EyeVEIL: Degrading Iris Authentication in Eye Tracking Headsets

Brendan John
University of Florida
brendanjohn@ufl.edu

Sanjeev Koppal
University of Florida
sjkoppal@ece.ufl.edu

Eakta Jain
University of Florida
ejain@cise.ufl.edu

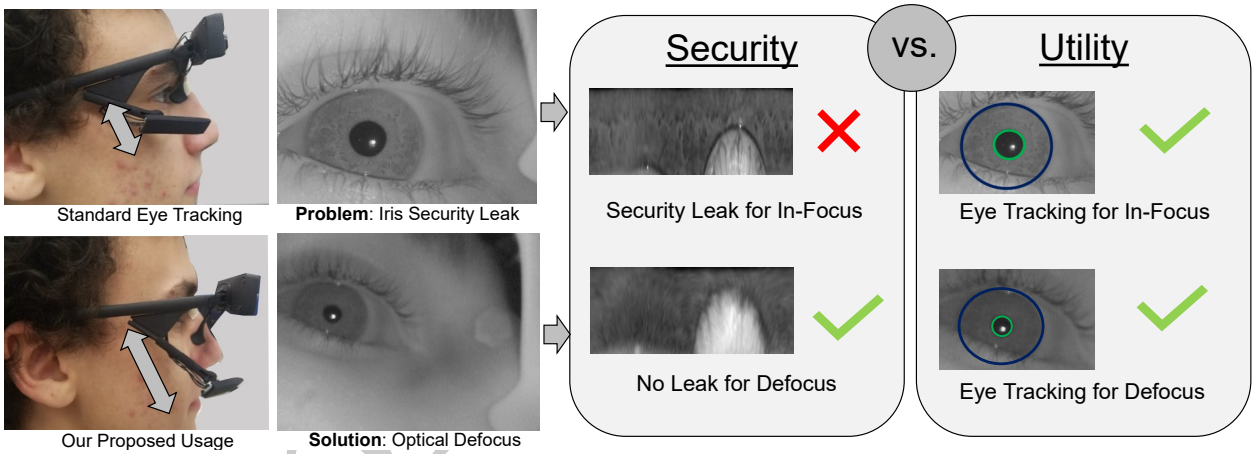

Figure 1: We find that the image quality from many eye tracking cameras is sufficient for iris authentication. While this can enable eye tracking apps such as mobile banking, it creates an opportunity for a hacker to steal a user's identity. We propose a low cost solution by deliberately introducing optical defocus in the eye tracking setup. In this paper, we show that such defocus degrades the iris textures significantly, while still preserving eye tracking utility. Practically, the defocus security option is simple: increase the distance of the IR camera from the eye, as shown in the figure. Further, the user can easily toggle this protection to reveal their biometric. The method's ease of use for a desirable security/utility trade-off suggest that it will impact almost every eye tracking use-case.

ABSTRACT

Mixed reality headsets are being designed with integrated eye trackers: cameras that image the user's eye to infer gaze location and pupil diameter. While the intent is to improve the quality of experience, built-in eye trackers create a security vulnerability for hackers – high resolution images of the user's iris. Anyone stealing an iris image has effectively captured a gold standard biometric, relied on for secure authentication in applications such as banking and voting. We present a low cost solution to degrade iris authentication while still permitting the utility of gaze tracking with acceptable accuracy. By demonstrating this solution on a commodity eye tracker, this paper urges the community to think about iris based authentication as a byproduct of eye tracking, and create solutions that empower a user to control this biometric.

CCS CONCEPTS

• **Security and privacy** → **Privacy protections**; • **Human-centered computing** → *Ubiquitous computing*;

KEYWORDS

Security, Eye Tracking, Iris Recognition, Mixed Reality

ETRA 2019, June 2019, Denver, Colorado USA
2019. ACM ISBN 123-4567-24-567/08/06…$15.00
https://doi.org/10.475/123_4

ACM Reference Format:
Brendan John, Sanjeev Koppal, and Eakta Jain. 2019. *EyeVEIL:* Degrading Iris Authentication in Eye Tracking Headsets. In *Proceedings of ACM ETRA (ETRA 2019).* ACM, New York, NY, USA, 5 pages. https://doi.org/10.475/123_4

1 INTRODUCTION

The successful deployment of mixed reality depends on a number of technological pieces, the most significant of which is an integrated camera that records gaze and the user's eye images. This camera is part of an eye tracking system that computes parameters of the user's gaze to serve critical functions such as reducing the vergence-accommodation conflict, which alleviates user discomfort [13], foveated rendering [1, 16], which optimizes compute and power resources, and user interaction, which allows the system to infer user intent and attention [2]. Eye tracking will transition from being research equipment to a commodity user interaction technology, as ubiquitous as the touch screen.

With a camera that looks at the user's eye, the headset can potentially use iris based authentication to log the user in. The iris is a gold standard biometric, used in sensitive applications such as voting and banking. When users wear cameras pointed at their eyes, an iris image will be captured at tens of frames per second. If a hacker accesses even one reasonable quality frame, they have the user's biometric, and therefore, their identity.

This paper presents three contributions:

(1) first, we show that iris based identification is possible with the resolution and quality of images captured by eye cameras in current mixed reality headsets,
(2) second, we demonstrate that a simple filter can blur the eye image and degrade the accuracy of iris authentication while still permitting gaze tracking,
(3) third, we show that a user can be empowered to easily perform focus manipulation of the eye tracking camera.

Scope: The findings and demonstrations of this paper do not guarantee anonymity for the user when they use a headset with an IR camera based eye tracker; rather they only alleviate the most conspicuous and strongest method of identity determination. There remain other features that machine learning methods could collect and combine into a probabilistic estimate of user identity or preferences. These features include the defocused iris texture, eye shape, patterns of attention allocation, tremors and micro-saccades. These features are weaker and more changeable, and are thus comparable to traces that users leave while browsing the Internet such as IP addresses, session histories, clicks, etc.

2 PROBLEM: SECURITY VULNERABILITY

Survey of commodity headsets: We surveyed popular commodity headsets on the market to understand if iris based authentication is possible with developer access to the eye image stream (see supplementary document). We found that the most common eye image resolution is 320×240, and several APIs allow this type of access.
Eye tracking headsets leak iris signatures: Although biometric standard guidelines state that iris authentication systems should aim for an iris diameter of 200 pixels or more for high quality iris images, recently, Philips and Komogortsev [17] found that the iris diameter can be as low as 50-60 pixels and still yield 90% and higher recognition rates. We conducted our own experiments with five participants wearing a Pupil Labs head-mounted monocular eye tracker [9] shown in Fig. 1. The eye image resolution was set to 320x240 to be representative of current VR and mixed reality eye trackers.

For each participant we selected three images from the recording that were in focus, and had minimal occlusion from the eye lid. Iris authentication was performed by treating each image from a participant as a source, comparing it with all of the images of the same participant and others as targets.

The process of performing iris based authentication is well established, with high frequency features of the iris pattern encoded into a binary code [5]. We use a robust iris segmentation method [6], and standard encoding procedure [15] to create iris codes for each image and compute the Hamming distance between them [4]. A positive match is returned if the Hamming distance is less than a threshold, which typically ranges from 0.2 to 0.4 [4, 15, 18].

Fig. 2 (Left) visualizes the average Hamming distance between source and target eye images for each participant in our experiment. We set the authentication threshold to 0.37. Values lower than this threshold will result in a match between source and target. As seen in Fig. 2 (Left) diagonal values are consistently < 0.37. We conclude that existing recognition algorithms used with commodity eye trackers can successfully perform iris authentication.

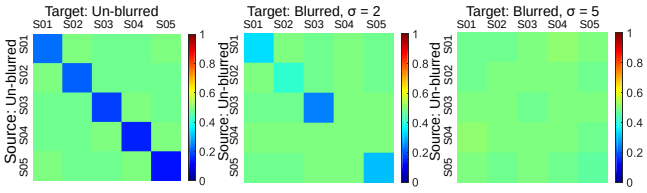

Figure 2: Eye trackers create a security vulnerability: Dark blue values represent valid iris authentication, while lighter values depict a failed match. We performed this experiment on five participants, using three images per participant. The first matrix shows perfect matching, while the next two show that simulated defocus increases failed matches. In particular, $\sigma = 5$ degrades iris matching significantly.

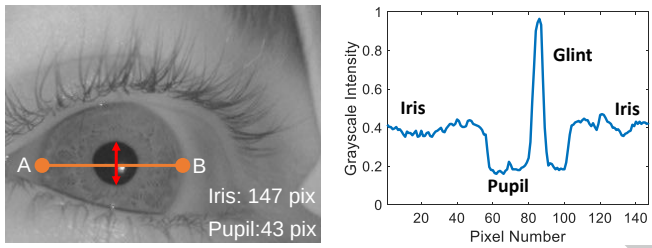

(a) The iris spans ≈ 150 pixels along $\overline{AB}$. **(b) Grayscale values along line $\overline{AB}$.**

Figure 3: We show real data comparing iris intensity values with both specularity (i.e. glint) and pupil edge brightness fall. This suggests that a low pass filter $F(x)$ can degrade the high frequency iris patterns while retaining low frequency features (glint or pupil edge) needed for gaze estimation.

Since eye trackers record eye imagery, mixed reality headsets could use iris authentication (say, for logging in) since this biometric is widely considered to have many security benefits. For example, Jain et al. [8] surveyed established biometric recognition modalities, including DNA, fingerprint, face and retina, and concluded that iris imagery is comparable to DNA with respect to universality, distinctiveness, permanence, and difficulty to circumvent.

Despite this, iris authentication is rare in personal technology because it is inconvenient to require close-in eye images. In contrast, for the next wave of mixed reality technology, such imagery is easily obtained. *In fact, these platforms will stream the biometric equivalent of the user's social security number at video rate, the moment a user wears a headset.* This data stream is a security feature when the application requires strong authentication, whether it is banking apps, personal contacts, family photos, or digital voting. However, the same data stream is a security vulnerability when not in use. If a hacker accesses even one reasonable quality frame, they have the user's biometric, and therefore, their identity.

3 SOLUTION: EYE CAMERA DEFOCUS

In this section we formulate how a defocus filter can be used to degrade high frequency patterns of the iris necessary for authentication, while preserving low frequency features required for gaze estimation. Simulated results are presented here, and in Section 4 we show a demo on a real head-mounted eye tracker.

3.1 Eye tracking vs. iris authentication

Let us denote an image received from the eye camera as $I = I_C + I_R$. For simplicity, we assume only 1D functions by thinking about them as the variation in grayscale intensity along a line that spans the iris and pupil from one side to the other (Fig. 3). All of our arguments generalize to 2D signals.

I_C is the component of the image that contains eye tracking signal, for example, the corneal reflection or the pupil, which we model as a Gaussian $I_C(\mu = 0, \sigma_C)$. The iris is I_R, and although iris texture has broad variation, we can at least assume that the highest frequency in the signal is limited by the size of each pixel of the camera. Let us denote the highest frequency as B. While I_C contains primarily low frequency content, I_R contains both low and high frequency content, with the higher frequencies being the identifying features (up to the maximum frequency B).

Consider a low-pass filter F of the form $F(x) = N(\mu = 0, \sigma)$, i.e., a Gaussian blur, which simulates optical defocus. When I is convolved with $F(x)$, the result is

$$I_D(x) = I(x) * F(x) = I_C(x) * F(x) + I_R * F(x) = I'_C(x) + I'_R(x). \quad (1)$$

Our claim: it is possible to select the filter $F(x)$ s.t. the eye tracking features are still detectable in $I'_C(x)$ while I'_R no longer contains the higher frequencies that enable iris based authentication.

3.2 Optical defocus reduces iris authentication accuracy: Simulation

Philips and Komogortsev [17] showed that convolving the eye image with a 2D Gaussian filter reduces authentication accuracy. We replicated their experiment, conducting our own simulations to understand the fall off in authentication accuracy for eye images received from a defocused head-mounted tracker. We simulated optical defocus by convolving with a Gaussian kernel and computed the Hamming distance between the blurred and unblurred images. Fig. 2 shows 40% degraded authentication for $\sigma = 2$ pixels and at $\sigma = 5$ pixels there are no positive matches.

3.3 Defocus retains gaze estimation capability: Simulation

In this section we show that defocus retains the utility of gaze estimation from eye tracking. This is done using a commodity glasses based monocular eye tracker. The eye tracker uses dark pupil tracking and a 320x240 resolution IR image of the eye.

Our experiment: Five participants with normal vision wore the Pupil Labs eye tracker (Pupil Pro 2016, 30Hz) in a lab environment and performed a 5-point calibration. The eye tracker reported a validation accuracy of <1.5 degrees. Participants were seated and using a chin rest were asked to keep their head still and just move their eyes to look at five targets that appeared on a computer monitor. We recorded the eye image stream, scene camera video stream, and gaze data. This experiment took approximately 32 seconds. See the supplementary video for a visualization of the data.

Each frame of the original eye image stream is denoted $f_i, i = 1, \ldots, N$. Each frame f_i defocused in simulation by convolving with a Gaussian kernel of $\sigma = 0, 1, 2, 3, \ldots, 10$ pixels respectively. The blurred frame was denoted f'_i. As pupil detection is a prerequisite for gaze tracking, it is only on those frames where the pupil was

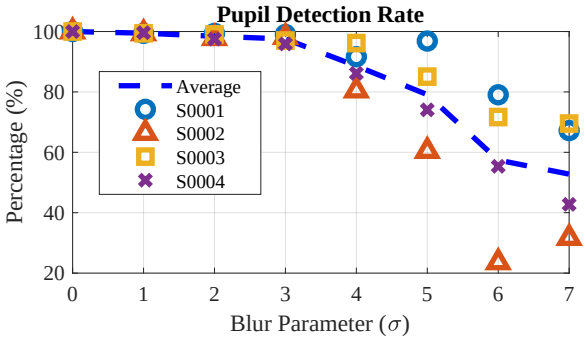

Figure 4: Detection rate for participants S0001-S0004 with simulated Gaussian blur up to $\sigma = 7$ pixels. Note that at $\sigma = 5$, where iris degradation is significant for this eye tracker, detection rate is still at $\approx$ 80%.

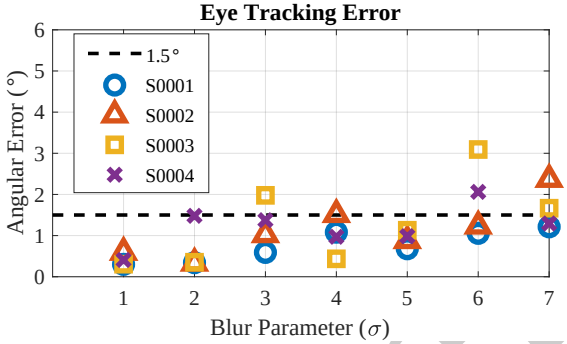

Figure 5: Error in visual degrees of eye tracking data produced with simulated Gaussian blur up to $\sigma = 7$ pixels. Note that at $\sigma = 5$, where iris degradation is significant for this eye tracker, the tracking error is at or below 1.5°.

detected that we can compute eye tracking accuracy. For those frames $f_j, j = 1, \ldots, M$ where the pupil was detected[1], the 3D gaze vector reported by the Pupil Labs eye tracker was denoted x_j. These gaze points formed the ground truth data. We ran the blurred frames f'_j through the Pupil Labs pipeline, and for those frames where the pupil was detected, the recomputed 3D gaze vector was denoted x'_j. Data from participant S0005 was removed from analysis due to large inconsistencies in pupil detection and gaze accuracy.

Fig. 4 shows the pupil detection rate of the frames f'_i for different levels of σ in pixels. Detection rate is computed as the percentage of frames f'_i where the pupil is detected, relative to the number of detected pupils in the frames f_i. Even with kernels as large as $\sigma = 5$, $\approx$ 80% of the frames have the pupil detected on average. Default parameters for pupil detection were used in this experiment.

Fig. 5 shows the average angular error between the gaze vector x_j and x'_j. The angular error was computed as $\theta = cos^{-1}(x_j \cdot x'_j)$, where both x_j and x'_j were normalized. For blur up to $\sigma = 5$ there is only one instance where gaze error exceeded 1.5°.

4 OPTICAL DEFOCUS DEMO

We asked a participant to wear the Pupil Labs head-mounted eye tracker and look at the presented gaze targets. Five targets were

[1]Note that $M < N$

presented, four at the corners of a screen, and one in the center, as in previous experiments. Data was recorded in two configurations of the eye tracker: first, when the eye tracker was configured so that the eye image was perfectly in-focus, and second, when the eye tracker was configured so that the eye image was significantly defocused. Fig. 1 shows the eye tracker on the participant, as well as an example of the eye image in each of these configurations. In our supplementary video, we show the gaze position on the scene camera video for each of these configurations. The eye tracker was calibrated to an accuracy of <1.5°.

4.1 Optical defocus reduces iris authentication accuracy: Demo

In addition to viewing targets on the screen the participant was instructed to look directly at the IR camera for a duration of five seconds. Our goal was to determine if the iris pattern captured in the defocus configuration could successfully authenticate the user. Fig. 6(b) visualizes the Hamming distance between the source and target frame for each frame in the defocus configuration. The average Hamming distance was 0.485 which was greater than our authentication threshold of 0.37. We also compared every in-focus frame with every other in-focus frame to determine the effectiveness of iris authentication in the vulnerable configuration. Visualized in Fig. 6(a), 91% of in-focus frames produced a successful match, providing a security vulnerability for a large number of frames that could identify the participant if a database containing at least one of their iris images was also hacked.

4.2 Defocus retains gaze estimation capability: Demo

In each configuration we computed the distance between the gaze location reported by the eye tracker and the presented target within the scene camera view. In Fig. 7, we show the gaze data plotted on a reference frame of the scene camera video, and the mean error from the presented targets for each configuration. Mean gaze error was calculated in pixels within the scene camera view, where the computer monitor spanned approximately 676 pixels horizontally, and 433 pixels vertically. The mean error for the in-focus and defocus configurations was 10 and 19 pixels respectively. This implies that the increase in error from our hardware solution was negligible with respect to on screen targets, as the difference in gaze position was less than 3% of the screen height and 5% of the screen width.

5 DISCUSSION

Much research has been done in the eye-tracking and biometrics community on using signatures within eye movements to identify a person [7, 10, 11, 14]. Specifically in the context of head-mounted eye tracking in virtual reality, Lohr et al. [14] recently proposed a software framework to authenticate end users using eye movements. While there is no doubt that we will continue to locate individual signatures in eye movements, hand gestures, head movements, etc., these signatures are indicators of who the user might be, and thus are useful for building profiles, or providing user-specific recommendations. These probabilistic inference tasks can absorb false positives, whereas biometric authentication for ultra-sensitive scenarios such as voting and banking have no room for such errors.

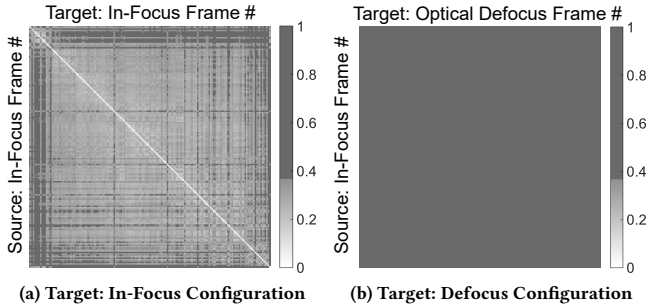

(a) Target: In-Focus Configuration (b) Target: Defocus Configuration

Figure 6: Hamming distances computed between the iris patterns from all recorded frames in our real experiment. Dark gray indicates a bad match, with white indicating a successful match and a security vulnerability.

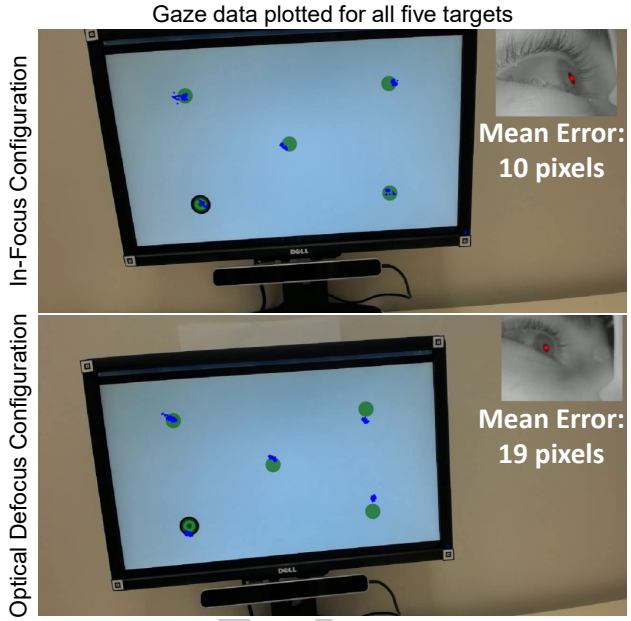

Figure 7: Our experiments show that reasonable eye tracking accuracy can be achieved with an eye tracker configuration that defocuses the eye image. This data was recorded using the configurations shown in Fig. 1.

Our key insight is that users need a way to *toggle iris authentication on or off* depending on their situation. The method we have proposed is a convenient solution; all the user has to do is to defocus their eye camera to prevent a hacker from stealing their in-focus iris image. We note that the biometrics community has considered the problem of fraudulent iris authentication, using contact lenses or artificial eyes, for example. Methods to prevent these attacks include "live-ness" detection to prevent fraudulent authentication [3, 12], which makes stealing live-action eye camera feeds, with real jitter and pupillary tremor, even more attractive. We hope that future work in the eye tracking community will create tools and techniques for the user to control how their personally identifying information is transmitted.

ACKNOWLEDGMENTS

Authors acknowledge funding from the National Science Foundation (Awards IIS-1566481 and IIS-1514154), and the NSF Graduate Research Fellowship (Awards DGE-1315138 and DGE-1842473).

REFERENCES

[1] Rachel Albert, Anjul Patney, David Luebke, and Joohwan Kim. 2017. Latency requirements for foveated rendering in virtual reality. *ACM Transactions on Applied Perception (TAP)* 14, 4 (2017), 25.

[2] Roman Bednarik, Hana Vrzakova, and Michal Hradis. 2012. What do you want to do next: a novel approach for intent prediction in gaze-based interaction. In *Proceedings of the symposium on eye tracking research and applications*. ACM, 83–90.

[3] Yangyu Chen and Weigang Zhang. 2018. Iris Liveness Detection: A Survey. In *2018 IEEE Fourth International Conference on Multimedia Big Data (BigMM)*. IEEE, 1–7.

[4] John Daugman. 2009. How iris recognition works. In *The essential guide to image processing*. Elsevier, 715–739.

[5] John G Daugman. 1993. High confidence visual recognition of persons by a test of statistical independence. *IEEE transactions on pattern analysis and machine intelligence* 15, 11 (1993), 1148–1161.

[6] Abhishek Gangwar, Akanksha Joshi, Ashutosh Singh, Fernando Alonso-Fernandez, and Josef Bigun. 2016. IrisSeg: A fast and robust iris segmentation framework for non-ideal iris images. In *Biometrics (ICB), 2016 International Conference on*. IEEE, 1–8.

[7] Corey D Holland and Oleg V Komogortsev. 2014. Software framework for an ocular biometric system. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. ACM, 365–366.

[8] Anil K Jain, Arun Ross, and Salil Prabhakar. 2004. An introduction to biometric recognition. *IEEE Transactions on circuits and systems for video technology* 14, 1 (2004), 4–20.

[9] Moritz Kassner, William Patera, and Andreas Bulling. 2014. 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*. ACM, 1151–1160.

[10] Tomi Kinnunen, Filip Sedlak, and Roman Bednarik. 2010. Towards task-independent person authentication using eye movement signals. In *Proceedings of the 2010 Symposium on Eye-Tracking Research & Applications*. ACM, 187–190.

[11] Oleg V Komogortsev, Sampath Jayarathna, Cecilia R Aragon, and Mechehoul Mahmoud. 2010. Biometric identification via an oculomotor plant mathematical model. In *Proceedings of the 2010 Symposium on Eye-Tracking Research & Applications*. ACM, 57–60.

[12] Oleg V Komogortsev, Alexey Karpov, and Corey D Holland. 2015. Attack of mechanical replicas: Liveness detection with eye movements. *IEEE Transactions on Information Forensics and Security* 10, 4 (2015), 716–725.

[13] Gregory Kramida. 2016. Resolving the vergence-accommodation conflict in head-mounted displays. *IEEE Transactions on Visualization & Computer Graphics* 1 (2016), 1–1.

[14] Dillon Lohr, Samuel-Hunter Berndt, and Oleg Komogortsev. 2018. An implementation of eye movement-driven biometrics in virtual reality. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications*. ACM, 98.

[15] Libor Masek and Peter Kovesi. 2013. MATLAB source code for a biometric identification system based on iris patterns, The University of Western Australia, 2003.

[16] Anjul Patney, Marco Salvi, Joohwan Kim, Anton Kaplanyan, Chris Wyman, Nir Benty, David Luebke, and Aaron Lefohn. 2016. Towards foveated rendering for gaze-tracked virtual reality. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 179.

[17] Clark Phillips and Oleg V Komogortsev. 2011. Impact of Resolution and Blur on Iris Identification. *Technical Report* (2011).

[18] S Rakshit and DM Monro. 2007. Robust iris feature extraction and matching. In *Digital Signal Processing, 2007 15th International Conference on*. Citeseer, 487–490.