

Synthetic ID Card Image Generation for Improving Presentation Attack Detection

Daniel Benalcazar, *Member, IEEE*, Juan E. Tapia^{ID}, *Member, IEEE*, Sebastian Gonzalez, *Member, IEEE*, and Christoph Busch^{ID}, *Senior Member, IEEE*

Abstract—Currently, it is ever more common to access online services for activities which formerly required physical attendance. From banking operations to visa applications, a significant number of processes have been digitised, especially since the advent of the COVID-19 pandemic, requiring remote biometric authentication of the user. On the downside, some subjects intend to interfere with the normal operation of remote systems for personal profit by using fake identity documents, such as passports and ID cards. Deep learning solutions to detect such frauds have been presented in the literature. However, due to privacy concerns and the sensitive nature of personal identity documents, developing a dataset with the necessary number of examples for training deep neural networks is challenging. This work explores three methods for synthetically generating ID card images to increase the amount of data while training fraud-detection networks. These methods include computer vision algorithms and Generative Adversarial Networks. Our results indicate that databases can be supplemented with synthetic images without any loss in performance for the print/scan Presentation Attack Instrument Species (PAIS) and a loss in performance of 1% for the screen capture PAIS.

Index Terms—Biometrics, ID card, tampering, presentation attack detection, synthetics images, GANs.

I. INTRODUCTION

THE GLOBAL pandemic of COVID-19 accelerated the adoption of remote biometric authentication for online services such as e-commerce, digital banking, fintech and document signing. This allowed people to carry on with their normal business activities from home without the risk of spreading the virus. Some services include remotely opening a bank account, something that required physical attendance only a few years back. To access this service, a user only has to validate his/her identity by capturing a selfie and a picture of their ID document. Therefore, remote services

Manuscript received 14 July 2022; revised 3 February 2023; accepted 6 March 2023. Date of publication 10 March 2023; date of current version 20 March 2023. This work was supported in part by the European Union's Horizon 2020 Research and Innovation Program under Agreement 883356; and in part by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science, and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE and TOC Biometrics, Research and Development Center. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Fernando Alonso-Fernandez. (*Corresponding author: Juan E. Tapia.*)

Daniel Benalcazar and Sebastian Gonzalez are with the Research and Development Center, TOC Biometrics Company, Santiago 7520424, Chile (e-mail: daniel.benalcazar@tocbiometrics.com; sebastian.gonzalez@tocbiometrics.com).

Juan E. Tapia and Christoph Busch are with the da/sec-Biometrics and Internet Security Research Group, Hochschule Darmstadt, 64295 Darmstadt, Germany (e-mail: juan.tapia-farias@h-da.de; christoph.busch@h-da.de).

Digital Object Identifier 10.1109/TIFS.2023.3255585

made commercial activities more dynamic and accessible for everybody while diminishing the chances of infection.

The downside of remote services is twofold. Firstly, in regions like South America, the accelerated technological leap was too quick for some countries, resulting in difficulties for the national identification systems to catch up with the advancements. For instance, identity documents issued by many countries neither comply with ICAO 9303 standard¹ nor have NFC chips with the personal information of the individual embedded in their government-issued ID cards. This makes it difficult to automatically verify an ID card's genuine nature. Secondly, certain individuals might attempt to defraud the remote service by tampering with the ID card's photo ID or other fields to conceal their identities or impersonate somebody else. In those cases, it is of utmost importance to rely on a system that can check whether an ID card digital image has been manipulated. Even today, many templates of ID cards, passports, driver's licenses, and others are available to help an attacker produce a high-quality fake document.²

Methods that detect fraud in personal ID documents have been presented in recent years [1], [2], [3], [4]. Those systems rely on deep learning methods, such as Convolutional Neural Networks (CNN), in order to achieve great detection accuracy [5]. Deep learning systems in general require a large number of examples to train successfully; however, the sensitive nature of ID cards and passports makes it very difficult to acquire the number of images needed. For that reason, in this work, we propose to create synthetic examples of ID card images to enhance the dataset on which fraud-detection networks are trained. We hypothesise that with the additional synthetic samples, a CNN will produce comparable results to having more real ID card samples to train with. We explored state-of-the-art Generative Adversarial Networks (GAN), as well as image processing techniques for synthetic image generation.

In this work, we consider three kinds of Presentation Attack Instruments (PAI) which are named “composite (modified manually or automatically)”, “print” and “screen”.

In this way, the four species used in this work for classification, illustrated in Figure 1, are bona fide, composite, print, and screen:

- **Bona Fide:** Real Chilean ID cards were captured. We used the same dataset as [3]. An example can be seen in Figure 1a.

¹https://www.icao.int/publications/Documents/9303_p4_cons_en.pdf

²<https://gotempl.cc/product-category/id/>



Fig. 1. Examples of captured Chilean ID cards of the bona fide class, as well as the presentation attack classes: composite, print and screen.




Fig. 2. Process of device-noise texture isolation for artificial generation of print and screen attack species.

- *Composite*: Bona Fide ID cards were printed, and the face and/or text fields of one ID card were manually cut with scissors and glued on top of another printed ID card (Figure 1b). Additionally, image processing methods were used to mimic this process automatically.
- *Print*: Bona Fide ID cards were printed on both normal and glossy paper. Then, smartphones were used to capture the printed images in a variety of backgrounds and illumination conditions (Figure 1c).
- *Screen*: Bona Fide ID cards were depicted on computer monitors, as well as tablet and smartphone screens. Then, a second smartphone was used to capture the screen images (Figure 1d).

To properly increase the number of captured ID card image samples in our datasets for fraud detection using synthetic images, examples of the four classes must be generated: bona fide, composite, print and screen. We explore GAN models to create more samples of the bona fide class, which can draw more examples of the same distribution of the input images [6], as well as synthesising ID card images from templates with image processing techniques. On the other hand, we propose adding texture noise from the printing/scanning and screen capture processes to create more examples of attack presentation ID card images. For this purpose, we used GANs, as well as an image processing-based noise isolation and addition technique. Later, we also used GANs and automatic splicing techniques for the composite scenario that combined two ID card images. Finally, To evaluate the impact of the synthetic images, we trained several MobileNetV2 [7] networks (based on [3]) using images of Chilean ID cards. It compared the training performance with captured images only against training with a combination of captured and synthetic images. It is essential to note that this work is one of the first to synthesise ID card images directly.

In this work, we explore using GAN and template-generated images from the Bona fide distribution as data supplements for authentic Bona fide images. If the experiments showed a positive impact with this approach, this would be a cost-effective way of obtaining new data, while not asking new subjects for

sensitive data. Other research fields, such as Morphing Attack Detection also use a similar approach [8].

The contributions of this work are as follows:

- *ID card generation from templates*: A traditional image processing algorithm is developed, capable of generating fake ID card images from a clear template that resembles the original bona fide and presentation attack scenarios.
- *Texture transfer-based presentation attack ID card image generation*: A traditional image processing algorithm capable of generating ID card images faster than GAN-based methods is proposed. This algorithm transfers the noise textures from presentation attack species images to bona fide images. These noise texture templates will be available to other researchers (Figure 2).
- *Synthetic ID card generation*: GAN models capable of generating synthetic ID cards that resemble the original bona fide and presentation attack images are developed. Generated images will be made available to other researchers.
- *Benchmark evaluation*: The influence of having synthetic ID card images as part of the training dataset of a state-of-the-art ID card image fraud detection network [3] is evaluated.
- *Analysis*: A comprehensive analysis of the difficulties of synthetic ID card generation and the proposed methods' benefits are provided.

The rest of the article is organised as follows: Section II summarises the related works on generative adversarial networks and image tampering. New methods for generating synthetic images are described in Section III. The experimental framework and results of this work are then presented in Section VI. We conclude the article in Section VII.

II. RELATED WORK

The GAN algorithm was first introduced by Ian Goodfellow et al. [9]. It approaches the problem of unsupervised learning by simultaneously training two deep networks, called Generator G and Discriminator D , respectively. These networks compete and cooperate with

each other. While the generator creates new instances of the data, the discriminator evaluates them for authenticity. In the course of training, both networks learn to perform their tasks. To learn the generator distribution p_g over data x , the generator builds a mapping function from a prior noise distribution $p_z(z)$ to data space as $G(z; \theta_g)$. The discriminator $D(x; \theta_d)$, on the other hand, outputs a single scalar representing the probability that x came from training data rather than p_g . During training, the parameters of G to minimise $\log(1 - D(G(z)))$ and the parameters of D to minimise $\log(D(x))$ are simultaneously adjusted as if they were following a two-player min-max game with value function $V(G, D)$, specified in Equation 1.

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

Karras et al. developed StyleGAN2 [6], [10] as an extension of the progressive, growing GANs. This approach enables training face generator models capable of synthesising large high-quality images via the incremental expansion of discriminator and generator models from small to large images during the training process. In addition to the gradual growth of the models during training, StyleGAN2 changes the architecture of the generator significantly. The StyleGAN2 generator no longer takes a point from the latent space as input; instead, two new sources of randomness are used to generate a synthetic image: a standalone mapping network and noise layers. StyleGAN2 introduces the mapping network f to transform z into this intermediate latent space w using eight fully connected layers. This intermediate latent space w can be viewed as the new z , (z'). Through this network, a 512-D latent space z is transformed into a 512-D intermediate latent space w applied to face images.

CycleGAN [11] is a technique for training unsupervised image translation models via the GAN architecture using unpaired collections of images from two different domains. Its utility has been demonstrated in a range of applications, including translating seasons in photographs, object transfiguration, style transfer, colorization, and generating photos from paintings. The discriminator models classify 70×70 overlapping patches of input images as belonging to the domain or having been generated; the discriminator output is then taken as the average prediction for each patch.

There are some works in the literature that have tackled the problem of fake ID cards detection. For instance, Shi and Jain [1], [2] proposed DocFace and DocFace+ to determine the authenticity of a personal ID document by comparison of a face image (selfie) with the photo ID in the ID document. In DocFace+, a kiosk scans the ID document photos or reads the photo from the embedded chip using an NFC reader.

Zheng et al. [5] present a survey and provide an overview on typical image tampering methods, released image tampering datasets, and recent tampering detection approaches. It presents a distinct perspective to rethink various assumptions about tampering clues, which can be discovered by different detection approaches. This further encourages the research community to develop general tampering localisation methods

in the future instead of adhering to single-type tampering detection. Most of the analysis was realised using handcrafted images in different domains. Deep learning methods have not been explored in detail.

Albiero et al. [12] present a method for compare selfie images to photo ID images from Chilean ID cards, across adolescence, employing fine-tuning techniques using a private dataset.

Stokkenes et al. [13], proposed an online banking authentication system based on features extracted from faces using bloom filters. This information is encoded and used as a key for accessing banking services.

Perera and Patel [14] proposed an active authentication system that attempts to continuously monitor user identity after access has been initially granted. A similar approach has recently been reported by Fathy et al. [15].

Arlazarov et al. [16] presented a tiny ID cards dataset containing 500 video clips of 50 different identity document types. This dataset was one of the first made publicly available for identity document analysis and recognition in video stream. Additionally, the paper presents three experimental baselines obtained using the dataset: face detection accuracy, separate text fields OCR precision for four major identity document field types, and identity document data extraction from video clips.

On a different approach, Gonzalez et al. [3] proposed a hybrid two-stage classification system that checks if the entire Chilean ID card was tampered or modified. In that work, the authors analysed the primary sources of fraud as image composition and image source tampering. This paper is relevant because it presents results on genuine transactions of a remote verification system with 24,778 images distributed on bona fide, composite, printed and display ID card images.

Zhu et al. [17] propose a method that indicates whether the photo ID or fields in the ID card image have been altered or replaced by digital or handcrafted means. On the other hand, if the source of the ID card comes from a printing/scanning process or the image was captured from a digital screen, it means that it does not come from the original plastic document, and alterations could have been made beforehand.

Mudgalgundurao et al. [4] proposed a method to detect fake German ID cards and residence permits using pixel-wise supervision based on DenseNet. This technique enables the method to leverage minute cues on various artefacts, such as moiré patterns and artefacts left by the printers. The authors present the baseline benchmark using different handcrafted and deep learning models on a newly constructed in-house database obtained from an operational system consisting of 886 users with 433 bona fide, 67 print and 366 display attacks.

Regarding the training and testing data, due to the privacy concerns of the ID cards, the databases for research and commercial purposes are difficult to obtain. So far, there is a limited number of databases available: the Chilean Sequestered ID card database used in [3], Public-IvS where the ID card images are cropped from the existing images from CASIA-IvS [17], and the German ID card and resident permit images [4]. Furthermore, the MIDV-500 dataset [16], which contains video clips of 50 different identity document types,

including 17 types of ID cards, 14 types of passports, 13 types of driving licenses, are also available.

III. METHODS

This section describes the implementation of two MobileNetV2 networks used as an ID card fraud-detection system and the generation of synthetic images using three different methods: templates, texture transfer and GAN models for ID cards. The entire ID card with the background removed [18] is the input to the MobileNetV2 networks. In this scheme, the networks have to learn how to identify forgery traits from any part of the ID card [3].

The goal of the synthetic images is to be as close as possible to the captured images of the Chilean ID card database; however, these techniques can be easily extended to other countries. The Chilean ID card dataset used in this work is explained in Section V, including the train, test and validation partitions of all the methods. Each image on the dataset was segmented using the network proposed by Lara et al. [18] in order to remove the background pixels and force the fraud-detection network to learn features only from the pixels inside of the ID card. Figure 1 show examples of the segmented ID card images.

A. Fraud-Detection Network

To detect tampering in ID cards and evaluate the impact of synthetic images, we implemented two MobileNetV2 [7] approaches based on [3]. The first network detected composition tampering and was trained to perform a binary classification task between bona fide and composite images. The network input size is 224×224 pixels, and the number of filters in the network was increased using an alpha parameter of 1.4. For the second network, the system must identify the source of the image among bona fide presentations and print/screen attack presentations. In this case, the network input is 448×448 pixels, and the alpha parameter was also raised to 1.4. The greater the resolution better represents the fine-grain device noise, which would be lost in a lower resolution.

We trained the two types of MobileNetV2 networks from ImageNet weights [19] on an Amazon AWS instance with 32 cores of CPU, 128GB of RAM and a GPU of 15GB. We used the following training hyperparameters for all the experiments: $bs=40$, $epochs=300$, and $lr = 1e - 4$ for the first network and $lr = 1e - 5$ for the second. Several of these networks were trained for the experiments in this work. The datasets and partitions for each experiment are explained in detail in Section VI.

B. Image Generation From Templates

This paper also analysed deterministic image-processing techniques that can automatically generate ID cards from templates. To create the templates, we started by selecting one high-quality bona fide example of a Chilean ID card. Then we straightened out the image and emptied all the fields, signatures and photos using Photoshop’s cloning tool. The result is a clean background template ready to be automatically filled, as shown in Figure 3. After that, a series of algorithms



Fig. 3. Example of the Chilean ID card template.

were developed to select random faces and signatures, as well as random names, dates and alphanumeric characters, to fill the templates with random information. The face images are selected from the bona fide probe FERET dataset,³ while the signatures originated from the Kaggle dataset of handwritten signatures.⁴ During the information generation process, we took special care to match the gender of the photo with that of the first name and the gender field on the ID card. We also generated plausible dates and numbers for each field. We draw from a dictionary containing Latin American countries’ most familiar names to create random names and surnames.

Afterwards, we applied a few post-processing steps. We first used a small random variation in the Hue, Saturation, and Value (HSV) channels to slightly change the colour of the ID cards. This method is intended to mimic changes in the light source. Then, we applied a random perspective change to resemble the differences in perspective found in bona fide ID card images. This is to simulate the difficulty of holding the camera perfectly parallel to the document when the user takes pictures of their ID cards. This is replicated using a random protective transformation. Finally, we confirm that the background of the image is entirely black ($R=0$, $G=0$, $B=0$), resembling the result from the segmentation network [18].

Since the face image dataset contains 2,069 images, we decided to create synthetic ID cards in batches of this number. In this way, every batch will have the same faces in the same order but with different signatures, text fields, colours and perspectives. In the remainder of this paper, we will refer to this method as “Templates”, for short.

C. Image Generation From Transferable Textures

This proposed image-processing technique isolates an image’s texture and transfers it to another. This process was applied to generate the print and screen scenarios. In this way, an automatic method could replace the labour-intensive process of capturing bona fide images, printing or displaying them on a screen, and re-capturing them to create replay attacks.

The primary motivation behind this method is the following: if we have a bona fide image, as well as the presentation attack version of the same image aligned pixel by pixel, in that case, the mathematical subtraction between the two images will cancel out the image information leaving only the differences caused by the PAI. However, the print/screen

³<https://www.nist.gov/itl/products-and-services/color-feret-database>

⁴<https://www.kaggle.com/datasets/divyanshrai/handwritten-signatures>




Fig. 4. Example of the 50 colours palette used to create artificial textures.

version is unavailable; therefore, we must capture this artefact from other images as a colour palette. After that, it is possible to apply this new texture to the print and screen scenarios. This technique will help us to isolate the noise texture left by the printing/screen or displaying processes. Then, the isolated texture can be added to any other image to tamper with the desired attack noise scenario artificially. To simplify the alignment process, we chose a palette of 50 solid colour images, as shown in Figure 4. Each colour has a QR code on the left side that identifies it.

To isolate the texture on each image of that dataset, we first read the included QR code to identify the original colour. Then, the colour rectangle is segmented, eliminating the QR code and the background. This process is illustrated in Figure 2. To obtain the segmentation mask (Figure 2b), we first sampled a 200×150 pixel patch around the central pixel of the image and computed the mean and standard deviation in the HSV space. Then, all the pixels in the image with colour values around the mean HSV colour of the central patch plus-minus four standard deviations were masked. A morphological closing operation is applied to the mask to fill any gaps. Then, the borders of the mask are obtained and the inscribed rectangle is computed, which is used to crop the image. Therefore, at this point, we have an image that only contains pixels with the main colour and the texture of interest (Figure 2e). Finally, we isolate the texture by subtracting the colour indicated by the QR code and save the texture to create a dataset of 5,000 print and 5,000 screen-isolated textures. In Figure 2f, a constant RGB value of (128, 128, 128) was added to the isolated texture for visualisation purposes only.

In the end, a dataset of 10,000 isolated textures is obtained, where any texture can be added to any bona fide image to generate the tampering effect artificially. This process consists of cropping the isolated texture with the image size of the bona fide image, and then applying a pixel-wise mathematical addition between the texture and the bona fide image, as illustrated in Figure 5. In this work, we applied the transferable texture technique to generate attack presentations from captured ID card images and Templates. Since the textures come from captured images, it is expected that fraud-detection networks would learn to identify the noise patterns as the corresponding attack scenario. We will refer to this method in the sections below as “Textures”.

D. Image Generation From GANs

The following GAN networks were used in this work: StyleGAN2-ADA [6], [10] and CycleGAN [11]. The later translates images of domain A into domain B, so examples of the two domains are needed during training. In other words, one must guide the training so the network can translate base images into the desired domain. For CycleGAN, datasets A

and B are unpaired. On the other hand, StyleGAN2 is capable of generating new examples of the desired domain without the need for base images [6]. To be more precise, StyleGAN2 takes a random vector in the latent space and transforms it into an image of the desired domain, using the distributions learnt during training. Therefore, these two networks are useful for different purposes, expanded below.

StyleGAN2-ADA was used to generate more examples of bona fide ID card images, as well as the three attack presentation species considered in this work: composite, print and screen. Therefore, the class-conditional version of SyleGAN2-ADA was used to generate examples of the four classes in a single training session. We trained this network using the Ch12-A dataset described in Table III, on a machine with a Ryzen 5 processor with 16GB of RAM and an Nvidia Titan-X GPU of 12GB. The best hyperparameters for training were: $\text{size}=256 \times 256$, $lr = 1e - 4$, $\text{bs}=20$, and $\text{kimgs}=2,000$. The main advantage of SyleGAN2-ADA is that it is optimised to learn from a dataset with limited examples [10]. However, synthesising ID card images from scratch is a difficult task to learn since the network has to create faces, signatures, alphanumeric characters, symbols, device noise and perspective changes all at once.

Conversely, CycleGAN was trained to create more examples of the print and screen classes from the bona fide class, using domain adaptation. In this way, the network is meant to learn and mimic the fine-grain textures left by paper, inks and printer devices, as well as pixels and aliasing left by screen displays. Therefore, the difficulty of this is smaller than the previous one since faces, signatures, and characters come from the base images, and the network has to focus only on generating device noise textures. We also trained this network on the Ch12-A dataset (Table III), using a Ryzen 5 computer with 16 GB of RAM and an Nvidia Titan-X GPU of 12GB. The best hyperparameters for training were: $\text{size}=224 \times 224$, $lr = 1e - 4$, $\text{bs}=8$, and $\text{epochs}=40$.

IV. METRICS

This section describes all the metrics that are used to compare the performance of the proposed methods, both for the image generation and Presentation Attack Detection (PAD) task.

A. Frechet Inception Distance

One of the difficulties with GAN algorithms, and in particular when applied to ID card images or biometrics in general, is how to meaningfully assess the quality of the resulting (synthesised) images. Only recently, a suite of qualitative and quantitative metrics have been developed to assess the performance of a GAN model based on the quality and diversity of the generated synthetic images [20], [21], [22]. Such proposed of such metrics are: The Inception Score (IS) [20], Frechet Inception Distance (FID) [21] and Perceptual Path Length (PPL) [10]. These metrics allow us to compare results from different GAN models. The FID score was used in this work to measure the objective quality of the ID card’s synthetic images.

Frechet Inception Distance (FID) compares the similarity between two groups of images A and B. First, to compute




Fig. 5. Process of device-noise texture addition. A random isolated texture is selected from the set of 5,000 available textures per PAI. The texture is cropped with the shape of the input image at a random location. The output textured ID card is the pixel-wise addition between the input ID card and the cropped texture. In this figure, an ID card generated by the Templates method is used (no real data), thus it is not blacked out.

the FID, all images from set A and set B have to be processed by an Inception v3 network [23], pre-trained on ImageNet [19]. Then the 2,048 feature vector of the Inception-v3 pool3 layer is stored for each image. Finally, the distributions of A and B in the feature space are compared using Equation 2, where μ_A and μ_B are the mean values of the distributions A and B, respectively, and Σ_A and Σ_B are the covariances of the two distributions.

$$FID = \|\mu_A - \mu_B\|^2 + Tr \left(\Sigma_A + \Sigma_B - 2(\Sigma_A \cdot \Sigma_B)^{1/2} \right) \quad (2)$$

B. Detection Performance Evaluation

The detection performance of biometric PAD algorithms is standardised by ISO/IEC 30107-3.⁵ The most relevant metrics for this study are: Attack Presentation Classification Error Rate (APCER), Bonafide Presentation Classification Error Rate (BPCER) and BPCER_{AP}. Those metrics determine the error rates when classifying an instance between bona fide and the different Presentation Attack Instrument Species (PAIS).

The APCER metric measures the percentage of attack presentations incorrectly classified as bona fide, for each different PAI. To evaluate an entire system, the worst-case scenario is considered. The computation method is detailed in Equation 3, where the value of N_{PAIS} corresponds to the number of attack presentation images, RES_i is 1 if the i th image is classified as an attack, or 0 if it was classified as a bona fide presentation.

$$APCER_{PAIS} = 1 - \frac{1}{N_{PAIS}} \sum_{i=1}^{N_{PAIS}} RES_i \quad (3)$$

On the other hand, the BPCER metric measures the proportion of bona fide presentations wrongly classified as attacks. BPCER can be computed using Equation 4, where N_{BF} is the amount of bona fide presentation images, and RES_i takes the same values described in the APCER metric. The two metrics together determine the performance of the system, and they are subject to a specific operation point.

$$BPCER = \frac{\sum_{i=1}^{N_{BF}} RES_i}{N_{BF}} \quad (4)$$

Finally, to analyse the system performance on a specific operating point, BPCER_{AP} and the Equal Error Rate (EER)

are used. The later is the operating point where APCER and BPCER are equal. This operating point corresponds to the intersection with the diagonal line in a Detection Error Trade-off (DET) curve, which is also reported for all the experiments. On the other hand, the BPCER_{AP} is the BPCER value when the APCER is 100/AP. In this work, we evaluate BPCER₁₀, BPCER₂₀ and BPCER₁₀₀, which correspond to APCER values of 10%, 5% and 1% respectively.

V. DATABASES

This section describes all the datasets used in this work. The TOC Biometrics company provided the images utilised in [3] for research purposes only. A dataset of 9,286 bona fide Chilean ID cards was used. This database was used as a baseline and the starting point to create composite, print and screen images. As mentioned in Section III a dataset of 10,000 transferable textures was also created. Additionally, we synthesised a significant number of ID card images for the purposes of data augmentation. These datasets are detailed in this section.

A. Database Organisation

For this work, a significant amount of bona fide images of Chilean ID cards were captured, which complies with ICAO standards. Then, we manually created presentation attack versions of those images for the composite, print and screen species. For the screen species, images of the bona fide ID cards were displayed on monitors, tablets and smartphones, and then images were re-captured using other Android and iPhone smartphones. The resulting images thus have the texture noise of the pixels and aliasing from the displayed devices. Additionally, for the print species, PDF documents containing eight bona fide ID cards images per page were printed in plain and glossy paper. Then, the individual ID cards were cut off with scissors and placed over various surfaces. Afterwards, each ID card was photographed using the same Android and iPhone smartphones used in the screen PAI. The resulting images contain artefacts from the type of paper, inks and the printing/scanning process. Finally, for the composite scenario, the face and fields of the print ID cards were cut and pasted over other print ID cards in order to create fake ID cards. This process was made manually with scissors, and automatically with digital splicing techniques.

Furthermore, as we mentioned before, the background of all captured images was removed using semantic

⁵<https://www.iso.org/standard/67381.html>

TABLE I
CAPTURED IMAGES OF CHILEAN ID CARDS

Class	Test	Train	Validation	Total
Bona fide	1,842	5,613	1,831	9,286
Composite	1,762	6,251	2,067	10,080
Print	2,322	6,893	2,299	11,514
Screen	1,538	4,542	1,517	7,597
Total	7,464	23,299	7,714	38,477

TABLE II
TRANSFERABLE TEXTURE DATASET

Class	Images	Textures
Bona fide	50	0
Print-A	2,500	2,500
Print-B	2,500	2,500
Screen	5,000	5,000
Total	10,050	10,000

segmentation. This background removal helps fraud-detection networks to focus on the ID card’s contents instead of its surroundings [3], [18].

Table I shows the final number of captured images for the bona fide, composite, print and screen classes, as well as test, train and validation partitions. For consistency, the test and validation partitions in this table will be used to evaluate all the fraud-detection networks in this work. We will refer to these partitions as *Chl-test* and *Chl-validation* for short.

In addition to capturing ID card images, we captured a dataset to help isolate the texture noise from the source material, as in the attack species print and screen. The dataset for texture study was captured over a palette of 50 solid colours with an identifying QR code on the left side, as shown in Figure 4. PDF pages containing 10 colours each were printed in several printers, both in plain and glossy paper. Then, several images of each individual colour were captured using smartphones of different brands at three times of the day (day, afternoon and night) at distances between 15 and 30cm and with portrait and landscape orientations. We captured 2,500 images printed on plain paper and 2,500 images on high-quality glossy paper. The later is meant to resemble the light reflection of the plastic material of authentic ID cards.

Similarly, the 50 colours in the palette were displayed on several monitors, tablets and smartphone screens. Images of those projected colours were captured using the same smartphones as above and with the same orientations and times of the day. In total 5,000 images were captured for the screen species. The device noise texture of the 10,000 captured images was isolated with the method described in Section III-C to obtain a dataset of 10,000 transferable textures. Both images and textures will be made freely available.⁶ The diversity of printers, screens and phones used to capture this dataset makes transferable textures a general-purpose method applicable to other fields. Figure 5 shows an example of texture application.

Table II displays the number of images and isolated textures in this dataset. This dataset was used for image processing algorithms. Therefore train, test and validation partitions were not needed.

TABLE III
TRAIN SET OF CAPTURED AND SYNTHETIC IMAGES.
SGAN2 REPRESENTS STYLEGAN2

Class	Chl-A	Chl-B	Chl-C	SGAN2	C-GAN	Templ.	Text.
Bona fide	2,807	2,806	2,806	3,000	0	3,104	0
Composite	3,126	3,125	0	3,000	0	3,104	0
Print	3,447	3,446	0	3,000	2,806	3,104	2,806
Screen	2,271	2,271	0	3,000	2,806	3,104	2,806

B. Training Set of Captured and Synthetic Images

Table III shows the quantity of images used for the training of fraud-detection networks. This table includes captured and synthetic images since a combination of both is used to train the networks. Synthetic images were generated using all the methods described in Section III. For a fair evaluation, all trained fraud-detection networks use the same *Chl-test* and *Chl-validation* images described in Table I. Examples of captured and synthetic images generated with the proposed methods are presented in Figure 6.

The training set of the Chilean ID cards was split into two halves, *Chl-A* and *Chl-B*, in order to perform the proposed experiments in Section VI. Also, a copy of bona fide *Chl-B* was created and named as *Chl-C*, as the base images for the CycleGAN and texture transfer methods.

For synthetic images, we generated around 3,000 images per class. This represents the average number of images per class in *Chl-A*.

The Template images were created from a dataset of 3,104 faces. The Templates method can create the bona fide and composite classes by itself; however, for the print and screen classes we translated the template images using the Textures method.

The process for StyleGAN2 was simple, since it can produce any desired number of images. Thus, exactly 3,000 images per class were generated. CycleGAN and Textures translate bona fide images into the attack instrument species print and screen. The base images were the 2,806 bona fide images from *Chl-C*; they were translated to create 2,806 print and 2,806 screen images using CycleGAN and Textures separately.

Examples of all the types of generated images can be seen in Figure 6. For captured and StyleGAN2 images, random examples of each class are shown. This is also true for composite images of the Templates method. However, for CycleGAN, Templates, and Textures, the base image, shown on the bona fide row, is translated to generate the print and screen PAIs in the rows below. As mentioned above, neither CycleGAN nor Textures can generate bona fide images, so the bona fide images presented in those two columns are in fact, bona fide images from *Chl-C*.

VI. EXPERIMENTS AND RESULTS

The experiments described in this section aim to compare the similarity score between captured and synthetic images, as well as evaluating the impact of using synthetic images while training fraud-detection networks. In this work, we assess two MobileNetV2 networks, one that detects composition tampering, and another that detects source alteration.

All evaluations are oriented to test the predictive value of adding synthetic images to the training set under the following

⁶https://github.com/jedota/Synthetic_ID-Card_Image



Fig. 6. Examples of captured and synthetic images generated with all the proposed methods. The rows and columns are colour coded. Tampered scenarios are in red and untampered in green. Captured images are in blue, whereas synthetic images are in purple. All sensitive information has been blacked out. Fields of StyleGAN2 and Templates are not covered since they contain random data of nonexistent subjects. N/A represent methods not used to generate the particular images.

scenario: suppose there are N captured images in the train set, which we would like to double. We are interested in studying if a combination of N captured plus N synthetic images would yield similar results to capturing $2 \times N$ images. This is the reason why the train set of captured Chilean ID cards is divided into two halves $Chl-A$ and $Chl-B$, as described in Table III. All experiments are explained in detail in this section.

A. Experiment 1 - Generation Evaluation

This experiment assesses the similarity between captured and synthetic images using the FID metric. For this evaluation, we compare each class (bona fide, composite, print and screen) of the test set of captured Chilean ID cards ($Chl-test$) against the corresponding class of synthetic images generated by the proposed methods (StyleGAN2, CycleGAN, Templates and Textures). This evaluation determines how similar captured and synthetic images are for each generation method. Additionally, we compute the FID score between $Chl-test$ and $Chl-validation$. This evaluation gives a baseline FID value for captured images under the same conditions.

Table IV shows the results of the FID scores. The smallest values in the table were obtained when comparing two sets of captured images: $Chl-test$ and $Chl-validation$. This comparison has an average FID of 5.45. This represents the best possible achievable performance for a generation method—the further the FID from 5.45, the less resemblance between synthetic and captured images.

StyleGAN2, CycleGAN, and Textures obtained similar FID values around 20. Among them, CycleGAN got the best score despite the success of StyleGAN2 for other image

TABLE IV
FID SCORES COMPUTED BETWEEN THE PROPOSED GENERATION METHODS AND CHL-TEST

Class	Chl-val	StyleGAN2	CycleGAN	Templates	Textures
Bona fide	4.54	20.33	-	36.49	-
Composite	6.61	20.79	-	36.07	-
Print	4.02	21.28	13.07	30.56	13.75
Screen	6.63	27.31	22.79	55.17	25.20
Average	5.45	22.43	17.93	39.57	19.48

domains, such as human faces [6]. We attribute the excellent performance of CycleGAN to the fact that it starts from a base image with a face, a signature and alphanumeric characters already on it. Therefore, this network must only replicate the device noise. For the same reason, Textures got a good score as well. On the other hand, StyleGAN2 has to generate a face, a signature, letters and numbers from scratch, along with perspective changes and device-generated noise.

The worst performance was obtained by the Templates method. We hypothesise that the main reason for this is an excessive range in the random projective transformation. From Figure 6, it can be observed that the Templates method had more extreme perspective changes concerning captured and the other synthetic images. Additionally, the random colour shift used could not represent the nature of the distribution of captured ID cards. Further studies are needed to optimise the random projective transformation and colour shift in order to reduce FID.

When comparing the different classes in Table IV, the screen attack was the hardest to replicate by all methods. However, the print scenario produced the smallest FID scores among the synthetic images. CycleGAN achieved an FID



Fig. 7. Empty template of the Argentinian ID card and a fake sample image generated with the Templates method.

of 13.07, while Textures obtained an FID of 13.75. Both values are close to the FID range among captured images. This result means that for the print attack with CycleGAN and Textures, the distribution of synthetic images is very similar to that of captured images.

Finally, to test the generalisation of the Templates method, we generated 2,069 Argentinian ID cards and compared the FID similarity against 2,683 real Bona fide Government-issued Argentinian ID cards. Figure 7 shows the empty template, as well as a synthetic ID card produced with that template. The resulting FID was 44.83, which is in the same range of the values observed in Table IV. This demonstrates that the Templates method can be easily applied to other countries with similar performance.

B. Experiment 2 - Composition Detection Network

A binary MobileNetV2 network detects whether an ID card image presents signs of composition or not. We trained and tested this network using only those methods capable of generating composite images, which are StyleGAN2 and Templates. In total, we trained four versions of this network using the following combinations of datasets from Table III.

- 1) Captured: *Chl-A* + *Chl-B* (11,864 images)
- 2) StyleGAN2: *Chl-A* + StyleGAN2 (11,933 images)
- 3) Templates: *Chl-A* + Templates (12,141 images)
- 4) Combined: *Chl-A* + StyleGAN2 + Templates (18,141 images)

Each training was designed to have a similar number of images in the train set, except for the last one, which evaluates the effect of adding images from all the available domains. The trained models were evaluated with the validation and test sets of Table I using PAD metrics for each trained model.

C. Experiment 3 - Source Detection Network

The second multi-class MobileNetV2 network classifies whether the ID card image came from a bona fide source or was tampered with by printing/scanning or screen-capturing PAIS. Therefore, the network classifies into three classes: bona fide, print and screen. In this case, we tested all the methods described in Section III, since they can all mimic source tampering. We trained the network several times with the following combinations of the train sets described in Table III, trying to utilise a similar number of images.

- 1) Captured: *Chl-A* + *Chl-B* (17,048 images)
- 2) StyleGAN2: *Chl-A* + StyleGAN2 (17,525 images)
- 3) Templates: *Chl-A* + Templates (17,837 images)
- 4) CycleGAN: *Chl-A* + *Chl-C* + CycleGAN (16,943 images)

TABLE V
RESULTS FOR THE COMPOSITE-TAMPERING DETECTION MODELS. SGAN2 REPRESENTS STYLEGAN2

	Captured	SGAN2	Templ	Combined
EER	2.33%	6.75%	6.03%	6.30%
BPCER ₁₀	0.05%	2.39%	1.36%	2.33%
BPCER ₂₀	0.43%	12.70%	7.82%	11.45%
BPCER ₁₀₀	6.13%	57.98%	43.43%	52.06%

TABLE VI
RESULTS FOR THE SOURCE-TAMPERING DETECTION MODELS

	Captured	SGAN2	C-GAN	Templ	Text	Comb
EER	1.82%	3.77%	2.93%	2.93%	6.50%	6.18%
BPCER ₁₀	0.38%	1.68%	0.65%	0.76%	3.42%	3.42%
BPCER ₂₀	0.49%	3.26%	1.57%	1.95%	8.25%	8.14%
BPCER ₁₀₀	3.53%	10.59%	6.24%	6.46%	23.51%	25.30%

- 5) Textures: *Chl-A* + *Chl-C* + Textures (16,943 images)
- 6) Combined: *Chl-A* + *Chl-C* + StyleGAN2 + CycleGAN + Templates + Textures (43,673 images)

All trained models use the validation and test sets of Table I for consistency. Finally, the PAD scores are also computed for the evaluation and comparison of each model.

D. PAD Performance

In this work, we trained 10 networks in total, as described in Section VI-B and Section VI-C. The DET curves for all the networks are presented in Figure 8. Additionally, the ISO/IEC 30107-3 error rates of the ten networks are presented in Table V and Table VI.

Training with captured images produces far superior performance for the composite model than mixing captured and synthetic images. The EER obtained was 2.33%, while that of all the synthetic models was approximately 6%, according to Figure 8a and Table V. Using Templates was marginally better than using StyleGAN2 or combining images from both methods. However, they are not suitable replacements for capturing more composite images.

For the source model, Table VI reports the scores for the worst PAI, which for all the 6 trained networks was screen. The DET curves for this PAI can be seen in Figure 8c. For this PAI, using captured images produced better performance than the synthetic models; however, the performance difference is not as pronounced as in the composite model. In this case, Captured image reached an EER of 1.82%, followed by CycleGAN and Templates with 2.93%. Therefore, performance drops only by 1% when the dataset is supplemented with synthetic images instead of doubling the amount of captured images. However, this is only true for CycleGAN and Templates since the other methods had EER values between 3.77% and 6.5%.

Although ISO/IEC 30107-3 standard indicates that only the worst-performing PAI has to be reported, we present in Figure 8b the DET curves of the print PAIS for a richer comparison. In this case, captured images have a very similar performance to CycleGAN of approximately 1.5%. This means that the labour-intensive process of printing, cutting with scissors and photographing the tampered ID cards with a smartphone can be replaced with images produced by

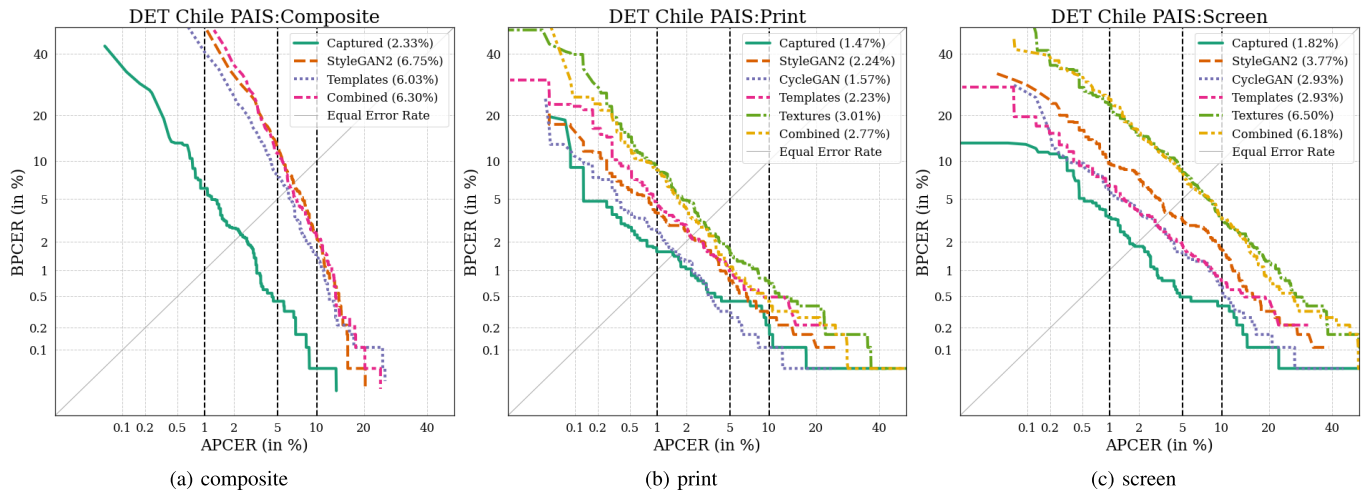


Fig. 8. Detection Error Trade-off curves for all the trained fraud-detection networks. Each plot corresponds to a different PAIS. The EER is shown in parenthesis for each scenario.

CycleGAN without any drop in performance. The EER of the other methods is low as well. StyleGAN2 and Templates have an EER of around 2.25%, while Textures and Combined were around 3%. These results correlate with the findings of the FID analysis in which the print PAIS obtained the best similarity between captured and synthetic images.

Table V and Table VI also present the $BPCER_{AP}$ values for three specific operating points. All those values follow the same trends observed with the EER.

A closer look at Figure 6 can reveal the subtle differences between all generation methods. StyleGAN2 could create the background, faces, signatures and composition tampering with no problems. However, it struggled with alphanumeric characters. For instance, despite proceeding from different random seeds, StyleGAN2 images of the four classes have the same first name: “Canora”. They also share the same middle name, surnames and RUN (National ID number); however, they present slight alterations that make them illegible. Reproducing legible alphanumeric characters is the most significant flaw of synthetic ID cards generated through StyleGAN2. On the other hand, Templates utilise handcrafted methods that place readable data, faces and signatures. However, the variability from image to image was handpicked and could not represent the actual distribution of captured ID cards.

The Textures method produced excellent reproductions of the print/scan textures and pixel noise and aliasing, as seen in the four images in the lower right corner. CycleGAN generated a subtle texture for the print PAI but a coarse square texture for the screen PAI. Lastly, StyleGAN2 produced subtle textures for both print and screen. Those are closer to bona fide images than their respective attack species.

VII. CONCLUSION

This work presented four different methodologies capable of generating synthetic ID cards, and evaluated the performance of each as a possible supplement for captured images. For this purpose, we trained two MobileNetV2 networks using different combinations of captured and synthetic images. Our results indicate that the composite PAIS was the hardest to

replicate. This attack scenario obtained the greatest EER and $BPCER_{AP}$ compared to other PAIS. The screen PAIS produced better PAD scores than composite. For CycleGAN and Templates, there is only 1% of performance reduction when supplementing the dataset with synthetic images instead of capturing more images. On the other hand, the print PAIS was the easiest to replicate according to the FID and PAD scores. In this case, CycleGAN achieved the same EER as using captured images only. This means supplementing with print images produced by CycleGAN would be a sound alternative to capturing more print images. Further, obtaining real ID cards from new people and manually simulating composite, print and screen attacks is very time-consuming and costly. Also, privacy regulations such as the GDPR assure individuals the right to withdraw their consent to use or store their private data, practically complicating using and distribution of large datasets. Therefore, achieving similar performance with synthetic data and fewer resources is valuable for future applications and extension to other countries and improving the fake-ID detection techniques.

The FID analysis demonstrated that for this application, CycleGAN and Textures produced images more similar to the captured ones than those produced by StyleGAN2. These results indicate that reproducing ID card images from scratch is very demanding and challenging, and it is more feasible to use domain adaptation from bona fide images. The images from Templates achieved the worst FID scores; however, they were still valuable for training fraud detection networks, achieving reasonable PAD scores.

Future work includes training CycleGAN using higher resolution images, improving the FID of Templates, and standardising the transferable textures for approved printers and scanners usable in Europe and USA.

DISCLAIMER

This work and the methods proposed are only for research purposes. Any implementation or commercial use modification must be analysed separately for each case to the email: juan.tapia-farias@h-da.de.

REFERENCES

- [1] Y. Shi and A. K. Jain, "DocFace: Matching ID document photos to selfies," in *Proc. IEEE 9th Int. Conf. Biometrics Theory, Appl. Syst. (BTAS)*, Oct. 2018, pp. 1–8.
- [2] Y. Shi and A. K. Jain, "DocFace+: ID document to selfie matching," *IEEE Trans. Biometrics, Behav., Identity Sci.*, vol. 1, no. 1, pp. 56–67, Jan. 2019.
- [3] S. Gonzalez, A. Valenzuela, and J. Tapia, "Hybrid two-stage architecture for tampering detection of chipless ID cards," *IEEE Trans. Biometrics, Behav., Identity Sci.*, vol. 3, no. 1, pp. 89–100, Jan. 2021.
- [4] R. Mudgalgundurao, P. Schuch, K. Raja, R. Ramachandra, and N. Damer, "Pixel-wise supervision for presentation attack detection on identity document cards," *IET Biometrics*, vol. 11, no. 5, pp. 383–395, Sep. 2022.
- [5] L. Zheng, Y. Zhang, and L. Vrizlynn, "A survey on image tampering and its detection in real-world photos," *J. Vis. Commun. Image Represent.*, vol. 58, pp. 380–399, Jan. 2019.
- [6] Y. Viazovetskiy, V. Ivashkin, and E. Kashin, "StyleGAN2 distillation for feed-forward image manipulation," in *Proc. Eur. Conf. Comput. Vis. Cham, Switzerland: Springer*, 2020, pp. 170–186.
- [7] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4510–4520.
- [8] M. Huber et al., "SYN-MAD 2022: Competition on face morphing attack detection based on privacy-aware synthetic training data," in *Proc. IEEE Int. Joint Conf. Biometrics (IJCB)*, Oct. 2022, pp. 1–10.
- [9] I. Goodfellow et al., "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, vol. 27, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, Eds. Red Hook, NY, USA: Curran Associates, 2014.
- [10] T. Karras, M. Aittala, J. Hellsten, S. Laine, J. Lehtinen, and T. Aila, "Training generative adversarial networks with limited data," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 12104–12114.
- [11] C. Chu, A. Zhmoginov, and M. Sandler, "CycleGAN, a master of steganography," 2017, *arXiv:1712.02950*.
- [12] V. Albiero et al., "Identity document to selfie face matching across adolescence," in *Proc. IEEE Int. Joint Conf. Biometrics (IJCB)*, Sep. 2020, pp. 1–9.
- [13] M. Stokkenes, R. Ramachandra, and C. Busch, "Biometric transaction authentication using smartphones," in *Proc. Int. Conf. Biometrics Special Interest Group (BIOSIG)*, Sep. 2018, pp. 1–5.
- [14] P. Perera and V. M. Patel, "Face-based multiple user active authentication on mobile devices," *IEEE Trans. Inf. Forensics Security*, vol. 14, no. 5, pp. 1240–1250, May 2019.
- [15] M. E. Fathy, V. M. Patel, and R. Chellappa, "Face-based active authentication on mobile devices," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Apr. 2015, pp. 1687–1691.
- [16] V. V. Arlazarov, K. B. Bulatov, and T. S. Chernov, "MIDV-500: A dataset for identity documents analysis and recognition on mobile devices in video stream," 2018, *arXiv:1807.05786*.
- [17] X. Zhu et al., "Large-scale bisample learning on ID versus spot face recognition," *Int. J. Comput. Vis.*, vol. 127, nos. 6–7, pp. 684–700, Jun. 2019.
- [18] R. Lara, A. Valenzuela, D. Schulz, J. Tapia, and C. Busch, "Towards an efficient semantic segmentation method of ID cards for verification systems," 2021, *arXiv:2111.12764*.
- [19] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 248–255.
- [20] T. Salimans, I. J. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training GANs," 2016, *arXiv:1606.03498*.
- [21] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local Nash equilibrium," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 1–10.
- [22] Q. Xu et al., "An empirical study on evaluation metrics of generative adversarial networks," 2018, *arXiv:1806.07755*.
- [23] S. Pouyanfar, S.-C. Chen, and M.-L. Shyu, "An efficient deep residual-inception network for multimedia classification," in *Proc. IEEE Int. Conf. Multimedia Expo (ICME)*, Jul. 2017, pp. 373–378.



Daniel Benalcazar (Member, IEEE) was born in Quito, Ecuador, in 1987. He received the B.S. degree in electronics and control engineering from Escuela Politecnica Nacional, Quito, in 2012, the M.S. degree in electrical engineering from The University of Queensland, Australia, in 2014, with a minor in biomedical engineering, and the Ph.D. degree in electrical engineering from Universidad de Chile, Chile, in 2020. From 2015 to 2016, he worked as a Professor with the Central University of Ecuador. Ever since, he has participated in research projects in biomedical engineering and biometrics. He is currently a Researcher with TOC Biometrics, Chile.



Juan E. Tapia (Member, IEEE) received the P.E. degree in electronics engineering from Universidad Mayor in 2004, the M.S. degree in electrical engineering from Universidad de Chile in 2012, and the Ph.D. degree from the Department of Electrical Engineering, Universidad de Chile in 2016. In addition, he spent one year of internship with the University of Notre Dame. From 2016 to 2017, he was an Assistant Professor with Universidad Andres Bello. From 2018 to 2020, he was the Research and Development Director of the Electricity and Electronics with Universidad Tecnologica de Chile and also the Research and Development Director of TOC Biometrics from 2019 to 2022. He is currently a Senior Researcher with Hochschule Darmstadt (HDA). His main research interests include pattern recognition and deep learning applied to iris biometrics, morphing attack detection, and feature selection. In 2016, he received the Award for Best Ph.D. Thesis.



Sebastian Gonzalez (Member, IEEE) received the B.Sc. degree in computer science from Universidad Andres Bello in 2019. He is currently pursuing the M.Sc. degree in computer science from Universidad de Santiago de Chile. He is currently a Researcher with TOC Biometrics, where he has been involved in different research projects encompassing computer vision, pattern recognition, and deep learning applied to biometrics. His main research interests include presentation attack detection, classification, segmentation, and applied research.



Christoph Busch (Senior Member, IEEE) is a member of the Department of Information Security and Communication Technology (IIK), Norwegian University of Science and Technology (NTNU), Norway. He holds a joint appointment with the computer science faculty with Hochschule Darmstadt (HDA), Germany. Since 2007, he has been a Lecturer of biometric systems with DTU, Denmark. On behalf of the German BSI, he has been the coordinator for the project series BioIS, BioFace, BioFinger, BioKeyS Pilot-DB, KBEinweg, and NFIQ2.0. In the European research program, he was an Initiator of the Integrated Project 3-D Face, FIDELITY, and iMARS. Further, he was/is partner in the projects TURBINE, BEST Network, ORIGINS, INGRESS, PIDaaS, SOTAMD, RESPECT, and TReSPaaS. He is also a Principal Investigator with the German National Research Center for Applied Cybersecurity (ATHENE). He has coauthored more than 600 technical articles and a speaker at international conferences. Moreover, he is a Co-Founder and a member of the board of the European Association for Biometrics (www.eab.org) which was established in 2011 and assembles in the meantime more than 200 institutional members. He is a member of the editorial board of the IET journal.