

Robust Partial Fingerprint Recognition

Yufei Zhang^{1*}, Rui Zhao², Ziyi Zhao², Naveen Ramakrishnan²
Manoj Aggarwal², Gerard Medioni², Qiang Ji¹
¹Rensselaer Polytechnic Institute, ²Amazon

{zhangy76, jiq}@rpi.edu, {zhaori, zhaziyi, navramk, manojagg, medioni}@amazon.com

Abstract

*Low quality capture and obstruction on fingers often result in partially visible fingerprint images, which imposes challenge for fingerprint recognition. In this work, motivated from the practical use cases, we first systematically studied different types of partial occlusion. Specifically, two major types of partial occlusion, including six granular types, and the corresponding methods to simulate each type for model evaluation and improvement were introduced. Second, we proposed a novel **Robust Partial Fingerprint (RPF)** recognition framework to mitigate the performance degradation due to occlusion. RPF effectively encodes the knowledge about partial fingerprints through occlusion-enhanced data augmentation, and explicitly captures the missing regions for robust feature extraction through occlusion-aware modeling. Finally, we demonstrated the effectiveness of RPF through extensive experiments. Particularly, baseline fingerprint recognition models can degrade the recognition accuracy measured in $FRR @ FAR=0.1\%$ from 14.67% to 17.57% at 10% occlusion ratio on the challenging NIST dataset, while RPF instead improves the recognition performance to 9.99% under the same occlusion ratio. Meanwhile, we presented a set of empirical analysis through visual explanation, matching score analysis, and uncertainty modeling, providing insights into the recognition model's behavior and potential directions of enhancement.*

1. Introduction

Fingerprint recognition refers to identifying or verifying a person based on their fingerprints. It is an active research area with numerous practical applications. One crucial step in fingerprint recognition is to effectively and efficiently extract representation from fingerprint images in order to perform matching. Recent advances in deep learning have made significant strides in vision-based biometrics. It not

only simplifies the traditional methods that heavily rely on hand-crafted features such as minutia map [6, 11, 12, 24], but also demonstrates promising accuracy and efficiency compared to feature engineering methods [18, 36, 37]. However, current deep learning-based models primarily consider complete biometrics data as input, limiting their performance in practice where partial observations due to occlusion or incomplete image capture are common.

Existing fingerprint recognition models typically consider three types of fingerprint data: contactless fingerprints, latent fingerprints, and rolled/plain fingerprints. In contactless fingerprint recognition, some methods consider partial fingerprints, but only those caused by misalignment [31, 45], while some other works neglect the possibility of missing regions [15, 19]. Latent fingerprints, usually collected from uncontrolled environments, can be low image quality and can miss part of fingerprints [6]. Some work [30, 48] aim to recover the low quality or missing regions for recognition, but obtaining prior knowledge for full image restoration is difficult due to significant variations in the collection condition. Among these three types of fingerprints, rolled/plain fingerprints are more widely used in commercial products. These rolled/plain fingerprint data are typically collected using a fingerprint scanner with well-controlled data acquisition environment, resulting in high quality impressions. Nevertheless, partial observations can still occur due to external occlusion, poor scanner quality, incomplete capture, *etc.* Some work [8, 14, 51, 52] present customized neural network architectures and loss functions to extract robust feature directly from partially observed fingerprints. However, there lacks a systematic study on the sources of partial observations and their impact on the model performance, leading to ineffective handling of partial fingerprints.

In this work, we developed robust partial fingerprint recognition model focusing on rolled/plain fingerprints as it is widely used in practice yet the partial fingerprint issue is often overlooked. We first identified different types of partial fingerprints that may occur in practical use cases. Then, we proposed a robust partial fingerprint recognition model

*Work done while as an intern at Amazon.

that builds upon occlusion-enhanced data augmentation and occlusion-aware modeling. We demonstrated its effectiveness through extensive experimentation and provided several empirical analysis to interpret the model behavior. To summarize, our main contributions include:

- A systematic characterization of different types of partial fingerprints with the corresponding methods to simulate them for modeling partial fingerprints.
- A robust partial fingerprint recognition framework *i.e.* RPF, that effectively utilizes the knowledge of partial fingerprints through occlusion-enhanced data augmentation and occlusion-aware modeling. Extensive experiments demonstrate the advantages of RPF over the-state-of-the-art methods, particularly in its robustness for partial fingerprints recognition.
- A set of empirical analysis that provide valuable insights into the model’s behavior, leading to more informed decisions on how to improve its performance.

2. Related Work

We begin by discussing the application of deep learning to fingerprint recognition, including earlier studies of partial fingerprint recognition. We then summarize closely related work in face recognition and other computer vision tasks that address input with occlusion.

Fingerprint Recognition with Deep Learning: Deep learning-based fingerprint recognition has gained increasing interest in recent years. Nguyen *et al.* [35] used a deep learning model to extract accurate minutiae maps. [5, 12, 44] adopted deep learning models to directly extract image features, which are further combined with traditional minutiae map estimates to produce more discriminatory features. These deep fingerprint recognition models, however, mainly consider fully observed fingerprints, and their robustness to partial fingerprints has not been well-studied.

When dealing with partial fingerprints, several methods relied on local-patch information for matching [8, 14, 25], ignoring the global information. Others designed neural network architectures [22, 45, 51] and loss function [52] for robust feature extraction. However, existing methods are limited to certain types of partial fingerprints, such as incomplete images captured by small-area fingerprint scanners [51] or partial observation due to misalignment [45]. In contrast, we systematically study different types of partial fingerprints (*e.g.* object blocking, corner cropping, *etc.*). To effectively handle partial fingerprints, we developed robust models considering different strategies inspired by the work as discussed below.

Related Research Handling Occlusion: There are several common approaches for handling occlusion in face recognition [9, 10, 42, 46, 47, 53] and other applications, such

as facial landmark detection [54], 3D human body reconstruction [29], and multi-view depth estimation [3]. One approach is to use *occlusion-enhanced data augmentation*, *i.e.* augmenting training data with synthetic occlusion. For example, augmenting training images with synthetic face masks can effectively improve the robustness to mask occlusion [32, 49]. However, exploration of such strategy in fingerprint domain remains limited. We are the first to systematically study the possible cases of partial fingerprints and effectively leverage these knowledge to improve model robustness. Another approach is *occlusion-aware modeling*, that explicitly captures the occlusion region and extracts an attention map to effectively weight image features [29, 38, 43, 50, 54]. Inspired by the occlusion-aware models proposed in other domains, we exploited occlusion-aware modeling for fingerprint recognition. Particularly, our approach predicts a flexible fingerprint area map to accommodate various occlusion patterns and effectively generate the attention map for weighting image features.

To gain deeper insights into the recognition behavior, [40] used a gradient-based visual localization method to show the importance region for classification, while [34] built a salience map to illustrate the matching behaviors of face recognition models. These insightful analysis, however, are missing for fingerprints. In this work, we conducted detailed empirical studies to characterize the behavior of fingerprint recognition models. Meanwhile, missing of information typically results in large uncertainty in recognition. Despite uncertainty modeling have been explored in various computer vision tasks [3, 7, 17, 26, 27, 33, 41], it has not been well-studied for fingerprints. Therefore, we leveraged uncertainty modeling to analyze the challenges associated with partial fingerprint recognition and identified opportunities for further improvements.

3. Methods

3.1. Problem Statement

Our focus in this paper is on image-based fingerprint verification in the presence of partial fingerprints. The goal is to determine if a template and a query belong to the same identity (genuine pair) or a different identity (imposter pair). For experimental evaluation, we focus on the rolled/plain fingerprints, and assume that both the template and query fingerprints have a similar alignment but part of the fingerprint could be missing due to various reasons.

3.2. Different Types of Partial Fingerprints

In practice, occlusion or being out of the field-of-view (FoV) during image acquisition often results in an incomplete fingerprint capture, losing biometric information. Below we summarize the common cases of these partial fingerprints. We use occlusion ratio (OR) to measure the loss of

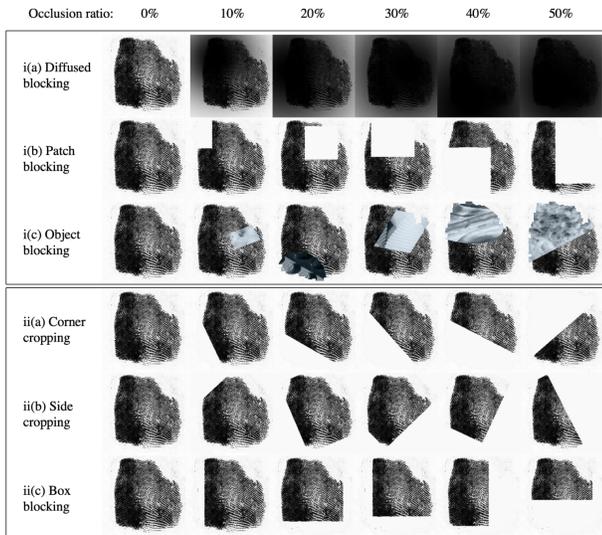


Figure 1. Simulation examples of different types of partial fingerprints at different ORs.

information. Specifically, we obtain the fingerprint area utilizing OpenCV [4]’s build-in edge detection and image morphology functions. Denoting the fingerprint area as s_f and the occluded fingerprint area as s_o , we define $OR = \frac{s_o}{s_f}$. We generate different partial fingerprints at different ORs with examples shown in Figure 1. We now discuss each type of partial fingerprints and our approach to simulate them.

Partial Fingerprints due to Occlusion: In practice, occlusion can occur due to dirt on fingers, poor quality of the fingerprint scanner, or intrusion of external objects. On this premise, we categorize the occlusion into three scenarios according to the cause:

- Diffused blocking (Figure 1-i(a)) simulates dirt, grease or sweat attached to the finger via blending the original image with black colors. We use 2D isotropic Gaussian distribution with mean μ , diagonal and equal variance σ^2 to describe the blending weight. During simulation, we randomly select μ and empirically specify the variance as $\sigma^2 = \frac{5s_o}{\pi}$, meaning the standard deviation is the radius of a circle with area equaling to $\sqrt{5}$ times of the desired occluded area.
- Patch blocking (Figure 1-i(b)) considers the occlusion due to dust or scratches on the fingerprint scanner and is simulated by blocking the pixels within a randomly selected patch. In details, the center of the patch to block is randomly selected, and the fingerprint pixels closest to the patch center within the occluded area s_o are erased (we use “inf” norm to measure the distance, resulting to rectangle patch).
- Object blocking (Figure 1-i(c)) simulates external object obstruction by occluding part of the image using the material textures from the ALOT dataset [1]. ALOT includes in total of 27500 texture images, vary-

ing illumination angle, color, and viewing angle. For object blocking, object center is randomly selected with the object size specified by OR. Meanwhile, the object shape can be ellipse or quadrilateral with random scale of the major axis to the minor axis and random direction of the major axis.

Partial Fingerprints due to Being Out of FoV: When the fingerprints are out of FoV, parts of the fingerprints that are close to the image boundary are cropped out. Based on our observation, there are three most common cases:

- Corner cropping (Figure 1-ii(a)) simulates the situation where one of the corners of the fingerprint image is out of FoV, and is generated by cropping out the corner using a straight line. Specifically, we first randomly select one of the four image corners to crop. For one side of the selected corner, we randomly select an anchor position. The cutting line is then obtained by finding the other anchor on the other side of the corner so that the cropped area equals to the specified OR.
- Side cropping (Figure 1-ii(b)) refers to the scenario where the finger is not well aligned with the capturing device. Such misalignment typically results in missing of multiple corner areas. We generate this type of partial fingerprints by cropping the fingerprint images using two perpendicular lines. The intersection and direction of the two lines are randomly specified, with the cropped area equal to s_o .
- Box cropping (Figure 1-ii(c)) represents incomplete image capture due to improper finger pressing or using small-area fingerprint scanners. It involves retaining a rectangular region of the original fingerprint. Specifically, we randomly select an anchor position and keep the pixels closest to the anchor position within the desired observed area $s_f - s_o$, while removing the other pixel values. The distance is measured via “inf” norm, leading to a rectangular partially observed area.

3.3. Robust Partial Fingerprint (RPF)

3.3.1 Baseline Fingerprint Recognition Model

We utilize Convolutional Neural Networks (CNN) to extract feature embedding from input fingerprint images. Matching is then performed based on a matching score *e.g.* cosine similarity computed between the feature embeddings of the template and query fingerprints. We illustrate such baseline framework in Figure 2. Specifically, we construct a baseline model using a ResNet50 [21] to extract feature embedding and a fully connected layer to predict identity. During training, the baseline model is trained in a supervised manner utilizing Arcface [9] as the classification loss \mathcal{L}_{cls} . Specifically, for a training sample i ,

$$\mathcal{L}_{cls,i} = \frac{e^{r \cdot \cos \theta_i + d}}{e^{r \cdot \cos \theta_i + d} + \sum_{j \neq i} e^{r \cdot \cos \theta_j}}, \quad (1)$$

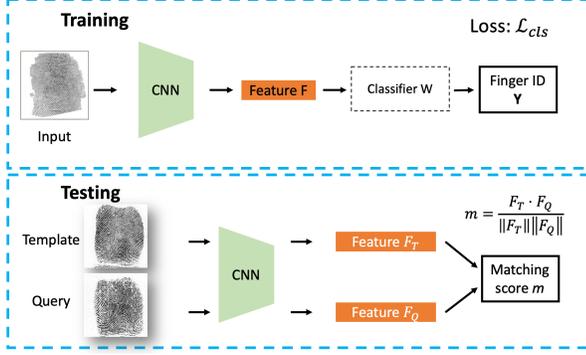


Figure 2. Overview of the baseline fingerprint recognition model.

where r and d are hyperparameters presenting the weight of the feature normalization factor and the desired margin between classification boundary, respectively. The overall training loss is $\mathcal{L} = \sum_{i=1}^N \mathcal{L}_{cls,i}$, where N is the total number of training samples. During testing, only the feature embedding is needed. If the cosine similarity between two embeddings is higher than a pre-defined threshold, we determine a pair of fingerprints is matched.

Despite the usage of the advanced model architecture and classification loss function, the performance of baseline model can suffer from the partial fingerprints given existing fingerprint data are mostly complete. In this paper, we introduce RPF, which combines novel occlusion-enhanced data augmentation and occlusion-aware modeling for improving model robustness to partial fingerprints. These two key components can be easily incorporated into different frameworks with details described in the following.

3.3.2 Occlusion-enhanced Data Augmentation

Data augmentation has been widely used in deep learning model training to improve model robustness. Traditional data augmentation (T-Aug) for fingerprint recognition includes the application of random affine transformations to training data, which, however, can not well compensate for the performance degradation due to various of partial fingerprints. In this work, we introduce a novel data augmentation method, Occlusion-enhanced AugMix (O-AugMix), that effectively utilized the knowledge of partial fingerprints to enhance the robustness against partial fingerprints.

O-AugMix improves on a state-of-the-art data augmentation technique, Augmix [23], by introducing additional occlusion related data augmentation. To begin with, Augmix improves T-Aug by mixing the augmented images with the original image and utilizing an additional Jensen-Shannon Divergence (JSD) loss to encourage feature consistency. Our proposed O-AugMix effectually enhances AugMix by further applying the comprehensive partial fingerprint simulation methods to augment data, allowing handle different types of partial fingerprints that may occur in practice. The pseudocode and illustration of O-AugMix is

Algorithm 1 Occlusion-enhanced AugMix.

```

1: Input: Original image  $x_{orig}$ , Operations
 $\mathcal{O} = \{\text{different partial fingerprint simulation methods, rotation, scaling, translation}\}$ .
2: function AUGMENTANDMIX( $x_{orig}, k, \alpha$ )
3:   Fill  $x_{aug}$  with zeros
4:   Sample adding weights:
5:      $(w_1, \dots, w_k) \sim \text{Dirichlet}(\alpha, \dots, \alpha)$ 
6:   for  $i=1, \dots, k$  do
7:     Apply one of operation  $op_i$  in  $\mathcal{O}$ 
8:      $x_{aug} \leftarrow x_{aug} + w_i op_i(x_{orig})$ 
9:   end for
10:  Sample weight  $m \sim \text{Beta}(\alpha, \alpha)$ 
11:  Obtain the mixed image
12:     $x_{augmix} \leftarrow m x_{orig} + (1 - m) x_{aug}$ 
13:  Return  $x_{augmix}$ 
14: end function
15:  $x_{augmix1} = \text{AugmentAndMix}(x_{orig})$ 
16:  $x_{augmix2} = \text{AugmentAndMix}(x_{orig})$ 

```

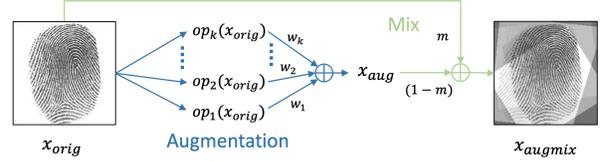


Figure 3. Illustration of O-AugMix.

in Algorithm 1 and Figure 3, respectively. The augmented and original images are then used to compute the JSD loss to encourage feature consistency between the two:

$$\mathcal{L}_{JS} = \frac{1}{3} (\text{JSD}(p_{x_{orig}} \| p_{x_{augmix1}}) + \text{JSD}(p_{x_{orig}} \| p_{x_{augmix2}}) + \text{JSD}(p_{x_{augmix1}} \| p_{x_{augmix2}})), \quad (2)$$

where $p_{\{\cdot\}}$ is the output probability distribution given different input images, $\text{JSD}(\cdot \| \cdot)$ measures the JSD between two probability distributions.

3.3.3 Occlusion-aware Modeling

Occlusion-aware models explicitly characterize the missing parts to extract more robust features. Inspired by a state-of-the-art occlusion-aware face recognition model [50], we effectively model occlusion by introducing an additional segmentation branch as illustrated in Figure 4. The segmentation branch predicts fingerprint area. We then use the segmentation features to generate attention masks and weigh image features at different feature levels. Instead of only predicting specific occlusion patterns, the proposed segmentation branch directly estimates visible fingerprint area to accommodate different types of occlusions, leading to more robust feature extraction. Incorporating the occlusion-aware modeling utilizes an additional segmentation loss:

$$\mathcal{L}_{seg} = \sum_i \sum_k -(y_{ik} \log \hat{y}_{ik} + (1 - y_{ik}) \log (1 - \hat{y}_{ik})), \quad (3)$$

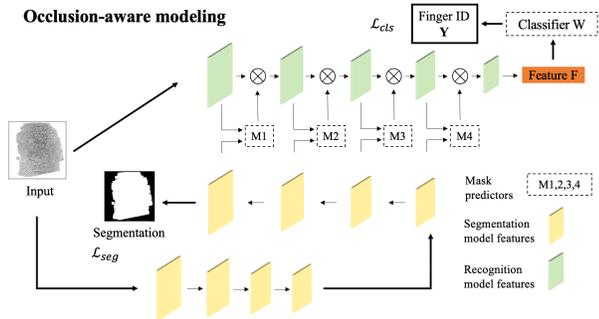


Figure 4. Overview of the proposed occlusion-aware model.

where y_{ik} and \hat{y}_{ik} are the fingerprint area label and prediction for image pixel k of training sample i , respectively.

Assembling the training losses utilized by O-AugMix (Section 3.3.2) and occlusion-aware modeling (Section 3.3.3), we train RPF with the total training loss:

$$\mathcal{L} = \mathcal{L}_{cls} + \lambda_{JS} \mathcal{L}_{JS} + \lambda_{seg} \mathcal{L}_{seg}, \quad (4)$$

where $\lambda_{JS} = 0.1$ and $\lambda_{seg} = 2$ are the corresponding weights, which we determined empirically in experiment.

During training, we augment training data via O-AugMix and train RPF using Equation 4. During testing, we use RPF to extract feature embeddings of a fingerprint pair and perform matching based on the cosine similarity.

4. Experiments

We first introduce the experiment protocol, including the datasets, evaluation metrics and implementation details. We highlight the robustness of RPF through the evaluation on different types of partial fingerprints as discussed in Section 4.1. We demonstrate RPF is efficient, and effective in addressing the partial fingerprint recognition with comparison to the state-of-the-art methods (SOTAs) in Section 4.2, followed by the ablation study in Section 4.3.

Datasets: We use **PrintsGAN** [13] and **NIST** [16] for training and evaluation. In addition, a partial fingerprint subset of NIST, **NIST-Partial**, is curated to further study the model robustness to partial fingerprints. Specifically, PrintsGAN is a large scale synthetic dataset containing 35,000 unique generated fingerprint identities. NIST is a real fingerprint dataset with data quality varying greatly among the sensors. We consider the following 17 sensors: A, B, C, E, F, G, J, K, L, M, N, P, Q, R, S, U, and V, resulting a total of 2,000 unique fingers. NIST-Partial includes the test images from challenging sensors A, B, C, E, F, and G. The fingerprints collected by these sensors are mostly partially observed, with blurred regions and out of FoV acquisitions. To split train/validation/test data, we use 80%/10%/10%. To generate the genuine and imposter pairs, we enumerate all the testing image combinations without repetition (for each

Table 1. Dataset splitting of PrintsGAN, NIST, and NIST-Partial datasets. ‘‘Avg. Impre.’’ denotes the average number of impressions of each fingerprint.

| Datasets | Train/Val/Test Image Number | Avg. Impre. | Number of Pairs | |
|--------------|-----------------------------|-------------|-----------------|----------|
| | | | Genuine | Imposter |
| PrintsGAN | 420K/52.5K/52.3K | 15 | 3672.9K | 6116.3K |
| NIST | 20K/2.5K/2.5K | 12 | 15.1K | 3229.7K |
| NIST-Partial | -/-/1.1K | 12 | 2.3K | 551.6K |

subject, we use one impression). The detailed dataset statistics is shown in Table 1. NIST-Partial is not used for training. We hence only report the number of its testing images.

Evaluation metrics: The evaluation metrics we use is False Rejection Rate (FRR) @ certain False Acceptance Rate (FAR), a typical metrics in biometric verification. Here FRR refers to the rate of mis-classifying genuine pairs as imposter pairs and FAR refers to the rate of mis-classifying imposter pairs as genuine pairs.

Implementation details: The image data is resized to 224×224 resolution while keeping the original aspect ratio. We use Adam optimizer [28] with weight decay of $5e-4$ and a third order learning rate decay. More implementation details and hyperparameters setting are in Supp. A.

4.1. Impact Analysis of Partial Fingerprints

We demonstrate that standard recognition models can have severe performance degradation due to partial observation, while RPF is robust to different types of partial fingerprints. Specifically, we evaluate the baseline model and RPF on different types of partial fingerprints at different ORs. During evaluation, we only consider query fingerprints are partially observed (we demonstrate in Supp. B that performance degradation is more severe when both the template and query fingerprints are partially observed). The evaluation results are illustrated in Figure 5.

For **baseline model** (blue lines in Figure 5), occlusion (Figure 5-i) leads to larger performance degradation compared to being outside FoV (Figure 5-ii). We hypothesize the reason is that the fingerprint data used to train the recognition models typically do not include blockages, but often exhibit mis-alignments. As a result, the trained models are less robust to occlusions compared to being out of FoV. Moreover, among all these partial fingerprints, object blocking (Figure 5-i(c)) causes the most severe performance degradation. This is due to object blocking not only removes fingerprint pixels, but also introduces new texture patterns, creating additional challenges for recognition. Furthermore, diffused blocking (Figure 5-i(a)) results in larger performance degradation even when OR is small as the blurred area in diffused blocking is larger than other types of partial fingerprints.

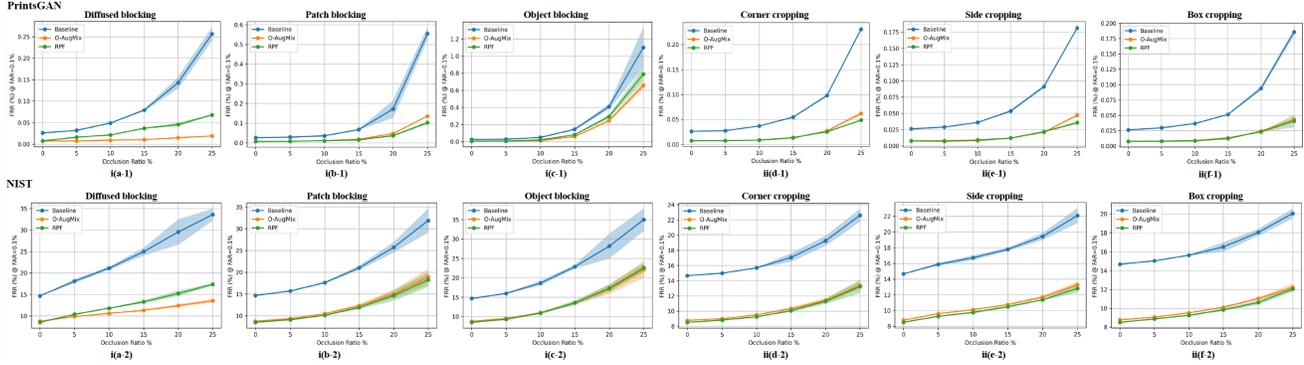


Figure 5. Evaluation of the baseline, O-AugMix, and RPF on different partial fingerprints (different columns) at different ORs on PrintsGAN (top row) and NIST (bottom row). O-AugMix denotes RPF without occlusion-aware modeling. For each type of partial fingerprints, we compute FRR (%) @ FAR=0.1% for 5 times (10 times for NIST) with randomly generated partial query fingerprints and report the mean and scaled standard deviation (error bar along each line). We omit the results after 25% OR as the degradation becomes too large.

Table 2. Comparison with SOTAs on NIST and NIST Partial datasets in model parameters, inference speed, and verification accuracy. For verification accuracy, the FRR numbers are in percentage with value the smaller the better. “ResNet50-Baseline” and “MobileNetV2-Baseline” are the baseline models with different backbones.

| Method | Number of Parameters | Inference Speed (Nvidia V100) | NIST, FRR@ | | NIST-Partial, |
|------------------------------------|----------------------|-------------------------------|------------|-----------|---------------|
| | | | FAR=0.1% | FAR=0.01% | FRR@FAR=0.1% |
| Verifinger ISO matcher [2] | - | 600ms | 7.67 | 9.84 | 12.49 |
| Verifinger proprietary matcher [2] | - | 600ms | 5.28 | 7.99 | 9.47 |
| Deepprints [12] | 76.93M | 40.4ms | 22.02 | 40.96 | 30.81 |
| ResNet50 + ViT Concat [20] | 47.56M | 35.7ms | 8.48 | 15.41 | 13.06 |
| ResNet50-Baseline [21] | 20.56M | 11.1ms | 14.67 | 22.41 | 19.98 |
| MobileNetV2-Baseline [39] | 2.26M | 8.4ms | 16.17 | 25.66 | 21.12 |
| Ours | 36.40M | 24.1ms | 8.52 | 14.37 | 12.23 |

Compared to the baseline method, **RPF** (green lines in Figure 5) not only successfully mitigates the performance degradation due to partial fingerprints but also improves the absolute model performance (0% OR). Particularly, the advantages of RPF become more significant when OR becomes larger. Additionally, the images in PrintsGAN are mostly complete, while NIST is a challenging dataset including partial fingerprints. RPF demonstrates a more significant mitigation of performance degradation on PrintsGAN (Figure 5 top row), while providing more substantial improvements in absolute performance on NIST (Figure 5 bottom row). Furthermore, RPF performs consistently well on different types of partial fingerprints. Comparing to only using O-AugMix, further exploiting occlusion-aware modeling provides better absolute performance and more performance mitigation on partial fingerprints with larger OR. But we are also aware that only using O-AugMix can better handle diffused blocking (Figure 5-i(a)). For diffused blocking, predicting effective feature masks is challenging as the diffused region is not totally erased but blended with black pixels. Given the success of the occlusion-aware modeling in handling other types of partial fingerprints, we believe

characterizing the diffused regions with more advanced segmentation models should provide further improvements.

4.2. Comparison with State-of-the-Arts (SOTA)

We compare RPF with SOTA commercial matchers and deep fingerprint recognition models as the results summarized in Table 2. Comparing to Verifinger (Table 2 row 1-2), a commercial system utilizing more complex feature extraction and matching models trained with additional data, RPF and other deep learning-based methods achieve competitive accuracy but with much better efficiency (average of ~ 20 ms relative to Verifinger’s ~ 600 ms for a single fingerprint pair). Moreover, FRR@FAR=0.1% of Verifinger ISO and Verifinger Proprietary degrades by $(12.49-7.67)/7.67=62.84\%$ and $(9.47-5.28)/5.28=79.36\%$, respectively. In contrast, ours degrades by $(12.23-8.52)/8.52=43.54\%$, highlighting that Verifinger is susceptible to partial fingerprints while our approach has better robustness. Moreover, RPF outperforms existing deep learning-based methods (Table 2 row 3-6) on both NIST (FRR@FAR=0.01%) and NIST-Partial, where the data are mostly partially observed. Existing deep fingerprint

Table 3. Ablation of RPF on PrintsGAN, NIST, and NIST-Partial. OM denotes occlusion-aware modeling. The FRR numbers are in percentage with value the smaller the better.

| Method | Training Loss | | | PrintsGAN, FRR@FAR=0.1% | NIST, FRR@ | | NIST-Partial, FRR@FAR=0.1% |
|-------------------------------|---------------------|--------------------|---------------------|----------------------------|------------|-----------|-------------------------------|
| | \mathcal{L}_{cls} | \mathcal{L}_{JS} | \mathcal{L}_{seq} | | FAR=0.1% | FAR=0.01% | |
| Baseline | ✓ | | | 0.046 | 14.67 | 22.41 | 19.98 |
| Baseline + T-Aug | ✓ | | | 0.017 | 8.97 | 15.24 | 13.32 |
| Baseline + AugMix | ✓ | | | 0.016 | 9.11 | 14.84 | 13.32 |
| Baseline + O-AugMix | ✓ | ✓ | | 0.016 | 8.78 | 14.80 | 12.71 |
| RPF: Baseline + O-AugMix + OM | ✓ | ✓ | ✓ | 0.015 | 8.52 | 14.37 | 12.23 |

recognition models, such as Deeprprints and ResNet50-ViT, also require combining the image features extracted by the deep models with minutiae features, while RPF achieves better performance by only using the occlusion-enhanced data augmentation and occlusion-aware modeling for robust feature extraction.

4.3. Ablation Study

We conduct an ablation study to demonstrate the effectiveness of each component of RPF in Table 3. Starting from the vanilla baseline model, we evaluate the model performance of further using T-Aug, AugMix, O-AugMix, and occlusion-aware modeling (Table 3 row 1 to 5). T-Aug and AugMix are degeneration of O-AugMix that only consider standard alignment-related data augmentation operations, such as rotation, translation, and scaling. Given the frequent presence of mis-aligned fingerprint data, employing T-Aug or AugMix results in improvements over the vanilla baseline model (Table 3 row 2,3 over 1). Nonetheless, O-AugMix (Table 3 row 4), by further effectually encoding the knowledge on different types of partial fingerprints, demonstrates more significant improvements, with substantial boosting on NIST-Partial (Table 3 column 6). It’s worth mentioning that O-AugMix does not include the photometric data augmentation used in AugMix as we found the photometric data augmentations do not provide extra benefits, probably because the fingerprint data is grayscale without large photometric variations.

Moreover, introducing RPF’s occlusion-aware modeling yields additional improvements (Table 3 row 5 over 4). Particularly, more advantages are demonstrated on NIST and NIST-Partial, where partial fingerprints are more common. For example, FRR@FAR=0.1% on NIST-Partial improves from 12.71% to 12.23%. The proposed occlusion-aware model enables more robust feature extraction by incorporating an additional efficient network to accurately capture the visible fingerprint areas (the segmentation accuracy evaluated on NIST’s test set is 95.89% in Intersection-Over-Union) and effectively predicting the image feature weighting masks. RPF achieves the best performance by combining the occlusion-enhanced data augmentation and occlusion-aware modeling, where both of the two methods

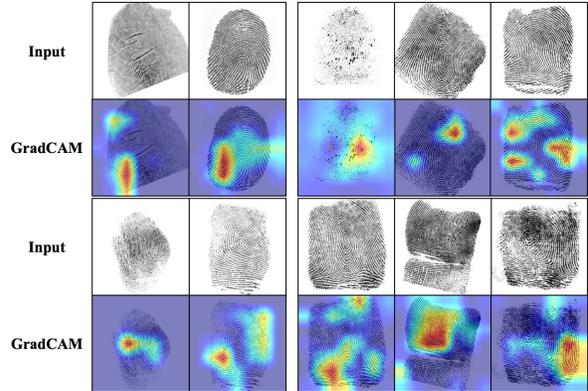


Figure 6. Information distribution captured by GradCAM. Images are from NIST (column 1-2) and PrintsGAN (column 3-5). More reddish color indicates higher importance for recognition.

can be easily incorporated into different models to improve model robustness to partial fingerprints.

5. Empirical Analysis

This section presents empirical analysis from three different perspectives to gain insights into the recognition behavior and identify areas for improvement.

Localization of Informative Regions. When finger impression is partially observed, the model must rely on different regions for recognition. Therefore, when developing recognition models that are robust to partial fingerprints, it is important to understand the information distribution captured by a deep recognition model on a fingerprint. Here we provide visual explanation using GradCAM [40] to illustrate the information distribution. GradCAM has been widely used for visualizing the importance region recognized by a deep learning model in classification tasks. We apply GradCAM to the baseline model with example results shown in Figure 6. We observe that the recognition model extract information from different finger regions without a specific focus on any dominant areas. In Supp. C, we also demonstrate that the performance degradation caused by missing of information at center position or random regions is similar. In other words, all the fingerprint region

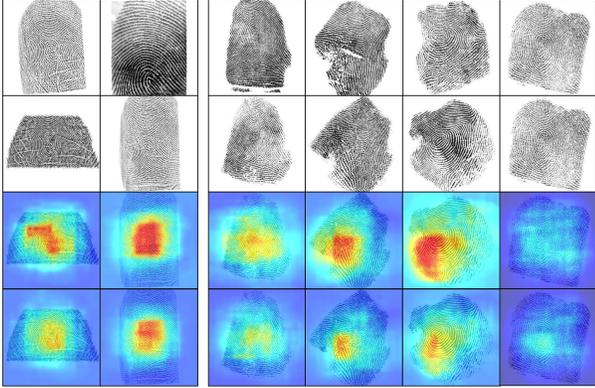


Figure 7. Images from top to bottom are template, query, query overlapped with the matching score map from baseline model and RPF, respectively. Fingerprint pairs are from NIST (column 1-2) and PrintsGAN (column 3-6). More reddish color in the matching score map represents larger matching score decrease.

can include important information for recognition purpose. Therefore, OR is a sensible measure of information loss and all visible finger regions should be exploited in developing a robust recognition model.

Illustration of Robust Matching Behavior. To gain deeper understanding towards the recognition robustness to partial fingerprints, we study the impact of missing certain regions on the final matching score. Specifically, for a genuine fingerprint pair with a template I_t and query I_q , we first compute their matching score m_f using the entire image as input. We then occlude I_q at image position (i, j) with patch blocking, obtaining an occluded query image I_q^{ij} . We compute the matching scoring m^{ij} using I_q^{ij} and I_t , and calculate score decreasing $m_f - m^{ij}$. Finally, we formulate the full matching score map by evenly introducing patch blocking on I_q , like sliding window with constant step size in the vertical and horizontal directions, and computing the corresponding matching score. For visualization, we resize the matching score map to the original fingerprint image size. The values in a matching score map represent the reduction in matching score due to missing different input regions. Intuitively speaking, the larger the value, the more significant loss of information for matching.

We compare the matching score map generated by the baseline model and RPF with example results shown in Figure 7. RPF achieves less matching score decreasing because RPF is more robust to the partial fingerprints. Moreover, RPF shows more evenly distributed matching score maps, illustrating that RPF learns discriminative features from different regions. This analysis further highlights the advantages of RPF in improving model robustness to partial fingerprints. Meanwhile, utilizing the matching score map as guidance to learn features from other regions to the occluded regions can be a promising future direction.

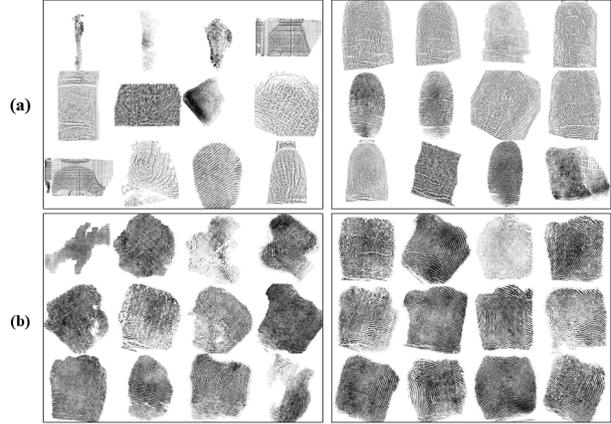


Figure 8. Samples with large (left column) and small (right column) data uncertainty in NIST (a) and PrintsGAN (b).

Estimation of Uncertainty. Partial fingerprints include certain regions being occluded or not captured, causing large uncertainty in recognition. Therefore, estimating the uncertainty is important in characterizing the problematic data and developing more robust recognition models. In this section, we model data uncertainty [26] using DUL [7]. We present the samples with small and large data uncertainty in Figure 8. The samples with large data uncertainty are mostly incomplete, indicating that the captured data uncertainty effectively characterizes the partial fingerprints. We hence believe exploiting uncertainty modeling for building robust partial fingerprint recognition framework is a promising future direction. In Supp. D, we provide details of implementing DUL and show the captured data uncertainty can be used to construct robust matching score. We also discuss the benefits of using model uncertainty to select out the partial fingerprints.

6. Conclusion

We conducted a thorough study of partial fingerprints that occur in practice and proposed RPF, a robust fingerprint recognition framework that effectively utilizes occlusion-enhanced data augmentation and occlusion-aware modeling to improve recognition performance. We demonstrated the effectiveness of RPF, particularly its advantages on handling partial fingerprints through extensive experiments using synthetic and real fingerprints datasets. In addition, we present multiple empirical analyses that provide insights into recognition behaviors and potential areas for improvement. Specifically, our results suggest that extracting robust features from different regions, such as regions with less matching score decreasing, or exploiting uncertainty modeling to characterize problematic data and build a robust matching score, could be promising future directions.

References

- [1] Amsterdam library of textures. [3](#)
- [2] Verifinger sdk. [6](#)
- [3] Gwangbin Bae, Ignas Budvytis, and Roberto Cipolla. Multi-view depth estimation by fusing single-view depth probability with multi-view geometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2842–2851, 2022. [2](#)
- [4] Gary Bradski. The opencv library. *Dr. Dobb's Journal: Software Tools for the Professional Programmer*, 25(11):120–123, 2000. [3](#)
- [5] Kai Cao and Anil K Jain. Fingerprint indexing and matching: An integrated approach. In *2017 IEEE International Joint Conference on Biometrics (IJCB)*, pages 437–445. IEEE, 2017. [2](#)
- [6] Kai Cao and Anil K Jain. Automated latent fingerprint recognition. *IEEE transactions on pattern analysis and machine intelligence*, 41(4):788–800, 2018. [1](#)
- [7] Jie Chang, Zhonghao Lan, Changmao Cheng, and Yichen Wei. Data uncertainty learning in face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5710–5719, 2020. [2](#), [8](#)
- [8] Shengjie Chen, Zhenhua Guo, Xiu Li, and Dongliang Yang. Query2set: Single-to-multiple partial fingerprint recognition based on attention mechanism. *IEEE Transactions on Information Forensics and Security*, 17:1243–1253, 2022. [1](#), [2](#)
- [9] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. [2](#), [3](#)
- [10] Siqi Deng, Yuanjun Xiong, Meng Wang, Wei Xia, and Stefano Soatto. Harnessing unrecognizable faces for improving face recognition. *arXiv preprint arXiv:2106.04112*, 2021. [2](#)
- [11] 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. [1](#)
- [12] 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. [1](#), [2](#), [6](#)
- [13] Joshua J Engelsma, Steven A Grosz, and Anil K Jain. Printsgan: synthetic fingerprint generator. *arXiv preprint arXiv:2201.03674*, 2022. [5](#)
- [14] Jude Ezeobijesi and Bir Bhanu. Patch based latent fingerprint matching using deep learning. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 2017–2021. IEEE, 2018. [1](#), [2](#)
- [15] Yulin Feng and Ajay Kumar. Detecting locally, patching globally: An end-to-end framework for high speed and accurate detection of fingerprint minutiae. *IEEE Transactions on Information Forensics and Security*, 18:1720–1733, 2023. [1](#)
- [16] 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. [5](#)
- [17] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016. [2](#), [12](#)
- [18] Luca Ghiani, Abdenour Hadid, Gian Luca Marcialis, and Fabio Roli. Fingerprint liveness detection using local texture features. *IET Biometrics*, 6(3):224–231, 2017. [1](#)
- [19] Steven A Grosz, Joshua J Engelsma, Eryun Liu, and Anil K Jain. C2cl: Contact to contactless fingerprint matching. *IEEE Transactions on Information Forensics and Security*, 17:196–210, 2021. [1](#)
- [20] Steven A Grosz, Joshua J Engelsma, Rajeev Ranjan, Naveen Ramakrishnan, Manoj Aggarwal, Gerard G Medioni, and Anil K Jain. Minutiae-guided fingerprint embeddings via vision transformers. *arXiv preprint arXiv:2210.13994*, 2022. [6](#)
- [21] 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. [3](#), [6](#)
- [22] Zhiyuan He, Eryun Liu, and Zhiyu Xiang. Partial fingerprint verification via spatial transformer networks. In *2020 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–10. IEEE, 2020. [2](#)
- [23] Dan Hendrycks, Norman Mu, Ekin D Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. *arXiv preprint arXiv:1912.02781*, 2019. [4](#)
- [24] Anil K Jain, Karthik Nandakumar, and Abhishek Nagar. Biometric template security. *EURASIP Journal on advances in signal processing*, 2008:1–17, 2008. [1](#)
- [25] Tsai-Yang Jea and Venu Govindaraju. A minutia-based partial fingerprint recognition system. *Pattern recognition*, 38(10):1672–1684, 2005. [2](#)
- [26] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30, 2017. [2](#), [8](#), [12](#)
- [27] Minchul Kim, Anil K Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18750–18759, 2022. [2](#)
- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [29] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11127–11137, 2021. [2](#)
- [30] Samuel Lee, Seok-Woo Jang, Dongho Kim, Hernsoo Hahn, and Gye-Young Kim. A novel fingerprint recovery scheme using deep neural network-based learning. *Multimedia Tools and Applications*, 80(26):34121–34135, 2021. [1](#)

- [31] Chenhao Lin and Ajay Kumar. Contactless and partial 3d fingerprint recognition using multi-view deep representation. *Pattern Recognition*, 83:314–327, 2018. 1
- [32] Boxiao Liu, Shenghan Zhang, Guanglu Song, Haihang You, and Yu Liu. Rectifying the data bias in knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1477–1486, 2021. 2
- [33] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. Magface: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14225–14234, 2021. 2
- [34] Domingo Mery and Bernardita Morris. On black-box explanation for face verification. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3418–3427, 2022. 2
- [35] Dinh-Luan Nguyen, Kai Cao, and Anil K Jain. Robust minutiae extractor: Integrating deep networks and fingerprint domain knowledge. In *2018 International Conference on Biometrics (ICB)*, pages 9–16. IEEE, 2018. 2
- [36] Timo Ojala, Matti Pietikainen, and Topi Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on pattern analysis and machine intelligence*, 24(7):971–987, 2002. 1
- [37] Ville Ojansivu and Janne Heikkilä. Blur insensitive texture classification using local phase quantization. *Lecture Notes in Computer Science*, 5099:236–243, 2008. 1
- [38] Haibo Qiu, Dihong Gong, Zhifeng Li, Wei Liu, and Dacheng Tao. End2end occluded face recognition by masking corrupted features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [39] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 6
- [40] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. 2, 7
- [41] Yichun Shi and Anil K Jain. Probabilistic face embeddings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6902–6911, 2019. 2
- [42] Yichun Shi, Xiang Yu, Kihyuk Sohn, Manmohan Chandraker, and Anil K Jain. Towards universal representation learning for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6817–6826, 2020. 2
- [43] Lingxue Song, Dihong Gong, Zhifeng Li, Changsong Liu, and Wei Liu. Occlusion robust face recognition based on mask learning with pairwise differential siamese network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 773–782, 2019. 2
- [44] Ai Takahashi, Yoshinori Koda, Koichi Ito, and Takafumi Aoki. Fingerprint feature extraction by combining texture, minutiae, and frequency spectrum using multi-task cnn. In *2020 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–8. IEEE, 2020. 2
- [45] Hanzhuo Tan and Ajay Kumar. Towards more accurate contactless fingerprint minutiae extraction and pose-invariant matching. *IEEE Transactions on Information Forensics and Security*, 15:3924–3937, 2020. 1, 2
- [46] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5265–5274, 2018. 2
- [47] Yandong Wen, Weiyang Liu, Adrian Weller, Bhiksha Raj, and Rita Singh. Sphreface2: Binary classification is all you need for deep face recognition. *arXiv preprint arXiv:2108.01513*, 2021. 2
- [48] Wei Jing Wong and Shang-Hong Lai. Multi-task cnn for restoring corrupted fingerprint images. *Pattern Recognition*, 101:107203, 2020. 1
- [49] Jun Yu, Xinlong Hao, Zeyu Cui, Peng He, and Tongliang Liu. Boosting fairness for masked face recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1531–1540, 2021. 2
- [50] Ge Yuan, Huicheng Zheng, and Jiayu Dong. Msml: Enhancing occlusion-robustness by multi-scale segmentation-based mask learning for face recognition. 2022. 2, 4
- [51] Fanfeng Zeng, Shengda Hu, and Ke Xiao. Research on partial fingerprint recognition algorithm based on deep learning. *Neural Computing and Applications*, 31(9):4789–4798, 2019. 1, 2
- [52] Fandong Zhang and Jufu Feng. High-resolution mobile fingerprint matching via deep joint knn-triplet embedding. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017. 1, 2
- [53] Fang Zhao, Jiashi Feng, Jian Zhao, Wenhan Yang, and Shuicheng Yan. Robust lstm-autoencoders for face de-occlusion in the wild. *IEEE Transactions on Image Processing*, 27(2):778–790, 2017. 2
- [54] Meilu Zhu, Daming Shi, Mingjie Zheng, and Muhammad Sadiq. Robust facial landmark detection via occlusion-adaptive deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3486–3496, 2019. 2