

Fingerprints: Fixed Length Representation via Deep Networks and Domain Knowledge

Joshua J. Engelsma, Kai Cao, Anil K. Jain
Michigan State University
East Lansing, MI, USA
{engelsm7, kaicao, jain}@cse.msu.edu

Abstract

We learn a discriminative fixed length feature representation of fingerprints which stands in contrast to commonly used unordered, variable length sets of minutiae points. To arrive at this fixed length representation, we embed fingerprint domain knowledge into a multitask deep convolutional neural network architecture. Empirical results, on two public-domain fingerprint databases (NIST SD4 and FVC 2004 DB1) show that compared to minutiae representations, extracted by two state-of-the-art commercial matchers (Verifinger v6.3 and Innovatrics v2.0.3), our fixed-length representations provide (i) higher search accuracy: Rank-1 accuracy of 97.9% vs. 97.3% on NIST SD4 against a gallery size of 2000 and (ii) significantly faster, large scale search: 682,594 matches per second vs. 22 matches per second for commercial matchers on an i5 3.3 GHz processor with 8 GB of RAM.

1. Introduction

Over 100 years ago, the pioneering giant of modern day fingerprint recognition, Sir Francis Galton, astutely commented on fingerprints in his 1892 book titled “Finger Prints”:

“They have the unique merit of retaining all their peculiarities unchanged throughout life, and afford in consequence an incomparably surer criterion of identity than any other bodily feature.” [1]

Galton went on to describe fingerprint *minutiae*, the small details woven throughout the papillary ridges on each of our fingers, which Galton believed provided uniqueness and permanence properties for accurately identifying individuals. In the 100 years since Galton’s ground breaking scientific observations, fingerprint recognition systems have become ubiquitous and can be found in a plethora of different domains [2] such as forensics [3], healthcare, mobile device

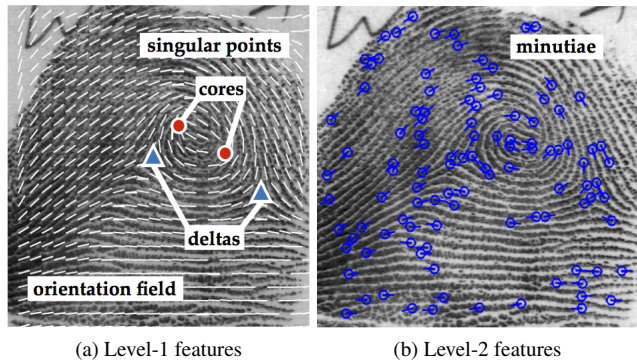


Figure 1: Traditional fingerprint representations consist of global level-1 features (a) and local level-2 features, called minutiae points, together with their descriptors (b).

security [4], mobile payments [4], border crossing [5], and national ID [6]. To date, virtually all of these systems continue to rely upon the location and orientation of minutiae within fingerprint images for recognition (Fig. 1).

Although fingerprint recognition systems based on minutiae (i.e. handcrafted features) have enjoyed significant success over the years, to our knowledge, not much effort has been devoted to augment handcrafted features with recent advances in deep learning to improve the recognition accuracy and reduce the complexity of large scale search. The argument for introducing deep learning models to fingerprint recognition is compelling given the following major limitations of prevailing minutiae-based fingerprint recognition systems.

1. Minutiae-based representations are of variable length (Table 1), since the number of extracted minutiae varies amongst different fingerprint images even of the same finger (Fig. 2 (a)). This causes two main problems: (i) pairwise fingerprint comparison is computationally demanding, accentuated when searching databases with an inordinate number of identities, e.g.,

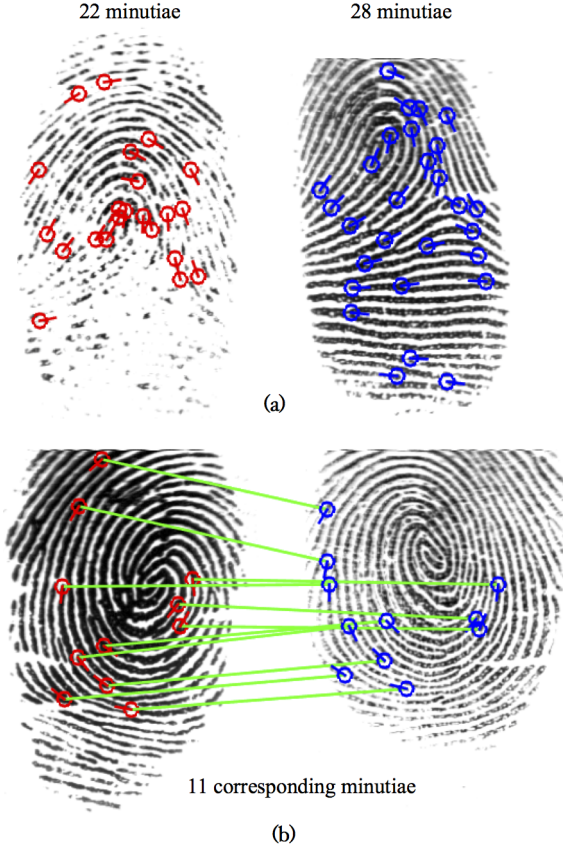


Figure 2: Failures of a minutiae-based matcher (COTS A). The genuine pair (two impressions from the same finger) in (a) was falsely rejected at 0.1% FAR due to inaccurate minutiae extraction. The imposter pair (impressions from two different fingers) in (b) was falsely accepted at 0.1% FAR due to the similar minutiae distribution in these two fingerprints.

India’s Aadhaar system with 1.25 billion identities [6] and (ii) template encryption, a necessity for user privacy protection, is a significant challenge [7].

2. Fingerprint recognition can be viewed as a **75 billion class problem** (≈ 7.5 billion living people, each with 10 fingers) with large intra-class variability and large inter-class similarity (Fig. 2). This necessitates devising extremely discriminative yet compact representations that go beyond just minutiae points.
3. Reliable minutiae extraction in low quality fingerprints (due to noise, distortion, finger condition) is problematic, causing false rejects in the recognition system (Fig. 2 (a)). See also NIST fingerprint evaluation FpVTE 2012 [8].

Given the limitations of minutiae-based matchers, we propose a custom deep network to extract discriminative




Figure 3: Proposed fixed-length, 512-dimensional fingerprint representations extracted from the same four fingerprints shown in Figure 2. Unlike COTS A, we correctly classify the pair in (a) as a genuine pair, and the pair in (b) as an imposter pair. The score threshold of our matcher @ FAR = 0.1% is 0.69

Table 1: Template Comparisons

Matcher	(Min, Max) # of Minutiae	(Min, Max) Template Size (kB)
COTS A	(8, 206)	(1.1, 22.6)
COTS B	(6, 221)	(0.1, 1.3)
Proposed	N.A. ²	2 [†]

¹ Statistics from NIST SD4 and FVC 2004 DB1.

² Template is not explicitly comprised of minutiae.

[†] Template size is fixed at 2 kilobytes, irrespective of the number of minutiae.

fixed-length representations (Fig. 3) for state-of-the-art fingerprint recognition and search. Unlike several prior attempts to arrive at a fixed length fingerprint representation [9, 10], our method incorporates fingerprint domain knowledge into the deep network architecture via a *minutiae-map*. More specifically, our architecture learns both a *texture based representation* and a *minutiae-map based representation*. The concatenation of these two representations comprises our fixed-length fingerprint representation. In addition to injecting domain knowledge into the network, we add a spatial transformer module [11] to crop and align the query (probe) fingerprint prior to extracting

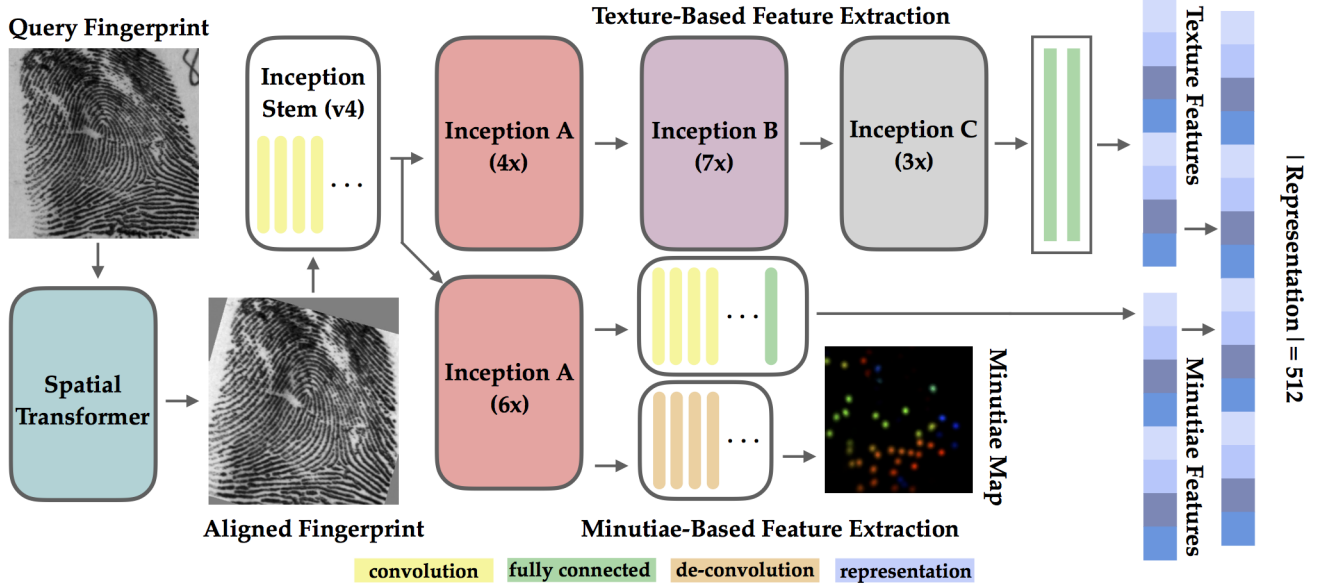


Figure 4: Flow diagram of the proposed deep learning based fingerprint recognition system: (i) a query fingerprint is cropped and aligned via the spatial transformer; (ii) the aligned fingerprint is passed to the Inception v4 stem which extracts both a textural representation and a minutiae-map based representation; (iii) The texture representation and minutiae representation are concatenated into a 512-dimensional compact representation of the input fingerprint.

its representation. Finally, we reduce the complexity of our network architecture and inference time via a variant of the teacher-student training paradigm.

The primary contributions of this work are as follows:

1. A custom deep network architecture that utilizes fingerprint domain knowledge (minutiae locations and orientations) to extract a discriminative fixed-length representation.
2. Empirical results demonstrating a significant improvement in fingerprint recognition speed (a **31,500 fold** improvement in the number of matches per second) over two fingerprint SDKs on two benchmark datasets (NIST SD4 and FVC 2004 DB1) while retaining comparable recognition accuracy.
3. A method for significantly reducing the memory consumption and inference time of the network using the teacher-student training paradigm.

2. Prior Work

While some attention has been given towards using deep learning models in fingerprint recognition, most of these works are focused on improving only a *sub-module* of an end-to-end fingerprint recognition system such as segmentation [12, 13, 14, 15], orientation field estimation [16, 17, 18], minutiae extraction [19, 20, 21], and minutiae descriptor extraction [22]. While they are able to improve the performance of the sub-modules, they still operate within the

conventional paradigm of extracting a variable length feature representation comprised of the minutiae and their associated descriptors in a given fingerprint image.

We completely reformulate the conventional fingerprint recognition paradigm. Rather than working towards an end goal of minutiae extraction for matching, we allow a deep network to *learn* a fixed-length feature vector that best discriminates fingerprints. While there have been two previous attempts along this direction [9, 10], both require computationally intensive global alignment of fingerprints prior to extracting fixed length representations via a black box deep network. The authors in [23] also proposed a fixed length fingerprint representation based upon handcrafted textural features. Both [9, 10] and [23] have inferior accuracy compared to minutiae-based systems. In contrast, our domain knowledge infused multitask network with built in alignment (via the spatial transformer module [11]) is competitive with state-of-the-art minutiae-based matchers in accuracy while being significantly faster. The network is also more interpretable than [9, 10] given the minutiae map output. The speed and accuracy of our algorithm will open up new opportunities for large scale search and matching.

3. Approach

We provide an overview and intuition of the proposed end-to-end, domain knowledge infused, deep learning based fingerprint recognition system. We then describe how incorporating a spatial transformer module into our

network enables fingerprint cropping and alignment as part of the same network inference used for representation extraction. Finally, we discuss the domain knowledge (minutiae map) that is injected into the network via a multitask learning objective.

Algorithm 1 Extract Fingerprint Representation

```

1:  $l(I)$ : shallow localization network
2:  $g(I, x, y, \theta)$ : bilinear grid sampler
3:  $s(I)$ : inception stem
4:  $m(F_{map})$ : minutiae branch
5:  $t(F_{map})$ : texture branch
6:
7: Input: Unaligned  $448 \times 448$  fingerprint image  $I_f$ 
8:  $(x, y, \theta) \leftarrow l(I_f)$ 
9:  $I_t \leftarrow g(I_f, x, y, \theta)$ 
10:  $F_{map} \leftarrow s(I_t)$ 
11:  $R_{minutiae} \leftarrow m(F_{map})$ 
12:  $R_{texture} \leftarrow t(F_{map})$ 
13:  $R_{fused} \leftarrow R_{minutiae} \oplus R_{texture}$ 
14: Output: fingerprint representation  $R_{fused} \in \mathbb{R}^{512}$ 

```

3.1. Overview

A high level overview of our proposed network architecture is provided in Figure 4 with pseudocode describing the process shown in Algorithm 1. The model is trained with a longitudinal database comprised of 440K rolled fingerprint images stemming from 37,410 unique fingers [24]. The primary task during training is to predict the finger class label $c \in [0, 37410]$ of each of the 440K training fingerprint images (≈ 12 fingerprint impressions / finger). Similar to prevailing face recognition systems, the last fully connected layer can be used as the representation for fingerprint matching.

The input to our model is a 448×448 grayscale fingerprint image which is first passed through the spatial transformer module. The spatial transformer acts as the network’s built-in fingerprint cropping and alignment module. After applying the spatial transformation to I_f , a cropped and aligned fingerprint I_t is passed to the base network.

The backbone of our base network is the Inception v4 architecture proposed in [25]. We specifically modified the network architecture to incorporate two different branches (Fig. 4) using the three Inception modules (A, B, and C) described in [25]. The first branch performs the primary learning task of predicting a finger class label c directly from the cropped, aligned fingerprint I_t and essentially learns the texture cues in the fingerprint image. The second branch again predicts the finger class label c from the aligned fingerprint I_t , but it also has a related side task of predicting the minutiae locations and orientations in I_t . In this

manner, we guide this branch of the network to extract representations influenced by fingerprint minutiae. The textural cues act as complementary discriminative information to the minutiae-guided representation. The final feature representation is the 512-dimensional concatenation of the two representations. Note that the minutiae set is not explicitly used in the final representation, we only use the minutiae-map to guide our network training.

In the following subsections, we provide details of the major subcomponents of the proposed network architecture.




Figure 5: An unaligned fingerprint (a) is cropped and aligned (b) by the spatial transformer module prior to representation extraction.

3.2. Cropping and Alignment

In nearly all minutiae-based fingerprint recognition systems, the first step is to crop a region of interest (ROI) and then perform global alignment based on some reference points (such as the core point). However, this alignment is computationally expensive. This motivated us to consider attention mechanisms such as the spatial transformers in [11].

The advantages of using the spatial transformer module in place of reference point based alignment algorithms are two-fold: (i) it requires only one forward pass through a shallow localization network (Table 2), followed by bilinear grid sampling. This reduces the computational complexity of alignment; (ii) The parameters of the localization network are tuned to minimize the classification loss of the base-network (representation extraction network). In other words, we let the base-network decide what a “good” transformation is, so that it can better classify the input fingerprints.

Given an uncropped, unaligned fingerprint image I_f , a shallow localization network first hypothesizes the parameters of an affine transformation matrix A_θ . Given A_θ , a grid sampler $T_\theta(G_i)$ samples the input image I_f pixels (x_i^f, y_i^f) for every target grid location (x_i^t, y_i^t) to output the cropped, aligned fingerprint image I_t in accordance with Equation 1.




Figure 6: Minutiae Map Extraction. The minutiae locations and orientations of an input fingerprint (a) are encoded as a 6-channel minutiae map. The “hot spots” in each channel indicate the spatial location of the minutiae points. The color of the hot spots indicate their orientations.

Table 2: Localization Network Architecture

Type	Output Size	Filter Size, Stride
Convolution	$128 \times 128 \times 24$	$5 \times 5, 1$
Max Pooling	$64 \times 64 \times 24$	$2 \times 2, 2$
Convolution	$64 \times 64 \times 32$	$3 \times 3, 1$
Max Pooling	$32 \times 32 \times 32$	$2 \times 2, 2$
Convolution	$32 \times 32 \times 48$	$3 \times 3, 1$
Max Pooling	$16 \times 16 \times 48$	$2 \times 2, 2$
Convolution	$16 \times 16 \times 64$	$3 \times 3, 1$
Max Pooling	$8 \times 8 \times 64$	$2 \times 2, 2$
Fully Connected	64	
Fully Connected	3	

$$\begin{pmatrix} x_i^f \\ y_i^f \end{pmatrix} = T_\theta(G_i) = A_\theta \begin{pmatrix} x_i^t \\ y_i^t \\ 1 \end{pmatrix} \quad (1)$$

Once I_t has been computed from the grid sampler $T_\theta(G_i)$, it is passed on to the base network for classification. Finally, the parameters for the localization network are updated based upon the classification loss of the base network.

The architecture used for our localization network is shown in Table 2 and images from before and after the spatial transformer module are shown in Figure 5. In order to get the spatial transformer to properly converge, (i) the learning rate for the localization network parameters was scaled by 0.035, (ii) the upper bound of the estimated affine matrix translation and rotation parameters was set to 224 pixels and ± 60 degrees, respectively, and (iii) the affine matrix scaling parameter was fixed to select a window of size 285×285 from the input fingerprint image. These constraints are based on our domain knowledge on the maximum extent a user would rotate / translate their fingers dur-

ing placement on the reader platen and the typical “high quality” area of a fingerprint image (usually the center portion).

3.3. Minutiae Map Domain Knowledge

To prevent overfitting the network to the training data, we incorporate fingerprint domain knowledge into the network to guide it to learn a more discriminative representation. The specific domain knowledge we incorporate into our network architecture is hereafter referred to as the *minutiae map* [26]. Note that the minutiae map is not explicitly used in the fingerprint representation, but the information contained in the map is indirectly embedded in the network during training.

A minutiae map is essentially a c -channel heatmap quantizing the locations (x, y) and orientations $\theta \in [0, 2\pi]$ of the minutiae within a fingerprint image. More formally, let h and w be the height and width of an input fingerprint image I_f and $T = \{m_1, m_2, \dots, m_n\}$ be its minutiae template with n minutiae points, where $m_t = (x_t, y_t, \theta_t)$ and $t = 1, \dots, n$. Then, the minutiae map $H \in \mathbb{R}^{h \times w \times c}$ at (i, j, k) can be computed by summing the location and orientation contributions of each of the minutiae in T . More formally,

$$H(i, j, k) = \sum_{t=1}^n C_s((x_t, y_t), (i, j)) \cdot C_o(\theta_t, 2k\pi/c) \quad (2)$$

where $C_s(\cdot)$ and $C_o(\cdot)$ calculate the spatial and orientation contribution of minutiae m_t to the minutiae map at (i, j, k) based upon the euclidean distance of (x_t, y_t) to (i, j) and the orientation difference between θ_t and $2k\pi/c$ as follows:

$$C_s((x_t, y_t), (i, j)) = \exp\left(-\frac{\|(x_t, y_t) - (i, j)\|_2^2}{2\sigma_s^2}\right) \quad (3)$$

$$C_o(\theta_t, 2k\pi/c) = \exp\left(-\frac{d\phi(\theta_t, 2k\pi/c)}{2\sigma_s^2}\right) \quad (4)$$

where $d\phi(\theta_1, \theta_2)$ is the orientation difference between angles θ_1 and θ_2 :

$$d\phi(\theta_1, \theta_2) = \begin{cases} |\theta_1 - \theta_2| & -\pi \leq -\theta_1 - \theta_2 \leq \pi \\ 2\pi - |\theta_1 - \theta_2| & \text{otherwise.} \end{cases} \quad (5)$$

In our approach, we extract minutiae maps of size $128 \times 128 \times 6$ to encode the minutiae locations and orientations of an input fingerprint image of size $448 \times 448 \times 1$. An example fingerprint image and its corresponding minutiae map are shown in Figure 6. In the following subsection, we will explain in detail how we inject the information contained in the minutiae maps into our custom architecture.

3.4. Multi-Task Architecture

The minutiae-map domain knowledge is injected into our Inception v4 network backbone via multitask learning. Multitask learning improves generalizability of a model since domain knowledge within the training signals of related tasks acts as an inductive bias [27, 28, 29, 30, 31, 32]. The multi-task branch of our network architecture is shown in Figures 4 and 7. The primary task of the branch is to classify a given fingerprint image into its “finger class” and the secondary task is to estimate the minutiae-map. In this manner, we guide the minutiae-branch of our network to extract fingerprint representations that are influenced by minutiae locations and orientations. A separate branch in our network aims to extract a complementary texture-based representation by directly predicting the class label of an input fingerprint without any domain knowledge (Fig. 4).

Note, we combine the texture branch with the minutiae branch in our architecture (rather than two separate networks) for the following several reasons: (i) the minutiae branch and the texture branch share a number of parameters (the Inception v4 stem), reducing the model complexity that two separate models would necessitate; (ii) the spatial transformer is optimized based on both branches (i.e. learned alignment benefits both the texture-based and minutiae-based representations) avoiding two separate spatial transformer modules.

We constructed a custom architecture $f_w(x)$ with parameters w to accommodate the multitask learning objective. Let $s_{w_3}(x)$ be the shared stem of our architecture and $m_{w_2}(x)$ and $t_{w_1}(x)$ be the minutiae and texture branches, respectively where $w = \{w_1, w_2, w_3\}$. The stem $s_{w_3}(x)$ provides a nonlinear mapping from the input fingerprint I_t to a shared feature map F_{map} . The first branch $t_{w_1}(x)$ takes F_{map} as input and outputs a feature representation $\mathbf{x}_1 \in \mathbb{R}^{256}$.

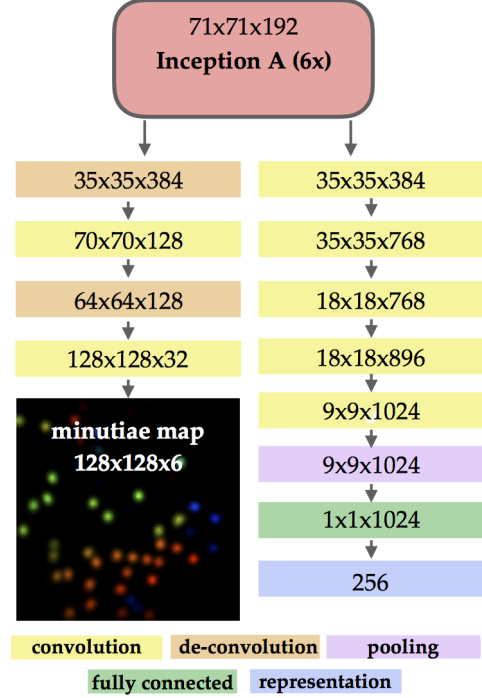


Figure 7: The custom multi-task **minutiae branch** m_{w_2} of our proposed network. The dimensions inside each box represent the input dimensions.

\mathbb{R}^{256} . The second branch $m_{w_2}(x)$ again takes F_{map} as input and outputs both a feature representation $\mathbf{x}_2 \in \mathbb{R}^{256}$ and also a minutiae map $\mathbf{H} \in \mathbb{R}^{128 \times 128 \times 6}$.

Given, \mathbf{x}_1 and \mathbf{x}_2 , fully connected layers are applied for identity classification logits, outputting $\mathbf{y}_1 \in \mathbb{R}^c$ and $\mathbf{y}_2 \in \mathbb{R}^c$ as follows:

$$\mathbf{y}_i = \mathbf{W}^T \mathbf{x}_i + \mathbf{b}, \quad i \in \{1, 2\} \quad (6)$$

where $\mathbf{W} \in \mathbb{R}^{d \times c}$, $\mathbf{b} \in \mathbb{R}^c$, d is the number of features, and c is the number of identities in the training set. Next, \mathbf{y}_1 and \mathbf{y}_2 are both passed to a softmax layer to compute the probabilities $\hat{\mathbf{y}}_1$ and $\hat{\mathbf{y}}_2$ of \mathbf{x}_1 and \mathbf{x}_2 belonging to each identity. Finally, $\hat{\mathbf{y}}_1$ and $\hat{\mathbf{y}}_2$ and their corresponding ground truth labels y_1^g and y_2^g can be used to compute the loss of the network from the cross-entropy loss function:

$$L_1(I_t, y_i^g) = -\log(p(\hat{\mathbf{y}}_i = y_i^g | I_t, w_3, w_i)) \quad i \in \{1, 2\} \quad (7)$$

For computing the loss of the minutiae map estimation side task, we employ the Mean Squared Error Loss between the estimated minutiae map \mathbf{H} and the ground truth minutiae map H_g as follows:

$$L_2(I_t, H_g) = \sum_{i,j} (\mathbf{H}_{i,j} - H_{g_{i,j}})^2 \quad (8)$$

Finally, using the addition of all these loss terms, and a dataset comprised of N training images, our model parameters w are trained in accordance with:

$$\operatorname{argmin}_w \sum_{i=1}^N L_1(I_t^i, y_1^i) + \sum_{i=1}^N L_1(I_t^i, y_2^i) + \sum_{i=1}^N L_2(I_t^i, H_g^i) \quad (9)$$

Note, during the training, we augment our dataset with random rotations, translations, brightness, and cropping. We use the RMSProp optimizer with a batch size of 30. Regularization included dropout with a keep probability of 0.8 and weight decay of 0.00004.

After the multitask architecture has converged, a fixed length feature representation can be acquired by extracting the fully connected layer before the softmax layers in both of the network’s branches. Let $\mathbf{x}_1 \in \mathbb{R}^{256}$ be the texture representation and $\mathbf{x}_2 \in \mathbb{R}^{256}$ be the minutiae representation. Then, the final feature representation is obtained by concatenation \mathbf{x}_1 and \mathbf{x}_2 into $\mathbf{x}_3 \in \mathbb{R}^{512}$ followed by normalization of \mathbf{x}_3 to unit length.

3.5. Matching

Given two 512-dimensional, unit length representations \mathbf{t}_1 and \mathbf{t}_2 , a match score s is computed as the cosine similarity between the two representations. In particular:

$$s = \mathbf{t}_1^\top \cdot \mathbf{t}_2 \quad (10)$$

3.6. Model Size Reduction

Because our multitask architecture is very large (305 MB), we use a variant of the teacher student training schema [33] to significantly reduce the model size to 91.3 MB (with minimal loss in accuracy) such that it can be run on small embedded devices. For a student architecture, we selected the Inception v3 architecture [34] due to its similarity to the Inception v4 architecture used as our backbone. We also experimented with smaller models for the student such as the MobileNet v2 architecture [35], but it was unable to adequately replicate the features output by the teacher model.

The goal of the student model is to replicate the behavior of the teacher. In particular, given N training images, we extract their 512-dimensional unit length feature representations $R = \{x_1, \dots, x_n\}$, $n = 1 \dots N$ from the teacher model. Then, the student model accepts the same N training images, and learns to estimate the representations R extracted by the teacher via the L2-loss function. Let $student_w(I_n) = x_n^s$ be the non-linear mapping of the student model from the input images I_n to the fingerprint

representations x_n^s where x_n^t is the ground truth representation extracted by the teacher. Then, the student is trained by optimizing the L2-loss function:

$$\operatorname{argmin}_w \sum_{i=1}^N L(x_n^s, x_n^t) \quad (11)$$

where


$$L(x_n^s, x_n^t) = \frac{1}{2} \|x_n^t - x_n^s\|_2^2 \quad (12)$$

4. Experimental Results

The baseline experimental results are comprised of two state-of-the-art commercial minutiae-based matchers¹ and the Inception v4 model trained in a manner similar to [9] (i.e. no alignment or domain knowledge is added to the network). Our proposed method includes three major steps: (i) adding a spatial transformer (Inception v4 + STN) for alignment, (ii) adding domain knowledge via multitask learning (Inception v4 + STN + MTL), and (iii) model size reduction via a smaller student model (Student). Recognition accuracy for all of these methods are reported in the following subsections.

4.1. Testing Datasets

We use the NIST SD4 [37] and FVC 2004 DB1 [38] databases for benchmarking our fixed-length matcher. The NIST SD4 dataset is comprised of 4000 *rolled* fingerprints (2000 unique fingers with 2 impressions per finger) (Fig. 8a) and the FVC 2004 DB1 is comprised of 800 *plain* fingerprints (100 unique fingers with 8 impressions per finger) (Fig. 8b). Both of these datasets are considered challenging datasets, even for commercial matchers.



(a) NIST SD4 rolled fingerprint



(b) FVC 2004 plain fingerprint

Figure 8: Examples from benchmark datasets. (a) A smudgy rolled fingerprint from NIST SD4 and (b) a distorted plain (slap) fingerprint from FVC 2004 DB1.

¹COTS used were Verifinger SDK v6.0 and Innovatrics [36], a top performer in NIST FpVTE 2012.

Table 3: Verification Results (1-to-1 comparison)

Database // Method	Inception v4 (unmodified)	COTS A	COTS B	Inception v4 (+STN)	Inception v4 (+STN + MTL)	Student
FVC 2004 ¹	58.4%	92.3%	95.6%	65.1%	93.2%	86.4%
NIST SD4 ²	92.0%	96.1%	97.0%	94.0%	96.9%	95.8%

¹ TAR @ FAR = 0.1% is reported due to the insufficient number of imposters for FAR = 0.01%

² TAR @ FAR = 0.01% is reported.

Table 4: NIST SD4 Rank-1 Search Results (background 2000 fingerprints)

Inception v4 (unmodified)	COTS A	COTS B	Inception v4 (+STN)	Inception v4 (+STN + MTL)	Student
94.8%	96.4%	97.3%	97.0%	97.9%	96.8%

4.2. Recognition Accuracy

From the results (Tables 3 and 4), we can observe that the unmodified off-the-shelf Inception v4 model performs worse than COTS matchers. Adding the spatial transformer module to the Inception v4 backbone (Inception v4 + STN) boosts the accuracy in all testing scenarios. With added domain knowledge via our proposed multitask learning architecture, accuracy is further boosted and becomes competitive with the two COTS matchers.

We posit the reason for the low performance of the Inception v4 model without fingerprint domain knowledge on the FVC 2004 database is because the fingerprint image characteristics in FVC 2004 are much different than the characteristics of our training data. In particular, our longitudinal training database [24] is comprised of large, inked, and scanned rolled fingerprints similar to NIST SD4 (Fig. 8a). Meanwhile, FVC consists of small area fingerprints acquired with a fingerprint reader (Fig. 8b). Therefore, the Inception v4 model without domain knowledge “overfits” to rolled ink fingerprints. As can be seen in Table 3, our approach for adding domain knowledge via the minutiae map in combination with our multitask architecture significantly improves the generalizability of the deep fingerprint representations on testing datasets different from the rolled fingerprints in the training dataset (increasing the TAR from 65.1% to 93.2%).

We also note that our use of the teacher student training paradigm enabled us to significantly reduce our model size with only a small loss in accuracy.

4.3. Speed

Extracting fixed length representations enables us to perform matching **several orders of magnitude** faster than minutiae-based matchers (Table 5). This has tremendous benefits for large scale searches against millions or even billions of fingerprints in the background. Admittedly our representation extraction time is slightly longer than minutiae extraction. However, representation extraction is only performed once whereas matching is performed many times.

Table 5: Speed Comparison

Matcher	Representation Extraction (ms)	Matches / Second
COTS A	142	22
COTS B	113	20
Inception v4 (+ STN + MTL)	313	682,594
Student	157	682,594

¹ Computed on an i5 3.3 GHz processor with 8 GB of RAM.

4.4. Large Scale Search

Having benchmarked our matcher against two minutiae-based COTS matchers, we further demonstrate its efficacy for large scale fingerprint search by computing the identification accuracy on NIST SD4 with a supplemented background (gallery) of 282,000 unique rolled fingerprints from a longitudinal database [24]. With a background of 282,000 fingerprint images, we achieve a Rank-1 Identification accuracy of 94.3%, only a slight drop from the search performance with a 2K background. A single search of a probe against the background takes only 416 milliseconds with our fixed-length matcher. Due to the limited speed of the COTS SDKs, we are unable to perform large scale retrieval experiments with them.

5. Summary

We have introduced a custom deep network which uses built-in alignment and embedded domain knowledge to extract discriminative fixed-length fingerprint representations. The strong advantages of our fixed-length representations over prevailing variable length minutiae templates are:

1. Orders of magnitude faster fingerprint matching (682,594 matches / second vs. 22 matches / second), invaluable for large scale search.
2. The ability to provide strong encryption (a significant challenge with variable length templates).
3. Additional discriminative textural cues beyond just minutiae, enabling competitive recognition accuracy with minutiae-based matchers.

References

- [1] F. Galton, *Finger prints*. Macmillan and Company, 1892. 1
- [2] D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar, *Handbook of Fingerprint Recognition*. Springer, 2nd ed., 2009. 1
- [3] “Next Generation Identification (NGI).” <https://www.fbi.gov/services/cjis/fingerprints-and-other-biometrics/ngi>. 1
- [4] “iOS Security.” https://www.apple.com/business/site/docs/iOS_Security_Guide.pdf. 1
- [5] “Office of Biometric Identity Management Identification Services.” <https://www.dhs.gov/obim-biometric-identification-services>. 1
- [6] “Unique Identification Authority of India, dashboard summary.” <https://portal.uidai.gov.in/uidwebportal/dashboard.do>. 1, 2
- [7] A. K. Jain, K. Nandakumar, and A. Nagar, “Biometric template security,” *EURASIP Journal on advances in signal processing*, vol. 2008, p. 113, 2008. 2
- [8] C. I. Watson, G. Fiumara, E. Tabassi, S. Cheng, P. Flanagan, and W. Salamon, “Fingerprint vendor technology evaluation, nist interagency/internal report 8034: 2015,” available at <https://dx.doi.org/10.6028/NIST.IR.8034>. 2
- [9] K. Cao and A. K. Jain, “Fingerprint indexing and matching: An integrated approach,” in *Biometrics (IJCB), 2017 IEEE International Joint Conference on*, pp. 437–445, IEEE, 2017. 2, 3, 7
- [10] D. Song and J. Feng, “Fingerprint indexing based on pyramid deep convolutional feature,” in *Biometrics (IJCB), 2017 IEEE International Joint Conference on*, pp. 200–207, IEEE, 2017. 2, 3
- [11] M. Jaderberg, K. Simonyan, A. Zisserman, et al., “Spatial transformer networks,” in *Advances in neural information processing systems*, pp. 2017–2025, 2015. 2, 3, 4
- [12] Y. Zhu, X. Yin, X. Jia, and J. Hu, “Latent fingerprint segmentation based on convolutional neural networks,” in *Information Forensics and Security (WIFS), 2017 IEEE Workshop on*, pp. 1–6, IEEE, 2017. 3
- [13] X. Dai, J. Liang, Q. Zhao, and F. Liu, “Fingerprint segmentation via convolutional neural networks,” in *Chinese Conference on Biometric Recognition*, pp. 324–333, Springer, 2017. 3
- [14] J. Ezeobiesi and B. Bhanu, “Latent fingerprint image segmentation using deep neural network,” in *Deep Learning for Biometrics*, pp. 83–107, Springer, 2017. 3
- [15] D.-L. Nguyen, K. Cao, and A. K. Jain, “Automatic cropping fingermarks: Latent fingerprint segmentation,” *arXiv preprint arXiv:1804.09650*, 2018. 3
- [16] K. Cao and A. K. Jain, “Latent orientation field estimation via convolutional neural network,” in *Biometrics (ICB), 2015 International Conference on*, pp. 349–356, IEEE, 2015. 3
- [17] P. Schuch, S.-D. Schulz, and C. Busch, “Deep expectation for estimation of fingerprint orientation fields,” in *Biometrics (IJCB), 2017 IEEE International Joint Conference on*, pp. 185–190, IEEE, 2017. 3
- [18] Z. Qu, J. Liu, Y. Liu, Q. Guan, R. Li, and Y. Zhang, “A novel system for fingerprint orientation estimation,” in *Chinese Conference on Image and Graphics Technologies*, pp. 281–291, Springer, 2018. 3
- [19] D.-L. Nguyen, K. Cao, and A. K. Jain, “Robust minutiae extractor: Integrating deep networks and fingerprint domain knowledge,” in *2018 International Conference on Biometrics (ICB)*, pp. 9–16, IEEE, 2018. 3
- [20] Y. Tang, F. Gao, J. Feng, and Y. Liu, “Fingernet: An unified deep network for fingerprint minutiae extraction,” in *Biometrics (IJCB), 2017 IEEE International Joint Conference on*, pp. 108–116, IEEE, 2017. 3
- [21] L. N. Darlow and B. Rosman, “Fingerprint minutiae extraction using deep learning,” in *Biometrics (IJCB), 2017 IEEE International Joint Conference on*, pp. 22–30, IEEE, 2017. 3
- [22] K. Cao and A. K. Jain, “Automated latent fingerprint recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 3
- [23] A. K. Jain, S. Prabhakar, L. Hong, and S. Pankanti, “Fingercode: a filterbank for fingerprint representation and matching,” in *Computer Vision and Pattern Recognition, 1999. IEEE Computer Society Conference on.*, vol. 2, pp. 187–193, IEEE, 1999. 3
- [24] S. Yoon and A. K. Jain, “Longitudinal study of fingerprint recognition,” *Proceedings of the National Academy of Sciences*, vol. 112, no. 28, pp. 8555–8560, 2015. 4, 8
- [25] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, “Inception-v4, inception-resnet and the impact of residual connections on learning,” in *AAAI*, vol. 4, p. 12, 2017. 4
- [26] K. Cao, L. Nguyen, C. Tymoszek, and A. K. Jain, “End to end latent fingerprint system,” tech. rep., MSU Technical report, MSU-PRIP-2018-11-16, 2018. 5
- [27] R. Caruana, “Multitask learning,” *Machine Learning*, vol. 28, pp. 41–75, Jul 1997. 6
- [28] X. Yin and X. Liu, “Multi-task convolutional neural network for pose-invariant face recognition,” *IEEE Transactions on Image Processing*, vol. 27, no. 2, pp. 964–975, 2018. 6
- [29] C. Zhang and Z. Zhang, “Improving multiview face detection with multi-task deep convolutional neural networks,” in *Applications of Computer Vision (WACV), 2014 IEEE Winter Conference on*, pp. 1036–1041, IEEE, 2014. 6
- [30] Z. Zhang, P. Luo, C. C. Loy, and X. Tang, “Facial landmark detection by deep multi-task learning,” in *European Conference on Computer Vision*, pp. 94–108, Springer, 2014. 6
- [31] Y. Tian, P. Luo, X. Wang, and X. Tang, “Pedestrian detection aided by deep learning semantic tasks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5079–5087, 2015. 6

- [32] A. H. Abdulnabi, G. Wang, J. Lu, and K. Jia, "Multi-task cnn model for attribute prediction," *IEEE Transactions on Multimedia*, vol. 17, no. 11, pp. 1949–1959, 2015. 6
- [33] J. Ba and R. Caruana, "Do deep nets really need to be deep?," in *Advances in neural information processing systems*, pp. 2654–2662, 2014. 7
- [34] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016. 7
- [35] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017. 7
- [36] "Fingerprint matching software." <https://www.innovatrics.com/idkit-fingerprint-sdk/>. 7
- [37] "NIST special database 4." <https://www.nist.gov/srd/nist-special-database-4>. 7
- [38] D. Maio, D. Maltoni, R. Cappelli, J. L. Wayman, and A. K. Jain, "Fvc2004: Third fingerprint verification competition," in *Biometric Authentication*, pp. 1–7, Springer, 2004. 7