

Resist : Reconstruction of irises from templates

Sohaib Ahmad
University of Connecticut
Storrs, CT USA
sohaib.ahmad@uconn.edu

Benjamin Fuller
University of Connecticut
Storrs, CT USA
benjamin.fuller@uconn.edu

Abstract

Iris recognition systems transform an iris image into a feature vector. The seminal pipeline segments an image into iris and non-iris pixels, normalizes this region into a fixed-dimension rectangle, and extracts features which are stored and called a template (Daugman, 2009). This template is stored on a system. A future reading of an iris can be transformed and compared against template vectors to determine or verify the identity of an individual.

As templates are often stored together, they are a valuable target to an attacker. We show how to invert templates across a variety of iris recognition systems. Our inversion is based on a convolutional neural network architecture we call RESIST (REconStructing IriSes from Templates).

We apply RESIST to a traditional Gabor filter pipeline, to a DenseNet (Huang et al., CVPR 2017) feature extractor, and to a DenseNet architecture that works without normalization. Both DenseNet feature extractors are based on the recent ThirdEye recognition system (Ahmad and Fuller, BTAS 2019). When training and testing using the ND-0405 dataset, reconstructed images demonstrate a rank-1 accuracy of 100%, 76%, and 96% respectively for the three pipelines. The core of our approach is similar to an autoencoder. To obtain high accuracy this core is integrated into an adversarial network (Goodfellow et al., NeurIPS, 2014)

1. Introduction

This work explores the vulnerability of storing the output of an iris recognition system. Irises are strong biometrics with a high entropy rate [25] and strong consistency [34, 49]. The seminal processing pipeline due to Daugman proceeds in three stages [9]. First, the iris region is *segmented* from the rest of the image. Second, this region is *normalized* into a rectangular representation. Lastly, a feature extractor, such as a 2-dimensional Gabor filter is applied with the resulting

values converted to binary to form a template. Daugman’s pipeline is known as the iriscode.

Due to the noise experienced when collecting irises in real environments, new pipelines are steadily proposed. In most systems, there are two main stored values. The first is the model or feature extraction mechanism (which may be a traditional set of feature extractors or a deep neural network architecture trained on irises). The second is a set of output feature vectors known as templates that are used to identify individuals. If a reading is collected the identity of this person is calculated (or verified) as the identity of the minimum distance stored iris provided that the distance is less than some predefined threshold.

If an attacker has a target’s iris they may be able to generate spoofed iris images [22, 47, 50] to fool an iris recognition system. There are multiple ways a spoofed iris can be *presented* to the iris recognition system including a printed iris or a textured contact lens; these attacks are known as presentation attacks. Defenses to these attacks [8, 12, 41] are usually trained on certain types of presentation attacks. As an example, Chen et al. [7] use deep convolutional neural networks (CNNs) to detect presentation attacks. Kohli et al. [32] extract Zernike moments and local binary pattern features which are used in a neural network classifier to detect presentation attacks.

A second attack vector inverts a template back into a corresponding an iris image [13]. An attack is successful if the resulting image is classified in the same class as the template by the recognition system. Templates are more accessible to an attacker than actual biometrics since recognition systems store templates to compare against any input to the system. Thus, if one can reconstruct an iris, one can execute a presentation attack in more settings. Galbally et al. [13] showed how to invert a Gabor filter feature extractor.

While templates are usually encrypted at rest, for authentication systems in use, templates will be unencrypted in memory. Template protection and cancellable biometrics protect against this threat [42, 55]. Bloom filters are widely used to protect Gabor filter based feature extractors [6, 15, 48]. Fuzzy extractors [11, 29, 30] have been used to provide secure

iris authentication [5, 18, 45].

Our Contribution The primary contribution of this work is a convolutional neural network architecture called RESIST for REConStructing IriSes from Templates. RESIST effectively reconstructs irises from a stored template. RESIST is a black-box attack which does not utilize the specifics of the feature extraction to train the reconstruction network. However, we do assume knowledge of the length of the feature vector. RESIST is able to accurately reverse a variety of iris template systems. RESIST has a *core* network which is similar to an autoencoder. We first train this *core* and then as the generator of the GAN. We apply RESIST to three iris processing pipelines which are summarized in Figure 1:

RESIST-G A traditional Gabor filter based processing pipeline (with segmentation and normalization before applying Gabor filters).

RESIST-DN A deep neural network feature extractor based on DenseNet [21] that works on normalized images.

RESIST-D A deep neural network feature extractor based on DenseNet that directly works on segmented images.

Both deep neural network approaches build on the recent ThirdEye architecture [2], see Section 2. Our goal is to directly invert the resulting feature vector. For the neural networks this is a vector of real values. It is a binary vector for RESIST-G. In many applications, the real valued vector resulting from a neural network is projected to a binary vector using locality sensitive hashes [26], we leave inverting this system as future work.

Our technical approach is drawn from the area of synthetic iris generation. Recent machine learning techniques such as generative adversarial networks or GANs [16] have made it possible to generate synthetic irises given access to a database. Yadav et al. [53] use RaSGAN (relativistic average standard GAN) [28] to generate synthetic irises for the purpose of studying their effects on presentation attack defence (PAD) algorithms. Irises from the RaSGAN perform well against PAD and follow real iris statistics well. Kohli et al. [33] use the DCGAN architecture to generate synthetic irises. Our reconstructed irises can be used in presentation attacks in a similar fashion as these works.

Our approach can be seen as creating a synthetic iris database with an important change. Instead of generating a database of (synthetic) irises we focus on the individual mapping from a single template to a corresponding iris. Reconstructing irises from corresponding templates is a harder task. Synthetic irises can be viewed as irises that must closely resemble bonafide irises as discussed in [33, 53]. The objective of RESIST is to create synthetic irises with the additional constraint that they match the source template, therefore it is not enough to learn the distribution of iris. RESIST must learn from the individual template.

Figure 2 shows a standard iris image on the left. As mentioned above, our approach has two training stages, first we build a core that is similar to an autoencoder. This core is then placed inside of GAN for a second training stage. The center image is a reconstructed iris using the core, this image has low pixel error (optimized by the core) with an average pixel error of 3%. However, when matched with real iris images, this technique only achieves rank-1 accuracy of 62% and true acceptance rate of only 18.5% at 1% FAR. Rank-1 accuracy is how frequently an iris template of the same class is the closest value in the dataset.

To deal with this inaccuracy, we introduce the second stage of training, the core is moved inside a GAN as the generator. When trained from scratch, this GAN collapsed. To deal with this issue, we use three techniques:

1. Pretraining the core before training in the GAN (pre-training was used in prior iris recognition networks [4]),
2. Adding noise to the real iris images during training, and
3. Using spectral normalization [38].

The right hand image in Figure 2 shows a characteristic image output by RESIST.

We report rank-1 accuracy as well as accuracy metrics from a recent facial reconstruction paper [36]. We report this for both the RESIST templates and the corresponding legitimate templates. The ND-0405 dataset for training and testing. We summarize our findings here with the rank-1 of reconstructed irises compared with rank-1 accuracy for legitimate irises (see Table 3):

- RESIST-D: 96.3% vs. DenseNet: 99.9%,
- RESIST-DN 89.3% vs. DenseNet Normal: 98.7%, and
- RESIST-G 97.1% vs. Gabor filter: 97.9%.

Underlying system accuracy is an upper bound for our accuracy. We explore the accuracy rates further in Section 5.

Related Work Galbally et al. [13] use a genetic algorithm which reconstructs irises from their Gabor filter templates. Their method is stochastic and generates multiple irises from a single template. Venugopalan et al. [50] show how given an Gabor filter template $f(x)$ for an individual x , another individual y can transform their iris image into one with a similar Gabor filter template as x . Mai et al. [36] recently considered reconstruction attacks on the facial biometric. Their attack relies on a large training dataset containing 2 million facial images. To our knowledge such a large iris dataset does not exist. Facial recognition is done using deep neural networks which work on a centered face image to generate a template. Iris as a biometric has multiple ways to generate templates. We primarily explore the different types of iris templates and use a single reconstruction network in contrast to Mai et al. where a GAN is used to augment their existing training dataset and train a reconstruction network

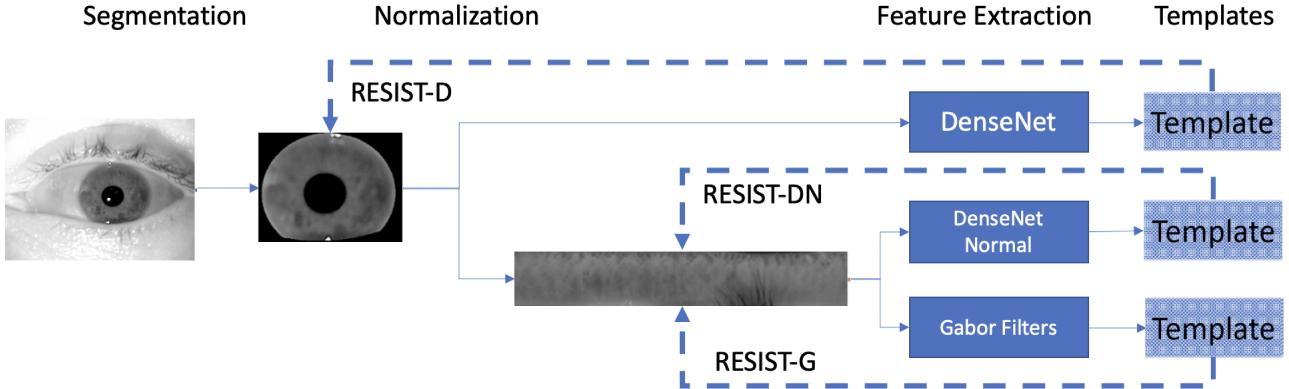


Figure 1: Our iris reconstruction system. Blue lines represent 3 different iris processing pipelines. The traditional Gabor filter pipeline is on the bottom. A DenseNet based pipeline on normalized images in the middle and a DenseNet pipeline on unnormalized images on the top. Both DenseNet pipelines follow the three stages shown in the top pipeline. These are drawn from prior work and discussed in Section 2. For the two normalized pipelines, RESIST reconstructs normalized images while for the pipeline without normalization it reverses to a segmented iris.

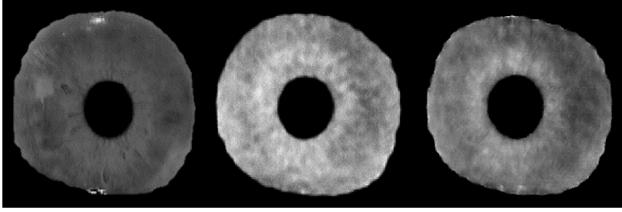


Figure 2: Iris reconstructions from templates. Original, network core and GAN reconstructions respectively.

on the augmented facial images. To the best of our knowledge, RESIST is the first attack that reconstructs irises from deep learning based processing templates.

Organization The rest of this work is organized as follows. In Section 2 we describe the three attacked pipelines, Section 3 describes the design of RESIST, Section 4 provides an overview of the evaluation strategy and the dataset used with results in Section 5. Section 6 concludes.

2. Iris Processing Pipelines

We apply RESIST to three iris processing pipelines. This section briefly introduces the attacked pipelines. We explore both Gabor filter and deep learning based systems.

In traditional iris recognition systems there are three stages 1) segmentation, 2) normalization and 3) feature extraction. The first stage segments the provided image into iris and non-iris pixels. Normalization then converts the iris region into a rectangular representation whose area is fixed. This allows correction for different possible locations of the iris within the image. Finally, features are extracted from this normalized rectangle. In the traditional Gabor filter template pipeline this extraction is two stages 1) convolutions with

Gabor filters and 2) binarization of the resulting complex numbers. Subsequent iris readings by the system will undergo the same procedure with Hamming distance used to match templates. Our goal is to take the binarized template and reconstruct back to the normalized iris representation. We form a traditional iris recognition pipeline using USIT software to do segmentation and normalization [20]. We use transform 5 from Osiris's [39] Gabor filter banks.

We consider two deep-learning based iris pipelines. We use the NotreDame 0405 [3] to train these two models. Both DenseNet networks use the first 25 left iris images of each subject for training. The test comprises first 10 right iris images of each subject. Both DenseNet networks are trained to convergence (98% rank-1 accuracy on testing set). We further report on the tradeoff between true and false accept rates of these models in Section 5.

For both pipelines, the training process of the DenseNet is to the specification of Ahmad and Fuller [2], however we change our network from ResNet [17] to DenseNet. We choose DenseNet-169 as our feature extractor. The ResNet based architecture achieves an EER of 1.32% on the ND-0405 test set specified by Ahmad and Fuller whereas changing to a DenseNet based architecture improves the EER to 1.16%. The resulting EER is close to State of Art EER of 0.99% [54]. The DenseNet learns feature embeddings per iris by minimizing the triplet loss [44]. The trained network outputs a 1024 dimensional feature vector per iris. A simple augmentation of flipping the iris image along the horizontal axis is done and another feature vector is extracted. This augmentation improves the recognition accuracy finally yielding a feature vector of 2048 dimensional feature vector per iris.

The first pipeline uses normalized iris images obtained by using the USIT software. This feature vector serves as

the input for RESIST-DN.

The second pipeline is trained directly on segmented iris images (again using DenseNet). We use the same DenseNet variant and train it to convergence using the same iris (segmented only) images. Since USIT directly outputs normalized images for this pipeline we use segmentation pipeline of Ahmad and Fuller [1]. The segmentation accuracy of the segmentation tool is comparable to state of art. We call these two pipelines DenseNetNormal and DenseNet respectively.

3. Design

The *core* of RESIST is a convolutional network similar to an autoencoder. We train the core standalone and inside a GAN as the generator. The objective of an autoencoder is to compress its input to form a feature representation and then recreate the input from the compressed feature representation. That is, the autoencoder learns two functions f and f' where the range of f has smaller dimension and $\hat{y} = f'(f(x)) \approx x$ for all trained values.

We slightly abuse this framework, we use templates as the input value x and try to reconstruct iris images \hat{y} . Our core is based on U-Net [43] which has been used in pix2pix GANs [24] for image to image translation and fits our template inversion use case. The full architecture is presented in Section 3.3. We first describe the loss functions for the network core when trained standalone. We then describe integration in the GAN, and finally present the overall architecture.

3.1. Standalone Network Core Training

Our *standalone* core training works in two stages. First, we train the core using the L1 loss. Our models minimize the absolute differences (of pixels) between the predicted and the ground truth iris image. The L1 loss or mae (mean absolute error) provides crisper details in reconstructed images than L2 loss [51]. The L1 loss function:

$$L_{mae} = \frac{\sum_{h,w} |y(h,w) - \hat{y}(h,w)|}{h * w}, \quad (1)$$

where h, w denote the height and width of the image, y is the original image and \hat{y} is the output of the core. Finally, *mae* is the mean absolute error or the L1 loss. To fine tune the reconstructions we add two new losses. The first optimizes the Structural similarity (SSIM) [52] between images. Structural similarity measures the product of two terms. The first term of SSIM is roughly the product of the two image means normalized by the sum of means. The second term is the covariance of the two images normalized by the sum of variance. The actual definition adds an additional term in the normalization to account for small values. We minimize structural dissimilarity between the real and reconstructed images:

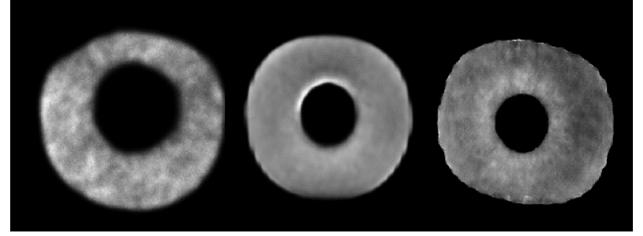


Figure 3: Iris reconstructions after addition of L1 loss, SSIM loss and perceptual loss from left to right.

$$L_{SSIM} = 1 - SSIM(y, \hat{y}) \quad (2)$$

The third loss function is the Perceptual loss [27]. This loss takes as input the reconstructed image from the first head (T2 in Figure 4) and passes it through a pre-trained VGG16 [46] network. This VGG16 network is trained on the ImageNet dataset [10] and not fine tuned on any iris dataset. The VGG16 network serves as a texture based feature extractor where an intermediary layer represents textures in the reconstructed iris image. Perceptual loss minimizes the L2 distance between textures of the real and reconstructed iris images.

$$L_{Perceptual} = \frac{(\|\phi_j(y) - \phi(\hat{y})\|)}{C_j H_j W_j} \quad (3)$$

where ϕ denotes activations of VGG16 after its 9th convolution operation, C, H, W are layer activation dimensions and y and \hat{y} are real and reconstructed iris images.

The need for the three loss functions To illustrate how the three loss functions work to improve the core, we discuss accuracy numbers of this standalone core using the DenseNet pipeline as an example. Reconstructions using just L_{mae} capture the boundary and some high level textures of the iris. However, such images have an average L1 distance of 8%. Furthermore, the quality of these images is not high, the left of Figure 3 shows a reconstruction using just L_{mae} . Adding L_{SSIM} improves the image quality and actually reduces L_{mae} to 4%. L_{SSIM} improves the reconstructions of the first head by extracting features from the output of the first head and trying to reconstruct an iris image by minimizing structural dissimilarity. The reconstructed irises now contain better textured iris patterns however they are not sharp relative to their equivalent real irises. This can be seen in the center iris of Figure 3. Finally, we employ $L_{Perceptual}$ loss [27]. Average L1 distance after using perceptual loss drops to 3%. Furthermore, the texture quality improves as well, see the right image in Figure 3.

To summarize, the standalone training for the core has two phases, first the core is trained on just the L1 loss to capture the iris structure. The second phase adds two losses and one extra head to the core. The two heads are T2 and T3 in Figure 4. Both L_{mae} and $L_{Perceptual}$ are on the T2

head. The T3 head minimizes L_{SSIM} . The standalone core is trained to minimize the loss:

$$L_{AE} = \alpha L_{Perceptual} + \beta L_{mae} + \gamma L_{SSIM} \quad (4)$$

Where α, β, γ are weight terms for the equation with values 2,1,1 respectively.

3.2. Improved Core Training in a GAN

We discussed training the core in a standalone fashion. Now we train the core as a generator inside a GAN. The idea for this comes from prior work which shows training autoencoders in an adversarial setting yields better results than training the standalone [37].

The final architecture is a GAN which takes as input a template and yields a reconstructed iris image (see Figure 4). In essence when RESIST is coupled with a feature extractor network for template inversion an autoencoder is formed where an iris image is converted to a template by the feature extraction network and back to an iris image by RESIST. A GAN has a generator which generates images and a discriminator which judges how good the generated images are. Our generator is the core discussed above. We use a recently proposed relativistic average discriminator [28] as our discriminator. To build up to the relativistic discriminator we first start with the original GAN loss functions:

$$L(D) = -\mathbb{E}_{y \sim \mathbb{P}_{real}} [\log(D(y))] \quad (5)$$

$$-\mathbb{E}_{\hat{y} \sim \mathbb{P}_{fake}} [\log(1 - D(\hat{y}))]$$

$$L(G) = -\mathbb{E}_{\hat{y} \sim \mathbb{P}_{fake}} [\log(D(\hat{y}))]. \quad (6)$$

$L(D)$ is called the discriminator loss and $L(G)$ is called the generator loss. y and \hat{y} are the original and reconstructed irises. In the original GAN a noise vector is sampled from a multivariate normal distribution with a mean of 0 and a variance of 1. This vector is transformed into an image by the generator while the discriminator compares the generated image and the ground truth image. \mathbb{P}_{real} is the distribution of real images and \mathbb{P}_{fake} is usually a multivariate normal distribution with mean of 0 and variance of 1, \mathbb{P}_{fake} is the distribution of the iris templates for RESIST.

Both losses are minimized using gradient descent. The generator and discriminator play a zero sum game. The generator weights are updated on how good its fake images are while the discriminator weights are updated on how well it differentiates between real and fake images. A discriminator outputs a probability, therefore the last layer of the discriminator is a sigmoid activation. To achieve this we add a sigmoid function σ as the last layer of a discriminator, we then define a discriminator D as:

$$D(y) : \sigma(C(y)) \quad (7)$$

Where σ is the sigmoid function and C is a layer in any discriminator preceding the final sigmoid output. This layer

is layer T4 in our architecture shown in figure 4. $C(x)$ is thus the non-transformed discriminator output (before the sigmoid function). The discriminator will ideally output 1 for real images and 0 for fake images.

We use the relativistic average discriminator in our work which judges the generated images relative to their real counterparts and vice versa. Then:

$$D_{Ra}(y) = \sigma(C(y - \mathbb{E}[C(y)])) \quad (8)$$

$$D_{Ra}(\hat{y}) = \sigma(C(\hat{y} - \mathbb{E}[C(y)])) \quad (9)$$

Where \mathbb{E} is an averaging operation which averages all fake data samples in a batch. With the definition of a discriminator above, the two original loss functions become:

$$L(D_{Ra}) = -\mathbb{E}_{y \sim \mathbb{P}_{real}} [\log(D_{Ra}(y))] \quad (10)$$

$$-\mathbb{E}_{\hat{y} \sim \mathbb{P}_{fake}} [\log(1 - D_{Ra}(\hat{y}))]$$

$$L(G_{Ra}) = -\mathbb{E}_{\hat{y} \sim \mathbb{P}_{fake}} [\log(D_{Ra}(\hat{y}))] \quad (11)$$

$$-\mathbb{E}_{y \sim \mathbb{P}_{real}} [\log(1 - D_{Ra}(y))]$$

\hat{y} is the generated output from the convolutional network core and y is the real iris image. The generator takes templates as input to generate reconstructed irises, \mathbb{P}_{fake} is therefore the distribution of the templates and \mathbb{P}_{real} is the distribution of the real iris images. The final loss for the generator of all variants of RESIST becomes:

$$L(G_{RESIST}) = L(G_{Ra}) + \alpha L_{Perceptual} + \beta L_{mae} + \gamma L_{SSIM} \quad (12)$$

To summarize, all RESIST variants are trained in two stages. The core is trained standalone using the process described in Section 3.1. This trained core is then used as the generator of a RESIST network which is trained as a GAN. Both training stages use the Adam optimizer [31].

Training In addition to using the relativistic discriminator and using an already trained generator we employ different techniques to stabilize the training of RESIST. Spectral normalization [38] is applied to both discriminator and generator. Convolution and dense operators are both normalized using their spectral norms. Spectral normalization provided sufficient stabilization for RESIST training, hence other stabilization techniques were not used. We also use two time scale update rule [19] for training, our learning rates are 1×10^{-5} and 1.5×10^{-5} for the generator and discriminator respectively. Gaussian noise sampled from a normal distribution of mean 0 and variance 1 is added to every pixel of real irises during training. RESIST variants are trained for 400 epochs (200 steps per epoch) with a batch size of 12. Each batch consists of y and \hat{y} pairs.

For RESIST-D and RESIST-DN the template is the output of the DenseNet based feature extractor and the DenseNet Normal feature extractors described in Section 2. From an architectural point of the view the only difference is the

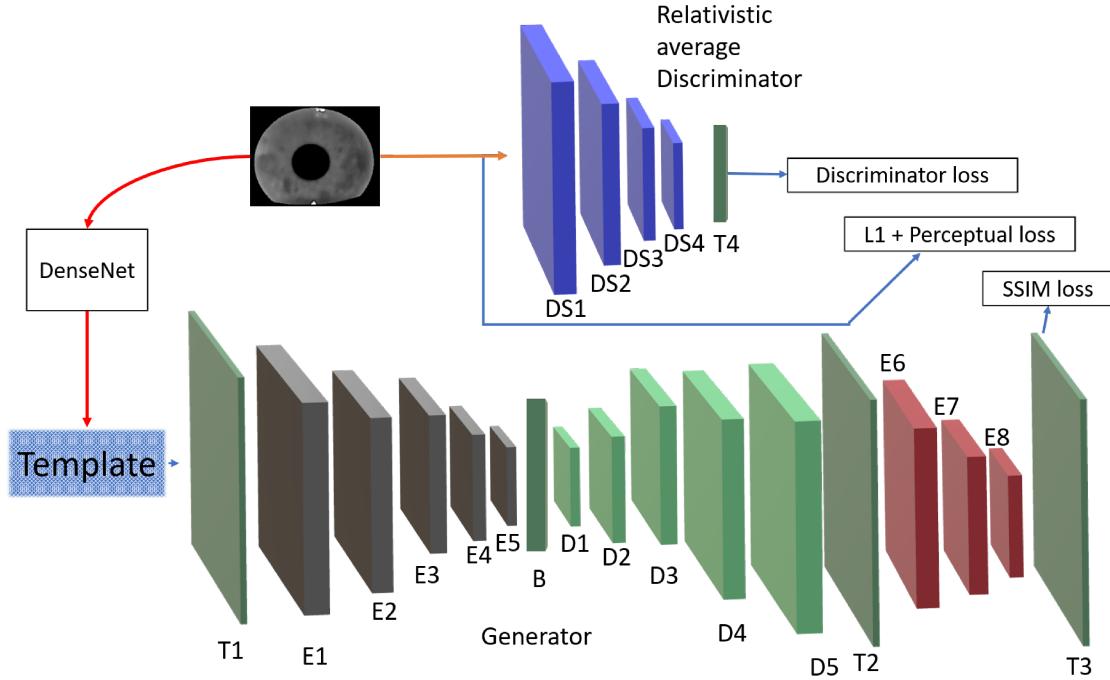


Figure 4: RESIST architecture. The generator reconstructs iris images from its corresponding template. The discriminator judges the reconstructed iris as real or fake. The DenseNet is not part of RESIST training. Perceptual loss is minimized by passing reconstructed iris from T2 to a VGG16 network which is not shown in the figure.

output dimension. RESIST-G differs from both RESIST-D and RESIST-DN in input templates because it is binary (instead of real-valued).

3.3. Architecture

RESIST architecture borrows concepts from U-Net [43] which has skip connections [17] from all its encoder layer to its decoder layers. The use of skip connections allows more gradient to flow through the network.

The layer and architecture details are in Tables 1 and 2 respectively. The architecture is shown using DenseNet as an example pipeline in Figure 4. First layer (T1) transforms the template to the same dimension of the iris image. This is done using a FC layer. The encoder consists of five layers (E1-5 in Table 2), each layer has a convolution operation with a LeakyReLU [35] activation followed by batchnormalization [23]. Dimensionality is reduced by strided convolution. The decoder also has five layers (D1-5 in table 2), each layer has deconvolution operation with a ReLU activation followed by batchnormalization. Each decoder layer accepts a skip connection from its corresponding encoder layer, as an example encoder layer 5’s gradient will be added to decoder layer 1’s gradient and so on for other layers in the encoder-decoder architecture in a U-Net fashion. Dimensionality is regained by strided deconvolution.

The first head of RESIST reconstructs the iris image us-

Layer	Elements
T1	Dense
E*	Conv(Stride 2),BN,LeakyReLU
D*	DeConv(Stride 2),BN,ReLU,Skip
T2	Conv
T3	Dense,Conv
DS*	Conv(Stride 2)
T4	Dense

Table 1: Layers

Layer	OutputSize	Kernels
Input	1x1024	-
T1	1x256x256	1
E{1-5}	Input Size/2	32,64,128,256,512,64,128,256
D{1-5}	Input Size*2	512,256,128,64,32
T2	1x256x256	1
T3	1x256x256	1
DS{1-4}	Input Size/2	128,128,128,128
T4	1	1

Table 2: GAN architecture

ing a convolution operation coupled with a sigmoid activation (T2) while minimizing the L1 loss from equation 1. This reconstructed image is passed through a VGG16 network to have its perceptual loss minimized shown in equation 3. The reconstructed image from the first head is also passed through another series of encoder layers and another reconstructed iris image is formed. These encoder layers accept skip connections from the encoder-decoder. This second

reconstruction is formed by using a dense layer with pixels equivalent to the dimension of the iris image followed by a convolution operation coupled with a sigmoid activation (T3) which minimizes the loss in equation 2.

The reconstructed image from the first head is output and fed to the discriminator along with its corresponding real image. The discriminator judges the reconstruction relative to the real iris image. The discriminator has four layers each comprising strided convolution operations (DS1-4). After the fourth convolution the features are flattened and passed to a dense layer which has a sigmoid activation outputting a probability (T4).

4. Evaluation

The ND-0405 [3] dataset contains 64,980 iris samples from 356 subjects and is a superset of the NIST Iris Challenge Evaluation dataset [40]. Iris images are captured using the LG 2200 (Near infrared) biometric system. The images have blurring from motion and some out of focus images.

Methodology We described how the two DenseNet models are trained in Section 2. All three RESIST variants are trained on features from all remaining left iris images from each subject (apart from the 25 images per subject used to train the feature extractors). The testing set for all three pipelines is 20% of right iris images. The training and testing set is kept same for all three pipelines. We consider the right and left irises of a subject as separate classes based on prior work showing statistical independent [9]. Training and testing sets are class disjoint. Furthermore, no training images for RESIST are used in training the DenseNet models.

DenseNet All images used to train DenseNet feature extractor and train RESIST-D have been segmented using a publicly available Mask R-CNN based iris segmentation tool [1]. The feature extractor (DenseNet) is trained on these images while minimizing the triplet loss and generating feature embeddings per iris. The segmented irises do not contain any non-iris occlusions. The segmented images are centered and resized to a resolution of 256x256 and directly fed to the feature extractor without any pre-processing.

DenseNetNormal For RESIST-DN the images are segmented and normalized using the USIT software. The images are of resolution 64x512. These images contain some non-iris occlusions (eyelid, eyelashes). We do not use the masks provided per iris image by the USIT software. The normalized iris images are directly fed to the normalized iris feature extractor.

Gabor Filter For the Gabor filter feature extractor, we use filter bank 5 from OSIRIS. The binarized output after convolution of normalized iris image with filter bank 5 of OSIRIS is our template. This is fed to RESIST-G which learns to reconstruct back the corresponding normalized iris image.

We explore two types of attacks defined by Mai et al. [36], **Type-1** attack where a reconstructed iris is matched with its corresponding real iris and **Type-2** attack where reconstructed irises are matched with real irises of the same class. Type-1 and Type-2 accuracy depend on a distance threshold. This distance threshold is varied to trade off between TAR/FAR. When we vary between TAR/FAR we are changing the threshold according to the feature extractors and observing the resulting change on RESIST. We also report **Rank-1** of our reconstructed irises. Rank-1 accuracy is an all-all matching and does not use a threshold. A reconstructed iris has its distance calculated across the entirety of the real dataset. A true positive in this case is if the lowest distance (between a reconstructed and a real iris) is with an iris of the same class. We use the cosine distance for our deep templates and the Hamming distance for Gabor filter templatebased template.

5. Results

We summarize our results in Table 3 and Figure 5. In Table 3, the legitimate row describes our three accuracy measures on legitimate irises. Note that Type-1 accuracy will always be 100% for a deterministic, legitimate transform (as the exact same template will be produced twice).

The first reconstruction is done using RESIST-D. The DenseNet based feature extractor is robust, achieving 99.9% rank-1 accuracy on legitimate irises on the test set. RESIST-D achieves a rank-1 accuracy of 96.29% when matching the reconstructed test set against the real test set. This means that in a recognition system a reconstructed iris would be classified as a member of the stored database 96% of the time a match is made. True acceptance rate at 1% false acceptance rate is 43.3% which represents the Type-2 accuracy rate for RESIST-D. Type-1 accuracy for RESIST-D is 76%. We also show results from our standalone network core.

Template inversion for deep templates for normalized irises works better than irises without normalization. The Rank-1 accuracy is lower at 89% for RESIST-DN vs 96% for RESIST-D. We attribute this to the normalized feature extractor model having worse accuracy compared to the segmented iris feature extractor (98% vs 99%). We also do not use any masks from the USIT software therefore the feature extractor has to learn extra information about occlusions. Eyelashes and eyelids are some common occlusions found in normalized iris images. RESIST-DN has higher Type-1 and Type-2 accuracy rates than RESIST-D. The Type-1 accuracy rate is especially high at 96% when compared to the Type-1 accuracy of RESIST-D at 76%.

Previous works have shown that Iriscode based templates can be inverted [14]. We achieve similar results with RESIST-G. Reconstruction accuracy across all three metrics is high. The Type-1 reconstruction accuracy for RESIST-G is 100%. This means that every iris in the test set reconstructed using

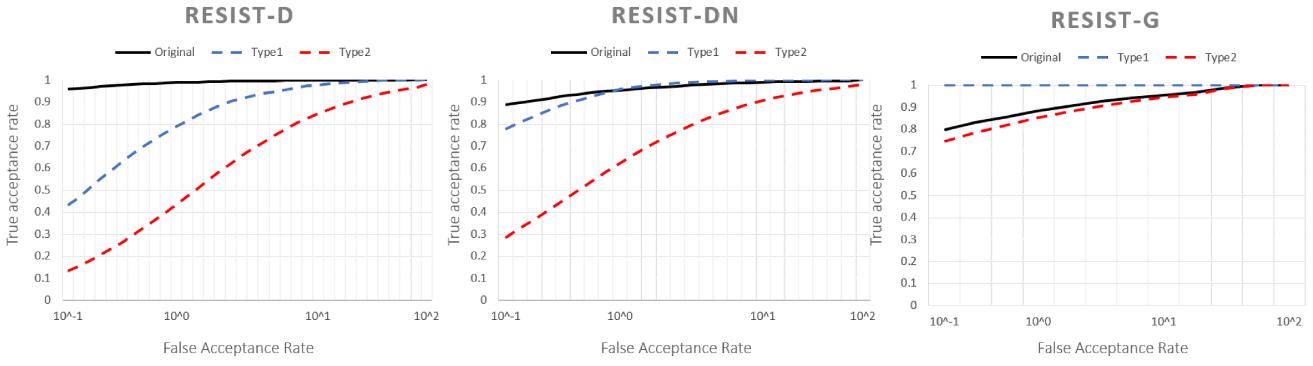


Figure 5: ROC curves for our three reconstruction pipelines. The x-axis is the false acceptance rate and the y-axis is the true acceptance rate. The Original curve is Type-2 error for the legitimate biometrics. For legitimate systems there is no Type-1 error.

Pipeline	Method	TAR		Rank-1
		Type1	Type2	
DenseNet	Legitimate	100%	98.7%	99.9%
	RESIST-D	75.9%	43.28%	96.3%
	Core	30.3%	18.5%	62.2%
DenseNet	Legitimate	100%	95.2%	98.7%
	Normal	96.7%	62.5%	89.3%
Gabor	Legitimate	100%	88.2%	97.9%
	RESIST-G	100%	85.2%	97.1%

Table 3: Reconstruction accuracy for 3 different pipelines. Legitimate shows the accuracy of the underlying pipeline, while the second row shows accuracy for RESIST images. True acceptance rate (TAR) is at 1.0% FAR.

RESIST-G matched with its real counterpart with a Hamming distance below the threshold. Type-2 and Rank-1 accuracy is comparable to accuracy numbers for the true dataset.

Discussion Reconstruction accuracy is proportional to recognition accuracy of a model. If all the important features are being captured by the feature extractor for comparison, then these features are also available for inversion.

Type-1 error validates RESIST networks learning features of a particular image while Type-2 error validates if the learnt features can accurately reconstruct irises of a specific person. The gap in these errors can be attributed to the variance among the irises of a single person.

The RESIST-D template is hardest to invert. We theorize this is due to limited spatial correlation between irises as the feature extraction network takes as input segmented only iris images. Only iris regions are present in the image which are hard to invert. Gabor filter template is the easiest to invert since the template contains texture information of the original iris. Occlusions in the Gabor filter template are well defined and are helpful in the inversion process. Among the deep templates normalized iris based templates are easier to invert than non-normalized iris based templates. We attribute this to spatial correlation in the irises and the presence of

occlusions. Some deep learning feature extractors use masks per iris image [54], we do not use any masks in our work. We hypothesize that masks would provide additional information to the reconstruction. We believe this would improve both recognition and reconstruction accuracy.

Type-1 accuracy for RESIST-DN is higher than the Type-2 original accuracy curve in Figure 5 after 1% percent FAR. We attribute this to sharp reconstructions since the reconstruction network behaves both as a reconstruction and super resolution network due to the presence of perceptual loss.

6. Conclusion

We study inverting three types of templates, Gabor filter template where prior work exists and has good inversion results and deep templates where there is limited prior work. Our iris reconstruction networks perform well on all three types of templates, with deep templates proving harder to invert than Gabor filter template. Within the two deep templates segmented only iris templates are harder to invert than normalized iris based templates. Our work reinforces the value of template protection mechanisms. In particular, it seems important to find template protection mechanisms that work on real valued feature vectors where comparison is via the cosine distance. Future work should look into template inversion across datasets, inversion of binary vectors produced from deep neural networks, and the feasibility of storing templates which have limited spatial correlation with their corresponding biometrics. Template inversion in the presence of template protection such as encryption or random sampling is also another direction for future work.

Acknowledgements

The authors thank the reviewers for their valuable help in improving the manuscript. This work was supported in part by Synchrony Financial Inc. and NSF Grant 1849904.

References

- [1] S. Ahmad and B. Fuller. Unconstrained iris segmentation using convolutional neural networks. In *Proceedings of the Asian Conference on Computer Vision*. Springer, 2018.
- [2] S. Ahmad and B. Fuller. Thirdeye: Triplet based iris recognition without normalization. In *IEEE International Conference on Biometrics: Theory, Applications, and Systems*, 2019.
- [3] K. W. Bowyer and P. J. Flynn. The ND-IRIS-0405 iris image dataset. *arXiv preprint arXiv:1606.04853*, 2016.
- [4] A. Boyd, A. Czajka, and K. Bowyer. Deep learning-based feature extraction in iris recognition: Use existing models, fine-tune or train from scratch? *arXiv preprint arXiv:2002.08916*, 2020.
- [5] J. Bringer, H. Chabanne, G. Cohen, B. Kindarji, and G. Zémor. Optimal iris fuzzy sketches. In *2007 First IEEE International Conference on Biometrics: Theory, Applications, and Systems*, pages 1–6. IEEE, 2007.
- [6] J. Bringer, C. Morel, and C. Rathgeb. Security analysis of bloom filter-based iris biometric template protection. In *2015 international conference on biometrics (ICB)*, pages 527–534. IEEE, 2015.
- [7] C. Chen and A. Ross. A multi-task convolutional neural network for joint iris detection and presentation attack detection. In *2018 IEEE Winter Applications of Computer Vision Workshops (WACVW)*, pages 44–51. IEEE, 2018.
- [8] J. Connell, N. Ratha, J. Gentile, and R. Bolle. Fake iris detection using structured light. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8692–8696. IEEE, 2013.
- [9] J. Daugman. Iris recognition border-crossing system in the uae. *International Airport Review*, 8(2), 2004.
- [10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [11] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM journal on computing*, 38(1):97–139, 2008.
- [12] J. S. Doyle and K. W. Bowyer. Robust detection of textured contact lenses in iris recognition using bsif. *IEEE Access*, 3:1672–1683, 2015.
- [13] J. Galbally, A. Ross, M. Gomez-Barrero, J. Fierrez, and J. Ortega-Garcia. From the iriscode to the iris: A new vulnerability of iris recognition systems. *Black Hat Briefings USA*, 1, 2012.
- [14] J. Galbally, A. Ross, M. Gomez-Barrero, J. Fierrez, and J. Ortega-Garcia. Iris image reconstruction from binary templates: An efficient probabilistic approach based on genetic algorithms. *Computer Vision and Image Understanding*, 117(10):1512–1525, 2013.
- [15] M. Gomez-Barrero, C. Rathgeb, J. Galbally, C. Busch, and J. Fierrez. Unlinkable and irreversible biometric template protection based on bloom filters. *Information Sciences*, 370:18–32, 2016.
- [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [18] F. Hernández Álvarez, L. Hernández Encinas, and C. Sánchez Ávila. Biometric fuzzy extractor scheme for iris templates. 2009.
- [19] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.
- [20] H. Hofbauer, F. Alonso-Fernandez, P. Wild, J. Bigun, and A. Uhl. A ground truth for iris segmentation. In *2014 22nd international conference on pattern recognition*, pages 527–532. IEEE, 2014.
- [21] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [22] Y. Huang, A. W. K. Kong, and K.-Y. Lam. From the perspective of cnn to adversarial iris images. In *2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–10. IEEE, 2018.
- [23] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML’15*, page 448–456. JMLR.org, 2015.
- [24] P. Isola, J.-Y. Zhu, T. Zhou, and A. 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.
- [25] G. Itkis, V. Chandar, B. W. Fuller, J. P. Campbell, and R. K. Cunningham. Iris biometric security challenges and possible solutions: For your eyes only? using the iris as a key. *IEEE Signal Processing Magazine*, 32(5):42–53, 2015.
- [26] Z. Jin, J. Y. Hwang, Y.-L. Lai, S. Kim, and A. B. J. Teoh. Ranking-based locality sensitive hashing-enabled cancelable biometrics: Index-of-max hashing. *IEEE Transactions on Information Forensics and Security*, 13(2):393–407, 2017.
- [27] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [28] A. Jolicoeur-Martineau. The relativistic discriminator: a key element missing from standard gan. *arXiv preprint arXiv:1807.00734*, 2018.
- [29] A. Juels and M. Sudan. A fuzzy vault scheme. *Designs, Codes and Cryptography*, 38(2):237–257, 2006.
- [30] A. Juels and M. Wattenberg. A fuzzy commitment scheme. In *Proceedings of the 6th ACM conference on Computer and communications security*, pages 28–36, 1999.
- [31] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [32] N. Kohli, D. Yadav, M. Vatsa, R. Singh, and A. Noore. Detecting medley of iris spoofing attacks using desist. In *2016*

- IEEE 8th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–6. IEEE, 2016.
- [33] N. Kohli, D. Yadav, M. Vatsa, R. Singh, and A. Noore. Synthetic iris presentation attack using idcgan. In *2017 IEEE International Joint Conference on Biometrics (IJB)*, pages 674–680. IEEE, 2017.
- [34] P. C. Kronfeld. The gross anatomy and embryology of the eye. In *Vegetative Physiology and Biochemistry*, pages 1–62. Elsevier, 1962.
- [35] A. L. Maas, A. Y. Hannun, and A. Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, page 3, 2013.
- [36] G. Mai, K. Cao, P. C. Yuen, and A. K. Jain. On the reconstruction of face images from deep face templates. *IEEE transactions on pattern analysis and machine intelligence*, 41(5):1188–1202, 2018.
- [37] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- [38] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018.
- [39] N. Othman, B. Dorizzi, and S. Garcia-Salicetti. Osiris: An open source iris recognition software. *Pattern Recognition Letters*, 82:124–131, 2016.
- [40] P. J. Phillips, K. W. Bowyer, P. J. Flynn, X. Liu, and W. T. Scruggs. The iris challenge evaluation 2005. In *Biometrics: Theory, Applications and Systems, 2008. BTAS 2008. 2nd IEEE International Conference on*, pages 1–8. IEEE, 2008.
- [41] K. B. Raja, R. Raghavendra, and C. Busch. Color adaptive quantized patterns for presentation attack detection in ocular biometric systems. In *Proceedings of the 9th International Conference on Security of Information and Networks*, pages 9–15, 2016.
- [42] N. K. Ratha, J. H. Connell, and R. M. Bolle. Enhancing security and privacy in biometrics-based authentication systems. *IBM systems Journal*, 40(3):614–634, 2001.
- [43] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [44] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [45] S. Simhadri, J. Steel, and B. Fuller. Reusable authentication from the iris. In *Information Security Conference*, 2019.
- [46] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
- [47] S. Soleymani, A. Dabouei, J. Dawson, and N. M. Nasrabadi. Adversarial examples to fool iris recognition systems. *arXiv preprint arXiv:1906.09300*, 2019.
- [48] M. Stokkenes, R. Ramachandra, M. K. Sigaard, K. Raja, M. Gomez-Barrero, and C. Busch. Multi-biometric template protection—a security analysis of binarized statistical features for bloom filters on smartphones. In *2016 Sixth International Conference on Image Processing Theory, Tools and Applications (IPTA)*, pages 1–6. IEEE, 2016.
- [49] M. Trokielewicz, A. Czajka, and P. Maciejewicz. Post-mortem iris recognition with deep-learning-based image segmentation. *Image and Vision Computing*, 94:103866, 2020.
- [50] S. Venugopalan and M. Savvides. How to generate spoofed irises from an iris code template. *IEEE Transactions on Information Forensics and Security*, 6(2):385–395, 2011.
- [51] X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y. Qiao, and C. Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018.
- [52] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [53] S. Yadav, C. Chen, and A. Ross. Synthesizing iris images using rasgan with application in presentation attack detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [54] Z. Zhao and A. Kumar. Towards more accurate iris recognition using deeply learned spatially corresponding features. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3809–3818, 2017.
- [55] J. Zuo, N. K. Ratha, and J. H. Connell. Cancelable iris biometric. In *2008 19th International Conference on Pattern Recognition*, pages 1–4. IEEE, 2008.