Accommodation and Comfort in Head-Mounted Displays

GEORGE-ALEX KOULIERIS, Inria, Université Côte d'Azur
BEE BUI, University of California, Berkeley
MARTIN S. BANKS, University of California, Berkeley
GEORGE DRETTAKIS, Inria, Université Côte d'Azur

Head-mounted displays (HMDs) often cause discomfort and even nausea. Improving comfort is therefore one of the most significant challenges for the design of such systems. In this paper, we evaluate the effect of different HMD display configurations on discomfort. We do this by designing a device to measure human visual behavior and evaluate viewer comfort. In particular, we focus on one known source of discomfort: the vergence-accommodation (VA) conflict. The VA conflict is the difference between accommodative and vergence response. In HMDs the eyes accommodate to a fixed screen distance while they converge to the simulated distance of the object of interest, requiring the viewer to undo the neural coupling between the two responses. Several methods have been proposed to alleviate the VA conflict, including Depth-of-Field (DoF) rendering, focus-adjustable lenses, and monovision. However, no previous work has investigated whether these solutions actually drive accommodation to the distance of the simulated object. If they did, the VA conflict would disappear, and we expect comfort to improve. We design the first device that allows us to measure accommodation in HMDs, and we use it to obtain accommodation measurements and to conduct a discomfort study. The results of the first experiment demonstrate that only the focus-adjustable-lens design drives accommodation effectively, while other solutions do not drive accommodation to the simulated distance and thus do not resolve the VA conflict. The second experiment measures discomfort. The results validate that the focus-adjustable-lens design improves comfort significantly more than the other solutions.

CCS Concepts: • **Computing methodologies** → **Perception**; *Virtual reality*;

Additional Key Words and Phrases: head-mounted displays, perception, vergence-accommodation conflict

1 INTRODUCTION

Head-mounted displays (HMDs) often provoke discomfort and even nausea. In this paper we evaluate discomfort in HMDs for different display configurations by designing and constructing a measurement device. We focus on one important cause of discomfort, which occurs because in an HMD the eyes focus (accommodate) to a fixed screen distance while they rotate and align (converge) to the simulated distance of the object of interest, which might be in front of or behind the screen [Kooi and Toet 2004; Lambooij et al. 2009; Urvoy et al. 2013]. In the real world, the eyes accommodate and converge

to the same distance, and understandably the two responses are neurally coupled. The difference in accommodative and vergence response is called the *vergence-accommodation (VA) conflict.* This difference requires the viewer to undo the neural coupling [Cumming and Judge 1986; Martens and Ogle 1959], and is known to cause discomfort in many situations [Hoffman et al. 2008; Shibata et al. 2011].

An obvious avenue to improve HMD design is to eliminate the VA conflict by designing a display with the goal of matching accommodation to vergence, in which case comfort should improve. To evaluate how well such a display achieves this, two – undoubtedly interdependent – questions must be answered: (i) Does the display design actually drive accommodation to the distance of the simulated object? (ii) If it does, does it actually improve comfort?

Several display techniques have been proposed to alleviate the VA conflict; of these, only a few can be applied to HMDs. One proposed solution is gaze-contingent, depth-of-field (DoF) rendering (e.g, [Mauderer et al. 2014]) in which the focal plane of the DoF rendering is determined by the eyes' vergence distance. Another approach employs focus-adjustable lenses that change focal power depending on the distance of the object being observed in the virtual scene [Johnson et al. 2016; Konrad et al. 2016]. Yet another approach is monovision in which the focal distances of the two eyes are set to quite different values in an attempt to expand the range of distances for which the VA conflict is manageable [Johnson et al. 2016; Konrad et al. 2016].

Our first question – whether these designs actually drive accommodation – has never been answered. Concerning the second question, existing subjective assessments of discomfort (e.g, [Duchowski et al. 2014; Hoffman et al. 2008; Johnson et al. 2016; Mauderer et al. 2014; Shibata et al. 2011]) do not provide a conclusive answer to which design provides more comfortable viewing; we hypothesize this is due to the different durations of the experiments.

In this paper we introduce a device that allows us to measure accommodation in an HMD for each of these designs: gaze-contingent DoF rendering, focus-adjustable lenses and monovision. Our device has a modular design with focus-adjustable lenses and an autorefractor (see Sec. 2.3) to measure accommodation. We address several challenges related to precise calibration of the device to allow accurate measurements. We use our device to perform two experiments to answer the above questions for the proposed designs. The first study measures accommodation for each of the proposed designs: namely, DoF rendering, focus-adjustable lenses, and monovision. Our results answer the first question by demonstrating that a focus-adjustable-lens design is the only one that drives accommodation effectively. The other designs do not drive accommodation effectively and thus do not alleviate the VA conflict.

(a) Splittable HMD   **(b) Measurement Device**   **(c) Experiments**   **(d) Results**

HMD Screen

RGB Signal

IR Signal

Hot Mirror

Focus-Adjustable Lens

Offset Lens

Autorefractor Camera

1. Which configurations drive accommodation ?
2. Which configurations improve comfort ?

Fig. 1. We evaluate the effect of different Head-Mounted Display (HMD) display configurations on discomfort by developing a splittable HMD with focus-adjustable lenses (a) that works with an autorefractor to objectively measure accommodation (b). (c) We run experiments to evaluate the effect of different display configurations. (d) Our results allow us to answer fundamental questions about each configuration.

To answer the second question, we perform a study based on a well established comfort-evaluation protocol [Shibata et al. 2011] that requires lengthy viewing sessions. The study shows that the focus-adjustable-lens solution provides the largest improvement in comfort with respect to the conventional HMD design, which we use as a baseline.

Our contribution is three-fold. First, we present the first precisely calibrated modular device with focus-adjustable lenses, allowing accommodation to be measured in an HMD. Second, we use this device to measure accommodation in several configurations being considered for HMDs. Third, we show that the ability of a given design to drive accommodation to the distance of the simulated object predicts how comfortable the resulting experience will be.

Our methodology and results demonstrate the effectiveness of grounding discomfort studies in a carefully calibrated measurement-based apparatus. We provide important insight into choices for future HMD design, both by indicating the strong potential of focus-adjustable lenses and by revealing the limitations of other designs.

2 RELATED WORK

In this section we review relevant principles from human vision, describe various discomfort-alleviating techniques that have been used in HMDs, and discuss the autorefractor. Please see the accompanying video for further explanation.

2.1 Human Vision

Binocular fixation involves two oculomotor functions: vergence and accommodation. Vergence is the rotation of the eyes in opposite directions to align the eyes and obtain a single fused image of the fixated object. When vergence is inaccurate, double vision occurs. Accommodation is the adjustment of the power of the eye's crystalline lens to obtain a sharp retinal image (Fig. 2(b) and (c)). When it is inaccurate, blurred vision results. Thus, accurate vergence and accommodation are both required to achieve a single, clear image of a fixated object. The primary stimulus that drives vergence is binocular disparity (disparity-driven vergence), and the primary stimulus that drives accommodation is retinal-image blur (blur-driven accommodation). But vergence and accommodation are also neurally coupled. Specifically, accommodative responses evoke vergence (blur-driven vergence), and vergence responses evoke accommodation (disparity-driven accommodation) [Fincham and Walton 1957; Martens and Ogle 1959; Schor 1992].

Vergence-accommodation coupling is helpful in the real world because vergence and accommodative distances are almost always the same no matter where one looks. However, in conventional stereoscopic 3D displays, the eyes must instead converge and accommodate to potentially different distances because the distance of the accommodative stimulus is the screen distance while the distance of the vergence stimulus is that of the simulated object of regard, which is frequently nearer or farther than the screen. The difference in distances forces the visual system to override the neural coupling between vergence and accommodation. The VA conflict has for years been believed to be a significant cause of visual discomfort in stereo displays [Sheedy et al. 1993] and this has been recently demonstrated experimentally [Hoffman et al. 2008; Johnson et al. 2016; Kim et al. 2014; Shibata et al. 2011]. Yang and Sheedy [2011] measured accommodation in viewers watching movies on a 3DTV. They found that stereoscopic content (i.e., binocular disparity) had a small effect on accommodation. In concurrent work [Padmanaban et al. 2017], an autorefractor was used to acquire measurements of accommodation in near-eye displays; however, extensive and long-duration, comparative discomfort studies were not performed.

2.2 Display Designs to Address the VA Conflict

Several previous solutions have been presented to reduce the VA conflict in stereo viewing systems. We review three categories that apply to HMDs, but first briefly discuss multi-plane and light-field displays that currently do not provide satisfactory solutions for HMDs.

Multi-Plane and Light-Field Displays. Multi-plane displays present stimuli on multiple focal planes so that the stimuli to accommodation and vergence can be matched in distance (e.g., [Hu and Hua 2014; Love et al. 2009]). These displays require very precise positioning of the viewer's eyes relative to the display making the approach impractical for practical applications, including HMDs. Light-field displays are designed to reproduce the four-dimensional light field, allowing appropriate stereoscopic and parallax cues (e.g., [Maimone et al. 2013]). These displays currently lack sufficient spatial and angular resolution to simultaneously create acceptable image quality and correct focus cues [Maimone et al. 2013; Narain et al. 2015]. One

study has attempted to measure accommodation [Takaki 2006] in a light-field display, but presented no details on how it was measured, on the number of subjects, and whether accommodation was voluntary or involuntary. Compared to previous multi-plane displays, our HMD uses continuous focal adjustment rather than discrete, avoiding the compromises in multi-plane systems due to the trade-off between inter-plane separation and total workspace range.

Gaze-Contingent Depth-of-Field Blur. Several researchers have created displays with gaze-contingent, depth-of-field (DoF) blur in an effort to reduce discomfort [Duchowski et al. 2014; Mauderer et al. 2014; Otani et al. 2008; Vinnikov et al. 2016]. Gaze direction is measured in real-time and the focal plane for the DoF rendering is set to the intersection between the gaze vector and the simulated scene. When the viewer's eyes are directed to a near point in the scene that part of the scene is rendered sharp and farther points are rendered blurred. When the eyes are directed to a far point, the far point becomes sharp and the near parts blurred. Relative to infinite DoF, gaze-contingent DoF increases the amount of perceived depth [Mauderer et al. 2014], but it also worsens perceived image quality and visual comfort [Duchowski et al. 2014; Vinnikov et al. 2016]. Indeed, many subjects report that they dislike gaze-contingent DoF [Duchowski et al. 2014; Vinnikov et al. 2016]. Investigating whether DoF rendering works is important for graphics because it has a non-negligible computational cost that should not be borne if such rendering provides no clear benefit [Hillaire et al. 2008].

Monovision. As people age, they gradually lose the ability to accommodate. By their mid 50's, they essentially cannot accommodate at all: a condition called presbyopia [Duane 1912]. Monovision is an optometric method for dealing with presbyopia. The optical correction for one eye is made appropriate for far distance while the correction for the other eye is made appropriate for near [Evans 2007]. The idea is that the patient's percept will be dictated by whichever eye is in better focus for the currently fixated object thus providing apparently sharp vision across a greater range of distances. A closely related approach has recently been applied to conventional stereoscopic displays [Johnson et al. 2016] and HMDs [Konrad et al. 2016]. Lenses of different powers are placed in front of the two eyes, creating a "near" eye and a "far" eye. Accommodation in humans is completely yoked between the two eyes [Ball 1952; Campbell 1960; Fisher et al. 1987], so when the viewer looks at a stereoscopic display, the retinal image in one eye will generally be in better focus than in the other. The hope is that when the disparity of the fixated object specifies a far object, both eyes will accommodate far; and when the object is near, that both eyes will accommodate near. If this occurred, the vergence and accommodative responses would be more similar than they are in conventional stereoscopic displays, and this would reduce the VA conflict. The results of our study of monovision will be very useful because if the VA conflict really is reduced with this technique, it would offer a simple and inexpensive solution to improve comfort in HMDs.

Focus-adjustable Lenses. In the real world, light rays emanating from a point at some distance z in front of a lens with a given focal power are focused on the opposite side of the lens, illustrated in Fig. 2(a). This results in a sharp retinal image for the cyan object in the figure and a blurry image for the orange object.

In Fig. 2(b), the eye focuses at the object at distance z. In 2(c), the eye focuses at a farther distance z' so the eye's focal power changes; notice the change in the shape of the lens. On the right, we see the situation with a focus-adjustable lens placed between the eye and screen. Initially (d), the focal distance of the adjustable lens is at infinity, and the eye focuses on the screen, placed here at distance z. In (e), we change the focal power of the focus-adjustable lens so that the retina receives an in-focus signal from the screen, and the *eye's* lens power is the same as was required to focus at z' in the real world (the eye lens shape is the same as (c)). Accommodation is thus consistent with the object distance and the image from the screen is sharp, eliminating the VA conflict. Details of the optics are presented in supplemental material.

This principle is used in the adjustable-lens design by adjusting the powers of lenses in front of each eye in synchrony with the simulated distance of fixated objects. This equates vergence and accommodation distances and thereby eliminates the VA conflict. The design requires a reasonably accurate estimate of the viewer's fixation distance, which has to date been accomplished by instructing the viewer to fixate a particular object as it moves through the simulated space [Johnson et al. 2016; Konrad et al. 2016]. In our experiments, we examined how effectively accommodation is driven with focus-adjustable lenses and how the design configuration affects visual comfort.

2.3 Autorefractor

An autorefractor is a computer-controlled device used during optometric eye examinations to measure a patient's refractive error (i.e., myopia, hyperopia, astigmatism). The device delivers infrared light into the eye, and the current accommodative state (i.e., the distance to which the eye is best focused) is measured from the refraction of the reflected light exiting the eye. We used a Grand Seiko WAM-5500 autorefractor that makes dynamic focus measurements at a rate of ~5Hz. In this device, a ring of invisible infrared light is projected onto the retina using an internal motor-driven neutralizing optometer. The size and shape of the reflected ring are analyzed to measure respectively the spherical state (overall focus distance) and cylindrical (astigmatic) state. We were only interested in the spherical state because that corresponds to the eye's accommodative state [Win-Hall et al. 2010]. Our results were not affected by a cylindrical component; none of the subjects had an uncorrected astigmatism greater than 0.5D.

3 HMD AND MEASUREMENT DEVICE

The first challenge we faced was to design a device that would allow us to measure accommodation for the various proposed designs, and also allow a discomfort study for HMDs. This required a modular design and a sequence of precise calibration steps to ensure high-quality rendering and measurements. During the accommodation experiments the device is mounted on a bench because the autorefractor is heavy and bulky and thus cannot be mounted on the head. However, the setup still maintains all essential properties of a head-mounted display. All HMD components except head tracking exist and function normally in our device: the user views a

Fig. 2. Accommodation in the real world and in a focus-adjustable-lens display. Left: Accommodation to a near object: the lens has a given focal power so the cyan object is in focus, while the orange object is out of focus and thus blurry. Middle: To maintain a sharp image on the retina, the power of the eye lens changes (see difference in eye lens shape from (b) to (c)). Right: Accommodation to a stimulus presented on the display screen with an external focus-adjustable lens in front of the eye. When the focus-adjustable lens changes power (e), the eye lens must change power to maintain sharp focus on the retina. As a result, the eye lens has the same power as in (c), "tricking" the eye to be in the same accommodative state as when viewing a real object placed at distance z' which is farther than the screen .

near-field display via lenses and as such light reaches the eye as in a conventional HMD. During the discomfort study the display was head-mounted, but not head tracked. Throughout this section we outline the procedures involved in our design and calibration. More details on different steps can be found in the supplemental material.

3.1 Modular Hardware Design

We had two main design goals: 1) create a modular device that can serve both as an HMD and an accommodation-measurement device, 2) use consumer-level components as much as possible. The display was a Koolertron 5.6" 1280x800 TFT-LCD panel with a 60Hz refresh rate (Model LED956), mounted on a plastic casing for HMDs (Andoer 3DVR). The only relatively expensive component is the pair of focus-adjustable lenses, specifically Optotune EL-10-30 VIS LD lenses with USB controllers, similar to [Konrad et al. 2016]. The horizontal and vertical fields of view were 30°. The controller sends an electrical current value to adjust the lens to a desired focal power. A key difference with Konrad et al. [2016] is the modular plastic-casing design, so that the display can be detached from the lens casing, allowing the autorefractor to be placed between them. We 3D-printed lens holders to embed the focus-adjustable lenses in the HMD casing. The wiring for the video signal and power of the LCD panel is independent of the head-tracking and lens-controller signals of the HMD, allowing us to split the device into two independent modules around the autorefractor. See Figs. 1 and 3 for overviews of the design.

The lenses are controlled via a USB interface; commands are sent in sync with each rendered frame to the lens controller. The controller software was written in C# around the original Google cardboard library, with a custom driver for the focus-adjustable lenses. We implemented the controller as an intuitive ad-hoc API, which is exposed to the graphics programmer with a "focus-at-distance" mode providing expected behavior (Sec. 2.2). We implemented our API in Unity3D, allowing its use in any Unity3D application or game.

When we were measuring accommodation, the autorefractor was placed between the screen and lenses as shown in Fig. 1. We mounted the lens assembly in the chin/head rest (Fig. 4, left) and mounted the display on the other side of the autorefractor (Fig. 4,


Fig. 3. Overview of the HMD design. (Left) Optotune lenses and the 3D printed holders. (Right) Detachable screen.

right). The autorefractor has a hot mirror above its beam sensor (Fig. 1). That mirror reflects the infrared retinal reflection to the sensor while transmitting the visible stimulus to the viewer.

Because of the focus-adjustable-lens assembly, the viewer's eyes were farther from the autorefractor camera than is customary. Thus, to focus the autorefractor camera on the viewer's cornea, we placed a -0.75D lens in the path and accounted for this in our calibration procedure (Sec. 3.2).


Fig. 4. Overview of the autorefractor setup. Back view (left) and subject measurement (right).

3.2 Lens Calibration and Adjustments

It is important to ensure that the overall pipeline with its various components is correctly calibrated and adjusted. A short description of our procedure follows, more details can be found in supplemental material.

Focus-adjustable Lens Calibration. We needed to know what electrical current value to send to each focus-adjustable-lens controller in order to achieve the desired focal power. We used a digital single-reflex (DSLR) camera to create a mapping between electric current and focal distance for both lenses. We found the relationship between current and focal power to be linear, so we fit the data via linear regression. The mappings differed for the two lenses due to imprecisions in the manufacturing process.

Measurement Pipeline Calibration. We also needed to calibrate the accommodation measurements for the whole pipeline. This required paralyzing the muscles that control accommodation in both eyes of one subject. With fixed accommodation, we could then obtain a precise mapping between autorefractor measurements and accommodation by placing lenses of known power in the path.

Lens Distortion. We used the Brown distortion model [Fryer and Brown 1986] to determine the radially distorted pixel coordinates $x_{\text{dist}}, y_{\text{dist}}$ due to a lens:

$$x_{\text{dist}} = x_{\text{proj}}(1 + k_1 * r^2 + k_2 * r^4) \tag{1}$$

$$y_{\text{dist}} = y_{\text{proj}}(1 + k_1 * r^2 + k_2 * r^4) \tag{2}$$

where $x_{\text{proj}}, y_{\text{proj}}$ are the original pixel positions, k_1, k_2 are the lens distortion coefficients, and r is pixel distance from the center of the image.

The focus-adjustable lenses had different distortions at different focal powers, so we estimated the geometric distortion at each lens power to be used in the experiments. To do this, we first estimated k_1 and k_2 for the DLSR camera itself using the standard procedure of OpenCV [OpenCV 2016]. We then took images with the same camera through the focus-adjustable lenses and applied the previously found camera-distortion parameters to those images. The distortion that remained was obviously from the focus-adjustable lenses. Then we estimated k_1 and k_2 for the focus-adjustable lenses for several focal powers. We use linear interpolation to obtain the $x_{\text{dist}}, y_{\text{dist}}$ radial distortion for powers in-between those measured in the calibration.

Lens Breathing and Depth-of-Field Rendering. Because the focus-adjustable lenses were positioned in front of the viewer's eyes, changes in focal power cause changes in retinal-image size. This magnification is often called "lens breathing". We eliminated this effect by adjusting the size of the rendered scene on the HMD display. To do DoF rendering, we employed a standard, high-performance approach using disc filters [Nguyen 2007; Potmesil and Chakravarty 1982] similar to [Konrad et al. 2016]. Our device simulates retinal blur for – the assumed – point of gaze; as a result, the flattened retinal image is correct for a single gaze point.

4 ACCOMMODATION EXPERIMENT

We first investigated the efficacy of different HMD designs for driving accommodation to the distance of the simulated object. We did so by testing the following viewing conditions: gaze-contingent DoF vs. no gaze-contingent DoF rendering, fixed vs. changing lenses, and monovision vs no lens offset. We included two additional viewing manipulations, monocular vs binocular viewing and low vs

high stimulus speed to evaluate additional assumptions from vision science that are relevant to the VA conflict.

4.1 Apparatus and Design

Data Acquisition. The autorefractor generates a file of spherical power values 4-6 times a second. We used a Unity3D application for rendering and it generated an SQL database with the current experimental settings and stimulus location. There were five database entries per second. We manually started the autorefractor recording and stimulus motion on their respective computers, which yielded approximate synchronization. Occasionally, the autorefractor stopped recording because the subject's cornea moved relative to the autorefractor's optical path. When this happened, we restarted the measurement. Blinks during measurements were assigned an undefined value.

Stimulus. The experimental stimulus was a 3D forest scene with near and far elements (Fig. 5) and a target that was a 0.1m^3 cube with Maltese crosses on its sides. The target translated in depth sinusoidally in diopters from 0.33m (3D) to 6m (0.17D) and back over a few periods. The range of motion was 2.83D. The motion occurred at 0.1 or 0.25Hz. The higher frequency is effective in driving vergence-accommodation coupling (so that vergence can drive accommodation) and the lower one is less effective in driving the coupling [Schor 1992]. We used sinusoidal motion to promote reflexive, natural accommodation (i.e., involuntary as opposed to voluntary accommodation; [Marg 1951]). Stimulus translation was along the line of sight for each subject's preferred eye, the eye in which accommodation was measured. With this type of translation, the measured eye did not have to rotate to follow the stimulus as it moved in depth; this simplified the alignment of the autorefractor with the optical axis of that eye.

Fig. 5. The experimental stimulus; a forest scene and a cube with a Maltese cross on it.

Subjects. We tested five naive subjects (three males, mean age 25, SD 2.1) who were 30 years of age or younger and had normal visual acuity and stereoacuity. We did not include older subjects because after age 30 most subjects have reduced accommodative range [Duane 1912]. Accommodative behavior was very similar across subjects (as evidenced by the statistical results in Sec. 4.3 and supplemental material), so we adopted the common practice in accommodation research of testing only a handful of subjects

[Kruger et al. 1993; MacKenzie et al. 2010]. The experiment lasted about an hour per subject. Each subject provided enough data to allow statistical comparisons of responses across conditions. The study was approved by the institutional ethics board at UC Berkeley.

4.2 Conditions

There were three main experimental conditions: monocular viewing, binocular viewing, and monovision. In monocular viewing, images were presented to one eye only. In binocular viewing, images were shown to both eyes and had appropriate binocular disparity; the powers of the lenses before the two eyes were always equal to one another. In monovision, images were presented to both eyes and again had appropriate disparity; the powers of the lenses before the two eyes differed by 1 or 2D and were fixed. For each of these main conditions, there were four sub-conditions that were various combinations of DoF rendering, fixed- vs. changing-lens power, and target speed. These combinations generated the viewing conditions of the conflict-alleviating designs being tested.

In the context of evaluating different HMD designs, comparing responses with and without DoF rendering will inform us about the contribution of changes in distance of the rendered focal plane to driving accommodation, while comparing responses with the lenses changing power or not will inform us about the contribution of blur to driving accommodation. Additionally, comparing accommodative responses in the monocular and binocular conditions will inform us about the contribution of disparity to driving accommodation.

Condition	DoF	Changing	f
Fixed/No-DoF			0.77D
Fixed/DoF	✓		0.77D
Changing/No-DoF		✓	3-0.17D
Changing/DoF	✓	✓	3-0.17D

Table 1. Gaze-contingent DoF blur and focus-adjustable lenses combinations in the four conditions. f is focal distance. The Fixed/No-DoF condition simulates current HMDs, e.g. Oculus Rift.

For the monocular viewing and binocular viewing conditions, we had four conditions (Table 1). When gaze-contingent DoF was implemented, the focal plane for DoF rendering moved with the simulated target distance. The subject's pupil diameter was measured *in situ* during a 3-minute pretest. We used that measured diameter for the DoF rendering. When the lenses were changing in focal power, the focal distance was consistent with the simulated target distance. In the binocular conditions, we stimulated both eyes and measured accommodation in the preferred eye. In the monocular conditions, we stimulated the preferred eye and measured accommodation in the same eye. For monovision, all conditions were stereoscopic by definition. DoF rendering was always implemented as in [Konrad et al. 2016]. The focal powers of the lenses were set to the values shown in Table 2. Two frequencies of sinusoidal motion were presented in all cases: 0.1 and 0.25Hz.

4.2.1 Protocol. When the subject arrived for an experimental session, we first described the experiment, requested that they read and sign the consent form, and answered any questions they had. We

Condition	Measured Eye	Non-measured Eye
1	1.77D	0.77D
2	2.77D	0.77D
3	0.77D	1.77D
4	0.77D	2.77D

Table 2. Monovision Conditions

then measured their refractive error in both eyes to adjust the focus-adjustable lenses to cancel the spherical component (i.e., myopia or hyperopia) for each eye. Two subjects had no refractive error; the others required such an adjustment. We measured the subject's inter-pupillary distance and used this value to set the disparities presented on the HMD screen individually for each subject. We showed the stimuli for three minutes at 0.77D and instructed subjects to fixate it. During those three minutes, we measured pupil diameter using the autorefractor.

We then began the main experiment. The subjects were instructed to fixate the center of the Maltese cross at all times. All subjects were tested in all conditions: a within-subjects design. We randomized main condition/sub-condition/trial succession to avoid order effects. Each subject was exposed to at least 90 seconds (> 450 measurements) of the stimulus in each condition generating a sufficient amount of accommodation data for analysis.


Fig. 6. Accommodative stimuli and responses in two representative conditions. Left: Binocular viewing with changing lenses and DoF rendering. Blue curve shows changes in distance specified by vergence, size, blur (due to the changing lenses), and focal plane (due to DoF rendering). Red curve represents accommodative responses from one subject. Green curve is best-fitting sinusoid to those responses. Right: Binocular viewing with fixed lenses and DoF rendering. Blue, red, and green curves represent stimulus, response, and best fit, respectively, for the same subject.

4.2.2 Data Processing. To analyze the data, we registered the stimulus (rendering) database with the autorefractor data files using peak-to-peak synchronization. After registration, we applied corrections to the autorefractor data as determined from the calibration procedure described earlier. Consistent with normal practice in accommodation research [Kruger et al. 1993; MacKenzie et al. 2010], we determined the gain of the accommodative response from the

measurement time series. Gain is the amplitude of the response at the stimulus frequency divided by the amplitude of the stimulus at that frequency. We estimated the gain by fitting the response data with a sinusoid at the stimulus frequency with amplitude, phase and DC offset as free parameters (using the Levenberg-Marquardt damped least-squares method). Gain was then the ratio of the amplitude of the fitted sinusoid and the stimulus sinusoid. Fig. 6 shows the stimulus, response, and fitted curves for two representative conditions. In accommodation research, the highest observed gains are 0.8-0.9D under optimal conditions [Kruger et al. 1993; MacKenzie et al. 2010]. Many of the responses we observed had similar high gains (such as in the left panel) indicating accurate accommodation. Many had, however, much lower gain (such as in the right panel) indicating inaccurate accommodation.

Determining the Floor Value of Gain Estimates. With our fitting technique, the estimated gain will never be zero because there will always be some modulation at some phase at the stimulus frequency even when accommodation is not being driven. To determine the floor for gain estimates, we used the *stand-by* accommodation data that were acquired in the first three minutes of the experiment (see Sec. 4.2.1) when the accommodation stimulus was fixed at 0.77D. We fit the resulting data with sinewaves at both the low and high frequencies. The average gain was 0.06 (STD = 0.04), which is thus the floor value: the expected gain when accommodation is not being driven.

4.3 Accommodation Study Results

Accommodative responses among the five subjects were very similar to one another (see supplemental material), so instead of showing individual data, we averaged across subjects. Fig. 7 plots the average gain for the monocular and binocular viewing conditions. The available cues in the monocular and binocular conditions were change in target size (present in all conditions), vergence (present in the binocular but not monocular conditions), blur (present in the changing-lens, but not fixed-lens conditions), and DoF (present in the conditions in which the focal plane for rendering moved in and out, but not present with pinhole rendering).

We performed appropriate tests to determine which effects were statistically significant. We had data for at least 16 stimulus cycles in each condition for each subject (>90 seconds or >450 data points in each condition). We performed a 4-way analysis of variance (ANOVA) with factors of a) monocular vs. binocular, b) low vs. high stimulus speed, c) DoF vs. pinhole rendering, and d) changing vs. fixed lens.

The first two of those factors are not directly relevant to our study of efficacy of different designs because HMDs are generally binocular and visual targets move with arbitrary speeds. We evaluated them nonetheless because these factors from the vision science literature have never been validated in an HMD. We found a small effect of binocular vs monocular viewing: gains were 0.50 and 0.40 respectively, a difference that was statistically significant $(F(1, 79) = 16.81, p \approx 1 \times 10^{-4})$. The effect of stimulus speed was statistically insignificant.

DoF Rendering. There was no significant effect of gaze-contingent DoF vs. pinhole rendering: Gains were 0.47 and 0.43, respectively.

Changing Lens. Changing lenses had a substantial effect on accommodative gain: Average gains were 0.72 and 0.18 in the changing- and fixed-lens conditions, respectively. The ANOVA yielded a statistically significant main effect for changing vs. fixed lens $(F(1, 79) = 484.6, p \approx 6 \times 10^{-6})$.

Thus changes in the blur stimulus created by changing lenses had a large effect on accommodation, while the other stimulus changes had much smaller effects.


Fig. 7. Accommodative response gains for monocular and binocular conditions. The gain of accommodative responses, averaged across subjects, is plotted for each viewing condition. "Low" and "High" represent low and high target speed respectively. Error bars indicate the standard error for each condition. This graph clearly shows that only changing lenses drive accommodation effectively.
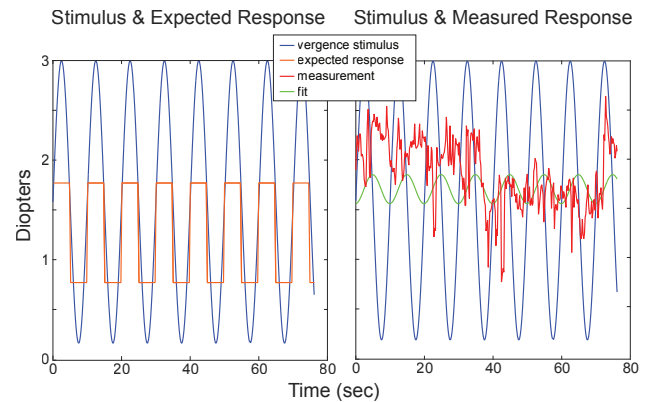

Fig. 8. Accommodative stimuli and responses in the monovision experiment. Left: Stimulus and expected response. The dark blue curve represents the stimulus and the orange curve the expected response if accommodation is driven as intended in the monovision approach. Right: Stimulus and response for one subject when the temporal frequency was $0.1Hz$ and lens offset was $1D$. The dark blue curve represents the stimulus specified by vergence and target size, the red curve is the response, and the green curve is the best-fitting sinusoid to the response.

Monovision. Fig. 8 shows expected and observed responses in the monovision experiment. The orange curve in the left panel represents the responses one would expect if both eyes accommodate far when the disparity of the stimulus specifies far and both eyes accommodate near when the disparity specifies near. The expected response is some waveform with the same frequency as the vergence stimulus, not necessarily a squarewave as shown in the figure. The red curve in the right panel shows the observed response for one representative condition and subject. As you can see, the responses were not well correlated with the expectation. Instead, accommodation seemed to "flip" at random intervals between the distance specified by the "near eye" (1.77D) and the distance specified by the "far eye" (0.77D) (in half the conditions, the near eye was at 2.77D and the far eye at 0.77D; not shown).

Fig. 9 shows the average accommodative gains for all conditions of the monovision experiment. Gain was uniformly low, meaning that accommodation was not well correlated with the change in simulated target distance. An ANOVA indicated that there was no effect on gain due to either stimulus speed or the magnitude of the lens offset (1 vs. 2D). The monovision results show quite clearly that accommodation is not driven in a way that corresponds with simulated distance.


Fig. 9. Average gain for all conditions in the monovision experiment. Error bars indicate the standard error for each condition. "Low" and "High" represent low and high target speed respectively.

4.4 Can Accommodation Predict Discomfort?

Our working hypothesis is that the VA conflict drives discomfort, so we wanted to see whether the accommodation data can be used to make predictions about which conditions should be more comfortable than others. We could have used accommodative gain as a predictor (e.g., lower gain predicts greater VA conflict and hence more discomfort), but it is not obvious how to do this for the monovision condition because the stimuli to the two eyes differed. Instead we computed the *accumulated VA conflict* for each eye in all conditions to see whether this predicts which conditions should be more comfortable.

Accumulated VA Conflict. From the accommodation experiment, we know the accommodative response and vergence stimulus over

time. We did not measure vergence responses, but they are generally quite accurate [Schor 1992], so we assumed that the vergence response was always equal to the vergence stimulus. We calculated the unsigned difference between the accommodative response and vergence stimulus: ~5 differences per second. We summed those values for the duration of a recording, and this is the *accumulated VA conflict.* As described earlier, there was a potential synchronization error between the stimulus presentation and accommodative measurement at the beginning of each recording, so we computed the accumulated VA conflict for different time offsets between stimulus and measurement for each individual recording. We found the offset that produced the smallest conflict and used that value for that recording. We computed the average accumulated VA conflict for each condition and subject. The conflict was the same in the two eyes for all conditions except monovision in which case we calculated the average value for the two eyes. The accumulated VA conflict differed significantly across conditions, but was very similar across subjects, so we report values averaged across subjects.

Fixed and Changing Lenses. The accumulated VA conflict for the fixed- and changing-lens conditions were respectively 1.06D and 0.36D. Thus, the conflict was nearly three times greater in the fixed-lens than in the changing-lens conditions because accommodation was driven much more effectively in the latter condition. We predict therefore that the fixed-lens condition should be considerably less comfortable than the changing-lens condition.

DoF and no DoF.. The accumulated VA conflict for the DoF and no-DoF conditions were 1.06D and 1.11D. Thus, the conflict was slightly greater in the no-DoF condition because accommodation was driven slightly more effectively in the DoF condition. We predict therefore that no-DoF should be slightly less comfortable or equally comfortable relative to DoF.

Monovision. In the monovision condition, the two eyes had different focal distances so accommodation in one eye was often closer to the vergence stimulus than in the other eye. Hence the conflict differed between eyes over time. We computed the accumulated VA conflict for both eyes and averaged those two values. The computation is described in greater detail in supplemental material. The resulting accumulated VA conflict for the 1D monovision condition was 1.23D (1.48D for the 2D condition), while the value for the fixed-lens baseline condition was 1.06D. We thus predict more discomfort with monovision than in the baseline fixed-lens condition.

5 DISCOMFORT EXPERIMENT

If the VA conflict is in fact a significant source of visual discomfort in HMDs, we should be able to predict viewer comfort from our accommodation results (Sec. 4.4), allowing us to answer our second question. Specifically, we predict that 1) employing changing lenses will reduce discomfort, 2) employing DoF rendering will slightly reduce discomfort or have no effect, and 3) employing monovision will increase rather than reduce discomfort. We next tested these predictions with a discomfort study.

5.1 Experimental Setup

Discomfort studies are long and complex, and as such we selected the conditions that would be most informative. The discomfort experiment consisted of three sessions, run on separate days. Each session had two experimental conditions each lasting 30 minutes. There was a mandatory 15-minute rest break between the two conditions to allow symptoms from the last condition to dissipate. All conditions were binocular. In every session, one of the conditions simulated a conventional HMD (e.g., Oculus Rift) which is our baseline: fixed lenses of the same power before the eyes and appropriate binocular disparity and DoF rendering. The other condition was either fixed lens with no DoF rendering, or changing lens with gaze-contingent DoF rendering, or monovision (1D offset) with gaze-contingent DoF rendering. We did not perform a 2D discomfort study because if a 1D offset had been uncomfortable for subjects, a 2D offset would have been even more disturbing [Evans 2007]. Since there was no head tracking we strongly discouraged subjects from moving their heads. The order of sessions and conditions was randomized across subjects.

The stimulus was again the 3D forest scene with an approaching and receding Maltese cross. The frequency of the approach and withdrawal was 0.1Hz. We added a Tumbling-E visual acuity task [Ferris et al. 1982]. An E with an orientation of 0, 90, 180, or 270° (where 0 corresponds to a regular E) was presented for 400msec at random times. Subjects were required to fixate the cross at all times and to press a key on the keyboard whenever a regular E appeared. We added this task to ensure that subjects maintained fixation and accommodation on the cross. Using Signal Detection [Nevin 1969], we excluded from further analysis subjects who had less than 60% correct.

After completing a session, the subject filled out a session - comparison questionnaire in which he/she compared their general fatigue, eye irritation, headache, and nausea at the end of the two conditions on a 7-point scale. Each session took about an hour and a half. The three sessions were completed on different days.

There were 16 naive subjects included in the analysis for each condition (mean age: 25.5, SD: 2.4). One subject from the changing-vs fixed-lens session and two from the monovision vs no-lens session were excluded because they did not meet our performance criterion of 60% correct in the Tumbling-E task. The age range was 21 to 30 years. The majority were male. All subjects had normal or corrected-to-normal visual acuity and stereo acuity. None were aware of the experimental hypotheses. Appropriate consent and debriefing were conducted according to the Declarations of Helsinki.

5.2 Results

The discomfort results are summarized in Figure 10. The upper, middle, and lower panels show the results respectively for DoF vs no-DoF, changing- vs fixed-lens, and monovision vs no lens.

DoF Rendering. The results reveal a slight preference for DoF over no-DoF rendering. Subjects reported consistently less eye irritation (Wilcoxon one-tailed test, $p < 0.05$) and headache ($p < 0.05$); upper panel in Figure 10. The differences for fatigue and nausea were statistically insignificant. This result is consistent with our prediction


Fig. 10. Discomfort results for the three experiments. DoF vs no-DoF rendering (top), changing- vs fixed-lens (middle), and monovision vs no lens offset (bottom). The questionnaire used a -3.0, 3.0 scale where -3.0 meant a strong preference for the first condition and 3.0 a strong preference for second condition as indicated by the labels at the top of each panel. ∗ indicates $p < 0.05$; one-tailed, Wilcoxon signed-rank test for zero median.

that DoF rendering will have either a small benefit or no benefit for viewer comfort.

Changing Lens. The results exhibit a systematic preference for changing-lens over fixed-lens. Subjects reported consistently less fatigue ($p < 0.05$), eye irritation ($p < 0.05$), headache ($p < 0.05$), and nausea ($p < 0.05$); middle panel in figure. This result is consistent with our prediction that the changing-lens condition leads to improved viewer comfort.

Monovision. The results reveal a systematic preference for no lens over monovision. Subjects reported consistently less fatigue ($p < 0.05$), eye irritation ($p < 0.05$) and headache ($p < 0.05$) with no lens offset between the eyes (i.e., in the non-monovision condition); lower panel in figure. This is consistent with the prediction derived from our accommodation experiment: Monovision did not produce

accommodation consistent with vergence distance and therefore did not improve viewer comfort.

The significance levels in the above analyses were not corrected for multiple comparisons because we treated each pairwise comparison as a planned comparison, and accordingly, reported which comparisons were significant and which were not.

6 DISCUSSION

Our experiments answered the two questions we asked at the outset (Sec. 1). Concerning how well each display configuration drives accommodation, our results reveal for the first time the properties of HMDs that enable effective stimulation of accommodation and thereby minimize the VA conflict. They show that accommodation is accurate when stimulus focal distance changes in concert with vergence distance; this was achieved by using focus-adjustable lenses coupled to the display. Thus, natural blur is a significant driver of accommodation in HMDs as it is in other viewing situations. Very importantly, our results show that the ability to drive accommodation did in fact improve viewer comfort, answering the second question. This finding establishes a link between accommodative accuracy and viewer comfort, demonstrating that – at least for the conditions we tested – the ability to drive accommodation is a good predictor of discomfort. We also found that vergence and DoF have a small effect on accommodation but fall well short of driving accommodation accurately to the simulated stimulus distance.

The next generation of HMDs could move productively in the direction of employing adjustable lenses coupled to the display, but this requires improvements in adjustable-lens technology, particularly maintenance of good optical quality with a large field of view.

DoF Rendering. We predicted that DoF would yield a small or no improvement in viewer comfort and our results are consistent with this prediction. Thus, gaze-contingent DoF rendering aids accommodation slightly and has a measurable benefit for viewer comfort. In retrospect, the finding that DoF rendering has but a small effect on accommodation makes great sense. Such rendering does not change the fact that the sharpest image occurs when the viewer accommodates to the screen distance. If the eye accommodates to a distance other than the screen distance, the retinal image becomes blurrier whether DoF rendering is employed or not. Because DoF rendering does not assist accommodation very much, display designers should weigh the benefits when deciding whether the computational cost associated with such rendering is justified. However, appropriate DoF rendering may be important for achieving other goals, e.g., creating a realistic sense of depth [Mauderer et al. 2014], or adjusting perceived scale [Held et al. 2010].

Monovision. Monovision did not produce accommodative responses consistent with the simulated distance specified by vergence and target size. Instead accommodation seemed to oscillate at random intervals between the focal distance associated with the "far eye" and the "near eye" (Sec. 4.4). We also found that monovision increased discomfort significantly compared to conventional stereo methods. Previous evaluations of the monovision technique have yielded somewhat inconsistent findings. Subjects in Konrad et al. [2016] had a slight preference for monovision over a baseline condition.

Subjects in Johnson et al. [2016] and the current study reported significantly greater visual discomfort in monovision than in the baseline condition. We attribute the difference in results to the duration of exposure to the various conditions. In Konrad et al. [2016], subjects were presented five conditions and toggled between them. Hence they were exposed to monovision for only a few seconds at a time. In Johnson et al. [2016] and the current study, subjects were exposed to each condition for 10-30 minutes before discomfort was assessed. Thus, subjects were reporting on very different experiences in the Konrad et al. [2016] study than in the Johnson et al. [2016] and current study.

Binocular vs Monocular. We found that binocular viewing yields somewhat more accurate accommodation in HMDs than monocular viewing. We did not assess comfort in the monocular conditions because there is little concern about discomfort with monocular viewing. Our finding is consistent with observations in vision science that vergence is a somewhat effective driver of accommodation [Cumming and Judge 1986; Martens and Ogle 1959].

7 CONCLUSION

We attempted to answer two important questions: whether proposed designs to alleviate the VA conflict in HMDs (i) do actually drive accommodation and (ii) if so, do they improve user comfort? We answered these questions by designing a modular HMD with focus-adjustable lenses that allowed us to measure accommodation directly. We conducted a set of experiments to determine how different display configurations affect accommodation. The results show that the most effective way to eliminate the VA conflict is also the most obvious: One must drive accommodation by getting focus cues correct or at least nearly correct. The results also show that other proposed solutions – most notably DoF rendering and monovision – do not drive accommodation accurately and therefore do not minimize the VA conflict as much as one would hope. We used the accommodation data to make specific predictions about which conditions should be more comfortable. The results from our discomfort study were completely consistent with those predictions. The one technique (changing-lens) that drove accommodation accurately and therefore reduced VA conflict, enabled significantly more comfortable viewing than any of the other protocols including the conventional HMD protocol. We also showed that vergence drives accommodation, but inaccurately. So using vergence alone to drive accommodation, as is done in conventional stereo displays, is not very effective in reducing the VA conflict.

Given our results, improving adjustable-lens designs for HMDs is a promising avenue for future research. Our current setup requires an assumption of where the viewer is fixating from moment to moment. One could avoid this limitation by using eye tracking [Stengel et al. 2015] to estimate current fixation distance and then drive the focus-adjustable lenses to the focal distance that is consistent with fixation distance. This would allow the development of a practical system for widespread use. Our results indicate that such a design, together with improved focus-adjustable lens technology will provide a much more comfortable HMD experience, and could be important in more widespread adoption of this technology.

In future designs, it may be necessary to measure accommodation in an HMD. Clearly, autorefractors are too slow and bulky to be useful for this purpose. Wavefront sensors, such as the Hartmann-Shack sensor, are generally faster and more accurate than autorefractors. There has been recent work on integrating wavefront sensors with eye tracking that could prove very beneficial for next-generation HMDs [Chirre et al. 2015]. It would be useful therefore to incorporate such a sensor in an HMD for *in situ* measurement of accommodation.

ACKNOWLEDGEMENTS

REFERENCES

Ernest A Walton Ball. 1952. A Study in Consensual Accommodation. *Optometry & Vision Science* 29, 11 (1952), 561–574.

Fergus W Campbell. 1960. Correlation of Accommodation Between the Two Eyes. *Journal of the Optical Society of America* 50, 7 (1960), 738–738.

Emmanuel Chirre, Pedro Prieto, and Pablo Artal. 2015. Dynamics of the Near Response Under Natural Viewing Conditions with an Open-view Sensor. *Biomedical Optics Express* 6, 10 (2015), 4200–4211.

Bruce G Cumming and Stuart J Judge. 1986. Disparity-induced and Blur-induced Convergence Eye Movement and Accommodation in the Monkey. *Journal of Neurophysiology* 55, 5 (1986), 896–914.

Alexander Duane. 1912. Normal Values of the Accommodation at all Ages. *Journal of the American Medical Association* 59, 12 (1912), 1010–1013.

Andrew T Duchowski, Donald H House, Jordan Gestring, Rui I Wang, Krzysztof Krejtz, Izabela Krejtz, Radosław Mantiuk, and Bartosz Bazyluk. 2014. Reducing Visual Discomfort of 3d Stereoscopic Displays with Gaze-contingent Depth-of-Field. In *Proc. ACM SAP*. 39–46.

Bruce JW Evans. 2007. Monovision: a Review. *Ophthalmic and Physiological Optics* 27, 5 (2007), 417–439.

Frederick L Ferris, Aaron Kassoff, George H Bresnick, and Ian Bailey. 1982. New Visual Acuity Charts for Clinical Research. *American Journal of Ophthalmology* 94, 1 (1982), 91–96.

Edgar F Fincham and John Walton. 1957. The Reciprocal Actions of Accommodation and Convergence. *The Journal of Physiology* 137, 3 (1957), 488–508.

S Kay Fisher, Kenneth J Ciuffreda, and Steven Hammer. 1987. Interocular Equality of Tonic Accommodation and Consensuality of Accommodative Hysteresis. *Ophthalmic and Physiological Optics* 7, 1 (1987), 17–20.

John G Fryer and Duane C Brown. 1986. Lens Distortion for Close-range Photogrammetry. *Photogrammetric Engineering and Remote Sensing* 52, 1 (1986), 51–58.

Robert T Held, Emily A Cooper, James F O'Brien, and Martin S Banks. 2010. Using Blur to Affect Perceived Distance and Size. *ACM Transactions on Graphics* 29, 2 (2010), 19:1–16.

Sébastien Hillaire, Anatole Lécuyer, Rémi Cozot, and Géry Casiez. 2008. Using an Eye-tracking System to Improve Camera Motions and Depth-of-Field Blur Effects in Virtual Environments. In *IEEE Virtual Reality Conference*. IEEE, 47–50.

David M Hoffman, Ahna R Girshick, Kurt Akeley, and Martin S Banks. 2008. Vergence–Accommodation Conflicts Hinder Visual Performance and Cause Visual Fatigue. *Journal of Vision* 8, 3 (2008), 33–63.

Xinda Hu and Hong Hua. 2014. Design and Assessment of a Depth-fused Multi-focal-plane Display Prototype. *Journal of Display Technology* 10, 4 (2014), 308–316.

Paul V Johnson, Jared AQ Parnell, Joohwan Kim, Christopher D Saunter, Gordon D Love, and Martin S Banks. 2016. Dynamic Lens and Monovision 3D Displays to Improve Viewer Comfort. *Optics Express* 24, 11 (2016), 118:08–27.

Joohwan Kim, David Kane, and Martin S Banks. 2014. The Rate of Change of Vergence–Accommodation Conflict Affects Visual Discomfort. *Vision Research* 105 (2014), 159–165.

Robert Konrad, Emily A Cooper, and Gordon Wetzstein. 2016. Novel Optical Configurations for Virtual Reality: Evaluating User Preference and Performance with Focus-tunable and Monovision Near-eye Displays. In *SIGCHI 2016*. ACM, 1211–1220.

Frank L Kooi and Alexander Toet. 2004. Visual Comfort of Binocular and 3D Displays. *Displays* 25, 2 (2004), 99–108.

Philip B Kruger, Steven Mathews, Karan R Aggarwala, and Nivian Sanchez. 1993. Chromatic Aberration and Ocular Focus: Fincham Revisited. *Vision Research* 33, 10 (1993), 1397–1411.

Marc Lambooij, Marten Fortuin, Ingrid Heynderickx, and Wijnand IJsselsteijn. 2009. Visual Discomfort and Visual Fatigue of Stereoscopic Displays: A Review. *Journal of Imaging Science and Technology* 53, 3 (2009), 1–14.

Gordon D Love, David M Hoffman, Philip JW Hands, James Gao, Andrew K Kirby, and Martin S Banks. 2009. High-speed Switchable Lens Enables the Development of a Volumetric Stereoscopic Display. *Optics Express* 17, 18 (2009), 15:716–725.

Kevin J MacKenzie, David M Hoffman, and Simon J Watt. 2010. Accommodation to Multiple-focal-plane Displays: Implications for Improving Stereoscopic Displays and for Accommodation Control. *Journal of Vision* 10, 8 (2010), 1–22.

Andrew Maimone, Gordon Wetzstein, Matthew Hirsch, Douglas Lanman, Ramesh Raskar, and Henry Fuchs. 2013. Focus 3D: Compressive Accommodation Display. *ACM Transactions on Graphics* 32, 5 (2013), 1–13.

Elwin Marg. 1951. An Investigation of Voluntary as Distinguished from Reflex Accommodation. *Optometry & Vision Science* 28, 7 (1951), 347–356.

Theodore G Martens and Kenneth N Ogle. 1959. Observations on Accommodative Convergence; Especially its Nonlinear Relationships. *American Journal of Ophthalmology* 47, 1 Pt 2 (1959), 455–462.

Michael Mauderer, Simone Conte, Miguel A Nacenta, and Dhanraj Vishwanath. 2014. Depth Perception with Gaze-contingent Depth of Field. In *SIGCHI 2014*. ACM, 217–226.

Rahul Narain, Rachel A Albert, Abdullah Bulbul, Gregory J Ward, Martin S Banks, and James F O'Brien. 2015. Optimal Presentation of Imagery with Focus Cues on Multi-plane Displays. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 59:1–12.

John A Nevin. 1969. Signal Detection Theory and Operant Behavior: A Review of David M. Green and John A. Swets' Signal Detection Theory and Psychophysics. 1. *Journal of the Experimental Analysis of Behavior* 12, 3 (1969), 475–480.

Hubert Nguyen. 2007. *Gpu Gems 3*. Addison-Wesley Professional.

OpenCV. 2016. Camera calibration With OpenCV. http://docs.opencv.org/2.4/doc/tutorials/calib3d/camera_calibration/camera_calibration.html. (2016). [Online; accessed 5-April-2017].

Keiichiro Otani, Takumi Yoshida, Hideaki Nii, Naoki Kawakami, and Susumu Tachi. 2008. Active Focus Control with Eye Gaze Detection for Depth Cognition. *Asiagraph2008* (2008).

Nitish Padmanaban, Robert Konrad, Tal Stramer, Emily A. Cooper, and Gordon Wetzstein. 2017. Optimizing virtual reality for all users through gaze-contingent and adaptive focus displays. *Proceedings of the National Academy of Sciences* 114, 9 (2017), 2183–2188. DOI:https://doi.org/10.1073/pnas.1617251114 arXiv:http://www.pnas.org/content/114/9/2183.full.pdf

Michael Potmesil and Indranil Chakravarty. 1982. Synthetic Image Generation with a Lens and Aperture Camera Model. *ACM Transactions on Graphics (TOG)* 1, 2 (1982), 85–108.

Clifton M Schor. 1992. A Dynamic Model of Cross-coupling Between Accommodation and Convergence: Simulations of Step and Frequency Responses. *Optometry & Vision Science* 69, 4 (1992), 258–269.

James E Sheedy, Michael G Harris, and Cheslyn M Gan. 1993. Does the Presbyopic Visual System Adapt to Contact Lenses?. *Optometry & Vision Science* 70, 6 (1993), 482–486.

Takashi Shibata, Joohwan Kim, David M Hoffman, and Martin S Banks. 2011. The Zone of Comfort: Predicting Visual Discomfort with Stereo Displays. *Journal of Vision* 11, 8 (2011), 11:1–29.

Michael Stengel, Steve Grogorick, Martin Eisemann, Elmar Eisemann, and Marcus A Magnor. 2015. An Affordable Solution for Binocular Eye Tracking and Calibration in Head-mounted Displays. In *Proceedings of the 23rd Annual ACM Conference on Multimedia*. ACM, 15–24.

Yasuhiro Takaki. 2006. High-density Directional Display for Generating Natural Three-dimensional Images. *Proc. IEEE* 94, 3 (2006), 654–663.

Matthieu Urvoy, Marcus Barkowsky, and Patrick Le Callet. 2013. How Visual Fatigue and Discomfort Impact 3D-TV Quality of Experience: A Comprehensive Review of Technological, Psychophysical, and Psychological Factors. *Annals of Telecommunications* 68, 11-12 (2013), 641–655.

Margarita Vinnikov, Rob S Allison, and S Fernandes. 2016. Impact of Depth of Field Simulation on Visual Fatigue: Who are Impacted? And how? *Int. J. of Human-Computer Studies* 91 (2016), 37–51.

Dorothy M Win-Hall, Jaime Houser, and Adrian Glasser. 2010. Static and Dynamic Measurement of Accommodation Using the Grand Seiko WAM-5500 Autorefractor. *Optometry and Vision Science* 87, 11 (2010), 11:1–19.

Shunnan Yang and James E Sheedy. 2011. Effects of Vergence and Accommodative Responses on Viewer's Comfort in Viewing 3D Stimuli. In *SPIE Electronic Imaging*. International Society for Optics and Photonics, 1–13.