

Fingerprint Synthesis: Search with 100 Million Prints

Vishesh Mistry
Michigan State University
East Lansing, MI, USA
mistryvi@msu.edu

Joshua J. Engelsma
Michigan State University
East Lansing, MI, USA
engelsm7@cse.msu.edu

Anil K. Jain
Michigan State University
East Lansing, MI, USA
jain@cse.msu.edu




Figure 1: Comparison of real fingerprints with synthesized fingerprints. Images in (a) are rolled fingerprint images from a law enforcement operational dataset [17] and fingerprints in (b) are synthesized using the proposed approach. One example for each of the four major fingerprint types is shown (from left to right: arch, left loop, right loop, whorl). All the fingerprints have dimensions of 512×512 pixels and a resolution of 500 dpi. A COTS minutiae extractor was used to extract the minutiae points overlaid on the fingerprint images. Fingerprints in (a), from left to right, have NFIQ 2.0 quality scores of (47, 64, 77, 73), respectively, while those in (b), from left to right, have quality scores of (55, 65, 60, 69), respectively.

Abstract

Evaluation of large-scale fingerprint search algorithms has been limited due to lack of publicly available datasets. To address this problem, we utilize a Generative Adversarial Network (GAN) to synthesize a fingerprint dataset consisting of 100 million fingerprint images. In contrast to existing fingerprint synthesis algorithms, we incorporate an identity loss which guides the generator to synthesize fingerprints corresponding to more distinct identities. The characteristics of our synthesized fingerprints are shown to be more similar to real fingerprints than existing methods via eight different metrics (minutiae count - block and

*template, minutiae direction - block and template, minutiae convex hull area, minutiae spatial distribution, block minutiae quality distribution, and NFIQ 2.0 scores). Additionally, the synthetic fingerprints based on our approach are shown to be more distinct than synthetic fingerprints based on published methods through search results and imposter distribution statistics. Finally, we report for the first time in open literature, search accuracy against a gallery of **100 million fingerprint images** (NIST SD4 Rank-1 accuracy of 89.7%).*

Algorithm	Orientation Model	Ridge Structure Model	Minutiae Model	Analysis ¹
Cappelli <i>et al.</i> [21]	Zero-pole [40]	Gabor Filter [27]	Gabor Filter [27]	• Generated 10K plain fingerprints (240 × 336)
SFinGe [10, 19]	Zero-pole [40]	Gabor Filter [27]	Gabor Filter [27]	• Generated 100K plain-prints per 24 hours
Johnson <i>et al.</i> [32]	Zero-pole [40]	Gabor Filter [27]	Spatial [23]	• Generated 2,400 plain-prints
Zhao <i>et al.</i> [44]	Zero-pole [40]	AM-FM [35]	Spatial [23]	• Generated 5,000 rolled-prints.
Bontrager <i>et al.</i> [16]		Wasserstein GAN (WGAN) [13]		• Generated 128 × 128 fingerprint patches
Cao & Jain [18]		IWGAN [31] and Autoencoder [36]		• Synthesized 10 million rolled-prints • Fingerprints lack “uniqueness”
Attia <i>et al.</i> [14]		Variational Autoencoder [34, 39]		• Insufficient training data (800 unique fingerprints) • No evaluation of “uniqueness”
Proposed Approach		IWGAN [31] and Autoencoder [36] with Identity Loss [25]		• Synthesized & searched 100 million fingerprints • Improved realism and uniqueness

¹ Generally speaking, model-based approaches lack “realism” in comparison to the more recent learning-based approaches.

Table 1: Summary of Fingerprint Synthesis Methods in the Literature

1. Introduction

Large-scale Automated Fingerprint Identification Systems (AFIS) have been widely deployed throughout law enforcement and forensics agencies, international border crossings, and national ID systems. A few notable applications, where large-scale AFIS are being used include: (i) India’s Aadhaar Project [11] (1.25 billion ten-prints), (ii) The FBI’s Next Generation Identification system (NGI) [4] (145.3 million ten-prints), and (iii) The DHS’s Office of Biometric Identity Management (OBIM) system [9] (fingerprints and face images of non-US citizens at ports of entry to the United States).

While these large-scale search systems are seemingly quite successful, little work has been done in the academic literature to understand the performance of these search systems at scale (search performance on a gallery of 100 million or larger). In particular, a number of fingerprint search algorithms have been proposed [41, 17, 20, 28, 15], but nearly all of them were evaluated on a gallery of at most 2,700 rolled fingerprints from NIST SD14 [5]. A few exceptions include [41, 25] and [17] with gallery sizes of 1 million and 250,000, respectively. Given this limitation in our understanding of large-scale search performance, a primary contribution of this paper is the evaluation of fingerprint search performance (accuracy and efficiency) against a gallery at a scale of 100 million fingerprints.

To evaluate against a gallery of 100 million requires one to either (i) collect large-scale fingerprint records with Institutional Review Board (IRB) approval, (ii) obtain 100 million fingerprint records from an existing forensic or govern-

ment dataset, or (iii) synthesize fingerprint images.

Collecting 100 million fingerprints is practically infeasible in terms of cost, time, and IRB regulations. Furthermore, governmental privacy regulations¹ prohibit sharing of existing fingerprint datasets (say from National ID systems). Indeed, NIST has even taken down some of their previously public fingerprint datasets (NIST SD27 [6], NIST SD14 [5], and NIST SD4 [8]) due to stringent privacy regulations. Therefore, to adequately evaluate large-scale fingerprint search algorithms, we propose to first synthesize 100 million fingerprints with characteristics close to those of real fingerprint images and later, to use our synthetic fingerprints to augment the gallery for search performance evaluation at scale.

More concisely, the contributions of this research are:

- A fingerprint synthesis algorithm built upon GANs [29] which is trained on a large dataset of real fingerprints [43]. We show that our *learning-based* synthesis approach generates more realistic fingerprints (both plain-prints and rolled-prints) than traditional *model-based* synthesis algorithms [44, 21, 10, 32].
- In contrast to existing learning-based synthesis algorithms [18, 14, 16, 37], we add an identity loss to generate fingerprints of more *unique* identities.
- Synthesis of 100 million fingerprint images. This synthesis required code optimization and resources (51 hours of run time on 100 Tesla K80 GPUs).

¹<https://bit.ly/2YD4e4A>




Figure 2: Comparison of synthetic fingerprints. Fingerprints are synthesized with methods in (a) [21], (b) [44], (c) [18], (e) [32], (f) [16], and (g) [14]. The plain-print (d) and rolled-print (h) were synthesized using the proposed approach. Qualitatively speaking, our synthetic fingerprints demonstrate more realism than existing approaches. This is further supported based on our quantitative evaluation.

- Large-scale fingerprint search evaluation (accuracy and efficiency) against the 100 million synthetic-prints gallery. To the best of our knowledge, this is the largest-scale fingerprint search results ever reported in the literature.

2. Related Work

A plethora of approaches have been proposed for synthetic fingerprint generation which are concisely summarized in Table 1. In addition, Figure 2 shows a qualitative comparison between the synthetic fingerprints generated using [14, 16, 44, 21, 18, 32] and our proposed approach. These approaches can be broadly categorized into either *model-based* approaches or *learning-based* approaches. While the learning-based methods require large amounts of training data, they are capable of generating more realistic fingerprints than the model-based approaches (shown in our experiments). The limitations of published model-based approaches are further elaborated upon below:

- Approaches in [21, 32, 44, 10, 19] assume an independent ridge orientation and minutiae model. As such, minutiae distributions could be generated without having a valid ridge orientation field.
- Approaches in [21, 32, 10, 19] assume a fixed fingerprint ridge-width. However, real fingerprints have

varying ridge widths. In fact, [22] reported a 98% accuracy in discriminating real fingerprints from synthetic fingerprints using only two ridge spacing features.

- The ridge structure model in [21, 32, 44, 10, 19] uses a masterprint. Due to Gabor filtering and the AM/FM models, the generated masterprints have non-consistent ridge flow patterns which lead to unrealistic fingerprint images.
- The approaches in [21, 32, 44, 10, 19] are susceptible to generating unrealistic minutiae configurations as their minutiae models do not consider local minutiae configurations. In particular, authors in [30] showed that they could perfectly discriminate between real and synthetic [21] fingerprints using minutiae configurations alone.
- None of the approaches generated a fingerprint dataset of size comparable to large-scale operational datasets.

More recent learning-based approaches to fingerprint synthesis [18, 14, 16, 37] have utilized generative adversarial networks (GANs) [29]. While several of these approaches are capable of improving the realism of fingerprint images using large training datasets of real fingerprints, they do not consider the uniqueness or individual-

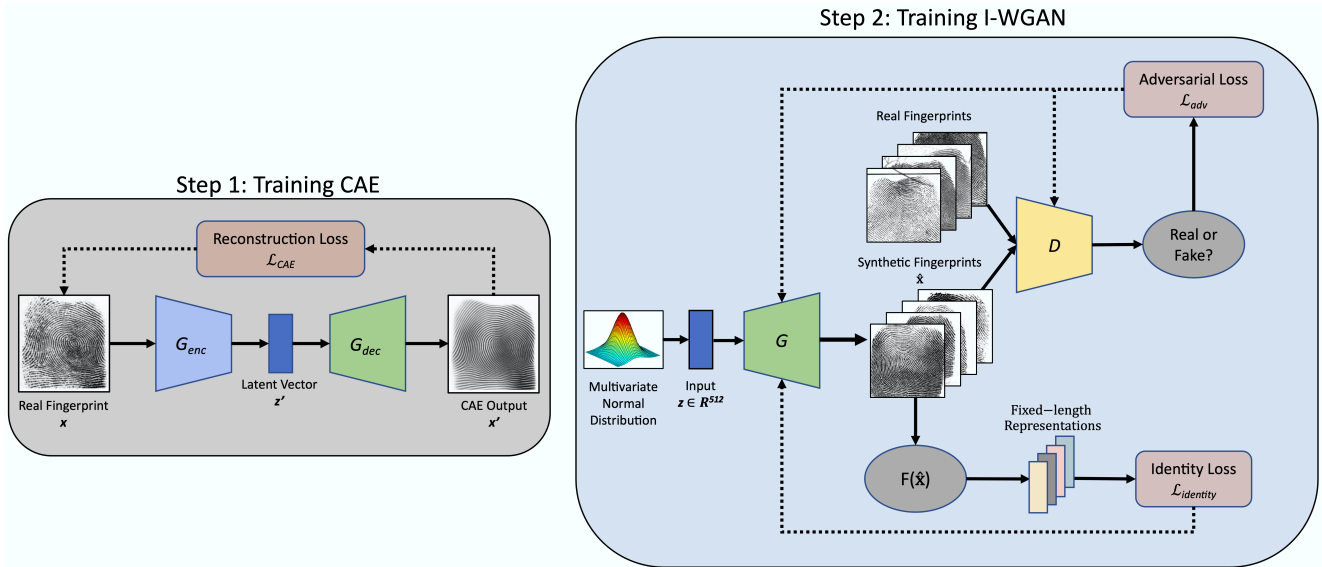


Figure 3: Schematic of the proposed fingerprint synthesis approach. Step 1 shows the training flow of the CAE, which includes the encoder G_{enc} and the decoder G_{dec} . Step 2 illustrates the training of the I-WGAN (G and D) along with the introduced identity loss. The weights of G were initialized using the trained weights of G_{dec} . The solid black arrows show the forward pass of the networks while the dotted black lines show the propagation of the losses.

ity of the generated fingerprints. Indeed, without any additional supervision, a generator may continually synthesize a small subset of fingerprints, corresponding to only a few identities. This was quantitatively demonstrated in [18] where the search accuracy against real fingerprint galleries was lower than that of synthetic fingerprint galleries. In this work, we build upon the state-of-the-art learning-based synthesis algorithms to exploit on the realism they offer, however, we also incorporate additional supervision (identity loss) to encourage the generation of fingerprints which are of distinct identities. After improving the uniqueness and realism of synthetic fingerprints, we synthesize 100 million fingerprints (the largest gallery reported in the literature) to conduct our large-scale fingerprint search evaluation.

3. Proposed Approach

Improved WGAN [31] (I-WGAN) is the backbone of the proposed fingerprint synthesis algorithm. As was done in [18], we initialize the weights of the I-WGAN generator using the trained weights of the decoder of a convolutional autoencoder (CAE). Our main contribution to the synthesis process is the introduction of an identity loss which guides the generator to synthesize fingerprint images with more distinct characteristics during training.

The schematic of the proposed approach is shown in Figure 3. The following sub-sections explain the major components of the algorithm in detail.




Figure 4: Fingerprint image in (a) was generated by G without CAE weight initialization and fingerprints image in (b) was generated by G with CAE weight initialization from G_{dec} .

3.1. Convolutional Autoencoder

The first step of our proposed approach involves training a convolutional autoencoder (CAE) [36] in an unsupervised fashion. The encoder (G_{enc}) takes a 512×512 grayscale input image (x) and transforms it into a 512-dimensional latent vector (z'). The decoder (G_{dec}) takes the latent vector (z') as input and attempts to reconstruct back the input image (x'). More formally, the CAE minimizes the reconstruction loss shown in Equation 1.

$$\mathcal{L}_{CAE} = \|x - x'\|_2^2 \quad (1)$$

After training the CAE, we use the weights of G_{dec} to initialize the Generator G of the I-WGAN. This initialization procedure injects domain knowledge into G enabling it to produce fingerprint images which are of higher quality (as was shown in [18]). Figure 4 shows a comparison between fingerprint images generated using G with and without weight initialization by G_{dec} .

3.2. Improved-WGAN

The architecture of I-WGAN [31] used in our approach follows the implementation of [38]. Both the generator G and the discriminator D consist of seven convolutional layers with a kernel size of 4×4 and a stride of 2×2 (architecture deferred to supplementary). G takes a 512-dimensional latent vector $z \sim \mathbb{P}_z$ as input, where \mathbb{P}_z is a multivariate Gaussian distribution, and outputs a 512×512 dimensional grayscale image. All the convolutional layers in G have ReLU as their activation function, except the last layer which uses \tanh . The discriminator D has the inverse architecture of the generator G . All the convolutional layers in D use LeakyReLU as the activation function. Moreover, batch normalization is applied to all the layers of G and D .

3.3. Identity Loss

The authors in [18] demonstrated that synthetic fingerprints are less unique than real fingerprints by showing that the search accuracy on NIST SD4 [8] (2,000 rolled fingerprint pairs) against 250K synthesized rolled fingerprints is higher than when 250K real rolled fingerprints are used as the gallery. One plausible explanation for this difference in search accuracies is that no distinctiveness measure was used during training of the synthetic fingerprint generator, leading to lower “diversity” among the synthesized fingerprints than among real fingerprints. To rectify this, we introduce an identity loss to encourage synthesis of more distinct or unique fingerprint images.

The authors in [25] introduced a deep network, called DeepPrint, to extract a fixed-length representation from a fingerprint image, encoding the minutiae as well as the texture information (i.e. the identity) of the fingerprint into a 192-dimensional vector. Since the DeepPrint feature extractor is a differentiable function (i.e. it is a deep network), we can use the trained DeepPrint [25] network ($F(x)$) during the training process (computing the loss) of the I-WGAN (unlike DeepPrint, COTS minutiae matchers are not differentiable and therefore cannot be used to train the network). In particular, for each mini-batch, we use DeepPrint to extract the fixed-length representations of all the fingerprint images (x) synthesized by the generator G . The similarity between each pair of representations in the mini-batch is then minimized in order to guide G to produce fingerprints with different identities, i.e. more distinct fingerprints.

More formally, for each pair of latent vectors (z_i, z_j) in

a mini-batch, the identity loss, which is given by Equation. 2, is minimized.

$$\mathcal{L}_{identity} = \frac{1}{\sum \|F(G(z_i)) - F(G(z_j))\|_2}, (z_i \neq z_j) \quad (2)$$

3.4. Training Details

The CAE and the I-WGAN were both trained using 280,000 rolled fingerprint images from a law enforcement agency [17]. The fingerprints in the training set range in size from 407×346 pixels to 750×800 pixels at 500 dpi. We coarsely segment them to 512×512 pixels since an area of 512×512 pixels with a resolution of 500 dpi is sufficient to cover the entire friction ridge pattern. We acquired the weights for DeepPrint from the authors in [25].

3.5. Fine-Tuning for Plain-Prints

The CAE and I-WGAN are trained to synthesize rolled fingerprints. However, we show that by fine-tuning I-WGAN on an aggregated dataset of 84K plain-prints from [2, 24], we can also generate high-quality plain-prints. This enables us to compare the realism of our learning-based synthetic fingerprints more fairly with existing, publicly available plain-print synthesis algorithms (SFinGe [10]).

3.6. Synthesis Process

After training I-WGAN with the identity loss, we generate 100 million rolled fingerprint images. Generating 100 million fingerprint images is non-trivial and requires significant memory and processor resources. Our synthesis was expedited via a High Performance Computing Center (HPCC) on our campus. In particular, we ran 100 jobs in parallel on the HPCC server, with each job generating 1 million fingerprint images. As each fingerprint was generated, we simultaneously extracted its DeepPrint [25] fixed-length representation (to be used in our later search experiments). Each job was assigned an Intel Xeon E5-2680v4@2.40GHz CPU, 32GB of RAM, and a NVIDIA Tesla K80 GPU. This took 51 CPU hours². Finally, we also generated 2,000 plain-prints using our fine-tuned plain-print model for comparison with other synthetic plain-print approaches.

4. Experimental Results

In our experimental results, we first demonstrate that our synthetic fingerprints are more realistic and unique than those from published methods. Then, we show the fingerprint search results of COTS and DeepPrint [25] matchers against a background of 1 million and 100 million synthetic fingerprints respectively.

²Since the HPCC follows a queue-based approach for running jobs, all the 100 jobs were initiated within 4 days.



Figure 5: Examples of fingerprint images (512×512 @ 500 dpi) synthesized by our approach. The top row corresponds to synthetic rolled prints while the bottom row shows synthetic plain fingerprints.

4.1. Fingerprint Realism

To evaluate the similarity of synthetic fingerprints to real fingerprints, we used the fingerprint datasets enumerated in Table 2. In particular, we compare real plain-prints to synthetic plain-prints and real rolled-prints to synthetic rolled-prints.

Metrics: Four template-level and three block-level comparison indicators used for this evaluation are taken from [19], including, minutiae count - template and block, direction - template and block, convex hull area, spatial distributions (minutiae spatial distributions represented as a 2D minutiae histogram [30, 18]), and block minutiae quality. In addition, we utilize the NFIQ 2.0 [42] quality scores distribution. Minutiae points (necessary to compute these metrics) were extracted using state-of-the-art COTS SDKs: VeriFinger v11.0 [12] and Innovatrics v.7.6.0.627 [3].

Statistical Test: We use the Kolmogorv-Smirnov test [26] to compute the difference between the distributions of each of the aforementioned metrics extracted from a set of real fingerprints (benchmark distribution), and synthetic fingerprints (comparison distribution). A low test statistic value indicates high similarity between real fingerprints and the synthetic fingerprints. We compared each synthetic fingerprint dataset with three real fingerprint datasets on all the above-mentioned metrics and calculated the average KS test statistic from the three comparisons.

Realism Results: Figure 6 shows the KS test statistics between our synthetic rolled and plain prints and real rolled and plain prints as well as the comparison with state-of-the-art methods for plain-print synthesis - SFinGe v5.0 [10] ($21.1\text{mm} \times 28.4\text{mm}$ acquisition area, 500 DPI resolution,

416×560 image dimensions, natural class distribution), and rolled-print synthesis [17]. We note that our synthetic plain-prints are more similar to real plain-prints on 7 out of the total 8 indicators than state-of-the-art synthetic plain-prints [10]. Similarly, our synthetic rolled-prints are more similar to real rolled-prints than state-of-the-art synthetic rolled-prints [18] on 7 out of the 8 indicators.

4.2. Fingerprint Uniqueness

4.2.1 Imposter Scores Distribution

To determine the diversity of our synthetic fingerprints (in terms of identity), we first computed 500K imposter scores using Verifinger in conjunction with our synthetic rolled-prints and also the synthetic rolled-prints from [18] (Figure 7). We note that the mean (3.47) and standard deviation (2.13) of the imposter scores computed when using our synthetic rolled-prints are both lower than the mean (3.48) and standard deviation of (2.18) of the imposter scores computed with synthetic rolled-prints from [18]. We also note a higher peak in the imposter distribution from our synthetic rolled-prints at lower similarity values (Figure 7). This indicates that our addition of an identity loss has helped guide our synthetic rolled-prints to be more unique.

4.2.2 DeepPrint Search Against 1 Million

We also demonstrate uniqueness by conducting large-scale search experiments. More uniqueness in the gallery leads to a lower search accuracy. First, we obtain an “upper-bound” on search performance using NIST SD4 [8] in conjunction with 10 different subsets of 1 million synthetic rolled-prints

	Plain Fingerprint datasets	Rolled Fingerprint datasets
Real	<ul style="list-style-type: none"> • CASIA-Fingerprint V5 [1] (2000 fingerprints) • NIST SD302L [7] (1951 fingerprints) • NIST SD302M [7] (1979 fingerprints) 	<ul style="list-style-type: none"> • NIST SD4 [8] (2000 enrollment fingerprints) • NIST SD14 [5] (last 2000 enrollment fingerprints) • NIST SD302U [7] (2000 fingerprints)
Synthetic	<ul style="list-style-type: none"> • SFinGe [10] (2000 fingerprints) • Proposed Approach (2000 fingerprints) 	<ul style="list-style-type: none"> • Cao and Jain [18] (2000 fingerprints) • Proposed Approach (2000 fingerprints)

Table 2: Real and synthetic fingerprint datasets that were used for comparison in our experiments.

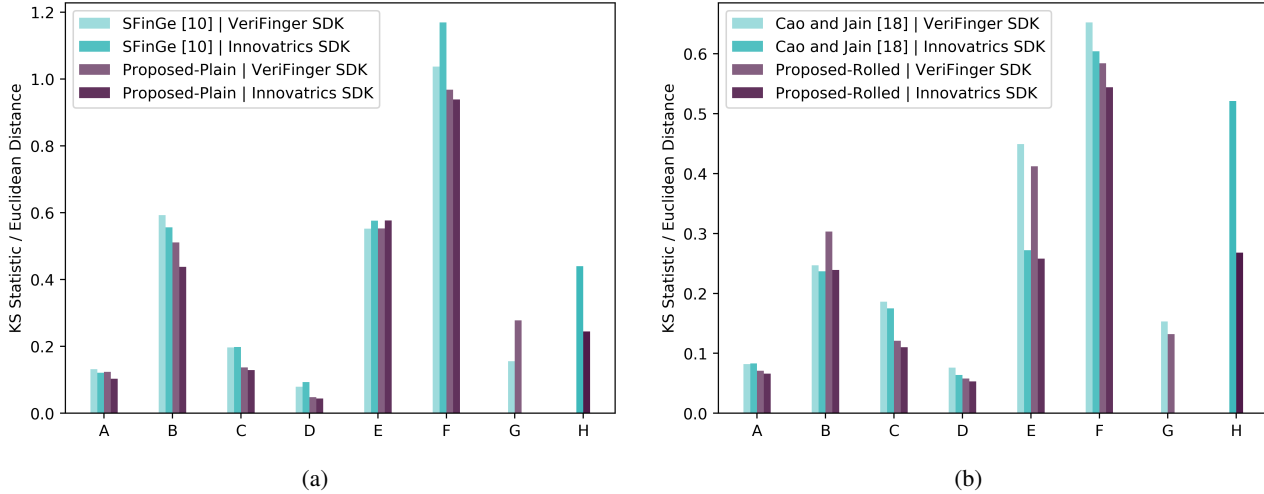


Figure 6: The comparison of synthetic plain (a) and rolled fingerprints (b) to real fingerprints using 8 metrics (minutiae count - block [A] and template [B], direction - block [C] and template [D], convex hull area [E], spatial distributions [F], block minutiae quality [G], and NFIQ 2.0 [H]) and a Kolmogorov-Smirnov test. A lower value indicates higher similarity of synthetic fingerprints to real fingerprints.




Figure 7: Imposter score distributions computed using the proposed synthetic rolled-prints, synthetic rolled-prints from [18] and real fingerprints from NIST SD4 [8]. Scores were computed using VeriFinger SDK v11.

from our proposed approach. We use DeepPrint [25] as the matcher. Confidence intervals for rank- N search accuracies are shown in Figure 8. The mean rank-1 search accuracy is 95.53% with confidence interval of [95.1, 95.8].

4.2.3 COTS Search Against 1 Million

We also evaluate the distinctiveness of our synthetic fingerprints using the Innovatrics SDK. In particular, we perform fingerprint search using NIST SD4 [8] in conjunction with our synthetic rolled-prints and synthetic rolled-prints from [18]. The rank-1 search accuracies for both the datasets (at gallery sizes of 250K and 1 Million) are shown in Table 3. From Table 3 we note that although the search accuracy is lower when using rolled-prints from [17] at a lower gallery size of 250K gallery, the search performance drops more rapidly as the synthetic gallery scales to 1 Million when using our proposed synthetic rolled-prints. This suggests that the uniqueness of our synthetic rolled-prints becomes more evident at large gallery sizes [17]. We also acknowledge that there is still room for further improvement in the “uniqueness” of synthetic fingerprints as the




Figure 8: Confidence interval for rank-N search accuracy on NIST SD4 [8]. 10-folds of 1 million synthetic rolled fingerprints were used to compute confidence intervals.

	Proposed Approach	Cao & Jain [18]
250K Gallery	91.45%	90.85%
1M Gallery	90.35%	90.40%

Table 3: Innovatrics rank-1 fingerprint search accuracies on NIST SD4 [8] using synthetic rolled-prints as gallery.

rank-1 search performance of Innovatrics against 1.1 Million real rolled-prints was reported to be 89.2% in [25] (vs. 90.35% against 1 Million of our synthetic rolled-prints).

4.2.4 Search Against 100 Million

To date, no work has been done in the literature to evaluate fingerprint search performance at a scale of 100 Million. Two main challenges preventing this important evaluation are (i) the computational challenges in generating 100 million fingerprints and (ii) finding a matcher fast enough to search 100 million. Our use of a high-performance computing center solves the former, and the recently proposed fast fixed-length matcher, DeepPrint [25] addresses the latter. Therefore, in this paper, we not only synthesize fingerprints (like others in the literature), but we also scale the synthesis to 100 million **and** we actually show the search performance against the large-scale synthesized gallery.

Figure 9 shows the rank-N search accuracies on NIST SD4 [8] in conjunction with 100 million of our synthetic rolled-prints (rank-1 accuracy of 89.7%). We also show the search performance when using 100 million synthetic rolled-prints from [17] as a benchmark (rank-1 accuracy of 93.55%). We note that at a gallery size of 100 million, the search performance is much lower when using our

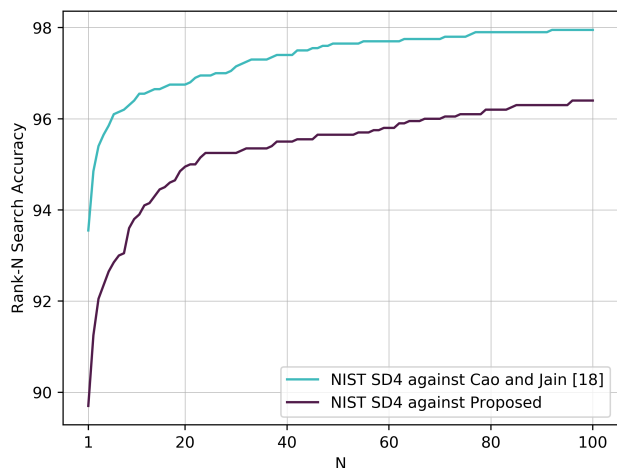


Figure 9: Rank-N fingerprint search accuracies on NIST SD4 [8] for the augmented gallery of 100 million fingerprints using DeepPrint [25] as the fingerprint matcher.

synthetic fingerprints than when using synthetic fingerprints from [17]. This is strong evidence that our fingerprints are more unique than the fingerprints from [17].

Since DeepPrint performed similarly to top-COTS matchers on a gallery of 1 million real fingerprints in [25], we posit that DeepPrint and by extension COTS matchers have room for improvement with galleries of hundreds of millions or even billions of fingerprints (e.g. Aadhaar). This is our ongoing research direction. Note that we could not evaluate COTS against 100 million due to the inadequate speed of the SDK licenses that we have for the matchers (about 400 days for a rank-1 search on a gallery size of 100 million images as compared to 11s for DeepPrint).

5. Conclusions

Given the difficulty in obtaining large-scale fingerprint datasets for evaluating large-scale search, we propose a fingerprint synthesis algorithm which utilizes I-WGAN and an identity loss to generate diverse and realistic fingerprint images. We show that our method can be used to synthesize realistic rolled-prints as well as plain-prints. We demonstrate that our synthetic fingerprints are more realistic and diverse than current state-of-the-art synthetic fingerprints. Finally, we show, for the first time in the literature, fingerprint search performance at a scale of 100 million. Our results suggest that state-of-the-art fingerprint matchers must be further improved to operate at a scale of 100 million. Our ongoing work will (i) further scale our search experiments to a gallery of 1 billion, and (ii) further improve the realism and uniqueness of synthesized fingerprints.

References

- [1] CASIA Fingerprint Image Database Version 5.0. <http://www.idealtest.org/dbDetailForUser.do?id=7>.
- [2] IARPA ODIN Program. <https://www.iarpa.gov/index.php/research-programs/odin>.
- [3] Innovatrics SDK. <https://www.innovatrics.com/innovatrics-abis/>.
- [4] Next Generation Identification (NGI). <https://www.fbi.gov/services/cjis/fingerprints-and-other-biometrics/ngi>.
- [5] NIST Special Database 14. <https://www.nist.gov/srd/nist-special-database-14>.
- [6] NIST Special Database 27. <https://www.nist.gov/itl/iad/image-group/nist-special-database-2727a>.
- [7] NIST Special Database 302. <https://www.nist.gov/itl/iad/image-group/nist-special-database-302>.
- [8] NIST Special Database 4. <https://www.nist.gov/srd/nist-special-database-4>.
- [9] Office of Biometric Identity Management. <https://www.dhs.gov/obim>.
- [10] SFinGe (Synthetic Fingerprint Generator). <http://biolab.csr.unibo.it/research.asp?organize=Activities&select=&selObj=12&pathSubj=111%7C%7C12&>.
- [11] Unique Identification Authority of India. <https://uidai.gov.in>.
- [12] VeriFinger SDK. <https://www.neurotechnology.com>.
- [13] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein GAN. *CoRR*, abs/1701.07875, 2017.
- [14] M. Attia, M. H. Attia, J. Iskander, K. Saleh, D. Nahavandi, A. Abobakr, M. Hossny, and S. Nahavandi. Fingerprint Synthesis Via Latent Space Representation. In *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*, pages 1855–1861.
- [15] B. Bhanu and X. Tan. Fingerprint indexing based on novel features of minutiae triplets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(5):616–622, 2003.
- [16] P. Bontrager, A. Roy, J. Togelius, N. D. Memon, and A. Ross. DeepMasterPrints: Generating MasterPrints for Dictionary Attacks via Latent Variable Evolution. In *9th IEEE International Conference on Biometrics: Theory, Applications and Systems, BTAS, 2018*, pages = 1–9.
- [17] K. Cao and A. K. Jain. Fingerprint indexing and matching: An integrated approach. In *IEEE International Joint Conference on Biometrics, IJCB, USA, 2017*, pages 437–445.
- [18] K. Cao and A. K. Jain. Fingerprint Synthesis: Evaluating Fingerprint Search at Scale. In *International Conference on Biometrics, ICB, 2018*, pages 31–38.
- [19] R. Cappelli, M. Ferrara, and D. Maltoni. Generating synthetic fingerprints. In *Hand-Based Biometrics: Methods and technology, IET, 2018*.
- [20] R. Cappelli, M. Ferrara, and D. Maltoni. Fingerprint indexing based on minutiae cylinder-code. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(5):1051–1057, 2011.
- [21] R. Cappelli, D. Maio, and D. Maltoni. Synthetic Fingerprint-Database Generation. In *16th International Conference on Pattern Recognition, ICPR 2002*, pages 744–747.
- [22] S. Chen, S. Chang, Q. Huang, J. He, H. Wang, and Q. Huang. SVM-Based Synthetic Fingerprint Discrimination Algorithm and Quantitative Optimization Strategy. *PLOS ONE*, 9, 2014.
- [23] Y. Chen and A. K. Jain. Beyond Minutiae: A Fingerprint Individuality Model with Pattern, Ridge and Pore Features. In *Advances in Biometrics, Third International Conference, ICB, Italy*, volume 5558 of *Lecture Notes in Computer Science*, pages 523–533. Springer, 2009.
- [24] D. Deb, T. Chugh, J. Engelsma, K. Cao, N. Nain, J. Kendall, and A. K. Jain. Matching fingerphotos to slap fingerprint images. *arXiv preprint arXiv:1804.08122*, 2018.
- [25] J. J. Engelsma, K. Cao, and A. K. Jain. Learning a fixed-length fingerprint representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2019.
- [26] F. J. Massey Jr. The Kolmogorov-Smirnov Test for Goodness of Fit. *Journal of the American Statistical Association*, 46(253):68–78, 1951.
- [27] I. S. Fogel and D. Sagi. Gabor filters as texture discriminator. *Biological Cybernetics*, 61:103–113, 2004.
- [28] R. S. Germain, A. Califano, and S. Colville. Fingerprint Matching Using Transformation Parameter Clustering. *IEEE Computational Science and Engineering*, 4(4):42–49, Oct. 1997.
- [29] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative Adversarial Nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in NeurIPS 27, Canada*, pages 2672–2680, 2014.
- [30] C. Gottschlich and S. Huckemann. Separating the real from the synthetic: minutiae histograms as fingerprints of fingerprints. *IET Biometrics*, 3(4):291–301, 2014.
- [31] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved Training of Wasserstein GANs. *CoRR*, abs/1704.00028, 2017.
- [32] P. A. Johnson, F. Hua, and S. A. C. Schuckers. Texture Modeling for Synthetic Fingerprint Generation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 154–159, 2013.
- [33] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR San Diego, Conference Track Proceedings*, 2015.
- [34] D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. 2014.
- [35] K. G. Larkin and P. A. Fletcher. A coherent framework for fingerprint analysis: are fingerprints holograms? *Opt. Express*, 15(14):8667–8677, Jul 2007.
- [36] J. Masci, U. Meier, D. C. Ciresan, and J. Schmidhuber. Stacked Convolutional Auto-Encoders for Hierarchical Fea-

- ture Extraction. In *21st International Conference on Artificial Neural Networks*, volume 6791 of *Lecture Notes in Computer Science*, pages 52–59. Springer, 2011.
- [37] S. Minaee and A. Abdolrashidi. Finger-gan: Generating realistic fingerprint images using connectivity imposed GAN. *CoRR*, abs/1812.10482, 2018.
- [38] A. Radford, L. Metz, and S. Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. In *4th International Conference on Learning Representations, ICLR, Puerto Rico, Conference Track Proceedings*, 2016.
- [39] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *Proceedings of the 31th International Conference on Machine Learning, ICML, China*, volume 32 of *JMLR Workshop and Conference Proceedings*, pages 1278–1286, 2014.
- [40] B. G. Sherlock and D. M. Monro. A model for interpreting fingerprint topology. *Pattern Recognition*, 26(7):1047–1055, 1993.
- [41] Y. Su, J. Feng, and J. Zhou. Fingerprint indexing with pose constraint. *Pattern Recognition*, 54:1–13, 2016.
- [42] E. Tabassi. NFIQ 2.0: NIST Fingerprint image quality. *NISTIR 8034*, 2016.
- [43] S. Yoon and A. K. Jain. Longitudinal study of fingerprint recognition. *Proceedings of the National Academy of Sciences*, 112(28):8555–8560, 2015.
- [44] Q. Zhao, A. K. Jain, N. G. P. Jr., and M. Taylor. Fingerprint image synthesis based on statistical feature models. In *IEEE International Conference on Biometrics: Theory, Applications and Systems, BTAS*, pages 23–30. IEEE, 2012.

Appendix

A. Improved-WGAN Architecture

The architecture of I-WGAN [31] used in our approach follows the implementation of [38]. Both the generator G and the discriminator D consist of seven convolutional layers with a kernel size of 4×4 and a stride of 2×2 which are detailed in table 4.

Generator G/G_{dec} architecture		
Operation	Output size	Filter size
Input	512×1	
Fully Connected	$16,384 \times 1$	
Reshape	$4 \times 4 \times 1024$	
Deconvolution	$8 \times 8 \times 512$	4×4
Deconvolution	$16 \times 16 \times 256$	4×4
Deconvolution	$32 \times 32 \times 128$	4×4
Deconvolution	$64 \times 64 \times 64$	4×4
Deconvolution	$128 \times 128 \times 32$	4×4
Deconvolution	$256 \times 256 \times 16$	4×4
Deconvolution	$512 \times 512 \times 1$	4×4

Table 4: Architecture of G/G_{dec} in I-WGAN. The discriminator architecture D is an inverse of the architecture G .

B. Training Details

The training of both CAE and I-WGAN were implemented using Tensorflow on an Intel Core i7-8700K @ 3.70GHz CPU with a RTX 2080 Ti GPU. The weights were randomly initialised from a gaussian distribution with mean 0 and standard deviation of 0.02. The optimizer function for the CAE was Adam [33] with a fixed learning rate of 0.0002, β_1 as 0.5, and β_2 as 0.9. The batch size was 128 and the number of training steps were 39,000. The Adam optimizer [33] was also used in the training of the I-WGAN with a fixed learning rate of 0.0001, β_1 as 0.0 and β_2 as 0.9. For I-WGAN, the batch size was set to 32 with 54,000 training steps. However, when I-GAN was finetuned to the dataset of plain fingerprints, it was trained for 37,000 steps.