

Fingerprint Spoof Buster: Use of Minutiae-centered Patches

Tarang Chugh*, *Student Member, IEEE*, Kai Cao, and Anil K. Jain, *Life Fellow, IEEE*

Abstract—The primary purpose of a fingerprint recognition system is to ensure a reliable and accurate user authentication, but the security of the recognition system itself can be jeopardized by spoof attacks. This study addresses the problem of developing accurate, generalizable, and efficient algorithms for detecting fingerprint spoof attacks. Specifically, we propose a deep convolutional neural network based approach utilizing local patches centered and aligned using fingerprint minutiae. Experimental results on three public-domain LivDet datasets (2011, 2013, and 2015) show that the proposed approach provides state-of-the-art accuracies in fingerprint spoof detection for intra-sensor, cross-material, cross-sensor, as well as cross-dataset testing scenarios. For example, in LivDet 2015, the proposed approach achieves 99.03% average accuracy over all sensors compared to 95.51% achieved by the LivDet 2015 competition winners. Additionally, two new fingerprint presentation attack datasets containing more than 20,000 images, using two different fingerprint readers, and over 12 different spoof fabrication materials are collected. We also present a graphical user interface, called Fingerprint Spoof Buster, that allows the operator to visually examine the local regions of the fingerprint highlighted as live or spoof, instead of relying on only a single score as output by the traditional approaches.

Index Terms—Fingerprint spoof detection, presentation attack detection, convolutional neural networks, minutiae-based local patches, liveness detection

I. INTRODUCTION

WITH the ubiquitous deployment of fingerprint recognition systems in many day-to-day applications, such as financial transactions, international border security, unlocking a smartphone, etc., the vulnerability of the system security to *presentation attacks* is of growing concern [2], [3]. The ISO standard *IEC 30107-1:2016(E)* [4] defines presentation attacks as the “*presentation to the biometric data capture subsystem with the goal of interfering with the operation of the biometric system*”. These attacks can be realized through a number of methods including, but not limited to, use of (i) *gummy fingers* [5], *i.e.* fabricated finger-like objects with accurate imitation of another individual’s fingerprint ridge-valley structures, (ii) *2D or 3D printed fingerprint targets* [6], [7], [8], (iii) *altered fingerprints* [9], *i.e.* intentionally tampered or damaged real fingerprint patterns to avoid identification, and (iv) *cadaver fingers* [10]. Among these, fingerprint spoof attacks (*i.e.* gummy fingers and printed targets) are the most common form of presentation attacks, with a multitude of

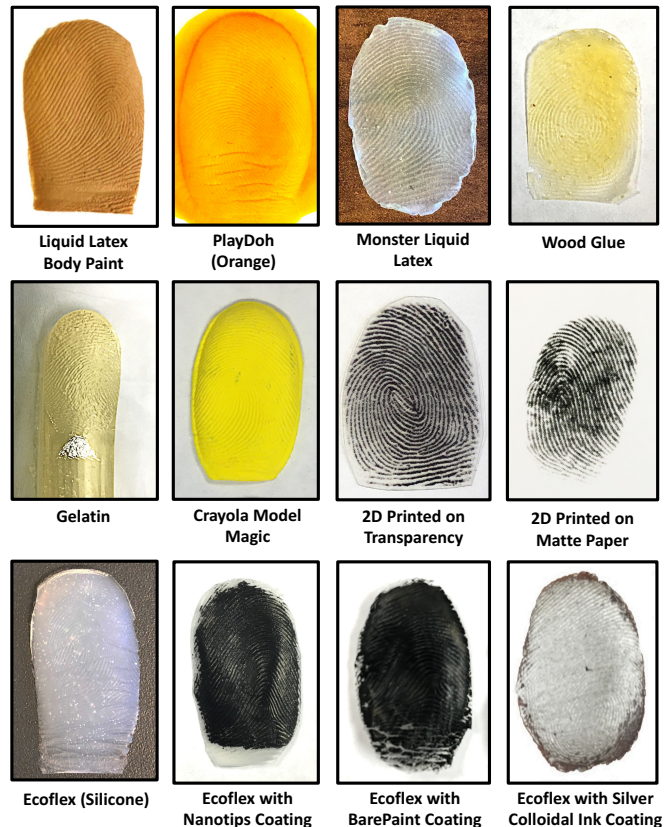


Fig. 1. Fingerprint spoof attacks can be realized using various readily available fabrication materials, such as PlayDoh, WoodGlue, Gelatin, etc.

fabrication processes ranging from basic *molding and casting* to utilizing sophisticated 2D and 3D printing techniques [5], [6], [7], [8].

Commonly available materials, such as gelatin, silicone, play-doh, etc., have been utilized to generate fingerprint spoofs (see Figs. 1 and 2), capable of circumventing a fingerprint recognition system security with a reported success rate of more than 70% [11]. For instance, in July 2016, researchers at Michigan State University unlocked a fingerprint secure-smartphone using a 2D printed fingerprint spoof to help police with a homicide case¹, using a technique proposed in [6]. In Sept. 2013, shortly after Apple released iPhone 5s with inbuilt TouchID fingerprint technology, Germany’s Chaos Computer Club² hacked its capacitive sensor by utilizing a high resolu-

T. Chugh, K. Cao, and A. K. Jain are with the Department of Computer Science and Engineering, Michigan State University, East Lansing, MI, 48824. E-mail: {chughtar, kaicao, jain}@cse.msu.edu

*Corresponding Author

A preliminary version of this paper was presented at the International Joint Conference on Biometrics (IJCB), Denver, CO, October 1-4, 2017 [1].

¹<http://statenews.com/article/2016/08/how-msu-researchers-unlocked-a-fingerprint-secure-smartphone-to-help-police-with-homicide-case>

²<http://www.ccc.de/en/updates/2013/ccc-breaks-apple-touchid>




Fig. 2. Visual comparison between (a) a live Fingerprint, and (b) the corresponding spoofs (of the same live finger) made with different materials. Images are taken from LivDet-2011 dataset (Biometrika sensor) [12]. Our method can successfully distinguish between live and spoof fingerprints. The spoofness scores for live fingerprint is 0.00, and for spoof fingerprints are 0.95, 0.97, 0.99, 0.99, and 0.95 for Ecoflex, Gelatin, Latex, Silgum, and Wood Glue, respectively.

TABLE I
PERFORMANCE COMPARISON (AVERAGE CLASSIFICATION ERROR [%]) OF SOFTWARE-BASED SPOOF DETECTION STUDIES [27], [29] ON LIVDET 2011, 2013, AND 2015 DATASETS.

Study	Approach	LivDet 2011	LivDet 2013*	LivDet 2015
Ghiani et al. [22]	Local Phase Quantization (LPQ)	11.1	3.0	N/A
Gragniello et al. [24]	Weber Local Descriptor (WLD)	7.9	N/A	N/A
Ghiani et al. [23]	Binarized Statistical Image Features (BSIF)	7.2	2.1	N/A
Gragniello et al. [25]	Local Contrast-Phase Descriptor (LCPD)	5.7	1.3	N/A
Nogueira et al. [27]	Transfer Learning + CNN-VGG + Whole Image	4.5	1.1	4.5
Pala et al. [28]	Custom CNN with triplet loss + Randomly selected local patches	3.33	0.58	N/A
Proposed Approach	CNN-MobileNet-v1 + Minutiae-based local patches	1.67	0.25	0.97

*LivDet 2013 includes results for Biometrika and Italdata sensors.

tion photograph of the enrolled user's fingerprint to fabricate a spoof fingerprint with wood glue. In another incident, in March 2013, a Brazilian doctor was arrested for using spoof fingers made of silicone to fool the biometric attendance system at a hospital in Sao Paulo³. Since then, many successful spoof attacks have highlighted the vulnerabilities of fingerprint biometric systems^{4,5}.

Fingerprint spoof detection methods are urgently needed to thwart such attacks on fingerprint authentication systems, thereby increasing the security and user confidence in such systems. The various anti-spoofing approaches proposed in the literature can be broadly classified into hardware-based and software-based solutions [2], [10], [13]. The hardware-based solutions typically require the fingerprint reader to be augmented with sensor(s) to detect the characteristics of vitality, such as blood flow [14], skin distortion [15], odor [16] and so on. There are also special types of fingerprint sensors, such as Lumidigm's multispectral scanner [17] and Compact Imaging's multiple reference optical coherence tomography (OCT) [18], that capture sub-dermal ridge patterns in the finger. An open-source fingerprint reader with a two-camera design provides two complementary streams of information useful for spoof detection [19]. Software-based solutions, on

the other hand, extract features from the presented fingerprint image (or a sequence of frames) acquired by the fingerprint sensors, without incurring any additional hardware cost, to differentiate between live and spoof fingers.

The software-based solutions published in the literature typically utilize one of the following approaches: (i) anatomical features (e.g. pore locations and their distribution [20]), (ii) physiological features (e.g. perspiration [21]), or (iii) texture-based features (e.g. Local Phase Quantization (LPQ) [22], Binarized Statistical Image Features (BSIF) [23], and Weber Local Descriptor [24]). Gragniello et al. [25] proposed a 2D local contrast-phase descriptor (LCPD), utilizing both spatial and frequency domain information. In contrast to the custom-tailored anti-spoof features, Menotti et al. [26], Nogueira et al. [27], and Pala et al. [28] have proposed convolutional neural network (CNN) based solutions whose performances were shown to surpass many published spoof detection algorithms (see Table I).

One of the limitations of many of the published anti-spoof methods is their poor generalization performance across spoof materials. Studies in [27], [30], [31] have shown that when a spoof detector is evaluated on spoofs fabricated using materials that were not seen during training, there can be up to a three-fold increase in the spoof detection error rates. To generalize an algorithm's effectiveness across spoof fabrication materials, called *cross-material* performance, some studies

³<http://www.bbc.com/news/world-latin-america-21756709>

⁴<http://fortune.com/2016/04/07/guy-unlocked-iphone-play-doh/>

⁵<https://srlabs.de/bites/spoofing-fingerprints/>

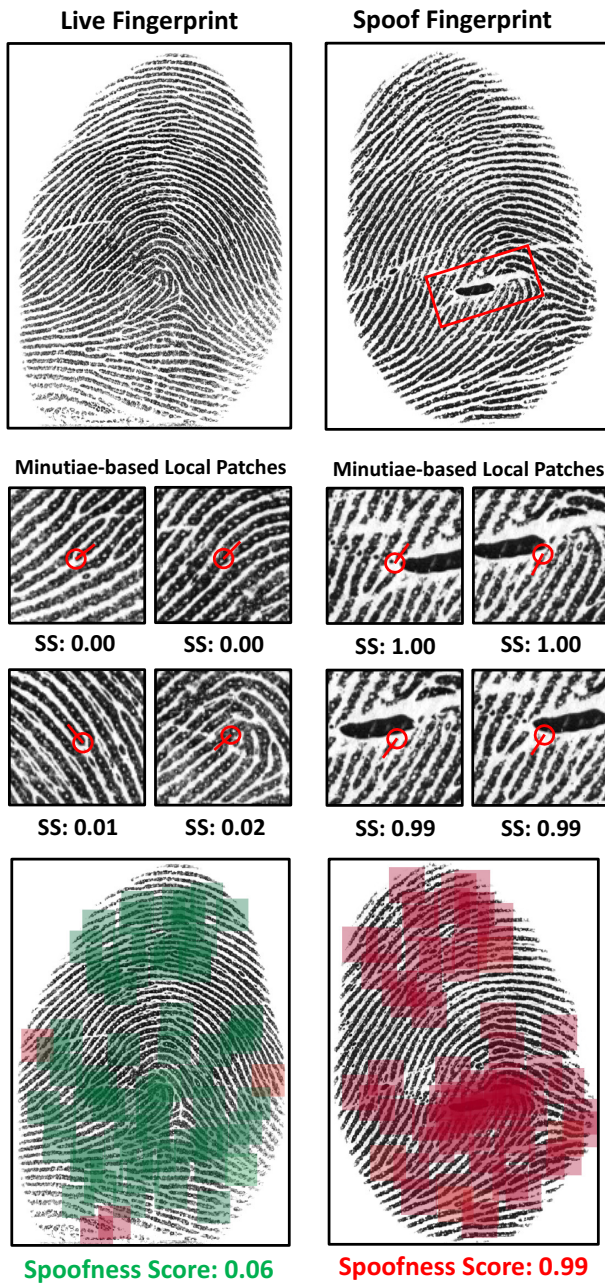


Fig. 3. Example of a live fingerprint and corresponding spoof fingerprint, with the artifacts introduced in the spoofs highlighted in red. The minutiae based local patches extracted around the artifacts are also presented. The images are taken from MSU Fingerprint Presentation Attack Dataset (MSU-FPAD) - CrossMatch Sensor and the spoof material used is Silicone (Ecoflex). The spoofness scores output by the proposed approach for the live and spoof fingerprints are 0.06 and 0.99, respectively.

have approached spoof detection as an *open-set problem*⁶. Rattani et al. [31] applied the Weibull-calibrated SVM (W-SVM), a variant of SVM based on properties of statistical extreme value theory, to detect spoofs made of new materials. Ding and Ross [32] trained an ensemble of multiple one-class SVMs using textural features extracted from only live fingerprint images.

⁶Open-set problems address the possibility of spoof classes during testing, that were not seen during training. Closed-set problems, on the other hand, evaluate only those spoof classes that the system was trained on.

A series of fingerprint Liveness Detection (LivDet) competitions have been held since 2009 to advance state-of-the-art and benchmark the proposed anti-spoofing solutions [33]. The best performer in the most recent⁷ LivDet 2015 [13], Nogueira et al. [27], utilized transfer learning, where deep CNNs originally designed for object recognition and pre-trained on ImageNet database [34], were fine-tuned on fingerprint images to differentiate between live and spoof fingerprints. In their approach, the networks were trained on whole fingerprint images resized to 227×227 pixels for VGG [35] and 224×224 pixels for AlexNet [36] as required by these networks. However, there are three disadvantages of using this approach: (i) fingerprint images from some of the sensors used in LivDet datasets, such as Crossmatch L Scan Guardian (640×480), have a large blank area ($\geq 50\%$) surrounding the friction ridge region. Directly resizing these images, from 640×480 to 227×227 , eventually results in the friction ridge area occupying less than 10% of the original image size; (ii) resizing a rectangular image of size, say $w \times h$, to a square image, say $p \times p$, leads to different amounts of information retained in the two spatial image dimensions; (iii) downsizing an image, in general, leads to a significant loss of discriminatory information.

It is important to consider various sources of noise involved in the spoof fabrication process itself that can introduce some artifacts, such as missing friction ridge regions, cracks, air bubbles, etc., in the spoofs. The primary consequence of such artifacts is the creation of spurious minutiae in the fingerprint images sensed from spoofs. The local regions around these spurious minutiae can, therefore, provide salient cues to differentiate a spoof fingerprint from live fingerprints (see Fig. 3). We utilize this observation to train a two-class CNN using local patches around the extracted minutiae, as opposed to the whole fingerprint images or randomly selected local patches, to design a fingerprint spoof detector. We show that this approach is more robust to novel fabrication materials than earlier approaches that utilize the whole image [27] or randomly selected local patches [28].

The proposed approach, utilizing local patches of size $p \times p$, ($p = 96$), centered at minutiae, (i) addresses the previously mentioned drawbacks of downsizing whole fingerprint images to train the CNN, (ii) provides large amount of data (an average of 48 patches/fingerprint image), sufficient to train the deep CNN architectures from scratch without overfitting, (iii) learns salient textural information from local regions, robust to differentiate between spoof and live fingerprints, and (iv) provides a fine-grained representation of the fingerprint images capable of localizing partial spoof fingerprints and fingerprint alterations. The output of the CNN is a confidence score in the range $[0 - 1]$, defined as *Spoofness Score*; higher the spoofness score, more likely the patch is extracted from a spoof fingerprint. For a given image, the spoofness scores corresponding to the minutiae-based local patches are averaged to give the global spoofness score for the whole image. Furthermore, a fusion of CNN models trained on multi-scale patches (ranging in size from from 64×64 to

⁷Results of LivDet 2017 are expected to be presented at the 11th IAPR International Conference on Biometrics (ICB), Brisbane, Feb. 20-23, 2018.




Fig. 4. An overview of the proposed approach for fingerprint spoof detection using CNNs trained on local patches centered and aligned using minutiae location and orientation, respectively.

128×128), centered and aligned using minutiae, is shown to further boost the spoof detection performance.

The main contributions of this study are enumerated below. They also highlight the differences between this paper and our preliminary study [1]:

- Utilized fingerprint domain-knowledge to design a robust fingerprint spoof detector, where local patches centered and aligned using fingerprint minutiae are utilized for training a CNN model. This differs from other published approaches which have generally used the whole fingerprint image for spoof detection.
- A fine-grained representation of fingerprint images using local patches capable of detecting partial spoof fingerprint regions and fingerprint alterations.
- Experimental results on publicly available datasets (LivDet 2011, 2013, and 2015), including intra-sensor, cross-material, cross-sensor, and cross-dataset scenarios, show that the proposed approach outperforms the state-of-the-art results published on these three datasets. For example, in LivDet 2015, our algorithm achieves 99.03% average accuracy over all sensors compared to 95.51% achieved by the LivDet 2015 winner [13].
- Real-time spoof detection by using MobileNet-v1 [37], compared to Inception-v3 [38] used in the preliminary study [1]. The average classification time for a single input image on a NVIDIA GTX 1080 Ti GPU is 100ms compared to 800ms for Inception-v3 network.
- Collected two new fingerprint presentation attack datasets containing more than 20,000 fingerprint images, using two different fingerprint readers and over 12 different spoof fabrication materials.
- Developed a graphical user interface (GUI) for real-time fingerprint spoof detection, called Fingerprint Spoof Buster, which allows the fingerprint reader operator to visually examine the local regions of the fingerprint highlighted as live or spoof, instead of relying on a single score, as output by the traditional approaches.

II. PROPOSED SPOOF DETECTION APPROACH

The proposed approach includes two stages, an offline training stage and an online testing stage. The offline training stage involves (i) detecting minutiae in the sensed fingerprint image (live or spoof), (ii) extracting local patches centered and aligned using minutiae location and orientation, respectively,

and (iii) training MobileNet models on the aligned local patches. During the testing stage, the spoof detection decision is made based on the average of spoofness scores for individual patches output from the MobileNet model. An overview of the proposed approach is presented in Fig. 4.

A. Datasets

The following datasets have been utilized in this study:

1) *LivDet Datasets*: In order to evaluate performance of the proposed approach, we utilized LivDet 2011 [12], LivDet 2013 [39], and LivDet 2015 [13] datasets. Each of these datasets contains over 16,000 fingerprint images, acquired from four different fingerprint readers, with equal numbers of live and spoof fingerprints that are equally split between training and testing sets. However, the CrossMatch and Swipe readers from LivDet 2013 dataset were not utilized for evaluation purposes because the (a) LivDet competition organizers found anomalies in the fingerprint data from CrossMatch reader and discouraged its use for comparative evaluations [33], and (b) the resolution of fingerprint images output from Swipe reader is very low, *i.e.* 96 dpi. Unlike other LivDet datasets, spoof fingerprint images from Biometrika and Italdata readers in LivDet 2013 dataset [39] are fabricated using the *non-cooperative method i.e.* without user cooperation. It should be noted that in LivDet 2015, the testing set included spoofs fabricated using new materials, that were not known in the training set. These new materials included liquid ecoflex and RTV for Biometrika, Digital Persona, and Green Bit readers, and OOMOO and gelatin for Crossmatch reader. Table II presents a summary of the LivDet datasets used in this study.

2) *MSU Fingerprint Presentation Attack Dataset*: In addition to utilizing LivDet Datasets, we collected a large dataset, called the MSU Fingerprint Presentation Attack Dataset (MSU-FPAD), using two different fingerprint readers, namely, CrossMatch Guardian 200 and Lumidigm Venus 302. There are a total of 9,000 live images and 10,500 spoof images captured using these two readers and 4 different spoof fabrication materials, namely, ecoflex, playhoh, 2D printed on matte paper, and 2D printed on transparency film. The selection of the fingerprint readers and the spoof materials is based on the requirements of IARPA ODIN program⁸ evaluation. This dataset will be further augmented with additional subjects, spoof fabrication materials, and fingerprint readers.

⁸<https://www.iarpa.gov/index.php/research-programs/odin/odin-baa>

TABLE II
SUMMARY OF THE LIVENESS DETECTION (LIVDET) DATASETS UTILIZED IN THIS STUDY.

Dataset	LivDet 2011 [12]				LivDet 2013 [39]		LivDet 2015 [13]			
Fingerprint Reader	Biometrika	ItalData	Digital Persona	Sagem	Biometrika	ItalData	GreenBit	Biometrika	Digital Persona	CrossMatch
Model	FX2000	ET10	4000B	MSO300	FX2000	ET10	DactyScan26	HiScan-PRO	U.are.U 5160	L Scan Guardian
Image Size	315 × 372	640 × 480	355 × 391	352 × 384	315 × 372	640 × 480	500 × 500	1000 × 1000	252 × 324	640 × 480
Resolution (dpi)	500	500	500	500	569	500	500	1000	500	500
#Live Images Train / Test	1000/1000	1000/1000	1000/1000	1000/1000	1000/1000	1000/1000	1000/1000	1000/1000	1000/1000	1510/1500
#Spoof Images Train / Test	1000/1000	1000/1000	1000/1000	1000/1000	1000/1000	1000/1000	1000/1500	1000/1500	1000/1500	1473/1448
Cooperative Subject*	Yes	Yes	Yes	Yes	No	No	Yes	Yes	Yes	Yes
Spoof Materials	Ecoflex, Gelatine, Latex, Silgum, Wood Glue		Gelatine, Latex, Play Doh, Silicone, Wood Glue		Ecoflex, Gelatine, Latex, Modasil, Wood Glue		Ecoflex, Gelatine, Latex, Wood Glue, Liquid Ecoflex, RTV			Body Double, Ecoflex, Play Doh, OOMOO, Gelatin

*In the cooperative subject scenario, a subject willingly provides a negative impression of the fingerprint as a mold, while in the non-cooperative subject scenario, the fingerprint mold is created by using the latent fingerprint lifted off a surface touched by the subject. LivDet datasets are available to interested researchers at <http://livdet.org/registration.php>.

TABLE III
SUMMARY OF THE MSU FINGERPRINT PRESENTATION ATTACK DATASET (MSU-FPAD) AND PRECISE BIOMETRICS SPOOF-KIT DATASET (PBSKD).

Dataset	MSU-FPAD		Precise Biometrics Spoof-Kit	
Fingerprint Reader	CrossMatch	Lumidigm	CrossMatch	Lumidigm
Model	Guardian 200	Venus 302	Guardian 200	Venus 302
Image Size	750 × 800	400 × 272	750 × 800	400 × 272
Resolution (dpi)	500	500	500	500
#Live Images Train / Test	2, 250 / 2, 250	2, 250 / 2, 250	250 / 250 [†]	250 / 250 [†]
#Spoof Images Train / Test	3, 000 / 3, 000	2, 250 / 2, 250	250 / 250	200 / 200 [‡]
Cooperative*	Yes	Yes	Yes	Yes
Spoof Materials	Ecoflex, PlayDoh, 2D Print (Matte Paper), 2D Print (Transparency)		Ecoflex, Gelatin, Latex body paint, Ecoflex with silver colloidal ink coating, Ecoflex with BarePaint coating, Ecoflex with Nanotips coating, Crayola Model Magic, Wood glue, Monster Liquid Latex, and 2D printed fingerprint on office paper	

[†]1000 randomly sampled live fingerprint images from MSU-FPAD are selected for Precise Biometrics Spoof-Kit Dataset.

[‡] Lumidigm fingerprint reader does not image Silicone (EcoFlex) spoofs with NanoTips and BarePaint coatings.

Fig. 5 presents some example fingerprint images, and Table III presents a summary of the MSU Fingerprint Presentation Attack Dataset.

3) *Precise Biometrics Spoof-Kit Dataset*: We also collected another dataset containing high quality 900 spoof fingerprint images fabricated using 10 different types of spoof materials, namely, (i) Ecoflex, (ii) Gelatin, (iii) Latex body paint, (iv) Ecoflex with silver colloidal ink coating, (v) Ecoflex with BarePaint coating, (vi) Ecoflex with Nanotips coating, (vii) Crayola Model Magic, (viii) Wood glue, (ix) Monster Liquid Latex, and (x) 2D printed fingerprint on office paper. The spoof specimens used for this dataset are taken from Precise Biometrics⁹ Spoof-Kit containing 10 specimens per spoof type, for a total of 100 spoof specimens. Each spoof specimen is imaged 5 times using two fingerprint readers, namely, CrossMatch Guardian 200 and Lumidigm Venus 302. Note that Lumidigm reader does not image Silicone (EcoFlex) spoofs with NanoTips and BarePaint coatings. An additional 900 randomly sampled live fingerprints from MSU-FPAD are selected for a total of 1,800 fingerprint images in Precise

Biometrics Spoof-Kit Dataset. Fig. 6 presents some example fingerprint images, and Table III presents a summary of the Precise Biometrics Spoof-Kit Dataset.

B. Minutiae Detection

The fingerprint minutiae are extracted using the algorithm from [40]. The three LivDet datasets (LivDet 2011, 2013, and 2015) comprise of fingerprint images captured at varying resolutions, ranging from 500 dpi to 1000 dpi. Since the minutiae detector in [40] was designed for 500 dpi images, all fingerprint images are resized to ensure a standard resolution of 500 dpi. A standard resolution for all the fingerprint images is also crucial to ensure similar amount of friction ridge area in each local patch, irrespective of the fingerprint reader used. An average of 46 minutiae (s.d. = 6.2) and 50 minutiae (s.d. = 6.9) are detected per live image and spoof image, respectively, for the LivDet datasets.

C. Local Patch Extraction

For a given fingerprint image I with k detected minutiae points $M = \{m_1, m_2, \dots, m_k\}$, where $m_i = \{x_i, y_i, \theta_i\}$, i.e.

⁹<https://precisebiometrics.com/>




Fig. 5. Example images from MSU Fingerprint Presentation Attack Dataset (MSU-FPAD) acquired using (a) CrossMatch Guardian 200, and (b) Lumidigm Venus 302 fingerprint readers. Note that Lumidigm reader does not image PlayDoh (orange) spoofs.

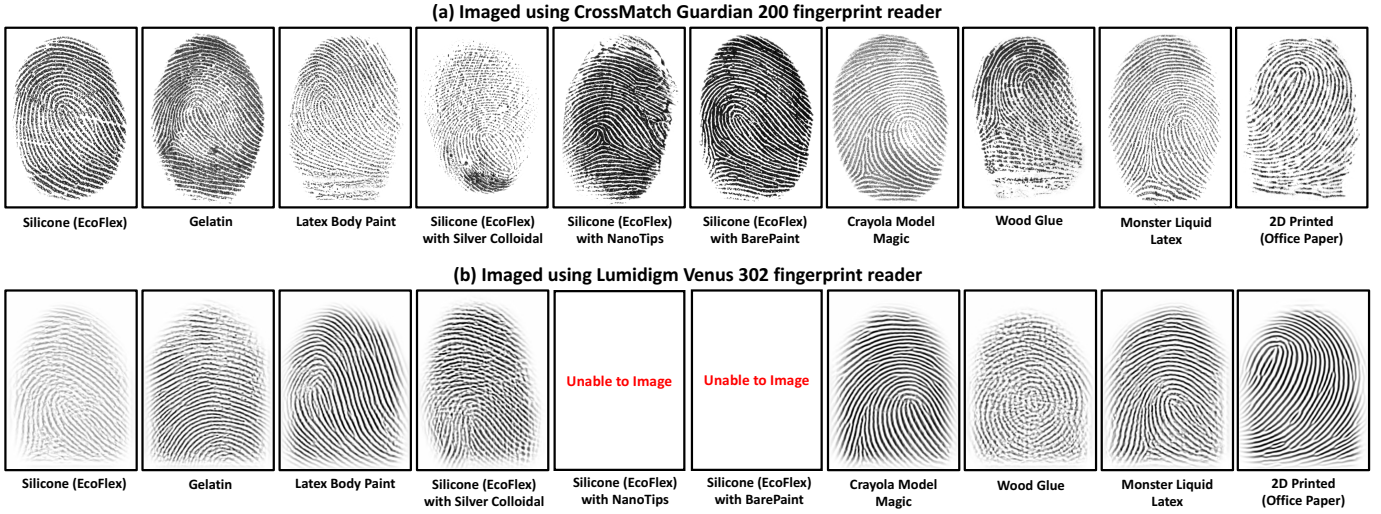


Fig. 6. Example images from Precise Biometrics Spoof-Kit Dataset (PBSKD) acquired using (a) CrossMatch Guardian 200, and (b) Lumidigm Venus 302 fingerprint readers. Note that Lumidigm reader does not image Silicone (EcoFlex) spoofs with NanoTips and BarePaint coatings.

the minutiae m_i is defined in terms of spatial coordinates (x_i, y_i) and orientation (θ_i) , a corresponding set of k local patches $L = \{l_1, l_2, \dots, l_k\}$, each of size $[q \times q]$ where $(q = \sqrt{2}p)$, are extracted. Each local patch l_i , centered at the corresponding minutia location (x_i, y_i) , is aligned¹⁰ based on the minutiae orientation (θ_i) . After alignment, the central region of size $[p \times p]$ ($p = 96$) is cropped from the rotated patch and used for training the CNN model (see Fig. 4). The size of larger patch is fixed to $[\sqrt{2}p \times \sqrt{2}p]$ to prevent any loss of information during patch alignment. Fig. 8 presents examples of real and spoof fingerprint images and the corresponding local patches centered around minutiae points before alignment.

For evaluating the impact of local patch size on the spoof detection performance, we also explore use of multi-resolution patches of size $p \in \{64, 96, 128\}$ for training independent CNN models and their fusion. All the local patches are resized¹¹ to 224×224 as required by the Mobilenet-v1 model.

D. MobileNet CNN

Since the success of AlexNet [36] for object detection in ILSVRC-2012 [34], different CNN architectures have been proposed in literature, such as VGG, GoogleNet (Inception v1-v4), ResNets, MobileNet, etc. Nogueira et al. [27], winner of LivDet 2015, utilized a pre-trained VGG architecture [35] to achieve the best performance in LivDet 2015 [13]. In this study, we utilize the MobileNet-v1 architecture [37] because it offers the following advantages over VGG and Inception-v3 architectures: (i) MobileNet-v1 is designed using depth-wise separable convolutions, originally introduced in [41], providing drastic decrease in model size and training/evaluation times while providing better spoof detection performance, (ii) it is a low-latency network requiring only 100ms to classify an input fingerprint image as live or spoof compared to 800ms required by Inception-v3 network [1], (iii) the number of model parameters to be trained in MobileNet-v1 (4.24M), is significantly smaller than the number of model parameters in Inception-v3 (23.2M) and VGG (138M), requiring significantly lower efforts in terms of regularization and data augmentation, to prevent overfitting [37].

We utilized the TF-Slim library¹² implementation of the MobileNet-v1 architecture. The last layer of the architecture,

¹⁰MATLAB's *imrotate* function with bilinear interpolation is used to rotate the local patch for alignment.

¹¹TensorFlow's *resize* utility with bilinear interpolation was used; available at https://www.tensorflow.org/api_docs/python/tf/image/resize_images

¹²<https://github.com/tensorflow/models/tree/master/research/slim>




Fig. 7. The proposed approach provides a fine-grained representation for spoof detection by using minutiae-based local patches. A fingerprint spoof fabricated using silicone which conceals only a partial region of the live finger is shown in (a) and the imaged fingerprint in (b) (enclosed in red). The proposed approach extracts and evaluates the minutiae-based local patches, and highlights the local regions as live (in green) or spoof (in red) as shown in (c) and (d). It can also highlight the regions of fingerprint alterations as shown for a “Z” cut altered fingerprint in (e), (f) and (g). The proposed approach detected (b) and (e) as spoofs with the spoofness scores of 0.78 and 0.65, respectively.




Fig. 8. Local patches extracted around the fingerprint minutiae for (a) real fingerprint, and (b) spoof fingerprint (gelatin). The spoofness score for each patch is in the range $[0-1]$; higher the score, more likely the patch is extracted from a spoof fingerprint. For a given input test image, the spoofness scores corresponding to the local patches are averaged to give a global spoofness score. The final decision is made based on a classification threshold learned from the training dataset; an image with a global spoofness score below the threshold is classified as live, otherwise as spoof.

a 1000-unit softmax layer (originally designed to predict the 1,000 classes of ImageNet dataset), was replaced with a 2-unit softmax layer for the two-class problem, *i.e.* live vs. spoof. The optimizer used to train the network was RMSProp with asynchronous gradient descent and a batch size of 100. Data augmentation techniques, such as brightness adjustment, random cropping, vertical flipping are employed to ensure the trained model is robust to the possible variations in fingerprint images. For the multi-resolution local patches, a separate network is trained for each patch size with the same parameters as mentioned above.

E. Fine-grained Fingerprint Image Representation

Partial spoofs and fingerprint alterations are not a rare means to avoid re-identification¹³, by masking the true identity from a fingerprint biometric system [9], [42]. Spoof detectors trained

on the whole fingerprint images are inadvertently ineffective against localizing partial spoof fingerprints, that conceal only a limited region of the live finger. Moreover, in many smartphones and other embedded systems that only sense a partial region (friction ridge area) of the fingerprint due to small sensor area (typically 150×150), it is very crucial to have a detailed representation of the sensed fingerprint region. One of the key advantages of employing a patch-based approach is the fine-grained representation of input fingerprint image for spoof detection. Fig. 7 (a) presents an example of a fingerprint spoof fabricated using silicone, concealing only a partial region of the live finger and Fig. 7 (b) presents the imaged partial spoof fingerprint using a CrossMatch Guardian 200 fingerprint reader. The proposed approach, utilizing minutiae-based local patches, highlights the local regions as live or spoof (shown in Figs. 7 (c) and (d) in green and red, respectively), providing a fine-grained representation of the fingerprint image. Fingerprint alterations, such as cuts, mutilations, stitches, etc., performed using surgical and chemical procedures (see Fig. 7 (e)), create spurious minutiae as shown in Figs. 7 (f) and (g). The proposed approach is able to highlight the regions of fingerprint alterations despite not being trained specifically on altered fingerprint database, indicating the generalizability of the proposed approach. The proposed approach detected both fingerprint images in Figs. 7 (b) and (e) as spoofs with the spoofness scores of 0.78 and 0.65, respectively.

F. Spoofness Score

The output from the softmax layer of the trained MobileNet-v1 model is in the range $[0-1]$, called as the *Spoofness Score*. The larger the spoofness score, the higher the support that the input local patch belongs to the spoof class (see Fig. 8). For an input test image I , the spoofness scores $s_{i \in \{1, 2, \dots, k\}}^I$ corresponding to the k minutiae-based local patches of size $p \times p$, extracted from the input image, are averaged to give a global spoofness score S^I . In case of multi-resolution local patches, the global spoofness scores ($S_{p_i}^I$) based on each local patch size, $p_i \in \{64, 96, 128\}$, are averaged to produce a final spoofness score. The threshold that minimizes the average classification error on training dataset is learned and utilized

¹³<http://abcnews.go.com/Technology/GadgetGuide/surgically-altered-fingerprints-woman-evade-immigration/story?id=9302505>

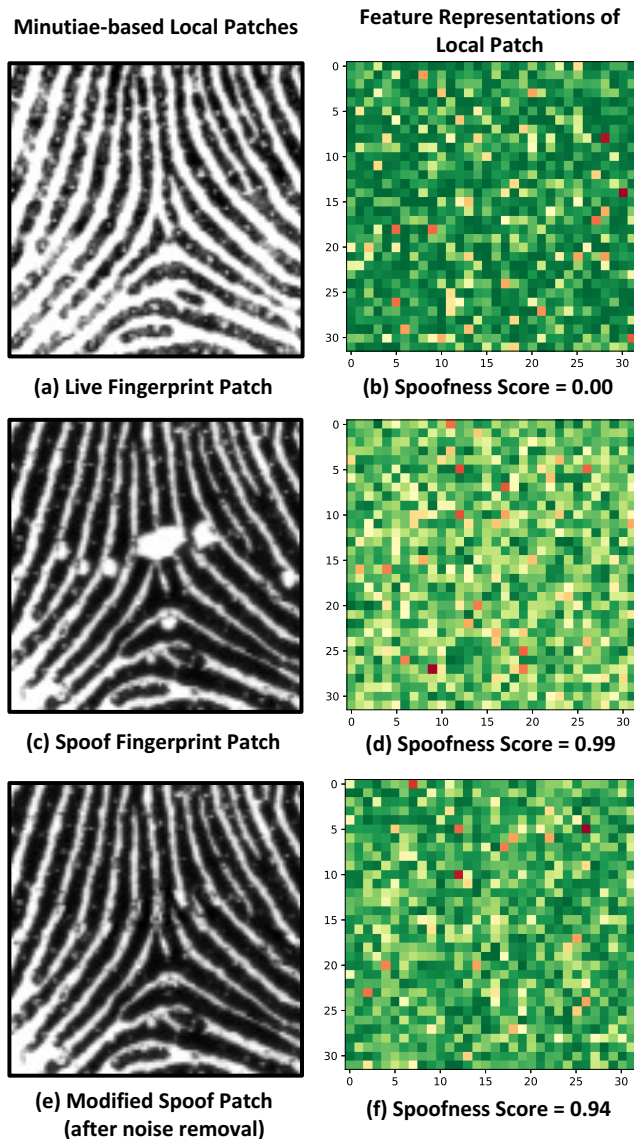


Fig. 9. Illustrating the embeddings of minutiae-based local patches (96×96), for (a) live patch, (c) spoof patch, and (e) modified spoof patch, in 1024-dimensional feature space from MobileNet-v1 bottleneck layer, transformed to 32×32 heat maps, (b), (d), and (e), respectively, for visualization. A high spoofness score for the modified spoof patch despite removal of artifacts indicates the robustness of the proposed approach.

as the classification threshold. An image with a spoofness score below the threshold is classified as live, otherwise as spoof. The learned threshold performed slightly better in spoof detection than selecting a pre-defined threshold of 0.5.

G. On Robustness of Patch-based Representation

As mentioned in Section-1, the proposed patch-based approach addresses several drawbacks of using whole fingerprint images for training deep CNN architectures. While the proposed approach is based on the premise that it is capable of capturing discriminatory information from local patches, such as missing ridges, air bubbles, etc., from spoof fingerprints, we also examine the robustness of patch-based representation

by evaluating it in absence of such artifacts. Figs. 9 (a) and (c) present minutiae-based local patches from a live fingerprint and the corresponding spoof fingerprint (fabricated using EcoFlex), respectively, for the same minutia point. Figs. 9 (b) and (d) present the feature representations of the local patches shown in Figs. 9 (a) and (c), respectively, obtained from the bottleneck layer of the MobileNet-v1 architecture. The 1024-dimensional feature representation is transformed to 32×32 heatmap for visualization. The spoofness scores for the two patches are 0.00 (Fig. 9 (b)) and 0.99 (Fig. 9 (d)), respectively. The spoof patch (Fig. 9 (c)) is modified, by the authors, using an open-source photo-editing utility, called *GIMP*¹⁴, to remove the visible artifacts and produce the modified spoof fingerprint patch as shown in Fig. 9 (e). The feature representation for the modified patch is shown in Fig. 9 (f). A high spoofness score for the modified spoof patch (0.94) despite removal of artifacts indicates the robustness of the proposed approach.

III. EXPERIMENTAL RESULTS

A. Performance Evaluation Metrics

The performance of the proposed approach is evaluated following the metrics used in LivDet [33].

- $F_{errlive}$: Percentage of misclassified live fingerprints.
- $F_{errfake}$ ¹⁵: Percentage of misclassified spoof fingerprints.

The average classification error (ACE) is defined as:

$$ACE = \frac{F_{errlive} + F_{errfake}}{2} \quad (1)$$

Additionally, we also report the $F_{errfake} @ F_{errlive} = 1\%$ for each of the experiments as reported in [33]. This value represents the percentage of spoofs able to breach the biometric system security when the reject rate of legitimate users $\leq 1\%$.

B. Results

The proposed approach is evaluated under the following four scenarios of fingerprint spoof detection, which reflect an algorithm's robustness against new spoof materials, use of different sensors and/or different environments.

1) *Intra-Sensor, Known Spoof Materials*: In this setting, all the training and testing images are captured using the same sensor, and all spoof fabrication materials utilized in the test set are known a priori. Our experimental results show that training the MobileNet-v1 model from scratch, using minutiae-based local patches, performs better than fine-tuning a pre-trained network, as reported in [27]. The large amount of available data, in the form of local fingerprint patches, is sufficient to train the deep architecture of MobileNet-v1 model without over-fitting. Also, a score level fusion of the proposed MobileNet-v1 model trained on local patches, with an independent MobileNet-v1 model trained on whole

¹⁴<https://www.gimp.org/>

¹⁵When all the spoof fabrication materials are known during the training, this metric is referred to as $F_{errfake_known}$, and in case all the spoof fabrication materials to be encountered during testing are not known during training, this metric is referred to as $F_{errfake_unknown}$.

TABLE IV

PERFORMANCE COMPARISON BETWEEN THE PROPOSED APPROACH (BOTTOM) AND STATE-OF-THE-ART (TOP) REPORTED ON LIVDET 2015 DATASET [13]. SEPARATE NETWORKS ARE TRAINED ON THE TRAINING IMAGES CAPTURED BY EACH OF THE FOUR FINGERPRINT READERS. *Ferrfake known* AND *Ferrfake unknown* CORRESPOND TO KNOWN SPOOF MATERIALS AND CROSS-MATERIAL SCENARIOS, RESPECTIVELY.

State-of-the-Art [13]	LivDet 2015	Ferrlive (%)	Ferrfake [†] (%)	Ferrfake known (%)	Ferrfake unknown* (%)	ACE (%)	Ferrfake (%) @ Ferrlive= 1% [33]
	GreenBit	3.50	5.33	4.30	7.40	4.60	17.90
	Biometrika	8.50	3.73	2.70	5.80	5.64	15.20
	Digital Persona	8.10	5.07	4.60	6.00	6.28	19.10
	Crossmatch	0.93	2.90	2.12	4.02	1.90	2.66
	Average	4.78	4.27	3.48	5.72	4.49	13.24
Proposed Approach	LivDet 2015	Ferrlive (%)	Ferrfake [†] (%)	Ferrfake known (%)	Ferrfake unknown* (%)	ACE (%)	Ferrfake (%) @ Ferrlive = 1%
	GreenBit	0.50	0.80	0.30	1.80	0.68	0.53
	Biometrika	0.90	1.27	0.60	2.60	1.12	1.20
	Digital Persona	1.97	1.17	0.85	1.80	1.48	1.96
	Crossmatch	0.80	0.48	0.82	0.00	0.64	0.28
	Average	1.02	0.93	0.64	1.48	0.97	0.96

† Ferrfake includes spoofs fabricated using both known and previously unseen materials. It is an average of Ferrfake-known and Ferrfake-unknown, weighted by the number of samples in each category.

*The unknown spoof materials in LivDet 2015 test dataset include Liquid Ecoflex and RTV for Green Bit, Biometrika, and Digital Persona sensors, and OOMOO and Gelatin for Crossmatch sensor.

TABLE V

PERFORMANCE COMPARISON BETWEEN THE PROPOSED APPROACH AND STATE-OF-THE-ART RESULTS REPORTED ON LIVDET 2011 AND LIVDET 2013 DATASETS FOR INTRA-SENSOR EXPERIMENTS IN TERMS OF AVERAGE CLASSIFICATION ERROR (ACE) AND FERRFAKE @ FERRLIVE = 1%.

Dataset	State-of-the-Art	Proposed Approach	
	ACE (%)	ACE (%)	Ferrfake @ Ferrlive = 1%
Biometrika	4.90 [25]	1.24	1.41
Digital Persona	1.85 [28]	1.61	3.25
ItalData	5.10 [28]	2.45	7.21
Sagem	1.23 [28]	1.39	4.33
Average	3.27	1.67	4.05
LivDet 2013			
Biometrika	0.65 [28]	0.20	0.00
ItalData	0.40 [27]	0.30	0.10
Average	0.53	0.25	0.05

fingerprint images does not offer any advantage in terms of performance improvement. This observation was also made in our preliminary study [1] which employed Inception-v3 model [38]. It was reported in [33] that most of the algorithms submitted to LivDet 2015 did not perform well on Digital Persona sensor due to the small image size. Our approach based on local patches does not suffer from this limitation. Tables IV and V present the performance comparison between the proposed approach and the state-of-the-art results for the LivDet datasets utilized in this study. Table VI presents the performance of the proposed approach on MSU Fingerprint Presentation Attack Dataset (MSU-FPAD) and Precise Biometrics Spoof-Kit Dataset (PBSKD). Independent MobileNet-v1 networks are trained for each evaluation. Note that in

TABLE VI

AVERAGE CLASSIFICATION ERROR (ACE), FERRFAKE @ FERRLIVE = 0.1% AND FERRLIVE = 1% ON THE MSU FINGERPRINT PRESENTATION ATTACK DATASET (MSU-FPAD) AND PRECISE BIOMETRICS SPOOF-KIT DATASET (PBSKD) FOR INTRA-SENSOR EXPERIMENTS.

Dataset	Proposed Approach		
	ACE (%)	Ferrfake @ Ferrlive = 0.1%	Ferrfake @ Ferrlive = 1%
MSU-FPAD			
CrossMatch Guardian 200	0.08	0.11	0.00
Lumidigm Venus 302	3.94	10.03	1.30
Average	2.01	5.07	0.65
PBSKD			
CrossMatch Guardian 200	2.02	5.32	0.65
Lumidigm Venus 302	1.93	3.84	0.33
Average	1.98	4.66	0.51

LivDet 2015 (Table IV), this scenario is represented by the *Ferrfake known*. For LivDet 2011 and 2013, MSU-FPAD, and PBSKD datasets (Table V), all spoof materials in the test set were known during training. Fig. 10 presents example fingerprint images for Biometrika sensor from LivDet 2015 dataset that were correctly and incorrectly classified by the proposed approach.

We also evaluate the impact of local patch size on the performance of the proposed approach, by comparing the performance of three CNN models trained on minutiae-centered local patches of size $[p \times p]$ where $p = \{64, 96, 128\}$, extracted from the fingerprint images captured by Biometrika sensor for LivDet 2011 dataset. Among these three models, the one trained on local patches of size $[96 \times 96]$ performed the best. However, a score-level fusion, using average-rule, of the three models reduced the average classification error (ACE) from 1.24% to 0.88%, and Ferrfake from 1.41% to 0.58% @ Ferrlive = 1%. Similar performance gains were observed for other




Fig. 10. Example live and spoof fingerprints for Biometrika sensor from LivDet 2015 dataset, correctly and incorrectly classified by our proposed approach.




Fig. 11. ROC curves for live v. spoof classification of fingerprint images from LivDet 2011 Dataset (Biometrika sensor) utilizing (i) whole image, (ii) randomly selected patches $[96 \times 96]$, (iii) minutiae-based patches of size $[p \times p]$, $p \in \{64, 96, 128\}$, (iv) score-level fusion of multi-resolution patches.

sensors, but there is a trade off between the performance gain and the computational requirements for spoof detector. In order to evaluate the significance of utilizing minutiae locations for extracting local patches, we trained independent MobileNet-v1 models on a similar number of local patches, extracted

randomly from LivDet 2015 datasets. It was observed that the models trained on minutiae-centered local patches achieved a significantly higher reduction (78%) in average classification error, compared to the reduction (33%) achieved by the models trained on randomly sampled local patches. Fig. 11 illustrates that (i) features extracted from local patches provide better spoof detection accuracy than the whole image, (ii) patches selected around minutiae perform better than random patches of the same size, (iii) 96×96 patch performs the best among the three patch sizes considered, and (iv) score-level fusion of multi-resolution local patches boosts the spoof detection performance.

2) *Intra-Sensor, Cross-Material*: In this setting, the same sensor is used to capture all training and testing images, but the spoof images in the testing set are fabricated using new materials that were not seen during training. For the first set of cross-material experiments, we utilize the LivDet 2015 dataset which contains two new spoof materials in the testing set for each sensor, *i.e.* Liquid Ecoflex and RTV for Green Bit, Biometrika, and Digital Persona sensors, and OOMOO and Gelatin for Crossmatch sensor. The performance of the proposed approach on cross-material experiments for LivDet 2015 dataset is presented in Table IV (column $F_{err\ fake_unknown}$) and is compared with the state-of-the-art performance reported in [13]. A significant reduction in the error rate is achieved by the proposed method. For better generalizability, a second set of cross-material experiments are performed on LivDet

TABLE VII

PERFORMANCE COMPARISON BETWEEN THE PROPOSED APPROACH AND STATE-OF-THE-ART RESULTS REPORTED ON LIVDET 2011 AND LIVDET 2013 DATASETS FOR CROSS-MATERIAL EXPERIMENTS, IN TERMS OF AVERAGE CLASSIFICATION ERROR (ACE) AND FERRFAKE @ FERRLIVE = 1%.

Dataset	Spoof Materials		State-of-the-Art	Proposed Approach	
	Materials - Training	Materials - Testing	ACE (%)	ACE (%)	Ferrfake @ Ferrlive = 1%
Biometrika 2011	EcoFlex, Gelatine, Latex	Silgum, WoodGlue	10.10 [27]	4.60	8.15
Biometrika 2013	Modasil, WoodGlue	EcoFlex, Gelatine, Latex	2.10 [28]	1.30	0.34
ItalData 2011	EcoFlex, Gelatine, Latex	Silgum, WoodGlue, Other	7.00 [28]	5.20	7.80
ItalData 2013	Modasil, WoodGlue	EcoFlex, Gelatine, Latex	1.25 [28]	0.60	0.68
Average			5.11	2.93	4.24

TABLE VIII

PERFORMANCE COMPARISON BETWEEN THE PROPOSED APPROACH AND STATE-OF-THE-ART RESULTS [27] REPORTED ON LIVDET 2011 AND LIVDET 2013 DATASETS FOR CROSS-SENSOR EXPERIMENTS, IN TERMS OF AVERAGE CLASSIFICATION ERROR (ACE), AND FERRFAKE @ FERRLIVE = 1%.

Training Dataset (Testing Dataset)	State-of-the-Art	Proposed Approach	
	ACE (%)	ACE (%)	Ferrfake (%) @ Ferrlive = 1%
Biometrika 2011 (ItalData 2011)	29.35 [28]	25.35	50.81
ItalData 2011 (Biometrika 2011)	27.65 [28]	25.21	76.20
Biometrika 2013 (ItalData 2013)	1.5 [28]	4.30	12.73
ItalData 2013 (Biometrika 2013)	2.30 [27]	3.50	70.35
Average	15.2	14.59	52.52

TABLE IX

PERFORMANCE COMPARISON BETWEEN THE PROPOSED APPROACH AND STATE-OF-THE-ART RESULTS [28] REPORTED ON LIVDET 2011 AND LIVDET 2013 DATASETS FOR CROSS-DATASET EXPERIMENTS, IN TERMS OF AVERAGE CLASSIFICATION ERROR (ACE) AND FERRFAKE @ FERRLIVE = 1%.

Training Dataset (Testing Dataset)	State-of-the-Art	Proposed Approach	
	ACE (%)	ACE (%)	Ferrfake (%) @ Ferrlive = 1%
Biometrika 2011 (Biometrika 2013)	14.00 [28]	7.60	89.60
Biometrika 2013 (Biometrika 2011)	34.05 [28]	31.16	78.84
ItalData 2011 (ItalData 2013)	8.30 [28]	6.70	16.70
ItalData 2013 (ItalData 2011)	44.65 [28]	26.16	75.09
Average	25.25	17.91	65.06

2011 and LivDet 2013 datasets, following the protocol adopted by the winner of LivDet 2015 [27]. Table VII presents the achieved error rates on these experiments, along with the spoof fabrication materials used in training and testing sets.

3) *Cross-Sensor Evaluation*: In this evaluation, the training and the testing images are obtained from two different sensors but from the same dataset. This setting reflects the algorithm's strength in learning the common characteristics used to distinguish live and spoof fingerprints across fingerprint

acquisition devices. For instance, using LivDet 2011 dataset, images from Biometrika sensor are used for training, and the images from ItalData sensor are used for testing. We follow the protocol for selection of training and testing sets for cross-sensor and cross-dataset experiments as adopted by Nogueira et al. [27]. Table VIII compares the average classification error and Ferrfake @ Ferrlive = 1% for the proposed approach with the state-of-the-art results obtained by [27] and [28] on cross-sensor experiments.

4) *Cross-Dataset Evaluation*: In this scenario, the training and the testing images are obtained using the same sensor, but from two different datasets. For instance, training images are acquired using Biometrika sensor from LivDet 2011 dataset and the testing images are acquired using Biometrika sensor from LivDet 2013. This set of experiments captures the algorithm's invariance to the changes in environment for data collection. Table IX presents the average classification error and Ferrfake @ Ferrlive = 1%. Results in Table IX show that the proposed local patch based approach achieves a reduction of 29% in the average classification error from 25.25% in [28] to 17.91% in our approach. However, the average Ferrfake @ Ferrlive = 1% that we report is 52.52% and 65.06% for cross-sensor and cross-dataset scenarios respectively, indicating the challenges, especially in applications where a high level of spoof detection accuracy is needed.

C. Processing Times

The MobileNet-v1 CNN model takes around 6-8 hours to converge using a single Nvidia GTX 1080 Ti GPU utilizing approximately 96,000 local patches for a training set with 2,000 fingerprint images (2,000 images \times 48 patches/fingerprint image). The average classification time for a single input image, including minutiae detection, local patch extraction and alignment, inference of Spoofness Scores for local patches, and producing the final decision, on a single Nvidia GTX 1080 Ti GPU is 100ms.

IV. FINGERPRINT SPOOF BUSTER

A graphical user interface for real-time fingerprint spoof detection, called Fingerprint Spoof Buster, is developed. This interface allows the operator to select a specific fingerprint reader and a trained MobileNet-v1 model for evaluation. The operator can perform the evaluation in either *online* or *batch* mode. In the *online* mode, a fingerprint is imaged using the

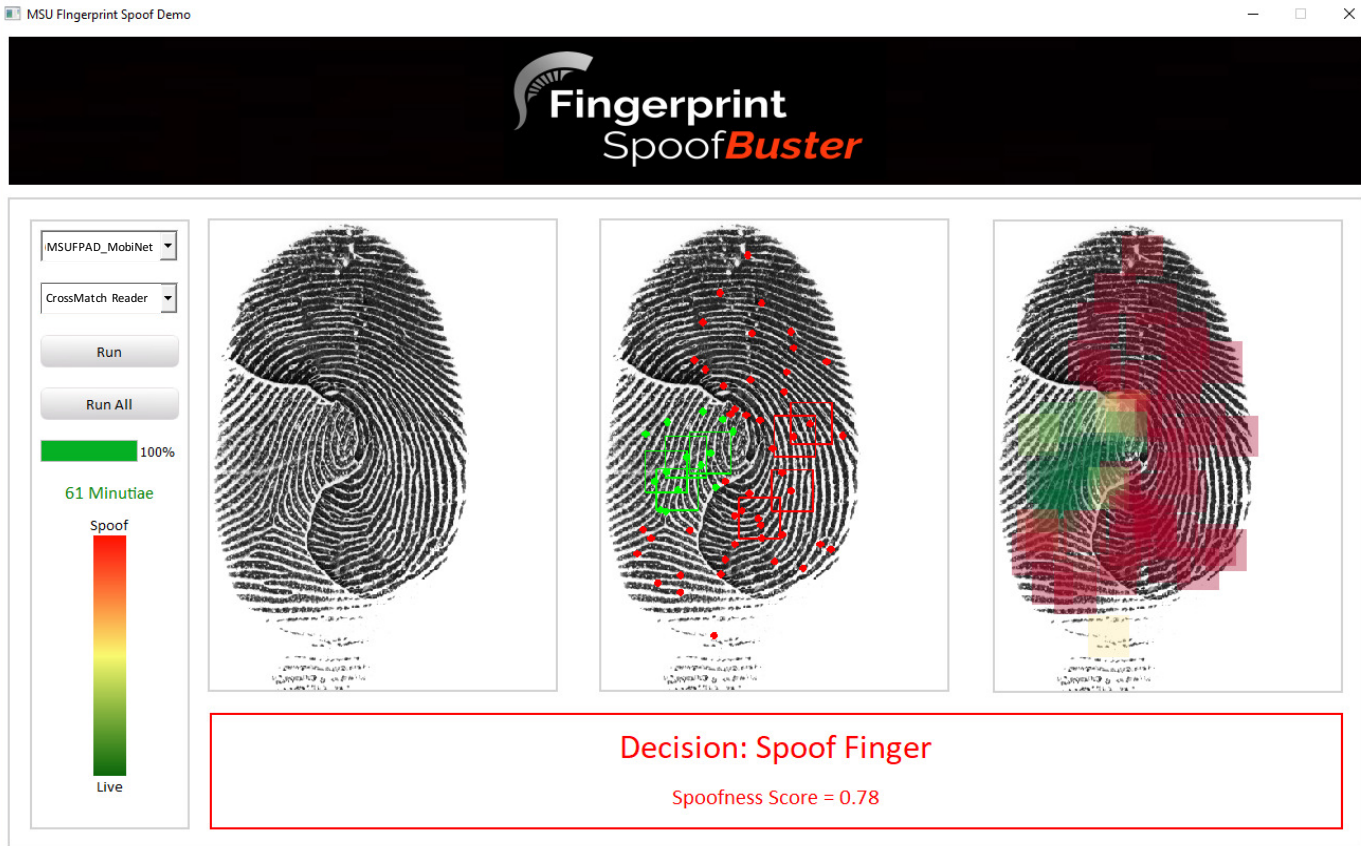


Fig. 12. Interface of the real-time fingerprint spoof detector, called Fingerprint Spoof Buster. It allows selection of the fingerprint reader and CNN model.

selected reader and displayed on the interface (see Fig. 12). The extracted fingerprint minutiae and the corresponding local patches are presented and color coded based on their respective spoofness scores (green for live and red for spoof). The global spoofness score and the final decision for the input image is also presented on the interface. In the batch mode, all fingerprint images within a specified directory are evaluated, and global spoofness scores for each file are output in a file. The graphical user interface allows the operator to visually examine the local regions of the fingerprint highlighted as live or spoof, instead of relying on only a single score as output by the traditional approaches. In the future, this interface will be extended to display output from multiple CNN models for an easy visual comparison.

V. CONCLUSIONS

A robust and accurate method for fingerprint spoof detection is critical to ensure the reliability and security of the fingerprint authentication systems. In this study, we have utilized fingerprint domain knowledge by extracting local patches centered and aligned using minutiae in the input fingerprint image for training MobileNet-v1 CNN models. The local patch based approach provides salient cues to differentiate spoof fingerprints from live fingerprints. The proposed approach is able to achieve a significant reduction in the error rates for intra-sensor (63%), cross-material (43%), cross-sensor (4%)

as well as cross-dataset scenarios (29%) compared to state-of-the-art on public domain LivDet datasets.

ACKNOWLEDGMENT

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA R&D Contract No. 2017 – 17020200004. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

REFERENCES

- [1] T. Chugh, K. Cao, and A. K. Jain, "Fingerprint Spoof Detection using Minutiae-based Local Patches," in *IEEE International Joint Conference on Biometrics (IJCB)*, 2017.
- [2] S. Marcel, M. S. Nixon, and S. Z. Li, *Handbook of Biometric Anti-Spoofing*. Springer, 2014.
- [3] ODNI, IARPA, "IARPA-BAA-16-04 (Thor)," <https://www.iarpa.gov/index.php/research-programs/odin/odin-baa>, 2016.
- [4] International Standards Organization, "ISO/IEC 30107-1:2016, Information Technology—Biometric Presentation Attack Detection—Part 1: Framework," <https://www.iso.org/standard/53227.html>, 2016.
- [5] T. Matsumoto, H. Matsumoto, K. Yamada, and S. Hoshino, "Impact of artificial gummy fingers on fingerprint systems," in *Proc. SPIE*, vol. 4677, 2012, pp. 275–289.

- [6] K. Cao and A. K. Jain, "Hacking mobile phones using 2D Printed Fingerprints," *MSU Tech. report, MSU-CSE-16-2*, 2016.
- [7] S. S. Arora, K. Cao, A. K. Jain, and N. G. Paulter, "Design and Fabrication of 3D Fingerprint Targets," *IEEE TIFS*, vol. 11, no. 10, pp. 2284–2297, 2016.
- [8] S. S. Arora, A. K. Jain, and N. G. Paulter, "Gold Fingers: 3D Targets for Evaluating Capacitive Readers," *IEEE TIFS*, vol. 12, no. 9, pp. 2067–2077, 2017.
- [9] S. Yoon, J. Feng, and A. K. Jain, "Altered fingerprints: Analysis and detection," *IEEE TPAMI*, vol. 34, no. 3, pp. 451–464, 2012.
- [10] E. Marasco and A. Ross, "A survey on antispooing schemes for fingerprint recognition systems," *ACM Computing Surveys*, vol. 47, no. 2, p. 28, 2015.
- [11] B. Biggio, Z. Akhtar, G. Fumera, G. L. Marcialis, and F. Roli, "Security evaluation of biometric authentication systems under real spoofing attacks," *IET Biometrics*, vol. 1, no. 1, pp. 11–24, 2012.
- [12] D. Yambay, L. Ghiani, P. Denti, G. L. Marcialis, F. Roli, and S. Schuckers, "LivDet 2011-Fingerprint liveness detection competition 2011," in *Proc. 5th IAPR ICB*, 2012, pp. 208–215.
- [13] V. Mura, L. Ghiani, G. L. Marcialis, F. Roli, D. A. Yambay, and S. A. Schuckers, "LivDet 2015 - Fingerprint liveness detection competition 2015," in *Proc. IEEE 7th Intl. Conf. BTAS*, 2015, pp. 1–6.
- [14] P. D. Lapsley, J. A. Lee, D. F. Pare Jr, and N. Hoffman, "Anti-fraud biometric scanner that accurately detects blood flow," US Patent 5,737,439, 1998.
- [15] A. Antonelli, R. Cappelli, D. Maio, and D. Maltoni, "Fake finger detection by skin distortion analysis," *IEEE TIFS*, vol. 1, no. 3, pp. 360–373, 2006.
- [16] D. Baldisserra, A. Franco, D. Maio, and D. Maltoni, "Fake fingerprint detection by odor analysis," in *Proc. ICB*. Springer, 2006, pp. 265–272.
- [17] R. K. Rowe and D. P. Sidlauskas, "Multispectral biometric sensor," Dec. 12 2006, uS Patent 7,147,153.
- [18] J. N. Hogan, "Multiple reference oct system," Aug. 25 2015, uS Patent 9,113,782.
- [19] J. J. Engelsma, K. Cao, and A. K. Jain, "Respireader: An open source fingerprint reader facilitating spoof detection," *arXiv preprint arXiv:1708.07887*, 2017.
- [20] G. L. Marcialis, F. Roli, and A. Tidu, "Analysis of fingerprint pores for vitality detection," in *Proc. 20th ICPR*, 2010, pp. 1289–1292.
- [21] E. Marasco and C. Sansone, "Combining perspiration-and morphology-based static features for fingerprint liveness detection," *Pattern. Reco. Letters*, vol. 33, no. 9, pp. 1148–1156, 2012.
- [22] L. Ghiani, G. L. Marcialis, and F. Roli, "Fingerprint liveness detection by local phase quantization," in *Proc. 21st ICPR*, 2012, pp. 537–540.
- [23] L. Ghiani, A. Hadid, G. L. Marcialis, and F. Roli, "Fingerprint liveness detection using Binarized Statistical Image Features," in *Proc. IEEE 6th Intl. Conf. BTAS*, 2013, pp. 1–6.
- [24] D. Gragnaniello, G. Poggi, C. Sansone, and L. Verdoliva, "Fingerprint liveness detection based on Weber Local Image Descriptor," in *Proc. IEEE Workshop on Biometric Meas. Syst. Secur. Med. Appl.*, 2013, pp. 46–50.
- [25] —, "Local contrast phase descriptor for fingerprint liveness detection," *Pattern Recognition*, vol. 48, no. 4, pp. 1050–1058, 2015.
- [26] D. Menotti, G. Chiachia, A. Pinto, W. R. Schwartz, H. Pedrini, A. X. Falcao, and A. Rocha, "Deep representations for iris, face, and fingerprint spoofing detection," *IEEE TIFS*, vol. 10, no. 4, pp. 864–879, 2015.
- [27] R. F. Nogueira, R. de Alencar Lotufo, and R. C. Machado, "Fingerprint Liveness Detection Using Convolutional Neural Networks," *IEEE TIFS*, vol. 11, no. 6, pp. 1206–1213, 2016.
- [28] F. Pala and B. Bhanu, "Deep Triplet Embedding Representations for Liveness Detection," in *Deep Learning for Biometrics. Advances in Computer Vision and Pattern Recognition*. Springer, 2017, pp. 287–307.
- [29] D. Gragnaniello, G. Poggi, C. Sansone, and L. Verdoliva, "An investigation of local descriptors for biometric spoofing detection," *IEEE TIFS*, vol. 10, no. 4, pp. 849–863, 2015.
- [30] E. Marasco and C. Sansone, "On the Robustness of Fingerprint Liveness Detection Algorithms against New Materials used for Spoofing," in *Proc. Intl. Conf. Bio-Insp. Syst. Sign. Process.*, 2011, pp. 553–558.
- [31] A. Rattani, W. J. Scheirer, and A. Ross, "Open set fingerprint spoof detection across novel fabrication materials," *IEEE TIFS*, vol. 10, no. 11, pp. 2447–2460, 2015.
- [32] Y. Ding and A. Ross, "An ensemble of one-class SVMs for fingerprint spoof detection across different fabrication materials," in *Proc. IEEE WIFS*, 2016, pp. 1–6.
- [33] L. Ghiani, D. A. Yambay, V. Mura, G. L. Marcialis, F. Roli, and S. A. Schuckers, "Review of the Fingerprint Liveness Detection (LivDet) competition series: 2009 to 2015," *Image and Vision Computing*, vol. 58, pp. 110–128, 2017.
- [34] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *Proc. IJCV*, vol. 115, no. 3, pp. 211–252, 2015.
- [35] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [36] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. NIPS*, 2012, pp. 1097–1105.
- [37] 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.
- [38] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the Inception Architecture for Computer Vision," in *Proc. IEEE CVPR*, 2016, pp. 2818–2826.
- [39] L. Ghiani, D. Yambay, V. Mura, S. Tocco, G. L. Marcialis, F. Roli, and S. Schuckers, "LivDet 2013 Fingerprint Liveness Detection Competition 2013," in *Proc. ICB*, 2013, pp. 1–6.
- [40] K. Cao, E. Liu, L. Pang, J. Liang, and J. Tian, "Fingerprint matching by incorporating minutiae discriminability," in *Proc. IJCB*, 2011, pp. 1–6.
- [41] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," *arXiv preprint arXiv:1610.02357*, 2016.
- [42] A. Roy, N. Memon, and A. Ross, "Masterprint: Exploring the vulnerability of partial fingerprint-based authentication systems," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 9, pp. 2013–2025, 2017.



Tarang Chugh received the B. Tech. (Hons.) degree in Computer Science and Engineering from the Indraprastha Institute of Information Technology, Delhi (IIIT-D) in 2013. He was affiliated with IBM Research Lab, New Delhi, India as a research engineer during 2013-2015. He is currently a doctoral student in the Department of Computer Science and Engineering at Michigan State University. His research interests include biometrics, pattern recognition, and machine learning.



Kai Cao received the Ph.D. degree from the Key Laboratory of Complex Systems and Intelligence Science, Institute of Automation, Chinese Academy of Sciences, Beijing, China, in 2010. He is currently a Post Doctoral Fellow in the Department of Computer Science & Engineering, Michigan State University. He was affiliated with Xidian University as an Associate Professor. His research interests include biometric recognition, image processing and machine learning.



Anil K. Jain is a University distinguished professor in the Department of Computer Science and Engineering at Michigan State University. His research interests include pattern recognition and biometric authentication. He served as the editor-in-chief of the IEEE Transactions on Pattern Analysis and Machine Intelligence and was a member of the United States Defense Science Board. He has received Fulbright, Guggenheim, Alexander von Humboldt, and IAPR King Sun Fu awards. He is a member of the National Academy of Engineering and foreign fellow of the Indian National Academy of Engineering.