Morphing Detection Using a General-Purpose Face Recognition System

Lukasz Wandzik*, Gerald Kaeding†, Raul Vicente Garcia‡
Security Technology Department
Fraunhofer Institute for Production Systems and Design Technology IPK
10587 Berlin, Germany
Email: *lukasz.wandzik@ipk.fraunhofer.de, †gerald.kaeding@ipk.fraunhofer.de, ‡raul.vicente@ipk.fraunhofer.de

***Abstract*—Image morphing has proven to be very successful at deceiving facial recognition systems. Such a vulnerability can be critical when exploited in an automatic border control scenario. Recent works on this topic rely on dedicated algorithms which require additional software modules deployed alongside an existing facial recognition system. In this work, we address the problem of morphing detection by using state-of-the-art facial recognition algorithms based on hand-crafted features and deep convolutional neural networks. We show that a general-purpose face recognition system combined with a simple linear classifier can be successfully used as a morphing detector. The proposed method reuses an existing feature extraction pipeline instead of introducing additional modules. It requires neither fine-tuning nor modifications to the existing recognition system and can be trained using only a small dataset. The proposed approach achieves state-of-the-art performance on our morphing datasets using a 5-fold cross-validation.**

***Index Terms*—face recognition, biometric anti-spoofing, face morphing, deep learning**

I. Introduction

The detection of biometric counterfeits, commonly known as anti-spoofing, is a very active field of research. In recent years a variety of methods have been proposed for protecting facial recognition systems [1], [2]. In this work, we consider the so-called *morphing attack* using machine-readable travel documents (eMRTDs) [3]. It can be performed by enrolling with a tampered image obtained from morphing the face of a legitimate document owner (accomplice) with the face of a reasonably similar looking impostor. When successful, the attacked facial recognition system positively matches the tampered template, stored in the eMRTD, with the live image of an impostor.

Although some research has been done on morphing detection using dedicated algorithms [4]–[7], there is no work addressing the capabilities of current face recognition algorithms for explicit protection against morphing. The advantage of this approach is that it can be deployed with minimal effort by reusing precomputed high-level features from face recognition or verification tasks.

We examine four publicly available face recognition methods that achieve state-of-the-art performance on the challenging Labeled Faces in the Wild (LFW) [8] dataset. Three of them are based on modern deep neural network architectures and one uses hand-crafted image features in a multi-scale approach.

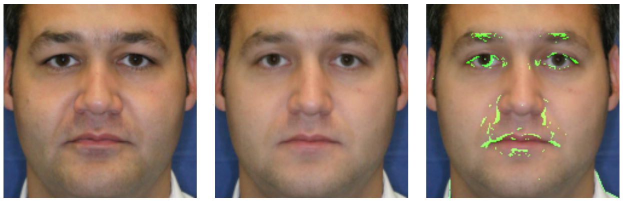

(a) Original image (b) Morphed image (c) Morphing artifacts

Fig. 1. Most prominent morphing artifacts reveled by subtracting the original destination image (a) from the morphed image (b). Source: [9].

We make two contributions. First, we show that facial features computed by a general-purpose face recognition system can be reused for morphing detection. Second, we propose a combined system for morphing detection and face verification which achieves high performance on our own automatically generated morphing dataset.

II. Related work

The threat of a morphing attack in a border control scenario using eMRTD and commercial face recognition software was first identified by Ferrara et al. [3]. Several studies have been done on the vulnerability of modern face recognition methods to morphing attacks [10]–[12].

More recent publications [5], [13] have addressed the specific problem of morphing detection and proposed methods for both the generation of morphs and for the detection of manipulation traces in images. Asaad et al. [7] showed that Local Binary Patterns (LBP) can be successfully applied to morphing detection. A similar approach was proposed by Raghavendra et al. [5] using a Support Vector Machine (SVM). However, there is no information available about the face verification performance of those methods. Deep learning based morphing detectors have been proposed independently by Seibold et al. [6] and Raghavendra et al. [4]. However, they require fine-tuning of the neural network, a fairly large dataset and special hardware to be deployed. Alternative approaches based on image degradation and Benford features extracted from quantized DCT coefficients of JPEG-compressed morphs were proposed by Neubert [14] and Makrushin et al. [13] respectively.

ISBN 978-90-827970-1-5 © EURASIP 2018

TABLE I
PERFORMANCE OF FOUR STATE-OF-THE-ART FACE RECOGNITION AND VERIFICATION METHODS ON THE LFW DATASET UNDER THE UNRESTRICTED PROTOCOL.

LFW Benchmark	
Method	Accuracy [%]
DLib [15]	99.1
FaceNet [16]	**99.2**
VGG-Face [17]	97.3
High-Dim LBP [18]	95.2

III. METHODS

In this section, we will describe four publicly available feature extraction algorithms which achieve state-of-the-art performance on the challenging LFW dataset. Next, we will shortly describe the classifier used in our experiments. Finally, we will propose an architecture that combines face verification and morphing detection.

A. Feature Extraction

1) FaceNet: FaceNet is a framework for face recognition and clustering originally proposed by Schroff et al. [19]. The implementation used in our experiments is based on the *Inception-ResNet-v1* architecture introduced by Szegedy et al. [16]. Instead of a pure inception model, as described in the original paper, Szegedy et al. combined inception modules and residual connections in one network. The feature extractor expects a 160×160 pixel input image and returns an 1792-dimensional embedding. The network was trained on a subset of the MS-Celeb-1M dataset [20] containing about 4 million images distributed over 51k classes.

2) Dlib: DLib is a machine learning toolkit designed for real-world applications [15]. It contains a face matching module based on the *ResNet-34* architecture introduced by He et al. [21]. The version used in our experiments has only 29 convolutional layers as opposed to 34 in the original network. In addition, the number of filters per layer has been reduced by half and the input layer was resized to 150×150 pixels. The output layer was also resized and returns a 128-dimensional embedding. The network was trained using metric learning on a dataset of about 3 million images and 7485 identities. Whereby a large number of images was acquired using the Internet and nearly a half stemmed from FaceScrub [22] and VGG-Face [17] datasets.

3) VGG-Face: VGG-Face is a face descriptor introduced by Parkhi et al. [17] along with a large-scale dataset of the same name. It is based on the *VGG-Very-Deep-16* convolutional neural network (CNN) as described in [23]. Unlike the previous two networks, the *VGG-Very-Deep-16* has a linear topology with 16 convolutional layers. The input layer of the network expects a 224×224 pixels image, whereas the output layer, which is the last layer before the softmax, returns a feature vector of 4096 dimensions. The network was trained on the VGG-Face dataset which contains about 2.6 million images distributed over 2622 classes. The version used in our experiments was trained as a classifier using a softmax layer.

4) High-Dim LBP: In contrast to the above methods, we also evaluate a non-deep-learning approach based on Local Binary Patterns [24]. Chen et al. [18] introduced a high-dimensional face feature extraction method which achieves state-of-the-art performance on the LFW dataset (Tab. I). The implementation used in our experiments is based on the DLib toolkit and can be summarized in five steps:

1) Rectify the input image based on five facial landmarks (eyes, nose, and mouth corners).
2) Build an image pyramid of the normalized facial image.
3) At each scale, extract fixed-size image patches centered around each landmark.
4) Divide each patch into a grid of 4×4 cells and encode each cell by an LBP descriptor.
5) Finally, concatenate all descriptors to form a high-dimensional feature vector.

This method outputs a 99120-dimensional feature vector and does not require any training or parameter adjustment.

B. Classification

Morphing detection is a binary classification problem. Given an input image, we want to decide whether it was manipulated by morphing with another image or not. We train a Support Vector Machine [25] on the high-level facial features to solve this classification problem. Since our training data has a high dimensionality, mapping it to an even higher dimensional space will not improve the performance [26]. We, therefore, use the linear kernel in all our experiments, which leaves us with only one free parameter to adjust.

C. Combined Approach

As we will show later in the results section, there is an inverse relationship between the performance in morphing detection and face verification. Methods which performed poorly on the morphing detection task were very good at face verification and vice versa. Therefore, we make the assumption that those tasks might complement each other and propose a combined architecture. Figure 2 shows one possible implementation of this approach. We first compute the facial features by passing reference and query images to one of the four described methods. Then, we use the output vectors to compute the Euclidean distance for the face verification task.

TABLE II
OVERVIEW OF FOUR FACE FEATURE EXTRACTION METHODS, THEIR ARCHITECTURES AND INPUT/OUTPUT SIZES.

Method	Architecture	Input	Output
DLib [15]	ResNet-29	150 × 150	128
FaceNet [16]	Inception-ResNet-v1	160 × 160	1792
VGG-Face [17]	VGG-16	224 × 224	4096
High-Dim LBP [18]	Multi-Scale LBP	40 × 40*	99120

*Size of the image patch centered around each landmark.

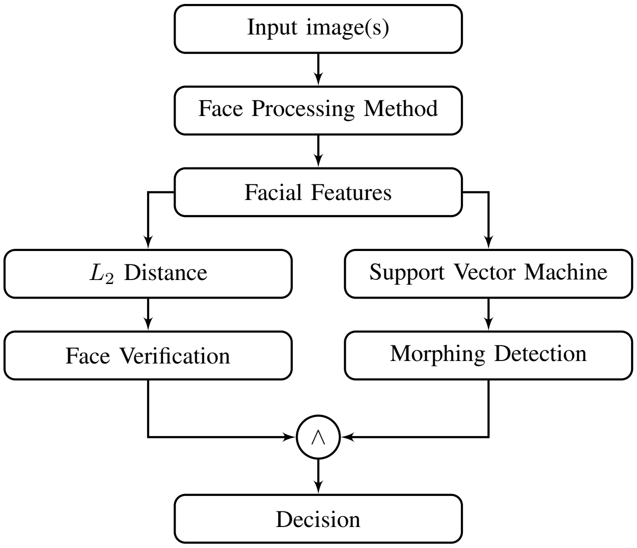

Fig. 2. The figure shows our approach for combining face verification and morphing detection in one system. *Face Processing Method* can be replaced by one of the introduced methods: FaceNet [19], DLib [15], VG-Face [17] or High-Dim LBP [18].

In addition, the reference image vector is passed to the SVM which outputs the decision for the morphing detection task. The binary decisions from both these tasks are combined to produce the final decision, either an acceptance or rejection of the individual.

IV. Experiments

In this section, we will describe the preparation procedure of our dataset, as well as the methodology and metrics used for evaluation.

A. Dataset

We use the *Multi-PIE* [9] dataset as the basis for our own automatically generated facial morphs. The exact procedure will be described in the next section. We chose the *Multi-PIE* dataset because the conditions under which its data was collected were similar to those prescribed for eMRTDs. The original dataset contains images of 337 subjects from four different sessions - including variations in viewpoint, illumination, and expression. In this work, we only consider images with a neutral pose, facial expression, and frontal illumination.

We generate the facial morphs automatically by using the splicing-morph method described in [11]. It relies on the DLib shape predictor for localizing facial landmarks and Poisson blending for the refinement of boundary regions. The morphs are generated with a constant blending factor of 0.5. We only use images from session *one* and restrict ourselves to subjects that appear in at least one more session. Furthermore, we only consider subjects not wearing glasses and those belonging to the same gender, in order to avoid unnatural image artifacts.

The evaluation dataset is organized into 5 train-test splits such that undesired correlations between train and test sets are avoided. The generation process can be summarized in the following steps.

1) Divide the dataset D into 5 disjoint sets

$$D = \{S_i | S_i \cap S_j = \varnothing\}, \forall\, i, j \in \langle 1, 5 \rangle \wedge j \neq i$$

2) Select and morph identity pairs (I_i, I_j) within each set such that $i \neq j$
3) Create test set Test$_i$ such that Test$_i = S_i$
4) Create training set Train$_i$ such that

$$\text{Train}_i = \{S_j | 1 \leq j \leq 5 \wedge j \neq i\}$$

Additionally, we create a dataset for the combined approach as described in the *Methods* section. We use the remaining images from session *two* to *four* for this purpose and replace the facial morphs with the corresponding destination image. This dataset is aimed to simulate the query image in a genuine access control scenario.

B. Methodology

We conduct four different experiments, morphing detection, its combination with face verification and repeat those two for features from lower layers of each CNN. Those layers represent lower level image features, which might be better suited for detecting morphing artifacts. They are extracted from the conv7, Mixed_6a and conv_5_2 layer of DLib, FaceNet and VGG-Face respectively.

C. Performance Metrics

We evaluate the proposed approaches in the context of morphing detection by using the following evaluation metrics.

$$\text{FAR} := \frac{|\text{Accepted morphs}|}{|\text{All morphed images}|} \tag{1}$$

The False Acceptance Rate (FAR) is synonymous to the Morph Acceptance Rate (MAR) and therefore the most important measure from the security point of view.

$$\text{FRR} := \frac{|\text{Rejected genuine individuals}|}{|\text{All genuine individuals}|} \tag{2}$$

The False Rejection Rate (FRR), on the other hand, measures the usability of a biometric system.

$$\text{ACC} := \frac{|\text{Correctly classified images}|}{|\text{All classified images}|} \tag{3}$$

The last measure is the Accuracy (ACC), which we mainly use for comparing face verification performance with the LFW benchmark.

V. Results

We present our results for morphing detection and its combination with face verification separately. All results were obtained using a 5-fold cross-validation. We report the average performance and standard deviation for each method. The evaluation data was preprocessed by performing an L_2 normalization of the feature vectors. We set the penalty parameter C of the SVM to 1 since other values did not improve the overall results.

ISBN 978-90-827970-1-5 © EURASIP 2018

A. Morphing Detection

Table III summarizes the results for morphing detection while reusing the features originally computed for face verification. The best result was achieved by the LBP descriptor, which is the worse method in terms of face verification performance (see Table I for comparison). In the next experiment, we only evaluated the CNN based methods. Table IV shows the results obtained by extracting features from lower layers of the particular network. The best results were achieved by the VGG-Face descriptor. Although we observe an improvement over the results in Table III, and the number of features is now comparable to the high-dimensional LBP approach, there is still a large performance gap between these two methods.

B. Face Verification and Morphing Detection

In this experiment, we combined the decisions from both tasks, and present our findings for a hybrid system. In order to perform face verification on the MultiPIE dataset, we first have to find an optimal threshold for each method. Table V summarized the results and thresholds obtained using the original dataset. Table VI and VII show the results for two different image sets. The first set (session 1) covers a pure morphing detection task, whereas the second set (session 2 to 4) simulates a real access control scenario in the presence of a morphing attack. The best results were once again achieved by the high-dimensional LBP descriptor, but the performance of the CNN based methods improved significantly on both sets.

We also evaluated our hybrid approach by using features from lower layers of each CNN and combine the results with face verification decisions. These experiments are carried out only on the CNN based methods. The findings are once again presented for two different image sets, from session 1 (Tab. VIII) and sessions 2 to 4 (Tab. IX). The best results are now achieved by FaceNet, closely followed by the remaining methods, DLib and VGG-Face.

VI. Conclusions

We showed that a general-purpose face recognition algorithm, as deployed in many biometric systems, can be used to detect morphing attacks. The advantage of this approach is that it can be deployed with minimal effort by reusing precomputed high-level features. The combination of face verification and morphing detection achieved even better results, especially for the deep learning based methods, which can be attributed to their better face verification performance.

TABLE III
Results using a 5-fold cross-validation with linear kernel and penalty term $C = 1$. The features were obtained from the output layer of the respective method.

Morphing Detection Performance			
Method	ACC [%]	FAR [%]	FRR [%]
DLib [15]	77.2 ± 5.2	23.0 ± 5.7	20.9 ± 5.8
FaceNet [16]	79.5 ± 9.1	20.6 ± 10.2	18.3 ± 7.0
VGG-Face [17]	83.2 ± 3.5	16.0 ± 4.1	26.1 ± 7.3
High-Dim LBP [18]	$\mathbf{99.2 \pm 0.8}$	$\mathbf{0.3 \pm 0.7}$	$\mathbf{7.0 \pm 4.4}$

TABLE IV
Results using a 5-fold cross-validation with linear kernel and penalty term $C = 1$. The features were obtained from the conv7, Mixed_6a and conv_5_2 layer respectively.

Morphing Detection Performance				
Method	ACC [%]	FAR [%]	FRR [%]	$\mathrm{dim}(\cdot)$
DLib [15]	87.3 ± 4.6	12.4 ± 5.3	15.7 ± 8.1	~ 40k
FaceNet [16]	90.3 ± 1.3	9.3 ± 1.5	13.0 ± 6.1	~ 25k
VGG-Face [17]	$\mathbf{93.6 \pm 1.6}$	$\mathbf{6.1 \pm 1.6}$	$\mathbf{10.4 \pm 2.1}$	~ 100k

Although convolutional neural networks performed significantly worse, in comparison to the LBP based method, they are still able to provide some protection against morphing attacks. Moreover, by extracting features from lower layers the results could be further improved at the expense of higher data dimensionality. On the other hand, the LBP descriptor achieved impressive results and almost met the requirements for biometric systems prescribed by Frontex [27]. It also achieved respectable performance on the face verification task, and therefore constitute a good trade-off between morphing protection and face verification performance.

Acknowledgment

This work has been partially funded by the German Federal Ministry of Education and Research (BMBF) under contract number FKZ: 16KIS 0512.

References

[1] J. Galbally, S. Marcel, and J. Fierrez, “Biometric antispoofing methods: A survey in face recognition,” *IEEE Access*, vol. 2, pp. 1530–1552, 2014.

TABLE V
Performance of the four face recognition methods on the MultiPIE dataset and their optimal thresholds. The false rejection rate was calculated at 0.1% false acceptance rate.

MultiPIE Dataset			
Method	FRR [%]	Accuracy [%]	Threshold$_{L_2}$
DLib [15]	0.62	99.66	0.29
FaceNet [16]	$\mathbf{0.00}$	$\mathbf{99.98}$	0.75
VGG-Face [17]	6.70	96.39	0.79
High-Dim LBP [18]	8.18	95.79	0.52

TABLE VI
Results using a 5-fold cross-validation with linear kernel and penalty term $C = 1$. Evaluated on images from session 1.

Face Verification & Morphing Detection Performance			
Method	ACC [%]	FAR [%]	FRR [%]
DLib [15]	96.3 ± 1.6	2.2 ± 2.1	20.9 ± 5.8
FaceNet [16]	97.6 ± 0.7	1.1 ± 1.2	18.3 ± 7.0
VGG-Face [17]	95.3 ± 1.3	2.9 ± 1.6	26.1 ± 7.3
High-Dim LBP [18]	$\mathbf{99.4 \pm 0.4}$	$\mathbf{0.1 \pm 0.2}$	$\mathbf{7.0 \pm 4.4}$

TABLE VII
RESULTS USING A 5-FOLD CROSS-VALIDATION WITH LINEAR KERNEL AND PENALTY TERM $C = 1$. EVALUATED ON IMAGES FROM SESSIONS 2 TO 4.

Face Verification & Morphing Detection Performance			
Method	ACC [%]	FAR [%]	FRR [%]
DLib [15]	96.5 ± 1.2	1.9 ± 1.7	22.6 ± 7.5
FaceNet [16]	98.1 ± 0.6	0.3 ± 0.4	20.9 ± 8.4
VGG-Face [17]	95.8 ± 0.7	1.7 ± 0.9	33.9 ± 11.5
High-Dim LBP [18]	$\mathbf{98.8 \pm 0.5}$	$\mathbf{0.0 \pm 0.0}$	$\mathbf{15.7 \pm 6.5}$

TABLE VIII
RESULTS USING A 5-FOLD CROSS-VALIDATION WITH LINEAR KERNEL AND PENALTY TERM $C = 1$. EVALUATED ON IMAGES FROM SESSION 1.

Face Verification & Morphing Detection Performance				
Method	ACC [%]	FAR [%]	FRR [%]	$\dim(\cdot)$
DLib [15]	98.0 ± 0.4	0.8 ± 0.5	15.7 ± 8.1	~ 40k
FaceNet [16]	$\mathbf{98.8 \pm 0.5}$	$\mathbf{0.2 \pm 0.3}$	13.0 ± 6.1	~ 25k
VGG-Face [17]	98.2 ± 0.6	1.0 ± 0.5	$\mathbf{10.4 \pm 2.1}$	~ 100k

TABLE IX
RESULTS USING A 5-FOLD CROSS-VALIDATION WITH LINEAR KERNEL AND PENALTY TERM $C = 1$. EVALUATED ON IMAGES FROM SESSIONS 2 TO 4

Face Verification & Morphing Detection Performance				
Method	ACC [%]	FAR [%]	FRR [%]	$\dim(\cdot)$
DLib [15]	98.0 ± 0.6	0.6 ± 0.4	18.3 ± 7.0	~ 40k
FaceNet [16]	$\mathbf{98.8 \pm 0.7}$	$\mathbf{0.0 \pm 0.0}$	$\mathbf{14.8 \pm 8.1}$	~ 25k
VGG-Face [17]	97.9 ± 1.1	0.5 ± 0.5	20.9 ± 8.9	~ 100k

[2] A. Hadid, N. Evans, S. Marcel, and J. Fierrez, “Biometrics systems under spoofing attack: An evaluation methodology and lessons learned,” *IEEE Signal Processing Magazine*, vol. 32, no. 5, pp. 20–30, 2015.

[3] M. Ferrara, A. Franco, and D. Maltoni, “The magic passport,” in *2014 IEEE International Joint Conference on Biometrics (IJCB)*, 2014, pp. 1–7.

[4] R. Raghavendra, K. B. Raja, S. Venkatesh, and C. Busch, “Transferable deep-cnn features for detecting digital and print-scanned morphed face images,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops, Honolulu, HI, USA, July 21-26, 2017*, 2017, pp. 1822–1830.

[5] R. Raghavendra, K. B. Raja, and C. Busch, “Detecting morphed face images,” in *8th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2016, Niagara Falls, NY, USA, September 6-9, 2016*, 2016, pp. 1–7.

[6] C. Seibold, W. Samek, A. Hilsmann, and P. Eisert, “Detection of face morphing attacks by deep learning,” in *Digital Forensics and Watermarking: 16th International Workshop , IWDW 2017*, ser. Lecture Notes in Computer Science, vol. 10431. Springer International Publishing, 2017, pp. 107–120.

[7] A. Asaad and S. Jassim, “Topological data analysis for image tampering detection,” in *Digital Forensics and Watermarking*, C. Kraetzer, Y.-Q. Shi, J. Dittmann, and H. J. Kim, Eds. Cham: Springer International Publishing, 2017, pp. 136–146.

[8] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, “Labeled faces in the wild: A database for studying face recognition in unconstrained environments,” University of Massachusetts, Amherst, Tech. Rep. 07-49, October 2007.

[9] Ralph Gross, Iain Matthews, Jeffrey Cohn, Takeo Kanade, Simon Baker, “Multi-pie,” in *Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition*. IEEE Computer Society, 2008. [Online]. Available: https://www.microsoft.com/en-us/research/publication/multi-pie/

[10] A. Mohammadi, S. Bhattacharjee, and S. Marcel, “Deeply vulnerable: a study of the robustness of face recognition to presentation attacks,” *IET Biometrics*, vol. 7, no. 1, pp. 15–26, 2018.

[11] L. Wandzik, R. V. Garcia, G. Kaeding, and X. Chen, *”CNNs Under Attack: On the Vulnerability of Deep Neural Networks Based Face Recognition to Image Morphing”*. ”Cham”: ”Springer International Publishing”, ”2017”, pp. ”121–135”.

[12] U. Scherhag, A. Nautsch, C. Rathgeb, M. Gomez-Barrero, R. N. J. Veldhuis, L. J. Spreeuwers, M. Schils, D. Maltoni, P. Grother, S. Marcel, R. Breithaupt, R. Raghavendra, and C. Busch, “Biometric systems under morphing attacks: Assessment of morphing techniques and vulnerability reporting,” in *International Conference of the Biometrics Special Interest Group, BIOSIG 2017, Darmstadt, Germany, September 20-22, 2017*, 2017, pp. 1–7.

[13] A. Makrushin, T. Neubert, and J. Dittmann, “Automatic generation and detection of visually faultless facial morphs,” in *Proceedings of the 12th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 6: VISAPP, (VISIGRAPP 2017)*, INSTICC. ScitePress, 2017, pp. 39–50.

[14] T. Neubert, “Face morphing detection: An approach based on image degradation analysis,” in *Digital Forensics and Watermarking*, C. Kraetzer, Y.-Q. Shi, J. Dittmann, and H. J. Kim, Eds. Cham: Springer International Publishing, 2017, pp. 93–106.

[15] D. E. King, “Dlib-ml: A machine learning toolkit,” *Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 2009.

[16] C. Szegedy, S. Ioffe, and V. Vanhoucke, “Inception-v4, inception-resnet and the impact of residual connections on learning,” *CoRR*, vol. abs/1602.07261, 2016.

[17] O. M. Parkhi, A. Vedaldi, and A. Zisserman, “Deep face recognition,” in *British Machine Vision Conference*, 2015.

[18] D. Chen, X. Cao, F. Wen, and J. Sun, “Blessing of dimensionality: High-dimensional feature and its efficient compression for face verification,” in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, June 2013, pp. 3025–3032.

[19] F. Schroff, D. Kalenichenko, and J. Philbin, “Facenet: A unified embedding for face recognition and clustering,” *CoRR*, vol. abs/1503.03832, 2015.

[20] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao, “MS-Celeb-1M: A dataset and benchmark for large scale face recognition,” in *European Conference on Computer Vision*. Springer, 2016.

[21] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *CoRR*, vol. abs/1512.03385, 2015.

[22] H. Ng and S. Winkler, “A data-driven approach to cleaning large face datasets,” in *2014 IEEE International Conference on Image Processing, ICIP 2014, Paris, France, October 27-30, 2014*, 2014, pp. 343–347.

[23] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *CoRR*, vol. abs/1409.1556, 2014.

[24] T. Ojala, M. Pietikäinen, and H. D. Harwood, “Performance evaluation of texture measures with classification based on kullback discrimination of distributions,” in *Proceedings of the 12th IAPR International Conference on Pattern Recognition (ICPR 1994)*, vol. 1, 1994, pp. 582–585.

[25] C. Cortes and V. Vapnik, “Support-vector networks,” *Mach. Learn.*, vol. 20, no. 3, pp. 273–297, Sep. 1995.

[26] C. wei Hsu, C. chung Chang, and C. jen Lin, “A practical guide to support vector classification,” 2010.

[27] E. A. f. t. M. o. O. C. a. t. E. B. European Union, *Best Practice Technical Guidelines for Automated Border Control (ABC) Systems*. FRONTEX, 2015.