Title

ChromaBlur: Rendering chromatic eye aberration improves accommodation and realism

Journal

ACM Transactions on Graphics, 36(6)

ISSN

07300301

Authors

Cholewiak, Steven A
Love, Gordon D
Srinivasan, Pratul P
et al.

Publication Date

2017-11-20

DOI

10.1145/3130800.3130815

Data Availability

The data associated with this publication are in the supplemental files.

Peer reviewed

ChromaBlur: Rendering Chromatic Eye Aberration Improves Accommodation and Realism

STEVEN A. CHOLEWIAK, University of California, Berkeley
GORDON D. LOVE, Durham University
PRATUL P. SRINIVASAN, University of California, Berkeley
REN NG, University of California, Berkeley
MARTIN S. BANKS, University of California, Berkeley

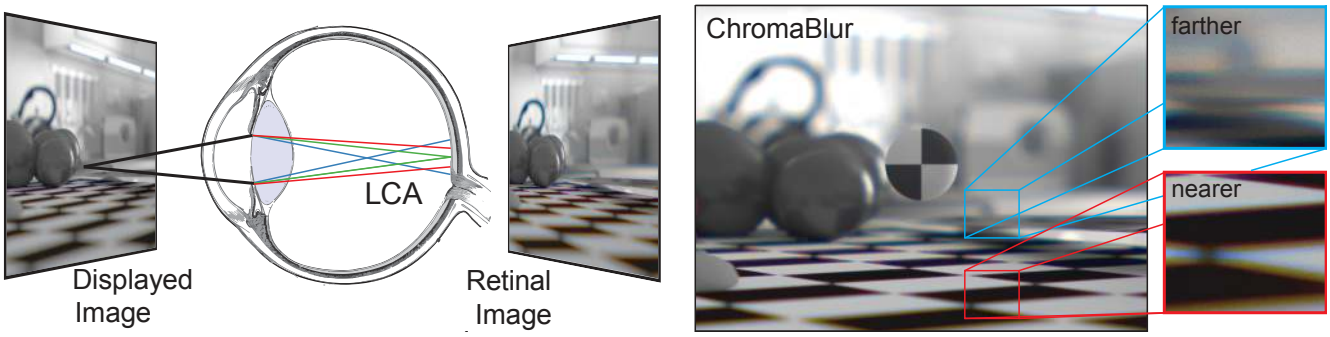

Fig. 1. Left: The pipeline for ChromaBlur rendering. The human eye has chromatic aberration, which produces depth-dependent chromatic effects on the retina: blue is focused in front of the retina and red behind. We calculate the image to be displayed that, when processed by the viewer's in-focus eye, will produce the correct chromatic effects on the viewer's retina for the simulated 3d scene. Right: Displayed image produced by ChromaBlur. The viewer is focused on the ball. Objects at other distances are blurred, but differently depending on whether they are farther (blue outline) or nearer (red) than the current focus distance.

Computer-graphics engineers and vision scientists want to generate images that reproduce realistic depth-dependent blur. Current rendering algorithms take into account scene geometry, aperture size, and focal distance, and they produce photorealistic imagery as with a high-quality camera. But to create immersive experiences, rendering algorithms should aim instead for perceptual realism. In so doing, they should take into account the significant optical aberrations of the human eye. We developed a method that, by incorporating some of those aberrations, yields displayed images that produce retinal images much closer to the ones that occur in natural viewing. In particular, we create displayed images taking the eye's chromatic aberration into account. This produces different chromatic effects in the retinal image for objects farther or nearer than current focus. We call the method *ChromaBlur*. We conducted two experiments that illustrate the benefits of ChromaBlur. One showed that accommodation (eye focusing) is driven quite effectively when ChromaBlur is used and that accommodation is not driven at all when conventional methods are used. The second showed that perceived depth and realism are greater with imagery created by ChromaBlur than in imagery created conventionally. ChromaBlur can be coupled with focus-adjustable lenses and gaze tracking to reproduce the natural relationship between accommodation and blur in HMDs and other immersive devices. It may thereby minimize the adverse effects of vergence-accommodation conflicts.

CCS Concepts: •**Computing methodologies** → **Perception;** *Rendering;* Virtual reality; Mixed / augmented reality;

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

ACM Reference format:
Steven A. Cholewiak, Gordon D. Love, Pratul P. Srinivasan, Ren Ng, and Martin S. Banks. 2017. ChromaBlur: Rendering Chromatic Eye Aberration Improves Accommodation and Realism. *ACM Trans. Graph.* 36, 6, Article 210 (November 2017), 12 pages.
DOI: 10.1145/3130800.3130815

0730-0301/2017/11-ART210 $15.00
DOI: 10.1145/3130800.3130815

1 INTRODUCTION

Seeing in three dimensions is a construction based on several depth cues. Rendering and display techniques are designed to reproduce those cues and thereby enable enriching 3d experiences. The methods for presenting perspective-based depth cues (e.g., linear perspective, texture gradient, relative size) and light-transport-based cues (e.g., shading, aerial perspective, occlusion) are well understood. Advances have also been made in reproducing triangulation-based depth cues (binocular disparity, motion parallax, and focus cues). Stereoscopic displays and direct-view, autostereoscopic displays enable effective presentation of disparity; head-mounted displays (HMDs) with head tracking enable motion parallax as well. The

most challenging remaining problem is the reproduction of focus cues: blur and accommodation.

Blur by itself is often considered a weak depth cue because it is said to be unsigned [Mather and Smith 2002]: i.e., the same blur arises whether an object is farther or nearer than the eye's current focus distance. But natural blur actually does contain sign information and can be used to drive accommodation in the correct direction [Kruger et al. 1993] and to perceive depth order [Nguyen et al. 2005; Zannoli et al. 2016]. One such optical signal that specifies sign is chromatic aberration, and we capitalize on it here.

In natural viewing, the rate at which blur changes with distance from the focal plane–i.e., depth-of-field (DoF) blur–depends on focal distance, object distances relative to that distance, and pupil diameter. DoF blur has powerful effects on depth perception and perceived realism. If DoF is increased, the scene appears magnified; if it is decreased, the scene looks miniature. Thus, setting DoF to the appropriate value enables accurate perception of scale and depth [Held et al. 2010; Vishwanath and Blaser 2010] and increases perceived realism [Mauderer et al. 2014; Zhang et al. 2015].

Traditionally, rendering has focused on "photorealism": simulating images from a camera with a pinhole or idealized lenses without aberrations. Even simulation of realistic optics has focused on ray-tracing camera lenses [Kolb et al. 1995; Ng and Hanrahan 2006] with well-corrected aberrations. In contrast, rendering for virtual reality (VR) should emphasize "perceptual realism" to enable immersion in a virtual environment. In this quest for perceptual realism, an important but neglected property of the visual system is the imperfect optics of the human eye. Real images on the retina are aberrated and are therefore not close to photorealistic images. Here we describe how to implement rendering that incorporates natural aberrations; this involves calculating what the retinal image should be and then computing the displayed image that, when processed through the viewer's aberrated eye, creates the intended, natural retinal image. We concentrate on two significant and universal aberrations[1]: defocus and chromatic aberration.

Our contribution is three-fold. 1) We develop a color-correct rendering method–*ChromaBlur*–that incorporates defocus and the eye's natural chromatic aberration to generate retinal images that are close to those created in viewing the real world. 2) We show that this rendering method drives accommodation effectively. 3) We show that the method yields a greater impression of depth than conventional rendering.

2 RELATED WORK

Realistic image synthesis in graphics has focused on computing the images produced by camera systems. Kolb et al. [1995] introduced realistic camera models by physically-based ray tracing [Pharr and Humphreys 2014] of optical formulas for camera lenses. Steinert et al. [2011] extended this approach to include spectral effects such as chromatic aberration. Our work takes a different approach by creating images to display that will create realistic retinal images rather than minimally aberrated images like those inside a camera.

[1]By significant we mean that the magnitude of the effect on the retinal image is large enough to be perceived. By universal we mean that there is little, if any, variation across individuals.

Some graphics researchers [Barsky 2004; Kakimoto et al. 2007; Mostafawy et al. 1997; Nießner et al. 2012; Tang and Xiao 2015; Wu et al. 2011] have pursued rendering of retinal images with the goal of visualizing optical effects due to human eye aberrations. Most used ray-tracing in the style of Kolb, employing accurate eye models from the vision science literature [Polans et al. 2015]. Barsky used Shack-Hartmann wavefront measurements of highly aberrated eyes. These graphics algorithms have many uses including giving customers an impression of how different types of correction with spectacles, contact lenses, or LASIK surgery will aid their vision. But the goal in that previous work was quite different from ours.

Deering's eye model [2005] accounts for chromatic aberration because he wanted to calculate an accurate photon count at each photoreceptor. One of his goals was principled evaluation of display system designs, but his model stopped at photometric calculations before considering human perception or response. In contrast, a central contribution of our work is measuring perceptual and oculomotor responses due to the eye's defocus and chromatic aberration.

Important perceptual and ergonomic issues arise from the way blur is reproduced and whether accommodation is enabled or not. Many of the issues are due to the vergence-accommodation conflict. Vergence and accommodation are neurally coupled [Schor 1992], which is beneficial in the real world where the distances to which the eyes should converge and focus are always the same. But the coupling is broken by conventional stereoscopic displays because such displays require the viewer to converge to one distance (that of the virtual object) while accommodating to another (the display screen). The resulting vergence-accommodation conflict causes visual discomfort [Hoffman et al. 2008; Koulieris et al. 2017; Lambooij et al. 2009; Shibata et al. 2011], reductions in performance [Akeley et al. 2004; Johnson et al. 2016; Konrad et al. 2016; Maiello et al. 2014], and distortions of perceived depth [Watt et al. 2005]. Our work provides an opportunity to control accommodation more effectively and thereby minimize the vergence-accommodation conflict.

3 OPTICAL ABERRATIONS OF THE EYE

Although the human eye has a variety of field-dependent optical imperfections, we restrict our analysis to on-axis effects because optical imperfections are much more noticeable near the fovea and because optical quality is reasonably constant over the central 10° of the visual field [Navarro et al. 1993]. In this section, we describe why we choose to incorporate only defocus and chromatic aberration in our rendering method and why we ignored other imperfections that could have been incorporated.

3.1 Defocus

Defocus is caused by the eye being focused at a different distance than the object. In most eyes defocus constitutes the great majority of the total deviation from an ideal optical system [Cheng et al. 2004; Porter et al. 2001]. The function of accommodation is to minimize defocus. The point-spread function (PSF) due to defocus alone is a disk whose diameter depends on the magnitude of defocus and diameter of the pupil. The disk diameter is given to close approximation by:

$$\beta \approx A\left|\frac{1}{z_0} - \frac{1}{z_1}\right| = A\,|\Delta D| \tag{1}$$

where β is in angular units, A is pupil diameter, z_0 is distance to which the eye is focused, z_1 is distance to the object creating the blurred image, and ΔD is the difference in those distances in diopters [Held et al. 2010]. Importantly, the PSF due to defocus alone is identical whether the object is farther or nearer than the eye's current focus. Thus, rendering of defocus is the same for far and near parts of the scene.

3.2 Chromatic Aberration

The eye's optical elements have different refractive indices for different wavelengths. Short wavelengths (e.g., blue) are refracted more than long (red), so blue and red images tend to be focused, respectively, in front of and behind the retina (Fig. 2). The wavelength-dependent difference in focal distance is longitudinal chromatic aberration (LCA). In diopters it is:

$$D(\lambda) = 1.731 - \frac{633.46}{\lambda - 214.10} \quad (2)$$

for an eye in-focus at 580nm, where λ is measured in nanometers [Marimont and Wandell 1994]. The magnitude of LCA is the same in all adult eyes [Nakajima et al. 2015; Thibos et al. 1992].

When the eye views a depth-varying scene, LCA produces different color effects (e.g., colored fringes) for different object distances relative to the current focus distance. For example, when the eye is focused on a white point, green is sharp in the retinal image and red and blue are not, so a purple fringe is seen around a sharp greenish center. But when the eye is focused nearer than the white point, the image has a sharp red center surrounded by a blue fringe. For far focus, the image has a blue center and red fringe. Thus, LCA can in principle indicate whether the eye is well focused and, if it is not, in which direction it should accommodate to restore sharp focus.

These color effects are generally not consciously perceived, but they definitely affect visual function. Kruger and colleagues examined LCA's role in accommodation [Aggarwala et al. 1995; Kruger et al. 1993]. They presented stimuli of constant retinal size to one eye and measured accommodative responses to changes in focal distance. Using special lenses, they manipulated LCA. Accommodation was accurate when LCA was unaltered and much less accurate when LCA was nulled or reversed. Some subjects even accommodated in the wrong direction when LCA was reversed. There is also evidence that LCA affects depth perception. Zannoli et al. [2016] presented two broadband abutting surfaces monocularly at different focal distances. Subjects perceived depth order correctly. But when the wavelength spectrum of the stimulus was made narrower (making LCA less useful), performance declined significantly. These accommodation and perception results are good evidence that LCA contributes to visual function even though the resulting color fringes are often not perceived.

3.3 Other Aberrations

Spherical aberration and uncorrected astigmatism have noticeable effects on the retinal image [López-Gil et al. 2007; Peters 1961] and could signal in which direction the eye must accommodate to sharpen the image. Our rendering method can in principle incorporate those optical effects, but we did not do so because these effects vary across individuals so no universal rendering solution is feasible for them [Cheng et al. 2004; Salmon and van de Pol 2006]. Diffraction is universal, but has negligible effect on the retinal image except when the pupil is very small. Our consideration of on-axis effects only is justified because LCA and the eye's other aberrations are fairly invariant over the central 30° of the visual field [Liu and Thibos 2017; Rynders et al. 1998]. The exception is astigmatism which increases in magnitude with increasing retinal eccentricity [Liu and Thibos 2017].

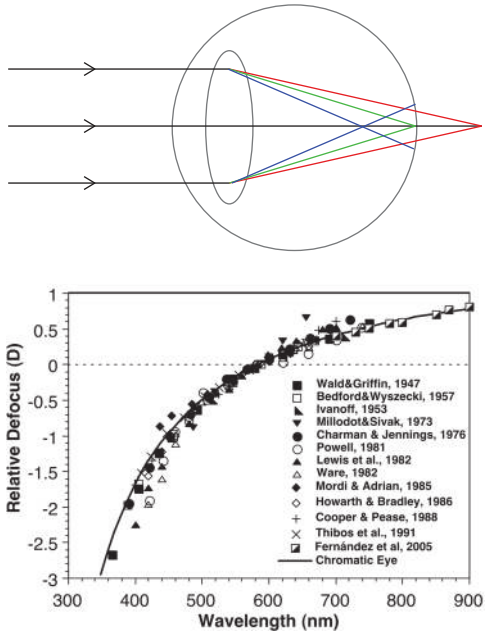

Fig. 2. Longitudinal chromatic aberration of the human eye. Top: Short wavelengths (blue) are refracted more than long (red). Medium wavelengths are generally in best focus for broadband lights. Bottom: Relative defocus in diopters as a function of wavelength. The data were adjusted such that defocus is zero at 589nm. From Thibos et al. [1992].

4 RENDERING METHOD

ChromaBlur aims to compute the pixel values to be shown on a display that will best recreate a realistic retinal image.

4.1 Optical Model (Forward Algorithm)

The conventional procedures for computing blur are quite different from ours. In vision science, defocus is almost always approximated by convolving parts of the scene with a 2d Gaussian [Mather and Smith 2002; Subedar and Karam 2016; Watson and Ahumada 2011]. In graphics, defocus is usually approximated by an ideal lens without chromatic aberration.

Our model for rendering incorporates defocus and LCA. It could include other optical effects, but we ignore these here in the interest of simplicity and universality (Sec. 3.3).

4.2 Conceptual Rendering Algorithm

The approach can be thought of as a two-step process. The first step is a forward calculation: Determine the target 2d image representing

the signal that would have appeared on the retina in response to a given scene. The second step is an inverse problem: Determine the pixel values for an image to display that, when viewed by the eye, will produce a retinal image that best approximates the target image. In this step we assume that the viewing eye is accommodated on the display screen, so that the green primary is focused on the retina while red and blue are not. We created two algorithms (2d and 3d), which are quite different in operation, but are both based on the two-step process.

4.3 2d ChromaBlur Algorithm

The accommodation experiments presented here use a 2d scene composed of a textured plane at different distances. We render these scenes with the 2d ChromaBlur algorithm, which is a physical (wave) optics simulation incorporating defocus, LCA, and diffraction. Given a target retinal image $I_{\{R,G,B\}}(x,y)$, we compute the image $D_{\{R,G,B\}}(x,y)$ to display on the screen. For each color primary, we have a wavelength-dependent blur kernel, $K_{\{R,G,B\}}(x,y)$, which is a PSF calculated from the square of the Fourier transform of the eye's complex aperture function (which takes into account the amplitude and phase of light). The target retinal image is therefore the 2d convolution ($**$) of the display image with the eye's PSF for each of the three color channels ($\{R,G,B\}$):

$$I_{\{R,G,B\}}(x,y) = D_{\{R,G,B\}}(x,y) **K_{\{R,G,B\}}(x,y). \quad (3)$$

We generated values of $D_{\{R,G,B\}}$ by varying induced defocus and compared these with the forward-model solution until we found the optimum image to display.

This model works well for 2d scenes and includes diffraction. But it is not based on conventional graphics techniques and is not suitable for complex 3d scenes. So we produced a second method based on alternating direction method of multipliers (ADMM) deconvolution.

4.4 3d ChromaBlur Algorithm

The 3d algorithm enables application to general graphics systems for rendering complex scenes. It was used for rendering the kitchen scene [Jay-Artist 2012] in Fig. 1 and the chessboard scenes [Hoyt 2016] in Figs. 6 and 12 and for one of the experiments (Sec. 7).

4.4.1 Forward Step: Monte Carlo Ray Tracing. We compute the target retinal image using Monte Carlo ray-tracing. This is a simple matter of using a finite-aperture lens in the camera model, where the focal length of the lens varies according to wavelength as defined by Eqn. 2. We approximate each primary with one representative wavelength. We used the physically based renderer Mitsuba [Jakob 2010] to implement this step for the results shown in the paper.

4.4.2 Inverse Step: ADMM Optimization. We again use Eqn. 3 as the basic model. For each color primary, we compute the display image $D(x,y)$ by solving the following optimization equation, where we omit R,G,B for brevity:

$$\min_{D(x,y)} ||D(x,y) **K(x,y) - I(x,y)||_2^2 + \psi||\nabla D(x,y)||_1$$
$$\text{such that } 0 \le D(x,y) \le 1 \quad (4)$$

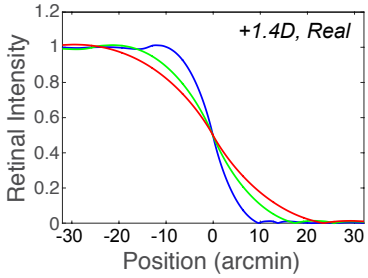

Fig. 3. Retinal image created by a real step edge. A high-contrast, white-black edge is viewed with the eye defocused by +1.4D. Red, green, and blue curves represent cross-sections for the R, G, and B primaries, respectively. Pupil diameter is 6mm.

The first term in Eqn. 4 is a data term that is the L_2 norm of the error between the target retinal image and the retinal image produced by the displayed image. The deconvolution problem defined by the data term alone is generally ill-posed due to zeros in the Fourier transform of the kernels. Therefore, the second term in Eqn. 4 is an L_1 regularization term comprising a total variation image prior, which corresponds to a prior belief that the displayed image solution is sparse in the gradient domain. Finally, we constrain the estimated displayed image to be between 0 and 1, the minimum and maximum display intensities. In our case, the residual will not be zero due to the constraint that the displayed image must be bounded by 0 and 1, and due to the regularization term, which reduces unnatural artifacts such as ringing.

The blur kernels K are cylinder functions, but in solving Eqn. 4, we smooth them slightly to minimize ringing artifacts.

The regularized deconvolution optimization problem in Eqn. 4 is convex, but it is not differentiable everywhere due to the L_1 norm. As a result, there is no straightforward analytical expression for the solution. We therefore solve the deconvolution using ADMM [Boyd et al. 2011], a standard algorithm for such problems. ADMM splits the problem into subproblems that are solved iteratively. For many problems, including ours, the subproblems have closed-form solutions that are efficient to compute. Furthermore, both the data and regularization terms in Eqn. 4 are convex, closed, and proper, so ADMM is guaranteed to converge to a global solution. In our implementation, we use a regularization weight of $\psi = 1.0$, ADMM hyperparameter $\rho = 0.001$, and run the algorithm for 100 iterations.

5 RENDERING RESULTS

We next consider the retinal images created by a defocused eye viewing real objects and displayed images created by different rendering techniques including ours. We first examine images formed by a white-black step edge.

5.1 Step Edges

Fig. 3 shows the variation in intensity across the retina due to an edge that is 1.4D nearer than the eye's current focus (positive defocus). The red, green, and blue curves represent the variation in retinal intensity for the R, G, and B primaries. With the real edge nearer than current focus, blue is sharper than green and red, so a blueish fringe is created. When the object is farther than current

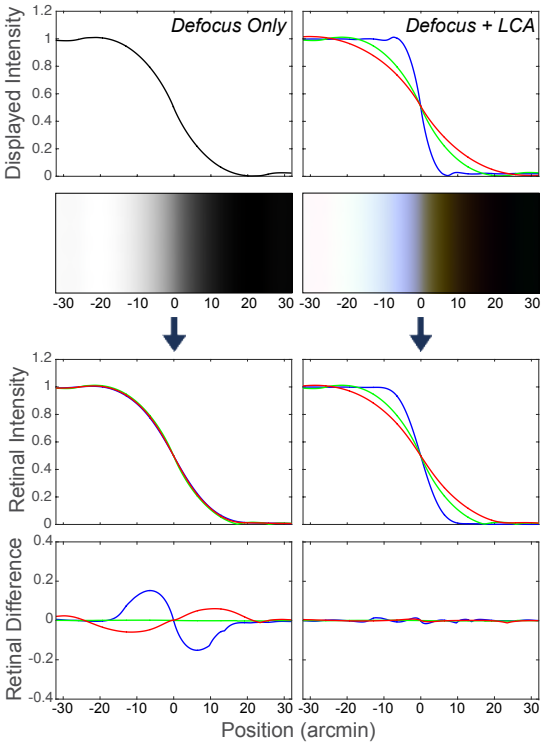

Fig. 4. Displayed images and retinal images created by conventional and ChromaBlur rendering. A white-black step edge is rendered to create the appearance that it is 1.4D closer than current focus. Left: Displayed and retinal images based on conventional rendering. The upper panel is a cross-section of the displayed image. R, G, and B are superimposed. The next panel is the corresponding 2d image. The next panel shows cross-sections of the retinal image for R, G, and B. G is slightly sharper than R and B. The bottom panel shows the differences between the retinal images for a real edge and the simulation of that edge for R, G, and B. Right: Displayed and retinal images based on ChromaBlur. The errors in the retinal image are now much smaller than with conventional rendering.

focus (negative defocus), the opposite occurs. Fig. 4 illustrates the displayed and retinal images for simulations of the same situation. The viewer's eye is focused on the display screen. The left column shows the displayed image and resulting retinal image based on conventional rendering in which all wavelengths in the displayed image are blurred by the same amount. Green is slightly sharper in the retinal image than blue and red due to the viewer's LCA, but this effect is swamped by the rendered blur. The bottom panel shows the differences for red, green, and blue between the retinal image from the real edge in Fig. 3 and the retinal image from the conventionally rendered displayed image. Errors occur in blue and red; blue because the retinal image from conventional rendering is not sharp enough at that wavelength and red because the conventional image is too sharp. The right column shows the displayed and retinal images based on ChromaBlur in which LCA is taken into account. The bottom panel again shows the retinal-image differences between real and simulated edges. The differences are very small at all three

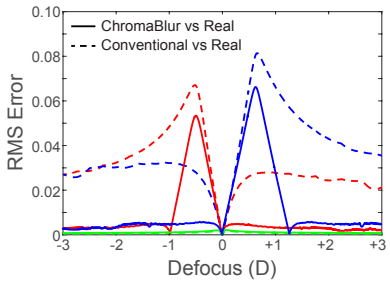

Fig. 5. Errors in retinal image for conventional and ChromaBlur rendering. RMS error between the correct retinal intensities and the intensities created by the two rendering techniques are plotted as a function of defocus. Red, green, and blue curves represent the errors for the R, G, and B primaries, respectively. Solid curves are the errors for ChromaBlur and dashed curves those for conventional.

wavelengths, which shows that ChromaBlur creates retinal images that are very close to natural ones.

We quantified the errors between real and rendered retinal images by computing the RMS error for R, G, and B as a function of defocus. For each defocus value and wavelength we computed: $\sqrt{\frac{1}{N}\sum_{x=-N/2}^{N/2}[R_{real}(x)-R_{sim}(x)]^2}$, where x is position, $R_{real}(x)$ and $R_{sim}(x)$ are the retinal intensities for real and simulated edges, respectively, and N is the number of sampled positions (N = 65 for these computations). Fig. 5 shows the results. Conventional rendering produces quite erroneous results at all defocus magnitudes greater than ~0.25D. ChromaBlur produces much more accurate results when the defocus magnitude is greater than 1D and when defocus is ~0. It yields somewhat erroneous results at small defocus values although the error is still smaller than with conventional rendering. The cause of that error is obvious. With small positive defocus in real scenes, blue becomes more sharply focused in the retinal image than when defocus is 0. Likewise, with small negative defocus, red becomes sharper. To simulate this, we would have to over-sharpen the displayed image for blue and red to simulate positive and negative defocus, respectively. This cannot be done when the object contrast is ~100%. It can of course be solved for lower contrasts.

There is another approximation in our method. Real color primaries, such as those in LCDs, have broad spectra. We approximate this with one wavelength for each primary. This does not reproduce the correct blur for off-peak wavelengths, but the errors are generally quite small (details in supplemental material).

5.2 Complex 3d Scenes

We next consider the displayed and retinal images for complex scenes. The upper part of Fig. 6 shows from left to right the retinal image associated with a real 3d scene, a magnified version of that image, the displayed image for that scene when rendered conventionally, and the displayed image for the scene when rendered with ChromaBlur. The focus distance is 2.5D; some parts of the scene are farther than that and some are nearer. The lower left panel shows the set of distances in the scene for a horizontal cross-section. The focus distance is indicated by the gray line. The lower middle panel shows the errors with conventional rendering for R, G, and B. The errors

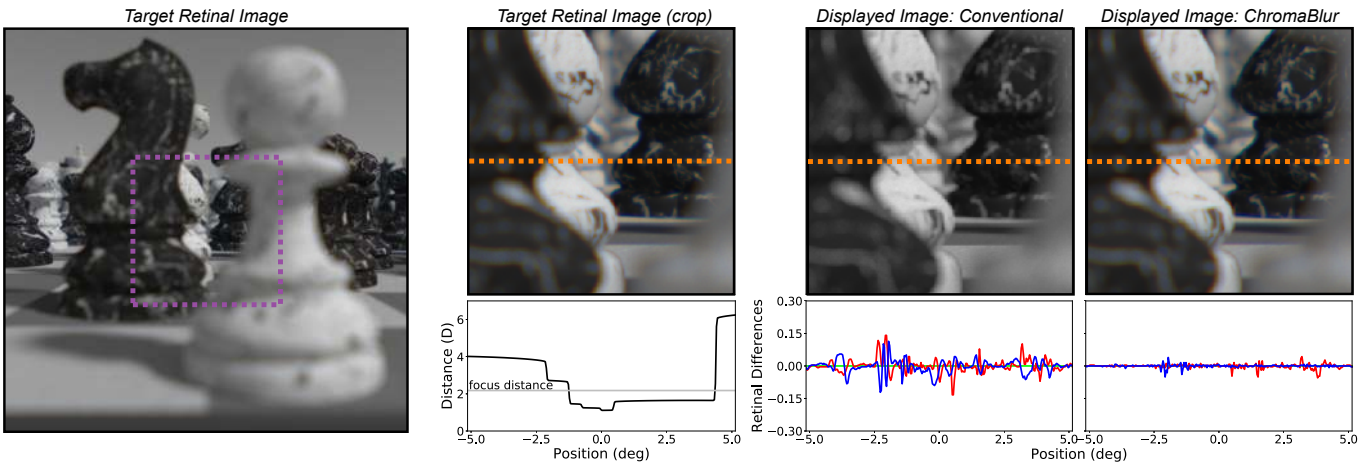

Fig. 6. Complex scenes and conventional and ChromaBlur rendering. Left: Retinal image for real 3d scene. Upper left: Crop of retinal image from the purple dashed region in full scene on left. Upper middle: Displayed image using conventional rendering for the same scene. Upper right: Displayed image using ChromaBlur. Lower left: Distance map for the scene. Distance in diopters is plotted for a horizontal cross-section through the scene. The gray line indicates the simulated focus distance. Lower middle: Differences in intensity between retinal images produced by conventional rendering vs retinal images produced by the real scene. The red, green, and blue curves represent the differences for R, G, and B, respectively. The difference is zero for G because the viewer is focused at the display screen. Lower right: Differences between retinal images produced by ChromaBlur minus retinal images produced by the real scene.

are the intensity differences between the retinal image produced by the real scene vs the retinal image produced by conventional rendering of that scene. The lower right panel shows the same for ChromaBlur vs real. Clearly, ChromaBlur produces a retinal image that is much more similar to the desired retinal image.

In summary, ChromaBlur produces excellent results for small and large defocus magnitudes. It can also produce excellent results for other magnitudes at medium and low contrasts.

6 DOES CHROMABLUR DRIVE ACCOMMODATION?

We want to know how effective ChromaBlur is in creating natural focus cues. We first investigated this by seeing if it helps drive the eye's accommodative response.

6.1 Methods

6.1.1 Subjects. Five naïve subjects (18-29 years) participated. All were female to ensure normal color vision [Sharpe et al. 1999]. Three were myopic and wore their contact-lens correction during the experiment. The other two did not require correction. All had normal acuity.

6.1.2 Apparatus. Fig. 7 shows the experimental setup. The left panel is a schematic of the light paths to the two eyes; the right panel shows a subject in the experiment. We stimulated the left eye while measuring accommodation in the right eye. Accommodation is completely yoked between eyes [Campbell 1960; Fisher et al. 1987], so this method is well justified. Stimuli were projected onto a screen using a DLP projector with a resolution of 1920×1080. The R, G, and B primaries were LEDs with relatively narrow spectra (supplemental material for details). The projection screen was 1.38m (0.72D) from the subject's eye and subtended 35.7×20.1°. Nyquist frequency was 27cycles/deg. The room was dark except for the projected stimulus.

A focus-adjustable lens (Optotune EL-10-30-VIS-LD) was placed just in front of the stimulated eye. We varied the power of this lens to manipulate the focal distance of the stimulus. The lens has nearly zero LCA (Abbe number = 100). We placed a -10D achromatic lens in the optical path to give a range of potential focal distances at the eye of -3.8 to +10.2D (supplemental material for details).

We measured accommodation with an autorefractor (Grand Seiko WV-500). Autorefractors are used in optometric eye examinations to measure refractive error (e.g., myopia, hyperopia, astigmatism). The device projects infrared light into the eye and records the image reflected from the retina. In its normal operating mode, sampling rate is 1Hz. But the composite video signal provides a much higher rate. Using a method similar to the one described by Wolffsohn [2004] and MacKenzie [2010], we were able to measure accommodation at 30Hz by processing the video offline (supplemental material for details). We removed data corrupted by eye blinks or eye movements.

6.1.3 Procedure. Subjects viewed the stimuli with their left eye. They first fixated and accommodated to a textured plane at 1.5D. For convenience, we refer to that distance as 0D and then describe the distances of the experimental stimuli as changes relative to 0. Subjects were instructed to fixate and focus carefully on the fixation stimulus whenever it appeared. The experimental stimulus was presented for 3sec at another optical or simulated distance. The screen then went blank (uniform gray) for 1.5sec before the fixation stimulus reappeared.

There were three conditions: *Real Change* in which the focal distance of the stimulus changed due to a change in the power of the focus-adjustable lens; *Defocus Only* in which focal distance did not change, but rendered blur changed by the same amount for all three color primaries (conventional rendering); *Defocus plus LCA* in which focal distance did not change, but rendered blur changed appropriately for each primary (ChromaBlur rendering). Stimulus size at the retina never varied. In all conditions, the changes in stimulus distance (real or simulated) were from 0D to ±0.6, ±1, or ±1.4D. Pupil diameter in the simulated conditions was assumed to

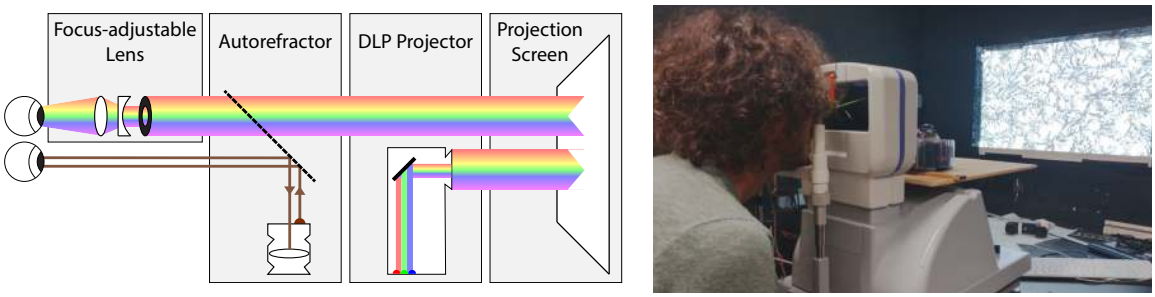

Fig. 7. Experimental apparatus. Left: Schematic. The DLP projector delivered images to the projection screen. The color primaries were three LEDs. The subject viewed the stimulus on the projection screen with the left eye. A focus-adjustable lens, fixed lens, and aperture were placed just in front of that eye, so the subject viewed the stimuli through the lenses and aperture. An autorefractor delivered infrared light to the right eye. The light was reflected by a hot mirror that reflects infrared but transmits visible light. The reflection of that light from the right eye's retina was recorded by the autorefractor's video camera. Right: The subject was positioned with a chin-and-head rest in order to see the stimulus with the left eye while accommodation of the right eye was measured.

be 6mm, which is close to the actual diameters (supplemental material). Conditions and distance changes were presented in random order. 252 trials were presented to each subject: 144 *Real Change*, 36 *Defocus Only*, and 72 *Defocus plus LCA*.

6.2 Results

Fig. 8 shows the accommodative responses in the *Real Change* condition for a representative subject. The stimulus changed from 0 to +1.4D at 0sec. The traces are responses on individual trials. This subject made consistent, though variable responses to the change in focal distance.

The data from all conditions are shown in Fig. 9. They have been subjected to a running median. Data for positive and negative stimulus changes are represented respectively by red and green. Responses in the *Real Change* condition were appropriate: positive when the stimulus change was positive and negative when the change was negative. As is generally observed [Kruger et al. 1993], response magnitude was somewhat less than stimulus magnitude. Essentially no accommodative responses were observed in the *Defocus Only* condition. This means that conventional rendering of defocus blur does not stimulate accommodation. Accommodative responses in the *Defocus plus LCA* conditions were robust and closely tied to the change in simulated distance. Indeed, responses were very similar to those in the *Real Change* condition. Subjects accommodated in the direction specified by simulated LCA despite the fact that the focal distance of the stimulus did not change. This shows rather remarkably that appropriate rendering of LCA actually drives accommodation. The other four subjects yielded similar data (see supplemental material).

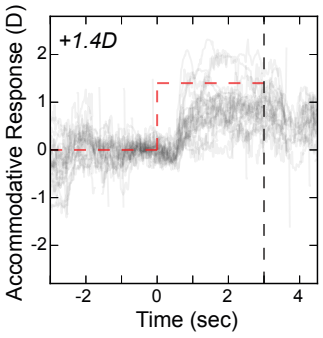

Fig. 8. Accommodative responses in the Real Change condition for one subject. Change in focal distance was +1.4D. The stimulus is indicated by the red dashed line. The vertical dashed line indicates when the screen went blank. Each trace is the response recorded on one trial.

To assess statistical reliability, we subjected the data from 0.5–3sec to a repeated-measures analysis of variance (ANOVA) with factors of subject, condition, and distance change. All three factors yielded statistically significant effects. We then conducted multiple pairwise comparisons using Tukey Contrasts with Bonferroni adjustment for multiple comparisons. Those results showed that accommodation was significantly greater with *Real Change* than with *Defocus Only* for stimulus changes of -1.4, -1.0, -0.6, +1.0, and +1.4D (i.e., not for +0.6D). Accommodation was significantly greater with *Defocus plus LCA* than with *Defocus Only* for changes of -1.4, -1.0, +1.0, and +1.4D (i.e., not ±0.6D). Accommodation was greater with *Real Change* than with *Defocus plus LCA* for a change of -1.4D (i.e., not for all the other values). Thus, responses were consistently larger in the *Real Change* and *Defocus plus LCA* conditions than in the *Defocus Only* condition and were generally not different in the *Real Change* and *Defocus plus LCA* conditions.

In summary, ChromaBlur rendering produced consistent accommodative responses that were remarkably similar to those produced by real changes in focal distance. And it did so despite our approximation of each primary with one wavelength (Fig. S3) and despite errors when simulating small amounts of defocus (Fig. 5).

The results are a novel contribution not only to graphics but also to the vision science literature. In previous accommodation research in vision science, the focal distance of the stimulus was manipulated optically to create two conditions: one in which all focus cues (LCA, defocus, higher-order aberrations, and micro-fluctuations) were informative and specified the change in focal distance, and one in which LCA was made uninformative, but the other focus cues remained informative [Aggarwala et al. 1995; Kruger et al. 1993]. Accommodative responses in some individuals became less accurate when LCA was uninformative, so we can conclude from that work that LCA plays a role in driving accommodation in those individuals. In our experiment, all focus cues were informative in the *Real Change* condition. But in the *Defocus plus LCA* condition, only LCA specified

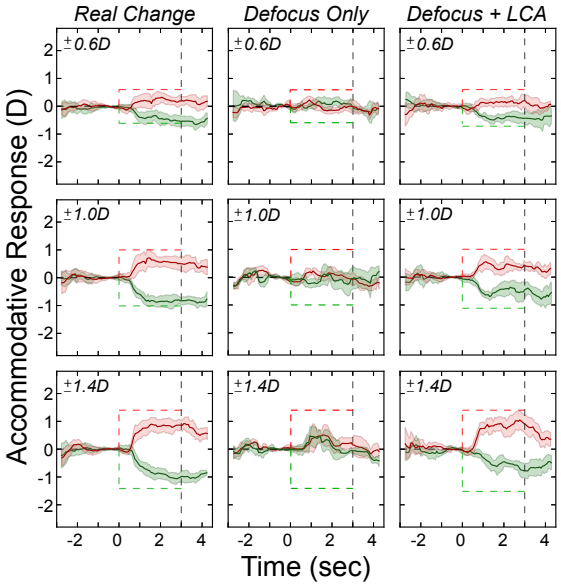

Fig. 9. Smoothed accommodative responses for one subject. The response recordings were subjected to a running median with a window of 500msec. Dashed lines represent the stimulus; red for positive and green for negative. The shaded regions are median absolute deviations. Thick curves in the middle of the shaded regions are the medians. The columns from left to right are for *Real Change*, *Defocus Only*, and *Defocus plus LCA*. The rows from top to bottom are for ±0.6, ±1.0, and ±1.4D changes.

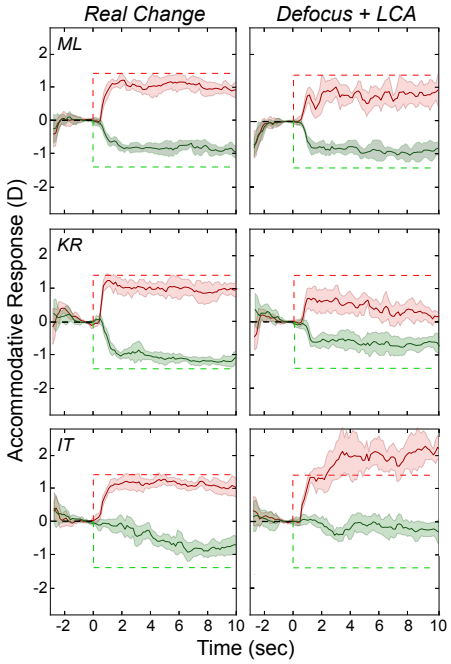

Fig. 10. Accommodative responses over 10sec. Responses were subjected to a running median. Dashed lines represent the stimulus; red for positive and green for negative. Shaded regions are median absolute deviations. Thick curves are medians. Left and right columns are for *Real Change* and *Defocus plus LCA*. Each row shows the data from one subject.

a change in focal distance; the other cues specified no change. Thus, our results show for the first time that one can drive accommodation effectively by manipulating LCA alone despite counter-information from defocus, higher-order aberrations and micro-fluctuations.

6.2.1 Persistence of response. It is perhaps not surprising that subjects initially accommodated in the direction of the simulated change in the *Defocus plus LCA* condition; we say this because the retinal images produced by ChromaBlur are quite similar to those produced by real changes in focal distance (Figs. 4 and 5). But it is quite surprising that subjects maintained the response for 3sec because accommodating caused an increase in the overall blur of the retinal image. Specifically, the optical distance of the stimulus in the *Defocus plus LCA* condition was always 0D. Applying ChromaBlur yielded an increase in the rendered blur for at least two of the color primaries. By accommodating away from 0D, subjects caused additional blur in the retinal image. It would have been less surprising if subjects responded initially in the specified direction, and then "backed up" to restore image sharpness. To see if such "backing up" occurs, we lengthened stimulus duration to 10sec and repeated the experiment in three subjects. We presented the *Real Change* and *Defocus plus LCA* conditions for stimulus changes of ±1.4D.

Fig. 10 shows the results. Each row shows the data from one subject. Left and right columns show the results from the *Real Change* and *Defocus plus LCA* conditions, respectively. Accommodative responses to the simulated change in distance (i.e., ChromaBlur) persisted for at least 10sec. Thus, subjects did indeed accommodate to the changes in stimulus distance specified by our rendering technique and maintained the response well after the visual system could ascertain that overall blur had increased. Our result raises the possibility that the visual system uses the relative amount of blur at short and long wavelengths, rather than overall sharpness, to guide accommodation.

6.2.2 Resolution. Most HMDs have limited resolution. The HTC Vive, Oculus DK1, Oculus DK2, Sony Playstation VR, and Samsung Gear VR have pixels subtending 3.2–6.4minarc (Nyquist frequencies of 4.7–9.4cycles/deg). We examined whether such low resolution prohibits the effectiveness of ChromaBlur in driving accommodation.

Native pixel size in our apparatus was 1.1×1.1minarc. By resampling (supplemental material for details), we generated images with simulated pixels that were 2.2, 4.4, 8.8, 17.6, and 35.2minarc in height and width. We presented the *Real Change* and *Defocus plus LCA* conditions with these resolutions for real and simulated defocus values of ±0.6, ±1.0, and ±1.4D. Three young adults were tested. The apparatus and procedure were the same as before.

The images with pixels of 4.4–35.6minarc created visible boundaries between adjacent simulated pixels, which creates an interesting set of signals. In the *Real Change* condition, the pixel boundaries became blurred along with the scene content when the focus-adjustable lens changed power. Thus, the boundaries and scene content provided consistent signals to drive accommodation. This

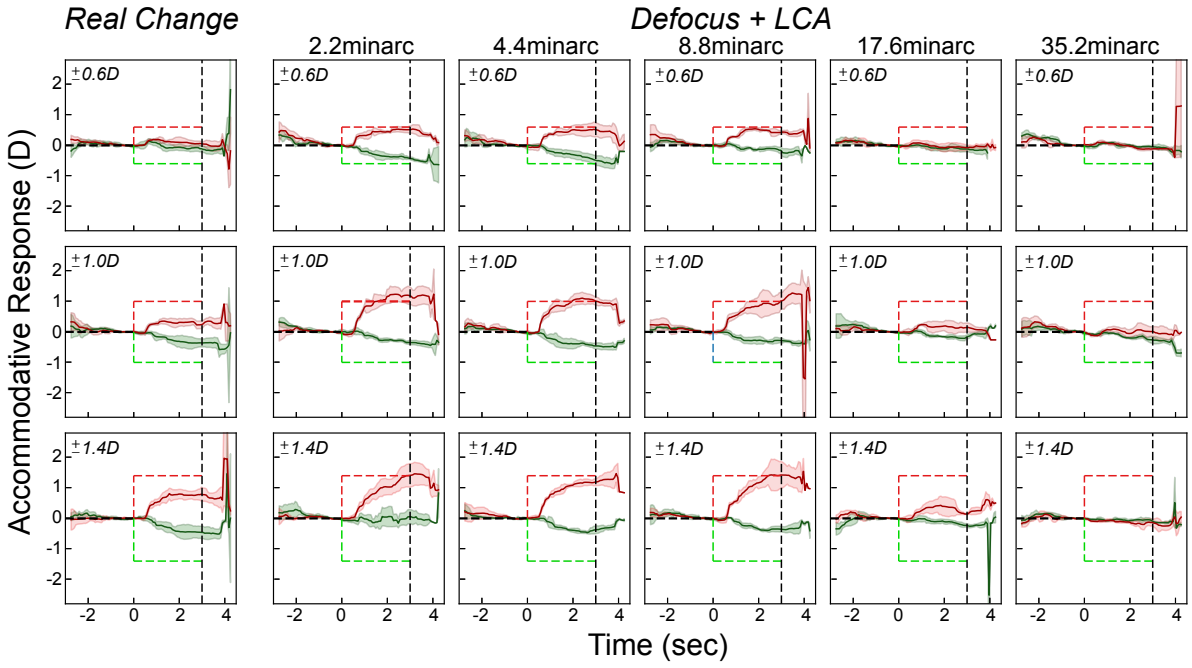

Fig. 11. Accommodative responses and resolution in one subject. The first column shows the data for the *Real Change* condition, collapsed across resolutions. Remaining columns show the data for the *Defocus plus LCA* condition for resolutions of 2.2, 4.4, 8.8, 17.6, and 35.2minarc. The rows show the data for real or simulated defocus values of ±0.6, ±1.0, and ±1.4D. Shaded regions are median absolute deviations; red for positive changes and green for negative. The large deviations after stimulus offset are due to blinks.

was not the case in the *Defocus plus LCA* condition. Here the focus-adjustable lens did not change power so the pixel boundaries signaled that the eye should remain accommodated to the screen distance while the scene content, which was subjected to ChromaBlur, signaled that the eye should accommodate to the distance specified by defocus and color. Thus, the two signals provided conflicting information akin to the screen-door effect in which focus tends to become locked at a screen door even though the viewer is attending to the scene beyond the door [Owens 1979].

Fig. 11 shows the results for one subject. (Results for the others in supplemental material.) The 1st column shows the data from the *Real Change* condition collapsed across resolutions. The 2nd through 7th columns show the data from the *Defocus plus LCA* condition with pixel size doubling for each successive column. The rows from top to bottom show the data for increasing amounts of real or simulated defocus. Accommodation was driven quite effectively in the *Real Change* condition because all signals (pixel boundaries and scene content) indicated how the eye should accommodate. Interestingly, accommodation was driven quite effectively in the *Defocus plus LCA* condition at resolutions of 2.2, 4.4, and 8.8minarc. This is somewhat surprising because the pixel boundaries were clearly visible at 4.4 and 8.8minarc and they provided a signal to hold accommodation constant at the screen distance.

We conclude that ChromaBlur applied in the current generation of HMDs (many of which have pixels of 5minarc or smaller) should effectively stimulate accommodation. In supplemental material, we examine how down-sampling ought to affect the usefulness of LCA, and then compare that analysis to our experimental results.

7 DOES CHROMABLUR ENHANCE REALISM?

We next investigated whether ChromaBlur enhances the impression of real depth in complex 3d scenes.

7.1 Methods

7.1.1 Subjects. Five naïve subjects (23-31 years, four females) participated. All had normal or corrected-to-normal acuity, and normal color vision. One subject wore corrective contact lenses during the experiment; the others did not because their refractive errors were small.

7.1.2 Apparatus. We used the same apparatus as in the accommodation experiments, but we did not measure accommodation.

7.1.3 Stimuli. The stimuli were generated with three rendering methods: *Conventional* in which all colors are treated the same, *ChromaBlur*, and *Reverse ChromaBlur* in which LCA is inverted. The stimuli were close-in views of a chessboard [Hoyt 2016] and contained a rich variety of monocular depth cues (e.g., perspective, occlusion, shading, etc.) (Fig. 12). Subjects viewed them with the left eye.

The target retinal images (forward model) were generated in Mitsuba [Jakob 2010] as described in Sec. 4.4. The *Conventional* images derived from the forward model were also the displayed images for that condition. The *Reverse ChromaBlur* images were

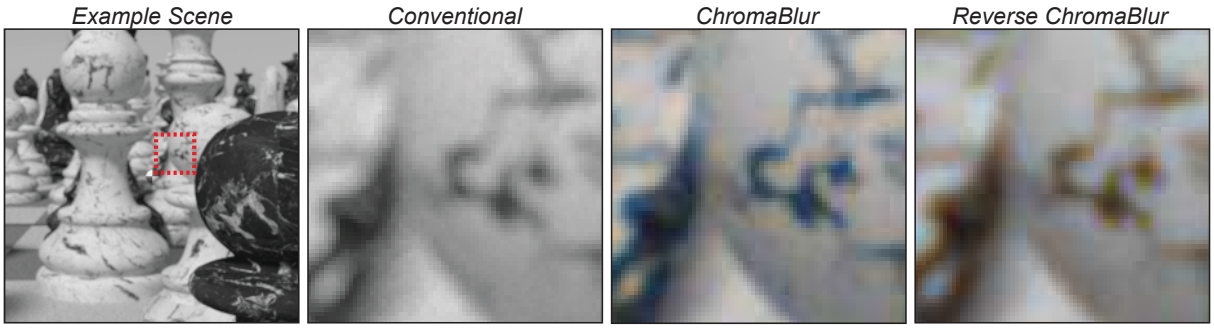

Fig. 12. Stimuli from the perceived depth experiment. Left: Example of the stimulus from Conventional rendering with the ball 3.5D from the viewer. The next three panels are enlargements of the red dashed area on the left. From left to right, they are Conventional, ChromaBlur, and Reverse ChromaBlur rendering. Note the color fringing in the latter two.

generated by inverting the eye model in Eqn. 2 to produce reversed defocus for R and B. They were then deconvolved in the same manner as the *ChromaBlur* images to produce the desired displayed images.

Each stimulus contained a small ball with a high-contrast texture; this served as the fixation target. On each trial, it was presented at one of five focal distances: 0.5, 1.25, 2.0, 2.75, or 3.5D. Subjects were told to fixate and focus on the ball throughout a trial.

7.1.4 Procedure. On each trial, the focus-adjustable lens was first set to the focal distance of the ball, which was shown for 1.5sec while the rest of the scene (rendered conventionally) was presented at low contrast. Then the first stimulus appeared for 750ms by raising the contrast of the scene to the appropriate value. The rendering method for the scene was either *Conventional*, *ChromaBlur*, or *Reverse ChromaBlur*. The contrast of the scene (rendered conventionally) was then again reduced for 250ms while the ball remained clearly visible. The second stimulus then was presented for 750ms by raising the contrast of the scene to the appropriate value. Scene rendering was again *Conventional*, *ChromaBlur*, or *Reverse ChromaBlur*, but not the same method as in the first stimulus. The contrast of the scene was then again reduced (rendered conventionally) for 250ms while the ball remained clearly visible. This marked the end of a trial. The subject responded by indicating, in a forced-choice judgment, which of the two stimuli had created a stronger impression of realistic depth. No feedback was provided. Each trial thus consisted of one of three pairwise comparisons: *Conventional* vs. *ChromaBlur*, *ChromaBlur* vs. *Reverse ChromaBlur*, or *Reverse ChromaBlur* vs. *Conventional*. 120 trials were presented to each subject. Order of conditions and focal distances was random.

7.2 Results

Fig. 13 shows the results averaged across subjects and focal distances. Subjects reported more realistic depth for *ChromaBlur* over *Conventional* on 59.5% of the trials. This percentage is significantly greater than 50% ($p = 0.004$, one-tailed binomial test). Subjects reported more realistic depth for *ChromaBlur* over *Reverse ChromaBlur* on 57.0% of the trials, which is also significantly greater than 50% ($p = 0.027$). And they reported more realistic depth for *Reverse ChromaBlur* over *Conventional* on 53.5% of the trials, which was not significantly greater than 50% ($p = 0.179$).

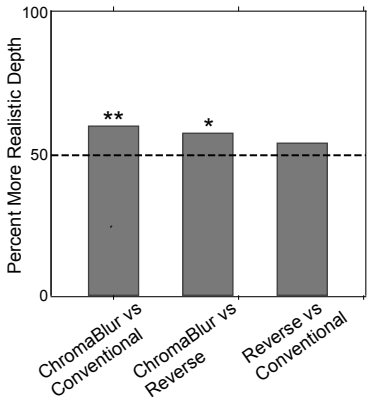

Fig. 13. Results of perceived depth experiment. The data were averaged across focal distances and subjects. The pairwise comparisons are represented on the horizontal axis: from left to right, *ChromaBlur* vs. *Conventional*, *ChromaBlur* vs. *Reverse ChromaBlur*, and *Reverse ChromaBlur* vs. *Conventional*. Percentage of trials in which subjects indicated that the first member of the pair yielded more realistic depth is plotted on the vertical axis. Percentages above the dashed line indicate more responses to the first member of the pair. Asterisks indicate statistically significant deviations from 50% ($*$ for $p<0.05$; $**$ for $p<0.01$).

The results show that ChromaBlur has a small but significant effect on perceived depth compared to conventional rendering. It is not surprising that the effect is small because, despite the strong influence of LCA on accommodation, LCA is thought to have little effect on conscious perception [Kruger et al. 1993]. Interestingly, reversing the chromatic effect did not adversely affect perceived depth relative to conventional rendering.

8 DISCUSSION

8.1 Summary

We described a method to compute displayed images that when viewed by the human eye create close approximations to the depth-dependent optical effects associated with defocus and chromatic aberration. We showed that the method–ChromaBlur–is remarkably effective in driving accommodation and has a beneficial effect on perceived realism.

8.2 Coupling ChromaBlur and Focus-Adjustable Lenses

As we said in Sec. 1, a very challenging problem is how to create a display that enables the appropriate relationship between accommodation and the blur observed at the retina. This relationship is, of course, not enabled in conventional displays because an accommodative response away from the screen distance simply adds blur to the whole retinal image. Instead, we want accommodative responses to cause the retinal images of objects at the new fixation distance to become sharp while the images of objects at other fixation distances become blurred. The effectiveness of ChromaBlur in driving accommodation opens an important opportunity to solve this challenging problem.

One can couple the ChromaBlur rendering technique and focus-adjustable lenses [Johnson et al. 2016; Konrad et al. 2016; Koulieris et al. 2017] to create a display in which accommodation can be driven reliably to the intended distance. Eye tracking could be used to estimate the viewer's fixation distance from moment to moment. When the eyes fixate a distance in the virtual scene, ChromaBlur can create the appropriate signals at the retinas for farther and nearer parts of the scene. And focus-adjustable lenses can be set so that when the eyes accommodate to the specified distance in the virtual scene the retinal image for that distance becomes sharp. In other words, when fixation changes to another distance, ChromaBlur rendering and the focus-adjustable lenses can be changed accordingly. With a binocular system, this rendering method allows the viewer to maintain consistency between vergence and accommodation and with DoF effects.

8.3 Applications

In binocular systems, such as VR and AR, greater ability to control accommodation and create realistic focus cues would provide significant benefits. It would allow one to minimize the vergence-accommodation conflict thereby improving visual comfort and performance. It would also allow a greater sense of realism by creating retinal images that are nearly the same as those created in natural viewing. Presbyopes (individuals who cannot accommodate due to aging) will not benefit from the accommodation-driving capability but should benefit from improved the realism. Focus-adjustable lenses (Sec. 8.2) could also be used to correct a viewer's refractive error (e.g., myopia, hyperopia, or astigmatism), making the display usable without spectacles or contact lenses. Thus, a software algorithm (ChromaBlur) plus a hardware implementation (focus-adjustable lenses) should engender a significantly better viewing experience. Indeed, any system that uses one display screen (by far the most common arrangement currently) can usefully employ ChromaBlur to push accommodation in a desired direction and to enhance depth realism.

8.4 Nulling Viewer's LCA

ChromaBlur requires deconvolution or some similar process (Eqn. 4). This is compute-intensive and has a fundamental limitation at high contrast (Fig. 5). One could avoid this limitation and greatly reduce computation time by undoing the LCA of the viewer's eye. This can be done with a static lens designed to null LCA [Kruger et al. 1993] or by synchronizing the power of a high-speed, adjustable lens [Hu and Hua 2015; Love et al. 2009] with the presentation of the color primaries. By nulling natural LCA optically, we eliminate the need to perform deconvolution. Specifically, K in Eqn. 4 becomes delta functions for R, G, and B so ChromaBlur becomes simply the forward problem, which can be done in real time.

8.5 Rendering Speed

Our emphasis has been on the rendering output's quality and its effectiveness in driving accommomdation and creating realistic depth appearance. As a consequence, the rendering times for our reference-quality ChromaBlur are in minutes. One can, however, greatly shorten these times to support real-time rendering for next-generation displays as outlined in Sec. 8.2. OpenGL shader-based approximations that utilize depth-varying disk blur for the three color channels may be sufficient for driving accommodation. Such approximations will, however, produce incorrect results where the depth gradient is large (e.g., occlusions, reflections). For performance and accuracy, one could use a real-time ray-tracer (Embree, Optix; [Parker et al. 2010; Wald et al. 2014]) rather than the off-line renderer we used (Mitsuba). We used ADMM optimization with 100 iterations to solve the inverse problem (Sec. 4.4.2). One could instead use Tikhonov regularized deconvolution, which has a closed-form, least-squares solution thereby requiring only one iteration or alternatively an efficient non-blind deconvolution [Fortunato and Oliveira 2014]. The inverse step is also amenable to efficient implementation with a high-performance, image-processing system such as Halide [Ragan-Kelley et al. 2013].

ACKNOWLEDGEMENTS

Research was funded by NSF grant BCS-1354029 to Banks. The authors thank George Koulieris and Rahul Narain for comments on a previous version of the manuscript. And they thank Cliff Schor for loaning the autorefractor.

REFERENCES

Karan R Aggarwala, Ekaterina S Kruger, Steven Mathews, and Philip B Kruger. 1995. Spectral bandwidth and ocular accommodation. *Journal of the Optical Society of America A* 12, 3 (1995), 450–455.

Kurt Akeley, Simon J Watt, Ahna Reza Girshick, and Martin S Banks. 2004. A stereo display prototype with multiple focal distances. In *ACM Transactions on Graphics (TOG)*, Vol. 23. ACM, 804–813.

Brian A Barsky. 2004. Vision-realistic rendering: simulation of the scanned foveal image from wavefront data of human subjects. In *Proceedings of the 1st Symposium on Applied Perception in Graphics and Visualization*. ACM, 73–81.

Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* 3, 1 (2011), 1–122.

Fergus W Campbell. 1960. Correlation of accommodation between the two eyes. *Journal of the Optical Society of America* 50, 7 (1960), 738–738.

Han Cheng, Justin K Barnett, Abhiram S Vilupuru, Jason D Marsack, Sanjeev Kasthurirangan, Raymond A Applegate, and Austin Roorda. 2004. A population study on changes in wave aberrations with accommodation. *Journal of Vision* 4, 4 (2004), 272–280.

Michael F Deering. 2005. A photon accurate model of the human eye. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 649–658.

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.

Horacio E Fortunato and Manuel M Oliveira. 2014. Fast high-quality non-blind deconvolution using sparse adaptive priors. *The Visual Computer* 30, 6-8 (2014), 661–671.

Robert T Held, Emily A Cooper, James F OfiBrien, and Martin S Banks. 2010. Using blur to affect perceived distance and size. *ACM Transactions on Graphics* 29, 2 (2010), 1–33.

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), 1–30.

Miriam Hoyt. 2016. Standard (Stauton) Chess Set. (2016). Retrieved March 3, 2017 from http://www.blendswap.com/blends/view/86207

Xinda Hu and Hong Hua. 2015. Design and tolerance of a free-form optical system for an optical see-through multi-focal-plane display. *Applied Optics* 54, 33 (2015), 9990–9999.

Wenzel Jakob. 2010. Mitsuba renderer. (2010). Retrieved January 6, 2017 from http://www.mitsuba-renderer.org

Jay-Artist. 2012. Country-Kitchen Cycles. (2012). Retrieved January 6, 2017 from http://www.blendswap.com/blends/view/42851

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), 11808–11827.

Masanori Kakimoto, Tomoaki Tatsukawa, Yukiteru Mukai, and Tomoyuki Nishita. 2007. Interactive simulation of the human eye depth of field and its correction by spectacle lenses. In *Computer Graphics Forum*, Vol. 26. Wiley Online Library, 627–636.

Craig Kolb, Don Mitchell, and Pat Hanrahan. 1995. A realistic camera model for computer graphics. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*. ACM, 317–324.

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 *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 1211–1220.

George-Alex Koulieris, Bee Bui, Martin S. Banks, and George Drettakis. 2017. Accommodation and Comfort in Head-Mounted Displays. *ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)* 36, 4 (July 2017), 11. http://www-sop.inria.fr/reves/Basilic/2017/KBBD17

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), 30201–1–30201–14.

Tao Liu and Larry N Thibos. 2017. Variation of axial and oblique astigmatism with accommodation across the visual field. *Journal of Vision* 17, 3 (2017), 1–23.

Norberto López-Gil, Frances J Rucker, Lawrence R Stark, Mustanser Badar, Theodore Borgovan, Sean Burke, and Philip B Kruger. 2007. Effect of third-order aberrations on dynamic accommodation. *Vision Research* 47, 6 (2007), 755–765.

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), 15716–15725.

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), 1–20.

Guido Maiello, Manuela Chessa, Fabio Solari, and Peter J Bex. 2014. Simulated disparity and peripheral blur interact during binocular fusion. *Journal of Vision* 14, 8 (2014), 1–14.

David H Marimont and Brian A Wandell. 1994. Matching color images: the effects of axial chromatic aberration. *Journal of the Optical Society of America A* 11, 12 (1994), 3113–3122.

George Mather and David RR Smith. 2002. Blur discrimination and its relation to blur-mediated depth perception. *Perception* 31, 10 (2002), 1211–1219.

Michael Mauderer, Simone Conte, Miguel A Nacenta, and Dhanraj Vishwanath. 2014. Depth perception with gaze-contingent depth of field. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 217–226.

Sina Mostafawy, Omid Kermani, and Holger Lubatschowski. 1997. Virtual eye: retinal image visualization of the human eye. *IEEE Computer Graphics and Applications* 17, 1 (1997), 8–12.

Masashi Nakajima, Takahiro Hiraoka, Yoko Hirohara, Tetsuro Oshika, and Toshifumi Mihashi. 2015. Verification of the lack of correlation between age and longitudinal chromatic aberrations of the human eye from the visible to the infrared. *Biomedical Optics Express* 6, 7 (2015), 2676–2694.

Rafael Navarro, David R Williams, and Pablo Artal. 1993. Modulation transfer of the human eye as a function of retinal eccentricity. *Journal of the Optical Society of America A* 10, 2 (1993), 201–212.

Ren Ng and Pat Hanrahan. 2006. Digital correction of lens aberrations in light field photography. In *Proceedings of the International Optical Design Conference*. Optical Society of America.

Vincent A Nguyen, Ian P Howard, and Robert S Allison. 2005. Detection of the depth order of defocused images. *Vision Research* 45, 8 (2005), 1003–1011.

Matthias Nießner, Roman Sturm, and Günther Greiner. 2012. Real-time simulation and visualization of human vision through eyeglasses on the GPU. In *Proceedings of the 11th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and its Applications in Industry*. ACM, 195–202.

D Alfred Owens. 1979. The Mandelbaum effect: Evidence for an accommodative bias toward intermediate viewing distances. *Journal of the Optical Society of America* 69, 5 (1979), 646–652.

Steven G Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke, David McAllister, Morgan McGuire, Keith Morley, Austin Robison, and others. 2010. Optix: a general purpose ray tracing engine. In *ACM Transactions on Graphics (TOG)*, Vol. 29. ACM, 66.

Henry B Peters. 1961. The relationship between refractive error and visual acuity at three age levels. *Optometry & Vision Science* 38, 4 (1961), 194–198.

Matt Pharr and Greg Humphreys. 2014. *Physically based rendering: From theory to implementation* (3rd ed.). Morgan Kaufmann.

James Polans, Bart Jaeken, Ryan P McNabb, Pablo Artal, and Joseph A Izatt. 2015. Wide-field optical model of the human eye with asymmetrically tilted and decentered lens that reproduces measured ocular aberrations. *Optica* 2, 2 (2015), 124–134.

Jason Porter, Antonio Guirao, Ian G Cox, and David R Williams. 2001. Monochromatic aberrations of the human eye in a large population. *Journal of the Optical Society of America A* 18, 8 (2001), 1793–1803.

Jonathan Ragan-Kelley, Connelly Barnes, Andrew Adams, Sylvain Paris, Frédo Durand, and Saman Amarasinghe. 2013. Halide: a language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines. *ACM SIGPLAN Notices* 48, 6 (2013), 519–530.

Maurice C Rynders, Rafael Navarro, and M Angeles Losada. 1998. Objective measurement of the off-axis longitudinal chromatic aberration in the human eye. *Vision Research* 38, 4 (1998), 513–522.

Thomas O Salmon and Corina van de Pol. 2006. Normal-eye Zernike coefficients and root-mean-square wavefront errors. *Journal of Cataract & Refractive Surgery* 32, 12 (2006), 2064–2074.

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.

Lindsay T Sharpe, Andrew Stockman, Herbert Jägle, and Jeremy Nathans. 1999. Opsin genes, cone photopigments, color vision, and color blindness. *Color vision: From genes to perception* (1999), 3–51.

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), 1–29.

KH Spring and Walter S. Stiles. 1948. Variation of pupil size with change in the angle at which the light stimulus strikes the retina. *The British Journal of Ophthalmology* 32, 6 (1948), 340–346.

Benjamin Steinert, Holger Dammertz, Johannes Hanika, and Hendrik PA Lensch. 2011. General spectral camera lens simulation. In *Computer Graphics Forum*, Vol. 30. Wiley Online Library, 1643–1654.

Mahesh M Subedar and Lina J Karam. 2016. 3D Blur Discrimination. *ACM Transactions on Applied Perception (TAP)* 13, 3 (2016), 1–13.

Ning Tang and ShuangJiu Xiao. 2015. Real-time human vision rendering using blur distribution function. In *Proceedings of the 14th ACM SIGGRAPH International Conference on Virtual Reality Continuum and its Applications in Industry*. ACM, 39–42.

Larry N Thibos, Ming Ye, Xiaoxiao Zhang, and Arthur Bradley. 1992. The chromatic eye: a new reduced-eye model of ocular chromatic aberration in humans. *Applied Optics* 31, 19 (1992), 3594–3600.

Dhanraj Vishwanath and Erik Blaser. 2010. Retinal blur and the perception of egocentric distance. *Journal of Vision* 10, 10 (2010), 1–16.

Ingo Wald, Sven Woop, Carsten Benthin, Gregory S Johnson, and Manfred Ernst. 2014. Embree: a kernel framework for efficient CPU ray tracing. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 143.

Andrew B Watson and Albert J Ahumada. 2011. Blur clarified: A review and synthesis of blur discrimination. *Journal of Vision* 11, 5 (2011), 1–23.

Simon J Watt, Kurt Akeley, Marc O Ernst, and Martin S Banks. 2005. Focus cues affect perceived depth. *Journal of Vision* 5, 10 (2005), 834–862.

James S Wolffsohn, Bernard Gilmartin, Edward AH Mallen, and Sei-ichi Tsujimura. 2001. Continuous recording of accommodation and pupil size using the Shin-Nippon SRW-5000 autorefractor. *Ophthalmic and Physiological Optics* 21, 2 (2001), 108–113.

James S Wolffsohn, Clare O'Donnell, W Neil Charman, and Bernard Gilmartin. 2004. Simultaneous continuous recording of accommodation and pupil size using the modified Shin-Nippon SRW-5000 autorefractor. *Ophthalmic and Physiological Optics* 24, 2 (2004), 142–147.

Jiaze Wu, Changwen Zheng, Xiaohui Hu, and Fanjiang Xu. 2011. Realistic simulation of peripheral vision using an aspherical eye model. *Eurographics* (2011).

Marina Zannoli, Gordon D Love, Rahul Narain, and Martin S Banks. 2016. Blur and the perception of depth at occlusions. *Journal of Vision* 16, 6 (2016), 1–25.

Tingting Zhang, Louise OfiHare, Paul B Hibbard, Harold T Nefs, and Ingrid Heynderickx. 2015. Depth of field affects perceived depth in stereographs. *ACM Transactions on Applied Perception (TAP)* 11, 4 (2015), 18.

S1 SUPPLEMENTAL MATERIALS

S1.1 Apparatus

S1.1.1 Color primaries. The stimuli were projected onto a screen by a DLP projector (Texas Instruments LightCrafter 4710 Evaluation Module). The LED light sources were OSRAM; Red: LE A P1W, 459nm peak; Green: LE CG P1A, 520nm peak; Blue: LE B P1W, 617nm peak. We measured the emission spectra of the three primaries through the autorefractor's hot mirror using a Photo Research PR-650 SpectraScan Colorimeter. Spectral curves are shown in the top plot of Figure S1. Table S1 shows the peak wavelength, bandwidth (full width at half height), luminance, and radiance (at highest intensity for each primary).

S1.1.2 Perceived depth experiment. For the perceptual experiment (Sec. 7), we narrowed the spectra, particularly that of the G primary. We did this by placing a Chroma 69002m triple-bandpass filter between the projector and screen. We again measured the spectra through the hot mirror from the subject's viewpoint. Spectral curves are shown in the bottom plot of Figure S1 and other values in Table S2.

S1.1.3 Focus-adjustable lens. The focus-adjustable lens was an Optotune EL-10-30-VIS-LD. It has a range of +6.2 to +20.2D (measured with a Topcon CL-200 lensmeter) at the room temperature of ~22° C. We placed a -10D achromatic lens (Edmund Optics #62-494) in the optical path to shift the range of the lens system to -3.8 to +10.2D.

S1.1.4 Autorefractor. We measured accommodative state (specifically, spherical power) using a Grand Seiko WV-500 autorefractor. The WV-500 samples at ~1Hz. We were able to greatly increase this rate by modifying a technique described by [MacKenzie et al. 2010; Wolffsohn et al. 2001, 2004].

The WV-500 projects bars arranged in a square pattern onto the retina and uses the separations of the bars in the reflected image to measure accommodative state [Wolffsohn et al. 2004]. The near-infrared LED that projects the bars is disabled during normal use and only activated when the machine makes a measurement. But the WV-500 has a "Sales Mode" in which one can force the LED to always be active [Wolffsohn et al. 2001]. By turning on the LED continuously, we could record the measurement bars from the autorefractor's video feed and thereby measure accommodative state at 30Hz.

Two computers were used, one for stimulus presentation and one for recording from the autorefractor. The stimulus computer was synchronized with the recording computer via an RS-232 serial connection at 57.6kbaud. The recording computer digitized the WV-500's NTSC video signal using an EasyCap DC60 composite video USB capture dongle with a usbtv007 chipset. The resulting video files were analyzed using a custom set of image-processing routines in Python to localize the measurement bars and to estimate the diameter of the pupil. Figure S2 shows example video frames and cross-sections of the measurement bars and the pupil estimator.

Table S1. Spectra in Accommodation Experiment

Primary Color	Peak (nm)	Bandwidth (nm)	Luminance (cd / m^2)	Radiance ($W \cdot sr^{-1}$ m^{-2} nm^{-1})
Red	630	19.4	49.5	0.257
Green	530	90.1	313.6	0.596
Blue	456	24.4	11.9	0.270

Table S2. Spectra in Perceptual Experiment

Primary Color	Peak (nm)	Bandwidth (nm)	Luminance (cd / m^2)	Radiance ($W \cdot sr^{-1}$ m^{-2} nm^{-1})
Red	631	19.1	46.1	0.240
Green	529	33.9	128.1	0.253
Blue	455	21.3	8.2	0.215

S1.2 Monochromatic vs Broadband Primaries

In ChromaBlur, we represent each primary by one wavelength. This approximation should in practice introduce error because most displays have broadband primaries. To investigate the significance of such errors, we calculated RMS errors for two situations. First, we compared the retinal images generated by a real broadband edge (equal intensity at all wavelengths from 400-700nm) versus an edge rendered with ChromaBlur when the displayed stimuli were composed of three monochromatic primaries (459, 520, and 617nm). We computed the RMS error of the image created by ChromaBlur rendering for a variety of defocus values. Second, we compared the retinal images generated by the same real edge versus an edge rendered with ChromaBlur when the displayed images were composed of three broadband primaries (each uniform with a bandwidth of 100nm; centered at 459, 520, or 617nm). The results are shown in Figure S3, dashed curves represent the errors when the primaries are monochromatic and solid curves those when the primaries are broadband. The monochromatic approximation does not in general worsen the error; indeed, in some cases assuming monochromatic

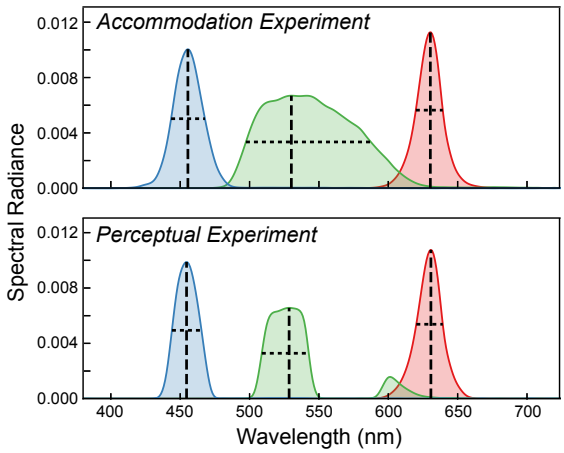

Fig. S1. Spectral radiance as a function of wavelength measured from the subject's viewpoint in the accommodation (top) and perception (bottom) experiments. Dashed vertical lines are peak radiances for each primary and dotted horizontal lines are bandwidths (50% of peak radiance).

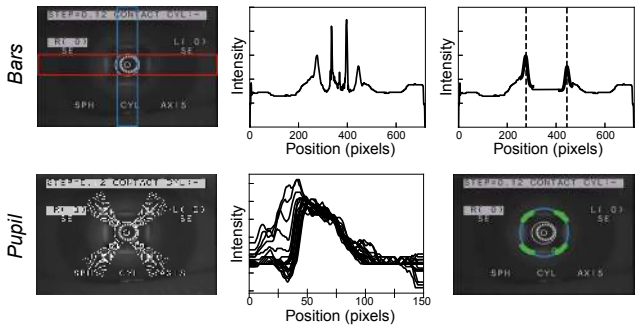

Fig. S2. Measurement bar (top row) and pupil (bottom row) localization. Top-left: Measurement bars are localized by first sampling the image intensities in horizontal and vertical regions for each video frame (red and blue boxes). Top-center: The median intensity profile for the horizontal region has peaks where the measurement bars are located and noise due to the autorefractor's information overlay, which cannot be disable. Top-right: The central reticle peaks are filtered and piecewise asymmetric Gaussian contours are fit to the curves (dark segments), localizing the peaks with sub-pixel accuracy (dashed vertical lines). Bottom-left: Pupil location and diameter are measured by first sampling linear segments of the image intensity, radiating from the center of the image (white lines). Bottom-center: Intensity gradients are computed from individual intensity profiles with maxima localizing the pupil/iris border. Bottom-right: Estimates of the border (green points) are used to fit a circle using RANSAC (blue circle), robustly localizing the pupil.

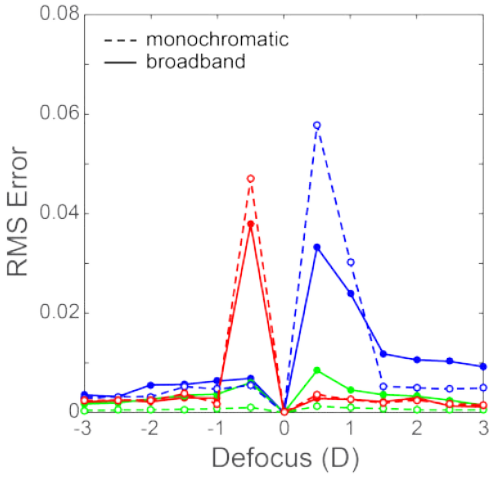

Fig. S3. Errors in retinal image due to monochromatic approximation. RMS error between the correct retinal-image intensities and those created by ChromaBlur are plotted as a function of defocus in diopters. Red, green, and blue curves represent the errors for R, G, and B primaries, respectively. Solid curves represent errors when the primaries are broadband (uniform in intensity with 100nm bandwidth). Dashed curves represent the error when the primaries are monochromatic. Pupil diameter = 6mm.

primaries but using broadband primaries actually improves the accuracy of the retinal image. We note that the emission spectra of most displays are narrower and smoother than the ones assumed in this analysis, so we conclude that ChromaBlur produces reasonably accurate retinal images for most practical displays.

S1.3 Pupil Diameter

How important is it to know the viewer's pupil diameter when implementing our technique? It might not be important for two reasons. First, steady-state pupil diameter under room lighting does not vary much. It is typically 3-5mm [Spring and Stiles 1948], so one can assume 4mm and not be far off. Second, when LCA is doubled or halved in magnitude, accommodative responses are largely unaffected [Kruger et al. 1993]. We can understand this from Eqn 1 which shows that the magnitude of blur is proportional to both pupil diameter and the difference in diopters between the object distance and where the eye is focused. The focal difference between two wavelengths (Eqn. 2) can be expressed in diopters from which one can see that the same chromatic effect at the retina occurs with various combinations of change in focal distance and pupil diameter: i.e., halving or doubling the magnitude of LCA is equivalent to halving or doubling pupil diameter. For this reason, we think that ChromaBlur will be effective even if the designer does not incorporate the viewer's current pupil diameter into the rendering pipeline.

S1.4 3d ChromaBlur Pseudocode

There are two steps to generating a ChromaBlur display image. First, a target retinal image is produced by ray tracing a scene with a perspective camera and thin lens model. For a scene that only contains a planar object, this method of rendering defocus is equivalent to convolving a sharply rendered image with cylinder kernels with diameters determined by Eq. 2. Three grayscale images are produced, one for each color channel. Focal distances for each color channel are offset from the target distance via Eq. 2. The images are then combined to produce a 3-channel target retinal image $I(x, y)$. Algorithm 1 describes this method.

Algorithm 1 Forward Step

```
function RENDER_SCENE(dist)
    for each color_channel do
        focus_dist ← LCA_SHIFT(dist, λ_color_channel)
        image ← Render with thins lens model at focus_dist
        Save image to I_{color_channel}(x, y)
    end for
    return I_{R,G,B}(x, y)
end function
function LCA_SHIFT(dist, λ)
    chromatic_eye_error ← D(λ)                ▷ See Sec. 3.2 Eq. 2
    focus_dist ← dist - chromatic_eye_error
    if focus_dist < 0 then
        Negative focus distance, raise error
    else
        return focus_dist
    end if
end function
```

The second step is to produce the display image $D(x, y)$ from the target retinal image I. We deconvolve each color channel with the PSF of an in-focus screen, defined as the circle of confusion/defocus blur spot with radius determined by Eq. 2. We assume the eye is

focused on green [Thibos et al. 1992], so its deconvolution kernel is a Dirac delta function; however, the function in Fig.2 (Eq. 2) can be shifted if the eye focuses on a saturated object. For example, if the object is mostly red, we would shift the function such that red rather than green would be in focus.

For a plane simulated at the same depth as the display screen, with no defocus, the convolution and deconvolution kernels will be identical for each color channel, so $D(x, y)$ would be rendered sharp and would not require deconvolution. We deconvolve the target retinal image with 100 ADMM update iterations. Algorithm 2 describes this method. Note that with conventional rendering (our *Defocus Only* condition), $D(x, y)$ is rendered sharp for all wavelengths when we create an in-focus object. Likewise, $D(x, y)$ is rendered equally blurred at all wavelengths when we create an out-of-focus object.

Algorithm 2 Inverse Step

Optimization problem:

$$\min_{D(x,y)} ||D(x,y) * *K(x,y) - I(x,y)||_2^2 + \psi||\nabla D(x,y)||_1$$
$$\text{such that } 0 \leq D(x,y) \leq 1$$

Rewritten to simplify future notation:

$$\min_x \tfrac{1}{2}||Ax - b||_2^2 + \lambda||\nabla x||_1$$
$$\text{such that } 0 \leq x \leq 1$$

Equivalent:

$$\min_x \tfrac{1}{2}||Ax - b||_2^2 + I_{[0,1]}(x) + \lambda||z||_1$$
$$\text{such that } Dx - z = \nabla x - z = 0$$
$$\text{where } I_{[0,1]}(x) = \begin{cases} 0, & 0 \leq x \leq 1 \\ \infty, & \text{otherwise} \end{cases}$$

Augmented Lagrangian:

$$\mathcal{L}_\rho(x, z, y) = \tfrac{1}{2}||Ax - b||_2^2 + I_{[0,1]}(x) + \lambda||z||_1 + y^T(Dx - z) + \tfrac{\rho}{2}||Dx - z||_2^2$$

ADMM Updates:

for each color_channel **do**
  for each update **do**
    $x \leftarrow \underset{x}{\text{argmin}}\ \tfrac{1}{2}||Ax - b||_2^2 + I_{[0,1]\{color\_channel\}} + \tfrac{\rho}{2}||Dx + u - z||_2^2$
    $z \leftarrow \underset{z}{\text{argmin}}\ \lambda||z||_1 + \tfrac{\rho}{2}||Dx + u - z||_2^2$
    $u \leftarrow u + Dx - z$
  end for
end for
where $u = \frac{1}{\rho}y$

S1.5 Does ChromaBlur Drive Accommodation in Color-deficient Viewers?

The use of LCA requires photoreceptors with different spectral sensitivities. To further examine the use of LCA in driving accommodation, we tested two color-deficient subjects (males, 19 and 20 years). Both were protanopes (missing the L cones, but not the S and M cones) as determined by the HRR Pseudochromatic Color Test. They were tested with the same procedure as in the main experiment.

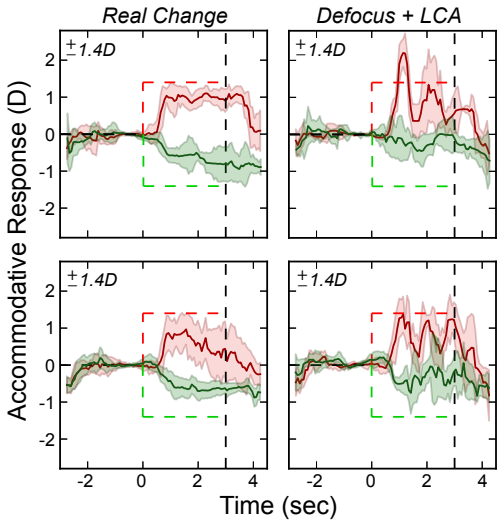

Fig. S4. Accommodative responses in two color-deficient subjects. Dashed lines represent the stimulus; red for positive and green for negative. Shaded regions are median absolute deviations; red for positive changes and green for negative. The thick curves are the medians. The left and right columns are for *Real Change* and *Defocus plus LCA*. Each row shows the data from one protanopic subject.

Figure S4 shows the responses of the color-deficient subjects in the *Real Change* and *Defocus plus LCA* conditions. Responses in the *Real Change* condition were reasonably normal, but responses in the *Defocus plus LCA* condition were not: Although they responded in the appropriate direction, the dichromats exhibited oscillations that we did not observe in the color-normal subjects. These protanopic subjects are not color-blind because they have two cone types and that enables them in principle to use the LCA signal, just less reliably than color-normal subjects with three cone types. These results are further evidence that our rendering technique provides a useful signal for accommodation, particularly for the great majority of people who are color-normal.

S1.6 Accommodative Responses from All Subjects

In the main experiment, we measured accommodative responses in five young adults. Figure 9 shows the results from one. Figure S5 shows the results from the others. Subject JB (upper left) had somewhat variable responses but clearly accommodated in the specified direction in the *Real Change* and *Defocus plus LCA* conditions; the response was generally larger in the *Defocus plus LCA* condition. Subject MR (upper right) accommodated by smaller amounts than the other subjects, particularly in the *Defocus plus LCA* condition. Subject IT (lower left) responded very similarly in the *Real Change* and *Defocus plus LCA* conditions; she responded consistently to positive defocus (i.e., focusing nearer) and did not respond consistently to negative defocus (focusing farther). This inability to accommodate in one direction is somewhat common. Subject AS (lower right)

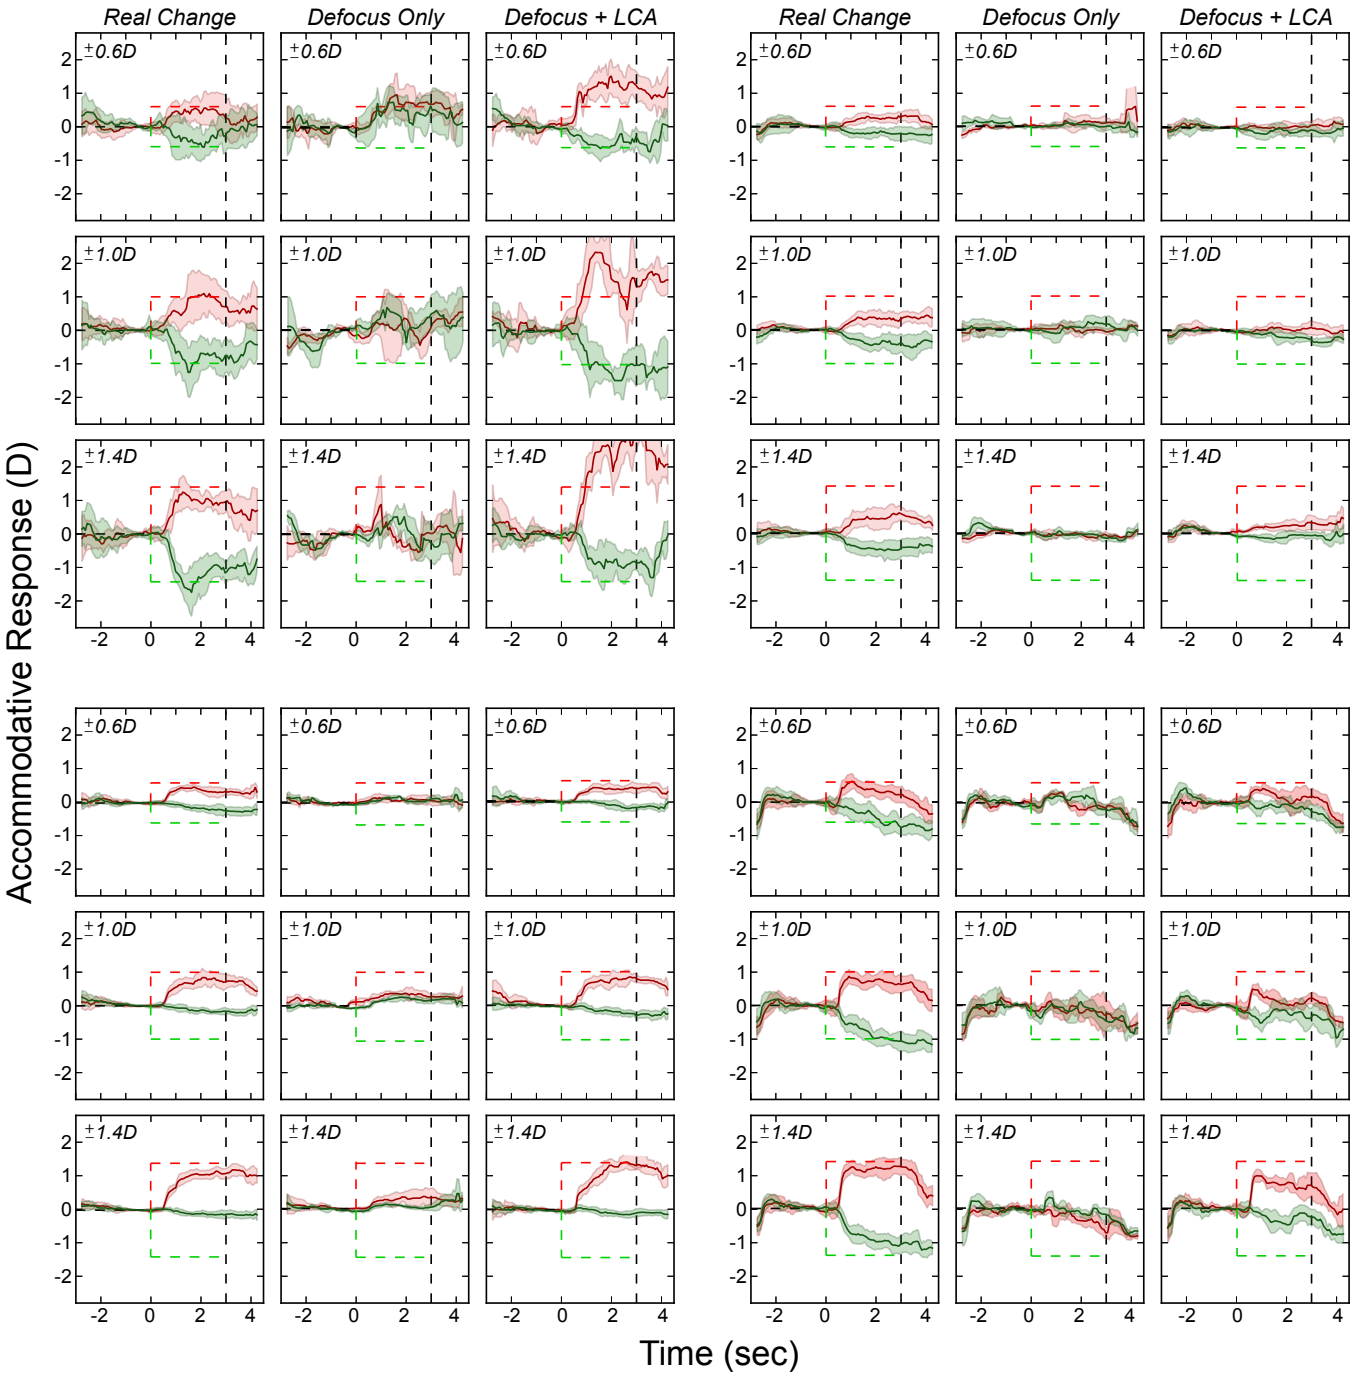

Fig. S5. Accommodation results from the other four subjects. The format is the same as Figure 9. Upper left is JB. Upper right is MR; this subject yielded rather small responses even when the focal distance was actually changed. Lower left is IT. Lower right is AS.

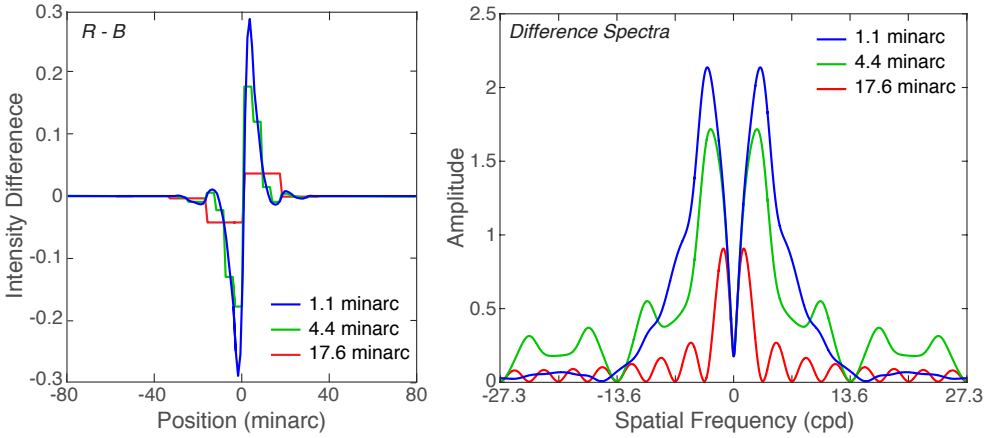

Fig. S6. Effect of sampling on the ChromaBlur signal. Left: The difference between the displayed images for the R and B primaries. Blue, green, and red curves represent those differences for 1.1, 4.4, and 17.6minarc pixels, respectively. Right: Amplitude spectra for the difference in displayed image for R and B. Again blue, green, and red represent the results for 1.1, 4.4, and 17.6minarc, respectively.

responded in the specified direction in the *Real Change* and *Defocus plus LCA* conditions; the response was somewhat smaller, however, in the *Defocus plus LCA* condition.

S1.7 Resolution Experiment

As we said in Sec 6.2.2, we wanted to know what the resolution of a display needs to be for ChromaBlur to drive accommodation effectively. To investigate this, we conducted the accommodation experiment with different simulated resolutions. We created different effective resolutions by down-sampling the stimuli with bi-linear decimation and then up-sampling to produce larger effective pixels with nearest-neighbor interpolation. Quantitatively, this process is:

$$o(x,y) = \left[\left(i(x,y) **r(\frac{x}{x_0},\frac{y}{x_0})\right)\cdot s(\frac{x}{x_0},\frac{y}{x_0})\right] **r(\frac{x}{x_0},\frac{y}{x_0}) \quad \text{(S1)}$$

where $o(x,y)$ is the output image, $i(x,y)$ is the input image, $r(\frac{x}{x_0},\frac{y}{x_0})$ is the 2d rectangle function with width and height x_0, $s(\frac{x}{x_0},\frac{y}{x_0})$ is the 2d comb sampling function with samples separated by x_0, and $**$ is two-dimensional convolution. $x_0 \times x_0$ pixels are created. x_0 was 2.2, 4.4, 8.8, 17.6, or 35.6 minutes of arc at the subject's eye.

In the frequency domain, Equation S1 becomes:

$$O(\omega,\nu) = [(I(\omega,\nu)\cdot R(x_0\omega,x_0\nu)) **S(x_0\omega,x_0\nu)]\cdot R(x_0\omega,x_0\nu) \quad \text{(S2)}$$

where $O(\omega,\nu)$, $R(\omega,\nu)$, and $S(\omega,\nu)$ are the Fourier transforms of $o(x,y)$, $r(x,y)$, and $s(x,y)$, respectively. $R(\omega,\nu)$ is the 2d sinc function and $S(\omega,\nu)$ the 2d comb function with samples separated by $\frac{1}{x_0}$.

To investigate how the ChromaBlur signal is affected by display resolution, we considered the rendering and imaging of a high-contrast step edge as in Figure 4. The ChromaBlur signal can be approximated by the difference between $o_R(x,y)$ (the displayed image for the red primary) and $o_B(x,y)$ (displayed image for blue). Figure S6 shows the amplitude spectra $O_R(\omega,\nu) - O_B(\omega,\nu)$ for pixel sizes of 1.1, 4.4, and 17.6minarc. The spectrum for 1.1minarc shows that the ChromaBlur signal is in the frequency range of 3–10 cycles/deg. As pixel size increases, the amplitude of the signal decreases and becomes distorted by aliases due to coarse sampling. The signal is still observable at 4.4minarc, but not at 17.6minarc where it is reduced and distorted. This analysis suggests, therefore, that the ability to drive accommodation with ChromaBlur should dissipate between 4.4 and 17.6minarc.

The results in Figure 11 confirm this expectation. Figures S7 and S8 show the results for the other two subjects. As with the subject in Figure 11, accommodation was driven effectively at pixel sizes of 1.1–8.8minarc and not driven at larger sizes.

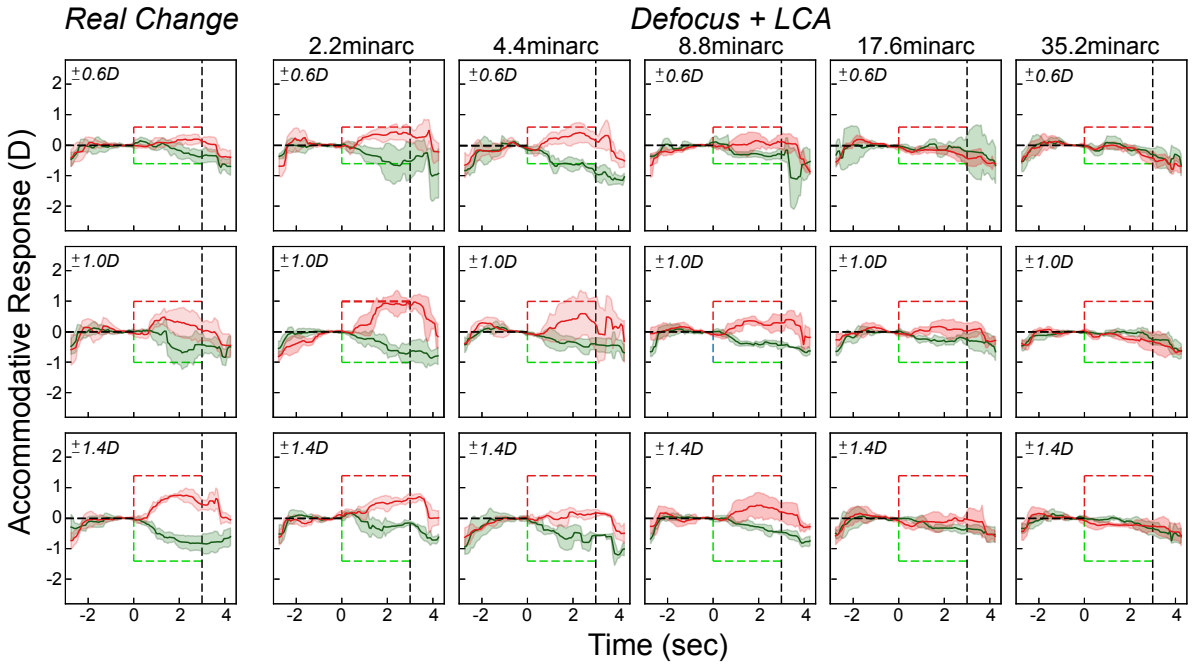

Fig. S7. Accommodative responses as a function of image resolution for subject AC. The upper three rows are from one subject and the lower three rows from the other. The first column shows the data for the *Real Change* condition, collapsed across resolutions. The remaining columns show the data for the *Defocus plus LCA* condition for resolutions, left to right, of 2.2, 4.4, 8.8, 17.6, and 35.2 minutes of arc. Dashed lines in each panel represent the stimulus; red for positive and green for negative. Shaded regions are median absolute deviations; red for positive defocus changes and green for negative. The thick curves are the medians. The large deviations after stimulus offset at 3sec are due to blinks.

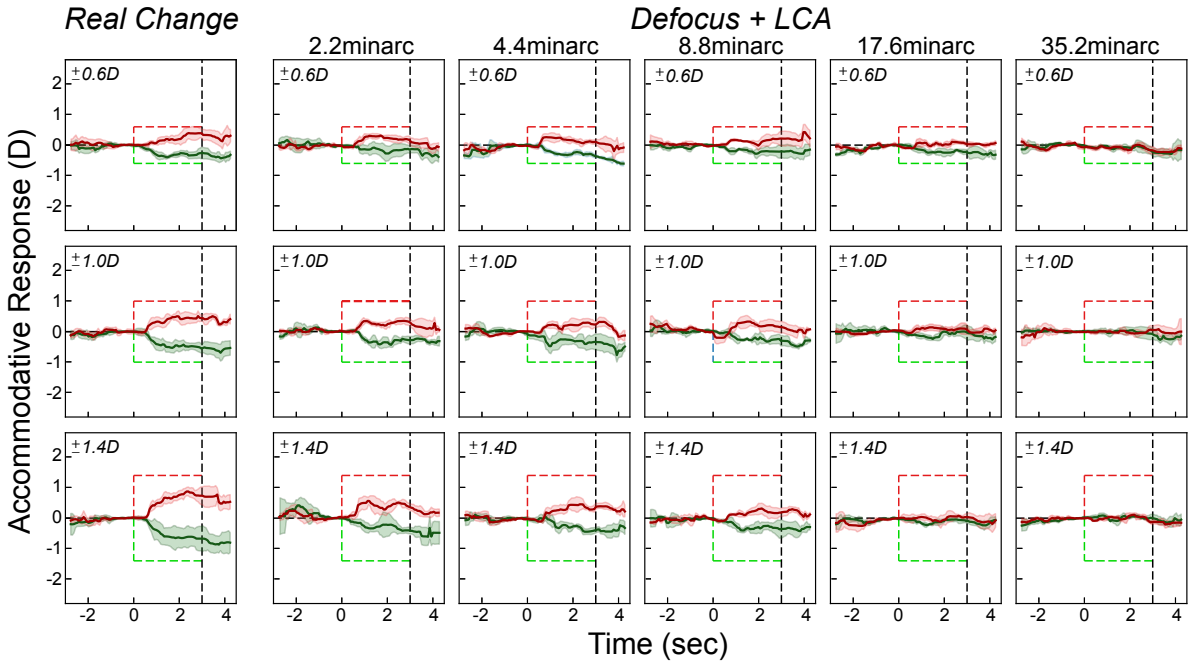

Fig. S8. Accommodative responses as a function of image resolution for subject KR. Same format as S7