

Robust and Interoperable Fingerprint Spoof Detection via Convolutional Neural Networks

Emanuela Marasco

Dept. of Computer Science
Univ. of North Carolina at Charlotte
28223 Charlotte, NC, USA.
Email: emarasco@uncc.edu

Peter Wild

Digital Safety & Security Department
AIT Austrian Institute of Technology GmbH
2444 Seibersdorf, Austria.
Email: peter.wild@ait.ac.at

Bojan Cukic

Dept. of Computer Science
Univ. of North Carolina at Charlotte
28223 Charlotte, NC, USA.
Email: bcukic@uncc.edu

Abstract—Fingerprint recognition for automated border control and other high-security applications needs robust integrated anti-spoofing capability. Facing the threat of presentation attacks, two key challenges to be solved are sensor interoperability and robustness versus new fabrication materials. This paper proposes convolutional neural networks for this task and presents an exhaustive comparison on latest LivDet 2011 and 2013 databases. Apart from classical classification nets, also metric-based deep siamese networks are evaluated learning a distance metric enforcing live-spoof pairs to be of higher distance than live-live pairs. This is useful for attended enrollment scenarios where a live gallery image is available (e.g. trusted-source fingerprint reference on the passport chip). Experiments reveal remarkable accuracy for all Convolutional Neural Networks (CNNs) CaffeNet (96.5%), GoogLeNet (96.6%), Siamese (93.1%), good material robustness (max. 5.6% diff.) but weak sensor-interoperability.

I. INTRODUCTION

With their widespread use for biometric identification, fingerprint systems are a key target for presentation attacks. Presentation Attack Detection (PAD) modules classify biometric samples as either live (non-spoof) or fake (spoof), trying to detect attempts to interfere with biometric systems' operation (ISO/IEC 30107-1). Due to covert acquisition and large variety of materials for presentation attack instruments (fake fingers) fingerprints are largely exposed. Since PAD algorithms are typically based on features specific for the fabrication material and sensor used for training, new fabrication materials and sensor diversity generally degrade spoof detection performance drastically [9]. Fabrication materials and live images likely appear different across sensors, affecting PAD cross-sensor performance and making robustness with regards to device diversity a highly challenging problem [8]. Recent works are looking at options to automatically adapt the PAD module to unseen new spoofing materials [12]. For robustness to new materials and sensor diversity, hand crafted features are not easy to generalize and prone to overfitting.

This paper investigates and compares deep convolutional neural networks (CNNs) [1] for PAD. Compared to previous work in this direction [4], networks with boosted accuracy are identified (CaffeNet [7] and GoogLeNet [14]), transfer-learning (fine-tuning using a pre-trained model) concepts to deal with limited amounts of training data are employed, and new siamese configurations using evidence about the template




Fig. 1: CaffeNet-based PAD Module mode of operation.




Fig. 2: Siamese PAD Module integrating gallery sample.

in question (an idea illustrated in [15]) are examined, see Figs.1,2. The paper presents an unseen evaluation focused on assessing generalization of methods: tests with unseen materials and sensor interoperability assessment.

The remainder is organised as follows: An introduction to related work is given in Section II. Tested CNNs are outlined in Section III. Exhaustive experiments conducted on multiple sensor-/material-specific sets of LivDet are presented in Section IV. Section V forms the conclusion.

II. RELATED WORK

Fake fingerprints can be made from a series of different materials, including silicone, latex, gelatin, play-doh, waxes, and wood glue, or even cadavers [8]. Presentation attack success depends on the type of sensing technology (e.g., optical vs. capacitive). In the effort to improve the state-of-the-art presentation attack detectors, several liveness detection competitions (LivDet 2009, 2011 2013 and most recently, 2015) have been conducted [6]. Generally, presentation attacks can be detected by either gathering further evidence of the liveness of the subject (e.g. sensing blood circulation, or fluids

TABLE I: Public databases for fingerprint liveness detection.

Name	Dataset	Sensor	Resolution	Train		Test		Subjects	Best Performance		
				Live	Spoof	Live	Spoof		APCER	NPCER	Ref
Biometrika'13	LivDet13	Biometrika Fx2000	569 dpi	1000	5x200	1000	5x200	75	0.1%	3.3%	[6]
Italdata'13	LivDet13	Italdata ET10	500 dpi	1000	5x200	1000	5x200	50	1.0%	0.2%	[6]
Sagem'11	LivDet11	Sagem MSO300	500 dpi	1000	5x200	1000	5x200	56	13.1%	13.8%	[16]
Digital'11	LivDet11	Digital Persona 400B	500 dpi	1000	5x200	1000	5x200	100	11.6%	6.2%	[16]
Identix'09	LivDet09	Identix DFR2100	686 dpi	2250	3x750	750	3x250	160	1.4%	1.1%	[5]

- perspiration patterns - secreted when touching surfaces) or by passive methods detecting the presence of known materials (e.g. material structure, lack of high-resolution detail). This work focuses on the latter software-based methods, for which several methods, including Fourier Transform (FT), Local Binary Patterns (LBP), or Histograms of Invariant Gradients (HIG), have been investigated. Good overviews of PAD methods in fingerprint recognition can be found in [13], [8].

Although deep learning approaches have been applied for fingerprint liveness detection [4], presented techniques were largely hybrid (combined with other classification techniques) and no pure CNNs have been evaluated for this task shedding light onto cross-sensor and robustness to new fabrication materials. Previously, Menotti et al. derived an efficient spoof detection system through deep representations [11]. They first learn a suitable CNN architecture, which is determined through a random search procedure involving hyperparameter optimization of the network. The candidate architecture is evaluated by executing linear SVM on the deep representation obtained by the considered net. Filter weights are optimized via back-propagation. Recently, Frassetto et al. [4] have examined fingerprint liveness detection using CNNs and LBPs, however they employ a hybrid approach feeding the net's output into an SVM rather than exploiting the power of deep networks only with best reported accuracy of 95.2% using 50,000 samples for training using LivDet 2009, 2011 and 2013 datasets.

The approach in this paper is even further enhancing accuracy with much less available training samples (adhering to the strict LivDet 2013 protocol).

III. CNNs FOR PAD

CNNs as deep instances of neural networks use multiple layers of non-linear processing units, each of which (single perceptron) estimate an output hypothesis $h_{w,b}(x) = f(w^T x) = f(\sum_{i=1}^n w_i x_i + b)$ using output vector x of the preceding layer. Within a training phase, using k Training samples (x^i, y^i) , $x^i \in \mathbb{R}^n, y^i \in \mathbb{R}^m$ a Matrix W (and vector b) is learnt, such that $h_{W,b}(x) \approx y$ when applying forward propagation. Weights are learned through back-propagation, iteratively updating weights optimizing a cost function $J(W, b) = \frac{1}{m} \sum_{i=0}^m (\frac{1}{2} \|h_{W,b}(x^{(i)}) - y^{(i)}\|^2) + \frac{\lambda}{2} \delta(W)$, where $\delta(W)$ is the weight decay term. CNNs exploit stationarity in natural images using patch-wise convolution thereby

drastically reducing the number of weights. We follow a fine-tuning strategy re-using pre-trained layers from these caffe¹ reference models:

- **CaffeNet** [7]: This adapted Alexnet implementation follows the idea of establishing a large, deep CNN through 5 convolutional and 3 fully connected layers on 227x227 RGB image patches. The raw input image is transformed into gradually higher levels of representation (e.g., edges, local shapes, object parts). Higher levels of the hierarchy are formed by composition of lower level features, see Fig. 1. The layers of the CNN are organized such that the first two are subdivided into four sublayers each: convolution (conv), max(x,0) rectified linear units (RELU), max pooling, and local response normalization.
- **GoogLeNet** [14]: This complex network of 22 layers (27 with pooling) for multi-scale Hebbian principle image processing is the best-performing net of the ImageNet² Large-Scale Visual Recognition Challenge 2014, taking 224x224 RGB image patches. It employs spatially stacked inception modules with dimension reductions, approximating an optimal sparse structure by dense building blocks.

A. Fine-tuning

Weights $W^{(i)}$ for different layers i can be learned using different models, re-using weights as initial values for a different model. Thereby learning speed is increased and overfitting avoided, which is crucial for the limited amount of spoofing images. Presented networks were pre-trained using 1.2 million ImageNet natural images using the original 1000 classes output layer (which was replaced to binary classification, as in the Kaggle challenge³) and stochastic gradient descent to optimize $J(W, b)$. As key adaption we used a low learning rate (10^{-4} to 10^{-6}) running a stochastic gradient descent on the target loss function for 5000 iterations. Supervised pre-training and fine-tuning are especially effective when training data is scarce.

B. Siamese Network

Siamese networks are neural networks where two branches processing image patches share exactly the same setup and weights [3], see also Fig. 2. The two branches' outputs

¹<http://caffe.berkeleyvision.org/>

²<http://www.image-net.org/challenges/LSVRC/>

³<https://www.kaggle.com/c/dogs-vs-cats>

are connected (concatenated) and processed in yet another layer learning the similarity metric. Siamese networks are particularly interesting for PAD as they allow to learn a distance metric forcing low pairwise distance between live-live samples and higher distance between live-spoof samples. Input patterns are mapped into a target space by the function $G_w(X)$ parameterized by w , such that the similarity metric $E_w(X_1, X_2) = \|G_w(X_1) - G_w(X_2)\|$ is small if X_1 and X_2 are both live fingerprints, and large if X_1 is a live fingerprint and X_2 is a spoof fingerprint. W is the shared parameter vector that is subject to learning (ground truth $Y = 0$ for live-live pairs (X_1, X_2) and $Y = 1$ for live-spoof pairs (X_1, X_2)). In training all possible combinations live-live and live-spoof are considered. Instead of using final probability layers indicating class memberships, Siamese networks can be built by replacing these layers with a linear "feature" layer that produces a 2 dimensional vector "comparing" the two vectors. The learning process minimizes a discriminative loss function that drives the similarity metric to be small for pairs of live fingerprints (in both intra- and cross-sensor) and large for pairs live-spoof.

This approach reformulates the typical classification problem adopted for static fingerprint liveness detection which processes a single fingerprint image. Specifically, the proposed method assumes an attended enrollment scenario in which there is access to live reference samples, for example the passport holder's fingerprint stored on the chip of an ICAO 9303 travel document. Note, this reduces the joint PAD-and-recognition problem [2] from 8-classes (live/spoof template vs. live-spoof sample with same/different source) to 4 classes.

IV. EXPERIMENTAL STUDY

This study focuses on (1) *comparing the performance of CNNs* for PAD; (2) *efficient parameter selection* finding best setup for learning rate, fine-tuning layers, and iteration count in presence of a limited number of training samples; (3) *robustness with unseen materials* leaving out materials at the training stage; and (4) *interoperability* evaluating the impact of changes in sensors.

A. Dataset

Our experiments were conducted on the publicly available databases provided in the Fingerprint Liveness Detection Competition⁴, LivDet 2009 [10], LivDet 2011 [16] and LivDet 2013 [6] keeping original Train and Test subdivisions, see Table I. Examples of fingerprint images and in particular material variability are illustrated in Fig. 3 and for sensor variability in Fig 4. These databases come with different sensors and spoofing materials, with a balanced 1:1 live/fake fingerprint ratio.

B. Evaluation Procedure

In order to evaluate robustness to new fabrication materials and device diversity, we closely followed the LivDet training

and test set partition and employed latest ISO/IEC 30107 metrics for evaluating presentation attacks at sensor-level:

- **APCER:** Attack presentation classification error rate - incorrectly as normal classified attack presentations (false acceptance of spoof samples, *ferrfake*).
- **NPCER:** Normal presentation classification error rate - incorrectly as attack classified normal presentations (false rejection of live samples, *ferrlive*).

Rates refer to test set accuracy by choosing a threshold where $APCER \approx NPCER$ for the training set. Further the AUC (area under the curve) of the Receiver Operating Characteristic (ROC) is reported. LivDet trains each sensor separately, and by representing the spoof class with same set of fabrication materials during both training and testing.

For robustness to new fabrication materials, we exclude each of the fabrication materials from training which will be used only for testing, i.e. for latex tests training includes all the fabrication materials from training sets except latex, while test involves either all or just the left-out material (latex) for the particular sensor and database (Biometrika'13). For interoperability assessment a similar approach is taken, however fixing evaluations to one material (Gelatin). Training uses 3 and 4 sensors, testing involves multiple or individual sensors using the sets of LivDet11, LivDet13 as listed in Table I.

C. Parameter Selection Results

With regards to parameter estimation, we employed the separate Identix'09 database as validation set for the *Siamese* net (CaffeNet in siamese configuration). Table III lists experimental results for this net: few features (25) in the last layer before comparison are optimal and base learn rate (LR) should ideally be larger than for the classical CaffeNet (best result for LR 10^{-5}), which makes sense given the much larger amount of available training images. CaffeNet and GoogLeNet




Fig. 3: Fabrication materials used to realize spoof samples with the Biometrika sensor (from LivDet13 [6]).

⁴<http://prag.dice.unica.it/>

TABLE II: Robustness to unseen materials (All = Ecoflex, Gelatin, Latex, Modasil, WoodGlue) on Biometrika’13.

Training	Testing	Siamese			CaffeNet			GoogLeNet		
		AUC	APCER	NPCER	AUC	APCER	NPCER	AUC	APCER	NPCER
All	All	0.971	6.1%	7.8%	0.991	4.2%	2.9%	0.992	4.3%	2.5%
All	Ecoflex	0.997	4.9%	7.8%	0.998	0.5%	2.9%	0.999	0%	2.5%
All	Gelatin	0.905	19.9%	7.8%	0.978	9%	2.9%	0.981	10%	2.5%
All	Latex	0.986	2.1%	7.8%	0.994	3.5%	3%	0.992	4.5%	2.5%
All	Modasil	0.991	1.9%	7.8%	0.995	2.5%	3%	0.991	3%	2.5%
All	WoodGlue	0.976	7.8%	5.9%	0.991	4.5%	2.9%	0.996	4%	2.5%
All without Ecoflex	All	0.982	6.1%	4.3%	0.999	1.1%	2.1%	0.993	3.2%	4.4%
All without Ecoflex	Ecoflex	0.996	2.5%	4.3%	0.999	0%	2.1%	0.999	0%	4.4%
All without Gelatin	All	0.986	4.7%	4.7%	0.998	4.6%	0.4%	0.991	7%	7.5%
All without Gelatin	Gelatin	0.947	18.3%	5%	0.993	18%	0.4%	0.948	21%	2.2%
All without Latex	All	0.938	17.7%	7.5%	0.997	1.3%	1.3%	0.995	5.5%	1.5%
All without Latex	Latex	0.972	11.5%	7.5%	0.997	2%	3.5%	0.994	8%	1.6%
All without Modasil	All	0.979	3.9%	2.8%	0.999	0%	4.7%	0.993	4.8%	2.6%
All without Modasil	Modasil	0.999	0%	2.8%	0.999	0%	4.7%	0.993	3.5%	2.7%
All without WoodGlue	All	0.960	10%	4.5%	0.997	1.3%	1.3%	0.995	3.4%	3.5%
All without WoodGlue	WoodGlue	0.941	14.2%	4.5%	0.997	1%	2.2%	0.998	1.5%	3.5%

TABLE III: Finding the best configuration for Siamese using validation set Identix’09.

Nr	Net	Base LR	Features	Iterations	AUC
1	Siamese	10^{-7}	50	7500	0.530
2	Siamese	10^{-4}	50	7500	0.947
3	Siamese	10^{-4}	50	5000	0.958
4	Siamese	10^{-4}	25	5000	0.970

configurations were also tested on a separate validation set (Biometrika’11). For CaffeNet the best performance on the validation set was obtained at a base LR of 10^{-5} (and 10^{-4} for GoogLeNet, respectively) and both nets used 2 output features indicating spoof/live probabilities. Snapshots were taken after 5,000 iterations, however this parameter turned out to have less impact than an inappropriate base LR. Note, that obtained learning rates were lower than in original configuration due to starting from pretrained versions from ImageNet (see Sct. III-A).

D. Performance of CNNs for PAD

We tested material robustness on the Biometrika’13 database and with overall classification accuracy of 93.1%



(a) Biometrika’13 (b) Italdata’13 (c) Sagem’11

Fig. 4: Gelatin Spoof samples recorded with different optical sensors (from LivDet11 [16] and LivDet13 [6]).

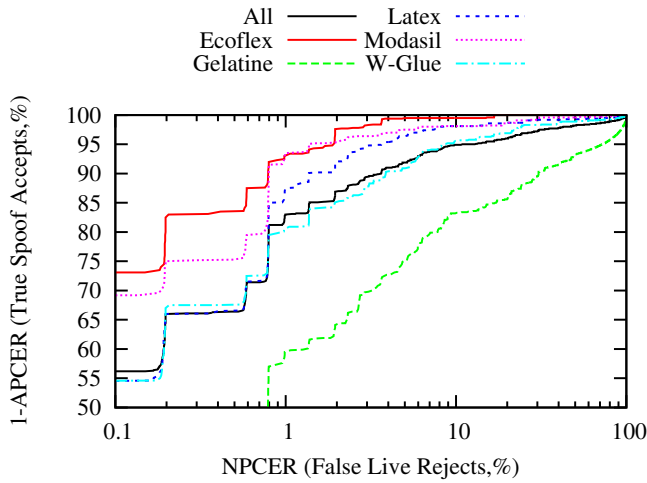
TABLE IV: Interoperability Evaluation (4 S = Biometrika’13, Italdata’13, Sagem’11, Digital’11; 3 S = 4 S w/o Digital’11).

Train	Test	CaffeNet			GoogLeNet		
		AUC	APCER	NPCER	AUC	APCER	NPCER
3 S	3 Sensors	0.814	41%	8.5%	0.953	18%	7.3%
4 S	4 Sensors	0.809	55.9%	3.9%	0.911	25.8%	11%
3 S	Biometrika’13	0.718	94.5%	5%	0.996	2.5%	3.5%
4 S	Biometrika’13	0.610	96.5%	5.5%	0.991	2%	13.5%
4 S	Italdata’13	0.972	4.5%	9.5%	0.932	8%	30.5%
4 S	Sagem’11	0.910	54.5%	1.0%	0.922	41%	0.5%
4 S	Digital’11	0.961	46.5%	0.5%	0.932	51.5%	0.5%

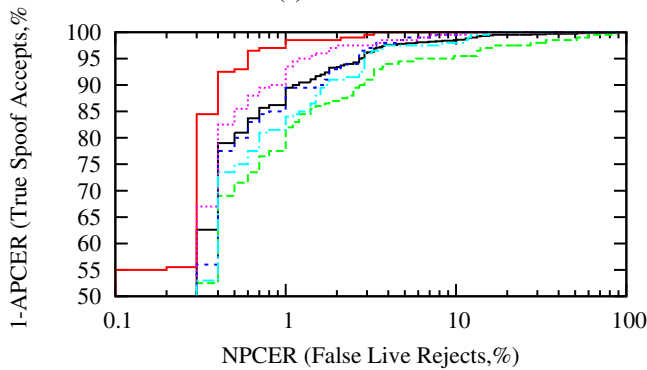
(Siamese), 96.5% (CaffeNet) and 96.6% (GoogLeNet) in the 2-class problem, all three tested networks achieved high performance in experimental tests, see Fig. 5. Table II lists experimental results in detail. For the baseline performance on different materials, it was interesting to see that there was a clear observable trend for a relative order of materials (Ecoflex, Modasil, Latex, WoodGlue, Gelatine) almost over the entire ROC. Interestingly, CaffeNet delivers very high EER performance while weakening a bit at the lower end of the ROC. With this regards GoogLeNet performed best and surprisingly even with relatively few images for finetuning (2,000 images for “All” instead of more than 2 million pairs for Siamese). This justified assumptions that nets adapt quite well to new representative samples limiting the additional benefit of pairwise-combinations.

E. Robustness to Unknown Materials


The test on unseen materials revealed, that AUC performance on testing all materials stayed the same or even slightly improved across different networks in 12/15 cases. However, tests on the left-out material performance stayed in the insignificant range of -3.6% to $+4.6\%$ of the original AUC values for all nets. It is worth mentioning, that obtained



(a) Siamese




(b) CaffeNet




(c) GoogLeNet

Fig. 5: Biometrika'13 ROCs for different nets (Training: All materials, Testing: specific material).

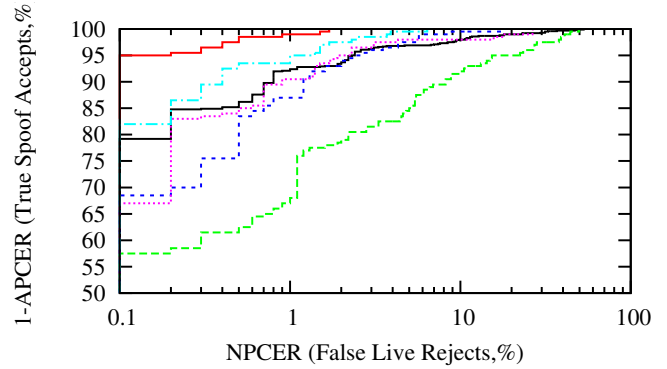
APCER and NPCER were further apart (comp. 21% APCER vs. 2.2% NPCER for Gelatin on GoogLeNet) despite the protocol trying to find a balanced threshold. When looking at ROC performance in Fig. 6 curves are more scattered and there is quite a drastic impact at low NPCER for GoogLeNet pushing its performance slightly below CaffeNet for this task (better or equal performance of CaffeNet in 4 of 5 cases).



(a) Siamese



(b) CaffeNet



(c) GoogLeNet

Fig. 6: Biometrika'13 ROCs for unseen materials (Training: All without specific material, Testing: specific material).

F. Sensor Interoperability

Finally, we tested sensor interoperability training networks with a single spoof-material for multiple sensors and testing on both the unseen sensor and all sensors, see Table IV. With AUCs as low as 0.814 for CaffeNet and 0.953 for GoogLeNet for the 3 sensors case and 0.809 and 0.911 for the 4 sensor case, respectively, it is evident that networks are less able to adapt to changes in sensors and raising this interoperability

issue. A further fact stressing the difficulty in this scenario is the limited number of training images that are generally available for the cross-sensor case (comp. Table I). Especially Biometrika 13 performance was much degraded in presence of other less challenging data sets and in particular the inability to detect attack presentations (high APCER) was evident for the trained CaffeNet network.

V. CONCLUSION AND FUTURE WORK

Among the 3 tested CNNs for PAD, GoogLeNet showed best performance learning fingerprint spoof materials with 4.3% APCER and 2.5% NPCER on Biometrika'13, closely followed by CaffeNet and Siamese - especially with regards to low NPCER. The successful ability to adapt to PAD problems using pre-training on ImageNet was shown, illustrating optimal parameters for the adaption. All networks exhibit high robustness to unseen materials (-3.6% to +4.6% AUC deviation). Sensor interoperability had a more pronounced effect illustrating a slight advantage of CaffeNet over GoogLeNet. Compared to the other two nets, Siamese' performance was inferior, yet it should be noted that for reasons of fairness the experiment did not consider live-live pairs of the same identity only, which could have boosted rates and is subject to future work. Further, it is planned to further customize CNNs with regards to processing speed.

ACKNOWLEDGMENT

The authors would like to thank Dr. Richard Souvenir and Junjie Shan for their support.

This work has been supported by the Modentity project and received funding from the Austrian Security Research Programme KIRAS (www.kiras.at), an initiative of the Federal Ministry for Transport, Innovation and Technology, Austria (www.bmvit.gv.at), under grant agreement no 845495.

REFERENCES

[1] Y. Bengio. *Statistical Language and Speech Processing (SLSP)*, chapter Deep Learning of Representations: Looking Forward, pages 1–37. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[2] I. Chingovska, A. Rabello dos Anjos, and S. Marcel. Biometrics Evaluation Under Spoofing Attacks. *IEEE Trans. Inf. Forensics and Sec.*, 9(12):2264–2276, Dec 2014.

[3] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *Proc. CVPR*, pages 539–546, June 2005.

[4] N. Frassetto, R. Nogueira, R. Lotufo, and R. Machado. Evaluating Software-based Fingerprint Liveness Detection using Convolutional Networks and Local Binary Patterns. *Proc. IEEE BIOMS Workshop*, pages 22–29, 2014.

[5] J. Galbally, J. Fierrez, J. Ortega-Garcia, and R. Cappelli. Fingerprint anti-spoofing in biometric systems. In S. Marcel, M. S. Nixon, and S. Z. Li, editors, *Handbook of Biometric Anti-Spoofing*, pages 35–64. Springer, 2014.

[6] 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. Int'l. Conf. on Biometrics*, pages 1–6. IEEE, 2013.

[7] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira et al., editors, *Advances Neural Information Processing Systems 25*, pages 1097–1105. Curran, 2012.

[8] E. Marasco and A. Ross. A survey on antispoofing schemes for fingerprint recognition systems. *ACM Comput. Surv.*, 47(2):28:1–28:36, Nov. 2014.

[9] E. Marasco and C. Sansone. On the Robustness of Fingerprint Liveness Detection Algorithms against New Materials used for Spoofing. *Proc. Biosignals*, pages 1–9, 2011.

[10] G. Marcialis, A. Lewicke, B. Tan, P. Coli, D. Grimberg, A. Congiu, A. Tidu, F. Roli, and S. Schuckers. First international fingerprint liveness detection competition - LivDet 2009. In *Proc. ICIAAP*, pages 12–23, Sept. 2009.

[11] D. Menotti, G. Chiachia, A. A. Pinto, S. Robson, H. Pedrini, F. Xavier, and A. Rocha. Deep Representations for Iris, Face, and Fingerprint Spoofing Detection. *IEEE Transactions on Information Forensics and Security*, 10(4):864–879, 2015.

[12] A. Rattani and A. Ross. Automatic adaptation of fingerprint liveness detector to new spoof materials. In *Proc. Int'l Joint Conf. Biometrics*, pages 1–8, Sept 2014.

[13] C. Sousedik and C. Busch. Presentation attack detection methods for fingerprint recognition systems: a survey. *Biometrics, IET*, 3(4):219–233, 2014.

[14] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR 2015*, 2015.

[15] P. Wild, P. Radu, L. Chen, and J. Ferryman. Robust multimodal face and fingerprint fusion in the presence of spoofing attacks. *Pattern Recognition*, 50(2):17–25, 2016.

[16] D. Yambay, L. Ghiani, P. Denti, G. Marcialis, F. Roli, and S. Schuckers. LivDet 2011 - fingerprint liveness detection competition 2011. In *Proc. Int'l. Conf. on Biometrics (ICB)*, pages 208–215, 2012.