SpoofGAN: Synthetic Fingerprint Spoof Images

Steven A. Grosz and Anil K. Jain, *Life Fellow, IEEE*

***Abstract*—A major limitation to advances in fingerprint spoof detection is the lack of publicly available, large-scale fingerprint spoof datasets, a problem which has been compounded by increased concerns surrounding privacy and security of biometric data. Furthermore, most state-of-the-art spoof detection algorithms rely on deep networks which perform best in the presence of a large amount of training data. This work aims to demonstrate the utility of synthetic (both live and spoof) fingerprints in supplying these algorithms with sufficient data to improve the performance of fingerprint spoof detection algorithms beyond the capabilities when training on a limited amount of publicly available "real" datasets. First, we provide details of our approach in modifying a state-of-the-art generative architecture to synthesize high quality live and spoof fingerprints. Then, we provide quantitative and qualitative analysis to verify the quality of our synthetic fingerprints in mimicking the distribution of real data samples. We showcase the utility of our synthetic live and spoof fingerprints in training a deep network for fingerprint spoof detection, which dramatically boosts the performance across three different evaluation datasets compared to an identical model trained on real data alone. Finally, we demonstrate that only 25% of the original (real) dataset is required to obtain similar detection performance when augmenting the training dataset with synthetic data.**

***Index Terms*—Generative Adversarial Networks, Fingerprint Presentation Attack Detection, Fingerprint Spoof Detection, Synthetic Fingerprint Generation**

I. Introduction

FINGERPRINT recognition has had a long history in person identification due to the purported uniqueness and permanence of fingerprints, originally pointed out by Sir Francis Galton in his 1982 book titled "Finger Prints" [1] and reaffirmed in many works over the last century; including the well-known studies on the individuality and longitudinal permanence of fingerprint recognition [2], [3]. Clearly, a significant contributor to their wide spread adoption is the high level of verification performance achieved by state-of-the-art (SOTA) algorithms for automated fingerprint recognition[1]. However, despite the impressive accuracy achieved to date by the top performing fingerprint recognition algorithms, there remains many on-going efforts to further improve the capabilities of fingerprint recognition systems - especially in terms of recognition speed and system security. As a result, there has been a recent push toward deep neural network (DNN) based models for fingerprint recognition [5], [6], [7], [8], [9], [10], [11]. These compact, fixed-length embeddings can be matched efficiently and combined with homomorphic encryption for added security [12]. For a more exhaustive account of existing deep learning approaches to fingerprint recognition and other biometric modalites, interested readers are encouraged to consult one of the many surveys on deep learning in biometrics (e.g., [13], [14]).

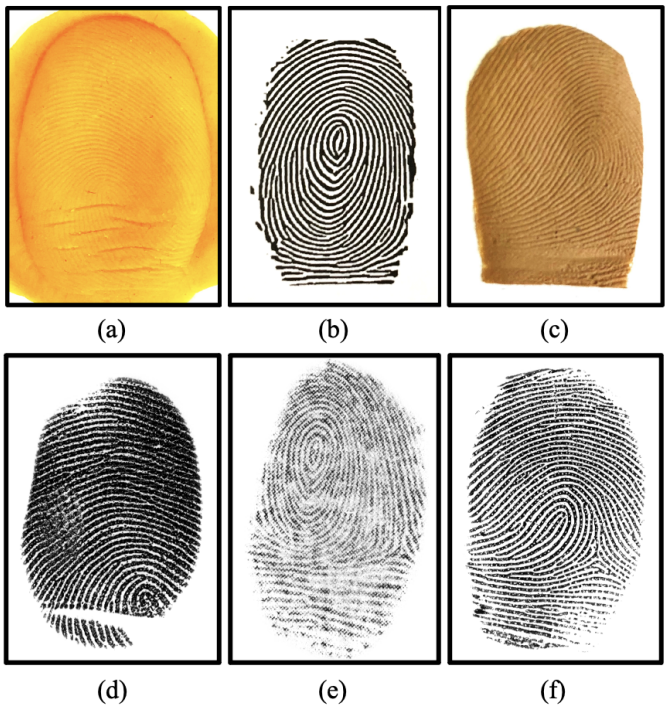

Fig. 1: Example fabricated fingerprint spoofs of various materials and corresponding fingerprint impressions captured on a CrossMatch Guardian200 fingerprint reader. (a) PlayDoh spoof, (b) printed paper spoof, (c) latex spoof, (d) fingerprint impression from the PlayDoh spoof, (e) fingerprint impression from the printed paper spoof, and (f) fingerprint impression from the latex spoof.

Indeed, this push toward DNN-based fingerprint recognition comes in the wake of the success demonstrated in the face recognition domain in applying DNN models to face recognition, which was aided by the availability of large-scale face recognition databases which were easily crawled from the web; despite the many ethical and privacy concerns which have led to many of these datasets to be recalled today. Arguably, at least in part, the reason for the delayed adoption of DNNs for fingerprint recognition has been the lack of publicly available, large-scale fingerprint recognition datasets and increased scrutiny over privacy of biometric data, which has led to many works to generate synthetic fingerprint images [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28].

Similarly, there has been an increased interest in DNN-based

S.A. Grosz and A.K. Jain are with the Department of Computer Science and Engineering, Michigan State University, East Lansing, MI, 48824 USA (e-mail: groszste@cse.msu.edu, jain@cse.msu.edu).

[1]Today's top performing algorithm in the FVCongoing 1:1 hard benchmark achieved a False Non-Match Rate (FNMR) of just 0.626% at a False Match Rate (FMR) of 0.01% [4]

models for fingerprint spoof, i.e., presentation attack (PA), detection; where the scale and amount of publicly available data is also limited. Table I gives a list of the publicly available fingerprint spoof datasets. Compared to the largest, public fingerprint recognition dataset, NIST Special Database 302 [29], which contains fingerprints from 2,000 unique fingers, the largest publicly available fingerprint spoof datasets, e.g., the LivDet competition datasets, contain at most 1,000 unique fingers (for the Swipe sensor in LivDet 2013 [30])[2]. Compounding the problem is the difficulty in collecting large-scale fingerprint spoof datasets due to the increased time and complexity in fabricating and collecting spoof fingerprints. All of which motivates the potential of synthetic data as a viable alternative; however, to the best of our knowledge, there does not exist a synthetic fingerprint spoof generator to fill the gap between the amount of publicly available fingerprint spoof data and training of data-hungry deep learning based models.

To address the lack of large-scale fingerprint spoof datasets, we propose SpoofGAN. Inspired by the impressive results of the recently proposed PrintsGAN [15], SpoofGAN is a multi-stage generative architecture to fingerprint generation. SpoofGAN is different from PrintsGAN in the following ways:

- Generation of plain print fingerprints, which compared to the rolled prints generated by PrintsGAN, are more representative of the publicly available spoof fingerprint datasets and exhibit different textural characteristics, distortions, etc.
- Ability to synthesize both live (i.e., bona fide) and spoof samples of the same fingerprint identity.
- Replacing the learned warping and cropping module with a statistical, controllable non-linear deformation model to synthesize multiple, realistic impressions per finger. This allows us to control the degree of distortion applied.

In this work, we use the terms *live* and *bona fide* interchangeably; however, live is not synonymous with real. We make the distinction that real fingerprint images are those captured on a fingerprint reader by either a real human finger or physical spoof artifact, whereas synthetic fingerprints are digital renderings of fingerprint like images. Thus, there can be both synthetic live fingerprint images and synthetic spoof fingerprint images that are generated, just as there are real live fingerprint images and real spoof fingerprint images.

We validate the realism of our synthetic live and spoof images through extensive qualitative and quantitative metrics including NFIQ2 [31], minutiae statistics, match scores from a SOTA fingerprint matcher, and T-SNE feature space analysis showing the similarity of real live and spoof embeddings to the embeddings of our synthetic live and spoof fingerprints. Besides verifying the realism of our synthetic spoof generator, we also show how SpoofGAN fingerprints can be used to train a DNN for fingerprint spoof detection. We show this by improving the performance of a spoof detection model by augmenting an existing fingerprint spoof dataset with additional samples from our synthetic generator. We also open the door to jointly optimizing for fingerprint spoof detection and recognition in an end-to-end learning framework with our ability to generate a large-scale dataset of multiple impressions per finger of both live and spoof examples.

More concisely, the contributions of this research are as follows:

- A highly realistic plain print synthetic fingerprint generator capable of generating multiple impressions per finger.
- The first, to the best of our knowledge, synthetic fingerprint spoof generator which is capable of producing both live and spoof impressions of the same identity. This opens the door to joint optimization of fingerprint spoof detection and recognition algorithms.
- Quantitative and qualitative analysis to verify the quality of our generated live and spoof fingerprints.
- Experiments showcasing improved fingerprint spoof detection when augmenting existing fingerprint datasets with our synthetic live and spoof fingerprints.

II. Related Work

A. Fingerprint Spoof Detection

One significant risk to the security of fingerprint recognition systems is that of presentation attacks, defined by ISO standard IEC 30107-1:2016(E) as a "presentation to the biometric data capture subsystem with the goal of interfering with the operation of the biometric system" [40]. The most common type of presentation attacks are spoof attacks, i.e., physical representations of finger-like structures aimed at either mimicking the fingerprint ridge-valley structure of another individual or subverting the user's own identity. Spoof attacks may come in many different forms and materials such as those shown in Figure 1.

Several hardware and software-based solutions to detecting spoof attacks have been proposed. Hardware based solutions include specialized sensors that leverage various "liveness" cues at the time of acquisition, such as conductivity of the material/finger, sub-dermal imaging, and multi-spectral lighting [41], [42], [43], [44], [45], [46]. On the other hand, software based solutions typically rely on only the information captured in the grayscale image acquired by the fingerprint reader [47], [48], [49], [50], [51], [52], [38]. Despite the limited publicly available fingerprint spoof data, many of the state-of-the-art software-based solutions to fingerprint spoof detection leverage convolutional neural networks to learn the decision boundary between live and spoof images. Some researchers have proposed training their algorithms on smaller patches of the fingerprint images as a way to deal with limited amounts of available training data, which roughly increases the number of training images by a factor proportional to the number of patches [38]. However, given the increased scrutiny over privacy concerns related to biometric datasets, it is not certain whether any spoof fingerprint datasets will remain available in the future. Thus, motivating the need for synthetic data to fill this gap.

Another challenge related to limited training data is that of unseen presentation attacks, or fingerprint images arising from never before seen PA instruments. This problem is also commonly referred to in the literature as cross-material generalization. Some strategies proposed to improve the cross-material

[2]Swipe sensors are no longer in vogue after Apple introduced the "area capacitive sensor" in Touch ID.

TABLE I: Publicly available fingerprint spoof datasets.

Name	# Train Images Live (Spoof)	# Test Images Live (Spoof)	Spoof types	Sensors
LivDet 2009 [32]	1000 (1000) 1000 (1000)	1000 (1000) 1000 (1000)	Ecoflex, Gelatine, Latex, Modasil, WoodGlue	Biometrika Italdata
LivDet 2011 [33]	1000 (1000) 1000 (1000) 1000 (1000) 1000 (1000)	1000 (1000) 1000 (1000) 1000 (1000) 1000 (1000)	Gelatine, latex, PlayDoh, Silicone, Wood Glue, Ecoflex	Biometrika DigitalPersona ItalData Sagem
LivDet 2013 [30]	1000 (1000) 1000 (1000) 1250 (1000) 1250 (1000)	1000 (1000) 1000 (1000) 1250 (1000) 1250 (1000)	Body Double, Latex, PlayDoh, Wood Glue, Gelatine, Ecoflex, Modasil	Biometrika ItalData CrossMatch Swipe
LivDet 2015 [34]	1000 (1000) 1000 (1000) 1000 (1000) 1510 (1473)	1000 (1500) 1000 (1500) 1000 (1500) 1500 (1448)	Ecoflex, Gelatine, Latex, Liquid Ecoflex, RTV, WoodGlue, Body Double, PlayDoh, OOMOO	Biometrika DigitalPersona GreenBit CrossMatch
LivDet 2017 [35]	1000 (1200) 1000 (1200) 999 (1199)	1700 (2040) 1700 (2676) 1700 (2028)	Body Double, Ecoflex, Wood Glue, Gelatine, Latex, Liquid Ecoflex	GreenBit Orcanthus DigitalPersona
LivDet 2019 [36]	1000 (1200) 1000 (1200) 1000 (1000)	1020 (1224) 990 (1088) 1019 (1224)	Body Double, Ecoflex, Wood Glue, Gelatine, Latex, Liquid Ecoflex	GreenBit Orcanthus DigitalPersona
LivDet 2021 [37]	1250 (1500) 1250 (1500)	2050 (2460) 2050 (2460)	Latex, RProFast, Nex Mix 1, Body Double, Elmer's Glue, GLS20, RFast30	GreenBit Dermalog
MSU FPAD [38]	2250 (3000) 2250 (2250)	2250 (3000) 2250 (2250)	Ecoflex, PlayDoh, 2D Matte Paper, 2D Transparency	CrossMatch
MSU FPAD v2 [39]	4743 (4912)	1000 (leave-one-out)	2D Printed Paper, 3D Universal Targets, Conductive Ink on Paper, Dragon Skin, Gelatine, Gold Fingers, Latex Body Paint, Monster Liquid Latex, Play Doh, Silicone, Transparency, Wood Glue	CrossMatch

[1] The dataset release agreement for all LivDet databases can be found at https://livdet.org/registration.php.
[2] Similarly, the dataset release form for the MSU FPAD dataset can be found at http://biometrics.cse.msu.edu/Publications/Databases/MSU_FPAD/

performance of spoof detection include learning a tighter boundary around the live class via one-class classifiers [53], [54], incorporating adversarial representational learning to encourage robustness to varying material types [55], [56], or applying style transfer to mix textures from some known PA materials to better fill the space of unknown texture characteristics that may be encountered [57], [58]. Similar ideas may apply to synthetic data generation, where new material types can be synthesized by mixing characteristics of known PAs.

III. Proposed Synthetic Spoof Fingerprint Generator

In this section, we detail the process of generating synthetic live and spoof fingerprints. Motivated by the success of previous multi-stage fingerprint generation methods (e.g., [15], [25], [59]), SpoofGAN generates highly realistic fingerprints in multiple stages. First, unique fingerprint identities are synthesized through generating binary master fingerprints which define the fingerprint ridge structure. Following the master print synthesis stage, perturbations such as random rotation, translation, and non-linear deformation are applied to simulate realistic, repeat impressions. Finally, each generated fingerprint impression is input to a second neural network to impart realistic textures which mimic a database of real fingerprints. An overview of the entire process is given in Figure 2.

A. Master Print Synthesis

The first step in generating synthetic fingerprints with SpoofGAN is generating binary master fingerprints $I_{id} \in \{0,1\}^{256\times256}$ from a random vector $z_{id} \in \mathcal{R}^{256}$ sampled from a standard normal distribution (i.e., $z_{id} \sim \mathcal{N}(0,1)$). In particular, we used a standard BigGAN [60] architecture for this task, consisting of a generator G_{id} and a discriminator D_{id}. Since many spoof impressions can exhibit non-realistic fingerprint ridge structures, either from artifacts introduced in the fabrication (e.g., bubbles in the ridges) or in the presentation process (e.g., smudges due to the high elasticity of some spoof material types), we chose to train G_{id} using a database of only live fingerprint impressions consisting of 38,164 images captured on a CrossMatch Guardian200 fingerprint reader. As we will show later, these artifacts for the spoof impressions can be introduced in the later texture rendering stage of our synthesis pipeline. The network is trained via an adversarial loss shown in Equation 1, where I is a binary fingerprint image extracted from a real fingerprint image using the Verifinger v12.0 SDK.

$$\mathcal{L}_{adv}(G,D) = \mathbb{E}_I\ [logD(I)] + \mathbb{E}_z\ [log(1 - D(G(z)))] \quad (1)$$

B. Generating Multiple Realistic Impressions

To generate multiple impressions from a single master print, we apply realistic rotation, translation, and non-linear deformation for each subsequent impression. Rotations are applied via a uniform random sampling in the range [-30°, 30°], whereas translations in both the x and y directions are uniformly sampled in the range [-25, 25] pixels. Finally, realistic non-linear deformations are applied via a learned, statistical deformation model proposed by Si et al. in [61]. The distortion parameters were learned from a database of 320

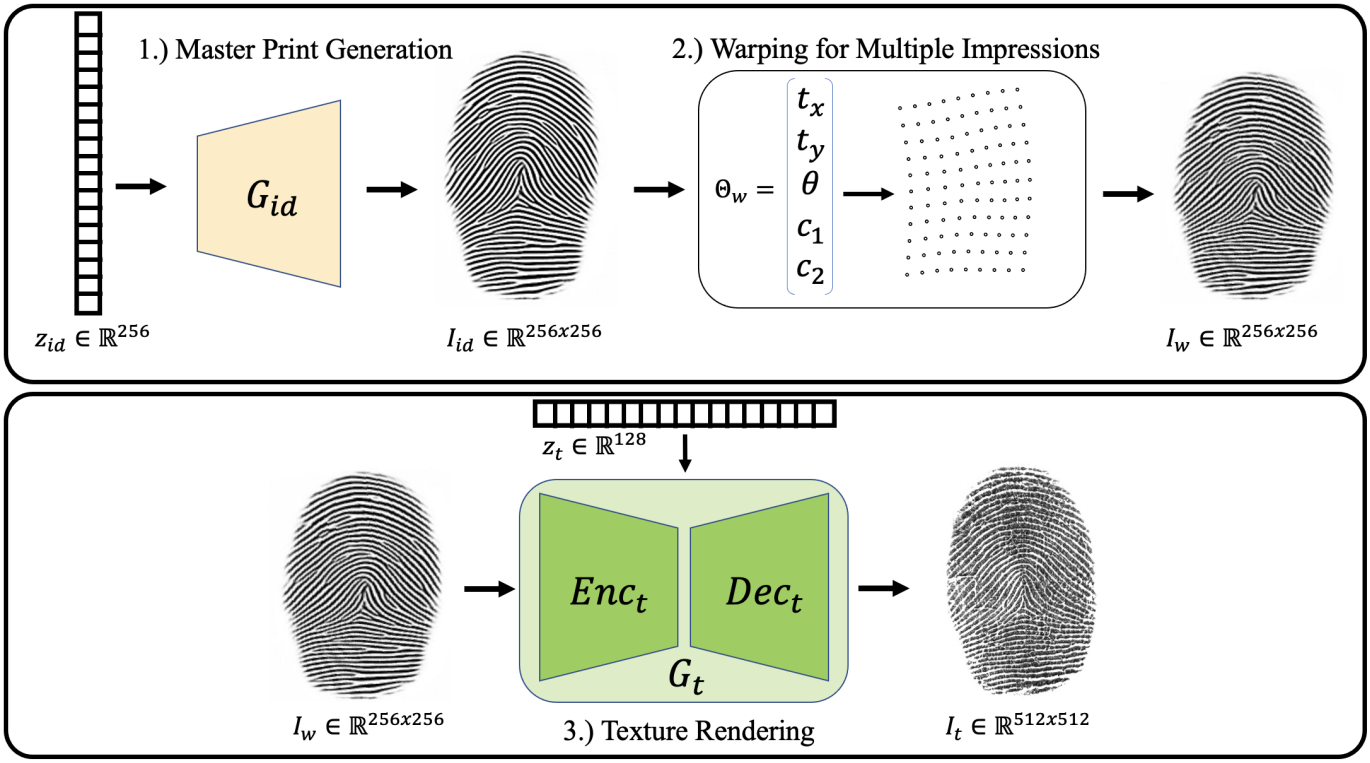

Fig. 2: Overview of the SpoofGAN Architecture.

distorted fingerprint videos in which the minutiae locations in the first and last frames were manually labeled and the displacements between corresponding minutiae points were used to estimate a distortion field via a Thin-Plate-Spline deformation model [62]. The distortion fields were condensed into a subset of eigenvectors, e_i, computed from a Principal Component Analysis of the covariance matrix estimated from the 320 videos. By varying the coefficients, c_i, multiplied to each of the eigenvectors, we vary the magnitude of distortion, d, applied to an input fingerprint along realistic distortion directions according to the equation2, where λ_i are the eigenvalues of each eigenvector and $\bar{d}$ is the average distortion field. For our implementation, we randomly sample the coefficients of the two largest eigenvectors from a normal distribution with mean 0 and standard deviation of 0.66, which were empirically determined to produce reasonable distortions.

$$d \approx \bar{d} \sum_{i=1}^{t} c_i \sqrt{\lambda_i e_i} \tag{2}$$

C. Texture Rendering

The final stage of our fingerprint generation process consists of imparting each fingerprint with a realistic texture that mimics the distribution of real live and spoof images. For the generator, G_t, we use an encoder-decoder architecture which translates an input, warped binary image, I_w, into a realistic fingerprint impression, I_t. To promote diversity in the rendered images, a random texture vector sampled from a standard normal distribution (i.e., $z_t \sim \mathcal{N}(0,1)$) is injected into the network and encoded into γ and β parameters for performing instance normalization on the intermediate feature maps of G_t. Finally, the discriminator, D_t utilizes the same architecture used in the binary master print synthesis network.

The goal of our texture renderer is two fold: i.) generate realistic texture details and ii.) maintain the identity of the rendered fingerprint between corresponding impressions of the same finger. Thus, we introduce two losses in addition to the conventional GAN loss (eq. 3) to maintain the identity of textured fingerprints. The first is an identity loss to minimize the L_2 distance between feature embeddings of corresponding fingerprint impressions using a SOTA fingerprint matcher DeepPrint [9] (eq. 4) and the other is an L_2 pixel loss between ground truth binary images and binary images extracted from the textured fingerprints (eg. 5). The L_2 pixel loss is computed on the binary images, rather than the grayscale images, to allow for the network to generate diverse "styles" in the generated fingerprints to simulate different pressure, moisture content, and contrast in subsequent impressions; all of which would lead to slightly different loss values compared to the ground truth image unless first converted to binary ridge images. To make the process of binarization of the generated fingerprints differentiable, we train a convolutional autoencoder to binarize input fingerprints which is trained on 38,164 grayscale/binary image pairs. The overall losses for G_t and D_t are given in equations 6 and 7, respectively.

1) GAN loss: Classical min-max GAN loss between the discriminator, $D_t(\cdot)$, trying to classify each original fingerprint image, I, as real and each synthetic fingerprint $I_t = G_t(I_w)$ as fake. Meanwhile, $G_t(\cdot)$ is trying to fool $D_t(\cdot)$ into thinking its outputs come from the original

image distribution.

$$\mathcal{L}_{adv} = \mathbb{E}_I\ [logD(I)] + \mathbb{E}_{I_w}\ [log(1 - D(G(I_w)))] \quad (3)$$

2) DeepPrint loss: L_2 distance between the DeepPrint embedding, R, extracted from the ground truth grayscale image and the DeepPrint embedding, $\hat{R}$, extracted from the synthesized grayscale fingerprint image.

$$L_{dp} = \frac{1}{2}\sum(R - \hat{R})^2 \quad (4)$$

3) Image/pixel loss: L_2 loss between the ground truth binary fingerprint image, I_w, and synthesized binary fingerprint image, $\hat{I_w}$.

$$L_i = \frac{1}{2}\sum_{x,y}(I_w(x,y) - \hat{I_w}(x,y))^2 \quad (5)$$

4) Overall loss for $G_t(\cdot)$: $\lambda_1 = 1$, $\lambda_2 = 2$, and $\lambda_3 = 10$ (determined empirically).

$$\mathcal{L}_{G_t} = \lambda_1\mathcal{L}_{adv} + \lambda_2\mathcal{L}_{dp} + \lambda_3\mathcal{L}_i \quad (6)$$

5) Overall loss for $D_t(\cdot)$:

$$\mathcal{L}_{D_t} = \mathcal{L}_{adv} \quad (7)$$

Unlike the binary master print synthesis and warping stages, an individual texture rendering network is trained for each material type (live, ecoflex spoof, PlayDoh spoof, etc.). Due to the limited number of images in our spoof dataset, we pretrained a texture rendering network on the 282K unique fingerprint database taken from the MSP longitudinal database introduced in [3][3]. Initially, following the pretraining, two texture rendering networks are trained further, one on the dataset of 38,164 live only impressions and the other on the 3,366 spoof fingerprint images consisting of all spoof types aggregated together. Finally, we further finetune the model trained on all spoofs for each of the individual spoof types to give more fine-grained control on the specific spoof style being generated.

Due to the limited number of images in our spoof dataset, we pretrained a texture rendering network on the 282K unique fingerprint database taken from the MSP longitudinal database introduced in [3][4]. Following the pretraining, two texture rendering networks are finetuned; one on the dataset of 38,164 live only impressions and the other on the 3,366 spoof fingerprint images consisting of all spoof types aggregated together. Finally, we further finetune the model trained on all spoofs for each of the individual spoof types to give more fine-grained control on the specific spoof style being generated. Thus, unlike the binary master print synthesis and warping stages which are shared, each individual spoof type has it's own rendering network. Alternatively, a conditional GAN structure could be used to generate spoof classes of each type within a single network; however, we found that due to the very limited amount of training images in some spoof types (e.g., 50 images), finetuning for just a few epochs on each spoof type individually produced higher quality images.

[3]This database is not publicly available, but the pretrained model can be made available upon request.

[4]This database is not publicly available, but the pretrained model can be made available upon request.

IV. Experimental Results

In this section, we aim to validate the realism of our synthetic live and spoof images via several qualitative and quantitative experiments. First, we provide details on the datasets involved in the following experiments, followed by some example fingerprint images generated by SpoofGAN to qualitatively compare with real fingerprint images. Finally, several quantitative metrics are used to compare the utility and distribution of SpoofGAN generated fingerprint images compared to the real fingerprint images.

A. Datasets

A main motivation for this paper is the lack of large-scale, publicly available fingerprint spoof datasets. Some of the largest datasets that are available have resulted from the annual LivDet competition series dating as far back as 2009 [33], [30], [34], [35], [36], [37]. A more comprehensive list of the fingerprint spoof datasets currently available to the research community is given in Table I, whereas the datasets used in this paper are given in Table II. In this paper we focus our experiments on fingerprint images obtained via the CrossMatch optical reader from LivDet 2013, LivDet 2015, and the Government Controlled Test (GCT) dataset of live and spoof fingerprints collected as part of the IARPA ODIN program[5]. Our training dataset for SpoofGAN, referred to as GCT 1-5, consists of 38,164 live fingerprint images and 3,366 spoof fingerprint images from 2,007 fingers and 11 different spoof types (Dragon Skin, Ecoflex, Paper, Silicone, Transparency, Gelatine, Glue, PDMS, Knox Gelatine, Gummy Overlay, and Tattoo). For our evaluations, we have followed the same train/test protocol referenced in LivDet2015 and LivDet2013, as well as reserved a fraction of the GCT dataset, GCT 6, as an evaluation dataset.

B. Qualitative Analysis of Synthetic Live and Spoof Images

Example synthetic live and spoof fingerprints of varying material types (i.e., presentation Attack Instruments) are shown in column (b) of Figure 3 with corresponding real examples in column (a) shown for reference. Visually, SpoofGAN is generating spoof examples that closely resemble the target material type, which can be seen in the different texture characteristics for each material seen in the synthetic examples (e.g., Paper vs. Tattoo, etc.). Additionally, looking across the rows of the synthetic SpoofGAN images, we notice that the underlying identity is successfully being preserved across the different spoof styles being generated. Figure 4 further highlights SpoofGAN's ability to successfully generate multiple impressions of each finger in both live and spoof styles.

C. Quantitative Analysis of Synthetic Live and Spoof Images

Similar to previous works on synthetic fingerprint generation [17], [15], we have evaluated the quality of SpoofGAN fingerprints on several quantitative metrics, including spoof

[5]We selected CrossMatch for our experiments since it is one of the most popular slap (4-4-2) capture readers used in law enforcement, homeland security and civil registry applications.

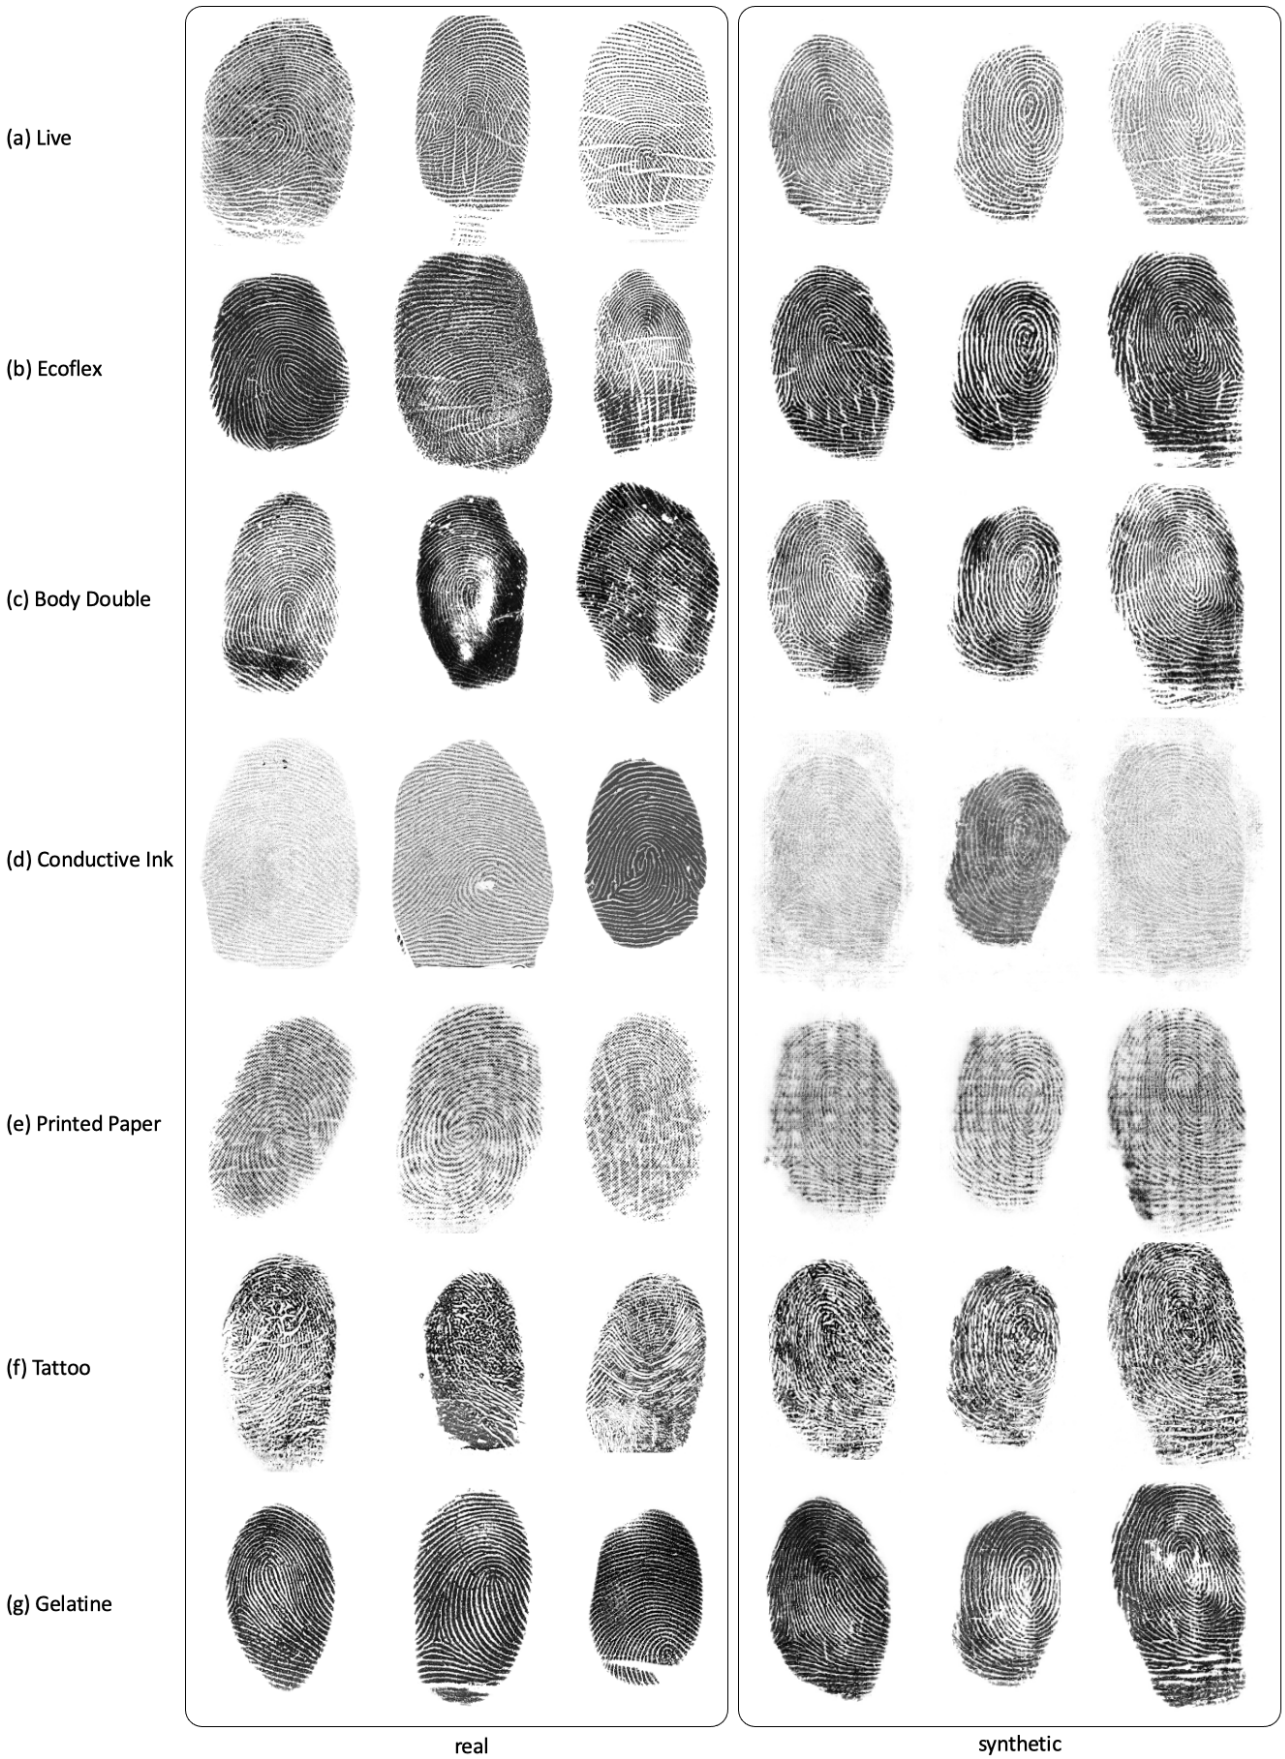

Fig. 3: Example real live and spoof images and synthetic live and spoof images generated by SpoofGAN of various material types: (a) live, (b) ecoflex, (c) body double, (d) conductive ink on paper, (e) 2d printed paper, (f) tattoo, and (g) gelatine.

TABLE II: Summary of the spoof datasets used in our experiments.

Dataset	LivDet 2013	LivDet 2015	GCT 1-5	GCT 6
Fingerprint Reader	CrossMatch	CrossMatch	CrossMatch	CrossMatch
Model	L Scan Guardian	L Scan Guardian	Guardian200	Guardian200
Resolution (dpi)	500	500	500	500
# Live Images (Train/Test)	1250 / 1250	1510 / 1500	38,164 / 0	7,357 / 14,236
# PA Images (Train/Test)	500 / 440	1473 / 1448	3,366 / 0	2,550 / 1,829
PA Materials	Body Double, Latex, PlayDoh, Wood Glue	Ecoflex, Gelatine, Body Double, PlayDoh, OOMOO	Dragon Skin, Ecoflex, Paper, Silicone, Transparency, Gelatine, Glue, PDMS, Knox Gelatine, Gummy Overlay, Tattoo	Ecoflex, Silicone, Gummy Overlay, Tattoo, Knox Gelatine

Fig. 4: Example images of multiple impressions of the same finger generated by SpoofGAN. (a) and (b) show three impressions each of two fingers rendered in a live style, whereas (c) and (d) show three impressions each of the same two fingers in a spoof style (ecoflex and body double, respectively).

detection performance by a pretrained spoof detector, the distribution of minutiae count, type, and quality extracted from SpoofGAN generated fingerprints compared to a real fingerprint dataset, NFIQ2 quality scores, match score distributions from a SOTA fingeprrint matcher, and identity leakage experiments.

1) Spoof Detection Performance of Real vs. Synthetic Fingerprints: Our first evaluation to verify the quality of our synthetic live and spoof fingerprints is to see whether a pretrained spoof detection algorithm trained on similar, real fingerprints performs equally well on our synthetic fingerprints. In particular, we pretrained an Inception v3 network on the GCT 1-5 data to classify between live and spoof fingerprint samples. Then, we evaluated the spoof detection performance on LivDet 2015 CrossMatch images compared to an equivalent sized database of synthetic fingerprints. As shown in Table III, the true detection rate (TDR) at a false detection rate (FDR) of 0.2% is similar across multiple spoof types of both datasets, supporting our hypothesis that the synthetic samples should be useful in training additional spoof detection models without access to a large database of real live and spoof fingerprints for training[6]. Lastly, the embeddings of both real and synthetic images in the T-SNE embedding space suggest high similarity between the embeddings of real and synthetic images (see Figure 5).

2) Feature Similarity Between Real and Synthetic Fingerprints: For synthetic fingerprint images to be useful as a substitute for real fingerprint images, the features between a database of real fingerprint images and synthetic fingerprint

[6]This metric is recommended by the IARPA ODIN program, one of the largest US government funded programs on spoof detection: https://www.iarpa.gov/research-programs/odin.

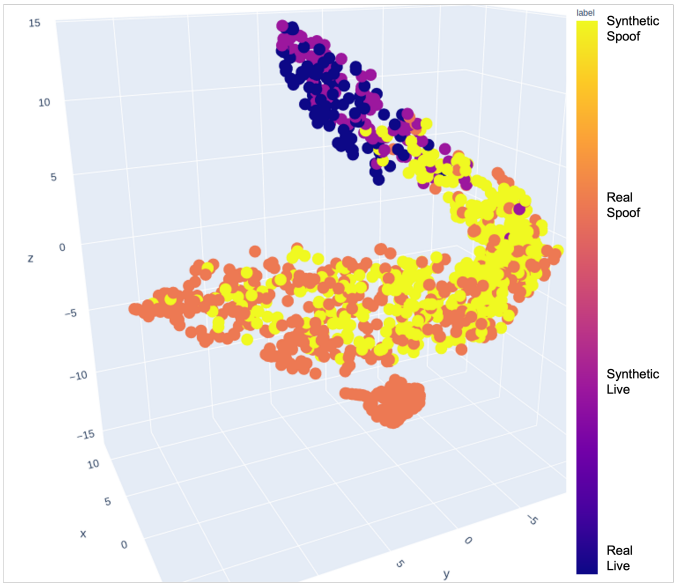

Fig. 5: 3D visualization of 2,048-dimensional embeddings of real live and spoof images from the LivDet2015 CrossMatch dataset compared with embeddings of our synthetic live and spoof images generated from SpoofGAN. Best viewed in color.

TABLE III: Spoof detection model trained on real fingerprints (GCT 1-5) and evaluated on LivDet 2015 CrossMatch images (top row) and an equivalent sized synthetic fingerprint dataset (bottom row)[1]. Results given in TDR @ FDR = 0.2%.

	BodyDouble	Ecoflex	PlayDoh	OOMOO	Gelatine
Real	100%	99.57%	98.58%	99.58%	100%
Synthetic	100%	100%	100%	99.58%	99.12%

[1] There are 1,500 live and 1,448 spoof fingerprint test images for CrossMatch in LivDet 2015, which we have replicated with synthetic data. Specifically, the spoof images consist of 300 Body Double, 270 Ecoflex, 297 OOMOO, 281 PlayDoh, and 300 Gelatine images.

images should closely align. For this analysis, we computed several statistics from the LivDet 2015 CrossMatch training dataset (lives only) and 1,500 SpoofGAN live fingerprint images which are shown in Table IV. In terms of fingerprint area, SpoofGAN images are, on average, smaller compared to the real fingerprint database. Since our training dataset consists of images of all 10 fingers, there is a bias toward smaller fingerprints considering the thumb as a minority class. Given this assumption, it is perhaps unsurprising that a GAN-based generation approach might exaggerate this class imbalance and generate smaller fingerprint area impressions. This problem is related to mode-collapse and has been noted in several GAN related works [63], [64], [65], with some recent papers proposing strategies to improve the generation process in class-imbalanced datasets [66], [67].

Next, we computed the NFIQ 2.0 quality metric [31] on both datasets. The NFIQ 2.0 scores for SpoofGAN are, on average, lower compared to LivDet. However, since one of the features considered in NFIQ 2.0 is the fingerprint area, we recomputed the scores on a 256×256 center crop of each of the fingerprints and observed that independent of fingerprint area, the NFIQ scores between SpoofGAN images and LivDet are much more aligned (36.88 ± 10.18 vs. 43.65 ± 14.12).

Lastly, we computed some additional metrics specific to the distribution of minutiae since many of the state-of-the-art fingerprint algorithms incorporate minutiae information. The average minutiae count and minutiae quality computed by Verifinger 12.0 are given in Table IV. The average number of minutiae found in SpoofGAN seems to be lower compared to the CrossMatch images from LivDet 2015; however, the minutiae per Megapixel is similar for both datasets (59.74 vs. 59.49). The minutiae quality given by Verifinger is also very similar between the two datasets (71.89 vs. 70.78).

3) Diversity in Generated Fingerprint Identities: To verify that SpoofGAN generated fingerprints mimic the similarity score distribution of real live and spoof fingerprints, we have computed genuine and imposter matches with Verifinger v12.0. In particular, we computed match scores (genuine and imposter) between the live impressions of the LivDet 2015 CrossMatch images as well as between the live samples of synthetic SpoofGAN fingerprints. These distributions are shown in Figure 6 (a). This figure highlights that SpoofGAN is generating diverse fingerprint identities with similar intraclass and interclass variation as the real LivDet 2015 CrossMatch dataset, albeit producing slightly lower genuine scores compared to the real dataset. However, in terms of recognition performance at a fixed false acceptance rate (FAR), the performance between the two datasets is quite similar, despite the slightly shifted genuine distribution of the SpoofGAN images (see Table V).

Furthermore, we computed genuine score distributions between the individual spoof types generated by SpoofGAN and their corresponding live impressions. These distributions are given in Figure 6 (b). Here we see that Verifinger is able to successfully match spoof and live images of corresponding identities, which we believe opens the door to synthesising a large-scale spoof and recognition dataset that can be used to train and evaluate joint PAD and recognition algorithms.

4) Identity Leakage: A major advantage of generating synthetic fingerprint data is that, theoretically, no fingerprint identity belongs to an actual user somewhere in the world. However, there remains a concern that synthesis methods, such as GANs, may inadvertently over-fit and leak private information from the training corpus [68]. Therefore, it is instructive to investigate whether, and to what degree, any of our SpoofGAN generated images are revealing, i.e., match with sufficient confidence, the identities present in our training database. Toward this end, we have computed match scores between 1,500 SpoofGAN generated fingerprint identities and each of the 38,164 live fingerprint images in our real training set. Out of the roughly 57.2 million $(1,500 \times 38,164)$ potential matches, only 50 comparisons exceeded the matching threshold of 48 set by Verifinger for a false acceptance rate of 0.01% with a maximum match score of 81. Furthermore, the 50 matches resulted from just 29 SpoofGAN generated fingers out of the 1,500 evaluated. Some example matched SpoofGAN and real fingerprint image pairs are shown in Figure 7 along with their corresponding match scores.

TABLE IV: Metrics for real and SpoofGAN fingerprint images. Minutiae quality and NFIQ2 scores have a range of [0, 100].

	LivDet 2015 CrossMatch		SpoofGAN	
Measure	Mean	Std. Dev.	Mean	Std. Dev.
Total Minutiae Count	55.56	18.43	40.45	11.57
Ridge Ending Minutiae Count	30.58	12.12	20.99	6.67
Ridge Bifurcation Minutiae Count	24.98	9.18	19.46	6.61
Verifinger Minutiae Quality	71.89	16.00	70.78	15.15
Fingerprint Area (Megapixels)	0.93	0.30	0.68	0.15
Fingerprint Image Quality (NFIQ2)	60.18	19.35	44.34	14.38

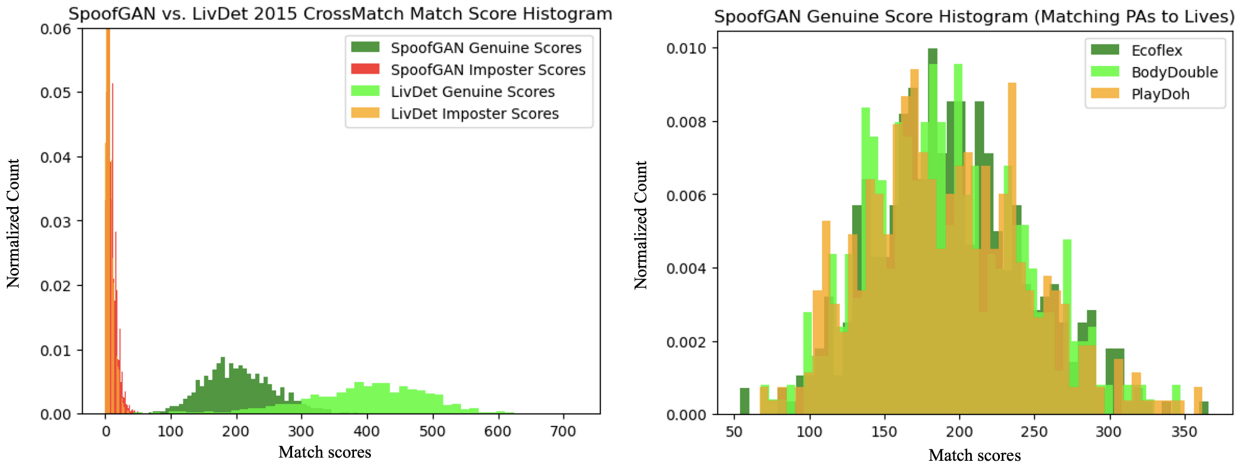

(a) (b)

Fig. 6: Verifinger 12.0 match score distributions of the real LivDet 2015 CrossMatch L Scan Guardian database vs. an equivalent sized synthetic SpoofGAN database. (a) match scores computed between the live impressions of each dataset and (b) match scores computed between live and spoof impressions for the SpoofGAN images.

TABLE V: TAR at Varying FAR Thresholds for LivDet 2015 CrossMatch Test Data vs. SpoofGAN Images.

FAR	LivDet	SpoofGAN
0.01% @ threshold=48	99.87%	100%
0.001% @ threshold=60	99.80%	100%
0.0001% @ threshold=72	99.74%	99.80%
1e-05% @ threshold=84	99.41%	99.47%
1e-06% @ threshold=96	99.35%	98.87%

D. Improved Spoof Classification with Synthetic Fingerprints

Ultimately, our synthetic fingerprints should offer some utility in advancing the training of fingerprint spoof detection algorithms. Toward this end, we have augmented three existing, publicly available datasets of live and spoof fingerprints with our synthetic fingerprints in an effort to improve the performance beyond that achievable when training on the real live and spoof images from each dataset alone. For this evaluation, we have trained several spoof detection models on the following training set compositions: i.) synthetic live and spoof images only, ii.) real live and spoof images only, iii.) synthetic live and spoof images plus only real live images, and iv.) synthetic live and spoof images plus real live and spoof images. We used the SpoofBuster model, which consists of two Inception v3 networks, one trained on the whole image input and the other trained on 96×96 minutiae centered patches [38]. The final spoof score is the weighted fusion of the two networks, with a minutiae patch score weight of 0.8 and a whole image score weight of 0.2. Each of the models are trained from scratch using Tensorflow on a single Nividia TitanX 1080 GPU on their respective datasets with identical hyper-parameters (learning rate of 0.01, step decay learning rate schedule, adam optimizer with default parameters, and total training updates of 200,000 steps).

Shown in Table VI, we have evaluated each of the models on their respective test sets. Despite the lower performance when training on synthetic data alone compared to training on real data, we see improvement in the overall spoof classification performance when the real training data is augmented with samples from our synthetic live and spoof generator. For example, the error of the minutiae patch model trained on real data from LivDet 2013 is reduced by 91.03% (from 15.60% to 1.40%) when augmented with synthetic data. Similarly, the error on LivDet 2015 is reduced from 0.48% to 0.0%, where the performance on GCT 6 remained the same at 100%.

Lastly, it may be the case the researchers and practitioners have access to a large, private database of real live fingerprint images, whereas collecting an equivalent database of spoof fingerprints is more difficult and costly; therefore, they may wish to augment their database of real live images with synthetic spoof images. As seen in Table VI, mixing synthetic

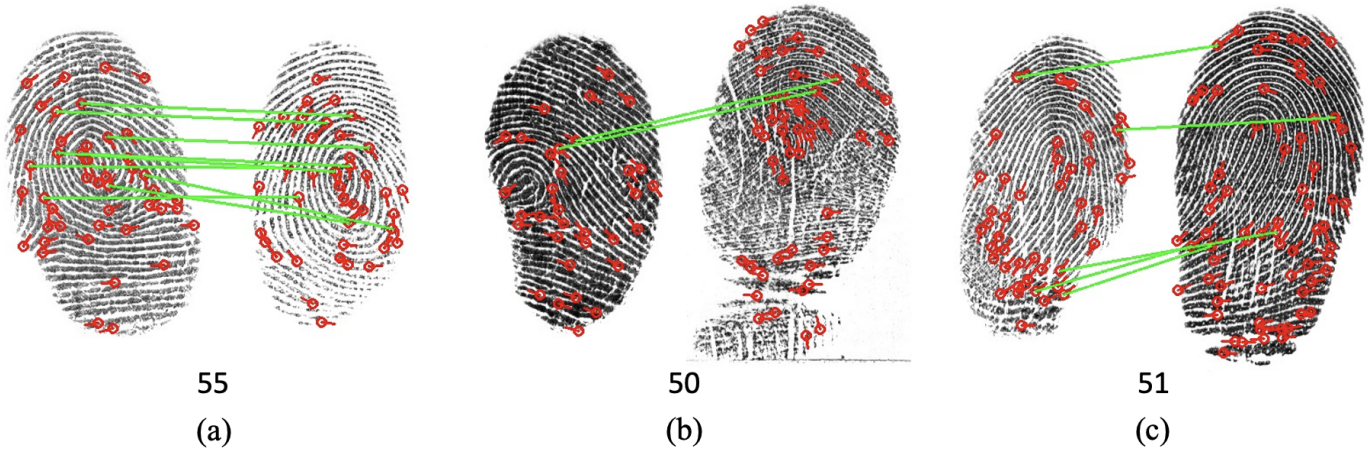

Fig. 7: Example SpoofGAN and real training database image pairs with corresponding match scores given by Verifinger. For each pair, the left image is a SpoofGAN fingerprint image and the right image is a real fingerprint image. The threshold for a genuine match at an FAR of 0.01% for Verifinger is 48, indicating that the identity of these training images have been "leaked" by SpoofGAN in the corresponding generated images.

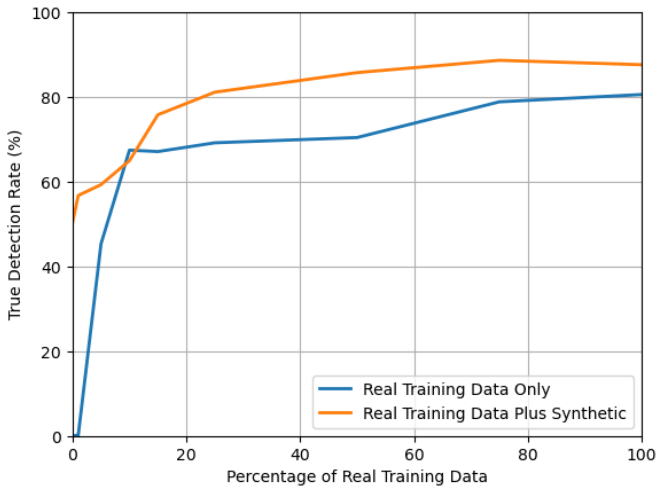

Fig. 8: Spoof detection performance as we keep the number of synthetic training samples fixed and vary the percentage of real data included in training. Performance reported as TDR @ FDR = 0.2% on the LivDet 2015 dataset.

data with only real live fingerprint examples improves the performance over training on just the synthetic examples alone; however, synthetic spoof data is still not a substitute for collecting real spoof examples as the performance still lags quite significantly. Alternatively, Figure 8 shows the trend in performance on LivDet 2015 as we keep the number of synthetic training samples fixed but varying the percentage of real data included when training the whole image-based spoof detection model. This figure suggests that when augmenting the training set with synthetic data, just 25% of the original (real) data is required to obtain similar performance to training on 100% of the real data alone, which significantly reduces the time and resources required for data collection to obtain similar performance.

V. Conclusion and Future Work

In this work, we presented a GAN-based synthesis method for generating high quality 512×512 plain fingerprint impressions of both live and spoof varieties. We demonstrated the utility of our synthetically generated spoof fingerprints in improving the performance of a spoof detector trained on real spoof fingerprint datasets. Additionally, our synthetic fingerprints closely resemble a database of real fingerprints both qualitatively and quantitatively in terms of various statistics, such as distribution of minutiae, NFIQ 2.0 image quality, spoof classification, and match score distributions computed by the state-of-the-art Verifinger 12.0 SDK. Finally, since our method is capable of generating multiple genuine and imposter fingerprints of unique identities of both live and spoof types, we open the door for large-scale training and evaluation of joint fingerprint spoof detection and recognition algorithms; overcoming a current limitation given the existing scale of publicly available spoof fingerprint datasets.

Despite the demonstrated utility of our synthetic spoof images, there remains several limitations that will be addressed in future work. First, given the lower standard deviation across the various fingerprint metrics presented in Table IV, the diversity of the generated images could be improved. A related issue, specific to fingerprint spoof detection, is the ever increasing novelty of spoof materials and types that may be encountered in the future; thus, instilling the generation process with the ability to adapt to novel spoof types is a promising future research direction. Similarly, leveraging synthetic data for improved cross-dataset and cross-material generalization is a promising future research direction. Lastly, training the synthetic fingerprint generator in an online fashion with a spoof detection network may provide additional supervision to generate more useful live and spoof examples to improve the spoof detection performance.

References

[1] F. Galton, *Finger prints*. Macmillan and Company, 1892.

TABLE VI: Spoof detection performance when trained on various combinations of real and synthetic data. We have followed the established protocols of train/test split for each of the evaluation datasets, which are shown in Table II

Training Data Composition				TDR @ 0.2% FDR		
Real Live	**Real Spoof**	**Synthetic Live**	**Synthetic Spoof**	**LivDet2015**	**LivDet2013**	**GCT6**
		✓	✓	36.53%	42.30%	56.08%
✓		✓	✓	80.11%	73.00%	86.86%
✓	✓			99.52%	84.40%	**100.0%**
✓	✓	✓	✓	**100.0%**	**98.60%**	**100.0%**

[2] S. Pankanti, S. Prabhakar, and A. K. Jain, "On the individuality of fingerprints," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 24, no. 8, pp. 1010–1025, 2002.

[3] S. Yoon and A. K. Jain, "Longitudinal study of fingerprint recognition," *Proceedings of the National Academy of Sciences*, vol. 112, no. 28, pp. 8555–8560, 2015.

[4] B. Dorizzi, R. Cappelli, M. Ferrara, D. Maio, D. Maltoni, N. Houmani, S. Garcia-Salicetti, and A. Mayoue, "Fingerprint and on-line signature verification competitions at icb 2009," in *International Conference on Biometrics*, pp. 725–732, Springer, 2009.

[5] K. Cao and A. K. Jain, "Fingerprint indexing and matching: An integrated approach," in *2017 IEEE International Joint Conference on Biometrics (IJCB)*, pp. 437–445, IEEE, 2017.

[6] D. Song and J. Feng, "Fingerprint indexing based on pyramid deep convolutional feature," in *2017 IEEE International Joint Conference on Biometrics (IJCB)*, pp. 200–207, IEEE, 2017.

[7] D. Song, Y. Tang, and J. Feng, "Aggregating minutia-centred deep convolutional features for fingerprint indexing," *Pattern Recognition*, vol. 88, pp. 397–408, 2019.

[8] R. Li, D. Song, Y. Liu, and J. Feng, "Learning global fingerprint features by training a fully convolutional network with local patches," in *2019 International Conference on Biometrics (ICB)*, pp. 1–8, IEEE, 2019.

[9] J. J. Engelsma, K. Cao, and A. K. Jain, "Learning a fixed-length fingerprint representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.

[10] C. Lin and A. Kumar, "A cnn-based framework for comparison of contactless to contact-based fingerprints," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 3, pp. 662–676, 2018.

[11] S. A. Grosz, J. J. Engelsma, E. Liu, and A. K. Jain, "C2CL: Contact to contactless fingerprint matching," *IEEE Transactions on Information Forensics and Security*, 2021.

[12] J. J. Engelsma, A. K. Jain, and V. N. Boddeti, "Hers: Homomorphically encrypted representation search," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2022.

[13] S. Minaee, A. Abdolrashidi, H. Su, M. Bennamoun, and D. Zhang, "Biometrics recognition using deep learning: A survey," *arXiv preprint arXiv:1912.00271*, 2019.

[14] K. Sundararajan and D. L. Woodard, "Deep learning for biometrics: A survey," *ACM Computing Surveys (CSUR)*, vol. 51, no. 3, pp. 1–34, 2018.

[15] J. J. Engelsma, S. A. Grosz, and A. K. Jain, "PrintsGAN: Synthetic fingerprint generator," *arXiv preprint arXiv:2201.03674*, 2022.

[16] R. Bouzaglo and Y. Keller, "Synthesis and reconstruction of fingerprints using generative adversarial networks," *arXiv preprint arXiv:2201.06164*, 2022.

[17] K. Bahmani, R. Plesh, P. Johnson, S. Schuckers, and T. Swyka, "High fidelity fingerprint generation: Quality, uniqueness, and privacy," *arXiv preprint arXiv:2105.10403*, 2021.

[18] A. B. V. Wyzykowski, M. P. Segundo, and R. d. P. Lemes, "Level three synthetic fingerprint generation," *arXiv preprint arXiv:2002.03809*, 2020.

[19] V. Mistry, J. J. Engelsma, and A. K. Jain, "Fingerprint synthesis: Search with 100 million prints," in *2020 International Joint Conference on Biometrics (IJCB)*, 2020.

[20] K. Cao and A. Jain, "Fingerprint synthesis: Evaluating fingerprint search at scale," in *2018 International Conference on Biometrics (ICB)*, 2018.

[21] M. A. I. Fahim and H. Y. Jung, "A lightweight gan network for large scale fingerprint generation," *IEEE Access*, vol. 8, pp. 92918–92928, 2020.

[22] M. S. Riazi, S. M. Chavoshian, and F. Koushanfar, "Synfi: Automatic synthetic fingerprint generation," *arXiv preprint arXiv:2002.08900*, 2020.

[23] M. Attia, M. H. Attia, J. Iskander, K. Saleh, D. Nahavandi, A. Abobakr, M. Hossny, and S. Nahavandi, "Fingerprint synthesis via latent space representation," in *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*, pp. 1855–1861, IEEE, 2019.

[24] S. Minaee and A. Abdolrashidi, "Finger-gan: Generating realistic fingerprint images using connectivity imposed gan," *arXiv preprint arXiv:1812.10482*, 2018.

[25] R. Cappelli, D. Maio, and D. Maltoni, "Synthetic fingerprint-database generation," in *Object recognition supported by user interaction for service robots*, vol. 3, pp. 744–747, IEEE, 2002.

[26] Q. Zhao, A. K. Jain, N. G. Paulter, and M. Taylor, "Fingerprint image synthesis based on statistical feature models," in *2012 IEEE Fifth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pp. 23–30, IEEE, 2012.

[27] P. Johnson, F. Hua, and S. Schuckers, "Texture modeling for synthetic fingerprint generation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 154–159, 2013.

[28] P. Bontrager, A. Roy, J. Togelius, N. Memon, and A. Ross, "Deepmasterprints: Generating masterprints for dictionary attacks via latent variable evolution," in *2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pp. 1–9, IEEE, 2018.

[29] G. P. Fiumara, P. A. Flanagan, J. D. Grantham, K. Ko, K. Marshall, M. Schwarz, E. Tabassi, B. Woodgate, and C. Boehnen, "Nist special database 302: Nail to nail fingerprint challenge," Tech. Rep. NIST.TN.2007, National Institute of Standards and Technology, Gaithersburg, MD, 2019.

[30] L. Ghiani, D. Yambay, V. Mura, S. Tocco, G. L. Marcialis, F. Roli, and S. Schuckcrs, "Livdet 2013 fingerprint liveness detection competition 2013," in *2013 International Conference on Biometrics (ICB)*, pp. 1–6, IEEE, 2013.

[31] E. Tabassi, M. Olsen, O. Bausinger, C. Busch, A. Figlarz, G. Fiumara, O. Henniger, J. Merkle, T. Ruhland, C. Schiel, and M. Schwaiger, "NIST fingerprint image quality 2," Tech. Rep. NIST.IR.8382, National Institute of Standards and Technology, Gaithersburg, MD, 2021.

[32] G. L. Marcialis, A. Lewicke, B. Tan, P. Coli, D. Grimberg, A. Congiu, A. Tidu, F. Roli, and S. Schuckers, "First international fingerprint liveness detection competition—livdet 2009," in *International Conference on Image Analysis and Processing*, pp. 12–23, Springer, 2009.

[33] D. Yambay, L. Ghiani, P. Denti, G. L. Marcialis, F. Roli, and S. Schuckers, "Livdet 2011—fingerprint liveness detection competition 2011," in *2012 5th IAPR international conference on biometrics (ICB)*, pp. 208–215, IEEE, 2012.

[34] V. Mura, L. Ghiani, G. L. Marcialis, F. Roli, D. A. Yambay, and S. A. Schuckers, "LivDet 2015 - Fingerprint liveness detection competition 2015," in *Proc. IEEE 7th Intl. Conf. BTAS*, pp. 1–6, 2015.

[35] V. Mura, G. Orrù, R. Casula, A. Sibiriu, G. Loi, P. Tuveri, L. Ghiani, and G. L. Marcialis, "Livdet 2017 fingerprint liveness detection competition 2017," in *2018 International Conference on Biometrics (ICB)*, pp. 297–302, IEEE, 2018.

[36] G. Orrù, R. Casula, P. Tuveri, C. Bazzoni, G. Dessalvi, M. Micheletto, L. Ghiani, and G. L. Marcialis, "Livdet in action-fingerprint liveness detection competition 2019," in *2019 International Conference on Biometrics (ICB)*, pp. 1–6, IEEE, 2019.

[37] R. Casula, M. Micheletto, G. Orrù, R. Delussu, S. Concas, A. Panzino, and G. L. Marcialis, "Livdet 2021 fingerprint liveness detection competition-into the unknown," in *2021 IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–6, IEEE, 2021.

[38] T. Chugh, K. Cao, and A. Jain, "Fingerprint Spoof Buster: Use of Minutiae-centered Patches," *IEEE Transactions on Information Forensic and Security*, vol. 13, pp. 2190 – 2202, September 2018.

[39] T. Chugh and A. K. Jain, "Fingerprint presentation attack detection: Generalization and efficiency," in *2019 International Conference on Biometrics (ICB)*, pp. 1–8, IEEE, 2019.
[40] International Standards Organization, "ISO/IEC 30107-1:2016, Information Technology—Biometric Presentation Attack Detection—Part 1: Framework." https://www.iso.org/standard/53227.html, 2016.
[41] D. Baldisserra, A. Franco, D. Maio, and D. Maltoni, "Fake fingerprint detection by odor analysis," in *Proc. ICB*, pp. 265–272, Springer, 2006.
[42] P. D. Lapsley, J. A. Lee, D. F. Pare Jr, and N. Hoffman, "Anti-fraud biometric scanner that accurately detects blood flow," 1998. US Patent 5,737,439.
[43] J. J. Engelsma, K. Cao, and A. K. Jain, "Raspireader: Open Source Fingerprint Reader," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
[44] T. Chugh and A. K. Jain, "OCT Fingerprints: Resilience to Presentation Attacks," *arXiv preprint arXiv:1908.00102*, 2019.
[45] R. Tolosana, M. Gomez-Barrero, J. Kolberg, A. Morales, C. Busch, and J. Ortega-Garcia, "Towards fingerprint presentation attack detection based on convolutional neural networks and short wave infrared imaging," in *2018 International Conference of the Biometrics Special Interest Group (BIOSIG)*, pp. 1–5, IEEE, 2018.
[46] L. Spinoulas, H. Mirzaalian, M. E. Hussein, and W. AbdAlmageed, "Multi-modal fingerprint presentation attack detection: Evaluation on a new dataset," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 3, no. 3, pp. 347–364, 2021.
[47] G. L. Marcialis, F. Roli, and A. Tidu, "Analysis of fingerprint pores for vitality detection," in *Proc. 20th ICPR*, pp. 1289–1292, 2010.
[48] E. Marasco and C. Sansone, "Combining perspiration-and morphology-based static features for fingerprint liveness detection," *Pattern Recognition Letters*, vol. 33, no. 9, pp. 1148–1156, 2012.
[49] L. Ghiani, G. L. Marcialis, and F. Roli, "Fingerprint liveness detection by local phase quantization," in *Proc. 21st ICPR*, pp. 537–540, 2012.
[50] L. Ghiani, A. Hadid, G. L. Marcialis, and F. Roli, "Fingerprint liveness detection using Binarized Statistical Image Features," in *Proc. IEEE 6th Int. Conf. BTAS*, pp. 1–6, 2013.
[51] R. F. Nogueira, R. de Alencar Lotufo, and R. C. Machado, "Fingerprint Liveness Detection Using Convolutional Neural Networks," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 6, pp. 1206–1213, 2016.
[52] F. Pala and B. Bhanu, "Deep Triplet Embedding Representations for Liveness Detection," in *Deep Learning for Biometrics. Advances in Computer Vision and Pattern Recognition*, pp. 287–307, Springer, 2017.
[53] J. J. Engelsma and A. K. Jain, "Generalizing Fingerprint Spoof Detector: Learning a One-Class Classifier," *IEEE International Conference on Biometrics (ICB)*, 2019.
[54] Y. Ding and A. Ross, "An ensemble of one-class SVMs for fingerprint spoof detection across different fabrication materials," in *Proc. IEEE WIFS*, pp. 1–6, 2016.
[55] S. A. Grosz, T. Chugh, and A. K. Jain, "Fingerprint presentation attack detection: a sensor and material agnostic approach," in *2020 IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–10, IEEE, 2020.
[56] J. A. Pereira, A. F. Sequeira, D. Pernes, and J. S. Cardoso, "A robust fingerprint presentation attack detection method against unseen attacks through adversarial learning," in *BIOSIG 2020-Proceedings of the 19th International Conference of the Biometrics Special Interest Group*, Gesellschaft für Informatik eV, 2020.
[57] T. Chugh and A. K. Jain, "Fingerprint spoof detector generalization," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 42–55, 2020.
[58] R. Gajawada, A. Popli, T. Chugh, A. Namboodiri, and A. K. Jain, "Universal Material Translator: Towards Spoof Fingerprint Generalization," in *IEEE International Conference on Biometrics (ICB)*, 2019.
[59] A. B. V. Wyzykowski, M. P. Segundo, and R. de Paula Lemes, "Level three synthetic fingerprint generation," in *2020 25th International Conference on Pattern Recognition (ICPR)*, pp. 9250–9257, IEEE, 2021.
[60] A. Brock, J. Donahue, and K. Simonyan, "Large scale gan training for high fidelity natural image synthesis," *arXiv preprint arXiv:1809.11096*, 2018.
[61] X. Si, J. Feng, J. Zhou, and Y. Luo, "Detection and rectification of distorted fingerprints," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 3, pp. 555–568, 2015.
[62] F. L. Bookstein, "Principal warps: Thin-plate splines and the decomposition of deformations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 6, pp. 567–585, 1989.
[63] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in Neural Information Processing Systems*, vol. 27, 2014.
[64] A. Srivastava, L. Valkov, C. Russell, M. U. Gutmann, and C. Sutton, "Veegan: Reducing mode collapse in gans using implicit variational learning," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
[65] K. Roth, A. Lucchi, S. Nowozin, and T. Hofmann, "Stabilizing training of generative adversarial networks through regularization," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
[66] G. Mariani, F. Scheidegger, R. Istrate, C. Bekas, and C. Malossi, "Bagan: Data augmentation with balancing gan," *arXiv preprint arXiv:1803.09655*, 2018.
[67] G. Huang and A. H. Jafari, "Enhanced balancing gan: Minority-class image generation," *Neural Computing and Applications*, pp. 1–10, 2021.
[68] S. M. Bellovin, P. K. Dutta, and N. Reitinger, "Privacy and synthetic datasets," *Stan. Tech. L. Rev.*, vol. 22, p. 1, 2019.

Steven A. Grosz received his B.S. degree with highest honors in Electrical Engineering from Michigan State University, East Lansing, Michigan, in 2019. He is currently a doctoral student in the Department of Computer Science and Engineering at Michigan State University. His primary research interests are in the areas of machine learning and computer vision with applications in biometrics.

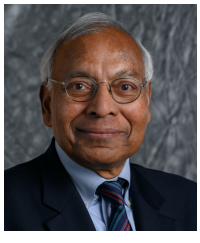

Anil K. Jain Anil K. Jain is a University distinguished professor in the Department of Computer Science and Engineering at Michigan State University. His research interests include pattern recognition and biometric authentication. He served as the editor-in-chief of the IEEE Transactions on Pattern Analysis and Machine Intelligence and was a member of the United States Defense Science Board. He has received Fulbright, Guggenheim, Alexander von Humboldt, and IAPR King Sun Fu awards. He is a member of the National Academy of Engineering, the Indian National Academy of Engineering, the World Academy of Sciences, and the Chinese Academy of Sciences.