PrivateEyes: Gaze-Preserving Anonymization for Data Sharing

Surabhi Gupta[1*] Dinesh Prabhu Muthumariappan[1*] Biplab Das[1]
Anoop Kolar Rajagopal[1] Kiran Nanjunda Iyer[1] Donghwan Seo[2]

[1]Samsung R&D Institute-Bangalore (India) [2]Samsung Electronics

Abstract

Eye images captured from wearable devices such as head-mounted displays (HMDs) contain identifiable biometric cues, posing significant challenges for safe data sharing. Existing eye anonymization techniques often degrade downstream performance, particularly gaze estimation while still retaining iris-recognizable features. Although these methods aim to anonymize the iris, they introduce noticeable visual artifacts that reduce image fidelity. To address these limitations, we propose ***PrivateEyes****, a privacy-preserving framework that synthesizes anonymized yet gaze-consistent eye images. Our approach employs a three-stage pipeline: (1) a deep segmentation network that isolates semantic eye regions and provides structural control signals for synthesis, (2) a pose estimation network (PEN) trained on anatomically accurate synthetic eye renders to infer precise eye pose, and (3) a conditional diffusion model that reconstructs realistic, anonymized eye images conditioned on segmentation and pose. Extensive experiments across multiple benchmark datasets show that PrivateEyes achieves superior gaze-estimation accuracy compared to state-of-the-art anonymization baselines, improving performance by over 10% while reducing iris-recognition accuracy by ∼50%. Our method also produces higher-fidelity images compared to other existing approaches. By enabling task-preserving and privacy-secure sharing of eye images, PrivateEyes supports responsible research and development in AR/VR and other gaze-driven applications.*

1. Introduction

Eye gaze serves as a crucial modality for foveated rendering [3, 21], gaze-driven interfaces, and immersive interactions [5, 20] in wearable systems such as head-mounted displays (HMDs). As these applications scale, the demand for large, diverse, and publicly shareable eye-image datasets has become essential for training robust AI models. However, privacy concerns remain a major bottleneck, as eye images inherently encode rich biometric iris textures that can uniquely identify individuals. This raises compliance challenges under regulations such as General Data Protection Regulation (GDPR), which restrict the sharing of personally identifiable biometric data. While large-scale face datasets have fueled rapid progress in facial analysis, eye-related research continues to be hindered by the scarcity of publicly shareable eye datasets due to privacy risks. This directly impacts reproducibility and slows advancements in gaze estimation, eye segmentation, and HMD-based interaction research.

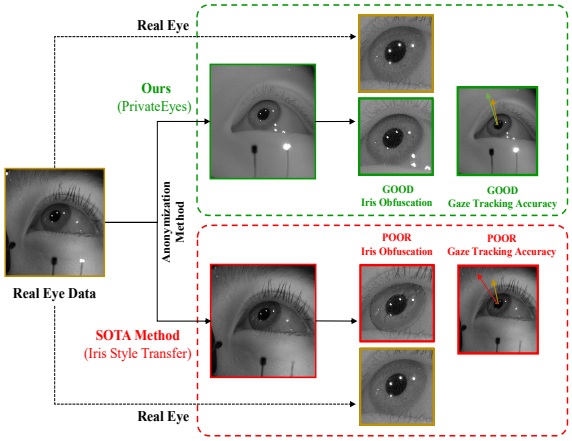

Figure 1. **PrivateEyes:** Our proposed approach (highlighted in green) synthesizes a gaze-preserving eye image from an input eye image while removing iris-specific identity information. In contrast, latest SOTA anonymization method (highlighted in red) degrade gaze accuracy while still retaining identifiable iris features.

Despite extensive research in facial anonymization [7, 14], anonymization of eye-region imagery remains significantly underexplored. Recent works in ocular privacy adopt either geometry-based or appearance-based strategies. Geometry-based methods such as the Rubber Sheet Model (RSM) [6] replace the iris via polar coordinate warping, removing some identity cues but introducing boundary artifacts and distorting geometric structures critical for gaze

*Equal contribution.

inference. Appearance-based approaches such as Iris Style Transfer (IST) [30] modify iris texture using neural style transfer; however, these methods struggle to maintain photorealism, illumination consistency, and accurate gaze geometry. As a result, they exhibit reduced fidelity and degrade performance on downstream eye-related tasks. Moreover, these methods typically rely on end-to-end frameworks, which limits controllability and interpretability.

We address these challenges by introducing a gaze-guided, diffusion-based anonymization framework specifically designed for safe data sharing in eye-tracking applications. We adopt a geometry-grounded modular pipeline over an end-to-end generator to provide explicit control over gaze interpretability and enable targeted improvements without full retraining. A key component of our approach is the ***Pose Estimation Network (PEN)*** that enables gaze-aware conditioning during image synthesis. As illustrated in Figure 1, our method synthesizes high-fidelity, gaze-preserving eye images that obfuscate identity-revealing eye information. We first segment key eye regions such as iris, pupil, and sclera and estimate eye pose information using our Pose Estimation Network (PEN). We then employ a diffusion-based ControlNet [34] conditioned on the segmentation map and estimated pose to generate photorealistic anonymized eye images. A major challenge in training PEN lies in obtaining ground-truth eye pose for real eye images. To overcome this, we develop a two-stage synthetic data pipeline: (1) an anatomically accurate 3D eye model incorporating optical effects such as refraction, and (2) a ray-traced rendering system that generates precise semantic segmentations across a wide distribution of gaze directions.

Our proposed method effectively anonymizes eye images while preserving gaze-estimation accuracy enabling safe data release for third-party research and model training. To further mitigate privacy risks associated with sharing trained models such as training sample memorization and membership inference [9, 12, 18], we emphasize the release of anonymized images rather than model parameters. By addressing the intersection of biometric privacy, gaze fidelity, and task utility, our work contributes to responsible and ethical deployment of eye-tracking technologies. To the best of our knowledge, our framework is the first diffusion-based model to enable large-scale generation of gaze-preserving, anonymized eye data. In contrast to geometry-based methods such as the Rubber Sheet Model and neural style transfer approaches like IST, our method synthesizes photorealistic, high-fidelity eye images suitable for a wide range of downstream tasks. Our key contributions are as follows:

- **Privacy-Compliant Gaze Data Sharing:** We propose ***PrivateEyes***, a framework that anonymizes eye images by removing biometric iris cues while preserving gaze information.
- **Anatomically Accurate 3D Eye Modeling:** We design a 3D eye model and employ inverse ray tracing to generate physically grounded gaze annotations, enabling accurate and robust pose estimation.
- **Extensive Evaluation:** We conduct comprehensive experiments across 3 benchmark datasets, demonstrating superior gaze preservation, strong privacy protection, and improved image fidelity compared to existing methods.

2. Related Works

In this section, we review prior work on privacy concerns in eye-tracking, eye-image anonymization, and diffusion-based generative models.

2.1. Privacy Concerns in Eye-Tracking Data

Eye tracking is integral to AR/VR and HCI systems, but gaze data encodes sensitive biometric, behavioral, and even medical information. Prior studies show that users remain concerned about misuse [19, 28], highlighting the need for privacy-preserving methods that retain utility. Statistical approaches such as differential privacy [19, 28] provide protection at an aggregate level but are insufficient for image-based tasks like gaze estimation or segmentation, where per-image information must be preserved. A recent survey [5] emphasizes the lack of solutions that jointly maintain privacy and task performance which is directly addressed in this work. Complementary efforts generate synthetic eye data for privacy-preserving training [4]; however, our approach focuses on scenarios where real images must be shared while maintaining downstream task fidelity.

2.2. Anonymization Techniques for Eye Images

Classical geometry-based approaches such as the Rubber Sheet Model (RSM) [6] and its extensions [2] replace the iris via polar unwrapping and re-projection, effectively removing biometric identity but often producing boundary distortions and inconsistent reflections. Appearance-based methods like Iris Style Transfer (IST) [30] anonymize iris texture through neural style transfer, yet frequently yield color mismatches, texture artifacts, and blurred pupil-sclera boundaries, which degrade photorealism and negatively impact downstream performance. A benchmarking study [31] quantitatively analyzes the privacy-utility trade-off of common iris obfuscation techniques and shows that naive perturbations such as RSM and IST significantly degrade gaze and segmentation accuracy. Eskildsen et al. [10] propose a standardized evaluation protocol for privacy and utility in eye-image obfuscation; we adopt a similar methodology in our benchmarking.

2.3. Image Diffusion Models

Diffusion models [16, 26, 27] have emerged as state-of-the-art generative frameworks producing high-fidelity images

through iterative denoising. Conditional variants guided by text [24, 25] or structural priors such as edges, depth, or segmentation maps [17, 34] enable fine-grained control over synthesis. In this work, we adopt ControlNet [34] to condition diffusion on both eye segmentation and gaze direction, enabling identity-free yet gaze-consistent synthesis. This demonstrates the potential of diffusion-based models for privacy-preserving eye-image generation with strong support for downstream computer vision tasks.

3. Methodology

In this section, we detail **PrivateEyes**, a privacy-preserving generative pipeline that synthesizes identity-free, gaze-consistent eye images to facilitate data sharing.

3.1. Eye Segmenter

We adopt the GhostNet architecture [13] to segment the input eye image into three regions: iris, pupil, and sclera. Ellipse parameters for the iris and pupil are then estimated and passed to the Eye Pose Estimator (details in Section 3.2). The segmentation map also serves as a structural guide for the ControlEye module.

3.2. Eye Pose Estimator

ControlEye can generate anonymized images using only the semantic segmentations from the Eye Segmenter as the control image. However, the 2D segmentations alone are insufficient to effectively retain the 3D optical and pose information of the eye during generation. Hence, we feed the 3D pose information of the eye to the controlEye as an additional gaze control to retain the 3D information of the eye. However, estimating the 3D eye pose from real captured eye data is challenging. Hence, we generate synthetic eye data that closely mimics the real eye using ray tracing and train a model to estimate the 3D eye pose from the semantic mask. This trained model is used on real data captures to estimate the eye pose information.

3.2.1. Anatomically Accurate 3D Eye Model

To generate the synthetic eye data, we base our eye model on the work of Aguirre [1], as it closely matches the structure of the real human eye. We implement a simplified and computationally efficient version of the model by representing the aperture stop as circular and omitting the crystalline lens, still accurately capturing anatomical and biometric parameters without compromising the anatomical accuracy. The cornea is modeled following the "canonical representation" established by Navarro et al. [22]. This anatomical modeling allows us to represent individual variations in ocular geometry effectively.

3.2.2. Eye Model simulation with Backward Ray Tracing

We employ a backward ray-tracing technique that projects the anatomical eye model onto a virtual camera's imaging plane generating the synthetic eye data. This approach allows us to trace rays from the camera back to the modeled eye surfaces, which are represented as quadric surfaces defined by their optical and geometric properties.

Ray Propagation. The ray-tracing process starts by defining rays that originate from the camera sensor and propagate backward through the optical system of the eye. We can express the position of a ray in 3D space as:

$$\mathbf{R}(t) = \mathbf{R}_0 + t \cdot \mathbf{r} \tag{1}$$

where $\mathbf{R}_0$ is the camera position, $\mathbf{r}$ is the ray direction vector, and t is the distance along the ray.

Intersection with Quadric Surfaces. The eye surfaces, including the cornea and lens, are modeled as quadric surfaces defined by the equation:

$$\frac{(x-h)^2}{a^2} + \frac{(y-k)^2}{b^2} + \frac{(z-d)^2}{c^2} = 1 \tag{2}$$

where (h, k, d) are the surface centers, and a, b, c are the semi-axis lengths. To calculate the ray-surface intersections, we substitute the ray's parametric equation (1) into the quadric surface equation and solve for the distance, t at which the ray intersects the surfaces, as given below:

$$\frac{(x_0 + t \cdot r_x - h)^2}{a^2} + \frac{(y_0 + t \cdot r_y - k)^2}{b^2} + \frac{(z_0 + t \cdot r_z - d)^2}{c^2} = 1 \tag{3}$$

Refraction at the Surface. When the ray intersects a refractive surface, such as the cornea, we compute the refracted ray direction $\mathbf{r}'$ using the normal vector $\mathbf{N}$ and the refractive indices n_1 and n_2 of the media, as:

$$\mathbf{r}' = \frac{n_1}{n_2}\mathbf{r} - \left(\frac{n_1}{n_2}\cos(\theta_1) - \cos(\theta_2)\right)\mathbf{N} \tag{4}$$

θ_1,θ_2 are angles of incidence and refraction, respectively.

Image Formation. Our backward ray tracer iteratively computes ray intersections and refractions until terminating at the aperture stop or a non-refractive surface. The originating pixel is labeled with the final intersected structure (e.g., eyeball, iris, pupil), enabling accurate simulation of ocular optics and generation of eye structure images. This approach produces 2D segmentation maps directly from 3D geometry, without appearance modeling or shading. The resulting images are purely structural, like manually labeled masks shown in Figure 2. This step is computationally optimized, taking $< 150ms/$image. The rendered structural images are then used to estimate the pupil and iris ellipse parameters, which serve as input to the next pipeline stage.

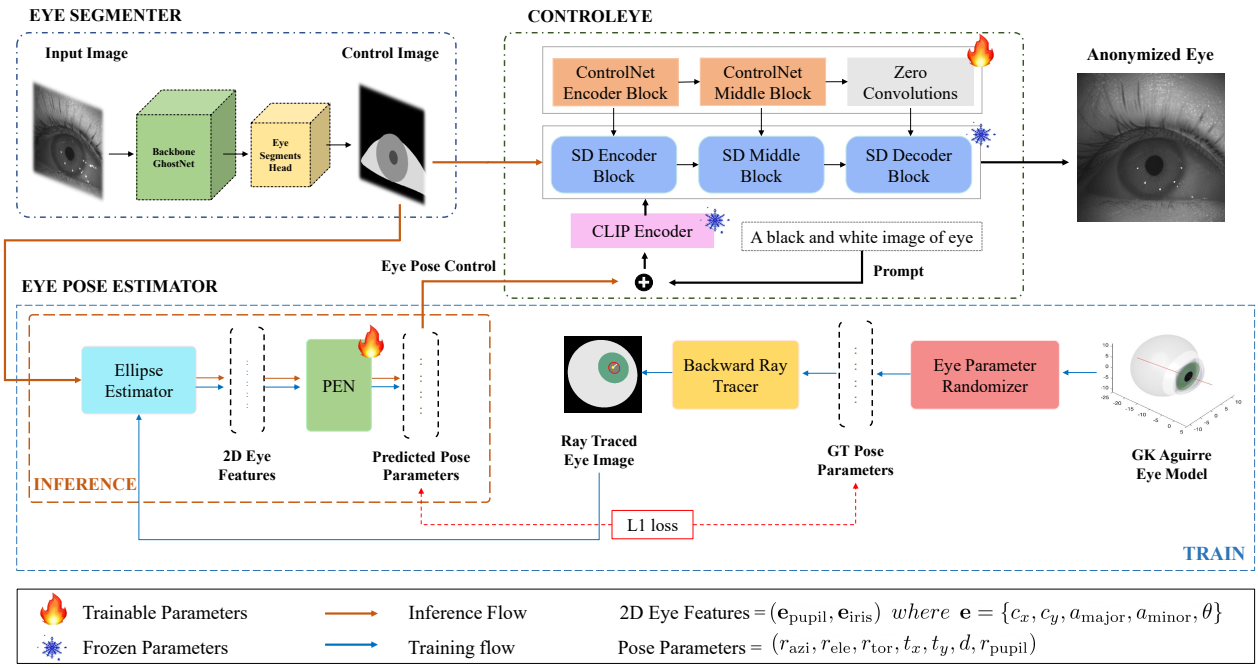

Figure 2. **Proposed Method.** The pipeline consists of three key modules: **1) Eye Segmenter**: Utilizes a GhostNet backbone to segment the eye into pupil, iris, and sclera regions. **2) Eye Pose Estimator**: Constructs an anatomically accurate 3D eye model [1] and applies backward ray tracing to generate images with varying continuous eye poses. Ellipse features extracted from these images are fed into an PEN to predict different eye poses, forming the training step. **3) ControlEye**: During inference, ellipse parameters are estimated from the segmented image to generate an eye pose embedding. This embedding is then processed by a CLIP encoder to obtain a control vector, which is used by a fine-tuned ControlNet [34] to synthesize a realistic eye image that preserves gaze while anonymizing identity.

3.2.3. Eye Pose Estimation via Neural Inversion

In this section, we present a method for estimating eye pose by combining ray tracing and neural networks. Based on the ray tracing method and the eye model described earlier, the 3D pose of the eye model (rotational angles and positions) are randomly set using a eye parameter randomizer to construct a segmentation map of the eye features. This segmentation map is then used to estimate the pupil and iris ellipses. Leveraging our approach, we generate a large-scale dataset of 6 million synthetic eye images with rich geometric and pose variations, along with ellipse annotations. This dataset enables the training of our Pose Estimation Network (PEN), which learns to invert ray tracing and infer 3D eye pose from the iris and pupil ellipse information, crucial for gaze-preserving anonymization.

We employ the PEN, a multi-layer perceptron, to map 2D eye features in the image space to the 3D pose of the eye with respect to the camera. This mapping is given by:

$$\mathbf{pose} = \mathbf{PEN}(\mathbf{f}_{\text{eye, 2D}})$$

where $\mathbf{f}_{\text{eye, 2D}} = (\mathbf{e}_{\text{pupil}}, \mathbf{e}_{\text{iris}})$ represents the input features derived from the pupil and iris ellipses. Each ellipse is parameterized as $\mathbf{e} = \{c_x, c_y, a_{\text{major}}, a_{\text{minor}}, \theta\}$, where (c_x, c_y) are the ellipse center coordinates, a_{major} and a_{minor} are the lengths of the major and minor axes, and θ is the orientation of the major axis. The output of the PEN is the 3D pose of the eye, represented as:

$$\mathbf{pose} = (r_{\text{azi}}, r_{\text{ele}}, r_{\text{tor}}, t_x, t_y, d, r_{\text{pupil}})$$

Here, $(r_{\text{azi}}, r_{\text{ele}}, r_{\text{tor}})$ are the azimuthal, elevational, and torsional rotation angles, respectively. (t_x, t_y, d), are the horizontal, vertical, and depth positions of the eye relative to the camera and r_{pupil} is the radius of the pupil.

Once trained on synthetic data, the PEN is applied to real eye images, providing 3D pose estimates improving gaze-preserving anonymized image generation.

3.3. ControlEye

ControlEye is built on ControlNet [34], leveraging eye-specific control signals - segmentation maps and gaze pose parameters to guide diffusion-based synthesis while ensuring gaze retention. Our control signal is identity-free by design (see Figure 2). The segmentation map from Section 3.1 provides spatial constraints, while a control term C enriches the generation with gaze intent. This term combines eye pose embeddings from our Eye Pose Estimator with high-level descriptors (e.g., eye type), processed via a CLIP [33] text encoder and injected into the Stable Diffusion (SD) blocks. Following ControlNet, feature modulation is applied across 12 trainable U-Net encoder blocks

and one middle block, spanning four resolutions (64×64, 32×32, 16×16, 8×8), each with three sub-blocks. These outputs are merged with their decoder counterparts via zero convolutions and skip connections.

We further introduce a 3D gaze-alignment constraint, encoded using CLIP and injected as a text-conditioning signal, to enforce consistency between the generated gaze direction and the original input. The denoising function is modified to include the control term C, as:

$$x_{t-1} = x_t - \epsilon_\theta(x_t, t, y, C) + \sqrt{\beta_t} z \tag{5}$$

Here, x_t is the noisy input at timestep t, y is the text condition, $z \sim \mathcal{N}(0, I)$ is Gaussian noise, and β_t controls the variance schedule.

The corresponding training loss ensures gaze fidelity while allowing identity removal through texture variation:

$$\mathcal{L}_{ControlEye} = \mathbb{E}_{x_0, C, t, \epsilon} \left[|\epsilon - \epsilon_\theta(x_t, t, y, C)|^2 \right] \tag{6}$$

where ϵ is the sampled noise, and ϵ_θ is the model's prediction. This formulation ensures privacy-preserving yet gaze-consistent image generation.

3.4. Handling Corneal Reflections (Glints)

Corneal reflections (glints) serve as important photometric cues in many eye-tracking pipelines, particularly for model-based gaze estimators. To ensure compatibility with such systems, we first inpaint glints using off-the-shelf inpainting methods during pre-processing and feed the glint-free images to our diffusion model. After synthesis, the removed glints are reintroduced to preserve gaze-relevant illumination cues. Figure 3 illustrates the glint extraction and reinsertion process.

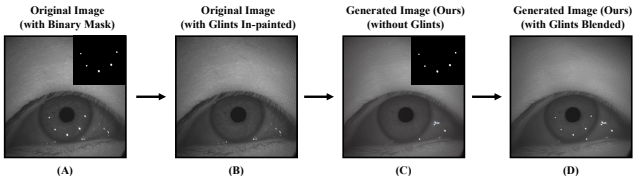

Figure 3. Overview of the glint-handling pipeline: (A) Input eye image, (B) Glint inpainted image, (C) Glint-free generated image by our method, and (D) Final output with blended glints.

4. Experiments

This section provides the experimental setup, datasets, preprocessing steps, and implementation details.

4.1. Benchmark Datasets

We evaluate our method on three widely used real-world eye-tracking datasets: OpenEDS2019 [11], EV-Eye [36], and LPW [29]. These datasets provide high-quality infrared (IR) eye imagery suitable for assessing both gaze estimation and iris recognition. Although OpenEDS2020 was considered, it is no longer publicly accessible[1].

- **OpenEDS2019** [11] consists of 12,759 IR images collected from 152 subjects using a VR headset with dual eye cameras, along with pixel-level semantic labels. Following prior work, we use 110 subjects for training and 42 for testing. Images are cropped to 400×400 for all experiments. For iris recognition and image-quality evaluation, iris regions are further cropped to 180×180.
- **EV-Eye** [36] contains over 2 million IR eye images from 48 subjects across multiple sessions, capturing broad gaze and appearance variability. We adopt the official split of 30 subjects for training and 18 for testing. Images are resized to 256×256, and iris patches are cropped to 90×90 for recognition and fidelity evaluations.
- **LPW** [29] (Labelled Pupils in the Wild) provides over 130,000 in-the-wild eye images from 22 subjects recorded with a head-mounted tracker. It includes diverse lighting, motion, and gaze variations. Excluding 2 subjects wearing glasses, we use 12 subjects for training and 8 for testing. Images are cropped to 480×480, and iris regions to 300×300 for evaluation consistency.

4.2. Implementation Details

Training and Inference Settings. The diffusion model is implemented in PyTorch and trained on an NVIDIA Tesla V100 GPU (32 GB). It is optimized with a learning rate of 10^{-4} and a batch size of 4. We adopt the DDPM [16] scheduler using a linear noise schedule from 1×10^{-4} to 2×10^{-2}. The number of diffusion steps is set to $T_{\text{diff}} = 1000$ during training and reduced to 20 at inference for faster sampling. A guidance scale of 3.0 is used, chosen based on ablation studies *(details in the supplementary material)*. ControlNet [34] is integrated into the Stable Diffusion U-Net via zero-convolution layers for conditional modulation. For consistency across experiments, we use fixed CLIP-encoded text prompts. For ray-traced synthetic data generation, we set the intrinsic parameters of the virtual camera to closely approximate the camera parameters for each dataset.

Baseline Methods. We evaluate PrivateEyes against existing anonymization methods that operates on the eye region. Rubber Sheet Model (RSM) [6] uses geometric warping to replace iris textures while preserving gaze, but often introduces boundary artifacts and illumination inconsistencies. Iris Style Transfer (IST) [30] applies neural style transfer to anonymize iris appearance while retaining geometry, yet can cause structural distortions. Conventional approach like Blur and Downsample [23] methods obscure identity but substantially degrade gaze accuracy and image quality.

[1] As confirmed by the dataset authors.

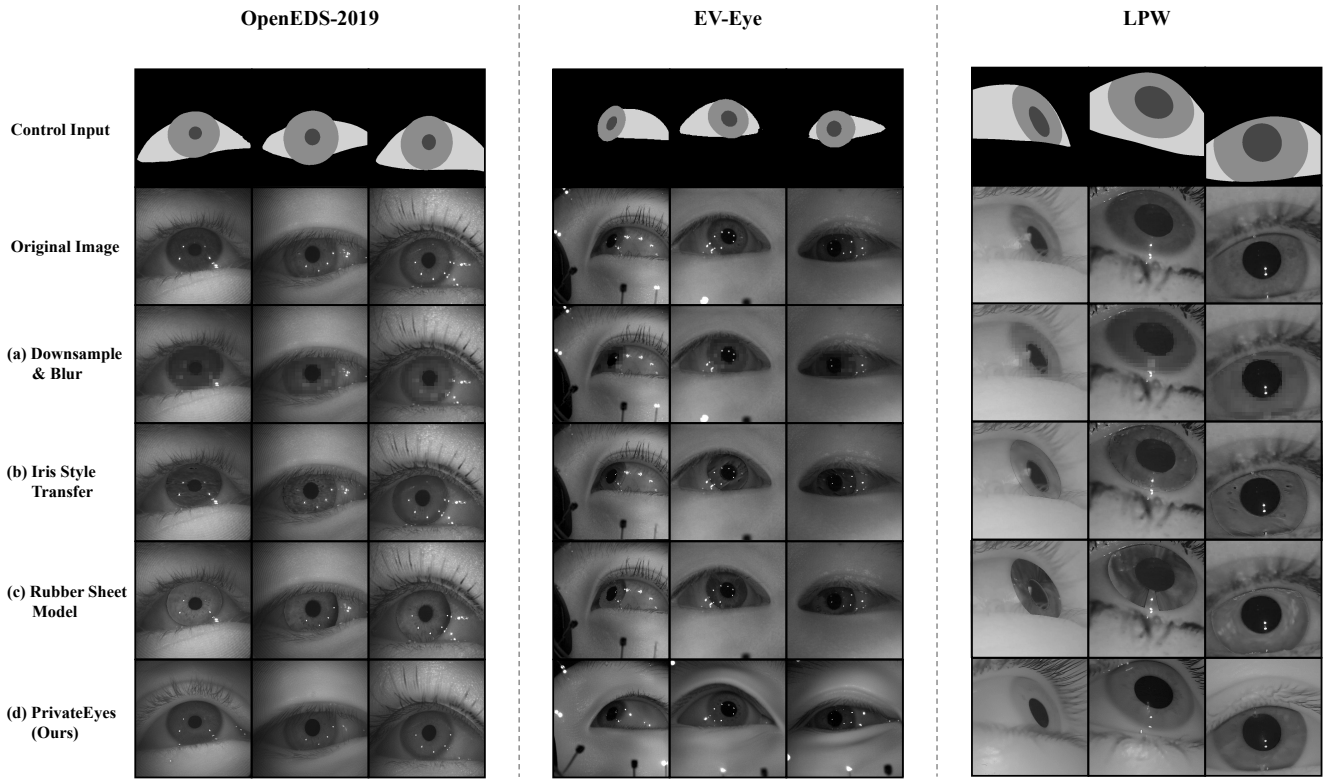

Figure 4. **Image Quality.** Qualitative comparison of anonymized eye images across multiple evaluation datasets. From top to bottom: Control input image, Original image, and baseline methods (a) Downsampling with Gaussian blur [23], (b) Iris Style Transfer [30], (c) Rubber Sheet Model [6], and (d) **PrivateEyes (Ours)**. **Zoom in** for improved visibility.

Metrics. To evaluate image quality, we use standard metrics: Fréchet Inception Distance (FID) [15], Kernel Inception Distance (KID) [32], and Learned Perceptual Image Patch Similarity (LPIPS) [35]. Task utility is assessed via gaze estimation error (angular deviation in °) and eye-region segmentation (mean Intersection-over-Union, mIoU). Identity preservation is measured using iris recognition accuracy (%), reflecting the degree of anonymization.

Gaze Ground-Truth Generation and Metric. We generate the ground-truth gaze data using the geometric-based Dierkes model [8], which estimates gaze vectors by leveraging key eye parameters including the eyeball center, pupil center, and eyeball rotation. We evaluate the accuracy of our method using the standard angular error metric, which measures the deviation between the predicted, y, and ground-truth gaze vectors, x.

$$Gaze\ Error(x, y) = \cos^{-1} \frac{x^{\mathrm{T}} y}{||x||\ ||y||} \tag{7}$$

Iris Recognition and Identity Obfuscation. To assess identity obfuscation, we implement an iris recognition model using a classification-based approach, where each user's iris pattern is mapped to a unique class. We use GhostNet [13] as the backbone for its efficient feature extraction, enabling lightweight yet high-performing biometric identification. We preprocess eye images by automatically segmenting and cropping the iris region to minimize background noise. This ensures that the model focuses on discriminative iris features, improving recognition accuracy while isolating identity-relevant patterns.

5. Results

This section presents both qualitative and quantitative evaluation results across multiple benchmark datasets, comparing PrivateEyes with existing state-of-the-art anonymization methods. We also perform detailed ablations on different aspects that are relevant to our work.

5.1. Qualitative and Quantitative Evaluations

Quantitative Results. Effective privacy-preserving gaze tracking requires maintaining a balance between strong anonymization (low iris recognition accuracy) and high task utility (low gaze error). Table 1 quantitatively evaluates this trade-off across various perturbation-based baselines, demonstrating that PrivateEyes achieves superior gaze preservation while effectively suppressing identity information. Our method surpasses baselines on EV-Eye for

Method	OpenEDS [11]			EV-Eye [36]			LPW [29]		
	Gaze Error [°] ↓	Iris Recog. (C) [%] ↓	Iris Recog. (F) [%] ↓	Gaze Error [°] ↓	Iris Recog. (C) [%] ↓	Iris Recog. (F) [%] ↓	Gaze Error [°] ↓	Iris Recog. (C) [%] ↓	Iris Recog. (F) [%] ↓
Gaussian Blur ($\sigma = 11$)	**3.88 ± 6.0**	47.2 ± 28.8	99.1 ± 0.2	4.01 ± 5.6	71.0 ± 16.9	99.3 ± 1.1	3.01 ± 3.9	98.8 ± 2.7	99.9 ± 0.2
Gaussian Gray Noise ($\sigma = 21$)	6.23 ± 8.0	46.1 ± 30.9	99.8 ± 0.2	4.12 ± 5.4	81.2 ± 21.7	99.8 ± 0.8	2.74 ± 3.6	95.1 ± 13.8	99.7 ± 0.4
Spatial Downsample (scale = 8)	7.86 ± 8.7	57.8 ± 22.5	99.8 ± 0.6	7.55 ± 8.6	95.9 ± 24.2	99.5 ± 0.3	3.19 ± 3.9	99.5 ± 13.1	100 ± 0.1
Gaussian Blur + Gaussian Gray Noise	7.91 ± 8.6	27.7 ± 27.1	99.4 ± 0.6	4.85 ± 6.0	63.7 ± 14.1	99.5 ± 1.6	2.52 ± 3.4	89.0 ± 24.5	99.7 ± 0.2
Gaussian Blur + Spatial Downsample [23]	7.77 ± 8.5	38.9 ± 25.9	99.8 ± 1.6	7.44 ± 8.2	85.4 ± 25.1	99.6 ± 1.7	3.11 ± 3.9	99.1 ± 3.0	99.5 ± 0.9
Gaussian Gray Noise + Spatial Downsample	12.60 ± 11.7	13.1 ± 11.5	99.9 ± 2.7	13.75 ± 12.3	93.0 ± 23.4	98.9 ± 2.2	3.76 ± 4.3	54.1 ± 36.4	99.1 ± 1.3
All Combined	13.42 ± 12.1	**11.4 ± 10.6**	99.1 ± 4.5	14.41 ± 12.2	68.1 ± 25.3	98.9 ± 2.2	3.86 ± 4.5	53.6 ± 37.1	99.2 ± 1.1
Iris Style Transfer [30]	7.65 ± 8.2	78.1 ± 17.6	90.5 ± 14.3	4.8 ± 6.1	90.0 ± 19.8	96.3 ± 12.1	**1.91 ± 2.59**	81.1 ± 23.8	97.5 ± 8.5
Rubber Sheet [6]	5.69 ± 7.4	38.6 ± 20	87.0 ± 14.2	3.52 ± 6.2	69.0 ± 16.8	89.9 ± 13.6	2.62 ± 3.76	79.4 ± 25.4	95.6 ± 7.1
PrivateEyes **(Ours w/o PEN)**	4.94 ± 6.5	16.1 ± 15.1	14.4 ± 14.1	3.33 ± 3.6	18.3 ± 14.4	16.8 ± 15.5	3.27 ± 3.1	19.8 ± 16.5	**17.2 ± 13.5**
PrivateEyes **(Ours w/ PEN)**	4.43 ± 6.2	16.5 ± 15.2	**14.2 ± 14.3**	**2.89 ± 3.4**	**17.7 ± 14.2**	**16.7 ± 15.1**	2.51 ± 3.0	**18.1 ± 15.9**	17.3 ± 14.2

Table 1. **Image Utility.** Quantitative comparison of our method against existing eye-anonymization baselines on the OpenEDS [11], EV-Eye [36], and LPW [29] datasets. PrivateEyes achieves substantially lower gaze-estimation error while also reducing iris-recognition accuracy (indicating stronger anonymization), demonstrating superior task utility across all benchmarks. Note that the gaze evaluation is performed on full eye images, whereas iris-recognition accuracy is computed on both cropped (represented in C) and full iris regions (represented in F). Best results are shown in **bold**, and second-best results are underlined.

Method	OpenEDS [11]			EV-Eye [36]			LPW [29]		
	FID ↓	KID ↓	LPIPS ↓	FID ↓	KID ↓	LPIPS ↓	FID ↓	KID ↓	LPIPS ↓
Gaussian Blur ($\sigma = 11$)	102.11	0.117	0.273	54.51	0.058	0.156	**52.83**	**0.028**	0.153
Gaussian Gray Noise ($\sigma = 21$)	130.81	0.157	0.291	157.20	0.207	0.261	138.54	0.124	0.335
Spatial Downsample (scale = 8)	176.80	0.232	0.371	154.03	0.182	0.257	54.80	0.032	**0.138**
Gaussian Blur + Gaussian Gray Noise	153.35	0.193	0.333	180.98	0.247	0.290	143.10	0.126	0.351
Gaussian Blur + Spatial Downsample [23]	184.39	0.243	0.346	149.16	0.176	0.249	62.86	0.038	0.152
Gaussian Gray Noise + Spatial Downsample	212.47	0.292	0.459	155.20	0.183	0.321	104.49	0.078	0.295
All Combined	212.14	0.291	0.458	150.96	0.178	0.324	101.67	0.077	0.294
Iris Style Transfer [30]	86.01	0.095	0.181	60.91	0.065	**0.125**	89.92	0.066	0.203
Rubber Sheet [6]	91.90	0.109	**0.171**	69.53	0.074	0.145	98.45	0.068	0.220
PrivateEyes **(Ours w/o PEN)**	**68.11**	0.085	0.292	43.93	0.037	0.311	105.78	0.075	0.321
PrivateEyes **(Ours w/ PEN)**	69.25	**0.083**	0.272	**42.98**	**0.034**	0.314	103.31	0.076	0.323

Table 2. **Image Quality.** Evaluation of proposed method against existing eye anonymization baselines on the OpenEDS [11], EV-Eye [36], and LPW [29] datasets. For fair comparison, all metrics are computed on cropped iris regions only. Best numbers are highighted in **bold** while second-best are underlined.

gaze estimation and is on par on OpenEDS and LPW, while for iris recognition it outperforms on EV-Eye and LPW and is comparable on OpenEDS. Our method overall provides a good balance between gaze accuracy and iris recognizability. Table 2 further reports image quality metrics (FID, KID, and LPIPS), confirming that our method produces highly realistic and perceptually faithful anonymized eye images that preserve gaze cues while concealing biometric details. Although high-fidelity realism is not the primary objective, our method consistently generates visually coherent images suitable for downstream tasks outperforming all baseline methods. It is worth noting that baselines methods operate solely on the iris region, keeping the surrounding areas unchanged, whereas our model synthesizes the entire eye image. The results in Table 1 clearly indicate that our method surpasses all baseline approaches in effectively obfuscating user identity features, extending beyond iris regions.

Qualitative Results. Existing methods struggle to balance privacy with visual realism and gaze fidelity. Figure 4 illustrates that baseline methods introduce noticeable artifacts, particularly around the pupil–iris boundary, leading to visually unrealistic results. Rubber Sheet Model (RSM) [6] replaces iris textures through geometric warping, but often introduces severe boundary artifacts and gaze distortion, limiting its photorealism. Iris Style Transfer (IST) [30] anonymizes identity using neural style transfer, yet produces illumination inconsistencies, texture bleeding, and shape deformation around the pupil-sclera boundary, leading to poor structural coherence. In contrast, our method produces smooth boundaries and coherent structures that better align with geometric eye models. Unlike existing methods, our approach also anonymizes all surrounding ocular structures, including eyelashes, eyebrows, and periocular skin yielding fully privacy-preserving eye imagery without compromising visual realism.

5.2. Ablation Study

To further analyze the impact of different anonymization techniques, we conduct an ablation study focusing on gaze estimation and segmentation accuracy.

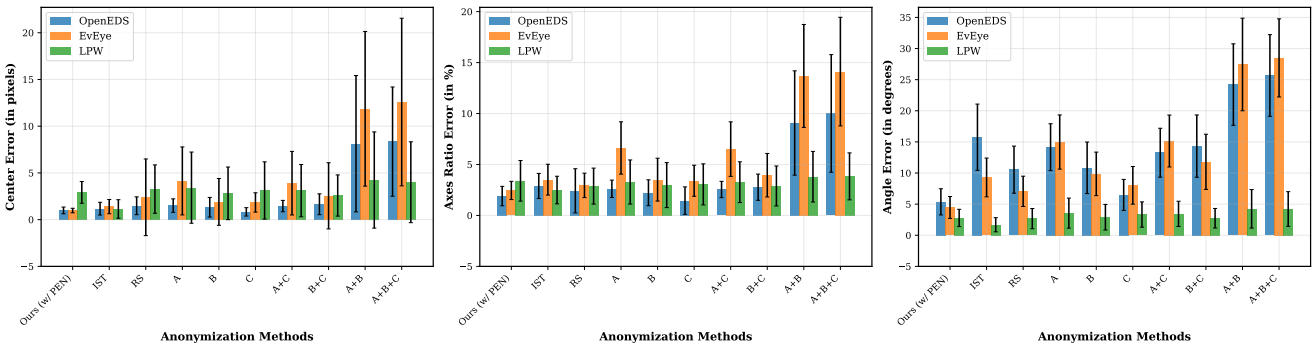

Figure 5. **Geometric Parameter Errors.** Comparison of evaluation metrics across baseline methods and datasets. Methods include: **A** (Downsample), **B** (Gaussian Noise), **C** (Gaussian Blur), RSM [6], IST [30]. Error bars represent standard deviation across multiple images.

5.2.1. Effect of Eye Pose Embedding in Gaze Synthesis

We evaluate the effect of incorporating eye pose embedding in gaze synthesis by comparing the gaze accuracy with and without pose embedding, as shown in Table 1. The results demonstrate that integrating pose embeddings into the ControlEye module enhances model guidance, leading to more precise gaze synthesis while preserving iris anonymization.

5.2.2. Segmentation Accuracy for Downstream Tasks

Figure 6 shows segmentation accuracy across various methods for anonymization. The proposed approach achieves segmentation performance comparable to conventional techniques while preserving privacy.

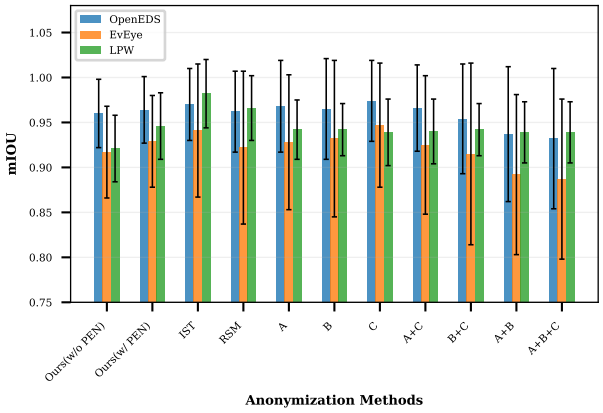

Figure 6. **Segmentation Accuracy:** Comparison of segmentation accuracy across baseline methods on 3 datasets: OpenEDS [11], EV-Eye [36], and LPW [29]. Methods include: **A** (Downsample), **B** (Gaussian Noise), **C** (Gaussian Blur), RSM [6], IST [30]. Error bars represent standard deviation across multiple images.

5.2.3. Geometric Parameter Errors

Figure 5 summarizes the quantitative errors in geometric eye parameter estimation across three datasets. Our method, PrivateEyes, achieves the lowest error rates in nearly all metrics, outperforming both traditional augmentations and prior anonymization approaches. Baseline methods such as gaussian blur, noise, and downsampling exhibit substantial degradation in pupil localization and shape estimation, with errors exceeding 8-12 px and angular deviations above 25° when combined. Iris Style Transfer [30] and Rubber Sheet Model [6] methods also introduce geometric distortions, yielding angle errors up to 15.7° and 10.5°, respectively. In contrast, our approach attains significantly lower center errors (< 1 px) and minimal angular deviation ($\approx 5^\circ$), maintaining strong geometric consistency across all datasets. These results demonstrate that PrivateEyes effectively anonymizes ocular identity while preserving the fine-grained geometry critical for gaze estimation.

6. Conclusion and Future Works

In this work, we introduce a gaze-preserving anonymization framework that enables privacy-compliant eye image data sharing without compromising user identity. Our method anonymizes personally identifiable features across the eye region while retaining critical gaze dynamics necessary for accurate gaze tracking. By leveraging anatomically grounded 3D eye modeling and diffusion-based image synthesis guided by ControlNet, we generate realistic anonymized eye images that effectively decouple biometric identity from gaze. Comprehensive experiments show that our approach outperforms baseline anonymization methods, striking a strong balance between privacy preservation and usability in downstream vision tasks. Additionally, the proposed method facilitates secure data sharing in gaze-driven systems, enabling ethical and privacy-aware dataset release for applications in mixed reality, healthcare, and beyond.

Limitations. The quality of the generated eye images is influenced by the control input, where extreme illumination conditions and gaze angles can adversely affect both the segmentation guidance and the final output. Moreover, the framework processes each frame independently without enforcing temporal coherence across video sequences, highlighting key directions for future work.

References

[1] Geoffrey K Aguirre. A model of the entrance pupil of the human eye. *Scientific reports*, 9(1):9360, 2019. 3, 4

[2] Mauro Barni, Ruggero Donida Labati, Angelo Genovese, Vincenzo Piuri, and Fabio Scotti. Iris deidentification with high visual realism for privacy protection on websites and social networks. *IEEE Access*, 9:131995–132010, 2021. 2

[3] Behnam Bastani, Eric Turner, Carlin Vieri, Haomiao Jiang, Brian Funt, and Nikhil Balram. Foveated pipeline for AR/VR head-mounted displays. *Information Display*, 33(6): 14–35, 2017. 1

[4] Efe Bozkir, Ali Burak Ünal, Mete Akgün, Enkelejda Kasneci, and Nico Pfeifer. Privacy preserving gaze estimation using synthetic images via a randomized encoding based framework. In *ETRA '20*, 2020. 2

[5] Efe Bozkir, Süleyman Özdel, Mengdi Wang, Brendan David-John, Hong Gao, Kevin Butler, Eakta Jain, and Enkelejda Kasneci. Eye-tracked virtual reality: A comprehensive survey on methods and privacy challenges, 2023. 1, 2

[6] 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, 2, 5, 6, 7, 8

[7] Brendan David-John, Kevin Butler, and Eakta Jain. Privacy-preserving datasets of eye-tracking samples with applications in xr. *IEEE Transactions on Visualization and Computer Graphics*, 29(5):2774–2784, 2023. 1

[8] Kai Dierkes, Moritz Kassner, and Andreas Bulling. A fast approach to refraction-aware eye-model fitting and gaze prediction. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, pages 1–9, 2019. 6

[9] Jinhao Duan, Fei Kong, Shiqi Wang, Xiaoshuang Shi, and Kaidi Xu. Are diffusion models vulnerable to membership inference attacks? In *International Conference on Machine Learning*, pages 8717–8730. PMLR, 2023. 2

[10] Anton Mølbjerg Eskildsen and Dan Witzner Hansen. Analysis of iris obfuscation: Generalising eye information processes for privacy studies in eye tracking. In *ETRA '21*, pages 2:1–2:10, 2021. 2

[11] 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. 5, 7, 8

[12] Xiangming Gu, Chao Du, Tianyu Pang, Chongxuan Li, Min Lin, and Ye Wang. On memorization in diffusion models. *arXiv preprint arXiv:2310.02664*, 2023. 2

[13] 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. 3, 6

[14] Yeon Gyu Han, Jung Im Na, Seong Hwan Kim, Sunje Kim, Seung-Han Song, and Dongheon Lee. Anonymize anyone: Toward race fairness in text-to-face synthesis using simple preference optimization in diffusion model. *Available at SSRN 5181873*. 1

[15] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems (NIPS)*, 2017. 6

[16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2, 5

[17] Lianghua Huang, Di Chen, Yu Liu, Shen Yujun, Deli Zhao, and Zhou Jingren. Composer: Creative and controllable image synthesis with composable conditions. 2023. 3

[18] Fei Kong, Jinhao Duan, RuiPeng Ma, Hengtao Shen, Xiaofeng Zhu, Xiaoshuang Shi, and Kaidi Xu. An efficient membership inference attack for the diffusion model by proximal initialization. *arXiv preprint arXiv:2305.18355*, 2023. 2

[19] Ao Liu, Lirong Xia, Andrew Duchowski, Reynold Bailey, Kenneth Holmqvist, and Eakta Jain. Differential privacy for eye-tracking data. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, pages 1–10, 2019. 2

[20] A MacQuarrie and A Steed. Perception of volumetric characters' eye-gaze direction in head-mounted displays. In *Proceedings of 2019 IEEE Virtual Reality (VR)*. IEEE, 2019. 1

[21] Xiaoxu Meng, Ruofei Du, and Amitabh Varshney. Eye-dominance-guided foveated rendering. *IEEE Transactions on Visualization and Computer Graphics*, 26(5):1972–1980, 2020. 1

[22] 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. 3

[23] Clark Phillips and Oleg V Komogortsev. Impact of resolution and blur on iris identification. Technical report, Technical Report. Technical Report, 2011. 5, 6, 7

[24] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1 (2):3, 2022. 3

[25] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3

[26] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2

[27] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 2

[28] Julian Steil, Inken Hagestedt, Michael Xuelin Huang, and Andreas Bulling. Privacy-aware eye tracking using differential privacy. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, pages 1–9, 2019. 2

[29] 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. 5, 7, 8

[30] 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. 2, 5, 6, 7, 8

[31] Mengdi Wang, Efe Bozkir, and Enkelejda Kasneci. Trade-offs in privacy-preserving eye tracking through iris obfuscation: A benchmarking study. *arXiv preprint arXiv:2504.10267*, 2025. 2

[32] Zixiao Wang, Farzan Farnia, Zhenghao Lin, Yunheng Shen, and Bei Yu. On the distributed evaluation of generative models, 2024. 6

[33] An Yan, Jiacheng Li, Wanrong Zhu, Yujie Lu, William Yang Wang, and Julian McAuley. Clip also understands text: Prompting clip for phrase understanding. 2022. 4

[34] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3836–3847, 2023. 2, 3, 4, 5

[35] Richard Zhang, Phillip Isola, Alexei A. Efros, endeavour Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018. 6

[36] 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. 5, 7, 8