

Transferable Deep-CNN features for detecting digital and print-scanned morphed face images

R. Raghavendra, Kiran B. Raja, Sushma Venkatesh, Christoph Busch
Norwegian Biometrics Laboratory, NTNU - Gjøvik, Norway

{raghavendra.ramachandra; kiran.raja; sushma.venkatesh; christoph.busch} @ntnu.no

Abstract

Face biometrics is widely used in various applications including border control and facilitating the verification of travellers' identity claim with respect to his electronic passport (ePass). As in most countries, passports are issued to a citizen based on the submitted photo which allows the applicant to provide a morphed face photo to conceal his identity during the application process. In this work, we propose a novel approach leveraging the transferable features from a pre-trained Deep Convolutional Neural Networks (D-CNN) to detect both digital and print-scanned morphed face image. Thus, the proposed approach is based on the feature level fusion of the first fully connected layers of two D-CNN (VGG19 and AlexNet) that are specifically fine-tuned using the morphed face image database. The proposed method is extensively evaluated on the newly constructed database with both digital and print-scanned morphed face images corresponding to bona fide and morphed data reflecting a real-life scenario. The obtained results consistently demonstrate improved detection performance of the proposed scheme over previously proposed methods on both the digital and the print-scanned morphed face image database.

1. Introduction

Face biometric systems are widely deployed in various access control applications that include law enforcement, physical and logical access control, surveillance, national identity card schemes, border control, to name a few. The non-intrusive capture and the low-cost operational setting has further elevated the benefits of face recognition systems. The most prevalent application of face biometrics is verification of a traveller's identity claim with respect to his electronic passport (ePass) that can facilitate the border control using Automatic Border Control (ABC) systems [6]. Following the specifications from the International Civil Aviation Organisation (ICAO) regarding the mandatory inclu-




Figure 1: Example of face morphing

sion of a facial reference image in the ePass, more than 800 million ePassports have been issued over the last twelve years.

Although the ePass is expected to establish a strong link between the facial image stored securely in the passport and the owner of the document, a recent study [4] has revealed the vulnerability of the ePass enrolment process with respect to face morphing attacks. Freely accessible morphing techniques allow any attacker to combine two different images from two data subjects to get a single (composite) image that resembles both face images that were used to generate the morphed image. Considering the fact that most countries to date still accept printed photos (by physically providing the printed passport photo) and other countries (e.g. New Zealand, Estonia, Ireland) even allow uploading a digital image through a web-portal, the attacker has a large window of opportunity to present the morphed face image instead of the actual face. This opportunity enables individuals with a criminal background to conceal their identity and to lower the risk of being scrutinized at border control as the morphed face matches equally to probe images stemming from both of the constituent subjects in the morphed image. Moreover one has to anticipate that a morphed face image will not be detected by a human observer, which further increases the complexity of the problem. The problem scale was well illustrated by an extensive study where

morphed images were presented to human observers including trained experts, who failed to detect morphed images in general [12, 4, 5]. Further two recent works investigated the problem from the perspective of a biometric system, and established the fact that for a morphed image even a commercial face recognition system provides similarity scores above threshold for probe images from both subjects [10], [14].

These factors underline the importance of detecting morphed face image to avoid the illegitimate passport holder from exploiting this vulnerability and passing through an ePass based automated gate system. While the severity of the problem is high, the complexity of creating a morphed face image is fairly low due to the large number of freely available morphing software (Example: GIMP and GAP tools) that can generate a high quality morphed facial image with minimal efforts in indicating the facial key points.

Figure 1 shows the example of a morphed face sample (middle face image in Figure 1) that is generated using two different subjects (left and right face images in Figure 1). As noticed from Figure 1, the morphed face can show a strong visual resemblance to both face used to generate the morphed face image.

1.1. Related work

The first work on the facial morphing was introduced in [4]. The vulnerability of the face recognition system was studied on a small set of the morphed face images. The analysis of morphed face identification using human subjects was presented in [12] and [5]. Those works showed the failure of human experts to accurately detect morphed images. The first work on automated morphed face detection was proposed in [10] that explores the utility of the micro-texture pattern using Binarised Statistical Image Features (BSIF) and Support Vector Machine (SVM). The experiments are carried out on 450 morphed face images (digital samples) that are generated manually with the intention to achieve high quality morphed images. The detection method has indicated a reasonable detection performance. It has to be noted that this early work focused on a solution for digital images only.

Considering the real-life ePass application process, morphed face images can be printed and subsequently be scanned again (in the passport application office). In this process for the attack scenario, the re-digitized version of the morphed images will discard the pixel level information aiding an easy approach towards detecting morphed images [10]. The digitally morphed images generated using the face morphing tools can be retouched to enhance the visual quality and remove the ghosting appearance before printing using a photo printer. The applicant will then submit this photo to the passport officer who will then scan the photo with 300dpi resolution [6] that is further used

in the e-passport. Thus, the real-life challenge is to detect the print-scanned version of the digitally morphed face images rather than a digital version of the morphed face image. Further it has to be noted that the process of printing and scanning will further reduce the quality of face sample due to the noise introduced in the print and scan process. Most recently, the first attempt on print-scanned morphed face images is presented in [14] that details the vulnerability of face recognition systems towards print-scanned morphed face images on both Commercial-off-the-shelf face systems from Neurotechnology [9] and also on freely available academic system - Open-Face based on the deep CNN for face recognition [2]. The vulnerability analysis carried out at an operating threshold of FAR = 0.1% (as recommended by FRONTEx for border control operations) has indicated vulnerability with respect to both digital and print-scanned morphed face image. The evaluation of baseline detection methods carried out using micro-texture based features has reported high error rates in detecting the print-scanned morphed face images.

1.2. Our contributions

In this work we explore the transfer learning approach using the Deep Convolutional Neural Networks (D-CNN) to detect both digital and print-scanned versions of morphed face images. To this extent, we explore feature level fusion of the first fully connected layer from the pre-trained VGG-19 [15] and AlexNet [8]. The fused features from the D-CNN networks are further classified using the Probabilistic Collaborative Representation Classifier (P-CRC) [3] to reliably detect a morphed face image. The pre-trained D-CNN (both AlexNet and VGG19) are fine-tuned independently on digital and print-scanned morphed face images. Extensive experiments are carried out using the morphed face database available from [14], but due to the limited number of bona fide samples available (around 104), we further augmented the database to have the relatively equal number of bona fide and morphed images. Thus, the key contributions of this work are outlined below:

- New framework based on feature level fusion of the first fully connected layer from pre-trained D-CNNs such as AlexNet and VGG19 to detect digital and print-scanned morphed face images.
- Extension of the earlier morphed face database from [14] to contain 352 bona fide (104 earlier) and 431 morphed images that correspond to both digital and print-scanned version. Further, to generate the print-scanned version, we have used both photo scanner from HP and office scanner from RICOH scanning at 300dpi.
- Extensive experiments are carried out on the newly constructed database to evaluate the performance of

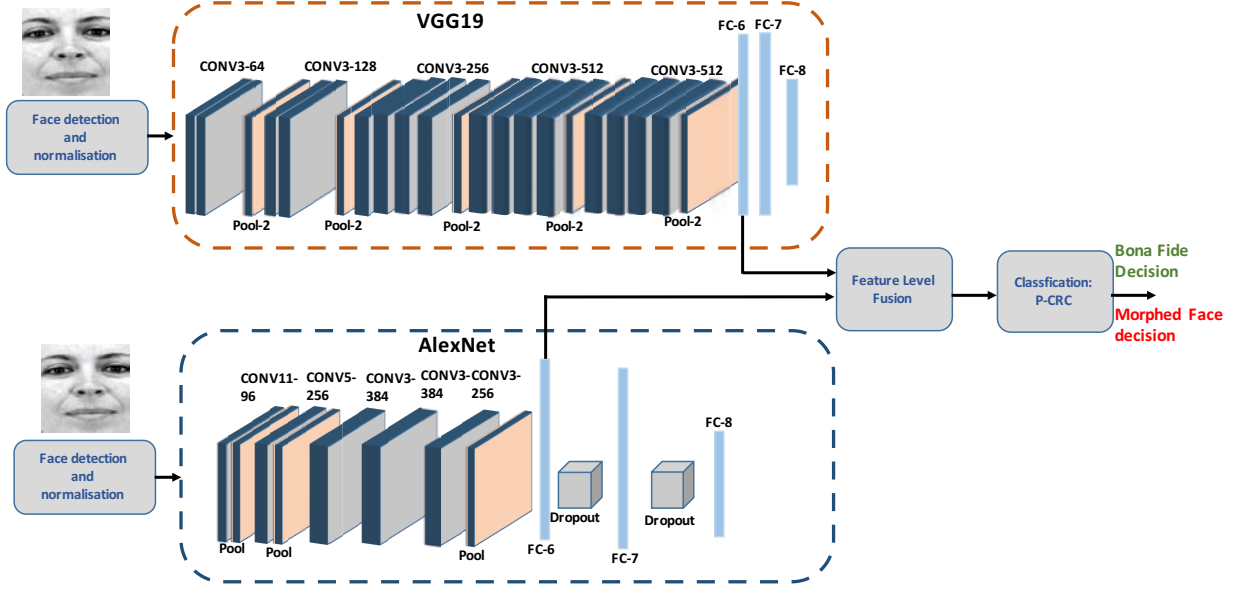


Figure 2: Block diagram of the proposed scheme

the proposed scheme. A benchmark of the proposed scheme with four different state-of-the-art techniques is also presented.

The rest of the paper is organised as follows: Section 2 presents the proposed scheme for reliable morphed face detection, Section 3 discuss the extension of the database, quality control between morphed and bona fide face images and the quantitative performance of the proposed method and it's benchmark with the four different state-of-the-art methods. Section 4 draws the conclusion.

2. Proposed Scheme

Figure 2 shows the block diagram of the proposed scheme based on feature level fusion of the D-CNN networks to detect morphed face images. The underlining idea of the proposed scheme is to fine-tune the pre-trained network for the application of morphed face detection. The proposed scheme is structured using three main functional blocks that include: (1) Pre-processing (2) Feature level fusion of D-CNNs and (3) Classification. The functionality of each block is discussed below.

2.1. Pre-processing

Given the face image, we first carry out face detection using the Viola-Jones algorithm [17]. In the next step, the detected face region is normalized to compensate rotation using the affine transform as introduced in [16] and the image is resized to 227×227 pixels. We are choosing the size

227×227 pixels to make it adaptable to the size of the input layer of the D-CNNs that are also pre-trained with images of size of 227×227 pixels.

2.2. Feature level fusion of D-CNNs

The primary objective of the proposed scheme is to explore the transferable features from the pre-trained D-CNN networks. To this extent, we have employed two popular pre-trained D-CNNs namely VGG19 and AlexNet that are trained on the large-scale ImageNet database. Since these D-CNNs are based on different configurations (or architecture) despite being trained on the same database (ImageNet), the fine-tuning of these two D-CNNs with morphed our face database and the combination of the two networks on the feature level can provide complementary features suitable for detecting morphed face images. In this work, we explore the features from the first fully connected layer (FC-6) from both fine-tuned networks VGG19 and AlexNet. We select features from FC-6 layer to utilise the full-scale features before the drop-out layers employed especially in AlexNet. Then, these features are concatenated to form a single vector before providing this feature vector to the classification block. In this work, the fine-tuning of VGG19 and AlexNet is carried out independently on the training dataset of the morphed face image database (please refer Table 2 for more details on training and testing samples). While fine-tuning the network, we have applied a high value of the learning rate on the last layer when compared to the rest of the layers in the network. The learning parameters used

in this work are as follows: *weight learning rate factor* as 10 and *bias learning rate factor* as 20.

In order to visualize the network features from the fine-tuned VGG19 and AlexNet, we have considered the features from the third Convolution layer (Conv3) from both D-CNNs. Figure 3 (a) shows the Conv3 features from the fine-tuned AlexNet with a convolution filter size of 3×3 and length 384. It is interesting to observe that each convolution filter exhibits the facial features and textural features. Further, the highlighted regions (in different colour) indicate the face specific features that are learned by the fine-tuned AlexNet. A similar observation can also be noted with the fine-tuned VGG19 with a Conv3 layer with a convolutional filter size of 3×3 and a length of 256 as illustrated in the Figure 3 (b). The highlighted regions in the Figure 3 (b) show the face landmark features learned by the fine-tuning of VGG19. Further to quantitatively estimate the complementary information between features from two fine-tuned D-CNN networks, we employed the cross-correlation coef-

ficients (CCC) [13] defined as follows:

$$CC = \frac{\sum_m \sum_n (F_{Alex} - \bar{F}_{Alex})(F_{VGG} - \bar{F}_{VGG})}{\sqrt{\left(\sum_m \sum_n (F_{Alex} - \bar{F}_{Alex})^2\right) \left(\sum_m \sum_n (F_{VGG} - \bar{F}_{VGG})^2\right)}} \quad (1)$$

where \bar{F}_{VGG} and \bar{F}_{Alex} represent mean values.

Table 1: Cross-Correlation Coefficients (CCC) between features of AlexNet and VGG19

Data Type	Feature	CCC
Digital	Train Alex - Train VGG	0.047
	Test Alex - Test VGG	0.045
Print-Scan (HP)	Train Alex - Train VGG	0.024
	Test Alex - Test VGG	0.009
Print-Scan (RICOH)	Train Alex - Train VGG	0.007
	Test Alex - Test VGG	0.012

As indicated in the Table 1, the lower values of the CCC between the features obtained from D-CNNs on both training and testing dataset indicates the complementary nature.

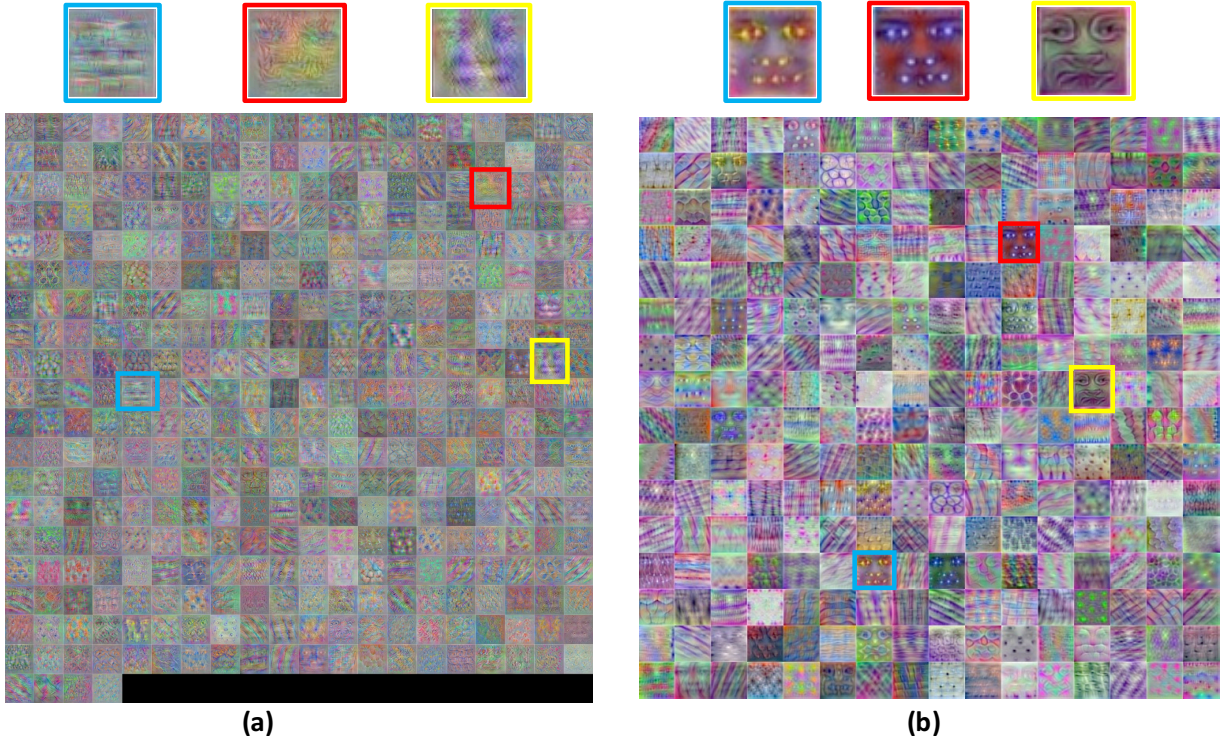


Figure 3: Illustration of the weights from Convolution 3 layer from the fine-tuned (a) AlexNet (b) VGG19; Clear structures of facial features can be observed from the highlighted blocks in different colors.

These facts further motivated us to explore the feature level fusion in the propose scheme.

2.3. Classification

In this work, we employed the Probabilistic Collaborative Representation Classifier (P-CRC), which maximizes the likelihood ratio of a test sample jointly with other classes to perform the classification [3]. Given the training set face images corresponding to both bona fide and morphed face image, the pre-trained D-CNN networks are fine-tuned, and the features from the first fully connected layer (FC-6) are extracted from both AlexNet and VGG19. Let the FC-6 features extracted from AlexNet be F_A and VGG19 be F_V . We then combine the extracted features by concatenating them to form a single feature vector $Tr_F = [F_A || F_V]$ that was used to train the P-CRC. The test face image F_{Te} is then projected independently on the FC-6 layers of AlexNet and VGG19 D-CNN to obtain the corresponding features and let these be F_{Te_A} and F_{Te_V} . These features are then combined using feature concatenation to form a single vector $Te_F = [F_{Te_A} || F_{Te_V}]$ that is used as test feature vector with P-CRC to obtain the morphing detection scores. The P-CRC used in this work utilizes the Regularized Least Square Regression (LSR) on the learned feature vectors versus the probe feature vectors [3] formulated as:

$$\hat{F} = \underset{\alpha}{\operatorname{argmin}} \|Te_F - \mathcal{D}\alpha\|_2^2 + \lambda \|\alpha\|_2^2 \quad (2)$$

where the Te_F is the feature vector of the test image, \mathcal{D} is the learned collaborative subspace dictionary using Tr_F , α is coefficient vector and λ is the regularization parameter. The distance obtained is used as the morphing scores to obtain the morphed face detection performance.

3. Experiments and Results

In this section, we discuss the quantitative results of the proposed scheme, which was evaluated on our morphed face image database. The newly built database is an extension of the previous database [14], which is created using the publicly available face database [11] comprised of 104 subjects. The morphed face images are generated using the freely available software packages such as GNU Image Manipulation Program v2.8 (GIMP) and GIMP Animation Package (GAP) tools [1] with the manual intervention to align the landmark points to achieve high quality morphed images. Then to simulate the real-life scenario of the passport issuance procedure, we generate the print-scanned version of the digitally morphed images. Thus, each of the digitally morphed image is printed using the high-quality photo paper (300 grams.)

Using a *HP Photosmart 5520* with a print resolution of 1200 dpi. In the next step, the printed images are scanned using two different kinds of scanners such as (1) HP Photo

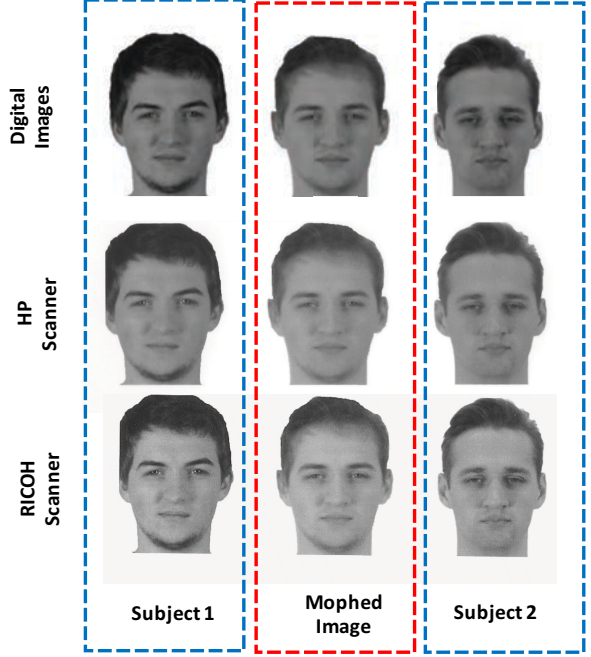


Figure 4: Example of the images from our morphed face database depicting mere digital images (top row), print-scanned images using HP Scanner (middle row) and print-scanned images using RICOH scanner (bottom row)

Scanner and (2) RICOH office scanner. The printed photos are scanned to have 300dpi following the specification from ICAO [6] regarding the face image reference in the ePass. Figure 4 illustrates an example of the digitally morphed image (see Figure 4 - top row), print-scanned images using HP scanner (see Figure 4 - middle row) and print-scanned images using RICOH scanner (see Figure 4 - bottom row). The middle image in each row of Figure 4 corresponds to the morphed face image generated using the two different subjects that are shown in either side of the morphed face image. As one can notice from Figure 4, the perceptual image quality degrades with the print-scanned process. This effect makes the detection more challenging. The database employed in this work has 352 bona fide and 431 morphed face images (for all three versions such as digital, print-scanned using HP and print-scanned using RICOH) that are partitioned into two disjoint (or non-overlapping) sub-sets used for training and testing. For the training set and the testing set the distribution of samples are presented in Table 2.

The process of generating the morphed face images involves a series of pre-processing and post-processing operations on the images that result in the different quality when compared to that of the bona fide images. Such differen-

Table 2: Composition of database used in this work

Type	Number of samples		
	Training set	Testing set	Total
Bona fide image	206	146	352
Morphed image	225	206	431

tial quality measures may bias the morph detection performance as the classifier may detect the quality/compression differences rather than the effects of morphing. Thus, we took the additional care to have the same quality of the morphed and bona fide face images by applying an equivalent sequence of pre- and post-processing operations on both bona fide and morphed face images.

3.1. Results and discussion

The performance of the morphed face detection algorithms are reported utilizing Detection Error Trade-off (DET) plots depicting the correlation between Attack Presentation Classification Error Rate (APCER) and Bona Fide Presentation Classification Error Rate (BPCER), which are defined following ISO/IEC 30107-3 [7] as:

APCER: proportion of attack presentations (i.e. morphed face samples) incorrectly classified as bona fide presentations in a specific scenario.

BPCER: proportion of bona fide presentations incorrectly classified as presentation attacks (i.e. morphed image) in a specific scenario.

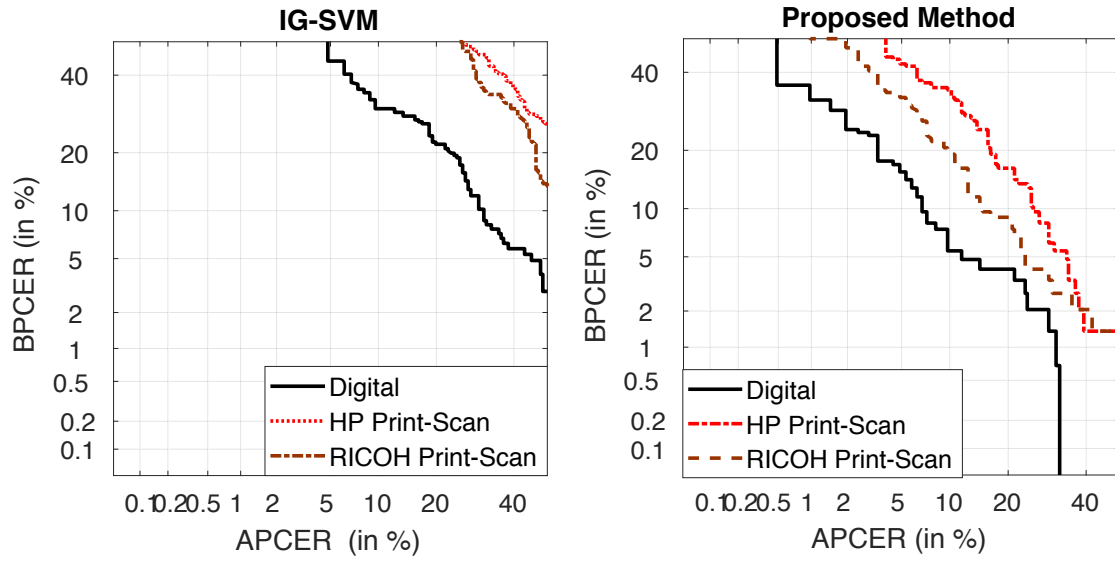
In addition, we will present the results by setting a fixed APCER at 5% and 10% and reporting the corresponding BPCER. Lower values of both APCER and BPCER achieved by the morphed face detection algorithm will represent the best detection performance.

In this work, we perform two different experiments to evaluate the robustness of the proposed method in our benchmark with state-of-the-art morphed face detection algorithms. **Experiment 1:** In this experiment, the detection performance of the proposed method together with the state-of-the-art methods are evaluated independently on the digital and print scanned morphed face image database. This experiment will analyse the performance of the detection schemes such that the training and the test data are from the same source (either digital or print-scanned (HP) or print-scanned (RICOH)). Thus, the detection methods will have prior knowledge of the type of the printer/scanner used to generate print-scanned morphed face image to be detected. **Experiment 2:** This experiment is designed to test the robustness of the morphed face detection algorithms including the proposed method. The morph detection methods are trained and tested on the different source of the

data. For example, the morphed face detection algorithms are trained on the digital version of the database and tested on the print scanned (either HP or RICOH). This experiment will indicate the generalizing capability of the morph detection algorithms. Further, the Experiment 2 will also address the realistic scenario of detecting the morphed face image as different scanners are used at different passport issuing authorities around the globe.

Table 3 shows the quantitative results of the proposed scheme in a benchmark with four different state-of-the-art algorithms evaluated independently on the digital subset, the HP-print-scanned subset and the RICOH-print-scanned subset. Figure 5 illustrates the DET curve showing the performance of the proposed method and the state-of-the-art methods based on the Image Gradients (IG-SVM) on all three morphed face data types. For the simplicity, we have considered to include only two DET curves while all the results are detailed in the Table 3. The following are the main observations from Experiment 1:

- The proposed method has demonstrated the best detection performance on all three independent evaluations on both digital and print-scanned version of the morphed face database. The proposed method has shown the least Detection Equal Error Rate (D-EER (%)) of 8.23% on the digital image database, the D-EER of 17.64% on print-scanned (HP) and a D-EER of 12.47% on print-scanned (RICOH). The second best detection performance is noted for the BSIFSVM scheme.
- In the general, the detection performance of the state-of-the-art schemes based on micro-texture schemes shows a degraded performance on all three version of the database.
- The detection performance of the morph detection algorithms including the proposed system shows the degraded performance on the print-scanned version of the database when compared with the digital version. It is still interesting to note the performance of the proposed scheme especially on the print-scanned database generated using RICOH has indicated a reasonable performance with an D-EER of 12.47% and a BPCER = 28.76% @ APCER = 5%. This shows the robustness of the proposed scheme on detecting both digital and print-scanned morphed face images.
- The type of the scanner also influences the performance of the morphed face detection algorithms. Our experiments revealed that the use of the high quality printed photo together with the photo scanner make the morph detection more challenging. This fact is demonstrated in our experiments as the performance of the morphed face detection algorithms including



(a) Morphing detection performance of the IG-SVM Method (b) Morphing detection performance of the proposed Method

Figure 5: Morphing detection performance on both digital and print-scanned morphed face database

Table 3: Quantitative detection performance of the proposed algorithm on both digital and print-scanned morphed face images

Datasets	Algorithms	D-EER (%)	BPCER (%) at	
			APCER = 10%	APCER = 5%
Digital	LBP-SVM [10]	29.28	44.52	56.84
	LPQ-SVM [10]	26.12	42.46	52.73
	IG-SVM [10]	21.88	30.13	41.78
	BSIF-SVM [10]	22.70	38.25	49.31
	Proposed Scheme	8.23	7.53	14.38
HP Print-Scan	LBP-SVM [10]	34.35	61.64	69.17
	LPQ-SVM [10]	49.82	92.46	94.52
	IG-SVM [10]	38.35	73.28	80.13
	BSIF-SVM [10]	26.12	45.89	55.47
	Proposed Scheme	17.64	32.87	41.78
RICOH Print-Scan	LBP-SVM [10]	22.70	58.21	42.46
	LPQ-SVM [10]	39.18	74.65	82.87
	IG-SVM [10]	34.35	52.05	63.69
	BSIF-SVM [10]	23.29	43.83	54.79
	Proposed Scheme	12.47	16.43	28.76

the proposed method has indicated high error rates on the print-scanned database generated using the photo scanner from HP when compared to the print-scanned

database created using the office scanner from RICOH.

Table 4 shows the performance of the proposed and state-of-the-art schemes on the cross database evaluation follow-

Table 4: Quantitative performance of the proposed scheme in the cross-data evaluation

Training Set	Testing set	Algorithm	D-EER (%)	BPCER (%) at	
				APCER = 10%	APCER = 5%
Digital	Print-Scan (HP)	BSIF-SVM	26.70	48.63	56.16
		Proposed Scheme	15.05	24.65	39.72
Digital	Print-Scan (RICOH)	BSIF-SVM	27.53	57.53	65.75
		Proposed Scheme	15.05	17.80	28.08
Print-Scan (HP)	Digital	BSIF-SVM	21.29	40.41	52.05
		Proposed Scheme	19.88	34.24	45.89
Print-Scan (HP)	Print-Scan (RICOH)	BSIF-SVM	28.11	58.90	67.80
		Proposed Scheme	13.06	21.23	28.76
Print-Scan (RICOH)	Print-Scan (HP)	BSIF-SVM	27.53	60.95	65.75
		Proposed Scheme	19.05	28.08	40.41
Print-Scan (RICOH)	Digital	BSIF-SVM	23.29	34.93	52.05
		Proposed Scheme	20.71	30.13	37.67

ing the Experiment 2. For the simplicity, we have presented the results of the proposed scheme with the second best method based on the BSIF-SVM [10]. The following are the main observations deduced from Experiment 2:

- The performance of the morph detection algorithms is degraded when compared to that of the results obtained in Experiment 1.
- The best detection performance is noted with the proposed method that significantly indicates the improved performance over the state-of-the-art schemes.
- The best detection performance of the proposed scheme is observed when the training data corresponds to the digital version of the morphed face database and testing the print-scanned version of the database. This indicates the robustness of the proposed scheme when compared to that of the state-of-the-art schemes.

Thus, based on our extensive experiments, the proposed method has emerged as the best method in detecting the morphed face images in both digital and print-scanned version generated from the different kinds of scanner. The proposed method also shows the best performance on the cross database evaluation following the protocol indicated in the Experiment 2 indicating the applicability of proposed method compared to other state-of-art methods.

4. Conclusion

Face morphing is emerging as a potential threat, especially for the passport control application. As many countries issue the electronic passport based on the photo submitted by the applicant, a wide range of face-morphing

technologies can be effectively applied to conceal the identity. The availability of freely available face morphing software provides an ample opportunity to the attacker to generate the morphed face photo, which is hard to detect by the human operator. In this work, we proposed a novel approach and framework to detect the morphed face images using the transferable features from the pre-trained Deep Convolutional Neural Networks (D-CNN) to detect both digital and print-scanned morphed face images. Extensive experiments are carried out on the newly constructed morphed image database comprised of 352 bona fide and 431 morphed face images created using 104 unique subjects. Furthermore, the morphed face images are generated with manual intervention to achieve best quality morph images and further post processing is carried out to remove the misaligned edges resulting in ghost images. To realistically analyze the performance of the proposed scheme, the database is created using two different kinds of scanners that include photo scanner and the office scanner (or document scanner) that can scan high quality printed morphed and bona fide face images at 300dpi. Extensive experiments that are carried out on this database show the lowest errors on both digital and print-scanned morphed face images. Further the improved detection performance achieved by the proposed method, especially in Experiment 2 on the cross database evaluation shows the robustness in detecting the morphed face images.

Acknowledgment

This work is carried out under the funding of the Research Council of Norway (Grant No. IKTPLUSS 248030/O70).

References

- [1] Gnu image manipulation program (gimp). <https://www.gimp.org>, 2016. Accessed: 2014-08-19.
- [2] Free and open source face recognition with deep neural networks. <https://cmusatyalab.github.io/openface/.html>, 2017. Accessed: 2017-04-10.
- [3] S. Cai, L. Zhang, W. Zuo, and X. Feng. A probabilistic collaborative representation based approach for pattern classification. 2016.
- [4] M. Ferrara, A. Franco, and D. Maltoni. The magic passport. In *IEEE International Joint Conference on Biometrics*, pages 1–7. IEEE, sep 2014.
- [5] M. Ferrara, A. Franco, and D. Maltoni. *Face Recognition Across the Imaging Spectrum*, chapter On the Effects of Image Alterations on Face Recognition Accuracy, pages 195–222. Springer International Publishing, 2016.
- [6] International Civil Aviation Organisation. Machine Readable Travel Documents - Part 9: Deployment of Biometric Identification and Electronic Storage of Data in eMRTDs. Technical report, ICAO, Montreal, 2006.
- [7] International Organization for Standardization. Information Technology – Biometric presentation attack detection – Part 3: Testing and reporting. ISO/IEC DIS 30107-3:2016, JTC 1/SC 37, Geneva, Switzerland, 2017.
- [8] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [9] Neurotechnology. VeriLook SDK.
- [10] R. Raghavendra, K. B. Raja, and C. Busch. Detecting Morphed Face Images. In *8th IEEE International Conference on Biometrics: Theory, Applications, and Systems (BTAS)*, pages 1–8, 2016.
- [11] R. Raghavendra, K. B. Raja, and C. Busch. Exploring the usefulness of light field cameras for biometrics: An empirical study on face and iris recognition. *IEEE Transactions on Information Forensics and Security*, 11(5):922–936, may 2016.
- [12] D. Robertson, R. S. Kramer, and A. M. Burton. Fraudulent id using face morphs: Experiments on human and automatic recognition. *IEEE Transactions on Image Processing*, 12(3):1–12, 2017.
- [13] J. L. Rodgers and W. A. Nicewander. Thirteen ways to look at the correlation coefficient. *The American Statistician*, 42(1):59–66, 1988.
- [14] U. Scherhag, R. Raghavendra, K. Raja, M. Gomez-Barrero, C. Rathgeb, and C. Busch. On the vulnerability of face recognition systems towards morphed face attack. In *International Workshop on Biometrics and Forensics (IWBF 2017)*, pages 1–6, 2017.
- [15] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [16] V. Struc. *The PhD face recognition toolbox : toolbox description and user manual*. Faculty of Electrical Engineering Ljubljana, 2012.
- [17] P. Viola and M. J. Jones. Robust real-time face detection. *International Journal of Computer Vision*, 57(2):137–154.