PRNU Variance Analysis for Morphed Face Image Detection

Luca Debiasi*, Christian Rathgeb†, Ulrich Scherhag†, Andreas Uhl* and Christoph Busch†

*WaveLab – The Multimedia Signal Processing and Security Lab, Universität Salzburg, Austria
†da/sec – Biometrics and Internet Security Research Group, Hochschule Darmstadt, Germany
{ldebasi,uhl}@cs.sbg.ac.at,
{ulrich.scherhag,christian.rathgeb,christoph.busch}@h-da.de

Abstract

In this work, a method to detect morphed face images based on Photo Response Non-Uniformity (PRNU) is presented. More specifically, the variance of PRNU-based features across image cells is estimated to distinguish bona fide from morphed and potentially post-processed morphed face images. The proposed morph detector is shown to be robust against post-processing techniques, which are likely to be applied to conceal the morphing process, e.g. histogram equalisation or image sharpening. Tested on a database of 961 bona fide and 2,414 automatically morphed face images, a detection equal error rate (D-EER) of 10.5% is obtained over all investigated attacks, including unaltered morphed images and various post-processing techniques.

1. Introduction

Automated face recognition [36, 17] represents a long-standing field of research in which a major break-through has been achieved by the introduction of deep neural networks [33, 24]. Resulting performance improvements paved the way for deployments of face recognition technologies in diverse application scenarios, ranging from mobile device access control to Automated Border Control (ABC). However, recently researchers found that the intended generalisability of deep face recognition systems also increases their vulnerability against attacks, e.g. spoofing attacks (a.k.a. presentation attacks) [22]. Most notably, a specific attack against face recognition systems based on morphed face images has been proposed in [3].

Morphing techniques can be used to create artificial biometric samples, which resemble the biometric information of two (or more) individuals in image and feature domain. In order to morph two face images, an attacker usually defines corresponding landmarks and a triangulation of landmarks is done on both images. The landmarks are then averaged to a single set of landmarks and both images are warped according to the resulting triangulation. Finally, alpha-blending is performed. Realistic morphed face images can be generated by non-experts employing easy-to-use face morphing software which can be purchased at a reasonable price, e.g. FantaMorph[1]. Fig. 1 depicts an example of morphing two face images.

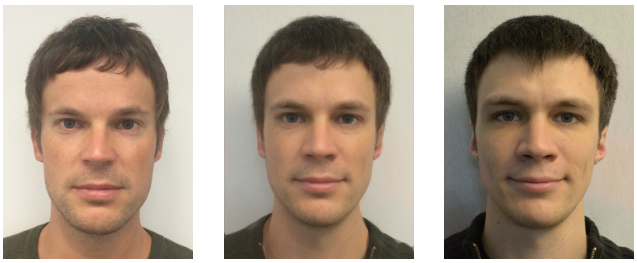

(a) Subject 1 (b) Morph (c) Subject 2

Figure 1: Examples for bona fide and morphed face images

It has been shown that morphed face images are realistic enough to fool human examiners [4]. This means, there is a risk that morphed biometric images are infiltrated to a biometric recognition system at enrolment, e.g. during the issuance process of electronic travel documents. In [3] commercial face recognition software tools have been exposed to be highly vulnerable to attacks based on morphed face images. This means that the subjects contributing to the morphed image were both (or all) successfully matched against that single enrolled morphed image. These findings have been confirmed by other researchers, e.g. in [32]. In their vulnerability analysis, researchers used decision thresholds yielding a False Match Rate (FMR) of 0.1%, following the guidelines provided by the European Agency for the Management of Operational Cooperation at the External Borders (FRONTEX) [1].

[1]FantaMorph: http://www.fantamorph.com/

978-1-5386-7180-1/18/$31.00

In the recent past, researchers have presented different approaches to distinguish bona fide from morphed face images, see Sect. 2. Proposed approaches either processes a single potentially morphed image, i.e. *no-reference* morph detection, or a potential morph together with a trusted live capture from an authentication attempt, i.e. *differential* morph detection. In the no-reference scenario different media forensic concepts have been applied [23, 16, 2]. Adaptations of such techniques, which are designed to detect digital forgeries, revealed promising results for the detection morphed face images. In particular, a PRNU-based detection of morphed face images was introduced in [2]. The extraction of the PRNU and an analysis of its distributions across image cells has been reported to reliably detect morphed face images, while the approach fails if image post-processing, e.g. histogram equalisation, is applied to generated morphs.

The work presented in this paper was inspired by the approach of [2] and proposes a PRNU variance analysis for morphed face image detection. It is shown that an increased variance of different PRNU statistics across image cells is a reliable indicator for image morphing. Further, the improved PRNU-based morph detector is shown to be resistant against common image post-processing methods. Finally, the presented approach is expected to be more robust against arbitrary post-processings, since it analyses image block interrelations rather than image features which might specifically result from a distinct morphing process applied to a certain face database.

This paper is organized as follows: related works are briefly discussed in Sect. 2. Fundamentals of PRNU extraction are explained in Sect. 3. The proposed morph detection method is described in detail in Sect. 4. Experimental results are reported in Sect. 5. Finally, conclusions and future works are summarized in Sect. 6.

2. Related Work

The topic of face morph detection has sparked the interest of numerous research laboratories working in the field of biometrics. Efforts to define evaluation metrics for morph detection and vulnerability analysis have already been made [28, 10], see Sect. 5. A recent overview on conducted vulnerability analyses and morph detection methods can be found in [20]. Presented approaches can be coarsely categorized with respect to the considered morph detection scenario. The majority of works assume the challenging no-reference scenario while some implement a differential morph detection which is motivated by the fact that trusted live captures are available in ABC scenarios.

A differential morph detection method referred to as demorphing was proposed in [5]. Within this approach a trusted live capture is aligned to a potential morph and subtracted from it in the image domain. The resulting image is then compared against the trusted live capture. A morph is detected if the biometric decision changes from "accept" to "reject". Robustness of de-morphing against slight face pose variations has been confirmed in [6]. Nevertheless, the authors note that in an ABC scenario the performance of de-morphing might degrade due to potential variations of quality and environmental conditions.

Several researchers have suggested the use of general purpose texture descriptors, e.g. Local Binary Patterns (LPB) or Binarized Statistical Image Features (BSIF), which have been employed widely for biometric recognition. Machine learning-based classifiers, e.g. Support Vector Machines (SVMs), are either trained directly on extracted feature vectors for no-reference morph detection [25, 29, 14] or differences between feature vectors can (additionally) be employed in a differential scenario [32]. Also, face-specific features such as differences between landmark positions or angles could be employed in a differential scenario which so far has been shown to reveal rather moderate detection performance [27]. Depending on the feature representation of texture descriptors the inputs of classifiers have to be adapted, e.g. for Scale-Invariant Feature Transform (SIFT) the number of extracted keypoints has been shown to be suitable for the task of morph detection [16, 32]. Score level fusions of different types of features have been proposed, too [30]. In particular, in the no-reference scenario classifiers may overfit to distinct micro texture features. These can be dataset-specific features which are altered or introduced by the applied morphing process. It has been shown that the performance of morph detectors based on general purpose texture descriptors might significantly decrease if training and test images stem from a different source, i.e. face database [31].

The use of convolutional neuronal networks for no-reference morphed face detection has been proposed by different researchers reporting promising results [26, 35]. Again, with these methods there is potentially a problem of overfitting. In particular, resulting deep classifiers may favour image locations where artefacts, e.g. shadows around the iris region, are likely to appear due to an imperfect automated morph creation process. Further, published approaches have been trained and tested for a single morph generation method, i.e. generalizability still has to be evaluated.

Focusing on the no-reference scenario diverse approaches related to media forensics have been presented. In different works, the detection of JPEG double-compression artefacts has been suggested for the purpose of morph detection [19, 10, 20]. However, the presence of such artefacts implies a strong assumption on the image format of face images used for morph generation as well as the resulting morphed face image. The International Civil Aviation Organization (ICAO) suggests face image data to be

stored in accordance with the specifications established by the International Organization for Standardization (ISO) in [12]. More specifically, the ICAO recommends face images to be stored in electronic travel documents at an average compressed sizes of 15kB to 20kB in JPEG or JPEG 2000 format [11]. Hence, depending on the image size and the employed compression algorithm the detection of JPEG double-compression artefacts might not be feasible. In [34] a morph detection method based on reflection analysis in face images is presented. The lightning direction is estimated based on reflections detected in the eyes of a potentially morphed image. Subsequently, reflections on the nose of the face are analysed. However, ISO requires hot spots and specular reflections to be absent in face images used in electronic travel documents. In particular, diffused lighting, multiple balanced sources or other lighting methods shall be used, i.e. a single bare "point" light source like a camera mounted flash is not acceptable for imaging [12]. Morph detection methods based on continuous image degradation have been proposed in [23, 16]. The basic idea behind these methods is to continuously degrade the image quality, e.g. by using JPEG compression, to create multiple artificial self-references of a face image. The distances from these references to the original image are then analysed for morph detection. Additionally, PRNU-based morph detection has been proposed in [2]. This approach is described in more detail in Sect. 4.

Despite promising results reported in many works a reliable detection of morphed face images still represents an open research challenge. Note that the generalizability/robustness of published approaches has not been shown, as these have been mostly trained and tested on single databases using a single morph generation algorithm. Further, the likely application of image post-processing techniques, e.g. image sharpening, is neglected in most works. Lastly, so far there are no publicly available database of bona fide and morphed face images and no publicly available morph detection algorithms.

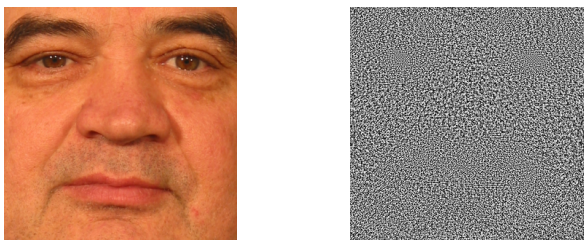

(a) Original (b) PRNU

Figure 2: Extracted and enhanced PRNU for an exemplary face image.

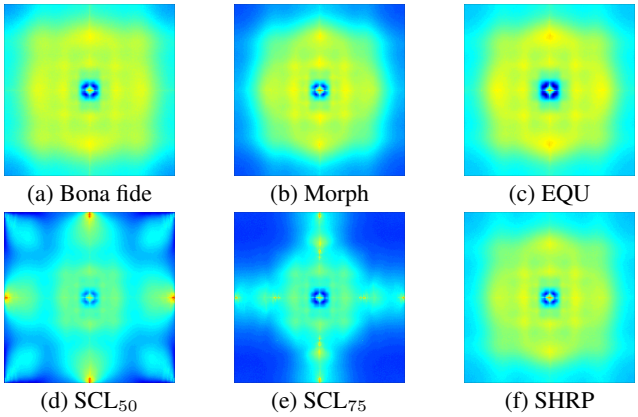

(a) Bona fide (b) Morph (c) EQU

(d) SCL_{50} (e) SCL_{75} (f) SHRP

Figure 3: Averaged PRNU DFT magnitude spectra of bona fide images (a), morphed images (b) and post-processed morphed images (c to f).

3. PRNU Extraction and Characteristics

Digital image forensics aims at acquiring knowledge on visual contents and acquisition devices by evaluating the traces that are left on the data during the acquisition and in the subsequent processing. The PRNU of imaging sensors [7] emerged as an important forensic tool. It can be used for a variety of important tasks, such as device identification, device linking, recovery of processing history, and detection of digital forgeries. The PRNU is an intrinsic property of all digital imaging sensors, which is characterised by slight variations among individual pixels in their ability to convert photons to electrons. Consequently, every sensor casts a weak noise-like pattern onto every image it captures. This noise-like pattern can be considered as an unintentional stochastic spread-spectrum watermark.

In [7] Fridrich presents an approach on how to extract the PRNU noise residual from an image. For each image I the noise residual W_I is estimated as described in Eq. (1),

$$W_I = I - F(I) \tag{1}$$

where F is a denoising function which filters out the sensor pattern noise. In this work, the denoising filter proposed by Mihcak *et al.* [21] is used in conjunction with a Filtering Distortion Removal (FDR) PRNU enhancement proposed by Lin *et al.* [18]. Said enhancement aims at improving the SNR of the extracted PRNU noise residual W_I in a two step process by abandoning certain components that are severely contamined by filtering errors introduced during the denoising of images. For further details on the denoising filter and FDR PRNU enhancement we refer to [21, 18]. Fig. 2 shows the extracted and enhanced PRNU for an exemplary face image.

The PRNU offers some essential advantages for the detection of morphed face images. First of all, as stated by Fridrich *et al.* [8], all digital image sensors exhibit PRNU, which makes this sensor noise virtually present in every

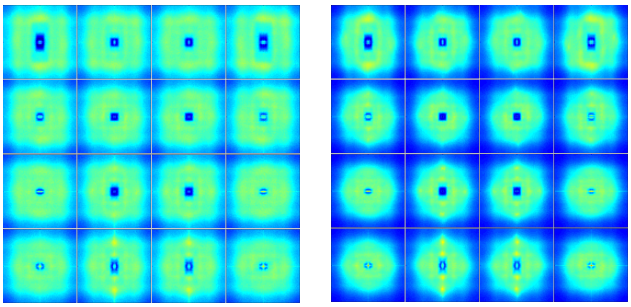

(a) Bona fide images (b) Morphed images

Figure 4: Illustration of variations across DFT magnitude spectra in morphed images compared to bona fide ones for 4×4 image cells (average of all images in dataset).

captured image. Furthermore, it is independent from the scene content and even robust against typical processing procedures like lossy compression or gamma correction, and it is even reported to be robust against high quality printing and scanning [9].

These criteria make the PRNU well suited for the morph detection scenario investigated in this work, because it offers significant advantages over analysing other high-frequency image components: First and foremost the PRNU is present in every image acquired with a digital camera, hence virtually every face image. In addition, in principle the PRNU is unrelated to the image content, but its high-frequency components might interfere with the PRNU. However, this interference can be attenuated by different PRNU enhancement approaches.

The spectral characteristics of the PRNU reveal whether an image has been subject to further processing [7]. Since face morphing usually comprises different non-linear warping and averaging operations, the distribution of the PRNU values is affected by these operations, as previously shown in [2]. The PRNU's DFT magnitude spectrum of morphed images shows a reduction of the high-frequency components as well as a compression of the whole spectrum, which is illustrated in Fig. 3b.

Debiasi *et al*. [2] furthermore investigated the effects of various post-processings on the PRNU's DFT magnitude spectrum. They applied four different post-processings to the morphed face images: Histogram equalisation (*EQU*), downscaling and subsequent upscaling (SCL_{50}, SCL_{75}) and sharpening (*SHRP*). More details are given in Sect. 5, while the effects of these operations are presented in Fig. 3. One can observe that the DFT spectra of SCL_{50} and SCL_{75} are clearly discriminable from bona fide images, whereas the spectra of *SHRP* and especially *EQU* show a high similarity to bona fide images.

4. Detection of Morphed Face Images

The PRNU-based morph detection system proposed by Debiasi *et al*. in [2] aims at exploiting the spectral alterations of the PRNU introduced by the non-linear warping during the face morphing process and therefore discriminate between bona fide and morphed images. Furthermore, the discrimination is performed in no-reference manner.

The morph detection system consists of five major components: (A) *PRNU extraction*, (B) *PRNU splitting*, (C) *cell-wise feature extraction*, (D) *cell aggregation* and the (E) *decision*. In short, the PRNU is extracted from a face image and divided into cells. Thereafter, the DFT magnitude spectrum is computed for each cell, whereof different features P are derived. By averaging the extracted features for each cell an aggregated score S is obtained. Finally, the system performs a binary decision (bona fide or morphed) based on a simple threshold, which can be determined by analysing the score distribution of bona fide images.

4.1. Variance Analysis

In this work, the approach of [2] is extended by proposing an analysis of the PRNU variance for morphed face image detection. Due to the morphing process's nature of producing inhomogeneous alterations across different image regions, an increased variance of the PRNU signal is expected across image cells. Fig. 4 shows the variations of the DFT magnitudes across different image cells of bona fide and morphed images. These local variations can be useful as a reliable indicator for image morphing. In order to analyse the variance of the PRNU, we propose some adaptations to Debiasi *et al*.'s [2] approach, which are presented in the remainder of this section. The proposed system is illustrated in Fig. 5.

4.1.1 Feature Extraction

In this work we propose to analyse the variance of two distinct features: P_{pos} and P_{en}. The first one, P_{pos}, has been proposed in [2] and is based on the PRNU's DFT magnitude histogram. It represents the peak's position (bin) within the histogram and is obtained as follows:

$$P_{pos} = \underset{n=1\ldots b}{\arg\max}\, H(n), \quad P_{en} = \sum_{x \in M} |x|^2 \tag{2}$$

where b is the number of bins and H is the magnitude histogram of a cell. As the second feature, P_{en}, we propose to compute the energy of the PRNU's DFT magnitudes, as defined in Eq. 2, where M are the DFT magnitudes within a cell and x their respective values. Both features lead to a scalar value P for each PRNU cell.

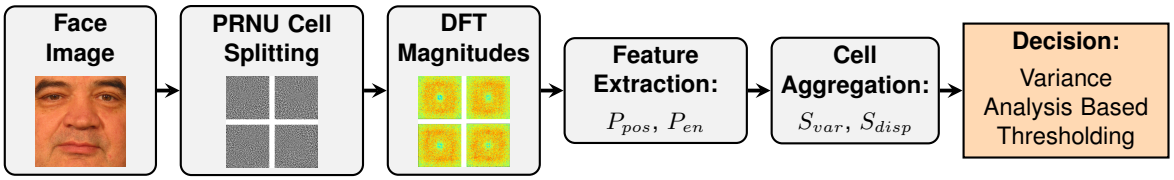

Figure 5: Processing steps of the proposed PRNU-based morph detection system.

4.1.2 Cell Aggregation

In order to perform the variance analysis across all image cells, we make use of two measures of dispersion. The variance, S_{var}, is given by

$$S_{var} = Var(P) = \frac{1}{N}\sum_{n=1}^{N}(P_n - \bar{P})^2 \tag{3}$$

$$\bar{P} = \frac{1}{N}\sum_{n=1}^{N} P_n, \quad S_{disp} = \frac{Var(P)}{\bar{P}} \tag{4}$$

The index of dispersion, S_{disp}, or variance to mean ratio, is given in Eq. 4, where N is the number of total PRNU cells, P_n is the feature (scalar value) obtained for the PRNU cell C_n, as described previously, and $\bar{P}$ is the average feature value for all PRNU cells C. In both cases, we obtain a single scalar value S for each image.

4.2. Decision

As mentioned above, the PRNU-based morph detection system proposed in [2] makes use of a simple thresholding to determine if the presented image is a bona fide one or not. It was shown that with this one dimensional decision it was not possible to reliably detect some of the post-processed morphed images, i.e. *SHRP* and in particular *EQU*.

Due to the large variety of possible unknown post-processings, we decided to focus on the known properties of bona fide images and to use this knowledge to our advantage by simply deriving the mean variation $\bar{B}$ from the bona fide images. With this characteristic of bona fide images, we are able to calculate the distance D of an investigated image to bona fide images as

$$D = |S - \bar{B}|, \quad \bar{B} = \frac{1}{N_B}\sum_{n=1}^{N_B} S \tag{5}$$

where S is the result of the cell-aggregation, $\bar{B}$ is the mean variation of all bona fide images N_B. The variation is either S_{var} or S_{disp}, whichever is used in the cell aggregation processing step. The final decision for a presented face image is taken by thresholding the distance D.

5. Experiments

In the following section, we describe the morphed face data set investigated in this work. In addition, we report experimental results which comprise a morph detection performance estimation and robustness under the presence of common post-processing techniques.

5.1. Face Morphing Data Set

In order to allow a direct comparison of the morph detection performance with [2], experiments are performed on a subset of the FRGCv2 face database, where 961 frontal faces with neutral expression have been manually selected as bona fide samples, which are all ICAO compliant according to [12]. Two face images are morphed by applying the *dlib* facial landmark detector [15] to both images. Subsequently, a Delaunay triangulation is computed, which forms the basis for a subsequent affine transform to the sets of triangles in both face images. The final morphed image is generated by alpha blending of the two warped images using an alpha value of 0.5.

The face images are then segmented and normalized according to eye coordinates detected by the *dlib* landmark detector. The resulting normalised region of interest is cropped to 320×320 pixels, to ascertain that the morphing detection algorithm is only applied to the facial region.

In total, $2,414$ high quality morphed face images have been automatically generated for pairs of subjects of same gender using the *OpenCV* library, which are well within the quality limits defined by ICAO and ISO/IEC standards.

(a) Bona fide (b) Morph (c) EQU

(d) SCL_{50} (e) SCL_{75} (f) SHRP

Figure 6: Data set examples: bona fide image (a), morphed image (b) and post-processed morphs (c - f).

Table 1: Performance of proposed PRNU-based morph detectors

		D-EER					BPCER10					BPCER20				
Algorithm	**Cells**	**Morph**	**EQU**	**SCL$_{50}$**	**SCL$_{75}$**	**SHRP**	**Morph**	**EQU**	**SCL$_{50}$**	**SCL$_{75}$**	**SHRP**	**Morph**	**EQU**	**SCL$_{50}$**	**SCL$_{75}$**	**SHRP**
$P_{pos}\|S_{mean}$	4	2.9%	33.0%	1.5%	0.1%	12.2%	0.2%	59.9%	0.1%	0.0%	16.5%	1.1%	71.8%	0.4%	0.0%	39.5%
$P_{pos}\|S_{var}$		30.3%	49.6%	18.2%	42.0%	14.0%	75.1%	88.2%	45.5%	85.0%	24.9%	87.1%	94.0%	71.1%	91.6%	54.4%
$P_{pos}\|S_{disp}$		25.2%	49.7%	14.1%	33.3%	12.9%	64.1%	88.6%	27.6%	77.3%	21.4%	79.9%	93.9%	57.5%	87.3%	50.9%
$P_{en}\|S_{var}$		19.4%	29.5%	4.3%	9.0%	2.3%	47.8%	51.6%	1.3%	7.9%	0.1%	69.3%	64.6%	3.5%	19.0%	0.6%
$P_{en}\|S_{disp}$		15.3%	30.3%	3.4%	5.5%	2.5%	30.2%	53.5%	0.6%	2.9%	0.1%	54.6%	67.0%	2.1%	6.1%	0.8%
$P_{pos}\|S_{mean}$	8	2.2%	33.8%	0.7%	0.0%	10.8%	0.1%	60.2%	0.0%	0.0%	11.7%	0.6%	71.5%	0.1%	0.0%	30.8%
$P_{pos}\|S_{var}$		18.5%	49.8%	2.5%	34.9%	4.9%	36.7%	89.4%	0.7%	78.6%	1.3%	64.6%	94.4%	1.3%	88.3%	4.7%
$P_{pos}\|S_{disp}$		11.1%	49.8%	1.7%	16.5%	4.8%	12.4%	89.6%	0.1%	29.8%	1.2%	27.7%	94.5%	0.7%	51.3%	4.1%
$P_{en}\|S_{var}$		20.2%	15.8%	4.2%	11.0%	1.3%	44.6%	20.3%	1.5%	12.3%	0.0%	66.0%	30.1%	3.5%	24.7%	0.2%
$P_{en}\|S_{disp}$		12.7%	16.8%	2.9%	4.5%	1.6%	16.2%	23.0%	0.5%	2.4%	0.0%	33.9%	33.1%	1.6%	4.1%	0.4%
$P_{pos}\|S_{mean}$	10	2.4%	34.9%	0.6%	0.0%	10.5%	0.0%	61.7%	0.0%	0.0%	11.2%	0.7%	71.6%	0.0%	0.0%	28.4%
$P_{pos}\|S_{var}$		15.3%	50.0%	1.4%	32.2%	3.6%	25.9%	90.0%	0.1%	77.6%	0.9%	44.2%	95.0%	0.4%	88.7%	2.2%
$P_{pos}\|S_{disp}$		7.5%	50.0%	1.0%	11.9%	3.8%	5.4%	90.0%	0.1%	15.1%	1.0%	11.8%	95.0%	0.1%	27.8%	2.2%
$P_{en}\|S_{var}$		18.3%	14.5%	3.5%	9.2%	1.1%	36.5%	17.5%	0.6%	8.3%	0.0%	56.4%	24.5%	2.3%	17.7%	0.0%
$P_{en}\|S_{disp}$		**11.0%**	**15.9%**	**2.6%**	**3.8%**	**1.5%**	**11.9%**	**20.0%**	**0.1%**	**1.9%**	**0.0%**	**22.0%**	**29.0%**	**0.9%**	**3.1%**	**0.1%**

Furthermore, Debiasi *et al*. [2] reported that the morphed face images generated for this data set pose a severe risk for a COTS face recognition system, since probe face images from both contributing subjects can match with the morph at high success rate. They obtained a Relative Morph Match Rate (RMMR) and the ProdAvg Mated Morph Presentation Match Rate (ProdAvg-MMPMR) of > 0.99, which emphasises the necessity of a robust morph detection system. For more details on metrics for reporting the vulnerability of face recognition systems to morphed faces, the reader is referred to [28].

Moreover, the data set also includes a variety of different post-processing techniques applied to the morphed images: *EQU*, *SCL*$_{50}$, *SCL*$_{75}$ and *SHRP*. They aim at hampering the detection performance of the morph detection system. Some examples for post-processed morphs, which are part of the investigated data set, are shown in Fig 6.

5.2. Morph Detection Performance Evaluation

The morph detection performance is examined according to metrics defined in ISO/IEC 30107-3 [13]: Attack Presentation Classification Error Rate (APCER) and bona fide Presentation Classification Error Rate (BPCER). APCER reports the proportion of attack presentations incorrectly classified as bona fide presentations in a specific scenario. BPCER, on the other hand, reports the proportion of bona fide presentations incorrectly classified as presentation attacks in a specific scenario. The operation point of the system, where APCER = BPCER, is defined as detection equal error rate D-EER. Furthermore, two additional operation points, BPCER10 (where APCER $= 10\%$) and BPCER20 (where APCER $= 5\%$), are reported.

Tab. 1 summarises the obtained morph detection performance in form of D-EER, BPCER10 and BPCER20 for images with (*EQU*, *SCL*$_{50}$, *SCL*$_{75}$, *SHRP*) and without post-processing (*Morph*). The column *Algorithm* comprises the combinations of extracted features P and aggregation strategies S defined in Sect. 4. The column *Cells* contains the cell splits of the investigated images. We focused on cell splits of 4×4, 8×8 and 10×10 in this work due to the improved results with higher cell counts reported in [2].

The proposed algorithm by Debiasi *et al*. in [2], $P_{pos}|S_{mean}$ for 8×8 cells, serves as baseline and achieves a D-EER performance of 2.2% for unaltered morphs, but fails at detecting morphs post-processed with *EQU* at 33.8% and shows a high performance decrease for detecting sharpened morphs (*SHRP*) at 10.5%. Because the magnitude spectra of *SCL*$_{50}$, *SCL*$_{75}$ and *SHRP* post-processing are quite distinct bona fide image's ones, as it can be observed in Fig. 3, they can be detected quite reliably in general. The remaining algorithms are based on the variance analysis described in Sect. 4.1.

The proposed $P_{pos}|S_{var}$ and $P_{pos}|S_{disp}$ algorithms show rather inconsistent results among the different post-processings, especially they completely fail at detecting *EQU* morphs. Since they are based on the DFT magnitude histograms, they are highly vulnerable to histogram shifts such as those caused by histogram equalisation (*EQU*), leading to a D-EER of up to 50%. When looking at $P_{en}|S_{var}$ and $P_{en}|S_{disp}$, one can immediately note the degradation in unaltered morph detection of 11% in the best case (compared to the baseline of 2.2%), as shown in Fig. 7a. However, a more stable performance across all post-processed morphs is achieved. The highest performance gains are achieved for *EQU* and *SHRP* with a D-EER of 14.5% and 1.1% respectively, as compared to the baseline of 33.0% and 10.5%, which are illustrated in Fig. 7b and 7c. In general, the variance analysis based algorithms lead to a trade off between unaltered morph detection and post-processed morph detection. It enables the system to be more robust against different attacks, while also increasing the overall performance when all attacks are considered (*Morphs*, *EQU*, *SCL*$_{50}$, *SCL*$_{75}$, *SHRP*). This can mainly be attributed to the statistical variations caused by the morph-

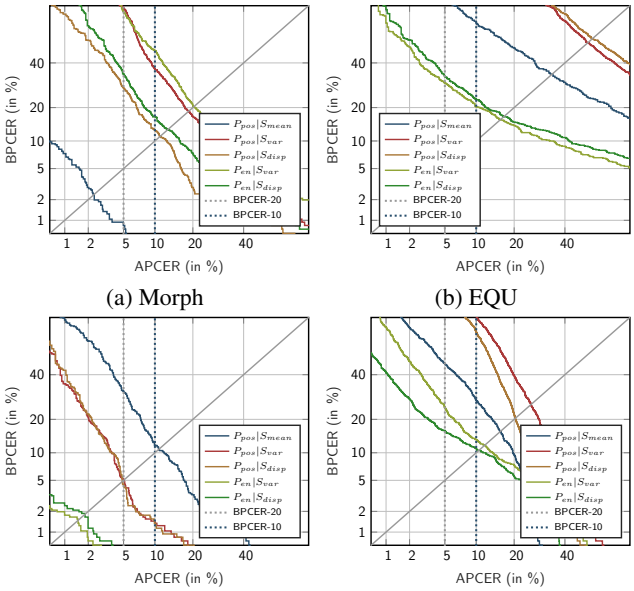

(a) Morph (b) EQU

(c) SHRP (d) ALL

Figure 7: DET curves for PRNU-based morph detectors (10×10 cells).

ing procedure across the image, which are most prominent in the P_{en} feature and can be captured best by the S_{disp} cell-aggregation strategy.

Thus, the overall best performing and most stable algorithm is based on the proposed variance analysis, $P_{en}|S_{disp}$ for 10×10 cells, is able to achieve respectable results across all altered and unaltered morphed images and is robust against a wide variety of post-processing attacks aiming at deteriorating the morph detection system. This robustness is a significant improvement over the baseline algorithm proposed in [2], which is much more vulnerable to post-processing attacks. Furthermore, the overall system performance is also improved from 15.7% average D-EER (baseline) to 10.5% D-EER (proposed algorithm), when all altered and unaltered morphs are considered. A direct comparison of both algorithms is presented in Fig. 7d and Tab. 2, where it can be observed that both algorithms have opposing strengths and weaknesses regarding the single post-processing techniques. Hence, a fusion of both approaches might be beneficial for the overall performance of the morph detection system.

Table 2: D-EER performance comparison of proposed PRNU variance analysis based detector ($P_{en}|S_{disp}$) with baseline ($P_{pos}|S_{mean}$) proposed in [2]. The column *ALL* reports the D-EER including all attacks (*Morph* to *SHRP*).

Algorithm	**Cells**	**D-EER**					
		Morph	**EQU**	**SCL$_{50}$**	**SCL$_{75}$**	**SHRP**	**ALL**
$P_{pos}\|S_{mean}$	8	2.2%	33.8%	0.7%	0.0%	10.8%	15.7%
$P_{en}\|S_{disp}$	10	11.0%	15.9%	2.6%	3.8%	1.5%	10.5%
Difference		+ 8.8%	-17,9%	+1,9%	+3,8%	-9,3%	-5,2%

6. Conclusion and Future Work

When infiltrated during the enrolment process of a face recognition system, morphed face images pose a serious security risk, in particular in the context of ABC. In this work, a morph detector, which analyses the variance of PRNU-based features across image cells, is proposed. In contrast to related work [2], the presented approach is shown to be robust to diverse image post-processing techniques and even improves the D-EER for all investigated attacks, which include unaltered morphed images and various post-processing techniques, to 10.5%.

Compared to many other schemes, the presented system is expected to achieve high robustness, as it analyses relative changes of PRNU-based features across images regions rather than distinct texture features. Such changes inevitably occur if image morphing is applied. In order to avoid artefacts, some morphing algorithms paste morphed face regions within the convex hull of averaged landmarks into the outer region of one of the contributing face images. This would cause an even higher variance of PRNU features across image regions resulting in improved detection performance.

Future work will be focused on a more thorough analysis of the proposed approach, i.e. detection performance will be evaluated for bona fide and morphed images created from different face image databases using different morph generation algorithms. A comparison of the presented system against published face morph detectors will also be performed in future work. Finally, the creation of a database of printed and scanned (morphed) face images and a corresponding evaluation of the presented morph detection methods in different scenarios is subject to future work.

Acknowledgements

This work was supported by the German Federal Ministry of Education and Research (BMBF) as well as by the Hessen State Ministry for Higher Education, Research and the Arts (HMWK) within the Center for Research in Security and Privacy (CRISP). Furthermore, this work was also supported by the Austrian Science Fund (FWF) under Project No. P26630.

References

[1] FRONTEX – Research and Development Unit: Best practice technical guidelines for automated border control (ABC) systems, 2012. Version 2.0.

[2] L. Debiasi, U. Scherhag, C. Rathgeb, A. Uhl, and C. Busch. PRNU-based detection of morphed face images. In *2018 6th Intl. Workshop on Biometrics and Forensics (IWBF)*. IEEE, 2018.

[3] M. Ferrara, A. Franco, and D. Maltoni. The magic passport. In *Proc. Int. Joint Conf. on Biometrics (IJCB)*, 2014.

[4] M. Ferrara, A. Franco, and D. Maltoni. On the effects of image alterations on face recognition accuracy. In T. Bourlai, editor, *Face Recognition Across the Imaging Spectrum*. Springer International Publishing, 2016.

[5] M. Ferrara, A. Franco, and D. Maltoni. Face demorphing. *IEEE Transactions on Information Forensics and Security*, 13(4), 2018.

[6] M. Ferrara, A. Franco, and D. Maltoni. Face demorphing in the presence of facial appearance variations. In *2018 26th European Signal Processing Conf. (EUSIPCO)*, 2018.

[7] J. Fridrich. Digital image forensic using sensor noise. *IEEE Signal Processing Magazine*, 26(2), 2009.

[8] J. Fridrich. Sensor defects in digital image forensics. In H. Sencar and N. Memon, editors, *Digital Image Forensics: There is more to a picture than meets the eye*, chapter 6. Springer Verlag, 2012.

[9] M. Goljan, J. Fridrich, and J. Lukas. Camera identification from printed images. In *Proc. of SPIE, Electronic Imaging, Forensics, Security, Steganography, and Watermarking of Multimedia Contents X*. SPIE, 2008.

[10] M. Hildebrandt, T. Neubert, A. Makrushin, and J. Dittmann. Benchmarking face morphing forgery detection: Application of stirtrace for impact simulation of different processing steps. In *2017 5th Intl. Workshop on Biometrics and Forensics (IWBF)*. IEEE, 2017.

[11] ICAO. *ICAO Doc 9303, Machine Readable Travel Documents – Part 9: Deployment of Biometric Identification and Electronic Storage of Data in MRTDs (7th edition)*, 2015.

[12] International Organization for Standardization. Information technology – Biometric data interchange formats – Part 5: Face image data. ISO/IEC 19794-5:2005 consolidated, JTC 1/SC 37, 2005.

[13] ISO/IEC JTC1 SC37 Biometrics. *ISO/IEC IS 30107-3:2017, IT – Biometric presentation attack detection – Part 3: Testing and Reporting*.

[14] S. Jassim and A. Asaad. Automatic detection of image morphing by topology-based analysis. In *2018 26th European Signal Processing Conf. (EUSIPCO)*, 2018.

[15] D. E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10, 2009.

[16] C. Kraetzer, A. Makrushin, T. Neubert, M. Hildebrandt, and J. Dittmann. Modeling attacks on photo-ID documents and applying media forensics for the detection of facial morphing. In *Proc. of the 5th ACM Workshop on Information Hiding and Multimedia Security - IHMMSec*. ACM Press, 2017.

[17] S. Z. Li and A. K. Jain. *Handbook of Face Recognition (2nd edition)*. Springer, 2011.

[18] X. Lin and C.-T. Li. Enhancing sensor pattern noise via filtering distortion removal. *IEEE Signal Processing Letters*, 23(3), 2016.

[19] A. Makrushin, T. Neubert, and J. Dittmann. Automatic generation and detection of visually faultless facial morphs. In *Proc. of the 12th Intl. Joint Conf. on Computer Vision, Imaging and Computer Graphics Theory and Applications*. SCITEPRESS - Science and Technology Publications, 2017.

[20] A. Makrushin and A. Wolf. An overview of recent advances in assessing and mitigating the face morphing attack. In *2018 26th European Signal Processing Conf. (EUSIPCO)*, 2018.

[21] M. Mihcak, I. Kozintsev, and K. Ramchandran. Spatially adaptive statistical modeling of wavelet image coefficients and its application to denoising. In *Proc. of the 1999 IEEE Intl. Conf. on Acoustics, Speech, and Signal Processing, ICASSP '99*. IEEE, 2009.

[22] A. Mohammadi, S. Bhattacharjee, and S. Marcel. Deeply vulnerable: a study of the robustness of face recognition to presentation attacks. *IET Biometrics*, 7(1), 2018.

[23] T. Neubert. Face morphing detection: An approach based on image degradation analysis. In *Digital Forensics and Watermarking*. Springer International Publishing, 2017.

[24] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *Proc. of the British Machine Vision Conf. 2015, BMVC*, 2015.

[25] R. Ramachandra, K. B. Raja, and C. Busch. Detecting morphed face images. In *2016 IEEE 8th Intl. Conf. on Biometrics Theory, Applications and Systems (BTAS)*. IEEE, 2016.

[26] R. Ramachandra, K. B. Raja, S. Venkatesh, and C. Busch. Transferable deep-CNN features for detecting digital and print-scanned morphed face images. In *2017 IEEE Conf. on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE, 2017.

[27] U. Scherhag, D. Budhrani, M. Gomez-Barrero, and C. Busch. Detecting morphed face images using facial landmarks. In *Lecture Notes in Computer Science*. Springer International Publishing, 2018.

[28] U. Scherhag, A. Nautsch, C. Rathgeb, M. Gomez-Barrero, R. Veldhuis, L. 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 *Int. Conf. of the Biometrics Special Interest Group (BIOSIG)*, 2017.

[29] U. Scherhag, R. Raghavendra, K. B. Raja, M. Gomez-Barrero, C. Rathgeb, and C. Busch. On the vulnerability of face recognition systems towards morphed face attacks. In *2017 5th Intl. Workshop on Biometrics and Forensics (IWBF)*. IEEE, 2017.

[30] U. Scherhag, C. Rathgeb, and C. Busch. Morph detection from single face images: a multi-algorithm fusion approach. In *Proc. of the 2018 Intl. Conf. on Biometrics Engineering and Application (ICBEA)*. ACM, 2018.

[31] U. Scherhag, C. Rathgeb, and C. Busch. Performance variation of morphed face image detection algorithms across different datasets. In *2018 6th Intl. Workshop on Biometrics and Forensics (IWBF)*. 2018.

[32] U. Scherhag, C. Rathgeb, and C. Busch. Towards detection of morphed face images in electronic travel documents. In *Proc. of the 13th IAPR Workshop on Document Analysis Systems (DAS)*, 2018.

[33] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *2015 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.

[34] C. Seibold, A. Hilsmann, and P. Eisert. Reflection analysis for face morphing attack detection. In *2018 26th European Signal Processing Conf. (EUSIPCO)*, 2018.

[35] C. Seibold, W. Samek, A. Hilsmann, and P. Eisert. Detection of face morphing attacks by deep learning. In *Digital Forensics and Watermarking*. Springer International Publishing, 2017.

[36] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld. Face recognition: A literature survey. *ACM Comput. Surveys*, 35(4), 2003.