

Trade-offs in Privacy-Preserving Eye Tracking through Iris Obfuscation: A Benchmarking Study

Mengdi Wang, Efe Bozkir, and Enkelejda Kasneci

Technical University of Munich, Munich, Germany

{mengdi.wang, efe.bozkir, enkelejda.kasneci}@tum.de

Abstract—Recent developments in hardware, computer graphics, and AI may soon enable AR/VR head-mounted displays (HMDs) to become everyday devices like smartphones and tablets. Eye trackers within HMDs provide a special opportunity for such setups as it is possible to facilitate gaze-based research and interaction. However, estimating users’ gaze information often requires raw eye images and videos that contain iris textures, which are considered a gold standard biometric for user authentication, and this raises privacy concerns. Previous research in the eye-tracking community focused on obfuscating iris textures while keeping utility tasks such as eye segmentation and gaze estimation accurate. Despite these attempts, there is no comprehensive benchmark that evaluates state-of-the-art approaches. Considering all, in this paper, we benchmark blurring, noising, downsampling, rubber sheet model, and iris style transfer to obfuscate users in the data to protect their privacy, and compare their impact on image quality, privacy, utility, and risk of imposter attack on two datasets. We use eye segmentation and gaze estimation as utility tasks, and reduction in iris recognition accuracy as a measure of privacy protection, and false acceptance rate to estimate risk of attack. Our experiments show that canonical image processing methods like blurring and noising cause a marginal impact on deep learning-based tasks. While downsampling, the rubber sheet model, and iris style transfer are effective in hiding user identifiers, iris style transfer, with higher computation cost, outperforms others in both utility tasks, and is more resilient against spoof attacks than the rubber sheet model. Our analyses indicate that there is no universal, optimal approach to balance privacy, utility, and computation burden. Therefore, we recommend practitioners consider the strengths and weaknesses of each approach, and possible combinations of those to reach an optimal privacy-utility trade-off.

I. INTRODUCTION

Recent advances in hardware, human-computer interaction, and AI will likely lead AR/VR and HMDs to become pervasive, similar to today’s smartphones and tablets. Previous works have extensively studied the interaction techniques while using HMDs [1], and gaze-based interaction is considered as one of the intuitive ways especially for hands- and voice-free interaction [2]. The use of gaze interaction and eye-tracking data provides many opportunities for AR and VR; however, it also raises privacy concerns [3] considering that it can reveal sensitive user characteristics, including race, sexual preference, or body mass index [4]. Furthermore, it is also possible to authenticate users with iris texture data [5] and with eye movements [6]. In fact, as each person’s iris texture is unique and as it is treated as visual fingerprints, users can be authenticated with very high accuracy [7]. Such an authentication process is useful when the purpose is user authentication, and

data protection is handled responsibly; otherwise, it may give adversaries a good opportunity for impersonation and access to sensitive resources when iris texture data is in hand.

To address this issue, some of the HMD manufacturers do not give access to raw eye image data to application developers and users through their HMDs (e.g., Apple Vision Pro [8], Microsoft HoloLens 2 [9]), others opt for providing access to raw images (e.g., Varjo XR-4 [10]). Since gaze estimation techniques mostly utilize eye images to determine the gaze direction, completely limiting access to raw data is not reasonable, especially for computer vision practitioners and researchers. Instead, privacy-preserving methods that degrade iris-based authentication while keeping utility tasks like gaze estimation and eye segmentation accurate, are necessary. Some of the previous works approached this by image blurring [11], [12], noising [13], spatial downsampling [14], replacing iris region through a rubber sheet model [15], [16], and style transfer [17]. However, there is no comprehensive benchmarking study that evaluates the proposed concepts in the literature and compares their privacy-utility trade-offs. We benchmark these concepts using widely recognized HMD datasets, particularly by using the OpenEDS2019 [18] and OpenEDS2020 [19] datasets for privacy protection, semantic segmentation, and gaze estimation tasks. We find that while traditional image processing techniques like noising and blurring do not degrade image utility, they also cannot hide user identity from deep learning-based iris authenticators. Spatial downsampling can notably obfuscate iris identifiable features but also degrades segmentation accuracy for pupil region and performance of the appearance-based estimator. While the rubber sheet model reduces the iris recognition rate to a large degree, it leads to a high risk of iris attack and undermines appearance-based gaze estimation. In comparison, iris style transfer works accurately for utility tasks and significantly reduces iris recognizability without escalating the false acceptance rate, whereas its vanilla pipeline is computationally intensive. In summary, we state that there is no universal best-performing approach to protecting iris in AR/VR settings that harmonizes all aspects including privacy protection, data utility, risk of attack, and computation overhead, and multiple approaches may be combined considering their strengths and weaknesses to find an optimal privacy-utility trade-off. We also publish our code¹ for reproducibility.

II. RELATED WORK

With eye tracking becoming ubiquitous with HMDs and smart glasses, despite all the advantages of using it in everyday settings, such as for hands-free interaction, foveated

¹<https://gitlab.lrz.de/hctl/Iris-Obfuscation-Benchmark>.

rendering, and understanding human perception and cognition, the use of eye-tracking data comes with privacy risks [3]. To mitigate these risks, researchers developed different techniques to protect users' privacy, with differential privacy [20], [21], secure multi-party computation [22], and signal processing-based approaches including spatial and temporal downsampling [23]. Yet, these approaches focus on the protection of eye movements and their features after gaze estimation takes place, considering eye movements and features can reveal sensitive information about users [4], [24]. However, when the raw eye data is considered, one of the most sensitive data types is iris texture, as it can uniquely identify individuals. To protect iris information, previous research has frequently focused on obscuring iris texture to impede iris-based user identification [25]. John et al. [11] proposed Gaussian blurring to achieve this while retaining gaze estimation capability. Further, John et al. [12] showed that with such a blurring mechanism to approach optical defocus in a conversational setting in VR, it is possible to find a trade-off that eye movement changes due to the privacy protection can go unnoticed. In another work [13], the authors introduced pixel noise on iris textures and stated similarly that gaze estimation errors due to privacy protection are in acceptable ranges. Chaudhary and Pelz [15] proposed to replace iris segments of an eye image with fake irises using a rubber sheet model and showed that such manipulated eye images can still be useful for eye semantic segmentation while protecting individuals' privacy. Later, the rubber sheet method has been studied in the opposite direction, namely constructing an iris attack by swapping in attackers' iris texture to fool liveness detection and iris authentication systems [16]. Recently, Wang et al. [17] proposed utilizing style features for iris recognition and suggested protecting privacy via iris style transfer accordingly. The authors showed their method maintains eye image utility for eye segmentation and for gaze estimation using both model- and appearance-based methods, while suppressing the risk of impersonation attack.

Despite the importance of preserving eye image privacy, prior works that systematically compare different iris obfuscation methods [25] are commonly limited to primitive image processing techniques like filtering and noising and their influence on traditional iris recognition systems and utility tasks, which commonly rely on handcrafted kernels and objectives [26]–[28]. A comprehensive benchmarking study in the literature that compares the state-of-the-art methods for iris obfuscation is missing. We address this in our work by comparing different iris manipulation methods regarding their effect on image quality, privacy, risk of attack, and utility, using deep learning-based models on two publicly available datasets collected in VR environments.

III. EXPERIMENTS

In this study, we extensively benchmarked five iris obfuscation methods, including blurring, noising, downsampling, rubber sheet model, and iris style transfer, regarding their impact on image quality, privacy, and utility. Particularly, the influence on image privacy involves the impact on iris recognition and the risk of malicious use, while the study on data utility covers the performance of eye semantic segmentation and gaze estimation using obfuscated eye images.

A. Methods

1) *Blurring*: Blurring, which is a specific type of filtering, is a classical image processing technique that is often used for noise reduction and edge softening. It operates by convolving the image with a kernel function K . Given an image $I(x, y)$, the image convolution can be represented as:

$$I_{\text{blurred}}(x, y) = \sum_{i=-k}^k \sum_{j=-k}^k K(i, j) I(x - i, y - j) \quad (1)$$

where k denotes kernel size. Blurring is also a low-pass filter. Phillips and Komogortsev [14] investigated the influence of Gaussian blurring on iris identification. Other work also hypothesized that eye images consist of low-frequency eye-tracking signals and high-frequency identifiable iris features, and used Gaussian blurring to preserve privacy [11], [12].

2) *Noising*: Image noising refers to the process of deliberately adding noise to an image. It can be modeled by:

$$I_{\text{noisy}}(x, y) = I(x, y) + N(x, y), N(x, y) \sim P_N(\theta) \quad (2)$$

with $N(x, y)$ being random noise sampled from a specific distribution $P_N(\theta)$ parametrized by θ . Noising can be used as a privacy-enhancing tool and is often employed in privacy mechanisms like differential privacy. Previous work proposed to intelligently add noise to specific image regions to protect privacy while preserving image aesthetics [29]. John et al. [13] introduced pixel-wise noise to eye images to obscure iris signature and achieved $(0, \delta)$ -differential privacy.

3) *Downsampling*: Image downsampling is the process of reducing the resolution of an image by decreasing the number of pixels while trying to retain visual information. For an image $I(x, y)$ and a scaling factor s , the reduced image can be represented as $I_{\text{down}}(x', y') = I(sx', sy')$. Since the decreased image resolution can distort sensitive information in the image, it can also be used to preserve privacy. For instance, Phillips and Komogortsev [14] showed that the resolution of eye images has a strong impact on iris identification.

4) *Rubber Sheet*: The rubber sheet model is an image warping and transformation technique. In iris recognition systems, the rubber sheet model can be used in combination with polar transformation to preprocess iris texture [28]. Recently, the rubber sheet model has been studied as a privacy-preservation method by swapping the real iris texture with a fake template [15]. Later, the same method was applied in the opposite direction, where the researchers switched the victims' iris texture with attackers' to fool liveness detection and iris authentication systems [16].

5) *Iris style transfer*: Iris style transfer [17] is a recently proposed eye privacy-enhancing technique inspired by neural style transfer [30], [31], which is a non-photorealistic image synthesis method [32], [33]. It works under the premise that images consist of content and style components, which can be separated and manipulated independently. By iteratively optimizing the content and style objectives, neural style transfer recombines the content and style of different images. Li et al. [34] proved that neural style transfer can be modeled as a domain adaptation problem in feature space, and proposed a




Fig. 1: Examples of raw irises (a and b) and obfuscated irises (c - g).

TABLE I: Impact of iris obfuscation methods on privacy, utility, and risk of attack. Baseline: no obfuscation. Metrics: classification accuracy for iris recognition, false acceptance rate (FAR) for attack success, intersection over union (IoU) for eye segmentation, and degree error for gaze estimation. s : scaling factor. e : style transfer iteration. c1: feed-forward CNN feature-driven classifier. c2: style feature-driven classifier. m1: model-based gaze estimator. m2: appearance-based gaze estimator.

Method	Parameter	Iris Recognition		Attack Success		Eye Segmentation				Gaze Estimation	
		c1	c2	c1	c2	skin	sclera	iris	pupil	m1	m2
Baseline	NA	0.962	0.978	0.001	0.000	0.995	0.932	0.956	0.945	6.970°	3.190°
Blurring	$\sigma = 1$	0.959	0.964	0.001	0.000	0.995	0.932	0.956	0.945	6.970°	3.193°
	$\sigma = 2$	0.953	0.963	0.001	0.000	0.995	0.932	0.956	0.944	6.971°	3.197°
	$\sigma = 3$	0.951	0.963	0.001	0.000	0.995	0.932	0.956	0.944	6.970°	3.198°
	$\sigma = 4$	0.951	0.963	0.001	0.000	0.995	0.932	0.956	0.944	6.971°	3.199°
	$\sigma = 5$	0.949	0.963	0.001	0.001	0.995	0.932	0.956	0.944	6.970°	3.199°
Noising	$\sigma = 0.01$	0.961	0.975	0.001	0.000	0.995	0.932	0.956	0.945	6.970°	3.189°
	$\sigma = 0.05$	0.950	0.964	0.001	0.002	0.995	0.932	0.956	0.945	6.970°	3.192°
	$\sigma = 0.1$	0.949	0.963	0.002	0.002	0.995	0.932	0.955	0.942	6.970°	3.199°
	$\sigma = 0.2$	0.948	0.963	0.002	0.002	0.995	0.931	0.953	0.936	6.975°	3.217°
	$\sigma = 0.5$	0.948	0.963	0.001	0.000	0.995	0.929	0.947	0.931	7.019°	3.259°
Down-sampling	$s = 1.5$	0.831	0.222	0.001	0.006	0.995	0.925	0.944	0.917	7.039°	3.425°
	$s = 2$	0.662	0.150	0.004	0.007	0.995	0.924	0.940	0.903	7.042°	3.872°
	$s = 3$	0.462	0.070	0.006	0.004	0.994	0.923	0.935	0.883	7.053°	4.353°
	$s = 4$	0.344	0.031	0.008	0.006	0.994	0.922	0.932	0.865	7.080°	4.854°
	$s = 5$	0.278	0.026	0.009	0.005	0.994	0.922	0.930	0.850	7.108°	5.287°
Rubber Sheet	NA	0.069	0.025	0.245	0.369	0.994	0.918	0.931	0.897	7.008°	19.574°
Iris Style Transfer	$e = 1$	0.835	0.175	0.005	0.041	0.995	0.928	0.949	0.933	7.016°	3.187°
	$e = 10$	0.835	0.175	0.005	0.041	0.995	0.928	0.949	0.933	7.016°	3.187°
	$e = 50$	0.773	0.112	0.008	0.054	0.995	0.927	0.948	0.930	6.976°	3.199°
	$e = 100$	0.742	0.093	0.011	0.054	0.995	0.927	0.947	0.928	6.988°	3.214°
	$e = 200$	0.696	0.064	0.012	0.062	0.995	0.925	0.945	0.924	7.026°	3.329°

novel way to compute style loss as statistical matching:

$$\mathcal{L}_{style}^l = \frac{1}{N^l} \sum_{i=1}^{N^l} \left((\mu_{F_S^l}^i - \mu_{F_I^l}^i)^2 + (\sigma_{F_S^l}^i - \sigma_{F_I^l}^i)^2 \right) \quad (3)$$

where F_S^l and F_I^l are, respectively, feature maps of the style and target images extracted by a convolution neural network (CNN) at layer l , μ and σ represent the mean and standard deviation of the feature maps, and N^l is number of channel in the feature maps. Wang et al. [17] proposed style feature-based iris recognition, which is robust against image rotation and perspective shift, and suggested transferring iris texture style to block iris recognition accordingly. Unlike the rubber sheet model which can be used for adversarial purposes [15], [16], the risk of malicious use of iris style transfer is low.

B. Datasets, models, and tasks

We employed two publicly available datasets in our experiments, namely OpenEDS2019 [18] and OpenEDS2020 [19], collected from VR HMDs with near-eye cameras. The former is from 152 users for eye semantic segmentation, whereas the latter is specialized for gaze estimation. We investigated the

influence of iris obfuscation methods on gaze estimation using OpenEDS2020, while we conducted the benchmarks regarding data privacy and eye segmentation on OpenEDS2019.

To explore the influence of iris obfuscation methods, we compared the outcome of different tasks before and after manipulation. To approach iris obfuscation, the iris regions were first segmented out from eye images. We utilized two publicly available pre-trained segmentation networks for this purpose, i.e., RITnet [35] and EfficientNet [36], for OpenEDS2019 and OpenEDS2020, respectively. Then, the glint blobs were detected with thresholding. Later, the iris regions went through different obfuscation techniques, and the glints were added back. In the end, the manipulated irises were placed back onto the eye images for post-obfuscation evaluation.

We emulated iris recognition with iris-based user classification among 152 users, and trained two classifiers. The first one uses the common CNN feed-forward features extracted by VGG19 [37] network, whilst the second one uses the style features [17]. For the impact on eye segmentation, we used the RITnet and EfficientNet models again. To examine the influence on gaze estimation, we trained both model-

and appearance-based gaze estimators [38]. The former relies on eye landmark features, while the latter utilizes a ResNet50 [39] to encode hidden representations. After being trained or loaded, all models were frozen for benchmarks.

Particularly, the rubber sheet model and iris style transfer may be used for malicious purposes, i.e., impersonation attacks by swapping iris texture or transferring iris style of an attacker's eye image into a victim's. To simulate such a scenario, for each eye image of user A, we randomly sampled another eye image of another user B, and swapped the iris texture or transferred the iris style of B's image into A's. The relevant measurement is FAR (false acceptance rate), a key metric for biometric authentication systems, and is computed as the ratio between the number of unauthorized but falsely accepted users and the total number of spoofing attacks.

In our benchmarks, we used Gaussian kernel for image blurring with σ in $\{1, 2, 3, 4, 5\}$ and kernel size $= 6\sigma + 1$, and Gaussian noise for image noising with $\mu = 0$ and σ from $\{0.01, 0.05, 0.1, 0.2, 0.5\}$. For downsampling, we set the interpolation strategy to the nearest neighbor and explored scaling factors from 1.5 to 5. For iris style transfer, we fixed the ratio between content and style objective weights to 1, and performed optimization for $\{1, 10, 50, 100, 200\}$ iterations.

C. Evaluation

We first analyzed the obscured iris images, as depicted in Figure 1. Visually, blurring ($\sigma = 2$) and downsampling (scaling factor 2) resulted in a similar distortion effect, while the latter also blurred the boundaries of glint blobs. In contrast, noising introduced noticeable white noise even with a relatively small variance ($\sigma = 0.05$). Different from the three computation-friendly methods that strongly degraded image quality, the rubber sheet model and iris style transfer (200 iters) generated images with decent quality and realism. While the rubber sheet model successfully deformed the attacker's iris according to the victim's iris shape, it also involved the glints, which could affect gaze estimation approaches that rely on them. Iris style transfer brought in subtle but sensible changes in iris patterns and maintained image visualization to a large extent.

The quantitative results, measured through iris recognition accuracy, FAR, eye segmentation accuracy, and gaze estimation error, were reported in Table I. We noticed that while blurring and noising caused negligible drop in image utility across different variance magnitudes, they did not particularly protect user identity, as iris identification using the manipulated iris images preserved its accuracy. This contradicts the previous findings in [25], where the impact of noising and blurring was remarkable. We believe the main reason for such discrepancy lies in that in [25] the tasks were mainly conducted with canonical image processing methods that commonly rely on handcrafted kernels, while our models are pre-trained CNNs, which are more robust against noise and distortion. Another lightweight manipulation technique, namely image downsampling, with increasing scaling factor, resulted in a remarkable decrease in recognition accuracy from 96.2% to 27.8% and from 97.8% to 2.6%, albeit sharing a similar visual effect with blurring. In the meanwhile, downsampling noticeably deteriorated segmentation accuracy for the pupil region, whilst other eye regions were slightly impacted. While downsampling

caused a subtle influence on the model-based gaze estimation approach, it resulted in a notable increase in prediction error for the appearance-based estimator. This can be attributed to that the model-based method necessitates eye segmentation as a preprocessing step, and since downsampling did not lower segmentation accuracy much, the model-based gaze estimator maintained its performance considerably.

For the iris swapping through the rubber sheet model, we observed a substantial drop in iris recognition accuracy from 96.2% to 6.9% for the CNN feature-based classifier and from 97.8% to 2.5% for the style feature-based identifier, indicating that rubber sheet model can hinder user identification. However, it also escalated the risk of impersonation attack, as FAR increased from 0.1% to 24.5% and from 0% to 36.9% in our two settings. This is in accordance with our expectation since the rubber sheet model substitutes the attacker's iris for the victim's. On the utility side, the rubber sheet model maintained the eye segmentation accuracy in essence, as it reshapes the attacker's iris according to the shape of the victim's iris. The model-based gaze estimation performance was preserved as well for this reason. On the contrary, as expected, we noticed a prominent degradation in appearance-based gaze estimation performance, indicated by an escalation of gaze error from 3.190° to 19.574° , because the whole iris region in the manipulated eye image belonged to someone else. The other computationally intensive method, namely iris style transfer, successively reduced the iris recognition accuracy of the style-based classifier to 17.5% - 6.4% with increasing iterations, while causing a lighter but noticeable impact on the CNN feature-based authenticator. Different from rubber sheet model which is prone to spoof attacks, iris style transfer merely caused a negligible growth in attack success chance, and the utility tasks' performance was preserved to a large extent.

From the perspective of computation overhead, primitive image processing methods, including blurring, noising, and scaling, are lightweight and can be performed easily in real time. Iris swapping using a rubber sheet model involves polar transformation and row-wise interpolation, and is hence more computationally demanding. Compared to other benchmarked methods, the vanilla style transfer operates through iterative optimization [40] and is thus the most resource-intensive. For a more efficient eye image and even video stylization in real time, advanced variants [41]–[43] should be applied.

IV. CONCLUSION

In this paper, we conducted extensive benchmarks of five iris obfuscation methods, including blurring, noising, downsampling, rubber sheet model, and iris style transfer, regarding their impact on eye image quality, privacy, utility, and risk of attack, in deep learning-based tasks. Our results indicate that there is no globally superior approach that balances all factors including privacy, utility maintenance, resilience against impersonation attacks, and computation overhead. Therefore, we suggest that the choice of iris manipulation method should be determined according to individual scenarios, with consideration to potential combinations of methods.

ACKNOWLEDGMENT

We acknowledge the funding by the Deutsche Forschungsgemeinschaft (DFG) – Project number KA 4539/5-1.

REFERENCES

- [1] B. Spittle, M. Frutos-Pascual, C. Creed, and I. Williams, "A review of interaction techniques for immersive environments," *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 9, 2023.
- [2] A. Plopski, T. Hirzle, N. Norouzi, L. Qian, G. Bruder, and T. Langlotz, "The eye in extended reality: A survey on gaze interaction and eye tracking in head-worn extended reality," *ACM Computing Surveys*, vol. 55, no. 3, 2022.
- [3] E. Bozkir, S. Özdel, M. Wang, B. David-John, H. Gao, K. Butler, E. Jain, and E. Kasneci, "Eye-tracked virtual reality: A comprehensive survey on methods and privacy challenges," 2023. [Online]. Available: <https://arxiv.org/abs/2305.14080>
- [4] D. J. Liebling and S. Preibusch, "Privacy considerations for a pervasive eye tracking world," in *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. ACM, 2014.
- [5] A. Kumar and A. Passi, "Comparison and combination of iris matchers for reliable personal authentication," *Pattern Recognition*, vol. 43, no. 3, 2010.
- [6] C. Katsini, Y. Abdrabou, G. E. Raptis, M. Khamis, and F. Alt, "The role of eye gaze in security and privacy applications: Survey and future hci research directions," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, 2020.
- [7] J. G. Daugman, "High confidence visual recognition of persons by a test of statistical independence," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no. 11, 1993.
- [8] Apple, "Apple Vision Pro Privacy Overview," https://www.apple.com/privacy/docs/Apple_Vision_Pro_Privacy_Overview.pdf, 2024, last accessed 06/03/2025.
- [9] Microsoft, "Eye tracking on HoloLens 2," <https://learn.microsoft.com/en-us/windows/mixed-reality/design/eye-tracking>, 2023, Last accessed 06/03/2025.
- [10] Varjo Developer, "Eye tracking," <https://developer.varjo.com/docs/native/eye-tracking>, n.d., last accessed 06/03/2025.
- [11] B. John, S. Koppal, and E. Jain, "EyeVEIL: degrading iris authentication in eye tracking headsets," in *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*. ACM, 2019.
- [12] B. John, S. Jörg, S. Koppal, and E. Jain, "The security-utility trade-off for iris authentication and eye animation for social virtual avatars," *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 5, 2020.
- [13] B. John, A. Liu, L. Xia, S. Koppal, and E. Jain, "Let it snow: Adding pixel noise to protect the user's identity," in *ACM Symposium on Eye Tracking Research and Applications*. ACM, 2020.
- [14] C. Phillips and O. V. Komogortsev, "Impact of resolution and blur on iris identification," Technical Report., Tech. Rep., 2011.
- [15] A. K. Chaudhary and J. B. Pelz, "Privacy-preserving eye videos using rubber sheet model," in *ACM Symposium on Eye Tracking Research and Applications*. ACM, 2020.
- [16] A. S. Narkar and B. David-John, "Swap it like its hot: Segmentation-based spoof attacks on eye-tracking images," in *Proceedings of the 2024 Symposium on Eye Tracking Research and Applications*. ACM, 2024.
- [17] M. Wang, E. Bozkir, and E. Kasneci, "Iris style transfer: Enhancing iris recognition with style features and privacy preservation through neural style transfer," *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, vol. 8, no. 2, 2025.
- [18] S. J. Garbin, Y. Shen, I. Schuetz, R. Cavin, G. Hughes, and S. S. Talathi, "OpenEDS: Open Eye Dataset," 2019. [Online]. Available: <https://arxiv.org/abs/1905.03702>
- [19] C. Palmero, A. Sharma, K. Behrendt, K. Krishnakumar, O. V. Komogortsev, and S. S. Talathi, "OpenEDS2020: Open Eyes Dataset," 2020. [Online]. Available: <https://arxiv.org/abs/2005.03876>
- [20] E. Bozkir, O. Günlü, W. Fuhl, R. F. Schaefer, and E. Kasneci, "Differential privacy for eye tracking with temporal correlations," *PLOS ONE*, vol. 16, no. 8, 2021.
- [21] J. Steil, I. Hagestedt, M. X. Huang, and A. Bulling, "Privacy-aware eye tracking using differential privacy," in *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*. ACM, 2019.
- [22] S. Özdel, E. Bozkir, and E. Kasneci, "Privacy-preserving scanpath comparison for pervasive eye tracking," *Proceedings of the ACM on Human-Computer Interaction*, vol. 8, no. ETRA, 2024.
- [23] B. David-John, D. Hosfelt, K. Butler, and E. Jain, "A privacy-preserving approach to streaming eye-tracking data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 5, 2021.
- [24] J. L. Kröger, O. H.-M. Lutz, and F. Müller, *What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking*. Springer International Publishing, 2020.
- [25] A. M. Eskildsen and D. Witzner Hansen, "Analysis of iris obfuscation: Generalising eye information processes for privacy studies in eye tracking," in *ACM Symposium on Eye Tracking Research and Applications*. Association for Computing Machinery, 2021.
- [26] R. P. Wildes, "Iris recognition: an emerging biometric technology," *Proceedings of the IEEE*, vol. 85, no. 9, pp. 1348–1363, 1997.
- [27] E. Ali, E. Ahmed, and A. Ali, "Recognition of human iris patterns for biometric identification," *Journal of Engineering and Applied Science-Cairo-*, vol. 54, no. 6, p. 635, 2007.
- [28] J. Daugman, "How iris recognition works," in *The essential guide to image processing*. Elsevier, 2009, pp. 715–739.
- [29] K. Rana, A. Pandey, P. Goyal, G. Singh, and P. Goyal, "A novel privacy protection approach with better human imperceptibility," *Applied Intelligence*, vol. 53, no. 19, pp. 21 788–21 798, 2023.
- [30] L. A. Gatys, "A neural algorithm of artistic style," *arXiv preprint arXiv:1508.06576*, 2015.
- [31] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2414–2423.
- [32] B. Gooch and A. Gooch, *Non-photorealistic rendering*. AK Peters/CRC Press, 2001.
- [33] T. Strothotte and S. Schlechtweg, *Non-photorealistic computer graphics: modeling, rendering, and animation*. Morgan Kaufmann, 2002.
- [34] Y. Li, N. Wang, J. Liu, and X. Hou, "Demystifying neural style transfer," *arXiv preprint arXiv:1701.01036*, 2017.
- [35] A. K. Chaudhary, R. Kothari, M. Acharya, S. Dangi, N. Nair, R. Bailey, C. Kanan, G. Diaz, and J. B. Pelz, "Ritnet: Real-time semantic segmentation of the eye for gaze tracking," in *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*. IEEE, 2019, pp. 3698–3702.
- [36] A. Zinollayev and I. Zhanseri. (2020) 1st place solution for facebook's opened 2020 sparse semantic segmentation challenge. [Online]. Available: https://github.com/zeenolife/openeds_2020
- [37] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [38] D. W. Hansen and Q. Ji, "In the eye of the beholder: A survey of models for eyes and gaze," *IEEE transactions on pattern analysis and machine intelligence*, vol. 32, no. 3, pp. 478–500, 2009.
- [39] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [40] Y. Jing, Y. Yang, Z. Feng, J. Ye, Y. Yu, and M. Song, "Neural style transfer: A review," *IEEE transactions on visualization and computer graphics*, vol. 26, no. 11, pp. 3365–3385, 2019.
- [41] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*. Springer, 2016, pp. 694–711.
- [42] X. Huang and S. Belongie, "Arbitrary style transfer in real-time with adaptive instance normalization," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 1501–1510.
- [43] A. Gupta, J. Johnson, A. Alahi, and L. Fei-Fei, "Characterizing and improving stability in neural style transfer," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 4067–4076.