

Optical Gaze Tracking with Spatially-Sparse Single-Pixel Detectors

Richard Li^{*1,2}, Eric Whitmire^{*1,2}, Michael Stengel², Ben Boudaoud², Jan Kautz², David Luebke², Shwetak Patel¹ and Kaan Akşit^{2,3}

¹ University of Washington

² NVIDIA Research

³ University College London



Figure 1: Gaze tracking, an essential component of next generation displays, needs to deliver several qualities such as accurate gaze estimation, low latency, small form-factor, low cost, low computational complexity, and a low power budget. To provide solutions for next generation displays, we demonstrate two single-pixel detector based gaze tracking prototypes. While the one shown on the left uses photodiodes and LEDs, the one shown in the middle uses only LEDs for both sensing and emitting light. As depicted on the right hand-side, we evaluate our gaze trackers with a series of subjective experiments.

ABSTRACT

Gaze tracking is an essential component of next generation displays for virtual reality and augmented reality applications. Traditional camera-based gaze trackers used in next generation displays are known to be lacking in one or multiple of the following metrics: power consumption, cost, computational complexity, estimation accuracy, latency, and form-factor. We propose the use of discrete photodiodes and light-emitting diodes (LEDs) as an alternative to traditional camera-based gaze tracking approaches while taking all of these metrics into consideration. We begin by developing a rendering-based simulation framework for understanding the relationship between light sources and a virtual model eyeball. Findings from this framework are used for the placement of LEDs and photodiodes. Our first prototype uses a neural network to obtain an average error rate of 2.67° at 400 Hz while demanding only 16 mW. By simplifying the implementation to using only LEDs, duplexed as light transceivers, and more minimal machine learning model, namely a light-weight supervised Gaussian process regression algorithm, we show that our second prototype is capable of an average error rate of 1.57° at 250 Hz using 800 mW.

Index Terms: Human-centered computing—Ubiquitous and mobile computing—Ubiquitous and mobile devices—Computer systems organization—Embedded and cyber-physical systems—Sensors and actuators

1 INTRODUCTION

Next generation displays [27] for virtual reality (VR) and augmented reality (AR) applications promise to improve our daily lives and routines. Gaze tracking is an essential and required component of these next generation displays, enhancing and enabling multiple methods and applications such as varifocal near-eye displays [18], foveated near-eye displays [25], super-resolution displays [3], and foveated computer graphics [52].

While gaze tracking has largely remained a research tool, we believe that several factors have hindered the deployability of gaze tracking systems: accuracy, latency, power consumption, cost, computational complexity, and form-factor. Improvements in gaze tracking hardware and software in one of these metrics often involves compromising other metrics. However, for gaze tracking technology to enable applications in next generation displays, such technology has to lead to a useful quality gaze tracker with a small form-factor, low latency, and low power consumption.

In this paper, we explore means of designing a useful-quality gaze tracker that accounts for all of these metrics. Toward this end, we investigate techniques for simplifying both the hardware and software components of a gaze tracking system. Concerned with the use of power and computationally demanding imaging sensors generally used for gaze tracking, we begin with a rendering-based simulation framework for exploring the possibility of decomposing cameras into individual single pixel sensors. Using findings from these simulations, we place pairs of photodiodes and LEDs around a human subject's eyes, modulating the light emission in a time multiplexed fashion and capturing the light reflected off of the eyes. We then use a fully connected neural network to process the recorded signals and estimate a gaze orientation, constituting our first standalone wearable gaze tracker prototype, NextGaze. To further minimize cost and increase flexibility in manufacturing, we remove the need for photodiodes in NextGaze by taking advantage

^{*}First and second authors contributed equally.

Table 1: A comparison of six systems across four dimensions. We position our work, the last two rows, as an unexplored middle ground across each of these dimensions.

Name	Modality	Rate	Error	Power
Scleral Coil	Magnetic	1000 Hz	0.1°	10+ W
SR Research	Camera	1000 Hz	0.33°	24 W
Pupil Labs	Camera	200 Hz	0.6°	1.5 mW
Tobii	Camera	100 Hz	0.5°	900 mW
LiGaze	Photodiode	128 Hz	6.1°	0.791 mW
CIDER	Photodiode	4 Hz	0.6°	7 mW
NextGaze	Photodiode	400 Hz	2.67°	16 mW
LED2Gaze	LED	250 Hz	1.57°	800 mW

of the bidirectional characteristics of LEDs to enable receiving and emitting light with a single component, reducing the number of components used and the number of signals processed. We show that this second prototype, LED2Gaze, can reduce the accuracy error by up to half using a Gaussian process regression (GPR) model. Our contributions are listed as the following:

1. A rendering-based framework for simulating gaze tracking devices with arbitrary single-pixel sensors with new insights into the behavior of eye reflections.
2. NextGaze, a wearable gaze tracking prototype equipped with photodiodes and LEDs. Our device obtains an average error of 2.67° at 400 Hz while consuming only 16 mW using a fully-connected neural network.
3. LED2Gaze, a standalone gaze tracking prototype that uses LEDs both for emitting and sensing light. Our device obtains an average error of 1.57° at 250 Hz consuming 800 mW using a lightweight GPR model.

2 RELATED WORK

We report and discuss the relevant literature on gaze tracking, and focus on three primary metrics: accuracy, sample rate, and power consumption.

2.1 Camera-Based Gaze Tracking

Video oculography is the most commonly used method for eye tracking. Most video-based eye trackers rely on infrared illumination of the eye and an infrared-sensitive video camera that detects either the location of the pupil or glints on the cornea. A calibration procedure is used to construct a mapping between glint/pupil locations and gaze orientation.

A high end commercial video eye tracking system, such as the SR Research EyeLink 1000 Plus [2], is capable of sampling at 1000 Hz with an average accuracy of 0.33°. More portable and affordable systems such as those produced by Tobii, SMI, and Pupil Labs operate at an order of magnitude lower sample rate, while maintaining a similar sub-degree accuracy. However, the power consumption of these devices is generally on the order of multiple watts [1].

In academic settings, the use of low resolution cameras for reducing the requirements of power consumption and computational resources needed for video oculography has seen promising results. Borsato et al. [6] was able to significantly reduce processing and power requirements by simply repurposing the optical flow sensor of a computer mouse for tracking the episcleral surface (the white part) of the eye. While they were able to obtain an error bound of 2.1°, the tracking was lost each time the user blinked, rendering it impractical for real-life use cases. Tonsen et al. [49] improved upon this accuracy by first simulating the different possible vantage points of cameras placed around a 3D model of the eye, eventually developing InvisibleEye, which leverages four millimeter-sized cameras of 5 x 5 pixels each to achieve a person-specific gaze estimation accuracy of 1.79°. Finally, Mayberry et al. [36] developed CIDER

to estimate the pupil location, but not the gaze orientation, hence we refrain from comparing their reported accuracy. However, we believe that CIDER was an exceptionally well-engineered system, with low power consumption (32 mW), high sampling rate (278 Hz), and detailed specifications given per component: camera, digitization, computation, and the near-infrared LEDs that were used for illumination.

In this paper, we describe a simulation framework similar to the one used by Tonsen et al. [49] for exploring the possibility of removing the need for a camera altogether by using discrete single pixel detectors instead. Ultimately, we use findings from these simulations to inform the design of a system that removes this significant component from CIDER’s bill of materials, saving power, cost, and computation.

2.2 Novel Sensing for Gaze Tracking

In addition to traditional camera-based gaze tracking techniques, a number of novel sensing techniques have been explored to leverage other properties of the eyes for tracking. On the high speed and invasive end, magnetic sensing has been employed to track scleral coils, wire coils embedded in a silicone ring that sits on the sclera of the eye [13, 45]. A voltage is induced in the wire coils when placed in a magnetic field, and that voltage is measured with thin physical connections to the coils [45]. This technique has been shown to offer a sampling rate of almost 10 kHz with an accuracy better than 0.1° [13]. This technique is generally used for lab-based studies and is not appropriate for consumer or mobile devices where power is a consideration.

At the other end of the spectrum, electrical sensing has been used to measure the voltage potential generated by the rotation of the eyes, which have an electric dipole between the cornea and retina, in a technique called electrooculography (EOG). Although these signals are generally not robust enough for continuous gaze tracking, they can be used to detect movements such as blinks as well as relative direction of movement. For example, Bulling and colleagues demonstrated wearable EOG goggles capable of detecting a set of eye gestures [9]. Furthermore, the eyeglasses company Jins produces a commercial product called the Meme with EOG electrodes embedded in the nose pad, which has been used to detect reading patterns [28], to measure fatigue [48], and to recognize facial expressions [46].

Finally, our work falls under the category of spatially sparse optical sensing. While the cameras described previously use image sensors with many pixels, sparse optical sensing approaches position single-pixel optical sensors sparsely around the region of interest. Such approaches have the primary potential advantage of requiring fewer pixel sensors, eliminating the capture of pixels that would be otherwise redundant and resulting in lower dimensional data that requires less computational power and bandwidth. In the case of head-mounted displays, sparse optical sensing also enables the possibility of moving the sensors out of the field of view for heightened immersion. For example, OLED-on-CMOS technology [54] has demonstrated the ability to capture images with single-pixel sensors placed in-line and alternating with single-pixel displays. While there is a recent ongoing effort to make a gaze tracker product at industry using single-pixel sensors with scanning micro electromechanical systems (MEMS) [57], usage of single-pixel sensors for gaze tracking largely remains in research. For greater sparsity, Topal et al. [50] developed EyeTouch, which consisted of only eight IR LED and sensor pairs around the lenses of a pair of glasses, but required the use of a bite bar or other means of stabilizing the head. LiGaze [33] showed that it was possible to use a solar cell indoors for harvesting the power needed to perform eye tracking using photodiodes, achieving a sample rate of 128 Hz with 6.1° accuracy and consuming 791 μ W. A sequel improved the power consumption to 395 μ W with a slightly reduced sample rate of 120 Hz [34].

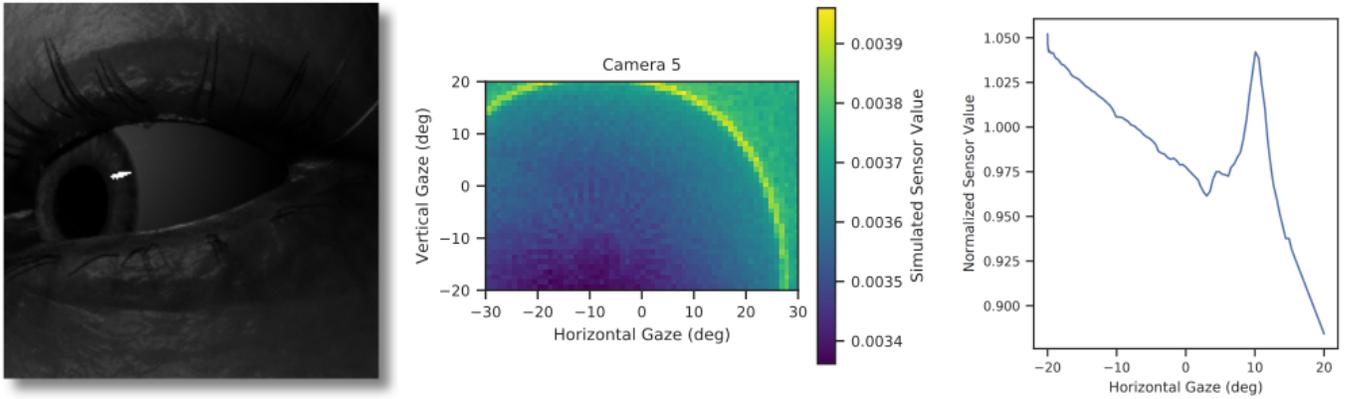


Figure 2: Left: one camera’s view of a simulated eye with corneal reflection. Center: the simulated output of a single-pixel detector as a function of gaze angle, obtained by rotating the model shown in the left panel. Right: real-life signal acquired by an actual photodiode as the eye rotated along the horizontal axis while staying constant vertically, verifying the distinct spike seen in the simulation shown in the center panel. The signal in the right panel matches one row of the center panel (horizontal sweep, vertical constant).

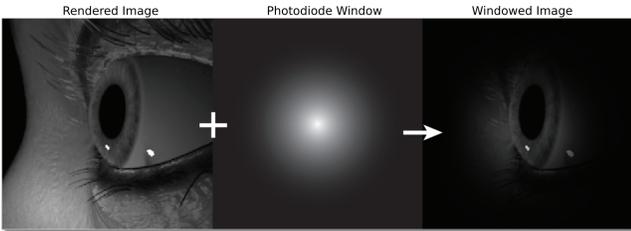


Figure 3: The simulation pipeline takes a rendered image from the point of view of a single-pixel detector, applies a windowing function to simulate the single-pixel detector lens, and sums the resulting image to simulate sensor output.

SynthesEyes [56] and UnityEyes [55] explored the possibility of using synthetic 3D models of human eyes for training models that can be applied to previously released, real-world datasets. Our work adopts a similar methodology of leveraging simulation findings for informing the design of our prototypes, with which we then implement and use to collect our own datasets for validation. Our functional prototypes gaze trackers are robust to head movement, removing the need for a bite bar, while improving the accuracy and sample rates given by LiGaze. We position our work at the intersection of high performance (accuracy and sample rate) and low power consumption.

3 SPARSE OPTICAL GAZE TRACKING

The idea of sparse optical gaze tracking is to use single pixel emitters and receivers, spread out in physical space, to detect the gaze orientation of the user’s eye. Depending on the orientation of the eye, the sensor will capture light directly reflected from the cornea (referred to as a “glint”) and scattered from the iris, sclera, and skin.

3.1 Rendering-Based Simulation Framework

In order to inform the design of our sparse optical gaze tracking systems, namely the placement of the optical sensors and emitters, we constructed a framework to simulate how a gaze tracker would perform under different configurations. This framework uses a realistic 3D model of a human face with a parametrically-defined eyeball from the work by Kim et al. [26]. The eyeball can be rotated to any horizontal and vertical gaze direction and the pupil size can be adjusted from 2 mm to 8 mm. Moreover, the top and bottom eyelids

can be adjusted from fully open to fully closed. The textures were adjusted to match the properties of skin under infrared illumination.

We place virtual cameras at the position and orientation of different proposed single-pixel sensor configurations following the guidance of Rigas et al. [44]. We use point light sources to simulate infrared emitters such as LEDs. For a given facial configuration (eye gaze in x and y , pupil size, eyelid position), we render an image from the viewpoint of each of the cameras. Because the sensitivity of a single-pixel sensor varies with the angle of incidence, we use a Gaussian distribution to represent sensitivity in accordance with the dataset of a sensor. For each image, we transform all pixels in the image using the following Gaussian window function:

$$g(x, y) = e^{-\frac{((x-x_0)^2 + (y-y_0)^2)}{2\sigma^2}},$$

where x and y represent the pixel coordinates, x_0 and y_0 represent the image centers, and σ is the standard deviation of the Gaussian distribution that represents angular sensitivity of a single-pixel sensor. Figure 3 summarizes the effect of this angular sensitivity of a single-pixel sensor. Finally, to simulate the accumulation of light at a single-pixel sensor, all pixels in the transformed image are summed as follows, where $i(x, y)$ represents a pixel from the original rendered image:

$$s = \sum_x \sum_y i(x, y) * g(x, y),$$

An important observation from this process is the importance of using 16-bit rendering and taking care not to saturate the image. The images in Figure 3 have been artificially brightened, for clarity, but note that the glints in these brightened images consist of saturated white pixels. If these images were used for the simulation, the signal due to the glint would be artificially weakened. To achieve a high-fidelity simulation, the simulated illumination must be decreased so that there are no clipped pixels from the direct reflections. Similar simulation techniques that either use rendered images [43, 44, 58] or images captured from a camera [22] could be prone to this issue.

A second observation concerns the interaction of the glints with the edge of the cornea as the eye changes concavity. As the eye moves and a glint approaches the edge of the cornea, the glint becomes stretched and the received signal strength at the sensor increases as depicted in Figure 2. The eye orientation at which this effect occurs depends on the position of both an emitter and a single-pixel sensor. Figure 2 shows the output of a single simulated single-pixel detector as a function of gaze directions, where the

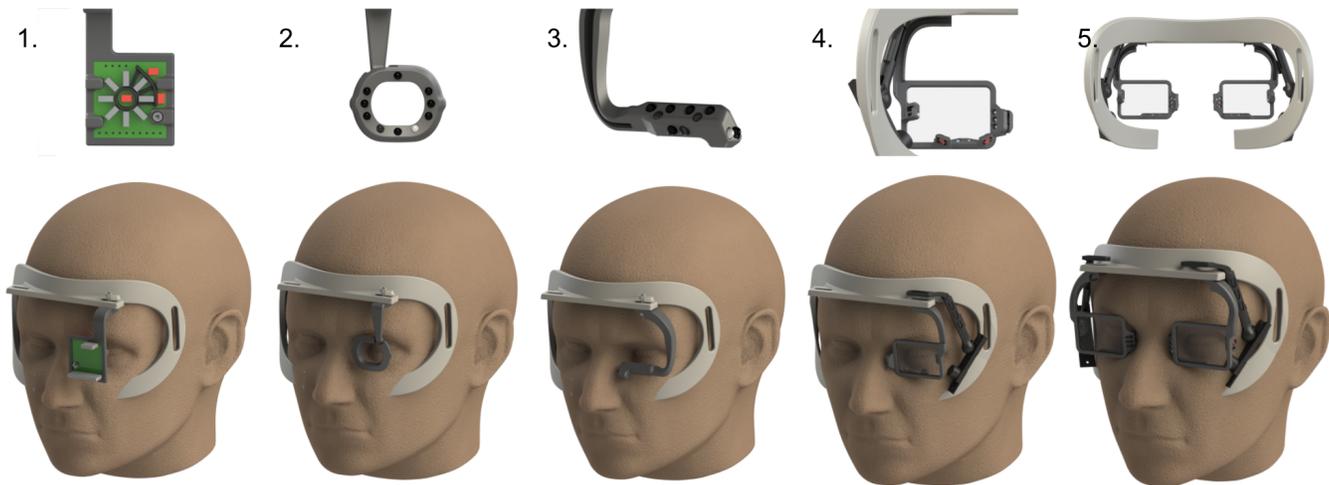


Figure 4: Examples of intermediate steps in our iterative co-design process between simulation and implementation. From left to right: 1. Sensors placed directly in front of the eye, 2. sensors placed around the eye enabling looking through the lens, 3. increasing the field of view by experimentally moving all the sensors to the bottom, 4. adding a camera and some sensors along the vertical axis, and 5. streamlining the design.

relatively smooth gradient along the lower left portion of the image is due to the pupil and glint moving within the sensitive region. The circular edge along the outsides of the image corresponds to gaze locations where the glint is positioned at the edge of the corneal bump. Within this bright edge, the glint is positioned on the cornea; outside of this edge, the glint is located on the sclera.

Center portion of Figure 2 shows an artificially brightened rendering corresponding to the gaze orientation of 27° horizontal and -20° vertical. The smearing of the glint along the edge of the cornea causes more light to enter the single-pixel detector. Compare these simulated results to actual collected data as the eye pursues a target along a horizontal line from -20° to 20° as shown in Figure 2 (right). We hypothesize that the spike around 10° corresponds to the glint aligning with the edge of the cornea. This effect has not been demonstrated or accounted for in prior work that relies on simulated data [24, 43, 44, 58]. A practical implication of this result is that it is best to minimize the number of point light sources active at a time. Multiple light sources will result in multiple glints that only complicate the tracking problem. Fewer light sources would maximize the smoothness of the transfer function. On the other hand, future work could consider using these direct reflections as a feature to model the geometry of the eye.

In addition to these two findings, a co-design process between simulated experiments and implementation iterations revealed that real life human faces varied so significantly that determining specific LED locations based on a simulation was not practical. We show some of the intermediate designs in Figure 4, and this iterative process helped reveal a number of design guidelines for working with spatially-sparse detectors:

- As mentioned previously, minimize the number of emitters on simultaneously to avoid saturated receivers.
- Co-locating emitters and receivers generally provides the best signal.
- For maximum information gain with sparse sensors, the more diversity in perspective the better.
- However, staying below the eye is recommended to avoid eyelash interference.

We used these guidelines as we iterated through a number of designs, two of which we have selected to describe in this paper.

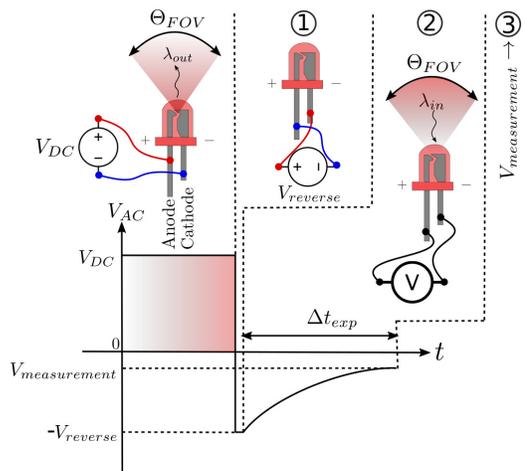


Figure 5: Different LED modes: (1) applying a forward voltage of V_{DC} , in which the LED emits light with a wavelength of λ_{out} and an emission cone angle, Θ_{FOV} ; (2) applying a reverse voltage pulse, $V_{reverse}$, for a short time duration, discharging LED with incoming light that has a wavelength of λ_{in} for a specific time, Δt_{exp} , with an reception cone angle of Θ_{FOV} ; and (3) measuring a voltage, $V_{measurement}$, from the LED.

3.2 Single-Pixel Sensing Hardware

In order to leverage the findings from the simulation framework for our hardware implementation, we consider the use of LED and photodiode pairs and duplexed LEDs as single-pixel sensors. In the LED-photodiode pairs case, the LEDs are used to illuminate the eye and may be modulated to provide a more strategic image of the eye. LEDs with infrared light emission are typically used in gaze tracking hardware for NEDs, since humans are insensitive to IR illumination [14]. A human eye's cornea has similar absorption and reflection characteristics in the near IR spectrum as in visible light [15]. Furthermore, IR LEDs have a narrow bandwidth (typically around 50 nm), avoiding cross-talk with other wavelengths. The photodiodes are used to capture signals related to gaze direction.

In the latter case, we also propose the use of LEDs to both illuminate the eye and capture light. LEDs provide illumination when a

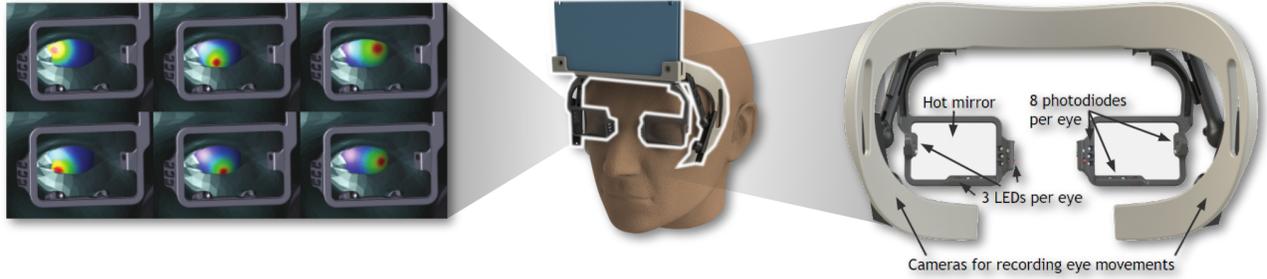


Figure 6: System architecture of NextGaze. Left: a Zemax optical simulation of six photodiode placements and their projected regions of sensitivity. Right: the arrangement of optical and electrical components within the prototype, configured according to findings from the simulation.

forward voltage is applied to their two electrical terminals. However, LEDs can also act as photodetectors [17]. This duplexing can be accomplished with three steps that are depicted in Figure 5. Typically, LEDs are most sensitive to wavelengths λ_{in} that are shorter than their emission spectrum (so $\lambda_{in} < \lambda_{out}$) [30]. Thus, larger exposure times are required if LEDs with the same emission spectrum are used. To achieve the lowest possible latency with a given configuration, we select different LEDs that have intersecting emission and sensing spectra in the IR range.

Eye safety is a very important aspect when a user is exposed to infrared radiation; Δt_{exp} and maximum irradiance of an LED must be considered according to safety regulations for infrared light sources. In our implementation, we followed a commonly accepted guideline [7], accordingly.

4 PROTOTYPES

Based on learnings from the simulations and described sensing approaches, we developed our first prototype, NextGaze. Then, by simplifying both the software and hardware components, we arrived at our second prototype, LED2Gaze, which demonstrated improved accuracy while using fewer physical hardware components and a simplified estimation algorithm.

4.1 Gaze tracking photodiodes

Our first system consists of a ring of LEDs and photodiodes around each eye. The system is designed to be used as a standalone gaze tracker with an external display, therefore it was constructed as part of a face mask. The mask attaches to the face using an elastic band. An external display is used to calibrate the gaze tracker prototype.

The full design is shown in Figure 6. The eye is illuminated by three LEDs embedded within the frame of the device. Two or three photodiodes are clustered around each LED. In total, there are eight photodiodes, placed strategically such that for a typical user, they will cover different parts of the eye. Figure 6 shows the results of a Zemax optical simulation for a subset of six photodiodes. The images highlight the intended sensitive region for each photodiode.

To facilitate development and proper placement of the optical elements, a Pupil Labs infrared camera is placed just outside the user’s field of view. A hot mirror in front of the eyes reflects infrared light from the eyes into the camera lens. The camera is used only for debugging purposes and is not part of the final sensing pipeline. No experiments were conducted using the camera since we can compare against the baselines already reported by camera-based gaze trackers as discussed in section 2.1, or specifically the specifications of the Pupil Labs camera [23].

The LEDs and photodiodes selected for our prototype are both optimized for 940 nm infrared light, which rejects most ambient illumination at other frequencies. In order to further improve signal robustness, modulated illumination is used to reject any ambient

light, even environmental 940 nm light (i.e. sunlight). The device uses an analog front-end (ADPD103) to synchronously modulate the LEDs and sample the photodiodes. Each LED is illuminated for 3 μ seconds every 24 μ seconds. The photodiode response is bandpass filtered and synchronously integrated. The result is a photodiode signal sensitive to changes in reflected light from the LEDs, but not from other light sources.

To minimize the effect of direct reflections of the cornea, only one LED is illuminated at a time. In this implementation, only two LEDs are used in a particular session; each LED is associated with the four nearest photodiodes. For a single frame of data, the first LED pulses four times while the four nearest photodiodes are integrated and sampled. This process repeats for the second LED and remaining photodiodes. The overall data rate is determined by the number of LED pulses and delay between LEDs. In this prototype, the output data rate was configured to be 400 Hz. The electrical current consumed by the LED is also configurable and determines the signal-to-noise ratio. In this prototype, it was set to 67 mA. Note that this is instantaneous current through the LEDs. At this current setting, the overall power consumption of the analog front-end and LEDs is only 16 mW.

4.2 Gaze tracking LEDs

Our second prototype consists of 6 LEDs per eye, with each LED functioning as both light sources and sensors, and an Arduino Nano microcontroller per eye for controlling those LEDs. We designed our second prototype to be a standalone platform to remove the need for a chin rest in our subjective experimentation, so we inserted a 2K resolution HDMI display (TopFoisson) into an off-the-shelf mobile VR headset (Samsung Gear VR), in which the sensing electronics were housed. LEDs are placed around the VR headset’s magnifier lenses.

The anodes of LEDs were attached to digital IO pins of a microcontroller, while their cathodes were attached to analog-to-digital converter (ADC) pins of the same microcontroller. Each time an LED is to be used in sensor mode, it follows the three steps described in Figure 5. LEDs have a soft-coded mechanism that adjusts exposure times, Δt_{exp} , on a per LED basis, so that saturation caused by varying light conditions can be avoided for each LED. The LEDs are sampled in a round-robin fashion, such that one LED records measurements while the remaining LEDs serve as emitters. In contrast to NextGaze, we chose to use the remaining LEDs as emitters to minimize the sensing time of an LED, and in turn minimize system latency.

The microcontroller communicates with the user interface application over a USB connection. This user interface application is developed using Kivy library [53]. Our user interface application handles a number of tasks: (1) collecting measurements from each LED by requesting them from the two microcontrollers used for

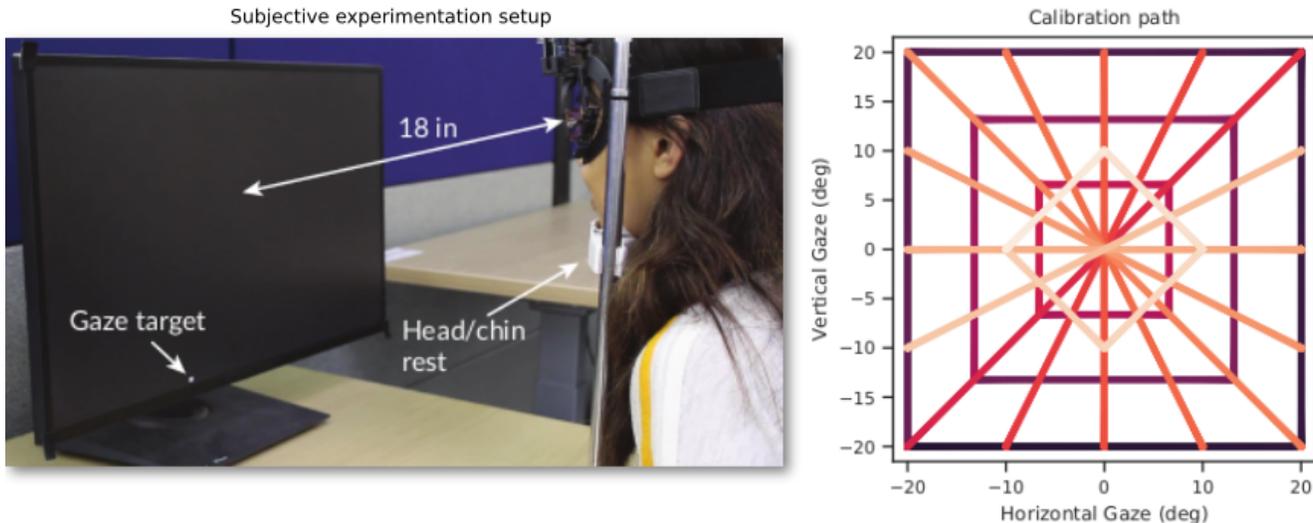


Figure 7: The calibration process requires a user to wear the prototype and follow a moving stimulus shown on an external computer screen with their gaze. Left: the user rests their chin in a head rest 18 inches from a fixed display. Right: the user is asked to follow a moving target that moves along multiple line segments with their gaze.

each eye, (2) updating the user interface, (3) estimating the gaze orientation, and (4) keeping logs related to captured data such as event timestamps, raw signals, estimations.

5 EVALUATION

5.1 Evaluating NextGaze

5.1.1 Procedure

We invited six participants (4 male, 2 female) to help us collect a dataset for evaluating the performance of NextGaze. Due to NextGaze’s design as a standalone device with no display, an external monitor was used to show the visual stimulus. Participants that wore glasses were asked to remove them, however, contact lenses were allowed to be worn. Participants were then asked to place their head in a desk mounted chin rest, putting his or her face 18 inches away from the monitor, as shown in Figure 7 (left). A circular gaze target with a radius of 10 pixels (1.2 mm) is placed on the screen for calibration. The x and y coordinates of the target in screen-space can be computed from the screen distance and desired angular gaze coordinates.

Participants first engaged in a smooth pursuit task, following the guidelines outlined by Pfeuffer et al. [41]. The gaze target travels in a series of linear paths over $\pm 20^\circ$ vertical and horizontal field of view as outlined in Figure 7 (right). For each segment, the target smoothly accelerates from rest over 1 second, up to a max speed of 6° per second and then decelerates over 1 second back to rest at the end of the segment. This produces a dense sample of the gaze space. Particular care was taken to optimize the display pipeline to minimize any jitter of the target as it moved, as this could cause undesired saccades during the pursuit.

Following the smooth pursuit task, participants were then asked to visually fixate on 25 static targets displayed on a grid within $\pm 20^\circ$ horizontal and vertical. The saccadic nature of this task helped to diversify the dataset collected. The first twenty segments of the smooth pursuit data were used for training. The last four segments as well as the 25 fixation targets were used for testing.

For evaluation, we report error on a per-frame basis, taking the difference in angle between the target’s position and the angle predicted by our modeling algorithm at every frame. Although this protocol is limited by the strong assumption that the participant’s

Table 2: Performance of the gaze model on different test sets. The model performs best on the smooth pursuit task.

Task	Mean Error	Standard Deviation
Smooth Pursuit Validation	1.68°	0.56°
Fixation Validation ($\pm 20^\circ$)	2.67°	0.98°
Central Fixation Validation ($\pm 10^\circ$)	2.35°	0.58°

gaze really is locked on the target, it is nonetheless standard procedure and we adopted it to produce results that can be compared with the related work.

5.1.2 Modeling and Results

First, a preprocessing step removes blinks from the signal by performing a Savitzky-Golay filter and removing areas where the filtered signal is different by more than a given threshold. The remaining signal is then downsampled to 100 points along each line segment to avoid overfitting to points along the line.

The gaze model maps the photodiode output to gaze coordinates in degrees. Because the mapping is highly nonlinear, we leverage a neural network model to map the eight photodiode signals to 2D gaze. We first scale all photodiode outputs to 0 to 1 and then apply Principal Component Analysis (PCA) to reproject the data. The transformed signals are input to a network with 4 hidden layers of 64 nodes each and tanh activation functions. The model uses scikit-learn’s implementation of a multi-layer perceptron neural network with a batch size of 4, up to 500 epochs, and an adaptive learning rate starting at 0.001, as depicted in Figure 8.

After calibration, the gaze is estimated and shown on screen in real time. Inferences from the model are post-processed using an exponential weighted moving average (EWMA) filter ($\alpha = 0.2$) to smooth the estimation results. This filter was designed empirically to improve the real-time experience provided to the user.

In addition to using the real-time inferences for supporting the interactive system, we also ran an offline evaluation producing the results shown in Table 2. Mean error on the last four segments of the smooth pursuit task was 1.68°. The error increased to 2.67° on the

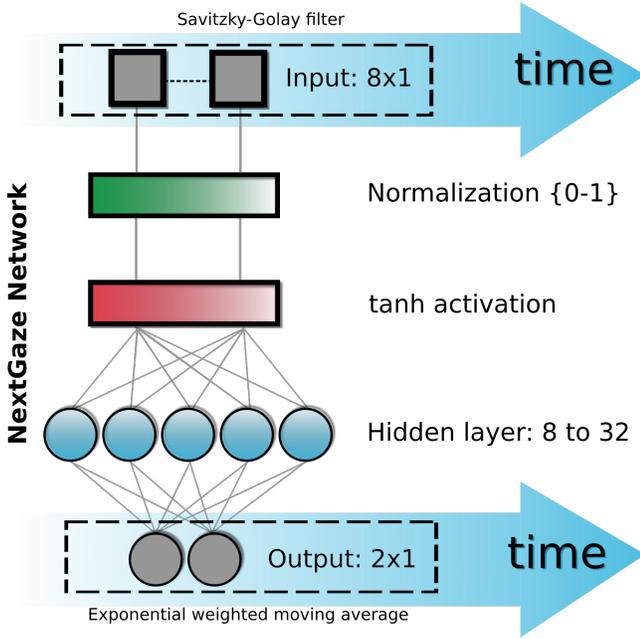


Figure 8: NextGaze Network. Measurements from photodiodes are passed through a normalization and activation step before going through a full connected neural network. Result is a two-dimensional output that represents gaze estimation along vertical and horizontal axes. To avoid temporal flickering or sudden variations in the results over the time, inputs of NextGaze network is filtered temporally using a Savitzky-Golay filter, while outputs of NextGaze network is temporally averaged using exponential weights.

fixation task, but is slightly better when limiting the data to $\pm 10^\circ$. Among the six trials, the best error on the $\pm 20^\circ$ fixation task was 1.1° and the worst was 4.2° .

5.2 Evaluating LED2Gaze

5.2.1 Procedure

We invited fourteen participants (10 male, 4 female) to help us collect a dataset for evaluating LED2Gaze. Since LED2Gaze featured its own built-in display, no chin rest was necessary for this study. Again, glasses were asked to be removed while contact lenses were allowed to be used. Participants were seated in an office chair and asked to wear the headset. Then they are engaged in a series of tasks similar to the procedure used in Study 1, again driven by a moving target. First, participants performed a smooth pursuit task consisting of 9 segments. The target guided the participant through two sessions of fixating on each of 16 targets forming a grid over the entire field of view of the headset (101°). Finally, 66 random points were presented for the participant to fixate on. In total, 9 segments of smooth pursuit and 98 saccadic fixation points were collected per participant. For evaluation, the dataset was split such that the smooth pursuit, one session of 16 grid points, and 66 random points were used for training the model, and the second session of 16 grid were used for testing.

5.2.2 Modeling and Results

In our previous prototype, we employed a fully-connected neural network to address the nonlinearity of the mapping from photodiode readings to gaze orientation. While the results were adequate, the complexity and abstract nature of neural networks can be tricky to interpret for humans and in turn difficult to improve the results of. For our second prototype, we explored the use of a simpler, more

Table 3: LED2Gaze’s performance on the fourteen participant dataset.

Participant	Mean Error	Median Error	Standard Deviation
Average	1.57°	1.12°	2.00°

interpretable approach, namely a GPR model. Such models take the following general form:

$$\begin{bmatrix} e_x \\ e_y \end{bmatrix} = k^T C^{-1} \begin{bmatrix} u_x \\ u_y \end{bmatrix}^T$$

$$k = \begin{bmatrix} \kappa(s(t), \bar{c}_1) \\ \dots \\ \kappa(s(t), \bar{c}_p) \end{bmatrix}$$

$$C = \begin{bmatrix} \kappa(\bar{c}_0, \bar{c}_0) & \dots & \kappa(\bar{c}_0, \bar{c}_p) \\ \dots & \dots & \dots \\ \kappa(\bar{c}_p, \bar{c}_0) & \dots & \kappa(\bar{c}_p, \bar{c}_p) \end{bmatrix}$$

Where e_x and e_y represents estimated gaze orientation along the x and y axes, respectively, k^T represents a vector that contains the similarity measures between the captured $s(t)$, and the calibration vectors \bar{c}_p . Finally, u_x and u_y represent vectors that correspond to the x and y position of each \bar{c}_p .

Comparing a vector with another vector can be accomplished in multiple ways. In evaluating multiple different distance measures (Cosine, Minkowski, Manhattan, Canberra) [8, 29, 42], we found that the Minkowski distance measure to be the most effective when used with the GPR algorithm.

The results from this evaluation are shown in Table 3. Mean error was improved to 1.57° when compared to our previous prototype. Among the fourteen participants, the best per-participant error was 1.1° (equal to the best performance of NextGaze) and the worst was 2.1° (improving the worst performance of NextGaze by half). Furthermore, we empirically found that the output was sufficiently smooth such that no post-processing such as the EWMA filter used by NextGaze was needed.

6 DISCUSSION

6.1 Towards Deploying Sparse Optical Gaze Trackers

We discuss some practical considerations needed for spatially-sparse optical gaze trackers to be deployed in devices to be used in natural conditions.

User and Session Dependence. In this work, our modeling and analysis reported per-session and per-user accuracy results. On the other hand, in practice, gaze trackers are ideally session and user independent, such that the system can be put on and immediately used. Further investigations are required to understand how much, if any, per-user or per-session calibration is required. For example, it might be only necessary that the user follows a calibration procedure the very first time they use the system (i.e. user registration), or briefly calibrate the system with a few points (i.e. session registration).

Synthetic Data Although our experiments only used real sensor data collected from participants for training and testing, there is an opportunity for making use of the synthetic data generated by the simulation framework to produce more robust models. For example, the synthetic data could be used in a Bayesian framework for augmenting the signals recorded from participants, increasing the size of the training set [51]. Alternatively, recent advances in neural network research have shown the potential for directly using synthetic data in the training procedure for more robust models. [47]

Wearable Accessories. Factors such as prescription glasses, contact lenses, eye color, eyelashes, and mascara influence data quality [39]. A common practice for avoiding usage of prescription glasses and contact lenses in VR/AR near-eye displays comes in the form of an add-on inset lens in the optical path of a near-eye display. On the other hand, next generation computational displays promises algorithmic approaches to the problem of eye prescriptions by using active components that can support various focus levels [11].

In our studies, participants were asked to remove prescription glasses. However, there were subjects that wore contact lenses and mascara during experiments. Contact lenses are known to form air bubbles in between the cornea of an eye and a contact lens resulting in unintended dynamic reflections, and mascara can create a false glint, introducing noise into our sensor signals and causing robustness issues in camera based gaze tracking hardware [39]. Although our experiments did not reveal any particular robustness issues against wearing contact lenses or mascara, we also did not specifically control for it, and this issue remains not only an open research question but also a major challenge for deploying gaze tracking systems.

Compensating for Slippage. In addition to per-session and per-user calibration, wearable devices also often face the additional challenge of within-session changes. Sparse optical gaze trackers are particularly sensitive to slippage of the device. Small millimeter-level shifts can cause significant changes in the observed sensor values. We conducted some initial experiments to explore the feasibility of compensating for slippage. We added four additional photodiodes in the device oriented toward the nose bridge and side of the face to collect signals corresponding to the position of the device on the face. Preliminary results suggest that a monolithic model that incorporates both eye-tracking sensors and face-tracking sensors may be capable of gaze tracking that is invariant to sensor shifts. Further investigation and additional training data is needed to demonstrate a full implementation.

6.2 Opportunities in Gaze-Based Interfaces

Gaze tracking has the potential to enable significant advances in interacting with mixed reality. While gaze tracking research has traditionally optimized for greater tracking accuracy, we suggest that other facets of gaze tracking can be just as important to the user experience. This subsection describes human perception and device interaction opportunities for gaze-based interfaces that span the spectrum of power, speed and latency, and accuracy requirements.

6.2.1 Human Perception

Virtual Social Avatars. Mixed reality offers the possibility of immersive telepresence through virtual avatars, or digital representations of oneself. Literature in psychology has shown that the eyes convey a significant amount of information in the interaction between two people. To improve the quality and immersiveness of telepresence interactions, gaze tracking is needed to drive the avatar [20]. As social cues, it is important to achieve a high sample rate and low latency for the animated avatar to seem responsive [35, 37]. In addition, low power consumption is needed as the networking requirements of a video call already consume significant amounts of power.

Foveated Rendering. Human eyes have maximum visual acuity in the fovea, a region in the retina of the eyes. Areas outside of the fovea are perceived with less clarity. Research in HMDs has explored the concept of “foveated rendering”, in which only the region the user is visually attending to is rendered with full quality, and has shown significant savings in computational requirements [21, 40]. However, foveated rendering requires understanding the gaze direction of the eye to begin with. As a technique for reducing computational requirements, it is natural that the sensing technique it relies on

should similarly use low computation and low power. Similar to animating avatars, the latency required of gaze tracking for foveated rendering needs to be low (less than 50 ms) [4].

6.2.2 Device Interaction

Activity Recognition. Related work has shown that both camera-based and EOG-based gaze trackers can be used for recognizing activities of daily living, such as detecting reading and counting how many words have been read [28]. Such signals can be indicators of visual attention, passively disabling distractions such as notifications or automatically turning pages. The use of EOG glasses in this related work exemplifies the notion that high accuracy is not needed to create a useful interface.

User Interfaces. With clever design of an user interface, varying degrees of gaze tracking error can be useful and effective. Bubble Gaze Cursor and Bubble Gaze Lens lessens the required gaze tracking accuracy by implementing an area cursor with a magnifying glass feature, essentially dynamically increasing the effective selection region [12]. Orbits further reduces the accuracy needed by presenting different moving visual stimuli, and simply confirming selection by measuring correlation between the gaze tracking results and the movements of the stimuli [19]. A similar system, implemented using EOG glasses, show how the lower accuracy requirements can also alleviate the power consumption of a system with an effective interface [16]. Naugle and Hoskinson [38] demonstrated that the coarsest gaze tracker, only recognizing whether a user is visually attending to a display or not, can be used to quickly interact with a head-mounted display while saving up to 75% of the head-mounted display’s normal power consumption.

6.3 Future Work

Exploring Multimodal Sensing. In this paper, we have explored the use of photodiodes and reverse-driven LEDs for sensing gaze as an alternative to camera-based approaches. While NextGaze featured two cameras, and videos were recorded, that data was never used for gaze inference. In the future, we plan to explore how sensor fusion might help leverage the camera data in conjunction with the signals from our single-pixel detectors. For example, the camera data might be used for periodic self-calibration, or as a higher accuracy fall-back when needed. We are also interested in exploring how the single-pixel signals can be used to fill in the gaps between camera frames, such as a way of informing an interpolation function.

In addition to cameras, other sensors could potentially be used in tandem with our single-pixel detectors. For example, strain gauges have been previously used to measure facial movements [31], and electrode arrays have been used to capture a variety of biosignals from the face [5].

Beyond Gaze Tracking. To make a compelling case of including a given sensor in future devices, such a technique should be able to serve multiple purposes. There has been prior work using single-pixel detectors for facial action tracking and recognition [32] and for vital sign monitoring, such as heart-rate and blood oxygenation [10]. We will explore such opportunities with our technique beyond gaze tracking in the future.

7 CONCLUSION

In this paper, we explore the design space of gaze trackers that leverage sparse single-pixel optical sensing techniques. Our rendering-based simulation framework enables accurate and rapid exploration of this design space. We present two wearable gaze tracking devices built on these techniques with designs grounded in insights gained from simulation results. NextGaze explores the use of infrared LEDs and photodiodes to estimate gaze in a low-power wearable device, while LED2Gaze builds on these ideas and introduces a path to further form-factor improvements by leveraging LEDs as both emitters

and sensors. These prototypes demonstrate the feasibility of using discrete optical elements to realize high-speed, low-power gaze tracking devices suitable for wearable use in virtual and augmented reality devices.

REFERENCES

- [1] Tobii pro glasses 2 user manual.
- [2] Eyelink 1000 plus - the most flexible eye tracker, Apr 2020.
- [3] K. Akşit. Patch scanning displays: spatiotemporal enhancement for displays. *Optics Express*, 28(2):2107–2121, 2020.
- [4] R. Albert, A. Patney, D. Luebke, and J. Kim. Latency requirements for foveated rendering in virtual reality. *ACM Transactions on Applied Perception (TAP)*, 14(4):1–13, 2017.
- [5] G. Bernal, T. Yang, A. Jain, and P. Maes. Physiohmd: a conformable, modular toolkit for collecting physiological data from head-mounted displays. In *Proceedings of the 2018 ACM International Symposium on Wearable Computers*, pp. 160–167, 2018.
- [6] F. H. Borsato and C. H. Morimoto. Episcleral surface tracking: challenges and possibilities for using mice sensors for wearable eye tracking. In *Proceedings of the Ninth Biennial ACM Symposium on Eye Tracking Research & Applications*, pp. 39–46, 2016.
- [7] A. Boucouvalas. IEC 825-1 eye safety classification of some consumer electronic products. *IEE Colloquium on Optical Free Space Communication Links*, pp. 13:1–13:6, 1996.
- [8] J. R. Bray and J. T. Curtis. An ordination of the upland forest communities of southern Wisconsin. *Ecological Monographs*, 27(4):325–349, 1957.
- [9] A. Bulling, D. Roggen, and G. Tröster. Wearable eog goggles: Seamless sensing and context-awareness in everyday environments. *Journal of Ambient Intelligence and Smart Environments*, 1(2):157–171, 2009.
- [10] G. Cennini, J. Arguel, K. Akşit, and A. van Leest. Heart rate monitoring via remote photoplethysmography with motion artifacts reduction. *Optics express*, 18(5):4867–4875, 2010.
- [11] P. Chakravarthula, D. Dunn, K. Akşit, and H. Fuchs. Focusar: Auto-focus augmented reality eyeglasses for both real world and virtual imagery. *IEEE transactions on visualization and computer graphics*, 24(11):2906–2916, 2018.
- [12] M. Choi, D. Sakamoto, and T. Ono. Bubble gaze cursor+ bubble gaze lens: Applying area cursor technique to eye-gaze interface. In *ACM Symposium on Eye Tracking Research and Applications*, pp. 1–10, 2020.
- [13] H. Collewijn, F. Van der Mark, and T. Jansen. Precise recording of human eye movements. *Vision research*, 1975.
- [14] H. Dai, S. Song, X. Zeng, S. Su, M. Lin, and M. Q.-H. Meng. 6-d electromagnetic tracking approach using uniaxial transmitting coil and tri-axial magneto-resistive sensor. *IEEE Sensors Journal*, 18(3):1178–1186, 2017.
- [15] G. Damian, T. Delbruck, and P. Lichtsteiner. Eye tracking using event-based silicon retina.
- [16] M. Dhuliawala, J. Lee, J. Shimizu, A. Bulling, K. Kunze, T. Starner, and W. Woo. Smooth eye movement interaction using eog glasses. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, pp. 307–311, 2016.
- [17] P. Dietz, W. Yerazunis, and D. Leigh. Very low-cost sensing and communication using bidirectional LEDs. In *UbiComp*, pp. 175–191, 2003.
- [18] D. Dunn, C. Tippets, K. Torell, P. Kellnhöfer, K. Akşit, P. Didyk, K. Myszkowski, D. Luebke, and H. Fuchs. Wide field of view varifocal near-eye display using see-through deformable membrane mirrors. *IEEE transactions on visualization and computer graphics*, 23(4):1322–1331, 2017.
- [19] A. Esteves, E. Velloso, A. Bulling, and H. Gellersen. Orbits: Gaze interaction for smart watches using smooth pursuit eye movements. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, pp. 457–466, 2015.
- [20] M. Garau, M. Slater, S. Bee, and M. A. Sasse. The impact of eye gaze on communication using humanoid avatars. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 309–316, 2001.
- [21] B. Guenter, M. Finch, S. Drucker, D. Tan, and J. Snyder. Foveated 3d graphics. *ACM Transactions on Graphics (TOG)*, 31(6):1–10, 2012.
- [22] K. Irie, B. A. Wilson, R. D. Jones, P. J. Bones, and T. J. Anderson. A laser-based eye-tracking system. *Behavior Research Methods, Instruments, & Computers*, 34(4):561–572, 2002.
- [23] M. Kassner, W. Patera, and A. Bulling. Pupil: An open source platform for pervasive eye tracking and mobile gaze-based interaction. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication, UbiComp '14 Adjunct*, p. 1151–1160. Association for Computing Machinery, New York, NY, USA, 2014. doi: 10.1145/2638728.2641695
- [24] D. Katrychuk, H. K. Griffith, and O. V. Komogortsev. Power-efficient and shift-robust eye-tracking sensor for portable vr headsets. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, pp. 1–8, 2019.
- [25] J. Kim, Y. Jeong, M. Stengel, K. Akşit, R. Albert, B. Boudaoud, T. Greer, J. Kim, W. Lopes, Z. Majercik, et al. Foveated ar: dynamically-foveated augmented reality display. *ACM Transactions on Graphics (TOG)*, 38(4):1–15, 2019.
- [26] J. Kim, M. Stengel, A. Majercik, S. De Mello, D. Dunn, S. Laine, M. McGuire, and D. Luebke. Nvgaze: An anatomically-informed dataset for low-latency, near-eye gaze estimation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2019.
- [27] G. A. Koulieris, K. Akşit, M. Stengel, R. K. Mantiuk, K. Mania, and C. Richardt. Near-eye display and tracking technologies for virtual and augmented reality. In *Computer Graphics Forum*, vol. 38, pp. 493–519. Wiley Online Library, 2019.
- [28] K. Kunze, K. Masai, M. Inami, O. Sacakli, M. Liwicki, A. Dengel, S. Ishimaru, and K. Kise. Quantifying reading habits: Counting how many words you read. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '15*, p. 87–96. Association for Computing Machinery, New York, NY, USA, 2015. doi: 10.1145/2750858.2804278
- [29] G. Lance and W. Williams. Computer programs for hierarchical polythetic classification (“similarity analyses”). *The Computer Journal*, 9(1):60–64, 1966.
- [30] V. Lange, F. Lima, and D. Kühlke. Multicolour led in luminescence sensing application. *Sensors and Actuators A: Physical*, 169(1):43–48, 2011.
- [31] H. Li, L. Trutoiu, K. Olszewski, L. Wei, T. Trutna, P.-L. Hsieh, A. Nicholls, and C. Ma. Facial performance sensing head-mounted display. *ACM Transactions on Graphics (ToG)*, 34(4):1–9, 2015.
- [32] R. Li and G. Reyes. Buccal: Low-cost cheek sensing for inferring continuous jaw motion in mobile virtual reality. In *Proceedings of the 2018 ACM International Symposium on Wearable Computers, ISWC '18*, p. 180–183. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3267242.3267265
- [33] T. Li, Q. Liu, and X. Zhou. Ultra-low power gaze tracking for virtual reality. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*, pp. 1–14, 2017.
- [34] T. Li and X. Zhou. Battery-free eye tracker on glasses. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pp. 67–82, 2018.
- [35] T. Louis, J. Troccaz, A. Rochet-Capellan, and F. Bérard. Is it real? measuring the effect of resolution, latency, frame rate and jitter on the presence of virtual entities. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*, pp. 5–16, 2019.
- [36] A. Mayberry, Y. Tun, P. Hu, D. Smith-Freedman, B. Marlin, C. Salt-house, and D. Ganesan. Cider: enhancing the performance of computational eyeglasses. In *Proceedings of the Ninth Biennial ACM Symposium on Eye Tracking Research & Applications*, pp. 313–314, 2016.
- [37] M. Nabiyouni, S. Scerbo, D. A. Bowman, and T. Höllerer. Relative effects of real-world and virtual-world latency on an augmented reality training task: an ar simulation experiment. *Frontiers in ICT*, 3:34, 2017.
- [38] E. Naugle and R. Hoskinson. Two gaze-detection methods for power reduction in near-to eye displays for wearable computing. In *Int. Conf.*

- on *Wireless and Mobile Comp., Net. and Comm.*, pp. 675–680, 2013.
- [39] M. Nyström, R. Andersson, K. Holmqvist, and J. Van De Weijer. The influence of calibration method and eye physiology on eyetracking data quality. *Behavior research methods*, 45(1):272–288, 2013.
- [40] A. Patney, J. Kim, M. Salvi, A. Kaplanyan, C. Wyman, N. Benty, A. Lefohn, and D. Luebke. Perceptually-based foveated virtual reality. In *ACM SIGGRAPH 2016 Emerging Technologies*, pp. 1–2, 2016.
- [41] K. Pfeuffer, M. Vidal, J. Turner, A. Bulling, and H. Gellersen. Pursuit calibration: Making gaze calibration less tedious and more flexible. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*, pp. 261–270, 2013.
- [42] C. E. Rasmussen. *Gaussian processes for machine learning*. MIT Press, 2006.
- [43] I. Rigas, H. Raffle, and O. V. Komogortsev. Hybrid ps-v technique: A novel sensor fusion approach for fast mobile eye-tracking with sensor-shift aware correction. *IEEE Sensors Journal*, 17(24):8356–8366, 2017.
- [44] I. Rigas, H. Raffle, and O. V. Komogortsev. Photosensor oculoigraphy: survey and parametric analysis of designs using model-based simulation. *IEEE Transactions on Human-Machine Systems*, (99):1–12, 2018.
- [45] D. A. Robinson. A method of measuring eye movement using a scleral search coil in a magnetic field. *IEEE Transactions on biomedical electronics*, 10(4):137–145, 1963.
- [46] S. Rostaminia, A. Lamson, S. Maji, T. Rahman, and D. Ganesan. W! nce: Unobtrusive sensing of upper facial action units with eog-based eyewear. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 3(1):1–26, 2019.
- [47] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb. Learning from simulated and unsupervised images through adversarial training. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2242–2251, 2017.
- [48] B. Tag, A. W. Vargo, A. Gupta, G. Chernyshov, K. Kunze, and T. Dinger. Continuous alertness assessments: Using eog glasses to unobtrusively monitor fatigue levels in-the-wild. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19*, p. 1–12. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3290605.3300694
- [49] M. Tonsen, J. Steil, Y. Sugano, and A. Bulling. Invisibleeye: Mobile eye tracking using multiple low-resolution cameras and learning-based gaze estimation. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)*, 1(3):106:1–106:21, 2017. doi: 10.1145/3130971
- [50] C. Topal, S. Gunal, O. Koçdeviren, A. Dogan, and O. N. Gerek. A low-computational approach on gaze estimation with eye touch system. *IEEE Transactions on Cybernetics*, 44(2):228–239, 2014.
- [51] T. Tran, T. Pham, G. Carneiro, L. Palmer, and I. Reid. A bayesian data augmentation approach for learning deep models. In *Advances in neural information processing systems*, pp. 2797–2806, 2017.
- [52] O. T. Tursun, E. Arabadzhyska-Koleva, M. Wernikowski, R. Mantiuk, H.-P. Seidel, K. Myszkowski, and P. Didyk. Luminance-contrast-aware foveated rendering. *ACM Transactions on Graphics (TOG)*, 38(4):1–14, 2019.
- [53] M. Virbel, T. Hansen, and O. Lobunets. Kivy—a framework for rapid creation of innovative user interfaces. In *Workshop-Proceedings der Tagung Mensch & Computer 2011. uberMEDIEN—UBERMorgen*. Universitätsverlag Chemnitz, 2011.
- [54] U. Vogel, D. Kreye, S. Reckziegel, M. Törker, C. Grillberger, and J. Amelung. OLED-on-CMOS integration for optoelectronic sensor applications. In , vol. 6477 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 647703, Feb. 2007. doi: 10.1117/12.704468
- [55] E. Wood, T. Baltrušaitis, L.-P. Morency, P. Robinson, and A. Bulling. Learning an appearance-based gaze estimator from one million synthesised images. In *Proceedings of the Ninth Biennial ACM Symposium on Eye Tracking Research & Applications*, pp. 131–138, 2016.
- [56] E. Wood, T. Baltrušaitis, X. Zhang, Y. Sugano, P. Robinson, and A. Bulling. Rendering of eyes for eye-shape registration and gaze estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3756–3764, 2015.
- [57] F. Yang, B. R. O’hanlon, N. Zahirovic, N. Sarkar, M. Olfat, and M. Mohsenzadeh. Timer-based eye-tracking, July 4 2019. US Patent App. 16/234,293.
- [58] R. Zemblys and O. Komogortsev. Making stand-alone ps-og technology tolerant to the equipment shifts. In *Proceedings of the 7th Workshop on Pervasive Eye Tracking and Mobile Eye-Based Interaction*, pp. 1–9, 2018.