

Photo-Realistic Monocular Gaze Redirection Using Generative Adversarial Networks

Zhe He^{1,2}, Adrian Spurr¹, Xucong Zhang¹, Otmar Hilliges¹

¹AIT Lab, ETH Zürich

²Institute of Neuroinformatics, ETH Zürich & University of Zürich

zhehe@student.ethz.ch, {adrian.spurr, xucong.zhang, otmar.hilliges}@inf.ethz.ch

Abstract

Gaze redirection is the task of changing the gaze to a desired direction for a given monocular eye patch image. Many applications such as videoconferencing, films, games, and generation of training data for gaze estimation require redirecting the gaze, without distorting the appearance of the area surrounding the eye and while producing photo-realistic images. Existing methods lack the ability to generate perceptually plausible images. In this work, we present a novel method to alleviate this problem by leveraging generative adversarial training to synthesize an eye image conditioned on a target gaze direction. Our method ensures perceptual similarity and consistency of synthesized images to the real images. Furthermore, a gaze estimation loss is used to control the gaze direction accurately. To attain high-quality images, we incorporate perceptual and cycle consistency losses into our architecture. In extensive evaluations we show that the proposed method outperforms state-of-the-art approaches in terms of both image quality and redirection precision. Finally, we show that generated images can bring significant improvement for the gaze estimation task if used to augment real training data.

1. Introduction

In the cognitive sciences it is well understood that gaze plays a crucial rule in social communication [16], since it conveys important non-verbal cues such as emotion, intention and attention. Hence, many applications such as videoconferencing and movies would benefit from the ability to redirect the gaze in images to establish eye-contact with the viewer. Furthermore, learning-based gaze estimation has recently made significant progress based on in-the-wild datasets [18, 33]. However, such data is difficult to acquire and datasets often only cover a restricted range of gaze angles due to the collection devices. A high-fidelity gaze redirection technique could be leveraged to alleviate this issue




Figure 1: Gaze redirection on Columbia Gaze dataset [26]. (a) Gaze of the subject in the source image is fully centered. (b) Sequence of eye images with different gaze directions synthesized by our method.

by synthesizing novel samples to augment existing datasets.

A reliable and robust gaze redirection approach must be able to (a) redirect the gaze precisely into any given direction, and (b) produce photo-realistic output images which preserve shape and texture details from the input images. Traditional solutions re-render the entire scene by performing 3D transformations, which requires heavy instrumentation to acquire the depth information [20, 31, 34, 5]. Recently, Ganin *et al.* directly rearranged the pixels of the input image to rotate the gaze direction via warping flow generated by a neural network [8]. However, their method fails to generate photo-realistic images for large redirection angles, especially in the presence of large dis-occlusions, such as large parts of the eyeball being covered by the eyelid in the source image. More importantly, such warping methods cannot be perceptually plausible in terms of gaze redirection, since it minimizes pixel-wise differences between the synthesized and ground-truth images without any geometric regularization.

To address the limitations of previous methods, we propose a novel gaze redirection method that builds upon generative adversarial networks (GANs) [9]. To the best of our knowledge, this is the first approach applying GANs to gaze redirection.

As shown in Fig. 1, the proposed method can output

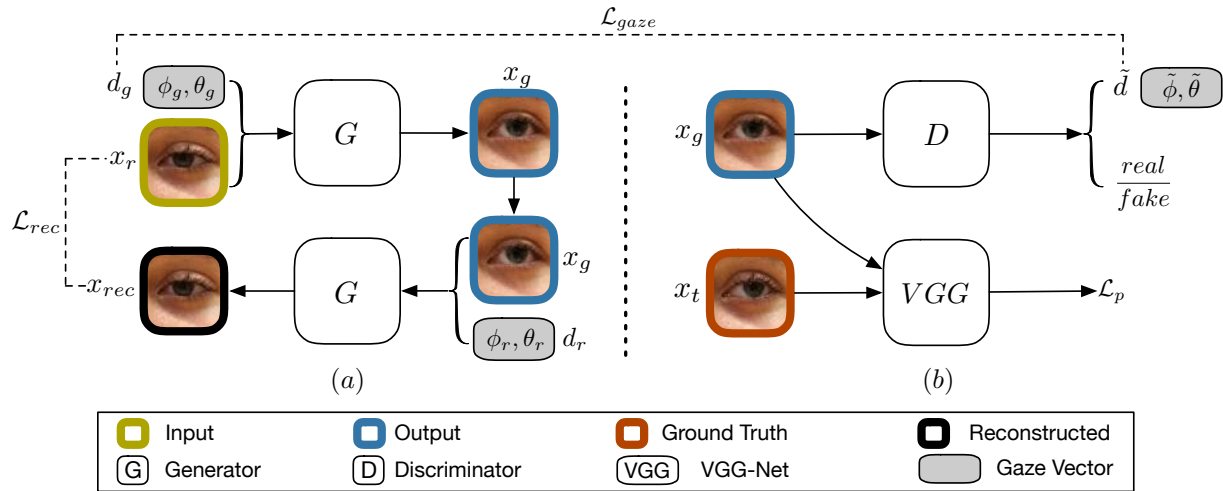


Figure 2: **Overview of our proposed method.** (a) Generator G takes the original eye image x_r and target gaze direction d_g as input, and outputs the synthesized redirected eye image x_g . And then x_g and the source gaze direction d_r are fed into G to reconstruct the x_{rec} . (b) Discriminator D is trained to discriminate real and synthesized images, and it also estimates the gaze direction to calculate the gaze estimation loss. VGG takes synthesized image x_g and ground-truth image x_t to produce perceptual losses for the refinement of generated images. Please refer to Sec. 3.1 for details.

photo-realistic eye images from a single monocular RGB image, while accurately preserving the desired gaze directions. More specifically, we use a conditional GAN [23] as backbone for our architecture shown in Fig. 2. The generator G takes a real eye image as input and generates a new synthetic eye image. Our main contribution is a novel discriminator D that serves the dual purpose of i) ensuring that generated images are realistic, as is common in many GAN formulations, and ii) ensuring that the gaze direction in the output coincides with the input gaze direction which was fed to the generator. This is achieved by incorporating a gaze estimator into the discriminator network. Furthermore, we seek to enhance the perceptual similarity between the generated patch and its ground-truth reference. To this end, we utilize a perceptual loss that penalizes discrepancies between features extracted from the generated images and the ground-truth images by a separate pre-trained neural network. Finally, to ensure that personalized features are not lost in the process of gaze redirection, we use a cycle-consistency loss that enforces consistency between the source image and the generated eye-patch.

We evaluate our method in quantitative experiments and via a qualitative user study. Furthermore, we argue that the pixel-wise difference as a metric of image quality is not suitable for the task of gaze redirection, since it does not correlate with visual perception. To address this, we propose to use LPIPS [32], image blurriness and gaze estimation error as metrics for our quantitative evaluations. Providing further evidence for the high-quality of the generated im-

ages, we show in a controlled experiment that the synthetic samples can be used to augment the training data for a gaze estimation network. Our results show significant improvements in terms of angular gaze error compared to training with real images only. This suggests that our method can be an important tool to further enhance the accuracy attained by deep-learning based gaze estimators.

Our main contributions can be summarized as follows:

- We propose a novel gaze redirection approach in monocular eye images. Technically this is achieved via a feature loss, gaze regularization, and adversarial training. To the best of our knowledge, it is the first GANs-based method for this task.
- We conduct thorough qualitative and quantitative evaluations on the gaze redirection task, showing that our method achieves state-of-the-art performance.
- Finally, we show the potential of leveraging gaze redirection to synthesize training data for the gaze estimation task via training data augmentation.

2. Related Work

Gaze Manipulation Approaches that redirect gaze can be divided into two groups: novel-view synthesis and monocular-gaze synthesis.

Novel-view synthesis methods [20, 31, 34, 5] render a scene containing the face of a subject from a given viewpoint to mimic gazing at the camera. These methods require a depth map of the face, and then synthesize a new

image of the subject with redirected gaze by performing 3D transformations. These approaches mainly serve the purpose of correcting gaze in video conferencing, where the camera is placed at a fixed distance from the screen. However, these methods require dedicated hardware to acquire depth. Furthermore, they alter the entire scene, which limits their applicability.

Monocular-gaze synthesis also aim to change the gaze within the eye region. Wolf *et al.* [28] proposes to replace the eyes in the image with eyes from the same person while looking into a different direction. Although this method retains the realism of eyes after editing, it requires collecting abundant eye images in advance. Furthermore, the movements of the eyelid are ignored in this approach. Recently, a number of warping-based methods have been proposed [8, 17]. These methods use random forests or deep neural networks to learn a flow field to move pixels from the input image to the output image with the desired gaze direction. However, such methods can not handle situations where part of the eye is occluded, since they only replace pixels with existing pixels from the original image without generating any new pixels. Euclidean distance is commonly used as error metric in warping-based methods [8, 17]. However, this does not accurately reflect the perceptual difference between images. A number of approaches based on 3D modeling have been proposed [3, 29]. A 3D model is used to fit both texture and shape of the source eye patch, and then the synthesized eyeballs are superimposed on the source image. However, modeling methods make strong assumptions that do not hold in practice. Therefore, they can not handle images with eyeglasses and other high-variability inter-personal differences.

Generative Adversarial Networks GANs [9] have successfully been applied to many computer vision tasks, such as image super-resolution [21] and image compression [1], and a myriad of further variants have been proposed in recent years (e.g., [22, 2, 4, 4, 10]). GAN-based approaches have also been proposed for the task of image-to-image translation, resulting in impressive results [23, 12]. However, these methods typically require paired data to train. Zhu *et al.* proposed CycleGAN which functions without such requirement [35]. Several derivatives of CycleGAN exist for various tasks [11, 30]. Our method is based on the GAN model while differing from these works in two aspects. First, we focus on a different task, namely that of gaze redirection. Second, we use a number of special purpose losses, including a perceptual loss between ground-truth and synthesized images and a gaze direction preservation loss for training, which we show experimentally to significantly impact the models performance.

3. Approach

3.1. Overview

Our goal is to learn a generator G which can redirect the eye gaze contained in an image into any direction. Given an RGB image of an eye patch $\mathbf{x}_r \in \mathbb{R}^{H \times W \times 3}$ and a target gaze direction vector $\mathbf{d}_g = [\phi_g, \theta_g]$, where $\phi_g \in \mathbb{R}$ and $\theta_g \in \mathbb{R}$ denote the target yaw and pitch angles respectively, the task is to redirect the gaze depicted in \mathbf{x}_r to correspond to the angles of the target vector \mathbf{d}_g , resulting in an output image \mathbf{x}_g . This output needs to satisfy two requirements. First, it needs to look real and consistent. This requires that both shape and texture of \mathbf{x}_g are indistinguishable from those of real data. To this end, we employ a discriminator D that discriminates between generated and real eye images. In order to refine the generated image more, we introduce a feature-based loss that penalizes discrepancies between generated images and ground-truth images. Second, the eye gaze direction in \mathbf{x}_g should look in the direction that the target gaze \mathbf{d}_g indicates. This is accomplished via an auxiliary eye gaze estimator D_{gaze} that enforces the gaze direction. Fig. 2 provides the full overview of the method. We discuss the components in more detail below.

3.2. Objectives

Our method extends the GAN framework via integration of novel loss terms discussed below. The backbone is formed by an existing conditional GAN framework.

Adversarial Loss We build upon WGAN-GP [10] due to its stable performance and adopt its adversarial loss to train the discriminator D and generator G , extending G to take conditional input:

$$\mathcal{L}_{adv} = \mathbb{E}_{\mathbf{x}_r \sim p_{\mathbf{x}_r}(\mathbf{x})} [D_{adv}(\mathbf{x}_r) - D_{adv}(G(\mathbf{x}_r, \mathbf{d}_g))] + \lambda_{gp} \mathbb{E}_{\hat{\mathbf{x}} \sim p_{\hat{\mathbf{x}}}(\hat{\mathbf{x}})} [(\|\nabla_{\hat{\mathbf{x}}} D_{adv}(\hat{\mathbf{x}})\|_2 - 1)^2] \quad (1)$$

In Eq. 1, $p_{\mathbf{x}_r}(\mathbf{x})$ denotes the probability distribution of real images. $D_{adv}(\mathbf{x})$ is the output of the discriminator. The last term is the gradient penalty, which is used to maintain the 1-Lipschitz continuity of D_{adv} . The hyperparameter λ_{gp} controls the strength of gradient penalty, and we use $\lambda_{gp} = 10$ in all experiments.

Gaze Estimation Loss One of our core contributions is the incorporation of an auxiliary gaze estimator D_{gaze} into the GAN framework. D_{gaze} is trained on real images and gaze direction pairs $(\mathbf{x}_r, \mathbf{d}_r)$ using MSE loss:

$$\mathcal{L}_{gaze}^D = \mathbb{E}_{\mathbf{x}_r \sim p_{\mathbf{x}_r}(\mathbf{x})} \|\mathbf{d}_r - D_{gaze}(\mathbf{x}_r)\|_2^2, \quad (2)$$

where in practice, D_{gaze} shares some layers with D_{adv} .

For training G , the generated image $\mathbf{x}_g = G(\mathbf{x}_r, \mathbf{d}_g)$ is fed into the gaze estimator D_{gaze} . Discrepancies between the estimated gaze $D_{gaze}(\mathbf{x}_g)$ and the target gaze \mathbf{d}_g are used as a loss to penalize G . More specifically, we add

the following loss function to the training objective of G , keeping the weights of D_{gaze} fixed:

$$\mathcal{L}_{gaze}^G = \mathbb{E}_{\mathbf{x}_r \sim p_{\mathbf{x}_r}(\mathbf{x})} \|\mathbf{d}_g - D_{gaze}(G(\mathbf{x}_r, \mathbf{d}_g))\|_2^2 \quad (3)$$

Reconstruction Loss The above two loss terms can force the generated eye patch images to be photo-realistic, and ensure redirection of the gaze directions simultaneously. However, none of these losses ensure that personalized features, such as eyeglasses, skin tone or eyebrow are maintained during the redirection process. This is an important feature in many of the envisioned application scenarios such as video-conferencing or interactive videos. Following [35] we enforce cycle consistency, penalizing bad reconstruction as follows:

$$\mathbf{x}_{rec} = G(G(\mathbf{x}_r, \mathbf{d}_g), \mathbf{d}_r) \quad (4)$$

$$\mathcal{L}_{rec} = \mathbb{E}_{\mathbf{x}_r \sim p_{\mathbf{x}_r}(\mathbf{x})} \|\mathbf{x}_r - \mathbf{x}_{rec}\|_1 \quad (5)$$

Here we ask the network to first redirect the gaze to a desired direction and consecutively we generate a third image with the original gaze as target. Above loss ensures that the input and twice-encoded image are as similar as possible.

By penalizing the reconstruction discrepancies, we force the generator to maintain personalized features of the eye, which otherwise would be lost. We use the $L1$ loss, since it empirically performed better in comparison to the $L2$ loss.

Perceptual Loss In our task, human gaze only depends on pitch and yaw angles, which makes it easy to attain a ground-truth gaze images by simply asking the subject to look at the target direction. These ground-truth images can also be incorporated into the training process. One possible approach is to use Mean Squared Error (MSE) between the ground-truth images and generated images as a penalty term. However, applying a MSE loss on generated images would be too strict, as it penalizes pixel-wise discrepancies in all aspects, where minor misalignment could lead to a large MSE while humans would hardly be able to tell the differences (see Table. 1). Alternatively, we adopt the perceptual losses proposed in [13] to penalize the generator G for generating images which do not match ground-truth images perceptually. For this purpose, we use a VGG-16 net [25] pre-trained on ImageNet [19].

Let ψ denote the pre-trained VGG-16 network, $\psi_j(\mathbf{x}) \in \mathbb{R}^{H_j \times W_j \times C_j}$ is the activation of j -th layer of ψ . Two perceptual losses, the content loss \mathcal{L}_c and style loss \mathcal{L}_s , are defined as follows,

$$\mathcal{L}_c = \mathbb{E}_{\mathbf{x}_r \sim p_{\mathbf{x}_r}(\mathbf{x})} \left[\frac{1}{H_j W_j C_j} \|\psi_j(G(\mathbf{x}_r, \mathbf{d}_g)) - \psi_j(\mathbf{x}_t)\|^2 \right] \quad (6)$$

$$\mathcal{L}_s = \mathbb{E}_{\mathbf{x}_r \sim p_{\mathbf{x}_r}(\mathbf{x})} \left[\sum_{j=1}^J \|f_j(G(\mathbf{x}_r, \mathbf{d}_g)) - f_j(\mathbf{x}_t)\|^2 \right] \quad (7)$$

In Equation 7, \mathcal{L}_s is the sum of all style losses from the 1-st layer to the J -th layer of the VGG net. $f_j(\mathbf{x})$ denotes the Gram matrix, which is defined as:

$$f_j(\mathbf{x})_{c,c'} = \frac{1}{N_j} \sum_h \sum_w \psi_j(\mathbf{x})_{h,w,c} \psi_j(\mathbf{x})_{h,w,c'} \quad (8)$$

$$N_j = H_j W_j C_j \quad (9)$$

Optimizing the content loss encourages \mathbf{x}_g to perceptually resemble \mathbf{x}_t in terms of overall structure and spatial relation. Meanwhile, by minimizing the style loss, the generator tries to refine the details of \mathbf{x}_g , such as color and texture, to increase the similarity to \mathbf{x}_t . The perceptual loss is the sum of content loss and style loss:

$$\mathcal{L}_p = \mathcal{L}_c + \mathcal{L}_s \quad (10)$$

Overall Objectives The final training objectives consists of two parts, one for G and D respectively:

$$\mathcal{L}_G = -\mathcal{L}_{adv} + \lambda_p \mathcal{L}_p + \lambda_{gaze} \mathcal{L}_{gaze}^G + \lambda_{rec} \mathcal{L}_{rec} \quad (11)$$

$$\mathcal{L}_D = \mathcal{L}_{adv} + \lambda_{gaze} \mathcal{L}_{gaze}^D \quad (12)$$

Where λ_p , λ_{gaze} and λ_{rec} are the hyperparameters that control the contribution of each loss term. In all experiments, we set them to $\lambda_p = 100$, $\lambda_{gaze} = 5$, $\lambda_{rec} = 50$.

4. Implementation

4.1. Network Architecture

Generator The generator takes an RGB eye patch image $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$ and a gaze direction vector $\mathbf{d} \in \mathbb{R}^2$ as input. \mathbf{d} is expanded into $\mathbb{R}^{H \times W \times 2}$ by channel-wise duplication, such that \mathbf{x} and \mathbf{d} can be concatenated depth-wise. We use a modified variant of the generator architecture introduced in [13], the details of which can be found in the supplementary.

Discriminator We modified the last layer of the discriminator architecture of WGAN-GP [10] to have two output branches: one performs real/fake discrimination and another one outputs gaze estimates respectively.

VGG-16 We use the standard architecture of VGG-16 introduced in [25]. We use the activation of the 5th layer to produce the content loss, and the 1st to 4th layers to produce the style loss.

4.2. Training Details

For all the following experiments, we use Adam [15] optimizer with $\beta_1 = 0.5$, $\beta_2 = 0.999$. Our model is trained for 300 epochs with batch size 32. The learning rate is set to 0.0002 for the first 150 epochs, and linearly decays to 0 during the next 150 epochs. For every update of the generator, we update the discriminator five times. The training process takes about 16 hours on a single NVIDIA[®] 1080Ti GPU.


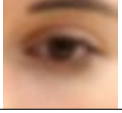

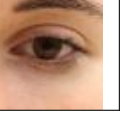
(a) Eye patch	(b) Blurred	(c) Noisy	(d) Shifted
			
MSE	69.57	155.36	176.06
LPIPS	0.122	0.106	0.016

Table 1: Examples of image degradations. (a) Eye patch from training set. (b) Blurred with Gaussian filter. (c) With random Gaussian noise. (d) Shifted up by one pixel.

5. Experiments

In this section, we detail the quantitative and qualitative experiments that were conducted to evaluate our approach.

5.1. Metrics

As mentioned before (see Sec. 1), gaze redirection models are required to be precise in redirecting and to produce photo-realistic and consistent images. Correspondingly, the evaluation metrics need to be able to assess these aspects. In previous work of monocular gaze manipulation [8, 29], the mean squared error (MSE) was used as the metric to measure the similarity between the generated eye images and ground-truth eye images. This was used as a quantitative measure of performance. However, we argue that MSE is not the ideal metric for this task, as has been observed previously in related work [27]. To illustrate the issue, we created three types of image degradations compared to the ground-truth as shown in Table. 1. Qualitatively, Table. 1 d) is the most similar to the ground-truth Table. 1 a). However, when calculating the MSE, we see that this does not correlate well with one’s qualitative judgment.

Instead, we propose to use the following three error metrics: LPIPS score, image blurriness and gaze estimation error.

LPIPS Score. We use the Learned Perceptual Image Patch Similarity (LPIPS) [32] metric to evaluate the visual quality of the generated gaze images. Different from traditional metrics, LPIPS is based on deep networks and aims to resemble human perception in image evaluation tasks. The LPIPS score is given as follows:

$$d(\mathbf{x}, \mathbf{x}_0) = \sum_l \frac{1}{H_l W_l} \sum_{h,w} \|\mathbf{w}_l \odot (\hat{\mathbf{y}}_{hw}^l - \hat{\mathbf{y}}_{0hw}^l)\|_2^2 \quad (13)$$

Where $d(\mathbf{x}, \mathbf{x}_0)$ denotes the LPIPS score between the images $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$ and $\mathbf{x}_0 \in \mathbb{R}^{H \times W \times 3}$. The variables $\hat{\mathbf{y}}^l \in \mathbb{R}^{H_l \times W_l \times C_l}$ and $\hat{\mathbf{y}}_0^l \in \mathbb{R}^{H_l \times W_l \times C_l}$ are the channel-wise unit-normalized activation from the l -th layer of the backbone network and $\mathbf{w}_l \in \mathbb{R}^{C_l}$ are the trainable weights

used for scaling the activations. In our work, we use the pre-trained Alex-Net [19] as a backbone,

When calculating LPIPS on the previous examples in Table. 1, we see that the scores agree more with human evaluation.

Image Blurriness (IB). To measure the blurriness of a generated gaze image, we use a Laplace filter k and perform convolution on the grayscale gaze image \mathbf{x}_{gray} . Image blurriness can be acquired by calculating the reciprocal variance of the filtered image as shown in the following equations:

$$k = \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 1 \end{bmatrix}, \text{IB} = \frac{1}{\text{Var}[k * \mathbf{x}_{gray}]}. \quad (14)$$

Gaze Estimation Error. For the assessment of gaze redirection accuracy, we employ a state-of-the-art gaze estimator proposed by Park *et al.* [24] which was pre-trained on MPIIGaze [33]. The estimator predicts the gaze direction of the generated gaze images. The angular error δ between the target gaze direction \mathbf{d}_g and the predicted gaze direction $\hat{\mathbf{d}}$ is used as the gaze estimation error. To attain δ , the yaw and pitch angles (ϕ, θ) need to be converted into three-dimensional Cartesian coordinates first:

$$\mathbf{v} = T(\mathbf{d}) = [\cos \phi \cos \theta, -\sin \phi, \cos \phi \sin \theta]. \quad (15)$$

where $T(\cdot)$ denotes the mapping between two coordinate systems. Then, δ can be obtained via the following calculations:

$$\mathbf{v}_g = T(\mathbf{d}_g), \hat{\mathbf{v}} = T(\hat{\mathbf{d}}) \quad (16)$$

$$\delta = \arccos \frac{\mathbf{v}_g^\top \cdot \hat{\mathbf{v}}}{\|\mathbf{v}_g\| \cdot \|\hat{\mathbf{v}}\|}. \quad (17)$$

5.2. Dataset

We used the Columbia Gaze dataset [26] for the evaluations, which is a high-resolution, publicly available human gaze dataset collected from 56 subjects. The head poses of each subject are discrete values in the set $[-30^\circ, -15^\circ, 0^\circ, 15^\circ, 30^\circ]$. For each head pose, there are 21 gaze directions, which are the combinations of three pitch angles $[-10^\circ, 0^\circ, 10^\circ]$, and seven yaw angles $[-15^\circ, -10^\circ, -5^\circ, 0^\circ, 5^\circ, 10^\circ, 15^\circ]$. Here, we only used the images with frontal faces, i.e. 0° head pose. Results on non-frontal faces are provided in the supplementary. We split the data into train and test set. The former set includes 50 subjects whereas the latter contains 6 subjects. We first run face alignment with `dlib` [14] by parsing the face with 68 facial landmark points. After that, a minimal enclosed circle with center (x, y) and radius R was extracted from the 6 landmark points of each eye. The cropping region of the eye patch is set as a square box with center (x, y) and side length $3.4R$. We flipped the right eye images horizontally to align with the left eye images. All eye patch images were

resized to 64×64 . Both the pixel values of images and gaze directions were normalized into the range $[-1.0, 1.0]$. Other publicly available gaze datasets, such as MPIIGaze [33] or EYEDIAP [7] only provide low-resolution images and would therefore introduce a bias towards low quality images. Therefore, these datasets were not suitable for our task.

5.3. Evaluation Protocol

We tested each model on the 6 subjects contained in the test set, which includes 252 eye patch images. For each image, we redirected the gaze into 20 gaze directions separately, excluding the gaze direction of the current image. Intuitively, it would be harder for the model to redirect the gaze if the target gaze direction is significantly different from the original gaze direction. Therefore, we defined the correction angle γ to indicate the angular difference between original and target gaze directions. It is calculated as follows:

$$\mathbf{v}_g = T(\mathbf{d}_g), \mathbf{v}_r = T(\mathbf{d}_r) \quad (18)$$

$$\gamma = \arccos \frac{\mathbf{v}_g^T \cdot \mathbf{v}_r}{\|\mathbf{v}_g\| \cdot \|\mathbf{v}_r\|} \quad (19)$$

Where $T(\cdot)$ is the aforementioned mapping in Eq. 15.

5.4. Comparison to State-of-The-Art

Baseline Model We adopt DeepWarp [8] as our baseline model. The original implementation uses 7 eye landmarks as input, including the pupil center. However, detecting the pupil center is very challenging task. Therefore we only used 6 landmarks as the input to DeepWarp. Unfortunately, evaluating the more recent work GazeDirector [29] with the proposed error metric is not possible, since their implementation is not available. Therefore, we did not compare GazeDirector in our paper.

Qualitative Evaluation Fig. 3 and Fig. 4 show the generated gaze images examples. Although both methods are capable of redirecting the gaze, we observe that the generated images of DeepWarp have several obvious defects. First, textures such as skin and eyebrows are more blurry. Second, the shapes of certain parts, such as the edges of eyelid (see Fig. 4), iris and eyeglasses (see Fig. 3), are distorted. In contrast, the generated gaze images of our proposed method are more faithful to the input images.

Quantitative Evaluation Fig. 5a plots the LPIPS scores of DeepWarp and our method. The range of correction angle is $[4.9^\circ, 35.9^\circ]$. From the figure we can see that our method achieves the lower LPIPS score than DeepWarp at every correction angle, which indicates that our method can generate gaze images that are perceptually more similar to the ground-truth images. This observation is consistent with the qualitative evaluation (Fig. 3 and Fig. 4).

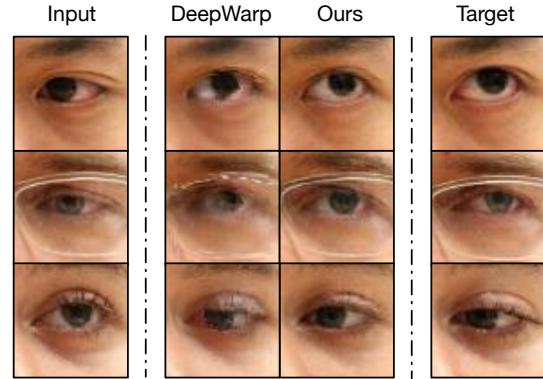


Figure 3: Gaze redirection comparison.




Figure 4: Zoom-in detail comparison.

Fig. 5b plots the blurriness of the produced images. Our method outperforms the related work by a large margin, being closer to the blurriness observed in real images.

Fig. 5c presents the results of gaze estimation error. The error of our method is much lower than DeepWarp, which indicates that our method can redirect the gaze with a higher precision.

User Study In addition, we conducted a user study to compare the performance of DeepWarp [8] and our method. As the overall range of correction angle is $[4.9^\circ, 35.9^\circ]$, we split the generated gaze images into three groups: $[4.9^\circ, 15.0^\circ]$, $(15.0^\circ, 25.0^\circ]$, $(25.0^\circ, 35.9^\circ]$, which represent the difficulty of gaze redirection from easy to hard. In each group, we randomly choose 19 pairs of images generated by both methods with the same input image and gaze direction. Two images in a pair were shown side by side to the user without any further information. The task for the users is to pick the gaze image that looks more realistic than the other.

In total, we have 16 users participated in our study. Table 2 shows the results of the user study. We can see that our method outperforms DeepWarp with a significant margin. The results of quantitative evaluations shown in Fig. 5 are consistent with the user assessment, which demonstrates the the metrics we used are effective in the evaluation of the gaze redirection task.

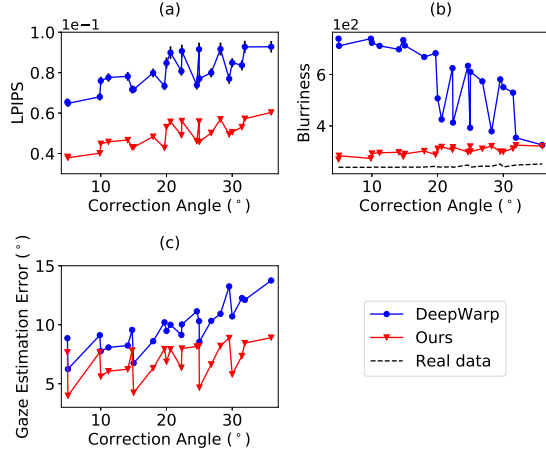


Figure 5: Quantitative evaluation results of DeepWarp and our method. (a) LPIPS score (lower is better). (b) Image blurriness (lower is better). (c) Gaze estimation error (lower is better)

Group	DeepWarp [8]	Ours
$[4.9^\circ, 15.0^\circ]$	21.9%	78.1%
$(15.0^\circ, 25.0^\circ]$	9.0%	91.0%
$(25.0^\circ, 35.9^\circ]$	13.4%	86.6%

Table 2: Voting results of user study, comparing DeepWarp with our method. Each row sums up to 100 %.

5.5. Ablation Study

To understand the effect of each component of our proposed model, we performed an ablation study. As mentioned in Sec. 3.1, besides the adversarial loss, there are three other loss terms: gaze estimation loss, reconstruction loss and perceptual loss. We trained a model for each one of these additional loss terms, where one of the terms was removed from the total loss.

Qualitative Results We show the results in Fig. 6. As can be seen from the second column of Fig. 6, the model is not able to maintain features from input images without \mathcal{L}_{rec} . The most significant example is in the first row, where the model without \mathcal{L}_{rec} does not preserve the rim of the eyeglasses.

When discarding \mathcal{L}_{gaze} , it can be observed that the model fails to redirect the gaze entirely. Therefore, we will not further consider models not using \mathcal{L}_{gaze} in the following quantitative evaluations.

Removal of \mathcal{L}_p causes reduced image quality as can be visually verified in the generated images. These show artifacts such as distortion of eyelid shape, iris and eyeglasses (Fig. 6).

Quantitative Results Fig 7a shows the LPIPS scores of the full model, a model without \mathcal{L}_{rec} and a model without \mathcal{L}_p . It is clear that the LPIPS score increases without either of \mathcal{L}_{rec} or \mathcal{L}_p , which indicates that both terms are essential for improving the visual quality of redirected gaze images.

The blurriness scores shown in Fig. 7b are also consistent with what has been observed in qualitative results, where the full model produces the sharpest images.

Fig. 7c presents the gaze estimation error. Removal of either \mathcal{L}_{rec} or \mathcal{L}_p does not significantly worsen the gaze estimation error, since the precision of redirection is mainly controlled by \mathcal{L}_{gaze} .

5.6. Augmenting Gaze Data

Lastly, we investigated the feasibility of leveraging our method for the purpose of data augmentation for eye gaze estimation tasks. This is motivated by the rapid progress in deep-learning based gaze estimation (e.g., [33, 24]). While appearance-based gaze estimation techniques that use Convolutional Neural Networks (CNN) have significantly surpassed classical ones [33] in in-the-wild settings, there still remains a significant gap towards applicability in high-accuracy domains. The currently lowest reported person-independent error of 4.3° [6] is roughly equivalent to 4.7cm at a distance of 60cm. One reason for this relatively high error is the lack of sufficient training data. In particular, it is known that many datasets only cover a relatively small range of gaze angles due to hardware limitations. Therefore we propose to leverage our model for augmenting existing datasets, in order to expand the range of gaze directions and leading to better gaze estimation performance. To the best of our knowledge, this is the first time that potential of gaze redirecting models to improve gaze estimation models have been explored.

To assess the applicability of our method in this setting, we performed a proof-of-concept experiment indicating that our technique can fill in unseen gaze angles. First, we constructed two datasets.

The **raw dataset** contains all the eye images with 10° pitch angles from the Columbia Gaze Dataset [26].

The **augmented dataset** contains the images from the raw dataset. Furthermore, we took the images of the 6 testing subjects (see Sec. 5.2), and used them to synthesize new gaze images with pitch angles -10° and 0° respectively.

We trained two gaze estimators on the raw and augmented datasets respectively. Both estimators were constructed by the same VGG-16 [25] architecture. Since augmented dataset contains more images, we trained the corresponding estimator for less epochs. Implementation details can be found in the supplementary.

To test the estimators, we used two test sets. (1) **Columbia Gaze**. Since the eye images in Columbia Gaze dataset with pitch angles -10° and 0° of the 50 training

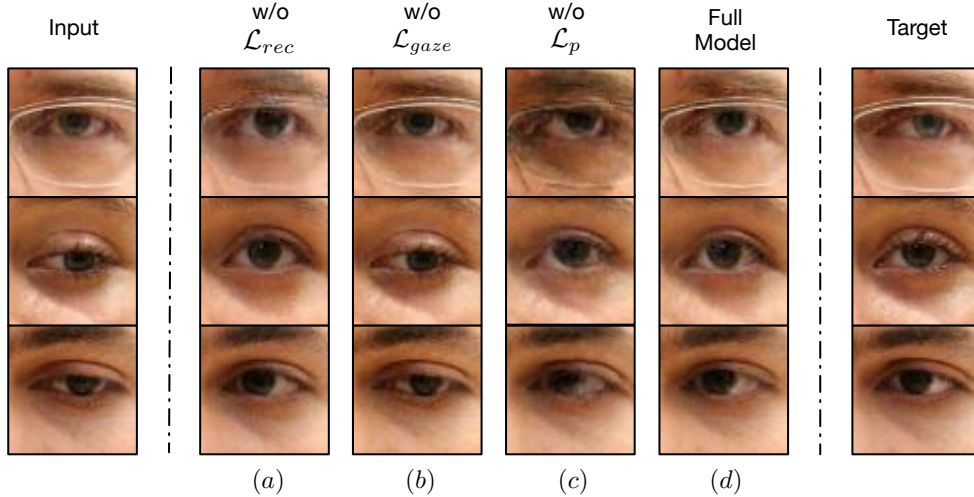


Figure 6: Gaze redirection results of (a) model without \mathcal{L}_{rec} and (b) model without \mathcal{L}_{gaze} , (c) model without \mathcal{L}_p and (d) full model.




Figure 7: Quantitative evaluation results of our full model, model without \mathcal{L}_{rec} and model without \mathcal{L}_p . (a) LPIPS score (lower is better). (b) Image blurriness (lower is better). (c) Gaze estimation error (lower is better)

subjects (see Sec. 5.2) have not been seen by the gaze estimators, we use these images as our test set without leaking information. (2) **MPIIGaze**. For cross-dataset evaluation, we take the test set of MPIIGaze [33], where the pitch angles are in the range $[-20^\circ, 1.5^\circ]$.

Results As shown in Table 3, the gaze estimator trained on the augmented dataset always performs better than the gaze estimator trained on the raw dataset. Intuitively, since the raw dataset only contains images with positive pitch an-

Dataset	Raw	Augmented
Columbia	14.3°	6.9°
MPIIGaze	20.2°	14.0°

Table 3: Gaze estimation errors. Column name is the training set, while row name is the testing set.

gles, the trained estimator is expected to generalize poorly on the test set, where most samples have different pitch angles. In contrast, the augmented images aid the estimator in generalizing better to unseen angles, improving the test set performance.

6. Conclusion

In this paper, we propose a novel monocular gaze redirection method leveraging generative adversarial networks. The proposed method can generate photo-realistic eye images while preserving the desired gaze direction. In order to further refine the generated images, we incorporate a perceptual loss into the adversarial training and include a cycle-consistent loss to preserve personalized features. Extensive evaluations show that our approach outperforms previous state-of-the-art methods in terms of both image quality and redirection precision. Finally, we show that our gaze redirection method can benefit gaze estimation tasks by generating additional training data with controlled gaze directions.

7. Acknowledgement

We thank the NVIDIA Corporation for the donation of GPUs used in this work.

References

- [1] Eirikur Agustsson, Michael Tschannen, Fabian Mentzer, Radu Timofte, and Luc Van Gool. Generative adversarial networks for extreme learned image compression. *arXiv preprint arXiv:1804.02958*, 2018.
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [3] Michael Banf and Volker Blanz. Example-based rendering of eye movements. In *Computer Graphics Forum*, volume 28, pages 659–666. Wiley Online Library, 2009.
- [4] David Berthelot, Thomas Schumm, and Luke Metz. Began: Boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717*, 2017.
- [5] Antonio Criminisi, Jamie Shotton, Andrew Blake, and Philip HS Torr. Gaze manipulation for one-to-one teleconferencing. In *ICCV*, volume 3, pages 13–16, 2003.
- [6] Tobias Fischer, Hyung Jin Chang, and Yiannis Demiris. RT-GENE: Real-Time Eye Gaze Estimation in Natural Environments. In *ECCV*, September 2018.
- [7] Kenneth Alberto Funes Mora, Florent Monay, and Jean-Marc Odobez. Eyediap: A database for the development and evaluation of gaze estimation algorithms from rgb and rgb-d cameras. In *Proceedings of the ACM Symposium on Eye Tracking Research and Applications*. ACM, Mar. 2014.
- [8] Yaroslav Ganin, Daniil Kononenko, Diana Sungatullina, and Victor Lempitsky. Deepwarp: Photorealistic image resynthesis for gaze manipulation. In *European Conference on Computer Vision*, pages 311–326. Springer, 2016.
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [10] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pages 5767–5777, 2017.
- [11] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 172–189, 2018.
- [12] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [13] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [14] Davis E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009.
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] Chris L Kleinke. Gaze and eye contact: a research review. *Psychological bulletin*, 100(1):78, 1986.
- [17] Daniil Kononenko and Victor Lempitsky. Learning to look up: Realtime monocular gaze correction using machine learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4667–4675, 2015.
- [18] Kyle Krafka, Aditya Khosla, Petr Kellnhofer, Harini Kannan, Suchendra Bhandarkar, Wojciech Matusik, and Antonio Torralba. Eye tracking for everyone. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2176–2184, 2016.
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [20] Claudia Kuster, Tiberiu Popa, Jean-Charles Bazin, Craig Gotsman, and Markus Gross. Gaze correction for home video conferencing. *ACM Transactions on Graphics (TOG)*, 31(6):174, 2012.
- [21] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew P Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photorealistic single image super-resolution using a generative adversarial network. In *CVPR*, volume 2, page 4, 2017.
- [22] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2794–2802, 2017.
- [23] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [24] Seonwook Park, Adrian Spurr, and Otmar Hilliges. Deep pictorial gaze estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 721–738, 2018.
- [25] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [26] Brian A Smith, Qi Yin, Steven K Feiner, and Shree K Nayar. Gaze locking: passive eye contact detection for human-object interaction. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*, pages 271–280. ACM, 2013.
- [27] Zhou Wang and Alan C Bovik. Mean squared error: Love it or leave it? a new look at signal fidelity measures. *IEEE signal processing magazine*, 26(1):98–117, 2009.
- [28] Lior Wolf, Ziv Freund, and Shai Avidan. An eye for an eye: A single camera gaze-replacement method. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 817–824. IEEE, 2010.
- [29] Erroll Wood, Tadas Baltrušaitis, Louis-Philippe Morency, Peter Robinson, and Andreas Bulling. Gazedirector: Fully articulated eye gaze redirection in video. In *Computer Graphics Forum*, volume 37, pages 217–225. Wiley Online Library, 2018.
- [30] Wayne Wu, Yunxuan Zhang, Cheng Li, Chen Qian, and Chen Change Loy. Reenactgan: Learning to reenact faces via boundary transfer. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 603–619, 2018.

- [31] Ruigang Yang and Zhengyou Zhang. Eye gaze correction with stereovision for video-teleconferencing. In *European Conference on Computer Vision*, pages 479–494. Springer, 2002.
- [32] Richard Zhang, Phillip Isola, Alexei A Efros, Eli 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*, pages 586–595, 2018.
- [33] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. Appearance-based gaze estimation in the wild. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4511–4520, June 2015.
- [34] Jiejie Zhu, Ruigang Yang, and Xueqing Xiang. Eye contact in video conference via fusion of time-of-flight depth sensor and stereo. *3D Research*, 2(3):5, 2011.
- [35] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2223–2232, 2017.