This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation.

Except for this watermark, it is identical to the accepted version;

the final published version of the proceedings is available on IEEE Xplore.

Synthetic Latent Fingerprint Generator

André Brasil Vieira Wyzykowski¹ and Anil K. Jain¹ ¹Michigan State University {wyzykow2, jain}@msu.edu

Abstract

Given a full fingerprint image (rolled or slap), we present CycleGAN models to generate multiple latent impressions of the same identity as the full print. Our models can control the degree of distortion, noise, blurriness and occlusion in the generated latent print images to obtain Good, Bad and Ugly latent image categories as introduced in the NIST SD27 latent database. The contributions of our work are twofold: (i) demonstrate the similarity of synthetically generated latent fingerprint images to crime scene latents in NIST SD27 and MSP databases as evaluated by the NIST NFIQ 2 quality measure and recognition accuracies obtained by a SOTA fingerprint matcher, and (ii) use of synthetic latents to augment small-size latent training databases in the public domain to improve the performance of DeepPrint, a SOTA fingerprint matcher designed for rolled to rolled fingerprint matching on three latent databases (NIST SD27, NIST SD302, and IIITD-SLF). As an example, with synthetic latent data augmentation, the Rank-1 retrieval performance of DeepPrint is improved from 15.50% to 29.07% on challenging NIST SD27 latent database. Our approach for generating synthetic latent fingerprints can be used to improve the recognition performance of any latent matcher and its individual components (e.g., enhancement, segmentation and feature extraction). https://prip-lab.github.io/ Synthetic-Latent-Fingerprint-Generator/

1. Introduction

Since the first use of fingerprints in criminal investigations in 1891, fingerprints have become the most widely used tool for accurately and quickly identifying perpetrators [5]. The frequent use of fingerprints in judicial courts worldwide has demonstrated their acceptance as scientific evidence in the conviction of criminals. Within this context, latent fingerprints - those collected from various surfaces in crime scene investigations, are essential for identifying suspects involved in a crime. The identification of these latent prints by comparing them against rolled/slap fingerprints in law enforcement databases continues to be crucial for criminal investigations.

Latent fingerprint images often contain various sources of noise, including blood, contaminants, and natural secretions. Furthermore, fingers leave distinct traces on surfaces, such as cellulose composites, glass, metal, and plastic, which lead to variations in the latent images captured in forensic data collection [41]. In addition, each surface where a person leaves their fingerprints provides distinct background for friction ridge impressions, such as printed or handwritten text, as illustrated in Figure 1.



Figure 1: Examples of a latent fingerprint (bottom row) and its rolled fingerprint mate (top row) from the NIST special database 27 (SD27) [13]. The second column shows the fingerprints with their minutiae points representation, and the last column shows the deep-network embeddings (192dimensional) by DeepPrint [9] represented as a heatmap.

Given the poor quality of latent fingerprints, one of the most challenging problem in fingerprint recognition is to

This material is based upon work supported by the U.S. Department of Homeland Security under Grant Award Number 17STCIN00001-05-00. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

perform identification using latent prints. Given a query latent print, it needs to be compared with large collections of rolled and plain fingerprints in forensic databases. Jain and Feng [19] designed a latent-to-rolled matching algorithm based on minutiae, singular points (delta and core), and ridge flow map. Later, Cao et al. [4] improved the recognition accuracy by implementing a deep-learningbased representation to enhance poor-quality latent images prior to minutiae extraction. These studies demonstrate that there are many open challenges, primarily related to matching poor-quality latent prints with relatively high-quality rolled/slap fingerprints captured under supervision during booking. Indeed, NIST ELFT evaluation [18] showed that the best performing latent matching algorithm was able to provide only 67% Rank-1 accuracy in retrieving the mate of a query latent with a gallery of 100K rolled prints. When we compare this with rolled-to-rolled matching accuracy of over 99%, then it is clear that latent fingerprint recognition offers fertile grounds for research [35].

Within this context, several factors make it difficult to conduct research on latent fingerprint recognition:

- 1. Sparsity of public domain databases available for training and evaluating a latent fingerprint recognition system. Table 1 contains a summary of the latent fingerprint databases which have been used in academic publications. Figure 2 shows latent fingerprint examples from some of these databases. Note that the only two latent databases that are in the public domain are NIST SD302 and IIITD-SLF which were collected in a laboratory setting and not from crime scenes. One of the most popular latent fingerprint database in academic publications is NIST SD27 [13], but it was withdrawn from the public domain by NIST. While NIST SD302 and IIITD-SLF are useful for research, they are relatively small and do not capture the type of distortion and background that is observed in operational databases such as NIST SD27 (see Figures 1 and 2).
- 2. Limited variations in the style and quality of the latent images in public domain databases. These factors include small friction ridge area, small no. of minutiae points, blurred regions, and background complexity. The study by Gonzalez et al. [23] demonstrates the challenges posed by such variations in latent fingerprint recognition accuracy; detection of even a small no. of spurious minutiae or failure to detect a few genuine minutiae can drastically impact the latent recognition accuracy. Given the complexity of the latent to rolled fingerprint matching problem, the lack of publicly available operational latent fingerprint databases has hindered the progress in latent fingerprint recognition despite its critical role in law enforcement and forensics [7, 32]. According to the Innocence Project, "Forensic science, or more specifically, problems in

forensic science, contributes to many wrongful convictions, as seen in nearly half (45%) of DNA exoneration cases and one-quarter (24%) of all exonerations in the United States." and this includes errors in latent fingerprints matching.

Database	Public domain	# of unique fingers (# rolled-latent pairs)	Collection details
NIST SD27[13]	No	258	Crime scene images
MSP Latent DB [42]	No	1,910	Crime scene images
NIST SD302 [12]	Yes	1,019*	Laboratory collection
IIITD-SLF [43]	Yes	150	Laboratory collection
*Olto 1 - ft - ft - ft		- 1 1 1 - + + +	· · · · · · · · · · ·

*Obtained after filtering the rolled and latent mates present in the finger position annotation in the SD302h subset. The total no. of latent fingerprints is 9,990.

Table 1: Latent Fingerprint databases used in this study.

The objective of this research is to develop a method to generate a large collection of realistic synthetic latent fingerprint images that would be publicly available to interested researchers. The purpose of latent fingerprint synthesis is to advance pre-processing of latent images, including: (i) enhancement and segmentation, (ii) data augmentation for learning a robust representation, and (iii) designing a robust matcher for latent to rolled matching.

With these objectives, this paper aims to design and develop models to generate a large collection of synthetic latent fingerprints from a given set of rolled prints. For each rolled print, we are able to generate multiple latent prints with different quality levels (Good, Bad and Ugly following the categorization of NIST SD27 database). We validate our synthesis approach by fine-tuning a state-of-the-art fingerprint matcher, DeepPrint [9] with synthesized images. We show that this data augmentation leads to improved performance of DeepPrint on latent fingerprint recognition on NIST SD27, NIST SD302 and IIITD-SLF latent databases.

Our contributions in this paper are as follows:

- 1. A method to generate synthetic latent fingerprints of different difficulties in terms of recognition.
- 2. Demonstrate, both qualitatively and quantitatively, the similarity of our synthetic latent fingerprints to real latents.
- Utility of synthetic latents for data augmentation to improve the performance of a state-of-the-art deep network model for latent fingerprint recognition.

Note that our objective here is not to build a state-ofthe-art latent fingerprint recognition system but to present models for synthetic latents generation to assist in designing a SOTA latent matcher. For this reason, we do not address issues related to pre-processing of latents, including enhancement and segmentation, typically done in COTS latent matchers.

https://innocenceproject.org/
forensic-science-problems-and-solutions/



Figure 2: Latent fingerprint examples from IIITD-SLF, NIST SD27, MSP latent and NIST SD302.

2. Background

We first describe general fingerprint synthesis approaches in Section 2.1, followed by methods for synthesizing latent fingerprints in Section 2.2.

2.1. Fingerprint synthesis

One of the earliest and most prominent fingerprint synthesis generator is SFinGe[6], which used handcrafted methods to generate ridge patterns, minutiae, and textures formed by the ridge-valley pattern. But due to the limitations of the model used as well as the computational resources at the time, the images lacked visual realism.

With the introduction of generative adversarial networks (GAN) [14], the next generation of synthetic fingerprint approaches [3, 25, 1] emerged which are capable of generating plain and rolled fingerprints. However, like SFinGe, they are not designed for latent fingerprint synthesis. Also, these generators only create a single impression per finger.

Engelsma et al. [10] proposed a state-of-the-art approach called PrintsGAN capable of generating a large number of distinct identities with multiple rolled/plain impressions per finger. Although PrintsGAN has shown promise for data augmentation, some minutiae are lost, or spurious ones are introduced during the creation of multiple impressions of an identity. Like PrintsGAN, Grosz and Jain [15] developed a model capable of synthesizing multiple fingerprint spoof impressions from the same identity.

Other fingerprint synthesis approaches use CycleGAN [44] to generate realistic textures. Wyzykowski et al. [37, 38] used CycleGAN to generate additional medium and high-resolution 500 ppi and 1250 ppi fingerprints images, corresponding to the SFinGe-based identity generator. Sams et. al [31] consolidate the use of CycleGAN, enabling the generation of new identities with StyleGAN 2 [20]. Despite these advances, all of the aforementioned approaches did not focus on latent fingerprint synthesis.

2.2. Latent fingerprint synthesis and reconstruction

Ozturk et al. [26] developed an algorithm for automated latent fingerprint recognition. Although the paper's focus is not on latent fingerprint synthesis, the authors used a private dataset of latent fingerprints to generate a model with CycleGAN capable of converting images from FVC databases of plain fingerprints [24] into latent prints. However, the authors did not use synthetic latent data for data augmentation purposes to show the utility of their synthesized latent database. Further, they did not (i) provide the number of identities generated, and (ii) evaluate the similarity between synthetic and real fingerprint latents.

Xu et al. [40] demonstrated that they can improve matching performance by using data augmentation in latent image reconstruction. However, the authors did not specify the metrics for evaluating latent fingerprints reconstructions, such as a minutiae detection analysis and their image quality.

3. Proposed Latent Fingerprint Generator

This study's main objective is to create latent fingerprints having control of the style generated by these models. This feature is essential since latents in different databases vary depending on the surface where the prints were left. We detail our method in Figure 3 with the pseudocode in Algorithm 1.

In Section 3.1 we detail the modifications and parameters used in CycleGAN in both stages of our method. In Section 3.2, we describe the first stage of the fingerprint synthesis method, where we create a generic model capable of transforming an image from the rolled domain into a generic latent domain. Then, Section 3.3 details the second stage, where we create specific models for each latent print style.

3.1. CycleGAN modification details

CycleGAN [44] is a neural network model focused on the process of unpaired image-to-image translation, i.e., a network able to map two distinct image domains and perform the transformation of image features from these domains. Through the Cycle-Consistency loss, CycleGAN can perform two-way style translation, transforming input images into styles of output images and vice versa. Furthermore, CycleGAN does not require paired examples to train, an important advantage when working with latent images, since there is no pairing between the latent and rolled impression in many databases. However, we noticed stability problems during training when we performed our initial training with the original CycleGAN architecture on latent



Figure 3: Steps to create synthetic latent fingerprints using the proposed approach.

Algorithm 1 Latent fingerprint synthesis

Input: M(i) set of rolled prints.

Output: Synthetic latent prints from M(i) denoted as SLP(i).

First stage:

Generate a coarse CycleGAN style model *GSM* using mated rolled and latent prints as training data. *output: GSM CycleGAN model*

Second stage: input: (GSM CycleGAN model)

Use ResNet152V2 [17] to extract features from the latent images.

Use K-Means for clustering the latent prints.

for each k cluster do

Generate a fine-tuned CycleGAN style model SM_k using rolled and clustered latent prints as training data. end for

for each CycleGAN SM_k model do for each rolled prints in M(i) do

Generate *SLP* synthetic latents prints.

```
end for
```

end for

output: SLP(i) synthetic latents prints.

prints. Other authors [33, 22] also report similar problems using CycleGAN on different image styles. Thus, we decided to use a global discriminator (already present in Cyle-GAN) and a patch discriminator, which we detail in Figure 4. The patch discriminator stabilized our training, generating better details of the fingerprint ridges. We use a public implementation of CycleGAN to perform the modifications described in this section.

We used Leaky Relu [39] as the activation function in the generator with a negative slope of $\alpha = 0.2$. For the global and patch discriminators, we keep the original CycleGAN

activation function, Relu. Additionally, we also used instance normalization.



Figure 4: CycleGAN architecture with discriminators.

There are six residual blocks [16] in our CycleGAN architecture. We train our model using the Adam optimizer [21] with $\beta_1 = 0.5$ and $\beta_2 = 0.999$. Our weight for cycleconsistency loss was $\lambda = 10$. We used a learning rate of 0.0003 for the generator and 0.0001 for the discriminator. We early-stop [28] the training if the loss does not improve after 50 epochs.

Our work aims to transform rolled fingerprints into latent fingerprints, not to reverse the process (rolled into latent). However, by utilizing Cycle-Consistent Loss, CycleGAN can improve latent images, which may further enhance the performance of fingerprint recognition systems, an investigation we plan to undertake in future work.

As described in Sections 3.2 and 3.3, we use our modified version of CycleGAN for stages 1 and 2.

3.2. First stage: rolled to latent generation

This first stage aims to create a CycleGAN "coarse" model capable of mapping latent and rolled domains. Using this "coarse" model as input, improved fine-tuned models are created in the second stage (see Section 3.3).

3.3. Second stage: clustering and model refinement

A latent fingerprint image can have different styles depending on where the fingerprint traces were left or if noise, dirt, blood, etc., were present. Therefore, reproducing this inherent variation is vital. Thus, the second stage involves

https://github.com/towardsautonomy/CycleGAN_
improved

clustering latent prints and generating custom CycleGAN models tailored to the styles of each cluster. First, in Section 3.3.1, we detail our approach to clustering latent fingerprints. Then, Section 3.3.2 describes the finetuning process in CycleGAN, generating models capable of creating styles corresponding to clusters.

3.3.1 Cluster latent images based on visual similarity

We first used a pre-trained public ResNet152V2 [17] model to extract feature vectors of the latent images. This network was trained using the ImageNet-21K dataset [29]. We use the last fully-connected layer of ResNet152V2 with 2048 output nodes as a feature vector to cluster the synthetic images with K-Means. As a result of this process, we obtain k sets of latent images, each containing distinct patterns and visual features. These k clusters then serve as input to perform a finetuning of CycleGAN in stage 2 (see Section 3.3.2).

3.3.2 Generating multiple latent print styles

We use the CycleGAN model developed in the first stage to generate multiple latent print styles and finetune the k clusters based on visual similarity. Each of the k CycleGAN models generates a different variation of the characteristics present in each latent group. Therefore, if the number of latent images per cluster is small, CycleGAN may not converge well. Section 4.1 details data augmentation operations that help the training converge better.

With k finetuned CycleGAN models. Greater control can be exerted over the distribution of synthetic latent fingerprint styles. Our approach uses identities (represented as rolled fingerprints) obtained from the NIST SD4 database for validation in our experiments. However, any other rolled fingerprint database can be used, for example, even the synthetic rolled images from PrintsGAN [10] for generating new synthetic identities. PrintsGAN is a neural method capable of generating unique identities by simulating parameters of distortions, textures, and finger pressure on a surface, among others. In combination, PrintsGAN and the ktexture models can generate a wide range of latent-rolled mated pairs of fingerprints with "virtual" identities.

4. Experimental Results

4.1. Synthetic latent fingerprint generation

During the first stage of our algorithm to create the latent fingerprint set, we train a CycleGAN model that receives as input 2,000 rolled fingerprints from the first acquisition of the NIST SD4 [36]. For our target texture transfer, we use 2,074 latent fingerprints from the MSP latent database. At the end of this first stage, we have a model capable of converting rolled images into latent images.

Although this model trained in the first stage is already capable of generating latent fingerprints, we aim to generate greater diversity in the visual characteristics of the latent images, providing different matching performance behaviors. In the second stage, we choose 3 clusters (k = 3) since the NIST SD27 latents have been partitioned into three categories, namely good, bad, and ugly, with three distinct levels of recognition difficulty and visual texture.

Given the three clusters in the first stage, we execute the stage 2 of our method, finetuning the coarse model with the styles for each cluster. We applied data augmentation operations on the fly during training by doing a random vertical and horizontal translation of the images (maximum of 100 pixels) and a random rotation with a maximum of 15 degrees. This data augmentation operation does not change the identity of the fingerprint. Finally, after training, we have three CycleGAN latent print models each with different characteristics, including distortion, texture and friction ridge area, a factor that provides more control in generating latent images.

To create our final set of synthetic latents, we used rolled images from the second acquisition of the NIST SD4 [36] and used the three CycleGAN models to generate 2,000 new latent images. We now have a database of 2,000 synthetic latents and their 2,000 rolled mated images from NIST SD4. Figure 5 shows a visual comparison between the good, bad, and ugly groups from the MSP latent database, NIST SD27, and our synthetic latent images.

4.2. Similarity of real and synthetic latent embeddings

To validate the similarities between our synthetic latents and real latent fingerprints in MSP latent and NIST SD27 databases, we first used NFIQ 2 [2] to compare the quality of latents. NFIQ 2 provides a score in the range [0, 100], where higher the score, better the image quality. Figure 6 contains the frequency histogram of the scores obtained from the MSP latents, NIST SD27, and our synthetic latent prints (SLP), showing that the quality of NIST SD27 latent images is, overall, better than MSP latents and our synthetic latents. The quality of our synthetic latents is similar to the MSP latent images, which is reasonable since we use the MSP latent in our training of CycleGAN.

To further illustrate the similarity of synthetic latents to MSP latents and NIST SD27 latents, we show a t-Distributed Stochastic Neighbor Embedding (t-SNE) plot in Figure 7. We see that the placement of 2D embeddings of latent fingerprints of real databases and our synthetic latents is similar. These results show that we are able to realistically generate latent fingerprints that replicate visual

https://www.tensorflow.org/api_docs/python/tf/ keras/applications/resnet_v2/ResNet152V2



Figure 5: Good, Bad and Ugly latents from NIST SD27, MSP database and our synthetic latent prints (SLP). NIST SD27 database comes with a subjective categorization into Good, Bad and Ugly classes. The same categorization is used for MSP latents and our synthetic fingerprints.

features of real latent images.



Figure 6: Histogram of NFIQ 2 scores. Note that the quality histogram of synthetic latents generated by our method overlaps the quality histograms of crime scene latents from NIST SD27 and MSP databases.

Finally, for each of the three CycleGAN models, one each for Good, Bad and Ugly latents, we matched the synthetic latent images with their mated rolled images from the NIST SD4 database using Verifinger V12.3 SDK. Fig 9. shows the ROC plots for the three groups of synthetic latents as well as the NIST SD27 database. This comparison shows that the recognition difficulty of Good, Bad and Ugly synthetic prints is similar to the difficulty levels of the corresponding three types of latents in NIST SD27 latents. The



Figure 7: 2D t-SNE plot of 192D embeddings showing the overlap of NIST SD27, MSP and synthetic latents images.

TDR of "good" synthetic images is 83.22% and the corresponding TDR value for "good" images in NIST SD27 is 92.05%. The difference in TDR can be explained as follows: the MSP latent database used for training CycleGAN has a different quality distribution than NIST SD27. Our goal is not to replicate NIST SD27 images, but to design a method capable of generating latents with distinct quality for data augmentation.

4.3. Minutiae analysis

To further analyze the visual differences between Good, Bad, and Ugly latent images generated by CycleGAN models, in Figure 8, we show Good, Bad and Ugly latent images generated from a single NIST SD4 identity, overlaid with the minutiae extracted by Verifinger V12.3. Note that the number of minutiae decreases in images generated by Good, Bad and Ugly models. To make a quantitative comparison of minutiae count between real and synthetic latent prints, we generated 258 latent prints, the same as in NIST SD27. The mean minutiae count for the NIST SD27 are: 68, 45, and 35 for Good, Bad and Ugly latents, respectively. The corresponding minutiae count mean for our synthetic prints are: 55, 39, and 35, respectively.



Figure 8: Good, Bad and Ugly latents synthesized from the rolled shown in (a). No. of minutiae in the images (a)-(d) by Verifinger V12.3 SDK are shown in parentheses. Generally, lower the minutiae count, lower the image quality.



Figure 9: ROC curves for synthetic latents prints (SLP) generated from NIST SD4 and NIST SD27. Good, Bad and Ugly ROC curves for the two comparisons appear to be grouped together in these plots. TDR (%) reported @ FAR = 0.1%.

4.4. Data augmentation to increase accuracy

This experiment aims to show that it is possible to increase the matching accuracy of the rolled-to-rolled matcher DeepPrint [9] by using the synthetic latent prints (SLP) as data augmentation for retraining DeepPrint. In performing this experiment, we seek to ascertain the relative performance increase within the scope of DeepPrint and not relative to other matchers, thus validating the usefulness of synthetic latent images to improve the recognition performance of any fingerprint matcher for latent recognition.

As a test set for our experimentation, we used NIST SD27 [13], NIST SD302 (N2N) [12] and IIITD-SLF [43], performing the matching between latent images and mated rolled/slap images from each of these databases. Furthermore, as our method allows the controlled synthesis of Good, Bad and Ugly latents, we only utilize Bad and Ugly latents for data augmentation as these provide higher matching challenge.

We designed the experiment as follows: (i) Use the original DeepPrint model [9] as a reference model. (ii) Finetune DeepPrint with MSP latent database images, finding an optimal value for the localization network hyperparameter LN in DeepPrint's architecture in the context of latent prints. Using the same LN value as in (ii) allows direct comparison and finetuning DeepPrint with our synthetic SLP latent prints. (iii) Generate a variation in the model that allows us to perform score level fusion, and perform finetuning using the same images as (ii), but with a different LN value that still provides satisfactory recognition results. The models used in our experiments are summarized in Table 2.

Model	Training data	Localization network hyper-parameter			
DeepPrint [9]	455K MSP rolled prints	0.035			
$DeepPrint_1$	518 MSP aligned pairs of rolled and latents	0.018			
$DeepPrint_2$	2K SD4 (rolled, synthetic latents) pairs (bad, ugly)	0.018			
$DeepPrint_3$	2K SD4 (rolled, synthetic latents) pairs (bad, ugly)	0.007			
$DeepPrint_2$ and $DeepPrint_3$ models do not use any real latent prints for					
retraining.					

Table 2: Models, training data and hyper-parameters.

In addition to the models in Table 2, we applied Min-Max normalization to the scores generated by the $DeepPrint_2$ and $DeepPrint_3$ models and performed simple score level fusion [30] as $ScoreFusion = \frac{(ns_2+ns_3)}{2}$, where ns_2 and ns_3 are the normalized scores from $DeepPrint_2$ and $DeepPrint_3$. The "fused" model is referred to as $DeepPrint_4$. Figure 10 shows the ROC curve of the DeepPrint [9], $DeepPrint_1$, $DeepPrint_2$, $DeepPrint_3$ and $DeepPrint_4$ models evaluated on NIST SD27 [13].



Figure 10: ROC curves of the DeepPrint [9] and its variants evaluated on NIST SD27 [13].

In all cases, our models trained with the synthetic images performed better than the original DeepPrint model and the finetuned models trained with the MSP latents on NIST SD27. We arrange all the experimental analysis results in Table 3. Only with the latent MSP latents was it possible to increase the system's performance accuracy. Still, it was necessary to perform alignment between the rolled and latent mates, which is also a challenging task. This alignment is needed because the Localization network in DeepPrint cannot perform such a precise alignment with latent images. However, the alignment problem does not exist with synthetic images since we can control the latent synthesis directly from a rolled print. These results corroborate that Verifinger V12.3 is one of the best-performing matchers, consistent with the results reported in FVC-onGoing [8] and NIST FpVTE [35].

Model	NIST SD27	NIST SD302 (N2N)*	IIITD-SLF
DeepPrint [9]	26.35	10.81	10.0
$DeepPrint_1$	29.45	13.15	14.16
$DeepPrint_2$	32.55	13.97	15.0
$DeepPrint_3$	37.20	14.07	17.5
$DeepPrint_4$	39.92	14.31	20.83
Verifinger V12.3	55.81	16.22	29.70

*N2N latent and rolled mates present in the finger position annotation in the SD302h subset. We reduced the resolution of N2N latent images to 500dpi and applied the Clahe filter [27] to highlight the fingerprint ridges.

Table 3: True detection rate (TDR (%)) @ FAR = 0.01%

4.4.1 Identification (1:N Comparison)

Our goal of this experiment is to perform a closed-set analysis of identification accuracy by matching latent prints in NIST SD27, NIST 302 (N2N) and IIITD-SLF against a galley augmented by 62,871 rolled fingerprints from the NIST SD300a (8,871k) [11] and NIST SD14 [34] (54K) databases. The results of this analysis are shown in Table 4 and the Cumulative Matching Characteristics (CMC) of the NIST SD27 is presented in Figure 11. We used the same protocol as NIST ELFT-EFS [18] to evaluate the performance of our models.

Model	NIST SD27	NIST SD302 (N2N)*	IIITD-SLF
DeepPrint [9]	15.50	4.18	9.8
$DeepPrint_1$	23.64	5.98	14.13
$DeepPrint_2$	26.74	6.56	14.7
$DeepPrint_3$	28.29	7.24	17.12
$DeepPrint_4$	29.07	7.93	18.21
Verifinger V12.3	45.45	10.81	22.71
*We used all 0 000 lat	ents in N2N_We	reduced the resolution of N	V2N latent image

to 500dpi and applied the Clahe filter [27] to highlight the fingerprint ridges.

Table 4: Rank-1 accuracies (%) of different DeepPrint models against a background set of 62,871 rolled fingerprints (54K from NIST SD14 and 8,871K from NIST SD300a), in addition to the true rolled mates of the latents prints of each database.



Figure 11: CMC curves of various DeepPrint models evaluated on NIST SD27 latent database.

5. Conclusion

We present a new approach capable of synthesizing latent fingerprints from any rolled fingerprint database. To accomplish this, we used a modified version of CycleGAN by adding a patch discriminator in addition to the global discriminator already present in the original CycleGAN. Our approach generates models capable of replicating Good, Bad and Ugly latent images present in the NIST SD27 latent database, thus allowing the generation of style-controlled latent images.

Additionally, by comparing NFIQ 2 quality values and t-SNE plots in two dimensions and the recognition performance of the latent to their mated rolled images we can claim that the synthetic latent images are realistic and similar to the real latent images.

Finally, we performed a finetuning on the original Deep-Print [9] model using our SLP latent images generated with our modified CycleGAN model. We achieved a performance boost, validating this on three latent databases. In future work, we plan to refine the synthesis of latent prints by adding more complexity to background and texture variations. Furthermore, CycleGAN is capable of reverting a latent into a rolled one. With this, we plan to improve the quality of latent prints, thus improving latent recognition in deep networks.

References

- Keivan Bahmani, Richard Plesh, Peter Johnson, Stephanie Schuckers, and Timothy Swyka. High fidelity fingerprint generation: Quality, uniqueness, and privacy. In 2021 IEEE International Conference on Image Processing (ICIP), pages 3018–3022. IEEE, 2021.
- [2] Oliver Bausinger and Elham Tabassi. Fingerprint sample quality metric nfiq 2.0. *BIOSIG 2011–Proceedings of the Biometrics Special Interest Group*, 2011.
- [3] Kai Cao and Anil Jain. Fingerprint synthesis: Evaluating fingerprint search at scale. In 2018 International Conference on Biometrics (ICB), pages 31–38. IEEE, 2018.
- [4] Kai Cao, Dinh-Luan Nguyen, Cori Tymoszek, and Anil K Jain. End-to-end latent fingerprint search. *IEEE Transactions on Information Forensics and Security*, 15:880–894, 2019.
- [5] Richard M Caplan. How fingerprints came into use for personal identification. *Journal of the American Academy of Dermatology*, 23(1):109–114, 1990.
- [6] R. Cappelli. Sfinge: an approach to synthetic fingerprint generation. *International Workshop on Biometric Technologies*, Jan 2004.
- [7] National Research Council et al. *Strengthening forensic science in the United States: a path forward.* National Academies Press, 2009.
- [8] Bernadette Dorizzi, Raffaele Cappelli, Matteo Ferrara, Dario Maio, Davide Maltoni, Nesma Houmani, Sonia Garcia-Salicetti, and Aurélien Mayoue. Fingerprint and on-line signature verification competitions at icb 2009. In *International Conference on Biometrics*, pages 725–732. Springer, 2009.
- [9] Joshua J Engelsma, Kai Cao, and Anil K Jain. Learning a fixed-length fingerprint representation. *IEEE Transactions* on pattern analysis and machine intelligence, 43(6):1981– 1997, 2019.
- [10] Joshua J. Engelsma, Steven A. Grosz, and Anil K. Jain. PrintsGAN: Synthetic fingerprint generator. *CoRR*, abs/2201.03674, 2022.
- [11] Gregory P Fiumara, Patricia A Flanagan, John D Grantham, Bruce Bandini, Kenneth Ko, John M Libert, et al. Nist special database 300: Uncompressed plain and rolled images from fingerprint cards. 2018.
- [12] Gregory P Fiumara, Patricia A Flanagan, John D Grantham, Kenneth Ko, Karen Marshall, Matthew Schwarz, Elham Tabassi, Bryan Woodgate, Christopher Boehnen, et al. Nist special database 302: Nail to nail fingerprint challenge. 2019.
- [13] Michael D Garris and Michael D Garris. NIST special database 27: Fingerprint minutiae from latent and matching tenprint images. US Department of Commerce, National Institute of Standards and Technology, 2000.
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. Advances in neural information processing systems. 27, 2014.
- [15] Steven A Grosz and Anil K Jain. Spoofgan: Synthetic fingerprint spoof images. arXiv preprint arXiv:2204.06498, 2022.

- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [18] Michael D Indovina, RA Hicklin, G I Kiebuzinski, et al. NIST evaluation of latent fingerprint technologies (ELFT) — extended feature sets evaluation 2. *National Institute of Standards and Technology, NIST*, 2012.
- [19] Anil K Jain and Jianjiang Feng. Latent fingerprint matching. *IEEE Transactions on pattern analysis and machine intelli*gence, 33(1):88–100, 2010.
- [20] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 8110–8119, 2020.
- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- [22] Taesung Kwon and Jong Chul Ye. Cycle-free cyclegan using invertible generator for unsupervised low-dose ct denoising. *IEEE Transactions on Computational Imaging*, 7:1354– 1368, 2021.
- [23] Octavio Loyola-González, Emilio Francisco Ferreira Mehnert, Aythami Morales, Julian Fierrez, Miguel Angel Medina-Pérez, and Raúl Monroy. Impact of minutiae errors in latent fingerprint identification: assessment and prediction. *Applied Sciences*, 11(9):4187, 2021.
- [24] Dario Maio, Davide Maltoni, Raffaele Cappelli, Jim L Wayman, and Anil K Jain. Fvc2004: Third fingerprint verification competition. In *International conference on biometric authentication*, pages 1–7. Springer, 2004.
- [25] Vishesh Mistry, Joshua J Engelsma, and Anil K Jain. Fingerprint synthesis: Search with 100 million prints. In 2020 IEEE International Joint Conference on Biometrics (IJCB), pages 1–10. IEEE, 2020.
- [26] Halil İbrahim Öztürk, Berkay Selbes, and Yusuf Artan. Minnet: Minutia patch embedding network for automated latent fingerprint recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1627–1635, 2022.
- [27] Stephen M Pizer, E Philip Amburn, John D Austin, Robert Cromartie, Ari Geselowitz, Trey Greer, Bart ter Haar Romeny, John B Zimmerman, and Karel Zuiderveld. Adaptive histogram equalization and its variations. *Computer vision, graphics, and image processing*, 39(3):355– 368, 1987.
- [28] Lutz Prechelt. Early stopping-but when? In Neural Networks: Tricks of the trade, pages 55–69. Springer, 1998.
- [29] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972*, 2021.

- [30] Arun Ross and Karthik Nandakumar. Encyclopedia of Biometrics. Fusion, Score-Level, pages 611–616. Springer US, Boston, MA, 2009.
- [31] Ataher Sams, Homaira Huda Shomee, and SM Rahman. Hqfingan: High-quality synthetic fingerprint generation using gans. *Circuits, Systems, and Signal Processing*, pages 1–16, 2022.
- [32] William Thompson, John Black, Anil Jain, Joseph Kadane, Deborah Runkle, Michelle Barretta, and Mark S Frankel. Latent fingerprint examination. *American Association for the Advancement of Science*, 2017.
- [33] Dmitrii Torbunov, Yi Huang, Haiwang Yu, Jin Huang, Shinjae Yoo, Meifeng Lin, Brett Viren, and Yihui Ren. Uvcgan: Unet vision transformer cycle-consistent gan for unpaired image-to-image translation. arXiv preprint arXiv:2203.02557, 2022.
- [34] Craig I Watson. NIST special database 14. Fingerprint Database, US National Institute of Standards and Technology, 1993.
- [35] Craig I. Watson, Gregory P. Fiumara, Elham Tabassi, Wayne J. Salamon, and Patricia A. Flanagan. Fingerprint vendor technology evaluation. Technical report, Dec. 2014.
- [36] Craig I Watson and Charles L Wilson. NIST special database 4. Fingerprint Database, National Institute of Standards and Technology, 1992.
- [37] André Brasil Vieira Wyzykowski, Mauricio Pamplona Segundo, and Rubisley de Paula Lemes. Level three synthetic fingerprint generation. In 2020 25th International Conference on Pattern Recognition (ICPR), pages 9250–9257. IEEE, 2021.
- [38] Andre Brasil Vieira Wyzykowski, Mauricio Pamplona Segundo, and Rubisley de Paula Lemes. Multiresolution synthetic fingerprint generation. *IET Biometrics*, 2022.
- [39] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *CoRR*, abs/1505.00853, 2015.
- [40] Ying Xu, Yi Wang, Jiajun Liang, and Yong Jiang. Augmentation data synthesis via gans: Boosting latent fingerprint reconstruction. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2932– 2936. IEEE, 2020.
- [41] Brian Yamashita, Mike French, et al. Latent print development. *The fingerprint sourcebook, NIJ Journal*, 1:155–222, 2011.
- [42] Soweon Yoon and Anil K Jain. Longitudinal study of fingerprint recognition. *Proceedings of the National Academy of Sciences*, 112(28):8555–8560, 2015.
- [43] Rozita Mohd Yusof and Norrozila Sulaiman. Multi-probes fusion framework for matching simultaneous latent impressions. In 2012 International Conference for Internet Technology and Secured Transactions, pages 564–567. IEEE, 2012.
- [44] J. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired imageto-image translation using cycle-consistent adversarial networks. *CoRR*, abs/1703.10593, 2017.