is 1 x 1 mm. ±90 nm color height scale.

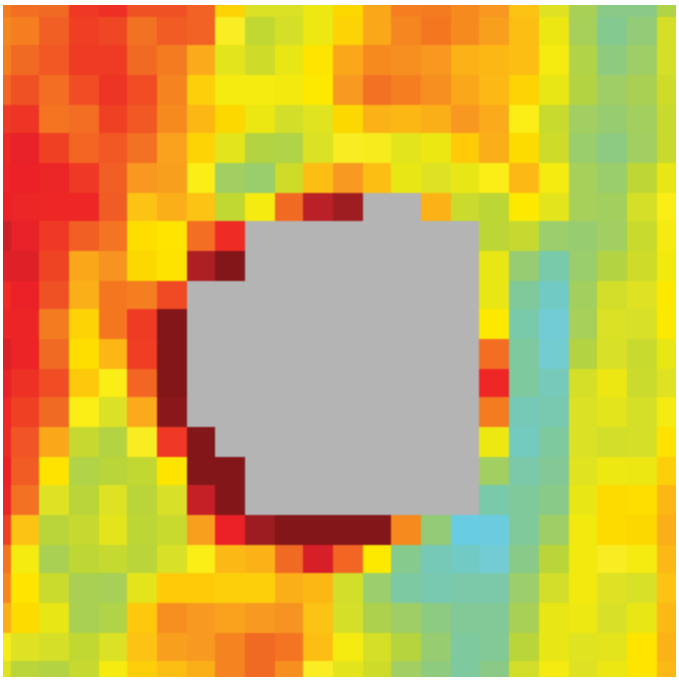

Figure 5.39: Distortion target feature on SPINALTOP measurement.
Display area is approximately 150 x 150 µm. Data dropout occurs at feature location (shown in grey).

5.2.2.6 DATA PRESENTATION FOR HUMAN SUBJECTS TESTING

Results from human subjects testing will be presented in the next section (Section 5.3). It will be easier to develop the presentation methods and demonstrate them with the brass verification surfaces in this section before they are applied to the human subjects testing results. First, the display of the raw corneal topography, or surface height data, is non-ideal. The display of any surface height information from the expected range of human topographies over a 6 mm area is at least 500 µm. The entire dynamic range would be dominated by the base shape and will not provide any sensitivity. An example of this can be seen with Figure 5.40 where the surface height map of surface with a 7 mm radius of curvature is shown on the left and 9 mm radius is shown on the right. The color height scale for both surfaces is 650 µm. Although there is a discernible difference, it is minimal and not very useful. A toric surface with would barely register any noticeable difference.

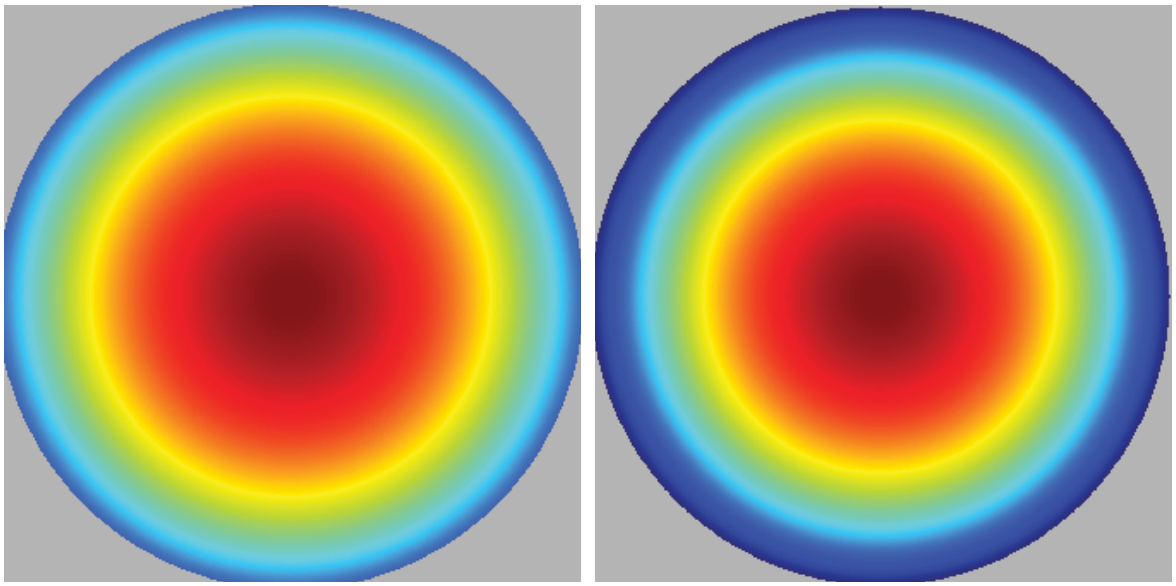

Figure 5.40: Simulated surface height measurements for simple aspheres.
(Left) Asphere with 7 mm radius. (Right) Asphere with 9 mm radius. Display area is 6 x 6 mm. ±325µm color height scale; red is away from eye.

The first set of data that can be displayed is the conic removed surface. A weighted least-squares fit of a conic surface is performed on the reconstructed corneal topography. The height, or sag, of the conic surface can be described as:

$$Z_{conic} = \frac{r^2/R}{1+\sqrt{1-(K+1)r^2/R^2}} \qquad (5.3)$$

Where R is the radius of curvature and K is the conic constant. The weighted lest-square fit is solved by use of the Levenberg-Marquadt algorithm (Section 4.3.2.4.1). The residual, or conic removed surface, is the raw set of data with the best fit conic subtracted:

$$Z_{conic\ residual} = Z_{meas} - Z_{conic} \qquad (5.4)$$

Applying this method to a measurement performed on part 102 results in the images shown in Figure 5.41. If a subject has enough corneal astigmatism, this method may be used to highlight it.

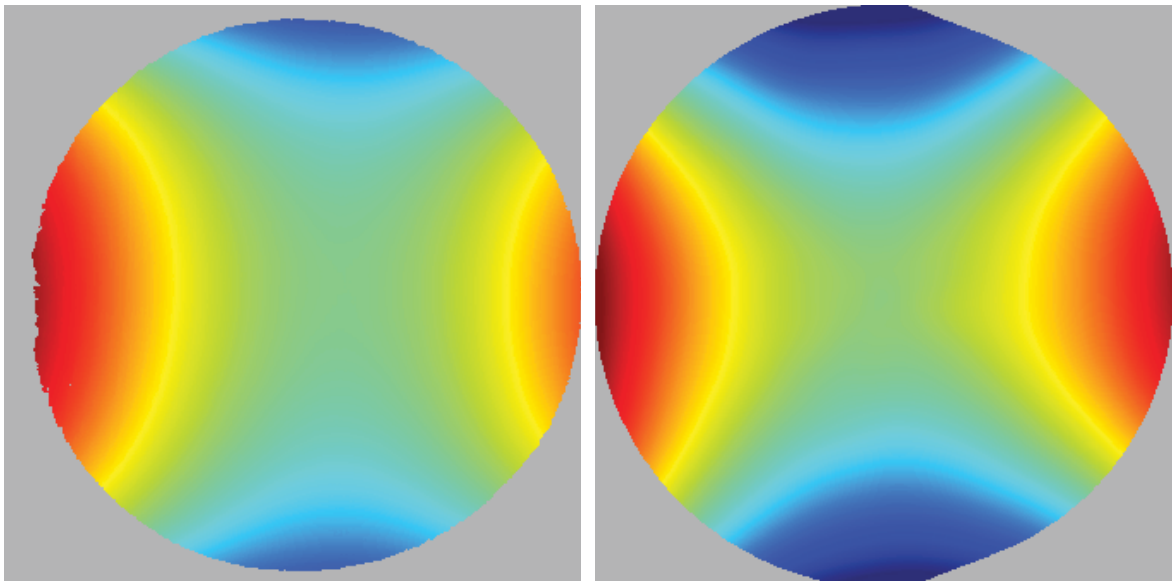

Figure 5.41: Conic removed surface height of part 102.
(Left) SPINALTOP measurement. (Right) Keratron Piccolo measurement.
Display area is 6.5 x 6.5 mm. ±11.0 µm color height scale; red is away from part surface.

The conic removed surface does not allow for a full description of the eye, which is better accomplished by a generalized biconic surface (Section 1.1). The process of subtracting the generalized biconic was previously described in Section 4.5.2 and when applied to part 102 results in the set of images in Figure 5.42. Removal of the best fit biconic surface backs out a significant portion of the subject's corneal topography, revealing mostly tear film structure. Although some of the topography may still remain, the display scale as demonstrated by Figure 5.42 now allows for tear film surface features to be resolved. The displays are set to the same color height scale which obscures information in the SPINALTOP measurement (left) and demonstrates the inaccuracy of the Keratron system (right). However, the tear film structure will contain features that are large enough for both instruments to resolve and the same display scale can be used. The fit parameters (e.g. radius of curvature) will be included with every measurement as shown in Table 5.8.

	Part 102	
	SPINALTOP	**Keratron**
Rx [mm]	-8.197	-8.274
Ry [mm]	-7.998	-7.971
Kx	-0.241	-0.123
Ky	-0.249	-0.240
θ [deg]	-0.32	-2.57
$SSIM_B$	4.1%	
$SSIM_Z$	8.9%	

Table 5.8: Part 102 biconic fit results.

A structural similarly (SSIM) index is included in Table 5.8, which provides a measure for the similarity between the two measurements (Z. Wang, Bovik, Sheikh, & Simoncelli, 2004). The metric is defined so that two identical images have an SSIM of 100% and differences between the images reduce the SSIM percentage. The method is similar to a correlation, but proved to be more reliable in detecting differences. For example, the SSIM index resulted in a 4.1% similarity between the SPINALTOP and Keratron measurement shown in Figure 5.42, whereas a correlation between the two measurements resulted in 75.8%. Similar issues were found with measurements from human subjects testing, so the SSIM index was chosen as a quantitative comparator. The SSIM index is defined as:

$$\mathrm{SSIM}(X,Y) = \frac{(2\bar{X}\bar{Y} + c_1)(2\sigma_{XY} + c_2)}{(\bar{X}^2 + \bar{Y}^2 + c_1)(\sigma_X^2 + \sigma_Y^2 + c_2)} \qquad (5.5)$$

Where $\bar{X}$ and $\bar{Y}$ are the mean of the two measurements (e.g. X is SPINALTOP, Y is Keratron Piccolo), σ_X and σ_Y are the standard deviations, and σ_{XY} is the covariance. The variables c_1 and c_2 are stabilisation variables to prevent the denominator from blowing up:

$$c_i = (k_i L)^2 \qquad (5.6)$$

Where $k_1 = 0.01$, $k_2 = 0.03$, and L is the dynamic range of the image. The SSIM index will be applied to the biconic ($SSIM_B$) and Zernike ($SSIM_Z$) removed data sets.

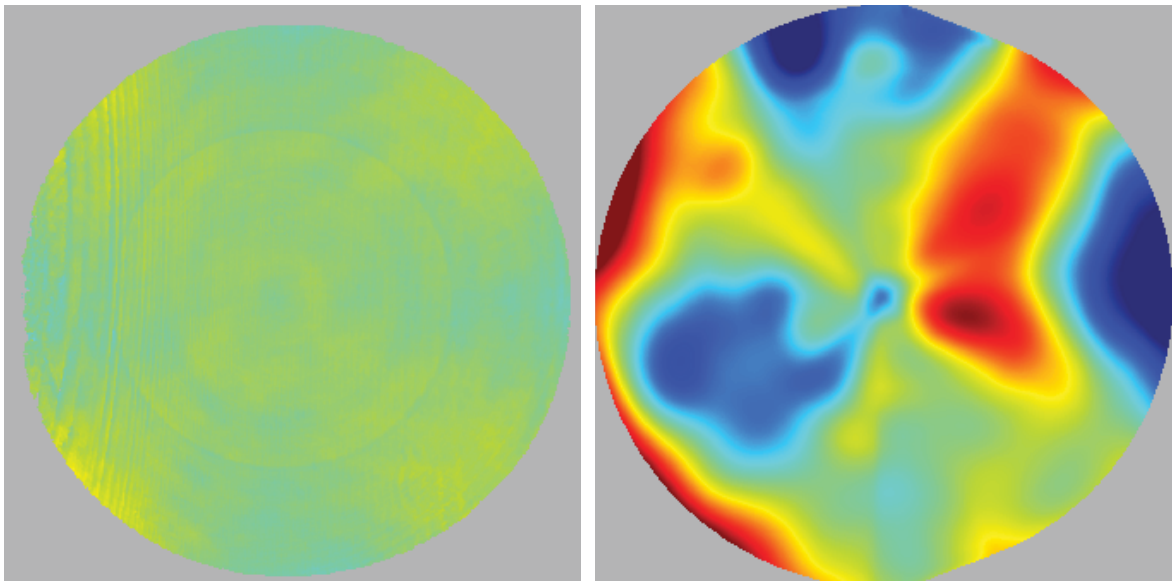

Figure 5.42: Biconic removed surface height of part 102.
(Left) SPINALTOP measurement. (Right) Keratron Piccolo measurement.
Display area is 6.5 x 6.5 mm. ±0.5 μm color height scale; red is away from part surface.

Residual corneal topography that remains after the biconic removed surface may still dominate the display scale. A final additional step may be performed on the data, which is to back out additional low-frequency spatial information by use of a Zernike polynomial fitting. The first 15 or 25 Zernike polynomial terms will be fit to the data and removed. The residual surface height information should be mostly representative of the tear film structure. The result of removing the first 25 Zernike terms from the biconic removed surface height of part 102 is shown in Figure 5.43. Although the Zernike removed surfaces do not appear to provide much change from the biconic removed surfaces shown in Figure 5.42, the human subjects measurements will benefit from this method.

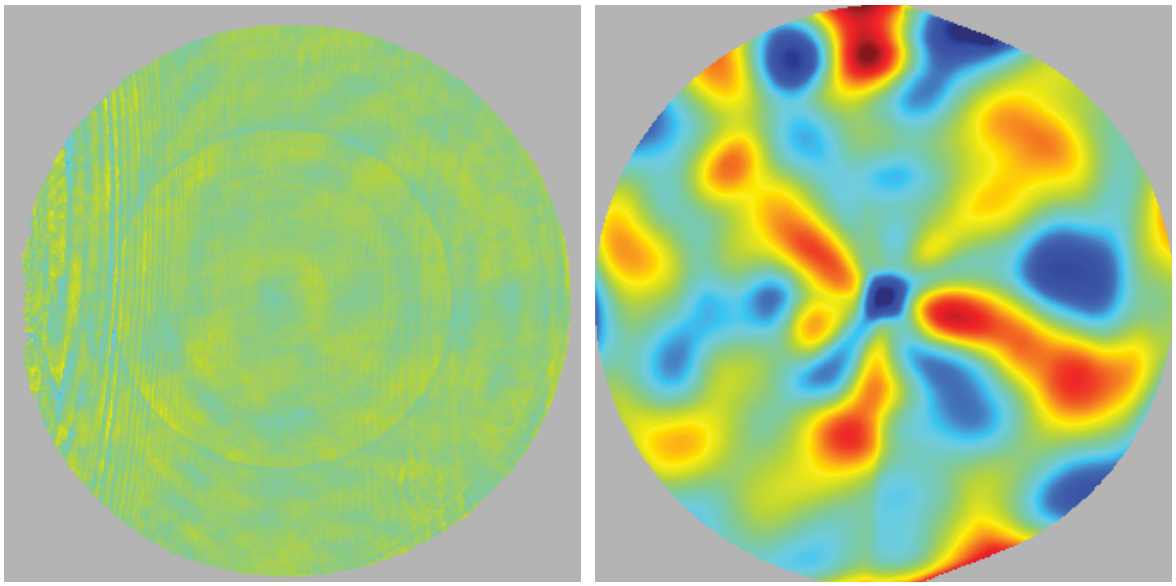

Figure 5.43: Zernike removed surface height of part 102.
(Left) SPINALTOP measurement. (Right) Keratron Piccolo measurement.
Display area is 6.5 x 6.5 mm. ±0.25 µm color height scale; red is away from part surface. First 15 Zernike terms removed.

The previous display methods represent surface height information of the corneal topography or tear film structure. Although these methods are familiar to those in optical metrology, they are a departure from what is commonly used in the ophthalmic community. Corneal topography maps are more commonly displayed in terms of axial power. Axial power is sometimes referred to as sagittal power, resulting from the fact that the derived power is proportional to the curvature of the surface along the sagittal direction (J. T. Schwiegerling, 2004). Describing the corneal surface height in cylindrical coordinates, $z = f(r, \theta)$, the axial power (Φ_a) is given by:

$$\Phi_a = \frac{(n_k - 1)\frac{dz}{dr}}{r\sqrt{1 + \left(\frac{dz}{dr}\right)^2}} \qquad (5.7)$$

Where $n_k = 1.3375$ which is the keratometric index of refraction; an effective index used when describing the total power of the cornea resulting from an anterior surface measurement by accounting for the negative power introduced by the posterior surface. An alternative display is the instantaneous, or

tangential power map. The instantaneous power (Φ_i) is proportional to the curvature of the surface along the tangential direction and is given by:

$$\Phi_i = \frac{(n_k - 1)\frac{d^2z}{dr^2}}{\left[1 + \left(\frac{dz}{dr}\right)^2\right]^{3/2}} = \frac{d}{dr} \mathrm{r}\Phi_a \qquad (5.8)$$

Calculation of the axial and instantaneous power from the raw SPINALTOP measurements will be sensitive to noise due to the numerical derivative. The instantaneous power measurement will be even more sensitive, requiring two derivative calculations. To minimize the impact of noise on the derivatives, the power is calculated from a low-pass filtered set of corneal topography data. A median filter with a width of at least 10 pixels will be required for most measurements.

The display of surface height and curvature (power) maps are standardized for corneal topographers by ANSI Z80.23 (ANSI, 2008). The standard requires that the curvature maps are displayed with one of the three diopter intervals: 0.5, 1.0, or 1.5 D. The total number of intervals has to be more than 21, but less than 25. The center interval is required to have a curvature of 44 D. Most of the data that will be measured is bound to a small interval around 44 D, so the minimum diopter interval (0.5 D) will be used. The resulting scale is shown in Figure 5.44, which shows both dioptric power (along the top) and radius of curvature (bottom).

Figure 5.44: Absolute scale from Keratron for instantaneous power.
Top row of numbers is power in Diopters, bottom row of numbers is radius in mm.

The calculated axial and instantaneous power maps for part 102 are shown in Figure 5.45 and Figure 5.46. The discretized height scale may limit the resolvability of some features when using the power maps; however, the power map is not the best form to study fine features of the topography. This is best left for

the surface height maps that have been shown previously. It should be noted that the ANSI standard also has a requirement for height elevation maps, but these were ignored due to the fact that the smallest allowed interval was 2 μm. The fact that SPINALTOP can resolve surface height features that are 0.02 μm or better makes this display requirement unsuitable for any display.

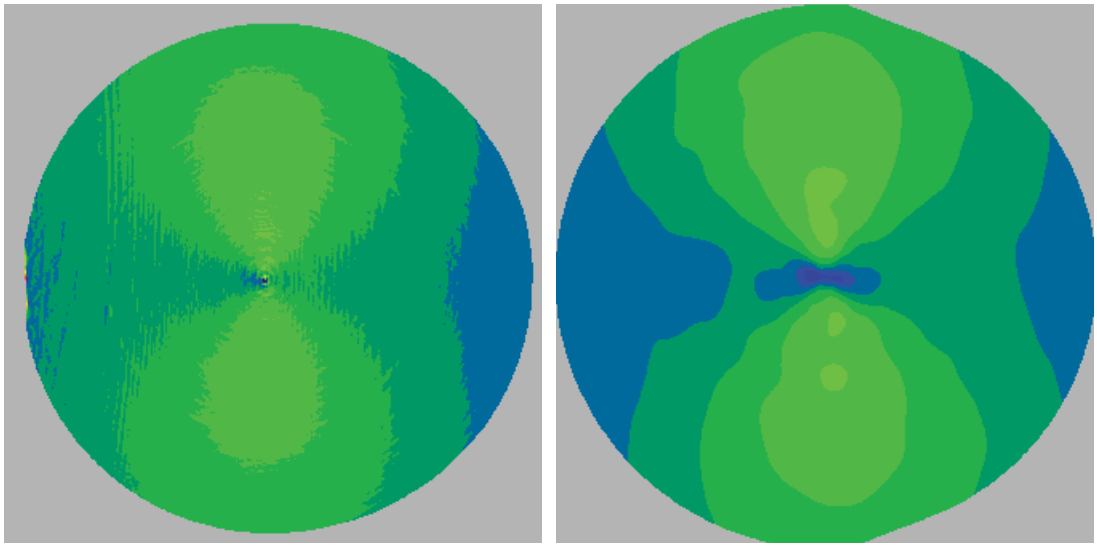

Figure 5.45: Part 102 axial power maps.
(Left) SPINALTOP measurement. (Right) Keratron Piccolo measurement.
Display area is 6.5 x 6.5 mm. Standardized power scale.

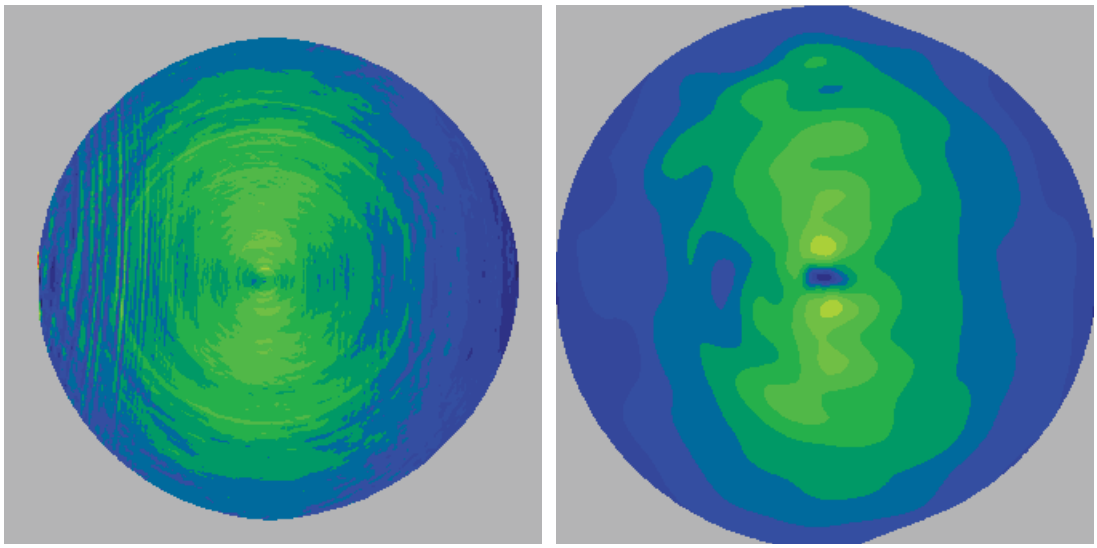

Figure 5.46: Part 102 instantaneous power maps.
(Left) SPINALTOP measurement. (Right) Keratron Piccolo measurement.
Display area is 6.5 x 6.5 mm. Standardized power scale.

5.3 HUMAN SUBJECTS TESTING

SPINALTOP was approved for human subjects testing by the University of Arizona Institutional Review Board and adheres to the tenants of the Declaration of Helsinki. Informed consent is obtained from a subject prior to testing. Testing was performed on healthy adults, aged 18 to 60 years. Subjects were not allowed to participate if they had any known eye disease or have undergone refractive surgery. Subjects outside of this age range or with an eye disease may introduce potential complications that inhibit the SPINALTOP system's ability to collect data.

Nine human subjects were examined using SPINALTOP. Functionally, the system operates similar to the TFI. A trained operator can align the system to a subject in under a minute. Acquisition times are limited only by the subjects' ability to comfortably maintain fixation. The software improvements that have been made since the TFI now allows the operator to adjust the system to the subject using live feedback while the system simultaneously captures and saves the data. Collection times are generally limited to 5 minutes before a break is given to the subject to minimize fatigue.

The results from human subjects testing are presented in this section. Temporal surface measurements (i.e. movies) are possible with this system, just as it was demonstrated with the TFI. However, the data is presented to focus on the absolute surface reconstruction capabilities of the system. Additional comparisons to measurements made on the Keratron Piccolo are also presented. The methods used to display the measurements were discussed in Section 5.2.2.6.

5.3.1 SUBJECT A

There is good agreement between SPINALTOP and the Keratron Piccolo for the measurements presented on this subject. The fit parameters for a generalized biconic surface are shown in Table 5.9. The subject has a unique enough topography to allow for a comparison between the fit residuals. In other words, the deviation of their topography from a generalized biconic surface is greater than a few microns, which can be resolved by the Keratron Piccolo. This can be seen with Figure 5.47, where the scale is 10 μm and a

strong correlation between the two measurements can be observed. This is supported by the SSIM index values for both eye measurements, which are both near 85%. Figure 5.48 shows a comparison between the two systems where low-spatial frequency information has been removed from the measurement by subtracting the first 15 Zernike terms. Although there appears to be some correlation between the two measurements, the SSIM index has been reduced to 58% and 24% for the left and right eyes, respectively. The reason for the reduced correlation is primarily due to two reasons. The first is that the residual surface height features that are being displayed are near the limits of the Keratron Piccolo (i.e. it starts to become difficult to distinguish noise from signal). Second, the Zernike removed surface contains more information about the tear film structure which is dynamic and would have changed between measurements on the two systems. The temporal difference between measurements on SPINALTOP and the Keratron Piccolo are sufficient enough to potentially obfuscate structural similarities. However, it is not expected that the Keratron can resolve any features of the tear film structure. The comparison between the two images is more to demonstrate the resolution capabilities of the SPINALTOP system.

Figure 5.49 is a comparison of the axial power maps from both systems. Although SPINALTOP has significantly more resolution, both spatial and height, the standardized scale used for power map displays somewhat obscures this. Despite this, a strong correlation between the two images can be seen with this example as well as the instantaneous power maps shown in Figure 5.50.

	Left		Right	
	SPINALTOP	Keratron	SPINALTOP	Keratron
Rx [mm]	-7.623	-7.668	-7.686	-7.697
Ry [mm]	-7.787	-7.826	-7.836	-7.857
Kx	-0.498	-0.320	-0.211	-0.189
Ky	-0.242	-0.075	0.107	0.067
θ [deg]	-16.01	-18.85	-42.74	-24.05
$SSIM_B$	85.6%		84.2%	
$SSIM_Z$	58.0%		23.8%	

Table 5.9: Subject A, biconic fit results.

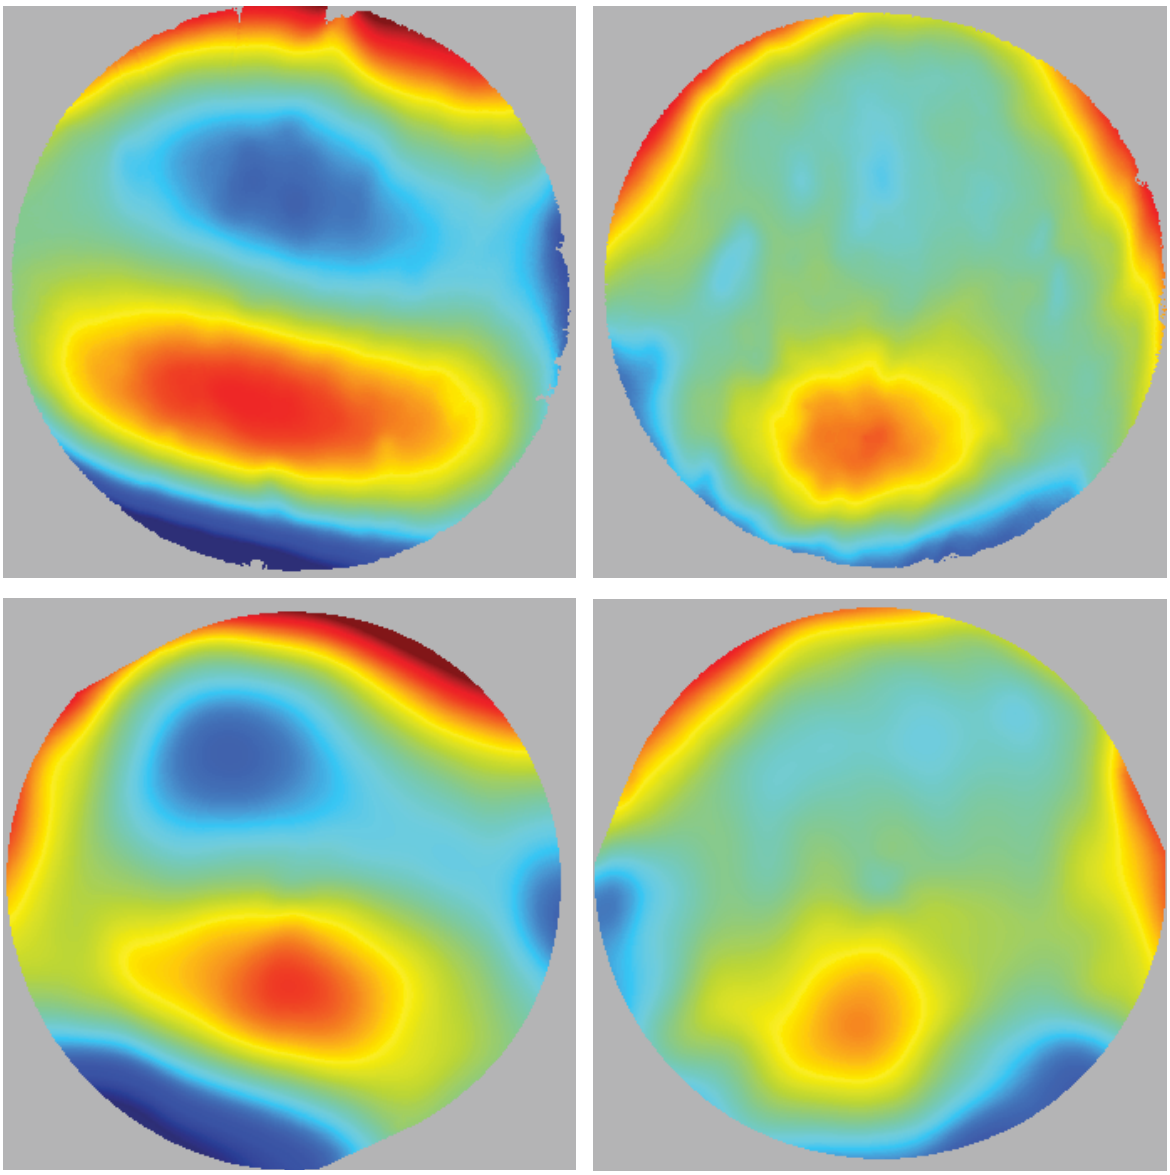

Figure 5.47: Subject A, biconic removed corneal surface height.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±5.0 μm color height scale; red is away from eye.

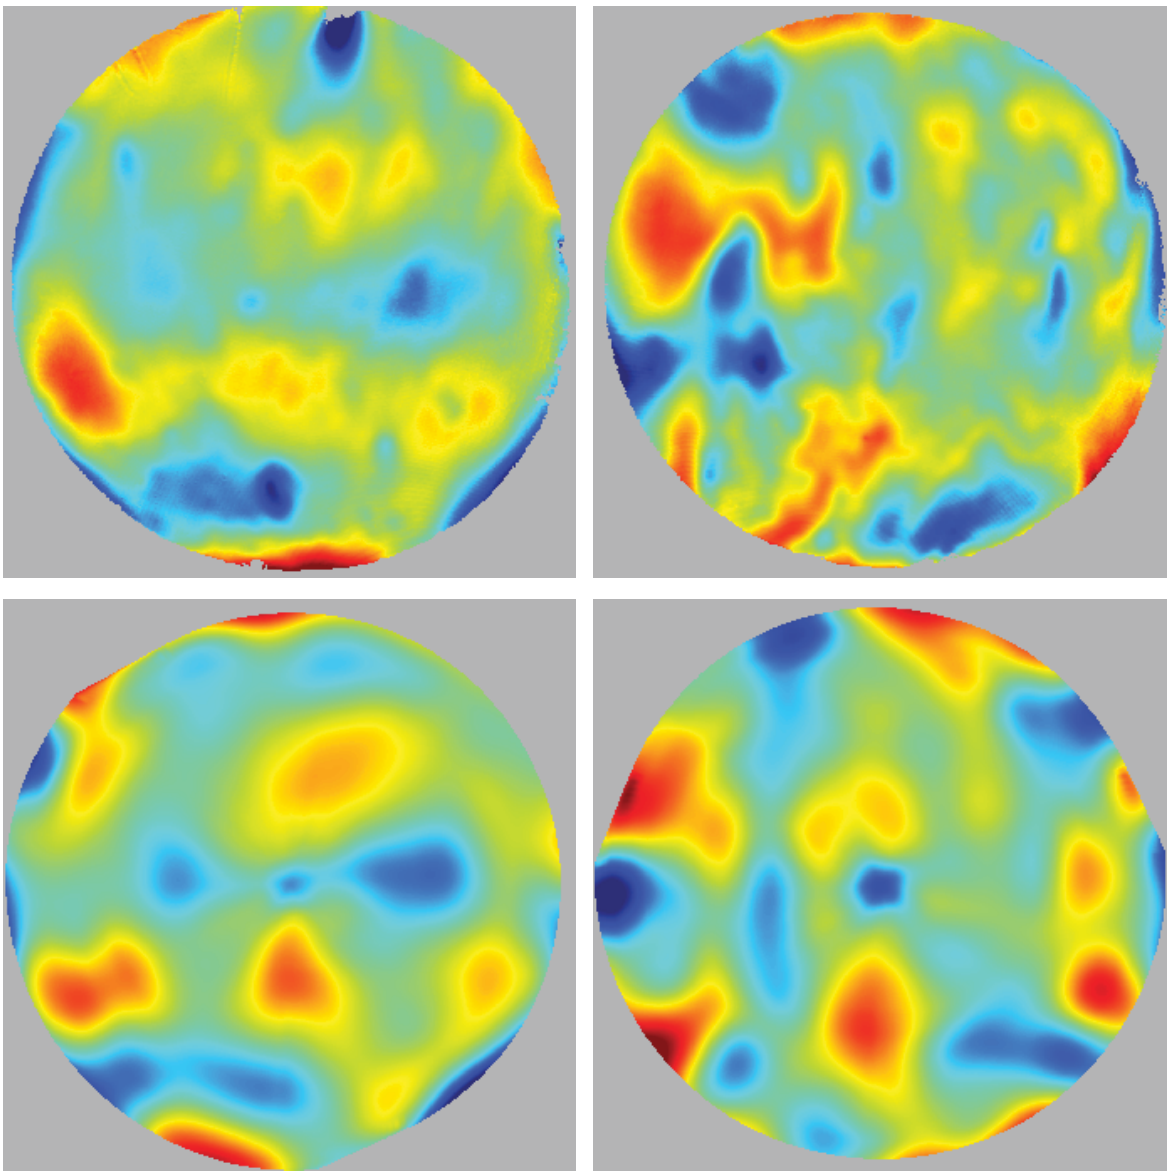

Figure 5.48: Subject A, Zernike removed corneal surface height.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±1.0 µm color height scale; red is away from eye. 15 Zernike terms removed from surface.

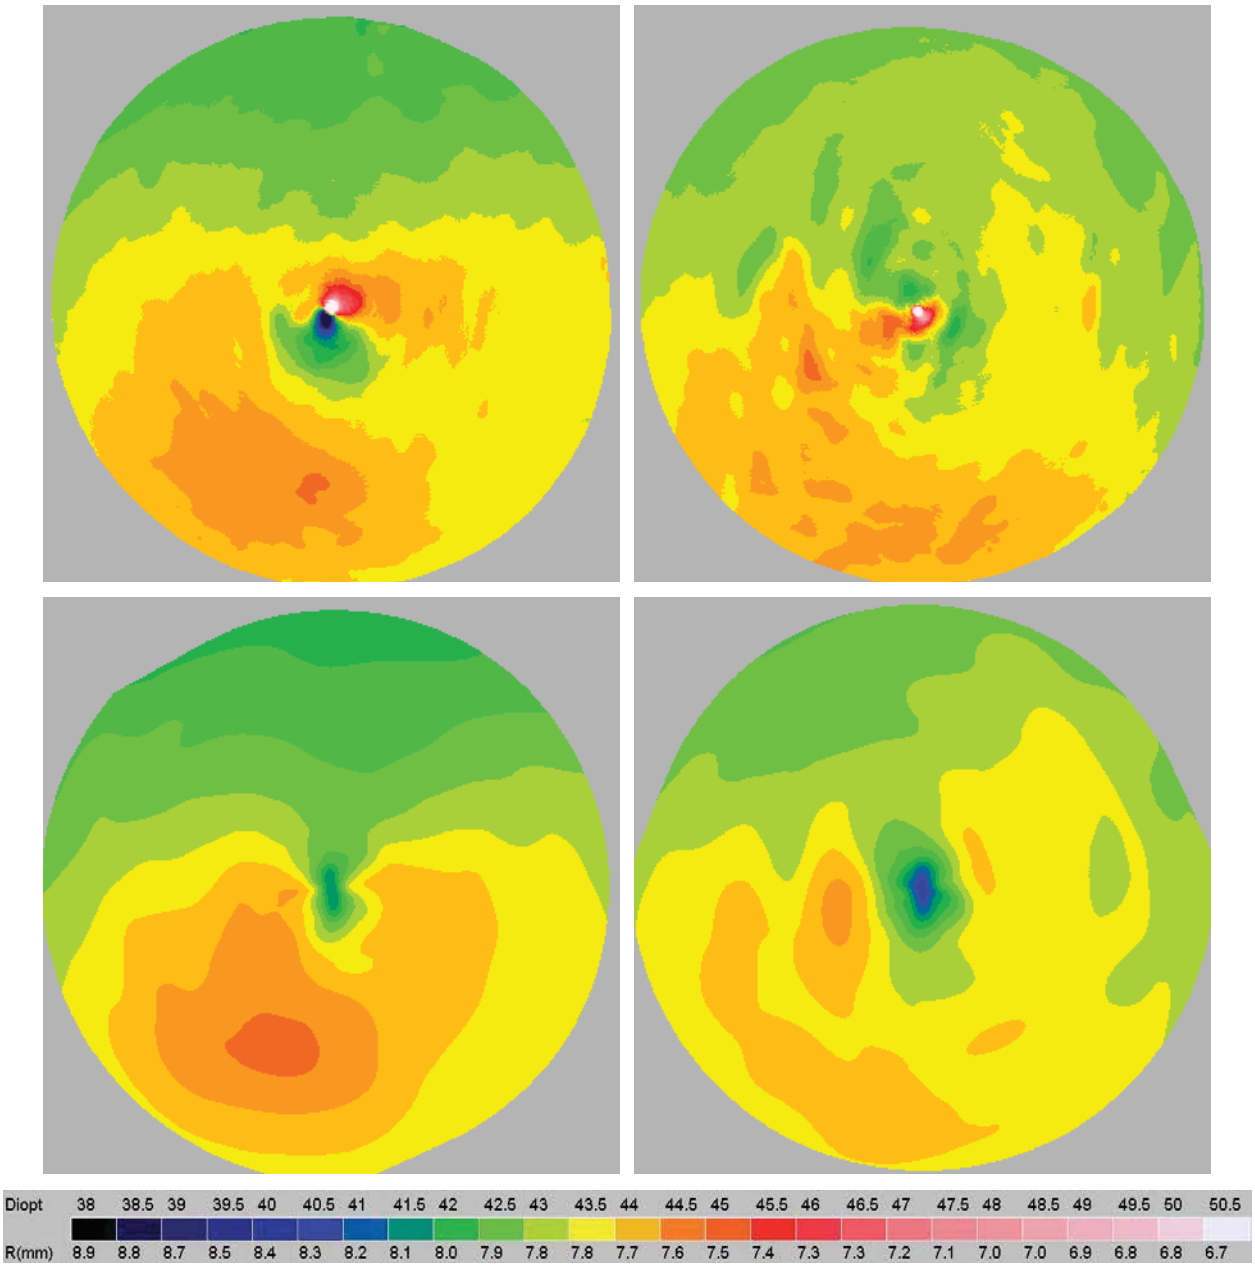

Figure 5.49: Subject A, axial power maps.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm.

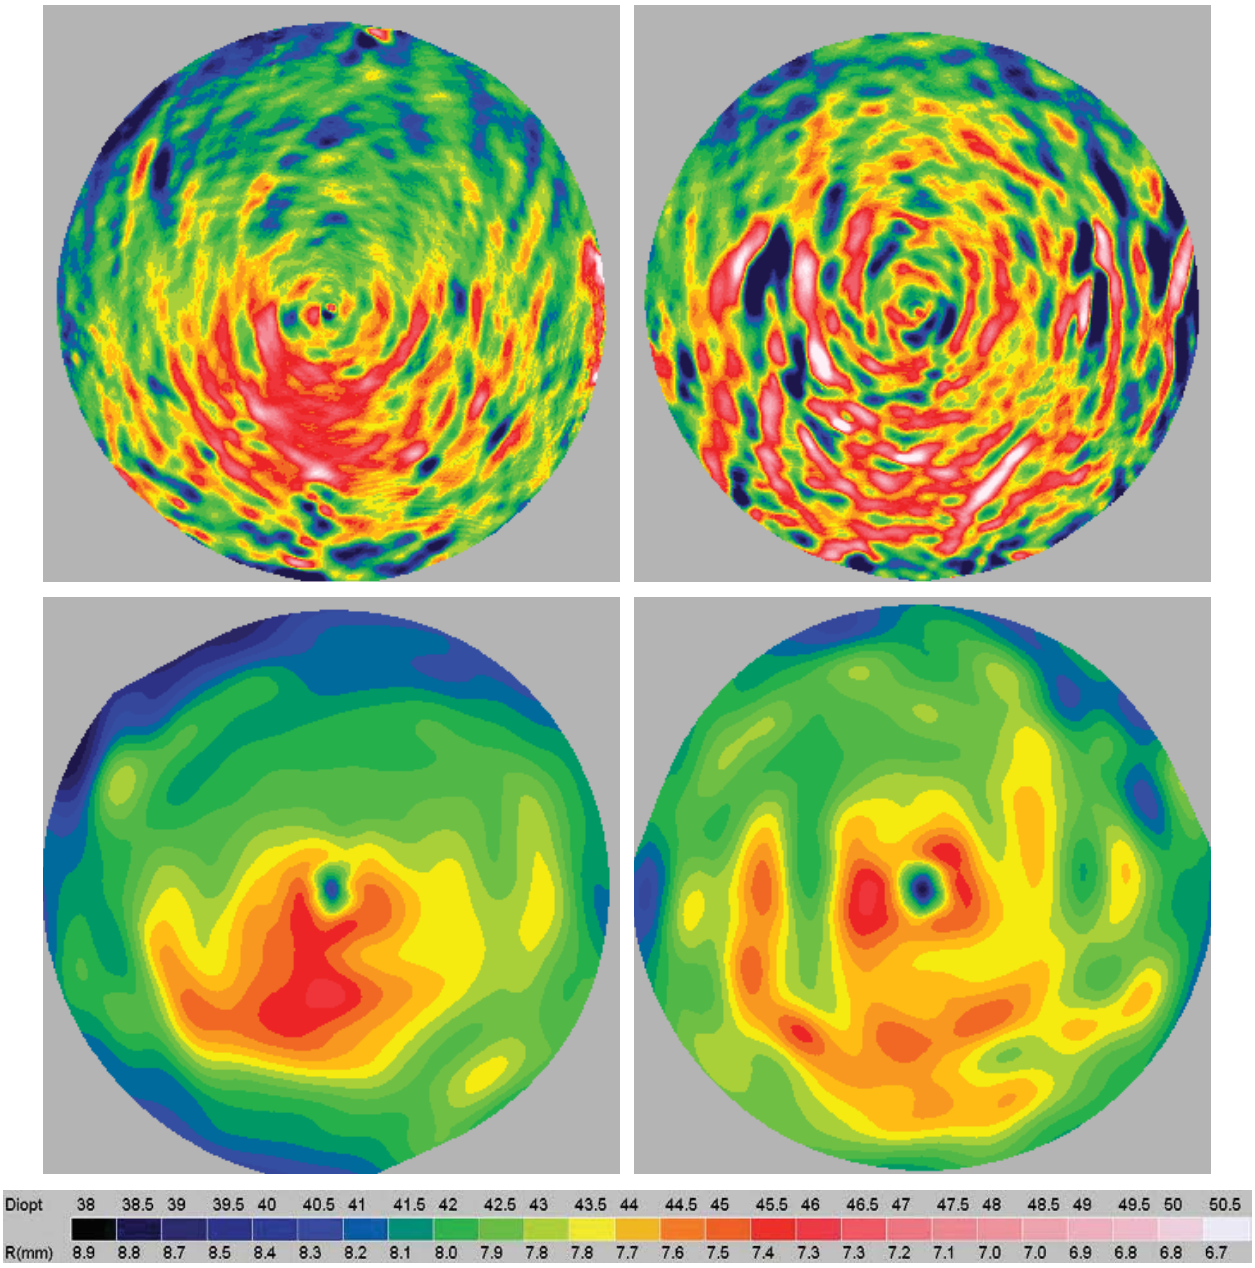

Figure 5.50: Subject A, instantaneous power maps.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm.

5.3.2 SUBJECT B

There is relatively good agreement between the two systems with this subject, which can be seen with the generalized biconic fit parameters shown in Table 5.10. However, the results present an issue that was hinted at with the previous subject: The subject's topography is fairly well represented by a generalized biconic surface that the residuals do not make a good comparison. This is demonstrated with Figure 5.52 and Figure 5.53, where the residuals from the SPINALTOP measurement are below the resolution of the Keratron Piccolo.

The subject had enough corneal astigmatism that a conic removed surface plot was included for comparison (Figure 5.51). Corneal astigmatism can also be viewed in terms of axial power in Figure 5.54. The general structure of the axial power maps are the similar, but the differences are most likely due to the lower spatial height resolution of the Keratron Piccolo. The instantaneous power maps in Figure 5.55 also show similarities, but the mid to high spatial-frequency structure becomes more pronounced with the SPINALTOP measurement due to its significantly higher spatial resolution capabilities.

Another interesting result from these measurements is the presence of mucin globules (approximately 100-200 μm in diameter, left eye) and the blink artifacts (horizontal ridge at the center of the image, right eye) that can be seen in Figure 5.52 and Figure 5.53. The SPINALTOP and Keratron Piccolo measurements are separated in time, which does not allow for the same exact structures to simultaneously appear in the two sets of data. Regardless, these types of features (e.g. mucin globules, blink artifacts) are never present in any of the measurements made with the Keratron Piccolo.

	Left		Right	
	SPINALTOP	**Keratron**	**SPINALTOP**	**Keratron**
Rx [mm]	-8.205	-8.226	-8.168	-8.015
Ry [mm]	-8.048	-7.886	-7.991	-7.941
Kx	-0.345	0.112	-0.218	-0.166
Ky	-0.253	0.043	-0.379	-0.036
θ [deg]	-3.68	3.28	23.00	20.02
$SSIM_B$	14.8%		15.1%	
$SSIM_Z$	26.1%		14.6%	

Table 5.10: Subject B, biconic fit results.

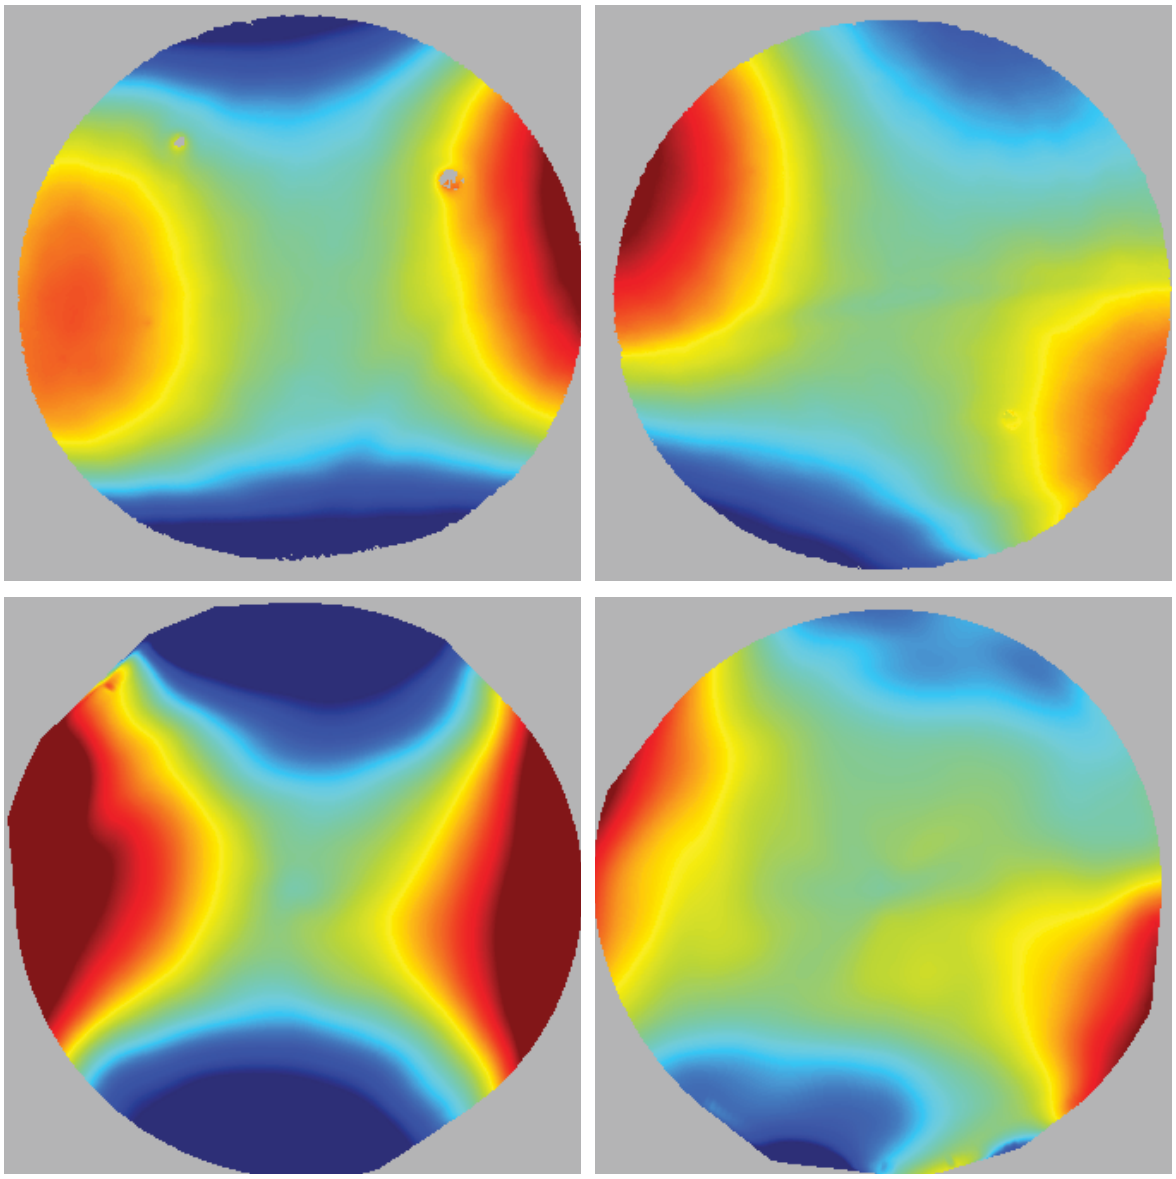

Figure 5.51: Subject B, conic removed corneal surface height.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm. ±12.0 µm color height scale; red is away from eye.

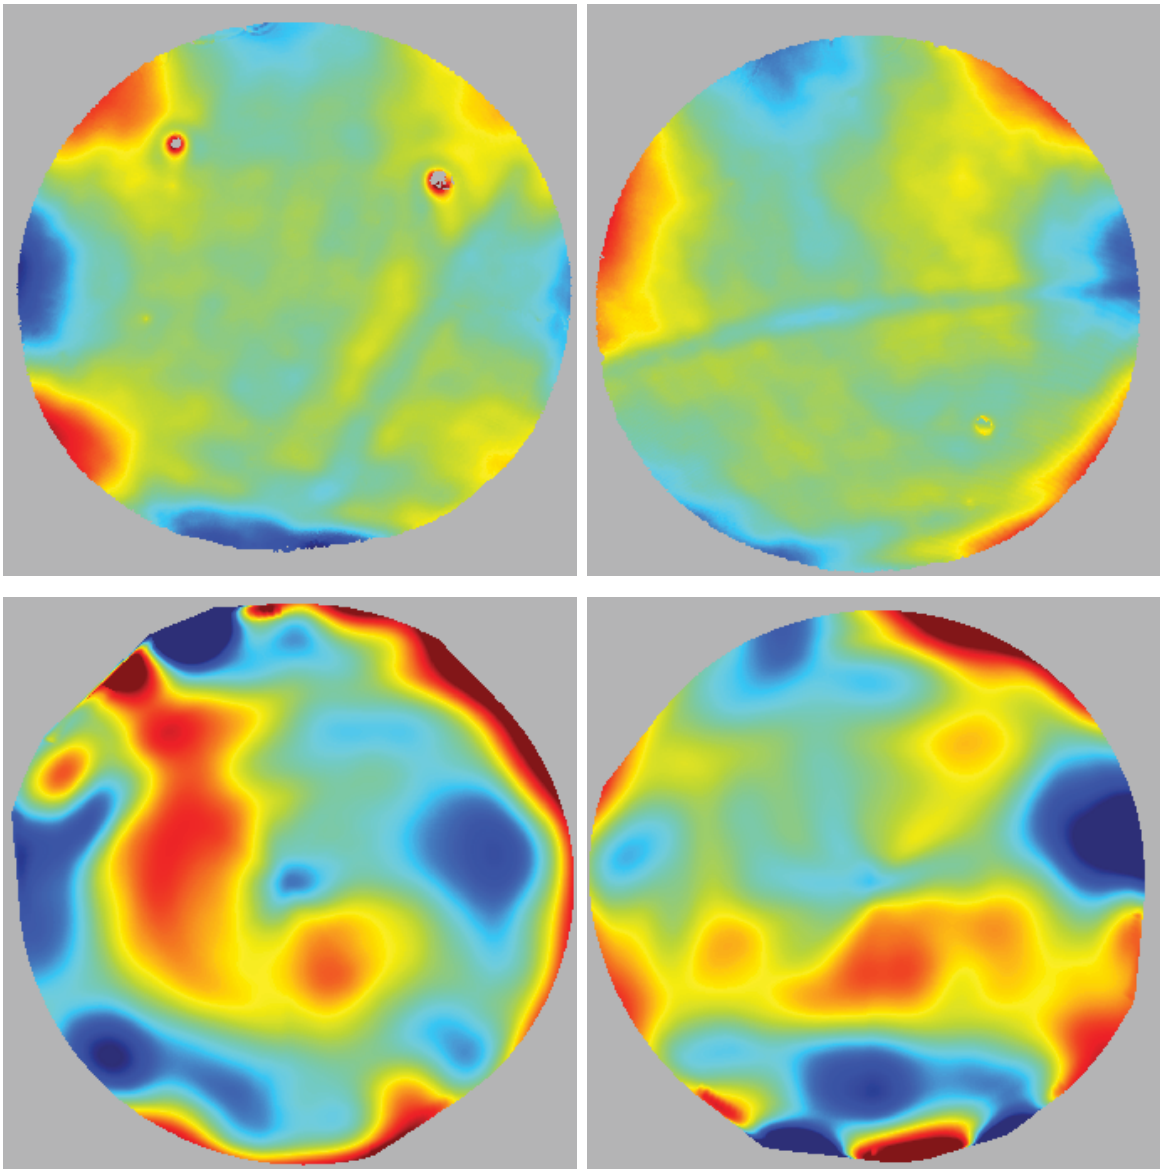

Figure 5.52: Subject B, biconic removed corneal surface height .
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm. ±1.5 µm color height scale; red is away from eye.

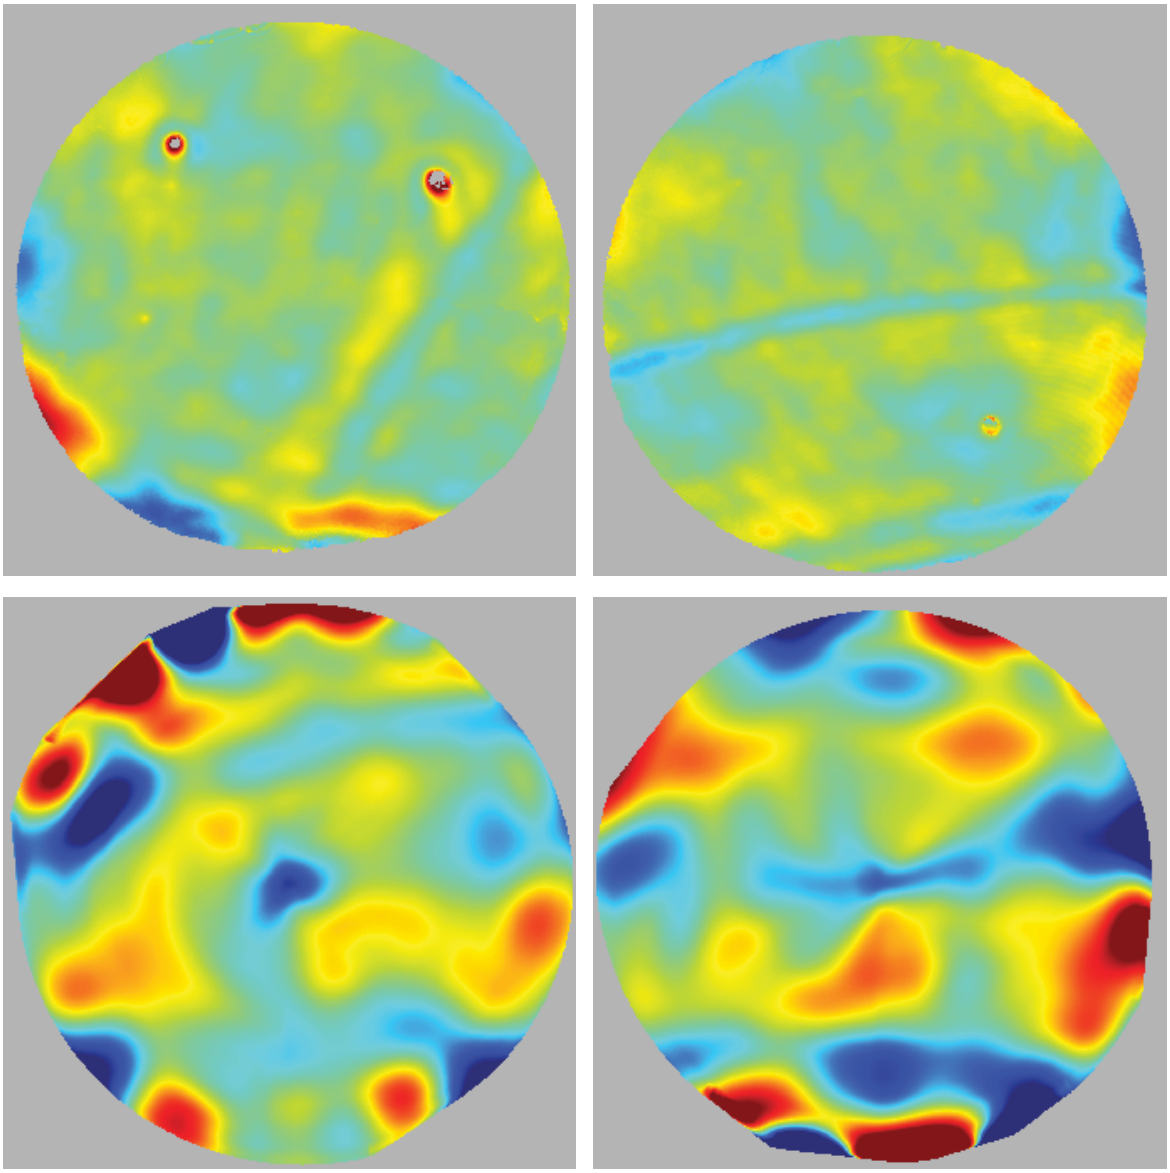

Figure 5.53: Subject B, Zernike removed corneal surface height.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm. ±1.0 µm color height scale; red is away from eye. 15 Zernike terms removed from surface.

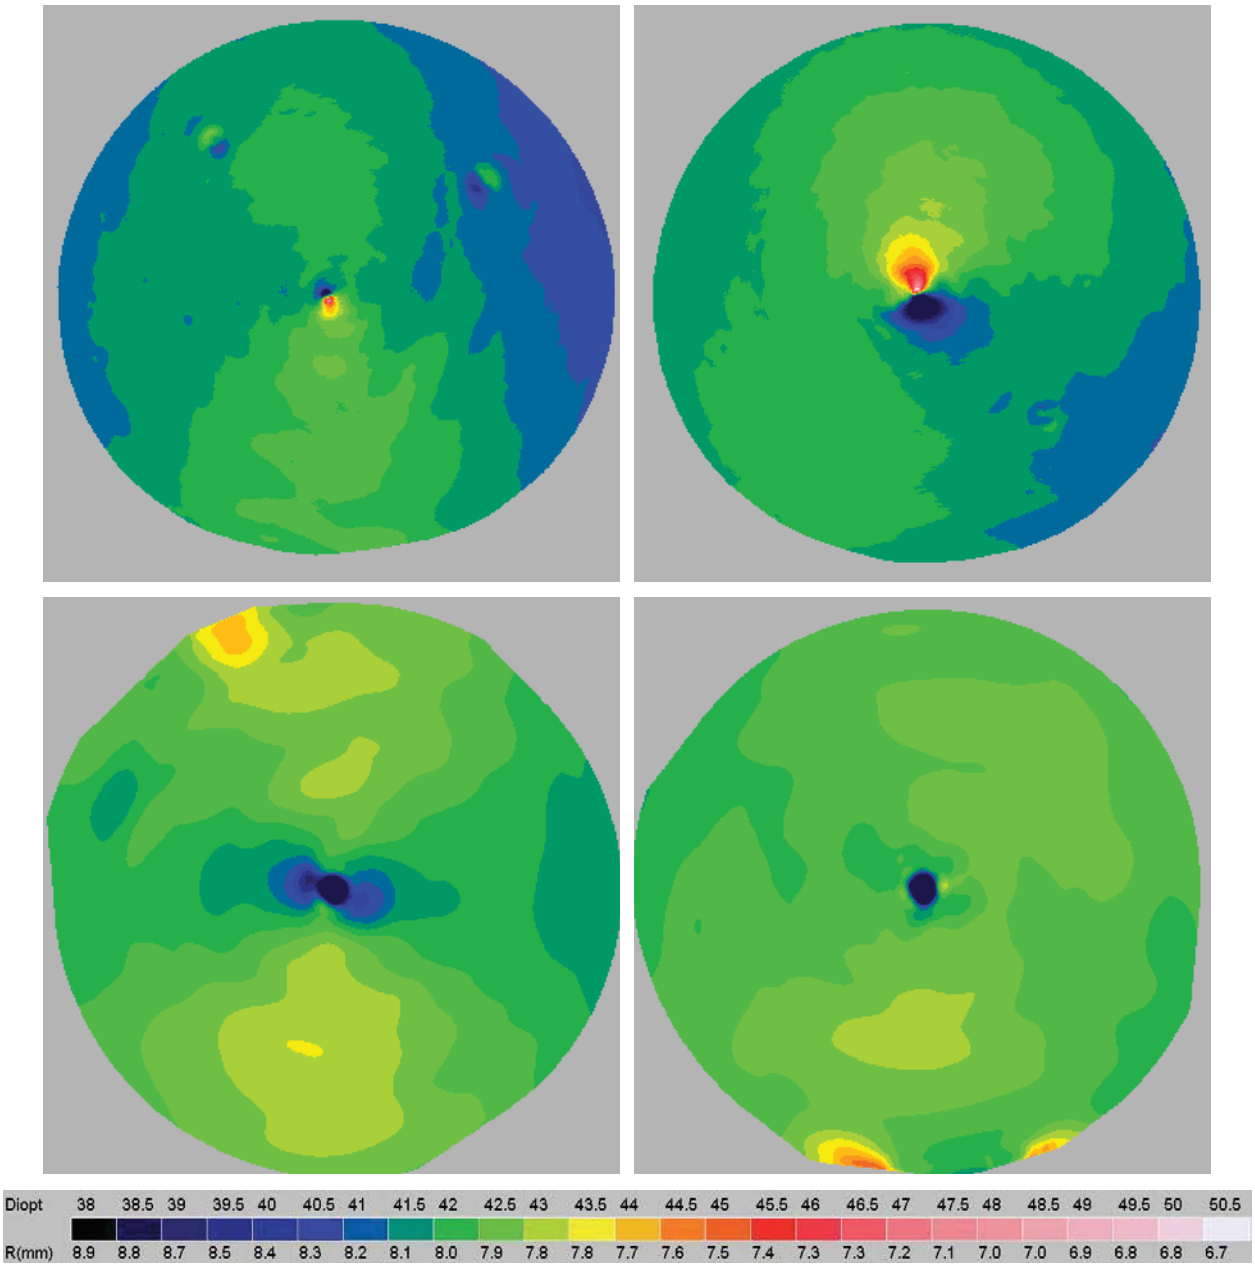

Figure 5.54: Subject B, axial power maps.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm.

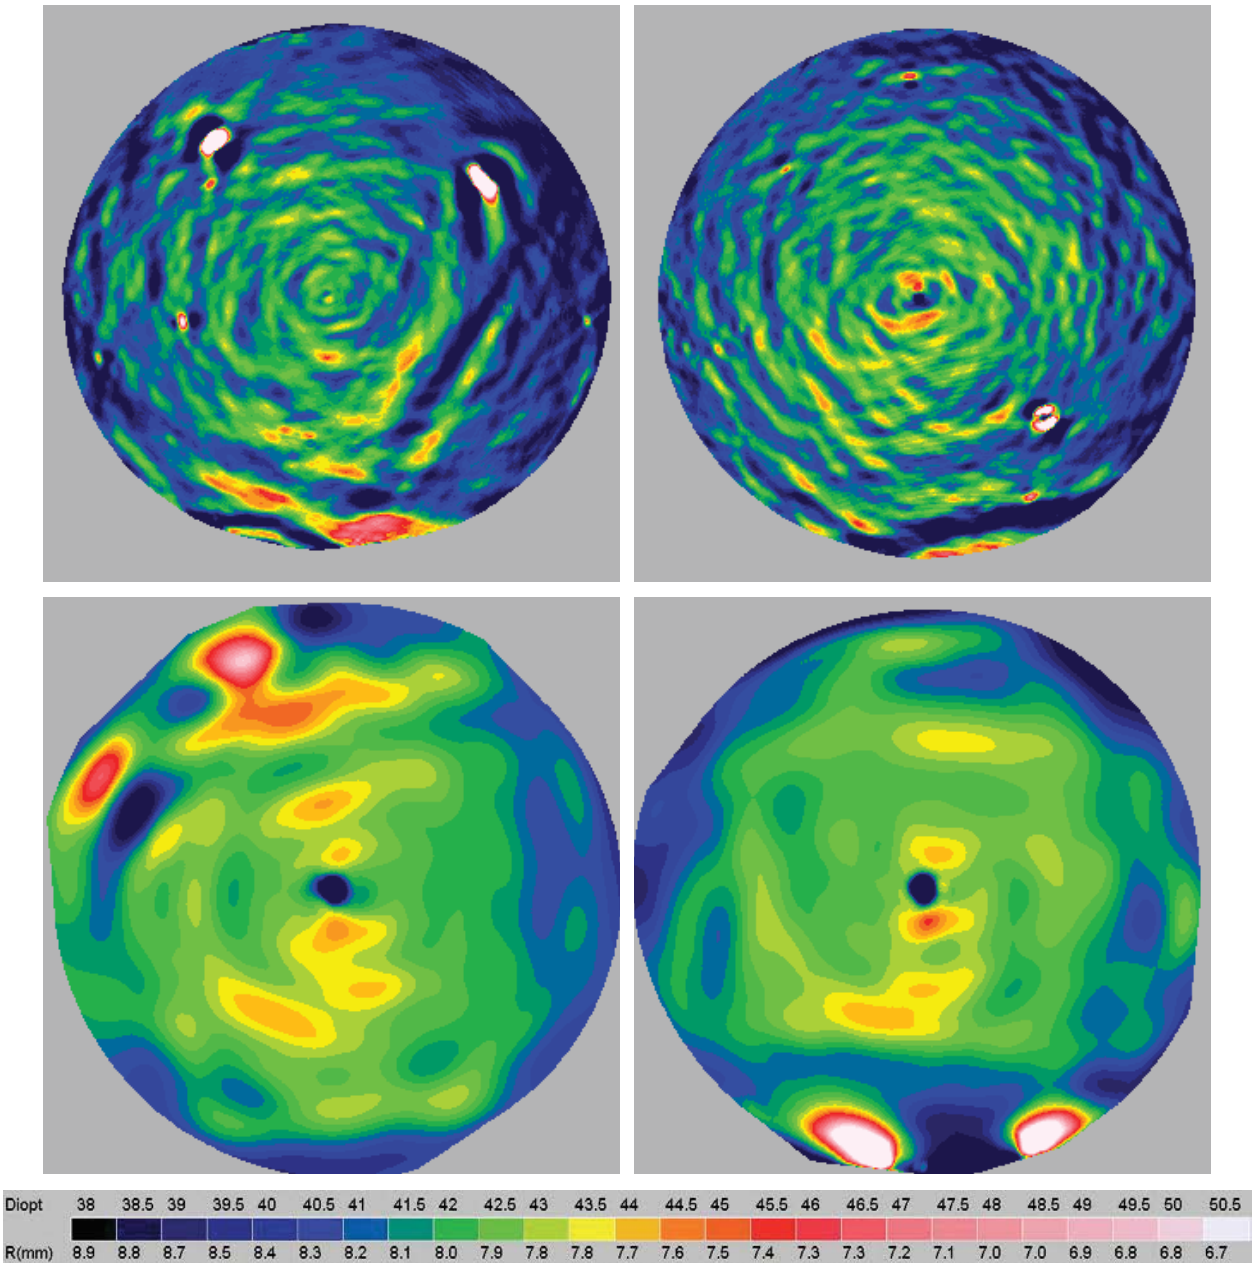

Figure 5.55: Subject B, instantaneous power maps.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm.

5.3.3 SUBJECT C

This subject demonstrated an issue with the Keratron Piccolo system. The circular targets of the Piccolo are back-illuminated by a bright source that puts out a lot of light and heat. This subject reported that the illumination was so bright that their eyes would tear up when attempting to make a measurement. The resulting measurements from the Keratron system were significantly different than SPINALTOP as shown in Table 5.11. The error lies with the Keratron Piccolo optical trip-line (Section 1.6.2). The trip-line can be prematurely tripped by structure on and around the eye which is not the corneal apex. Most often, it is tripped by the subject's eye-lashes. Despite having the subject hold their eyes open with their fingers to avoid this issue, it was difficult to capture valid data. This is not an issue with SPINALTOP.

The issue with the Keratron Piccolo resulted in an underestimation of the shape of the cornea (i.e. smaller radius of curvature). This affects the recovery of other parameters such as the conic constant or any of the low to mid spatial-frequency structure. This can be seen with the generalized biconic removed surfaces in Figure 5.56 and the Zernike removed surfaces in Figure 5.57. Additionally, the Keratron Piccolo results in Figure 5.57 show a structure at the bottom of the cornea which was a result of the tears pooling at the base of the subject's eye. The underestimated radius measurement with the Keratron Piccolo results in the skewed power maps in Figure 5.58 and Figure 5.59.

	Left		Right	
	SPINALTOP	Keratron	SPINALTOP	Keratron
Rx [mm]	-7.848	-6.640	-7.871	-7.114
Ry [mm]	-7.880	-6.854	-7.928	-7.233
Kx	-0.235	0.864	-0.180	0.597
Ky	-0.307	1.273	0.027	0.869
θ [deg]	9.18	7.27	21.50	19.08
$SSIM_B$	21.8%		20.1%	
$SSIM_Z$	5.8%		37.6%	

Table 5.11: Subject C, biconic fit results.

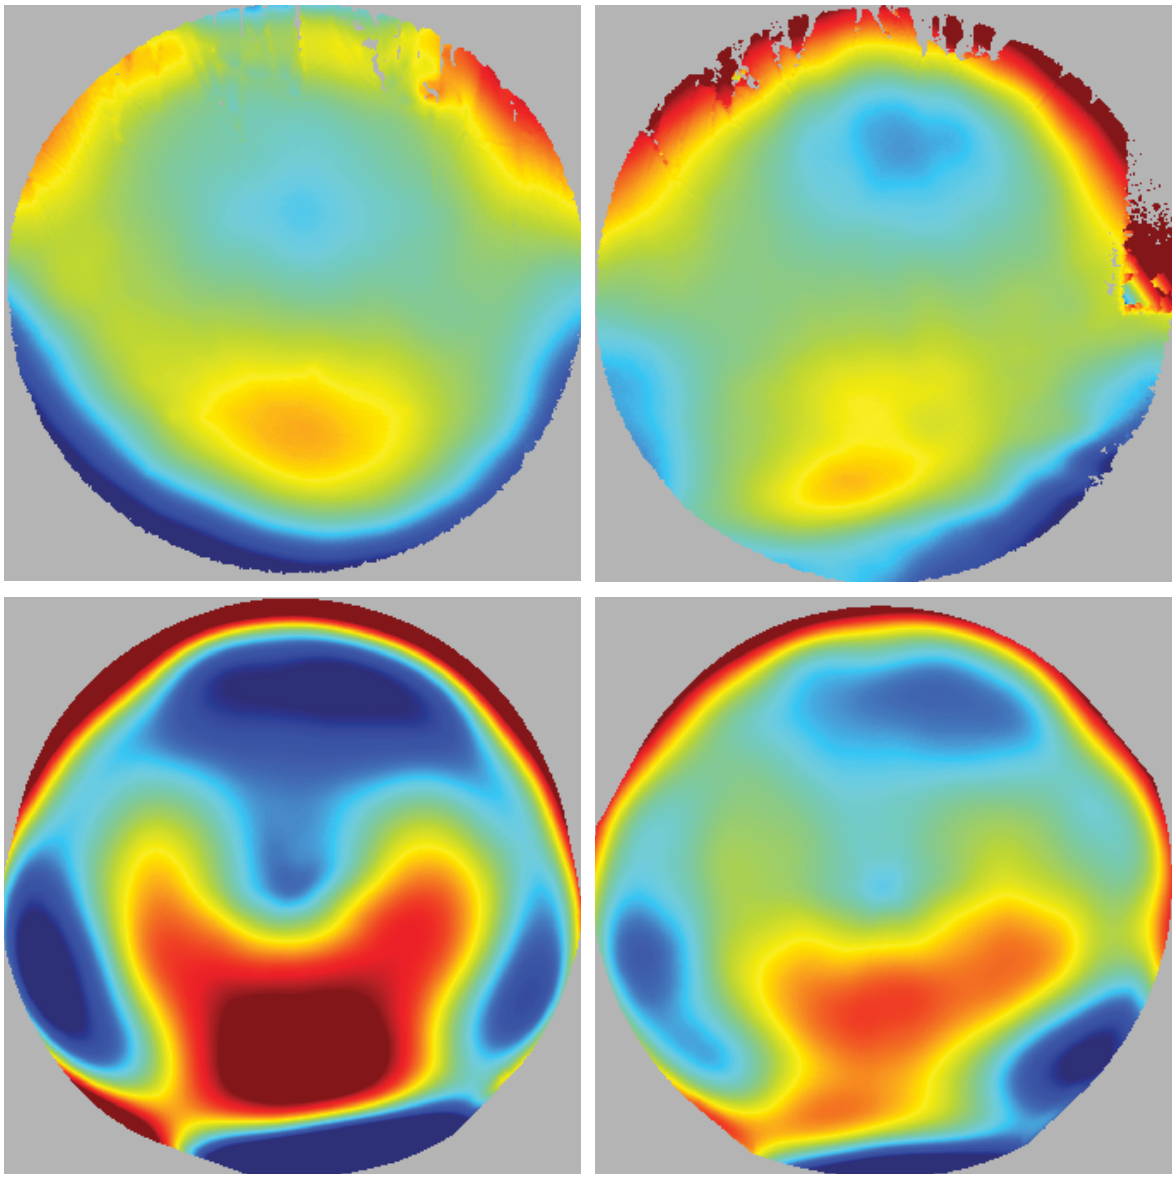

Figure 5.56: Subject C, biconic removed corneal surface height.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±3.0 μm color height scale; red is away from eye.

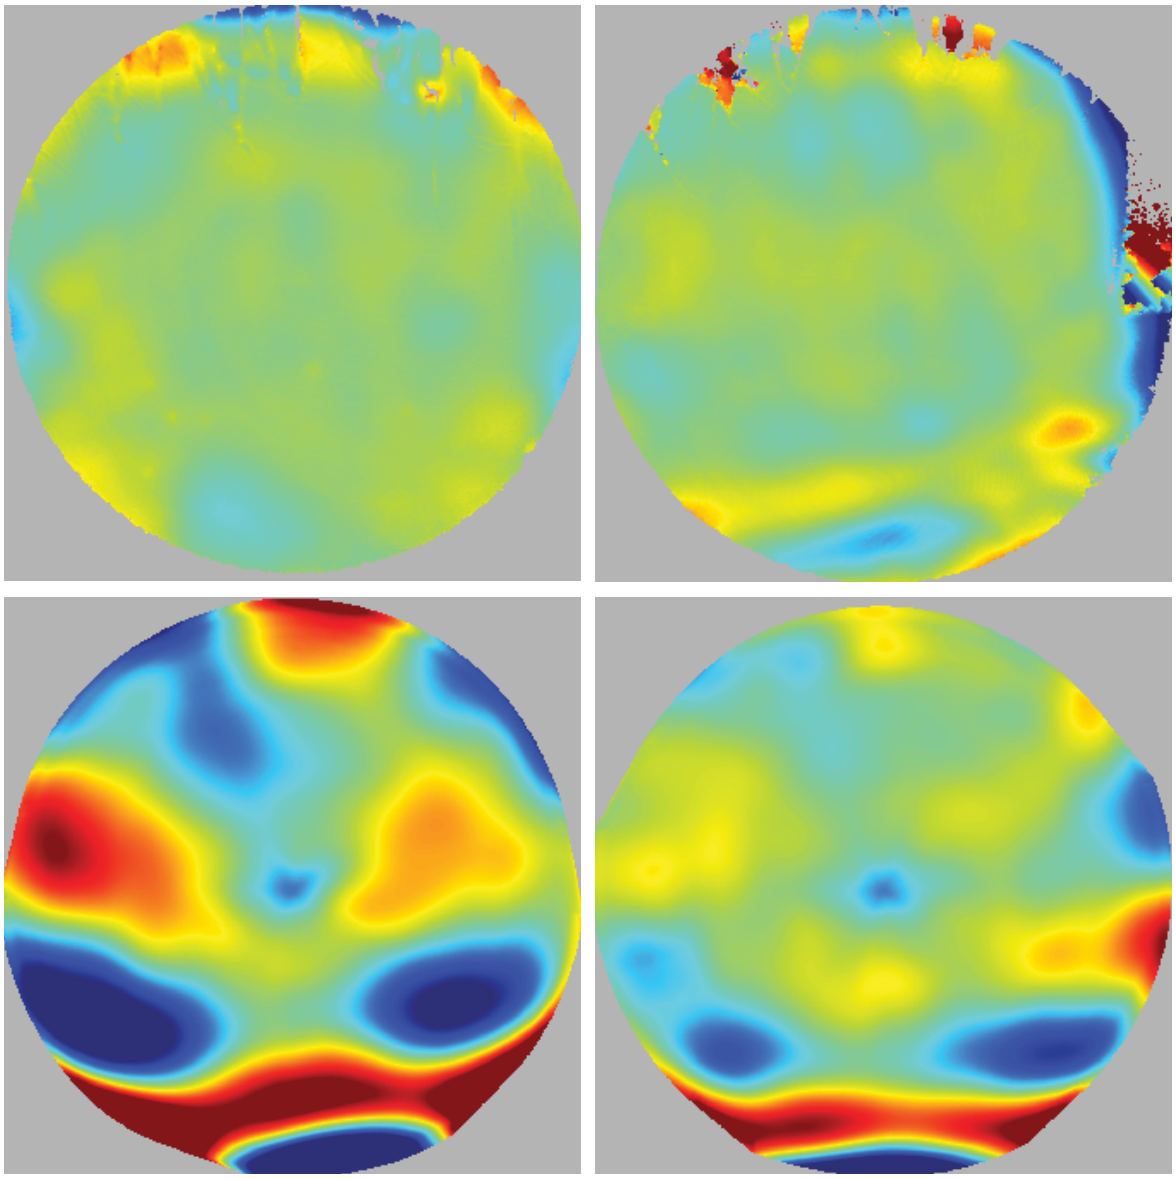

Figure 5.57: Subject C, Zernike removed corneal surface height.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±1.5 µm color height scale; red is away from eye. 15 Zernike terms removed from surface.

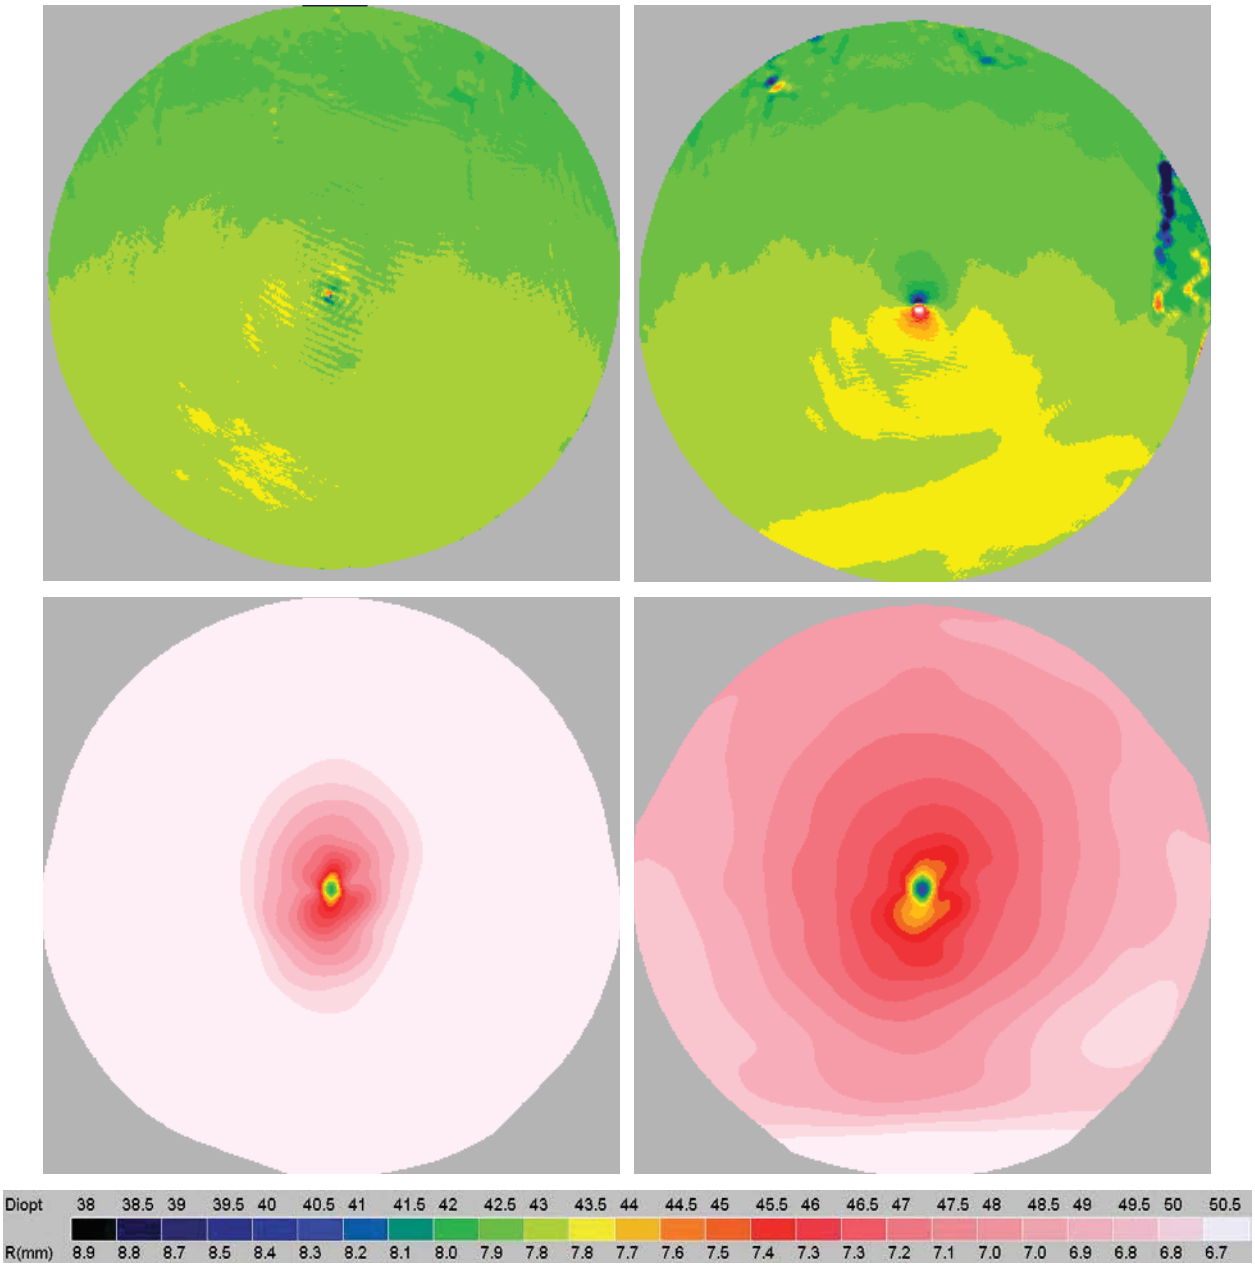

Figure 5.58: Subject C, axial power maps.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm.

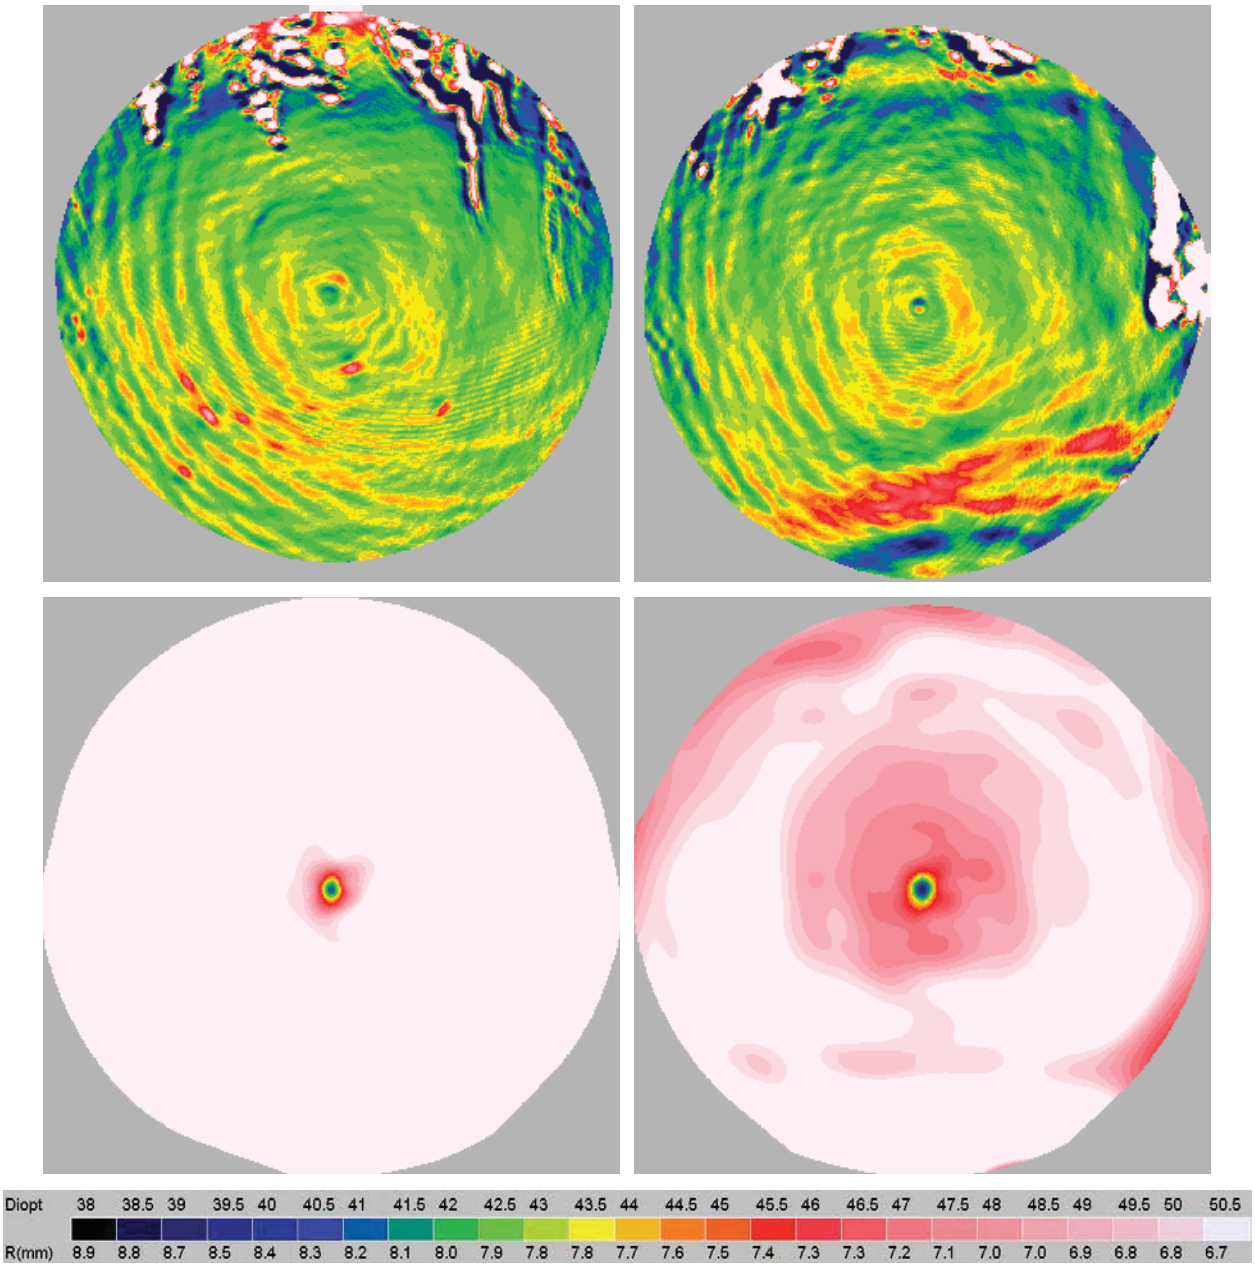

Figure 5.59: Subject C, instantaneous power maps.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm

5.3.4 SUBJECT D

This subject has enough unique topography that spans low to mid spatial-frequency structure to provide a good comparison between the two systems. The generalized biconic fit results (Table 5.12) are in good agreement. The conic removed surfaces from both systems shows comparable amounts of corneal astigmatism (Figure 5.60). The biconic removed surface also shows good agreement (Figure 5.61), but the mid to high-spatial frequency structure is only resolvable in the SPINALTOP measurement. The Zernike removed surface further demonstrates the resolution capabilities of the SPINALTOP system (Figure 5.62). The axial power maps are in good agreement, showing the presence of corneal astigmatism (Figure 5.63). The instantaneous power maps also show good agreement, but also show the influence of high spatial-frequency structure in the SPINALTOP measurement (Figure 5.64).

	Left		Right	
	SPINALTOP	Keratron	SPINALTOP	Keratron
Rx [mm]	-7.708	-7.594	-7.597	-7.577
Ry [mm]	-7.441	-7.339	-7.369	-7.339
Kx	-0.334	-0.248	-0.320	-0.174
Ky	-0.318	-0.285	-0.388	-0.237
θ [deg]	-8.08	-11.56	2.89	5.77
$SSIM_B$	86.2%		77.7%	
$SSIM_Z$	44.3%		75.1%	

Table 5.12: Subject D, biconic fit results.

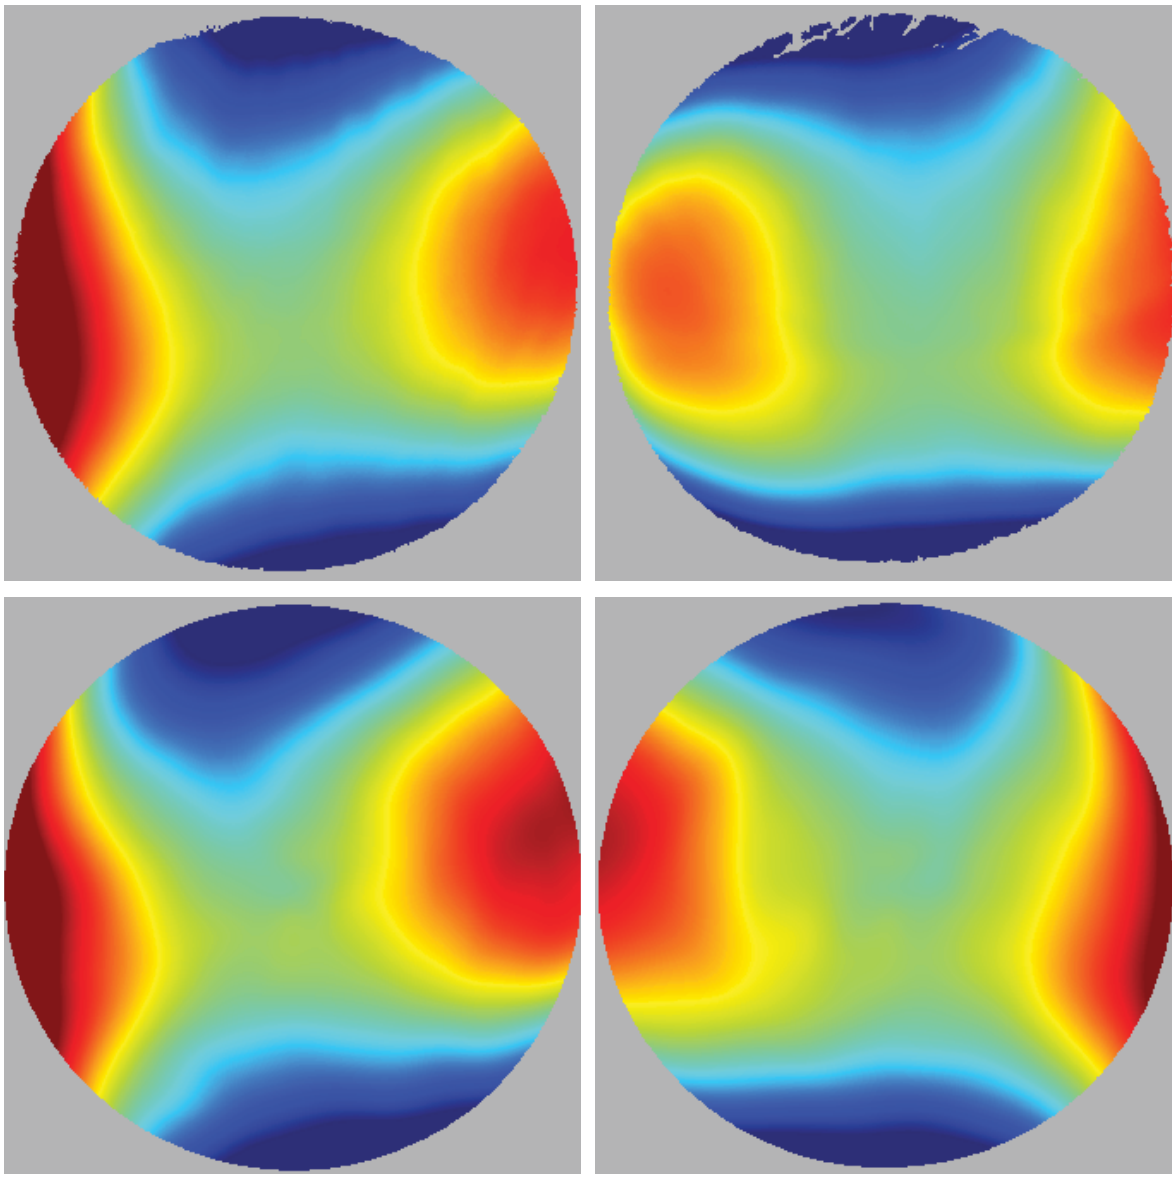

Figure 5.60: Subject D, conic removed corneal surface height.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±16.0 μm color height scale; red is away from eye.

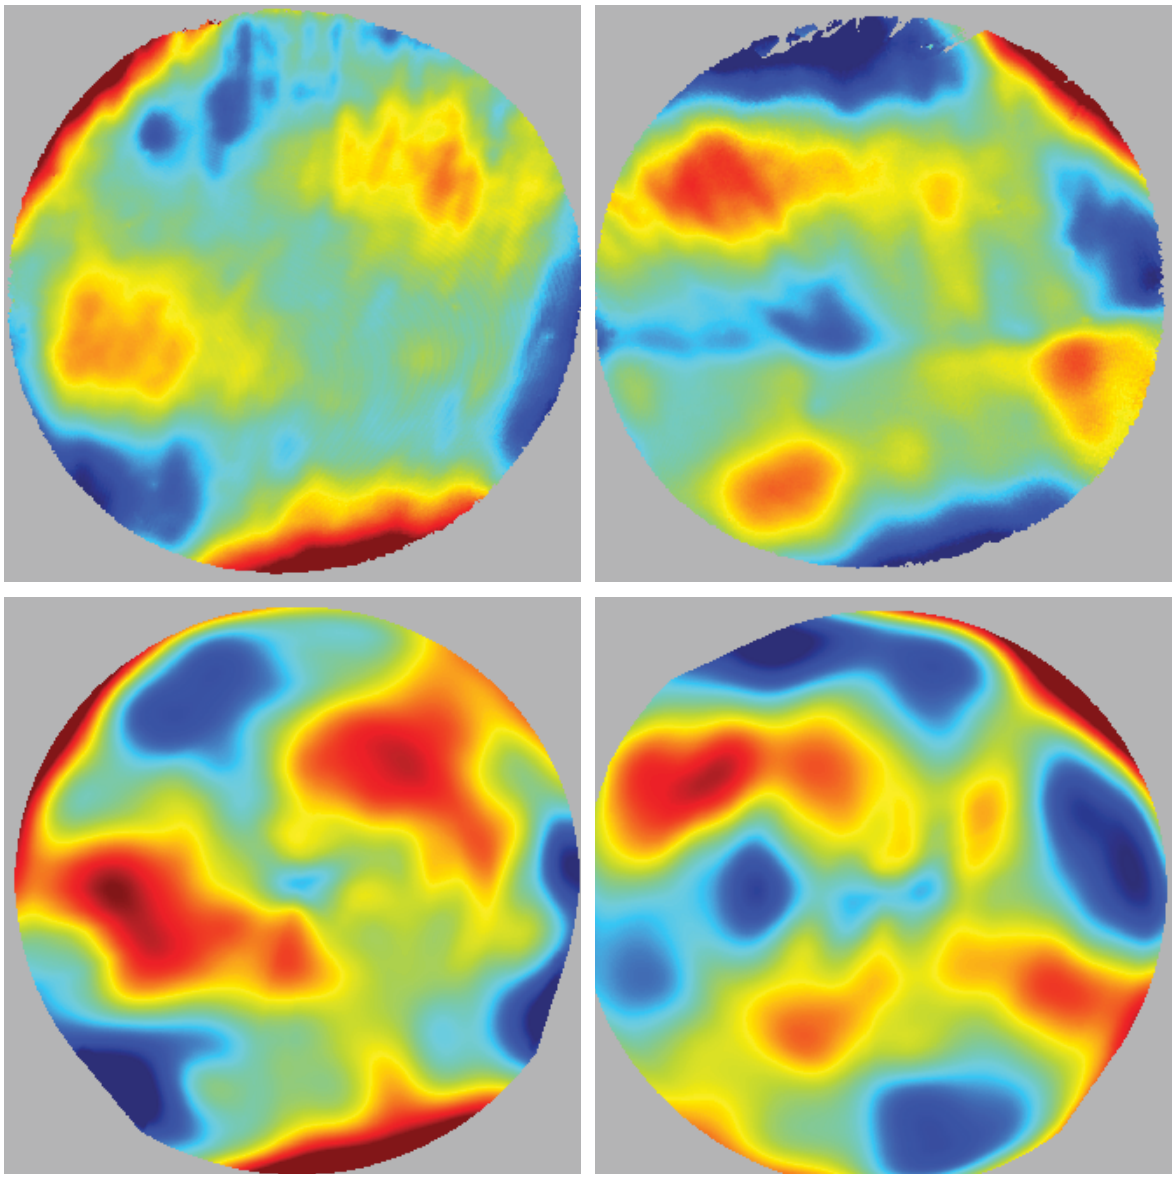

Figure 5.61: Subject D, biconic removed corneal surface height.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±1.25 μm color height scale; red is away from eye.

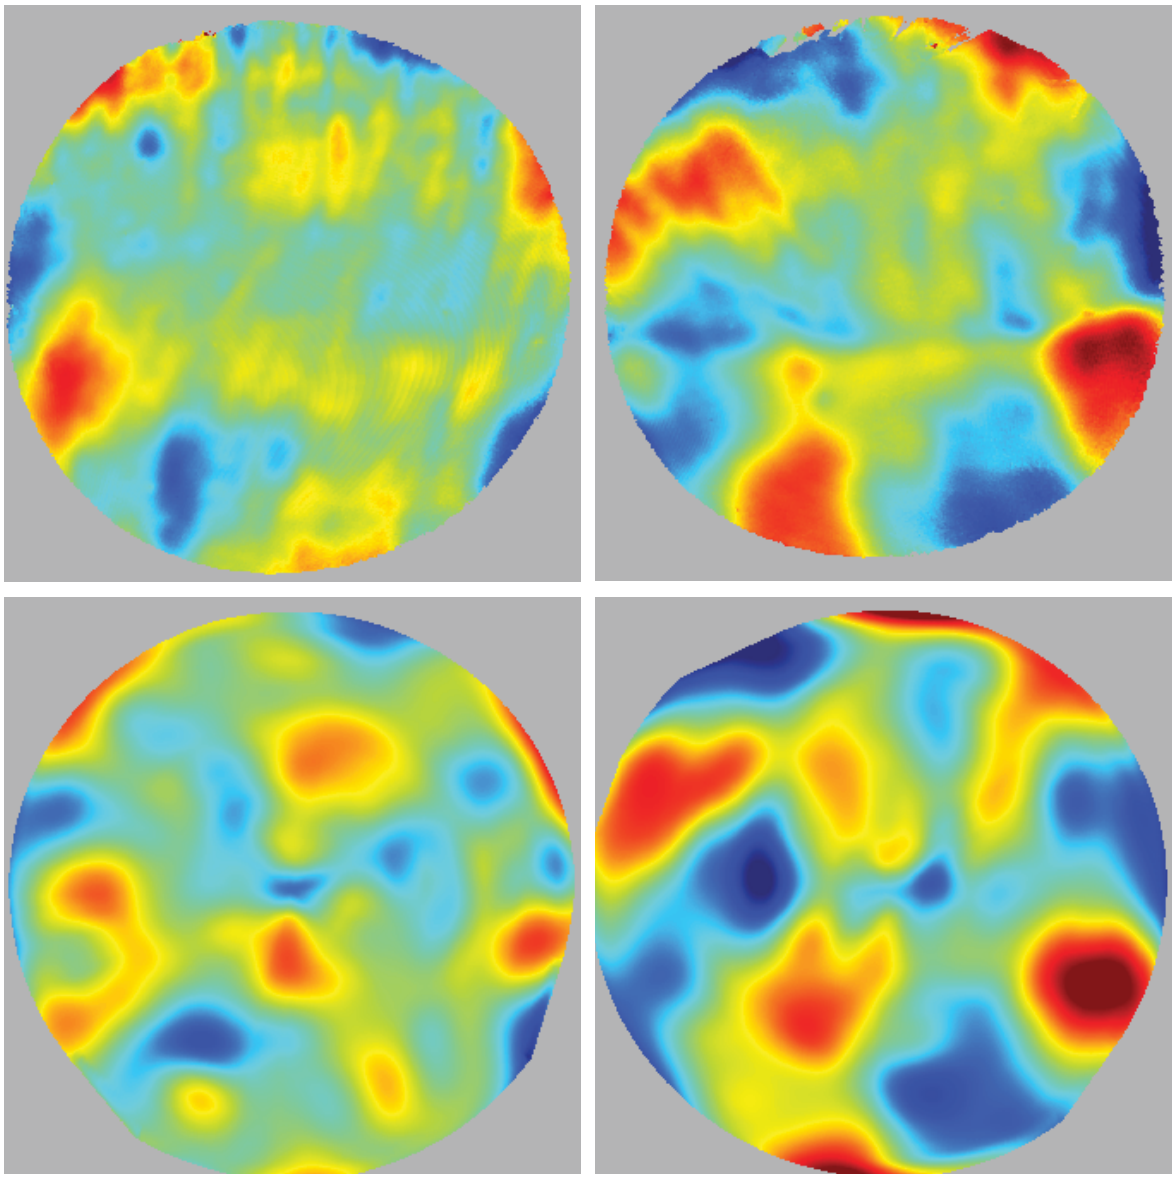

Figure 5.62: Subject D, Zernike removed corneal surface height.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±1.0 µm color height scale; red is away from eye. 15 Zernike terms removed from surface.

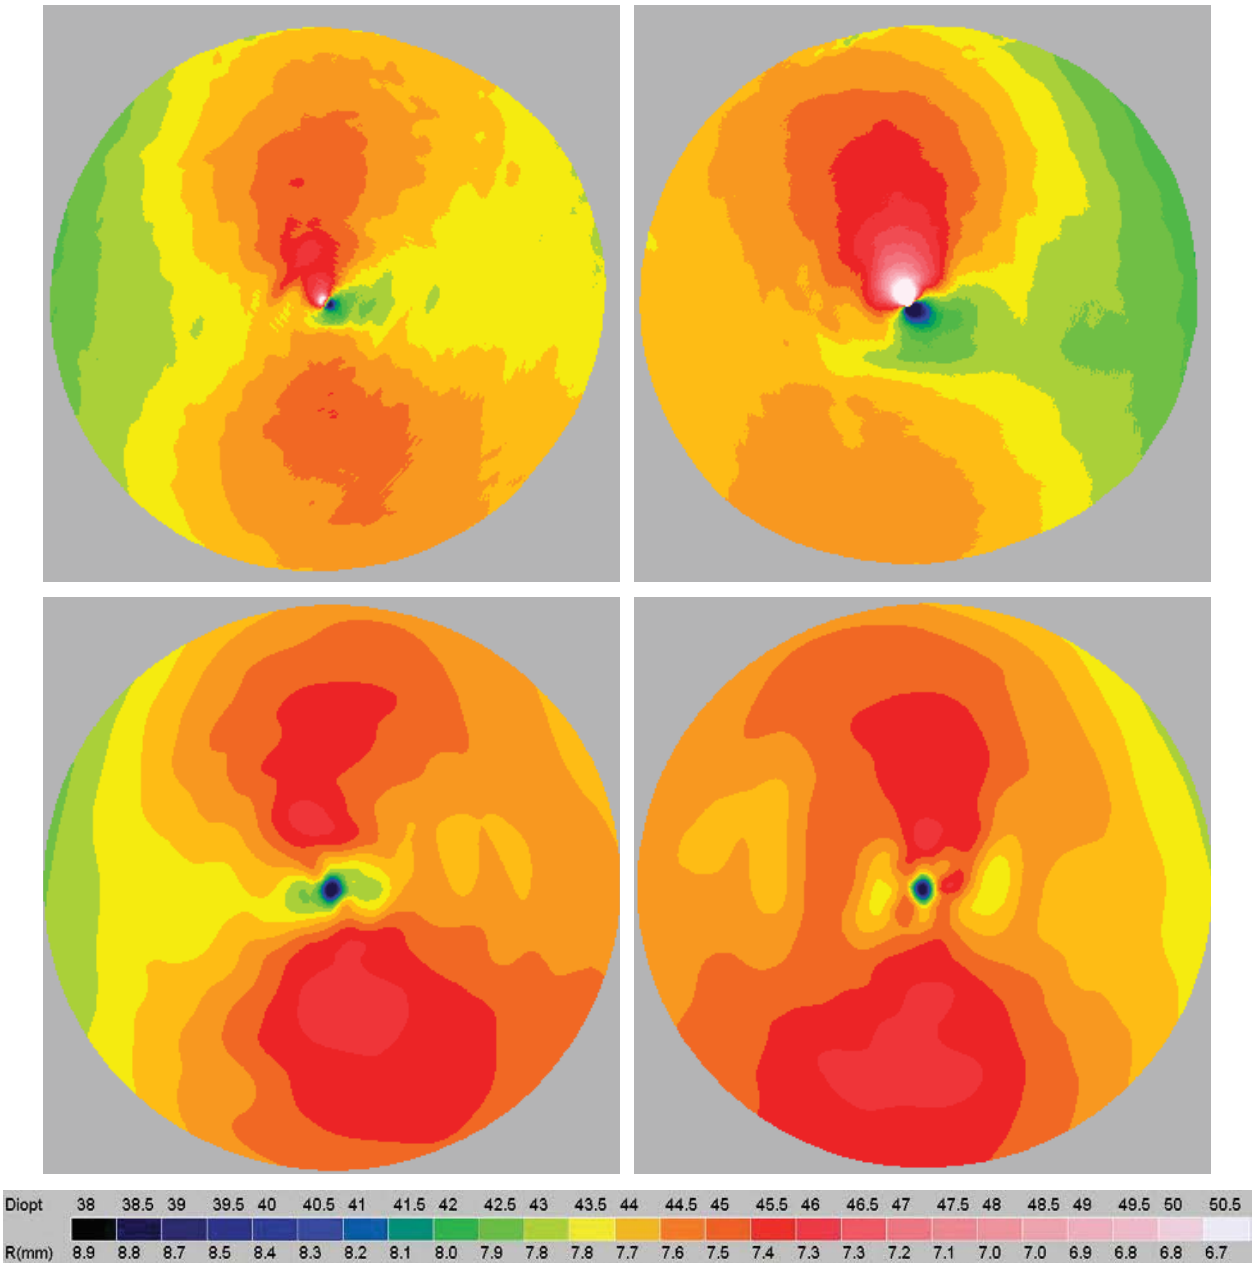

Figure 5.63: Subject D, axial power maps.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm.

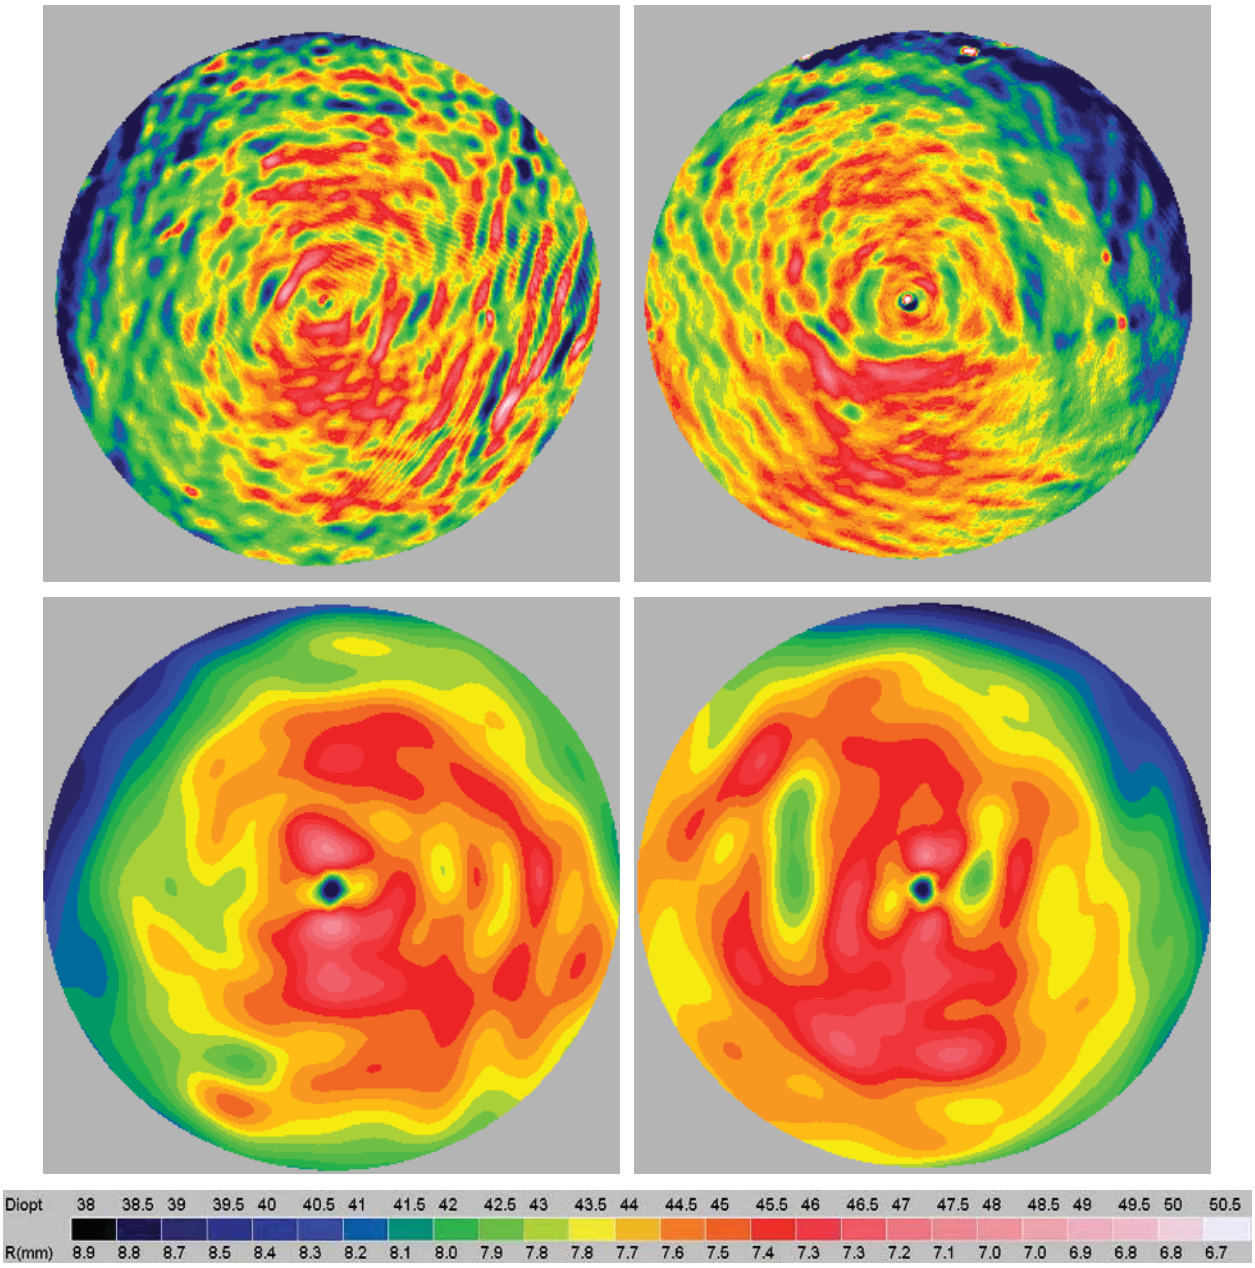

Figure 5.64: Subject D, instantaneous power maps.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm.

5.3.5 SUBJECT E

The biconic fit parameters and SSIM index values are given in Table 5.13. The subject has relatively little corneal astigmatism, so biconic fit residuals are presented (Figure 5.65). The low to mid spatial-frequency structure in the biconic residuals appears to correlate well for the left eye. The right eye has less structure (i.e. < 1 μm) which is near the limit of resolution for the Keratron Piccolo. The Zernike removed surfaces (Figure 5.66) further demonstrates the amount of high spatial-frequency information that cannot be resolved by the Keratron Piccolo. There is significant high spatial-frequency information that describes the tear film structure of this subject that the SPINALTOP system can resolve. Finally, power maps are presented in Figure 5.67 and Figure 5.68. The power maps are in good agreement between the two systems.

	Left		Right	
	SPINALTOP	Keratron	SPINALTOP	Keratron
Rx [mm]	-8.020	-8.041	-7.936	-8.031
Ry [mm]	-8.109	-8.009	-7.922	-7.956
Kx	-0.115	0.020	-0.133	-0.051
Ky	0.058	-0.022	-0.044	-0.081
θ [deg]	-36.58	-12.32	20.54	-12.24
$SSIM_B$	74.2%		28.3%	
$SSIM_Z$	24.5%		42.6%	

Table 5.13: Subject E, biconic fit results.

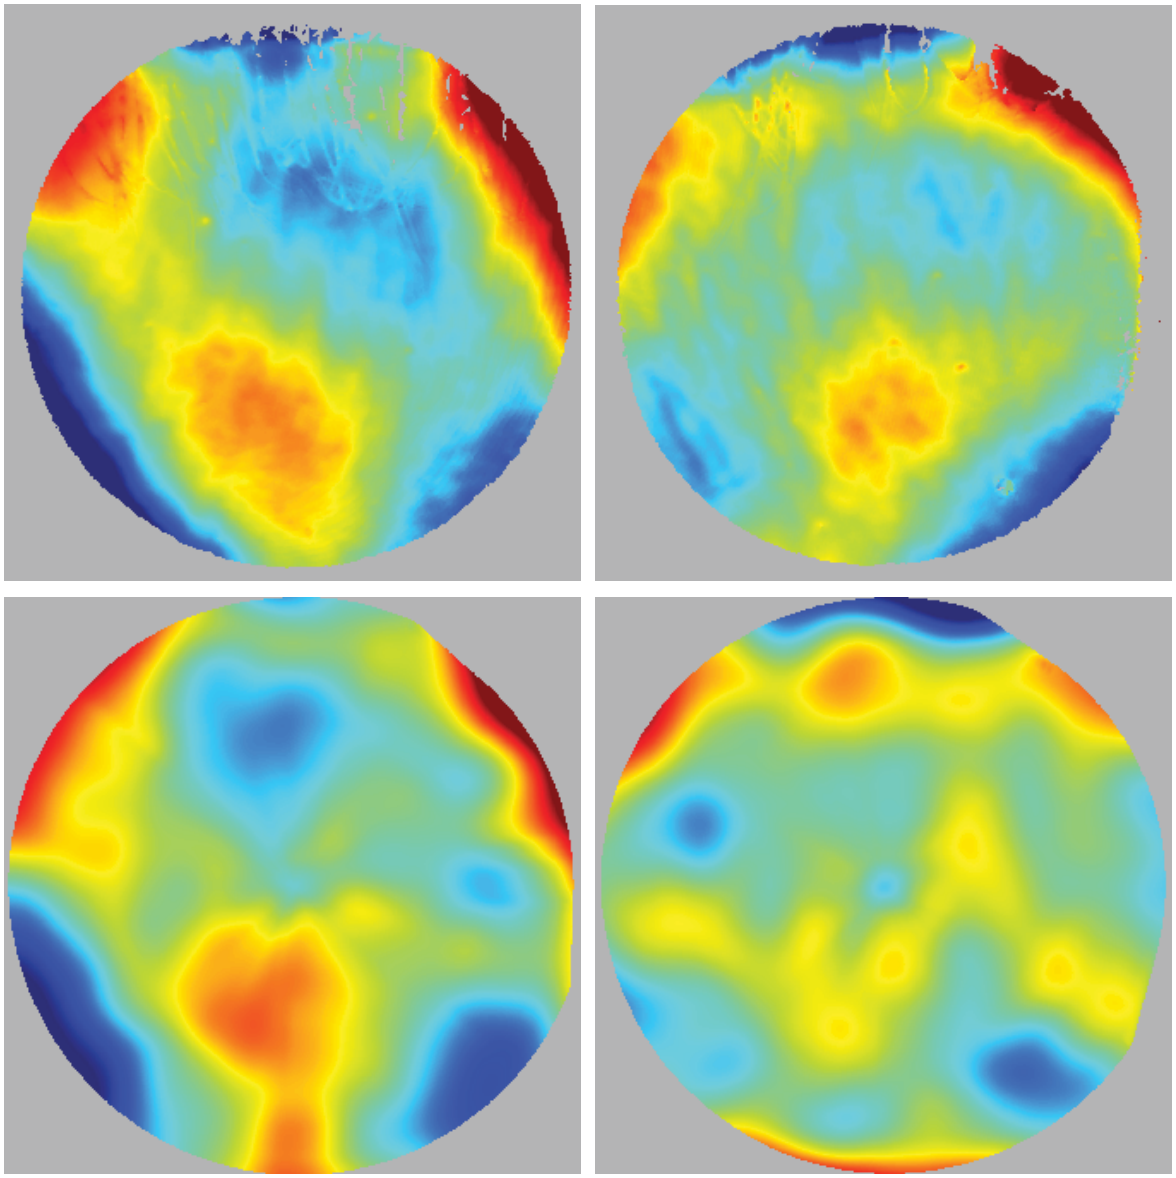

Figure 5.65: Subject E, biconic removed corneal surface height.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm. ±1.5 µm color height scale; red is away from eye.

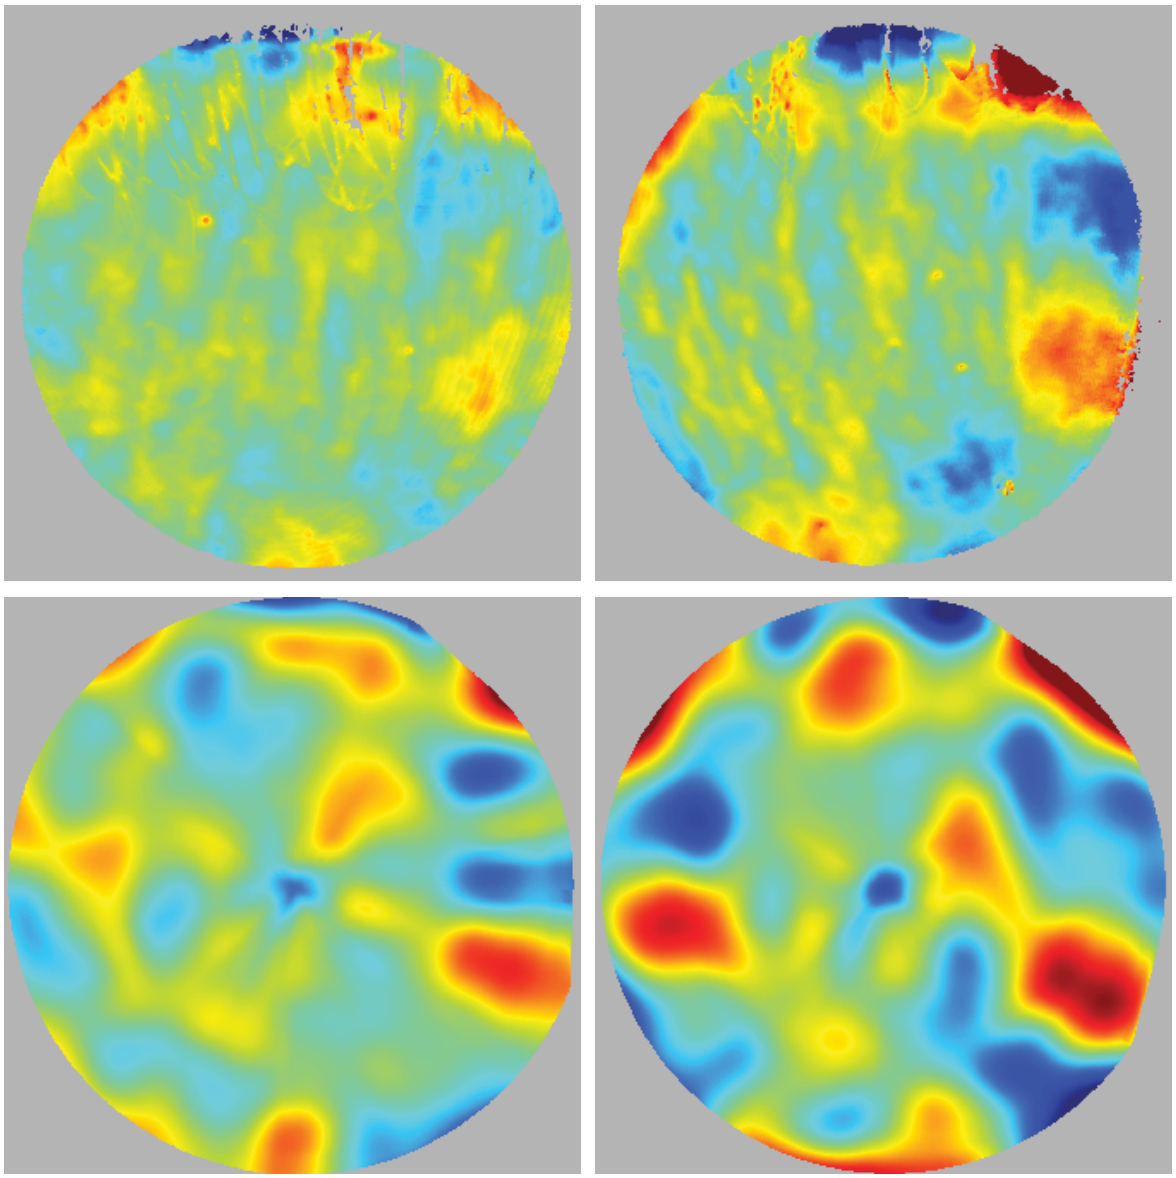

Figure 5.66: Subject E, Zernike removed corneal surface height.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm. ±0.75 µm color height scale; red is away from eye. 15 Zernike terms removed from surface.

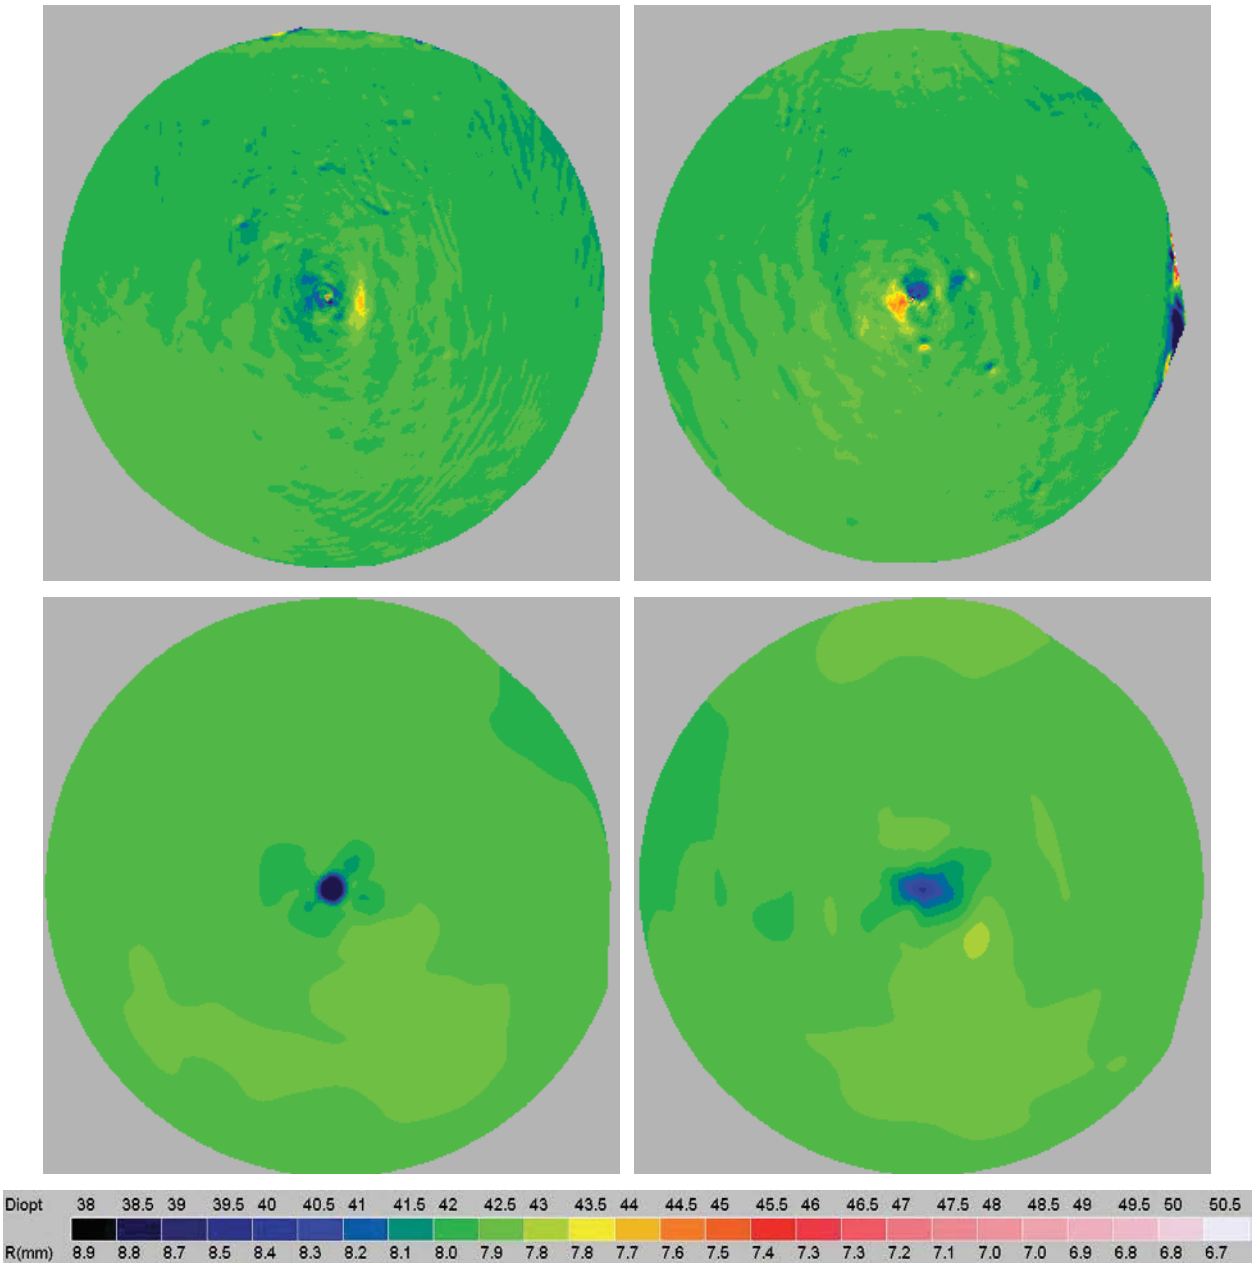

Figure 5.67: Subject E, axial power maps.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm.

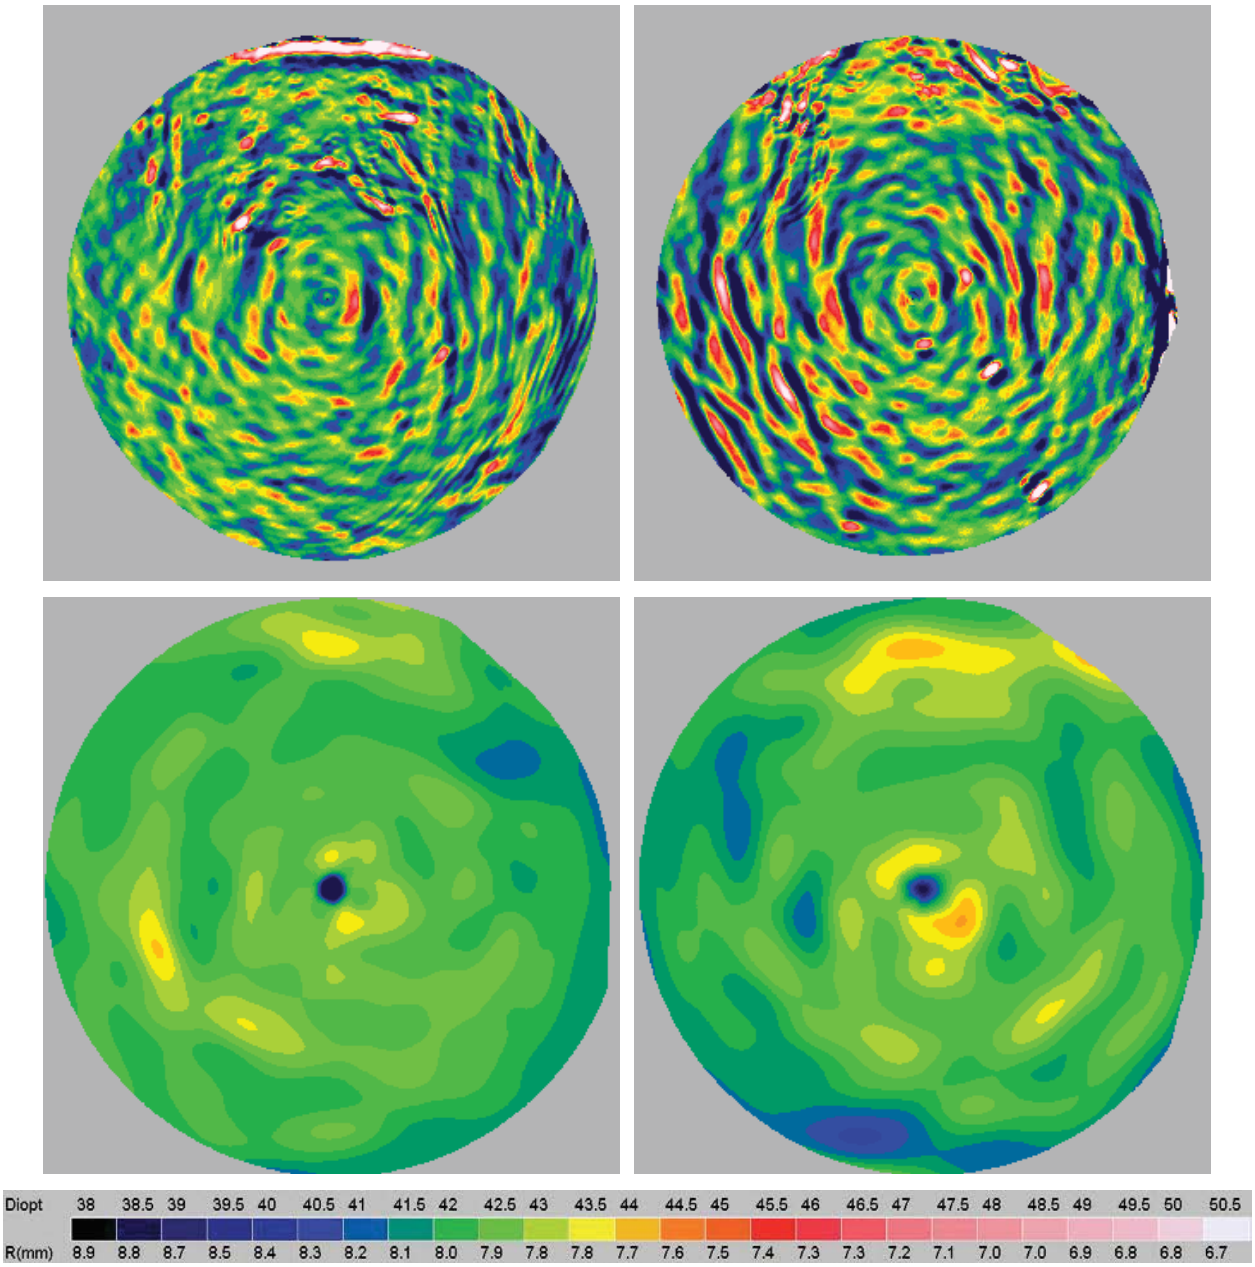

Figure 5.68: Subject E, instantaneous power maps.
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm.

5.3.6 SUBJECT F

The results from this subject are only shown because they contain measurements with a subject wearing contact lenses. This data was collected during an early learning phase, where one of the lessons learned from these measurements was to ask the subject to hold their eyes open. The live display option in the software now allows the operator to quickly determine if the subject needs to hold their eyes open. Data is presented on the subject's bare eyes followed by the subject wearing contact lens. The subject wore Oasys (Acuvue) contact lenses that have a prescription of -5.5 Diopters.

5.3.6.1 BARE EYE

The biconic fit parameters and SSIM index values are shown in Table 5.14. The low SSIM values are in part due to the differences in the measurements and the eyelashes that appear in the SPINALTOP measurements. Although eyelashes appear to corrupt some areas of the data in the SPINALTOP measurement, the remaining data is still provides valid information. This is provided as a comparison to the Keratron Piccolo measurement from Subject C (Section 5.3.3) where eyelashes corrupted the entire measurement due to the optical trip-line.

This subject has a small amount of corneal astigmatism, so conic removed residual surfaces are presented (Figure 5.69). Measurements from both systems appear to correlate well with the right eye, but not with the left eye. The SPINALTOP measurement reports that the left eye has approximately 0.3 D of corneal astigmatism, whereas the Keratron Piccolo measurements report almost none. However, the verification results (Section 5.2.2) show that the uncertainty in the recovery of corneal astigmatism from a Keratron Piccolo measurement is on the order of 0.5 D (1σ), whereas the uncertainty from a SPINALTOP measurement is 0.01 D (1σ). Therefore, it is most likely that the conic removed surface for the Keratron Piccolo measurement is dominated by noise and is therefore unable to resolve an amount of astigmatism that is reported by SPINALTOP.

The biconic removed surfaces (Figure 5.70) show the same issues as the conic removed, where the Keratron Piccolo measurements are close to the noise level and result in features that aren't as well correlated with SPINALTOP measurements. The Zernike removed surfaces (Figure 5.71) show the high spatial-frequency structure of the tear film for SPINALTOP, compared to the low/mid spatial-frequency noise that dominates most of the Keratron measurements. The Zernike removed SPINALTOP measurements (top row) show a ridge near the top of the display, just below the eyelashes, which is most likely a pool of tears held in place by surface tension with the upper eyelid.

The axial power measurements in Figure 5.72 show good agreement between the two systems. The same is true for the instantaneous power measurements in Figure 5.73, but the presence of eyelashes is now more pronounced in the SPINALTOP measurements.

	Left		Right	
	SPINALTOP	**Keratron**	**SPINALTOP**	**Keratron**
Rx [mm]	-8.035	-8.110	-8.086	-8.144
Ry [mm]	-7.979	-8.065	-8.026	-8.019
Kx	-0.191	0.124	-0.215	-0.111
Ky	-0.008	-0.250	0.027	-0.106
θ [deg]	29.20	16.03	-23.88	-32.44
$SSIM_B$	6.9%		33.0%	
$SSIM_Z$	4.8%		28.7%	

Table 5.14: Subject F, biconic fit results (bare eye).

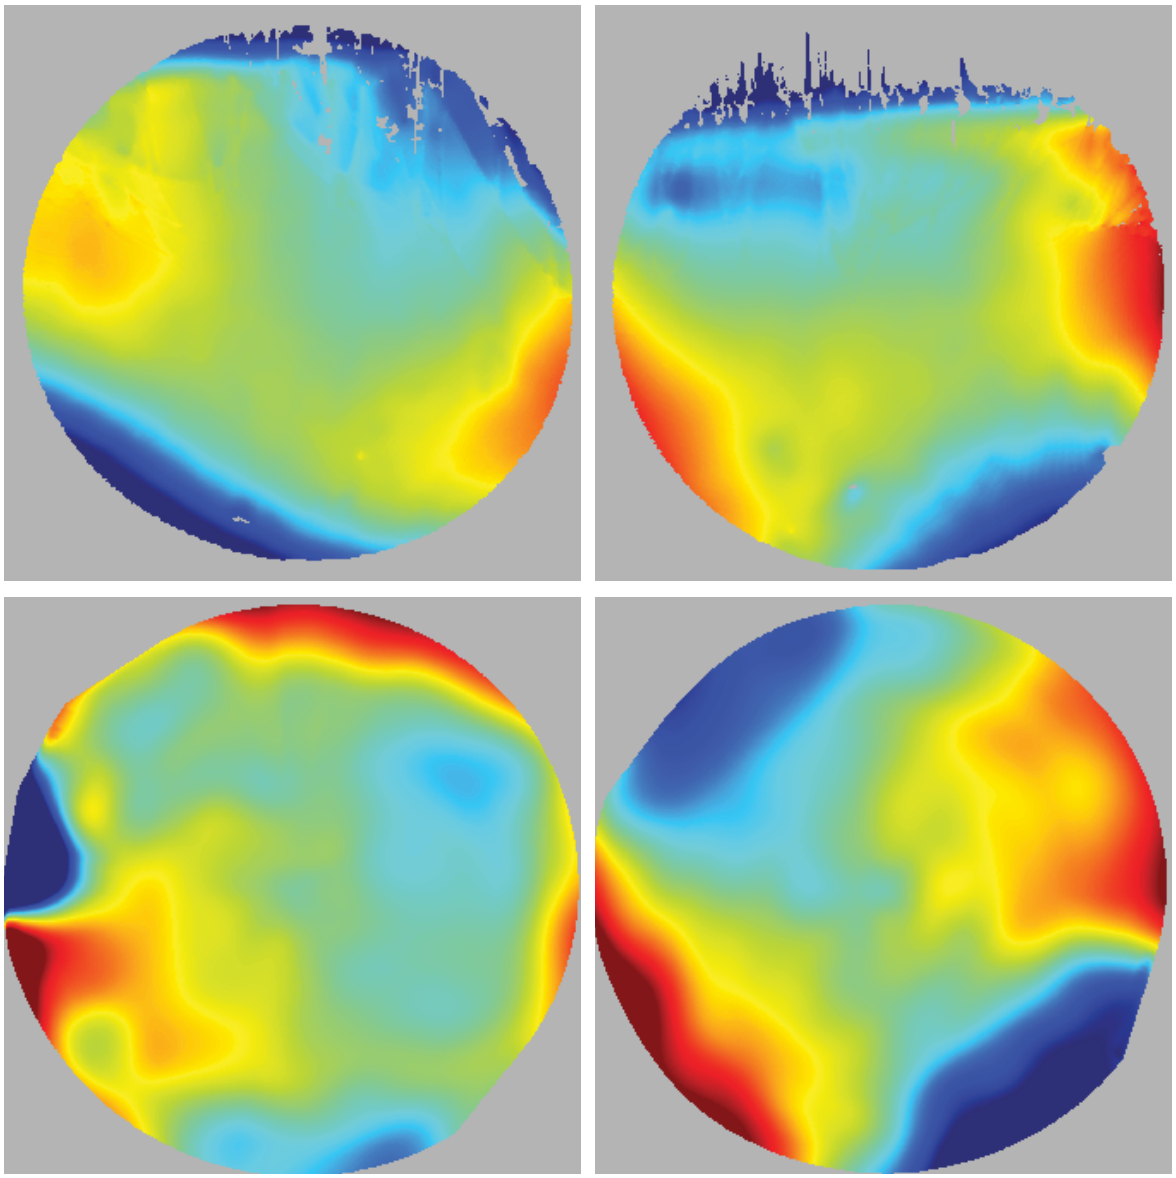

Figure 5.69: Subject F, conic removed corneal surface height (bare eye).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm. ±5.0 µm color height scale; red is away from eye.

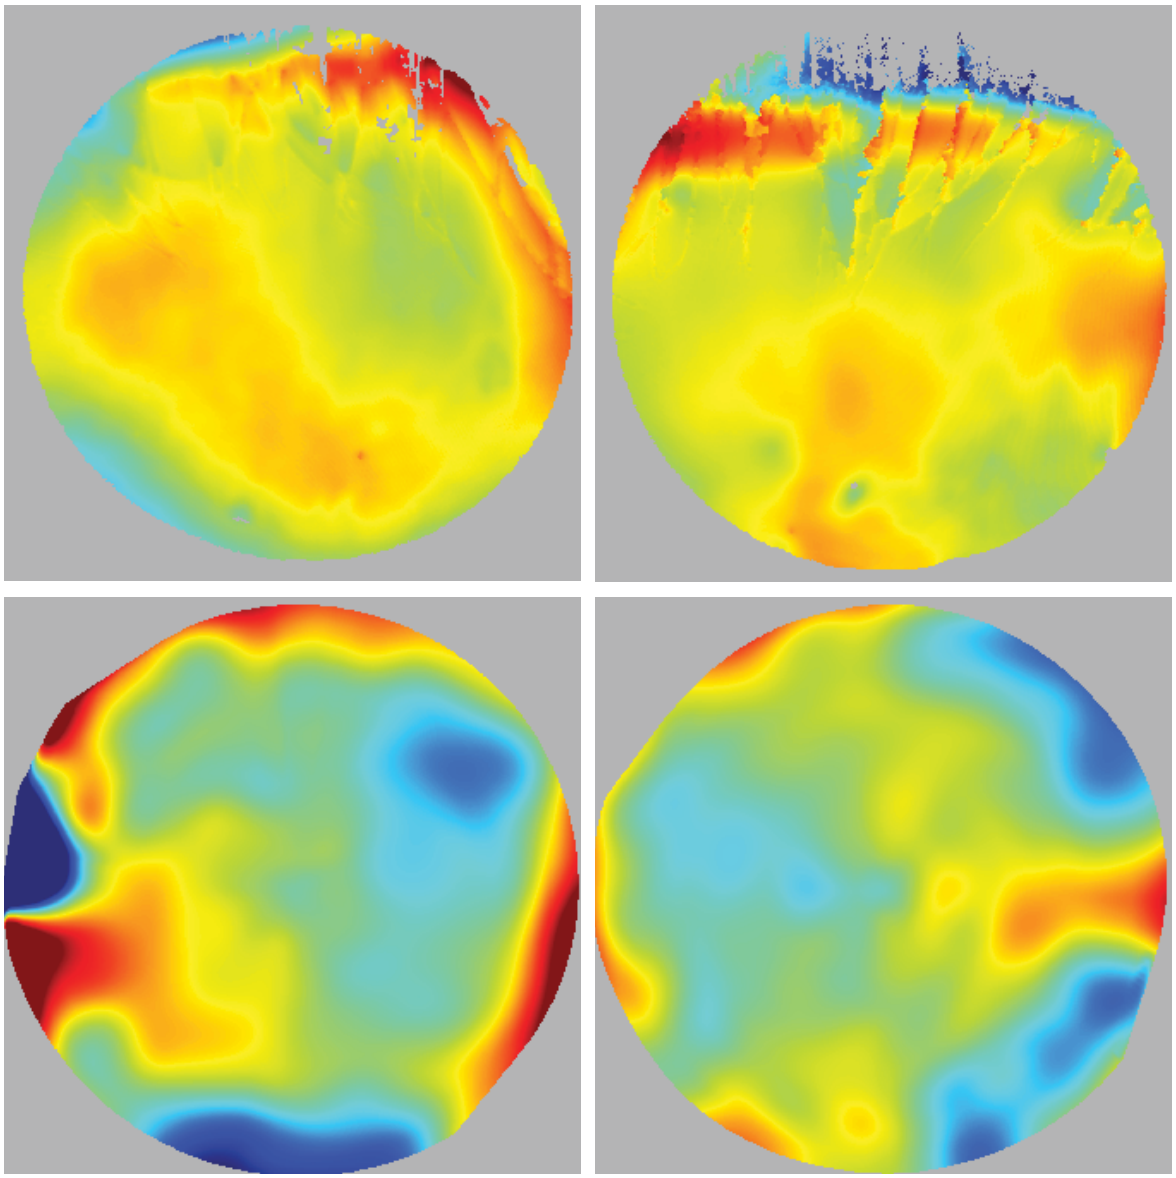

Figure 5.70: Subject F, biconic removed corneal surface height (bare eye).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm. ±4.0 µm color height scale; red is away from eye.

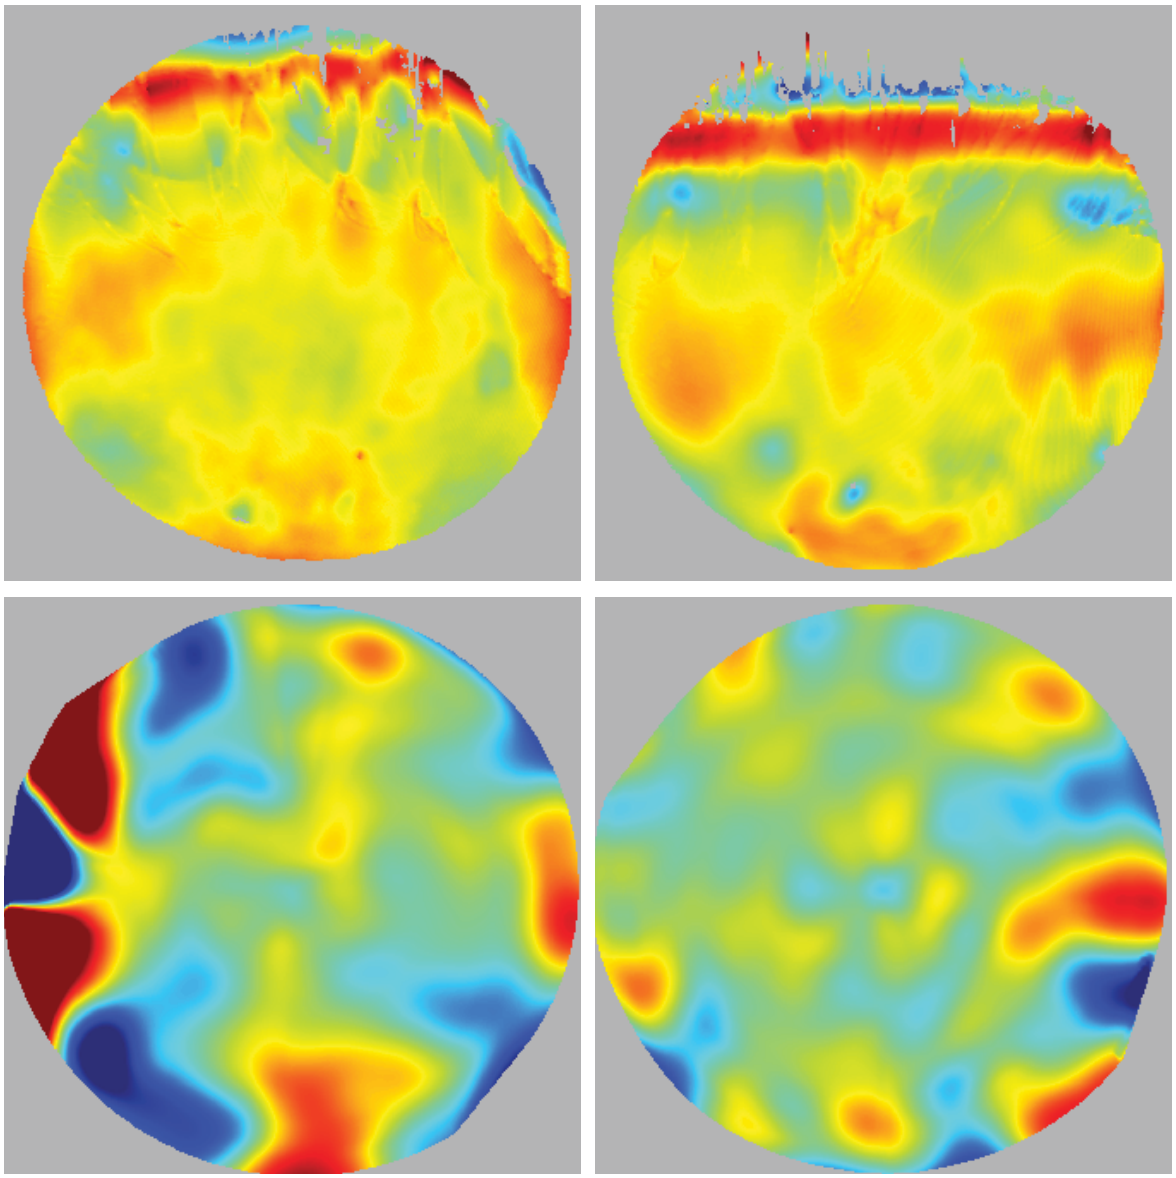

Figure 5.71: Subject F, Zernike removed corneal surface height (bare eye).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm. ±2.0 µm color height scale; red is away from eye. 15 Zernike terms removed from surface.

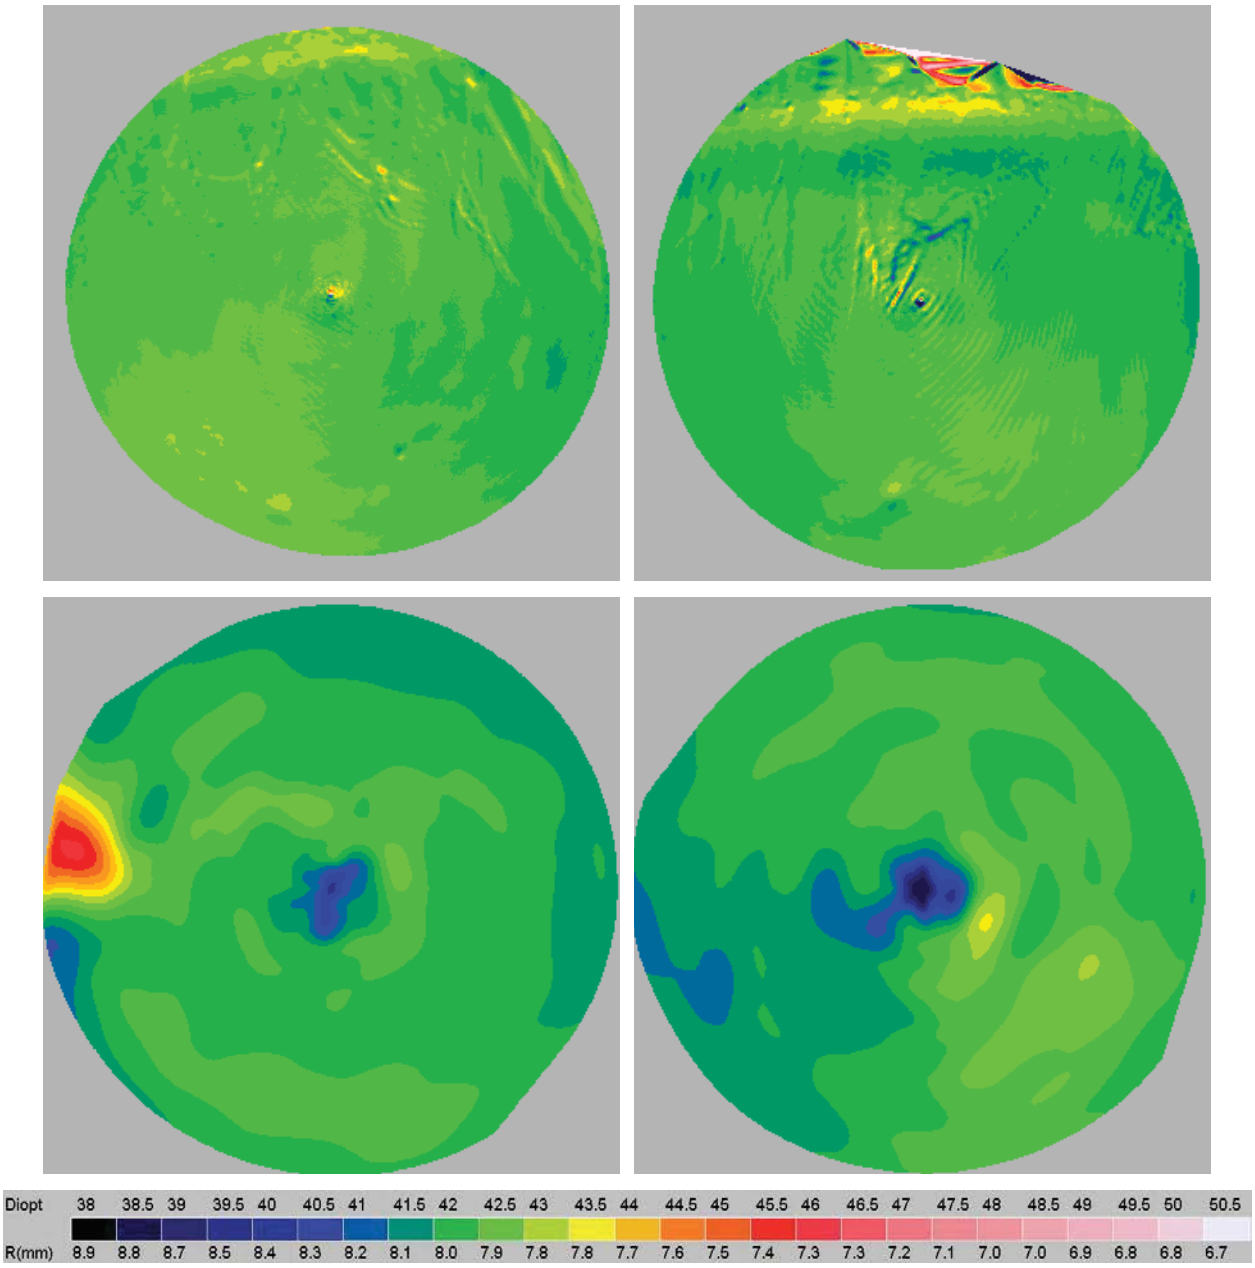

Figure 5.72: Subject F, axial power maps (bare eye).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm.

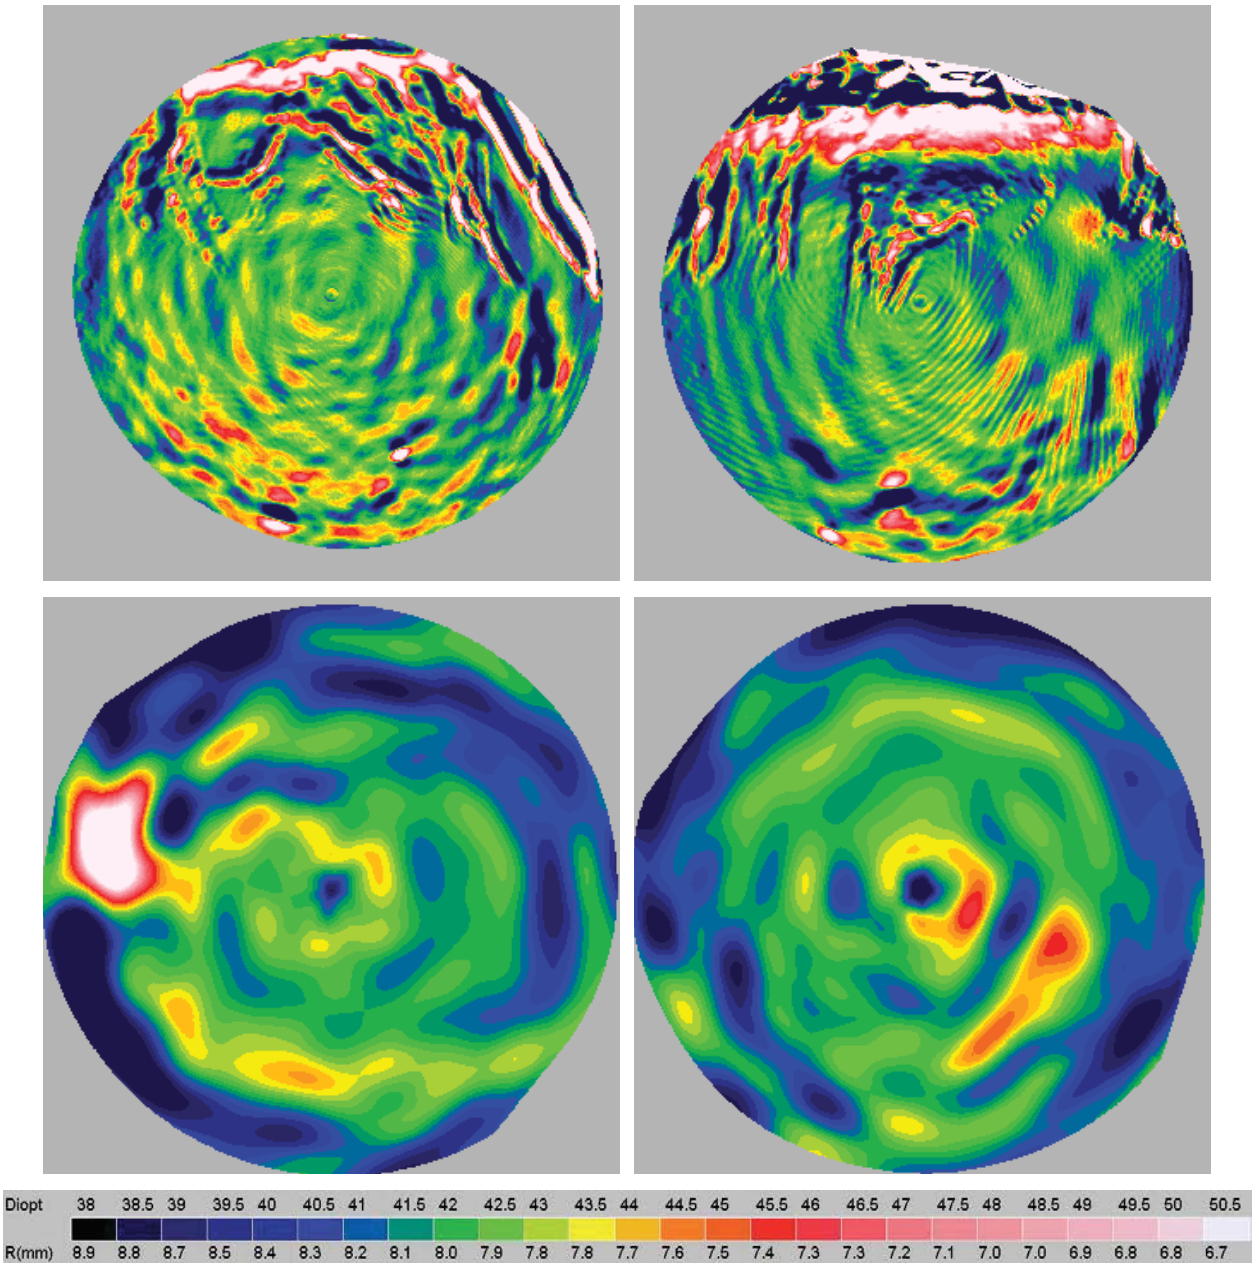

Figure 5.73: Subject F, instantaneous power maps (bare eye).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm.

5.3.6.2 CONTACT LENS

The contact lens disrupts the tear film enough to generate significant high-spatial frequency content that can be seen in all of the measurements. The tear film structure allows for a qualitative comparison of the mid to high spatial-frequency structure that is resolved by the two systems, despite the fact that the temporal dynamics does not allow for a proper comparison.

The biconic fit parameters in Table 5.15 shows the large change in radius from the bare eye measurements that is a result of the -5.5 D of correction from the contact lenses. The conic removed surfaces (Figure 5.74) shows that astigmatism is still present in the contact lens measurements. The magnitude of corneal astigmatism is the same between the bare eye and contact lens measurements, which seems to indicate that the contact lens conforms to the surface of the eye such that the underlying corneal topography maps to the front surface of the contact lens. The result is not surprising, as the thickness of a soft contact lens is on the order of 100 μm near the central 6 mm area that SPINALTOP measures. This example demonstrates that high resolution measurements of the corneal topography with and without a contact lens may provide additional information on the structural properties of contact lens materials.

The biconic removed surfaces (Figure 5.75) also show some correlation between the two systems, but now the mid to high-spatial frequency structure that results from the tear film interaction with the contact lens is only visible in the SPINALTOP measurements. The Zernike removed surfaces (Figure 5.76) further enhances the display of high spatial-frequency structure in the contact lens. At this point, the Keratron Piccolo is nearly useless when trying to characterize the tear film and contact lens interaction. A comparison of the spatial resolution capabilities of SPINALTOP is shown in Figure 5.77, which shows the central 1 x 1 mm area of the contact lens measurements. As demonstrated in Section 5.2.2.5, the spatial resolution of the Keratron Piccolo measurements makes it indeterminable if the structures are noise or tear film.

The -5.5 D of correction from the contact lenses will skew the power measurement as shown in Figure 5.78. The measurement on the left shows what the power map would like with the current standardized display scale. The measurements saturate the display at the low end and result in a constant color that provides no information. In order to bypass this issue, a temporary display scale is used for the contact lens power maps. The updated display scale is shown below the two power maps in Figure 5.78, which has been offset by +5 D to negate the power correction introduced by the contact lenses. This offset is temporarily applied to the axial power maps (Figure 5.79) and the instantaneous power maps (Figure 5.80).

	Left		Right	
	SPINALTOP	**Keratron**	**SPINALTOP**	**Keratron**
Rx [mm]	-9.091	-9.073	-9.260	-8.939
Ry [mm]	-9.024	-9.010	-9.164	-8.777
Kx	-0.192	-0.309	-0.057	0.217
Ky	0.006	-0.234	-0.052	0.087
θ [deg]	31.08	14.50	-23.74	-24.91
$SSIM_B$	60.5%		60.5%	
$SSIM_Z$	6.7%		6.1%	

Table 5.15: Subject F, biconic fit results (contact lens).

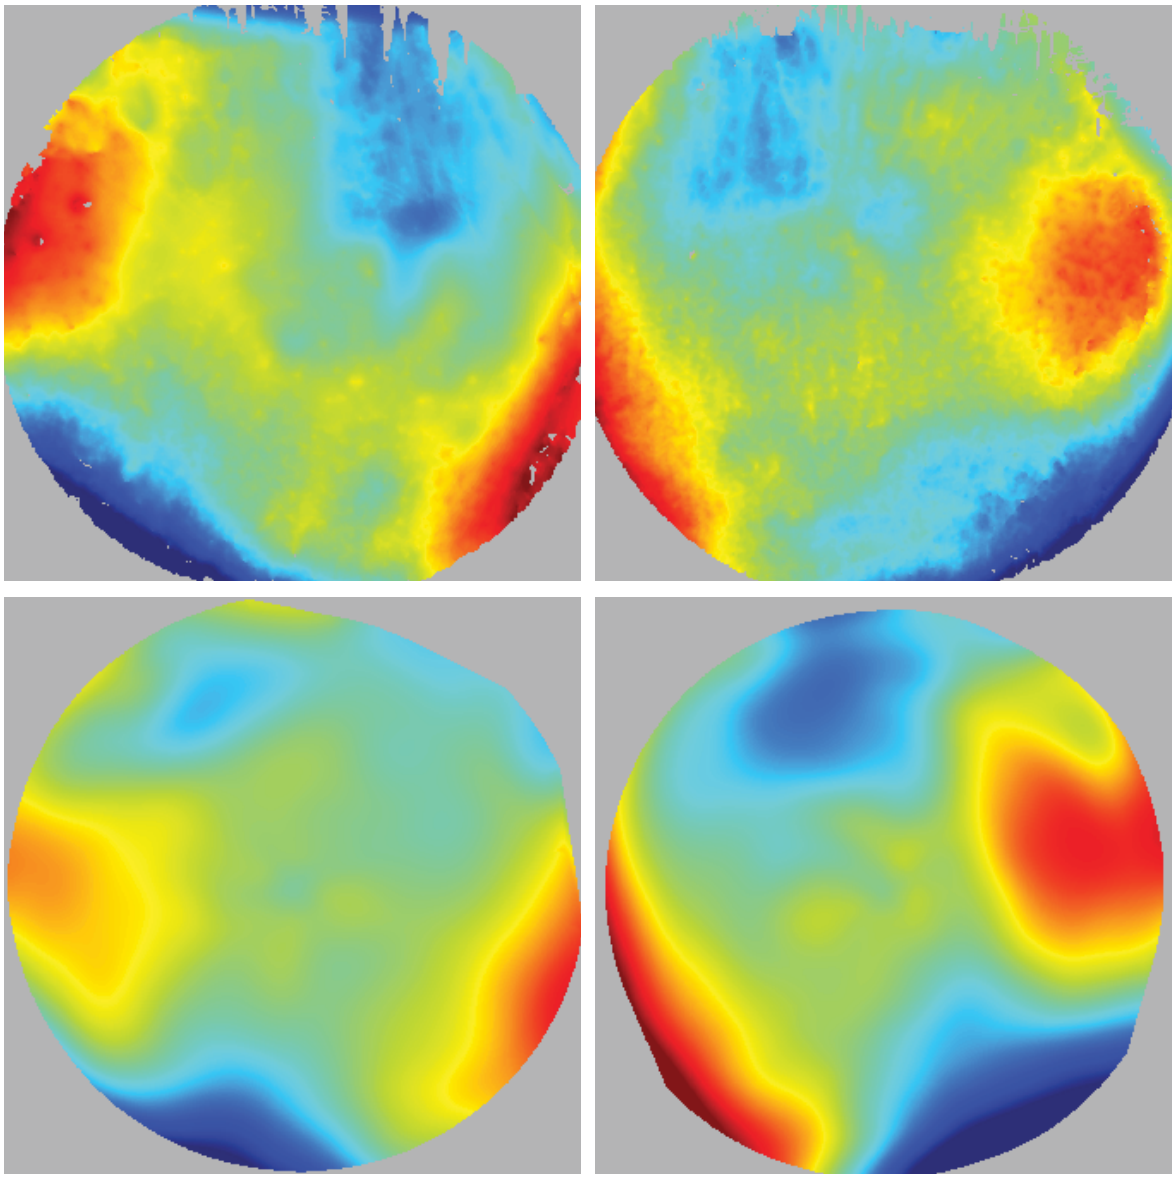

Figure 5.74: Subject F, conic removed corneal surface height (contact lens).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm. ±5.0 µm color height scale; red is away from eye.

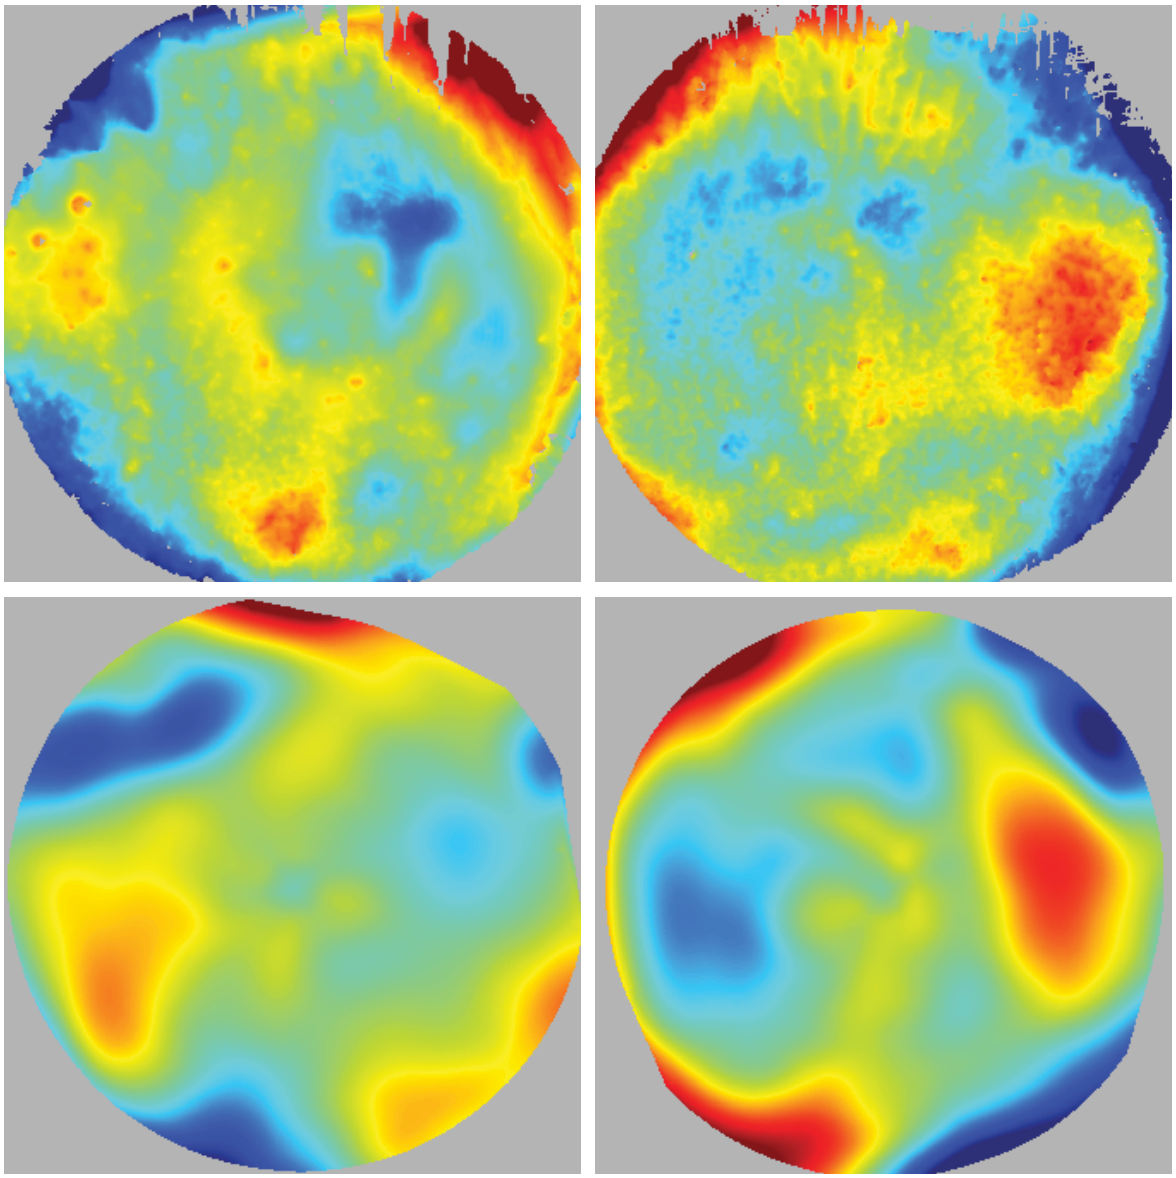

Figure 5.75: Subject F, biconic removed corneal surface height (contact lens).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm. ±2.25 µm color height scale; red is away from eye.

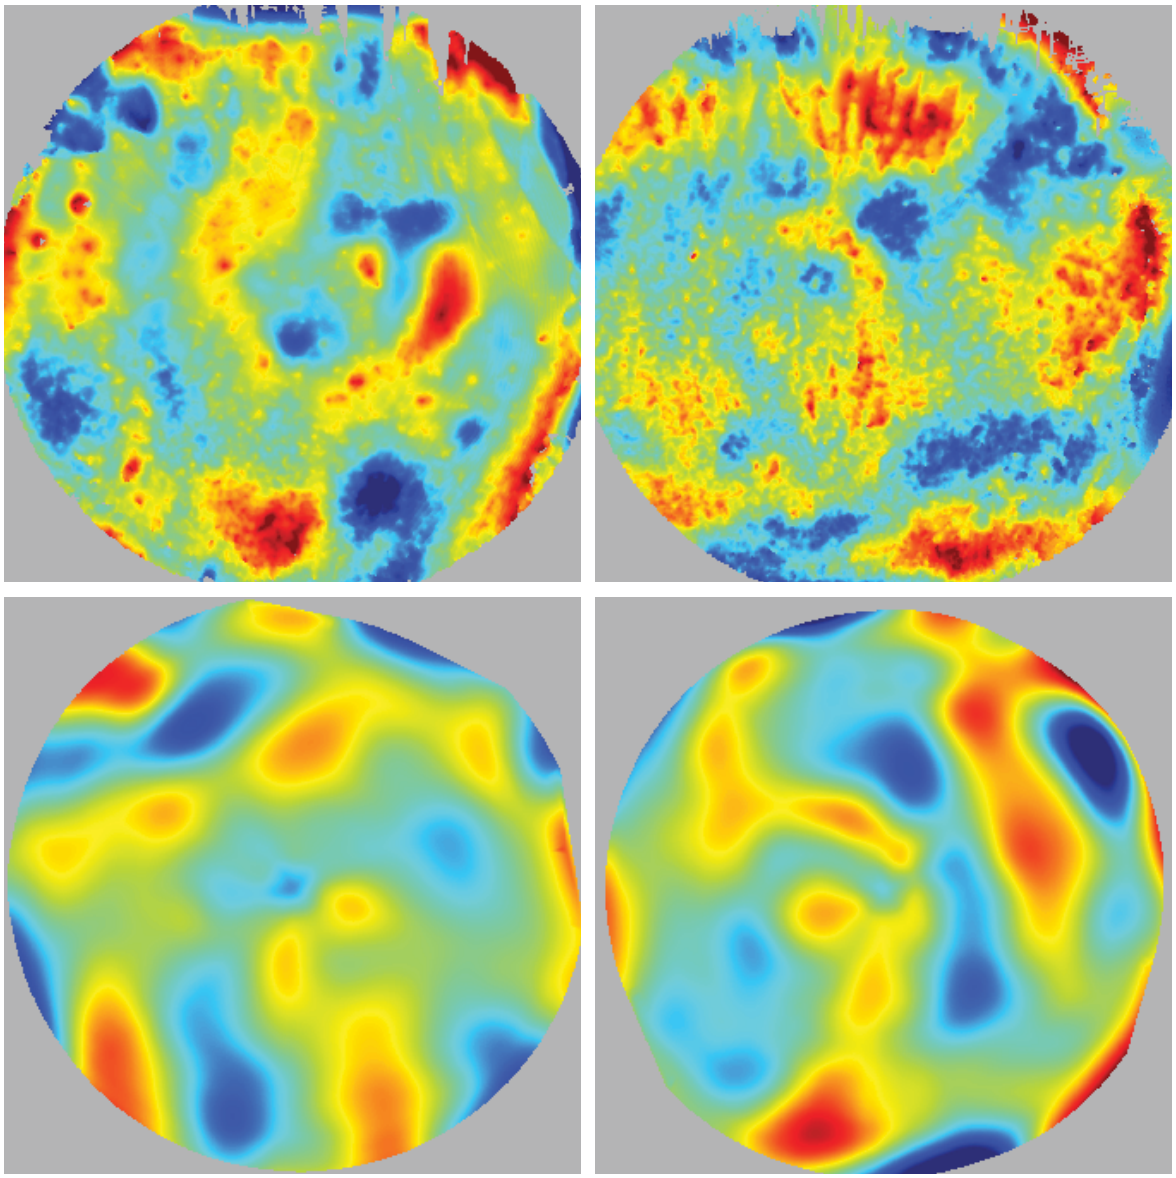

Figure 5.76: Subject F, Zernike removed corneal surface height (contact lens).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm. ±1.25 µm color height scale; red is away from eye. 15 Zernike terms removed from surface.

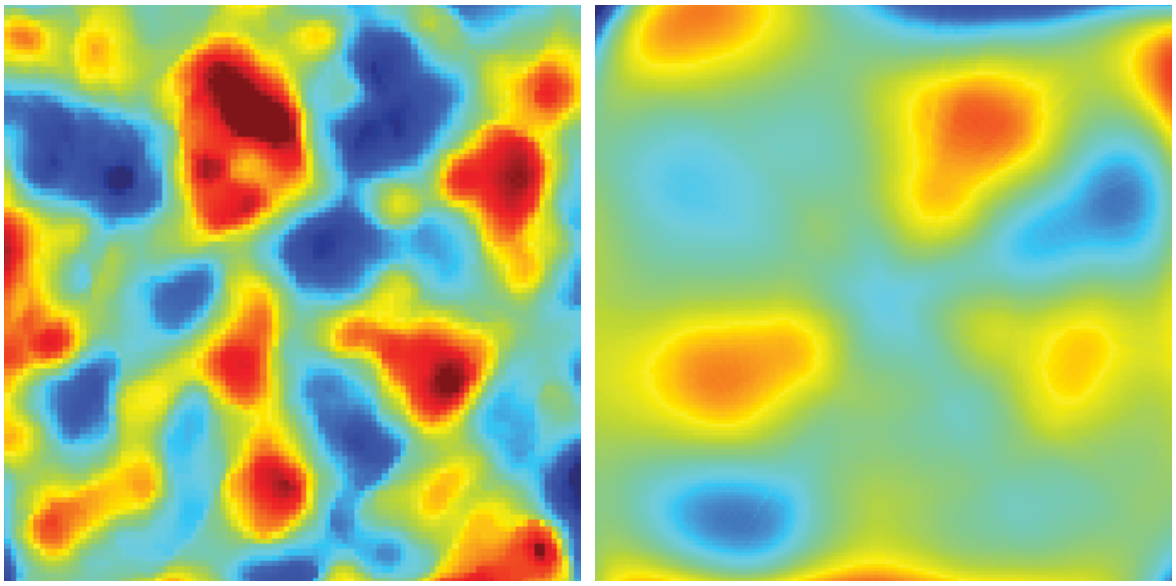

Figure 5.77: Subject F, central zone of contact lens measurement.
(Left) SPINALTOP measurement of the right eye. (Right) Keratron measurement of the right eye.
Display area is 1 x 1 mm. ± 200 nm color height scale; red is away from eye. 15 Zernike terms removed.

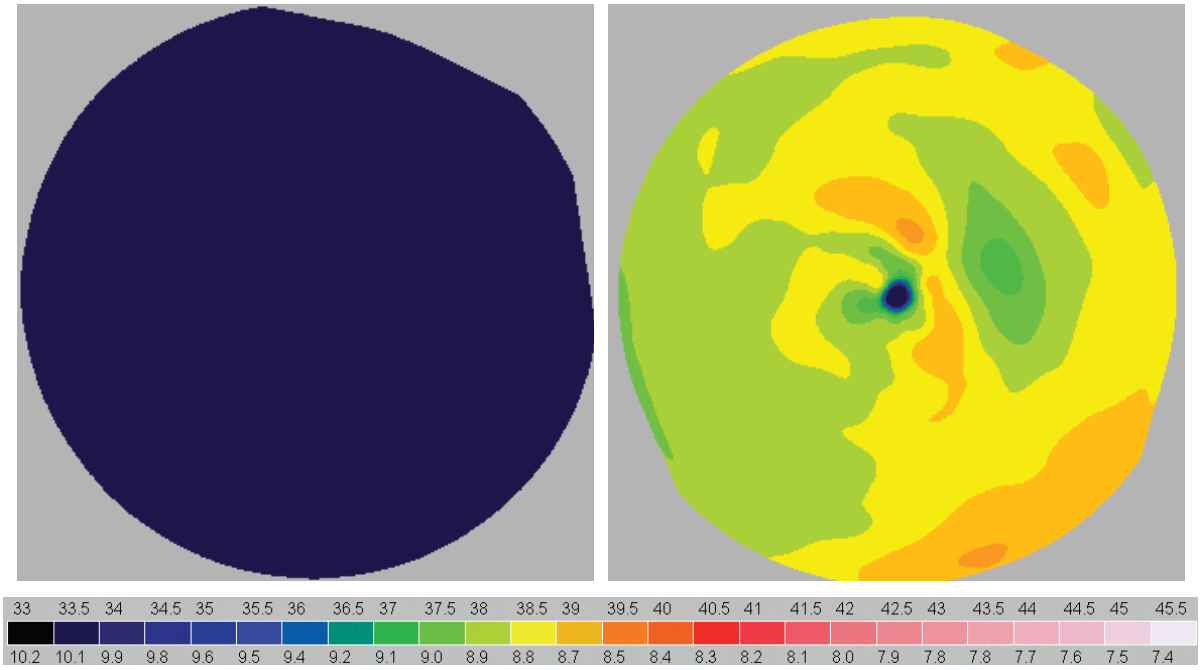

Figure 5.78: Power map display with offset (Keratron Piccolo).
(Left) Unaltered power map. (Right) Power map offset by +5D.
(Bottom) Updated power scale.

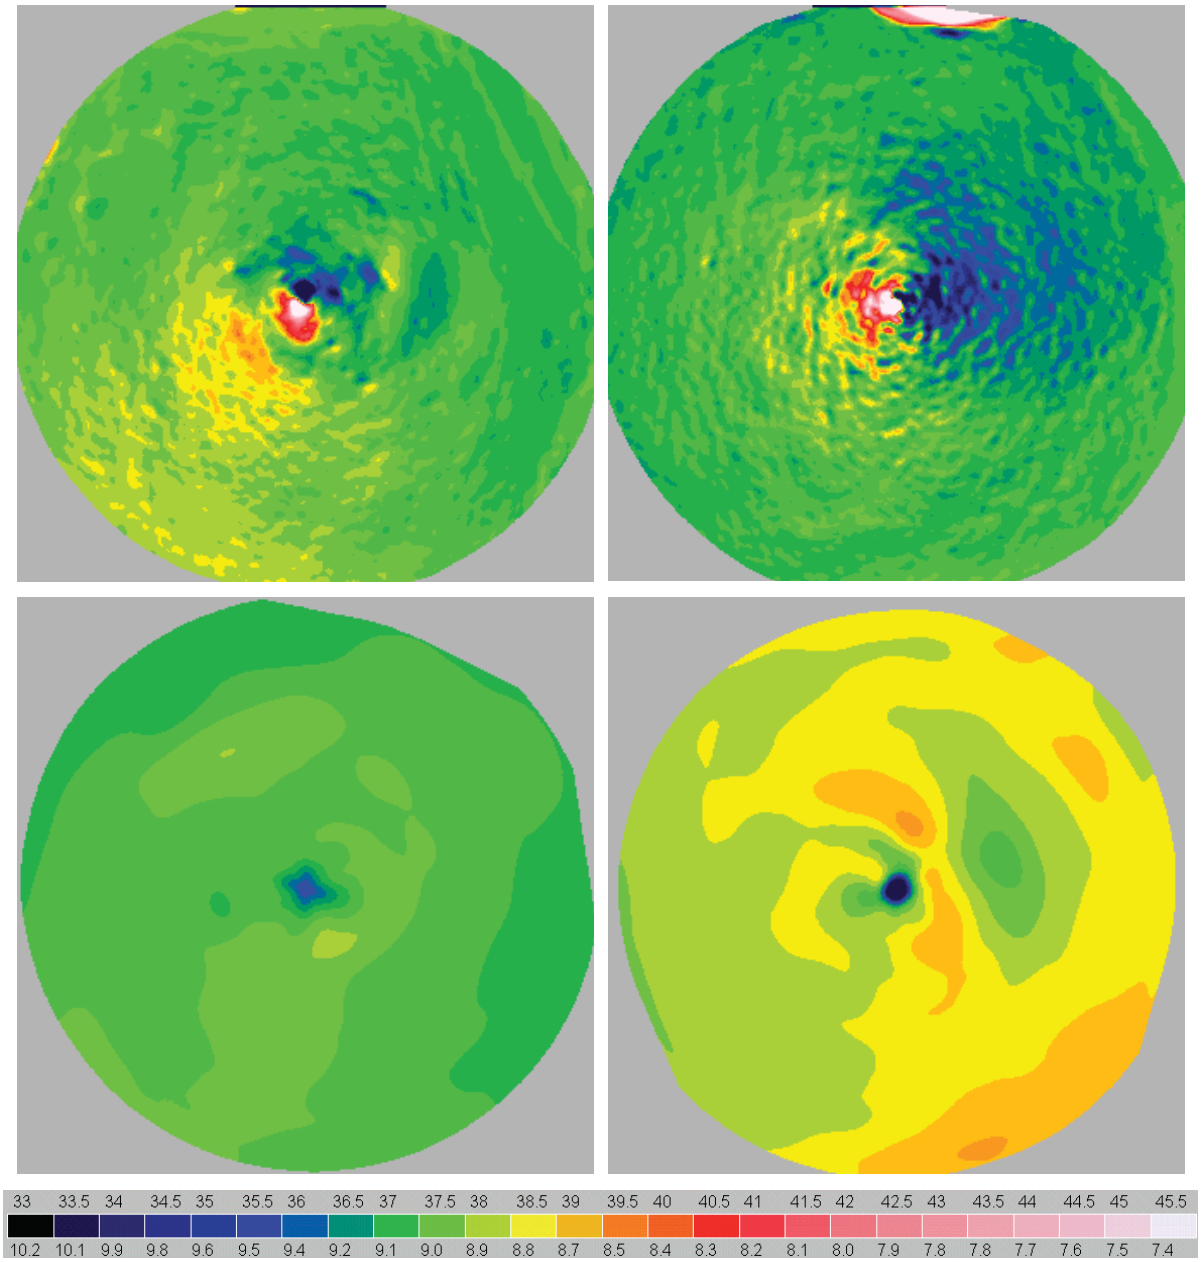

Figure 5.79: Subject F, axial power maps (contact lens).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm.

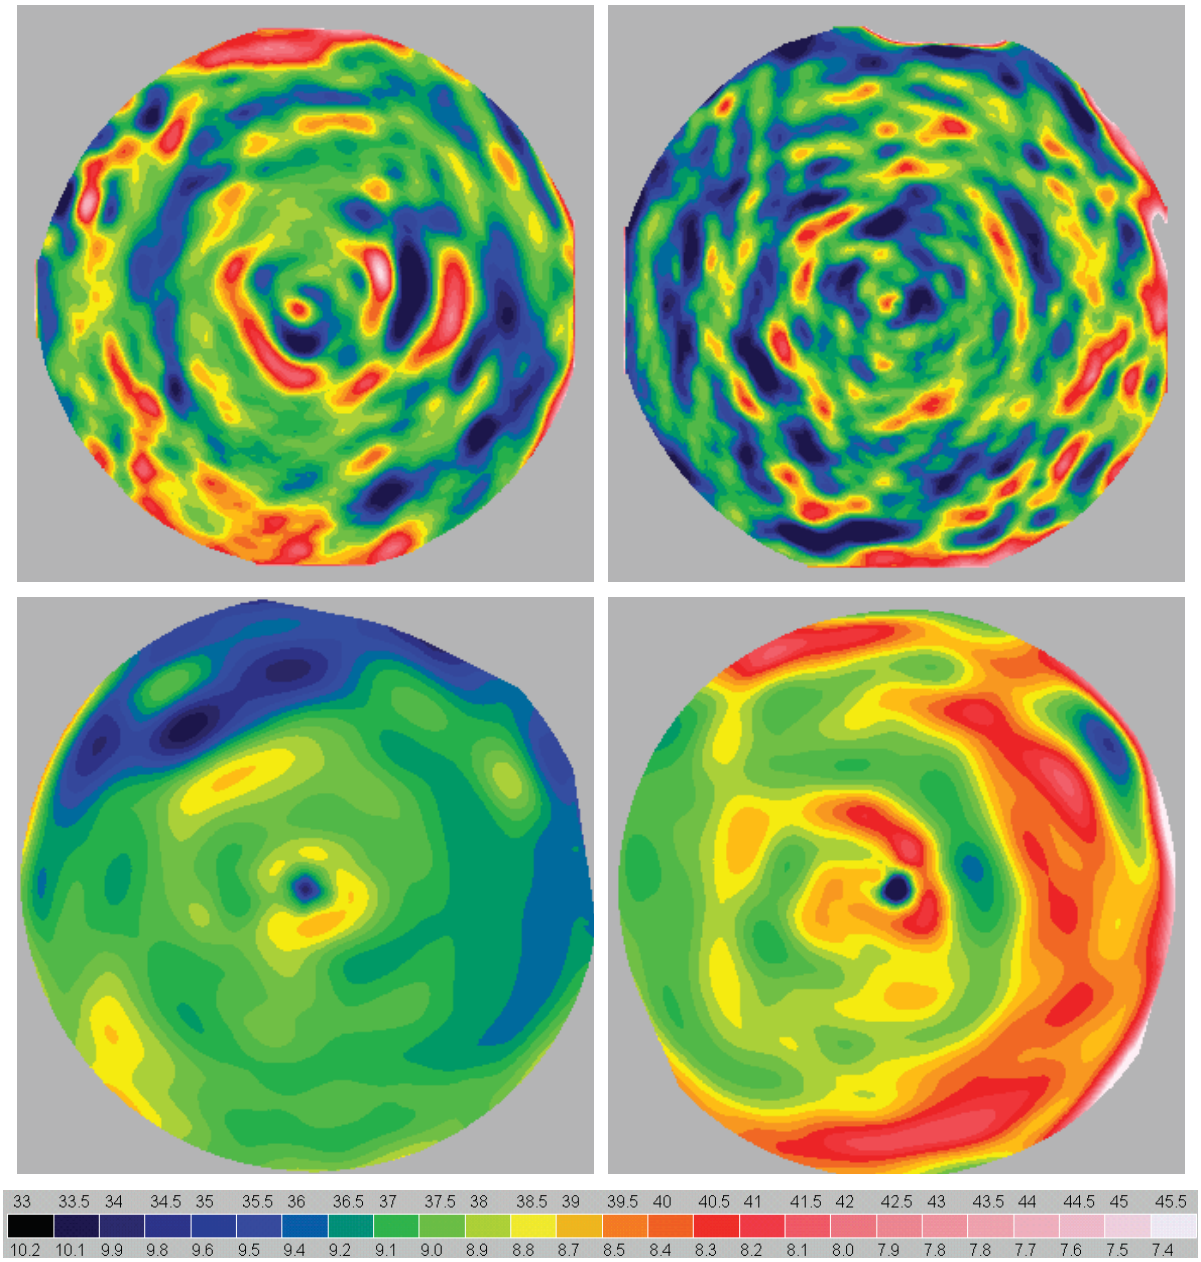

Figure 5.80: Subject F, instantaneous power maps (contact lens).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6.5 x 6.5 mm.

5.3.6.3 COMPARISON

Table 5.16 is a comparison of the estimated parameters from the bare eye and contact lens. As noted previously the prescription of the contact lenses were -5.5 Diopters in both eyes. The dioptric power was calculated from the radius measurement using the keratometric index of refraction $n_k = 1.3375$:

$$\Phi = \frac{337.5}{R\ [mm]} \qquad (5.9)$$

		Left		Right	
		SPINALTOP	Keratron	SPINALTOP	Keratron
Rx [mm]	Bare	-8.035	-8.110	-8.086	-8.144
	Contact	-9.091	-9.073	-9.260	-8.939
	Diff.	1.056	0.963	1.174	0.795
Ry [mm]	Bare	-7.979	-8.065	-8.026	-8.019
	Contact	-9.024	-9.010	-9.164	-8.777
	Diff.	1.045	0.945	1.138	0.758
Φ_x [m^{-1}]	Bare	-42.00	-41.62	-41.74	-41.44
	Contact	-37.12	-37.20	-36.45	-37.76
	Diff.	-4.88	-4.42	-5.29	-3.69
Φ_y [m^{-1}]	Bare	-42.30	-41.85	-42.05	-42.09
	Contact	-37.40	-37.46	-36.83	-38.45
	Diff.	-4.90	-4.39	-5.22	-3.63

Table 5.16: Subject F, Comparison of bare eye to contact lens.

A side-by-side comparison of the bare eye and contact lens, left eye measurements made on SPINALTOP are presented. Figure 5.81 shows the conic removed surfaces for both measurements. The corneal astigmatism that is present in the bare eye measurement (left) can be seen to print through onto the front surface of the contact lens (right). The measurement area for the contact lens is increased due to the larger radius of curvature of the surface. The biconic removed surface (Figure 5.82) shows that there is still some low to mid spatial-frequency structure that correlates between the bare eye and contact lens

measurement, but now high spatial-frequency structure is starting to dominate. The Zernike removed surface (Figure 5.83) now contains primarily the tear film structure, which shows the interaction between the tear film and contact lens.

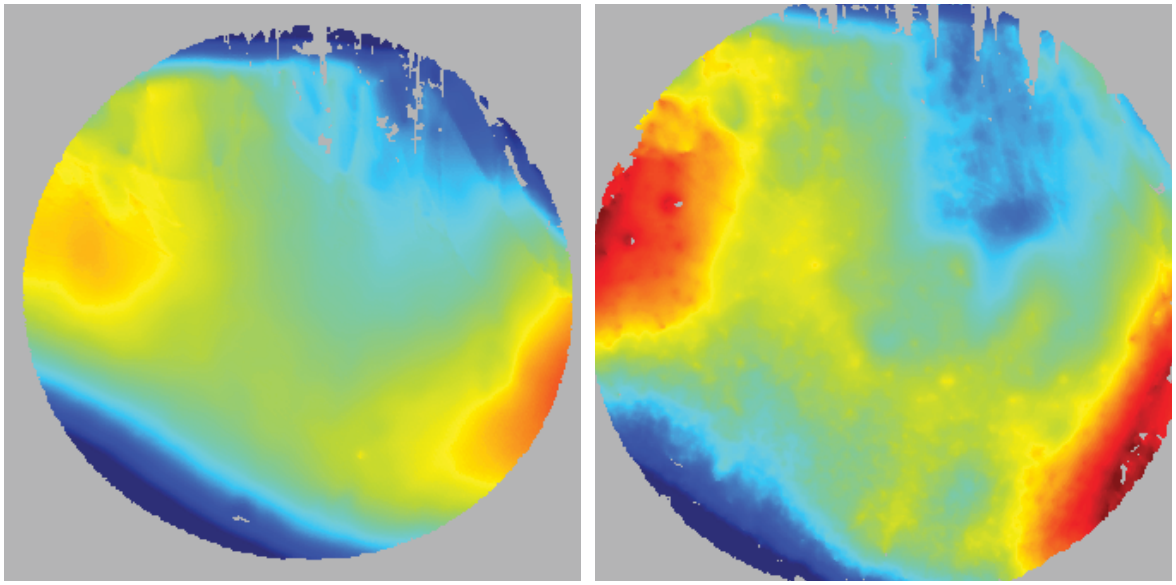

Figure 5.81: Subject F, left eye, conic removed surface comparison.
(Left) Bare eye. (Right) Contact lens. Measurements made on SPINALTOP. Display area is 6.5 x 6.5 mm. ±5.0 µm color height scale; red is away from eye.

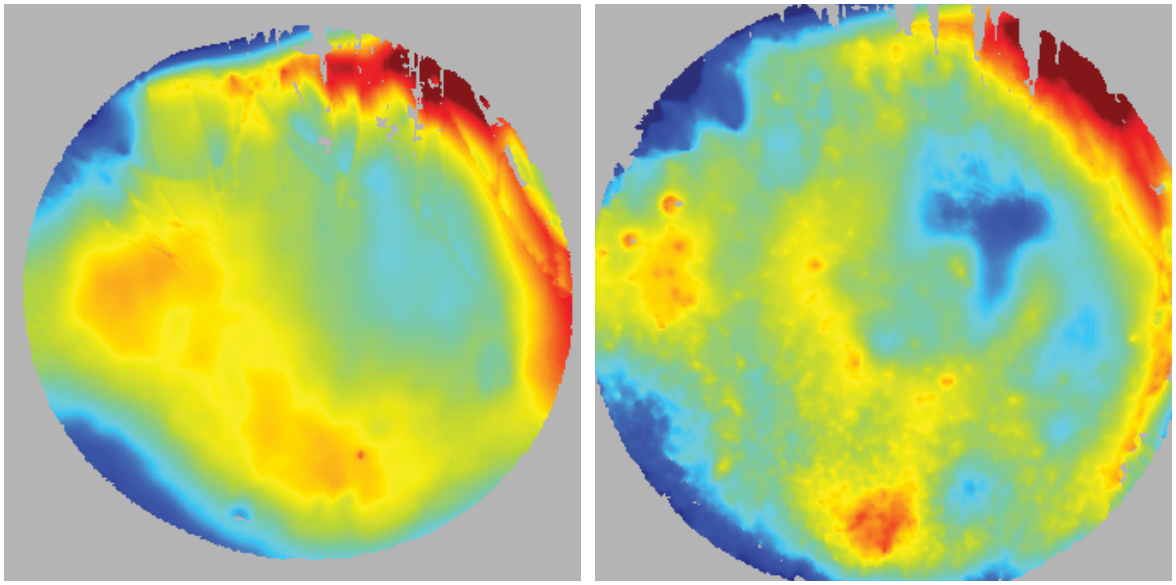

Figure 5.82: Subject F, left eye, biconic removed surface comparison.
(Left) Bare eye. (Right) Contact lens. Measurements made on SPINALTOP. Display area is 6.5 x 6.5 mm. ±2.25 µm color height scale; red is away from eye.

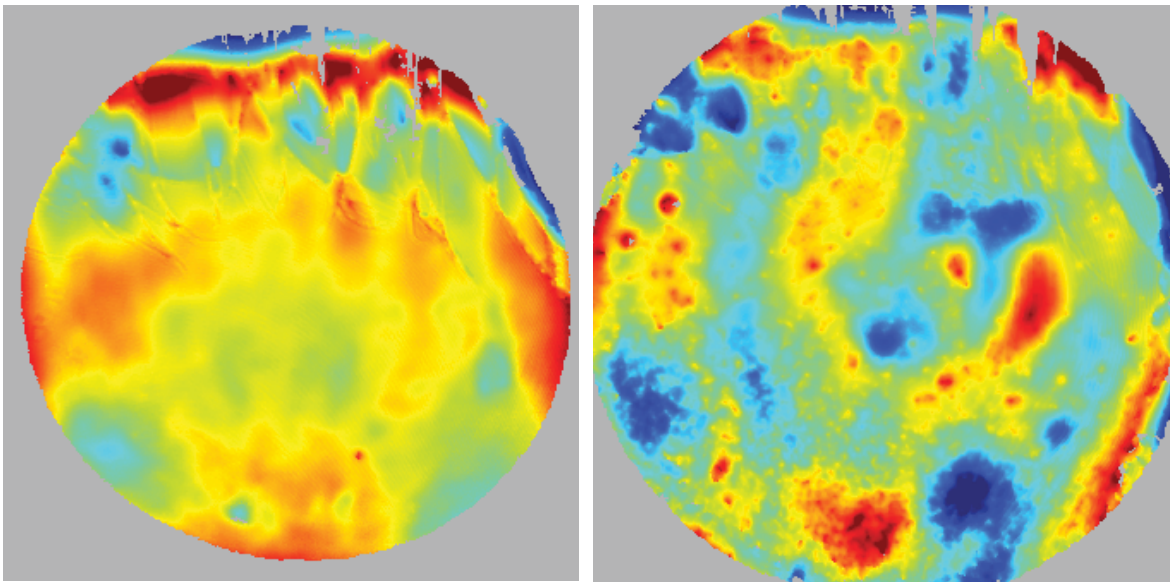

Figure 5.83: Subject F, left eye, Zernike removed surface comparison.
(Left) Bare eye. (Right) Contact lens. Measurements made on SPINALTOP. Display area is 6.5 x 6.5 mm. ±1.25 µm color height scale; red is away from eye.

5.3.7 SUBJECT G

The results presented here are from a subject that had a relatively high amount of corneal astigmatism (approximately 3 Diopters). The subject also wore contact lenses that corrected for the corneal astigmatism. The contact lenses were Biofinity toric by CooperVision. The left eye prescription was +0.5D / -2.25D x 10º and the right eye prescription was +0.25D / -1.75 x 170º.

5.3.7.1 BARE EYE

Biconic fit results for the bare eye measurement are presented in Table 5.17. The conic removed surfaces (Figure 5.84) shows the large amount of corneal astigmatism present with this subject. The general astigmatic shape appears slightly different between SPINALTOP and the Keratron Piccolo. However, it is known that the reconstruction method used by the Keratron Piccolo is susceptible to errors for non-rotationally symmetric surfaces (C. Roberts, 1994). Therefore, it is more likely that the SPINALTOP surfaces provide a more correct representation of the surface.

The biconic removed surfaces (Figure 5.85) shows better correlation between the two systems and are both dominated by low to mid spatial-frequency structure. The Zernike removed surfaces (Figure 5.86) displays what appears to be mostly tear film structure, but the residual pentafoil that appears in the SPINALTOP measurements is residual corneal topography. The orientation of the pentafoil appears aligned with astigmatic axis and therefore appears to be correlated to the astigmatic structure. This raises the question: Does the generalized biconic surface best represent the eye? It is possible that a better description of the corneal topography may include additional terms that account for the residual pentafoil. A study on a larger population may be able to determine whether this is the case or not.

The axial and instantaneous power maps are shown in Figure 5.87 and Figure 5.88, respectively. The large amount of corneal astigmatism dominates most of the measurement display.

	Left		Right	
	SPINALTOP	**Keratron**	**SPINALTOP**	**Keratron**
Rx [mm]	-8.011	-7.992	-8.078	-8.104
Ry [mm]	-7.471	-7.604	-7.622	-7.669
Kx	-0.143	-0.242	-0.329	-0.201
Ky	-0.118	0.146	0.102	0.042
θ [deg]	-4.79	-8.22	2.38	5.50
$SSIM_B$	76.7%		71.5%	
$SSIM_Z$	42.1%		3.7%	

Table 5.17: Subject G, biconic fit results (bare eye).

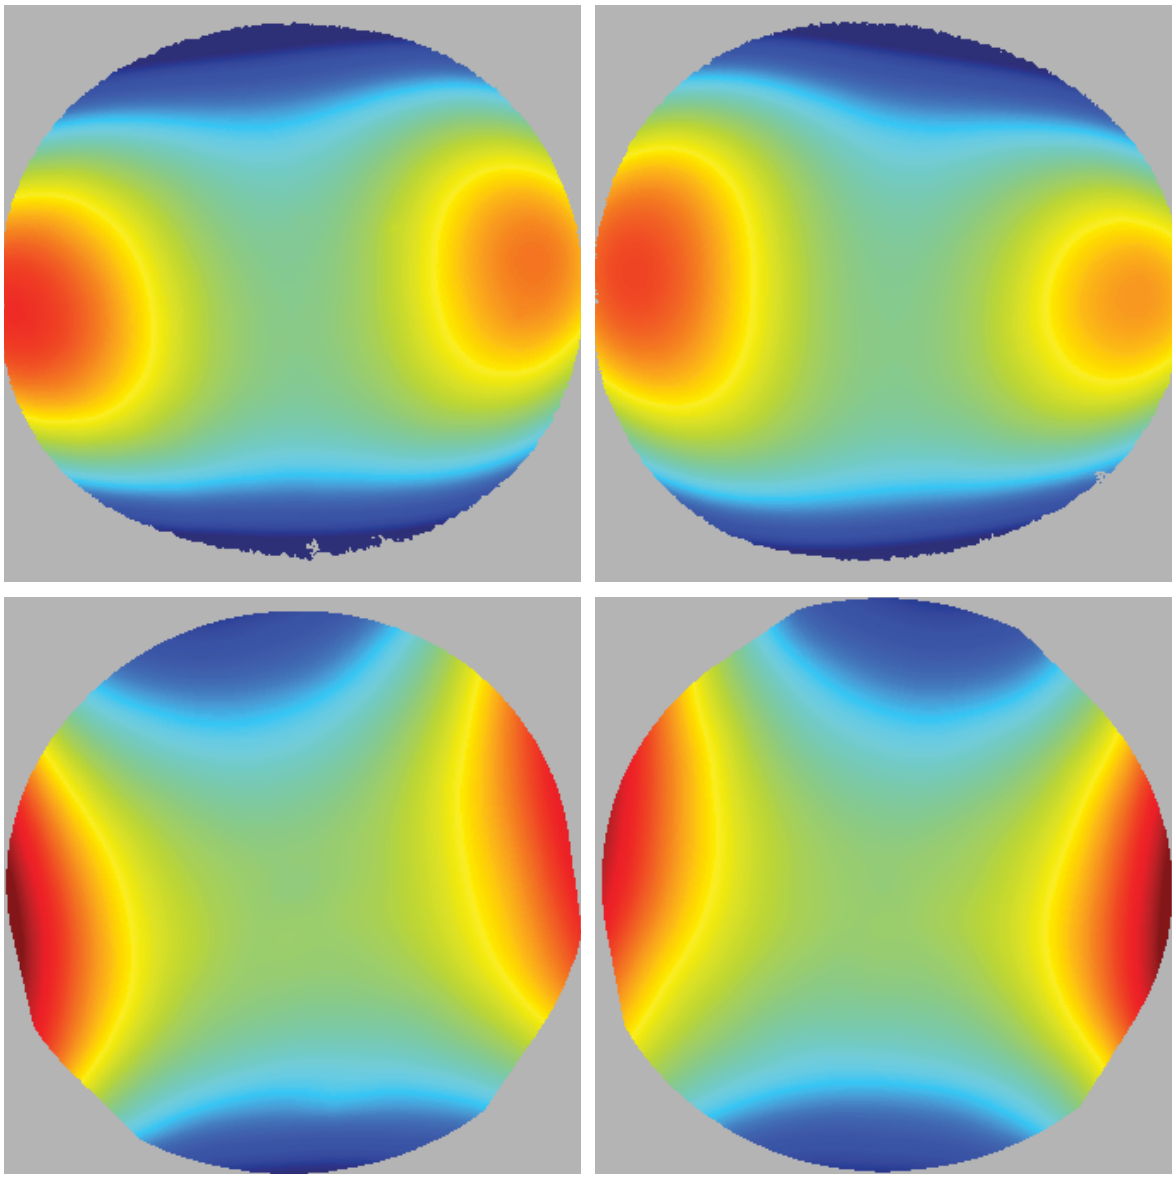

Figure 5.84: Subject G, conic removed corneal surface height (bare eye).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±20.0 μm color height scale; red is away from eye.

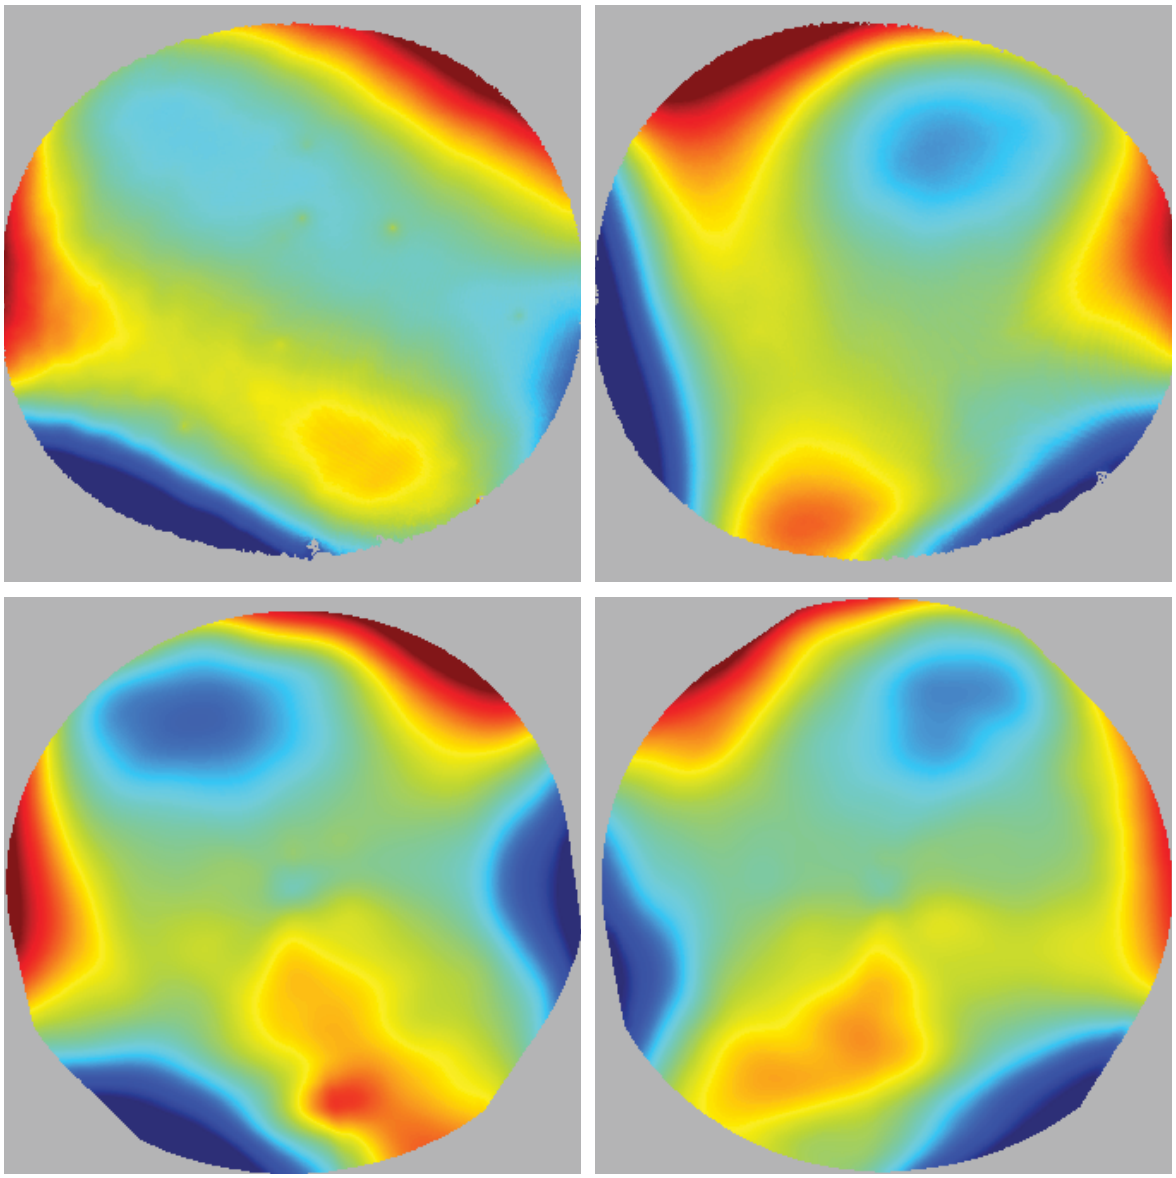

Figure 5.85: Subject G, biconic removed corneal surface height (bare eye).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±2.5 μm color height scale; red is away from eye.

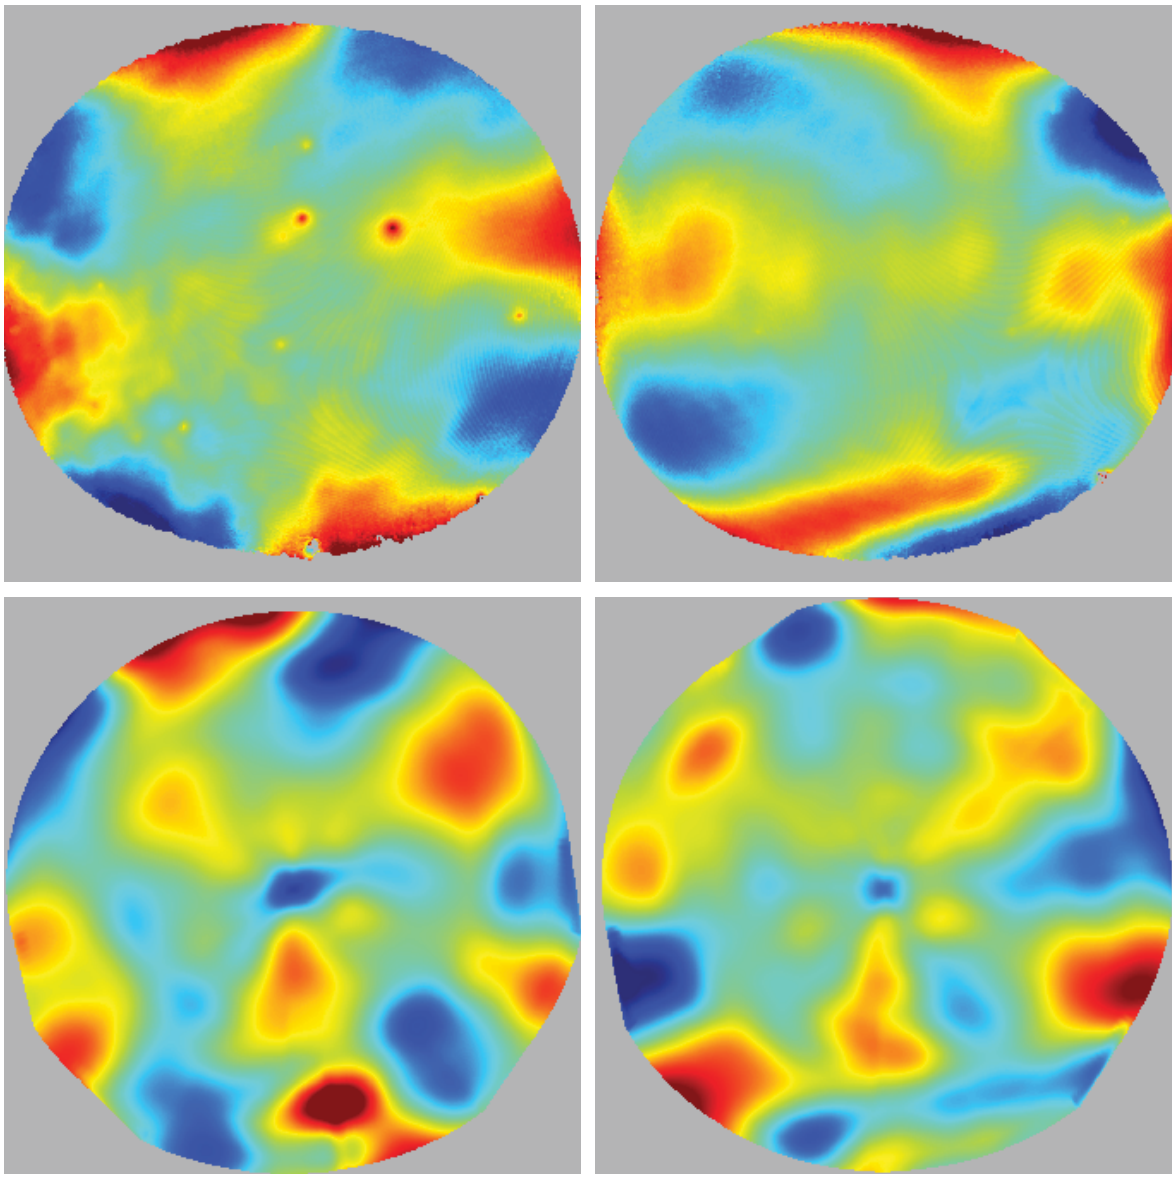

Figure 5.86: Subject G, Zernike removed corneal surface height (bare eye).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±0.6 µm color height scale; red is away from eye. 15 Zernike terms removed from surface.

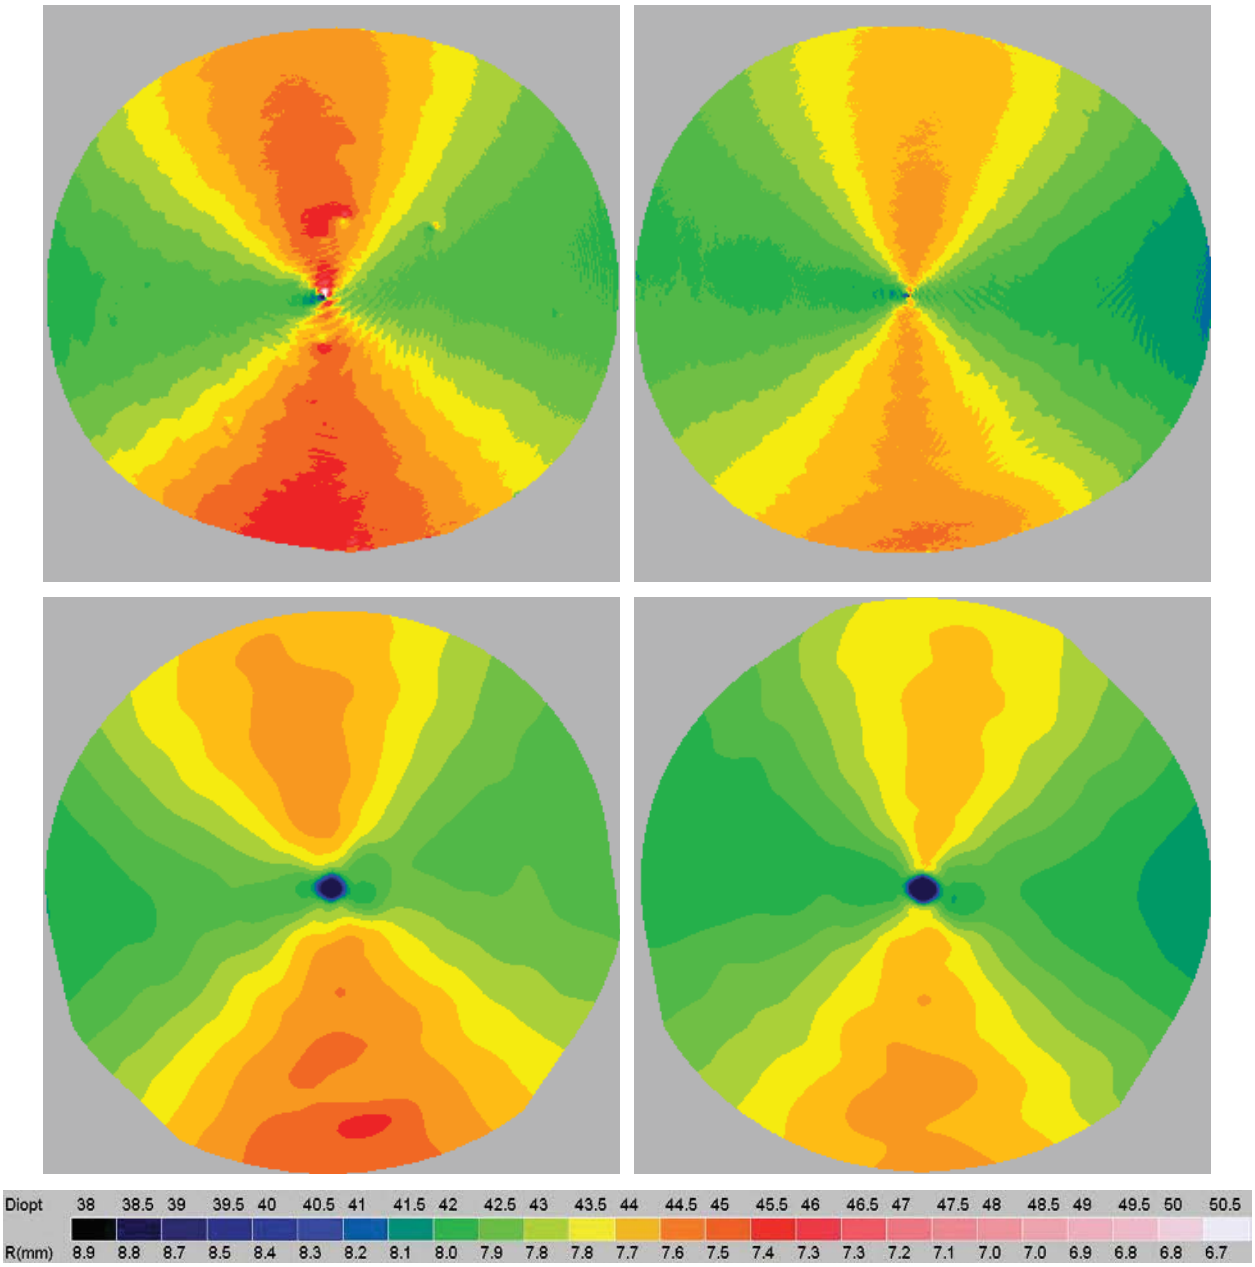

Figure 5.87: Subject G, axial power maps (bare eye).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm.

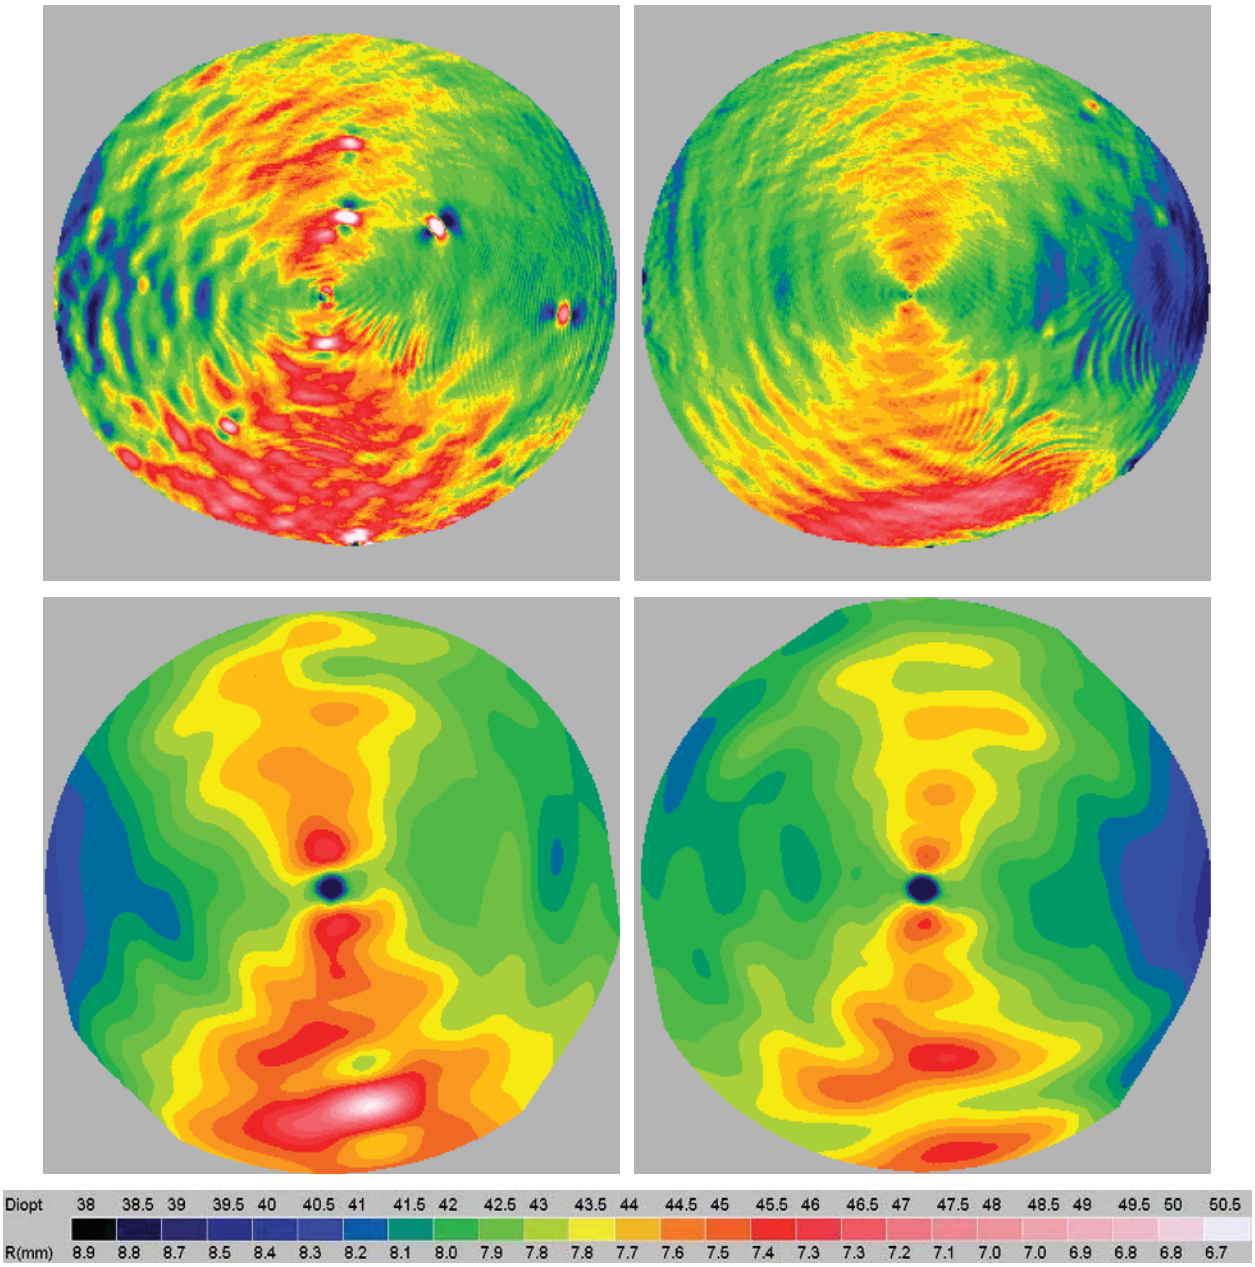

Figure 5.88: Subject G, instantaneous power maps (bare eye).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm

5.3.7.2 CONTACT LENS

The biconic fit results for Subject G wearing contact lenses is shown in Table 5.18. A large portion of the subject's corneal astigmatism has been adjusted for by the contact lens. The remaining astigmatism may be due to the fact that the subject has additional compensating astigmatism from other surfaces in the eye (e.g. the crystalline lens). The residual astigmatism can be seen in Figure 5.89 with the conic removed surface. The general astigmatic shape appears similar to the bare eye measurement, but the display scale has been reduced in half. A better comparison will be shown in the following section. The biconic removed surfaces (Figure 5.90) shows some of the residual topography, but again shows more of the mid to high spatial-frequency structure of the tear film. The Zernike removed surfaces (Figure 5.91) provide the best display for the higher spatial-frequency structure of the tear film. Some of the residual pentafoil from the bare eye measurement can still be seen in the contact lens measurement, which again raises the issue of how important is the contact lens structure in maintain the underlying corneal topography.

The axial power maps (Figure 5.92) shows that astigmatism is still present, but is no longer apparent in the instantaneous power maps (Figure 5.93).

	Left		Right	
	SPINALTOP	Keratron	SPINALTOP	Keratron
Rx [mm]	-7.899	-8.120	-8.076	-8.113
Ry [mm]	-7.723	-7.863	-7.846	-8.019
Kx	-0.283	-0.222	-0.231	-0.401
Ky	0.047	-0.002	0.006	-0.044
θ [deg]	-5.40	-5.06	1.90	7.82
$SSIM_B$	46.2%		45.2%	
$SSIM_Z$	4.4%		6.9%	

Table 5.18: Subject G, biconic fit results (contact lens).

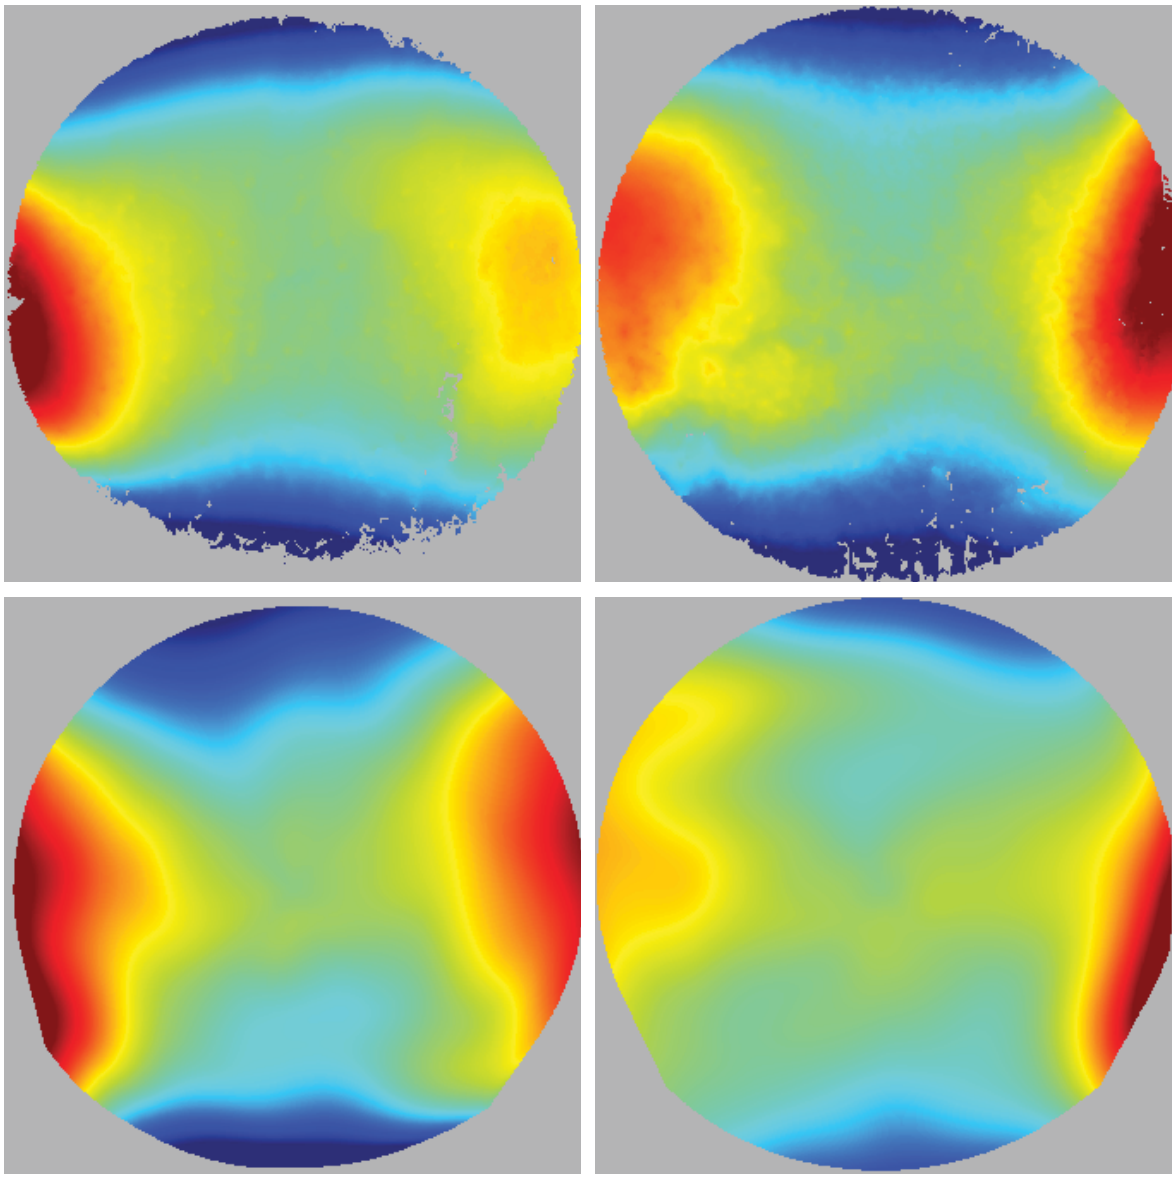

Figure 5.89: Subject G, conic removed corneal surface height (contact lens).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±10.0 μm color height scale; red is away from eye.

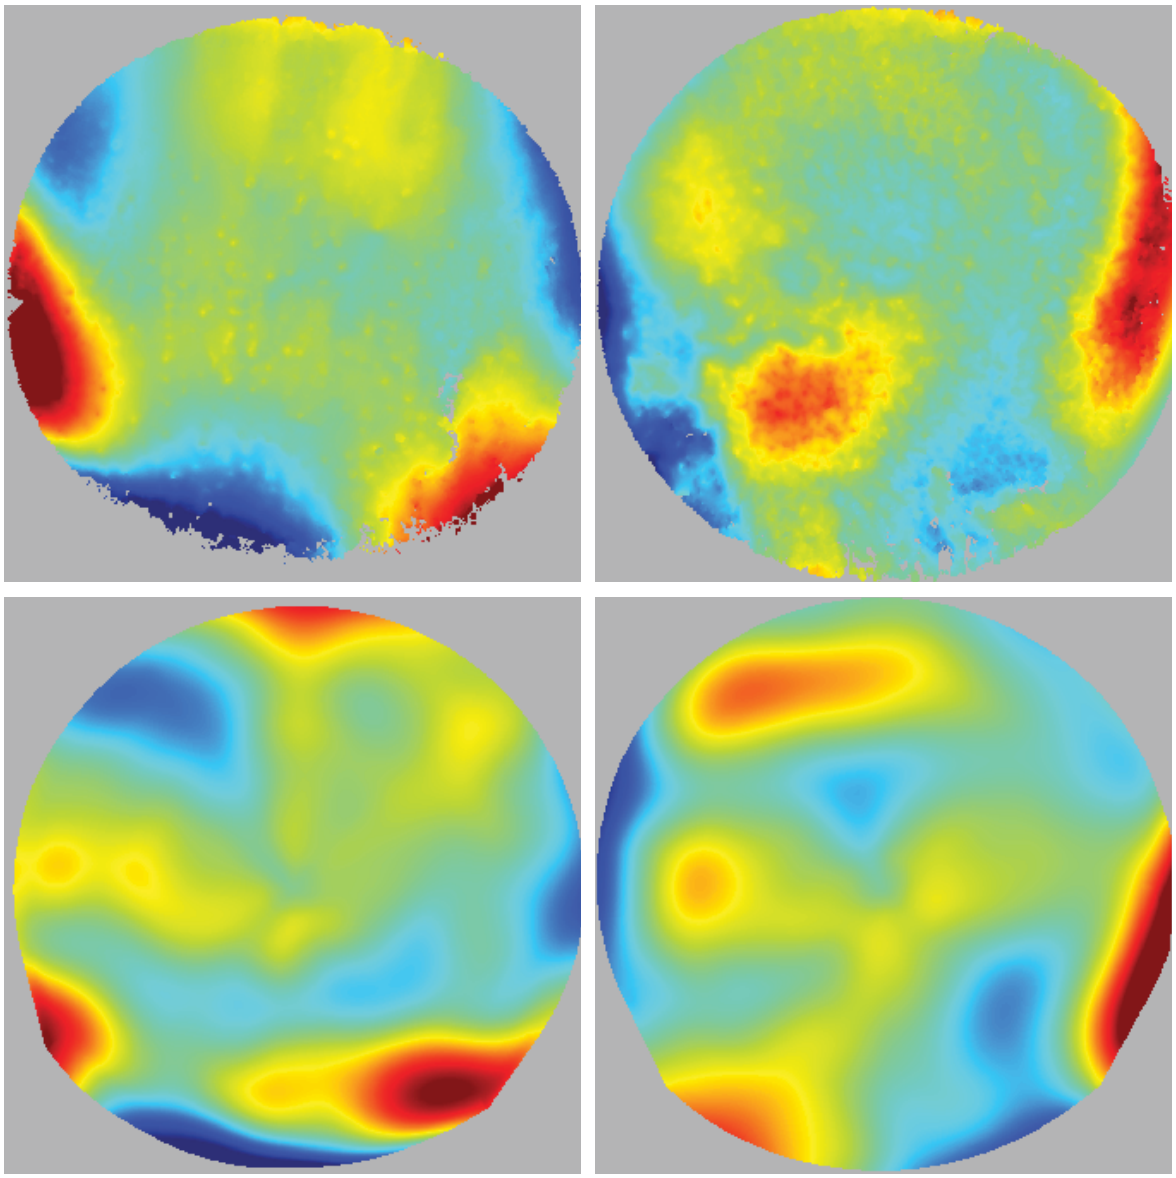

Figure 5.90: Subject G, biconic removed corneal surface height (contact lens).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±3.5 μm color height scale; red is away from eye.

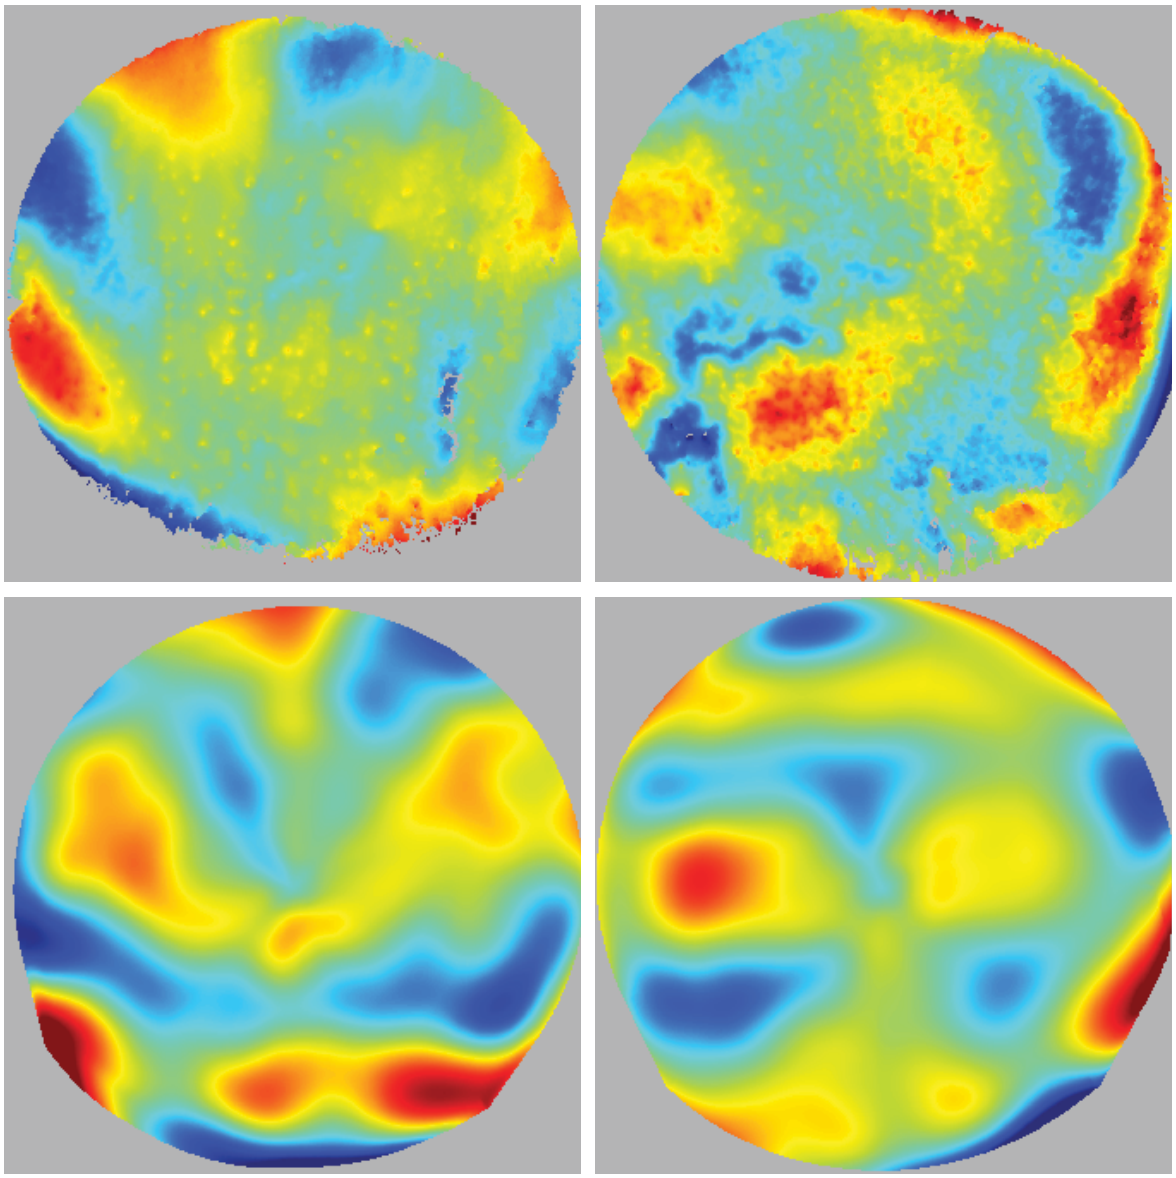

Figure 5.91: Subject G, Zernike removed corneal surface height (contact lens).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm. ±2 µm color height scale; red is away from eye. 15 Zernike terms removed from surface.

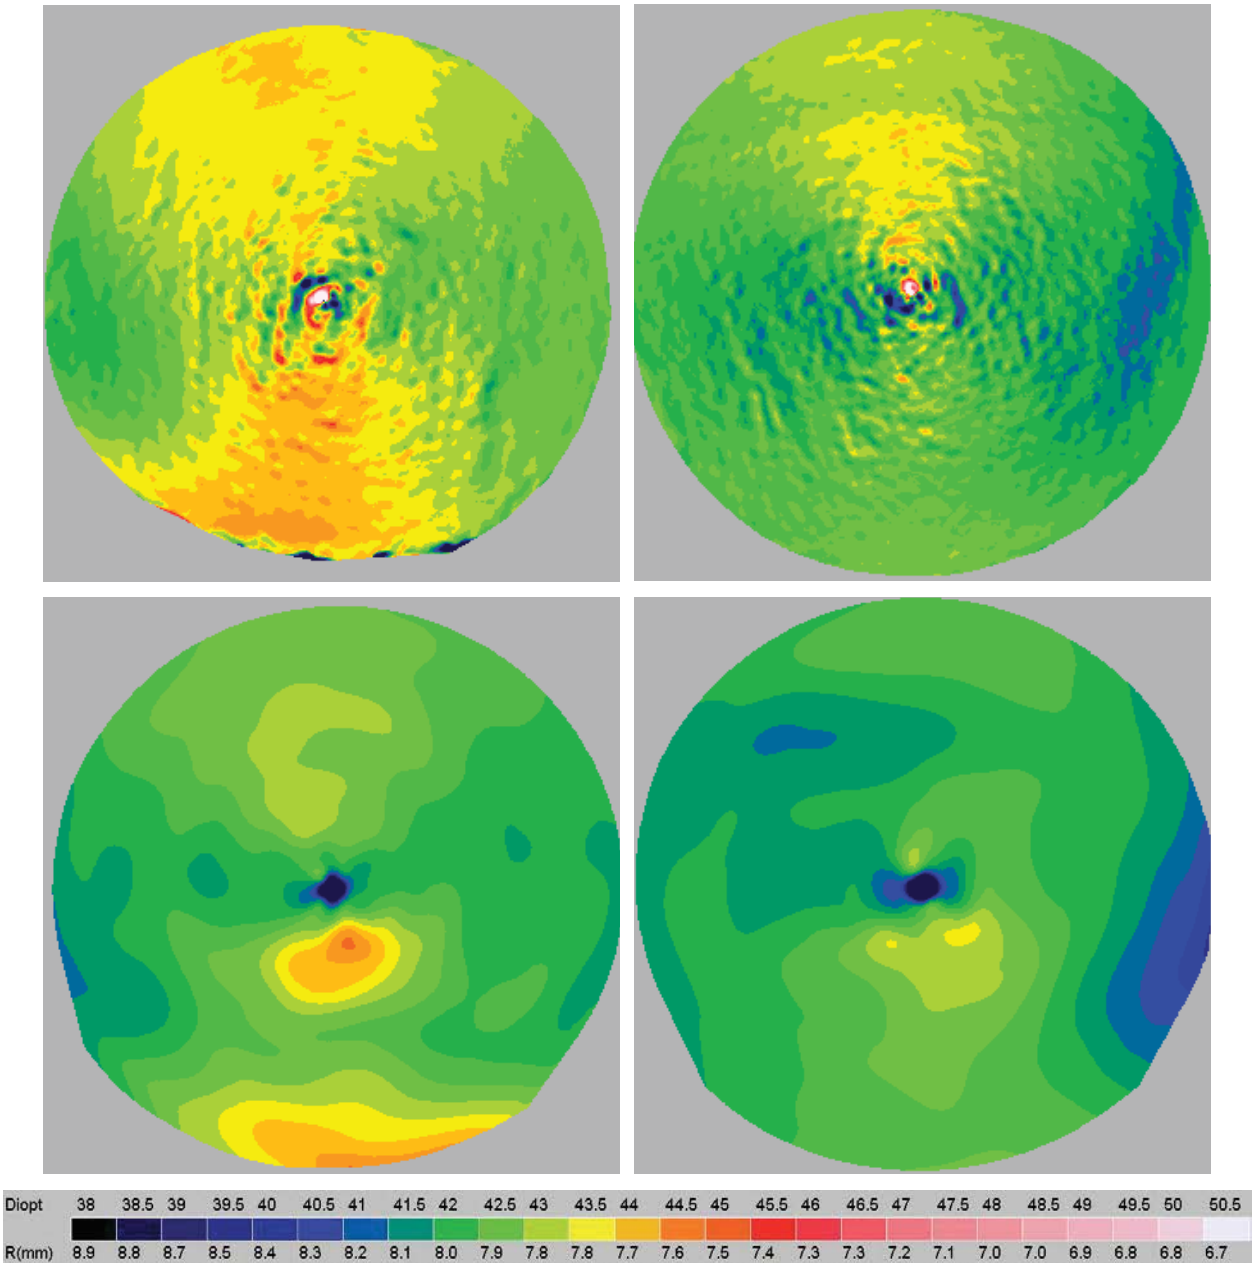

Figure 5.92: Subject G, axial power maps (contact lens).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm.

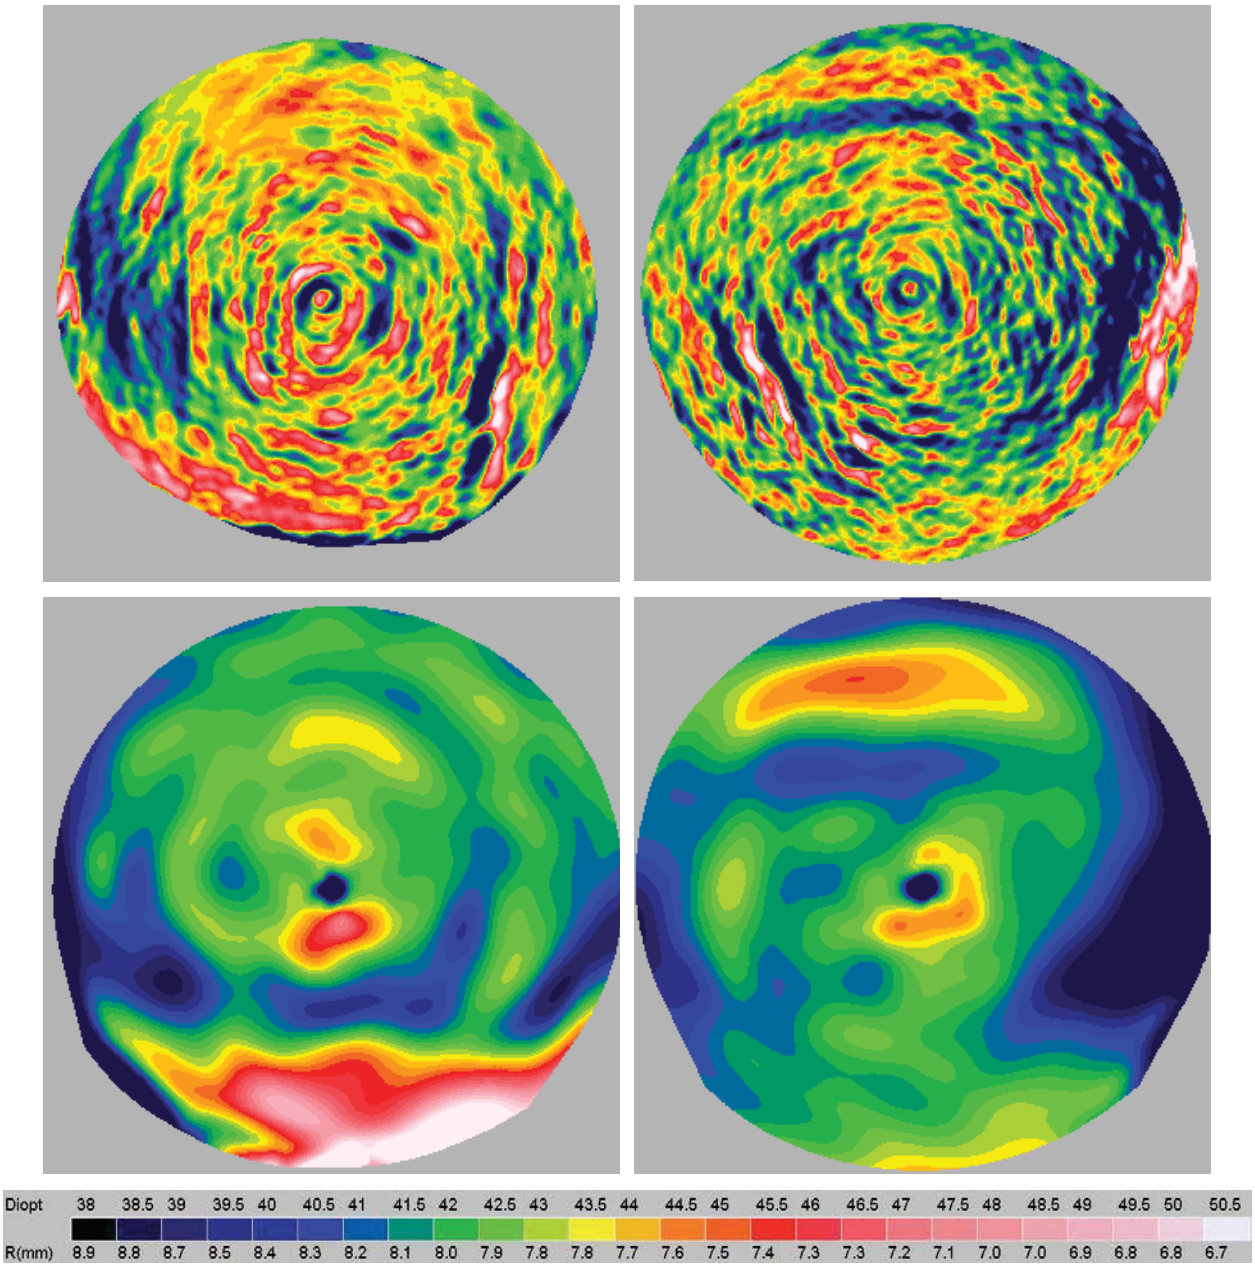

Figure 5.93: Subject G, instantaneous power maps (contact lens).
Left eye left column, right eye right column, SPINALTOP top row, Keratron bottom row.
Display area is 6 x 6 mm.

5.3.7.3 COMPARISON

Table 5.19 is a comparison of the measured parameters from the bare eye and contact lens. As previously noted, the left eye prescription was +0.5D / -2.25D x 10º and the right eye prescription was +0.25D / -1.75 x 170º. Power is calculated using Eq. 5.9.

		Left		Right	
		SPINALTOP	Keratron	SPINALTOP	Keratron
Rx [mm]	Bare	-8.011	-7.992	-8.078	-8.104
	Contact	-7.899	-8.120	-8.076	-8.113
	Diff.	-0.112	0.128	-0.002	0.009
Ry [mm]	Bare	-7.471	-7.604	-7.622	-7.669
	Contact	-7.723	-7.863	-7.846	-8.019
	Diff.	0.252	0.259	0.224	0.350
Φ_x [m^{-1}]	Bare	-42.13	-42.23	-41.78	-41.65
	Contact	-42.73	-41.56	-41.79	-41.60
	Diff.	0.60	-0.67	0.01	-0.05
Φ_y [m^{-1}]	Bare	-45.17	-44.38	-44.28	-44.01
	Contact	-43.70	-42.92	-43.02	-42.09
	Diff.	-1.47	-1.46	-1.26	-1.92

Table 5.19: Subject G, Comparison of bare eye to contact lens.

A side-by-side comparison of the bare eye and contact lens, left eye measurements made on SPINALTOP are presented. Figure 5.94 shows the conic removed surfaces for both measurements. This shows how the astigmatism has been reduced by the presence of the contact lens, but has only reduced it by half. The biconic removed surface comparison (Figure 5.95) shows that some of the residual low to mid spatial-frequency topography still maps through the contact lens. This comparison also starts to show the tear film structure that is being affected by the presence of the contact lens. The Zernike removed surface (Figure 5.96) further enhances the effect of the tear film structure, but also shows the residual corneal

topography that is in the form of pentafoil. It is interesting to note that the contact lens appears to enhance the amount of pentafoil that is present. It is possible that the contact lens material is magnifying this feature through by some material or structural phenomenon. For example, if there is a large enough shape difference between the cornea and posterior surface of the contact lens (i.e. the side that sits on the cornea), then it is possible that the contact lens is buckling or distorting as it sits on the subjects eye. This type of phenomenon could be easily studied with SPINALTOP, which could lead to the development of contact lenses that can further improve visual performance. These types of measurements were not possible with the Keratron Piccolo, nor are the possible with any other known type of system.

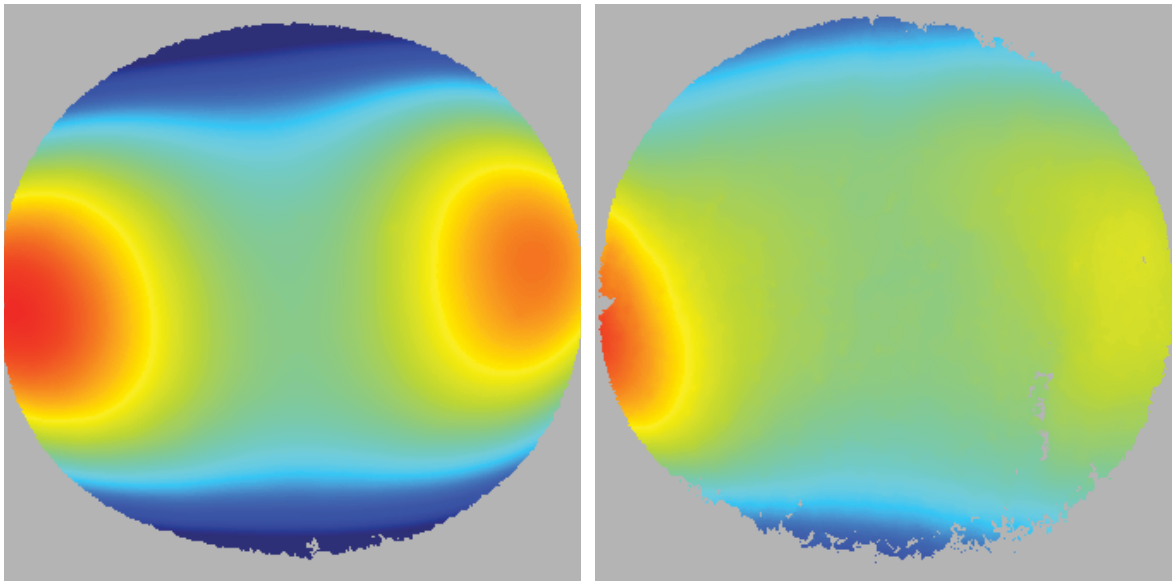

Figure 5.94: Subject G, left eye, conic removed surface comparison.
(Left) Bare eye. (Right) Contact lens. Measurements made on SPINALTOP. Display area is 6 x 6 mm. ±16.0 µm color height scale; red is away from eye.

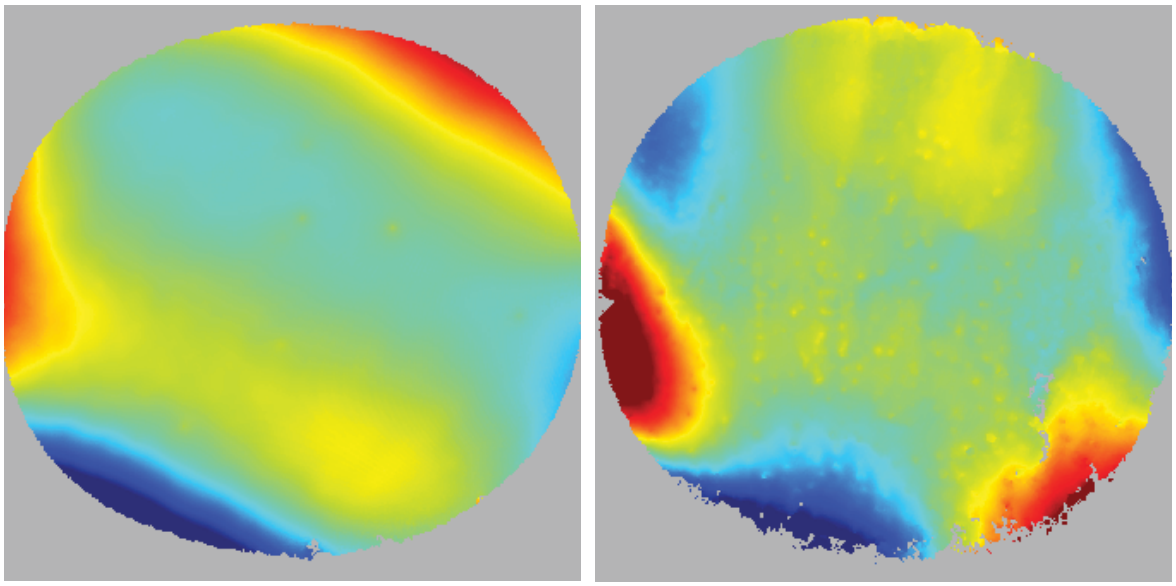

Figure 5.95: Subject G, left eye, biconic removed surface comparison.
(Left) Bare eye. (Right) Contact lens. Measurements made on SPINALTOP. Display area is 6 x 6 mm.
±3.5 µm color height scale; red is away from eye

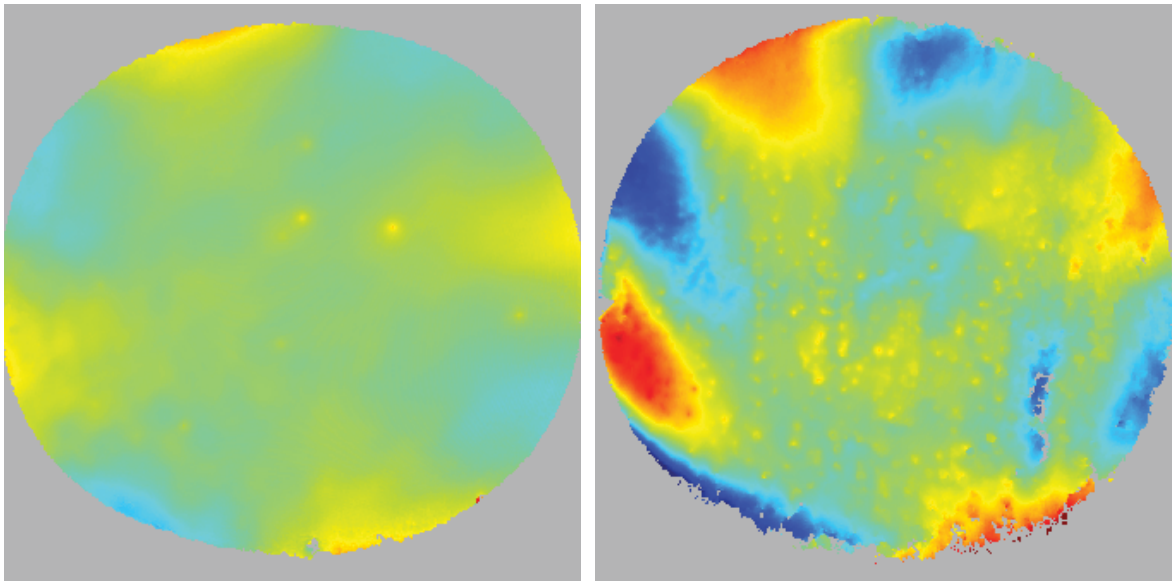

Figure 5.96: Subject G, left eye, Zernike removed surface comparison.
(Left) Bare eye. (Right) Contact lens. Measurements made on SPINALTOP. Display area is 6 x 6 mm.
±2.0 µm color height scale; red is away from eye.

5.3.8 CONCLUSION ON HUMAN SUBJECTS TESTING

The measurements presented in this section demonstrate the SPINALTOP system's ability to outperform the Keratron Piccolo. High resolution corneal topography measurements are able to provide feedback on the overall structure of the human cornea. These measurements could be related to improving our understanding of vision quality, developing better models for the eye and cornea, or provide feedback on the fitting of contact lenses on the eye. The contact lens measurements demonstrated how the underlying corneal topography maps through the contact lens material. Understanding this interaction could provide feedback for the development of newer contact lens materials or possibly the ability to develop specialized contact lenses that are better matched to the unique topography of individual subjects.

The SPINALTOP system is also capable of resolving the mid to high spatial-frequency structures of the cornea, which includes the tear film. The tear film is of interest for many reasons, including basic science, disease studies, and contact lens development. The interaction of the tear film with the contact lens is critical for maintaining visual quality but also ocular health. The high resolution measurements of the topography and tear film with and without a contact lens could be used to support the development of newer contact lens materials.

Additionally, SPINALTOP is capable of producing dynamic results – everything image that has been displayed in this section is effectively a single frame out of a 30 FPS movie. Similar to the TFI results in Section 2.7, every measurement result in this section could be displayed sequentially in time to show the dynamics of the absolute corneal topography and tear film. The combination of a non-invasive, high resolution spatial and height capabilities, absolute topographic reconstruction, and dynamic measurements results in a system that exceeds the capabilities of any known system.

5.4 SPINALTOP VS. TFI

A final topic of discussion for the results chapter is a comparison of the SPINALTOP system to the TFI. Although Chapters 4 and 5 discuss the changes that were involved in developing the SPINALTOP

system, the results that have been presented in this chapter may not have sufficiently demonstrated the differences between the two systems. The purpose of this section is to highlight the differences between the two systems.

The most significant change from the TFI was the addition of the cat's eye interferometer. The purpose of the cat's eye interferometer is to provide a secondary measurement that allows the absolute topography of the cornea to be recovered. The method to recover the absolute topography requires a reverse ray trace operation (Section 4.4). Some of the information that is captured by a corneal topography measurement includes the radius of curvature, asphericity, and corneal astigmatism, just to name a few parameters. None of this information could be determined from the TFI. As discussed in Section 2.7, TFI measurements require that, at a minimum, the first 12 Zernike polynomial terms have to be removed from the surface, destroying any potential information that related to the aforementioned parameters. There is little to no information contained within the removed terms that could be related back to the topography of the cornea. Additional information is required; hence, the cat's eye interferometer.

A number of other methods could have been used to track the position of the cornea, which in turn is required to reconstruct the absolute topography of the eye. The cat's eye interferometer provided one of the highest resolution methods to reconstruct the topography. However, without reverse ray tracing, any of the accuracy gained by using the cat's eye interferometer is lost. This is easily demonstrated with the measurements that have been previously made. The cat's eye interferometer provides a secondary measurement that is equivalent to translating the interferometer from a cat's eye to confocal measurement position. This means that the calibration process used to track the displacement of the cornea can be modified to instead track the absolute radius of curvature of the cornea (Section 4.3.3). In other words, the polynomial used to track the cornea (Eq. 4.114) is fitted as a function of radius vs. cat's eye/surface interferometer Zernike coefficients. The polynomial fitting error will be the same as the displacement; the error in estimating the radius of curvature error will be 1.5 μm RMS and 8.4 μm PV (Section 5.1.2). However, this error only applies to the specific shapes used in the calibration process, as will be shown.

A pseudo-reconstruction can then be performed. For every measurement made with the system, the cat's eye interferometer can be used to determine the radius of curvature of the surface under test. The radius of curvature is applied to a nominal eye model. The nominal eye model is the average corneal shape for which the converger assembly was designed for (Section 2.1). The measured optical path differences at the detector can be double-pass and parity corrected, then mapped onto the nominal eye model surface. The resulting surface is treated as the pseudo-reconstructed eye surface. The previous operations from Section 5.2 (Verification) are then run on the pseudo-surfaces. In other words, for every measurement that was made with SPINALTOP for verification, the same set of data can be re-used with the pseudo-reconstruction method described above. The same processes can be re-used (e.g. biconic fit parameters, surface fit residuals) and plotted with the box and whisker plots that were used for Figure 5.16 through Figure 5.26. The pseudo-reconstruction results are plotted as box and whisker plots, along with the previous verification data and are shown in Figure 5.97 through Figure 5.103.

What these results show is that without reverse ray tracing, the system performs worse than the Keratron Piccolo. Parts 001-009 appear to be somewhat comparable to the Keratron results, but the errors are deceptively low due to the fact that the surfaces are rotationally symmetric with a conic constant that matches the nominal eye model. Once the conic constant changes (parts 011-014) or the surface is non-rotationally symmetric (parts 101-105), the errors quickly blow up. Figure 5.104 shows a comparison of the biconic removed surface residuals for part 105 using the pseudo-reconstruction method as described above and the standard SPINALTOP method. The residual error, which does not include radius and other low-spatial frequency terms, is nearly 1 μm (PV). Therefore, a method is required to correct for retrace errors (e.g. reverse ray tracing) to recover any of the absolute corneal topography with any reasonable level of accuracy.

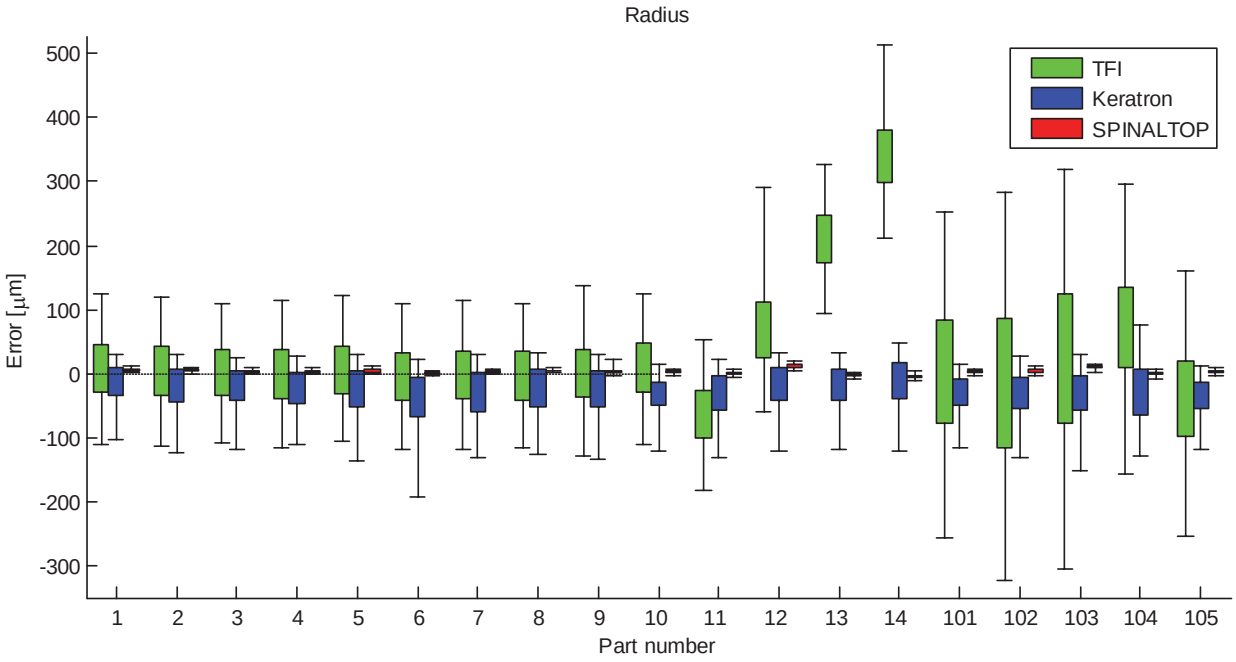

Figure 5.97: TFI radius estimation errors with reference and verification surfaces. Errors are in microns.

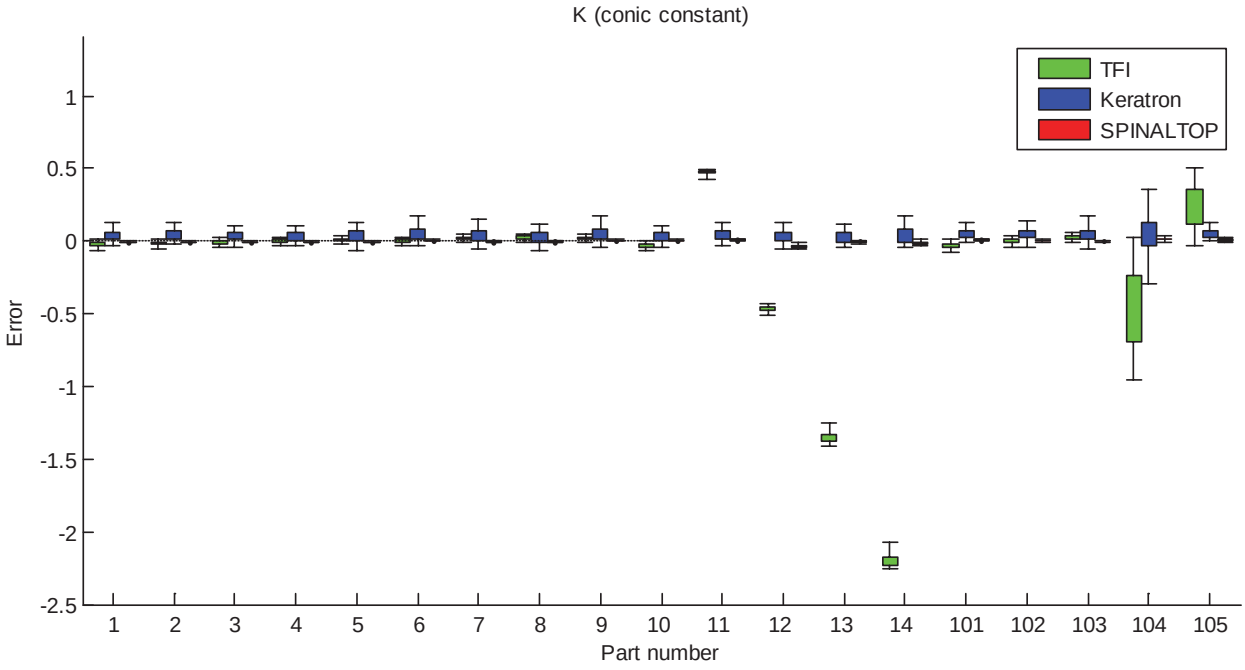

Figure 5.98: TFI conic constant estimation errors with reference and verification surfaces. Errors are unitless.

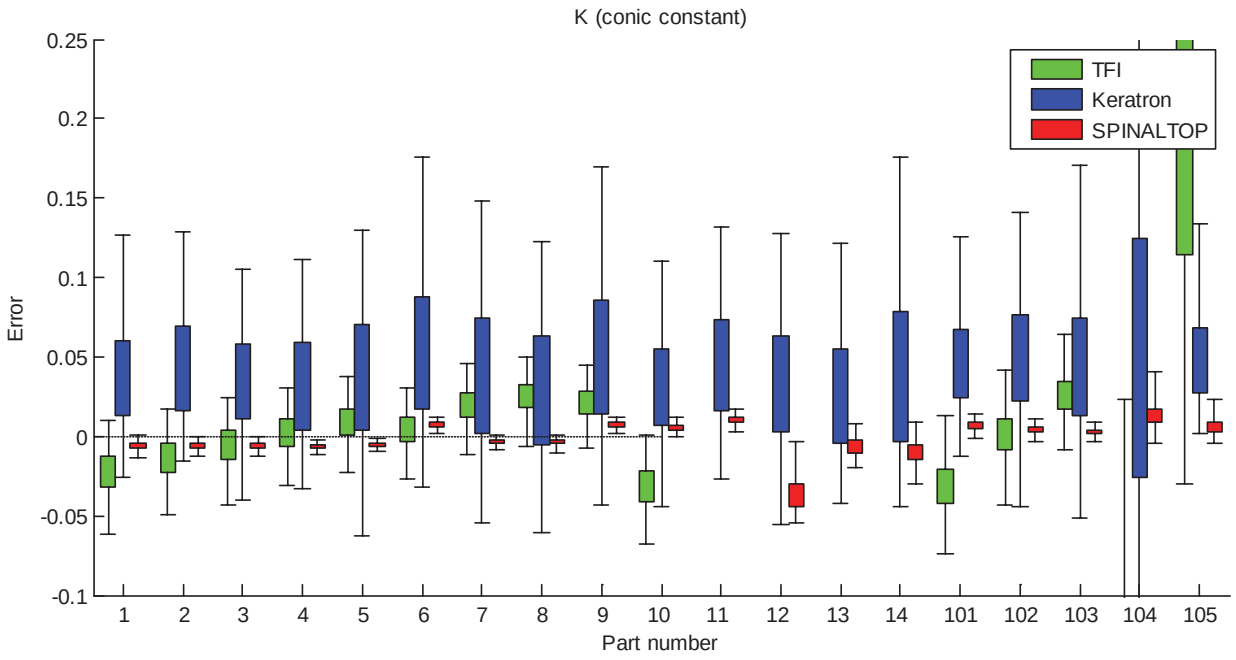

Figure 5.99: TFI conic constant estimation errors (enhanced).
Errors are unitless.

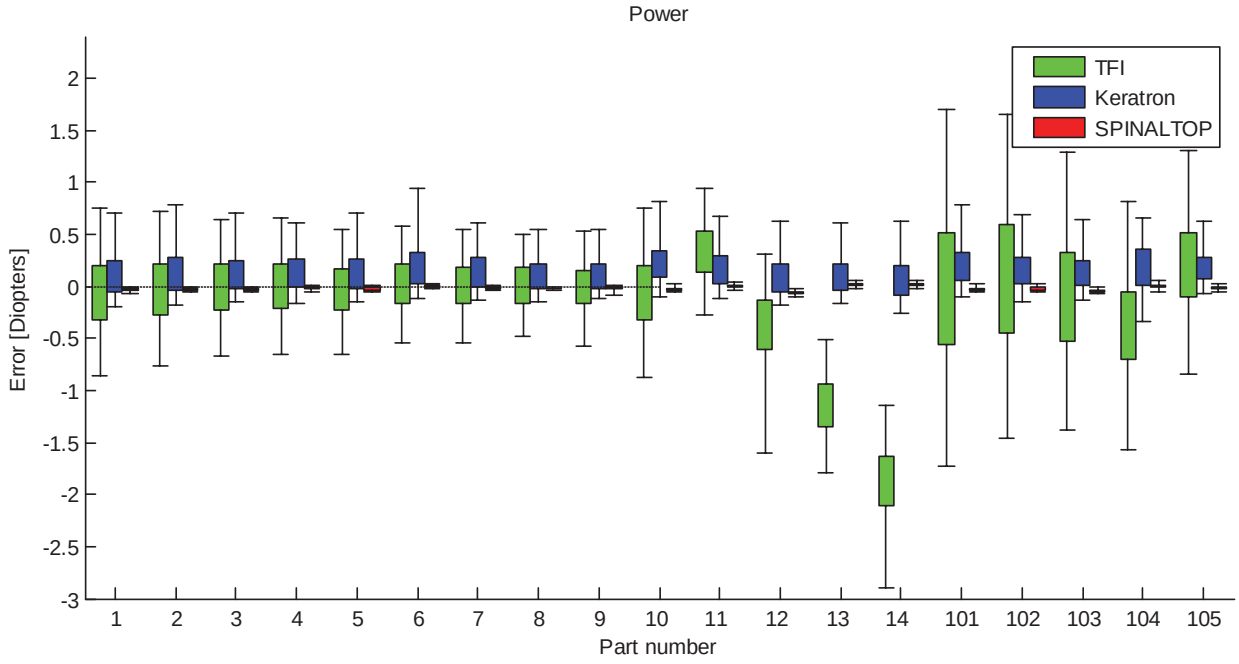

Figure 5.100: TFI power estimation errors with reference and verification surfaces.
Errors are in diopters.

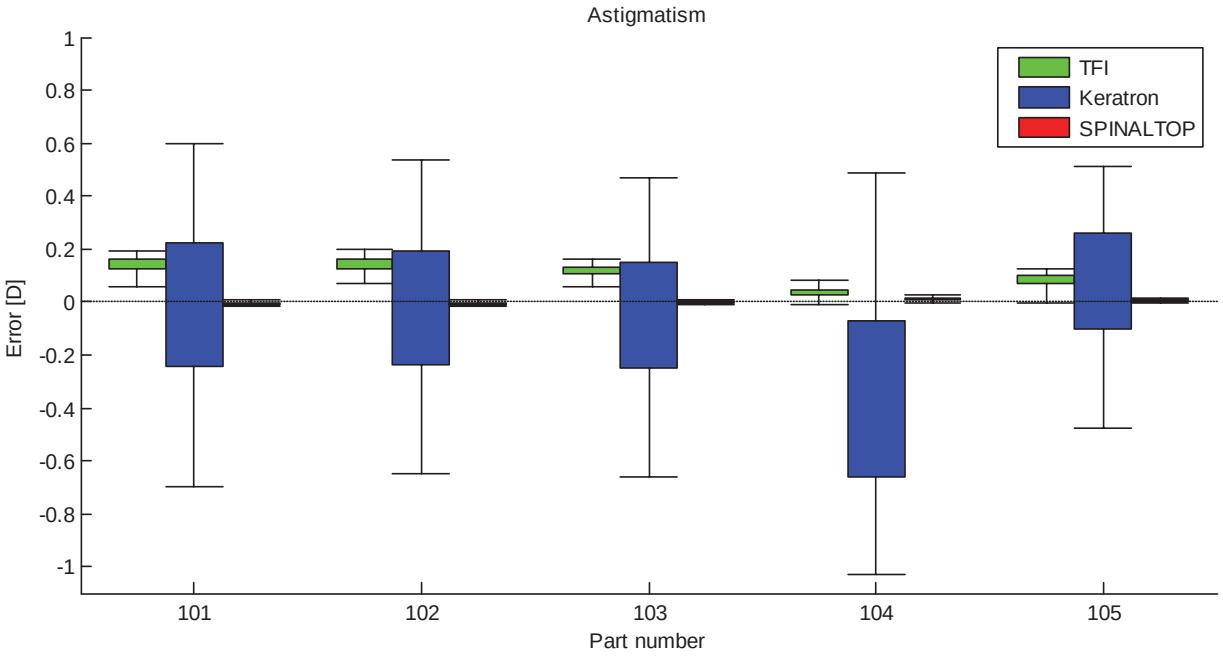

Figure 5.101: TFI astigmatism estimation errors for reference and verification surfaces. Errors are in diopters.

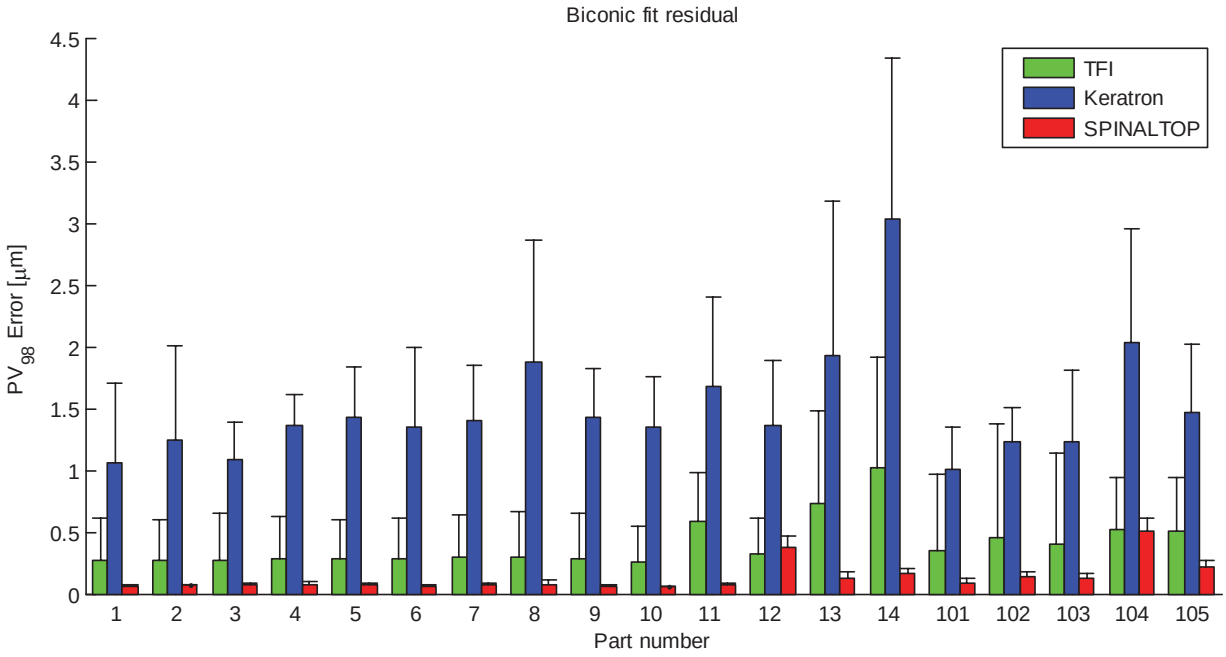

Figure 5.102: TFI biconic fit, PV_{98} residual errors with reference and verification surfaces. Errors are in microns.

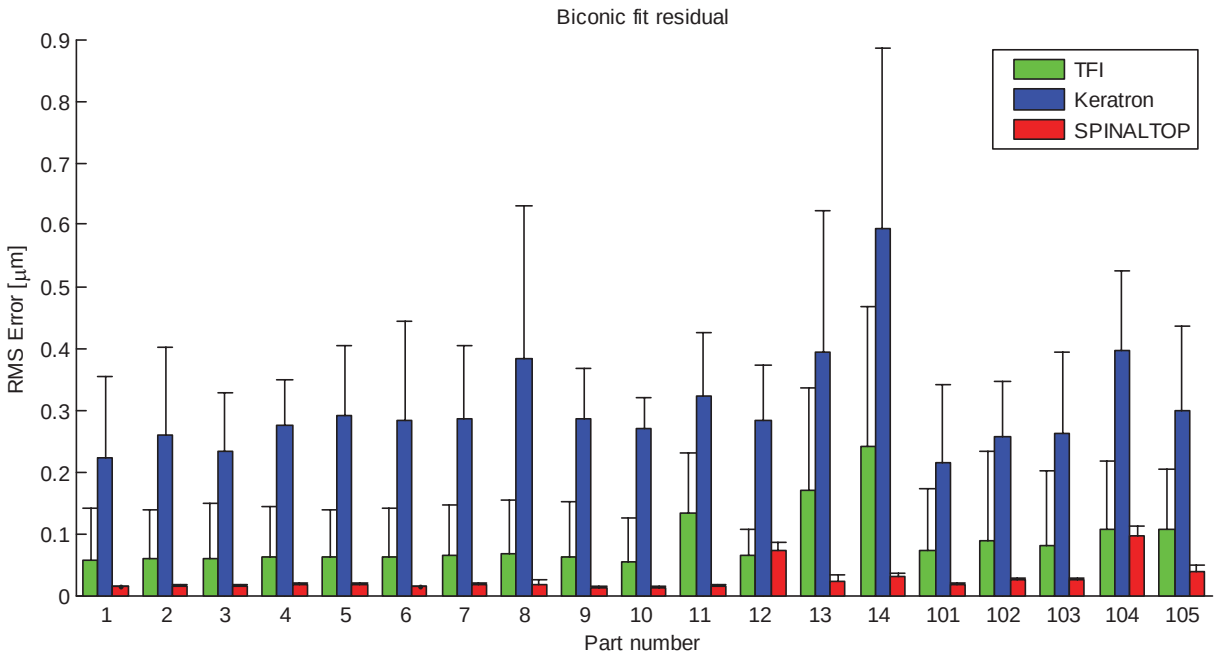

Figure 5.103: TFI biconic fit, RMS residual errors with reference and verification surfaces. Errors are in microns.

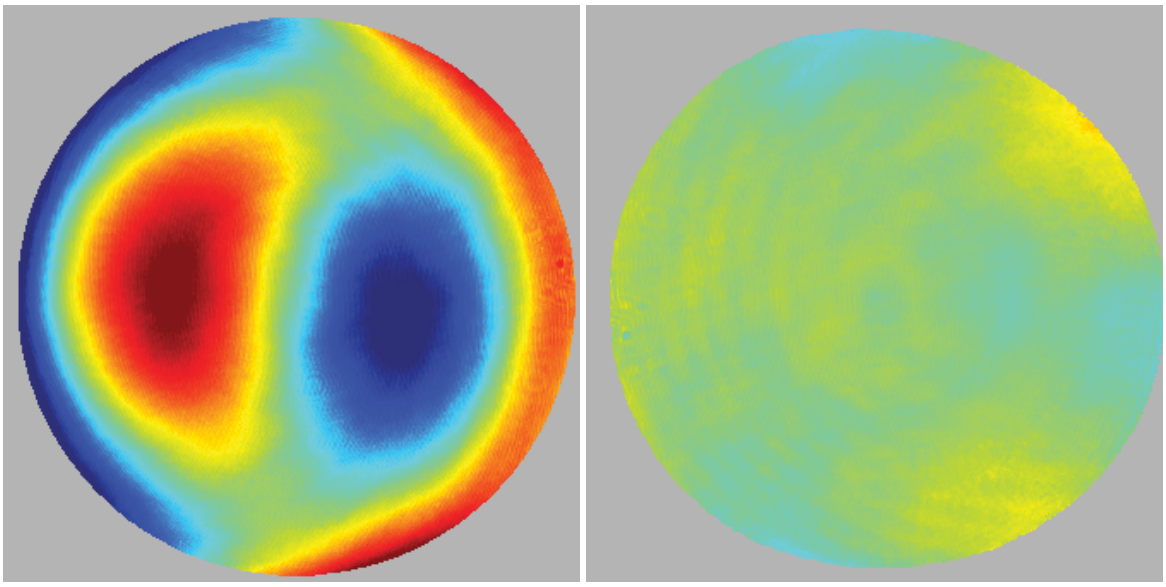

Figure 5.104: Part 105 biconic removed residual error comparison of pseudo-reconstruction to reverse ray trace correction. (Left) TFI with cat's eye data. (Right) SPINALTOP reconstruction. The surface was defocused by 150 μm and translated by 75 μm. Display area is 6.5 x 6.5 mm. ±0.5 μm color height scale; red is away from part surface.

An additional benefit that comes out of the reverse ray trace operation is distortion correction. The amount of distortion is defined as:

$$\%\text{Distortion} = \frac{\left(H_p - h_p\right)}{h_p} \times 100 \qquad (5.10)$$

Where H_p is a real ray height and h_p is the paraxial ray height (Smith, 2000). A real ray height was calculated from the model of the surface interferometer to be $H_p = 4.80\,mm$ and a corresponding paraxial ray height was $h_p = 6.02\,mm$ (Figure 5.105), resulting in 20% distortion for this system. A plate scale or distortion calibration could be attempted on the system, but as demonstrated in Section 4.4.4.1 this could introduce even greater retrace errors (Figure 4.67).

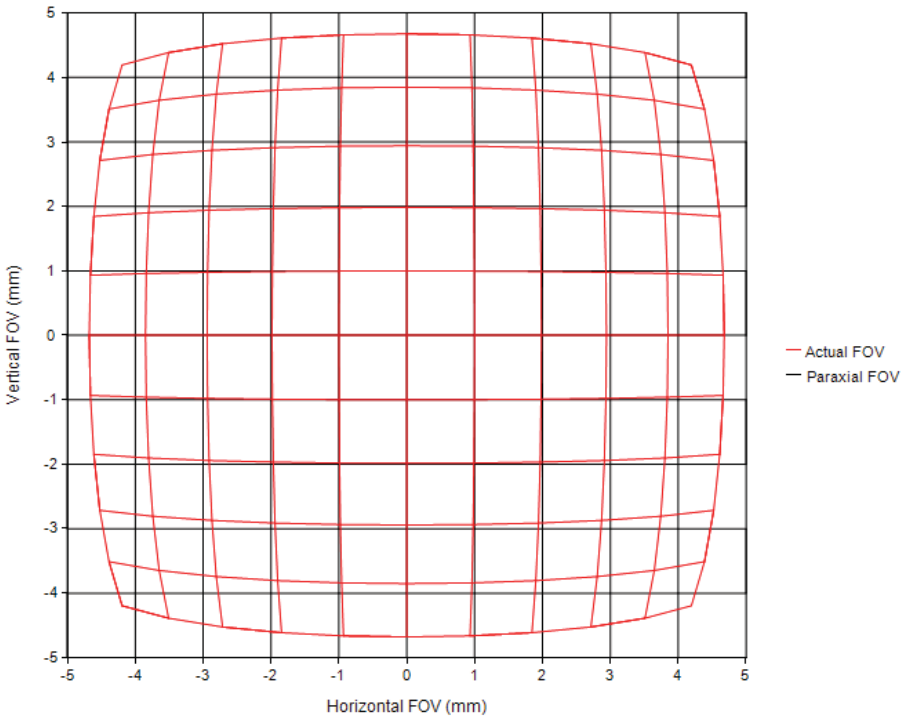

Figure 5.105: Surface imaging system distortion with a nominal alignment.

Figure 5.106 shows an example of parts 001 (left column) and 009 (right column) that were processed with SPINALTOP (top row) and the pseudo-reconstruction with the TFI (bottom row). The SPINALTOP surfaces have had the best fit biconic removed and the TFI surfaces have had the first 15 Zernike

polynomial terms removed. This example demonstrates how the TFI can over- or under-estimate the reconstruction area and spatial resolution of a surface. A more extreme example of a surface reconstruction can be seen with Figure 5.84, where a large amount of corneal astigmatism resulted in an asymmetric reconstruction area. Had this been tested with the TFI system, this information would have been unknown.

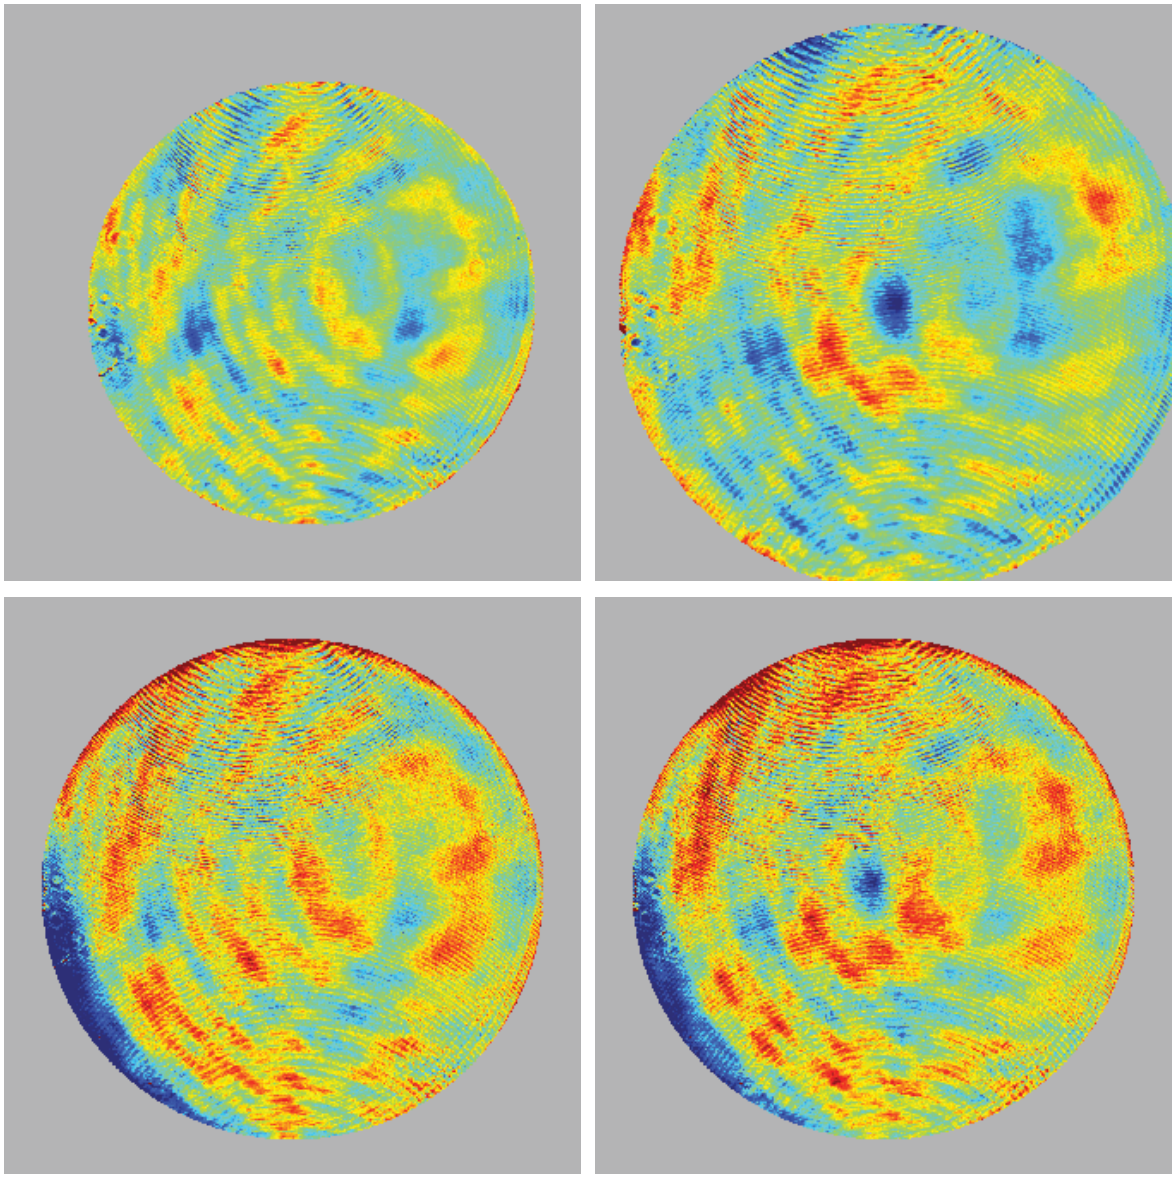

Figure 5.106: Distortion uncertainty with TFI compared to SPINALTOP.
(Left) Part 001. (Right) Part 009. (Top) SPINALTOP reconstruction with biconic fit removed. (Bottom) TFI measurement with 15 Zernike terms removed. Both surfaces have been defocused by 150 μm and translated by 75 μm. Display area is 7 x 7 mm. ±50 nm color height scale.

Another change that was made with the SPINALTOP system was the in-line fixation. The original intent to move the fixation internal to the system was to simplify the measuring process. With the external fixation of the TFI, the setup made testing on both eyes an unnecessarily difficult chore. The subject would first be placed in front of the TFI. The external fixation assembly had to be adjusted to the subject, which included translating and rotating the mirror that folded the fixation system (Section 2.1). The fold mirror was often difficult to reach and difficult to adjust. Once the fixation system was aligned, the interferometer could be aligned. If it was desired to test on the subject's other eye, the fixation assembly had to be torn down and moved to the other side. This was a lengthy and unnecessary process, so most testing was performed only on one eye. The internal fixation assembly fixed all of these issues. Subjects will place their head in the head-rest fixture and the operator will align the SPINALTOP system to the subject. Fixation is in-line with the system, so the subject will see the fixation target with the same eye that will be tested when the system is nominally aligned to the subject. The operation is similar to how most ophthalmic devices are used in any clinical setting. The system can easily move between both eyes and without the need for the subject to remove their head from the head-rest.

A second issue with the external fixation assembly was accommodation. The TFI fixation offered limited fixation, especially for subjects requiring more than 2 diopters of correction. The internal fixation does not have this issue. An internal focusing lens can be used to accommodate for up to 10 diopters or more. Additionally, a system aperture controls the working f/# of the imaging system, allowing the subject to view the fixation target with minimized aberration content.

A final issue is the alignment of the subject's eye to the interferometer. With the external fixation of the TFI, there was limited control on where to direct the subject's gaze. The direction of gaze was dictated by the off-eye fixation target and seemingly random with the fold mirror adjustments. It is estimated that a subject's head could be tilted as much as 10^{o}, including the subject's line-of-sight, with no impact on the system to collect the data. The in-line fixation ensures that the line-of-sight of the individual is co-incident with the interferometer optical axis. This is accomplished by aligning all of the sub-systems

(surface, cat's eye, and fixation) to the optical axis of the converger assembly. The slower working f/# of the fixation imaging system and the matched pupil design ensures that the subject must be aligned to the fixation assembly, and thus the interferometer. This has been verified during testing with the subject's head placed in the head-rest at an angle – the subject can no longer view the fixation target or the target appears vignetted and distorted to the point that it is unusable. Data could be collected with the system in this state, although finding an initial alignment with the subject and maintaining it becomes exponentially more difficult.

A final correction made with SPINALTOP that is noteworthy is the issue of vignetting. Section 2.4 originally discussed the issue with uncertainty as to what was the source of vignetting. A number of elements were swapped out with SPINALTOP, with the exception of the converger assembly. However, it appears as though the issue was resolved by replacing one of the other optical elements in the system. Figure 2.9 demonstrated a measurement on the TFI for a part that was defocused by 100 μm and decentered by 75 μm. However, all of the measurements shown in the verification section have had surface defocused and decentered by amounts equal to and greater than this value with minimal issues.

A number of significant changes have occurred between the TFI and SPINALTOP systems, which included both hardware and software. These changes were necessary for the ability to recover the absolute surface topography of the cornea and dynamic tear film. This section has demonstrated the importance of these changes.

6 CONCLUSION

6.1 SPINALTOP PERFORMANCE SUMMARY

A dual interferometer system for measuring the dynamic corneal topography has been designed, constructed, demonstrated, and tested on human subjects. The system, dubbed SPINALTOP, was demonstrated in Chapter 5 to outperform a commercial corneal topographer, the Keratron Piccolo. SPINALTOP was shown to resolve low-spatial frequency surface height information to 3 ± 4 μm (bias ± 1σ), whereas the Keratron could only resolve to -25 ± 51 μm. The ability of the system to resolve low-spatial frequency height structure determines how well it can estimate the shape parameters of the cornea, such as the radius of curvature, asphericity, or corneal astigmatism. SPINALTOP was further shown to resolve mid- to high-spatial frequency height information to 25 nm RMS (130 nm PV), whereas the Keratron could only resolve 306 nm RMS (1500 nm PV). Most of the structural information at these spatial frequencies relate to residual corneal topography and tear film structure. Spatial resolution of the SPINALTOP system was shown to be 6 μm. The spatial resolution of the Keratron could not be accurately determined and it is estimated to be >100 μm.

The SPINALTOP system is capable of making non-invasive, high resolution, dynamic measurements of the absolute surface topography of the cornea and tear film structure. This includes four separate requirements met by the SPINALTOP system that cannot be met by any known measurement system. The Keratron Piccolo is an example of the class of topographers (reflection-based Section 1.6.2) that is most capable of passively measuring the cornea topography, but at the great cost of low spatial and height resolution. Raster-stereography (Section 1.6.3) and Moiré-deflectometry (Section 1.6.4) provide high resolution measurements of the corneal topography, but require a disruptive substance to be added to the tear film. Slit-scanning (Section 1.6.6), confocal microscopy (Section 1.6.6), and optical coherence tomography (Section 1.6.7) all provide high resolution measurements of the cornea, but cannot measure the dynamics of the tear film. Holography (Section 1.6.8) and early interferometry (Section 1.6.9)

provided some of the first non-invasive and high-resolution measurements of the tear film, but could not determine any of the corneal topography. SPINALTOP has all of the advantages of the previously stated requirements, without any of the functional disadvantages.

SPINALTOP also presents a significant advancement in capabilities over the previous TFI system (Chapter 2). The TFI is not capable of determining the absolute topography of the cornea (e.g. radius, asphericity) and has no mechanism for distortion correction. Measurements made on the TFI are dominated by enough retrace error, that at a minimum the first 12 Zernike polynomial terms have to be removed. The addition of the cat's eye interferometer and development of reconstruction software enabled SPINALTOP system to recover the absolute topography. As demonstrated in Section 5.4, an eye tracking mechanism, such as the cat's eye interferometer, by itself is not sufficient to recover any low-spatial frequency surface information. A method to correct for retrace error was required. A final change from the TFI was the addition of the in-line fixation assembly. In-line fixation simplifies the alignment process, including the ability to test both eyes in a single session, and constrains the location of the cornea with respect to the interferometer.

The utility of a system with the capabilities of SPINALTOP may not be fully known at this time, but there are already a number of potential applications. A primary motivation for this instrument is the ability to characterize the *in vivo* tear film and its interaction with a contact lens. Preliminary data from human subjects testing has shown that the contact lens shape and tear film structure can be measured *in vivo*. Examples of what can be done with the measurement data include looking at how the tear film shape is altered by the presence of the contact lens to ensure fitting and proper refractive correction. The tear film structure can also be studied to better understand how the material properties effects tear film stability. Information from these types of measurements could be used as feedback in the development of new contact lens materials and designs.

The potential applications for this system can be used in other areas of ophthalmology, providing a better understanding of the corneal topography and dynamic tear film structure. Higher resolution mapping of the corneal structure may provide new information that can be used to improved current eye models. Improved eye models could account for the dynamic changes of the cornea, including the tear film, which could be used to predict visual performance of the eye. The corneal topographer aspect of the system may be used for studying eye diseases that affect the structure of the cornea, such as keratoconus. The tear film measurement capabilities could be used to study diseases that affect the tear film, such as dry eyes.

Methods developed for the SPINALTOP system design could be extended beyond ophthalmology to support areas of optical metrology. The reverse ray tracing and instantaneous surface reconstruction methods could be used for aspheric and freeform metrology.

6.2 FUTURE WORK

6.2.1 ACCURACY IMPROVEMENTS

SPINALTOP was demonstrated to have an absolute surface reconstruction accuracy of 3 ± 4 μm. The accuracy is predominantly limited by the resolution of the corneal displacement measurement, or more specifically the cat's eye interferometer. This can be attributed to the converger assembly re-use. First, the current working f/# of the cat's eye interferometer is f/5. Initial plans for the system (i.e. before budget changes) were to increase the coverage of the test surface to 8 mm from 6 mm. In the current configuration, this would decrease the working f/# and increase the sensitivity of the cat's eye system by almost 2X. The current system could be modified to place the cat null (Section 3.3.4) closer to the converger assembly, allowing for a faster working f/#. By placing priority on the cat null instead of the fixation assembly, the system could gain an additional 3X in displacement accuracy. A more advanced design for the converger assembly could result in an order of magnitude improvement. Figure 6.1 shows a preliminary converger design that uses dispersion to shift the focus of the cat's eye system to the cat's eye position. The converger accepts collimated light for two wavelengths. One wavelength produces a null

with the nominal eye model, similar to the current converger. The secondary wavelength is shifted to a focus 7.8 mm in front of the surface, producing a cat's eye measurement. The system produces an f/1 cat's eye beam that results in nearly 25X increase in sensitivity or potential improvement in surface reconstruction accuracy.

When the best fit biconic surface was removed (Eq. 4.162) the surface height errors were within 25 nm RMS (130 nm PV). The remaining accuracy could be improved by further characterization of the interferometer. As previously discussed, a portion of the residual errors has to do with the incomplete characterization of the converger assembly (Section 5.1.1). One option is to tear apart the converger assembly, characterize every element, and re-assemble it. Another option would be to build a new converger assembly and characterize it before assembly. Either of the two options would allow for a more complete interferometer model to be built up and these errors could be further reduced.

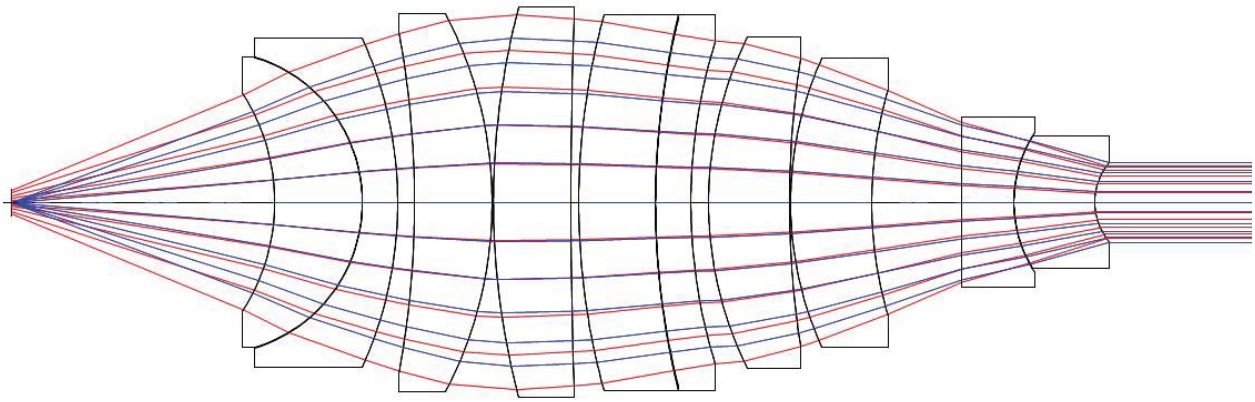

Figure 6.1: Achromatic converger for surface and cat's eye interferometers.
Red path is designed for 785 nm (surface), blue path is for 850 nm (cat's eye).

6.2.2 COMPACT SYSTEM DESIGN

The SPINALTOP system is an experimental system and is inefficiently designed for clinical or commercial use. For example, many optical components have adjustment capabilities that would not be necessary in a more compact system. However, these are obvious areas for improvement. The largest burden on space and weight in the system is the illumination system. Instability of the TFI laser source

required the addition of laser monitoring systems to SPINALTOP. Improved laser sources are available on the market. This would remove the Faraday isolators, Fabry-Pérot etalons, and necessary pick-offs. Additionally, the lasers sources could be offloaded to a separate assembly and fiber-coupled into the system. This would eliminate all of the monitoring systems, modulation, shutter, and spatial filter-objective systems. These changes would remove everything that is highlighted in red in Figure 6.2 from the main unit.

A more compact design would involve an achromatization of the surface and cat's eye sources. This would allow both interferometers to operate on the same beam-paths, up until the just before the imaging systems. Achromatization would require achromatic polarization optics and a custom converger assembly like the one shown in Figure 6.1. The achromatized converger assembly would accept two collimated beams of different color (e.g. 785 and 850 nm) and generate an aspheric wavefront that matches the nominal shape of the eye with one color, while the other color comes to focus on or near the cornea for the same eye model. Achromatization would eliminate most of the cat's eye interferometer system, highlighted in yellow in Figure 6.2.

More advanced optics could be used in the test path to provide a cat's eye measurement that could alternate between surface measurements. For example, the surface dichroic beamsplitter (Section 3.2.4) could be replaced by a deformable mirror. Every other frame would leave the mirror in an unpowered or non-deformed state which would provide the surface measurement. Alternating frames would actuate the mirror to shift the focus of the converger to the cat's eye position. The current configuration would require a maximum of an 80 μm stroke consisting of mostly power and some astigmatism to account for the oblique angle of incidence. Other similar methods could be used, such as a deformable lens or spatial light modulator directly in the test path. One down-side to this method is that the temporal resolution of the system would be halved (e.g. from 30 FPS to 15 FPS). However, such a method could greatly simplify the design while still being capable of producing high resolution measurements of the cornea and tear film.

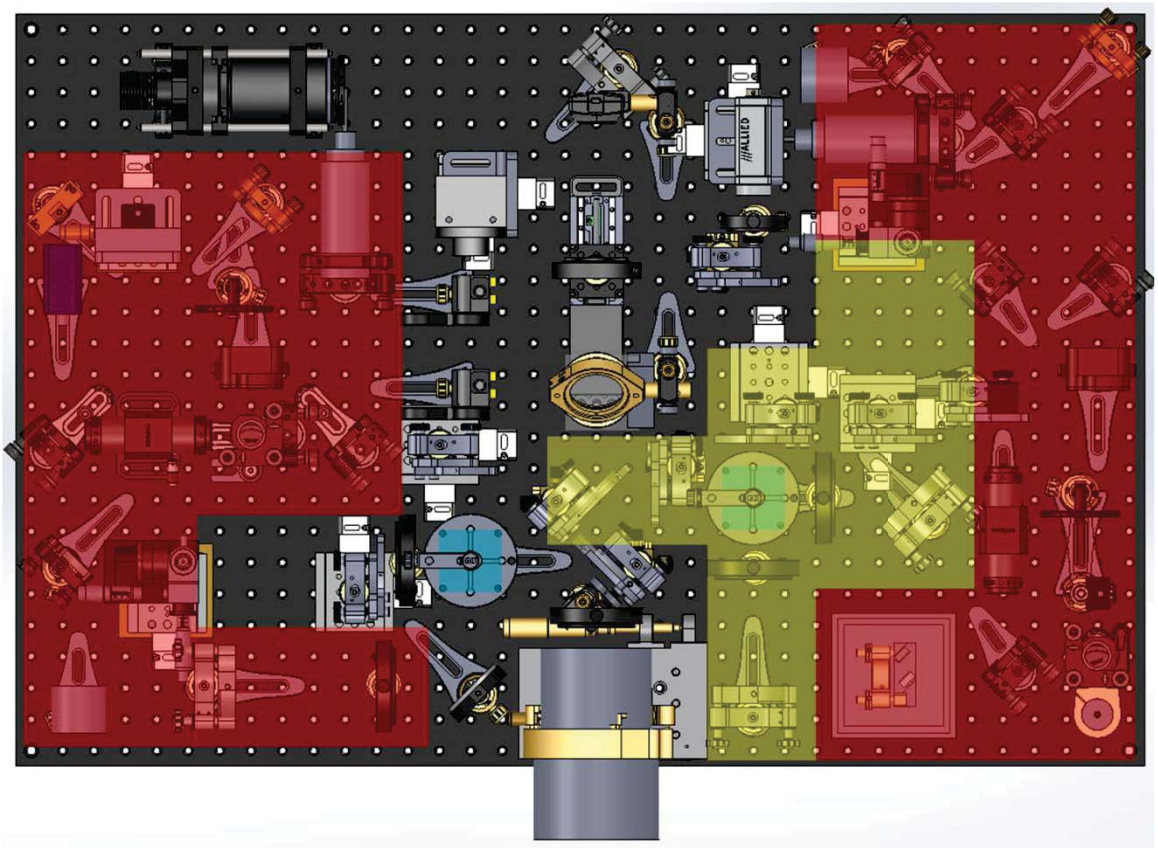

Figure 6.2: Hardware reduction improvements for SPINALTOP.

6.2.3 MISCELLANEOUS HARDWARE IMPROVEMENTS

In addition to achromatization of the converger, other systems could be considered. As noted in Section 3.4.2, designing a pupil tracking system with the pre-existing TFI converger presented serious issues. A custom converger design could also feed in a pupil or eye tracking camera system. Improvements to the fixation assembly could also be made by properly designing the converger. Although it may seem nearly impossible to design a single optical assembly for four separate systems, it should be considered that three separate systems were demonstrated to work with a design optimized for a single system.

Other basic improvements with the system can be made, such as increasing the dynamic range of the system to account for greater eye variation and/or eye motion. Additionally, the system could be pushed to a larger test area on the cornea. A more impressive addition would be to combine a system that can measure tear film thickness simultaneously with the corneal topography measurement capabilities of

SPINALTOP to provide a complete characterization of the tear film and corneal structure. Systems for high resolution tear film thickness measurements can be found in the literature (King-Smith, Fink, Nichols, Nichols, & Hill, 2006).

6.2.4 SOFTWARE IMPROVEMENTS

Although the intent of this project was not to develop software, the ability of the system to recover the surface topography is dependent on a number of algorithms (Chapter 4). Many of the algorithms and methods implemented were designed to be “good enough”. That is, no attempt was made to work on 100% of the data, nor was it made to operate efficiently. A fair amount of development could be put into optimizing these routines and determining what methods may or may not be necessary. One example in particular is the reverse ray tracing process. The process still relies on ray tracing software to counter-propagate the recovered test wavefront to a test surface. More efficient approximations could be developed, such as deriving an algorithm that maps a measured OPD map to test space through a polynomial description. Other additions to the software controls could be made to improve the alignment capabilities of the system. The instantaneous phase shifting method used on the surface interferometer allows for near instantaneous feedback on alignment. This could allow the system to be put into a closed-loop using feedback from the surface measurements maintain alignment of the instrument to the subject, potentially increasing collection duration and data fidelity.

REFERENCES

ANSI. (2007). American National Standard for Safe Use of Lasers. *ANSI Z136.1-2007*.

ANSI. (2008). ANSI Z80.23-2008 American National Standard for Ophthalmics – Corneal Topography Systems – Standard Terminology, Requirements.

Baker, P. C. (1990). Holographic Contour Analysis of the Cornea. In *Noninvasive Diagnostic Techniques in Ophthalmology* (pp. 82–98). Springer New York. doi:10.1007/978-1-4613-8896-8_6

Barlow, H. B. (1952). Eye movements during fixation. *The Journal of Physiology*, *116*(3), 290–306. doi:10.1113/jphysiol.1952.sp004706

Barrett, R., Berry, M., Chan, T. F., Demmel, J., Donato, J., Dongarra, J., … van der Vorst, H. (1994). *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods* (Second Edi.). Philadelphia: SIAM.

Belin, M. W., Cambier, J. L., Nabors, J. R., & Ratliff, C. D. (1995). PAR corneal topography system (PAR CTS): The clinical application of close-range photogrammetry. *Optometry and Vision Science*, *72*(11), 828–837.

Benedetto, D., Clinch, T. E., & Laibson, P. R. (1984). In vivo observation of tear dynamics using fluorophotometry. *Archives of Ophthalmology*, *102*(3), 410–412. doi:10.1001/archopht.1984.01040030328030

Bone, D. J. (1991). Fourier fringe analysis: The two-dimensional phase unwrapping problem. *Applied Optics*, *30*(25), 3627–3632. doi:10.1364/AO.30.003627

Boyce, P. R. (1967). The Effect of Change of Target Field Luminance and Colour on Fixation Eye Movements. *Optica Acta*, *14*(3), 213–217. doi:10.1080/713818033

Braun, R. J., King-Smith, P. E., Begley, C. G., Li, L., & Gewecke, N. R. (2015). Dynamics and function of the tear film in relation to the blink cycle. *Progress in Retinal and Eye Research*, *45*, 132–164. doi:10.1016/j.preteyeres.2014.11.001

Bruning, J. H., Herriott, D. R., Gallagher, J. E., Rosenfeld, D. P., White, A. D., & Brangaccio, D. J. (1974). Digital wavefront measuring interferometer for testing optical surfaces and lenses. *Applied Optics*, *13*(11), 2693–2703. doi:10.1364/AO.13.002693

Burge, J. H. (1995). Applications of computer-generated holograms for interferometric measurement of large aspheric optics. *Proceedings of SPIE*, *2576*, 258–269. doi:10.1117/12.215609

Cairns, G., McGhee, C. N. J., Collins, M. J., Owens, H., & Gamble, G. D. (2002). Accuracy of Orbscan II slit-scanning elevation topography. *Journal of Cataract & Refractive Surgery*, *28*(12), 2181–2187. doi:10.1016/S0886-3350(02)01504-3

Calkins, J. L., Hochheimer, B. F., & Stark, W. J. (1981). Corneal wound healing: Holographic stress-test analysis. *Investigative Ophthalmology & Visual Science*, *21*(2), 322–334.

Chang, W.-Y., Chen, K.-H., Chen, D.-C., Tseng, J.-K., Chen, S.-T., Sun, H.-Y., … Hsu, K. Y. (2014). Heterodyne moiré interferometry for measuring corneal surface profile. *Optics and Lasers in Engineering*, *54*, 232–235. doi:10.1016/j.optlaseng.2013.07.013

Cheng, Y.-Y., & Wyant, J. C. (1984). Two-wavelength phase shifting interferometry. *Applied Optics*, *23*(24), 4539–4543. doi:10.1364/AO.23.004539

Creath, K., Cheng, Y.-Y., & Wyant, J. C. (1985). Contouring Aspheric Surfaces Using Two-wavelength Phase-shifting Interferometry. *Optica Acta*, *32*(12), 1455–1464. doi:10.1080/713821689

Daily, L., & Coe, R. E. (1962). Lack of effect of anesthetic and mydriatic solutions on the curvature of the cornea. *American Journal of Ophthalmology*, *53*, 49–51. doi:10.1016/0002-9394(62)90395-1

Doane, M. G. (1980). Interaction of eyelids and tears in corneal wetting and the dynamics of the normal human eyeblink. *American Journal of Ophthalmology*, *89*(4), 507–516.

Doss, J. D., Hutson, R. L., Rowsey, J. J., & Brown, D. R. (1981). Method for calculation of corneal profile and power distribution. *Archives of Ophthalmology*, *99*(7), 1261–1265. doi:10.1001/archopht.1981.03930020135018

Drexler, W., Morgner, U., Ghanta, R. K., Kartner, F. X., Schuman, J. S., & Fujimoto, J. G. (2001). Ultrahigh-resolution ophthalmic optical coherence tomography. *Nature Medicine*, *7*(4), 10–15. doi:10.1038/86589

Dubra, A., Paterson, C., & Dainty, C. (2005). Double lateral shearing interferometer for the quantitative measurement of tear film topography. *Applied Optics*, *44*(7), 1191–1199. doi:10.1364/AO.44.001191

Engbert, R., & Kliegl, R. (2003). Microsaccades uncover the orientation of covert attention. *Vision Research*, *43*(9), 1035–1045. doi:10.1016/S0042-6989(03)00084-1

Evans, C. J. (1994). Software based improvements in the accuracy of measurement of aspherics using a Fizeau interferometer. *Optical Fabrication and Testing*, *13*, 259–262.

Evans, C. J., & Bryan, J. B. (1993). Compensation for Errors Introduced by Nonzero Fringe Densities in Phase-Measuring Interferometers. *CIRP Annals - Manufacturing Technology*, *42*(1), 577–580. doi:http://dx.doi.org/10.1016/S0007-8506(07)62513-X

Fazekas, Z., Soumelidis, A., Bodis-Szomoru, A., & Schipp, F. (2008). Specular surface reconstruction for multi-camera corneal topographer arrangements. *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*, *2008*, 2254–2257. doi:10.1109/IEMBS.2008.4649645

Fleming, M., & Mooradian, a. (1981). Spectral characteristics of external-cavity controlled semiconductor lasers. *IEEE Journal of Quantum Electronics*, *17*(1), 44–59. doi:10.1109/JQE.1981.1070634

Friedlander, M. H., Mulet, M., Buzard, K., & Granet, N. (1991). Holographic Interferometry of the Corneal Surface. *Proceedings of SPIE*, *1429*, 229–236. doi:10.1117/12.44675

Fry, G. A., & Hill, W. W. (1962). The center of rotation of the eye. *American Journal of Optometry and Archives of American Academy of Optometry, 39*, 581–595.

Gappinger, R. O., & Greivenkamp, J. E. (2004a). Iterative reverse optimization procedure for calibration of aspheric wave-front measurements on a nonnull interferometer. *Applied Optics*, *43*(27), 5152–5161. doi:10.1364/AO.43.005152

Gappinger, R. O., & Greivenkamp, J. E. (2004b). Non-null interferometer for measurement of aspheric transmitted wavefronts. *Proceedings of SPIE*, *5180*, 301–312. doi:10.1117/12.506047

Gaskill, J. D. (1978). *Linear systems, Fourier transforms, and Optics*. New York: Wiley.

Geary, J. M., Yoo, M., & Si, G. (1992). Retrace error: A case study. *Proceedings of SPIE, 1776*, 98–105. doi:10.1117/12.139232

Ghiglia, D. C., Mastin, G. A., & Romero, L. A. (1987). Cellular-automata method for phase unwrapping. *Journal of the Optical Society of America A*, *4*(1), 267–280. doi:10.1364/JOSAA.4.000267

Ghiglia, D. C., & Pritt, M. D. (1998). *Two-dimensional phase unwrapping: Theory, algorithms, and software*. New York, NY: Wiley.

Ghiglia, D. C., & Romero, L. A. (1994). Robust two-dimensional weighted and unweighted phase unwrapping that uses fast transforms and iterative methods. *JOSA A*, *11*(1), 107–117. doi:10.1364/JOSAA.11.000107

Goldstein, R. M., Zebker, H. A., & Werner, C. L. (1988). Satellite radar interferometry: Two-dimensional phase unwrapping. *Radio Science*, *23*(4), 713–720. doi:10.1029/RS023i004p00713

Gollub, G. H., & Van Loan, C. F. (1990). Iterative methods for linear systems. In *Matrix Computations* (2nd ed., pp. 516–538). Baltimore: John Hopkins University Press.

Goodman, J. W. (1985). *Statistical Optics*. New York: Wiley-Interscience.

Goodwin, E. P. (2007). *Dual interferometer system for measuring index of refraction*. University of Arizona.

Greivenkamp, J. E. (1987). Sub-Nyquist interferometry. *Applied Optics*, *26*(24), 5245–5258. doi:10.1364/AO.26.005245

Greivenkamp, J. E., & Gappinger, R. O. (2004). Design of a nonnull interferometer for aspheric wave fronts. *Applied Optics*, *43*(27), 5143–5151. doi:10.1364/AO.43.005143

Greivenkamp, J. E., Williby, G. A., Goodwin, E. P., Primeau, B. C., Heideman, K. C., Micali, J. D., & Spaulding, R. T. (2014). Interferometry and ophthalmics at the College of Optical Sciences. *Proceedings of SPIE*, *9186*, 91860W. doi:10.1117/12.2064333

Guillon, J. P. (1986). Tear film structure and contact lenses. In *The preocular tear film in health, disease and contact lens wear* (Vol. 85, pp. 914–939). Lubbock, Texas: Dry Eye Institute.

Guirao, A., & Artal, P. (2000). Corneal wave aberration from videokeratography: Accuracy and limitations of the procedure. *Journal of the Optical Society of America. A, Optics, Image Science, and Vision*, *17*(6), 955–965. doi:10.1364/JOSAA.17.000955

Halstead, M. A., Barsky, B. A., Klein, S. A., & Mandell, R. B. (1995). A spline surface algorithm for reconstruction of corneal topography from a videokeratographic reflection pattern. *Optometry & Vision Science*, *72*(11), 821–827.

Hariharan, P. (2003). *Optical Interferometry* (Second.). Academic Press.

Heideman, K. C. (2014). *Surface Metrology of Contact Lenses in Saline Solution*. University of Arizona.

Hercher, M. (1968). The spherical mirror Fabry-Pérot interferometer. *Applied Optics*, *7*(5), 951–966. doi:10.1364/AO.7.000951

Hestenes, M. R., & Stiefel, E. (1952). Methods of Conjugate Gradients for Solving Linear Systems. *Journal of Research of the National Bureau of Standards*, *49*(6), 409–436. doi:10.6028/jres.049.044

Heynen, J., & Kahn, D. A. (1988). Eye-position sensor. US4720189.

Holly, F. J. (1980). Tear film physiology. *American Journal of Optometry and Physiological Optics*, *57*(4), 252–257.

Holly, F. J., & Lemp, M. A. (1977). Tear physiology and dry eyes. *Survey of Ophthalmology*, *22*(2), 69–87. doi:10.1016/0039-6257(77)90087-X

Honda, T., Huang, J., Tsujiuchi, J., & Wyant, J. C. (1987). Shape Measurement Of Deep Aspheric Optical Surfaces By Radial Shear Interferometry. *Proceedings of SPIE*, 351–352. doi:http://dx.doi.org/10.1117/12.967294

Hooke, R., & Jeeves, T. A. (1961). "Direct Search" Solution of Numerical and Statistical Problems. *Journal of the ACM*, *8*(2), 212–229. doi:10.1145/321062.321069

Huang, C. (1993). Propagation errors in precision Fizeau interferometry. *Applied Optics*, *32*(34), 7016–7021. doi:10.1364/AO.32.007016

Iskander, D. R. (2009). Modeling videokeratoscopic height data with spherical harmonics. *Optometry and Vision Science*, *86*(5), 542–547. doi:10.1097/OPX.0b013e31819fa8ec

Itoh, K. (1982). Analysis of the phase unwrapping algorithm. *Applied Optics*, *21*(14), 2470. doi:10.1364/AO.21.002470

Jaycock, P. D., Lobo, L., Ibrahim, J., Tyrer, J., & Marshall, J. (2005). Interferometric technique to measure biomechanical changes in the cornea induced by refractive surgery. *Journal of Cataract & Refractive Surgery*, *31*(1), 175–184. doi:10.1016/j.jcrs.2004.10.038

JCGM. (2008). JCGM 200:2008 International vocabulary of metrology — Basic and general concepts and associated terms (VIM). doi:10.1016/0263-2241(85)90006-5

Jeong, H. J., & Lawrence, G. N. (1988). Simultaneous Determination Of Misalignment And Mirror Surface Figure Error Of A Three Mirror Off-Axis Telescope By End-To-End Measurements And Reverse Optimization: Numerical Analysis And Simulation. *Proceedings of SPIE*, *966*, 341–353.

Jeong, H. J., Lawrence, G. N., & Nahm, K. B. (1987). Auto-alignment of a three-mirror off-axis telescope by reverse optimization and end-to-end aberration measurements. *Proceedings of SPIE*, *818*, 419–430. doi:10.1117/12.978915

Jongsma, F. H. M., Brabander, J. de, & Hendrikse, F. (1997). Meten aan het hoornvlies. *Nederlands Tijdschrift Voor Natuurkunde*, *2*, 39–42.

Jongsma, F. H. M., De Brabander, J., Hendrikse, F., & Stultiens, B. A. T. (1998). Development of a Wide Field Height Eye Topographer: Validation on Models of the Anterior Eye Surface. *Optometry and Vision Science*, *75*(1), 69–77.

Jóźwicki, R. (1991). Influence of aberrations of Fizeau interferometer elements on measurement errors. *Applied Optics*, *30*(22), 3126–3132. doi:10.1364/AO.30.003126

Judge, T. R., & Bryanston-Cross, P. J. (1994). A review of phase unwrapping techniques in fringe analysis. *Optics and Lasers in Engineering*, *21*(4), 199–239. doi:10.1016/0143-8166(94)90073-6

Kasprzak, H. T., & Iskander, D. R. (2010). Ultrasonic Measurement of Fine Head Movements in a Standard Ophthalmic Headrest. *IEEE Transactions on Instrumentation and Measurement*, *59*(1), 164–170.

Kasprzak, H. T., Kowalik, W., & Jaroński, J. (1995). Interferometric measurements of fine corneal topography. *Proceedings of SPIE*, *2329*, 32–39.

Kasprzak, H. T., & Licznerski, T. (1999). Influence of the characteristics of tear film break-up on the point spread function of an eye model. *Proceedings of SPIE*, *3820*, 390–396.

Keratoconus NZ. (2015). Sodium fluorescein. Retrieved January 6, 2015, from http://www.kcnz.co.nz/uploads/6/4/9/1/6491489/5496202_orig.jpg

Kiely, P. M., Smith, G., & Carney, L. G. (1982). The mean shape of the human cornea. *Journal of Modern Optics*, *29*(8), 1027–1040. doi:10.1080/713820960

Kiely, P. M., Smith, G., & Carney, L. G. (1984). Meridional variations of corneal shape. *American Journal of Optometry and Physiological Optics*, *610*(10), 619–626.

Kimbrough, B. T. (2006). Pixelated mask spatial carrier phase shifting interferometry algorithms and associated errors. *Applied Optics*, *45*(19), 4554–4562. doi:10.1364/AO.45.004554

King-Smith, P. E., Fink, B. a, Nichols, J. J., Nichols, K. K., & Hill, R. M. (2006). Interferometric imaging of the full thickness of the precorneal tear film. *JOSA A*, *23*(9), 2097–2104. doi:10.1364/JOSAA.23.002097

King-Smith, P. E., Fink, B. A., Hill, R. M., Koelling, K. W., & Tiffany, J. M. (2004). The thickness of the tear film. *Current Eye Research*, *29*(4-5), 357–368. doi:10.1080/02713680490516099

Kojima, T., Ishida, R., Dogru, M., Goto, E., Takano, Y., Matsumoto, Y., … Tsubota, K. (2004). A new noninvasive tear stability analysis system for the assessment of dry eyes. *Investigative Ophthalmology & Visual Science*, *45*(5), 1369–1374. doi:10.1167/iovs.03-0712

Kopf, M., Yi, F., Robert Iskander, D., Collins, M. J., Shaw, A. J., & Straker, B. (2008). Tear Film Surface Quality with Soft Contact Lenses Using Dynamic Videokeratoscopy. *Journal of Optometry*, *1*(1), 14–21. doi:10.3921/joptom.2008.14

Korpel, A. (1981). Acousto-Optics - A Review of Fundamentals. *Proceedings of the IEEE*, *69*(1), 48–53. doi:10.1109/PROC.1981.11919

Küchel, M. F. (2006). Absolute Measurement of Rotationally Symmetrical Aspheric Surfaces. *Frontiers in Optics*, OFTuB5. doi:10.1364/OFT.2006.OFTuB5

Kurita, H. (1989). *Interferometric aspheric surface testing using ray tracing code*. University of Arizona.

Kurita, H., Saito, K., & Kato, M. (1987). Influence of System Aberrations on Interferometric Aspheric Surface Testing. *Proceedings of SPIE*, *680*, 47–52.

Kwon, O., Wyant, J. C., & Hayslett, C. R. (1980). Rough surface interferometry at 10.6 micron. *Applied Optics*, *19*(11), 1862–9. doi:10.1364/AO.19.001862

Lam, P., Gaskill, J. D., & Wyant, J. C. (1984). Two-wavelength holographic interferometer. *Applied Optics*, *23*(18), 3079–3081. doi:10.1364/AO.23.003079

Lee, D. T., & Schachter, B. J. (1980). Two algorithms for constructing a Delaunay triangulation. *International Journal of Computer & Information Sciences*, *9*(3), 219–242. doi:10.1007/BF00977785

Lemp, M. A. (1995). Report of the National Eye Institute/Industry workshop on Clinical Trials in Dry Eyes. In *The CLAO Journal* (Vol. 21, pp. 221–232).

Levenberg, K. (1944). A Method for the Solution of Certain Non-Linear Problems in Least Squares. *Quarterly of Applied Mathematics*, *2*, 164–168.

Lim, H., Xu, W., & Huang, X. (1995). Two New Practical Methods for Phase Unwrapping. *Geoscience and Remote Sensing Symposium, 1995. IGARSS '95. "Quantitative Remote Sensing for Science and Applications", International*, *1*(65), 196–198.

Lowman, A. E., & Greivenkamp, J. E. (1985). Modeling an interferometer for non-null testing of aspheres. *Proceedings of SPIE*, *2536*, 139–147. doi:doi:10.1117/12.218416

Lowman, A. E., & Greivenkamp, J. E. (1994a). Interferometer induced wavefront errors when testing in a non-null configuration. *Proceedings of SPIE, 2004*, 173–181. doi:10.1117/12.172590

Lowman, A. E., & Greivenkamp, J. E. (1994b). Modeling of interferometric errors in non-null configurations. *Optical Fabrication and Testing*, *13*, 263–266.

Maguire, L. J., & Bourne, W. M. (1989). Corneal topography of early keratoconus. *American Journal of Ophthalmology*, *108*(2), 107–112. doi:10.1016/0002-9394(89)90001-9

Malacara, D. (2007). *Optical shop testing* (3rd ed.). Hoboken, NJ: John Wiley & Sons.

Malacara, D., Creath, K., Schmit, J., & Wyant, J. C. (2007). Testing of Aspheric Wavefronts and Surfaces. In D. Malacara (Ed.), *Optical Shop Testing* (Third., pp. 435–497). Wiley-Interscience.

Malacara, Z., & Malacara, D. (1995). Design of lenses to project the image of a pupil in optical testing interferometers. *Applied Optics*, *34*(4), 739–742.

Malacara, Z., & Servín, M. (2010). *Interferogram analysis for optical testing* (Second Edi., Vol. 84). Boca Raton, FL: CRC press.

Marquardt, D. W. (1963). An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *Journal of the Society for Industrial and Applied Mathematics*, *11*(2), 431–441. doi:10.1137/0111030

MathWorks. (2014). MATLAB. Natick, Massachusetts, United States.

McDonald, J. E. (1968). Surface phenomena of tear films. *Transactions of the American Ophthalmological Society*, *66*, 905–939.

Mengher, L. S., Bron, A. J., Tonge, S. R., & Gilbert, D. J. (1985). A non-invasive instrument for clinical assessment of the pre-corneal tear film stability. *Current Eye Research*, *4*(1), 1–7.

Mengher, L. S., Pandher, K. S., & Bron, A. J. (1986). Non-invasive tear film break-up time: sensitivity and specificity. *Acta Ophthalmologica*, *64*(4), 441–444. doi:10.1111/j.1755-3768.1986.tb06950.x

Micali, J. D., Greivenkamp, J. E., & Primeau, B. C. (2015). Dynamic measurement of the corneal tear film with a Twyman-Green interferometer. *Journal of Biomedical Optics*, *20*(5), 055007. doi:10.1117/1.JBO.20.5.055007

Millerd, J. E., Brock, N. J., Hayes, J. B., North-Morris, M., Novak, M., & Wyant, J. C. (2004). Pixelated Phase-Mask Dynamic Interferometer. *Proceedings of SPIE*, *5531*(520), 304–314. doi:10.1117/12.560807

Mishima, S. (1965). Some physiological aspects of the precorneal tear film. *Archives of Ophthalmology*, *73*(2), 233–241. doi:10.1001/archopht.1965.00970030235017

Murphy, P. E., Brown, T. G., & Moore, D. T. (1999). Optical vernier interferometry for aspheric metrology. *Proceedings of SPIE*, *3676*(3), 643–652.

Murphy, P. E., Brown, T. G., & Moore, D. T. (2000a). Interference imaging for aspheric surface testing. *Applied Optics*, *39*(13), 2122–2129. doi:10.1364/AO.39.002122

Murphy, P. E., Brown, T. G., & Moore, D. T. (2000b). Measurement and calibration of interferometric imaging aberrations. *Applied Optics*, *39*(34), 6421–6429. doi:10.1364/AO.39.006421

Murphy, P. E., Fleig, J., Forbes, G., Miladinovic, D., DeVries, G., & O'Donohue, S. (2006). Subaperture stitching interferometry for testing mild aspheres. *Proceedings of SPIE*, *6293*, 62930J. doi:10.1117/12.680473

Nachmias, J. (1961). Determiners of the drift of the eye during monocular fixation. *JOSA*, *51*(7), 761–766. doi:10.1364/JOSA.51.000761

Németh, J., Erdélyi, B., Csákány, B., Gáspár, P., Soumelidis, A., Kahlesz, F., & Lang, Z. (2002). High-Speed Videotopographic Measurement of Tear Film Build-up Time. *Investigative Ophthalmology & Visual Science*, *43*(6), 1783–1790.

Nichols, J. J. (2015). 2014 Annual Report - Contact Lenses 2014. *Contact Lens Spectrum*, *30*, 22–27.

Nichols, J. J., & King-Smith, P. E. (2003). Thickness of the Pre- and Post-Contact Lens Tear Film Measured In Vivo by Interferometry. *Investigative Ophthalmology & Visual Science*, *44*(1), 68–77. doi:10.1167/iovs.02-0377

Patel, S., Murray, D., McKenzie, A., Shearer, D. S., & McGrath, B. D. (1985). Effects of Fluorescein on Tear Breakup Time and on Tear Thinning Time. *American Journal of Optometry and Physiological Optics*, *62*(3), 188–190.

Patel, D. V, & McGhee, C. N. J. (2007). Contemporary in vivo confocal microscopy of the living human cornea using white light and laser scanning techniques: a major review. *Clinical & Experimental Ophthalmology*, *35*(1), 71–88. doi:10.1111/j.1442-9071.2007.01423.x

Powell, I. (2000). Employment of reverse optimization to relax manufacturing tolerances imposed on system constructional parameters associated with complex optical systems. *Applied Optics*, *39*(13), 2174–2183. doi:10.1364/AO.39.002174

Press, W. H. (2007). *Numerical Recipes 3rd Edition: The Art of Scientific Computing* (Third Edit.). Cambridge University Press.

Priest, D., & Munger, R. (1998). Comparative study of the elevation topography of complex shapes. *Journal of Cataract & Refractive Surgery*, *24*(6), 741–750. doi:10.1016/S0886-3350(98)80125-9

Primeau, B. C. (2011). *Interferometric Characterization of Tear Film Dynamics*.

Primeau, B. C., Goldstein, G. L., & Greivenkamp, J. E. (2012). Laser exposure analysis for a near-infrared ocular interferometer. *Optical Engineering*, *51*(6), 064301. doi:10.1117/1.OE

Primeau, B. C., & Greivenkamp, J. E. (2012a). Interferometer and analysis methods for the in vitro characterization of dynamic fluid layers on contact lenses. *Optical Engineering*, *51*(6), 063601. doi:10.1117/1.OE.51.6.063601

Primeau, B. C., & Greivenkamp, J. E. (2012b). Interferometer for measuring the dynamic surface topography of a human tear film. *Proceedings of SPIE*, *8215*, 821504. doi:10.1117/12.907083

Primeau, B. C., Greivenkamp, J. E., & Sullivan, J. J. (2011). In vitro interferometric characterization of dynamic fluid layers on contact lenses. *Proceedings of SPIE*, *8133*, 81330J. doi:10.1117/12.892830

Pritt, M. D. (1996). Phase Unwrapping by Means of Multigrid Techniques for Interferometric SAR. *IEEE Transactions on Geoscience and Remote Sensing*, *34*(3), 728–738.

Rabinowitz, Y. S. (1998). Keratoconus. *Survey of Ophthalmology*, *42*(4), 297–319. doi:10.1016/S0039-6257(97)00119-7

Ratliff, F., & Riggs, L. A. (1950). Involuntary motions of the eye during monocular fixation. *Journal of Experimental Psychology*, *40*(6), 687–701. doi:10.1037/h0057754

Rattle, J. D. (1969). Effect of target size on monocular fixation. *Optica Acta*, *16*(2), 183–192. doi:10.1080/713818162

Renka, R. J. (1984). Interpolation of data on the surface of a sphere. *ACM Transactions on Mathematical Software*, *10*(4), 417–436. doi:10.1145/2701.2703

Roberts, C. (1994). Characterization of the inherent error in a spherically-biased corneal topography system in mapping a radially aspheric surface. *Journal of Refractive & Corneal Surgery*, *10*(2), 103–111.

Roberts, C. J. (1994). Comparison of the EyeSys corneal analysis system and the TMS topographic modeling system using a bicurve test surface. *Proceedings of SPIE*, *2126*, 168–173. doi:10.1117/12.178582

Roberts, C. J., & Züger, B. J. (2006). *The Advantage and Principle of Dual Scheimpflug Imaging for Analyzing the Anterior Segment of the Human Eye*. *White Paper*. Retrieved from http://www.ziemergroup.com/fileadmin/media/products/GALILEI/Roberts__Galilei_2006.pdf

Rolfs, M., Engbert, R., & Kliegl, R. (2005). Crossmodal coupling of oculomotor control and spatial attention in vision and audition. *Experimental Brain Research*, *166*(3-4), 427–439. doi:10.1007/s00221-005-2382-y

Rottenkolber, M., & Podbielska, H. (1996). Coherent methods for measuring of ophthalmic surfaces. *Proceedings of SPIE*, *2628*, 157–164.

Rudder, S. (2006). Hybrid ECL/DBR wavelength and spectrum stabilized lasers demonstrate high power and narrow spectral linewidth. *Proceedings of SPIE*, *6101*, 61010I. doi:10.1117/12.646791

Schaumberg, D. A., Dana, R., Buring, J. E., & Sullivan, D. A. (2009). Prevalence of dry eye disease among US men: estimates from the Physicians' Health Studies. *Archives of Ophthalmology*, *127*(6), 763–768. doi:10.1001/archophthalmol.2009.103

Schaumberg, D. A., Sullivan, D. A., Buring, J. E., & Dana, M. R. (2003). Prevalence of dry eye syndrome among US women. *American Journal of Ophthalmology*, *136*(03), 318–326. doi:10.1016/S0002-9394(03)00218-6

Schein, O. D., Muñoz, B., Tielsch, J. M., Bandee-Rohe, K., & West, S. (1997). *Prevalence of Dry Eye Among the Elderly*. *American Journal of Ophthalmology* (Vol. 124). doi:10.1016/S0002-9394(14)71688-5

Schein, O. D., Tielsch, J. M., Muñoz, B., Bandeen-Roche, K., & West, S. (1997). Relation between signs and symptoms of dry eye in the elderly. A population-based perspective. *Ophthalmology*, *104*, 1395–1401. doi:10.1016/S0161-6420(97)30125-0

Schwiegerling, J., Greivenkamp, J. E., & Miller, J. M. (1995). Representation of videokeratoscopic height data with Zernike polynomials. *JOSA A*, *12*(10), 2105–2113. doi:10.1364/JOSAA.12.002105

Schwiegerling, J. T. (2004). *Field Guide to Visual and Ophthalmic Optics*. SPIE Press. doi:10.1097/01.opx.0000178357.14193.94

Schwiegerling, J. T., & Greivenkamp, J. E. (1997). Using corneal height maps and polynomial decomposition to determine corneal aberrations. *Optometry and Vision Science : Official Publication of the American Academy of Optometry*, *74*, 906–916. doi:10.1097/00006324-199711000-00024

Selberg, L. A. (1991). Interferometer accuracy and precision. *Proceedings of SPIE*, *1400*, 24–32.

Seong, K., & Greivenkamp, J. E. (2008). Surface figure measurement based on the transmitted wavefront with reverse raytracing. *Optical Engineering*, *47*(4), 43602. doi:10.1117/1.2904030

Sliney, D. H. (2010). *Preliminary Optical Hazard Evaluation - Johnson & Johnson Vision Care , Inc. Laser Tear-Film Interferometer*.

Smith, W. J. (2000). *Modern Optical Engineering* (Third Edit.). McGraw-Hill.

Sonka, M., Hlavac, V., & Boyle, R. (1993). Segmentation. In *Image Processing, Analysis, and Machine Vision* (First Edit., pp. 112–191). Springer.

Steinman, R. M. (1965). Effect of target size, luminance, and color on monocular fixation. *JOSA*, *55*(9), 1158–1165. doi:10.1364/JOSA.55.001158

Steinman, R. M., Haddad, G. M., Skavenski, A. A., & Wyman, D. (1973). Miniature eye movement. *Science*, *181*(4102), 810–819. doi:10.1126/science.181.4102.810

Stenstrom, S. (1948). Investigation of the variation and the correlation of the optical elements of human eyes. *American Journal of Optometry & Archives of American Academy of Optometry*, *25*(8), 218–504.

Stone, B. D. (2000). Modeling interferometers with lens design software. *Optical Engineering*, *39*(7), 1748–1759. doi:10.1117/1.602554

Stone, B. D., & Thompson, K. P. (2009). Modeling interferometers with lens design software: beyond ray-based approaches. *Proceedings of SPIE*, *7427*, 74270A. doi:10.1117/12.826665

Sullivan, J. J., & Greivenkamp, J. E. (2007). Design of partial nulls for testing of fast aspheric surfaces. *Proceedings of SPIE*, *6671*, 66710W. doi:10.1117/12.734874

Takeda, M. (1990). Spatial-carrier fringe-pattern analysis and its applications to precision interferometry and profilometry: An overview. *Industrial Metrology*, *1*(2), 79–99. doi:http://dx.doi.org/10.1016/0921-5956(90)80019-R

Takeda, M., Ina, H., & Kobayashi, S. (1982). Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry. *JOSA*, *72*(1), 156–160.

Tang, W., Collins, M. J., Carney, L., & Davis, B. (2000). The accuracy and precision performance of four videokeratoscopes in measuring test surfaces. *Optometry and Vision Science*, *77*(9), 483–491. doi:10.1097/00006324-200009000-00009

Tavakoli, M., Hossain, P., & Malik, R. a. (2008). Clinical applications of corneal confocal microscopy. *Clinical Ophthalmology*, *2*(2), 435–445. doi:10.2147/OPTH.S1490

Thaler, L., Schütz, A. C., Goodale, M. A., & Gegenfurtner, K. R. (2013). What is the best fixation target? The effect of target shape on stability of fixational eye movements. *Vision Research*, *76*, 31–42. doi:10.1016/j.visres.2012.10.012

Tighe, B. J. (2013). A Decade of Silicone Hydrogel Development: Surface Properties, Mechanical Properties, and Ocular Compatibility. *Eye & Contact Lens-Science and Clinical Practice*, *39*(1), 4–12. doi:10.1097/ICL.0b013e318275452b

Tripoli, N. K., Cohen, K. L., Holmgren, D. E., & Coggins, J. M. (1995). Assessment of radial aspheres by the Arc-step algorithm as implemented by the Keratron keratoscope. *American Journal of Ophthalmology*, *120*(5), 658–664. doi:10.1016/S0002-9394(14)72213-5

Wang, J., Shousha, M. A., Perez, V. L., Karp, C. L., Yoo, S. H., Shen, M., … Li, M. (2011). Ultra-high resolution optical coherence tomography for imaging the anterior segment of the eye. *Ophthalmic Surgery, Lasers & Imaging*, *42*(4), S15–S27. doi:10.3928/15428877-20110627-02

Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, *13*(4), 600–612. doi:10.1109/TIP.2003.819861

Wieman, C. E., & Hollberg, L. (1991). Using diode lasers for atomic physics. *Review of Scientific Instruments*, *62*(1), 1–20. doi:10.1063/1.1142305

Wikimedia Commons. (2009a). Computerised Corneal Topography. Retrieved July 10, 2015, from https://en.wikipedia.org/wiki/File:Computerised_Corneal_Topography.jpg

Wikimedia Commons. (2009b). Keratoconus. Retrieved January 6, 2015, from https://en.wikipedia.org/wiki/File:Keratoconus.svg

Williby, G. A., Smith, D. G., Brumfield, R. B., & Greivenkamp, J. E. (2003). Interferometric testing of soft contact lenses. *Proceedings of SPIE*, *5180*, 329–339.

Wyant, J. C. (1987). Interferometric Testing Of Aspheric Surfaces. *Proceedings of SPIE*, *0816*, 19–39.

Wyant, J. C., & Creath, K. (1992). Basic wavefront aberration theory for optical metrology. *Applied Optics and Optical Engineering*, *Xl*, 1–53.

Xu, W., & Cumming, I. (1999). A region-growing algorithm for InSAR phase unwrapping. *IEEE Transactions on Geoscience and Remote Sensing*, *37*(1), 124–134. doi:10.1109/36.739143

Xu, Y., & Ai, C. (2003). Simple and effective phase unwrapping technique. *Proceedings of SPIE*, *253*, 254–263.