

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2017.DOI

Iris Deidentification with High Visual Realism for Privacy Protection on Websites and Social Networks

MAURO BARNI¹, (Fellow, IEEE), RUGGERO DONIDA LABATI², (Member, IEEE), ANGELO GENOVESE², (Member, IEEE), VINCENZO PIURI², (Fellow, IEEE), and FABIO SCOTTI², (Senior Member, IEEE)

¹Department of Information Engineering and Mathematics, Università degli Studi di Siena, 53100 Siena, Italy (e-mail: barni@dii.unisi.it)

²Department of Computer Science, Università degli Studi di Milano, 20133 Milano, Italy (e-mail: {ruggiero.donida, angelo.genovese, vincenzo.piuri, fabio.scotti}@unimi.it)

Corresponding author: Angelo Genovese (e-mail: angelo.genovese@unimi.it).

This work was supported in part by the EC within the H2020 program under project MOSAICrOWN by the Italian Ministry of Research within the PRIN program under project HOPE and by the Università degli Studi di Milano under the project “Artificial Intelligence for Image Analysis in Forensic Anthropology and Odontology”. We thank NVIDIA Corporation for donating a GPU under the project “Deep learning and CUDA for advanced and less-constrained biometric systems”. Due to copyright reasons, all the figures in the paper include only those released under the Creative Commons license.

ABSTRACT The very high recognition accuracy of iris-based biometric systems and the increasing distribution of high-resolution personal images on websites and social media are creating privacy risks that users and the biometric community have not yet addressed properly. Biometric information contained in the iris region can be used to automatically recognize individuals even after several years, potentially enabling pervasive identification, recognition, and tracking of individuals without explicit consent. To address this issue, this paper presents two main contributions. First, we demonstrate, through practical examples, that the risk associated with iris-based identification by means of images collected from public websites and social media is real. Second, we propose an innovative method based on generative adversarial networks (GANs) that can automatically generate novel images with high visual realism, in which all the biometric information associated with an individual in the iris region has been removed and replaced. We tested the proposed method on an image dataset composed of high-resolution portrait images collected from the web. The results show that the generated deidentified images significantly reduce the privacy risks and, in most cases, are indistinguishable from real samples.

INDEX TERMS Biometrics, Deidentification, GAN, Iris, Privacy

I. INTRODUCTION

THE number of high-resolution images and videos uploaded by users on social networks and web-based applications is constantly increasing. These images present a relevant privacy risk since biometric recognition could be performed by third parties without the explicit consent of the owners [1]. In fact, the need to protect high-resolution images posted on social media from the possibility of biometric recognition was proven in recent studies [2].

Although iris recognition algorithms have traditionally been designed for ocular images acquired from cooperative users using infrared light and dedicated acquisition devices,

recent studies have reached remarkable biometric recognition accuracy even for samples acquired in the wild, with images taken at long distances from sensors and under natural light conditions [3], [4]. Furthermore, images of faces captured using cameras integrated in recent smartphones frequently represent irises with a diameter of more than 300 pixels, which exceeds the value needed to obtain a satisfactory recognition accuracy [5]. Therefore, recent iris recognition techniques introduce the possibility of performing biometric recognition by using portrait pictures uploaded on websites or social networks [6]. Fig. 1 shows an example of a failed face recognition [7] for which the iris recognition method

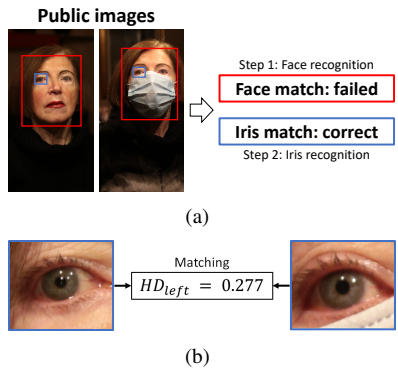


FIGURE 1. Privacy concerns in public images caused by the possibility of correctly achieving iris recognition. The figure shows an example of failed face recognition [7] using public-domain face images of the same person (a). Regardless of failure, the iris recognition method [8] obtained $HD_{left} = 0.277$ (b), which is below the threshold currently deployed in most iris recognition systems. Therefore, even when face recognition is not applicable, it is possible to be highly confident that two irises belong to the same person.

[8] obtained $HD_{left} = 0.277$, which is below the threshold currently deployed in most iris recognition systems.

Among the biometric characteristics visible in pictures uploaded on websites and social media, iris patterns represent one of the most sensitive biometric traits for several reasons: *i*) the iris is stable throughout a person’s lifetime, thus enabling individual recognition using images even when taken several years apart [9]; *ii*) the probability that two individuals will have iris traits that are recognized to pertain to the same individual is extremely low, enabling high-confidence matches even when dealing with millions of images [10]; *iii*) iris recognition can be successfully conducted in cases where face and periocular recognition algorithms fail due to the presence of thick makeup, occlusions, rotations, and unnatural expressions [11]; *iv*) the iris pattern visible in a face image could be stolen and used by ill-intentioned people to create synthetic traits usable in spoofing attacks [1]; *v*) the two iris patterns and other characteristics could be used by a multibiometric system (Fig. 2), which significantly increases the recognition capability [12] and, consequently, the associated privacy risks; and *vi*) people are wary of unauthorized uses of biometric traits traditionally acquired in a cooperative manner (e.g., via iris and fingerprint) because such traits are frequently used for governmental applications.

Fig. 2 shows the steps of the biometric process for recognizing irises in images downloaded from the web: *i*) face and eye detection; *ii*) iris segmentation; and *iii*) iris matching. The first two steps can be performed using automatic libraries or manually by a human operator to achieve higher accuracy.

Protecting the distinctive characteristics of the iris in images uploaded online is a topic that has not yet been properly addressed in the literature. To the best of our knowledge, no studies evaluating the privacy risks exist that are related to the use of iris regions extracted from online face images. Furthermore, only three works on the topic of protecting iris samples exist [13]–[15]; these obfuscate or blur the iris




FIGURE 2. Outline of the methodology used to perform iris recognition using public domain face images. It is possible to use robust methods to automatically detect the face, extract the eye positions, segment the irises, and then compare them to achieve high-confidence recognition.

region of the image but do not provide deidentified visually realistic images suitable for posting on the web in place of the original image.

In this study, we address the issue of protecting the distinctive characteristics of the iris in images uploaded online while maintaining a satisfactory level of realism and visual quality. Users of social networks and websites, in fact, usually desire that their images be recognizable by other people, but one would think it should be highly desirable to be able to post such images with automatically removed distinctive biometric information from the iris regions to reduce privacy risks. The new iris regions should be visually plausible, preserve the original eye color and be of sufficient quality to satisfy users. In fact, people would prefer to preserve the most distinctive human characteristics (e.g., eye color) while discriminative biometric details that are not usually noticed by human observers (e.g., the texture of the iris pattern) are removed. Therefore, the proposed method does not modify facial characteristics and it preserves the eye color nuances of the original image.

With the above ideas in mind, the contributions of this paper are twofold. First, we analyze the privacy risks related to the visibility of the iris pattern in high-resolution images. Second, we propose a novel method to deidentify the iris region in face images by replacing the iris region with a synthetic pattern¹. The deidentification process consists of removing or replacing personal identifiers with surrogate personal identifiers, with the aim of preventing the disclosure and use of data for purposes unrelated to the one for which the information was originally obtained [16]. In contrast to prior works, our method is designed to obtain visually plausible iris textures with high resolution and to leave other aspects of the original image unaltered. The proposed method replaces irises with synthetic biometric characteristics computed randomly using fractals or images produced by a generative adversarial network (GAN). The obtained images are highly realistic and visually plausible, preserving the visual aspect of the reflections, which is frequently used to discriminate between real and synthetic images [1], and

¹Source code available at <http://iebil.di.unimi.it/irisGan/irisGan.html>

preserving the color nuance of the original iris. The images do not include any biometric information originating from the original iris.

We evaluated the deidentification capability of the proposed method as well as the visual realism of the obtained samples using face and iris images collected from websites and social media. The performed tests are based on the analysis of the performance of state-of-the-art iris recognition methods and on the answers of volunteers to a questionnaire. Positive results were obtained for every aspect of the proposed iris deidentification method.

The remainder of this paper is organized as follows. Section II discusses the related works. Section III describes the proposed method to estimate the privacy risk associated with iris recognition using images downloaded from the web. Section IV illustrates the method proposed to perform iris region deidentification. Section V presents the experimental evaluation, and Section VI concludes the work.

II. RELATED WORKS

Most of the studies in the literature on the deidentification of biometric characteristics focus on the face trait. The earliest studies on face deidentification used simple strategies based on common image processing operations to modify the face region [17], such as “black box”, “pixelation”, and “blurring”. However, these methods remove information unrelated to the individual’s identity and degrade the overall realism of the image. More recent techniques try to overcome these limitations and provide formal guarantees regarding the anonymity of deidentified data by using the concept of κ -anonymity [18], [19], for example, the κ -same approach [20]. These methods preserve some of the original distinctive characteristics of the biometric trait to try to obtain an image as similar as possible to the original sample. In recent years, researchers have proposed face deidentification methods based on deep learning (DL) techniques, which frequently use a GAN to generate modified face images or mixtures of faces computed starting from a feature database [21], [22]. Since the distinctive characteristics of iris patterns are more complex for humans to memorize compared to other face traits, methods based on the concept of κ -anonymity are not convenient and it should be possible to use visually realistic patterns generated using pseudorandom approaches to compute synthetic iris patterns. Furthermore, the GANs used to generate face images are not directly applicable for creating synthetic iris regions due to the low image resolution and the low level of detail in the iris region.

To the best of our knowledge, only a few studies on iris deidentification techniques exist; these studies are intended only to protect the iris region. The method presented in [13] first searches for the iris region and then degrades that region using a JPEG extended range (XR) encoder. Another study [14] applied a cryptographic technique designed for JPEG 2000 images to protect iris images converted using the rubber sheet model (RSM) [23]. The method proposed in [15] removes distinctive biometric characteristics while preserv-

ing iris biological features from ocular images by using an algorithm that adds a controlled amount of Laplacian noise to blur the iris region. However, none of these methods attempt to preserve the visual realism of the deidentified irises. In [22], [24], the authors proposed methods to obtain visually realistic ocular regions. These methods were intended to be integrated into face portraiture software for inpainting closed eyes or enforcing a specific gaze direction. However, such methods can be applied only to very low-resolution images compared to images that are suitable for iris recognition. In addition, they do not generate detailed iris textures, and they replace the entire eye and eyelash region, which alters the original facial expression.

In this paper, we propose a novel method for generating synthetic iris textures that achieve visually pleasant results. The literature contains several studies involving methods to compute synthetic iris images. However, none of these methods can be directly used to create visually realistic iris textures to be embedded in input facial images. The existing methods can be grouped into algorithmic approaches [25]–[29] and methods based on DL and GAN models [30]–[33]. Table 1 presents a summary of the existing methods for generating synthetic iris textures.

III. PRIVACY RISK ESTIMATION

In this section, we describe a simple approach for estimating the privacy risks associated with distributing high-resolution facial images with visible iris regions. In our analysis, we focus on a monomodal recognition strategy based on a single iris. Since multimodal biometrics tend to achieve better recognition accuracy than do monomodal systems [12], this analysis should be considered an optimistic estimate.

For our analysis, we used I-SOCIAL-DB [6], containing 3,286 ocular images collected from websites and social media. For each ocular image, the dataset includes the corresponding iris segmentation mask and the parameters of the circles approximating the inner and outer iris boundaries. The average size of the face images is $\approx 3,000 \times 3,200$ pixels and the iris radii vary from ≈ 56 to ≈ 137 pixels. Fig. 3 shows examples of the iris images collected from websites and social media. Notably, it is not possible to obtain information about possible image enhancements performed by photographers, which can drastically reduce the accuracy of biometric recognition algorithms.

To estimate the privacy risk, we analyzed the cumulative distributions of the genuine and impostor matching scores obtained by comparing every possible pair of samples in the dataset, thus allowing the average privacy risk for the population in the database to be computed. As an example, Fig. 4 compares the results achieved by a public implementation [8] of a contrast-adjusted segmentation algorithm [34] and a well-known recognition method in the literature [35] for a subset of the Institute of Automation of the Chinese Academy of Sciences version 4 (CASIA-v4) interval dataset [36] and I-SOCIAL-DB. A comparison of Fig. 4 (a) to Fig. 4 (b) reveals some important differences between the iris images

TABLE 1. Summary of methods for generating synthetic iris textures

Ref.	Year	Type	Method	Approach
[25]	2006	Algorithmic approach	Feature agglomeration	First, this method uses a Markov random field to generate a random texture; then, it generates detailed features and embeds them in the texture.
[26]	2007	Algorithmic approach	Anatomy-based iris generation	First, this approach uses a generation process based on simulating a dense fiber structure; then, it applies image processing operators to refine the generated image.
[27]	2008	Algorithmic approach	Patch-based sampling	This method first creates a visual primitive of the iris texture using iris patch-based sampling; then, it generates pseudo-irises by introducing intraclass variations.
[28]	2010	Algorithmic approach	Multiresolution approach	This approach first decomposes the training iris image into lower-resolution components and then combines the components to generate random samples.
[29]	2013	Algorithmic approach	NOISYRIS	This approach first applies a stochastic method based on creating and grouping fibers to generate a synthetic iris; then, it applies rendering algorithms to simulate illumination effects and different nonideal conditions.
[30]	2017	DL	iDCGAN	This approach uses a deep convolutional GAN trained on iris images and a corresponding quality index to generate synthetic samples starting from a random vector.
[31]	2017	DL	S+U GAN	This model uses a GAN trained using a combination of simulated and unsupervised learning to generate highly realistic synthetic irises by starting from a synthetic iris with coarser realism.
[32]	2018	DL	Iris-GAN	It uses a GAN to generate synthetic images resembling the irises in public databases.
[33]	2019	DL	RaSGAN	This model trains a GAN composed of relativistic networks and uses quality-based metrics to improve the realism of synthesized irises.

Notes. DL = Deep Learning.



FIGURE 3. Examples of iris images collected from websites and social media. The iris radii are sufficiently large to perform iris recognition using state-of-the-art algorithms.

acquired using the biometric scanners of CASIA-IrisV4 and those downloaded from public websites. Fig. 4 (b) shows that while the privacy risks related to images collected from the web are less than those of databases of iris images collected using biometric scanners, the risks are still relevant. Notably, by setting a threshold $HD = 0.365$, we obtained a correct genuine identity comparison percentage of 25.86% at a false matching rate (FMR) of $\approx 10^{-4}$.

IV. PROPOSED IRIS DEIDENTIFICATION METHOD

Our proposed iris deidentification approach can work in different configurations by the use of heterogeneous algorithms for generating synthetic iris textures. Our approach extracts the iris region from the high-resolution face image, creates a synthetic iris texture, and finally inserts the synthetic iris into the original face image. During the generation of the synthetic iris, no biometric information from the original iris texture is used; we extract only appearance-based statistics on eye color nuances, which are not relevant for most of the state-of-the-art iris recognition technologies (the mean and standard deviation of the intensity values of the color channels). We use the extracted statistics to generate visually plausible synthetic iris textures that resemble those in the

original images.

Our approach can be divided into the following steps: *A)* eye region extraction and iris segmentation; *B)* computation of the RSM; *C)* computation of the synthetic texture; *D)* color domain adaptation; *E)* conversion to cartesian coordinates and blending. We repeat this procedure for both the left and right irises. Fig. 5 shows the outline of the proposed synthetic iris generation method.

A. EYE REGION EXTRACTION AND IRIS SEGMENTATION

This step first processes the face image to extract the eye region and then segments the iris to compute a binary segmentation mask. We considered two variants for completing this step. In the first variant, both the eye region extraction and the iris segmentation are performed manually by an expert user. In the second variant, we use state-of-the-art automatic algorithms to perform both tasks. For both variants, we adopt pixelwise segmentation.

In the remainder of this section, we describe the variant using automatic algorithms. To extract the eye region, we use the method described in [37], based on a convolutional neural network (CNN). We chose this method because it represents the state-of-the-art segmentation algorithm for high-resolution face images. The network automatically estimates the coordinates of the image corresponding to the centers of the eyes $(x_{left}, y_{left}), (x_{right}, y_{right})$. The ocular regions I_{left}, I_{right} are obtained by cropping the face image around each eye center using squared regions whose sides are equal to $1/3$ of the Euclidean distance between (x_{left}, y_{left}) and (x_{right}, y_{right}) .

To segment the iris from the ocular images I_{left}, I_{right} , we use an algorithm based on CNNs [38] because of its high accuracy in segmenting iris images acquired in visible light

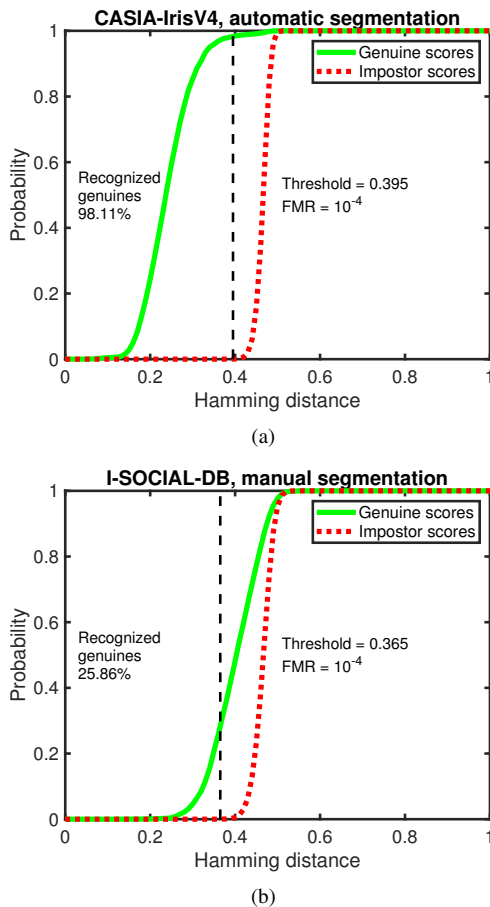


FIGURE 4. Privacy risk analysis for a dataset of iris images collected (a) using a biometric scanner (b) from face images downloaded from websites and social media and using segmentation masks created by a human expert. In our experiments, the matching scores are computed as the Hamming distance between two templates. The resulting graph reveals that although privacy risks are not comparable for samples acquired using a biometric scanner and samples extracted from face images downloaded from websites and social media, there are concrete privacy risks even for downloaded samples. As an example, by setting a threshold of $HD = 0.365$, we obtained 25.86% of correct genuine identity comparisons at a FMR of $\approx 10^{-4}$.

conditions.

A segmentation algorithm computes the binary segmentation masks B_{left} , B_{right} for the left and right irises, respectively. For each iris, the algorithm also computes the parameters describing two circles approximating the inner and outer iris boundaries. The inner boundary is described by the center coordinates (x_i, y_i) and the radius r_i , while the outer iris boundary is described by the parameters (x_o, y_o) and r_o .

In this paper, we consider the methods [37], [38] only as an example; it is possible to use any suitable algorithm from the literature to extract the eye region and segment the iris.

In the remainder of the section, because our deidentification method processes both the left and the right eyes in the same manner, we describe the remaining steps in the processing chain by referring to a single iris.

B. COMPUTATION OF THE RUBBER SHEET MODEL

This step aims at creating a normalized representation of the iris region invariant to image resolution, pupil dilation, and nonconcentricity of the pupil with respect to the iris. For normalization, we adopt the RSM, which is one of the simplest and most commonly used techniques in the literature.

The normalization algorithm converts the iris region of the ocular image I , described by the segmentation mask B , into a rectangular polar image P representing the pixels included between two circles approximating the inner and outer iris boundaries. Specifically, the cartesian coordinates (x, y) of every pixel of the iris region of I are converted to a double dimensionless nonconcentric polar coordinate system (ρ, θ) , where ρ belongs to the unit interval $[0, 1]$ and θ is an angle in the range $[0, 2\pi]$. The image I_R is obtained by quantizing ρ and θ into n_θ and m_ρ values, respectively. We set the parameters n_θ and m_ρ empirically.

C. COMPUTATION OF THE SYNTHETIC TEXTURE

This task creates a realistic synthetic iris texture represented as a rectangular image T with a fixed size of n_θ by m_ρ pixels. The goal is to obtain a synthetic texture as similar as possible to those obtained by normalizing a real iris region using the RSM algorithm but lacking any biometric information related to the original iris.

The main advantages of simulating the iris texture in the normalized domain with respect to performing the same computation in cartesian image space are as follows: *i*) the RSM is invariant to the image resolution; *ii*) the RSM is robust to pupil dilation; and *iii*) the RSM does not require that the pupil be concentric with respect to the iris. These advantages help in embedding the iris texture into the face image and thereby creating visually realistic deidentified images.

To compute the synthetic texture, we propose two techniques, optimized in terms of resources and visual realism. The first technique is based on a simple and fast fractal algorithm, while the second technique is based on a GAN and can achieve more visually realistic results. In our work, we consider both techniques and use the fractal algorithm as a baseline against which to compare the GAN-based technique. In fact, GANs can produce more visually realistic images taking advantage of large datasets in their training, while the fractal algorithm requires only a single random number for its initialization.

1) Fractal generation of the synthetic iris texture

To rapidly compute a pseudorandom representation of the iris texture, we approximate the texture as a plasma fractal and compute it using the diamond-square algorithm. This algorithm is frequently used to compute height maps for computer graphics [39].

The algorithm consists of n_f iterations, during which it divides an image T into local square regions and sets the center point of each region to the average of the four corner points plus a random value. In each iteration, the algorithm reduces the magnitude of the randomness. The algorithm generates

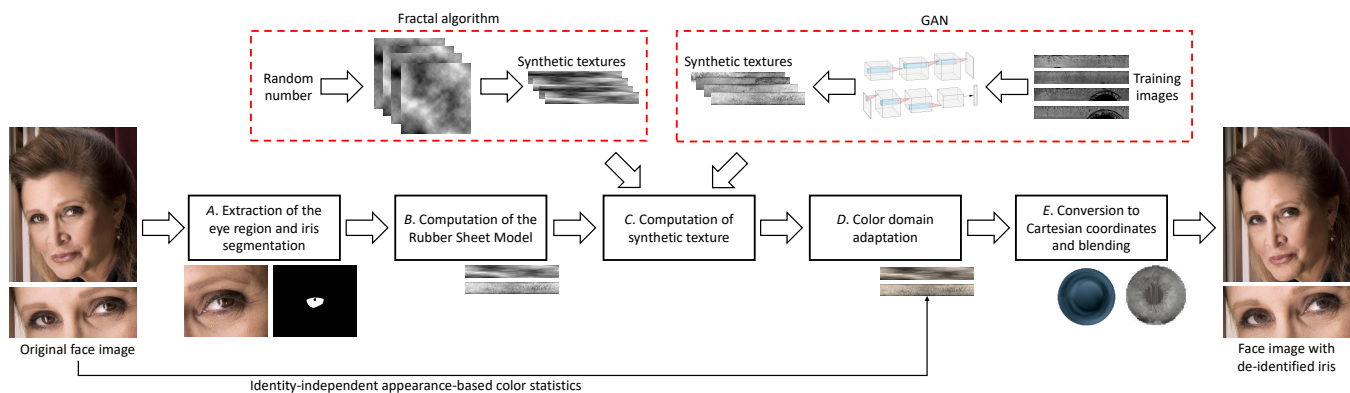


FIGURE 5. Outline of the proposed iris deidentification method. Iris deidentification is applied separately to both the left and the right irises in each image. As a result, we obtain images with visually plausible synthetic iris textures that resemble the original images. The outline shows the results of two alternatives for generating the synthetic pattern: a fractal algorithm and a GAN. We consider the fractal algorithm as a baseline against which to compare the GAN-based technique.




FIGURE 6. Examples of plasma fractals used to compute representations of the iris texture: (a,b) squared plasma fractals; (b,c) corresponding images after resizing the fractals to the proportions of the rubber sheet model (RSM). The images exhibit a pseudorandom pattern.

a squared plasma fractal T with a size of $2^{2+n_f} \times 2^{2+n_f}$. Then, region T is resized to a rectangular image with size n_θ by m_ρ , with n_θ, m_ρ selected to imitate the proportions of the RSM. Fig. 6 shows some examples of plasma fractals used to compute the pseudorandom representations of the iris texture and the corresponding image after resizing.

2) GAN-generated iris textures

Among computational intelligence approaches, owing to their advantage of being able to automatically learn data representations, techniques that use DL are being increasingly used in a wide variety of pattern recognition fields. In particular, DL methods based on GANs are emerging as state-of-the-art techniques for generating highly realistic synthetic images. They work by combining two machine learning models: a generator G , which generates synthetic data, and a discriminator D , which takes as input the data generated by G and classifies it as real or synthetic. Learning algorithms for GANs are based on adversarial training of G and D , which compete against each other to reach an equilibrium point [40]. When this equilibrium point is reached

(or training is terminated), the generator has been trained to create synthetic images from a vector of random numbers.

Fig. 7 shows the architecture of the GAN used in our work. Specifically, we use a deep convolutional GAN (DCGAN) in which G and D are implemented as CNNs. DCGANs have successfully been used to generate visually realistic images in different application scenarios [41], [42]. The DCGAN is trained to generate synthetic iris textures using a training set of iris RSMs².

The DCGAN consists of several layer types, which are described as follows:

- *Linear layer*: applies a linear transformation to the input data, according to the equation: $y = xA^T + b$, where x is the input data, A is the transformation matrix and b is the bias.
- *Hyperbolic tangent*: applies the hyperbolic tangent function $y = \tanh(x)$ to the input data.
- *Sigmoid*: applies the sigmoid function $y = \frac{1}{1 + e^{-x}}$ to the input data.
- *Convolutional layer*: computes its output by applying a convolution of the input data using a bank of two-dimensional filters. For each coordinate (i, j) , the output is computed according to the equation: $y(i, j) = b + \sum_{m=1}^H \sum_{n=1}^W f(m, n) \times x(i-m, j-n)$, where M and N are the horizontal and vertical dimensions of the filter f , respectively, and b is the bias. In this work, we set $M = N = 3$, and the padding = 1.
- *Leaky Rectified Linear Unit (LeakyReLU) layer*: applies the function $y = \max(0, x) + m \times \min(0, x)$.
- *Dropout layer*: randomly sets the input data to 0, with a probability of p_{drop} .

The architecture of the generator G CNN is shown in Fig. 8a. G is composed of linear, resizing, convolutional, and ReLU layers, arranged as shown in Table 2. We apply batch normalization after layers L_2, L_3, L_6 using the function described in [43]. The DCGAN performs the generation

²The source code is available at <http://iebil.di.unimi.it/irisGan/irisGan.html>

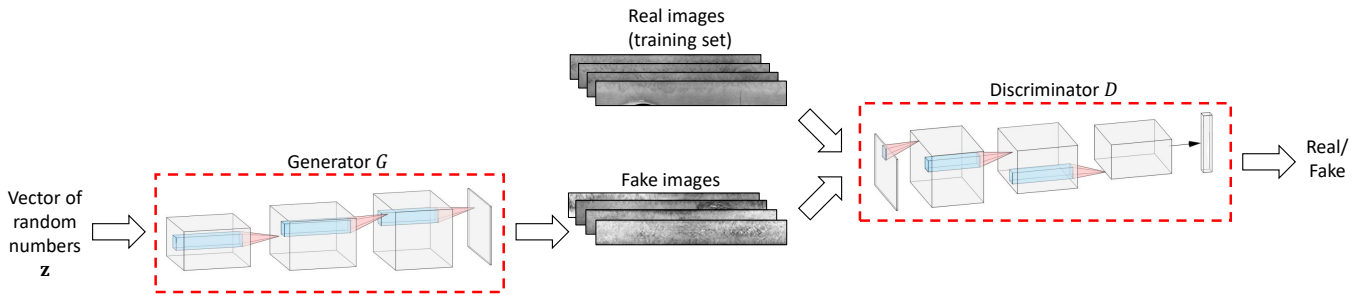


FIGURE 7. Architecture of the deep convolutional GAN (DCGAN) used in this work. To train the generator G and the discriminator D , we use a database of RSMs of irises. The generation process is performed by applying G on a vector \mathbf{z} composed by random numbers in the range $[0, 1]$, extracted following a normal distribution.

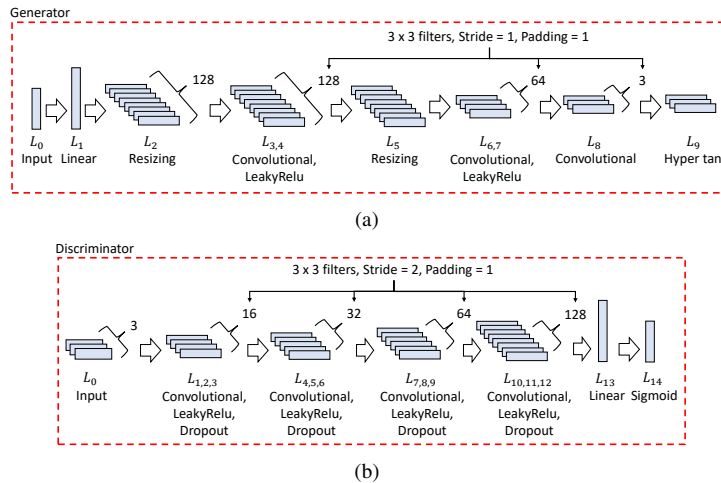


FIGURE 8. Architecture of the generator G and discriminator D CNNs used in the DCGAN: (a) generator G ; (b) discriminator D .

TABLE 2. Summary of the architecture of the CNN of generator G

Layer N.	Layer type	Feature size	Number of filters	Stride	Batch Norm.
L_0	Input	$[z = 100]$	-	-	N
L_1	Linear	$[n_\theta/4 \cdot m_\rho/4 \cdot 128]$	-	-	N
L_2	Resizing	$[n_\theta/2 \times m_\rho/2 \times 128]$	-	-	Y
L_3	Conv.	$[n_\theta/2 \times m_\rho/2 \times 128]$	128	1	Y
L_4	LeakyReLU	$[n_\theta/2 \times m_\rho/2 \times 128]$	-	-	N
L_5	Resizing	$[n_\theta \times m_\rho \times 128]$	-	-	N
L_6	Conv.	$[n_\theta \times m_\rho \times 64]$	64	1	Y
L_7	LeakyReLU	$[n_\theta \times m_\rho \times 64]$	-	-	N
L_8	Conv.	$[n_\theta \times m_\rho \times 3]$	3	1	N
L_9	Hyper. tan.	$[n_\theta \times m_\rho \times 3]$	-	-	N

Notes: Batch norm. = Batch normalization (Y = Yes; N = No);
Conv. = Convolutional; Hyper. tan. = Hyperbolic tangent.

process by applying G on a vector \mathbf{z} , with size $|z| = 100$, composed by random numbers in the range $[0, 1]$, extracted following a normal distribution [42]. As a result, G outputs an image with size $n_\theta \times m_\rho \times 3$.

The architecture of the discriminator D CNN is shown in Fig. 8b. D is composed of linear, convolutional, and ReLU layers, arranged as shown in Table 3.

We train the DCGAN with the adaptive moment estimation (Adam) algorithm, which optimizes the binary cross-entropy

TABLE 3. Summary of the architecture of the CNN of discriminator D

Layer N.	Layer type	Feature size	Number of filters	Stride	Batch Norm.
L_0	Input	$[n_\theta \times m_\rho \times 3]$	-	-	N
L_1	Conv.	$[n_\theta \times m_\rho \times 16]$	16	2	N
L_2	LeakyReLU	$[n_\theta/2 \times m_\rho/2 \times 16]$	-	-	N
L_3	Dropout	$[n_\theta/2 \times m_\rho/2 \times 16]$	-	-	N
L_4	Conv.	$[n_\theta/4 \times m_\rho/4 \times 32]$	32	2	N
L_5	LeakyReLU	$[n_\theta/4 \times m_\rho/4 \times 32]$	-	-	N
L_6	Dropout	$[n_\theta/4 \times m_\rho/4 \times 32]$	-	-	N
L_7	Conv.	$[n_\theta/4 \times m_\rho/4 \times 64]$	64	2	N
L_8	LeakyReLU	$[n_\theta/8 \times m_\rho/8 \times 64]$	-	-	N
L_9	Dropout	$[n_\theta/8 \times m_\rho/8 \times 64]$	-	-	N
L_{10}	Conv.	$[n_\theta/8 \times m_\rho/8 \times 128]$	128	2	N
L_{11}	LeakyReLU	$[n_\theta/16 \times m_\rho/16 \times 128]$	-	-	N
L_{12}	Dropout	$[n_\theta/16 \times m_\rho/16 \times 128]$	-	-	N
L_{13}	Linear	$[1]$	-	-	N
L_{14}	Sigmoid	$[1]$	-	-	N

Notes: Batch norm. = Batch normalization (Y = Yes; N = No);
Conv. = Convolutional.

function [44]:

$$\min_G \max_D V(G, D) = \min_G \max_D \mathbb{E}_{p \approx p_{data}} [\log D(x)] + \mathbb{E}_{z \approx p_z} [\log(1 - D(G(z)))] \quad (1)$$

After training the DCGAN, we generate an iris texture by




FIGURE 9. Examples of synthetic iris textures computed using the proposed DCGAN. The images exhibit high visual realism and resemble RSMs computed from real irises.

supplying a vector \mathbf{z} of random numbers to G . The result is a synthetic iris texture T with a size of $n_\theta \times m_\rho$. Fig. 9 shows some examples of synthetic iris textures created by the use of the proposed DCGAN.

D. COLOR DOMAIN ADAPTATION

This step aims at adapting the simulated texture T in the color domain to obtain an image C with color characteristics similar to those of the irises included in I . To perform this task, we also consider identity-independent appearance-based color statistics extracted from the iris image I but without including any biometric information originating from the real iris.

To perform the color domain adaptation, we first reduce the possible presence of visual incoherence at the extremes of T due to the transition of θ from 0 to 2π . To meet this goal, we apply a Gaussian filter to T using a kernel with an empirically estimated size of $s_k \times s_k$ pixels and with a standard deviation of σ_g . The filter is applied by considering the image T as continuous in the convolution operation, thus obtaining the smoothed image T' .

We then adapt the intensity range of T' for each color channel of the iris region. Starting from I and a binary mask B representing the segmented iris, we compute a vector of intensity values V_c , where $c \in \{R, G, B\}$, for each of the color channels of the red, green and blue (RGB) space. We compute each channel of the color texture image C as follows:

$$\begin{aligned} A &= T' - \text{mean}(T'), \\ C_c &= A \times [\text{std}(V_c) \times w_1 + \text{mean}(V_c) \times w_2] \\ \forall c &\in \{R, G, B\}, \end{aligned} \quad (2)$$

where w_1 and w_2 are two empirically estimated constants.

E. CONVERSION TO CARTESIAN COORDINATES AND BLENDING

The goal of this step is to create a deidentified image by creating the image S , representing the synthetic iris texture C in the cartesian coordinates of the eye image I , and then blending S into the real image, obtaining a face image with a deidentified iris \hat{I} .

First, we compute an image S that represents the synthetic iris texture in cartesian coordinates by considering the parameters that describe the inner and outer iris boundaries, which are computed using the method described in Section IV-A.



(a) (b) (c)

FIGURE 10. Example of the results of the proposed iris deidentification process: (a) original iris; (b) deidentified iris obtained using the fractal algorithm; (c) deidentified iris obtained via the GAN. The deidentified images (b,c) are highly visually realistic.

Specifically, we compute the cartesian coordinates X, Y as follows:

$$X(i) = x_i + \{[R(i) \times \cos \Theta(i)] \times [(r_o - r_i) / m_\rho]\}, \quad (3)$$

$$Y(i) = y_i + \{[R(i) \times \sin \Theta(i)] \times [(r_o - r_i) / m_\rho]\}, \quad (4)$$

where R and Θ are two matrices representing the polar coordinates of the texture T . The matrices X and Y are then used to compute the cartesian image C by performing a Poisson image interpolation.

Finally, we obtain the deidentified image \hat{I} by substituting the pixels of I in the region of interest defined by the binary mask B with the corresponding pixels from S by the use of a Poisson-based blending approach [45]. To achieve a more natural transition between the synthetic texture and the original image, during the blending process, we superimpose the external iris ring I_r of the original iris image I on S . We compute I_r by considering only the regions of I whose distance $\frac{9}{10}r_o \leq r_b \leq r_o$, where r_o is the radius of the external iris boundary. We used this simplification to obtain a more natural transition between iris and sclera and because it has been demonstrated that iris regions close to the border carry limited biometric information [46].

Fig. 10 shows an example of the results of the proposed iris deidentification process.

V. EXPERIMENTAL RESULTS

This section describes the procedures used to train the GAN, summarizes the method parameters, presents an analysis of the deidentification performance of the proposed approach, and presents a qualitative analysis of the generated images.

A. GAN TRAINING PROCEDURE

To train the GAN, we created a dataset of iris images obtained by applying a data augmentation procedure to sets of iris images acquired using traditional iris scanners. For this purpose, we considered portions of the public iris databases CASIA-IrisV4 [36] and the Indian Institute of Technology Delhi (IITD)-IrisV1 [47], captured with near-infrared light. We used only the images for which the corresponding segmentation masks are publicly available [48]. Specifically, we used 2,639 images from the CASIA-Iris-Interval subset (captured from 249 individuals and with a size of 320×280 pixels) and 2,240 images from the IIT Delhi Iris Database,

version 1.0 (captured from 224 individuals and with a size of 320×240 pixels).

The proposed procedure for training the GAN is based on the following steps.

- 1) Database merging to obtain a single database with 4,879 samples.
- 2) Data augmentation to inpaint occlusions in the normalized textures by replicating the iris pattern over the occluded areas. This step is necessary to teach the GAN to generate synthetic iris textures with no occlusions. The proposed inpainting procedure is based on the following steps.
 - a) *Selection*. We selected only the RSMs for which the percentage of occlusions in the corresponding mask is $\leq 30\%$.
 - b) *Extraction*. For each RSM, we extracted the longest portion P of images with no occlusions. The sizes of this portion are $Y_s = m_\rho$ and $X_s = x_{s,end} - x_{s,start}$, where m_ρ is the size of the RSM along the y -axis, computed via the weighted adaptive Hough and ellipsopolar transform (WAHET) algorithm, and $x_{s,start}, x_{s,end}$ are coordinates along the x -axis, computed as follows:
$$(x_{s,start}, x_{s,end}) = \underset{x_1, x_2}{\operatorname{argmax}} \sum_{x_i=x_1}^{x_2} \sum_{y_i=1}^{Y_s} B(x_i, y_i), \quad (5)$$
where B is the segmentation mask corresponding to the RSM, in which the occluded areas are set to 0.
 - c) *Replication*. The extracted portion of P was replicated along the x -axis on the areas of the RSM $\leq x_{s,start}$ and $\geq x_{s,end}$. For each replication of P , the image was mirrored to ensure the continuity of the iris pattern.
- 3) The data augmentation procedure is performed to increase the dimensionality of the database by performing horizontal and vertical flipping operations for each image along the x - and y -axes, respectively. As a result, we obtained a training set with $\approx 14,000$ images.
- 4) Training of DCGAN was implemented with the training set described above for $n_e = 200$ epochs, with a batch size of $s_b = 60$, a learning rate of $lr = 0.0002$, and exponential decay rates for the first and second gradient moment estimates of $b_1 = 0.5$ and $b_2 = 0.999$, respectively. The size of the random number vectors used as input to the network is $|z| = 100$. The number of trainable parameters is 26,699,137 for G and 113,985 for D .

B. PARAMETER TUNING

During RSM computation, we set the values of n_θ and m_ρ to $n_\theta = 512$ and $m_\rho = 64$. These values resulted in RSMs with dimensions similar to those used by the majority of iris recognition methods in the literature.

When applying the fractal algorithm to create synthetic textures, we adopted $n_f = 7$ iterations; this value resulted in a good compromise between visual realism and computational complexity.

During the color domain adaptation step, we used a Gaussian filter with a kernel size of $s_k \times s_k$ pixels and a standard deviation of σ_g , where $s_k = 5$ and $\sigma_g = 4$. We chose these values to smooth the representation without reducing the visual realism. In addition, we adopted $w_1 = 5$ and $w_2 = 2$ to obtain visually realistic synthetic textures with an average color intensity similar to those of the original irises.

C. DEIDENTIFICATION CAPABILITY

In this section, we evaluate the ability of the proposed method to generate deidentified irises. We applied the proposed method to the following 2 datasets of deidentified face images:

- *DB-DeIdent-Face_{fractal}*: database of 1,643 deidentified face images, in which the irises were generated using synthetic textures computed using the fractal method described in Section IV-C1.
- *DB-DeIdent-Face_{GAN}*: database of 1,643 deidentified face images, in which the irises were generated using synthetic textures computed using the GAN described in Section IV-C2 and trained using the procedure described in V-A.

We then extracted the iris regions from DB-DeIdent-Face_{fractal} and DB-DeIdent-Face_{GAN} using the coordinates estimated by a human operator for F-SOCIAL-DB. In this way, we obtained two datasets of ocular images, called DB-DeIdent-Iris_{fractal} and DB-DeIdent-Iris_{GAN}.

To analyze the deidentification capability of the proposed method, we evaluated the accuracy of different biometric recognition schemes for real images and deidentified images, analyzed the matching scores obtained by matching real iris images and deidentified images, and evaluated the capability of the proposed GAN to generate random textures.

- 1) Effect of the proposed deidentification method on the accuracy of biometric systems

The identity verification process is composed of a segmentation task and a recognition scheme that includes specific feature extraction and matching methods. To test the deidentification capability of the proposed method, we compared the identity verification accuracy achieved by different biometric recognition schemes for I-SOCIAL-DB and for the deidentified images of DB-DeIdent-Iris_{GAN} and DB-DeIdent-Iris_{fractal}. We considered the results achieved using the manually segmented masks provided by I-SOCIAL-DB and those obtained by automatically segmenting the iris images using a deep neural network (region-based CNN (R-CNN)) [38] in conjunction with a technique for estimating the limits of RSMs (cnn2rubber) [49]. We selected this segmentation algorithm since it achieved the best results in our tests (more details are reported in Section V-D).

The considered biometric recognition schemes are based on heterogeneous features, handcrafted as well as learned by using deep neural networks. In particular, we evaluated the accuracy of a neural network with a unified deep learning architecture (UNINET) [50], a method based on machine learning and binary statistical image features (BSIF) [51], and the following recognition methods implemented in the University of Salzburg Iris Toolkit (USIT) version 3.0 [8]: log Gabor (LG) [52], complex Gabor (CG) [23], local intensity variations (CR) [53], cumulative sums of grayscale blocks (KO) [54], and quadratic spline wavelet (QSW) [55]. Each test involved 3,286 iris images, including 11,092 genuine comparisons and 10,783,418 impostor comparisons. Table 4 summarizes the achieved results in terms of equal error rate (EER) [56]. Fig. 11 shows the receiver operating characteristic (ROC) curves obtained by the best performing recognition schemes (BSIF and LG). The results are compared with the ROC curve obtained from a vector of random numbers of size equal to the number of identity comparisons performed for I-SOCIAL-DB, DB-DeIdent-Iris_{GAN}, and DB-DeIdent-Iris_{fractal}.

Table 4 and Fig. 11 show that, using manually segmented masks, all the considered biometric recognition schemes achieved EER close to 50% for both DB-DeIdent-Iris_{GAN} and DB-DeIdent-Iris_{fractal}. Notably, an EER equal to 50% suggests that the distributions of the genuine and impostor scores are not substantially different, thus implying that the distinctive biometric information has been completely removed from the original samples. Furthermore, the ROC curves obtained for DB-DeIdent-Iris_{GAN} and DB-DeIdent-Iris_{fractal} by using manually segmented masks are similar to the ROC curve obtained from randomly generated numbers. This result proves that the proposed deidentification method is effective, removing distinctive features from the iris samples. Using automatic segmentation algorithms, the EER is slightly inferior because the things computed a limited number of distinctive information in the incorrectly segmented regions. Nevertheless, the achieved result is satisfactory for practical applications since all the considered biometric recognition schemes achieved EERs higher than 41%.

We also evaluated the separation between the genuine and impostor scores for I-SOCIAL-DB, DB-DeIdent-Iris_{fractal}, and DB-DeIdent-Iris_{GAN}. For this analysis, we used the segmentation masks provided by I-SOCIAL-DB and the recognition schema LG; this study can be considered a reference point in the literature on iris recognition systems. The more widely the impostor and genuine distributions are separated, the higher the privacy risk is. Fig. 12 shows the results. The deidentified ocular images (in both configurations) do not present distinctive information in the iris region; thus, they obtain a genuine score distribution comparable to the impostor score distribution.

As a further test, we evaluated the matching scores obtained by comparing the original ocular images and the deidentified images. We used the recognition schema LG.

Specifically, we performed 3,286 identity comparisons: one for each image in I-SOCIAL-DB. We performed this comparison using manually segmented masks. For DB-DeIdent-Iris_{fractal}, we obtained a mean score of 0.493 with a standard deviation of 0.030. For DB-DeIdent-Iris_{GAN}, we obtained a mean score of 0.490 with a standard deviation of 0.027. A comparison of these results with the distributions shown in Fig. 12 reveals that the deidentified images do not present sufficient distinctive information for comparisons to the original samples using the considered biometric recognition approach.

2) Capability of the proposed GAN to generate random textures

We evaluated the capability of the proposed GAN to generate textures that present no common biometric information among them. We used the recognition schema LG. Observing the genuine distribution in Fig. 12 (c) shows that the samples computed for each individual by the employed biometric recognition method are sufficiently different to the extent that they appear to belong to different individuals. Furthermore, a visual inspection confirms that the deidentified images generated for the same individual present relevant iris texture differences. As an example, Fig. 13 shows a real ocular image and two different deidentified images created by starting from the same real ocular image.

Furthermore, we analyzed the ability of the GAN to generate samples different from those used to train the network. We compared the RSM obtained from each sample of the training set with 1,000 RSMs generated by the proposed GAN by using the recognition schema LG. Fig. 14 contains a plot of the distribution of the obtained matching scores, showing that the RSMs created by the GAN are substantially different from those used for training the network. In fact, the shape of the matching score distribution is similar to the shape of the impostor distributions obtained using the same algorithm, as shown in Fig. 12 (the mean of the matching scores is 0.480, with a standard deviation of 0.026). When the employed recognition method is used, a matching score of 0.480 is usually obtained for impostor identity comparisons performed for samples with substantial differences. These results demonstrate the ability of the GAN to create images different from those in the training set. The obtained results also prove that the proposed deidentification method guarantees robustness to reidentification attacks even in cases in which the samples to be deidentified pertain to the training set because the textures generated by the GAN do not present distinctive biometric characteristics in common with the samples in the training set.

D. REALISM OF DEIDENTIFIED IMAGES

To analyze the realism of the deidentified images obtained by the proposed method, we performed a visual analysis, evaluated the results achieved by segmentation algorithms based on heterogeneous features (edge-based as well as texture-based features), analyzed the results of questionnaires, and

TABLE 4. Identity verification accuracy of different biometric recognition schemes for real and deidentified images.

Iris Recognition Library	I-SOCIAL-DB		DB-DeIdent-Iris _{GAN}		DB-DeIdent-Iris _{fractal}	
	Manual Segmentation EER (%)	R-CNN + cnn2rubber EER (%)	Manual Segmentation EER (%)	R-CNN + cnn2rubber EER (%)	Manual Segmentation EER (%)	R-CNN + cnn2rubber EER (%)
UNINET	31.89	36.38	48.81	42.69	49.78	42.68
KO	31.22	35.06	49.19	41.10	49.12	41.10
QSW	29.23	37.57	49.34	47.20	49.78	47.20
CR	28.48	34.86	50.14	43.09	50.09	43.00
CG	24.51	31.75	50.04	44.56	49.86	44.29
LG	21.71	27.42	48.28	46.80	49.01	46.80
BSIF	18.94	25.72	49.64	41.03	49.86	41.03

Notes. EER values close to 50% indicate that the distinctive information of the samples has been completely removed. U-Net uses only the parameters of the circles approximating the inner and outer iris boundaries and segments the iris region in the coordinate system of the RSM. The results of U-Net refer to the deep neural network trained for the IITD dataset, which achieved the best results in our tests.




FIGURE 11. ROC curves obtained comparing the biometric recognition schemes (a) BSIF and (b) LG with manually segmented masks and automatically computed segmentation masks (R-cnn + cnn2rubber) for I-SOCIAL-DB, DB-DeIdent-Iris_{GAN}, and DB-DeIdent-Iris_{fractal}. The results are compared with the ROC curve obtained from a vector of random numbers of size equal to the number of identity comparisons performed for I-SOCIAL-DB, DB-DeIdent-Iris_{GAN}, and DB-DeIdent-Iris_{fractal}. The proposed deidentification method effectively removes the biometric information present in the iris region since the curves obtained for DB-DeIdent-Iris_{GAN} and DB-DeIdent-Iris_{fractal} by using manually segmented masks are similar to the curve obtained from randomly generated numbers.

evaluated the performance of automatic face recognition methods.

1) Visual analysis

Fig. 15 shows a face image and a corresponding image with iris regions deidentified using the proposed method. Then, Fig. 16 and Fig. 17 show examples of images selected from DB-DeIdent-Face_{GAN}. Specifically, Fig. 16 shows a complete image of the face, while Fig. 17 shows only the iris region. We considered only the images of the DB-DeIdent-Face_{GAN} database because, from our visual examination, they exhibited a greater visual realism than did the images in DB-DeIdent-Face_{Fractal} (in agreement with the opinions

of the volunteers involved in our tests). It can be observed that the proposed method generates highly realistic images in which the iris patterns closely resemble the original patterns but contain synthetic information unrelated to the original biometric traits.

2) Applicability of iris segmentation methods

We compared the segmentation accuracy achieved via different methods for I-SOCIAL-DB and for deidentified images of DB-DeIdent-Face_{GAN} and DB-DeIdent-Face_{fractal}. The considered iris segmentation methods are based on heterogeneous features, including edge-based and texture-based features, as well as those learned by using deep neu-

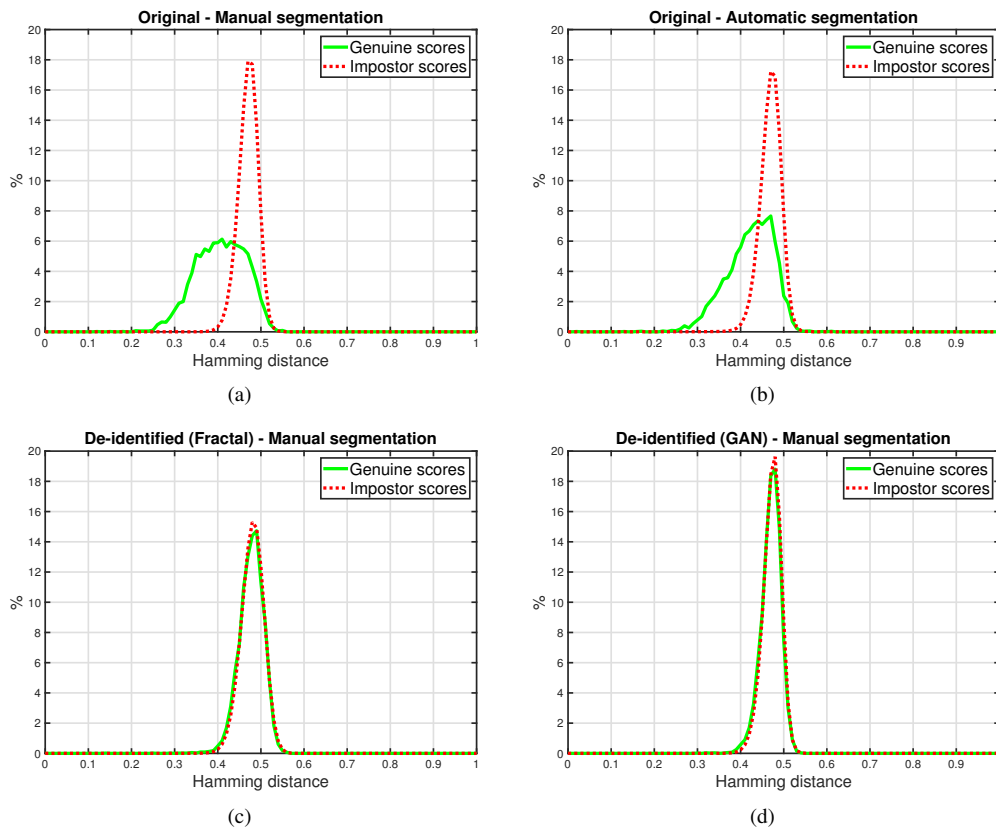


FIGURE 12. Distributions of the matching scores obtained by applying the proposed iris deidentification: (a) ocular images segmented by a human expert; (b) ocular images segmented using automatic segmentation software; (c) ocular images segmented by a human expert and deidentified using the fractal approach; and (d) ocular images segmented by a human expert and deidentified using the GAN approach. The deidentified ocular images (under both configurations) do not include distinctive information in the iris region; thus, they obtain a genuine score distribution comparable to the impostor score distribution. Furthermore, the results in part (b) show that the privacy risk is still relevant even when an automatic segmentation algorithm is used.




FIGURE 13. Examples of patterns generated by the proposed GAN for the same sample before blending the edges: (a) original ocular image; (b) deidentified image obtained by the first execution of the GAN; and (c) deidentified image obtained by a second GAN execution. To enhance the visibility of the differences between the images (b) and (c), in this example, we did not apply the blending algorithm (subsection IV-E), which is designed to smooth the transition between the iris and sclera. Examples of iris images obtained by applying the the blending algorithm (subsection IV-E) are shown in Fig. 17. A visual inspection shows that deidentified images generated for the same individual have substantial iris texture differences.

ral networks. In particular, we considered a segmentation method based on the total variation model (TVM) [57], a fast segmentation algorithm for nonideal images (FSA) [58], a segmentation technique based on deep learning (R-CNN) [38], and three segmentation algorithms included in USIT version 3.0 [8] (contrast-adjusted Hough transform (CAHT) [34], iterative Fourier-series push and pull (IFPP) [59], and WAHET [60]). We considered two figures of merit com-




FIGURE 14. Distribution of the matching scores obtained by comparison of the RSMs of the samples of the training set with 1,000 randomly generated RSMs. The shape of the matching score distribution is similar to the shape of the impostor distributions obtained using the same algorithm and shown in Fig. 12. Furthermore, the mean of the matching scores is 0.48; a similar value is usually obtained by the employed matcher when applied to impostor identity comparisons performed for samples with substantial differences. These results show the ability of the GAN to create images different from those in the training set. The obtained results also prove that the proposed deidentification method guarantees robustness to reidentification attacks even in cases in which the samples to be deidentified pertain to the training set because the textures generated by the GAN do not present distinctive biometric characteristics in common with the samples in the training set.




FIGURE 15. Examples of face images with iris regions deidentified using the proposed approach: (a) original image; (b) face image with deidentified iris regions. The proposed method for iris deidentification generates images with high visual realism.

only used in the literature, introduced for the Noisy Image Challenge Evaluation, Part 1 (NICE.I) competition [61]: the classification error rate (E1) and a metric for evaluating the disproportion between the false positive rate (FPR) and false negative rate (FNR) of the pixel classification (E2). Table 5 summarizes the obtained results.

Table 5 shows that the state-of-the-art segmentation methods based on heterogeneous features achieved similar accuracy for real samples and deidentified images. These results prove that the proposed deidentification method does not substantially affect the performance of iris segmentation methods.

3) Analysis of questionnaires

To further evaluate the capability of the proposed method to generate deidentified images with high visual realism, we used an evaluation procedure based on questionnaires compiled by volunteers. The questionnaires consisted of evaluating the visual aspect of real images and the deidentified images. The images were presented to nonexperts based on two criteria: 1) user appreciation and 2) visual realism. We considered the answers to questionnaires from participants shown complete face images as well as those shown only the ocular region. In the first questionnaire, we compared the results obtained by the fractal and GAN algorithms. We extracted the face/ocular region from 15 images randomly selected from F-SOCIAL-DB, DB-DeIdent-Face_{Fractal}, and DB-DeIdent-Face_{GAN} and asked users whether they preferred the samples from DB-DeIdent-Face_{Fractal} or those from DB-DeIdent-Face_{GAN}. The test was conducted with 16 volunteer participants, yielding 240 answers in total. We displayed the images to the participants on different screen types, such as those of laptops and smartphones, and showed the original image and the corresponding deidentified image on the same page (see examples in Fig. 16). The ocular regions (see examples in Fig. 17) are shown at a zoom factor of 100%. After considering the face images, 99.3%



FIGURE 16. Examples of faces with irises deidentified using the proposed approach (face image), selected from DB-DeIdent-Face_{GAN}. The proposed method for iris deidentification generates images with high visual realism.



FIGURE 17. Examples of images with irises deidentified using the proposed approach (only iris region), selected from DB-DeIdent-Iris_{GAN}. The proposed method for iris deidentification generates images with high visual realism.

of the users assigned a major or equal rate to the GAN-based configuration. After considering the ocular images, 91.2% of the users assigned a major or equal rate to the GAN-based configuration. Taken together, the results showed that the images generated with the GAN-based configuration received higher approval from users than did those produced by the fractal algorithm.

In the second questionnaire, we performed a Turing-like test by extracting the face/ocular region from 30 images randomly selected from F-SOCIAL-DB and from DB-DeIdent-Face_{GAN} and asking each user to decide whether the image was real or synthetic. In this test, we considered only the results obtained by the GAN algorithm, since this method yielded the best results in the previous test. This test was performed by 32 volunteers. We displayed the images on different screen types, such as laptops and smartphones. The ocular regions are shown at a zoom factor of 100%. The results showed that 62.5% of the real faces were not recognized as real samples and that 57.1% of deidentified faces were not recognized as synthetic samples. Similarly, the results showed that 41.2% of the real ocular images were not recognized as real samples and that 38.5% of deidentified eyes were not recognized as synthetic samples. These results indicate high error levels in the user judgments; they did not correctly identify many faces and ocular regions as showing deidentified or real iris patterns, thus demonstrating that users were not able to perceive relevant differences between real and deidentified images.

4) Face recognition performance

We also evaluated the effect of proposed iris deidentification approach on the performance of state-of-the-art face

TABLE 5. Accuracy of different iris segmentation methods for real and deidentified images.

Segmentation Library	I-SOCIAL-DB Segmentation error		DB-DeIdent-Iris _{GAN} Segmentation error		DB-DeIdent-Iris _{fractal} Segmentation error	
	E1	E2	E1	E2	E1	E2
WAHET	0.1347	0.2831	0.1571	0.3118	0.1571	0.3118
IFPP	0.1121	0.1855	0.1522	0.2557	0.1513	0.2539
FSA	0.0943	0.3226	0.1011	0.3315	0.1036	0.3343
CAHT	0.0862	0.4042	0.0871	0.4062	0.0871	0.4062
TVM	0.0316	0.1406	0.0302	0.1273	0.0300	0.1273
R-CNN	0.0146	0.0660	0.0144	0.0600	0.0142	0.0600

Notes. When needed, we set the parameters describing the minimum and maximum radii of the circles approximating the iris boundaries. We did not modify any other parameter of the segmentation methods. The R-CNN uses the configurations designed for the University of Beira Interior Iris version 2 (UBIRIS v.2) database.

recognition methods. Specifically, we evaluated the accuracy of the deep neural networks described in [62] for F-SOCIAL-DB, DB-DeIdent-Face_{fractal}, and DB-DeIdent-Face_{GAN}. The considered deep neural networks achieved similar performances for each dataset. As an example, the squeeze-and-excitation network (SeNet) achieved an EER of approximately 1.4% for the three datasets. The achieved results show that the proposed iris deidentification method preserves the original details of the face and maintains face pictures that are recognizable by both humans and state-of-the-art biometric recognition methods working effectively only on the iris pattern.

VI. CONCLUSION

In this paper, we raised a significant privacy problem caused by the possibility of applying state-of-the-art iris recognition techniques on images uploaded on websites and social media. First, we empirically demonstrated that the risk associated with iris-based identification is real. Second, we presented an iris deidentification method based on generative adversarial networks, which automatically generates novel images with high visual realism, in which all the distinctive biometric features of the iris textures are removed and substituted. We evaluated the deidentification capability of the proposed deidentification method as well as its ability to construct realistic images. The results showed that iris recognition algorithms are unable to extract distinctive features from the computed deidentified samples. Furthermore, a panel of interviewed volunteers was not able to correctly distinguish between the real and deidentified images. Based on the obtained results, our method can be used as an effective privacy-preserving tool when uploading high-resolution facial images to websites and social media. The use of our method guarantees that the iris visible in the uploaded images does not contain any identifiable biometric information and works without introducing modifications easily recognizable by humans.

REFERENCES

- [1] S. Marcel, M. S. Nixon, J. Fierrez, and N. W. D. Evans, Eds., *Handbook of Biometric Anti-Spoofing - Presentation Attack Detection, Second Edition*. Springer, 2019.
- [2] A. Malhotra, S. Chhabra, M. Vatsa, and R. Singh, "On privacy preserving anonymization of finger-selfies," in *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020, pp. 120–128.
- [3] K. Nguyen, C. Fookes, R. Jillela, S. Sridharan, and A. Ross, "Long range iris recognition: A survey," *Pattern Recognition*, vol. 72, pp. 123–143, 2017.
- [4] R. Donida Labati, E. Muñoz, V. Piuri, A. Ross, and F. Scotti, "Non-ideal iris segmentation using Polar Spline RANSAC and illumination compensation," *Computer Vision and Image Understanding*, 2019.
- [5] N. A. Schmid, J. Zuo, F. Nicolo, and H. Wechsler, "Iris quality metrics for adaptive authentication," in *Handbook of Iris Recognition*, J. M. Burge and W. K. Bowyer, Eds. London: Springer, 2013, pp. 67–84.
- [6] R. Donida Labati, A. Genovese, V. Piuri, F. Scotti, and S. Vishwakarma, "I-SOCIAL-DB: A labeled database of images collected from websites and social media for iris recognition," *Image and Vision Computing*, vol. 105, no. 104058, pp. 1–9, 2021.
- [7] Microsoft AI, "Face and emotion recognition," 2019. [Online]. Available: <https://aidemos.microsoft.com/face-recognition>
- [8] C. Rathgeb, A. Uhl, P. Wild, and H. Hofbauer, "Design decisions for an iris recognition SDK," in *Handbook of Iris Recognition*, K. W. Bowyer and M. J. Burge, Eds. Springer London, 2016, pp. 359–396.
- [9] R. Donida Labati, A. Genovese, V. Piuri, and F. Scotti, "Iris segmentation: state of the art and innovative methods," in *Cross Disciplinary Biometric Systems*, ser. Intelligent Systems Reference Library, C. Liu and V. Mago, Eds. Springer, 2012, vol. 37, pp. 151–182.
- [10] J. Daugman, "Information theory and the IrisCode," *IEEE Trans. on Information Forensics and Security*, vol. 11, no. 2, pp. 400–409, February 2016.
- [11] A. Dantcheva, C. Chen, and A. Ross, "Can facial cosmetics affect the matching accuracy of face recognition systems?" in *Proc. of the 2012 IEEE Fifth Int. Conf. on Biometrics: Theory, Applications and Systems (BTAS)*, 2012, pp. 391–398.
- [12] S. K. S. Modak and V. K. Jha, "Multibiometric fusion strategy and its applications: A review," *Information Fusion*, vol. 49, pp. 174–204, 2019.
- [13] D. Lee and K. N. Plataniotis, "A novel eye region based privacy protection scheme," in *Proc. of ICASSP*, 2012.
- [14] M. Rieger, J. Hammerle-Uhl, and A. Uhl, "Efficient iris sample data protection using selective JPEG2000 encryption of normalised texture," in *Proc. of IWBF*, 2018.
- [15] H. Zhang, H. Zhou, W. Jiao, J. Shi, Q. Zang, J. Sun, and J. Zhang, "Biological features de-identification in iris images," in *Proc. of the 15th International Symposium on Pervasive Systems, Algorithms and Networks (I-SPAN)*, 2018, pp. 67–71.
- [16] S. Ribaric, A. Ariyaecinia, and N. Pavesic, "De-identification for privacy protection in multimedia content: A survey," *Signal Processing: Image Communication*, vol. 47, pp. 131–151, 2016.
- [17] C. Neustaedter, S. Greenberg, and M. Boyle, "Blur filtration fails to preserve privacy for home-based video conferencing," *ACM Trans. Comput.-Hum. Interact.*, vol. 13, no. 1, pp. 1–36, March 2006.
- [18] S. De Capitani di Vimercati, S. Foresti, G. Livraga, and P. Samarati, "Data privacy: Definitions and techniques," *Int. Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 20, no. 6, pp. 793–817, December 2012.
- [19] P. Samarati, "Protecting respondents identities in microdata release," *IEEE Trans. on Knowledge and Data Engineering*, vol. 13, no. 6, pp. 1010–1027, November 2001.

- [20] E. M. Newton, L. Sweeney, and B. Malin, "Preserving privacy by de-identifying face images," *IEEE Trans. on Knowledge and Data Engineering*, vol. 17, no. 2, pp. 232–243, February 2005.
- [21] B. Meden, R. C. Malli, S. Fabijan, H. K. Ekenel, V. Struc, and P. Peer, "Face deidentification with generative deep neural networks," *IET Signal Processing*, vol. 11, no. 9, pp. 1046–1054, 2017.
- [22] B. Dolhansky and C. C. Ferrer, "Eye in-painting with exemplar generative adversarial networks," in *Proc. of CVPR*, June 2018, pp. 7902–7911.
- [23] J. Daugman, "How iris recognition works," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 14, pp. 21–30, 2002.
- [24] H. Kaur and R. Manduchi, "EyeGAN: Gaze-preserving, mask-mediated eye image synthesis," in *Proc. of WACV*, 2020, pp. 299–308.
- [25] S. Shah and A. Ross, "Generating synthetic irises by feature agglomeration," in *Proc. of ICIP*, 2006.
- [26] J. Zuo, N. A. Schmid, and X. Chen, "On generation and analysis of synthetic iris images," *IEEE Trans. on Information Forensics and Security*, vol. 2, no. 1, pp. 77–90, March 2007.
- [27] Z. Wei, T. Tan, and Z. Sun, "Synthesis of large realistic iris databases using patch-based sampling," in *Proc. of ICPR*, 2008.
- [28] L. Wecker, F. Samavati, and M. Gavrilova, "A multiresolution approach to iris synthesis," *Computers & Graphics*, vol. 34, no. 4, pp. 468–478, 2010.
- [29] L. Cardoso, A. Barbosa, F. Silva, A. M. G. Pinheiro, and H. Proença, "Iris biometrics: Synthesis of degraded ocular images," *IEEE Trans. on Information Forensics and Security*, vol. 8, no. 7, pp. 1115–1125, July 2013.
- [30] N. Kohli, D. Yadav, M. Vatsa, R. Singh, and A. Noore, "Synthetic iris presentation attack using iDCGAN," in *Proc. of IJCB*, 2017.
- [31] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb, "Learning from simulated and unsupervised images through adversarial training," in *Proc. of CVPR*, 2017.
- [32] S. Minaee and A. Abdolrashedi, "Iris-GAN: Learning to generate realistic iris images using convolutional GAN," *CoRR*, vol. abs/1812.04822, 2018.
- [33] S. Yadav, C. Chen, and A. Ross, "Synthesizing iris images using RaSGAN with application in presentation attack detection," in *Proc. of CVPRW*, 2019.
- [34] C. Rathgeb, A. Uhl, and P. Wild, *Iris Recognition: From Segmentation to Template Security*. Springer, Berlin, 2013.
- [35] L. Masek and P. Kovesi, "MATLAB source code for a biometric identification system based on iris patterns," 2003. [Online]. Available: <https://www.peterkovesi.com/studentprojects/libor/sourcecode.html>
- [36] CASIA, "CASIA Iris Image Database V4.0," 2010. [Online]. Available: <http://www.cbsr.ia.ac.cn/china/Iris%20Databases%20CH.asp>
- [37] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, October 2016.
- [38] S. Ahmad and B. Fuller, "Unconstrained iris segmentation using convolutional neural networks," in *Computer Vision – ACCV 2018 Workshops*, G. Carneiro and S. You, Eds. Cham: Springer International Publishing, 2019, pp. 450–466.
- [39] A. Fournier, D. Fussell, and L. Carpenter, "Computer rendering of stochastic models," *Commun. ACM*, vol. 25, no. 6, pp. 371–384, Jun. 1982.
- [40] Z. Wang, Q. She, and T. E. Ward, "Generative adversarial networks in computer vision: A survey and taxonomy," *ACM Comput. Surv.*, vol. 54, no. 2, Feb. 2021.
- [41] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," *CoRR*, vol. abs/1812.04948, 2018.
- [42] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," 2015.
- [43] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *CoRR*, vol. abs/1502.03167, 2015.
- [44] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. of ICLR*, 2014.
- [45] P. Pérez, M. Gangnet, and A. Blake, "Poisson image editing," *ACM Trans. on Graphics*, vol. 22, no. 3, pp. 313–318, July 2003.
- [46] K. P. Hollingsworth, K. W. Bowyer, and P. J. Flynn, "The best bits in an iris code," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 31, no. 6, pp. 964–973, June 2009.
- [47] IITD, "IIT Delhi Iris Database (Version 1.0)," 2007. [Online]. Available: https://www4.comp.polyu.edu.hk/~csajaykr/IITD/Database_Iris.htm
- [48] H. Hofbauer, F. Alonso-Fernandez, P. Wild, J. Bigun, and A. Uhl, "A ground truth for iris segmentation," in *Proc. of ICPR*, August 2014, pp. 527–532.
- [49] A. U. Heinz Hofbauer, Ehsaneddin Jalilian, "Exploiting superior cnn-based iris segmentation for better recognition accuracy," *Pattern Recognition Letters*, vol. 120, pp. 17–23, 2019.
- [50] Z. Zhao and A. Kumar, "Towards more accurate iris recognition using deeply learned spatially corresponding features," in *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 3829–3838.
- [51] A. Czajka, D. Moreira, K. Bowyer, and P. Flynn, "Domain-specific human-inspired binarized statistical image features for iris recognition," in *Proc. of the IEEE Winter Conf. on Applications of Computer Vision*, 2019, pp. 959–967.
- [52] L. Masek and P. Kovesi, "MATLAB source code for a biometric identification system based on iris patterns," School of Comput. Sci. and Software Eng., The University of Western Australia, 2003.
- [53] C. Rathgeb and A. Uhl, "Secure iris recognition based on local intensity variations," in *Image Analysis and Recognition*, A. Campilho and M. Kamel, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 266–275.
- [54] J.-G. Ko, Y.-H. Gil, J.-H. Yoo, and K.-I. Chung, "A novel and efficient feature extraction method for iris recognition," *ETRI Journal*, vol. 29, no. 3, pp. 399–401.
- [55] L. Ma, T. Tan, Y. Wang, and D. Zhang, "Efficient iris recognition by characterizing key local variations," *IEEE Trans. on Image Processing*, vol. 13, no. 6, pp. 739–750, June 2004.
- [56] A. K. Jain, P. Flynn, and A. A. Ross, *Handbook of Biometrics*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2007.
- [57] Z. Zhao and A. Kumar, "An accurate iris segmentation framework under relaxed imaging constraints using total variation model," in *2015 IEEE Int. Conf. on Computer Vision (ICCV)*, 2015, pp. 3828–3836.
- [58] A. Gangwar, A. Joshi, A. Singh, F. Alonso-Fernandez, and J. Bigun, "Irisseg: A fast and robust iris segmentation framework for non-ideal iris images," in *Proc. of the Int. Conf. on Biometrics*, June 2016, pp. 1–8.
- [59] J. Daugman, "New methods in iris recognition," *IEEE Trans. Syst., Man, Cybern., B, Cybern.*, vol. 37, pp. 1167–1175, October 2007.
- [60] P. Wild, H. Hofbauer, J. Ferryman, and A. Uhl, "Segmentation-level fusion for iris recognition," in *Proc. of the 2015 Int. Conf. of the Biometrics Special Interest Group*, September 2015, pp. 1–6.
- [61] H. Proença and L. A. Alexandre, "Introduction to the special issue on the recognition of visible wavelength iris images captured at-a-distance and on-the-move," *Pattern Recognition Letters*, vol. 33, pp. 963–964, 2012.
- [62] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "VGGFace2: A dataset for recognising faces across pose and age," in *Proc. of the 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, May 2018, pp. 67–74.

•••