

# Present and Future of Everyday-use Augmented Reality Eyeglasses

Praneeth Chakravarthula, *University of North Carolina at Chapel Hill, NC, 27599, USA*

*Abstract—Augmented reality (AR) is emerging as the next ubiquitous wearable technology and is expected to significantly transform various industries in the near future. There has been tremendous investment in developing AR eyeglasses in recent years, including about \$45 billion investment by Meta since 2021. Despite such efforts, the existing displays are very bulky in form factor and there has not yet been a socially acceptable eyeglasses-style AR display. Such wearable display eyeglasses promise to unlock enormous potential in diverse applications such as medicine, education, navigation and many more; but until eyeglass-style AR glasses are realized, those possibilities remain only a dream. My research addresses this problem and makes progress “towards everyday-use augmented reality eyeglasses” through computational imaging, displays and perception. My dissertation [1] made advances in three key and seemingly distinct areas: 1) Digital holography and advanced algorithms for compact, high-quality, true 3D holographic displays, 2) Hardware and software for robust and comprehensive 3D eye tracking via Purkinje Images and 3) Automatic focus adjusting AR display eyeglasses for well-focused virtual and real imagery, towards potentially achieving 20/20 vision for users of all ages.*

Enhancing Virtual Reality (VR) has been identified by the National Academy of Engineering (NAE) as one of the fourteen grand challenges for engineering in the 21st century, along with challenges such as “reverse engineer the brain”, “provide energy from fusion”, and “secure cyberspace”<sup>1</sup>. Augmented Reality (AR) is the form of VR in which users see the synthetic “virtual world” imagery overlaid and merged with their real-world surroundings. Many experts predict that VR, and especially AR, will be the next mass market platform taking the place of laptops and mobile phones.

Over the last few years, near-eye display approaches have been emerging at a rapid pace, promising practical and comfortable virtual reality (VR) and augmented reality (AR) in the future. However, they are still limited to large form factors and are unable to allow for continuous focus cues needed to avoid the vergence-accommodation conflict. If AR is indeed to

be the next platform, then AR systems will have to be comfortable enough to be worn for long periods, perhaps all day, like a pair of ordinary eyeglasses (see Figure 1). Otherwise, people will just continue to carry their smartphones in their pockets. Building an AR display eyeglasses that presents well-focused images, both real and virtual, and near and far, in a compact form factor requires overcoming three major challenges:

- A (holographic) augmented reality display in a compact eyeglasses form factor that shows imagery of high quality and resolution, comparable to the current day mobile phones or laptop displays.
- Robust and comprehensive eye tracking to determine the user’s current gaze fixation as well as focal accommodation depth, along with 3D pupil orientation and compensation for eyeglasses slippage.
- Designing the AR eyeglasses to dynamically adjust focus for both the internal display (showing rendered synthetic content) and the external real world scene.

---

XXXX-XXX © 2021 IEEE

Digital Object Identifier 10.1109/XXX.0000.0000000

<sup>1</sup><http://www.engineeringchallenges.org/9140.aspx>



**FIGURE 1.** *Vision for AR:* Future AR wearable displays need to be compact and like a pair of socially acceptable, everyday-use prescription eyeglasses.

### ***New opportunity with AR displays***

Future AR displays can take advantage of promising technology and capabilities such as powerful processors, and outward looking depth sensing for tracking and hand-based user interaction (e.g. Microsoft HoloLens 2, Meta Quest 3), while achieving a compact form factor and wide field of view. If rapid, accurate, and robust binocular eye tracking is added, a system could measure the user's object of attention in both the real and virtual world. Then adaptive focus could be added to the real world (external) view alongside a separate adaptive focus for the virtual world (internal AR display) view, ultimately bringing into focus both the real and virtual imagery. Such a display could also operate just as *auto-focus prescription eyeglasses*, with the virtual content turned off. This would allow for potentially providing 20/20 vision for users of all ages, a long-standing challenge since the century-old invention of prescription lenses by Benjamin Franklin.

## **HOLOGRAPHIC DISPLAYS**

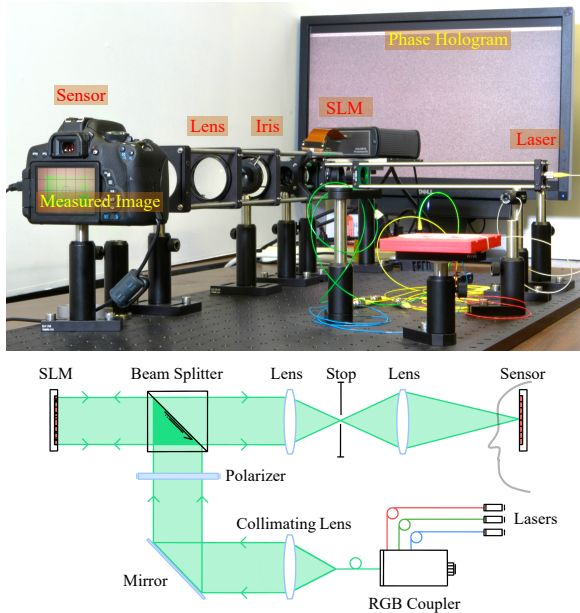
Mass adoption of AR near-eye displays (NEDs) will likely require displays that can seamlessly integrate into our day-to-day activities, like a pair of everyday eyeglasses. Of the several existing technologies, holography is perhaps the only demonstrated technology that can potentially enable everything needed for a near-eye display in an eyeglasses like form factor, i.e., high-resolution, wide field-of-view, compact form factor, per-pixel focus control, eye aberration correction and many more. This section briefly introduces holographic displays, current progress and future directions.

### **Computer-Generated Holography**

Computer-generated holography (CGH) is a technique for numerically simulating the process of recording a light field of an object. This recording, often called a hologram, when played back on a spatial light modulator (SLM) creates a 3D image that preserves depth and parallax (see Figure 2). Given that the existing hardware SLMs can not modulate both amplitude and phase of light simultaneously, typically phase-only modulating SLMs are used in holographic displays due to their high light efficiency. Computing the appropriate phase modulation by numerically simulating the light propagation, diffraction and interference effects, referred to as phase retrieval, forms the core of CGH and directly affects the displayed image quality.

#### *Optimizing Holograms using Gradient Descent*

Phase retrieval algorithms, however, are still in their infancy. Previous holographic display approaches resorted to heuristic encoding methods or iterative methods relying on various relaxations, which severely compromised the holographic display quality. In my dissertation, I introduced a way to directly optimize the hologram phase pattern using complex Wirtinger derivatives and first-order gradient-descent methods for the first time [3], resulting in over two orders of magnitude improvement in image quality compared to previous approaches (see Figure 3). This Wirtinger holography framework is flexible and facilitates the use of different loss functions, including learned perceptual losses parameterized by deep neural networks, as well as stochastic optimization methods. As a result, it integrated well with popular machine learning libraries and



**FIGURE 2.** *Prototype Holographic Display:* (Top) The display-prototype setup to generate true 3D holographic images. (Bottom) Setup schematic, with RGB lasers that are coupled into a single-mode fiber which illuminates the reflective SLM displaying the phase pattern. The modulated wave which carries the image is measured using a conventional intensity camera.

later inspired a variety of AI-driven neural holography methods [13], [8], [11].

#### Computing High-fidelity 3D Holograms

Holography is capable of projecting complex 3D objects with depth and parallax cues. However, most phase retrieval methods for holography treat 3D objects as a collection of discretely sliced layers in depth. Such multilayer 3D holography approaches fail to model wavefronts in the presence of partial occlusions. While 4D light fields offer a better representation to handle depth and parallax cues, light field-based holographic stereogram methods have to make a fundamental trade off between spatial and angular resolution. In addition, existing 3D holographic display methods rely on heuristic encoding of complex amplitude into phase-only pixels which results in holograms with severe artifacts. Fundamental limitations of the input representation, wavefront modeling, and optimization methods prohibited artifact-free 3D holographic projections in today's displays. The Hoge-free holography technique [5] I introduced lifts these limitations and enables optimizing true 3D holograms supporting both depth- and view-dependent effects, including those

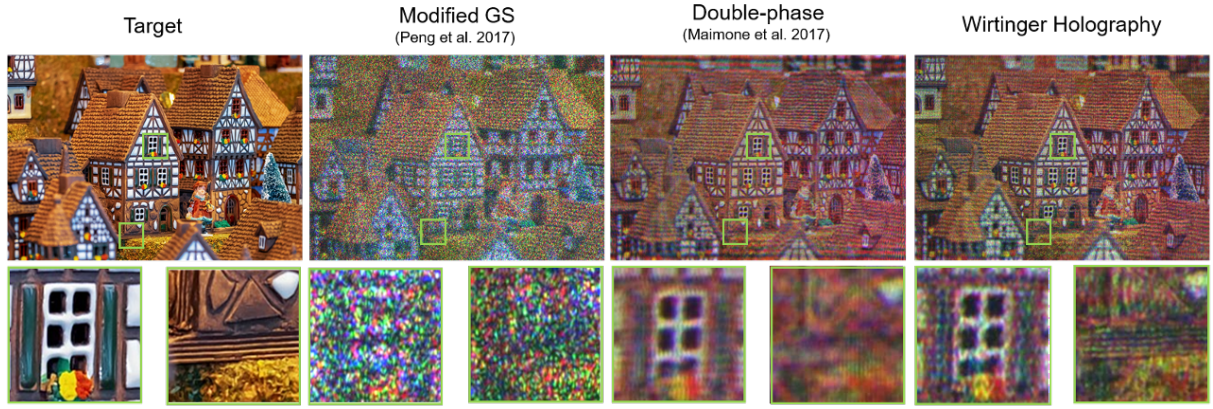
at challenging occlusion boundaries, for the first time (see Figure 4). This approach not only overcomes the fundamental spatio-angular resolution trade-off typical to stereogram approaches, but also avoids heuristic encoding schemes to achieve high image fidelity over a 3D volume.

#### Compensating for Real-world Errors

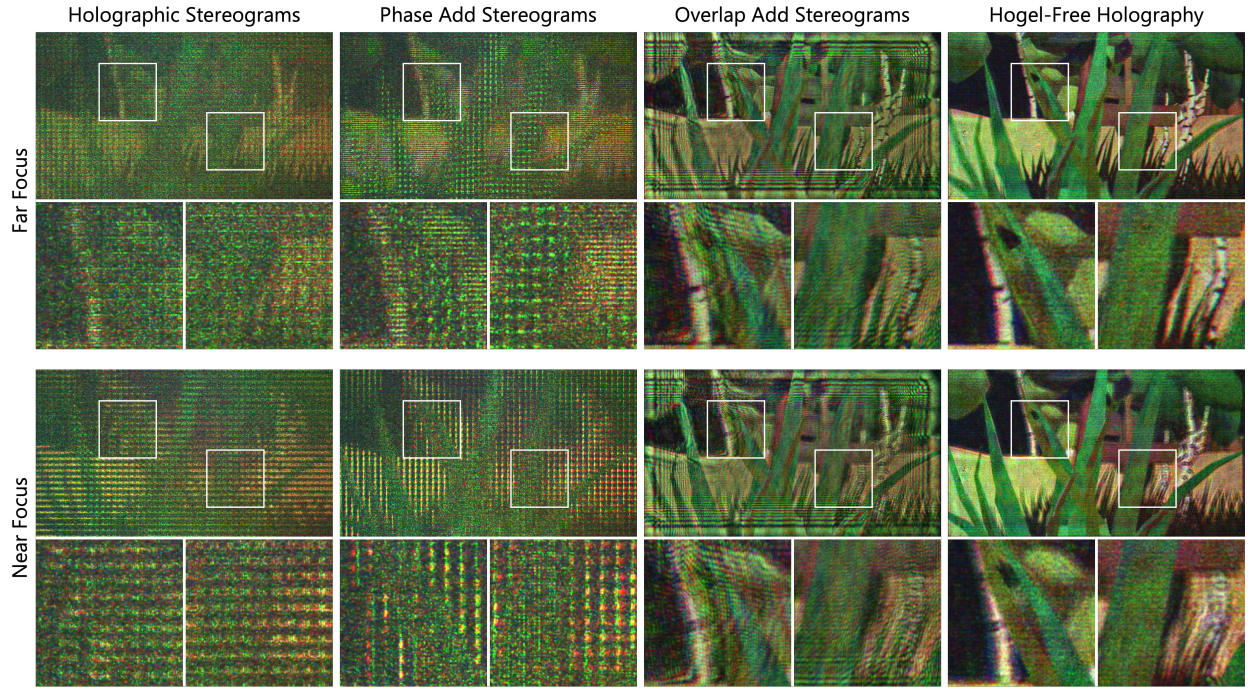
While these methods significantly improved the hologram quality, the image quality of existing holographic displays is far from that of current generation conventional displays, effectively making today's holographic display systems impractical. We realized that the severe real-world deviations from the idealized approximations of the light transport model used for computing holograms predominantly causes this gap in image quality. One can compensate for such deviations by using an active camera-in-the-loop calibration for computing each hologram [11]. However, this is not practical and enlarges the form factor of near-eye display systems. Instead, we learn the deviations of the real display from the ideal light transport model using the images measured with a display-camera hardware system. After this “unknown” light propagation is learned, we use it to compensate for severe aberrations in real holographic imagery. This approach can be practical and applied at end-of-manufacturing-line, akin to existing electronic devices. Moreover, this learned hardware-in-the-loop approach is robust to spatial, temporal and hardware deviations. The learned hardware-in-the-loop approach also improves the image quality of holographic displays qualitatively and quantitatively, both in SNR and perceptual quality by fully compensating unknown aberrations and erroneous, non-linear SLM phase delays, without explicitly modeling them (see Figure 5). I believe (a variation of) this method will be applied in future holographic displays. The learned hardware-in-the-loop framework is not only applicable to holography and has already shown significant improvements in other domains such as fabrication using lithographic techniques [18].

#### Future Directions

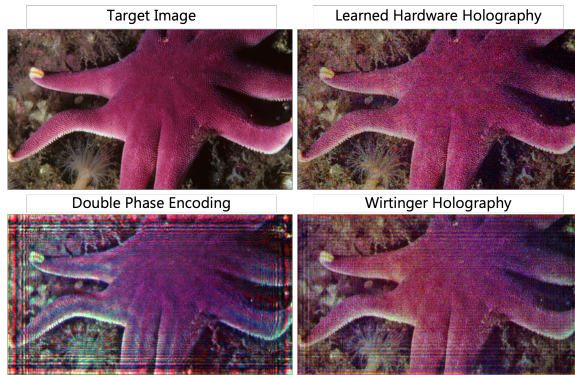
The ultimate requirement of a near-eye display system is to produce natural-looking 3D imagery for a realistic and comfortable all-day viewing experience. To this end, displays need to cater to the perceptual comfort of the user. While the above mentioned works improved the image quality of holographic displays as seen by a camera, several challenges need to be overcome to present comfortable and immersive viewing experiences to the user including low-latency compute, compression and transmission of holograms



**FIGURE 3.** *Experimental results on prototype display.* We present RGB color images with each color channel captured sequentially for existing phase retrieval methods and Wirtinger holography optimization. All images are captured with the same camera settings and the output power of three lasers are tuned before acquisition to approximately white-balance the illumination. Directly optimizing for hologram phase using Wirtinger holography outperforms traditional approaches.



**FIGURE 4.** *Experimental validation of 3D Hogel-free Holography.* Hogel-free generates high-quality 3D holograms with accurate depth- and view-dependent effects, and without sacrificing spatio-angular resolution. Existing approaches rely on computing sub-holograms, so-called hogels, for achieving such effects. However, the chosen size of hogels is typically scene-dependent and follows a trade-off between angular and spatial resolution of holographic imagery. We lift these limitations by formulating a holographic forward model and phase retrieval that takes RGB-D light fields as input and directly optimizes the target phase, without spatial segmentation into hogels or phase encoding approaches.



**FIGURE 5.** *Learned Hardware-in-the-loop Holography* Holographic displays often show poor image quality due to severe real world deviations in the light propagation compared to the simulated “ideal” light propagation model. We approximate the unknown light propagation in a real display and the resulting aberrations via a trained neural network. We use this trained network as a proxy to the real display and compute phase holograms that compensate for real world aberrations in a hardware-in-the-loop fashion. The proposed method produces images on real hardware that are aberration-free and close to the target image.

[19], [15]. Of these several remaining works, I discuss here the aspects I believe are important and need to be addressed immediately for practical holographic displays.

**Eyebox and Field-of-View** Holographic displays can create high quality 3D images while maintaining a small form factor. However, the limited spatial resolution of existing dynamic spatial light modulators imposes a tight bound on the diffraction angle due to its limited étendue. This creates a tradeoff between the eyebox size and the field-of-view. We have recently proposed solutions based on learned étendue expanding elements [14] and novel holographic optical elements [17], the design of which builds off of the complex field Wirtinger optimization, that effectively increased étendue by two orders of magnitude. Exploring novel hardware and software solutions to expand the étendue, and hence the field of view and eyebox, while maintaining high image quality of holographic displays is of paramount importance and needs further investigation.

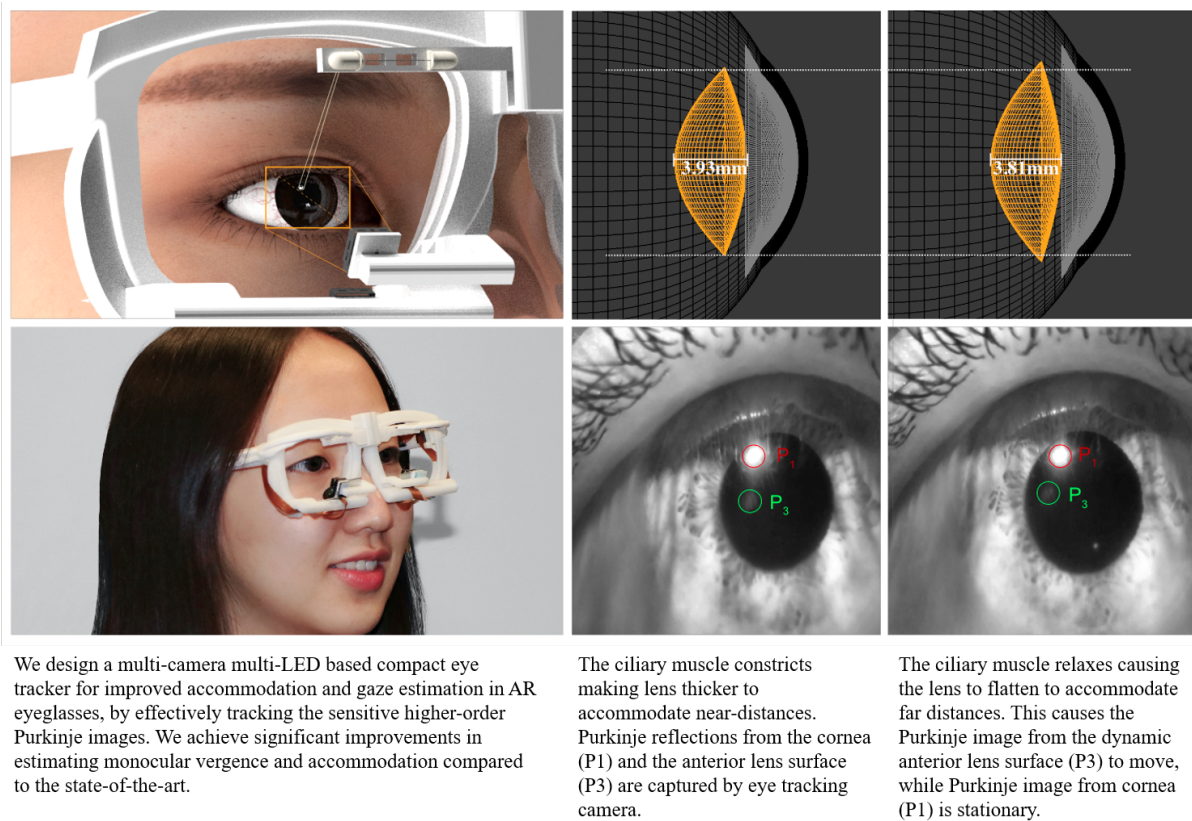
**Laser and Subjective Speckle Reduction** Holography relies on coherent illumination which creates speckle noise in the final image due to random interference. Moreover, human eyes are imperfect optical systems

and have aberrations of their own, resulting in further degradation of the perceived image quality due to subjective speckle. Handcrafting the final image phase to be smooth can create speckle-free images, but results in a tiny eyebox with a non-uniform energy distribution, making that approach impractical. On the other hand, speckle reduction with partially coherent sources reduces resolution. Averaging sequential frames can reduce speckle but requires high speed modulators consuming the temporal bandwidth of the system. Therefore novel systems and accompanying algorithms are needed to mitigate the effect of laser and subjective speckle in the final imagery [10]. Our preliminary investigation into subjective speckle reduction by employing point-spread function of the eye in the hologram computation pipeline [6] is a promising direction and needs further investigation with rigorous user studies.

**Perception-aware Algorithms and Displays** All existing holography approaches aim at generating high-quality images as seen by a camera. Given the significant differences between a camera and the human visual system, not only is this approach inaccurate, but it often results in sub-optimal utilization of SLM bandwidth, for example, to create high resolution imagery over the entire field of view despite the eye perceiving high resolution only in the fovea. Moreover, the holographic field as sampled by the eye pupil is highly varying and we need pupil-aware holographic displays that maximize the perceptual image quality irrespective of the size, location, and orientation of the eye pupil [2], [16], [12]. Therefore, perception-aware algorithms and display configurations that maximize perceptual image quality irrespective of the eye pupil's size, location, and orientation are necessary and represent an interesting future direction.

### 3D GAZE AND FOCUS TRACKING

In viewing objects at different depths, our eyes independently rotate to fixate on objects and perceive them as a single, unified image, a phenomenon called *vergence*. Simultaneously, the eye focuses on the object to produce a sharp image on the retina, a mechanism referred to as *accommodation*. By including rapid, accurate, and robust eye tracking in near-eye displays, a user's object of attention can be measured and virtual imagery can be displayed at correct focal depth, enabling viewing of both real and virtual imagery simultaneously. Accurate and rapid vergence tracking can potentially also enable gaze-contingent rendering and displays [7].



**FIGURE 6.** *Purkinje image-based comprehensive 3D eyetracking.* We track the reflections from the inner surfaces of the eye to independently and simultaneously measure the vergence and focal depth of the eyes.

While existing eye trackers can measure vergence, note that depth from vergence can only approximate the accommodation depth of the eye and typically occurs only in users with normal vision. Therefore, we must not only measure vergence accurately but also separately measure monocular accommodation. Such a display combined with robust monocular vergence and accommodation tracking of eyes can also potentially act as a pair of auto-focus “everyday” prescription eyeglasses, with the virtual display content turned off.

Estimating both monocular vergence and accommodation independently and simultaneously requires tracking of not just the pupil but also the continuous changes in eye lens shape. Sophisticated equipment like an auto-refractor or a Shack-Hartmann wavefront sensor can measure the change in lens accommodation by analyzing the reflected wavefront from the retina through the eye. However, these are too bulky to incorporate into eyeglasses form factor displays. I realized that the eyes are imperfect optical systems and part of the light incident on the eye gets reflected from every

surface encountered in the eye (i.e., front and back surfaces of cornea and lens). We looked at imaging the reflections from various surfaces of the eye, called the Purkinje reflections, using miniature cameras fitted into the eyeglasses frame. Since these reflections occur from curved surfaces of the eye including the front and back surfaces of the eye’s crystalline lens, they are sensitive to eye gaze changes and accommodation, and are well-suited for independent vergence and accommodation tracking (see Figure 6).

We believe that future AR eyeglasses will be individually owned and personalized akin to the current day personal smart-phones. Such personalized AR eyeglasses will have eyetrackers and prescription correction lenses, also catering to the individual user. Toward this vision, we demonstrated a compact eyeglasses-style eyetracker based on tracking higher-order Purkinje reflections from the eye, with the tracker customized to learn the nuances of eye movement for a specific user via a convolutional neural network (CNN). The higher-order Purkinje reflections occurring

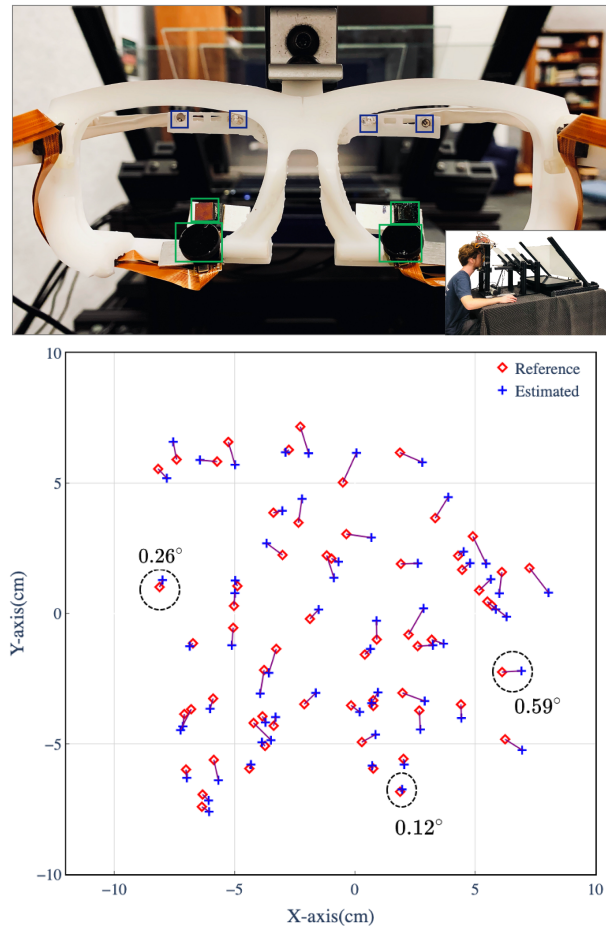
from the deeper layers of the eye are typically faint and are often difficult to capture robustly and consistently. To this end, we also introduced a sophisticated, physically-accurate, and anatomically-informed synthetic eye model to simulate the Purkinje reflections from the eye. We used this synthetic eye model to optimize for the positions of cameras and infrared LEDs to robustly and consistently capture the Purkinje reflections.

We designed a special room-sized multi-plane display setup with adjustable focal planes to collect the ground truth calibration data required to train the neural network (see Figure 7). The same setup is later used to also measure the accuracy of 3D eye tracking – tracking both vergence and accommodation independently and simultaneously. The trained network was then used to estimate the gaze and accommodation from the Purkinje images. Parameterizing and learning the relationship between the Purkinje images, and gaze vergence and accommodation depth allowed us to not only estimate the vergence and accommodation robustly, but also at faster frame rates. We experimentally found that this eye tracker outperforms the existing state-of-the-art (vergence-only) eyetracking methods. I anticipate that, in the future, the calibration data can be collected directly via focus-supporting near-eye display eyeglasses.

### Future Directions

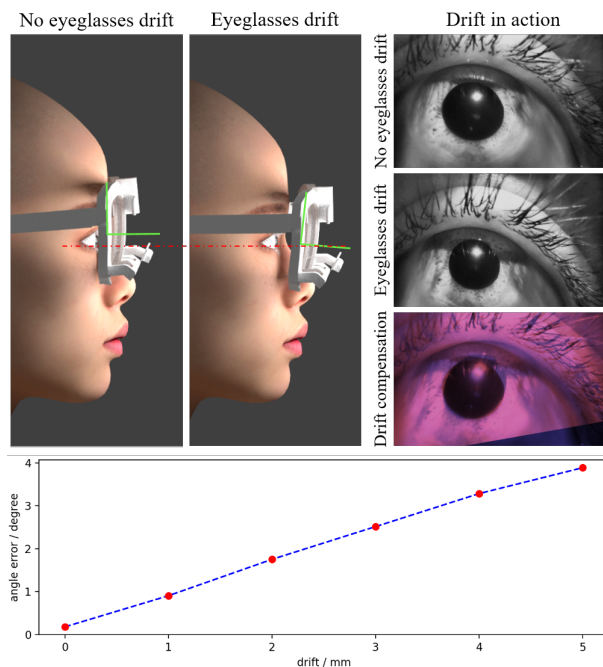
It is clear that the future eye trackers should be capable of independently and simultaneously measuring vergence and accommodation. Recent developments in ultra-thin and flat optics [4], [9] also show promise of miniaturizing such eye trackers to small form factors that seamlessly fit in the rim of eyeglasses. Creating an anatomically-accurate eye model in simulation as well as a hardware prototype will immensely benefit exploring novel eye tracking hardware setups. While the eye model we designed for our work on a Purkinje image tracker is perhaps the best available, it is only an approximate model that might not generalize well for a wide variety of users, potentially leading to minor discrepancies in the simulated and real positions of the Purkinje images. However, the approximations only result in a one-time, additional step to manually refine the optimized eye tracker configuration in simulation. A better approach involves modeling the physiological variability of eyes to ensure generalization across a diverse population, alongside designing optimization strategies that reduce or remove the need for any manual adjustments in the future.

Another major challenge that needs to be solved



**FIGURE 7.** (Top) *Eyeglasses with multi-plane setup.* The LEDs are marked with blue boxes and the cameras are marked in green boxes. The full-view of multi-plane room-sized data capture system is shown in the bottom-right corner insert. (Bottom) *2D Error visualization at 40cm depth.* Reference and estimated points are marked with a red square and a blue cross, respectively. Vergence error, i.e., error between the estimated gaze angle and the ground truth, for three pairs of measurements is also visualized.

is compensating for eyeglasses slippage. Eyeglasses sliding down the bridge of the nose is a common problem encountered by all people wearing eyeglasses. We anticipate that, in the future, when the tracker is attached to AR eyeglasses, this kind of slippage will be a significant problem. The slippage of eyeglasses would cause the attached eyetracking cameras to drift in space, causing errors in the gaze estimation. A pilot study we conducted demonstrated that even slippage of 1mm to 5mm along the nose bridge can significantly impact the accuracy of gaze estimation



**FIGURE 8.** (Top left) Drifting of the eyeglasses over the nose bridge. (Top right) Errors caused by drifting of the eyeglasses and hence the eye tracking cameras. Notice how the captured frames differ when the eyeglasses drift. We compensate for the eyeglasses drift using geometric techniques and validate by overlaying the non-drifted image and the warped, drifted image to show alignment. (Bottom) The error in gaze estimation caused due to the eyeglasses slippage over the nose bridge.

(see Figure 8). Such slippage could be severe in an eye tracker in-the-wild. However, such errors caused by drift can be compensated for by taking advantage of the multiple cameras and IR LEDs that are going to be present in eye trackers. The multiple cameras allow for estimating the eye's 3D position relative to the eyeglasses frame, making it possible to account for the eyeglasses slippage in the gaze estimation pipeline. Figure 8 shows drift compensation for one frame of a real experiment by overlaying the drift-compensated camera image with that of a non-drifted camera image. While this uses geometric techniques to estimate the camera drift via detecting landmark features on the captured eye frames, I expect this to be included as part of an end-to-end learning framework in the future.

## AUTO-FOCUS AR EYEGLASSES

Vergence and accommodation are neurally coupled, i.e., as the vergence angle changes, the eye adjusts its

accommodation depth, thus bringing the scene into focus. Proper matching between vergence distance and accommodation depth is important in AR and VR, the absence of which, called the vergence-accommodation conflict (VAC), causes fatigue and discomfort. Additionally, in AR, the virtual images need to be co-located with real world objects at the same focal depth so that both real and virtual objects appear sharp.

### VAC with prescription eyeglasses and adaptation

For users who need corrective lenses in their everyday lives ("near-sighted", "far-sighted"), the situation is even more complex, because these users already have to deal with VAC even without AR or VR. Consider a "near-sighted" user who can comfortably verge and accommodate to, say, 0.5 meters, but needs corrective lenses to focus clearly at objects at 10 meters. When he first uses the corrective "distance" lenses, an object at 10 meters appears in focus (because to his eyes, it is at 0.5 meters), but he will verge to 0.5 meters, giving him "double vision". Only after many hours, days or even weeks of wear, does his vision system gradually adapt to verging at 10 meters while still accommodating to 0.5 meters.

When donning a head mounted display (HMD), users requiring vision correction still need to wear their corrective lenses. In fact, as shown in Table 1, users of any vision type – normal or with refractive errors – need independent adjustments for the virtual image depth and real world vision correction for an AR display to support focus cues for both real and virtual imagery. Especially for presbyopes (people over 40 years of age), who account for about 40% of US population<sup>2</sup>, this does not solve the problem because the user's range of focus is restricted by the focus range of the lenses being worn at any moment - for instance "reading" glasses or "driving" glasses. Installing bifocals, trifocals, or progressive lenses merely puts a particular distance in focus at one vertical angle, forcing the user to tilt their head up or down to bring in focus a real-world object that is at a particular distance. Inventors, since at least Benjamin Franklin, have tried to solve the problem of getting objects at all distances to be in focus, but even the most recent offerings require the user to turn a focus knob on the lens (e.g., Alvarez lens) to adjust the depth of the focal plane - an unacceptably awkward requirement for most users.

For dynamically adjusting the external corrective lens power, we need a *tunable-focus lens* that can operate over a range of focal distances. With robust

<sup>2</sup><https://www.census.gov/prod/cen2010/briefs/c2010br-03.pdf>

User vision type	External real world correction	Internal AR display correction
Normal-vision	No correction	Dynamic
Myopic ('near-sighted')	Static	Dynamic (Offset by prescription)
Hyperopic ('far-sighted')	Static	Dynamic (Offset by prescription)
Presbyopic ('limited accommodation')	Dynamic	Static

**TABLE 1.** Comparison of focus adjustment requirements for different users that are viewing well focused imagery of both real and virtual objects at all distances.



**FIGURE 9.** (Left) The view of both real and virtual imagery for a presbyopic user with distant vision – with the accommodation depth fixed at 7 m – on a conventional AR HMD. The virtual bunny is at a mid-distance (1 m) together with the stamp (0.25 m), text book (1 m) and bicycle (5 m) arranged at near, medium and far distances respectively. Notice that both the real and virtual imagery appears blurred to the user as none of the objects are in the presbyopic user's accommodation range. (Middle) A presbyopic user with near-zero accommodation range looking through our auto-focus AR eyeglasses. Our prototype AR eyeglasses are capable of providing well-focused imagery of both real and virtual objects at all depths by independently adjusting for the user focus for viewing both real world and virtual imagery from the internal display, based on the user's current eye accommodation state. (Right) The well-focused view of both real and virtual objects of the same scene on our auto-focus AR eyeglasses, due to independent focus adjustments for both real and virtual. Notice that the letters on the textbook at the mid-distance (1 m) are in sharp focus, as well as the virtual bunny, which also is set to appear at the mid-distance.

3D eye gaze tracking, a multitude of outward looking cameras on the headset, and a prior knowledge of the degree of a user's refractive error in their eye, we can determine the depth of the object of interest and adjust the focus of the external corrective lens accordingly, so as to bring the real world target into sharp focus. The internal display should be capable of rendering objects at various depths while spatially registering them to the real world, providing the depth cues either statically or dynamically. Providing depth cues statically ensures the correct retinal blur, whereas providing dynamic depth cues requires rendering objects away from the focus plane with an appropriate amount of retinal blur.

We built a system that corrects dynamically for the focus of the real world surrounding the near-eye display of the user and simultaneously the internal

display for augmented synthetic imagery, with an aim of completely replacing the user prescription eyeglasses (see Figure 9). The ability to adjust focus for both real and virtual stimuli will be useful for a wide variety of users, but especially for users over 40 years of age who have limited accommodation range. Our solution employed a tunable-focus lens for dynamic prescription vision correction, and a varifocal internal display for setting the virtual imagery at appropriate spatially registered depths. Therefore, it can be seen that both real and virtual imagery are in focus at all near, medium and far distances. Also, with such auto-focus AR eyeglasses, the vision of myopic, hyperopic and presbyopic users can be significantly improved with the perceived image quality being close to that of a person with 20/20 vision at all depths.

## Future Directions

The design for auto-focus eyeglasses is limited by currently available off-the-shelf tunable-focus lenses with a severely diminished field of view, and sometimes gravity induced comatic effects. Liquid crystal-based tunable-focus lenses might be an interesting future direction to achieve compact and wide FoV dynamic real-world correction. Finally, all described components need to come together into a compact eyeglasses form factor, and miniaturization and design of such eyeglasses is an interesting future direction in itself. These eyeglasses need to be unencumbered, untethered and should be worn in a more natural position to allow a wider range of head movements. This requires incorporating low-power and/or distributed computing, and a full six degree-of-freedom tracker to allow for a wide range of spatially registered augmented reality applications. Exhaustive user studies should then be conducted to analyze issues with such eyeglasses.

## CONCLUSION

The discussed components and approaches such as holographic displays, comprehensive eye tracking and auto-focus displays can be instrumental to the success and mass adoption of augmented reality, and pave the way to achieving everyday-use eyeglasses-style augmented reality displays. Holographic displays have already improved with respect to their image quality and miniaturization of these displays with expanded étendue with perceptually-aware hologram generation algorithms will be available in future commercial products. While the eye tracking solution offers perhaps the best performance compared to existing solutions to the best of my knowledge, its user-specific design to efficiently capture Purkinje images needs to be improved to work for a general pool of users. I believe that the idea of an AR display that can utilize the onboard sensors and compute for completely replacing a user's prescription eyeglasses is very powerful and will be fundamental for future displays. I would also make a bold claim that only such auto-focus eyeglasses that avoid the awkward necessity of centuries-old, multiple prescription lenses will be successful in the future. Although it is unclear if all of the described approaches will find their way to commercialization, I believe that some version of the described methods will be adopted in future commercial displays. I am excited by the possibility of future AR displays employing dynamic vision correction and comprehensive eye tracking alongside focus supporting, compact holographic displays, enabling the promise of 20/20 vision at all distances, for all users, for both real world and virtual imagery.

## REFERENCES

1. P. Chakravarthula. *Towards Everyday-Use Augmented Reality Eyeglasses*. PhD thesis, The University of North Carolina at Chapel Hill, 2021.
2. P. Chakravarthula, S.-H. Baek, E. Tseng, A. Maimone, G. Kuo, F. Schiffrers, N. Matsuda, O. Cossairt, D. Lanman, and F. Heide. Pupil-aware holography. *arXiv preprint arXiv:2203.14939*, 2022.
3. P. Chakravarthula, Y. Peng, J. Kollin, H. Fuchs, and F. Heide. Wirtinger holography for near-eye displays. *ACM Transactions on Graphics (TOG)*, 38(6):213, 2019.
4. P. Chakravarthula, J. Sun, X. Li, C. Lei, G. Chou, M. Bijelic, J. Froesch, A. Majumdar, and F. Heide. Thin on-sensor nanophotonic array cameras. *ACM Transactions on Graphics (TOG)*, 42(6):1–18, 2023.
5. P. Chakravarthula, E. Tseng, H. Fuchs, and F. Heide. Hoge-free holography. *ACM Transactions on Graphics (TOG)*, 2022.
6. P. Chakravarthula, Z. Zhang, O. Tursun, P. Didyk, Q. Sun, and H. Fuchs. Gaze-contingent retinal speckle suppression for perceptually-matched foveated holographic displays. *IEEE Transactions on Visualization and Computer Graphics*, 27(11):4194–4203, 2021.
7. N. Deng, Z. He, J. Ye, B. Duinkharjav, P. Chakravarthula, X. Yang, and Q. Sun. Fovnerf: Foveated neural radiance fields for virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 28(11):3854–3864, 2022.
8. M. H. Eybposh, N. W. Cairra, M. Atisa, P. Chakravarthula, and N. C. Pégard. Deepcgh: 3d computer-generated holography using deep learning. *Optics Express*, 28(18):26636–26650, 2020.
9. J. E. Fröch, P. K. Chakravarthula, J. Sun, E. Tseng, S. Colburn, A. Zhan, F. Miller, A. Wirth-Singh, Q. A. Tanguy, Z. Han, et al. Beating bandwidth limits for large aperture broadband nano-optics. *arXiv preprint arXiv:2402.06824*, 2024.
10. A. Georgiou, J. Kollin, C. Hewitt, P. Chakravarthula, and B. Guenter. Visual perception of noise in a simulated holographic display—a user study. *Displays*, 76:102333, 2023.
11. Y. Peng, S. Choi, N. Padmanaban, and G. Wetzstein. Neural holography with camera-in-the-loop training. *ACM Transactions on Graphics (TOG)*, 39(6):185, 2020.
12. F. Schiffrers, P. Chakravarthula, N. Matsuda, G. Kuo, E. Tseng, D. Lanman, F. Heide, and O. Cossairt. Stochastic light field holography. In *2023 IEEE Inter-*

- national Conference on Computational Photography (ICCP)*, pages 1–12. IEEE, 2023.
13. L. Shi, B. Li, C. Kim, P. Kellnhofer, and W. Matusik. Towards real-time photorealistic 3d holography with deep neural networks. *Nature*, 591(7849):234–239, 2021.
  14. E. Tseng, G. Kuo, S.-H. Baek, N. Matsuda, A. Maimone, F. Schiffers, P. Chakravarthula, Q. Fu, W. Heidrich, D. Lanman, et al. Neural étendue expander for ultra-wide-angle high-fidelity holographic display. *Nature communications*, 15(1):2907, 2024.
  15. Y. Wang, P. Chakravarthula, Q. Sun, and B. Chen. Joint neural phase retrieval and compression for energy-and computation-efficient holography on the edge. *ACM Transactions on Graphics*, 41(4), 2022.
  16. Y. Wang, B. Chen, and P. Chakravarthula. Pupil-adaptive 3d holography beyond coherent depth-of-field. *arXiv preprint arXiv:2409.00028*, 2024.
  17. X. Xia, Y. Guan, A. State, P. Chakravarthula, T.-J. Cham, and H. Fuchs. Towards eyeglass-style holographic near-eye displays with statically expanded eyebox. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 312–319. IEEE, 2020.
  18. C. Zheng, G. Zhao, and P. T. So. Neural lithography: Close the design-to-manufacturing gap in computational optics with a'real2sim'learned photolithography simulator. *arXiv preprint arXiv:2309.17343*, 2023.
  19. M. Zhou, H. Zhang, S. Jiao, P. Chakravarthula, and Z. Geng. End-to-end compression-aware computer-generated holography. *Optics Express*, 31(26):43908–43919, 2023.

**Praneeth Chakravarthula** is an assistant professor at the University of North Carolina at Chapel Hill. At the intersection of optics, perception, graphics, optimization, and machine learning/AI, his research focus is on novel computational imaging and displays. Dr. Chakravarthula is a Senior Member of Optica, recipient of IEEE VR Best Dissertation Award, and several Best Paper and Demo Awards at premiere conferences like ACM SIGGRAPH, IEEE VR and ISMAR. He completed his postdoctoral research at Princeton University, obtained a PhD from UNC Chapel Hill, and B.Tech and M.Tech degrees from the Indian Institute of Technology Madras.

Contact him at [praneeth@unc.edu](mailto:praneeth@unc.edu)