PrivateEyes: Gaze-Preserving Anonymization for Data Sharing

Supplementary Material

This supplementary document provides additional information and experimental results to complement the main paper.

1. Additional Experiments

1.1. Influence of Guidance Scale and Diffusion Steps

To establish the optimal training and inference configuration for our diffusion model, we conducted systematic ablation studies analyzing the impact of guidance scale and diffusion steps, as presented in Section 4.2 of the main paper. Table 1 summarizes the performance comparison in terms of gaze errors and FID scores across three datasets. Based on our observations, a guidance scale of 3.0 combined with 20 diffusion steps provides an optimal balance between gaze accuracy and image fidelity.

1.2. Inference Runtime Evaluation

To evaluate runtime performance, we measure the per-image inference time required to generate an anonymized eye image. IST [9] requires approximately 14-15 seconds per image, while RSM [3] operates in under 1 second. Our method achieves a favorable balance between runtime and accuracy, requiring approximately 3-4 seconds per image.

1.3. Iris Recognition Under Varying Noise Levels

We further evaluated the robustness of our approach by testing task performance under different noise conditions, including Gaussian blur, Gaussian noise, and downsampling on the EV-Eye dataset [10]. As shown in Figure. 2, increasing the noise level degrade iris recognition performance (maintaining privacy), but this comes at the cost of reduced gaze estimation accuracy. Representative examples of these transformations and their combinations are shown in Figure 1.

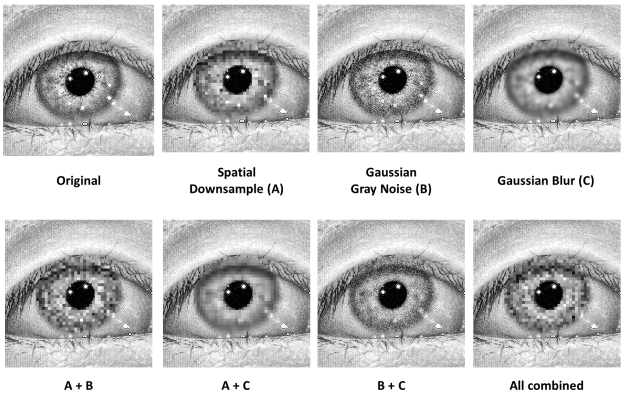

Figure 1. Images showing different traditional degradations for iris anonymization and their various combinations. Images are preprocessed to enhance brightness and contrast for better visualisations.

2. Additional Implementation Details

2.1. Eye Segmenter

The segmentation module is designed to generate a segmentation map for a given input eye image. Currently, it is trained to segment three distinct classes: pupil, iris, and sclera. However, the model can be extended to incorporate an additional class for segmenting glints, thereby improving its ability to capture finer details such as corneal reflections.

The architecture of the segmentation module is based on the GhostNet framework [5], which is optimized for computational efficiency through lightweight operations. This design choice ensures that the segmentation process remains fast while maintaining high accuracy.

One key challenge is the lack of pretrained segmentation models for the datasets used to evaluate eye anonymization methods such as IST [9], RSM [3], and our proposed method, PrivateEyes. To address this, we train a segmentation model using the ground-truth annotations available in OpenEDS [4] and LPW [8]. Because the EV-Eye [10] dataset does not provide segmentation labels, we further train the model on publicly available eye-segmentation datasets that include annotations for the pupil, iris, and sclera. For all experiments, segmentation masks are generated using our trained models and used consistently across all methods and datasets to ensure fair evaluation.

2.2. Anatomically Accurate 3D Eye Model

To have an eye model that closely matches with the structures of the real eye, we define anatomical properties of the eye as defined by GK Aguirre [1]. We implement a simplified and computationally efficient version of the Aguirre's eye model, accurately capturing anatomical and biometric parameters, including spherical refractive errors without degrading the anatomical accuracy of the eye model. The cornea is modeled as a tri-axial ellipsoid, following the "canonical representation" established by [6]. The anterior surface of the cornea is parameterized to vary according to the spherical refractive error, as described by [2]. The radii of curvature at the corneal vertex are adjusted linearly with the refractive error, yielding:

$$R_i = R_{i,0} \cdot (1 - k \cdot SR) \tag{1}$$

Guidance Scale	Diffusion Steps	OpenEDS [4]		EV-Eye [10]		LPW [8]	
		Gaze Error [°] ↓	FID ↓	Gaze Error [°] ↓	FID ↓	Gaze Error [°] ↓	FID ↓
3.0	5	5.89 ± 7.89	87.87	3.43 ± 3.70	52.29	4.27 ± 5.41	152.58
3.0	10	6.21 ± 7.80	85.21	4.19 ± 3.06	46.51	4.51 ± 5.62	135.24
3.0	15	6.11 ± 8.12	82.77	4.34 ± 4.20	44.00	4.27 ± 5.60	128.01
3.0	**20**	**6.06 ± 7.78**	**74.19**	**3.94 ± 3.52**	**41.73**	**4.2 ± 5.42**	**109.55**
7.5	5	6.77 ± 8.57	82.26	6.53 ± 4.92	86.34	5.47 ± 6.27	133.90
7.5	10	6.92 ± 8.86	79.06	5.58 ± 4.98	52.36	4.81 ± 5.72	139.71
7.5	15	6.17 ± 8.28	78.42	4.93 ± 4.10	47.74	4.67 ± 5.81	143.26
7.5	20	6.04 ± 7.85	112.32	4.36 ± 4.15	69.94	4.88 ± 6.05	255.50
12.0	5	7.20 ± 7.72	97.74	12.15 ± 9.29	196.00	8.68 ± 8.99	165.30
12.0	10	6.07 ± 6.98	88.32	8.71 ± 6.89	160.40	6.53 ± 7.87	126.05
12.0	15	7.39 ± 8.74	82.73	7.63 ± 6.05	115.15	5.25 ± 6.02	131.81
12.0	20	6.75 ± 8.34	80.23	7.22 ± 5.95	113.51	5.17 ± 6.07	260.85

Table 1. Ablation study evaluating the impact of guidance scale and diffusion steps on gaze accuracy and image fidelity. Experiments are conducted on a representative subset sampled from all datasets.

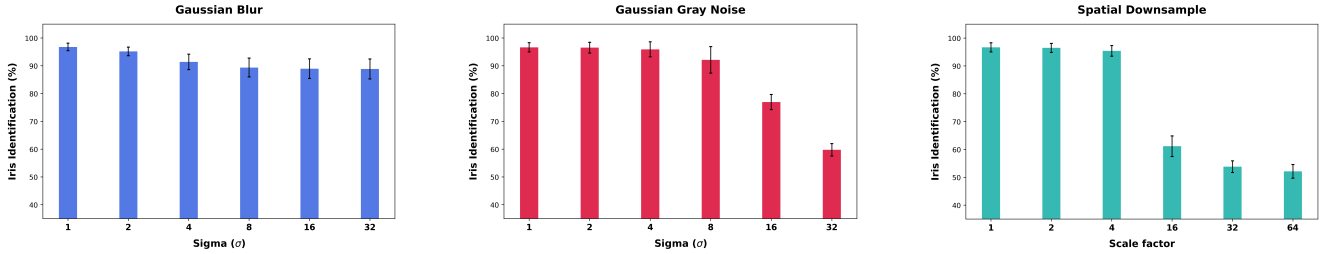

(a) Gaussian Blur (b) Gray Noise (c) Spatial Downsample

Figure 2. Effect of iris recognition accuracy on different parameters like sigma values and scale factors.

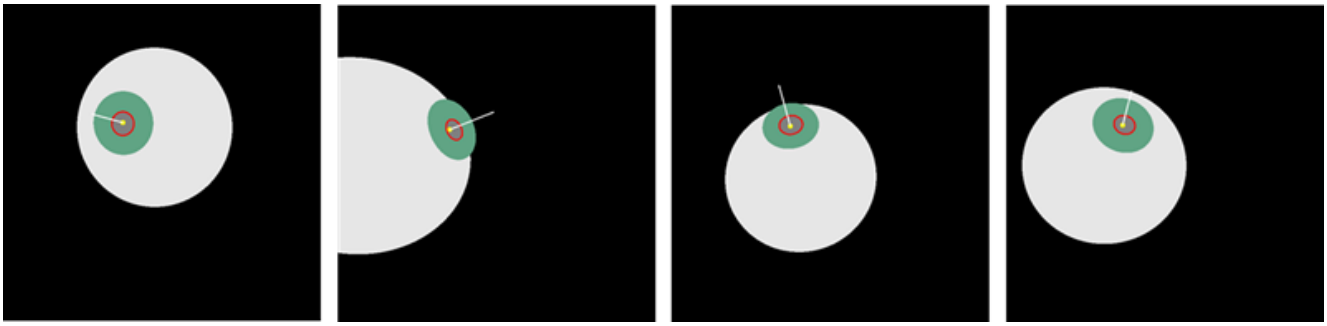

Figure 3. Synthetic data generated via ray tracing of an anatomically accurate eye model, illustrating the eyeball (white), iris (green), and pupil (gray), along with the gaze ray (white) originating from the pupil center (yellow).

where R_i represents the modified radius corresponding to the axial, horizontal and vertical directions, $R_{i,0}$ is the original radius for emmetropic eyes, k is a proportionality constant which is set to 0.0028, and SR is the spherical refractive error in diopters. For the posterior corneal surface, we adopt Atchison's parameters, keeping them constant across ametropias while scaling relative to the axial corneal radius defined by Navarro. The corneal thickness is set to 0.55 mm at the apex, ensuring anatomical consistency. This modeling approach allows for a realistic representation of individual ocular variations. [2] presents the radii of curvature and asphericity for a bi-conic model of the vitreous chamber, which is both de-centered and tilted relative to the visual axis. As negative refractive error (myopia) increases, the vitreous chamber undergoes dimensional changes, notably elongating in the axial direction and, to a lesser extent, in both the vertical and horizontal dimensions. The vitreous chamber is modeled as an ellipsoid that is centered

and aligned with the optical axis.

In Aguirre's model, the iris is represented as a plane that is perpendicular to the optical axis, corresponding to a zero "iris angle". It is situated 3.9 mm behind the anterior surface of the cornea, aligning with the anterior surface of the lens. The inner edge of the iris serves as the boundary for the entrance pupil's image. The model assumes that the aperture stop is both centered and fixed along the optical axis. The crystalline lens is modeled through a series of quadric surfaces. Specifically, its anterior and posterior surfaces are approximated as half of a two-sheeted hyperboloid.

2.3. Simulation of the Eye Model with Backward Ray Tracing

The backwards ray tracing module is designed for the generation of synthetic data for training the PEN (Pose Estimation Network) due to difficulty in capturing or recording the 3D pose of the eye from the real captured data. Using this method, we generate anatomically accurate images of the structures of the eye for the desired pose based on the simplified version of the GK Aguirre's eye model. Samples images generated using this module have been illustrated in Figure 3.

3. Discussion and Design Considerations

3.1. Assumed Privacy Threat Model

We consider a threat model where the attacker has access to the released anonymized eye dataset and non-anonymized images of known individuals, with the goal of linking anonymized samples to real identities. As reported in Table 1 of main paper, our re-identification results show very low matching rates for both iris and full-eye regions, effectively mitigating this threat. A similar threat model is commonly adopted in face anonymization literature [7].

3.2. Compatibility with other generative models

A key contribution of this work is PEN model, which captures geometric properties of the eye and guides generative model for eye image generation. While our experiments use Stable Diffusion (SD) with ControlNet for its explicit structural control and high-fidelity synthesis, our framework is adaptable to any better/newer SOTA generative models.

References

[1] Geoffrey K Aguirre. A model of the entrance pupil of the human eye. *Scientific reports*, 9(1):9360, 2019. 1

[2] David A Atchison. Optical models for human myopic eyes. *Vision research*, 46(14):2236–2250, 2006. 1, 2

[3] Aayush Kumar Chaudhary and Jeff B Pelz. Privacy-preserving eye videos using rubber sheet model. In *ACM Symposium on Eye Tracking Research and Applications*, pages 1–5, 2020. 1

[4] Stephan J Garbin, Yiru Shen, Immo Schuetz, Robert Cavin, Gregory Hughes, and Sachin S Talathi. Openeds: Open eye dataset. *arXiv preprint arXiv:1905.03702*, 2019. 1, 2

[5] Kai Han, Yunhe Wang, Qi Tian, Jianyuan Guo, Chunjing Xu, and Chang Xu. Ghostnet: More features from cheap operations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1580–1589, 2020. 1

[6] Rafael Navarro, Luis González, and José L Hernández. Optics of the average normal cornea from general and canonical representations of its surface topography. *Journal of the Optical Society of America a*, 23(2):219–232, 2006. 1

[7] Roland Stenger, Steffen Busse, Jonas Sander, Thomas Eisenbarth, and Sebastian Fudickar. Evaluating the impact of face anonymization methods on computer vision tasks: A trade-off between privacy and utility. *IEEE Access*, 13:11070–11079, 2024. 3

[8] Marc Tonsen, Xucong Zhang, Yusuke Sugano, and Andreas Bulling. Labelled pupils in the wild: a dataset for studying pupil detection in unconstrained environments. In *Proceedings of the ninth biennial ACM symposium on eye tracking research & applications*, pages 139–142, 2016. 1, 2

[9] Mengdi Wang, Efe Bozkir, and Enkelejda 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*, 8(2):1–21, 2025. 1

[10] Guangrong Zhao, Yurun Yang, Jingwei Liu, Ning Chen, Yiran Shen, Hongkai Wen, and Guohao Lan. Ev-eye: rethinking high-frequency eye tracking through the lenses of event cameras. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2023. Curran Associates Inc. 1, 2