Received: 14 March 2021 | Revised: 5 April 2022 | Accepted: 22 May 2022

DOI: 10.1049/bme2.12083

IET Biometrics

The Institution of Engineering and Technology WILEY

ORIGINAL RESEARCH

Multiresolution synthetic fingerprint generation

Andre Brasil Vieira Wyzykowski[1] | **Mauricio Pamplona Segundo**[2] | **Rubisley de Paula Lemes**[1]

[1]Institute of Mathematics and Statistics, Federal University of Bahia, Salvador, Brazil

[2]Computer Science and Engineering, University of South Florida, Tampa, Florida, USA

Correspondence

Andre Brasil Vieira Wyzykowski, Institute of Mathematics and Statistics, Federal University of Bahia, Av. Adhemar de Barros, s/n° - Ondina, Salvador, Brazil.

Email: abvwmc@gmail.com

Funding information

Conselho Nacional de Desenvolvimento Científico e Tecnológico

Abstract

Public access to existing high-resolution databases was discontinued. Besides, a hybrid database that contains fingerprints of different sensors with high and medium resolutions does not exist. A novel hybrid approach to synthesise realistic, multiresolution, and multisensor fingerprints to address these issues is presented. The first step was to improve Anguli, a handcrafted fingerprint generator, to create pores, scratches, and dynamic ridge maps. Using CycleGAN, then the maps are converted into realistic fingerprints, adding textures to images. Unlike other neural network-based methods, the authors' method generates multiple images with different resolutions and styles for the same identity. With the authors' approach, a synthetic database with 14,800 fingerprints is built. Besides that, fingerprint recognition experiments with pore- and minutiae-based matching techniques and different fingerprint quality analyses are conducted to confirm the similarity between real and synthetic databases. Finally, a human classification analysis is performed, where volunteers could not distinguish between authentic and synthetic fingerprints. These experiments demonstrate that the authors' approach is suitable for supporting further fingerprint recognition studies in the absence of real databases.

1 | INTRODUCTION

Fingerprint recognition is widely studied thanks to its compliance with the core premises of biometrics: permanence and distinctiveness [1, 2]. It is possible to analyse the ridge patterns on fingerprints at different scales and resolutions. These patterns can be classified as Level 1 (L1—global patterns, such as ridge orientation maps and fingerprint classes), Level 2 (L2—local patterns, such as minutiae) and Level 3 (L3—fine details, such as sweat pores, incipient ridges and dots).

With the development of high-resolution sensors that are able to capture L3 fingerprint images, researchers saw an opportunity to devise more accurate recognition approaches by using extra information, such as sweat pores [3–5]. Besides, L3-based approaches improve security by hindering spoof attempts [6–8].

Despite the recent improvements brought to the fingerprint recognition research area, L3 fingerprint databases are being discontinued. For instance, databases such as the NIST Special Database 30 [9] and the Hong Kong Polytechnic University High Resolution Fingerprint database (PolyU) [10] are no longer publicly available. More recently, Anand and Kanhangad [11, 12] created L3 databases for their recognition experiments. However, until the present moment, these databases were not released. The main reason for that is the existence of legal restrictions protecting the privacy of biometric data.

There are several procedures and limits imposed that must be followed to acquire and distribute fingerprint images. This is a legal trend all over the world, as observed in the Illinois Biometric Information Privacy Act (BIPA) in the United States, the General Data Protection Regulation (RGPD) 2016/679 in the European Union, and the General Data Protection (LGPD)—article 5, clause II of the Law n° 13.709/2018 in Brazil. Although they are important laws and regulations, they hinder the evolution of fingerprint biometric recognition algorithms in a time context of popularisation of fingerprint sensors in various devices.

Another evident problem is the lack of a fingerprint database containing high and medium resolution fingerprints from different sensors. This is a severe obstacle for new research to improve the interoperability of fingerprint recognition across different sensors and resolutions. This topic has been discussed in the state of the art [13]. The performance is not optimal in current algorithms and some methods were proposed to tackle this issue [14, 15]. However, they do not perform cross-sensor matching between a high-resolution L3 sensor and a regular L2 sensor.

Creating a multiresolution and multisensor synthetic database is a valid alternative to address these problems, still avoiding costly data collection campaigns. Synthetic generation techniques are widely exploited in several areas, such as optical flow computation [16], indoor robot navigation [17], and autonomous driving [18]. The same is true for the fingerprint recognition field.

Earlier studies, such as SFinGe [19] and Anguli [20], developed handcrafted frameworks for fingerprint generation based on the knowledge of an expert. Their design allows reasonable control over the identity of the output images but lack realism, especially when considering L3 features. Besides, the non-commercial SFinGe version restricts the generation of large datasets, difficulting the generation of multiple instances of a single identity. Finger-GAN [21] and Cao and Jain's approach [22] employed Generative Adversarial Networks (GAN) [23] to learn how to generate realistic fingerprints from a real training set. However, both of these studies present many unnatural artefacts and cannot generate multiple images for the same identity. The Clarkson Fingerprint Generator (CFG) method [24] uses the StyleGAN [25] to synthesise fingerprints. However, the database of real fingerprints captured by the authors is not publicly available. This fact makes it impossible to replicate and compare this work under identical conditions. Furthermore, the method also does not generate multiple fingerprints for a single identity, suffering from the lack of identity variability (see Figure 1).

We propose a novel approach to create realistic, multiresolution and multisensor synthetic fingerprints while maintaining control over the identity of the output images. Our goal is to foster further studies in this field without raising legal issues that come with real biometric data. Our contributions are:

1. A novel hybrid fingerprint generation approach that combines a handcrafted identity generator and a learnt texturiser across different sensors and resolutions. A visual comparison to existing approaches shows that our results are the new state of the art. Also, a perception experiment shows that humans can hardly differ between real and our synthetic images.
2. A new approach that can generate multiple instances of high and medium resolution fingerprints of a single identity. This ability is a game-changer in fingerprint synthesis because current neural network-based methods cannot perform this task, limiting performance recognition studies.
3. A public database[1] of L3 and L2 synthetic fingerprint images, each one containing five subsets of 148 identities, with 10 samples per identity, totalling 14,800 high-resolution fingerprint images. These images enable cross-matching researches for different purposes. Also, we include sweat pore annotations for 740 images to assist in pore detection research [26]. This is the largest publicly available database with L3 fingerprints nowadays.

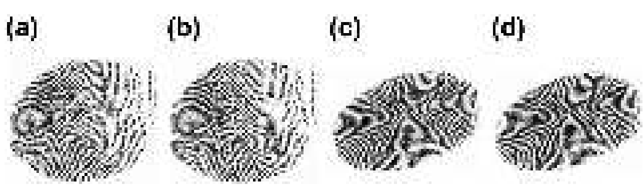

FIGURE 1 The Clarkson Fingerprint Generator (CFG) method [24] samples. Note the lack of variation between the different identities (a) and (b), and (c) and (d)

This article is an extension of our previous conference paper [27].[1] The main extension between the two versions is the addition of the generation of L2 fingerprints, where we performed a new experiment comparing fingerprint matching at different resolutions and sensors. The generation of multiresolution and multisensor fingerprints for a pre-defined set of identities is unique in the state of the art.

2 | MULTIRESOLUTION SYNTHETIC FINGERPRINT GENERATION

Fingerprints are used in many contexts. Whether in forensics, government control systems, or banking systems, fingerprints play an essential role in asserting a person's identity. However, there are many sensors developed by various manufacturers (see examples in Figure 2), and each sensor design will cause variations in the styles of the images created. These variations create differences in image resolutions, textures, image tonality, thus generating a significant challenge in biometric recognition. Optical and capacitive sensors are commonly used in many devices we use daily. However, a new diversity of sensors has been proposed and adopted in products in the industry. For example, ultrasonic sensors [28, 29] are started to be integrated into smartphones and tablets successfully. Nevertheless, there are a variety of transducer technologies that can be applied to the acquisition, which will also generate changes in the biometric signal [30]. More recently, new studies [31, 32] have further expanded the topic of fingerprint imaging using contactless acquisition techniques. However, challenges such as segmentation and distortion correction are essential factors for performing an appropriate matching, problems that researchers are still tackling these challenges [33].

Ideally, all these biometric recognition methods must work in an interoperable and ubiquitous way in a highly connected world. Over the years, several studies have tackled the problem of interoperability in fingerprint sensors [13, 34, 35], evidencing several challenges. Another challenge encountered occurs when

[1] https://andrewyzy.github.io/MR-SF/.

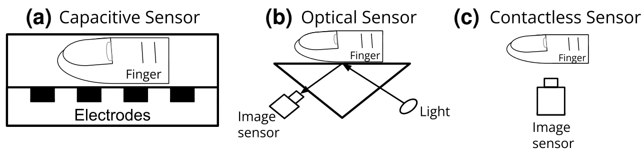

FIGURE 2 Fingerprint sensor examples. Each sensor will produce different fingerprint image styles

matching databases from different countries or regions that use different standards. This problem received much attention from the members of the European Union, where they use Interpol and Europol standards and still require interoperability with the FBI database.[2]

Another problem in fingerprint interoperability is the evidence of the lack of research involving medium and high-resolution sensors. Having a database containing images of the same identities in different styles and resolutions would help in biometric recognition in tasks where a high degree of accuracy is required. However, there is no public database with this feature to date, which results in a significant obstacle to research in this area.

To generate multiresolution synthetic fingerprints, we split our approach into two stages. The first stage concerns procedures required to create multiple instances of fingerprints, which we call **seed images**. The second stage consists of using CycleGAN [36] to translate seed images into realistic fingerprints. Figure 3 summarises the workflow of the proposed method.

The first stage consists of the following processes:

1. **Master fingerprint generation:** we extended Anguli [20], an open-source implementation of SFinGe [19], to create fingerprint ridge maps with random ridge-flow frequency. Also, we segment ridges and dynamically change their thicknesses. We call the resulting images master fingerprints (Section 2.1).
2. **Pore and scratch generation:** we add pores and scratches to each master fingerprint following a distribution learnt from real images to obtain L3 master fingerprints (Section 2.2).
3. **Fingerprint acquisition simulation:** to simulate acquisition, we randomly cut the L3 master fingerprints following a distribution of the displacement among images of the same person in a real database, thus creating different instances for each identity. These instances are called seed images (Section 2.3).

The second stage consists of performing CycleGAN translation to transform any seed image into a realistic multi-level fingerprint. The model employed in this step is obtained through the following processes:

1. **Training set generation:** to train CycleGAN, we create several seed images following the steps of the first stage of our approach. Data augmentation is then applied to real and seed images to compose the training set (Section 3.1.1).
2. **CycleGAN training:** uses the training set to create a CNN model that translates seed images to high or medium resolution fingerprints (Section 2.4).

2.1 | Master fingerprint generation

We use Anguli [20] to create the ridge map that composes a master fingerprint within one of the following classes: Whorl, Right Loop, Left Loop, Plain Arch, and Tented Arch. Figure 4 illustrates these classes. When creating a new identity, it is necessary to follow the proportion of these classes in a real population. Martijn van Mensvoort [37] gathered fingerprint distributions from 32 countries through a compilation of 28 published articles. Each country has its peculiarities and proportions, so we decided to use the global mean distribution. The mean distribution for 32 countries is as follows: Whorl: 41%, Right Loop: 50%, Left Loop: 3%, and Arch: 6%. Since Anguli can create two arch fingerprint types, we used the ratio provided by Wang [38]: Plain Arch: 72.22% and Tented Arch: 27.77%.

The Gabor filter plays an important role when generating ridges using Anguli. By changing its scale, we introduce variations to the ridge frequency in synthetic fingerprints. We changed Anguli to randomly modify the scale of the Gabor filter by up to 20%. To exemplify this effect, we show images in Figure 4 with different frequencies. As can be observed, the first image frequency is much higher than the second one.

To segment ridges, we upscale the modified Anguli images, which have 275 × 400 pixels, by a factor of 3 with FSRCNN [39] and smooth the result by applying a 3 × 3 mean filter. Then, we skeletonise those images using Zhang and Suen's thinning algorithm [40] and split continuous segments of pixels as individual ridges. To prevent spurious minutiae, we eliminate ridges smaller than 5 pixels.

To create a higher variability in the ridges' thicknesses, we dynamically change ridge thickness based on the sine function, generating a smooth transition among neighbouring ridges. We iteratively calculate $w_i = |3 \times \sin t|$, where w_i is the width of the ith ridge and t is a counter starting at a random value for each image (see Figure 5). After processing a ridge, t is incremented by 0.1. Ridges are processed from left to right and top to bottom. This approach avoids abrupt changes in thickness among neighbouring ridges and also adds a stochastic factor to the thickness generation. The outcome is a **master fingerprint** with 825 × 1200 pixels, as illustrated in Figure 6.

2.2 | Pore and scratch generation

This step consists in marking where the pores will be placed on the ridges and applying the scratches on the images. Sweat pores are essential information in high-resolution fingerprints. In medium-resolution fingerprints, they appear less frequently

[2]https://www.aware.com/blog-biometric-data-interoperability-challenges/.

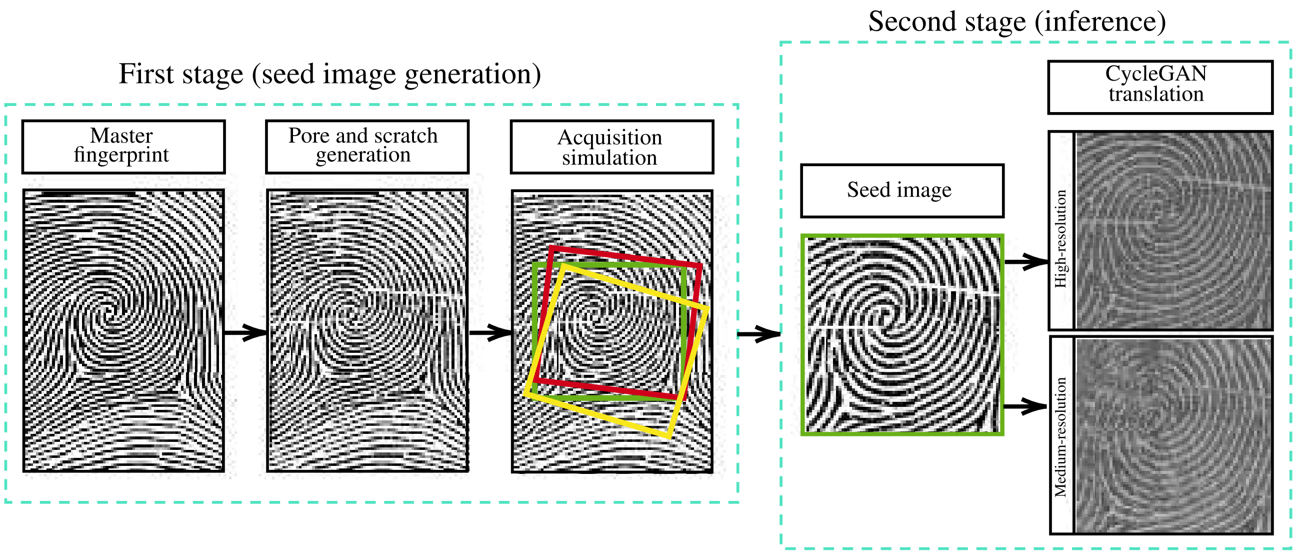

FIGURE 3 Flowchart with the steps to create a multiresolution synthetic fingerprint using the proposed approach

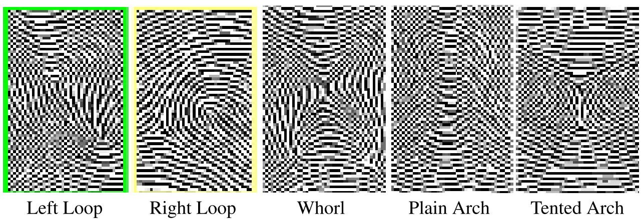

FIGURE 4 Five fingerprint class patterns created by our modified version of Anguli. Note that the first image (green) frequency is higher than the second one (yellow)

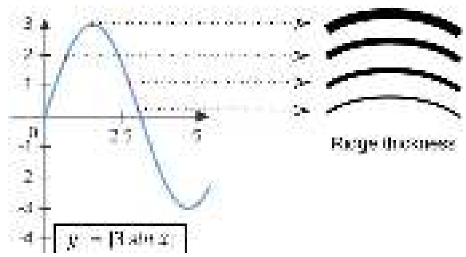

FIGURE 5 Variability in the ridges' thicknesses. Note that the sinusoidal function is used to generate changes in the thickness of the ridges

compared to a high-resolution image, and when it appears, they are in a lower definition. So regardless of the fact that resolution being required to generate the fingerprint, it is important to generate synthetic pores.

To measure the distance distribution from one pore to another, we use a pore-based ridge reconstruction approach [3]. Given a training set of real fingerprint images, we compute the average distance and the standard deviation among neighbouring pores as our reference distribution.

To add pores, we utilise the segmented ridges generated in Section 2.1. Starting from the beginning of a ridge, we iteratively sample distances d_i from the reference distribution. We follow the ridge pixels to add the ith pore d_i away from the previous pore until the end of the ridge is reached. This process of creating pores is illustrated in Figure 7.

To create scratches, we count the number of scratches in each fingerprint on a real database. With these values, we used the normalised cumulative density function, and we choose the number of scratches based on a uniform random number. Starting on a random point, for each scratch, we draw n consecutive line segments, where n is a random value between 1 and 4. The line segments have a random length with a maximum value of 150 pixels and a random angle ($-15° \leq \theta \leq 15°$) between them. This process of creating scratches is illustrated in Figure 8.

Figure 9 shows an **L3 master fingerprint**, the outcome of this step.

2.3 | Fingerprint acquisition simulation

Different finger positions on the imaging sensor produce distinct rotations and shifts on fingerprint images. Our approach aims to simulate these acquisition variations.

To measure rotation and shift variations between fingerprints of the same person from a training set, we extract SIFT [41, 42] and ORB [43] keypoints, perform an affine alignment [44] for each keypoint set separately, and select the one with the highest number of inliers. After, we use the RANSAC algorithm [45] to obtain a rigid transformation.

We assume that the centre of a L3 master fingerprint is the average of the centre of its samples aligned to each other in the same coordinate system. With this assumption, the average shift of the samples in relation to the L3 master fingerprint centre is zero in each axis, with standard deviations σ_x and σ_y independent from each other.

However, we cannot measure these displacements in a real dataset, as we have samples but do not have master fingerprints. What we can observe is the difference between two samples.

With x_i being the x-coordinate of the centre of the ith image, the expected difference between the centre of two samples i and j, aligned to the same coordinate system is given by:

$$\mathbb{E}\left[\left(x_i - x_j\right)^2\right] = \mathbb{E}\left[\mathcal{N}\left(0, {\sigma_x}^2\right) + \mathcal{N}\left(0, {\sigma_x}^2\right)\right]$$
$$= \mathbb{E}\left[\mathcal{N}\left(0, 2{\sigma_x}^2\right)\right]$$

Therefore, if we measure the average square difference $\overline{d}_x$ between pairs of samples from the same person, σ_x can be estimated as follows:

$$\sigma_x = \sqrt{\frac{\overline{d}_x}{2}}$$

The same can be done independently for σ_y and σ_θ (rotation). To create a seed image, we sample a random transformation from the normal distribution using σ_x, σ_y, and σ_θ. After that, we rotate and translate the L3 master fingerprint before cropping the centre region of size 512 × 512 from the resulting image. Examples of seed images from the same L3 master fingerprint are shown in Figure 10.

Before cropping, we also perform a random affine transformation on L3 master fingerprints. Assuming that (X, V, k) and (Z, W, k) are two affine spaces, where X and Z are point sets, we generate a random γ value between −10 and 10, summing γ to V and W (vector spaces over the field k). This is a way to simulate non-rigid deformations on fingerprints, which occur in real images due to distinct finger pressures during acquisition

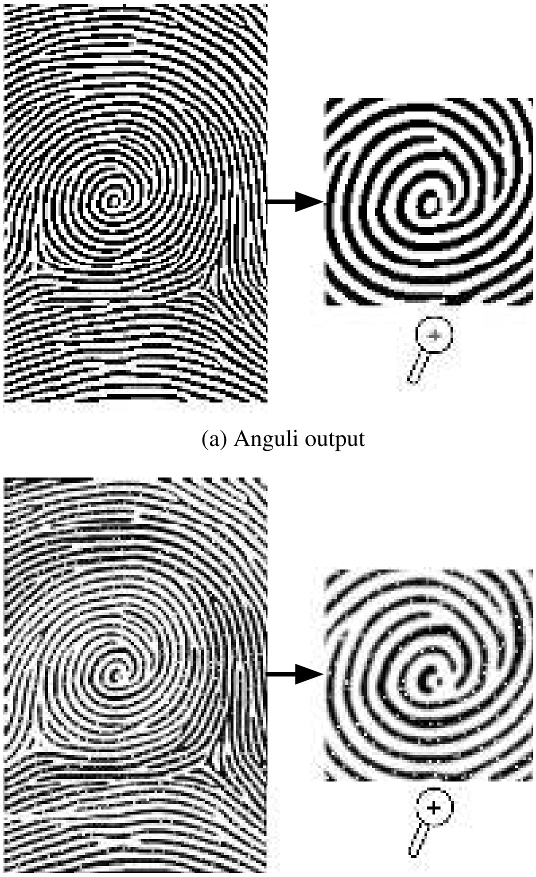

(a) Anguli output

(b) Master fingerprint

FIGURE 6 Example of (a) an output of our modified version of Anguli and (b) its respective master fingerprint after ridge thickness sinusoidal perturbation. Note the ridge thickness variability in the master fingerprint, where the ridges are thicker on the borders compared to the centre. Also, note our aliased result in comparison to Anguli's pixelated ridges

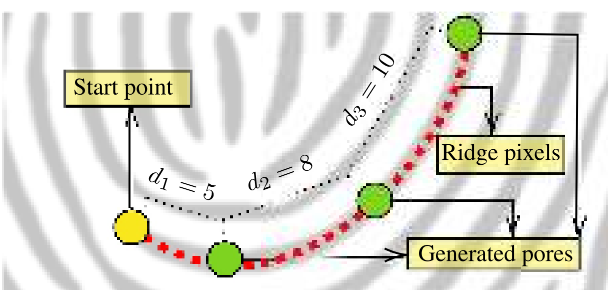

FIGURE 7 Illustration of the pore generation process. Red squares represent a sequence of pixels in a ridge, and the green circles are the generated pores. Given a start point, pore distances are sampled from the reference distribution. In this example, new pores are created with distances $d_1 = 5$, $d_2 = 8$ and $d_3 = 10$

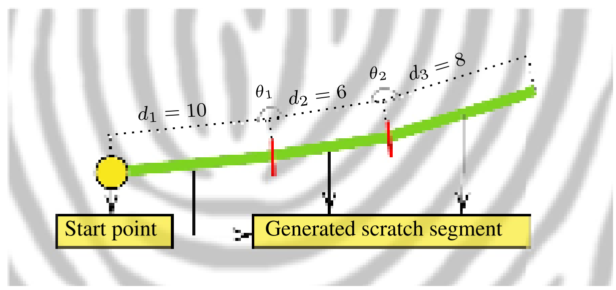

FIGURE 8 Illustration of the scratch generation process. Given a random start point, we draw n line segments with random length and random angle between them. In this example, there are three lines segments with lengths $l_1 = 10$, $l_2 = 6$ and $l_3 = 8$

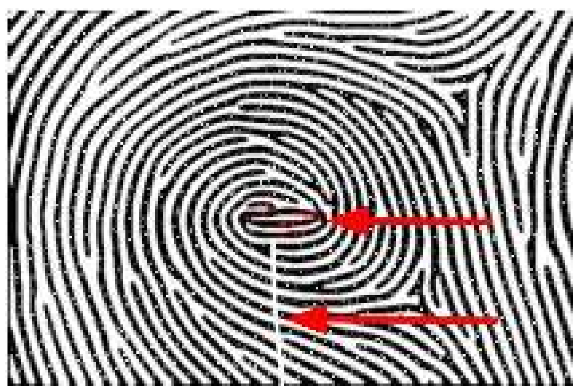

FIGURE 9 An L3 master fingerprint. Note the generated pores and scratch

or optical distortions caused by the acquisition sensor. Figure 11 illustrates this process.

Finally, we included a method to randomly drop some pores in different seed images from a single identity. This is a way of simulating the process of perspiration. In this work, we used a pore dropout rate of 3%. In Section 4.3 we show experimental results that clarify the dropout rate number.

2.4 | CycleGAN-based domain translation

To generate realistic fingerprint images, we learn to map the seed image domain into the real image domain of the chosen training database using CycleGAN [36].

CycleGAN is a viable solution for the task of translating two different domains as it does not require direct pairing between the training instances. Thus, our seed images do not need to be perfectly aligned to a real image in the training set.

When creating seed images for training, we seek to balance the number of real and synthetic samples. To increase the number of real training samples, besides performing horizontal flips, we take full real fingerprint images and apply the same acquisition simulation described in Section 2.3 (except pore dropout). Section 3.1.1 describes how we create the training set for this work.

We use the original CycleGAN architecture with 13 residual blocks [46] and input size 256 × 256 (seed images are resized to these dimensions). Besides CycleGAN's cycle consistency loss, we use the identity mapping loss [47] as it contains a regularising component that encourages the generator to map samples from the real fingerprint domain to themselves. We train our model using the Adam optimiser [48] with a learning rate of 0.0002 for 3 epochs.

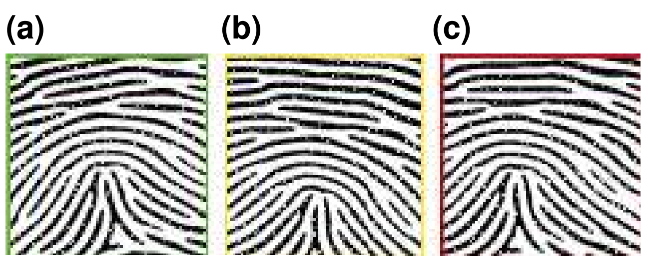

FIGURE 10 Seed images generated from a single L3 master fingerprint, presenting distinct shifts and rotations

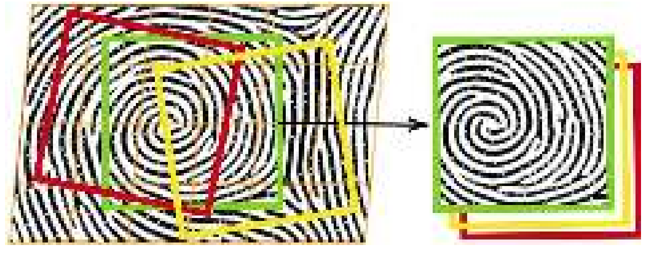

FIGURE 11 Illustration of an affine transformation over an L3 master fingerprint

At the beginning of the training, the weight for the identity mapping loss is 0 for high-resolution fingerprints and 0.1 for medium-resolution fingerprints. We iteratively increase the weight up to $0.7 \times \lambda$, where λ is the weight for the cycle consistency ($\lambda = 10$ in this work). We did this because CycleGAN tends to lose the master fingerprint identity, changing the ridge flow and the location of the minutiae.

After training, the outcome is a model that can translate any seed image into a realistic fingerprint, even if it was not seen during training. Examples of the inference using CycleGAN are presented in Figure 12.

3 | MULTIRESOLUTION SYNTHETIC FINGERPRINT DATABASE

In this session, we detail the creation of the Multiresolution Synthetic Fingerprint Generation database (MR-SF). Our database consists of L3 synthetic fingerprints (see Section 3.1) and L2 synthetic fingerprints (see Section 3.2). In addition to different resolutions, the two subsets also depict different acquisition sensors. The process of generating synthetic fingerprints using our approach can be applied with little effort for other resolutions and sensors.

3.1 | Level three synthetic fingerprint generation

We decided to replicate the structure of an existing high-resolution database to be able to perform a comparison in terms of image realism and fingerprint recognition performance. To this end, we used the PolyU database [10], which is divided in two subsets: DBI and DBII. DBI contains 320 × 240 images of cropped fingerprints, while DBII contains 640 × 480 images of full fingerprints. Both have 148 different identities, and each identity has 10 images acquired in two sessions (five per session). We use DBII for training purposes, as we need full fingerprints to perform the augmentation mentioned in

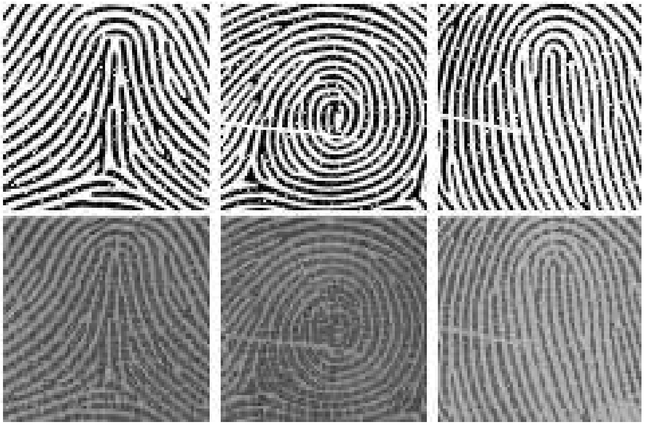

FIGURE 12 CycleGAN inference: seed images (top) and their respective high-resolution results (bottom)

Section 2.4. More details on how we create the training set for CycleGAN to generate synthetic L3 fingerprints are given in Section 3.1.1. We then use the obtained model to create synthetic datasets with the same structure as DBI, as described in Section 3.1.2.

3.1.1 | CycleGAN training set for L3 fingerprint generation

We created 296 L3 master fingerprints for training, which proved to be sufficient to map synthetic images to the real domain with CycleGAN.

We then flip these master fingerprints horizontally and perform the acquisition simulation described in Section 2.3 to create 5920 **seed images**. We create additional 5920 seed images by repeating the acquisition simulation with a larger cropping larger area (825 × 825 pixels) to cope with fingerprints with higher ridge frequency. Finally, we apply random elastic deformations [49] to the set of seed images to create 11,840 distorted images, totalling 23,680 training seeds.

For **real images**, we use the DBII subset of the PolyU database, which contains 1480 full fingerprint images. First, we flipped these images horizontally. After, we performed the acquisition simulation to create 10 samples per image, totalling 29,600 real training samples. Table 1 summarises the training images and the augmentation operations.

3.1.2 | L3-SF subset generation

We created five replicas of the PolyU DBI subset to establish a confidence interval in our recognition experiments. To do so, we generated 5 sets of 148 master fingerprints. After simulating fingerprint acquisition, we end up with 1480 seed images per set, totalling 7400 images for inference.

We then transform these seed images into real images using our CycleGAN model. To replicate the resolution and aspect ratio of PolyU DBI, we upscale the inferred images (256 × 256) by a factor of 2 using FSRCNN [39], crop the centre region of size 512 × 384 from the resized images (512 × 512), and resize the crops to the PolyU DBI resolution (320 × 240). These images, split into five subsets, compose our L3 synthetic fingerprint (L3-SF) database subset.

TABLE 1 Training set images for L3 fingerprint generation

	Synthetic	Real
Initial images	256 (master fingerprint)	1480
Flip horizontally	592 (master fingerprint)	2960
Shift, crop and rotation (10 variations)	5920 (seed image)	29,600
Higher ridge frequency cut	11,840 (seed image)	-
Elastic deformation [49]	23,680 (seed image)	-
Total images	23,680	29,600

3.2 | Level two synthetic fingerprint generation

To generate synthetic L2 fingerprints, we use our approach described in Section 2.4. We chose the NIST Special Database 300a (NIST SD300a) [50] with 500 pixels per inch (PPI) images as a training set. We also decided to replicate the structure of PolyU database, which allows direct comparisons between L3 and L2 subsets. We use all NIST SD300a images for training purposes. More details on how we create the training set for CycleGAN to generate synthetic L2 fingerprints are given in Section 3.2.1. We use the obtained model to create L2 synthetic fingerprints (L2-SF), as described in Section 3.2.2.

3.2.1 | CycleGAN training set for L2 fingerprint generation

The NIST SD300a is composed of rolled fingerprints, which capture a larger area than the L3-SF subset. Thus, we must cut subareas from NIST SD300a images so that the CycleGAN training is equivalent for L3-SF and L2-SF subsets. To do this, we upscale the 8842 NIST SD300a images by a factor of 3 using FSRCNN [39]. Then, we flip NIST SD300a images horizontally to create 17,684 augmented images. We crop the centre of these images to create 512 × 512 images. We create additional 17,684 images by cropping a larger area (700 × 700 pixels) to cope with fingerprints with higher ridge frequency, totalling 35,368 real training images. As the number of images on the NIST SD300a is larger than the PolyU DBII database, the acquisition simulation step was not necessary. For seed images, we use the same 23,680 augmented images described in Section 3.1.1. Table 2 summarises the training images and the augmentation operations.

3.2.2 | L2-SF subset generation

For the CycleGAN training, we use the same procedures described in Section 2.4, only modifying the weight for the identity mapping loss to 0.1 and iteratively increasing the weight up to $0.7 \times \lambda$, where λ is the weight for the cycle consistency ($\lambda = 10$ in this work). A higher identity mapping

TABLE 2 Training set images for L2 fingerprint generation

	Synthetic	Real
Initial images	256 (master fingerprint)	8842
Flip horizontally	592 (master fingerprint)	17,684
Shift, crop and rotation (10 variations)	5920 (seed image)	-
Higher ridge frequency cut	11,840 (seed image)	35,368
Elastic deformation [49]	23,680 (seed image)	-
Total images	23,680	35,368

loss was necessary because the ridge colours were becoming inverted in the inference.

After training, the outcome is a model that can translate any seed image into a realistic L2 fingerprint. We replicate the resolution and aspect ratio of the L3-SF subset by upscaling the inferred images (256 × 256) by a factor of 2 using FSRCNN [39], cropping the centre region of size 512 × 384 from the resized images (512 × 512), and resizing the crops to the L3-SF subset (320 × 240). These images, split into five subsets, compose our L2 synthetic fingerprint (L2-SF) database subset. A comparison between the inferences of medium and high-resolution fingerprints is presented in Figure 13.

4 | EXPERIMENTAL RESULTS

Section 4.1 presents a visual analysis of our high and medium resolution synthetic fingerprint images, including a visual comparison with other methods. A human perception experiment involving 60 volunteers is reported in Section 4.2. Section 4.3 presents a fingerprint recognition analysis using both PolyU DBI database and high-resolution images present in the L3-SF database subset. An experiment to validate multi-resolution cross-matching performance using our L2 and L3 synthetic fingerprints is reported in Section 4.4. Section 4.5 presents a quality analysis of our generated synthetic databases. In Section 4.6 we perform a minutiae experiment. Finally, Section 4.7 presents a sharpness experiment that we performed in real and synthetic fingerprints.

4.1 | Visual analysis

We visually inspected our results to evaluate our method's ability to create realistic high-resolution fingerprint images. We observed that our method creates pores at the indicated position in the seed image, thus preserving the fingerprint identity (see Figure 12). Open and closed pores were also observed. We noticed other L3 traits, such as distinct ridge contours and incipient ridges, and features that increase the reliability of fingerprint recognition [51, 52]. Figure 14 shows examples of L3 traits present in high-resolution images of L3-SF database subset.

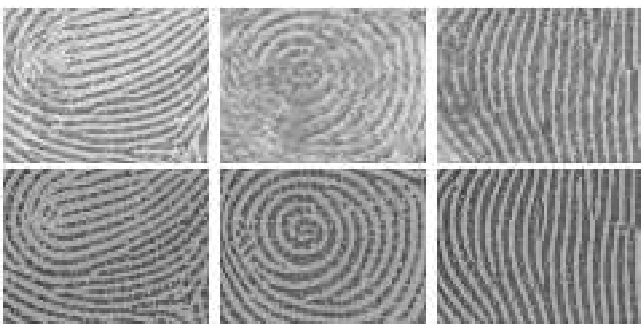

FIGURE 13 Synthetic medium-resolution fingerprints (top) and their respective high-resolution fingerprints (bottom)

Figure 15 shows a visual comparison between fingerprints generated by the proposed approach, by a publicly available SFinge demo [19], by a public model of Cao and Jain's method [22], by the Clarkson Fingerprint Generator (CFG) method [24], and by our implementation of Finger-GAN [21].

A visual comparison between real and our high-resolution synthetic fingerprints is shown in Figure 16. A visual comparison between our L2 synthetic fingerprints and cropped NIST SD300a fingerprints is presented in Figure 17. We selected these images randomly to provide an unbiased judgement. Note that high and medium synthetic images have different styles that are similar to the respective real datasets they were trained on. Finally, note that some Level 3 characteristics, such as the presence of pores, are visible in the images of the NIST SD300a and our L2 synthetic fingerprints. However, these characteristics are not visible in all images due to the variability in conditions present at the acquisition of rolled images on paper. Our CycleGAN model replicated this visual style distribution on L2 synthetic fingerprints.

Our high-resolution synthetic fingerprints contain different realistic aspects, such as pores with different sizes and shapes, and ridges with acute details and texturisation. Meanwhile, SFinGe [19] generates rectangular, single-sized pores and Finger-GAN [21] does not generate pores at all. Besides, Finger-GAN and Cao and Jain's method [22] produce unnatural ridge shapes. Cao and Jain's method also produces irregular minutiae patterns and irregular occlusions. Our medium-resolution synthetic fingerprints also contain different realistic aspects, such as different ridge contrasts similarly to a real fingerprint. As our L2 subset was created with a texture of an ink-and-paper database of rolled fingerprints, some ridge regions are joined together. This happens in a real acquisition due to a large ink quantity or a finger movement on the paper. Our new L2-SF and L3-SF subsets can be used in future research to improve fingerprint recognition algorithms using L3 and L2 sensors.

The PolyU database [10] is divided into two subsets: DBI and DBII, where both have 148 different identities, and each identity has 10 images acquired in two sessions (five per session). We noticed visual differences between PolyU sessions, such as contrast, variations in focus, and background colour. Although we replicate in our synthetic dataset the same structure as DBI, differences between sessions are not noticed. Figure 18 shows examples of PolyU and L3-SF distinct sessions.

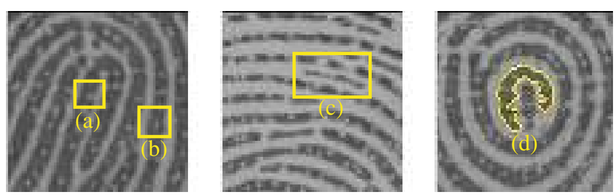

FIGURE 14 (a) and (b) show open and closed pores in an image of the L3-SF database subset. (c) shows an incipient ridge. (d) highlights an unique ridge shape

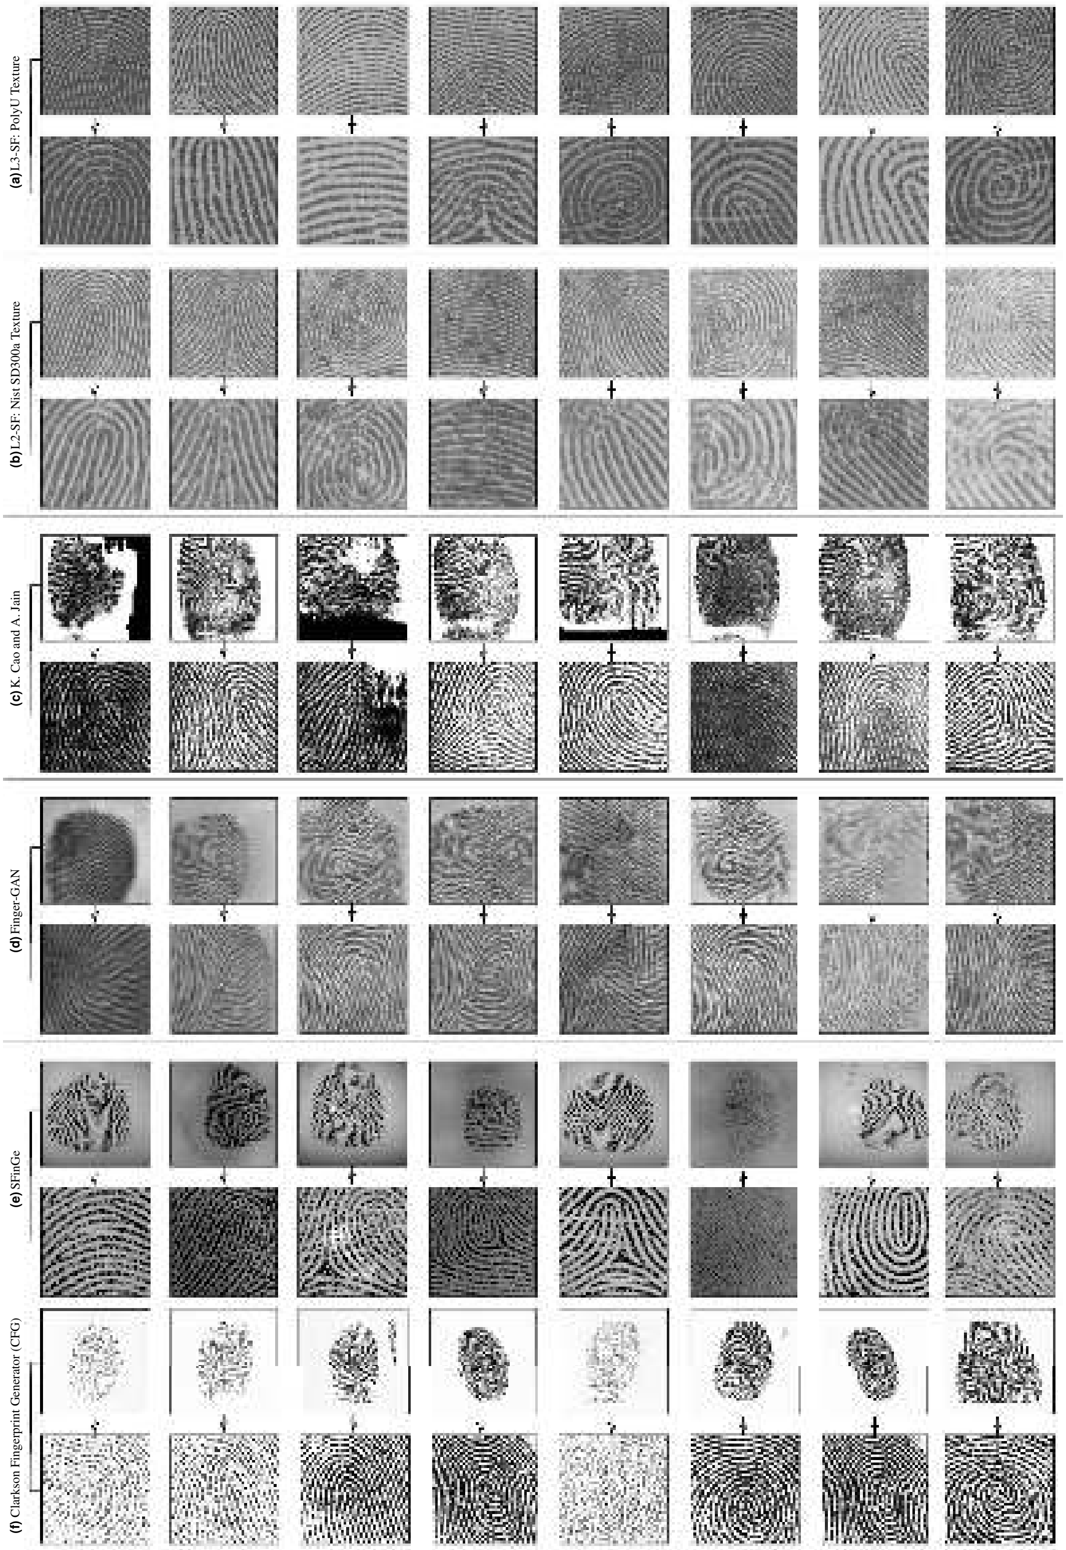

FIGURE 15 Visual comparison between (a) the proposed approach with the PolyU texture, (b) the proposed approach with the Nist SD300a texture, (c) Cao and Jain [22], (d) Finger-GAN [21], (e) SFinGe [19] and (f) Clarkson Fingerprint Generator (CFG) [24]

We include sweat pore annotations for 740 images to assist in pore detection research. Several studies are focussing on the detection of pores in L3 fingerprint images [53–55]. These require pore annotations to perform training with neural networks. However, they only focus on the PolyU image domain. Including other visual characteristics, a new dataset will help in

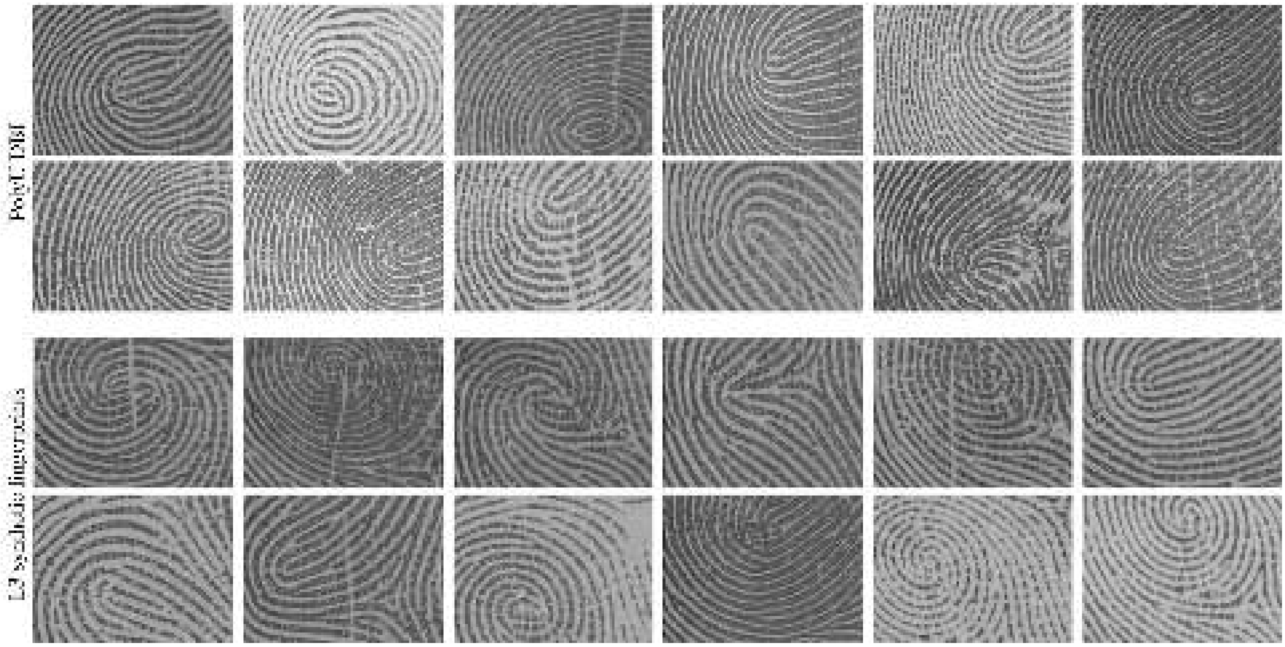

FIGURE 16 Visual comparison between real fingerprints from PolyU DBI (top) and high-resolution synthetic ones from L3-SF (bottom)

FIGURE 17 Visual comparison between real fingerprints from cropped NIST SD300a (top) and L2 synthetic fingerprints from L2-SF subset (bottom)

the pore detection task in different image domains, improving the biometric recognition process in multiple sensors. Figure 19 shows examples of L3-SF database annotations.

4.2 | Human perception experiment

Humans generate conscious and unconscious inferences from stimulus on the visual cortex [56]. These inferences provide a judgement mechanism about different characteristics in images, where each person perceives and interprets stimuli in different ways and strategies. Based on this fact, to evaluate people's visual perception of real and high-resolution synthetic images, we performed an experiment similar to the 'real versus fake' approach popularly used on Amazon Mechanical Turk (AMT). In this experiment, we randomly show five PolyU images and five L3-SF high-resolution fingerprints to the participants, which annotate the five images they consider false.

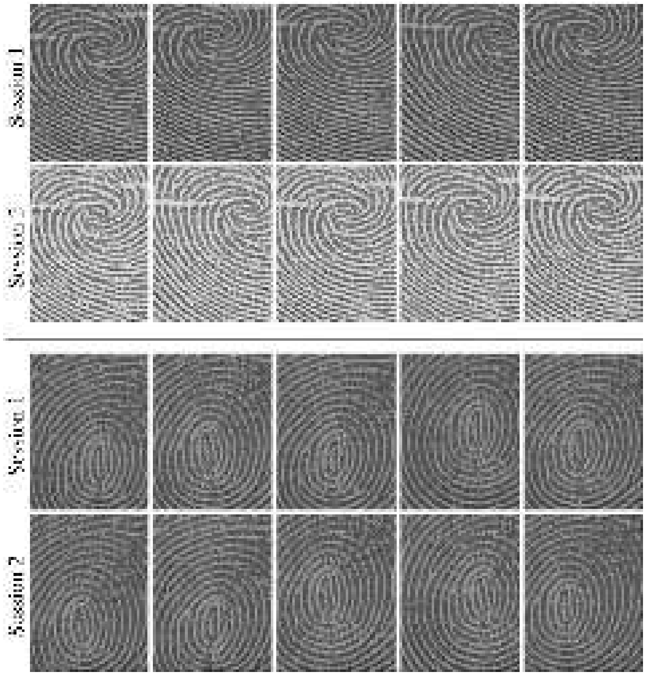

FIGURE 18 Visual comparison between two sessions of real fingerprints from PolyU DBI (top) and synthetic ones from L3-SF (bottom)

Each participant repeats this task 10 times. This methodology forces the participants to identify visual characteristics and develop strategies to classify the two image types.

Sixty volunteers participated in this experiment, which allowed us to analyse the overall human behaviour in this classification task. When participants fail to notice any visual characteristics that help discriminating real and high-resolution synthetic images, their classification accuracy is close to random (50%). Figure 20 shows the histogram of human classification error. As can be seen, the average misclassification rate for all participants was 45.65%. This result shows that the participants had difficulties distinguishing the two classes and highlights the realism of our synthetic fingerprints.

4.3 | Fingerprint recognition analysis

The goal of this experiment is to compare real and high-resolution synthetic images in terms of recognition performance. Ideally, both should have close results. To carry out this comparison, first we utilise Bozorth3 [57], which is a minutiae-based fingerprint matching approach. Our focus was not to get the best matching accuracy but to compare similarities between the databases. Bozorth3 proved to be useful in this task. To perform pore-based fingerprint matching, we utilise Segundo and Lemes' approach [3]. For these analyses, we use the same protocol proposed by Liu et al. [58] for PolyU DBI and all L3-SF replicas, as all of them have the same configuration (148 subjects, 2 sessions, 5 images per subject per session, 320 × 240 images).

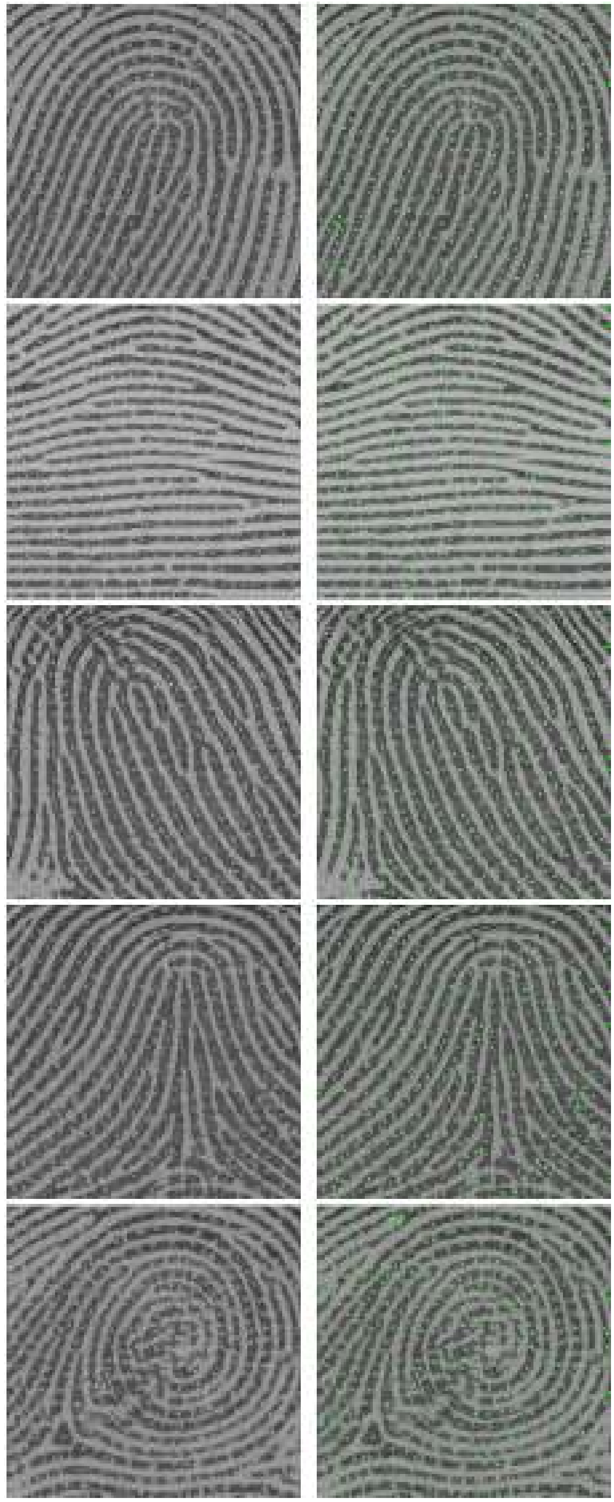

FIGURE 19 Pore annotation examples. On the left, images of the L3-SF dataset. On the right, the correspondent fingerprint with the annotation in green circles

We computed the False Non-Match Rate (FNMR) and the False Match Rate (FMR) using different threshold values for both matching approaches in all test databases. With the obtained FNMR and FMR values, we plot a Receiver Operating

Characteristic (ROC) curve for PolyU DBI. For our synthetic databases, we plot the average ROC with a 95% confidence interval. These plots are shown in Figures 21 and 22.

The Equal Error Rate (EER) values for minutiae-based matching are 23.84% for real images and 18.79% for synthetic ones. The EER values for pore-based matching are 3.37% for real images and 3.80% for synthetic ones. These results validate the L3-SF subset as a realistic database, since existing recognition methods were successful without any adjustments. Besides, the accuracy in both real and synthetic datasets was very close. The larger gap in minutiae-based matching can be caused by a difference in minutiae distribution in real and high-resolution synthetic images. At the moment, the minutiae distribution is only controlled by Anguli.

To study the variability of our synthetic fingerprint generation method and to check the matching results when pores close and open during perspiration, we repeated the experiment in our five distinct sets three times, randomly drooping 0%, 3%, and 5% pores on each experiment. The EER values for pore-based matching are 3.37% for real images and 3.28%, 3.80%, and 3.88% for synthetic ones (randomly drooping 0%, 3%, and 5% pores on each experiment). We plot a Receiver Operating Characteristic (ROC) curve for PolyU DBI and all L3-SF variations. This plot is shown in Figure 23.

The results confirm that the 3% dropout is a suitable factor to apply in our method. However, the results also show that this parameter can be used to increase or decrease the level of challenge when performing biometric recognition

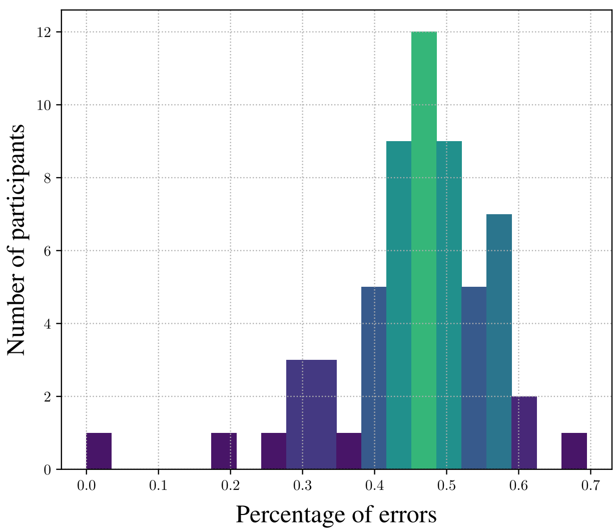

FIGURE 20 Histogram of the human perception experiment. Several participants had a classification accuracy close to random (50%)

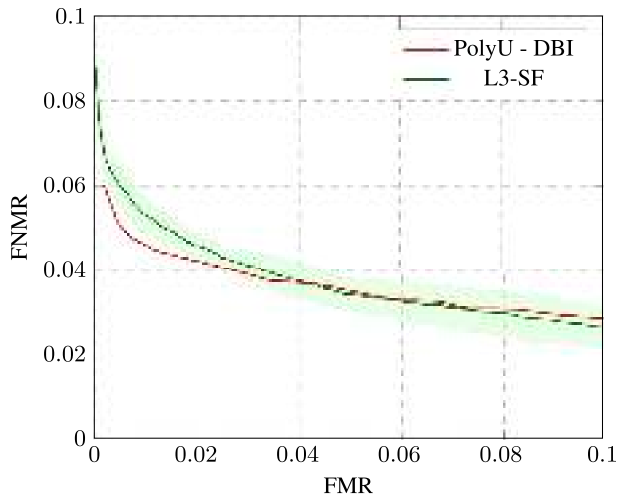

FIGURE 22 ROC curves from a pore-based matcher for PolyU DBI and L3-SF database. For the latter, we show the average ROC and a 95% confidence for its five subsets

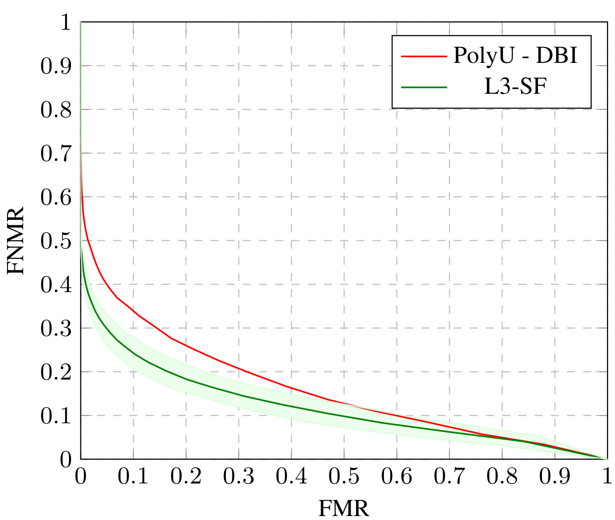

FIGURE 21 ROC curves from a minutiae-based matcher for PolyU DBI and L3-SF. For the latter, we show the average ROC and a 95% confidence for its five subsets

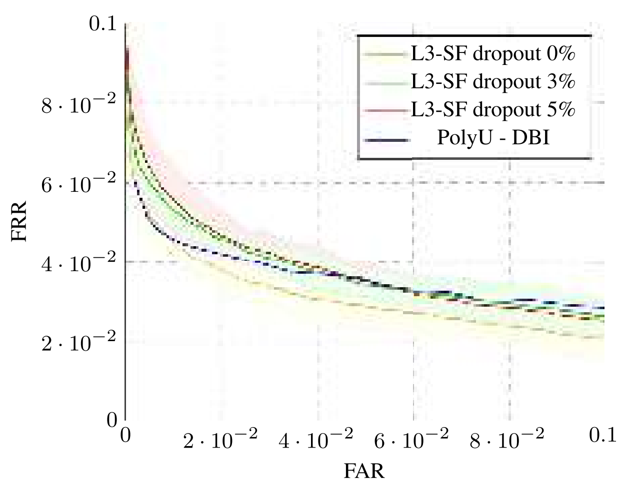

FIGURE 23 ROC curves from a pore-based matcher for PolyU DBI and L3-SF database. For L3-SF, we randomly remove 0%, 3%, and 5% pores on each experiment. For the latter, we show the average ROC and a 95% confidence for its five subsets

with pores, which can be helpful in other research studies with different objectives.

4.4 | Cross-matching multiresolution experiment

This experiment's objective is to perform the fingerprint recognition analysis, verifying the matching performance when using different sensors (cross-sensor) compared to matching with the same sensors (intra-sensor). When matching coarsely aligned images, it is possible to observe the performance differences caused by texture, optical distortions, skin elasticity and rotations. For this reason, we use images without translation variations (see Section 4.4.1 and 4.4.2).

4.4.1 | Real data preparation

For the matching analysis, we used the NIST Special Database 301a (NIST SD301a) [59]. This database contains subsets of several sensors. We use the subsets dryrun-A, dryrun-B, dryrun-F and dryrun-G. Dryrun-A and dryrun-B are L2 rolled fingerprints captured with the Crossmatch Guardian 300 sensor. Dryrun-F and dryrun-G are L2 plain fingerprints captured with the Jenetric LIVETOUCH QUATTRO sensor. We selected 93 fingerprints that are present in all of these subsets to perform the matching analysis. Table 3 shows the intra- and cross-sensor matchings that we performed in real data images.

Before matching any pair of images, we eliminate translation variations by finding keypoint correspondences using the ASIFT [60] algorithm and aligning them by the centre of mass of these keypoints. We also crop the centre of the aligned images so that the finger area is equivalent to the area in our MR-SF fingerprints. To do this, we scale the MR-SF image size (320 × 240) by the ratio of the mean ridge frequencies from NIST SD301a and PolyU DBI. The NIST SD301a images' cropping area was defined as 166 × 124.

4.4.2 | Synthetic data preparation

To compare medium and high-resolution synthetic fingerprints, we used the same master fingerprints created to generate the five replicas containing 148 identities described in Section 3.1.2. We use the models generated by CycleGAN to infer the same identities in the textures of the PolyU and NIST SD300a databases. The inference generated by the two texture styles was passed through the same procedures described in Section 2.3. However, during acquisition simulation, we set σ_x and σ_y to 0 to eliminate translation variations.

TABLE 3 Intra- and cross-sensor matching and alignment

Intra-sensor	Cross-sensor
Dryrun-A–dryrun-B	Dryrun-A–dryrun-F
Dryrun-F–dryrun-G	Dryrun-A–dryrun-G
-	Dryrun-B–dryrun-F
-	Dryrun-B–dryrun-G

4.4.3 | Multi-resolution cross-matching matching analysis

The goal of this experiment is to compare the intra-sensor and cross-sensor in terms of recognition performance. In this scenario, the cross-sensor is expected to have worse performance in fingerprint matching. We use Bozorth3 [57] to perform the matching, where Bozorth3 proved to be consistent in comparing the databases. We modified the protocol proposed by Liu et al. [58]. For genuine comparisons, the first fingerprint of the first session of the first sensor is compared with the first fingerprint of the same individual from the second session of the second sensor for all identities. For impostor comparisons, the first fingerprint of the first session of the first sensor is compared with all images of all other individuals in the second session of the second sensor. This protocol change was necessary because the NIST 301a contains only one image pair per sensor of the same identity. This protocol was used in the intra-sensor and cross-sensor matching on the MR-SF and NIST 301a images. Table 4 summarises the matching comparisons performed.

We computed ROC curve for each matching comparison. For our high and medium resolution synthetic fingerprints, we plot the average ROC with a 95% confidence interval. For the real NIST 301a images, we plot the average ROC with a 95% confidence interval for the cross-matching comparisons. These plots are shown in Figures 24 and 25. It is important to mention that the L2-SF subset was not made to be equivalent to NIST 301a images in matching performance. Our main goal is to illustrate the complexity of a cross-dataset experiment using our synthetic dataset.

In the synthetic images, the EER values are 12.69% for the L3-SF PolyU texturised images and 13.70% for L2-SF NIST SD300a texturised images. The EER on the cross-matching is 33.03%. The difference between cross- and intra-matching is close to 20%.

In the real NIST 301a images, the EER values are 27.46% for (dryrun-A and dryrun-B) matching and 28.04% for (dryrun-F and dryrun-G). The EER on the cross-matching is 31.53%. The difference between cross- and intra-matching was about 4%.

These results show that our synthetic database behaved as expected. The performance was worse on cross-matching for both synthetic and real images. Since all NIST 301a images are L2, the difference between cross- and intra-matching EER was much smaller than this difference in our synthetic database. This was also expected since the synthetic dataset consists of L2 and L3 subsets, making the cross-matching even more challenging. Still, this result indicates that our approach can be used to create challenging cross-sensor scenarios, which in turn can support fingerprint recognition research.

TABLE 4 Matching comparisons

Matching comparisons	Genuine	Impostor	Intra-sensor	Cross-sensor
Dryrun-A and dryrun-B	93	8556	x	
Dryrun-F and dryrun-G	93	8556	x	
Dryrun-A and dryrun-F	93	8556		x
Dryrun-A and dryrun-G	93	8556		x
Dryrun-B and dryrun-F	93	8556		x
Dryrun-B and dryrun-G	93	8556		x
L3-SF: PolyU texture session 1 and 2	148	21,756	x	
L2-SF: NIST 300a texture session 1 and 2	148	21,756	x	
L3-SF and L2-SF: PolyU and NIST 300a textures	148	21,756		x

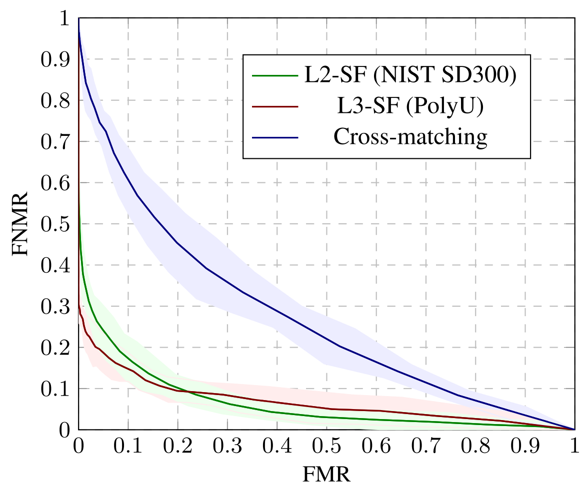

FIGURE 24 ROC curves from a minutiae-based matcher for L3-SF (PolyU), L2-SF (NIST 300a) and the cross-matching. We show the average ROC and a 95% confidence for its five subsets

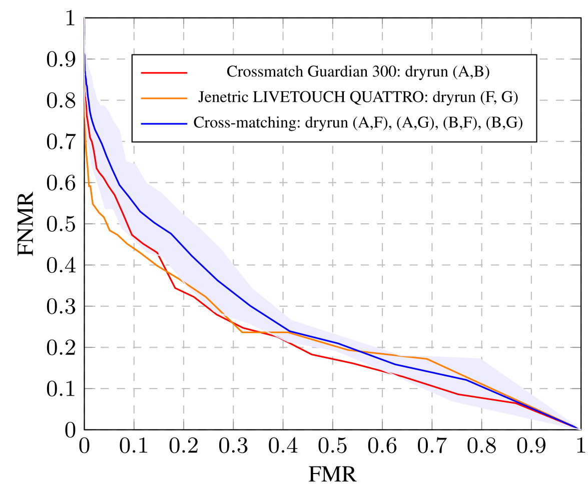

FIGURE 25 ROC curves from a minutiae-based matcher for Crossmatch Guardian 300 (dryrun-A and dryrun-B), Jenetric LIVETOUCH QUATTRO (dryrun-F and dryrun-G) and the cross-matching dryrun (A,F), (A,G), (B,F), and (B,G). For the latter, we show the average ROC and a 95% confidence for its four cross-matching comparisons

4.5 | Quality analysis

The Frechet Inception Distance (FID) metric [61] is a widely used technique used to evaluate the quality of generative models in machine learning. However, many FID implementations have problems performing resizing operations. The work of Parmar et al. [62] solves this problem by defining the best kernel interpolation to use. We used this implementation to directly compare the PolyU DBI base with the five variations of the L3-SF and the NIST 300a base with the five variations of the L2-SF. As the area of NIST 300a is larger than the area of L2-SF, even though both are 500 dpi, we performed crops in the centre of the NIST 300a images containing the same area as the L2-SF images. It was unnecessary to perform crops on the PolyU DBI images as it already contains the same area of interest as the L3-SF. With this, the FID value for L3-SF was **47.00**, and for L2-SF, it was **97.28**. The smaller this value, the more realistic these images are to images of real databases. We noticed that the FID for L2-SF was higher than L3-SF, indicating that the texturing task in low-resolution images is also challenging. It is important to mention that the FID level is not a directly impacting factor for recognising fingerprints. Comparing the face synthesis results of the current state-of-the-art GAN method StyleGAN 3 in a similar size dataset METFACES-U [63], the FID value was 18.75. This result shows that still exists a gap in fingerprint synthesis compared to face synthesis.

In addition to FID analysis, we used NFIQ [64] and NFIQ 2 [65] software to verify the quality of the generated synthetic images compared to the real images. These tools are specific to fingerprints. The original NFIQ generates outputs from 1 to 5, where 1 is the best possible fingerprint quality and 5 is the worst quality. We convert all images from all databases to 500 dpi, as required by NFIQ and NFIQ 2. NFIQ 2 generates a value from 0 to 100, where the higher this value, the higher the fingerprint quality.

For the analysis with the NFIQ, we generated five-folds without repetition containing 1480 images from the NIST 300a dataset, this number being the same as present in the PolyU DBI and L2-SF databases. We obtain the NFIQ metric information for the five NIST 300a folds, the five variations of L3-SF and L2-SF, and the PolyU DBI. We repeated the same operations for the NFIQ 2 analysis. We report average values for datasets with multiple subsets. The quantitative results are presented in Table 5 and in Figures 26–28.

Ideally, we would like to have a similar distribution between authentic and synthetic databases concerning the NFIQ and NFIQ 2 scores. This evaluation also provides indications regarding the style variation generated by the texturisation process using the generated models described in Sections 3.1 and 3.2. Although the mean values of the NFIQ and NFIQ 2 scores of the synthetic databases are close to the real databases, it is evident that there is still an opportunity for improvement in the texturing models to replicate the quality distribution of real databases.

4.6 | Minutiae experiment

To perform an analysis regarding the probability of minutiae occurrences in our synthetic fingerprints, we run mindtct [57] to extract the number of minutiae per image generated. For comparative purposes, we also performed this analysis with the PolyU DBI and NIST 300a real image databases (the last one, we cut the images to the same area as the L2-SF images). Ideally, synthetic databases should follow the same occurrence probability distribution as real databases. As is already known from the literature [66], the number of minutiae generated by SFinge, which Anguli is based on, is similar to the minutiae distribution of the NIST Special Database 4 (SD04) real fingerprint database [67]. In this way, this experiment aims to verify if our neural texturing model does not modify the seed images to alter the texturing distribution. The histogram of Figure 29 contains the results of our experimentation. The high-resolution databases obtained very similar results. However, we noticed a slight difference in the minutiae distribution in the medium-resolution databases.

In this experiment, it is evident that PolyU DBI and L3-SF have a very similar minutiae profile. This factor also validates the realism of the generated synthetic images. In medium-resolution images, the challenge of finding minutiae is more significant, as noise or the union of ridges is a difficult task for the detector. However, the medium resolution synthetic database still presented an adequate probability distribution than an authentic one.

4.7 | Sharpness experiment

The L3-SF and L2-SF use a CycleGAN model, where the inference generates 256 × 256 images. We improved the quality of the images by upscaling them with FSRCNN [39]. However, authentic images tend to be sharper than those generated by generative neural networks. Therefore, it is essential to assess the difference in sharpness between real and synthetic images. Figure 30 exemplifies this manner, where we compare our synthetic databases with real databases. In addition to this visual analysis, we used the cumulative probability of blur detection (CPBD) metric [68] to determine the sharpness levels. The results of this analysis are found in Table 6. We noticed that both

TABLE 5 Fingerprint metrics for the PolyU DBI, L3-SF, NIST 300a and L2-SF images

	PolyU DBI				L3-SF (5 variations)				NIST 300a				L2-SF (5 variations)			
Measure	Mean	STD	Kurtosis	Skewness	Mean	STD	Kurtosis	Skewness	Mean	STD	Kurtosis	Skewness	Mean	STD	Kurtosis	Skewness
NFIQ	2.23	0.85	−0.39	0.08	2.65	0.87	−0.30	−0.67	3.48	0.89	−0.58	1.05	3.43	0.83	−0.10	1.37
NFIQ 2.0	37.27	8.53	3.86	−0.93	35.17	6.92	2.48	−0.57	17.37	17.75	−1.73	0.19	24.31	18.28	−1.50	−0.38

Note: NFIQ score has a range of [1, 5] where 1 it's in a higher quality. NFIQ 2.0 score has a range of [0, 100], where the higher the number, better the quality.

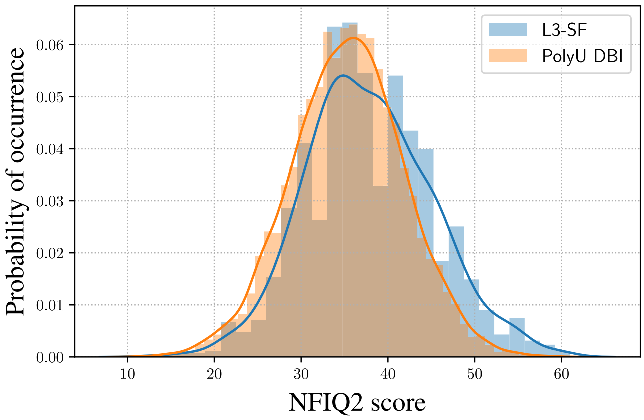

FIGURE 26 Probability histogram of L3-SF and PolyU DBI databases

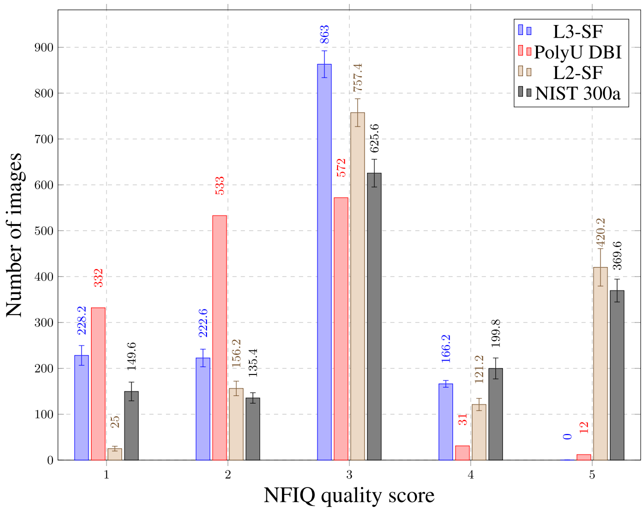

FIGURE 27 Quality assessment of the original images, binarised images and the reconstructed binarised images NFIQ

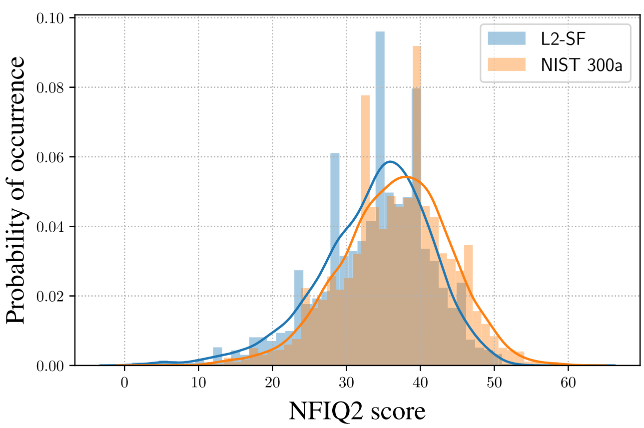

FIGURE 28 Probability histogram of L2-SF and NIST 300a databases

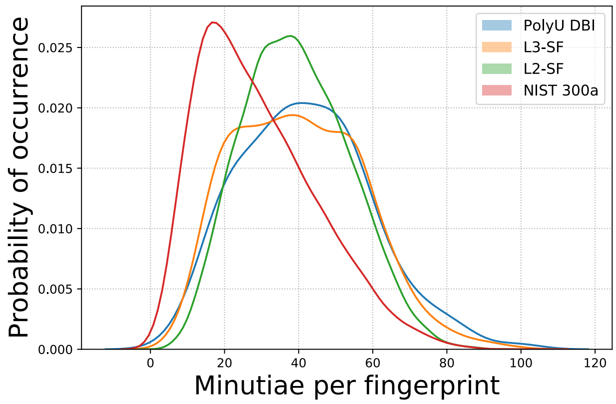

FIGURE 29 Probability distribution of occurrences of minutiae in the synthetic databases L3-SF, L2-SF and the real PolyU DBI and NIST 300a databases

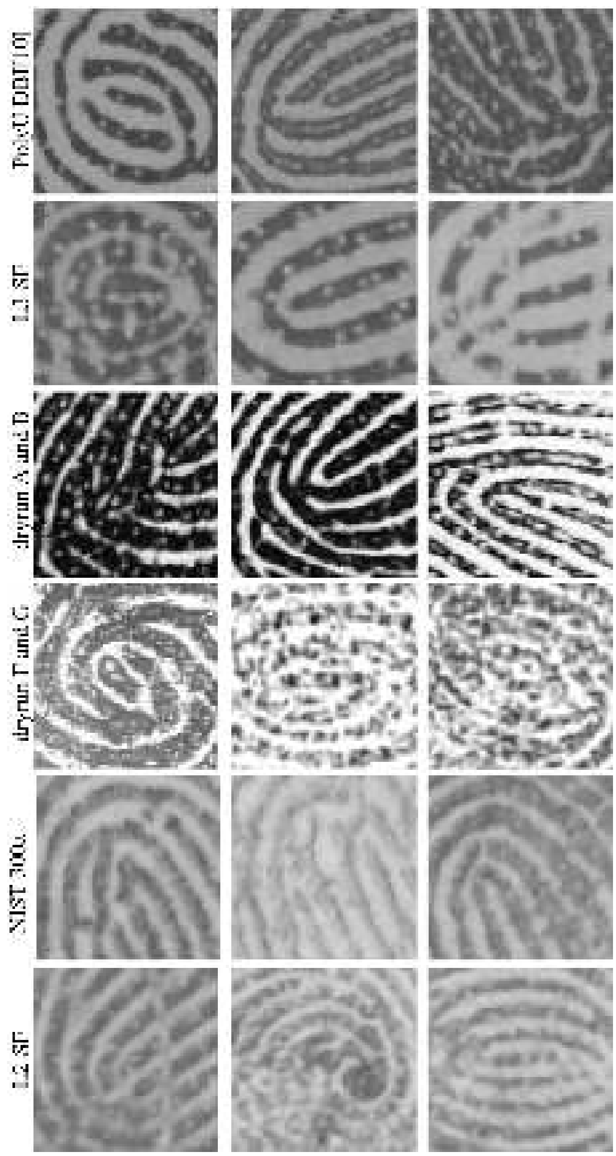

FIGURE 30 Cropped samples from PolyU DBI, L3-SF, NIST 301a (dryrun A, B, F, G), NIST 300a and L2-SF

TABLE 6 CPBD mean for our synthetic databases and for PolyU DBI, NIST 300a, and NIST 301a (dryrun-A,B,F,G) databases

	High-resolution				Medium-resolution			
	L3-SF	**PolyU DBI**	**L2-SF**	**NIST 300a**	**Dryrun-A**	**Dryrun-B**	**Dryrun-F**	**Dryrun-G**
CPBD	0.1662	0.4091	0.0732	0.1808	0.2131	0.2004	0.1585	0.0992

Note: The higher the CPBD value, the better.

synthetic databases are very different from PolyU DBI and NIST 300a and 301a concerning sharpness.

The best analysis we can do on medium-resolution images is to compare L2-SF with the NIST 300a images because we used this real database to generate our inference models. However, we noticed that the NIST 301a dryrun-G variation obtained a CPBD value similar to our synthetic database. These results also show sharpness variations in medium-resolution sensors, a topic that can be further explored in future studies.

5 | CONCLUSIONS

We presented an approach to generate realistic, multiresolution and multisensor synthetic fingerprints. We trained a CycleGAN using real fingerprint images and handcrafted seed images to create a model capable of translating between these two image domains while preserving all identification cues (e.g. ridges, minutiae, and pores). Using this approach, we created the MR-SF database containing two subsets. The first subset (L3-SF) consists of high-resolution images containing the PolyU DBI database's same characteristics. The second subset (L2-SF) consists of medium resolution images containing the same texture characteristics of the NIST SD300a database. More importantly, the MR-SF database allows further studies in the field of fingerprint biometrics without raising privacy-related legal issues. Our experimental results show that MR-SF images can be used by existing fingerprint recognition methods without any adjustments and achieve similar recognition performance. We performed a human perception experiment with 60 volunteers, which evidenced our high-resolution synthetic images' realism thanks to nearly random human classification performance. We visually compared our results with the literature's best performing studies to highlight the quality enhancement over existing studies. Finally, we generate L2 fingerprints to evaluate cross-sensor fingerprint recognition performance. Our experiment showed that our synthetic L2 and L3 databases could be used in recognition experiments using multisensor and multiresolution fingerprints. These images could be used to support the development of new cross-sensor matching algorithms capable of handling L2 and L3 fingerprints. In addition, our database can help reduce overfitting, especially in low training data scenarios. One example is the use of synthetic fingerprint images for pore detection. The PolyU-HRF DBI subset has only 30 images with annotated pores, a low quantity that can affect the performance in this task. Methods capable of generating synthetic samples can diversify the training.

Our main limitation is the lack of control over the distribution of minutiae, which is currently managed by Anguli. This limitation inhibited our ability to reduce the gap in recognition performance between real and synthetic datasets for minutiae-based fingerprint matchers.

ACKNOWLEDGEMENTS

This research study was funded by CNPq (http://cnpq.br/). The Titan Xp used for this research study was donated by the NVIDIA Corporation.

CONFLICT OF INTEREST

The authors declare no conflict of interest.

PERMISSION TO REPRODUCE MATERIALS FROM OTHER SOURCES

None: all materials were generated by the authors.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available in MR-SF/ at https://andrewyzy.github.io/MR-SF/, reference number MR-02.

ORCID

Andre Brasil Vieira Wyzykowski https://orcid.org/0000-0002-2279-4177

REFERENCES

1. Maltoni, D., et al.: Handbook of Fingerprint Recognition. Springer, London (2009). https://doi.org/10.1007/978-1-84882-254-2
2. Jain, A.K., Kumar, A.: Biometrics of next generation: an overview. Second Gen. Biometr. 12(1), 2–3 (2010)
3. Pamplona Segundo, M., de Paula Lemes, R.: Pore-based ridge reconstruction for fingerprint recognition. In: 2015 IEEE CVPR Workshops, pp. 128–133 (2015)
4. Anand, V., Kanhangad, V.: Porenet: CNN-based pore descriptor for high-resolution fingerprint recognition. CoRR abs/1905.06981 (2019). http://arxiv.org/abs/1905.06981
5. Nguyen, D., Jain, A.K.: End-to-end pore extraction and matching in latent fingerprints: going beyond minutiae. CoRR abs/1905.11472 (2019). http://arxiv.org/abs/1905.11472
6. Espinoza, M., Champod, C.: Using the number of pores on fingerprint images to detect spoofing attacks. In: 2011 International Conference on Hand-Based Biometrics, ICHB 2011—Proceedings (2011)
7. Marcialis, G.L., Roli, F., Tidu, A.: Analysis of fingerprint pores for vitality detection. In: 2010 20th International Conference on Pattern Recognition, pp. 1289–1292 (2010)
8. da Silva, M., Marana, A., Paulino, A.: On the importance of using high resolution images, third level features and sequence of images for fingerprint spoof detection. In: IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 1807–1811 (2015)
9. N. I. of Standards and Technology: Nist special database 30 (2002). https://www.nist.gov/srd/nist-special-database-30
10. Zhao, Q., et al.: Direct pore matching for fingerprint recognition. In: Tistarelli, M., Nixon, M. S. (eds.) Advances in Biometrics, pp. 597–606. Springer, Berlin (2009)

11. Anand, V., Kanhangad, V.: Cross-sensor pore detection in high-resolution fingerprint images using unsupervised domain adaptation. arXiv preprint arXiv:1908.10701 (2019)
12. Anand, V., Kanhangad, V.: Pore-based indexing for fingerprints acquired using high-resolution sensors. Pattern Anal. Appl. 23(1), 429–441 (2020). https://doi.org/10.1007/s10044-019-00805-3
13. AlShehri, H., et al.: A large-scale study of fingerprint matching systems for sensor interoperability problem. Sensors. 18(4), 1008 (2018). https://doi.org/10.3390/s18041008
14. Lin, C., Kumar, A.: Improving cross sensor interoperability for fingerprint identification. In: 2016 23rd International Conference on Pattern Recognition (ICPR), pp. 943–948. IEEE (2016)
15. Alshehri, H., et al.: Cross-sensor fingerprint matching method based on orientation, gradient, and gabor-hog descriptors with score level fusion. IEEE Access 6, 28951–28968 (2018). https://doi.org/10.1109/access.2018.2840330
16. Butler, D.J., et al.: A naturalistic open source movie for optical flow evaluation. In: European Conference on Computer Vision, pp. 611–625. Springer, Berlin (2012)
17. Wu, Y., et al.: Building generalizable agents with a realistic and rich 3D environment. arXiv preprint arXiv:1801.02209 (2018)
18. Wrenninge, M., Unger, J.: Synscapes: a photorealistic synthetic dataset for street scene parsing. arXiv preprint arXiv:1810.08705 (2018)
19. Cappelli, R.: Sfinge: an approach to synthetic fingerprint generation. In: International Workshop on Biometric Technologies (2004)
20. Ansari, A.H.: Generation and storage of large synthetic fingerprint database. ME Thesis (2011)
21. Minaee, S., Abdolrashidi, A.: Finger-gan: generating realistic fingerprint images using connectivity imposed gan. arXiv preprint arXiv:1812.10482 (2018)
22. Cao, K., Jain, A.: Fingerprint synthesis: evaluating fingerprint search at scale. In: 2018 International Conference on Biometrics (ICB) (2018)
23. Goodfellow, I., et al.: Generative adversarial nets. In: Advances in Neural Information Processing Systems, pp. 2672–2680 (2014)
24. Bahmani, K., et al.: High fidelity fingerprint generation: quality, uniqueness, and privacy. In: 2021 IEEE International Conference on Image Processing (ICIP), pp. 3018–3022 (2021)
25. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4401–4410 (2019)
26. de Paula Lemes, R., et al.: Dynamic pore filtering for keypoint detection applied to newborn authentication. In: 2014 22nd International Conference on Pattern Recognition, pp. 1698–1703. IEEE (2014)
27. Wyzykowski, A.B.V., Pamplona, M., de Paula Lemes, R.: Level three synthetic fingerprint generation. In: 2020 25rd International Conference on Pattern Recognition (ICPR). IEEE (2020)
28. Lu, Y., et al.: Ultrasonic fingerprint sensor using a piezoelectric micromachined ultrasonic transducer array integrated with complementary metal oxide semiconductor electronics. Appl. Phys. Lett. 106(26), 263503 (2015). https://doi.org/10.1063/1.4922915
29. Tang, H.-Y., et al.: 3-d ultrasonic fingerprint sensor-on-a-chip. IEEE J. Solid State Circ. 51(11), 2522–2533 (2016). https://doi.org/10.1109/jssc.2016.2604291
30. Iula, A.: Ultrasound systems for biometric recognition. Sensors 19(10), 2317 (2019). https://doi.org/10.3390/s19102317
31. Yin, X., Zhu, Y., Hu, J.: Contactless fingerprint recognition based on global minutia topology and loose genetic algorithm. IEEE Trans. Inf. Forensics Secur. 15, 28–41 (2019). https://doi.org/10.1109/tifs.2019.2918083
32. Tan, H., Kumar, A.: Towards more accurate contactless fingerprint minutiae extraction and pose-invariant matching. IEEE Trans. Inf. Forensics Secur. 15, 3924–3937 (2020). https://doi.org/10.1109/tifs.2020.3001732
33. Grosz, S.A., et al.: C2cl: contact to contactless fingerprint matching. IEEE Trans. Inf. Forensics Secur. 17, 196–210 (2021). https://doi.org/10.1109/tifs.2021.3134867
34. Ross, A., Jain, A.: Biometric sensor interoperability: a case study in fingerprints. In: International Workshop on Biometric Authentication, pp. 134–145. SpringerBerlin (2004)
35. Lugini, L., et al.: Interoperability in fingerprint recognition: a large-scale empirical study. In: 2013 43rd Annual IEEE/IFIP Conference on Dependable Systems and Networks Workshop (DSN-W), pp. 1–6. IEEE (2013)
36. Zhu, J., et al.: Unpaired image-to-image translation using cycle-consistent adversarial networks. CoRR abs/1703.10593 (2017)
37. van Mensvoort, M.: The fingerprints-world-map! (2015). http://fingerprints.handresearch.com/dermatoglyphics/fingerprints-world-map-whorls-loops-arches.htm
38. Wang, P.S.P. (ed.): Pattern Recognition, Machine Intelligence and Biometrics. Springer, Berlin (2011). https://doi.org/10.1007/978-3-642-22407-2
39. Dong, C., Loy, C.C., Tang, X.: Accelerating the super-resolution convolutional neural network. In: European Conference on Computer Vision, pp. 391–407. Springer, Berlin (2016)
40. Zhang, T.Y., Suen, C.Y.: A fast parallel algorithm for thinning digital patterns. Commun. ACM 27(3), 236–239 (1984). https://doi.org/10.1145/357994.358023
41. Lowe, D.G.: Object recognition from local scale-invariant features. In: Proceedings of the International Conference on Computer Vision-Volume 2 - Volume 2, Ser. ICCV '99, pp. 1150. IEEE Computer SocietyWashington (1999). http://dl.acm.org/citation.cfm?id=850924.851523
42. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. Int. J. Comput. Vis. 60(2), 91–110 (2004). https://doi.org/10.1023/b:visi.0000029664.99615.94
43. Rublee, E., et al.: Orb: an efficient alternative to sift or surf. In: Proceedings of the 2011 International Conference on Computer Vision, pp. 2564–2571 (2011). https://doi.org/10.1109/ICCV.2011.6126544
44. Yu, G., Morel, J.-M.: Asift: an algorithm for fully affine invariant comparison. Image Process. Line. 1, 11–38 (2011). https://doi.org/10.5201/ipol.2011.my-asift
45. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM 24(6), 381–395 (1981). https://doi.org/10.1145/358669.358692
46. He, K., et al.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
47. Taigman, Y., Polyak, A., Wolf, L.: Unsupervised cross-domain image generation. arXiv preprint arXiv:1611.02200 (2016)
48. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
49. Simard, P.Y., et al.: Best practices for convolutional neural networks applied to visual document analysis. In: Seventh ICDAR, 2003. Proceedings, pp. 958–963 (2003)
50. Fiumara, G., et al.: NIST special database 300: uncompressed plain and rolled images from fingerprint cards. Technical Report (2018). https://doi.org/10.6028/nist.tn.1993
51. Lee, W., et al.: Partial fingerprint matching using minutiae and ridge shape features for small fingerprint scanners. EXWA 87(C), 183–198 (2017). https://doi.org/10.1016/j.eswa.2017.06.019
52. Jain, A.K., Feng, J.: Latent fingerprint matching. IEEE TPAMI 33(1), 88–100 (2010). https://doi.org/10.1109/tpami.2010.59
53. Dahia, G., Segundo, M.P.: Improving fingerprint pore detection with a small FCN. arXiv preprint arXiv:1811.06846 (2018)
54. Shen, Z., et al.: Stable pore detection for high-resolution fingerprint based on a CNN detector. In: 2019 IEEE International Conference on Image Processing (ICIP), pp. 2581–2585. IEEE (2019)
55. Zhao, Q., et al.: Adaptive fingerprint pore modeling and extraction. Pattern Recogn. 43(8), 2833–2844 (2010). https://doi.org/10.1016/j.patcog.2010.02.016
56. Frégnac, Y., Bathellier, B.: Cortical correlates of low-level perception: from neural circuits to percepts. Neuron 88(1), 110–126 (2015). https://doi.org/10.1016/j.neuron.2015.09.041

57. N. I. of Standards and Technology: NIST biometric image software (NBIS) (2015). https://www.nist.gov/services-resources/software/nist-biometric-image-software-nbis
58. Liu, F., et al.: Fingerprint pore matching based on sparse representation. In: 2010 20th International Conference on Pattern Recognition, pp. 1630–1633 (2010)
59. Flanagan, P., et al.: Nist special database 301 nail to nail (N2N) fingerprint challenge dry run (2018). https://www.nist.gov/itl/iad/image-group/nist-special-database-301
60. Yu, G., Morel, J.-M.: A fully affine invariant image comparison method. In: 2009 IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 1597–1600. IEEE (2009)
61. Heusel, M., et al.: Gans trained by a two time-scale update rule converge to a local Nash equilibrium. In: Advances in Neural Information Processing Systems, vol. 30 (2017)
62. Parmar, G., Zhang, R., Zhu, J.-Y.: On buggy resizing libraries and surprising subtleties in fid calculation. arXiv preprint arXiv:2104.11222 (2021)
63. Karras, T., et al.: Training generative adversarial networks with limited data. arXiv preprint arXiv:2006.06676 (2020)
64. Tabassi, E., Wilson, C., Watson, C.: Nist fingerprint image quality. NIST Res. Rep. NISTIR7151, vol. 5 (2004)
65. Bausinger, O., Tabassi, E.: Fingerprint sample quality metric NFIQ 2.0. In: BIOSIG 2011–Proceedings of the Biometrics Special Interest Group (2011)
66. Zhao, Q., et al.: A generative model for fingerprint minutiae. In: 2013 International Conference on Biometrics (ICB), pp. 1–8. IEEE (2013)
67. Watson, C.I., Wilson, C.L.: Nist special database 4. Fingerprint Database 17(77), 5 (1992)
68. Narvekar, N.D., Karam, L.J.: A no-reference image blur metric based on the cumulative probability of blur detection (CPBD). IEEE Trans. Image Process. 20(9), 2678–2683 (2011). https://doi.org/10.1109/tip.2011.2131660

How to cite this article: Wyzykowski, A.B.V., Segundo, M.P., de Paula Lemes, R.: Multiresolution synthetic fingerprint generation. IET Biome. 11(4), 314–332 (2022). https://doi.org/10.1049/bme2.12083