

Challenges in Off-angle to Frontal Iris Image Conversion using Pix2Pix Generative Adversarial Networks

Jitendra Sai Kota and Mahmut Karakaya

Kennesaw State University

Marietta, GA 30060 USA

`jkota@students.kennesaw.edu`, `mkarakay@kennesaw.edu`

Abstract

Person identification using biometrics has become a safer and trustworthy mechanism with the advancement of technology. Among all biometric identification methods, iris recognition has achieved very low false acceptance rates due to its complex and unique patterns. The low acceptance rates apply only to frontal iris images. Capturing frontal iris images is not always possible, especially in uncontrolled environments, where most of the iris images captured tend to be non-ideal, such as off-angle images. Off-angle iris images suffer from several issues, including corneal refraction, limbus occlusion, the effect of gaze angle, and depth of field blur. These effects distort the iris patterns, causing the similarity scores between the same individual to widen and scores between different individuals to become closer. This also causes false acceptance rates to increase, as it increases the chances of misclassification. This highlights the need for improving the performance of off-angle iris recognition.

By leveraging the low false-acceptance rates of the frontal iris images, we build generated frontal version of the iris images using off-angle iris images and achieved better performance compared with the perspective transformation. We built a modified version of the Pix2Pix GAN to achieve the frontal projection of off-angle iris images. Instead of using a Mean Squared loss function in the Pix2Pix GAN, we use a combination of Mean Squared loss function, Matrix Multiplication loss, and SSIM loss function to generate sharper images that can capture the textural information of the original image better.

1. Introduction

The practice of biometric human identification is a longstanding concept. Over time, humans have utilized multiple sensory cues, including facial features, vocal nuances, and occasional tactile interactions, to facilitate the recognition and differentiation of individuals. These methods of identification are adequate when subjects are in close proximity. As human civilizations began expanding,

there was a need to identify other humans from a distance. This problem gave rise to other identification mechanisms, such as signatures, passwords, and encrypted messages. However, these methods are highly prone to manipulation and are not always reliable. This led humans to go back to using biometrics as an identification mechanism even when they are not in proximity. Fingerprints, handprints, and footprints were used for a long time to handle biometric identification.

Amongst these, fingerprints remain the most widely used biometric identification method, even in contemporary times. With continuous advancements in camera technology, facial recognition has emerged as a prominent area of research. Facial recognition relies on various facial features, including the shape and position of the eyes, ears, nose, and mouth, along with the skin. Although frontal facial recognition has achieved low false acceptance rates historically, the limitation lies in requiring subjects to face the camera directly [1]. Attempts have been made to address the limitation in a long-range biometric setup by algorithms reconstructing missing facial features in profile images, but asymmetries in the face remain challenging to restore accurately. Other biometric identification methods such as finger vein pattern and palm vein pattern have also been explored. Finger vein pattern uses the veins present underneath the skin of a finger, making them more secure than fingerprints. Palm vein pattern is a better version of finger vein pattern because it uses more reference points [2]. However, control systems that use this technology are expensive. The availability of high-quality cameras at low costs has also ensured that people can now use irises of humans for recognition. Unlike fingerprints, which are prone to duplication because of their high false acceptance rates, irises are more reliable because of their low false acceptance rates. Iris has a lot more unique characteristics than fingerprints, allowing for more comparisons between the images, reducing the possibility of misclassification [3].

Iris recognition, despite being one of the best biometric identification systems in the market, is not as widely accepted as fingerprints. This is attributed to several factors. One of the biggest reasons is the discomfort

humans face when their irises are being scanned. They must assume a certain posture and face the camera to let their eyes be scanned. Another big drawback is the speed at which the iris scanning is done. Since iris recognition works well only when the iris images are frontal in nature, that is, the person must be facing the camera and looking at it, it is hard to always capture such an image straight away. So, iris systems capture multiple images or a video stream to capture the correct image. This greatly impacts the time taken to successfully identify a person using this system.

The above problems created the necessity to build standoff iris recognition systems. Some of the unconstrained iris recognition systems capture images from long range in the visible spectrum, as traditional iris systems capture images in the near-infrared spectrum, resulting in noisy and blurry images. One such problem that could speed up iris recognition is the problem of recognizing off-angle iris images.

The off-angle iris images suffer from various challenges, such as reflections, gaze angle, corneal refraction, depth of field blur, limbus occlusion, and the complex 3D structure of the iris. All these challenges mean that intra-class hamming distance scores increase, and the inter-class hamming distance scores decrease when compared with the frontal iris images. This would make it harder for the iris recognition systems to differentiate individuals, when the iris images are captured at off-angles, which is the less intrusive way of capturing the iris images. This paper discusses a technique to convert the off-angle iris images into their frontal versions to speed up the process of iris recognition and make it more reliable. The rest of the paper is organized as follows: Section 2 explores the related work in this field. Section 3 discusses the methodology. Section 4 demonstrates the experiments conducted and the results. The paper is concluded in Section 5.

2. Related Work

The earliest work in frontal iris recognition was done by Daugman [4]. They used the Gabor-phase quadrant feature descriptor for iris recognition. These filters converted the normalized iris images into a binary image called an iris code. These iris codes were then used for identification. Hamming distances were calculated for the iris codes, and the intra-class and inter-class distances were used to distinguish one person from another.

Several challenges, such as corneal refraction, depth-of-field blur, three-dimensional structure of the iris, the limbus effect, and the gaze angle, did not play an important role when comparing iris images taken from a frontal angle [5]. However, these challenges significantly affected the iris images when they were taken from an off-angle. Daugman [6] suggested a preprocessing technique that used affine transformations to counter the geometric

deformations that occurred due to the difference in gaze angle. However, this method did not factor other challenges of off-angle images. Zuo et al. [7] used ellipses to fit iris boundaries to correct for the geometric deformations, and Li et al. [8] used support vector machines to classify frontal and off-angle iris images and built separate models for both sets of images. However, these methods did not work for off-angle images captured at a gaze angle of over 30° .

A few others have worked on the other set of challenges plaguing off-angle iris recognition. Santos-Villalobos et al. [9] used the ray tracing method to compute the effect of corneal refraction on images and provided a way to convert off-angle iris images to their frontal versions. This method worked well for synthetic images but failed for real images due to the presence of the limbus. Karakaya et al. [10,11] investigated the effect of limbus occlusion and the three-dimensional structure of the iris on off-angle iris recognition. They noticed that the presence of the limbus could significantly increase the intra-class hamming distance scores and bring them closer to the inter-class hamming distance scores. They also noticed that the effect of the limbus was compounded for images captured at higher gaze angles.

Generating a related image based on a seed image has been studied recently in different areas of interest. One such area of interest is the generation of face images. Zhang et al. [12] used a Conditional Adversarial Autoencoder to generate a person's face images corresponding to different ages, given one image with the corresponding age label. A Conditional Adversarial Autoencoder consists of an encoder that transforms the input face image into a latent space vector. This vector is coupled with an age label, and traversing along forward and backward directions in the latent space can produce faces corresponding to different ages. A generator is then used to map the latent space vector to an output face image. However, this methodology cannot be extended to iris images as the traversal alone is not sufficient to learn the iris patterns, and the traversal in latent space based on the off-angle does not lead to frontal image conversion. Another way to generate images is by using a Generative Adversarial Network (GAN) [13]. GAN has a generator and a discriminator. The generator takes random noise as input and outputs an image, whereas the discriminator distinguishes a fake image from a real one. They both are trained simultaneously to generate highly realistic images. Conditional Generative Adversarial Networks (CGANs) use label information to generate images corresponding to a particular class. The main difference between GAN and CGAN is that the generator in CGAN takes noise and a label as input. There are many versions of CGANs. Taherkhani et al. [14] worked on matching profile face images with their frontal versions in a database, by using couple of conditional generative adversarial networks.




Figure 1: Illustration of off-angle iris recognition pipeline using frontal image conversion.

Their model comprises two GAN modules. The generators use a U-Net architecture, whereas the discriminator is a patch-based discriminator. The two modules are connected by Contrastive loss function. Their main objective is to project the profile and frontal face images to a latent space using generative modeling and use it for face matching. They also suggest that the model can be used for frontal face generation.

Pix2Pix GAN [15] is a type of CGAN that does Image-to-Image translation, using a similar but simpler architecture. The label in CGAN becomes an image, enabling the operation of translation. The generator used in Pix2Pix GAN is a U-net with skip connections, and the discriminator is a Patch-based discriminator. This method uses an adversarial loss and an L-1 loss to generate realistic images. Pix2Pix GAN is domain independent and can be used for image translations of any nature.

Using GANs to convert off-angle iris images into frontal iris images has not been extensively studied in the literature. This paper is one such attempt at efficiently generating frontal iris images using their off-angle versions. The paper specifically uses variations of the Pix2Pix GAN networks.

3. Methodology

There exist several ways to convert off-angle iris image to frontal view. First of all, directly employing the captured off-angle image for the purpose of converting to its frontal images presents inherent challenges due to the coexistence of various facial and eye structures within the image, including the skin, eyelids, pupil, and sclera. Second, converting the segmented and masked iris image to its frontal view appears to be a practical solution. However, even the segmented iris image can suffer from problems such as pupil dilation that can lead to fluctuations in the iris size by contracting and expanding. To address such challenges, we utilized the normalized iris image as a preferable solution. Normalization standardizes the input image by allowing for more consistent and reliable processing. Therefore, the effect of pupil dilation can be minimized, and it will improve the accuracy and stability

of off-angle to frontal iris image conversion. In the proposed method, the iris texture in the eye image is segmented and normalized to a standard 64×512 image to be used as the input of the frontal conversion module. Since the inputs are normalized images, the generated outputs will be normalized images. Comparisons between iris images are done using the iris codes of these generated normalized images with the iris codes of the original frontal images.

The proposed methodology of the off-angle iris recognition using frontal image conversion is illustrated in Fig. 1. It is developed in two separate branches including frontal iris database generation and frontal image conversion. The frontal iris database is generated using the traditional pipeline where frontal iris image is captured, segmented, normalized, and encoded before including in the database for later comparison with probe images. The second branch of the proposed methodology is focused on the generation of the encoded iris code as the probe image. First, the captured off-angle iris image is segmented to fit ellipses to its inner and outer boundaries. After segmentation, the angle of the off-angle segmented iris image is predicted using the gaze estimation method [16]. At the same time, the segmented off-angle iris image is normalized. Then, the estimated gaze angle and normalized off-angle image are used for frontal iris conversion. The frontal iris conversion generates a normalized frontal version of the off-angle image. The generated normalized iris image is then converted into a binary code using Gabor filters [4]. Finally, the generated iris code is compared with other codes in the database using hamming distance. If the frontal and off-angle image belong to the same subject, the hamming distance score is expected to be low. If they belong to different subjects, the hamming distance score will be high.

One of the important initial steps of the off-angle to frontal iris image conversion is the estimation of the gaze angle, i.e., the angle of the iris with respect to the camera. Diab et al. [16] have used CNN models with a regression analysis for gaze estimation. Their method showed an average error of 4° in angle. Our proposed method utilizes

this gaze estimation model to achieve the task of off-angle to frontal iris image conversion.

This paper adopts the Pix2PixGAN architecture [15] for the frontal iris reconstruction. The generator of the architecture is comprised of a U-net with skip connections. The input to a generator is a normalized off-angle iris image. U-net comprises an encoder and a decoder. The encoder acts as a down sampler, and the decoder acts as an up sampler. A single block in the encoder consists of 3 layers – Convolution, Batch normalization, and ReLU. A single block in the decoder consists of 3 layers too including dropout for the first three layers, transposed convolution, and batch normalization. The encoder and decoder are also connected using skip connections. The discriminator is a patch-based discriminator [15].

3.1. Loss functions

The original Pix2Pix GAN generator has two loss functions for generating images: (i) Sigmoid Cross Entropy loss function (L_{CGAN}), (ii) L1 loss (L_{L1}). The Sigmoid cross entropy loss function compares the discriminator's output of the generated images to an array of ones, while the L1 loss calculates the mean absolute error between the generated and original images.

These Pix2Pix GAN loss functions can be expressed as follows:

$$L_{CGAN} = E_y[\log D(y)] + E_{x,z} \left[\log \left((1 - D(x, G(x, z))) \right) \right] \quad (1)$$

$$L_{L1} = E_{x,y,z} [|y - G(x, z)|] \quad (2)$$

where D is the discriminator, G is the generator, x is the input off-angle iris image, y is the expected real frontal iris image, z is random noise.

In addition to these two loss functions, we have added two more for the generator. The first is the Matrix Multiplication loss function, which is the dot product of the generated and original image pixel values. The average of the resulting matrix is used as the loss function value. This function helps to better understand the texture of the image by considering the interaction of pixel intensities from different parts of the image.

The Matrix Multiplication loss can be expressed as:

$$L_{MM} = E_{x,y,z} [y * G(x, z)] \quad (3)$$

where G is the generator, x is the input off-angle iris image, z is random noise, y is the expected real frontal iris image, L_{MM} is the Matrix Multiplication loss function.

The second loss function is the Structural Similarity Index (SSIM) loss function [17]. While the L1 loss only captures the difference in the value of the grayscale intensities at a given pixel, the SSIM loss extracts three key features from images - luminance, contrast, and structure.

For iris images, all three features play a significant role in distinguishing individuals. The expressions for luminance, contrast, and structure are given as follows:

$$\mu_x = \frac{1}{N} \sum_{i=1}^N x_i \quad (4)$$

$$\sigma_x = \left(\frac{1}{N-1} \sum_{i=1}^N (x_i - \mu_x)^2 \right)^{\frac{1}{2}} \quad (5)$$

$$s_x = \frac{(x - \mu_x)}{\sigma_x} \quad (6)$$

where μ_x (luminance) is the mean value of all pixels, σ_x (contrast) is the standard deviation, and s_x (structure) is the normalized version of intensity.

The formulae for the comparison of luminance, contrast, and structure for two different images are as follows:

$$l(x, y) = \frac{2\mu_x\mu_y + C_1}{\mu_x^2 + \mu_y^2 + C_1} \quad (7)$$

$$c(x, y) = \frac{2\mu_x\mu_y + C_2}{\mu_x^2 + \mu_y^2 + C_2} \quad (8)$$

$$s(x, y) = \frac{\sigma_{xy} + C_3}{\sigma_x\sigma_y + C_3} \quad (9)$$

where $l(x, y)$, $c(x, y)$, and $s(x, y)$ are the comparative measures for luminance, contrast, and structure for the off-angle input iris image x and the expected real frontal iris image y . C_1, C_2, C_3 are the constants.

The final expression for SSIM is a combination of luminance, contrast, and structure. This is given as follows:

$$SSIM(x, y) = [l(x, y)]^\alpha \cdot [c(x, y)]^\beta \cdot [s(x, y)]^\gamma \quad (10)$$

$$L_{SSIM} = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \quad (11)$$

where L_{SSIM} is the SSIM loss function.

Using (1), (2), (3) and (11), the new generator loss function for the conversion of off-angle iris images to their frontal versions is given as follows:

$$L_{total} = L_{CGAN} + \lambda \cdot L_{L1} + L_{MM} + L_{SSIM} \quad (12)$$

where λ is the regularization parameter set as 100 in the original Pix2Pix GAN paper.




Figure 2: Sample iris images (a) frontal, (b) off-angle.

4. Experimental Setup and Results

The dataset [10] used in the paper consists of frontal and off-angle iris images captured using a moving camera in a controlled environment. Images were captured from 100 subjects at angles ranging from -50° to $+50^\circ$ with an increment of 10° . At each angle, 10 images were captured for each subject. The negative effects of gaze angle on iris images are not as vivid at lower angles, i.e., at 10° and 20° , because the elliptical unwrapping normalization method can tolerate challenges such as refraction of light at cornea and occlusion at limbus [18]. Therefore, only the results of images captured at 30° , 40° , and 50° angles were included in this study. The images are further divided into five folds with an 80/20 train-test split to conduct different sets of experiments. Segmentation of the pupil, iris, and eyelids is done by fitting ellipses to boundaries of the iris and pupil and using two quadratic curves for the eyelids.

The first fold comprises 80 subjects as the training dataset and the last 20 subjects as the test dataset. Training dataset has everything except the first 20 subjects. The first 20 are used as the test dataset for this fold. Similarly, following a cyclical approach, five different training and test datasets were created. This ensured that every subject in the dataset is evaluated at least once. Since we considered six different sets of angles (i.e., 30° , 40° , 50° , -30° , -40° , and -50°), we further partitioned each dataset based on angles. The proposed Pix2Pix GAN model is applied on these 30 datasets separately, resulting in 30 different models. To compare the generated normalized images with the original normalized images, we used Gabor filters to generate iris codes and hamming distance to measure their similarity. Gabor filters are spatial image filters that convert the grayscale-normalized image into a binary image. Hamming distances are calculated using these binary images. The formula for hamming distance calculation is as follows:

$$HD = \frac{|| (c_A \oplus c_B) \cap (m_A \cap m_B) ||}{|| m_A \cap m_B ||} \quad (13)$$

where, c_A is the iris code of the first image, c_B is the iris code of the second image, m_A is the mask of the first image, m_B is the mask of the second image.




Figure 3: Normalized iris images of (a) frontal iris, (b) off angle iris, (c) generated iris image of the off-angle image in (b).

In the rest of the paper, we will refer to the normalized iris images as iris images, interchangeable. For images belonging to the same subject, the hamming distance scores should be lower than for images belonging to different subjects. These scores are referred to as intra-class hamming distance and inter-class hamming distance, respectively. Ideally, inter-class hamming distance scores should be high and intra-class hamming distance scores should be low. The baseline for comparison in our experiments is the hamming distance scores between off-angle and frontal iris images. We will compare them with hamming distance scores between the generated iris images and the frontal iris images. We analyze their performance using Receiver Operating Curve (ROC) and compare their Area Under Curve.

To demonstrate the effect of gaze angle, the frontal and the off-angle eye images of the same subject are shown in Fig. 2. We normalized iris images using elliptical unwrapping and showed their normalized iris images in Fig. 3(a-b), respectively. Since Pix2Pix GAN model requires the input size as 256×256 , we normalized the iris images as 256×256 . The generated normalized image of off-angle image in (b) using proposed Pix2Pix GAN is shown in Fig. 3(c). The dark abnormal structures at the bottom of each normalized image are the upper eyelid. Notice how the texture in (c) becomes closer to the original frontal image in (a). This is in stark contrast to the frontal image in (a) and the off-angle image in (b).




Figure 4: Distributions of hamming distance scores for (top) the frontal vs. off angle iris images (-50° in angle) (bottom) the frontal vs. generated frontal images.

The histogram plots of intra-class and inter-class hamming distance distributions for original off-angle iris images and the generated iris images are presented in Fig. 4. The test subject is compared with 19 other subjects for inter-class evaluations. Fig. 4(top) shows the hamming distance scores for traditional approach where frontal and off-angle images are compared with elliptical unwrapping. Fig. 4(bottom) shows the distribution of hamming distance scores for proposed method where generated frontal images compared with frontal images. We observed that the hamming distance plots reveal that the intra-class and inter-class hamming distances overlap differently at traditional and proposed methods. The mean intra-class hamming distance score for the off-angle images is 0.40, whereas the intra-class hamming distance score goes down to 0.36 for the generated images. In addition, we analyzed their performance using the ROC plots as shown in Fig. 5. It also confirms that the generated images perform better than the baseline traditional method with a higher accuracy. This demonstrates that the proposed method can identify the subject much better than the baseline model with off-angle images. On the other hand, the mean inter-class hamming distance score for the off-angle images is




Figure 5: Performance analysis using ROC plots for off-angle images at -50° in angle vs. generated frontal images for the test subjects.




Figure 6: Performance analysis using ROC plots for off-angle images at -50° in angle for different loss functions.

0.47, whereas the same metric for the generated frontal images is 0.49. This shows that the new model can distinguish the subject from the rest much better than the baseline model with off-angle images.

To evaluate the performance of this model against other loss functions, we considered different combinations of loss functions including Wasserstein loss [19], SSIM loss, L1 loss, and Matrix Multiplication. In this set of experiments, we considered -50° off-angle iris images in the dataset where 85% images were in training and 15% in test. Fig. 6 compares their performances using ROC plots. For the first combination, we considered the Wasserstein loss function applied at the pixel level, the L1 loss function, and the SSIM loss function. The Wasserstein loss function is a loss function that increases the gap between real and generated images. The second combination is the Wasserstein loss function applied at the pixel level without the calculation of the mean of the matrices, the L1 loss function, and the SSIM loss function. The third




Figure 7: Performance analysis of all subjects using ROC plots for off-angle images at (top) -40° and (bottom) -50° in angle.

combination is the base Pix2Pix model. The fourth combination is the loss function described in this paper as the Matrix Multiplication loss, in combination with the SSIM loss and L1 losses. We observed that the Matrix Multiplication loss with the combination of the SSIM loss function showed a better performance compared with other loss functions. It highlights the superiority of the proposed model over the existing ones.

Due to the heterogeneity of the dataset, this improvement over the traditional approach is not consistent. The overall performance comparison of the proposed Pix2Pix GAN model and baseline is shown in Fig. 7 for off-angle images at -40° and -50° in angle. For -40° off-angles images, improvement is limited where their equal error rate (EER) values are similar as 0.002 and the Area Under Curve (AUC) are 0.989 and 0.994, respectively. For off-angle images at -50° in angle, AUC improves from 0.907 to 0.959 where their EER are 0.014 to 0.008 for proposed and traditional method. The off-angle images at -30° and +30° in angle did not show any improvement, as the images were already being identified without any errors, where they have perfect ROC. This is the reason why smaller angles were also not considered in further analysis.

Although this does not account for most of the tested datasets, a closer look at each of the above models revealed many more cases where the improvised Pix2Pix GAN model is working better than the base off-angle iris recognition case. Some of them are detailed below.

The proposed method performed the best on off-angle images at -50° and -40° in angle. A closer inspection of each individual subject revealed that among the -50° images, 95 out of 100 test subjects had perfect ROC curves. The proposed method exceeded the baseline model in three of the imperfect cases, whereas the baseline was still the best in the other three cases. Some of the commonalities among the images on which the proposed method worked better were eyelids occluding the iris view and the presence of limbus to some extent. For those images with no improvement, the iris texture seemed to be of low contrast. The proposed method also performed well on off-angle images at -40° in angle. The performance of two out of all subjects had been improved. The performance of one out of the remaining 99 subjects could not be improved. The remaining 98 cases already exhibited a flawless ROC. The presence of eyelids with good contrast iris images was improved by the model, while those with low contrast were not improved. The decreasing number of challenging images at -50° to -40° in angle explains the reason behind the near-perfect recognition performance at -30° angles.

For positive angles, the proposed model did not perform as well as the baseline model with off-angle iris images at +40° and +50° in angle. However, closer individual observation revealed that the proposed method is performing better than what the cumulative ROC curves suggest. For 11 out of all subjects, the proposed method yielded better ROC curves. These images shared common issues such as pupil dilation and a blur effect. For another 11 out of the remaining 90 subjects, performance was not improved. These subjects suffered from low contrast iris texture, the 3D structure of the iris, uneven limbus, and a partial occlusion of iris portion by nose. Nose occlusion is not an issue in other angles because +50° is the most extreme of all angles. In addition, the visual and geometric axis of the eye do not overlap. When the subject looks to the camera, there is an 8° angle between the cameras and the eye. This makes the appearance of iris images captured at +50° angle as +58° [20]. This explains why the iris image conversion for positive gaze angles requires a more complex approach compared with the negative angles in left eyes. The remaining 79 subjects have a perfect ROC. Among the off-angle images at +40° in angle, only 3 out of all subjects achieved better performance than the baseline, while the remaining subjects have a perfect ROC.

5. Conclusion and Future Work

Frontal iris images have traditionally obtained very low false acceptance rates because of the huge gap between

intra-class and inter-class hamming distance scores. The iris images are distinguishable with high accuracy even in off-angle images. However, the thresholds of intra-class and inter-class hamming distances come closer, causing concern. The concern becomes a bigger problem when the two hamming distances overlap, making identification more challenging.

This paper suggests a method to convert off-angle iris images to frontal iris images, thereby broadening the gap between intra-class and inter-class hamming distance scores. The use of matrix multiplication as a loss function, in addition to L1 loss and SSIM loss functions in a Pix2Pix GAN network, has yielded significant improvements in performance for images captured at -50° and -40° angles. However, images captured at $+40^\circ$ and $+50^\circ$ angles did not show improvement cumulatively. A subject-wise analysis of the images revealed that the recognition performance is also improved for subjects with specific characteristics. The common characteristics for iris images that were improved were pupil dilation, eyelid occlusion with high contrast iris images, and limbus occlusion of iris pixels. The common characteristics for the iris images that were not improved were the nose or the eyelids occluding most of the pixels, low contrast iris texture, uneven limbus occlusion, and the distortion caused by the 3D structure of the iris images. Images captured at lower angles have very little room for improvement.

This study suggests that tweaking the loss functions on a Pix2Pix GAN-like architecture to learn the contrast better can yield better results for off-angle iris recognition than perspective transformations like affine, which could yield good performance for only up to images captured at a gaze angle of 30° . Furthermore, the problem gets harder with more extreme angles, giving more insights into the causes of the comparatively poor recognition performances at those angles. Capturing iris images at those angles could provide more information. Additionally, using the Iris Codes instead of the normalized iris images could be an alternate way to look at the problem.

Acknowledgments

This project was made possible by support from Secure and Trustworthy Cyberspace (SaTC) program of The National Science Foundation (NSF) under grant award CNS-1909276 and CNS-2100483.

References

- [1] Y. Taigman et al. Deepface: Closing the gap to human-level performance in face verification. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1701-1708, 2014.
- [2] C. Wilson. Vein pattern recognition: a privacy-enhancing biometric. *CRC press*, 2010.
- [3] S. Sanderson and J. Erbetta. Authentication for secure environments based on iris scanning technology, *IEEE Colloquium on Visual Biometrics*, 8/1 – 8/7, 2000.
- [4] J. Daugman. How iris recognition works, *IEEE Trans. Circuits and Syst. Video Technol.*, 14(1): 21–30, 2004.
- [5] M. Karakaya. A study of how gaze angle affects the performance of iris recognition, *Pattern Recognition Letters*, 2015.
- [6] J. Daugman. New Methods in iris recognition, *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 37(5): 1167-1175, 2007.
- [7] J. Zuo, N. D. Kalka, and N. A. Schmid. A Robust IRIS Segmentation Procedure for Unconstrained Subject Presentation, *Biometrics Symposium: Special Session on Research at the Biometric Consortium Conference*, 1-6., 2006.
- [8] X. Li, L. Wang, Z. Sun, and T. Tan. A feature-level solution to off angle iris recognition, *Proceedings of International Conference on Biometrics*, 1–6, 2013.
- [9] H. J. Santos-Villalobos, et al. ORNL biometric eye model for iris recognition, *IEEE Fifth International Conference on Biometrics: Theory Applications and Systems (BTAS)*, 176-182, 2012.
- [10] M. Karakaya, D. Barstow, H. Santos-Villalobos, and J. Thompson. Limbus impact on off angle iris degradation, *International Conference on Biometrics*, 1-6, 2013.
- [11] G. N. Cerme and M. Karakaya. Effects of 3D Iris Texture on off angle Iris Recognition, *Proceedings of IEEE 23rd Conference of Signal Processing and Communication Applications (SIU 2015)*, 2015.
- [12] Z. Zhang, Y. Song and H. Qi. Age Progression/Regression by Conditional Adversarial Autoencoder. *Computer Vision and Pattern Recognition(CVPR)*, 4352-4360, 2017.
- [13] I. Goodfellow et al. Generative adversarial nets, *Advances in neural information processing systems*, 2672–2680, 2014.
- [14] F. Taherkhani, V. Talreja, J. Dawson, M. Valenti, and N. M. Nasrabadi. Profile to frontal face recognition in the wild using coupled conditional generative adversarial network, *IET Biometrics*, 11, 2022.
- [15] I. Philip. Image-to-Image Translation with Conditional Adversarial Networks, *Computer Vision, and Pattern Recognition*, 2018.
- [16] K. Diab and M. Karakaya. CNN-Based Gaze Estimation for off angle Iris Recognition, *SoutheastCon*, 736-742, 2022.
- [17] M. Moshirfar, R.N. Hoggan, and V. Muthappan. Angle Kappa and its importance in refractive surgery, *Oman Journal of Ophthalmology*, 6(3):151-8, 2013.
- [18] E. Ehrlich and M. Karakaya. Comparison of elliptical unwrapping and frontal projection for off angle iris normalization, *Defense + Commercial Sensing*, 2021.
- [19] Z. Wang et al. Image quality assessment: from error visibility to structural similarity, *IEEE Transactions on Image Processing*, 13(4):600-612, 2004.
- [20] C. Frogner et al. Learning with a Wasserstein loss, *Advances in neural information processing systems*, 28, 2015.