

Supplementary Materials for

Autofocals: Evaluating gaze-contingent eyeglasses for presbyopes

Nitish Padmanaban, Robert Konrad, Gordon Wetzstein*

*Corresponding author. Email: gordon.wetzstein@stanford.edu

Published 28 June 2019, *Sci. Adv.* **5**, eaav6187 (2019)

DOI: [10.1126/sciadv.aav6187](https://doi.org/10.1126/sciadv.aav6187)

The PDF file includes:

Supplementary Materials and Methods

Supplementary Text

Legend for Preference Questionnaire

Fig. S1. A partially exploded view of the headset computer-aided design model.

Fig. S2. An image of the previous prototype, which had a glasses form factor.

Fig. S3. Optical characteristics of the Optotune EL-30-45 focus-tunable lenses captured using a camera.

Fig. S4. A wavefront map of the coma correctors, designed in Zemax.

Fig. S5. Focus accuracy at different positions before and after addition of the coma corrector.

Fig. S6. Measured optical lens power as a function of target lens power.

Fig. S7. Evaluations of the accuracy of the two main external sensors.

Fig. S8. A visual representation of sources of error in the estimated vergence.

Fig. S9. Error in the estimated vergence distance from various sources.

Fig. S10. An example recording of the sensor fusion algorithm.

Algorithm S1. Sensor fusion: Vergence + error.

Algorithm S2. Depth denoiser.

Legends for data S1 and S2

References (45–48)

Other Supplementary Material for this manuscript includes the following:

(available at advances.sciencemag.org/cgi/content/full/5/6/eaav6187/DC1)

Preference Questionnaire (.pdf format)

Data S1 (.zip format). A zip file containing comma-separated values (CSV) files with the raw data for participants for visual acuity, contrast sensitivity, and task performance.

Data S2 (.csv format). A CSV file containing the raw data for participants for the natural use questionnaire.

Supplementary Materials and Methods

Hardware Components

The ski goggles form factor of our prototype (fig. S1) makes the eye tracking more reliable, as an earlier glasses form factor (fig. S2) tended to lose calibration due to slippage.

With the optical axis oriented in the horizontal position, gravity causes the lens' liquid to pool towards the bottom, leading to noticeable coma when imaging through the full 30 mm aperture (fig. S3, right). At any given time, humans only look through a subset of the lens due to their smaller pupil/aperture. Under these conditions, the non-uniform distribution of the liquid is instead perceived as a spatially varying focal power across the lens. While this perceived variation in focal power could theoretically be corrected using the focus-tunable lens itself, we measured a focal power variation of ± 1.5 D over a ± 10 mm region from the lens center. This represents a prohibitively large fraction of the total tunable range of the optics. Therefore, to counteract the coma, we used a wavefront characterization provided by the manufacturer and diamond turned a freeform coma corrector (wavefront pictured in fig. S4) from polymethyl methacrylate (PMMA). We confirmed that this reduces focal power variance from ± 1.5 D over a ± 10 mm region from the lens center to only ± 0.3 D (worst case, at any focus distance), when simulated by placing a 4 mm aperture (see fig. S5) in front of the camera. To minimize any residual error, we also fit the autofocus headset on users such that their pupils are as close to the optical axis as possible.

To ensure that the lenses provide enough fidelity for vision correction purposes, we evaluated their optical characteristics. We measured the modulation transfer function (MTF) of the optical system at different distances using the slanted-edge algorithm based on the ISO 12233 standard (fig. S3, left). During measurements, a 4 mm aperture was set 2 cm away from the lens, mimicking the typical pupil size for photopic vision amongst presbyopes (45) and distance of the lenses from the eyes. A 1 D offset lens was also inserted during MTF measurement to match the offset used during wear. The dashed lines indicate the MTF50 (i.e., the resolution supported at 50% contrast) for each of the focusing distances, with all focusing distances having an MTF50 greater than 30 cpd, which corresponds to typical human visual acuity. We observe no consistent trend in MTF50 as a function of distance, so differences are likely noise in the measurement.

We also consider the possibility that the focus-tunable lenses may not be set to the programmatically specified lens power due to calibration errors. However, as seen in fig. S6, the lenses stay within ± 0.1 D of the correct value at the distances used for the quantitative user studies (up to 2.5 D), verifying that the temperature-dependent factory calibration of the lenses is sufficient for our operation. We do not expect the optical system to be the limiting factor for the visual acuity study in system evaluation.

Hardware Limitations The RealSense R200 has a rated depth range of 0.5–3.5 m (0.285–2 D). However, our application requires that we use it for both nearer and farther distances. We characterized the error of the R200 using a flat white wall at known distances of 0.167, 1.25, and 2.5 D. The R200 does not provide dense measurements and requires further processing to obtain a complete depth map. We fill in the unknown regions with a naïve Navier-Stokes inpainting

algorithm (46). The measured errors are given in fig. S7 (left) for both the raw and inpainted depth values. When considering both error and variance together, as expected, the R200 does best at the 1.25 D distance, which is within its supported range. The other distances suffer from higher variance in the measurements. The naïve inpainting also breaks down at these extremes, especially the closest distance. However, on average, the measurement stays within 0.3 D of the correct depth, implying that filtering the inpainted values is a good strategy.

The Pupil Labs eye trackers claim an accuracy of 0.6° per eye, which for a wearer with an interpupillary distance of 64 mm, limits the accuracy of accommodation estimates to about 0.33 D in ideal conditions, larger than an eyeglass prescription step size. In practice, our measured accuracy was even lower. To estimate accuracy, we calibrated a few users wearing autofocals, then had them look at the center of \times -shaped targets that subtended 3° of visual field. The targets were placed at the center of their visual field, and 20° above, below, to the left, and to the right of the center. Error measurements averaged over target positions, at distances of 0.167, 1.25, and 2.5 D, are given in fig. S7 (right). The average error is in the range of $3\text{--}5^\circ$. The standard deviation is roughly 3° , which may correspond simply to the size of the target itself. However, that still leaves a substantial source of error in the gaze estimate. Since we derive our vergence estimate from the gaze directions, an error of even 2° inwards or outwards per eye could lead to a little over 1 D of error in depth (see fig. S8 and fig. S9 center).

The depth map seems more accurate, but ultimately suffers from issues near depth discontinuities. A vergence estimate on the other hand would be ideal since vergence and accommodation are neurally linked processes, but we see that it potentially has unacceptable levels of error. Therefore, we propose a sensor fusion algorithm to leverage the strengths of both as best as possible.

Fixation Depth Estimation

The depth of the fixated object is dynamically estimated via sensor fusion of four “raw” inputs: the gaze point of the left eye, the gaze point of the right eye, a scene-facing depth camera, and the wearer’s interpupillary distance (IPD). We then process these into a vergence estimate and a single depth estimate, which are the inputs to the sensor fusion algorithm.

The gaze directions are given as (x, y) coordinates normalized to a range of 0 to 1 relative to the FOV of the R200 – this is done via Pupil Labs’ calibration software. Since the geometry of the calibration scene relative to the wearer is set ahead of time (screen 45 cm away, with calibration targets of known position), the wearer’s IPD and the separation distance of two gaze points is sufficient to determine the wearer’s vergence distance. Though not necessary, this estimate can then be refined by calibrating the vergence at two known distances and scaling the geometric estimate accordingly. Furthermore, we discard any values outside the -1 to 4 D range. These limits are based on the depth extremes in our study setup.

The depth map is represented as an image which is 320×240 pixels, but it has “holes” where the depth could not be estimated. To remove these holes, we first inpaint the depth map using a Navier-Stokes algorithm (46). Then, we use the left eye’s gaze coordinate to index into a 5×5 neighborhood of depth map, from which we take the median depth. We use the left eye because

it is approximately aligned with the camera of the R200; if the left eye is not available, we fall back to the right eye.

Given these processed values, the high-level operation of the sensor fusion algorithm can be described as vergence plus an error term (estimated as the difference between the vergence and depth). We choose this approach since vergence updates much faster (120 Hz) than the depth (30 Hz) and is not affected by depth discontinuities at edges. However, it is sensitive to inaccuracies in calibration, which manifests itself as a nearly constant offset in the range of operation (fig. S10). Since the depth map itself is accurate away from edges, this offset is corrected for with the gaze-indexed depth estimate.

The details of this algorithm are in Algorithm S1. Since the vergence estimate, v , is noisy (fig. S10, red), it is first smoothed with an exponential average to obtain v_{filt} (line 5). Then, to correct for the offset in vergence, we check for a new filtered depth estimate, d_{filt} (fig. S10, blue, see next paragraph). These two values are used to update the error estimate, v_{err} , with an exponential average, but such that fluctuations from depth discontinuities are not considered (line 8). We also allow error updates at either extreme of the depth range to control for initial conditions with large vergence error. Finally, the lens focal power (fig. S10, green) is set to the new error-compensated vergence estimate (line 13), with a minimum step size of 0.25 D to minimize jitter caused by limits of the eye tracking accuracy (see Hardware Limits section).

The value for d_{filt} , used in Algorithm S1, is calculated as seen in Algorithm S2. The raw depth estimate, d , exhibits a tendency to drop by nearly a diopter for some frames (fig. S10, gray). We consider a drop of more than 0.5 D to potentially be an artifact. However, some drops are clearly the users looking to a new distance. Therefore, we reject depth estimates until an exponential filter starting from the last valid depth (line 11) “catches up.”

User Study Design

The first study consisted of 27 participants. Of these, the first four were excluded due to a software error that set the wrong lens power offsets. Two more participants were excluded due to strong astigmatism that we were unable to accurately measure, and one participant was excluded due to strabismus. Another two participants were unable to run the task performance portion of the study because the eye trackers failed to consistently differentiate their mascara from their pupils. Finally, due to scheduling constraints, one participant only ran the task performance portion. This leaves the study population at 19 people for the acuity and contrast portions, and 18 for the task performance. Due to possible improvements in proficiency using presbyopic correction as a function of length of use, we only compare autofocals against the correction the users wear on a regular basis (this inherently puts autofocals at a slight disadvantage), specifically progressive addition lenses (progressives) and monovision. The summaries for the study subpopulations are thus: progressives ($n=14$, ages 55–70, mean age 61.3, 5 female) and monovision ($n=5$, ages 52–67, mean age 59.6, 4 female) wearers for acuity and contrast, and progressives ($n=14$, ages 55–70, mean age 60.9, 4 female) and monovision ($n=4$, ages 52–67, mean age 59.3, 3 female) wearers for task performance.

The second user study was conducted as part of a public demonstration at the SIGGRAPH 2018 conference, with an optional preference questionnaire. Users that were younger than 50 were excluded, leaving 37 participants (ages 50–66, mean age 56.5, 4 female). The phases of this study were threefold: first, a prescription measurement; second, a few minutes spent using autofocus in the eye-tracked mode; and third, using autofocus in the depth-tracked mode. Users were not constrained to a chin rest and were instead free to look around their environment as they saw fit.

The following sections for each study detail how the focal powers of the focus-tunable lenses are constrained to the known display distances for some studies. The rationale for the decision to use this alternative mode of operation is grounded in the need to isolate possible points of failure within the autofocus prototype. The measurements of acuity and contrast sensitivity serve to validate that the focus-tunable lenses maintain acceptable optical quality over a range of distances without a complicating variable of eye tracking. The letter matching task serves to verify that the overall latency of the 120 Hz eye tracking and 100 ms lens settling time does not impede refocusing speed. Finally, the natural use task makes use of the fully unconstrained system to determine how the remaining system-level or algorithm-dependent factors, such as comfort or jitter, affect user experience. Our choices in constraining various parts of the fixation estimation stem from a desire to evaluate the underlying technology, which is in its infancy, as opposed to limiting ourselves to our current prototype.

Acuity As is common practice, acuity corresponding to 20/20 vision is defined as 0.0 logMAR. The acuity tests were conducted at distances of 0.167, 1.25, and 2.5 D (6 m, 80 cm, and 40 cm), with three repeats at each distance, and with the order randomized across all trials.

Visual acuity charts were displayed using monitors with approximately 190 cd/m² of output when white, and with high enough resolution to support at least 20/10 vision at the intended distance (at 2.5 D and 1.25 D, a 5.98" display with at a 1440p resolution, and at 0.167 D a 24" display at 1080p). The monitors were all calibrated with measurements of light output at all gray levels with the SpectraScan PR670, a spectroradiometer. Letters were displayed using the Sloan font (37). Depth-dependent magnification of the letters was determined to be about a 5% change in size, corresponding to about 1 letter (0.02 logMAR) when tested experimentally. Furthermore, since the magnification depends only on lens power and distance of the lens to the eye, progressives and autofocus should both exhibit similar magnification, making it a fair comparison. This may affect results for monovision contacts, but our measured acuity improvement is far larger than 0.02 logMAR. It should be noted that for the acuity tests, the focus-tunable lenses are set to the known display distances.

Contrast Sensitivity Measurements were repeated twice each with their correction and autofocus. There were two charts, identical with the exception of the letters on each. The first chart was alternated between users, and charts were also alternated after each trial. As in the acuity section, contrast measurements were also done with the focus-tunable lenses fixed to the correct distance.

Task Performance For the task performance study, we used the full extended Sloan alphabet (with the exception of the letters 'G' and 'W' which had unusual shapes compared to typical

fonts). Users were encouraged to focus on accuracy while maintaining speed with the following prompt, “We are recording both speed and accuracy, so try to get as many correct as you can while still going quickly.” Users were given a training period of 20 pairs ahead of the timed portion to get accustomed to the controls. When wearing autofocals for this task, the user was reminded that they could simply look back and forth with their eyes, though they were still allowed to rotate their head. Similarly to acuity, the focus-tunable lenses are limited to the known display distances, for the same reasons. Unlike acuity, the correct focal power during task performance still depends on the wearer; we set the lenses to focus on the display closer (in diopters) to the wearer’s estimated fixation distance.

Natural Use Preference Finally, for the qualitative natural use questionnaire, eye-tracked mode was always the first mode they tried. When starting each mode, we explained to the user how the focus control works and how they could make it focus to closer or farther objects. Depth-tracked mode estimates fixation depth by using the distance of the object at the center pixels of the RealSense R200’s FOV. Since there is only one tethered pair of autofocals, the viewing period was necessarily limited to a few minutes instead of a few weeks of everyday use. After the users evaluated the scene, we had them rank (ties not allowed) the conditions on three metrics: first, physical comfort between their correction and the autofocal headset; second, which “made it easiest to switch between focusing clearly at multiple distances,” between their correction, eye-tracked autofocals, and depth-tracked autofocals; and third, which was the most convenient overall between the three. The exact text for each ranking was:

1. Rank (no ties) the corrections in order of which were the most physically comfortable (consider, would you wear them all day?).
2. Rank (no ties) the corrections in order of which made it easiest to switch between focusing clearly at multiple distances.
3. Rank (no ties) the type of correction which is the most convenient overall, remembering, for example, that progressives and monovision have initial adjustment periods, and that autofocal glasses would need regular recharging.

The full questionnaire form can be found in the other supplementary materials.

Supplementary Text

Related Work

Focus-Tunable Lenses The class of technologies that make true presbyopia correction a tangible possibility is focus-tunable lenses. One of the earlier works in this space was the Alvarez lens (16). These lenses consist of a pair of cubic phase lenses that have complementary correction to one another – when overlapped, they have no refractive power, and shifting one relative to the other can introduce either positive or negative powers. However, this design is not ideal since mechanical components tend to wear down faster, and more extreme powers inherently reduce field of view as the overlapping region of the lenses is reduced.

Since then, there has been much effort made to develop wide field-of-view lenses for use in presbyopic correction. Two main categories that these lenses fall into include liquid-crystal

lenses and liquid-filled membrane lenses. In the space of liquid lenses, there have been proposals such as using a liquid lens mounted on a polydimethylsiloxane (PDMS) contact lens (21), but also work put into creating functional larger lenses for eyeglasses-based correction (20). Liquid-crystal lenses, on the other hand, may be more flexible in that they can directly address phase across the field of view, allowing for more varied corrections with a single element, with much recent work on making them a more viable option for eyeglasses, improving focusing time and addressable range (17–19). Another possible approach involves using a deformable mirror. Papadatou et al. (47) for example used one to create and test simultaneous vision correction.

Focus-Tunable Glasses Some have taken these lenses one step further and incorporated them in eyeglasses. For example, Wang et al. (26) developed a set of liquid lenses that they tested on a single observer. Unfortunately, their system had a response time on the order of seconds and was partially table-mounted. Using a liquid-crystal diffractive lens, Li et al. (25) created a wearable prototype that had a response time of under a second. However, they did not have any method of focus control that depended on the viewer, making the solution somewhat incomplete. Finally, there is the prototype from Hasan et al. (23) which uses a time-of-flight sensor pointing directly ahead of the user to close the control loop and update the liquid lenses. While an important step, this is somewhat limiting and requires the wearer to move their head around to focus to different distances. Integrating an eye-tracker into the refocusing pipeline would remove this requirement. There have been steps taken to address this, and to make the eye trackers smaller and lower power, but without full integration into a system such that they drive the focus-tunable lenses (42, 48). Finally, it is important to note that none of these have been evaluated and verified to be more preferable or to improve task performance relative to traditional fixed-focus methods of correction.

Comparison of Presbyopia Treatments When it comes to the traditional forms of correction, each tends to fall short on some metric. In terms of acuity, the standard to beat is the correction afforded by simply carrying multiple pairs of single-vision glasses for each desired viewing distance – which is also the least convenient solution by far, giving us ample motivation to improve upon this solution. While bifocals at first glance seem like a reasonable fix – they are in some sense two stacked single-vision lenses – they, along with progressives and trifocals, have a decreased field of view at any given distance. This has been shown to impair contrast sensitivity and depth perception when walking, increasing risk of injury (5). Furthermore, progressives have been shown to perform worse in tasks requiring lateral head movement than those wearing single-vision lenses (7).

Monovision and simultaneous-vision contacts (both diffractive or concentric lenses) offer both a wider field of view as well as clear vision over a range of distances without requiring head motion. While monovision performs better than diffractive and concentric lenses in terms of visual acuity and contrast and tends to be preferred by wearers over concentric lenses, it tends to perform worse on stereoacuity (10). Neither of these corrections, however, outperform reading glasses, or even bifocals. Studies show that monovision (12) and concentric lenses (13) are worse for near-distance tasks and acuity than reading glasses. Also, in terms of contrast and stereoacuity, reading glasses (8, 11) and bifocals (9) outperform both monovision and concentric lenses.

What autofocals and other focus-tunable lens solutions aim to do is provide corrected vision that has comparable field of view to reading glasses, monovision, and simultaneous vision contacts, while also providing acuity and task performance that outperforms them in tasks requiring focus at various distances. The wider field of view should improve safety outcomes, while providing overall better vision based on acuity and task performance.

Results (Extended)

Numerical values for the mean contrast sensitivity, speed, and accuracy were excluded from the main paper for brevity. The exact numbers are listed here.

The mean contrast sensitivities are as follows: 2.02 logCS for autofocals with monovision wearers, 2.01 for monovision, 2.02 for autofocals with progressives wearers, and 1.98 for progressives.

The mean speed in each correction is: 0.55 letters/sec for autofocals with monovision wearers, 0.51 for monovision, 0.53 for autofocals with progressives wearers, and 0.50 for progressives. The mean accuracy in each correction is: 94.6% for autofocals with monovision wearers, 94.6% for monovision, 95.1% for autofocals with progressives wearers, and 87.9% for progressives.

Preference Study Feedback

A few themes emerge from the written and verbal feedback we received. Most users commented that the clicking sound emitted by our autofocal system during refocusing is distracting. The clicking occurs due to a voice coil that controls the lenses; Optotune states that this can be mitigated by low-pass filtering the control current. One user mentioned that the magnification changes during refocusing could be uncomfortable. This is an important effect that requires further investigation, but we note that the magnification becomes less pronounced the closer the lenses can be mounted to the eyes. We also had several users comment that they prefer their own correction over autofocals for the metric of ease of refocusing because they were used to their own. This speaks to the high adaptability of the human visual system, but we note that despite the disadvantage in familiarity, the natural accommodation mechanism of autofocals is preferred overall.

We also had some seemingly contradictory feedback about the field of view (FOV) of the lenses. On an earlier glasses-like form factor (fig. S2), we had feedback from users saying they enjoyed having “full field vision” wearing autofocals. On the other hand, with the new VR form factor, users felt like they had less peripheral vision. Though they seem to be opposing statements, it can be seen how both could hold true. In the glasses form factor, the periphery is only blocked by the glasses’ temples. The users then noticed how the FOV where they had sharp vision at each distance now encompassed the entire lens as opposed to just a section. However, in the VR form factor of the newer prototype, much of the area beyond the lenses is blocked by the headset. This leads to a scenario where the lenses themselves offer a bigger sharply focused FOV, but the headset form factor performs poorly in overall FOV.

Next, we discuss the issue of physical comfort, comments about which were an impetus for the second prototype's (fig. S1) form factor. Comfort is directly related to the weight of the lenses. There are several ways to reduce weight, including using higher refractive-index liquids to reduce the necessary lens material, or different types of focus-tunable technology such as diffractive liquid-crystal lenses (19). Of course, the lenses are not the only source of weight on potential gaze-contingent eyewear: batteries, control electronics, and cameras also need consideration.

Finally, we had a few users suggest improvements that would facilitate their personal use of automatic eyeglasses. As an example, one suggested that if they were looking at something at a fixed distance for an extended period of time, that the glasses could "lock" their focus state until some visual gesture was given. This would avoid refocusing jitter, and any momentary distractions from a passerby could be ignored unless important enough to resume "unlocked" refocusing.

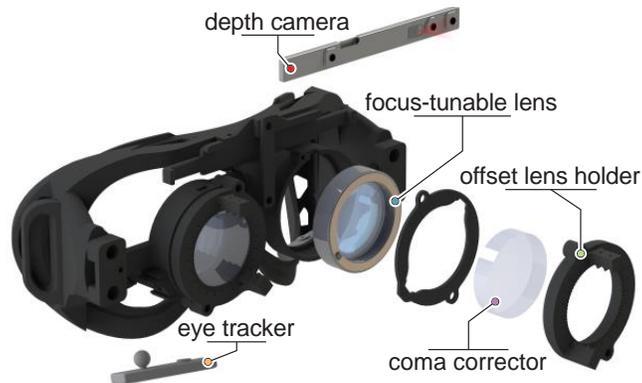


Fig. S1. A partially exploded view of the headset computer-aided design model. Models of the RealSense R200 depth camera, Pupil Labs eye trackers, Optotune EL-30-45 lens, coma corrector, and offset lens holder are illustrated.

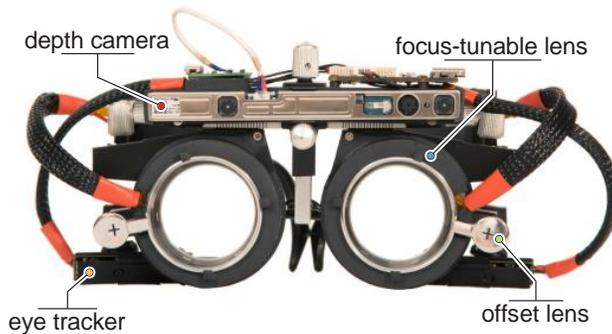


Fig. S2. An image of the previous prototype, which had a glasses form factor. It was retired in favor of the VR or ski goggles form factor prototype (fig. S1) due to slippage of the glasses, which caused eye-tracking performance to degrade. (Photo credit: Robert Konrad, Stanford).

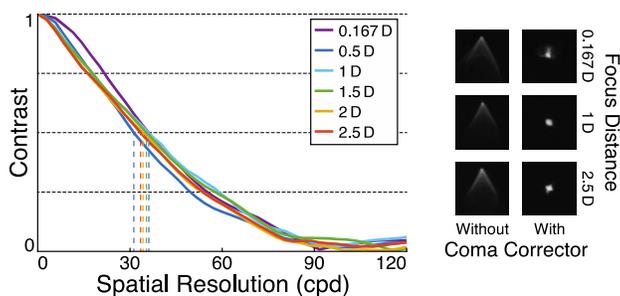


Fig. S3. Optical characteristics of the Optotune EL-30-45 focus-tunable lenses captured using a camera. (Left) The MTF curves of the lenses, focused at different distances. Dashed lines indicate the MTF50 for each distance. The lenses perform at or above 20/20 acuity, with the MTF50 above 30 cpd at all distances using a pupil-sized viewing aperture. (Right) Point spread functions captured through the lenses at different distances. When the optical axis is horizontal, gravity causes the liquid in the focus-tunable lenses to pool to the bottom, leading to coma (left column). We use a custom coma corrector element to remove the aberration (right column).

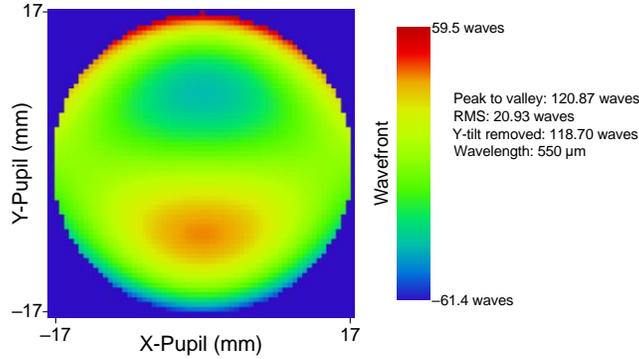


Fig. S4. A wavefront map of the coma correctors, designed in Zemax.

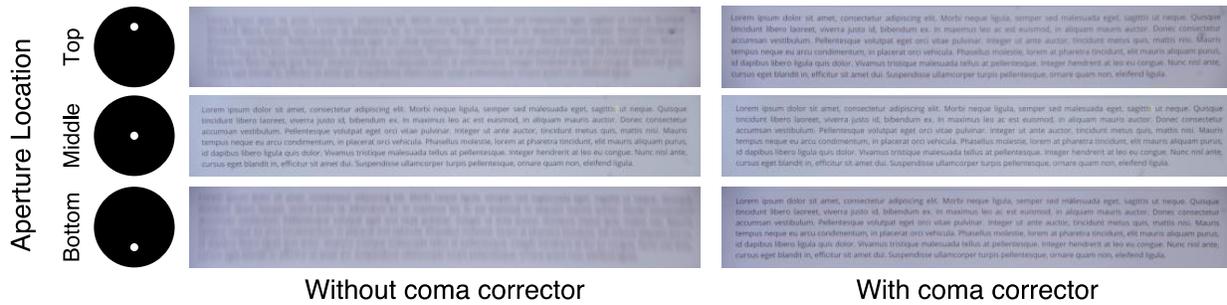


Fig. S5. Focus accuracy at different positions before and after addition of the coma corrector. We simulate a human viewer using a 4 mm sub-aperture at different vertical locations, and image using a camera focused to the same distance in all cases. Without the coma corrector, the resulting focal power of the lens varies across the full aperture. With the sub-aperture at the center, the image is in focus. However, with the sub-aperture shifted 10 mm above or below the center, the focal power changes by ± 1.5 D, resulting in observable blur. On the other hand, with the coma corrector in place, the variance in focal power is negligible, verifying its proper operation.

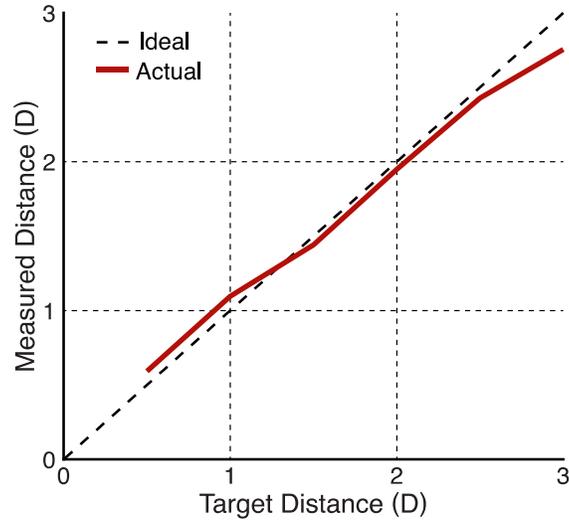


Fig. S6. Measured optical lens power as a function of target lens power. An ideal response would track the target distance exactly. Our measured results show that the Optotune EL-30-45 lenses stay within ± 0.1 D of the correct value at the distances used for the quantitative user studies (up to 2.5 D). The nearest distance of 3 D shows a slightly larger 0.245 D offset.

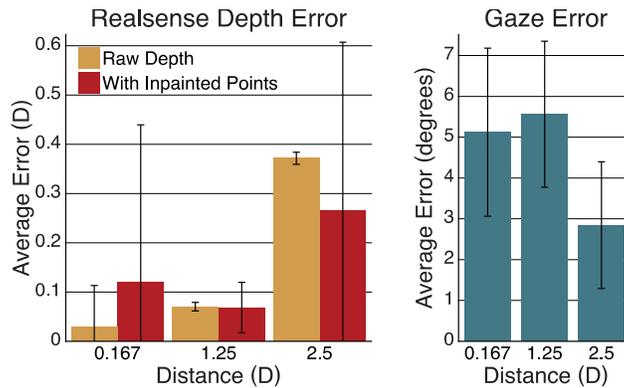


Fig. S7. Evaluations of the accuracy of the two main external sensors. (Left) The RealSense R200 has relatively little error, though naïve inpainting can greatly impact the noise in the measurement. (Right) On the other hand, the gaze estimate from the eye trackers seems to be off by about 3–5° in practice, a sizeable departure from the ideal 0.6° error. Error bars are standard deviation.

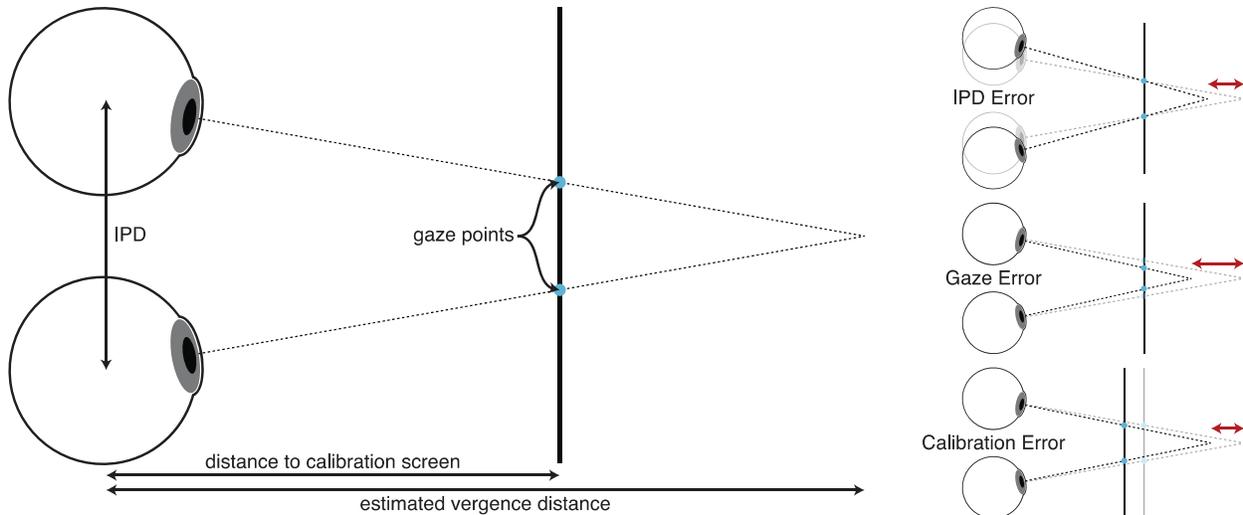


Fig. S8. A visual representation of sources of error in the estimated vergence. (Left) The geometry relating interpupillary distance (IPD), calibration distance, and the gaze points to the vergence estimate. (Right) Given the correct gaze and calibration, *IPD error* results in vergence error from the assumed eye positions being incorrect. If the gaze points were co-located, i.e. the viewer is looking at the calibration distance, IPD error has no effect. *Gaze error* results from the gaze points themselves being incorrect. If both gaze measurements are to the right, or both to the left, the errors may cancel. However, when the gaze measurements are too far inwards (as above) or outwards, it results in vergence error. *Calibration error* occurs when the user calibrates at a distance different from the assumed calibration distance. As the vergence distance approaches infinity (0 D), error due to calibration approaches 0 since the eyes point perpendicular to the calibration screen, reducing the importance of the screen's exact distance.

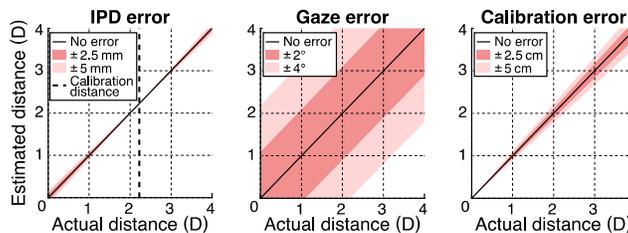


Fig. S9. Error in the estimated vergence distance from various sources. (Left) As seen above, even a large error in IPD measurement of 5 mm results in little vergence error. (Center) On the other hand, a gaze error of 2° in each eye can lead to over 1 D of error. (Right) Finally, the calibration distance is also a relatively minor source of error, most noticeable at near distances. Calibrating the eye trackers at a distance 5 cm away from the assumed 45 cm mark may lead to about 0.5 D of error. See fig. S8 for a visual interpretation of error sources.

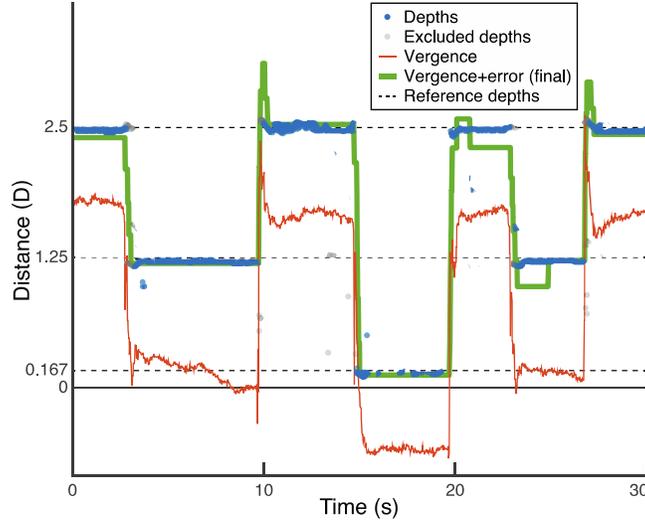


Fig. S10. An example recording of the sensor fusion algorithm. The raw vergence (red) has the right trends but is noisy and has an offset from the actual fixated depths. Depth values from the R200 (dots) are filtered to valid values (blue). The depth informs an error estimate, which is used to obtain the final lens output powers (green).

Algorithm S1. Sensor fusion: Vergence + error.

```

1   $v_{filt} = v_{err} = 0; \alpha_v = 0.2; \alpha_e = 0.2;$ 
2  while running do
3      if checkNewValidEyeTrackerVergence() then
4           $v \leftarrow \text{getNewValidEyeTrackerVergence}();$ 
5           $v_{filt} \leftarrow \alpha_v v + (1 - \alpha_v)v_{filt};$ 
6      if checkDepthUpdated() then
7           $d_{filt} \leftarrow \text{getLatestDepth}();$ 
8          if  $|d_{filt} - (v_{filt} + v_{err})| < 0.5$  or  $\max(v_{filt}, d_{filt}) < 0.75$ 
              or  $\min(v_{filt}, d_{filt}) > 2.0$  then
9               $v_{err} \leftarrow \alpha_e(d_{filt} - v_{filt}) + (1 - \alpha_e)v_{err};$ 
10     if  $|\text{getLensPower}() - (v_{filt} + v_{err})| > 0.25$  then
11         setLensPower( $v_{filt} + v_{err}$ );

```

Algorithm S2. Depth denoiser.

```
1  $d_{denoiser} = 0; \alpha_d = 0.2;$ 
2 while running do
3   setDepthUpdated(False);
4   if checkNewR200DepthMeasurement() then
5      $d \leftarrow$  getNewR200DepthMeasurement();
6     if  $d_{denoiser} - d < 0.5$  then
7        $d_{denoiser} \leftarrow d;$ 
8       setLatestDepth( $d$ );
9       setDepthUpdated(True);
10    else
11       $d_{denoiser} \leftarrow \alpha_d d + (1 - \alpha_d) d_{denoiser};$ 
```

Preference Questionnaire

A PDF file with the form used to obtain responses for the natural use preference study.

Data S1. A zip file containing comma-separated values (CSV) files with the raw data for participants for visual acuity, contrast sensitivity, and task performance. Order of data within each file corresponds to the order of trials for that measurement.

Data S2. A CSV file containing the raw data for participants for the natural use questionnaire. Each row is a separate users' response.