

Level Three Synthetic Fingerprint Generation

André Brasil Vieira Wyzykowski, Mauricio Pamplona Segundo, Rubisley de Paula Lemes
Federal University of Bahia, Institute of Mathematics and Statistics, Bahia, Brazil
Corresponding author: andre.wyzykowski@ufba.br

Abstract—Today’s legal restrictions that protect the privacy of biometric data are hampering fingerprint recognition researches. For instance, all high-resolution fingerprint databases ceased to be publicly available. To address this problem, we present a novel hybrid approach to synthesize realistic, high-resolution fingerprints. First, we improved Anguli, a handcrafted fingerprint generator, to obtain dynamic ridge maps with sweat pores and scratches. Then, we trained a CycleGAN to transform these maps into realistic fingerprints. Unlike other CNN-based works, we can generate several images for the same identity. We used our approach to create a synthetic database with 7400 images in an attempt to propel further studies in this field without raising legal issues. We included sweat pore annotations in 740 images to encourage research developments in pore detection. In our experiments, we employed two fingerprint matching approaches to confirm that real and synthetic databases have similar performance. We conducted a human perception analysis where sixty volunteers could hardly differ between real and synthesized fingerprints. Given that we also favorably compare our results with the most advanced works in the literature, our experimentation suggests that our approach is the new state-of-the-art.

I. 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 analyze 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 able to capture L3 fingerprint images, researchers saw an opportunity to devise new and more accurate recognition approaches by using the most exploited L3 feature in the literature, the presence of sweat pores [3], [4], [5]. Besides, L3-based approaches improve security by hindering spoof attempts [6], [7], [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, which hinders the evolution of fingerprint recognition research.

Creating a synthetic database is a valid alternative to solve this problem. Synthetic generation techniques are widely ex-

ploited in several areas, such as optical flow computation [13], indoor robot navigation [14], and autonomous driving [15]. The same is true for the fingerprint recognition field.

Earlier works, such as SFinGe [16] and Anguli[17], 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, SFinGe restricts the generation of large datasets, diffculting the generation of multiple instances of a single identity. Finger-GAN [18] and Cao and Jain’s approach [19] employed Generative Adversarial Networks (GAN) [20] to learn how to generate realistic fingerprints from a training set of real examples. However, both of these works present many unnatural artifacts and cannot generate multiple images for the same identity.

We propose a novel approach to create realistic, high-resolution 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 learned texturizer to achieve realistic results and generate multiple images of a single identity. 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 public database¹ of L3 synthetic fingerprint images with five subsets of 148 identities, with 10 samples per identity, totaling 7400 fingerprint images. Also, we include sweat pore annotations for 740 images to assist in pore detection research [21]. This is the largest publicly available database with L3 fingerprints nowadays.

II. SYNTHETIC FINGERPRINT GENERATION

To generate high-resolution 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 [22] to translate seed images into realistic L3 fingerprints. Figure 1 summarizes the workflow required to create a high-resolution synthetic fingerprint using the proposed method.

¹<https://andrewyzy.github.io/L3-SF/>

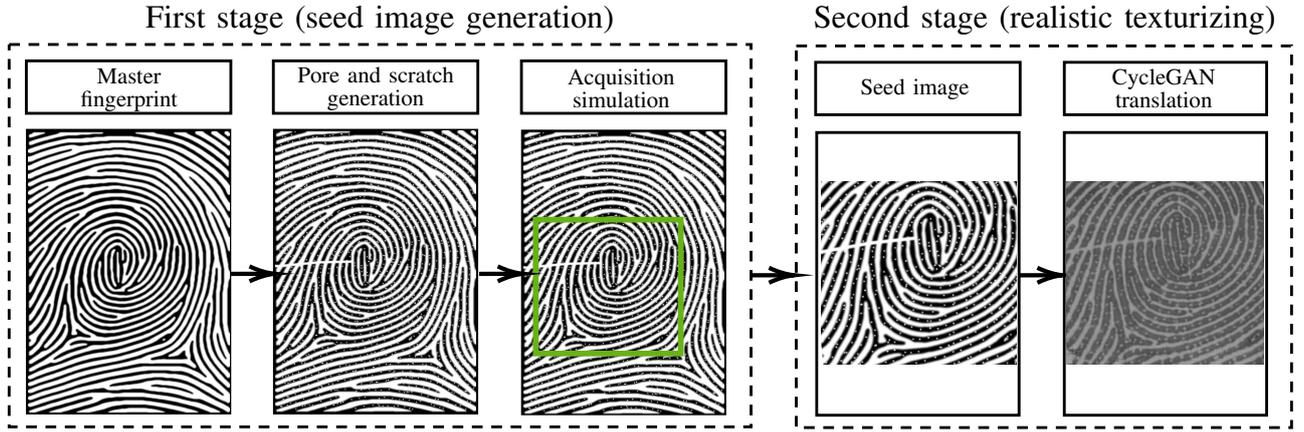


Fig. 1. Flowchart with the steps to create a high-resolution synthetic fingerprint using the proposed approach.

The first stage consists of the following processes:

- 1) **Master fingerprint generation:** we extended Anguli [17], an open-source implementation of SFinGe [16], 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 II-A).
- 2) **Pore and scratch generation:** we add pores and scratches to each master fingerprint following a distribution learned from real images to obtain L3 master fingerprints (Section II-B).
- 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 II-C).

The second stage consists of performing CycleGAN translation to transform any seed image into a realistic L3 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 III-A).
- 2) **CycleGAN training:** uses the training set to create a CNN model that translates seed images to high-resolution fingerprints (Section II-D).

A. Master fingerprint generation

We use Anguli [17] 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 2 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 [23] 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: Whorl: 41%, Right Loop: 50%, Left Loop: 3%, Arch: 6%. Since Anguli can create two arch fingerprint types, we used the ratio provided by Wang [24]: Plain Arch: 72.22% and Tented Arch: 27.77%.

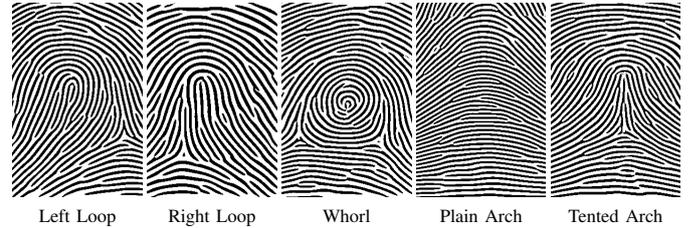


Fig. 2. Five fingerprint class patterns created by our modified version of Anguli.

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 2 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 [25] and smooth the result by applying a 3×3 mean filter. Then, we skeletonize those images using Zhang and Suen’s thinning algorithm [26] 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 neighboring ridges. We iteratively calculate $w_i = |3 \times \sin t|$, where w_i is the width of the i -th ridge and t is a counter starting at a random value for each image. After processing a ridge, t is incremented by 0.1. Ridges are processed from left to right, top to bottom. This approach avoids abrupt changes in thickness among neighboring 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 3.

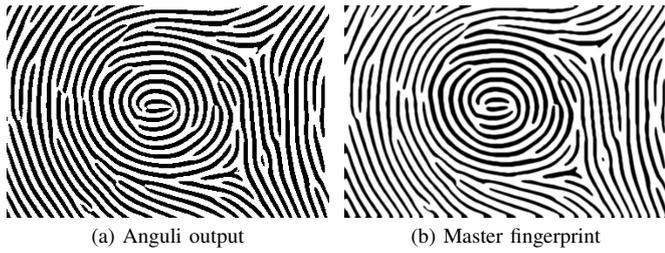


Fig. 3. Example of (a) an output of our modified version of Anguli and (b) its respective master fingerprint after ridge thickness sinusoidal perturbation. Note the ridge thickness variability in the master fingerprint, where the ridges are thicker on the center compared to the borders. Also, note our aliased result in comparison to Anguli’s pixelated ridges.

B. 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. 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 neighboring pores as our reference distribution.

To add pores, we utilize the segmented ridges generated in Section II-A. 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 i -th 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 4.

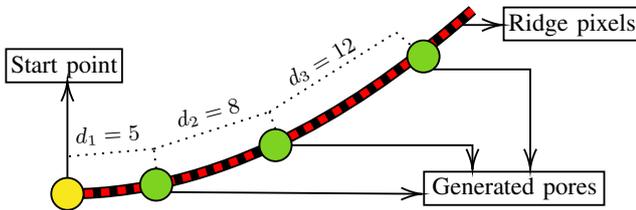


Fig. 4. Illustration of the pore generation process. Red squares represent a sequence of pixels in a ridge and the green circles are the generated pores. Given a start point, pore distances are sampled from the reference distribution. In this example, new pores are created with distances $d_1 = 5$, $d_2 = 8$ and $d_3 = 12$.

To create scratches, we count the number of scratches in each fingerprint on a real database. With these values, we used the normalized 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 ($-15^\circ \leq \theta \leq 15^\circ$) between them. This process of creating scratches is illustrated in Figure 5.

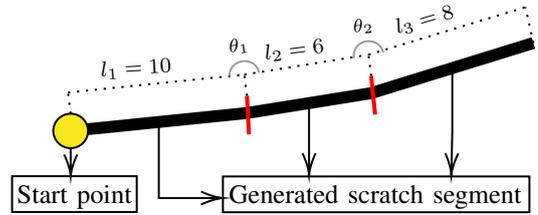


Fig. 5. Illustration of the scratch generation process. Given a random start point, we draw n line segments with random length and random angle between them. In this example, there are 3 lines segments with lengths $l_1 = 10$, $l_2 = 6$ and $l_3 = 8$.

Figure 6 shows an **L3 master fingerprint**, the outcome of this step.

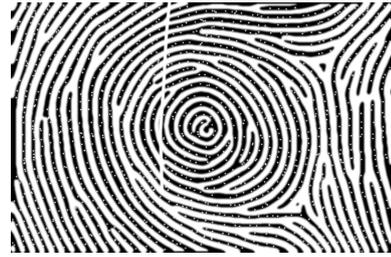


Fig. 6. An L3 master fingerprint. Note the generated pores and scratch.

C. Fingerprint acquisition simulation

Different finger positions on the image 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 [27], [28] and ORB [29] keypoints, perform an affine alignment [30] for each keypoint set separately, and select the one with the highest number of inliers. After, we use the RANSAC algorithm [31] to obtain a rigid transformation.

We assume that the center of a L3 master fingerprint is the average of the center 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 center 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 center of the i -th image, the expected difference between the center of two samples i and j , aligned to the same coordinate system, is given by:

$$\begin{aligned} \mathbb{E}[(x_i - x_j)^2] &= \mathbb{E}[\mathcal{N}(0, \sigma_x^2) + \mathcal{N}(0, \sigma_x^2)] \\ &= \mathbb{E}[\mathcal{N}(0, 2\sigma_x^2)] \end{aligned}$$

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

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

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 center region of size 512×512 from the resulting image. Examples of seed images from the same L3 master fingerprint are shown in Figure 7.

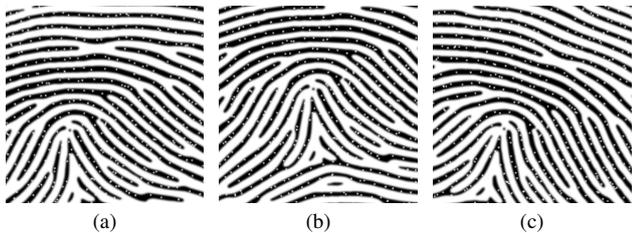


Fig. 7. Seed images generated from a single L3 master fingerprint, presenting distinct shifts and rotations.

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 or optical distortions caused by the acquisition sensor. Figure 8 illustrates this process.

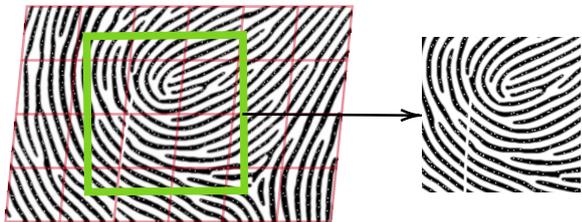


Fig. 8. Illustration of an affine transformation over an L3 master fingerprint.

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%.

D. 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 [22].

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 II-C (except pore dropout). Section III-A describes how we create the training set for this work.

We use the original CycleGAN architecture with 13 residual blocks [32] and input size 256×256 (seed images are resized to these dimensions). Besides CycleGAN's cycle consistency loss, we use the identity mapping loss [33] as it contains a regularizing component that encourages the generator to map samples from the real fingerprint domain to themselves. We train our model using the Adam optimizer [34] with a learning rate of 0.0002 for 3 epochs.

At the beginning of the training, the weight for the identity mapping loss is 0. 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 9.



Fig. 9. CycleGAN inference: seed images (top) and their respective results (bottom).

III. SYNTHETIC DATABASE CREATION

We decided to replicate the structure of an existing 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 Section II-D. More details on how we create the training set for CycleGAN are given in Section III-A. We then use the obtained model to create synthetic datasets with the same structure as DBI, as described in Section III-B.

A. CycleGAN training set

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 II-C 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 [35] to the set of seed images to create 11840 distorted images, totaling 23680 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, totaling 29600 real training samples. Table I summarizes the training images and the augmentation operations.

TABLE I
TRAINING SET IMAGES.

	Synthetic	Real
Initial images	256 (master fingerprint)	1480
Flip horizontally	592 (master fingerprint)	2960
Shift, crop and rotation (10 variations)	5920 (seed image)	29600
Higher ridge frequency cut	11840 (seed image)	-
Elastic deformation [35]	23680 (seed image)	-
Total images	23680	29600

B. Inference for database 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, totaling 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 [25], crop the center 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.

IV. EXPERIMENTAL RESULTS

Section IV-A presents a visual analysis of the L3-SF database, including a visual comparison with other methods. Section IV-B presents a fingerprint recognition analysis using both PolyU DBI and L3-SF databases. A human perception experiment involving 60 volunteers is reported in Section IV-C.

A. 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 9). Open and closed pores were also observed. We noticed other L3 traits, such as distinct ridge contours and incipient ridges, features that increase the reliability of fingerprint recognition [36], [37]. Figure 10 shows examples of L3 traits present in the L3-SF database.

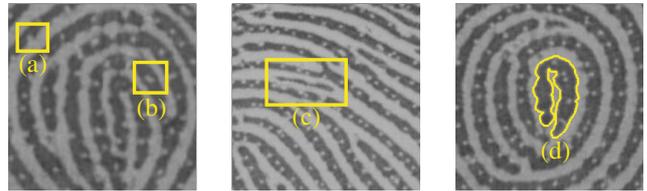


Fig. 10. (a) and (b) are open and closed pores in an image of the L3-SF database. (c) shows an incipient ridge. (d) highlights an unique ridge shape.

A visual comparison between real and our synthetic fingerprints is shown in Figure 12. Note that synthetic images have different styles similarly to real ones. We selected these images randomly to provide an unbiased judgment.

Figure 11 shows a visual comparison between fingerprints generated by the proposed approach, by a publicly available SFinge demo [16], by a public model of Cao and Jain's method [19], and by our implementation of Finger-GAN [18].

L3-SF fingerprints contain different realistic aspects, such as pores with different sizes and shapes, and ridges with acute details and texturization. Meanwhile, SFinGe [16] generates rectangular, single-sized pores and Finger-GAN [18] does not generate pores at all. Besides, Finger-GAN and Cao and Jain's method [19] produce unnatural ridge shapes. Cao and Jain's method also produces irregular minutiae patterns and irregular occlusions.

B. Fingerprint recognition analysis

The goal of this experiment is to compare real and synthetic images in terms of recognition performance. Ideally, both should have close results. To carry out this comparison, first we utilize Bozorth3 [38], 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 utilize Segundo and Lemes' approach [3]. For these analyses, we use the same protocol proposed by Liu et al. [39] for PolyU DBI and all L3-SF versions, as all of them have the same configuration (148 subjects, 2 sessions, 5 images per subject per session, 320×240 images).

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 13 and 14.

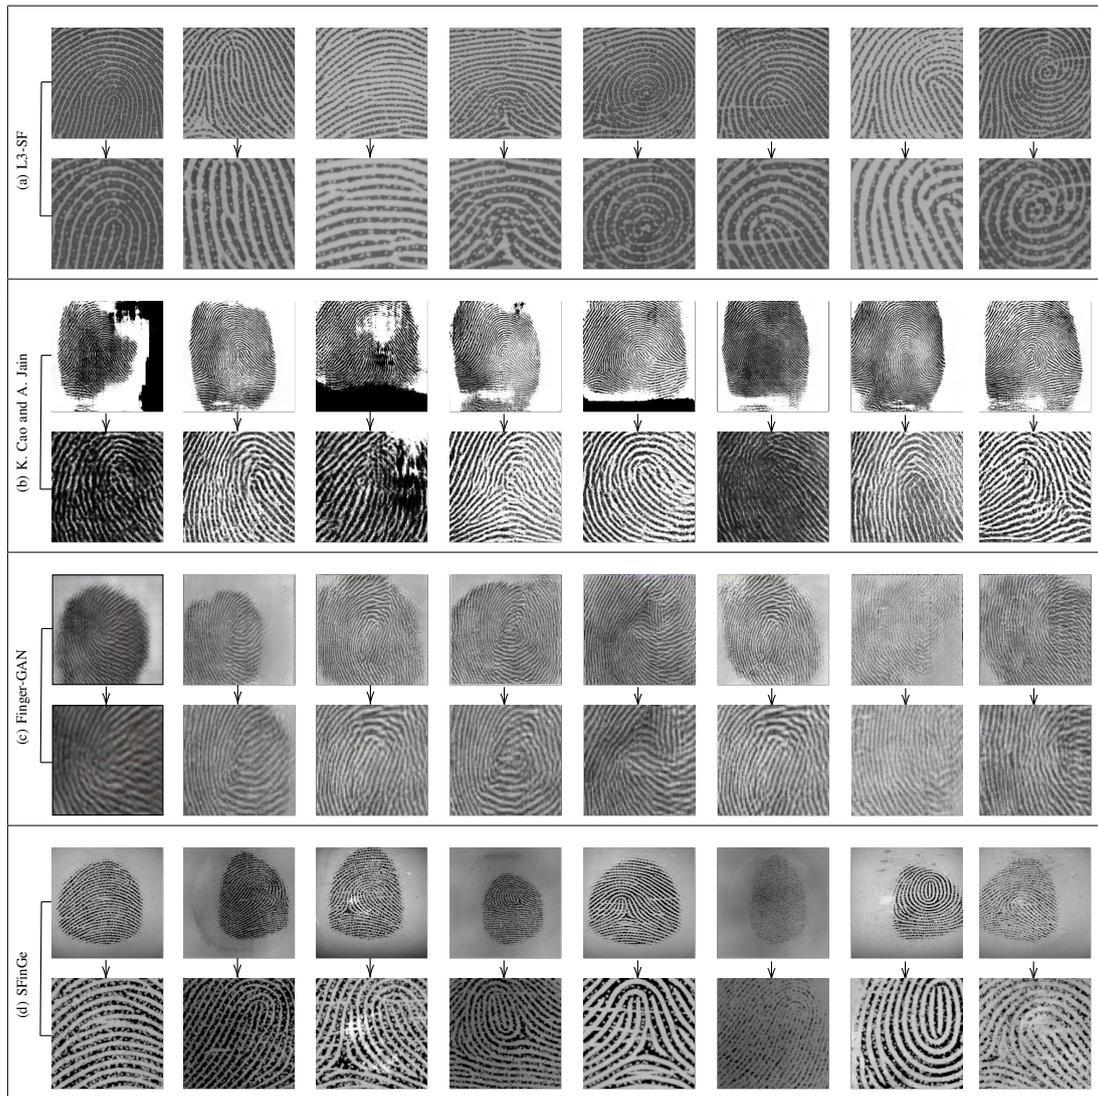


Fig. 11. Visual comparison between (a) the proposed approach, (b) Cao and Jain [19], (c) Finger-GAN [18] and (d) SFinGe [16].

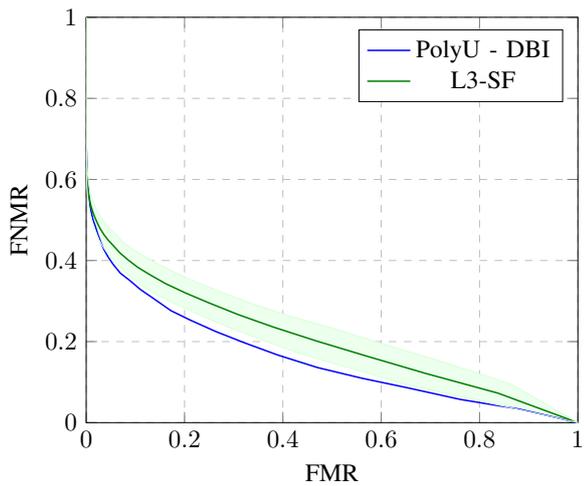


Fig. 13. ROC curves from a minutiae-based matcher for PolyU DBI and L3-SF. For the latter, we show the average ROC and a 95% confidence for its five subsets.

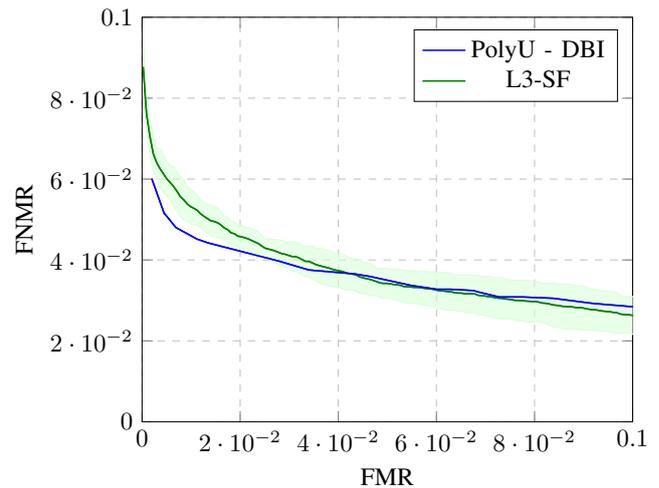


Fig. 14. ROC curves from a pore-based matcher for PolyU DBI and L3-SF database. For the latter, we show the average ROC and a 95% confidence for its five subsets.



Fig. 12. Visual comparison between real fingerprints from PolyU DBI (top) and synthetic ones from L3-SF (bottom).

The Equal Error Rate (EER) values for minutiae-based matching are 23.84% for real images and 28.12% 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 subsets as realistic databases, as 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 synthetic images. We leave this investigation as future work.

C. Human Perception Experiment

Humans generate conscious and unconscious inferences from stimulus on the visual cortex [40]. These inferences provide a judgment 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 synthetic images, we performed an experiment similar to the “real vs. fake” approach popularly used on Amazon Mechanical Turk (AMT). In this experiment, we randomly show five PolyU images and five L3-SF fingerprints to the participants, which annotate the five images they consider false. 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 analyze the overall human behavior in this classification task. When participants fail to notice any visual characteristics that help discriminating real and synthetic images, their classification accuracy is close to random (50%). Figure 15 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.

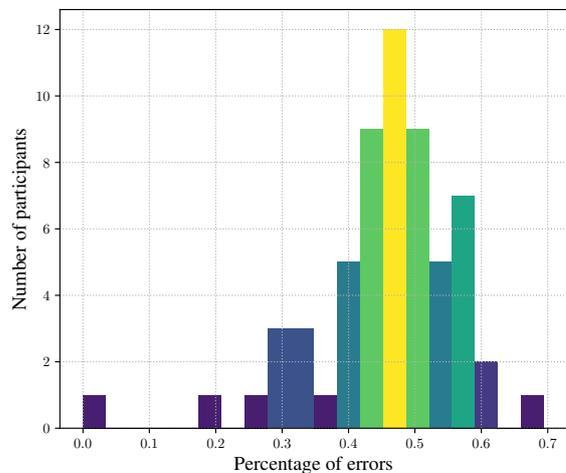


Fig. 15. Histogram of the human perception experiment. Several participants had a classification accuracy close to random (50%).

V. CONCLUSIONS

We presented an approach to generate realistic, high-resolution synthetic fingerprints containing different L3 traits. We trained a CycleGAN using real fingerprint images from PolyU and handcrafted seed images to create a model capable of translating these two image domains while preserving all identification cues (e.g., ridges, minutiae, pores). Using this approach, we created the L3-SF database with the same

characteristics of the PolyU DBI. More importantly, the L3-SF database allows further studies in the field of fingerprint biometrics without raising privacy-related legal issues. Our experimental results show that L3-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 the realism of our synthetic images thanks to a nearly-random human classification performance. Finally, we visually compared our results with the best performing works in the literature to highlight the quality enhancement over existing works. In future works, we intend to extend our approach to multiple databases and multiple fingerprint resolutions (L2 and L3) to investigate the generalization of our approach to different acquisition styles and support cross-sensor fingerprint matching research.

ACKNOWLEDGMENT

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

REFERENCES

- [1] D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar, *Handbook of Fingerprint Recognition*. Springer London, 2009. [Online]. Available: <https://doi.org/10.1007/978-1-84882-254-2>
- [2] A. K. Jain and A. Kumar, "Biometrics of next generation: An overview," *Second Generation Biometrics*, vol. 12, no. 1, pp. 2–3, 2010.
- [3] M. Pamplona Segundo and R. de Paula Lemes, "Pore-based ridge reconstruction for fingerprint recognition," in *2015 IEEE CVPR Workshops*, June 2015, pp. 128–133.
- [4] V. Anand and V. Kanhangad, "Porenet: Cnn-based pore descriptor for high-resolution fingerprint recognition," *CoRR*, vol. abs/1905.06981, 2019. [Online]. Available: <http://arxiv.org/abs/1905.06981>
- [5] D. Nguyen and A. K. Jain, "End-to-end pore extraction and matching in latent fingerprints: Going beyond minutiae," *CoRR*, vol. abs/1905.11472, 2019. [Online]. Available: <http://arxiv.org/abs/1905.11472>
- [6] M. Espinoza and C. Champod, "Using the number of pores on fingerprint images to detect spoofing attacks," *2011 International Conference on Hand-Based Biometrics, ICHB 2011 - Proceedings*, Nov 2011.
- [7] G. L. Marcialis, F. Roli, and A. Tidu, "Analysis of fingerprint pores for vitality detection," in *2010 20th International Conference on Pattern Recognition*, Aug 2010, pp. 1289–1292.
- [8] M. da Silva, A. Marana, and A. Paulino, "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*, Apr 2015, pp. 1807–1811.
- [9] N. I. of Standards and Technology, "Nist special database 30," <https://www.nist.gov/srd/nist-special-database-30>, 2002.
- [10] Q. Zhao, L. Zhang, D. Zhang, and N. Luo, "Direct pore matching for fingerprint recognition," in *Advances in Biometrics*, M. Tistarelli and M. S. Nixon, Eds. Springer Berlin Heidelberg, 2009, pp. 597–606.
- [11] V. Anand and V. Kanhangad, "Cross-sensor pore detection in high-resolution fingerprint images using unsupervised domain adaptation," *arXiv preprint arXiv:1908.10701*, 2019.
- [12] —, "Pore-based indexing for fingerprints acquired using high-resolution sensors," *Pattern Analysis and Applications*, Mar 2019. [Online]. Available: <https://doi.org/10.1007/s10044-019-00805-3>
- [13] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black, "A naturalistic open source movie for optical flow evaluation," in *European conference on computer vision*. Springer, 2012, pp. 611–625.
- [14] Y. Wu, Y. Wu, G. Gkioxari, and Y. Tian, "Building generalizable agents with a realistic and rich 3d environment," *arXiv preprint arXiv:1801.02209*, 2018.
- [15] M. Wrenninge and J. Unger, "Synscapes: A photorealistic synthetic dataset for street scene parsing," *arXiv preprint arXiv:1810.08705*, 2018.
- [16] R. Cappelli, "Sfinge: an approach to synthetic fingerprint generation," *International Workshop on Biometric Technologies*, Jan 2004.
- [17] A. H. Ansari, "Generation and storage of large synthetic fingerprint database," *ME Thesis*, Jul, 2011.
- [18] S. Minaee and A. Abdolrashidi, "Finger-gan: Generating realistic fingerprint images using connectivity imposed gan," *arXiv preprint arXiv:1812.10482*, 2018.
- [19] K. Cao and A. Jain, "Fingerprint synthesis: Evaluating fingerprint search at scale," in *2018 International Conference on Biometrics (ICB)*, 2018.
- [20] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," 2014.
- [21] R. de Paula Lemes, M. P. Segundo, O. R. Bellon, and L. Silva, "Dynamic pore filtering for keypoint detection applied to newborn authentication," in *2014 22nd International Conference on Pattern Recognition*. IEEE, 2014, pp. 1698–1703.
- [22] J. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," *CoRR*, vol. abs/1703.10593, 2017.
- [23] M. van Mensvoort, "The fingerprints-world-map!" <http://fingerprints.handresearch.com/dermatoglyphics/fingerprints-world-map-whorls-loops-arches.htm>, mar 2015.
- [24] P. S. P. Wang, Ed., *Pattern Recognition, Machine Intelligence and Biometrics*. Springer Berlin Heidelberg, 2011. [Online]. Available: <https://doi.org/10.1007/978-3-642-22407-2>
- [25] C. Dong, C. C. Loy, and X. Tang, "Accelerating the super-resolution convolutional neural network," in *European conference on computer vision*. Springer, 2016, pp. 391–407.
- [26] T. Y. Zhang and C. Y. Suen, "A fast parallel algorithm for thinning digital patterns," *Commun. ACM*, vol. 27, no. 3, pp. 236–239, Mar. 1984. [Online]. Available: <http://doi.acm.org/10.1145/357994.358023>
- [27] D. G. Lowe, "Object recognition from local scale-invariant features," in *Proceedings of the International Conference on Computer Vision - Volume 2 - Volume 2*, ser. ICCV '99. Washington, DC, USA: IEEE Computer Society, 1999, pp. 1150–. [Online]. Available: <http://dl.acm.org/citation.cfm?id=850924.851523>
- [28] —, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [29] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *Proceedings of the 2011 International Conference on Computer Vision*, 2011, pp. 2564–2571. [Online]. Available: <http://dx.doi.org/10.1109/ICCV.2011.6126544>
- [30] G. Yu and J.-M. Morel, "Asift: An algorithm for fully affine invariant alternative to sift or surf," *Image Processing On Line*, vol. 1, pp. 11–38, 2011.
- [31] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, Jun. 1981. [Online]. Available: <http://doi.acm.org/10.1145/358669.358692>
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [33] Y. Taigman, A. Polyak, and L. Wolf, "Unsupervised cross-domain image generation," *arXiv preprint arXiv:1611.02200*, 2016.
- [34] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [35] P. Y. Simard, D. Steinkraus, J. C. PLATT *et al.*, "Best practices for convolutional neural networks applied to visual document analysis," in *Seventh ICDAR, 2003. Proceedings.*, Aug 2003, pp. 958–963.
- [36] W. Lee, S. Cho, H. Choi, and J. Kim, "Partial fingerprint matching using minutiae and ridge shape features for small fingerprint scanners," *EXWA*, vol. 87, no. C, pp. 183–198, 2017.
- [37] A. K. Jain and J. Feng, "Latent fingerprint matching," *IEEE TPAMI*, vol. 33, no. 1, pp. 88–100, 2010.
- [38] N. I. of Standards and Technology, "Nist biometric image software (nbis)," <https://www.nist.gov/services-resources/software/nist-biometric-image-software-nbis>, 2015. [Online]. Available: <https://www.nist.gov/services-resources/software/nist-biometric-image-software-nbis>
- [39] F. Liu, Q. Zhao, L. Zhang, and D. Zhang, "Fingerprint pore matching based on sparse representation," *2010 20th International Conference on Pattern Recognition*, pp. 1630–1633, 2010.
- [40] Y. Frégnac and B. Bathellier, "Cortical correlates of low-level perception: from neural circuits to percepts," *Neuron*, vol. 88, no. 1, pp. 110–126, 2015.