ABSTRACT

Title of Thesis: IMPROVING NON-CONTACT TONOMETRY: A DEEP NEURAL NETWORK BASED METHOD FOR CORNEAL DEFORMATION MAPPING

Moshe Ackman, Lauren Cho, Kun Do, Aaron Green
Sam Klueter, Eliana Krakovsky, Jonathan Lin
Ross Locraft, James Muessig, Hongyi Wu

Glaucoma, a disease characterized by increased intraocular pressure (IOP), is one of the leading causes of preventable blindness worldwide. Accurate measurement of IOP is essential in monitoring glaucomatous progression in order to deliver treatment and prevent long-term vision loss. Currently, non-contact tonometry, known as an "air-puff test", is a common diagnostic method despite its inaccessibility, discomfort, high cost, and reliance on a trained professional. To improve upon these shortcomings, we designed a cheaper tonometer integrating a novel depth-mapping neural network with a custom air-puff generation system. We deformed porcine corneas with a controlled-intensity air-puff while imaging the deformation with a single stationary camera—a contrast to the standard Scheimpflug method. From the footage, our neural network predicted a three-dimensional map of corneal deformations. The network was able to predict a general negative trend between the IOP and the corneal deformation extracted. We compared our results to accepted literature deformation values and ground truth footage, allowing us to determine that the deformation amplitudes were physically plausible. With a more robust

imaging setup, we present a promising alternative to traditional IOP measurement methods. Future studies should make the simulated footage more representative of clinical conditions to increase the generalizability of the neural network. Additionally, anatomical differences between porcine and human eyes as well as corneal variability due to socio-demographic differences must be addressed for our results to be applied to clinical settings.

IMPROVING NON-CONTACT TONOMETRY: A DEEP NEURAL NETWORK BASED METHOD FOR CORNEAL DEFORMATION MAPPING

TEAM CONTACT

Authors:

Moshe Ackman, Lauren Cho, Kun Do, Aaron Green
Sam Kleuter, Eliana Krakovsky, Jonathan Lin
Ross Locraft, James Muessig, Hongyi Wu

Mentor: Giuliano Scarcelli

Thesis submitted in partial fulfillment of the requirements of the
Gemstone Honors Program, University of Maryland, 2021

Discussion Committee:
Associate Professor Giuliano Scarcelli, Chair/Advisor
Assistant Professor Amal Isaiah
Professor Peter Kofinas
Professor Yang Tao

Acknowledgments

We would like to thank our mentor, Dr. Scarcelli, for his guidance in our project; the graduate students in his lab, particularly Raymundo, for training us in lab skills and protocol; our librarian, Ms. Rachel Gammons, for helping us pursue and present research; and the Gemstone staff and faculty—Dr. Frank Coale, Dr. David Lovell, Dr. Kristan Skendall, Dr. Leah Kreimer Tobin, Dr. Vickie Hill, and Jessica Lee—for supporting our research endeavors.

Table of Contents

List of Figures

List of Abbreviations

BCE	Binary-Cross Entropy
CAE	Convolutional Autoencoder
CCT	Central Corneal Thickness
CLS	Contact Lens Sensor
CNN	Convolutional Neural Network
CST	Corneal Visualization Scheimpflug Tonometry
GAT	Goldmann Applanation Tonometer
GWF	Graphene Woven Fabric
IOP	Intraocular Pressure
OAG	Open Angle Glaucoma
POAG	Primary Open Angle Glaucoma
RNFL	Retinal Nerve Fiber Layer
SD-OCT	Spectral-Domain Optical Coherence Tomography
SES	Socioeconomic Status
ST	Scheimpflug Tonometry

Chapter 1: Introduction

Glaucoma is a chronic disease characterized by increased pressure within the eye and is the second leading cause of preventable blindness worldwide [14]. If diagnosed early, there exist treatments to slow the progression of glaucoma or lower the risk of developing it entirely, such as medication or surgery [15]. These treatments often involve consistent monitoring of intraocular pressure (IOP) as a way of tracking progression of the disease. Tonometers, devices that measure IOP, are traditionally only used in specialized healthcare settings. As a result, and despite the plethora of existing tonometer technologies, many people do not sufficiently monitor their glaucomatous progression nor take medication at a recommended frequency due to the prevalence of noncompliance among patients, especially of lower health literacy and education levels. It has been shown that there exist differences in the rate of glaucoma treatment compliance between socio-demographic groups in the United States. One factor that causes this discrepancy is inaccessibility to measurement devices [16] [17]. A health care provider requires special training to perform the examination, meaning that optometrists are often geographically distant from the communities that need the most help.

Many studies show that communities of color are disproportionately affected by glaucoma in both prevalence and lack of treatment. One study by the Baltimore Eye Research Group found that socioeconomic status (SES) was a significant factor in determining an individual's risk to developing glaucoma, and recent efforts distinguish prevalence of glaucoma between ethnic groups and communities with more granularity [16, 17]. In addition, studies show that type of insurance impacts the quality of glaucoma treatment, with Medicaid recipients 234% more likely not to receive any glaucoma testing within the first 15 months after diagnosis [18]. Because Medicaid often provides healthcare for those who are unable to access private insurance, disparities exist in treatment between communities of various SES. We sought to address this issue by developing a streamlined measurement tool that can be operated in small, general-practice clinics rather than requiring a system only found in specialized care facilities, with the goal of increasing the frequency of IOP examinations for the communities served by those clinics.

These differences in treatment are not exclusive to Medicaid or race, but are pervasive throughout the population and intricately related to many sociodemographic and institutional factors and barriers. Access to treatment and knowledge as well as immigration status, incarceration status, or any language barriers play a large impact on overall result of the disease, and it is crucial to understand the place of research and experimentation in light of these real issues. This contextualization is elaborated in our Racial Equity Impact Analysis in Appendix A, while the majority of the rest of this paper focuses on the technical aspects and outcomes of our experiment.

Current diagnostic tools can be broadly categorized into contact and non-contact tonometry. Contact tonometry, such as Goldmann Applanation Tonometry (GAT), requires the measurement device to make physical contact with the eye because the IOP is calculated from direct pressure measurements in this approach. Non-contact tonometry most commonly uses a puff of air to deform the cornea and measures the IOP from the biomechanical response. One limitation of the current standard tonometry techniques is that they only provide a measurement at a single moment, although IOP is known to fluctuate throughout the day [19]. This IOP fluctuation can be significant and is unique to each person, so many physicians recommend measuring IOP at certain times of the day, or at the same time of day for each measurement [20].

To ensure patients are regularly evaluated, so that any glaucomatic eye conditions can be detected early, a new diagnostic method will be proposed that will balance several considerations. First, the measurement must be precise and meaningful. In other terms, the result must be accurate enough to determine whether an individual is at elevated risk of developing glaucoma, with repeatability to add to the confidence in the diagnosis. Next, the devices must be inexpensive enough to be distributed in an accessible and equitable manner such that a broad population would be able to utilize it regularly. Third, the devices must be streamlined enough to be used with minimal training, making them more accessible to all populations. Altogether, these three goals lead us to our guiding research question.

Research Question

How can we leverage modern machine learning techniques to streamline and lower the cost of non-contact tonometry methods while maintaining precision and accuracy?

To answer our question, we developed a two-part system. The first part is the deformation and imaging of a cornea, and the second part is the analysis of those images. For our experiments, we used porcine eyes collected biweekly from a local butcher.

To deform the cornea, we deliver an air puff over a controlled time period similar to that of existing non-contact tonometers. We then record the cornea with a high frame rate camera while it is being deformed and collect a series of images. The images are cleaned and processed by two neural networks that output a three-dimensional corneal deformation map from the single view. By utilizing a neural network in connection with a camera, we can overcome the need for multiple or moving cameras in precise alignment, a limitation of current imaging techniques.

We will begin with a review of existing research on glaucoma monitoring and corneal imaging techniques, followed by a description of our methodology and experimental protocol. We then will analyze our experimental results and discuss the significance of our findings. Finally, we will examine the limitations of our study and explore future directions and applications for this research.

Chapter 2: Literature Review

Since the invention of the ophthalmoscope in 1850 by von Helmholtz and the invention of the Maklakoff applanation tonometer in the late nineteenth century, IOP has been known to vary between people. However, the first significant advancement towards glaucoma detection did not occur until 1950, with the creation of the GAT. Since then, advances in engineering and medicine have greatly improved glaucoma detection, which continues to spur on further research [21].

In the following sections, we describe notable advances and variations in tonometry. We divide existing tonometers into two classes: contact and non-contact tonometers. We discuss several variants of contact tonometers. Given the relative novelty of non-contact tonometry, we review the methods used to retrieve corneal movement without contacting the surface of the eye, which often rely on novel computer vision algorithms and other image processing techniques. Finally, we discuss clinical considerations and caveats in the measurement and monitoring of IOP.

2.1 Biomechanical Properties of the Cornea

Several biomechanical properties of the cornea have a significant effect on the diagnosis of glaucoma. Studies have found that the three most important properties are the intraocular pressure (IOP), corneal hysteresis, and central corneal thickness (CCT). Corneal hysteresis is a measure of the viscoelastic damping effect of the cornea [22]. A study by Medeiros et al. (2013) found that higher corneal hysteresis and lower IOP resulted in a lower rate of visual field loss over time. Mangouritsas et al. (2009) found that in patients with primary open-angle glaucoma, CCT had more of an effect on corneal hysteresis than in non-glaucomatous patients [23].

Corneal hysteresis is measured by an Ocular Response Analyzer. The operation of this device can be seen in Figure 2.1. Similar to non-contact tonometers, the Ocular Response Analyzer delivers a metered air puff to the eye over a period of roughly 20 ms. During the puff, the air pressure is gradually increased until the cornea is fully deformed, and then the air pressure is decreased. As the cornea deforms, it passes through two applanation points: once as the cornea deforms to concavity, and once more as it returns to its normal position. At these points, the reflected light is maximized. The Ocular Response Analyzer measures the peaks of the reflected light and records the air pressure at these points. The difference in pressure represents the corneal hysteresis.

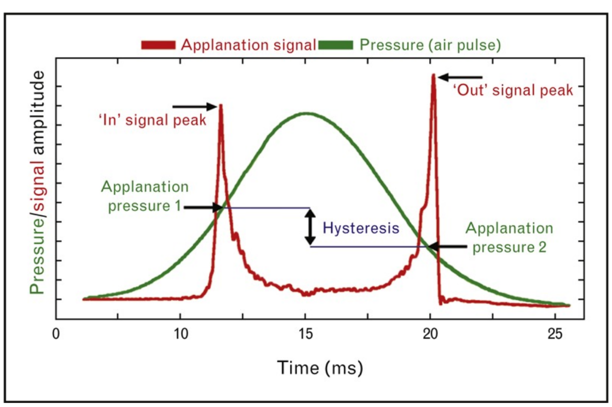

Figure 2.1: Operation of the Ocular Response Analyzer [1]

The gold standard for CCT measurement utilizes ultrasound pachymetry. It is important to know a patient's CCT when diagnosing glaucoma, as an abnormal CCT may affect the accuracy of a tonometry reading [24]. A study by Sadoughi et al. (2015) compared the precision of CCT measurements from ultrasound pachymetry and the Orbscan II, which utilizes scanning-slit topography. The study found that the two methods had good agreement for normal eyes (approximately 501-570 microns thick), but the Orbscan II tended to overestimate the CCT more than ultrasound pachymetry, and the discrepancy was even greater in thinner corneas. Furthermore, several studies have reported that a thinner CCT leads to underestimation of IOP measurement and has been identified as a good predictor of glaucomatous progression. It is unclear, however, if the increased risk of glaucoma is primarily due to greater IOP measurement error or if it is due to biomechanical differences [25].

When attempting to characterize these biomechanical properties, there are several confounding physiological factors that affect how they are assessed in both clinical and laboratory settings [26]. The stroma, which comprises over 90% of the cornea's thickness, is made up of a complex intertwining network of collagen fiber that provides structural support by inserting into an anterior membrane called the Bowman's layer [27]. These collagen structures vary substantially throughout the cornea. For example, the collagen across the central cornea in the mid to posterior stroma run in two orthogonal directions along the nasal-temporal and superior-inferior and provide high tensile strength to resist large ocular deformation and forces imposed during blinking [28]. The collagen in the peripheral cornea forms a circumferential arrangement and demonstrates higher circumferential strength, providing stability to the curvature of the central cornea and absorbing small changes in IOP [29]. Collagen crimp, a term used to describe the waviness of collagen fibers under tension, is also thought to contribute significantly to absorbing fluctuations in IOP [30]. While regional variability in collagen crimp has not been explored in detail, initial studies found that the limbus and peripheral cornea show higher waviness, tortuosity, and amplitude than the central cornea [31]. Like other biological tissue, the cornea is viscoelastic and hence exhibits hysteresis and creep during deformation [32]. In addition to its complex collagen orientation and structure, the cornea responds differently to different modes of corneal strain, demonstrating non-linear, hyperelastic stiffness behavior over a range of IOP [33]. Studies also suggest that age affects corneal stiffness, approximately doubling it from age 20 to 100 [34].

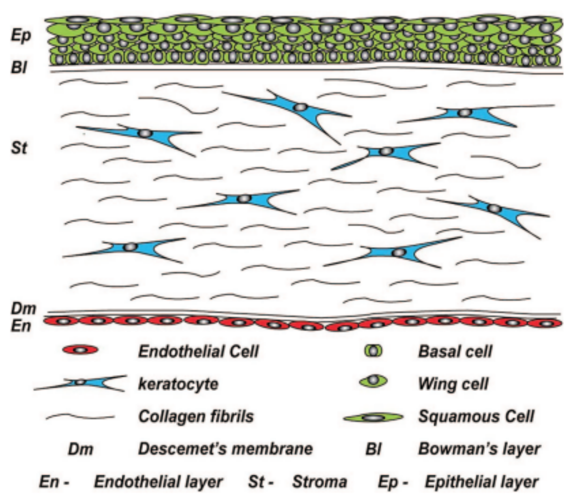

Figure 2.2: A cross section of the human cornea with labeled layers and components [2]

Studies have shown that the corneal endothelium is responsible for regulating corneal hydration to maintain high corneal transparency in fluctuating physiological environments [35]. There is significant variability in the hydration of the cornea both *in vivo* and in experiments performed with corneas *ex vivo*. Corneas that are analyzed *ex vivo* are unable to maintain physiological hydration levels, leading to loss of function in the endothelial pump and subsequent swelling [36]. Corneas *in vivo* also have variable hydration levels, with water content varying by 7.2% on average in healthy subjects throughout the day [37]. Factors including humidity, contact lens usage, and age all influence corneal hydration [38–40].

While important for properly diagnosing ocular diseases and facilitating effective, patient-specific treatment, quantification of corneal biomechanical properties poses a complicated challenge. Thus it is important to reconcile the laboratory techniques used *ex vivo*, the factors evidenced above that influence corneal variability, and the genetic and ethnic variability between patients when developing a reliable clinical method.

2.2 Contact Tonometry

2.2.1 Goldmann Applanation Tonometer

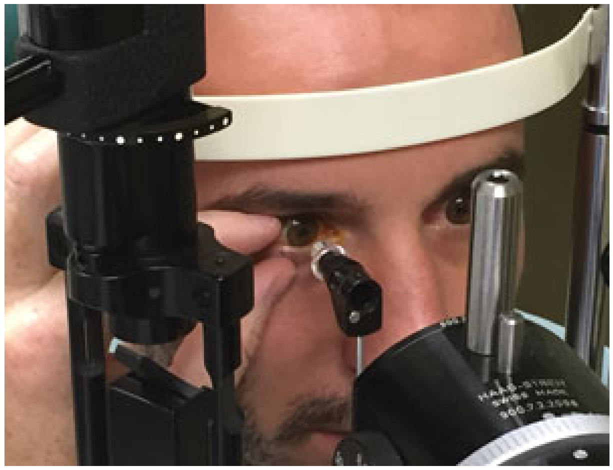

Figure 2.3: A patient getting their eye measured using the GAT [3].

The GAT quickly gained wide acceptance and became the "Gold Standard" for IOP measurements [21]. Typically, applanation tonometry requires both a small instrument and an ophthalmologist to support the eyelids of the patient [3]. In Figure 2.3, a patient's IOP is measured using the GAT.

The GAT measures the force required to applanate a given surface area of the cornea. When the bending resistance becomes equal to the surface tension, the applanation pressure is recorded and considered equal to the IOP based on the Imbert-Fick Law, expressed as

$$P = \frac{W}{A} \tag{2.1}$$

and shown visually in Figure 2.4.

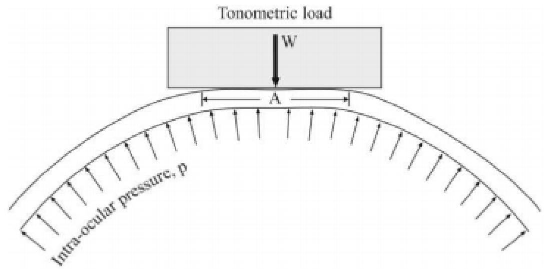

Figure 2.4: Schematic of the corneal deformation by contact from the GAT with influencing factors such as bending resistance of the cornea and surface tension from liquid tear drops depicted [3].

When a sufficient pressure is applied, the cornea will *applanate*, meaning the concavity of the cornea will flatten. It is only under this assumption that the Imbert-Fick law can be applied and (2.1) is used; otherwise, the effect of corneal thickness is no longer negligible so (2.1) is invalid. The pressure required to applanate the cornea is measured to determine the IOP [41].

Although robust, there is a risk of contamination from contact with the prism tip [42]. GAT has also been generally shown to significantly underestimate intracameral IOP, otherwise known as the IOP within the anterior chamber of the eye since there is a complex biomechanical response when the corneal surface is applanated [43].

Other drawbacks to GAT include the reliance on professionally trained operators. This means that IOP is almost never measured at night, when many researchers claim the IOP can be at its maximum. Additionally, only 2D cross-sectional images of the fluctuating IOP can be obtained. GAT is also complicated because it requires the administration of a topical anesthetic agent and sodium fluorescein on the corneal surface to prevent movement of the eye during measurement. Even if the method was used at night, sleep would be interrupted as GAT requires a specific head posture to determine the IOP of an individual [44].

Therefore, although GAT still remains the "gold standard" and foundation of tonometry, further directions of tonometry will need to be explored to mitigate the shortcomings of GAT such as lack of comfort, lack of portability, required expertise reliance, invasiveness, and measurement of a single-isolated value rather than continuous IOP monitoring.

2.2.2 Contact Lens Sensors

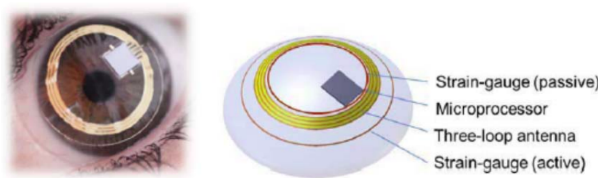

Figure 2.5: SENSIMED Triggerfish CLS in the eye, and components [4].

Not all newer methods for measuring IOP utilize non-contact tonometry. Contact lens sensors (CLS) represent a promising, less invasive technology to measure IOP, and often use piezoelectric, transparent, flexible sensor materials.

CLS methods were first approved in 2009 when the European Regulatory Authorities approved Triggerfish from SENSIMED, which was later approved by the United States Food and Drug Administration in 2016. Triggerfish is a disposable silicone contact lens with an embedded electromechanical sensor that can detect changes in the curvature of the cornea due to IOP variations, and is shown in Figure 2.5 [44].

The main problem that the Triggerfish CLS sought to address is the lack of continuous monitoring of IOP, as IOP can fluctuate widely over a day due to differences and variations in circadian rhythms, physiological postural changes, and blood pressure variations. Hence, the Triggerfish CLS has a 24-hour sensing mechanism to detect the fluctuations in IOP throughout the course of the day. CLS are promising because they can detect specific changes in the IOP through the course of time worn, when used in conjunction with software algorithms such as the B-spline smoothing transform or the Triggerfish software with an incorporated cosinor-based function. However, obtaining relevant biomechanical parameters is difficult, and is the main limitation for the practicability of this smoothing transform method [44].

In continuation, although Triggerfish CLS is a promising technique, it has several limitations. It is worn on the eye, so users may experience conjunctival hyperemia, which is an inflammation of the eye due to excess blood buildup in vessels, that causes vision obstruction, ocular pain, transient myopia and/or blurred vision. The major problem with CLS is that it not only causes these discomforts from long-term use, but can also alter IOP measurements. A study by Miki et al. in 2020 found that after 24 hours of wearing the CLS, participants' central cornea became steeper whilst the mid-peripheral cornea became flatter. The myopic spherical equivalent also increased by approximately 1 diopter [45]. Tojo et al. in 2019 discovered the corneal thickness increased and the corneal curvature was altered, which can impact IOP measurements, meaning that this method is limited in terms of accuracy [46]. Therefore, the Triggerfish CLS needs to be improved for larger practicability.

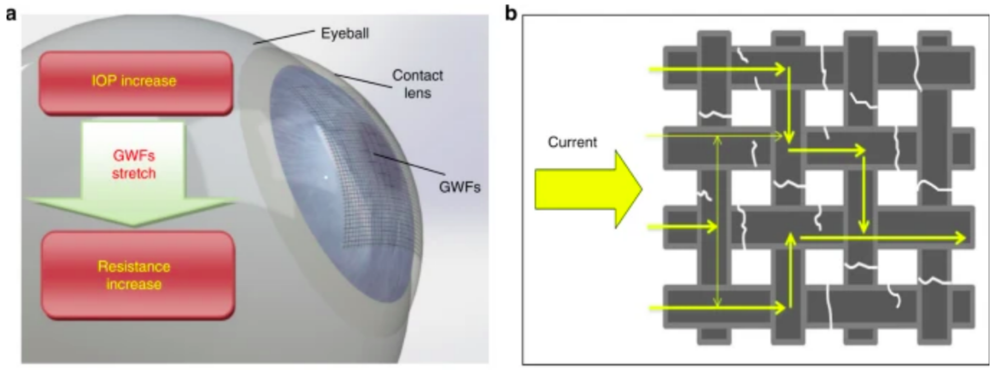

Figure 2.6: Schematic of a GWF bonded to a contact lens [5].

More recently, the integration of a graphene-woven fabric (GWF) into a contact lens has crated a high resolution, non-invasive tonometer that provides continuous monitoring [5]. GWFs are used in many other nano-optics and nano-electronic applications such as catalyst carriers, heat insulation materials, and radiation-resistant coatings. GWFs are electrically conductive, transparent, easy to handle, and have low fabrication costs, which are all valuable properties for this class of flexible photo-voltaic and composite materials [47].

The GWF is a graphene mesh tightly mounted on the cornea by homogeneous hydrophilicity that is very sensitive to changes in strain. When the IOP increases, the GWF stretches and creates micro-tears, increasing the electrical resistance of the material, or decreasing the current at a constant applied voltage by Ohm's law. The opposite also holds true, when the IOP decreases, the current will increase. The fluctuation in electrical current can be monitored by the device, and measure the effective change in IOP. The working principle of the GWF IOP measurement can be seen in Figure 2.6 [5].

This device has a device resistance change resolution of 6.8% per in the variation range of 0-10 . Previous contact lens tonometers have suffered from high production costs, but this sensor requires fewer components and is much cheaper. Finally, the GWF is roughly 80% transparent, so blocking light is of little concern [5]. Thus, the GWF CLS shows great promise in disadvantaged communities or countries, owing to its potential as a low-power, low-cost, high-resolution, disposable option.

However, the greatest limitation of this method is that it has to be calibrated before every use, and requires the aid of GAT for a baseline. This dependence on GAT means that GWFs are promising as a secondary method to detect IOP but cannot be reliable and stable as a standalone method. Also, factors such as corneal thickness and Young's modulus can vary between individuals, explaining why this method is better suited to determine IOP fluctuations per individual in a given environmental condition such as constant temperature. However, when the GWF is applied to different individuals, because the contact lenses and the piece of equipment varies, the requirement of the calibration makes the GWF not as robust as other methods, such as the "gold standard", GAT [5].

In conclusion, CLS methods are very promising as breakthroughs in electronics and micro-fabrication technology occur. The biggest benefit to CLS methods in comparison to other methods is the ability to continuously monitor IOP without disturbance to everyday activities, including sleep [44]. Further refinements and innovations are necessary, but it is an emerging technique and could be used in conjunction with other tonometers, such as the system we later propose in this paper.

2.2.2.1 Rebound Tonometry

Figure 2.7: The tonometer in use [6].

Rebound tonometry is a new type of contact tonometer that drops a magnetized, disposable probe towards the cornea and uses a solenoid to measure the deceleration of the probe. The device then employs the principle that higher IOP causes a quicker deceleration to calculate the IOP [48]. An example of this rebound tonometer is the iCare rebound tonometer, seen in Figure 2.7.

These tonometers are portable, easy to use, and are not discomforting for the patient, yet suffer from a high product cost. However, the iCARE rebound tonometer is very susceptible to changes in corneal hysteresis and corneal resistance factor. It also overestimated higher IOP and underestimated lower IOP when compared to GAT [49]. Rebound tonometry has also been shown to produce different results when the patient is sitting versus supine, and shows greater variation between different ages [50].

2.2.3 Non-Contact Tonometry

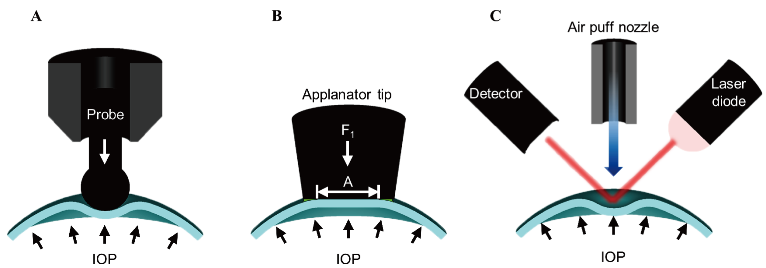

Figure 2.8: Schematics for the various methods for the applanation of the eye for corneal IOP measurements are shown. The workings of a rebound tonometer is shown in (A), of the GAT in (B), and the non-contact air puff tonometer in (C). From [7].

Contact-based methods of tonometry generally require sterile tips and anesthetic. As a less intrusive alternative to the GAT, non-contact methods have emerged as the clinical standard for tonometry. The most popular method of non-contact tonometry is the air-puff tonometer which uses bursts of air to apply pressure to the cornea. The diagram depicting the experimental principle of air puff tonometry is shown in Figure 2.8C.

When a sufficient pressure is applied, the cornea will applanate, meaning the concavity of the cornea will flatten [51]. The pressure required to applanate the cornea is measured to determine the IOP. Similar to GAT, the air-puff tonometer is susceptible to corneal variability. Additionally, the time scale of the air-puff tonometer is very short, which results in an increased susceptibility of measurements to momentary pressure variations [52]. One solution to this problem is to use an air puff of longer duration. Increasing the duration of the air puff is the method adopted by the commonly used Corvis Tonometer, which has been shown to agree well with GAT [53]. However, this method is somewhat aggressive and often induces blinking upon measurement. This blinking may itself affect the IOP and may be a significant source of error for air puff tonometry [54], which we discuss in Section 2.2.7.1.

2.2.3.1 Corvis and Air-Puff Tonometry

Figure 2.9: The OCULUS Corvis STL is a non-contact air-puff tonometer with a high-speed Scheimpflug camera [7].

One popular example of a non-contact tonometer is the Corneal Visualization Scheimpflug Technology (Corvis or CST) tonometer which applanates the cornea with an air-impulse [53, 55]. The Corvis then measures the deformation of the eye with an ultra-high speed Scheimpflug camera, capable of recording 4330 fps within a 100 ms time duration [53]. The Corvis instrument is shown in Figure 2.9.

The CST camera is able to detect excellent resolution on the order of 640 $\times$ 480 pixels, and replay eye movement in ultra-slow motion [53]. The device can measure pressures between 1 and 60 mmHg, making it an excellent candidate as an instrument to measure IOP, which typically ranges between 12 and 22 mmHg [56] in a normal eye and causes rapid vision loss above 40-50 mmHg [57].

Although the theory behind the CST is not as easily understood in relation to other commonly used tonometry techniques like GAT and rebound tonometry, it can also measure biomechanical properties in the eye such as corneal thickness. Additionally, the CST had the best intraobserver and interobserver variability in measurements of IOP, and was not as affected by the biomechanical properties in terms of measurement variability [53].

The deformation time from the applanation of the cornea largely influences the precision and decreased variability of the IOP measurements. For instance, many non-contact tonometers measure the IOP in less than 5 ms, and the CST puff measurements are separated by an average of 15 ms. The study by Hong et al. found that applanation time plays a significant factor in the variability of the calculated "read" IOP values yet the difference between the differences in IOP and the time it took to applanate the cornea was statistically marginal ($p = 0.067$) [53]. However, more research is needed to determine the effect of the applanation time on the accuracy of the instrument IOP readings.

2.2.4 Imaging Modalities in Ophthalmology

Increased IOP limits blood flow, causing degeneration in retinal ganglion cells. The death of these cells causes enlargement of the optical disk, or optical nerve head [58] and a thinning of both the retinal nerve fiber layer (RNFL) and macula. Since glaucoma is a progressive disease, it is vital to obtain qualitative information about these structural changes to best intervene in glaucoma progression. Additionally, these structural pathologies are evident well before vision loss [59, 60]. Three well-developed *in vivo* imaging techniques to achieve this are Scheimpflug imaging, confocal scanning laser ophthalmoscopy, scanning laser polarimetry, and Spectral-Domain optical coherence tomography (SD-OCT). These methods significantly vary on how they obtain data for monitoring purposes, but all have shown promising results for monitoring structural changes due to glaucoma. This section describes each method, names the devices that use it, and discusses its advantages and disadvantages.

2.2.4.1 Corneal Topographers

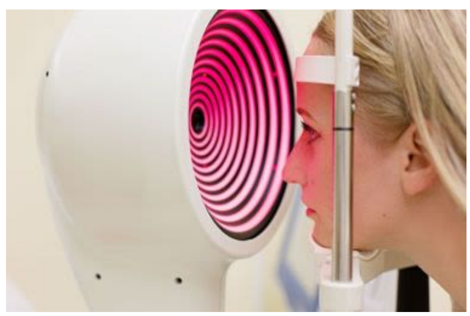

Figure 2.10: A patient being evaluated with the Placido topographer [8].

Most corneal topographers are based on the Placido topographer, which was introduced in 1880 by Antonio Placido and further developed by Allvar Gullstrand in 1896, who developed a numerical algorithm to accompany the method [61]. A set of high-contrast concentric rings are placed in front of the cornea (Figure 2.10), creating contour-like reflections whose displacement may be analyzed to calculate corneal wavefront parameters like astigmatism, trefoil, and coma [62]. More recent variations on the method sometimes employ other patterns and methods of projection. The Orbscan topographer, for example, uses a "scanning slit" pattern projected from multiple angles, which creates 240 reference points on each slit for analysis [63]. In general, reflection-based methods have been found to have excellent accuracy and repeatability [64–66].

Originally, the contours produced by such topographers were evaluated only qualitatively, in comparison to standardized images of "normal" corneas [67]. However, as computing power increased and the implementation of increasingly complex image processing algorithms became feasible, it became possible to produce quantitative intensity maps of corneal power using algorithms that detected and processed variations in some or all of the reflected rings (see Section 2.2.7.1 for a discussion of these algorithms).

More recently, corneal topographers have begun using the Scheimpflug imaging technique to capture corneal thickness and other corneal properties that are only accessible from a side perspective. The most well-known of these topographers is the OCULUS Pentacam [68] and the Ziemer Galilei [69]. In the following section, we describe the principles behind Scheimpflug imaging.

2.2.4.2 Scheimpflug Imaging

Scheimpflug imaging relies on the Scheimpflug principle. A photograph of an object plane that is not parallel to the image plane can be maximally focused given "certain angular relations among the object, the lens, and the image plane" [69]. By being able to vary the plane of focus by varying the camera angle, the composition of many images is able to produce a reliable image from the anterior cornea to the posterior lens surface.

Modern Scheimpflug corneal imaging systems use rotating cameras to image cross sections of the cornea that are illuminated by slit beams at different angles. These images are then interpreted by an algorithm that reconstructs a 3-D tomographic map of the cornea's anterior segment and also calculates other physical information such as the topographic corneal thickness and the corneal curvature [70].

The Pentacam (Oculus Optikgeraete GmbH, Wetzlar, Germany) uses a single rotating Scheimpflug camera illuminated by a 470 nm blue LED to capture images of the anterior segment, while a second stationary camera is used to detect and measure the pupil, which helps with aligning the images and compensates for ocular movement. In this device, Scheimpflug imaging takes approximately 2 seconds to image 25 to 50 different cross sections between 0 and 180 degrees in a single scan, accumulating approximately 25,000 data points to be used for reconstruction of the anterior segment using its proprietary software [71].

The newer Galilei (Ziemer Ophthalmic Systems AG, Port, Switzerland) dual Scheimpflug analyzer combines corneal tomography from Scheimpflug imaging and Placido disc topography to also obtain a comprehensive curvature profile of the central cornea [72].

2.2.4.3 Optical Coherence Tomography

Optical coherence tomography (OCT) is a non-contact, non-invasive method that uses low coherence interferometry to obtain high-resolution tomographic images of biological tissues [73]. Since (RNFL) health is a strong indicator of glaucomatous damage [74], OCT is frequently used to obtain measurable information about retinal layer thickness. The earliest variant, known as Time Domain-OCT, relied on the delay time from back-scattered light to construct a 2D image of the optic nerve. Time Domain-OCT provided an axial resolution of 10 μm and takes about two seconds to provide a 2D image of the optic nerve [75]. Recent advancements have led to SD-OCT, allowing for 3D imaging of the optic nerve [76] with an axial resolution of 5 μm. Currently, the standard device in clinical settings is Spatial Domain-OCT). There remains a significant difference in measurements taken by different SD-OCT devices. For example, one commonly used brand, the RTVue, gives a much higher peripapillary RNFL thickness than other brands [77]. Since comparing results over time is essential to make a proper treatment plan, switching imaging devices or going to another practice can lead to misdiagnosis. Additionally, cataracts and vitreous debris can lead to a false reading of an increase in RNFL thickness. Meanwhile, myelinated RNFL, epiretinal membrane, swelling of optical nerve head, and peripapillary retina can lead to false increase readings of the retinal nerve fiber thickness [78].

In summary, OCT is a valuable method for monitoring existing ocular damage and providing information to clinicians to help prevent further damage. Different pathologies can reduce the accuracy of OCT measurements. Additionally, different SD-OCT brands do not have a consistent baseline. With this in mind, OCT measurements are most valuable when used alongside other diagnostic methods such as visual field tests and other evaluation methods.

2.2.5 Computer Vision and Ophthalmology

2.2.5.1 Depth Estimation

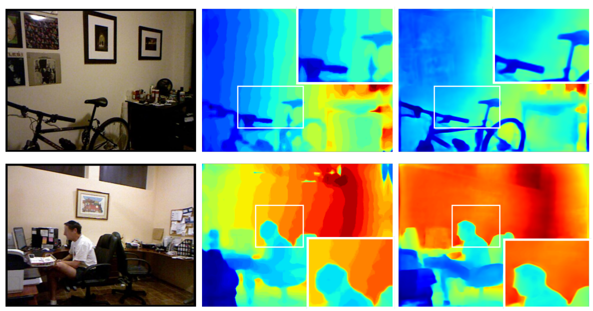

Figure 2.11: An example of a scene and its corresponding depth map. This depth map was creating using an algorithm for monocular depth estimation, AdaBins [9].

The practice of estimating corneal topography from camera images may be reduced to the classic computer vision problem of depth estimation. Solutions to the problem attempt to label each pixel of an input image with a value representing its distance from the camera; a completely labeled image is known as a depth map. To date, devices like Corvis or the Pentacam rely on moving camera technology and the composition of many separate images to solve this problem. Indeed, traditional solutions in computer vision have focused on the use of computer stereo vision to most accurately regress 3D geometry from images. Broadly, these approaches rely on feature extraction and matching using engineered features like SIFT or SURF

[79, 80], before applying calculations based on feature movement and parallax to calculate object depth [81,82]. The defining feature of these approaches is that they require multiple images or motion of the camera. Consequently, existing methods for depth estimation in ophthalmology require additional equipment to move the camera and acquire the necessary data.

However, more recent advances in computer vision have focused on a different solution to the problem, monocular depth estimation. As its name suggests, monocular depth estimation seeks to solve the depth regression problem using only a single input image. Although a monocular perspective introduces unavoidable scale ambiguities to the solution space, with low-resolution supplementary information and correct assumptions about geometry, it is possible to accurately solve the depth estimation problem [83–85]. This would eliminate the need for moving parts and multiple images currently demanded by computer stereo vision and other multi-image algorithms.

Monocular depth estimation has attracted particular attention since the advent of convolutional neural networks (CNNs). There have been a myriad of network architectures created to solve the problem, with increasingly promising performance and generalization capabilities [86–89]. Most approaches are based on the convolutional autoencoder (CAE) architecture, in which a CNN learns to extract specific features from input images. CAEs have also shown promise in image denoising [90–93], medical image segmentation [94–96], and feature extraction [97–99]. We describe the principles of neural networks and their current applications in ophthalmology in Section 2.2.6.

2.2.5.2 Structured Light and Ophthalmology

These algorithms are designed to work on a wide variety of scenes and perspectives. However, they become most effective when given access to supplementary information. An especially notable example is the structured light approach to depth estimation, which allows for precise reconstruction of an object's topography [100–102]. This approach projects a fixed, regular pattern onto an object. The distortions caused by the object's geometry are then used to compute highly precise representations of its surface. Aside from engineered algorithms, highly accurate reconstructions have also been produced by CAEs, with the added benefit that preprocessing and denoising steps taken by other algorithms are learned implicitly by the network. In addition, because CAEs mainly rely on the convolution operation, they are easily parallelizable, allowing for greater optimization and faster execution times compared to engineered algorithms [103].

Structured light-like algorithms that appear in ophthalmology are typically known as videokeratography. The Placido and Pentacam topographers may be considered examples of monocular structured light reconstruction, though neither method takes advantage of deep learning. A number of solutions have been developed to automate and enable quantitative analysis of the images generated by Placido-style topographers. Two prominent examples of current computer vision techniques in the field are the SimK and CorT algorithms. However, the SimK algorithm does not make use of the information from all Placido rings. Instead, it uses only one ring within the series to generate its estimate. The newer CorT algorithm takes advantage of the distortion of all visible Placido rings and has been shown to achieve greater accuracy as a result [104, 105].

2.2.6 Convolutional Neural Networks

In this section, we describe the basic mathematical structure of neural networks, beginning with the multilayer perceptron before moving to CNNs. We also describe the process by which these networks are trained, under the paradigm of supervised learning. We conclude by summarizing current applications of CNNs in ophthalmology.

2.2.6.1 Design and Architecture

Neural networks have undergone extensive modifications. Although their full history is beyond the scope of this review, we begin by describing the quintessential neural network, the multilayer perceptron (MLP), which motivates many of the concepts underlying modern neural networks.

Artificial neural networks use mathematical abstractions of biological neurons. An MLP is an extension of the perceptron, introduced by Rosenblatt in 1958 [106]. Given a set of inputs $\mathbf{x} = x_1, \ldots, x_n$, the perceptron outputs either 0 or 1. The output is determined using a "threshold" function on a linear combination of the elements in $\mathbf{x}$. Mathematically, the output y of a single neuron within the perceptron may be represented as

$$y = (\mathbf{w} \cdot \mathbf{x} + b) = (w_1 x_1 + \ldots + w_n x_n + b) = \left(\sum_{i=1}^{n} w_i x_i + b \right) \tag{2.2}$$

where b is some bias term learned by the perceptron and $\mathbf{w}$ is a set of learned weights.

The perceptron is an example of a linear classifier, which creates a hyperplane to classify its inputs. However, this creates an issue as few training sets are linearly separable. For instance, the binary XOR function is not linearly separable and therefore cannot be learned by a single perceptron, as famously shown by Minsky and Papert in their 1990 book *Perceptrons*. Although it is possible to learn such a function by stacking "layers" of perceptrons in series, such a system is difficult to train using Rosenblatt's learning rule, and relatively inefficient [107].

Modern perceptron neurons circumvent this issue by applying a nonpolynomial differentiable activation function σ to the output of each of their neurons. In fact, the modern perceptron neuron is nearly identical to that described by Rosenblatt in Equation (2.2), except that the function is replaced by σ:

$$y = \sigma(\mathbf{w} \cdot \mathbf{x} + b) = \sigma(w_1 x_1 + \ldots + w_n x_n + b) = \sigma \left(\sum_{i=1}^{n} w_i x_i + b \right) \tag{2.3}$$

The resulting neurons can learn non-linearly separable data on their own, may be stacked, and may learn using the much more powerful backpropagation algorithm because their output is fully differentiable, which is discussed in Section

2.2.6.2. When stacked in layers, we obtain the modern multilayer perceptron, which may approximate any continuous function to an arbitrary degree of accuracy given enough nodes using a nonpolynomial activation function [108].

See Figure 2.12 for a graphical representation of a MLP, where each circular node represents a single perceptron as described in Equation (2.3). MLPs have been used for a broad variety of tasks in classification and regression [109, 110].

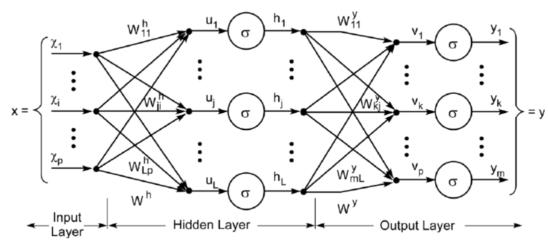

Figure 2.12: An example of a multilayer perceptron with hidden layers. Note that layers comprising multiple neurons allow for vector outputs [10].

Most neural network architectures employ the sigmoid, hyperbolic tangent (tanh), softplus or rectified linear unit (ReLU) functions as activation functions. The first three functions are more common in sequence-based or vector-based input networks. The ReLU function has found usage in modern deep learning networks, which can have up to thousands of layers and encounter issues when trying to backpropagate through the other functions [111].

From Perceptrons to Convolutions

Although MLPs can approximate any continuous function, as discussed above, in practice this often requires an impractical number of nodes and can be very difficult to train. In addition, the “fully-connected” architecture of an MLP does not lend itself naturally to images, as they often contain locally-dense features which may appear at any point within the image at different angles and scales [112, 113].

In 1989, Yann Lecun et al. introduced CNNs to process image data using less parameters and greater emphasis on spatially invariant local patterns, showing their effectiveness in recognizing handwritten digits [114]. The advent of graphical processing units (GPUs) and the discovery that deep CNNs with many layers had great representative power introduced the deep CNN, which has become commonplace in image processing tasks [115] and achieved superhuman performance on computer vision tasks [116].

CNNs are based on the convolution operation. This operation "slides" a small convolution kernel over the input image, using it to take a weighted sum of each pixel's neighbors. Let a given image of height H pixels and width W pixels be represented as a 2D matrix I with $H \times W$ elements, and a given convolution kernel be represented as a 2D matrix K with R rows and C columns (assume R and C are odd). Then mathematically, the convolution operation may be represented as

$$I'_{x,y} = \sum_{i,j} K_{i,j} I_{x+i-\frac{R-1}{2},y+j-\frac{C-1}{2}} \tag{2.4}$$

where (x, y) index I, (i, j) zero-index K, and I' is the resulting convolved image. One caveat is that there are not enough "neighbors" at the edge of I to use in the convolution operation. This is often addressed either by padding the image with zeros or some other constant values, image reflection, "wrapping", or cropping pixels that do not have enough neighbors. However, we omit the specifics of these workarounds in our formulation.

Figure 2.13: Sample application of the Sobel convolution kernel on an image. Original image is on left, convolved result is on right.

As an example, Figure 2.13 shows the results of the Sobel convolution kernel, an engineered kernel created by Irwin Sobel to detect high-frequency features (like edges) in an image. The Sobel kernels have the entries

$$G_x = \begin{pmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{pmatrix} \qquad G_y = \begin{pmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{pmatrix} \tag{2.5}$$

which are analogous to the discrete derivatives of the image intensity [117].

CNNs, however, begin with randomly initialized convolution kernels and learn to update their entries to extract the most salient features in an image—in effect, they are automated feature extractors. By applying convolutions in series, a CNN can learn to recognize higher and higher-level features [115]. Figure n shows an example CNN architecture.

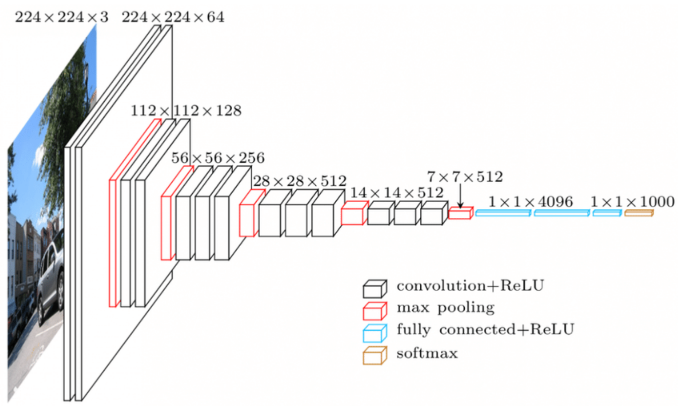

Figure 2.14: An example of a CNN used to classify images, VGG16. Note the stacked convolution maps, "pyramid" structure, and high number of feature maps [11].

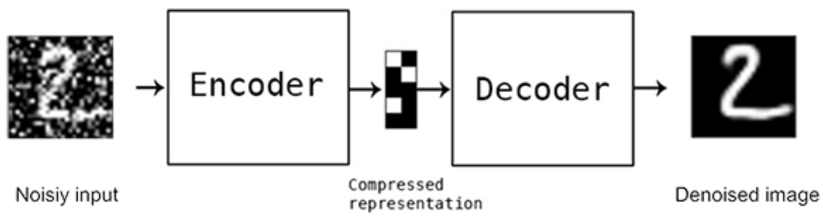

Figure 2.15: High-level structure of an autoencoder used for denoising handwritten digits [12].

Convolutional Autoencoders

CAEs are a special subset of CNNs. However, every autoencoder performs dimensionality reduction on input data, before learning to reconstruct it in some way. In the case of images, this often refers to the operation of subsampling, in which an image is reduced in size by a constant factor. The autoencoder is then forced to reconstruct the original input image as completely as possible using only its compressed representation. This makes autoencoders extremely useful for feature extraction, as the network must learn to preserve the most relevant features of the image [118].

Autoencoders can also be used for other purposes. Figure 2.15 shows an example of an autoencoder used for denoising. Once a compressed representation has been learned by the encoder, it is also possible to train a decoder to reconstruct this representation in such a way that features in the input image are omitted or modified. This allows autoencoders to output denoised versions of input images, segmentation masks, and depth maps as discussed in Section 2.2.5.1.

2.2.6.2 Supervised Learning and Optimization

In this section, we briefly describe the core principles of supervised learning and backpropagation, which we use to train our network. Although we have described the mechanism by which neural networks output their predictions, we have not described how they learn to adjust their weights. In supervised learning, examples are annotated with ground truth labels to help guide predictions. Outputs from the network are then compared with these labels to determine how the weights of the network should be adjusted.

Loss Functions

Loss functions are the method by which labels are compared to network outputs. Let a network be f, such that for some $\mathbf{x}$ in the training set we have $f(\mathbf{x}) = \hat{\mathbf{y}}$, and our ground truth label be $\mathbf{y}$. Then our loss function $L(\mathbf{y}, \hat{\mathbf{y}})$ outputs some value, the magnitude of which represents the loss of the network output. To enable backpropagation, L must be differentiable. The mean squared error and cross-entropy functions are the two most common loss functions used in supervised learning [119].

Backpropagation

Once the loss function has output a loss value, the network's parameters are adjusted to reduce it. Backpropagation is the most popular method to compute the necessary adjustments [120].

We may view a feedforward neural network as the composition of a series of differentiable processes. Because the loss function is itself differentiable, we can calculate the partial derivative of any weight in a network with respect to its output. In other words, given some $\mathbf{y}, \hat{\mathbf{y}}$ as described in the previous section, let the output of the loss function be $C = L(\mathbf{y}, \hat{\mathbf{y}})$. Then our goal is to find $\frac{\partial C}{\partial W}$ for every weight W in the network.

To calculate this quantity, we employ backpropagation. This algorithm calculates the derivatives of each layer, iterating backwards from the output. Intuitively, it takes advantage of the fact that because a feedforward neural network is a composition of functions, one may iteratively apply the chain rule to continuously differentiate its output.

Once this gradient has been computed, different gradient descent algorithms and optimizers are used to decide the exact calculation used to adjust each weight. Examples of these algorithms include stochastic gradient descent, Adagrad, Adadelta, or Adam [121].

2.2.6.3 Applications in Ophthalmology

Prior applications of CAEs and neural networks in ophthalmology seem to focus on classifying corneal diseases like keratoconus based on existing corneal topography data [122, 123], or detecting the presence of disease from images of the retinal fundus [96].

Lavric and Valentin classified corneal topography maps in 2019, searching for keratoconus, and achieved a test accuracy of 99.33%. Similarly, Zéboulon et al. classified corneal topography measurements with keratoconus and histories of refractive surgery, achieved a similar accuracy of 99.3% ($n = 300$), and concluded that "using combined corneal topography raw data with a CNN is an effective way to classify examinations and probably the most thorough way to automatically analyze corneal topography" [122, 123]. These results for keratoconus demonstrate promise in extending CAE and similar techniques to other ophthalmic diseases such as glaucoma.

2.2.7 Clinical Considerations and Caveats

2.2.7.1 Effect of Blinking on IOP

The increase in IOP due to blinking, both voluntarily and involuntarily, has been observed in humans. Pressure measurements taken immediately after blinking show an increase of 10 mmHg for normal, unforced blinking, while pressures can be much higher (on the order of 90 mmHg) for "squeezing," or forceful, conscious closing of the eyelids [124].

These increases represent momentary spikes in the IOP lasting around 60 seconds [125]. Therefore, the effect of blinking-induced pressure increases is substantial for fast measurements lasting less than a minute, such as those involved in standard air-puff tonometry. In the case of air-puff tonometry, the effect is especially significant since the air burst tends to cause forceful blinking. Therefore, when measuring IOP by air-puff tonometry, the measurement is increased by subsequent failed attempts marred by blinking.

In order to reduce error caused by blinking, two possible avenues may be pursued. One is continuous monitoring of IOP. If the IOP is continuously monitored, blinking can be observed as a pulse in pressure. The baseline pressure before and after blinking may still be observed [125, 126].

Another solution is to decrease the occurrence of blinking during measurement. Since blinking occurs as a response to anticipated trauma from a forceful, large air puff, blinking can be largely prevented by decreasing the size and forcefulness of the air puff.

Because blinking can produce a large pressure increase without the application of force by an outside device, this was considered as a potential research avenue. The movement of the cornea could be measured by a camera after a patient blinks. Ultimately, this direction was not pursued in this project because of the complications caused by potential anatomical and behavioral differences. The pressure increase due to blinking depends both on the forcefulness of the blink and the physical properties of the eye and eyelid. Therefore, standardization of the pressure applied by blinking across different patients is difficult.

Chapter 3: Methodology

In this section, we describe our experimental setup and the architecture and training process for our proposed algorithm. Our setup comprised of an applanation system and an imaging system, which delivered air puffs and imaged the resulting deformation respectively (Figure 3.1). We name the components used in the setup, describe our design process, and elaborate upon our proposed calibration method for our tonometer.

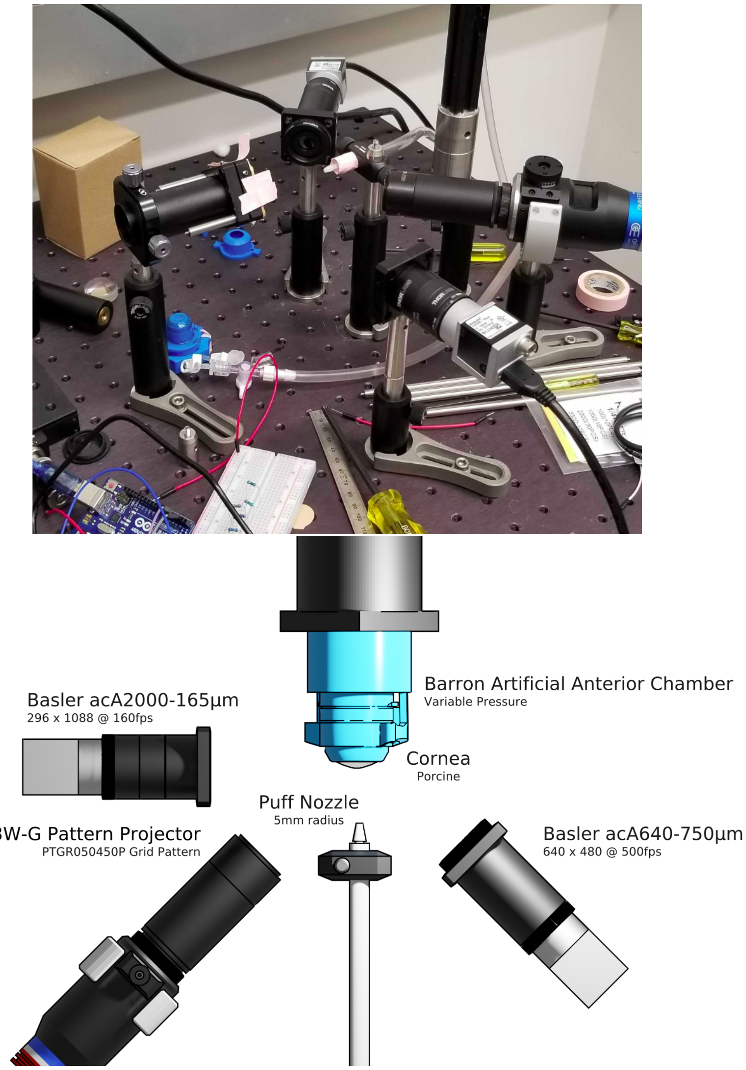

Figure 3.1: A photo (top) and schematic (bottom) of our experimental setup.

3.1 Applanation System

The applanation system consisted of a mechanical system which allowed for air flow from a pressurized air source through a nozzle, generating puffs whose intensities and time intervals are controlled by an electronic microcontroller. These puffs then applanated a mounted porcine cornea.

3.1.1 Mechanical Design

To deliver an air puff with consistent pressure, we used a compressed air tank as a gas source with a single stage regulator to control the output pressure. The air flowed through a length of polyurethane tubing. Between puffs, a three-way, two-position solenoid valve exhausts the gas away from the eye. An analog pressure gauge was situated between the nozzle and the solenoid valve to provide a quantitative readout during this time. The tubing terminated at a 2 mm plastic nozzle to concentrate the puff onto a small area of the cornea.

Iteration 0: Mechanical Pumps for Puff Generation

Initially, we investigated the use of mechanical pumps to generate an air puff. The first pump was the Diaphragm Electric Operated Positive Displacement Pump from RS Pro, acquired from Allied Electronics Automation. This pump has a flow rate of 150 mL/min at 1 bar, with an operating voltage range of 1.5 to 4.5 V. The next pump we used was the Vacuum Pump from SparkFun. It has a flow rate between 9 and 15 L/min at 12 V.

Neither of these pumps were able to provide adequate pressure to deform the corneas tested in our setup. As a result, this design was not pursued further.

Iteration 1: Compressed Air Source

Rather than purchase a large pump to generate puffs, we decided to try using a compressed air source. We acquired a tank of industrial-grade compressed nitrogen gas from Airgas. This tank had a two-stage pressure regulator which controlled the gas pressure released. This method generated enough pressure to deform a cornea.

Iteration 2: Two-Valve System

Our first system used two NITRA two-way, two-position solenoid valves. The solenoid valves actuate at 15 A, so an external power supply was required. The valves consisted of a magnet attached to a diaphragm inside a solenoid. Below the diaphragm was a channel to allow air flow. In a normally-open valve, the resting position of the magnet is up, allowing air flow through the valve. When the solenoid actuates, the magnet is forced down such that the diaphragm blocks the passage of air through the valve. In a normally-closed valve, the magnet is in the low position when no current is applied, blocking air flow through the valve. The solenoid valves we acquired had a 12-14 ms complete cycle response time, which denotes the time duration required to go from the rest state to the active state, then back to the rest state. This cycle time enabled us to create puffs with a similar duration to the industry standard Corvis, which has a puff duration of roughly 15 ms.

Two valves were connected to the air source in series with flexible tubing. In this configuration, the input valve, or the valve closer to the air source, was opened, allowing the space between the two valves to pressurize. When the appropriate pressure was reached, the input valve closed and the output valve opened. This opening caused air to flow through the nozzle, puffing the cornea.

Iteration 3: One-Valve System

When the two-valve system was tested, the pressure between the two valves could not be precisely measured and controlled. As a result, this chamber often over-pressurized and the fittings disconnected. To mitigate this, we changed our system to only use one NITRA three-way, two-position valve.

The three-way valve has one input and two outputs. In the ground state, the valve connects the input to an exhaust output, which vents the pressurized gas to the atmosphere. When power is supplied to the valve, it changes the output such that the compressed air is sent through the tubing length and out the nozzle. This valve allowed us to control the intensity of the puff by adjusting the pressure of the air source, and control the puff duration by changing the amount of time that the relay is energized.

Final Design: Fume Hood Air Source

By our third iteration significantly more gas was used, causing the nitrogen tank to deplete quickly. To avoid any delays related to replacing the tank, we changed our air source to a pressurized fume hood air line.

3.1.2 Electrical Design

A normally-open electromagnetic relay was used to operate the solenoid valves. The relay was connected to a DC power supply with variable voltage and current modes, and to one of the digital pins on an Arduino Uno microcontroller. The relay output was connected to the solenoid valve. A schematic diagram of the configuration can be seen in Figure 3.2.

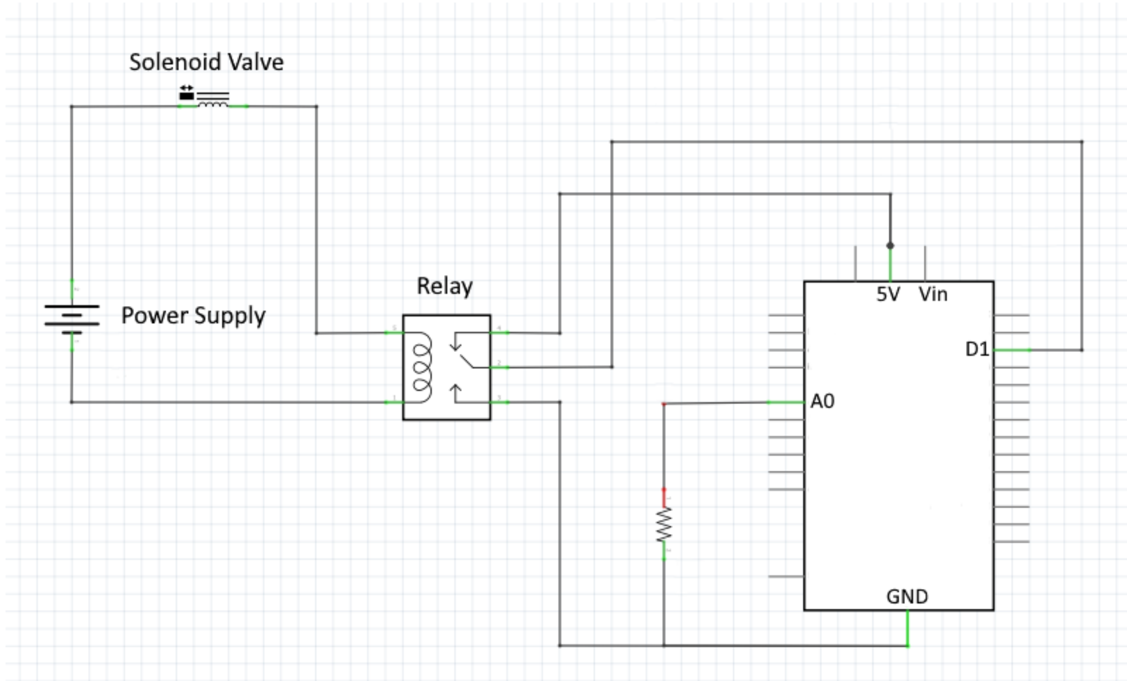

Figure 3.2: A schematic diagram of the electronic system.

A program running on the microcontroller set the voltage on the digital pin connected to the relay. A square pulse was applied to the pin, causing the relay to close. The duration and frequency of pulses could be changed by the software. When a pulse was sent to the digital pin, the circuit between the power supply and the solenoid valve closed, causing the valve to open and air to flow. Through this system, pulse lengths could be controlled with a resolution of one ms, however the lower limit of actuation time was constrained by the power supply. At the valve's recommended operating voltage of 15 V, the minimum cycle time was found to be roughly 35 ms, and somewhat dependent on the pressurized air within the valve. Increasing the voltage to the maximum output of the power supply, 18.5 V, reduced this time to roughly 25 ms.

Through a combination of manual pressure adjustment using the regulator, pulse duration adjustment using the software controls, and nozzle size selection, the properties of the air puff could be varied widely to replicate the puff found in the Corvis, other commercially available tonometers, as well as any other variations we were interested in testing.

3.1.3 Eye Mounting

Porcine eyes were used to test our system. The eyes were stored in a refrigerated container until they were ready to be used. After the cornea was removed from the eye tissue with a razor, the lens was removed, and the remainder was cleaned using Dulbecco's phosphate-buffered saline solution. The cornea was then mounted onto a Barron Anterior Chamber, which allowed the IOP to be artificially set. A water column was used to manually vary the intraocular pressure over a range of IOP values on a given sample.

3.2 Imaging System

Our imaging system was comprised of two cameras and a projector. The projector was used to apply the necessary structured light pattern, while the cameras were used to create input footage and ground truth respectively.

3.2.1 Imaging Equipment

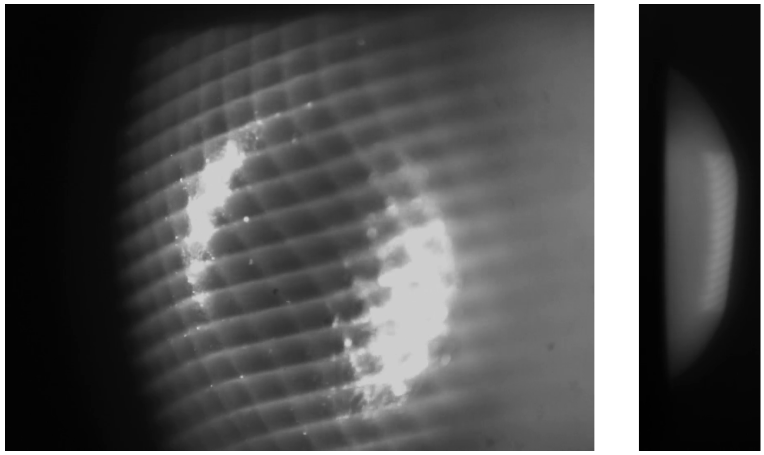

Figure 3.3: Deformation (left) with accompanying ground truth (right).

We imaged the cornea simultaneously from two different angles to produce input footage and a corresponding ground truth. The input footage was taken at 500 fps using a Basler acA640-750 μm camera, while the ground truth angle was taken at 160 fps using a Basler acA2000-165 μm camera. Without equipment restrictions, the same type of camera should be used for both angles; however, instead we synchronized the ground truth with the input footage by slowing it down by a factor of 0.32. The input footage was taken at a 45° angle to the cornea, while the ground truth was taken at a 90° angle to the cornea, directly from the side. To obtain 160 fps performance from the high-resolution ground truth camera, we used the Basler Pylon software to crop the camera's output resolution to 296×1088 pixels. See Figure 3.1 for a schematic of our experimental setup. Figure 3.3 depicts an example of a deformation and its corresponding ground truth.

3.2.2 Pattern Projector

We used an Opto Engineering 3D LTPRSMHP3W-G pattern projector that was designed for use in structured light applications, and use the PTGR050450P grid pattern as a projection template. The pattern grid has 22 rows and columns, with a line spacing of 0.45 mm and line thickness of 0.05 mm. The projector is pointed at a 45° angle to the cornea.

3.2.3 Choice of Projection Template

We chose to project a regular grid onto the eye, in contrast to the Placido topographer's concentric rings and the Orbscan's scanning slit pattern. We believe that a regular grid has advantages over both varieties.

The concentric circles used by the Placido topographer implicitly define distortion radially, with respect to a definite center; however, it is not always clear where this "center" may be, and centering the pattern requires care and time. Additionally, it is often necessary to use multiple rings to understand the full extent of a deformation [104, 105], requiring that more of the pattern be analyzed. In contrast, we expect that a grid will have a higher density of information at every intersection.

Rather than concentric circles, the Orbscan topographer employs 40 scanning slits projected from two angles, creating 240 intersection points on each slit [68]. However, creating more points of reference using a grid is not difficult; for instance, our grid has 16×16 lines with 256 points of intersection. Denser grids could be employed to create more information as necessary.

Existing methods may use these radial or low-resolution patterns in the interest of algorithmic tractability. These methods use explicit selection of features, either by ring or by intersection, which places an upper bound on the number of features that may be processed in a reasonable time. Algorithms must compute astigmatism scores or compare the movement of feature points between fixed camera perspectives [104, 105]. By contrast, the high information density within a regular grid used in combination with a CNN, does not rely on manually-extracted features and may instead rely on more general patterns of grid deformation. Additionally, because the features extracted by a CNN are largely spatially invariant, these networks are more robust to a variable number of intersection points and allow for scaling of grid density.

3.2.4 Calibration Procedure

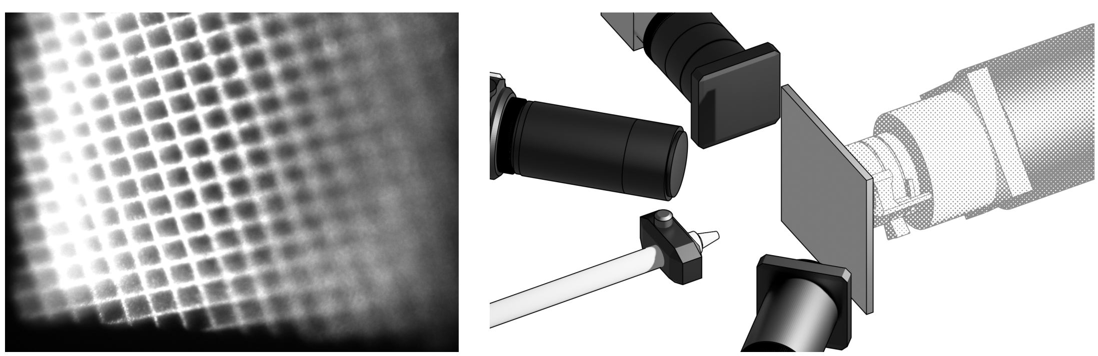

Figure 3.4: Calibration plate pattern sample photograph, and suggested position of calibration plate relative to Barron chamber(right).

To resolve the ambiguity caused by variations in focal length and the positions of the projector and cornea, we imaged a calibration shot to establish scene priors regarding depth of field, camera viewpoint, and projection angle. An example calibration plate is shown in Figure 3.4, along with its position relative to the Barron chamber. To create this image, we substituted the Barron mount for a featureless flat metal plate while keeping the projector and camera intact. We attempted to position the plate exactly at the base of the cornea on our mount. Our algorithm then used this calibration plate as an additional input when considering corneal deformation footage taken with the same scene priors, which is discussed in more detail in Section 3.5.2.

3.3 Data Collection

Many tests were involved in tuning the parameters of the experimental setup. This includes changing the frame rate and exposure settings of the cameras, adjusting the focus of the cameras and projector, varying the puff intensity and duration, and varying the IOP of the cornea. Once the optimal parameters were achieved, data were collected for use in the final analysis.

During data collection, the knob on the fume hood air source was opened an appropriate amount and then kept in position for the duration of the test. This ensures that the intensity of the puff is as consistent as possible throughout the different measurements. The IOP is then set to 10 mmHg according to the water column manometer. The microcontroller began puffing the cornea for 5 seconds at 500 ms increments. Once the cameras record these puffs, the IOP was raised by 2 mmHg, and the process was repeated. Once data were collected at 28 mmHg, or the limit of the manometer column, the pressure was decreased to 10 mmHg, and the process was repeated. In total, we completed three of these cycles.

3.4 Data Simulation

With the resources and time available to us, only simulated data would provide ground truth data with enough quantity and resolution to train our network for this task. We rendered 20,000 frames of simulated deformation and then augmented this simulated data using a variety of transformations (as described in the next subsection) to increase its size to 80,000 samples. We used the free and open-source 3D computer graphics software Blender, version 2.91.2. Figure 3.5 shows an example of a simulated cornea, along with its corresponding calibration plate.

3.4.1 Rendering Procedure

The render starts with spherical mesh sectioned to approximate the corneal surface. A Displace modifier is added to the mesh, in which the texture is a 2D gradient that falls of spherically from the origin. When applied to the mesh, this results in a "puffed" surface, and the strength of the modifier can be varied to achieve appearances between resting and "puffed".

In the first two iterations of our simulated datasets, we modeled the cornea using an infinitely thin mesh with a Glass BDSF (bidirectional scattering distribution function) node. Our third iteration focuses on greater physical accuracy, and includes thickness and scattering effects within the corneal surface.

Iteration 0: Simulated Iris and Grid Projection

This iteration was focused on simulating the effect of projecting a pattern on an iris. We modeled the iris using a circular disk, textured using a sample image of a real iris. We then projected a checkerboard onto the cornea using a Blend shader with a coefficient determined using the Fresnel node, resulting in a "mixed" image, which is shown in Figure.

Iteration 1: Basic Grid Distortion

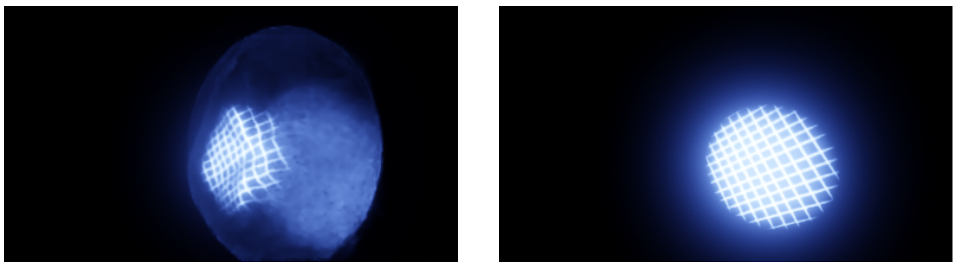

Figure 3.5: A simulated cornea with projected pattern (left) and corresponding calibration plate (right). Simulated topography is complex to encourage model robustness.

In our first iteration of simulation, we focused on reproducing the pattern distortion. Consequently, we ignored physical scattering effects and thickness, instead modeling the cornea as an ideal refractive element with an IOR of 1.45 using the Glass BSDF node. The results are shown in Figure 3.5. Note that no scattering or defocus is present in the resulting patterns.

Iteration 2: Simulated Iris and Grid Projection

In the second iteration of our data generation, we attempted to improve the similarity between our rendered images and the captured images. This was accomplished in two main ways: adding thickness and volumetric scattering to the cornea, and improving the lighting in the scene.

The mesh is given both a surface and volumetric shader, which controls how light is reflected off of the mesh. The surface shader uses the Principled BSDF node with maximum metallic and specular character. The roughness is decreased to a low value, and the alpha of the surface is decreased slightly to allow some light to pass through. The volumetric shader consists of both volumetric scattering and volumetric absorption nodes. The densities of the absorption and scattering nodes is adjusted to match the experimental conditions. The combination of these properties creates a material that both reflects light off of the surface and within the volume of the cornea.

The lighting of the scene consists of several spot lamps with a grid pattern applied to them. The size of each lamp is adjusted to control the sharpness of the pattern on the cornea, and the brightness is adjusted to simulate a gradual blurring of the pattern. Additionally, the lamps are divided into groups that reflect only off of the surface of the cornea, and a group that reflects only off of the volume of the cornea. In this way, the relative intensity of the specular reflection can be adjusted to match the collected footage.

3.4.2 Physical Accuracy of Rendering.

When rendering simulated data, our goal was to reproduce experimentally-gathered images as accurately as possible. That being said, some details of the scene, such as the mount and the ambient environment, are difficult to accurately reproduce. As such, we focused on reproducing the deformation of the grid projected onto the cornea, while the environment only bears a loose resemblance. This reproduction allowed us to augment the data during training which, in turn, helped the network identify that background details can be discarded.

Blender provides several rendering engines. In the interest of achieving a realistic render, we chose the path tracing-based Cycles rendering engine, which aims to reproduce optical phenomena as faithfully as possible by simulating the bounces of individual rays of light emanating from light sources. In this way, the rendering engine could most accurately capture complex reflection and refraction even on surfaces with complex topography.

A detailed description of the modeling process is included in Appendix C.

3.5 Image Processing and Algorithms

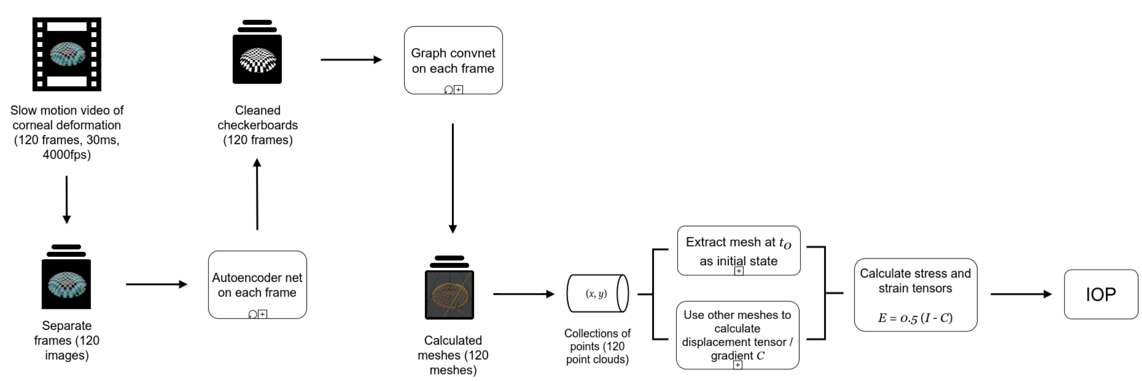

Figure 3.6: Our process for IOP extraction from footage of an illuminated applanation.

Our algorithm attempted to predict corneal topography from a monocular video. To make it more generalizable, we separated the operation of the algorithm into two parts: one which extracted the projected pattern from an image of the eye (CheckMark), and one which predicted the topography corresponding to a "pure" pattern (DeepSquish).

CheckMark is a CAE, while DeepSquish is a deep CNN following the ResNet skip connection architecture [127] that ultimately outputs points on a 2D surface. In the following subsections we detail the architecture, training, and performance of these networks. We also demonstrate that our network displays promising performance on real footage using only simulated data as training input. We describe our process for verifying our network's predictions using our ground truth footage.

3.5.1 Autoencoder Checkerboard Extraction ("CheckMark")

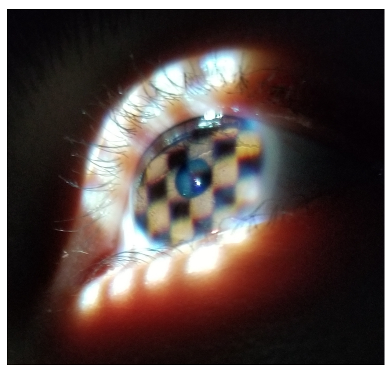

Figure 3.7: Real sample image of projected checkerboard on eye.

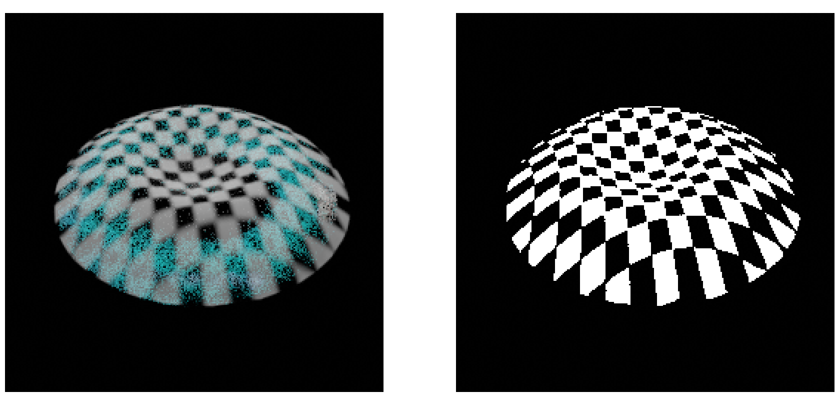

Figure 3.8: Samples from the dataset used to train "CheckMark". Simulated projection (left) and ground truth (right).

Introduction

Any algorithm attempting to make use of reflected light from the cornea will face significant difficulty from the patterns of the eye itself, particularly from the iris (Figure 3.7). To circumvent this difficulty for following stages of the algorithm, we extracted a clean pattern from the image using a CAE.

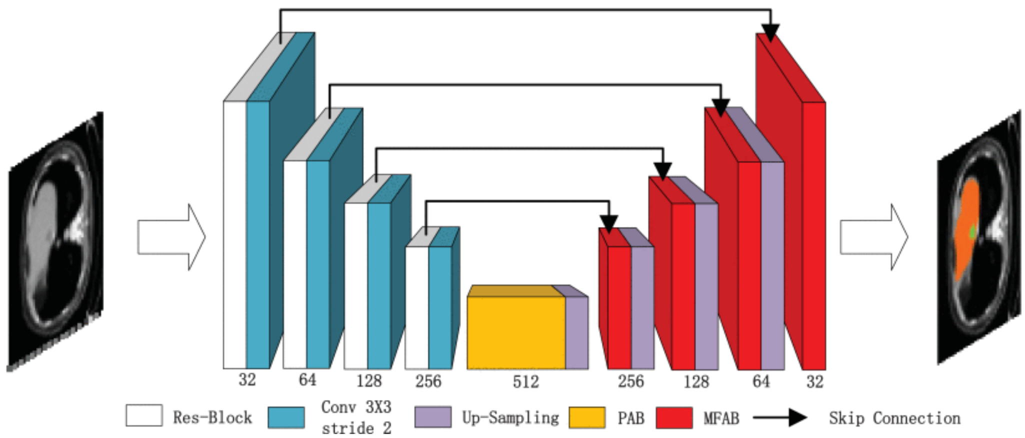

Figure 3.9: The MA-Net architecture [13].

Architecture

We performed transfer learning using the MA-Net architecture, first suggested as a method for segmenting liver and tumor images [13]. The architecture is designed to extract features at multiple scales, and learn relationships among pixels using a spatial attention module. We show a broad overview of the MA-Net architecture in Figure 3.9, although we do not show the custom modules for spatial attention and multi-scale feature fusion. Note the Position-Wise Attention Block (PAB) and custom ResBlocks, which differentiate this architecture from a U-Net autoencoder [128].

Loss Function

We used sum of the binary-cross entropy (BCE) and Jaccard distance loss functions as our loss function to minimize. These loss functions were well-suited for this task because the problem of extracting a clean checkerboard from an image was essentially binary segmentation.

The Jaccard distance term was especially important for our task, as our ground truth labels were imbalanced (i.e. there was a high ratio of “foreground” or checkerboard pixels to “background” pixels). The Jaccard distance term is known for its ability to adequately represent loss in such imbalanced situations and preserve pixel-wise detail, where other loss functions may inaccurately average away these nuances [129]. We could also have used the Dice score, which is closely related to the Jaccard distance, but found better performance when training with Jaccard.

Given a ground truth image y and network prediction $\hat{y}$, our loss function L was formulated as

$$L(y, y') = \mathrm{BCE}(y, y') + (1 - J(y, y')) = \sum_{i \in |y|} (y_i \log \hat{y}_i + (1 - y_i) \log y_i) + \left(1 - \frac{|y \cap \hat{y}|}{|y \cup \hat{y}|}\right) \tag{3.1}$$

where i indexes the image pixels, and the intersection and union operations are done element-wise over y and $\hat{y}$.

Training

Our dataset was 250 256×256 simulated renders of checkerboard patterns of varying density (i.e. number of squares) and angle projected on eyes with varying pupil size and iris color. Figure 3.8 shows an example of an input and corresponding ground truth label from this dataset. We normalized images by dividing them by 255 (0-1 scaling). In addition, we augmented the images by applying random rotations, Gaussian noise, and adjusting the brightness and contrast of the images to help prevent the model from memorizing the train set.

We implemented CheckMark using PyTorch 1.8.0 and the `segmentation-models` library, and train our model on an NVIDIA GeForce RTX 2070 SUPER. The objective function was optimized using the Adam optimizer with a weight decay of 1×10^{-5}. We used a learning rate of 1×10^{-3} for the first 10 epochs, and used a learning rate of 1×10^{-4} for remaining epochs. Weights were initialized using weights from a ResNet-34 architecture trained on the Imagenet dataset. Finally, we trained our model using a batch size of 6 for 50 epochs.

3.5.2 2D Surfaces With Convolutional Neural Network ("DeepSquish")

Introduction

Once a grid had been extracted from our image, we used a second network to process this grid and output a corresponding topography. To output this topography, we present "DeepSquish", a residual CNN, and output the corneal topography as a 2D surface. We also present a loss function that preserved surface normals and detail without explicit calculation of normal vectors or surface derivatives.

Related works in the literature generally regress either depth maps, meshes, or point clouds. Because we are regressing a continuous, thin 2D surface, we elect to predict depth maps. However, our method differs from most approaches to depth map prediction in our choice of viewpoint. Most, if not all, approaches output depth maps from the viewpoint of the camera. This viewpoint is a natural choice given the wide variety of scene geometries; solutions are aimed at generalizing across many possible shapes and sizes.

However, our network outputs a depth map that is treated as a top-down view of the imaged cornea, represented as 2D surface intensity values on a fixed grid. We took this approach because corneas are mostly constant in radius; moreover, the size of the area we are predicting is bounded by the size of the pattern being projected on it. By having our network output a top-down depth map rather than a camera-perspective map, the network represented every part of the corneal surface with equal resolution, as well as parts of the cornea that may be occluded from the camera viewpoint.

Architecture

As with CheckMark, we employed a model architecture that was very similar to the ResNet in its use of residual blocks and skip connections. The network was not quite an autoencoder, as it only had a downsampling path and dis not upsample the input again.

To make the network fully convolutional and support variable-size inputs, we avoided using any fully-connected layers and instead used the fractional max pooling layer introduced by Graham (2014) [130] to downsample the final feature maps to the desired resolution. In our case, we downsampled the final feature maps to a 17×17 pixel map to achieve reasonable convergence times. However, if given more data, there is no reason that the output depth map could not be a different size.

We also employed the Convolutional Block Attention Module (CBAM) to allow the model to employ both spatial and channel-wise information at each feature layer. This module gave the network a larger effective receptive field, an important capability for this class of problem [131]. We employed a CBAM module in each residual block.

Loss Function

We minimized the sum of the mean-squared error (MSE) term and a custom loss function that was designed to encourage the preservation of fine surface detail.

This custom loss function was designed to specify the similarity of a point's elevation to its surroundings, a metric we call the "peak" loss. Specifically, let y and $\hat{y}$ represent the ground truth and predicted depth maps, respectively. Let each depth map be an $N \times N$ depth map of values on a evenly spaced grid from -1 to 1. Then the equation for our custom loss function L_{PEAK} is

$$L_{\text{PEAK}}(y, \hat{y}) = \sum_{i \in |y|} \sum_{j \in |y| \neq i} y_i - y_j \tag{3.2}$$

and the equation for our complete loss function L is

$$L(y, \hat{y}) = ||y - \hat{y}||^2 + L_{\text{PEAK}}(y, \hat{y}) \tag{3.3}$$

where subtraction is done element-wise over y and $\hat{y}$.

Compared to other surface normal-focused losses that explicitly calculate surface normals and are based on the cosine distance or dot product, we found that L_{PEAK} is effective and more efficient.

Training

In this subsection, we describe our training process, including model hyperparameters, training hardware, and data augmentation.

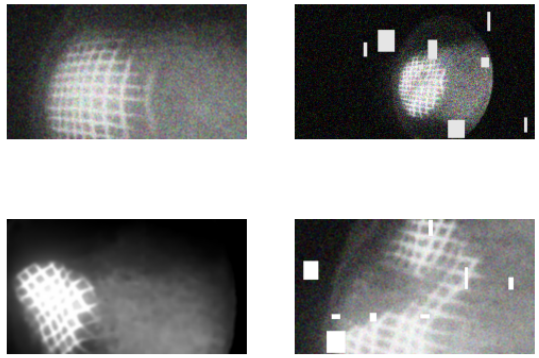

Figure 3.10: Sample augmented images in the train set.

Data Augmentation

Aside from simulating different projection and camera angles, we also augmented the data to make the model robust to changes in lighting, rotation, texture, and scale. To achieve this, we applied the following transformations to the images with random probability and parameters:

- Affine transformations (scaling, rotation, shifting)
- Brightness, saturation, contrast
- Gaussian noise (variance up to 0.05)
- Gamma correction
- Coarse dropout (cutting random-sized holes out of the image)

We applied coarse dropout during data augmentation to help make the model more robust to loss of information in some sections of the pattern, e.g. because of intense specular highlights or irregular surface normals. See Figure 3.10 for example image augmentation, and Appendix B for a link to our repository and our exact augmentation pipeline.

Hyperparameters and Execution

We implemented DeepSquish using PyTorch 1.8.0 and trained it on an NVIDIA GeForce RTX 2070 SUPER. The objective function was optimized using the Adam optimizer with a weight decay of 1×10^{-5}. We use a learning rate of 1×10^{-3} for the first 10 epochs, and use a learning rate of 1×10^{-4} for remaining epochs. Convolution layers were initialized using Kaiming random initialization. Finally, we trained our model using a batch size of 3 for 20 epochs.

3.5.3 Extracting Displacement from Ground Truth

In addition to the footage used by the neural network, we also processed the ground truth footage in order to extract the deformation of the cornea.

The ground truth data was collected as a series of bitmap images which are grouped into folders for each trial, with the IOP and time of measurement recorded. Mathematica was used to analyze these images. First, the images were imported to a multidimensional list organized by the group of puffs and the time step within each group. In order to smooth the image, a Gaussian Filter was applied to the images with a kernel radius of 3 pixels. The image was then binarized with a threshold of .1. This modified the brightness of each pixel such that any pixel with a brightness above the threshold is assigned a value of 1, and all other pixels are assigned a value of 0. A selected example can be seen in 3.11 The images were then converted to image data, which turned each image object into a two dimensional list of 1s and 0s.

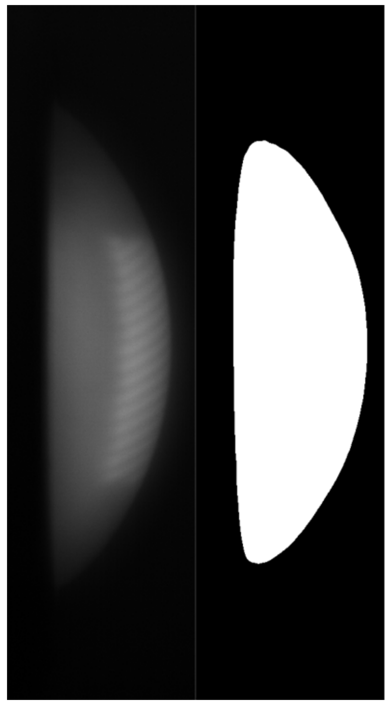

Figure 3.11: A. The output of the ground truth Basler camera and B. the same image with the Gaussian filter and binarization applied.

The front edge of the cornea was then extracted from each image. Each image in the image data list was sorted through row by row. In each row, the x-coordinate of the rightmost bright pixel was stored in a new list. The result of this operation was a new set of lists that contains a list of values corresponding to the position of the front edge of the cornea for each set of puffs. This list was then exported to a set of CSV files, split by each trial, for further analysis and visualization.

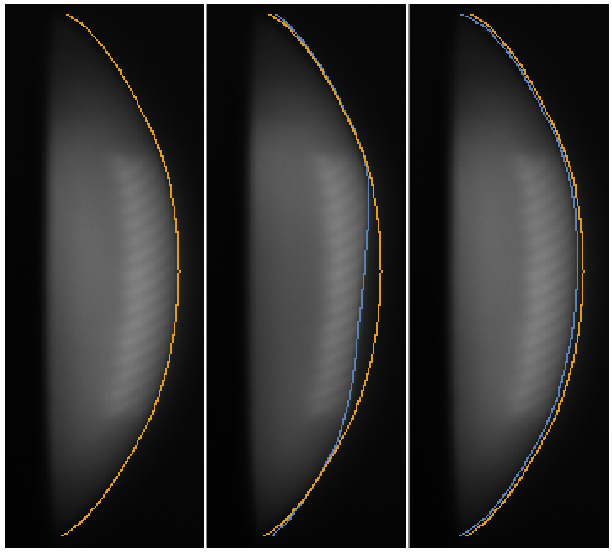

Figure 3.12: Visualization of our analysis process on sample frames.The orange curve shows the edge at $t! = 0$, while the blue curve shows the edge at a non-zero time. The Chamfer distance is taken between the orange and blue contour at every frame.

Another Mathematica notebook was used to calculate the relative deformation from these lists of absolute position. The CSVs were imported, each processed row by row. The Chamfer distance was then calculated between the undeformed curve at the first time step of the trial, t_0, and the deformed curve, $t > t_0$ (Figure 3.14). Rather than search the entire curve, only points with x-values within 30 pixels were considered. The two-dimensional distance was then computed between the point of interest and each of the closest points on the curve at t=0, and the minimum value was stored in a new list. When this was done for every point in every curve, it showed the relative displacement between all curves at $t > 0$ and the initial curve at t_0. This showed only noise when there was no deformation of the cornea and a clear, roughly parabolic, deformation during the puff. This list was also exported to CSVs for further analysis and visualization.

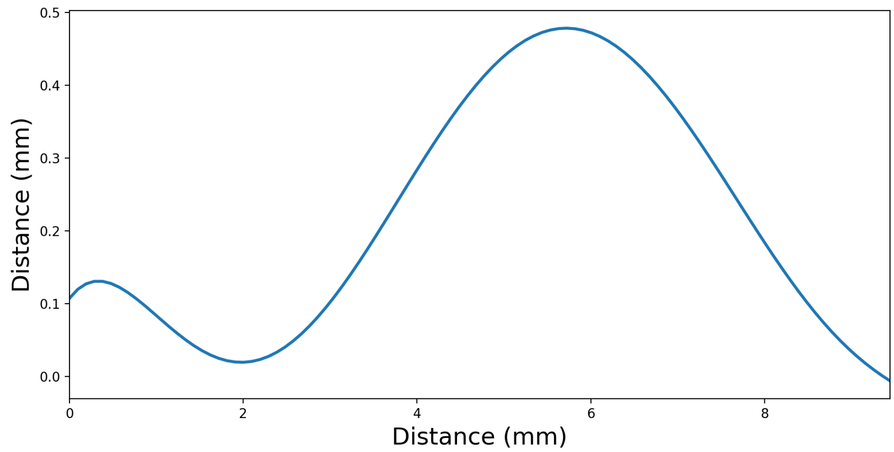

Figure 3.13: Example plot of the magnitude of displacement for some time $t > t_0$

Finally, another Mathematica notebook was used to extract the maximum deformation at each timestep. From the footage and from this analysis, we observed that the cornea deforms in a complex manner: the center of the cornea is formed downward, while the base of the cornea is pushed out. Upon the rebound, this is inverted, and once the puff has stopped entirely the whole surface oscillates about its rest position. While this information was valuable, in order to calibrate our system we are only interested in the deformation of the spot on the cornea where the puff strikes it. As such, we pruned the data by removing points on either side of the cornea. This resulted in a list for each series of puffs that shows, at each timestep, the maximum deformation from rest of the center of the cornea.

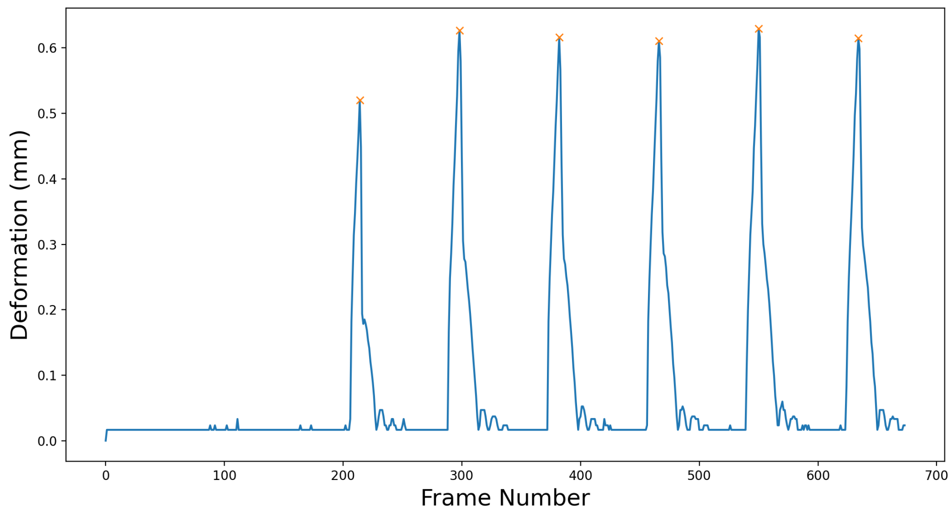

Figure 3.14: Corneal Response for IOP 10 (cm) H_2O.

From these lists, the maximum displacement for each puff was extracted by first setting a threshold on the minimum peak value and then finding the maximum through a comparison with neighboring values. The maximum deformation for each trial was then averaged. The displacement was plotted versus IOP for each puff.

Chapter 4: Results

We evaluated the performance of our algorithm on simulated data before testing its ability to generalize to real footage. In our experiments, we measured the network's ability to accurately reproduce surface topography and capture the amplitude and period of deformations.

4.1 Simulated Data

4.1.1 CheckMark

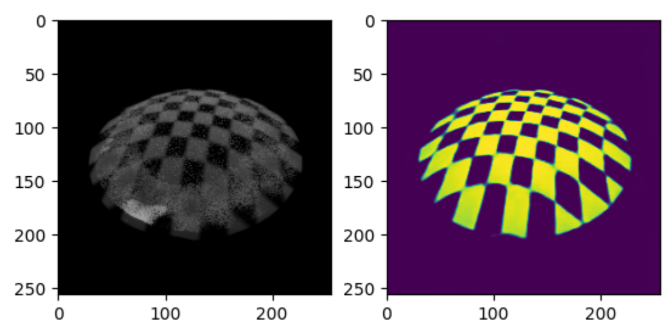

Figure 4.1: Sample performance from CheckMark. Input image is on left, predicted "clean" pattern is on right.

CheckMark achieved excellent performance on the test set, successfully displaying the ability to extract a complete checkerboard pattern with sharp edges from images containing simulated irises (Figure 4.1). The network converges to a train error of 0.25 and a validation error of 0.1946 in 50 epochs (Figure 4.4).

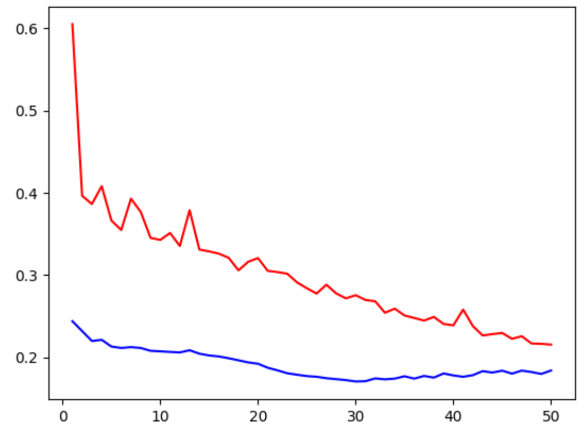

Figure 4.2: Training (red) and test performance (blue) for CheckMark.

RMSE	Jaccard
0.1946	.031

4.1.2 DeepSquish

Performance

DeepSquish also achieved excellent performance on the simulated dataset, achieving a test error of 0.031 in 20 epochs of training. As seen in Figure 4.3, the network is able to reproduce even complex surface topography with very little qualitative difference.

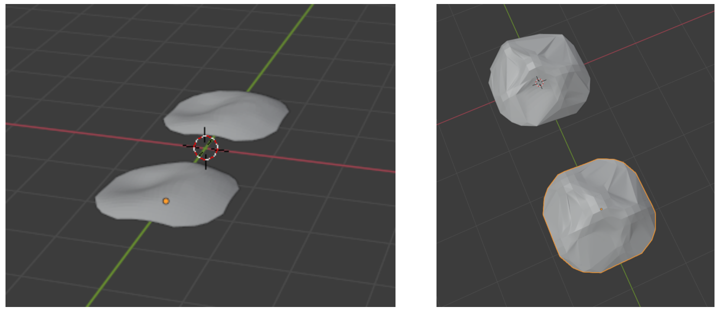

Figure 4.3: Sample performance from DeepSquish. Predicted meshes are in front, ground truth meshes are in back.

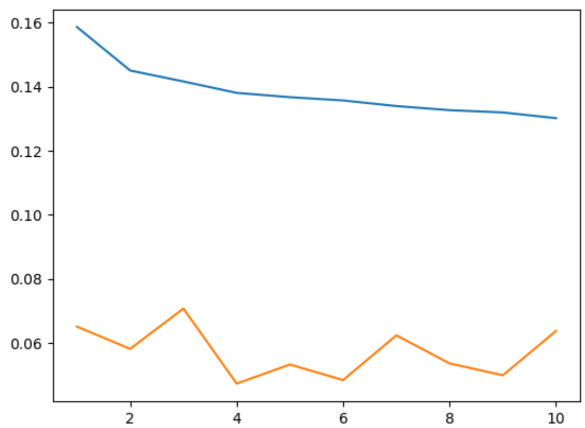

Figure 4.4: Training (blue) and test performance (orange) for DeepSquish (left). Relevant performance metrics (right).

MSE	L_{PEAK}
0.00523	0.09939

Model Interpretability

It is important to understand what features our network uses to make its predictions. To this end, we analyzed our network using the Grad-CAM method [132], which is often used to create "visual explanations" of a network's prediction process. To achieve this, Grad-CAM uses backpropagation to compute a weighted sum of the activations for some feature layer. Feature map pixels with high positive

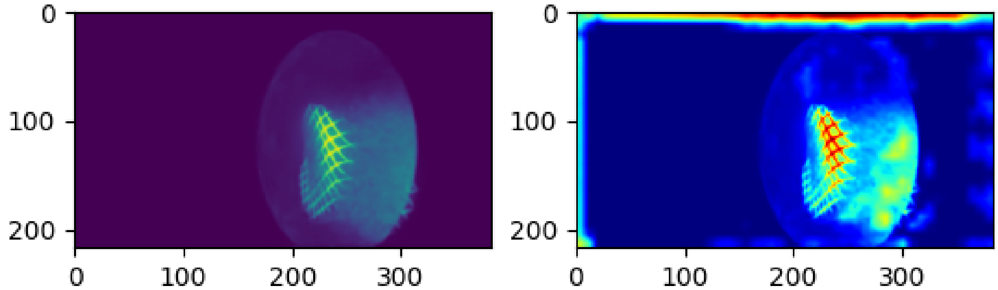

Figure 4.5: Localization heatmap generated using the Grad-CAM method for our model. Note the focus on the pattern and distance from cornea edge.

gradients are considered "important" to the network, allowing for the creation of an localization heatmap. In our case, we would expect to see a localization map that is strongly focused on the projected pattern, which would greatly decrease the probability that the model is operating on memorization or bias. We ran Grad-CAM on an intermediate feature layer in our model (layer 11, BatchNorm2d), following the intuition that deeper layers provide higher-level visualizations, and show a sample heatmap in Figure 4.5.

Our results show heavy focus on the grid lines, as well as the region between the edge of the pattern and cornea, as indicated by the concentration of red near the grid lines. This strongly suggests that the model is indeed learning the location and shape of deformations based on pattern shapes and locations.

4.2 Real Footage

In this section, we discuss the extracted deformation data from our ground truth footage, and compare it to the output of the network. We analyze how the deformation of the cornea correlates to increases in IOP, and compare trends observed in the ground truth to trends predicted by our algorithm.

4.2.1 Ground Truth

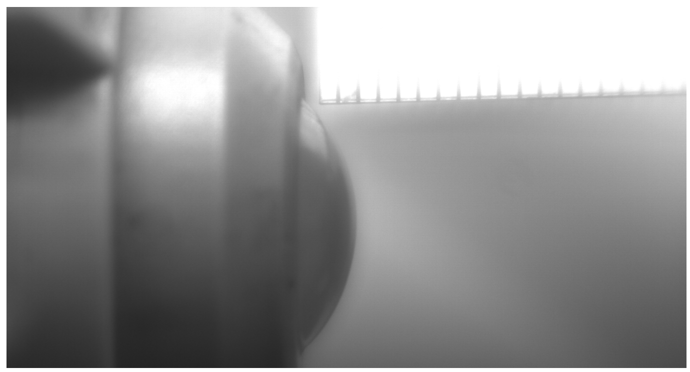

Figure 4.6: An image from the ground truth view showing a ruler for calibration. Each tick mark represents one millimeter, with each pixel being equal to 16.56 micrometers.

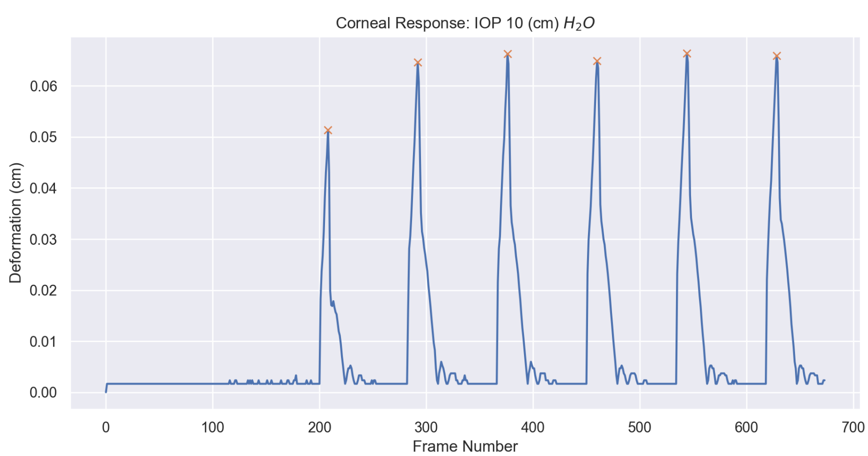

Figure 4.7: Example plot of predicted deformation over time.

The ground truth footage consisted of 196 puffs over a range of IOP values between 10 and 28 mmHg. The extracted deformation for each puff, calculated in pixels, was converted to a distance in millimeters using an image captured with a ruler, shown in Figure 4.6, for scale. The measured deformation versus IOP for each puff is shown in Figure 4.8. A linear regression of the data yielded a line of best fit with the equation $0.820856 - 0.0184785x$ and an R^2 value of 0.913358. This reveals a clear negative correlation between the deformation and IOP; however, inspection of the data also reveals that there is a group of outliers below each group of data points.

For many of the groups of puffs, the first puff resulted in a significantly different amount of deformation. As is the case with many test of the mechanical properties of organic tissue, it seemed that a pre-stressing step is required to probe the true properties mechanical properties of the biological specimen. Once these pre-stressing data points were removed from the set, as shown in Figure 4.8, the line of best fit changed only slightly to $0.840319 - 0.0189514x$, while the R^2 value increases to 0.968.

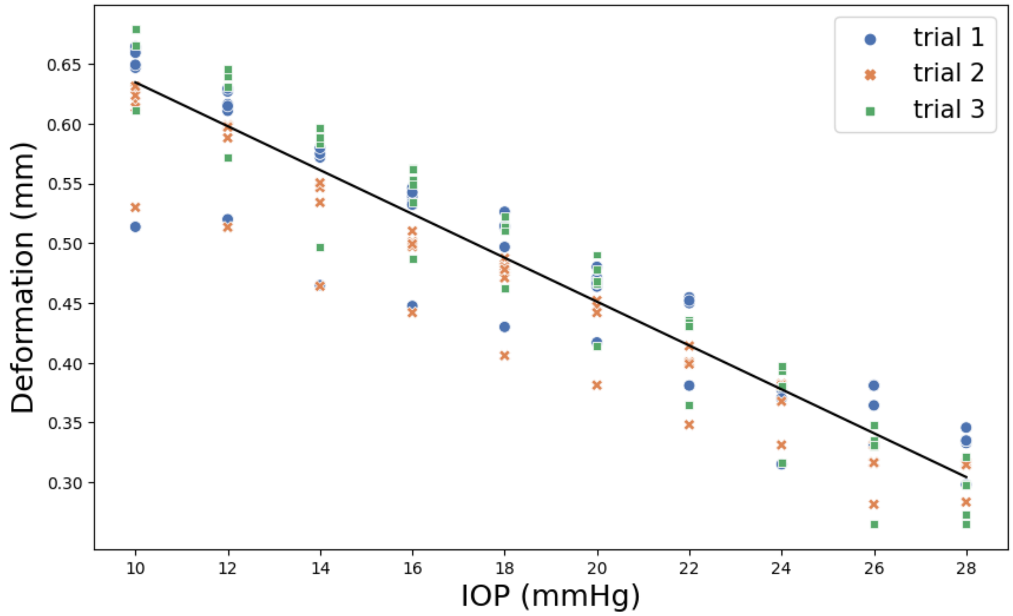

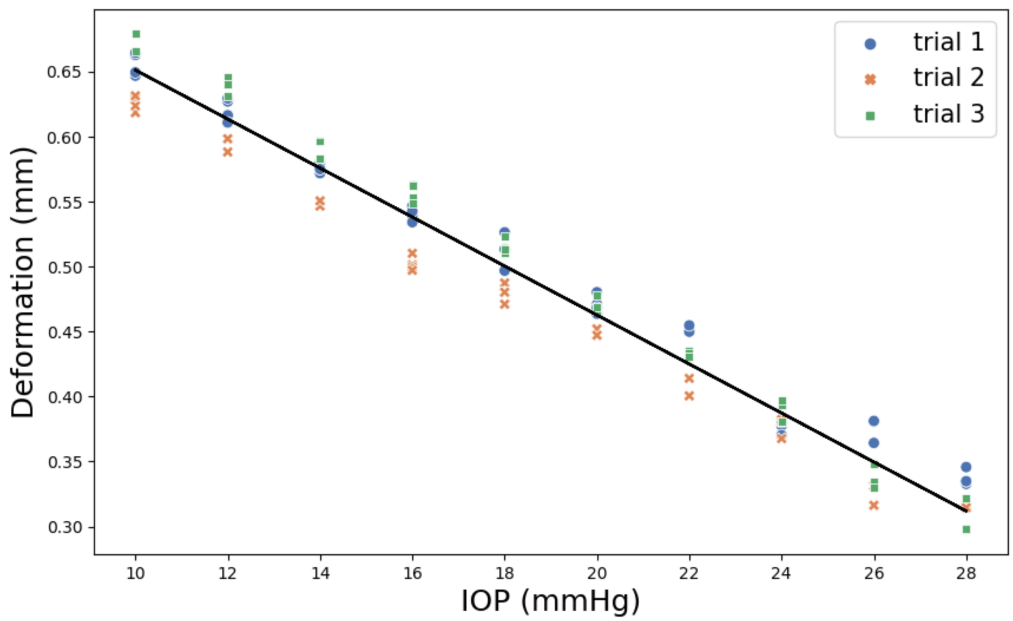

Figure 4.8: The maximum deformation for each puff versus IOP for all trials (top), and trials one through three (bottom)

4.2.2 Comparison with Network Prediction

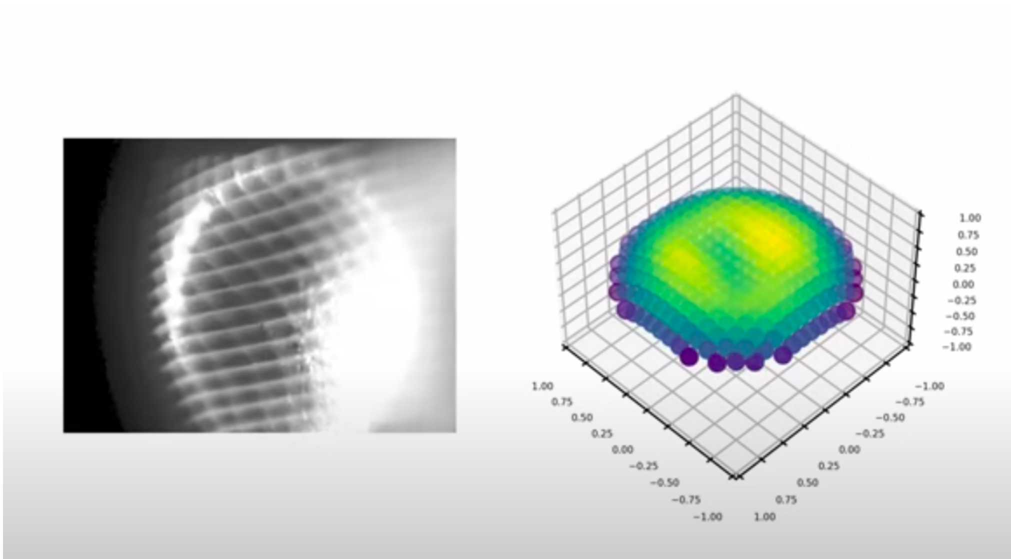

Figure 4.9: An example of a deformation map generated from our real footage.

Iteration 1: Basic Grid Distortion

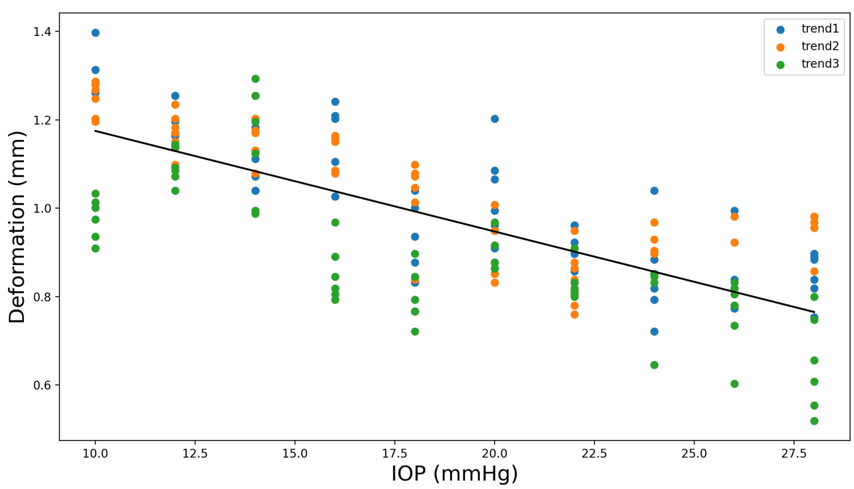

Figure 4.10: Example plot of predicted deformation over time.

In this test, we used a network that was trained on the first simulated dataset as described in Section 3.4.1, which did not include optical imperfections or scattering effects from the cornea. The resulting trend is shown in Figure 4.10. We observed a weak negative linear correlation between IOP and deformation amplitude. A linear regression produced a line of best fit with an equation $y = -0.023x + 1.402$ and an R^2 of 0.555.

This shows that although the data simulated in this iteration bears some resemblance to our real footage, it is still quite dissimilar to the real footage we gathered. In the next section, we describe the substantial improvement obtained by training the same network on the dataset created in Iteration 2.

Iteration 2

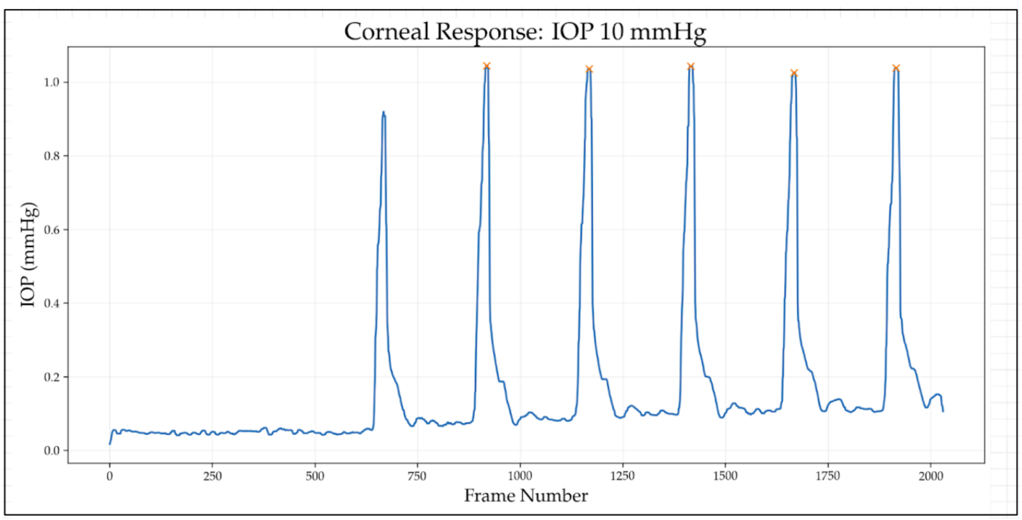

Figure 4.11: Example plot of predicted deformation over time.

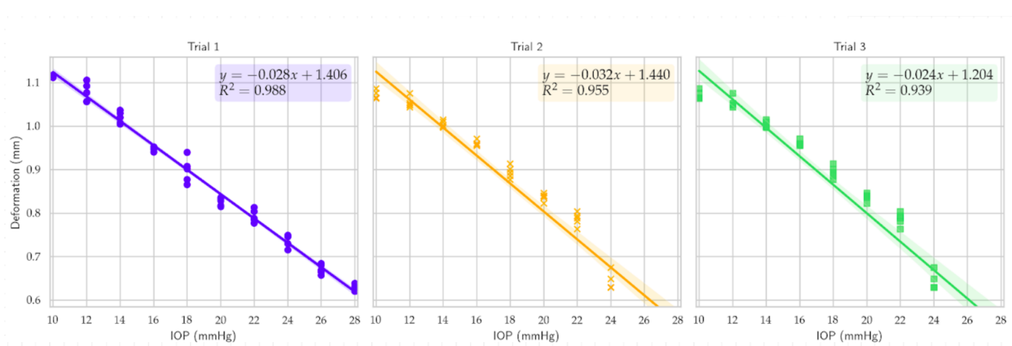

Figure 4.12: Model correlation for each trial.

We applied a similar process to the network output, storing the output mesh for the first frame of the footage and then calculating the Chamfer distance between it and predictions made for following frames. We did this for the two-dimensional surface output by the network, in contrast to the one-dimensional contour extracted from the ground truth. An example of the resulting deformation plot is shown in Figure 4.11.

It is important to note that the ground truth does not show the concavity of the deformation, making it difficult to confirm our results for the predicted amplitude. In addition, we could not quantitatively locate the puff in the axis moving away from the camera. Therefore, we compared the trend of predicted maximum deformation versus increases in IOP to that found in the ground truth, and compared the profile of the predicted deformations to the ground truth footage.

To extract the trend for maximum deformation versus IOP, we extracted the peak deformation amplitudes in each time series. We then plotted the resulting points against IOP, and investigated the correlation obtained. In Figure 4.13, we show that we find strong negative linear correlations in each trial, with R^2s between 0.939 and 0.988.

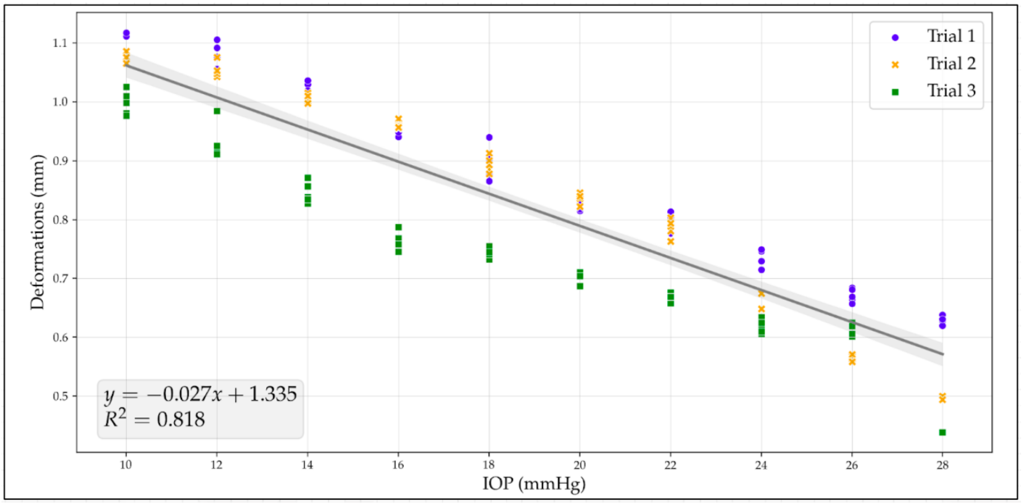

Figure 4.13: Combined predicted datapoints from all trials.

When combined, the datapoints from the resulting trials have an R^2 of 0.818.

Chapter 5: Discussion and Conclusion

5.1 Meaning and Significance of Results

The foundation of our imaging and analysis system is our DeepSquish neural network, so we begin this section by discussing the results of our neural network training using simulated data and then discuss the meaning and significance of our results of using our network on footage from real corneal deformations.

5.1.1 Simulated Results

Our neural network receives an input image of a cornea, extracts a cleaned version of the projected checkerboard mapping (CheckMark), maps a predicted three dimensional topography of that pattern (DeepSquish), and outputs that mapping.

Our first tests for the network involved only the extraction of clean checkerboard patterns from realistic images. In order to train and test this portion of the network, we created simulated images given limited experimentation time and access. We simulated differences in illumination, focus, and angle. We found that our CAE is capable of extracting patterns from mixed images, achieving a Jaccard loss index of 0.03. This result means that, when given an imperfect simulated image, our first network is able to output a clean mapping to be interpreted. It also is significant enough that we are able to test this network on real images, the results of which are discussed in the next section.

The tests for the next part of the network were aimed at finding the effectiveness of our neural network for transforming a clean checkerboard pattern to a three dimensional topographical mapping. Our tests for this network showed that the network learned simulated footage very well, reproducing complex surface detail, giving us the confidence to test our network on real-world footage and see if it was possible to generalize using only simulation.

5.1.2 Real Results

Due to the COVID-19 pandemic, we were unable to perform a large quantity of tests of real footage, but we were still able to establish a proof of concept using a network trained only on simulated footage.

The main test involving our network and real imaging system was a series of videos taken at several different IOPs. While holding the puff intensity constant, we show that our neural network is able to predict a clear correlation between deformation and increase in IOP. The network predicts a deformation of 1.3 mm at an IOP of 10 mmHg, decreasing to 0.78 mm at 28 mmHg. This presents an expected amount of deformation; other studies have reported a deformation of 1.3 mm at an IOP of 14 mmHg [133]. This test means that our network can both clearly distinguish between different IOPs when puffed at the same air pressure, and provide a physically appropriate prediction of the deformation. This is the basis for tonometry, and it shows that our system has potential as a method of determining IOP.

When using the first batch of simulated data, we conclude that the poor R^2 value within the results is largely due to the fact that the distribution within the simulated footage is simply too different from the real footage. Our Grad-CAM analysis and performance on simulated data suggests that the network is able to successfully learn the relationship between deformation and pattern distortion, but real-world data will be needed to make our model robust to the different varieties of optical distortion. Although we were able to simulate the basic phenomenon of pattern deformation, we were unable to capture more complex optical effects like projector depth of field and specular highlight. A potential remedy for this issue is to obtain more precise ground truth data for real footage, allowing us to train our network on real data and avoid using simulated data entirely.

The improvement when using a version of our network trained on the second batch of our simulated data is further support for this hypothesis. We observe a rise in R^2 from 0.555 to 0.812 among all combined trials, with a per-trial R^2 of 0.939 to 0.988. This substantial improvement suggests that it is important, but still computationally tractable, to create a version of our neural network that can generalize to the real world using improved simulation capabilities.

At the same time as we recorded images of the projected pattern, we also measured the profile of the cornea from a 90° angle, which captures the profile view of the cornea during deformation. We compared our predicted deformations to the measured ground truth and found that, while the deformations show a similar slope, the magnitude of deformation predicted by the neural network is about twice as high as the ground truth measurement. This is explained by the fact that once the cornea is deformed past applanation, the concavity cannot be seen from the profile view. At this point, the ground truth deformation will slowly increase while the deformation at the apex will continue. This also means that the difference between the two measurements should be greater for larger deformations and decrease as the deformation shrinks towards applanation. This can also be seen in the data, further supporting the efficacy of the system.

Although the ground truth shows a strong correlation between IOP and deformation magnitude, there is still a large variance within the data at each IOP. Aside from the initial pre-stressing puffs showing consistently less deformation, the rest of the variance appears to be random, indicating that it is caused by variability within the instrumentation, such as variation in the flow rate of air or small differences in the cycle time of the valve, and not drying of the cornea or another degradation mechanism within the tissue.

5.2 Comparison to Existing Works

The use of structured light to recover corneal topography is not new. However, we believe our method is unique algorithmically. Neural networks have not yet been used to analyze anterior corneal reflections to recover topography. Instead, as mentioned, applications of deep learning in ophthalmology have focused on the classifying of corneal topography data rather than predicting it [134]. Existing algorithms for the recovery of corneal topography rely on hand-engineered features, can be computationally intensive, and are sometimes forced to ignore some information for the sake of computational tractability, as in the case of the SimK algorithm. In contrast, we suggest that experiments on simulated and real data show that neural networks learn to take advantage of minute pattern deformations across the corneal surface, and are much easier to parallelize.

In addition, our tonometer is different from existing non-contact tonometers in several respects. Aside from the algorithmic difference, which we have discussed, there are also differences in equipment requirements. Other non-contact tonometers have taken advantage of the Scheimpflug imaging technique and thin-slit illumination to get high-resolution information about a cross-section of the cornea. However, this method necessitates the use of a rotating camera, and can only recover a thin slice of the corneal surface, which requires the ophthalmologist to carefully center the resulting puff in order to achieve near-radial symmetry. By contrast, our method has the potential to recover the entire corneal surface in 3D, possibly removing the need for careful centering and allowing for more comprehensive analysis of corneal deformation. Consequently, we believe that our presented method may be incorporated in a device with fewer moving parts, fewer operational restrictions, and more robust analysis compared to existing non-contact tonometers.

5.3 Future Directions

We have created an air puff generator and analysis system which clearly models varying deformation of porcine eyes. We now look toward ways to develop this novel proof of concept into a widely accessible product. Due to time and resource constraints, as well as research restrictions throughout a large portion of 2020, several aspects of this experiment would benefit from more robust development, experimentation, and analysis.

The nearest future goals are those which further test the validity and usefulness of our experimental setup. These shorter term research goals would require either the same supplies currently used in the setup with additional time for experimentation or a small number of additional measurement tools such a pressure sensor calibrated for a specific range of pressures.

Beyond these shorter goals are tasks which will compare our system to modern tonometry techniques and display our method's unique advantages or disadvantages. Several tonometers require deformation to be precisely in the center of the cornea to isolate a specific two-dimensional cross section of the eye. Our method should be able to classify non-centered deformations through tensor analysis of the produced deformation mapping, but will require extensive research to confirm.

Finally, other steps need to be taken to solidify our method as a realistic and equitable diagnostic tool. These steps include further research into the tolerance of varying air puff intensities as well as in vivo experimentation.

5.3.1 Puff Generation

As seen in our methodology, we have used two sources of air puffs with varying puff monitors. In the setup used for the majority of our analysis, we simply connected to a variable air supply and did not monitor the actual air pressure that reaches the eye. It is essential to know the pressure of the air impacting the eye in order to properly characterize IOP. Solely regulating output air pressure of an air puff does not provide enough quantitative control for this purpose, as the air pressure exiting a nozzle is proportional to but vastly different from the force and pressure of that air hitting a cornea.

Because we will never be able to directly measure the force on the eye, we need to finely control output pressure of air and create a relationship between that output pressure and the pressure felt by a cornea at a specific distance. By purchasing and integrating a pressure sensor correctly calibrated to the range of pressures we work with, one can create this relationship between output air pressure, deformation, and IOP, and further validate the precision of our air puff generation.

5.3.2 Robust Modeling

Our goal is to lay the foundation for a product that will help reduce the economic barrier between patients and modern tonometry devices, and to introduce new technological advances in machine learning into ophthalmic practices. There are several branches of this experiment that lay directly ahead which will help place our experiment in context with current tonometry techniques and drive it forward into the realm of realism.

5.3.2.1 Establishing Agreement with Existing Tonometers

Using air puffs to deform eyes for the purpose of measuring IOP is not new; there are several commonly used tonometers that all use this method, notably the Corvis series of tonometers. These tonometers are all proprietary and have not published detailed writings about why they use the puff pressure and shape that they use. Because air-puff tonometers have been found to be a suitable method for wide-scale screenings of IOP, it is prudent to compare our imaging and analysis system to other tonometers under the same air puff and corneal conditions [135].

We began this process while designing our air puff system; we used a microphone to imprecisely classify the intensity and duration of a Corvis air puff as a model for our system. Once we acquire a pressure sensor and validate the precision of our own puffs, a close next step would be to again measure the Corvis air puff in an attempt to replicate its intensity and duration with increased accuracy and precision. By calibrating our system to release the same puffs as the Corvis system, we will be able to compare our deformation mapping with the deformation classified by the Corvis software to compare our output IOPs.

This process will also allow us to more effectively compare our model's advantages with advantages of other air puff tonometers. For example, if our system is able to accurately model deformation and accurately determine the IOP of patients using lower pressures than in existing tonometers, our system can decrease discomfort and invasiveness and can be tailored to sensitive patients.

5.3.2.2 Realism Tests

This subsection discusses potential research endeavors to begin testing efficacy and application of our model to more realistic scenarios and would help us determine which parts of our neural network to rework in order to best account for variance in human eyes. There are two types of these future goals. The first type are research endeavors involving testing on humans to better understand comfort in relation to our experimental setup and air puff. The second type involves researching more abstract human conditions such as age, gender, race, corneal thickness, and other specific eye attributes that may change the results of our system.

One common motivation for developing new tonometry techniques is limiting any potential invasiveness or discomfort during the measurement. Contact tonometry may result in small scratches on the cornea which cause discomfort for several hours, and receiving a strong enough puff of air to displace an eye is uncomfortable. One of our goals is to devise a device which has advanced enough imaging and processing to permit lower air pressures in order to reduce discomfort as much as possible during measurement. In order to accurately judge the needs and wants of potential users, it would be beneficial to survey community members throughout a range of air puff pressures to determine at which point if any air puffs become uncomfortable in order to maximize cost, comfort, and efficacy.

Broader than just reaching out to community members in order to judge discomfort from air puffs is the acknowledgement of different needs of different people. It is important to create this device head on with the idea of increasing equity, so it will be crucial to research further differences in parameters which may impact IOP output of our model from age, gender, and race, which may be more easily known, to corneal thickness, curvature (astigmatism), and eye size, which may be unknown prior to measurement.

One way our model may have a flexibility advantage is its ability to train on specific data. If this type of imaging system is used with many people and types of eyes, different models could be trained for different communities using transfer learning in order to maintain best results.This specificity of modeling would be beneficial for several reasons. The usefulness from the model comes from it transforming many images to a deformation map of the cornea throughout the air puff, but if factors such as age, gender, or corneal thickness impact the images we record, then the network will incorrectly map them. In addition, for eyes of different shapes from, for example, astigmatism, our model will not recognize which eyes are astigmatic and which are not, and will create deformation mappings which would most likely assume that the eye is not astigmatic if the majority of its training samples are not from patients with astigmatism.

5.3.2.3 Device Implementation and Design

Our experimental setup was designed with the potential of easily being converted into a distributable tonometer product while still maintaining its core properties. At its core, our design is composed of one stationary camera, one stationary projector, and one air pressure source. Our current experimental setup has the air source positioned in line with the eye, at a distance of 40 mm. The camera and projector are positioned at roughly 45° angles to either side of the air source at a distance of approximately 80 mm each. As such, this kind of setup could easily be assembled into a single device containing all three elements that could rest on a tabletop.

We will briefly compare our tonometer to other tonometers, both contact and non-contact. One of the most state of the art and widely used non-contact tonometers is the Corvis tonometer, which uses a more complex imaging system. Corvis tonometers use rotating cameras which quickly move on a circular track around the eye, a feature that adds potential failure points as well as cost and complexity. Corvis tonometers also are of similar dimension to our setup. Corvis systems also use a thin slit illuminator in conjunction with their imaging system. This thin illuminator is in contrast to our compact pattern projector, which is larger in scale than the illuminator. In order to make our design as compact as the Corvis, we would need to shrink our projector while still meeting the constraints that it projects a bright pattern while lacking a color wheel. The Corvis is a significant investment for any establishment to obtain. Thus, the large price results in many ophthalmologist places renting the machines rather than buying them outright. A more user-friendly tonometer option is the iCare Home tonometer, which costs $1000 . This product, though much smaller and lighter than our design, uses contact tonometry and constantly requires new materials, which leads to increasing cost of ownership over time.

When compared to our design, it is much more expensive.

The method for obtaining measurements will also differ in clinical settings from our porcine tests. A fully operational unit will need to include ergonomic features and calibration steps. Similar to current non-contact tonometers, an adjustable chin-face stage is required to properly position the cornea of the patient. Because our neural network-based approach is more robust against off-center puffs and differences in the camera and projector alignment, the alignment of the cornea of a patient requires much less accuracy than traditional methods. This should enable us to both reduce the cost of our overall design, as well as make it more user-friendly and easy to use. Our current system requires calibration with a flat plate to ensure the camera and projector are in focus.

A future design could include a more flexible lens configuration which would reduce the depth-of-field effect and make the requirements for focusing the device more flexible as well. Furthermore, the strength of the network in extracting a pattern from a distorted image may also be leveraged to bring a slightly out of focus image into focus, even further reducing the calibration requirements.

In addition, the final product will need a user interface so it can be controlled by a technician or user. Traditionally, the physician will adjust settings on the opposite side of the tonometer, either on an attached computer or directly through the tonometer, as is in the Corvis ST. Our product could include several options: first, similar to the Corvis, the tonometer could have a screen and knobs on the opposite side for a physician or other person to use. This would increase the cost of the machine, but prevent the requirement of additional purchases to fully utilize the equipment. Next, the product could connect to a computer or phone via a wired or wireless connection. Studies have shown that similar networks to ours can easily run on cell phones, and these controls would remove the need to include expensive non-essential software in each tonometer.

5.4 Conclusion

The experimental and simulated results demonstrate the viability of neural networks for use in tonometry. We constructed two neural networks, which together are able to predict a three dimensional corneal deformation map based on real footage. The first neural network cleanly extracts a checkerboard pattern as projected onto the cornea. Using simulated footage, we demonstrate that this network can extract the pattern successfully in spite of imperfections such as reflections and image focusing errors. The second neural network converts this pattern into a three dimensional topographical mapping. This process was also verified using simulated data, justifying experimental analysis of the model. Experimentally, the deformation predictions were shown to behave similarly to ground truth measurements. We are able to obtain a clear correlation between the deformation predicted by the model and the IOP, meaning that this method could realistically be applied to extract IOP measurements in the future.

Further experimental analysis must be done to improve the robustness of the experiments and to compare the advantages and disadvantages of this method to current tonometry techniques. However, our existing implementation suggests that neural networks can be used to determine corneal topography with comparatively low levels of computation. Additionally, we are able to construct a three dimensional mapping, potentially improving the flexibility of individual measurements. We believe this to be a unique and promising new method of tonometry with the potential to improve the accessibility of glaucoma screening technology, acting as one step to equitably reduce the impact of this preventable disease.

Appendix A: Racial Equity Impact Analysis

In addition to building a solely technical solution, our research project, like many others, has a goal of establishing a lasting positive impact on society. In order to systematically examine how different groups will be impacted by our research, we have conducted a Racial Equity Impact Analysis (REIA) [136].

Our project seeks to remove barriers, both cost and convenience, from ocular tonometry with the goal of making it easier to more frequently monitor IOP. In particular, our method is designed to help monitor risk for Primary Open-Angle Glaucoma (POAG), the most common form of Open-Angle Glaucoma (OAG), itself the most common form of glaucoma, accounting for approximately 70 to 90% of all glaucoma cases [137, 138]. However, notably there are other less common forms of primary and secondary glaucoma such as normal-tension glaucoma that disproportionately affects individuals of Japanese ancestry, angle-closure glaucoma, congenital glaucoma which occur in newborn babies, and pigmentary glaucoma that primarily impacts young, Caucasian males [138].

In this analysis we will examine differences in POAG for people of various heritages, ages, gender, as well as other factors and consider possible impacts that our research can have. Though we briefly discuss the state of OAG in nations outside of the United States, much of our focus in this assessment will be on racial equity impact within the United States. We also discuss potential applications of our project in real-world communities; however, much of this discussion is beyond the scope of our project.

A.1 Background

A.1.1 Identifying Stakeholders

Currently in 2020, among the age pool with the greatest prevalence in glaucoma (ages 40 to 80) there were approximately 76.0 million cases of glaucoma worldwide, up from an estimated 64.3 million in 2013, and projected to increase to 11.8 million in the year 2040 [139], impacting every nation, with certain demographic populations at greater risk than others. Because these statistics do not include individuals younger than forty years old or older than eighty years old, we expect a greater number in truth than aforementioned. Part of identifying stakeholders for our project relies on knowing the change in risk factors between different racial, ethnic, gender, age, and habitation groups, as our research will have more impact on groups with higher risk for POAG. In addition, as our research seeks to lower barriers for all individuals in accessing regular IOP checkups, we will focus on high-risk populations such as low-income communities and underdeveloped or socioeconomically disadvantaged communities within high-income areas. Many of these disproportionately impacted communities do not have reliable access to tonometers or ophthalmologists.

A.1.1.1 Risk Factor Impact On Prevalence of POAG/OAG

Many studies on the prevalence of glaucoma, OAG, and POAG exist, each with a slightly different focus or result. Many studies commonly recognize that increasing age increases risk for OAG anywhere from 50 to 150% per decade [139–141] for all groups. In fact, studies that examine the influences of socio-demographic factors such as age, gender, race, or ethnicity begin measuring prevalence at age forty [139, 140, 142]. Studies generally agree that glaucoma occurs in 2 to 5% of adults around age forty and 8 to 10% of adults aged eighty and above [139, 142–144].

Many studies agree that OAG disproportionately affects racial minorities within the United States [18]. Within the past 20 years, more effort has been made to distinguish aspects of glaucoma development in racial minority populations throughout the United States, typically the Black, Hispanic, and more recently, Asian populations. For example, one study showed that although Black populations demonstrate a higher prevalence for glaucoma overall than Caucasian populations, Caucasian populations have a faster increase of risk in glaucoma than Black and Asian populations, with the odds ratio of OAG increasing by about 2 per decade for Caucasian populations and about 1.6 per decade for Black and Asian populations [140]. Other studies have shown that blindness from glaucoma is at least six times more prevalent in Black Americans than in Caucasian Americans, and that glaucoma prevalence is three times higher in Black Americans and Hispanic Americans with Mexican ancestry compared to non-Hispanic Caucasian Americans [141]

Socio-demographic factors aside from race should be considered that also may influence the odds of developing glaucoma. Two frequently studied factors are sex and geographic area. Many studies agree that men are more likely to develop glaucoma than women, by about 1.5 times [139,140]. Also, a correlation has been found between urban and rural populations, where individuals living in urban areas are 1.58 times more likely to develop glaucoma than those living in rural areas [139].

A.1.1.2 Other Stakeholders To Consider

There are many other stakeholders to consider aside from socio-demographic groups that have a higher risk of having or developing glaucoma.

One of the most important aspects of managing glaucoma after diagnosis are follow-up appointments to monitor glaucomatous progression and administer medicine [145]. Patients more adherent to a determined follow-up medical schedule have more favorable outcomes than those who are not [146].

One study conducted in Baltimore, Maryland found three of the largest factors correlated with inconsistent follow-ups were Black race, Latino ethnicity, and unfamiliarity with treatment duration [146]. Similarly, one study conducted across many sites throughout Canada found that the presence of misunderstandings, hesitancy, and lack of formal education led to increases in noncompliance with following prescribed medication and treatment procedures, such as patients improperly self-administering eye drops [145].

Therefore, we perceive that the problem is on a much grander scale than simply a genetic, socioeconomic, or geographic predisposition to developing glaucoma.

Many people do not have enough education to understand the treatment process, and unknowingly allow their glaucoma to progress. This lack of education or understanding can often be associated with the disproportionate disadvantages certain socio-demographic groups and communities often have. For instance, socioeconomic status (SES) is a significant factor that influences the accessibility of resources, level of education, level of cognitive competency, financial security, and other privileges and opportunities [147]. In fact, low SES does not only contribute to fewer opportunities for wealth accumulation and financial security in afflicted individuals, but also lower education attainment, poverty, and poor health [147]. In addition, deep gaps between SES in both the United States and at the global level must be addressed and mitigated to benefit society as a whole. SES collectively contributes to and allows for the lingering resource, financial, quality of life, and lastly but utmost, health distribution inequities that negatively affect the human population globally [147].

These factors are deeply intertwined with many other factors such as demographics, politics, economics, education, psychological health and well-being, and many others, making it an extremely difficult singular issue to tackle [147]. In fact, many adolescents raised in a household with lower SES will have a (1) slower cognitive development and lower literacy rate, reducing realization that glaucoma leads to vision loss, as well as the prevalence of noncompliance in following ophthalmologists' or health professionals' recommendations, (2) greater accrued financial debt and lower financial income, posing an economic barrier to diagnostic and treatment options, (3) reduced accessibility to resources and/or leisure time due to working to survive, leading to high levels of psychological stress, potentially making one hesitant to receiving regular check-ups, or simply ignoring symptoms because they do not have the time, knowledge, or finances to make an appointment and receive healthcare [145–147].

Additionally, some people do not have the means to access health care or other societal infrastructures, such as undocumented immigrants, uninsured patients, and incarcerated individuals. [148] These indigent individuals have been deprived of or denied access to healthcare. For example, Jacob Wilensky, MD mentions

> "Care of glaucoma patients with limited financial resources is both challenging and frustrating. These patients are more time-consuming and costly to treat, but generate less or even no compensation. These patients need care, and we need to provide that care"

during his address to the joint meeting of the American Academy of Ophthalmology and the Pan-American Association of Ophthalmology [148]. The major problem is a disconnect between the needs of the ophthalmologists and these patients. Ophthalmologists should provide care to all glaucomatous patients in need, yet even in this country, there are individuals who have an additional barrier blocking them from the healthcare that they require.

According to Teresa C. Chen, MD, an associate professor of ophthalmology at the Harvard Medical School, Massachusetts Eye and Ear Infirmary Glaucoma Service, "The current global economic crisis underscores the urgent need for a better solution for both the physician and, most importantly, the patient" [148]. Hence, we would like to bring to attention this deeply ingrained problem in the United States where doctors' economic interests and professional responsibilities may be at odds when it comes to treating patients who cannot afford care. Tackling this problem would require an overhaul of the current global and national healthcare system, which adds an additional challenge should our proposed system be implemented in the future.

Problems with healthcare accessibility are not limited to indigent individuals, but also apply to the general American population. According to a 2016 Commonwealth Fund International Health Policy Survey, around 33% of Americans went without recommended care, did not see a doctor when sick, or did not buy prescription pills [149]. Furthermore, the United States has an incarceration rate of 716 out of 100,000 individuals, which is approximately five times greater than the United Kingdom with the second highest incarceration rate of 147 out of 100,000 individuals [150]. The United States also has a large population of undocumented immigrants, which make up "10.5 million to 12 million, or approximately 3.2%–3.6% of the population." [150] Incarcerated individuals and undocumented immigrants do not have adequate access to healthcare. Most incarcerated individuals are disproportionately black, when compared with the racial distribution in the general American population. The Bureau of Justice Statistics reports that of state prisoners, 35% are Caucasian, 38% are Black, and 12% are Hispanic [151]. Often undocumented immigrants have lower educational attainment, as a majority of the undocumented population are from Mexico or Central America, which have by far the smallest proportion of individuals with a bachelor's degree: 7% and 11% respectively, in

2018 [152]. Additionally, nearly all were non-white, as "of 9.75 million people, about 700,000 were white (not Hispanic)", according to Pew [153]. Therefore, we see race play a role in who makes up the population of incarcerated individuals or undocumented immigrants.

There are clear disparities in healthcare treatment between different socio-demographic communities around the United States. And beyond the general social inequities that are found throughout all aspects of healthcare in the United States, several studies have demonstrated specific inequities regarding glaucoma care and blindness in different populations, spanning from race to insurance coverage, all of which are part of the intricate web of identity.

Moving in this direction of healthcare accessibility, a study performed in 2018 analyzed the testing and treatment of people with commercial health insurance or Medicaid recipients [18]. For those with newly-diagnosed OAG, the study found those with Medicaid were 234% more likely to not receive any glaucoma testing in the 15 months following initial diagnoses, a portion of time which is crucial for monitoring the disease [154]. This difference was observed across all races and ethnicities in the study, but the researchers found that it was most notable for Black people when accounting for confounding factors, who had 291% higher odds of receiving no glaucoma testing compared to Black people with commercial health insurance. This is compared to the 198% higher odds of receiving no glaucoma testing for Caucasian people. This inequity demonstrates the great differences in ophthalmologic care not only between races and ethnicities, but also between insurance providers, exacerbating the difference in care between SES.

Many of these indigent individuals are also members of other socio-demographic groups such as by race, gender, sexual orientation, and ableism, all of which will be elaborated [147].

There are multiple groups to consider in terms of lower SES and their internalized racial identity such as African Americans, Latinx Americans, Native Americans, and Asian Americans. Bluestein elaborates in his book, “The Psychology of Working: A New Perspective for Career Development, Counseling, and Public Policy”, that the first two aforementioned groups often face challenges in which many employers and service providers will have a “sense of fear, social distance, and at times pure antagonism” towards these employees or consumers [155]. As for the latter two groups often there is a perceived difference in cultural attributes and views towards Western society and lifestyles. [155]. In the case of Asian Americans, there is specifically a model minority myth, but case studies have shown that contrary to this ideality, “people from Asian backgrounds face racism, prejudice, and stereotyping in their educational and working lives” [155].

Additionally, individuals of handicap or disabilities often can be associated with lower SES. Many of these individuals also are perceived differently externally, just like previously mentioned with racial identity and face similar difficulties in navigating the American social infrastructure [155]. Also, members of the LGBTQIA+ communities will not only have similar challenges but also a perceived “discrimination and stigma in the psychological experience of one’s working life” [155]. Many of them are trying to navigate their identity and may be conflicted between their professional work life and their personal identities. In essence, many socio-demographic factors impact peoples' access to social infrastructure such as healthcare and are intimately related to SES.

There are also several studies that display the correlation between region and visual impairment and blindness [156, 157]. In addition to broader disparities in equality across large communities like nations or the world, there are regional differences in prevalence of visual impairment and blindness. One discussed correlation for these differences is economic development in each region, which itself may be correlated with factors such as race [158].

To conclude, there are many groups of people that would benefit from lowered barriers to measurement of IOP—barriers ranging from long wait times to the much more restrictive obstacles such as decreased literacy, inaccessibility to resources, and financial instability that face individuals from a lower socioeconomic class [146]. These inequalities are well documented, both quantitatively and qualitatively. In the United States, racial groups are impacted by OAG differently and race is a factor in received care for glaucoma, yet institutional barriers also play a role in the accessibility of glaucoma diagnosis and treatment.

A.1.2 Engaging Stakeholders

Recently, more studies have been done in order to particularly analyze the relationship between glaucoma and traditionally under-studied populations. However, due to the constraints of this research project, we have been unable to consult many stakeholders due to our larger focus on the implementation of new technology.

It is our hope that, should our technology become fully fleshed out, we are able to work with stakeholders in order to identify ways to best help communities whether through outreach or collecting training data or other methods, but much of engaging stakeholders has remained beyond the scope of our project.

A.1.3 Examining the Causes

As we have discussed in the previous section, there is agreement that there are inequalities in the treatment and prevalence of visual impairment, blindness, and OAG. Several pieces on this topic discuss the idea of a Social Gradient, which "refers to the fact that inequalities in population health status are related to inequalities in social status" [159].

There are many complex and related causes for the development and impact of a Social Gradient with respect to glaucoma and eye-related care. Especially for issues such as diseases, where genetic disposition interacts with environmental factors for development, progression, and harm, many causes are present. As we have described in previous sections, several studies do believe that there are biological differences between different groups, like racial and sex, which place people at greater risk for glaucoma, taking other factors into account [140]. However, even when considering genetic factors, environmental factors controllable by society act as causes as well. For example, the fact that the rate of glaucoma is higher in urban settings than rural settings can indirectly impact already marginalized populations that are more likely to live in urban settings [139].

A.2 Our Project

Our project has the goal of being applicable in a variety of settings. We hope that with a streamlined and forgiving design, our device helps expand the amount of opportunities to monitor IOP.

A.2.1 Clarifying The Purpose

Our project seeks to help discover methods to accurately and inexpensively monitor IOP in a way that does not require expertise. One benefit of developing cheaper and less invasive methods of measuring IOP is that it would not require a physician to be present in order to measure IOP. This would allow for communities which do not have access to many physicians to have access to machines, and lower the barrier to access these tools. This should reduce disparities in healthcare in communities with less access because it would provide them with more opportunity to monitor their health.

One way that our neural network may help reduce any existing discrimination in this field is in its adaptable analysis. As an example, if an ophthalmologist is trained to measure IOP from one ethnic group, then they risk a biased or inaccurate measurement for the eyes of other ethnic groups due to genetic differences in the structure of the cornea. Our neural network approach could offer adaptability to different corneal parameters.

A.2.2 Considering Adverse Impacts

Every research project contains impacts, both foreseen and unforeseen. In addition, the adoption and implementation of a new technology such as this one could have a large influence on how it impacts society, so we will explore adverse impacts of both the technology itself and the possible implementations of the technology.

The main technologically differentiating factor of our product is its use of machine learning via two neural networks. The first network "cleans up" input images, and the second network models the shape of the eye during the deformation as a series of points in three dimensions. The main inputs that can change the results of our neural networks are the shape of the eye, the intensity and placement of the air puff, and the deformation of the eye from the air puff, which is tracked using a green lattice projected onto the cornea. Again, although this method of tracking deformation is novel, its output is functionally the same as current tonometers. With that in mind, one adverse impact this could have is the way in which light reflects off the eye. We have assumed throughout this process that although populations may have different eye properties and propensities for diseases like glaucoma, in general the amount of light reflected from eyes stays constant enough for us to map the shape of eyes. This of course may not necessarily be the case. Different eye colors may reflect different amounts of light, which in turn may skew the results for different ethnic groups who are more likely to have a specific eye color, however this was beyond the scope of our project. Additionally, there are documented racial biases in AI that arise from the networks being trained on a narrow data set [160].

We believe that our system is more robust to such biases because it was trained on only corneas, which display lower variance between racial and ethnic groups.

It is our goal to create a technology that is more accessible and usable than existing tonometry systems, ideally removing the need for an ophthalmologist at every tonometry reading. However, this could potentially cause adverse impacts.

Results from a survey of glaucoma patients in North India demonstrate that many people who receive a diagnosis for glaucoma after getting their IOP measured through tonometers do not have enough education or literacy to adequately understand their condition [161]. If people feel that with measurement tools, ophthalmologists are not needed, then they risk not fully understanding their condition and receiving improper treatment. It should be noted, however, that consultation with an ophthalmologist is still necessary. Though our proposed method allows for glaucoma screening without the direct, hands-on support of a medical professional, care should be taken to ensure that professionals remain available to assist in interpretation of the results and subsequent treatment decisions. This system does not replace the need for an ophthalmologist to diagnose and treat glaucoma, it merely simplifies the screening process so that regular checkups are more widely available to populations with the greatest needs.

A.2.3 Advancing Equitable Impacts

We will explore possible advanced equitable impacts in the same form as we explored adverse impacts, first through the technology itself and next through implementation.

As mentioned, deep neural networks have already shown great promise in corneal disease classification and detection. Our project attempts to apply neural networks to corneal topography and tonometry, and presents potential advantages for accuracy and efficiency. To date, corneal topographers like the Placido topographer use algorithms that do not take all the information produced into account, as in the case of the widely-used SimK algorithm, which attempts to reconstruct corneal topography from only one of the several Placido rings. Current non-contact tonometers also fail to consider all information available, as they rely on the two-dimensional cross-sections produced by the Scheimpflug imaging technique and consequently incur assumptions of radial symmetry to accurately analyze apex motion. The Scheimpflug technique also introduces the need for a moving camera, which adds additional complexity and cost to any device implementing it. As our method produces a three-dimensional map of the cornea using a stationary camera, it removes all assumptions about the geometry of a corneal deformation, and allows for potential analysis methods that may accurately locate the apex of an applanation, and incorporate the motion of the surrounding cornea to aid analysis. In addition, neural networks select spatially invariant, detailed features across an image, and therefore may take all details about available patterns into account. This creates the potential for much more accurate analysis using our methods, with fewer assumptions than in previous methods. In addition, the removal of the rotating camera track allows for greater efficiency and less manufacturing complexity.

In order to truly provide positive impacts, the method of implementation of this technology is crucial. We have proposed and developed a cheaper method to more flexibly measure IOP in a way that may be easier to more equitably distribute than currently available tonometers.

Currently we have only developed a bare-bones camera and projector system, which requires a housing in the future. However, some components of this setup can be adapted. The form of the projector grid is crucial, but the camera itself and imaging software are able to be adapted to the needs of a given community up to a certain point. As we have mentioned before, this can help to close the healthcare gap for communities of color who are most at-risk of developing glaucoma. In the future, if our research were to continue to develop, then it could eliminate the need for a trained specialist to deliver a tonometry test, and this important health screening practice could be given in a widely used general clinic or a smaller private practice that serves those communities. The lowered cost of our proposed measurement device also means that appointments could be cheaper and more facilities could acquire the device, both leading to more accessible visual health care.

In addition, though the fitting of the neural network requires a large amount of computing power, once fit, the network requires much less power to classify a single image. Either way, the goal is to remove barriers to measuring IOP, and perhaps it can be developed to the point where a user can operate the tonometer without a medical care provider at all, if the computing system provides feedback.

Even if a medical care provider is still required, the added flexibility of measurement removes some inequities in implementation: each measurement potentially takes less time because the imaging system does not require careful centering, and it is less likely to need repetition. Therefore, any given medical care provider can accommodate a larger number of patients per day.

A.2.4 Examining Alternatives or Improvements

As previously mentioned, unity between technology and implementation is crucial in order to make the most positive impact. With this in mind, there are several improvements that can be considered.

The first step is to consult with people who specialize in racial equity and inclusion in the healthcare field, particularly with eye care, such as ophthalmologists. Additionally, we can talk to public healthcare and social workers that bridge the gap between the patients and health professionals. Additionally, better educating oneself on glaucoma and understanding its risks in causing preventable permanent blindness and necessary treatments, as well as sharing this information with others can help bring attention to glaucoma being a real issue. We want to increase the awareness and health literacy, specifically around glaucoma so that even patients who have been ignored and mistreated by the current American healthcare and social infrastructure such as those from racial minority group to those who are undocumented or on Medicaid can decide to get diagnosed and treated for glaucoma earlier. Additionally, increasing awareness would allow for these patients to be more compliant with all the treatment plans prescribed. Taking the step to lower the barriers for these individuals to find ophthalmologists and have access to healthcare would help lower the global crisis of people going permanently blind when glaucoma is treatable when diagnosed early.

Lobbying the government and current American and global healthcare infrastructures to do more for people with greater obstacles to healthcare will help also with enabling ophthalmologists to treat indigent individuals who cannot access or afford healthcare. This would benefit society as a whole, so although a first-world country, the United States would have a lower gap between people of different socioeconomic and socio-demographic backgrounds, and that the United States does not brag of a large institutional barrier to healthcare, even if the United States does have many medicinal and technological research advancements in the world.

Next, another improvement would be to continually analyze our system in relation to various socio-demographic groups, in order to provide scientific results on the accuracy of our system as a function of race. Currently there is a wealth of studies about occurrences of POAG between groups, however, there are not many studies on the efficacy of measurement techniques on socio-demographic groups.

Appendix B: Relevant Code and Repositories

We have several GitHub repositories under the name `https://github.com/teamcontact`. Our code for CheckMark may be found at `https://github.com/teamcontact/extract-checkers`, and our code for DeepSquish may be found at `https://github.com/teamcontact/checker2cornea`.

Diagrams of neural network architectures in this paper were generated using Haris Iqbal's `PlotNeuralNet` library, which may be found at `https://github.com/HarisIqbal88/PlotNeuralNet`.

Appendix C: Modeling Process

C.1 Modeling

We scaled a hemisphere along the z-axis by a factor of 0.3, producing a flattened surface with approximate similarities to a cornea. We neglected to give the surface any thickness, as it would negligibly affect its reflective properties. The final mesh had 241 vertices.

C.2 Lighting

Blender allows for the simulation of projected images using its shader node editor. Specifically, one can link an image texture to the Strength input of the Emission node of a Spot lamp. The scaling and offset of the projected texture can be changed using a Mapping node, and the aperture size of the projector can be simulated by changing the Size property of the lamp, for varying depths of field.

C.3 Deformations

We chose to create surface deformations using the Displace modifier. While other methods exist to deform meshes, the Displace modifier modifies actual mesh geometry, allowing for the export of Wavefront OBJ files and therefore a usable ground truth. The Displace modifier deforms a surface based on the intensity of an input texture, which may be either procedurally generated or image-based. We chose to use procedurally generated noise textures, as they would allow us to create arbitrary amounts of complex deformation data. In particular, we took advantage of 4D Perlin noise. By sampling along the time axis, we obtain a smooth, continuous series of deformations.

C.4 Simulated Movement and Parameters

We trained the network to deal with variations in angle and distance. We neglected to vary focal length, as the resulting changes in image structure are negligible at such a shallow depth of field and can simply be simulated via data augmentation. Our simulated camera has a focal length of 50 millimeters, and begins at a 45 angle to the corneal surface, 3.68 meters away from the cornea (the cornea is modeled as a circular surface with a radius of 1 meter). Using Blender's Noise modifier, we apply random rotations to the camera.

In a similar manner, we also randomly change the angle of projection onto the cornea by randomizing the rotation of the simulated "projector" relative to the cornea. Additionally, we vary the scale of the projected grid.

Bibliography

[1] Madhvi Deol, David A. Taylor, and Nathan M. Radcliffe. Corneal hysteresis and its relevance to glaucoma. *Current Opinion in Ophthalmology*, 26(2):96–102, March 2015.

[2] Genevieve A. Secker and Julie T. Daniels. Limbal epithelial stem cells of the cornea. In *StemBook*. Harvard Stem Cell Institute, Cambridge (MA), 2008.

[3] OD Natalie A. Townsend. The do's and don'ts of measuring iop, May 2015.

[4] De novo classification request for sensimed triggerfish, May 2014.

[5] Yushi Zhang, Yufeng Chen, Tianxing Man, Dong Huang, Xiao Li, Hongwei Zhu, and Zhihong Li. High resolution non-invasive intraocular pressure monitoring by use of graphene woven fabrics on contact lens. *Microsystems & Nanoengineering*, 5(1), August 2019.

[6] icare products and services for modern and reliable ocular diagnostics, Jan 2021.

[7] Hyung Jin Kim, Yeong Ho Seo, and Byeong Hee Kim. New intraocular pressure measurement method using reflected pneumatic pressure from cornea deformed by air puff of ring-type nozzle. *PLOS ONE*, 12(12):e0186738, December 2017.

[8] M.D. William Trattler. Keratoconus causes, symptoms and 10 treatment options, Apr 2020.

[9] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins, 2020.

[10] Aram Faghfouri and Michael Frish. Robust discrimination of human footsteps using seismic signals. *Proc SPIE*, May 2011.

[11] Manolis Loukadakis, José Cano, and Michael O'Boyle. Accelerating deep neural networks on low power heterogeneous architectures. January 2018.

[12] Andrey, Adrian Rosebrock, W.-C. Cheng, Mingxing wang, Johnny, Sumit Vaise, Jijou Farah, and Aghiles. Denoising autoencoders with keras, tensorflow, and deep learning, Apr 2020.

[13] T. Fan, G. Wang, Y. Li, and H. Wang. Ma-net: A multi-scale attention network for liver and tumor segmentation. *IEEE Access*, 8:179656–179665, September 2020.

[14] H. A. Quigley and A. T. Broman. The number of people with glaucoma worldwide in 2010 and 2020. *The British Journal of Ophthalmology*, 90(3):262–267, March 2006.

[15] Robert N. Weinreb, Tin Aung, and Felipe A. Medeiros. The Pathophysiology and Treatment of Glaucoma: A Review. *JAMA*, 311(18):1901–1911, May 2014.

[16] James M. Tielsch. Socioeconomic status and visual impairment among urban americans. *Archives of Ophthalmology*, 109(5):637, May 1991.

[17] Anne L Coleman and Gergana Kodjebacheva. Risk Factors for Glaucoma Needing More Attention. *The Open Ophthalmology Journal*, 3:38–42, September 2009.

[18] Angela R. Elam, Chris Andrews, David C. Musch, Paul P. Lee, and Joshua D. Stein. Large Disparities in Receipt of Glaucoma Care Between Enrollees in Medicaid and Those with Commercial Health Insurance. *Ophthalmology*, 124(10):1442–1448, October 2017.

[19] Ji Hyun Kim and Joseph Caprioli. Intraocular Pressure Fluctuation: Is It Important? *Journal of Ophthalmic & Vision Research*, 13(2):170–174, 2018.

[20] J. T. Wilensky, D. K. Gieser, M. L. Dietsche, M. T. Mori, and R. Zeimer. Individual variability in the diurnal intraocular pressure curve. *Ophthalmology*, 100(6):940–944, June 1993.

[21] Robert L. Stamper. A history of intraocular pressure and its measurement. *Optometry and Vision Science*, 88(1):E16–E28, January 2011.

[22] Felipe A. Medeiros, Daniel Meira-Freitas, Renato Lisboa, Tung-Mei Kuang, Linda M. Zangwill, and Robert N. Weinreb. Corneal hysteresis as a risk factor for glaucoma progression: A prospective longitudinal study. *Ophthalmology*, 120(8):1533–1540, August 2013.

[23] George Mangouritsas, George Morphis, Spyridon Mourtzoukos, and Elias Feretis. Association between corneal hysteresis and central corneal thickness in glaucomatous and non-glaucomatous eyes. *Acta Ophthalmologica*, 87(8):901–905, December 2009.

[24] Mohammad Mehdi Sadoughi, Bahram Einollahi, Neda Einollahi, Javad Rezaei, Danial Roshandel, and Sepehr Feizi. Measurement of Central Corneal Thickness Using Ultrasound Pachymetry and Orbscan II in Normal Eyes. *Journal of Ophthalmic & Vision Research*, 10(1):4–9, March 2015.

[25] Sophia Y. Wang, Ronald Melles, and Shan C. Lin. The impact of central corneal thickness on the risk for glaucoma in a large multiethnic population. *Journal of Glaucoma*, 23(9):606–612, December 2014.

[26] Abby Wilson and John Marshall. A review of corneal biomechanics: Mechanisms for measurement and the implications for refractive surgery. *Indian Journal of Ophthalmology*, 68(12):2681–2683, 2020.

[27] Moritz Winkler, Dongyul Chai, Shelsea Krilling, Chyong Jy Nien, Donald J. Brown, Bryan Jester, Tibor Juhasz, and James V. Jester. Nonlinear optical macroscopic assessment of 3-d corneal collagen organization and axial biomechanics. *Investigative Ophthalmology Visual Science*, 52(12):8818, 2011.

[28] Sally Hayes, Craig Boote, Jennifer Lewis, Jack Sheppard, Mohammad Abahussin, Andrew J. Quantock, Christine Purslow, Marcela Votruba, and Keith M. Meek. Comparative study of fibrillar collagen arrangement in the corneas of primates and other mammals. *The Anatomical Record*, 290(12):1542–1550, 2007.

[29] Abby Wilson, John Jones, John R. Tyrer, and John Marshall. An interferometric ex vivo study of corneal biomechanics under physiologically representative loading, highlighting the role of the limbus in pressure compensation. *Eye and Vision*, 7(43):1–3, 2020.

[30] J.S. Bell, S. Hayes, C. Whitford, J. Sanchez-Weatherby, O. Shebanova, C. Vergari, C.P. Winlove, N. Terrill, T. Sorensen, A. Elsheikh, and K.M. Meek. The hierarchical response of human corneal collagen to load. *Acta Biomaterialia*, 65:3–4, 2018.

[31] Ning-Jiun Jan, Bryn L. Brazile, Danielle Hu, Garrett Grube, Jacob Wallace, Alexandra Gogola, and Ian A. Sigal. Crimp around the globe; patterns of collagen crimp across the corneoscleral shell. *Experimental Eye Research*, 172:163–166, 2018.

[32] Brad L. Boyce, J. Mark Grazier, Reese E. Jones, and Thao D. Nguyen. Full-field deformation of bovine cornea under constrained inflation conditions. *Biomaterials*, 29(28):3903, 2008.

[33] Ahmed Elsheikh, Defu Wang, and David Pye. Determination of the modulus of elasticity of the human cornea. *Journal of Refractive Surgery*, 23(8):812–813, 2007.

[34] Nathaniel E. Knox Cartwright, John R. Tyrer, and John Marshall. Age-related differences in the elasticity of the human cornea. *Investigative Ophthalmology Visual Science*, 52(7):4326, 2011.

[35] Stuart Hodson. The regulation of cornea1 hydration to maintain high transparency in fluctuating ambient temperatures. *Experimental Eye Research*, 20(4):378, 1975.

[36] Sabine Kling and Susana Marcos. Effect of hydration state and storage media on corneal biomechanical response from in vitro inflation tests. *Journal of Refractive Surgery*, 29(7):493–494, 2013.

[37] C.L. Harper, M.E. Boulton, D. Bennett, B. Marcyniuk, J.H. Jarvis-Evans, A.B. Tullo, and A.E. Ridgway. Diurnal variations in human corneal thickness. *British Journal of Ophthalmology*, 80(12):1070–1071, 1996.

[38] Stephen R. Cohen, Kenneth A. Polse, Richard J. Brand, and Robert B. Mandell. Humidity effects on corneal hydration. *Investigative Ophthalmology Visual Science*, 31(7):1285–1286, 1990.

[39] Fan Lu, Suzhong Xu, Jia Qu, Meixiao Shen, Xiaoxing Wang, Haizhen Fang, and Jianhua Wang. Central corneal thickness and corneal hysteresis during corneal swelling induced by contact lens wear with eye closure. *American Journal of Ophthalmology*, 143(4):620–621, 2007.

[40] Kenneth A. Polse, Richard Brand, Robert Mandell, David Vastine, David Demartini, and Roanne Flom. Age differences in corneal hydration control. *Investigative Ophthalmology Visual Science*, 30(3):620–621, 1989.

[41] Elliot M. Kirstein, Ahmed Elsheikh, and Pinakin Gunvant. Tonometry – Past, Present and Future. *Glaucoma - Current Clinical and Research Aspects*, November 2011.

[42] Tarannum Mansoori and Nagalla Balakrishna. Effect of central corneal thickness on intraocular pressure and comparison of topcon CT-80 non-contact tonometry with goldmann applanation tonometry. *Clinical and Experimental Optometry*, 101(2):206–212, October 2017.

[43] Sean McCafferty, Jason Levine, Jim Schwiegerling, and Eniko T. Enikov. Goldmann applanation tonometry error relative to true intracameral intraocular pressure in vitro and in vivo. *BMC Ophthalmology*, 17(1), November 2017.

[44] Xiang Chen, Xingdi Wu, Xueqi Lin, Jingwen Wang, and Wen Xu. Outcome, influence factor and development of CLS measurement in continuous IOP monitoring: A narrative review. *Contact Lens and Anterior Eye*, October 2020.

[45] Atsuya Miki, Miho Kumoi, Naoyuki Maeda, Shizuka Koh, Kenji Matsushita, and Kohji Nishida. Transient changes in refractive error and corneal tomography after 24-h continuous monitoring of intraocular pressure patterns with a contact lens sensor. *Japanese Journal of Ophthalmology*, 64(2):127–133, February 2020.

[46] Naoki Tojo and Atsushi Hayashi. Influence of ocular dimensional change on 24-hour intraocular pressure measurement with contact lens sensor. *Journal of Glaucoma*, 28(9):808–810, September 2019.

[47] Xiao Li, Pengzhan Sun, Lili Fan, Miao Zhu, Kunlin Wang, Minlin Zhong, Jinquan Wei, Dehai Wu, Yao Cheng, and Hongwei Zhu. Multifunctional graphene woven fabrics. *Scientific Reports*, 2(1), May 2012.

[48] A Cervino. Rebound tonometry: new opportunities and limitations of non-invasive determination of intraocular pressure. *British Journal of Ophthalmology*, 90(12):1444–1446, August 2006.

[49] Wan sang Chui, Andrew Lam, Davie Chen, and Roger Chiu. The influence of corneal properties on rebound tonometry. *Ophthalmology*, 115(1):80–84, January 2008.

[50] Shunsuke Nakakura. Icare rebound tonometers: review of their characteristics and ease of use. *Clinical Ophthalmology*, Volume 12:1245–1253, July 2018.

[51] Shalini Mohan, Satyaprakash Tiwari, Arvind Jain, Jaya Gupta, and Surendra Kumar Sachan. Clinical comparison of pulsair non-contact tonometer and goldmann applanation tonometer in indian population. *Journal of Optometry*, 7(2):86–90, April 2014.

[52] I. Hashimoto. Somatosensory evoked potentials elicited by air-puff stimuli generated by a new high-speed air control system. *Electroencephalography and Clinical Neurophysiology*, 67(3):231–237, September 1987.

[53] Jiaxu Hong, Jianjiang Xu, Anji Wei, Sophie X. Deng, Xinhan Cui, Xiaobo Yu, and Xinghuai Sun. A new tonometer—the corvis ST tonometer: Clinical comparison with noncontact and goldmann applanation tonometers. *Investigative Opthalmology & Visual Science*, 54(1):659, January 2013.

[54] Parag D Gandhi, Rabia Gürses-Özden, Jeffrey M Liebmann, and Robert Ritch. Attempted eyelid closure affects intraocular pressure measurement. *American Journal of Ophthalmology*, 131(4):417–420, April 2001.

[55] Masato Matsuura, Kazunori Hirasawa, Hiroshi Murata, Mieko Yanagisawa, Yoshitaka Nakao, Shunsuke Nakakura, Yoshiaki Kiuchi, and Ryo Asaoka. The relationship between corvis ST tonometry and ocular response analyzer measurements in eyes with glaucoma. *PLOS ONE*, 11(8):e0161742, August 2016.

[56] James C. Tsai. High eye pressure and glaucoma, Oct 2017.

[57] P T Khaw. Glaucoma–1: Diagnosis. *BMJ*, 328(7431):97–99, January 2004.

[58] Harry A. Quigley and W. Richard Green. The histology of human glaucoma cupping and optic nerve damage: Clinicopathologic correlation in 21 eyes. *Ophthalmology*, 86(10):1803–1827, October 1979.

[59] Vinay Kansal, James J. Armstrong, Robert Pintwala, and Cindy Hutnik. Optical coherence tomography for glaucoma diagnosis: An evidence based meta-analysis. *PLOS ONE*, 13(1):e0190621, January 2018.

[60] Andrew J. Tatham and Felipe A. Medeiros. Detecting structural progression in glaucoma with optical coherence tomography. *Ophthalmology*, 124(12):S57–S65, December 2017.

[61] A. Gullstrand. Procedure of the rays in the eye. imagery-laws of first order. In *Helmholtzs treatise on physiological optics, Vol 1 (Trans. from the 3rd German ed.).*, pages 301–358. Optical Society of America.

[62] Mohamed Anbar, Engy Mohamed Mostafa, Ashraf Mostafa Elhawary, Islam Awny, Mahmoud Mohamed Farouk, and Amr Mounir. Evaluation of corneal higher-order aberrations by scheimpflug–placido topography in patients with different refractive errors: A retrospective observational study. *Journal of Ophthalmology*, 2019:1–5, June 2019.

[63] Gerard Cairns and Charles McGhee. Orbscan computerized topography: Attributes, applications, and limitations. *Journal of cataract and refractive surgery*, 31:205–20, 02 2005.

[64] Giacomo Savini, Piero Barboni, Michele Carbonelli, and Kenneth J. Hoffer. Repeatability of automatic measurements by a new scheimpflug camera combined with placido topography. *Journal of Cataract and Refractive Surgery*, 37(10):1809–1816, October 2011.

[65] Michael W Belin and Stephen S Khachikian. An introduction to understanding elevation-based topography: how elevation data are displayed - a review. *Clinical & Experimental Ophthalmology*, 37(1):14–29, January 2009.

[66] Evangelos Pateras and A Nousi. The reliability between placido topography and scheimpflug topography in normal corneas with astigmatism less than 0.50 dc. *Ophthalmology Research: An International Journal*, pages 27–35, 04 2020.

[67] John R. Levene. The true inventors of the keratoscope and photo-keratoscope. *The British Journal for the History of Science*, 2(4):324–342, 1965.

[68] Cristina M Oliveira, Celina Ribeiro, and Sandra Franco. Corneal imaging with slit-scanning and scheimpflug imaging techniques. *Clinical and Experimental Optometry*, 94(1):33–42, August 2010.

[69] Alexander Nguyen, Tiffany Liu, and Ji Liu. Applications of scheimpflug imaging in glaucoma management: Current and potential applications. *Journal of Ophthalmology*, 2016:1–9, 01 2016.

[70] Imaging the anterior segment: Scheimpflug analysis.

[71] Lyndon W. Jones, Sruthi Srinivasan, Alison Ng, and Marc Schulze. Diagnostic instruments. In *Contact Lens Practice*, pages 327–345.e5. Elsevier, 2018.

[72] Moshirfar Majid, Mahsaw Motlagh, Michael Murri, Hamed Momeni Moghaddam, Yasmyne Ronquillo, and Phillip Hoopes. Galilei corneal tomography for screening of refractive surgery candidates: A review of the literature, part ii. *Medical hypothesis, discovery innovation ophthalmology journal*, 8:204–218, 10 2019.

[73] Joel S. Schuman, Michael R. Hee, Adarsh V. Arya, Tamar Pedut-Kloizman, Carmen A. Puliafito, James G. Fujimoto, and Eric A. Swanson. Optical coherence tomography: A new tool for glaucoma diagnosis. *Current Opinion in Ophthalmology*, 6(2):89–95, April 1995.

[74] Joel S. Schuman. Quantification of nerve fiber layer thickness in normal and glaucomatous eyes using optical coherence tomography. *Archives of Ophthalmology*, 113(5):586, May 1995.

[75] Michael R. Hee. Optical coherence tomography of the human retina. *Archives of Ophthalmology*, 113(3):325, March 1995.

[76] Christopher K.S. Leung, Shi Lam, Robert N. Weinreb, Shu Liu, Cong Ye, Lan Liu, Jing He, Gilda W.K. Lai, Taiping Li, and Dennis S.C. Lam. Retinal nerve fiber layer imaging with spectral-domain optical coherence tomography. *Ophthalmology*, 117(9):1684–1691, September 2010.

[77] Mauro T. Leite, Harsha L. Rao, Robert N. Weinreb, Linda M. Zangwill, Christopher Bowd, Pamela A. Sample, Ali Tafreshi, and Felipe A. Medeiros. Agreement among spectral-domain optical coherence tomography instruments for assessing retinal nerve fiber layer thickness. *American Journal of Ophthalmology*, 151(1):85–92.e1, January 2011.

[78] Lucy Q. Shen Chhavi Saini. Monitoring glaucoma progression with OCT, 6 2020. (Accessed on 03/24/2021).

[79] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, November 2004.

[80] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. SURF: Speeded up robust features. In *Computer Vision – ECCV 2006*, pages 404–417. Springer Berlin Heidelberg, 2006.

[81] Nalpantidis Lazaros, Georgios Christou Sirakoulis, and Antonios Gasteratos. Review of stereo vision algorithms: From software to hardware. *International Journal of Optomechatronics*, 2(4):435–462, November 2008.

[82] Elena Bebeselea-Sterp, Raluca Brad, and Remus Brad. A comparative study of stereovision algorithms. *International Journal of Advanced Computer Science and Applications*, 8(11), 2017.

[83] J. Lee and C. Kim. Monocular depth estimation using relative depth maps. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9721–9730, 2019.

[84] Mark Nishimura, David B. Lindell, Christopher Metzler, and Gordon Wetzstein. Disambiguating monocular depth estimation with a single transient. In *Computer Vision – ECCV 2020*, pages 139–155. Springer International Publishing, 2020.

[85] Pierluigi Zama Ramirez, Matteo Poggi, Fabio Tosi, Stefano Mattoccia, and Luigi Di Stefano. Geometry meets semantics for semi-supervised monocular depth estimation. In *Computer Vision – ACCV 2018*, pages 298–313. Springer International Publishing, 2019.

[86] Shariq Farooq Bhat, I. Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. *ArXiv*, abs/2011.14141, 2020.

[87] Chang Shu, Kun Yu, Zhixiang Duan, and Kuiyuan Yang. Feature-metric loss for self-supervised learning of depth and egomotion. In *Computer Vision – ECCV 2020*, pages 572–588. Springer International Publishing, 2020.

[88] Clement Godard, Oisin Mac Aodha, Michael Firman, and Gabriel Brostow. Digging into self-supervised monocular depth estimation. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, October 2019.

[89] Ibraheem Alhashim and Peter Wonka. High Quality Monocular Depth Estimation via Transfer Learning. *arXiv e-prints*, page arXiv:1812.11941, December 2018.

[90] Lovedeep Gondara. Medical image denoising using convolutional denoising autoencoders. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*. IEEE, December 2016.

[91] Mizuho Nishio, Chihiro Nagashima, Saori Hirabayashi, Akinori Ohnishi, Kaori Sasaki, Tomoyuki Sagawa, Masayuki Hamada, and Tatsuo Yamashita. Convolutional auto-encoder for image denoising of ultra-low-dose CT. *Heliyon*, 3(8):e00393, August 2017.

[92] Diangang Wang, Wei Gan, Chenyang Yan, Kun Huang, and Hongyi Wu. Inception model of convolutional auto-encoder for image denoising. In *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 174–186. Springer International Publishing, 2020.

[93] Swarup Kr Ghosh, Biswajit Biswas, and Anupam Ghosh. SDCA: a novel stack deep convolutional autoencoder – an application on retinal image denoising. *IET Image Processing*, 13(14):2778–2789, November 2019.

[94] Hans E. Atlason, Askell Love, Sigurdur Sigurdsson, Vilmundur Gudnason, and Lotta M. Ellingsen. Unsupervised brain lesion segmentation from MRI using a convolutional autoencoder. In Elsa D. Angelini and Bennett A. Landman, editors, *Medical Imaging 2019: Image Processing*. SPIE, March 2019.

[95] L. Chen, Y. Xie, J. Sun, N. Balu, M. Mossa-Basha, K. Pimentel, T. S. Hatsukami, J. Hwang, and C. Yuan. 3d intracranial artery segmentation using a convolutional autoencoder. In *2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 714–717, 2017.

[96] Abhishek Pal, Manav Rajiv Moorthy, and A. Shahina. G-eyenet: A convolutional autoencoding classifier framework for the detection of glaucoma from retinal fundus images. In *2018 25th IEEE International Conference on Image Processing (ICIP)*. IEEE, October 2018.

[97] Jonathan Masci, Ueli Meier, Dan Cireşan, and Jürgen Schmidhuber. Stacked convolutional auto-encoders for hierarchical feature extraction. In *Lecture Notes in Computer Science*, pages 52–59. Springer Berlin Heidelberg, 2011.

[98] M. Chen, X. Shi, Y. Zhang, D. Wu, and M. Guizani. Deep features learning for medical image analysis with convolutional autoencoder neural network. *IEEE Transactions on Big Data*, pages 1–1, 2017.

[99] H. F. Pardede, E. Suryawati, R. Sustika, and V. Zilvan. Unsupervised convolutional autoencoder-based feature learning for automatic detection of plant diseases. In *2018 International Conference on Computer, Control, Informatics and its Applications (IC3INA)*, pages 158–162, 2018.

[100] Sam Van der Jeught and Joris J.J. Dirckx. Real-time structured light profilometry: a review. *Optics and Lasers in Engineering*, 87:18–31, December 2016.

[101] Jason Geng. Structured-light 3d surface imaging: a tutorial. *Advances in Optics and Photonics*, 3(2):128, March 2011.

[102] Song Zhang. High-speed 3d shape measurement with structured light methods: A review. *Optics and Lasers in Engineering*, 106:119–131, July 2018.

[103] Sam Van der Jeught and Joris J. J. Dirckx. Deep neural networks for single shot structured light profilometry. *Optics Express*, 27(12):17091, June 2019.

[104] Noel Alpins, James K.Y. Ong, and George Stamatelatos. New method of quantifying corneal topographic astigmatism that corresponds with manifest refractive cylinder. *Journal of Cataract and Refractive Surgery*, 38(11):1978–1988, November 2012.

[105] Noel Alpins A.M, James Ong, and George Stamatelatos. Corneal topographic astigmatism (cort) - a better measure of corneal astigmatism that corresponds to manifest refractive cylinder. *Clinical and Experimental Ophthalmology*, 41:28–28, 11 2013.

[106] F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958.

[107] Marvin Minsky and Seymour Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA, USA, 1969.

[108] Moshe Leshno, Vladimir Ya. Lin, Allan Pinkus, and Shimon Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Networks*, 6(6):861–867, January 1993.

[109] M.W Gardner and S.R Dorling. Artificial neural networks (the multilayer perceptron)—a review of applications in the atmospheric sciences. *Atmospheric Environment*, 32(14-15):2627–2636, August 1998.

[110] Fionn Murtagh. Multilayer perceptrons for classification and regression. *Neurocomputing*, 2(5-6):183–197, July 1991.

[111] Li Deng. A tutorial survey of architectures, algorithms, and applications for deep learning. *APSIPA Transactions on Signal and Information Processing*, 3:e2, 2014.

[112] S. Albawi, T. A. Mohammed, and S. Al-Zawi. Understanding of a convolutional neural network. In *2017 International Conference on Engineering and Technology (ICET)*, pages 1–6, 2017.

[113] S. Ben Driss, M. Soua, R. Kachouri, and M. Akil. A comparison study between MLP and convolutional neural network models for character recognition. In Nasser Kehtarnavaz and Matthias F. Carlsohn, editors, *Real-Time Image and Video Processing 2017*. SPIE, May 2017.

[114] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989.

[115] Asifullah Khan, Anabia Sohail, Umme Zahoora, and Aqsa Saeed Qureshi. A survey of the recent architectures of deep convolutional neural networks. *Artificial Intelligence Review*, 53(8):5455–5516, April 2020.

[116] D. Ciregan, U. Meier, and J. Schmidhuber. Multi-column deep neural networks for image classification. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3642–3649, 2012.

[117] Irwin Sobel. An isotropic 3x3 image gradient operator. *Presentation at Stanford A.I. Project 1968*, 02 2014.

[118] Weibo Liu, Zidong Wang, Xiaohui Liu, Nianyin Zeng, Yurong Liu, and Fuad E. Alsaadi. A survey of deep neural network architectures and their applications. *Neurocomputing*, 234:11–26, April 2017.

[119] Qi Wang, Yue Ma, Kun Zhao, and Yingjie Tian. A comprehensive survey of loss functions in machine learning. *Annals of Data Science*, 04 2020.

[120] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2017.

[121] Sebastian Ruder. An overview of gradient descent optimization algorithms. *CoRR*, abs/1609.04747, 2016.

[122] Alexandru Lavric and Popa Valentin. KeratoDetect: Keratoconus detection algorithm using convolutional neural networks. *Computational Intelligence and Neuroscience*, 2019:1–9, January 2019.

[123] Pierre Zéboulon, Guillaume Debellemanière, Magalie Bouvet, and Damien Gatinel. Corneal topography raw data classification using a convolutional neural network. *American Journal of Ophthalmology*, 219:33–39, November 2020.

[124] D. Jackson Coleman and Stephen Trokel. Direct-recorded intraocular pressure variations in a human subject. *Archives of Ophthalmology*, 82(5):637–640, 1969.

[125] Kaweh Mansouri and Tarek Shaarawy. Continuous intraocular pressure monitoring with a wireless ocular telemetry sensor: Initial clinical experience in patients with open angle glaucoma. *The British Journal of Ophthalmology*, 95(5):627–9, 2011.

[126] Matteo Leonardi, Peter Leuenberger, Daniel Bertrand, Arnaud Bertsch, and Philippe Renaud. First steps toward noninvasive intraocular pressure monitoring with a sensing contact lens. *Investigative Ophthalmology Visual Science*, 45(9):3113–7, 2004.

[127] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015.

[128] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. volume 9351, pages 234–241, 10 2015.

[129] T. Eelbode, J. Bertels, M. Berman, D. Vandermeulen, F. Maes, R. Bisschops, and M. B. Blaschko. Optimization for medical image segmentation: Theory and practice when evaluating with dice score or jaccard index. *IEEE Transactions on Medical Imaging*, 39(11):3679–3690, 2020.

[130] Benjamin Graham. Fractional max-pooling. *CoRR*, abs/1412.6071, 2014.

[131] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. CBAM: convolutional block attention module. *CoRR*, abs/1807.06521, 2018.

[132] Ramprasaath R. Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra. Grad-cam: Why did you say that? visual explanations from deep networks via gradient-based localization. *CoRR*, abs/1610.02391, 2016.

[133] Robert Koprowski. Corvis ST tonometer and the possibility of analysing corneal deformation dynamics during intraocular pressure measurement. In *Air-Puff Tonometers: Challenges and insights*. IOP Publishing, 2019.

[134] Bernardo T. Lopes, Ashkan Eliasy, and Renato Ambrosio. Artificial intelligence in corneal diagnosis: Where are we? *Current Ophthalmology Reports*, 7(3):204–211, July 2019.

[135] Qasim Farhood. Comparative evaluation of intraocular pressure with an air-puff tonometer versus a goldmann applanation tonometer. *Clinical Ophthalmology*, page 23, December 2012.

[136] Racial Equity Impact Assessment Toolkit, March 2009.

[137] Primary Open-Angle Glaucoma, March 2017.

[138] Types of Glaucoma | National Eye Institute.

[139] Yih-Chung Tham, Xiang Li, Tien Y. Wong, Harry A. Quigley, Tin Aung, and Ching-Yu Cheng. Global prevalence of glaucoma and projections of glaucoma burden through 2040: A systematic review and meta-analysis. *Ophthalmology*, 121(11):2081 – 2090, 2014.

[140] Alicja R. Rudnicka, Shahrul Mt-Isa, Christopher G. Owen, Derek G. Cook, and Deborah Ashby. Variations in Primary Open-Angle Glaucoma Prevalence by Age, Gender, and Race: A Bayesian Meta-Analysis. *Investigative Ophthalmology & Visual Science*, 47(10):4254–4261, October 2006. Publisher: The Association for Research in Vision and Ophthalmology.

[141] Bruce E. Prum, Lisa F. Rosenberg, Steven J. Gedde, Steven L. Mansberger, Joshua D. Stein, Sayoko E. Moroi, Leon W. Herndon, Michele C. Lim, and Ruth D. Williams. Primary Open-Angle Glaucoma Preferred Practice Pattern® Guidelines. *Ophthalmology*, 123(1):P41–P111, January 2016. Publisher: Elsevier.

[142] Glaucoma Tables | National Eye Institute.

[143] Naira Khachatryan, Maxwell Pistilli, Maureen G. Maguire, Rebecca J. Salowe, Raymond M. Fertig, Tanisha Moore, Harini V. Gudiseva, Venkata R. M. Chavali, David W. Collins, Ebenezer Daniel, Windell Murphy, Jeffrey D. Henderer, Amanda Lehman, Qi Cui, Victoria Addis, Prithvi S. Sankar, Eydie G. Miller-Ellis, and Joan M. O'Brien. Primary open-angle african american glaucoma genetics (poaagg) study: gender and risk of poag in african americans. *PLOS ONE*, 14(8):1–11, 08 2019.

[144] Joshua D. Stein, Denise S. Kim, Leslie M. Niziol, Nidhi Talwar, Bin Nan, David C. Musch, and Julia E. Richards. Differences in Rates of Glaucoma among Asian Americans and Other Racial Groups, and among Various Asian Ethnic Groups. *Ophthalmology*, 118(6):1031–1037, June 2011.

[145] Ramin Kholdebarin, Robert J. Campbell, Ya-Ping Jin, and Yvonne M. Buys. Multicenter study of compliance and drop administration in glaucoma. *Canadian Journal of Ophthalmology*, 43(4):454–461, August 2008.

[146] Yohko Murakami, Bradford W. Lee, Martin Duncan, Andrew Kao, Jehn-Yu Huang, Kuldev Singh, and Shan C. Lin. Racial and Ethnic Disparities in Adherence to Glaucoma Follow-up Visits in a County Hospital Population. *Archives of Ophthalmology*, 129(7):872–878, July 2011.

[147] Education and Socioeconomic Status Factsheet.

[148] Treating indigent glaucoma patients is financially difficult but necessary, clinician says.

[149] States of Incarceration: The Global Context | Prison Policy Initiative.

[150] Elaine Kamarck and Christine Stenglein. How many undocumented immigrants are in the United States and who are they?, November 2019.

[151] Ashley Nellis. The Color of Justice: Racial and Ethnic Disparity in State Prisons, June 2016.

[152] Abby Budiman. Key findings about U.S. immigrants, August 2020.

[153] Miriam Valverde PolitiFact.com. Fact check: Has ICE detained or deported any white immigrants? *Austin American-Statesman*, September 2019.

[154] Bruce E. Prum, Lisa F. Rosenberg, Steven J. Gedde, Steven L. Mansberger, Joshua D. Stein, Sayoko E. Moroi, Leon W. Herndon, Michele C. Lim, and Ruth D. Williams. Primary Open-Angle Glaucoma Preferred Practice Pattern® Guidelines. *Ophthalmology*, 123(1):P41–P111, January 2016. Publisher: Elsevier.

[155] David Blustein. *The Psychology of Working: A New Perspective for Career Development, Counseling, and Public Policy.* Routledge, January 2013. Google-Books-ID: Iwyh_d9m_i0C.

[156] Anna Rius Ulldemolins, Van C Lansingh, Laura Guisasola Valencia, Marissa J Carter, and Kristen A Eckert. Social inequalities in blindness and visual impairment: A review of social determinants. *Indian Journal of Ophthalmology*, 60(5):368–375, 2012.

[157] Anna Rius, Lucía Artazcoz, Laura Guisasola, and Joan Benach. Visual impairment and blindness in spanish adults: geographic inequalities are not explained by age or education. *Ophthalmology*, 121(1):408–416, January 2014.

[158] and Stacy A. Jay Shambaugh Anderson, Ryan Nunn. How racial and regional inequality affect economic opportunity, February 2019.

[159] M. G Marmot. Understanding Social Inequalities in Health. *Perspectives in Biology and Medicine*, 46(3):S9–S23, 2003. Publisher: Johns Hopkins University Press.

[160] Shawnasty Bankovich. Racial Bias in Machine Learning and Artificial Intelligence.

[161] Devindra Sood, Ishaana Sood, Shivam Sood, Dinesh Kumar, Rajeev Sood, and Narender N. Sood. Patients' perception of glaucoma diagnosis practice: Results from a survey of glaucoma patients in North India. *International Ophthalmology*, 40(7):1797–1805, July 2020.