Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2017.DOI

Face Recognition Systems under Morphing Attacks: A Survey

ULRICH SCHERHAG[1], CHRISTIAN RATHGEB[1 2], JOHANNES MERKLE[2], RALPH BREITHAUPT[3], CHRISTOPH BUSCH[1]

[1]da/sec - Biometrics and Internet Security Research Group, Hochschule Darmstadt, Germany, {ulrich.scherhag, christian.rathgeb, christoph.busch}@h-da.de
[2]secunet Security Networks AG, Essen, Germany, johannes.merkle@secunet.com
[3]Federal Office of Information Security (BSI), Bonn, Germany, ralph.breithaupt@bsi.bund.de

Corresponding author: Ulrich Scherhag (e-mail: ulrich.scherhag@h-da.de).

ABSTRACT Recently, researchers found that the intended generalisability of (deep) face recognition systems increases their vulnerability against attacks. In particular, attacks based on morphed face images pose a severe security risk to face recognition systems. In the last few years, the topic of (face) image morphing and automated morphing attack detection has sparked the interest of several research laboratories working in the field of biometrics and many different approaches have been published. In this work, a conceptual categorisation and metrics for an evaluation of such methods is presented, followed by a comprehensive survey of relevant publications. Additionally, technical considerations and trade-offs of the surveyed methods are discussed along with open issues and challenges in the field.

INDEX TERMS Biometrics, face morphing attack, face recognition, image morphing, morphing attack detection.

I. INTRODUCTION

Automated face recognition [1], [2] represents a longstanding field of research in which a major break-though has been achieved by the introduction of deep neural networks [3], [4]. Due to the high generalization capabilities of deep neural networks specifically and recognition systems in general, the performance of operational face recognition systems in unconstrained environments, e.g., regarding illumination, poses, image quality or cameras, improved significantly. Resulting performance improvements paved the way for deployments of face recognition technologies in diverse application scenarios, ranging from video-based surveillance and mobile device access control to Automated Border Control (ABC). However, recently researchers found that the generalizability of (deep) face recognition systems increases their vulnerability against attacks, e.g., spoofing attacks (also referred to as presentation attacks) [5]. An additional attack vector enabled by the high generalization capabilities is a specific attack against face recognition systems based on morphed face images, as introduced by Ferrara et al. [6].

A. FACE MORPHING ATTACK

Image morphing has been an active area of image processing research since the 80s [7], [8] with a wide variety of application scenarios, most notably in the film industry. 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. An example of a *morphed* face image as the result of two non-morphed, i.e., *bona fide* [9], face images, is depicted in Fig. 1. The created morphed face image will be successfully verified against probe samples of both contributing subjects by state-of-the-art face recognition systems. This means, if a morphed face image is stored as reference in the database of a face recognition system, both contributing subjects can be successfully verified against this manipulated reference. Thus, morphed

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

FIGURE 1: Example for a morphed face image (b) of subject 1 (a) and subject 2 (c). The Morph was manually created using FantaMorph.

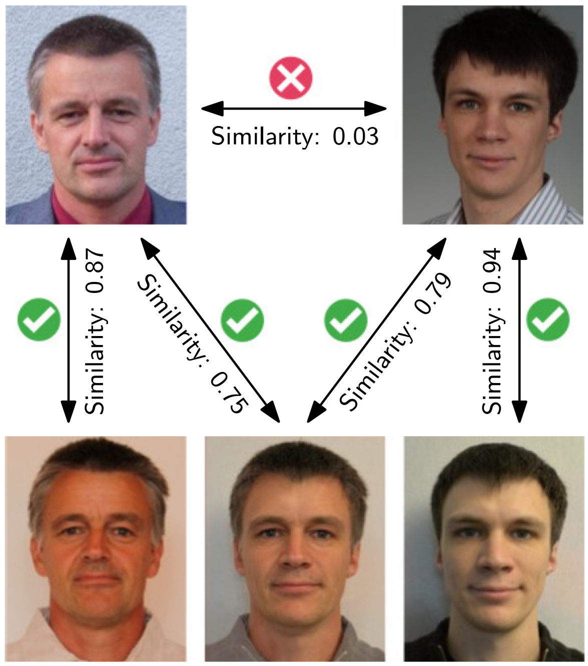

FIGURE 2: Example for the face morphing attack: different instances of face images of both subjects contributing to a face morph are successfully matched against it using a COTS face recognition software with a default decision threshold of 0.5, resulting in an FMR of 0.1%.

face images pose a severe threat to face recognition systems, as the fundamental principal of biometrics, the unique link between the sample and its corresponding subject, is violated.

In many countries, the face image used for the ePassport issuance process is provided by the applicant in either analog or digital form. In a face morphing attack scenario, a wanted *criminal* could morph his face image with one of a lookalike *accomplice*. If the accomplice applies for an ePassport with the morphed face image, he will receive a valid ePassport equipped with the morphed face image. It is important to note, that morphed face images can be realistic enough to fool human examiners [10], [11]. Both, the criminal and the accomplice could then be successfully verified against the morphed image stored on the ePassport, as visualized in Fig. 2. This means, the criminal can use the ePassport issued to the accomplice to pass ABC gates (or even human inspections at border crossings). The risk posed by this attack, referred to as *face morphing attack*, is amplified by the fact that realistic morphed face images can be generated by non-experts employing easy-to-use face morphing software which is either freely available or can be purchased at a reasonable price, e.g., FaceMorpher[1], WinMorph[2] or FantaMorph[3].

[1]FaceMorpher, Luxand: http://www.facemorpher.com/
[2]WinMorph, DebugMode: http://www.debugmode.com/winmorph/
[3]FantaMorph, Abrasoft: http://www.fantamorph.com/

B. CONTRIBUTION AND ORGANIZATION

In 2014, Ferrara et al. [6] were the first to thoroughly investigate the vulnerability of commercial face recognition systems to attacks based on morphed face images. Up to now, a significant amount of literature related to face morphing attacks and their detection has already been published, while only a rather brief overview has been given in [12]. This survey provides a comprehensive overview and critical discussion of published literature related to said topics. This survey primarily addresses biometrics researchers and practitioners. The remainder of this article is organized as follows: the fundamentals of (face) image morphing and quality assessment of face morphs are described in Sect. II and Sect. III, respectively, along with an overview of available software tools in Sect. IV. Subsequently, relevant metrics to assess the vulnerability of face recognition systems against said attack and the performance of morphing attack detection methods are summarized in Sect. V. Proposed approaches for automated morphing attack detection are surveyed and discussed in Sect. VI. Open issues and challenges are outlined in Sect. VII. Finally, a conclusion is given in Sect. VIII.

II. MORPHING OF FACE IMAGES

Image morphing in general represents a well-investigated field of research, for comprehensive surveys the reader is referred to [7], [8]. In this section, surveyed approaches are limited to morphing techniques, which have been explicitly applied to (frontal) face images. Face images used to create a morph should meet certain requirements. The best results can be achieved with frontal images exhibiting a neutral facial expression. In the context of the face morphing attack it should be expected that not only for the input face images (provided by the photographer) but also for the resulting morph the prerequisites of the International Civil Aviation Organization (ICAO) [13] for the production of passport portrait photos have to be met. These specifications ensure that all faces are represented equally with respect to resolution, exposure, etc. Semi-profile recordings can indeed be partially corrected, but then there is usually information missing of the far side of the face. Furthermore, the quality of the source images has a direct influence on the result. The quality of the morph cannot be expected to be higher than that of the source images. Distortions and scaling usually negatively affect quality during the process chain. The quality of morphed face images is further discussed in Sect. III.

In general, the morphing process of face images can be divided into three steps. First, a *correspondence* between the contributing samples is determined. In a second step, called *warping*, both images are distorted, such that the corresponding elements of both samples are geometrically aligned. Finally, the colour values of the warped images are merged, referred to as *blending*, in order to create the morphed face image. Said processing steps are described in detail in the following subsections, along with post-processing, studies on human perception of morphed face images and a summary of available research resources.

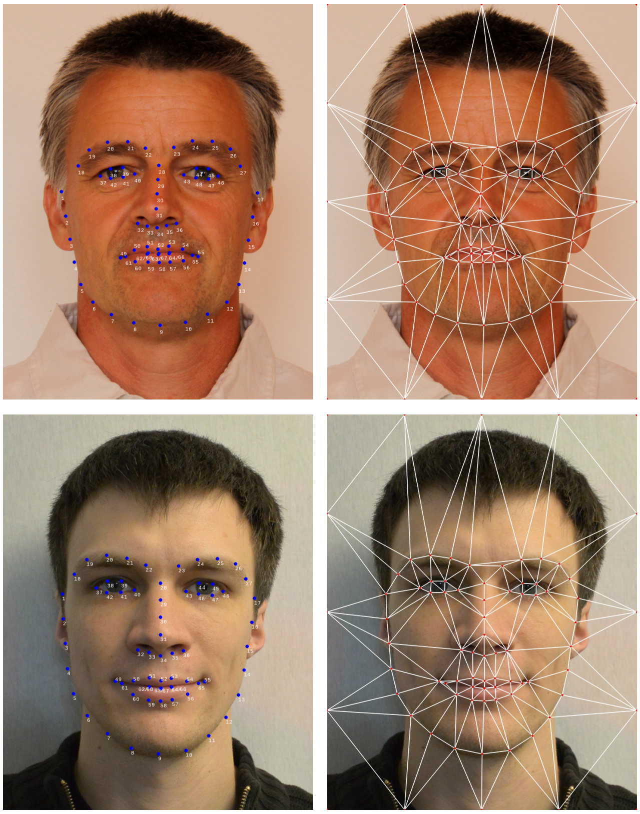

FIGURE 3: Examples of detected landmarks (using dlib landmark detector) and corresponding Delaunay triangles.

A. CORRESPONDENCE

The most common way of determining correspondences between face images is by determining salient points in both images, so-called *landmarks*. The simplest way is to manually define the coordinates of prominent characteristics, e.g., eyes, eyebrows, tip of the nose, etc., as for instance done in the morphing process of [6] and [14]. The manual annotation of images is very accurate (if done properly), but time consuming. More convenient is the automated detection of landmarks. The established approach for landmark detection is to detect each point separately, e.g., utilizing geometric features [15]. A more sophisticated solution is to fit a predefined model, e.g., active shape models [16] or elastic bunch graph models [17], [18] to the face image, whereas the fitting of the model is the key issue. Zanella and Fuentes propose an untrained generic model, which is fit to the contours of a binary image using evolutionary strategies [19]. Saragih et al. [20] propose a principled optimization strategy where a non-parametric representation of the landmark distributions is maximized within a hierarchy of smoothed estimates. Further algorithms train multiple regression trees for landmark detection [21], [22], of which the method of Kazemi and Sullivan [22] was further implemented in the widely used dlib landmark detector [23]. For detailed information and benchmarks of different automated landmark detection approaches the reader is referred to [24].

B. WARPING

If the landmarks are determined, the image should be distorted in a manner, that corresponding landmarks are aligned. A straight forward method for morphing is scattered data interpolation [25]. The landmarks, also called control points, are moved to a new position, the new position of all intervening pixels is interpolated based on the nearby control points. More advanced morphing techniques take the correlation between the landmarks into account. For example, Sederberg et al. [26] propose a grid or mesh-based warping technique called Free-Form Deformation (FFD), which was extended by Lee et al. to multi-level FFD [27]. The whole image is considered as a grid, which is deformed by the flow of the landmarks. Another approach is field morphing introduced by Beie and Neely [28], where grid lines are controlling the metamorphosis of the image in the transformation. In particular, for manual morphing this approach has advantages, as the user can position lines instead of points. For automatic morphing the lines can be derived from detected landmarks. In the work of Schäfer et al. [29] the moving least squares are minimized in order to estimate the optimal affine transformation. This approach can be employed to optimize different warping methods based on landmarks or lines. Choi et al. proposes a morphing process by simulating the image as a mass spring system [30]. Thus, each translated landmark influences nearby pixels and landmarks.

Most state-of-the-art morphing algorithms, e.g., as used for the morph-creation in [31]–[38], do not consider the image as a grid, but apply a Delaunay triangulation on the landmarks in order to determine non overlapping triangles, as depicted in Fig. 3. Delaunay triangulations maximize the minimum angle of each triangle in the triangulation and can be calculated efficiently. Subsequently, the triangles of both contributing images are distorted, rotated and shifted until an alignment is achieved.

The first step in traditional approaches for creating a morph between a pair of face images I_0 and I_1 is to define a map ϕ from I_0 to I_1. The contribution of each subject to the warping process is defined by an α_w-value, whereas an $\alpha_w = 0$ would be the landmark-position of the first subject, $\alpha_w = 1$ the landmark-position of the second subject and an α_w between 0 and 1 any combination of both. The impact of different α_w-values on the resulting face morph can be seen by analysing the first versus the last row of Fig. 4. One issue that might occur are *disocclusions* which refers to regions in the object space that are visible in I_0, but disappear in I_1 as described by Liao in [39]. For disocclusions in I_0, the map ϕ is typically undefined, for disocclusions in I_1 it is discontinuous. To obtain a more complete representation, one can introduce a second map from I_1 back to I_0. Maintaining consistency between the two maps during an optimization process becomes quite expensive [39]. One approach solving this issue is proposed by Wu et al. [40]. The images are warped forward and backward in order to obtain a complete mapping ϕ. In addition, to obtain a more natural warping, the face images are projected into a 3D space and an energy

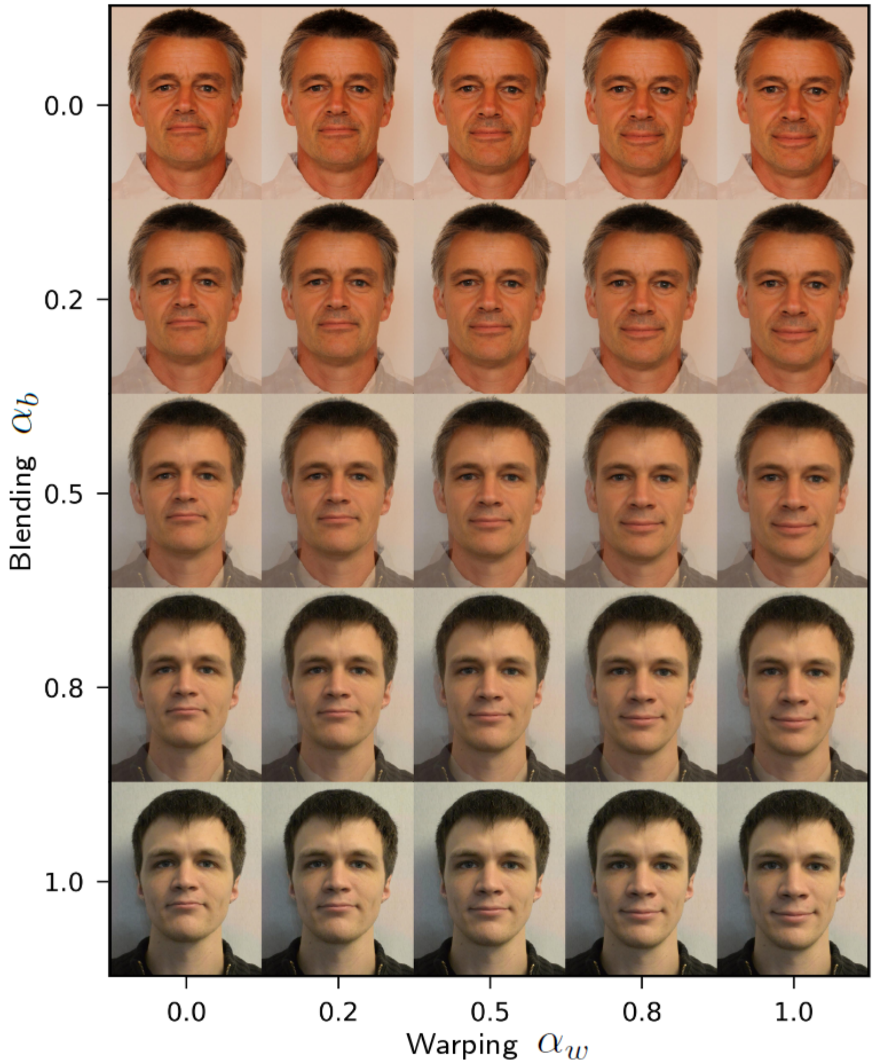

FIGURE 4: Matrix of the two variables in a morphing process (blending and warping). This morph sequence was created using dlib for landmark detection, Delaunay triangulation and linear affine transformation for warping and linear blending.

function is minimized to avoid ghost and blur artefacts. Seitz et al. [41] also proposes a projection into 3D-space, in order to consider perspective effects during the morphing process. Another technique for morphing in 3D-space is given by Yang et al. in [42]. In order to recover the face geometry, the 2D face image is projected on a pre-learned 3D face mask. In particular, for variances in pose and expression this approach promises a higher quality.

Further, some warping algorithms do not need previously detected landmarks. Bichsel et al. [43] propose to employ the Bayesian framework in order to determine the optimal mapping function.

C. BLENDING

After the alignment of the two contributing images, the two arranged textures are combined using blending, usually over the entire image region. The most frequent way of blending for face morph creation is linear blending, i.e. all colour values at same pixel positions are combined in the same manner. Similar to the warping process the contribution to the blending of each image can be weighted by an α_b-value, e.g. $\alpha_b = 0.5$ for averaging. The impact of a changing α_b-value to the morphed image can be seen in Fig. 4 on the vertical axis.

D. FURTHER APPROACHES

There are, however, some morphing algorithms, where a subdivision into the steps described above is not feasible. In [44], a morphing approach is proposed using generative morphing to combine warping and blending. The resulting morphed image is regenerated from small pieces of the source images. Korshunova et al. [45] propose to train a Convolutional Neural Network (CNN) to swap the face image of one subject with the face of a second one. A huge disadvantage of this method is, that a new network has to be trained for each subject.

Beside the morphing of samples in image domain, it is possible to morph in feature domain, as e.g., shown in [46] for minutiae sets and in [47] for iris-codes. It would be feasible to also morph face representations in feature domain, e.g., by averaging the feature vector of a CNN [48]. In order to use the morphed feature vector in a face recognition system, a face image can be reconstructed from the feature domain, as shown in [49]. However, it is most likely, that the reconstructed morphed face image only works for the same feature space, meaning an attack against the same face recognition system, as used for creation of the morphed feature vector.

E. POST-PROCESSING

After the creation of the morphed face image, the image might be further processed and altered. In order to obscure the image manipulation, the image quality might be enhanced or reduced on purpose.

In particular, the automated creation of morphed face images can lead to morphing artefacts. Missing or misplaced landmarks might cause shadow or ghost artefacts, as they can be seen in Fig. 5 (a). This issue can be tackled by swapping the facial area of the morphed face image with an adapted outer area of one of the subjects [35], [50]. Artefacts in the hair region can be concealed by an interpolation of the hair region as proposed by Weng et al. [51]. Further, unnatural colour gradients and edges might occur, due to inappropriate interpolation methods, which can be removed by blurring or sharpening. Due to the averaging during the blending process, the histograms of the colour values might get narrow. This artefact can be avoided by an adaptation of the colour histogram, e.g. by using histogram equalization or an adaption of lumination, in order to achieve realistic histogram shapes. Examples for sharpening and histogram equalization are depicted in Fig. 5 (b) and (c).

In addition to the removal or reduction of morphing artefacts, further post-processing steps might be carried out, which can sometimes be unavoidable, i.e., printing and scanning of the image, in order to use it as a passport photo. Even with high-end photo printer in the processing pipeline, some information contained in the face image signal will always be lost in the process, masking or reducing morphing artefacts, as described in [36]. Once the image has been submitted to a passport application office, it has to be scanned again. Again, information can be lost, helping to hide or reduce erroneous artefacts.

Further, information from or trace of the morphing process can be lost when the image format is changed. By storing the image in a lossy format, high-frequency information is eliminated from the signal permanently. If the image is loaded and stored multiple times as part of the process chain, the accumulated compression error can significantly degrade the image quality.

(a) original morph (b) sharpness (c) hist. equlization

FIGURE 5: Examples of different post-processing methods likely to be applied by an attacker to conceal the morphing process.

III. QUALITY ASSESSMENT OF FACE MORPHS

Generally speaking, automatically generated databases of morphed face images are expected to differ in quality from real world attack scenarios. Automatically generated morphs might reveal artefacts, which can be avoided when the attacker is producing only one single high quality morph between himself and his accomplice and manually optimising the resulting image. When aiming to develop a robust detection algorithm on such an automatically generated database, it is crucial to assure high quality of morphed face images. Otherwise, it is likely that a trained classifier might strongly rely on these specific artefacts.

As described by Scherhag et al. [52] it is difficult to define objective metrics for quality assessment of face morphs due to the large number of contributing factors. Basically, the output image of the algorithms can be evaluated according to the criteria summarized in the following subsections.

A. IMAGE QUALITY

Each processing step affects the quality of an image. In particular, factors such as image size, sharpness, colour saturation, aspect ratio and the overall natural appearance of the face image should be influenced as little as possible by the morphing algorithm. The minimum requirements for these factors can be found in the specifications for passport images of the ICAO [13]. Thus, for example, the minimum resolution of the facial image is set to an inter eye distance of 90 pixels. If a picture deviates from these minimum requirements, it is no longer accepted in countries that comply with ICAO recommendations to produce a passport or other machine readable travel documents (e.g., citizen cards). Furthermore, the image quality may be affected by compression of the image. In the case of lossy compression, the storage of high-frequency information is deliberately omitted in order to increase the compression rate. At high compression rates, however, this can lead to elimination of details and compression artefacts in the image. Since poor image quality usually results from lack of information, for example, too few pixels or too little high-frequency information, it is often difficult to improve the quality later.

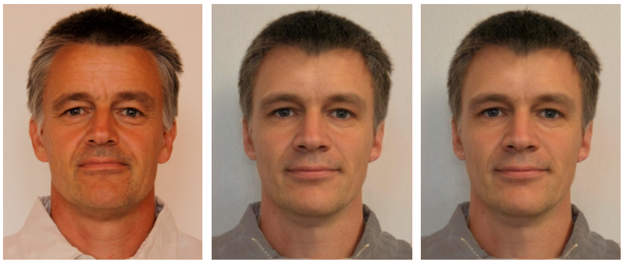

(a) BRISQUE: 21 (b) BRISQUE: 29 (c) BRISQUE: 50

FIGURE 6: Examples of BRISQUE scores for quality estimation (low values indicate high image quality). The BRISQUE score of bona fide (a) and uncompressed morphed images (b) are close to each other, the score of a JPEG compressed morphed image (c) is noticeably higher.

FIGURE 7: Comparison between a manually created high quality (left) and an automatically created low quality face morph (right).

Quality metrics for images can be used to objectively evaluate the output images based on quality measures derived from the signal. Since no reference image is available in the evaluation of the output image, the classical image quality determination methods, such as signal-to-noise ratio or mean square deviation, are not feasible. For the selection of the quality metric, the quality properties to be considered have to be determined. The metric proposed by Farias and Mitra [53] evaluates the occurrence of image artefacts, such as block artefacts, blur or noise. If the authentic appearance of a submitted passport image is to be evaluated for the human observer, then metrics are recommended that take into account the human perception, i.e., factors like sharpness [54] or perceptual quality [55] of the image. Another option is the automated assessment of the naturalness of the image using some no-reference image quality metrics, e.g., Blind / Referenceless Image Spatial Quality Evaluator (BRISQUE) [56]. Fig. 6 shows examples of BRISQUE values where low values indicated high quality and vice versa. On the left a non-morphed face image is shown, the associated BRISQUE value of 21 corresponds to a high quality. The middle image is a high quality morph without compression, the BRISQUE value is slightly worse. The image on the right shows the same morph with JPEG compression. Even if no artefacts are visible, the BRISQUE value is strongly influenced by the compression.

B. MORPHING ARTEFACTS

Morphing artefacts as illustrated in Fig. 7 (right) can appear in the image during the multi-step morph process. Within landmark-based methods artefacts are usually caused by the absence or misplacement of landmarks. As a result, the corresponding image areas are not transformed correctly so that they do not completely overlap. This creates shadow-like, semi-transparent areas, so-called ghost artefacts. Fig. 7 depicts a manual morphed face image and an automatically generated morph comprising said artefacts. On the right, one can see a morphed facial image with poorly placed landmarks. Especially, in the region of the neck, but also on the hair and ears, strong ghost artefacts can be observed. The iris proved to be particularly susceptible to artefacts because algorithms for automatic landmark determination are usually not able to provide the iris with correct landmarks. As a workaround, the located left and right eye corner could approximate the iris center half way between the two corners. Furthermore, shadow effects may occur in facial hair (e.g., beards and eyelashes), in differently pigmented areas (e.g., liver spots, tattoos), or by glasses and jewellery. Morph artefacts, which are caused by landmark-based morphing, can usually be remedied by manual post-processing in image processing programs as shown by Ferrara et al. [6]. An additional cause of artefacts may be the differences in the source images or inappropriate interpolation methods, which can lead to unnatural colour gradients and overly hard edges in the target images. Further artefacts induced by morphing may be low contrast and blur of the images, which may result from the averaging and interpolation of pixel positions and colour values. Another type of morph artefact may be generated using machine learning to create the morphed facial images. Due to the opacity of the process of the training algorithms, the errors might be difficult to narrow down or classify. Some of the potential mistakes are missing or deformed facial features, blurred areas and ghost artefacts. The emergence of such artefacts can be reduced by appropriate learning methods and a large number of training data. Due to the high agility of the relevant research area, a rapid improvement in the quality of morph images that can be achieved by the application of machine learning can also be expected.

C. PLAUSIBILITY OF FACE MORPHS

The quality of a morph can also be assessed by how plausible the image appears as a facial image. Here, on the one hand, the natural appearance of the produced image plays a role, and on the other hand, the similarity of the morph with the contributing data subject. The natural appearance can be adversely affected by strong artefacts. In addition, the

TABLE 1: Overview of publicly available morphing tools.

Developer	Software	Platform	Method	Automatic	Manual Effort	Required Skills	Parameters	Expected Quality
Commercial Software								
Morpheus	Photo Morpher	Win 7 / MacOS	landmarks	no	medium	positioning of landmarks	α	no outer region minor shadow artefacts
DebugMode	WinMorph	Win 7	probably landmarks	no	medium	positioning of landmarks	α	minor shadow artefacts issues in hair regions
Abrasoft	FantaMorph	Win 7-10 MacOS	landmarks	landmark detection	low	no	α	high quality for manual morphs minor shadow artefacts
Luxand Inc.	FaceMorpher	Win 7-10	landmarks	landmark detection	low	no	α	shadow artefacts blurry
Adobe	After Effects + RE:flex	Win 7-10 MacOS	lines	no	very high	operate Adobe After Effects	α_b, α_w	very high quality minor shadow artefacts
Adobe	Photoshop + Morph Animation	Win 7-10 MacOS	landmarks	rough shape	high	operate Adobe Photoshop	α_b, α_w	high quality minor shadow artefacts
PiVi & Co.	MixBooth	Android / iOS	swapping	no	low	no	no	low resolution unrealistic morphs
Moment Media	FaceFusion	iOS	landmarks	morph process	very low	no	no	limited by landmark detection issues e.g., for pupils
Open Source Software								
The blender project	blender	Win 7-10 MacOS Linux	manual mesh warping	via plugins	high	operate blender	inf.	nearly faultless (manual morphing)
OpenCV team	OpenCV	Win 7-10 MacOS Linux	landmarks + triangulation + warping	full automatic	implementation	Python commandline OpenCV	inf.	limited by landmark detection issues e.g., for pupils
Alyssa Quek	Face Morpher	MacOS / Linux	landmarks + triangulation + warping	full automatic	low	Python commandline	α, blur	good quality no outer region
The GIMP-Team	GIMP + GAP	Win 7-10 MacOS Linux	landmarks + triangulation + warping	automation via API	medium (if manual)	positioning of landmarks (if manual)	α	good quality easy to postprocess
Atsushi Nitanda	VAEGAN	Theano + Python	DNN	full automatic	very low	Theano + Python	no	deep learning artefacts low resolution
Michael Gourlay	gtkmorph	Win 7-10 MacOS Linux	landmarks + mesh warping	no	medium	positioning of landmarks	α	very detailed partly too sharp

similarity of the contributing subjects, e.g., with respect to gender, ethnicity or age group, influences the plausibility of the resulting morph. e.g., the morph depicted in Fig.1 appears less plausible since the age gap between the two contributing subjects is more than 20 years. Thus, it is recommended to select similar subjects as a basis. An approach for an automatic selection of suitable subjects is given in [57].

D. HUMAN PERCEPTION OF MORPHED FACE IMAGES

The issue of morphed face images in face comparison scenarios (e.g., border control) does not only affect automated face recognition systems, but also human observers. In general, humans are rather weak in recognizing unfamiliar faces as reported by Megreya and Burton [58] and Bruce et al. [59], independent of comparing two face images or a face image to a live data subject [60]–[62]. In particular, for border control scenarios, it is of relevance, that the difficulty to successfully verify a subject against its reference face image increases with the age of the taken image [63]. Depending on the individual, the face comparison capabilities vary. Hereby it is not relevant, if the human examiner is a border guard or an untrained student, the ratio of false negative to false positive remains the same [64], thus, it is uncertain whether a human expert can effectively detect morphed face images unless he is explicitly trained on morphing attacks. Recently, it has been shown that training makes a huge difference for a human observer. In [65], Robertson et al. showed, that without the knowledge of the morphing issues, a human observer would accept 68% of morphed images created with an α factor of 0.5. After a briefing, the false acceptance rate of morphed images dropped as low as 21%. Further, examiners that are better in distinguishing faces have a higher success chance to detected morphed face image [11]. Another parameter to consider is the weight (α) of the two subjects contributing to the morphed face image which represents a key factor in morphing attack scenario [66]. The role of the two subjects could be asymmetric, since the accomplice has to fool a human examiner, e.g., at the passport application office, and the criminal must fool the face verification algorithm, e.g., at an ABC gate. A higher weight of the accomplice is expected to hamper a successful detection of the morphed face image by a human examiner during presentation at enrolment, e.g., at the time of the passport issuance.

IV. MORPHING SOFTWARE

Table 1 lists available proprietary/open source morphing software and their properties. Applications were considered for the common desktop operating systems (Windows, Linux, Mac) and mobile operating systems (Android, iOS). Excluded from the list are web services available on the internet. These web services provide an easy way to manually create morphed images. However, firstly, an automated generation of face morphs is difficult and secondly, it is unclear how the uploaded images are processed and stored, which would make it impossible for researchers to upload face images of their models/volunteers and to comply with privacy regulations at the same point in time.

In order to enable well-founded and efficient experiments, it is generally advisable to use applications that can produce morphs in an automated manner in good quality without manual post-processing. Open source algorithms have the advantage that they can be much better automated and adapted to the needs than commercial applications. For commercial programs, automation is generally more difficult to achieve.

A. MORPHING MORE THAN TWO FACE IMAGES

The procedures described above for morphing two face images are easily extended to any number of source images. The contributing images may be weighted similarly to the α-factor, each image having its own factor such that the sum of the factors is 1. The more images included in an equally weighted morph, the smaller the weights will be. The more subjects are contributing to the image, the higher is the risk of quality issues described in Sect. III. Furthermore, the morphing of more than two images can also be done iteratively in pairs, i.e. the morphs are used as source images for the next morph process. Generally, no difference is visually discernible between the morphs created by both methods, i.e., the difference between artefacts resulting from a direct or iterative morphing process is below the perception threshold of a human observer. For this reason, the representation of sample images is omitted.

V. METRICS FOR MORPHING ATTACK EVALUATIONS

Standardized metrics are vital to enable direct benchmarks and comparative assessments of proposed methods. Regarding the topic of face morphing attacks efforts to define evaluation metrics for morphing attack detection and vulnerability analysis have already been made, e.g., in [33], [52]. Metrics suggested by Scherhag et al. [52] are briefly summarized in the following subsections.

A. VULNERABILITY ASSESSMENT

In their well-established guidelines Mansfield and Wayman [67] recommended that all comparisons in a biometric system's evaluation should be uncorrelated. That is, the samples compared to the morphed face images should not be the same as the ones used for the morphing process since such a comparison would ignore the natural biometric variance.

Regarding evaluation metrics the Impostor Attack Presentation Match Rate ($IAPMR$) introduced in ISO/IEC 30107-3 on Presentation Attack Detection evaluation [9] represents a standardized metric for attack success evaluation:

> *IAPMR*: in a full-system evaluation of a verification system, the proportion of impostor attack presentations using the same Presentation Attack Instrument (PAI) species in which the target reference is matched.

However, for the evaluation of face morphing attacks, the aforementioned IAPMR metric presents some drawbacks, as a morphing attack might only be considered successful if all contributing subjects are successfully matched against the morphed face image. The comparison of a morphed sample to another independent sample of one contributing subject is referred to as *mated morph comparison*. Motivated by the ISO/IEC 30107-3 [9], the impact of a morphing attack in a full-system evaluation is referred to as Mated Morph Presentation Match Rate ($MMPMR$) as introduced in [52].

As the morphing attack succeeds if all contributing subjects are verified successfully, only the minimum (for similarity scores) or maximum (for dissimilarity scores) of all mated morph comparisons of one morphed sample are of interest. The $MMPMR$ for similarity scores is defined as:

$$MMPMR = \frac{1}{M} \cdot \sum_{m=1}^{M} \left\{ \left[\min_{n=1,\ldots,N_m} S_m^n \right] > \tau \right\}, \qquad (1)$$

where τ is the decision threshold, S_m^n is the mated morph comparison score of the n-th subject of morph m, M is the total number of morphed images and N_m the total number of subjects constituting to morph m. Decisions of human examiners could be integrated to the above equation to evaluate a scenario with human inspection in the loop. Further, Scherhag et al. [52] proposed adaptations of the metric for evaluations where multiple samples of one subject are compared to one morphed face image.

$MMPMR$, as well as $IAPMR$, are directly dependent on the threshold τ of the biometric system. In order to achieve a more generalized metric in relation to the False Non-Match Rate ($FNMR$) of the system, Scherhag et al. propose to compute the difference between $1 - FNMR$ and $MMPMR$ or $IAPMR$, respectively. The Relative Morph Match Rate ($RMMR$) is defined as follows:

$$\begin{aligned} RMMR(\tau) &= 1 + (MMPMR(\tau) - (1 - FNMR(\tau))) \\ &= 1 + (MMPMR(\tau) - TMR(\tau)). \end{aligned} \qquad (2)$$

Different relevant examples for combinations of score distributions, thresholds and resulting $RMMR$ values are depicted in Fig. 8.

Gomez-Barrero et al. [68], [69] proposed a theoretical framework to predict the vulnerability of biometric systems to attacks based on morphed biometric samples. Further, key factors which take a major influence on a system's vulnerability to such attacks have been identified, e.g., the shape of mated (genuine) and non-mated (impostor) score distributions or the False Match Rate (FMR) the system is operated at.

B. DETECTION PERFORMANCE REPORTING

Given multiple procedures for preparing morphed images and/ or multiple morph detectors these can be benchmarked

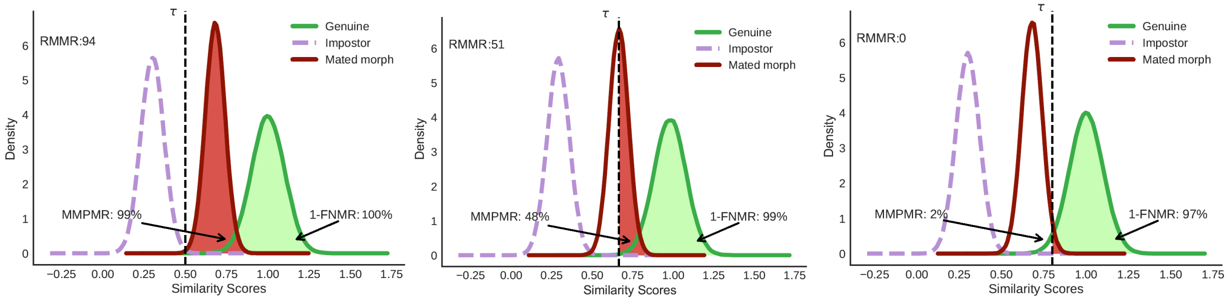

FIGURE 8: Behaviour of RMMR for different decision thresholds and score distributions.

employing metrics defined in [9], in particular, Attack Presentation Classification Error Rate ($APCER$) and Bona Fide Presentation Classification Error Rate ($BPCER$). The $APCER$ is defined as the proportion of attack presentations using the same presentation attack instrument species incorrectly classified as bona fide presentations in a specific scenario. The $BPCER$ is defined as the proportion of bona fide presentations incorrectly classified as presentation attacks in a specific scenario. Further, the $BPCER\text{-}10$ and $BPCER\text{-}20$ representing the operation points related to an $APCER$ of 10% and 5%, respectively, can be used to rank the tested morphing attack detection mechanisms. Additionally, it is recommend to plot the $BPCER$ over the $APCER$ in a Detection Error Tradeoff (DET) curve. In order to achieve reproducible and comparable performance evaluations of morphing attack detection systems, a common comprehension of the training and testing methodology is needed. In general, the standards defined in ISO/IEC 19795-1 on biometric performance testing and reporting [70] should be followed, e.g., a disjoint subdivision of the data into training and testing set. In particular a strict separation of the morphed samples with respect to the originating subjects is important, in order to avoid an unrealistic high detection performance. It should be noted, that one morphed sample is related to at least two subjects and each subject might contribute to several morphing samples.

VI. FACE MORPHING ATTACK DETECTION

Proposed approaches can be coarsely categorized with respect to the considered morphing attack detection scenario. The two classes of detection methods, i.e., *no-reference* and *differential*, are described in the following subsection. Subsequently, the state-of-the-art with respect to morph detection algorithms is surveyed.

A. DETECTION SCENARIOS

Two automated morph detection scenarios depicted in Fig. 9 can be distinguished:

- *No-reference morphing attack detection*: the detector processes a single image, e.g., an off-line authenticity check of an electronic travel document (this scenario is also referred to as single image morphing attack detection or forensic morphing attack detection);

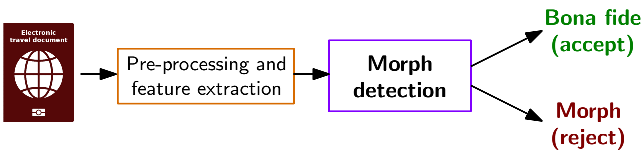

(a) no-reference morphing attack detection

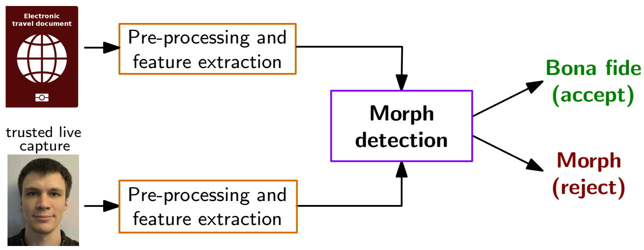

(b) differential morphing attack detection

FIGURE 9: Morphing attack detection scenarios.

- *Differential morphing attack detection*: a trusted live capture from an authentication attempt serves as additional source of information for the morph detector, e.g., during authentication at an ABC gate (this scenario is also referred to as image pair-based morphing attack detection). Note that all information extracted by no-reference morph detectors might as well be leveraged within this scenario [38].

B. STATE-OF-THE-ART

In the past years, numerous approaches to automated face morphing attack detection have been proposed. Published methods and their properties are summarized in Table 2. In some works, more than one system was presented, in such cases only those approaches, which were reported to reveal best morphing attack detection performance are listed. The majority of works assume the challenging no-reference scenario while some implement a differential morphing attack detection. Despite promising results reported in many works, a reliable detection of morphed face images still represents an open research challenge. It is important to note that the generalizability/robustness of published approaches has not been shown. So far, there are no publicly available large-scale databases of bona fide and morphed face images and no publicly available morph detection algorithms, which allow for a comprehensive experimental evaluation. Hence, the vast majority of methods has been mostly trained and tested on different in-house databases. In addition, face morph detection methods are mostly trained and tested on a single database using a single morph generation algorithm. Further, the likely appliance of image post-processing techniques by an attacker, e.g., image sharpening, is neglected in most works. Due to these facts, a comparison of published approaches in terms of reported detection performance would potentially be misleading and is purposely avoided in this survey. However, planned benchmark tests, e.g., by the National Institute of Standards and Technology (NIST) [71], are expected to facilitate a meaningful quantitative comparison of published approaches in the near future.

1) No-reference morphing attack detection

Several researchers have suggested the use of general purpose image descriptors, e.g., Local Binary Patterns (LBP) [102] or Binarized Statistical Image Features (BSIF) [103], which have been employed widely for biometric recognition. Ramachandra et al. [14] proposed a no-reference detection system based on a Support Vector Machine (SVM) trained on extracted BSIF-features of grayscale images. For training and evaluation of the SVMs an in-house database of morphed face images was created. On a derivate version of the same database, Scherhag et al. [36] investigated the accuracy of morphing detection on printed and scanned images employing the proposed algorithm. Further, a Probabilistic Collaborative Representation Classifier (ProCRC) [104] trained on LBP-feature extracted from the colour channels was proposed in [72]. As database an in-house database based on FRGCv2 [73] was used. The authors focus on the differences between morphed and averaged images in the evaluation. In [48] the suitability of LBP features for the detection of morphs generated by Generative Adversarial Networks (GANs) was tested.

The features extracted by texture descriptors can be further processed. A more complex method for morphing detection is proposed in [75], [76], where a Vietoris–Rips complex is built of the responses of uniform LBP extractors on the image. In [100], a high detection performance was shown for a linear SVM trained on high-dimensional LBP features [105] extracted from the FEI database [?]. Agarwal et al. [74] propose to train an SVM with Weighted Local Magnitude Pattern. Similar to LBP, the proposed descriptor encodes the differences between a center pixel and it's neighbors. However, instead of binarizing them, it assigns the weights inversely in proportion to the difference from the center pixel. Depending on the feature representation of texture descriptors the inputs of classifiers have to be adapted. E.g., for Scale-Invariant Feature Transform (SIFT) [106] the number of extracted keypoints has been shown to be suitable for the task of morph detection [38], [78]. A score-level fusion of multiple image descriptors might even improve the detection rate [79]. Therefore, LBP, BSIF, SIFT, Speeded Up Robust Features (SURF) [107], Histogram of Oriented Gradients (HOG) [108] and the deep features of Openface [109] were fused and evaluated in [79].

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. In particular the combination of features reflecting different information, e.g., LBP and SIFT, leads to improvements. It has been shown that the performance of morph detectors based on general purpose image descriptors might significantly decrease if training and test images stem from a different source, i.e., face database [37], [82]. In order to adapt the no-reference general purpose image descriptors a differential scenario, differences between feature vectors can (additionally) be employed [38].

During the morphing process, not only the texture, but the whole signal of the image is manipulated. Thus, a further detection approach is to analyze the changes in noise patterns, e.g., Photo Response Non-Uniformity (PRNU) [84]. Therefore, the PRNU-patterns, that are originating from the camera and which are distinct not only for each model but for each single camera, are extracted from a face image, the discrete Fourier magnitudes are computed. Subsequently, the mean and variance are derived from the resulting histogram. A very similar approach was presented in [86]. Recently, an improved version of this scheme based on PRNU variance analysis across image blocks was proposed in [85]. Morphing attack detection methods based on continuous image degradation have been proposed in [78], [110], [111]. 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. Ramachandra et al. [89] proposes the analysis of high frequencies in grayscale images. Therefore, the images are converted to grayscale according their luminance, a steerable pyramid is build and a Collaborative Representation Classifier (CRC) is trained on the high frequencies. The employed database was printed and scanned, but no further post-processing was tested. An alternative to handcrafted feature extractors is to employ statistical machine learning on the unprocessed image in order to distinguish between morphed and bona fide images. Ramachandra et al. [94] proposed to adapt two CNNs (VGG19 [112] and AlexNet [113]) by transfer-learning and combine the intermediate features to train a CRC. In [35], three CNNs, namely VGG19, AlexNet and GoogLeNet [114], are benchmarked as pre-trained and non-pre-trained models regarding their morph detection capabilities. Again, with these methods there is a potential 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, as described in Sect. III-B. As an attempt to avoid overfitting, Seibold et al. [95] trained a VGG19-net on a set of diverse images with two different databases, morphing algorithms and post-processings (mo-

TABLE 2: Overview of published morph detection algorithms.

Publication	Approach	Scenario	Morph Algorithms	Source Face Database	Post-Processing	Remarks
[14]	BSIF + SVM	no-reference	GIMP/GAP	in-house	-	-
[36]	BSIF + SVM	no-reference	GIMP/GAP	in-house	print and scan	fixed database of [14]
[72]	multi-channel-LBP + Pro-CRC	no-reference	OpenCV	FRGCv2 [73]	print and scan	-
[74]	WLMP + SVM	no-reference	Snapchat	in-house	-	-
[75], [76]	ULBP + RIPS + KNN	no-reference	[32]	Utrecht [77]	-	-
[78]	image degradation	no-reference	triangulation + blending (+ swapping)	in-house, Utrecht [77]	-	-
[38]	BSIF + SVM	no-reference, differential	triangulation + blending	FRGCv2 [73]	-	-
[79]	general purpose image descriptors + score-level fusion	no-reference	triangulation + blending	FRGCv2 [73]	-	-
[37]	HOG + SVM	no-reference	triangulation + blending	FRGCv2 [73], FERET [80], ARface [81]	-	cross database performance evaluation
[82]	LBP + SVM	no-reference	triangulation + blending	FRGCv2 [73], FERET [80]	-	cross database performance evaluation
[48]	LBP + SVM	no-reference	MorGan [48]	CelebA [83]	-	-
[84], [85]	PRNU analysis	no-reference	triangulation + blending	FRGCv2 [73]	hist. equalization, scaling, sharpening	-
[86]	SPN analysis	no-reference	triangulation + blending (+ swapping)	Utrecht [77], FEI [87]	-	-
[32]	double-compression artefacts	no-reference	triangulation + blending (+ swapping)	Utrecht [77], FEI [87]	-	-
[33]	double-compression artefacts	no-reference	[32]	Utrecht [77], FEI [87]	-	-
[88]	reflection analysis	no-reference	triangulation + blending (+ swapping)	in-house	-	-
[89]	luminance component +steerable pyramid + ProCRC	no-reference	unclear	[72] extended	print and scan	-
[90]	landmark angles	differential	OpenCV	ARface [81]	-	-
[91]	Demorphing	differential	GIMP/GAP	ARface [81]	-	-
[92]	Demorphing	differential	GIMP/GAP	ARface [81], CAS-PEAL-R1 [93]	-	CAS-PEAL-R1 contains images with pose variations
[94]	VGG19 + AlexNet + ProCRC	no-reference	[36]	in-house	print and scan	-
[95]	VGG19	no-reference	triangulation + blending (+ swapping)	BU-4DFE [96], CFD [97], FEI [87], FERET [80], PUT [98], scFace [99], Utrecht [77], in-house	motion blur, Gaussian blur, salt-and-pepper noise, Gaussian noise	trained on all combinations (no unseen attack classes)
[100]	high-dim. LBP + SVM	no-reference	triangulation + blending + swapping	Multi-PIE [101]	-	-

tion blur, Gaussian blur, salt-and-pepper noise, Gaussian noise). To avoid a focusing of the CNN on specific regions, images with specific regions covered (eyes, nose, mouth) were added to the training set. As the CNN was trained on all kind of databases, morphing algorithms and post-processings a statement about the resulting robustness of the classifier is difficult. Wandzik et al. [100] proposed to employ pre-trained face recognition networks, e.g. VGG-Face [4] or FaceNet [3], for morphing attack detection. The high-level features generated by the networks are classified using a linear SVM.

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 [32], [33]. 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 ICAO suggests face image data to be stored in accordance with the specifications established by the International Standard ISO/IEC 19794-5 [115]. More specifically, the ICAO requires face images to be stored in electronic travel documents at an average compressed sizes of 15kB to 20kB in JPEG or JPEG 2000 format [13]. However, JPEG 2000 is the de-facto-standard for electronic travel documents, as it maintains a higher quality when compressing face images to 15 kB. Hence, depending on the image size and the employed compression algorithm the detection of JPEG double-compression artefacts might not be feasible. In [88], 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/IEC standard 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 [115].

2) Differential morphing attack detection

Morphing detection algorithms based on general purpose image descriptors, signal or quality analysis are mostly no-reference algorithms, but can be adapted to differential morphing attack detection scenarios. However, there are some algorithms, that can solely be used in differential scenarios, as they require a trusted live capture. In [90], a morph detection algorithm based on landmark positions and angles is introduced. Therefore, the landmarks between both, the passport image and the trusted live capture are determined, the angle between all combinations of landmarks per image are computed and compared over both images. Due to the high intra-class variance of landmarks, the detection performance of this algorithm is rather moderate.

Another differential morph detection method referred to as de-morphing was proposed by Ferrara et al. [91]. In this approach a trusted live capture is aligned to a potential morph and "subtracted" from it in the image domain by applying a reverse morphing operation. The resulting image is then compared against the trusted live capture. The assumption is, that, if two subjects are morphed into one image, and one of the subjects is subtracted, the second subject remains. If there is only one subject in the image, this subject will remain after the subtraction. Thus, a morph is detected if the biometric decision changes from "accept" to "reject" when using the de-morphed image as reference. Robustness of de-morphing against slight face pose variations has been confirmed in [92]. Nevertheless, the authors indicate that in an ABC scenario the performance of de-morphing might degrade due to potential variations of quality and environmental conditions.

VII. ISSUES AND CHALLENGES

Several open issues and challenges exist in research related to face morphing and face morphing attack detection. The most relevant issues and challenges, which have already been pointed out throughout this survey, can be briefly summarized as follows:

- *Quality*: the automated generation of high-quality face morphs remains a challenging issue and of utmost importance in order to enable statistically significant testing of developed morphing attack detection methods under realistic conditions, see Sect. III.
- *Comparability/benchmarks*: the lack of publicly available large-scale databases comprising bona fide as well as morphed face images and open-source face morphing attack detection software prevents from a meaningful comparative benchmark of the current state-of-the-art in this field, see Sect. V.
- *Result reporting*: while first efforts have been made to apply standardized metrics for reporting the performance of morphing attack detection mechanisms equivalent measures for the vulnerability of face recognition systems w.r.t morphing attacks are non-existent; however, these would be vital in order to enable an unambiguous comparisons of proposed approaches, see Sect. V.
- *Over-fitting/robustness analysis*: like any other image-based classification task, approaches to morphing attack detection are prone to overfitting, i.e., rigorous evaluations including face morphs from unseen databases created by unseen morphing techniques are necessary, see Sect. VI.
- *Print-scan databases*: to simulate real-world scenarios where potentially morphed portrait images are printed and scanned, publicly available large-scale databases of printed and scanned bona fide and morphed face images are required, see Sect. VI.

VIII. CONCLUSION

This survey provides a comprehensive overview of published literature in the field of (face) image morphing and face morphing attack detection as well as a detailed discussion of open issues and challenges. The research in this important field is only in its infancy while not being limited to face recognition systems. The feasibility of morphing biometric samples has also been shown for other biometric characteristics, e.g. fingerprint [46], [116] or iris [47], which might as well be morphed in feature domain. The possibility of morphing biometric features and subsequently reconstructing a biometric sample from morphed feature vectors underlines the importance of data protection mechanisms, i.e. biometric template protection [117], [118] or conventional cryptographic techniques [119], [120]. Similar to face, for other characteristics certain aspects require more in-depth analysis, e.g., biometric quality estimation of (morphed) fingerprint [121], [122] or iris samples [123], [124], respectively. The reported face image morphing attack detection accuracy is yet not reflecting generalization to datasets incorporating the real world variety of capture conditions. This will change, once benchmark portals such as the NIST Face Recognition Vendor Test (FRVT) MORPH competition [71] are established. Nevertheless, robust algorithms must also anticipate the large variety of image post-processing as well as printing and scanning technology that could be used in the governmental procedures for the application of electronic travel documents. Morphing attack detection mechanisms that are robust against all those factors, will require a significant amount of future research.

ACKNOWLEDGMENT

This work was partially supported by the German Federal Ministry of Education and Research (BMBF), by the Hessen State Ministry for Higher Education, Research and the Arts (HMWK) within the Center for Research in Security and Privacy (CRISP) as well as by the Federal Office of Information Security (BSI) within the FACETRUST project.

REFERENCES

[1] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld, "Face recognition," ACM Computing Surveys, vol. 35, no. 4, pp. 399–458, dec 2003.
[2] S. Z. Li and A. K. Jain, Eds., Handbook of Face Recognition. Springer London, 2011.

[3] F. Schroff, D. Kalenichenko, and J. Philbin, “FaceNet: A unified embedding for face recognition and clustering,” in Proceedings of the 2015 Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, jun 2015.
[4] O. M. Parkhi, A. Vedaldi, and A. Zisserman, “Deep face recognition,” in Procedings of the British Machine Vision Conference 2015. British Machine Vision Association, 2015.
[5] 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, jan 2018.
[6] M. Ferrara, A. Franco, and D. Maltoni, “The magic passport,” in Proceedings of the 2014 International Joint Conference on Biometrics (IJCB). IEEE, sep 2014.
[7] G. Wolberg, “Image morphing: a survey,” The Visual Computer, vol. 14, no. 8-9, pp. 360–372, dec 1998.
[8] A. Patel and P. Lapsiwala, “Image morphing algorithm: A survey,” International Journal of Computer Applications (IJCA), vol. 5, no. 3, pp. 156–160, 2015.
[9] ISO/IEC JTC1 SC37 Biometrics, “Information technology – biometric presentation attack detection – part 3: Testing and reporting,” International Organization for Standardization, Geneva, Switzerland, ISO ISO/IEC IS 30107-3:2017, 2017.
[10] M. Ferrara, A. Franco, and D. Maltoni, “On the effects of image alterations on face recognition accuracy,” in Face Recognition Across the Imaging Spectrum. Springer International Publishing, 2016, pp. 195–222.
[11] D. J. Robertson, A. Mungall, D. G. Watson, K. A. Wade, S. J. Nightingale, and S. Butler, “Detecting morphed passport photos: a training and individual differences approach,” Cognitive Research: Principles and Implications, vol. 3, no. 1, jun 2018.
[12] A. Makrushin and A. Wolf, “An overview of recent advances in assessing and mitigating the face morphing attack,” in Proceedings of the 26th European Signal Processing Conference (EUSIPCO), 2018.
[13] International Civil Aviation Organization, “ICAO doc 9303, machine readable travel documents – part 9: Deployment of biometric identification and electronic storage of data in MRTDs (7th edition),” ICAO, Tech. Rep., 2015.
[14] R. Ramachandra, K. B. Raja, and C. Busch, “Detecting morphed face images,” in Proceedings of the 8th International Conference on Biometrics Theory, Applications and Systems (BTAS). IEEE, sep 2016.
[15] I. Craw, D. Tock, and A. Bennett, “Finding face features,” in Computer Vision – ECCV’92. Springer Berlin Heidelberg, 1992, pp. 92–96.
[16] T. Cootes, C. Taylor, D. Cooper, and J. Graham, “Active shape models-their training and application,” Computer Vision and Image Understanding, vol. 61, no. 1, pp. 38–59, jan 1995.
[17] L. Wiskott, J.-M. Fellous, N. Krüger, and C. von der Malsburg, “Face recognition by elastic bunch graph matching,” in Computer Analysis of Images and Patterns. Springer Berlin Heidelberg, 1997, pp. 456–463.
[18] T. Cootes, G. Edwards, and C. Taylor, “Active appearance models,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 23, no. 6, pp. 681–685, jun 2001.
[19] V. Zanella and O. Fuentes, “An approach to automatic morphing of face images in frontal view,” in MICAI 2004: Advances in Artificial Intelligence. Springer Berlin Heidelberg, 2004, pp. 679–687.
[20] J. M. Saragih, S. Lucey, and J. F. Cohn, “Face alignment through subspace constrained mean-shifts,” in Proceedings of the 12th International Conference on Computer Vision (ICCV). IEEE, sep 2009.
[21] X. Zhu and D. Ramanan, “Face detection, pose estimation, and landmark localization in the wild,” in Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, jun 2012.
[22] V. Kazemi and J. Sullivan, “One millisecond face alignment with an ensemble of regression trees,” in Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, jun 2014.
[23] D. E. King, “Dlib-ml: A machine learning toolkit,” J. Mach. Learn. Res., vol. 10, pp. 1755–1758, Dec. 2009.
[24] O. Çeliktutan, S. Ulukaya, and B. Sankur, “A comparative study of face landmarking techniques,” EURASIP Journal on Image and Video Processing, vol. 2013, no. 1, mar 2013.
[25] D. Ruprecht and H. Muller, “Image warping with scattered data interpolation,” IEEE Computer Graphics and Applications, vol. 15, no. 2, pp. 37–43, mar 1995.
[26] T. W. Sederberg and S. R. Parry, “Free-form deformation of solid geometric models,” in Proceedings of the 13th annual conference on Computer graphics and interactive techniques – SIGGRAPH ’86. ACM Press, 1986.
[27] S.-Y. Lee, K.-Y. Chwa, and S. Y. Shin, “Image metamorphosis using snakes and free-form deformations,” in Proceedings of the 22nd annual conference on Computer graphics and interactive techniques - SIGGRAPH ’95. ACM Press, 1995.
[28] T. Beier and S. Neely, “Feature-based image metamorphosis,” in Proceedings of the 19th annual conference on Computer graphics and interactive techniques - SIGGRAPH ’92. ACM Press, 1992.
[29] S. Schaefer, T. McPhail, and J. Warren, “Image deformation using moving least squares,” in ACM SIGGRAPH 2006 Papers on - SIGGRAPH ’06. ACM Press, 2006.
[30] D. W. Choi and C. J. Hwang, “Image morphing using mass-spring system,” in Proceedings of the International Conference on Computer Graphics and Virtual Reality (CGVR), 2011, p. 1.
[31] J. Wu, “Face recognition jammer using image morphing,” Boston University, Tech. Rep., 2011.
[32] 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. SCITEPRESS - Science and Technology Publications, 2017.
[33] 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 Proceedings of the 5th International Workshop on Biometrics and Forensics (IWBF). IEEE, apr 2017.
[34] 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,” in Digital Forensics and Watermarking. Springer International Publishing, 2017, pp. 121–135.
[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, pp. 107–120.
[36] U. Scherhag, R. Ramachandra, K. B. Raja, M. Gomez-Barrero, C. Rathgeb, and C. Busch, “On the vulnerability of face recognition systems towards morphed face attacks,” in Proceedings of the 5th International Workshop on Biometrics and Forensics (IWBF). IEEE, Apr. 2017.
[37] U. Scherhag, C. Rathgeb, and C. Busch, “Performance variation of morphed face image detection algorithms across different datasets,” in Proceedings of the 6th International Workshop on Biometrics and Forensics (IWBF). IEEE, Jun. 2018.
[38] ——, “Towards detection of morphed face images in electronic travel documents,” in Proceedings of the 13th IAPR Workshop on Document Analysis Systems (DAS), 2018.
[39] J. Liao, R. S. Lima, D. Nehab, H. Hoppe, P. V. Sander, and J. Yu, “Automating image morphing using structural similarity on a halfway domain,” ACM Transactions on Graphics, vol. 33, no. 5, pp. 1–12, sep 2014.
[40] E. Wu and F. Liu, “Robust image metamorphosis immune from ghost and blur,” The Visual Computer, vol. 29, no. 4, pp. 311–321, sep 2012.
[41] S. M. Seitz and C. R. Dyer, “View morphing,” in Proceedings of the 23rd annual conference on Computer graphics and interactive techniques - SIGGRAPH ’96. ACM Press, 1996.
[42] F. Yang, E. Shechtman, J. Wang, L. Bourdev, and D. Metaxas, “Face morphing using 3d-aware appearance optimization,” in Proceedings of Graphics Interface 2012, ser. GI ’12. Toronto, Ont., Canada, Canada: Canadian Information Processing Society, 2012, pp. 93–99.
[43] M. Bichsel, “Automatic interpolation and recognition of face images by morphing,” in Proceedings of the Second International Conference on Automatic Face and Gesture Recognition. IEEE Comput. Soc. Press, 1996.
[44] E. Shechtman, A. Rav-Acha, M. Irani, and S. Seitz, “Regenerative morphing,” in Proceedings of the 2010 Computer Society Conference on Computer Vision and Pattern Recognition. IEEE, jun 2010.
[45] I. Korshunova, W. Shi, J. Dambre, and L. Theis, “Fast face-swap using convolutional neural networks,” in Proceedings of the 2017 International Conference on Computer Vision (ICCV). IEEE, oct 2017.
[46] M. Ferrara, R. Cappelli, and D. Maltoni, “On the feasibility of creating double-identity fingerprints,” IEEE Transactions on Information Forensics and Security, vol. 12, no. 4, pp. 892–900, apr 2017.

[47] C. Rathgeb and C. Busch, "On the feasibility of creating morphed iris-codes," in Proceedings of the 2017 International Joint Conference on Biometrics (IJCB). IEEE, oct 2017.

[48] N. Damer, Y. Wainakh, V. Boller, S. von den Berken, P. Terhörst, A. Braun, and A. Kuijper, "MorGAN: Recognition vulnerability and attack detectability of face morphing attacks created by generative adversarial network," in Proceedings of the 9th IEEE International Conference on Biometrics: Theory, Applications, and Systems (BTAS). IEEE, 2018.

[49] G. Mai, K. Cao, P. C. Yuen, and A. K. Jain, "On the reconstruction of face images from deep face templates," IEEE Transactions on Pattern Analysis and Machine Intelligence, 2018.

[50] D. Bitouk, N. Kumar, S. Dhillon, P. Belhumeur, and S. K. Nayar, "Face swapping: Automatically replacing faces in photographs," in ACM SIGGRAPH 2008 papers on - SIGGRAPH '08. ACM Press, 2008.

[51] Y. Weng, L. Wang, X. Li, M. Chai, and K. Zhou, "Hair interpolation for portrait morphing," Computer Graphics Forum, vol. 32, no. 7, pp. 79–84, oct 2013.

[52] U. Scherhag, A. Nautsch, C. Rathgeb, M. Gomez-Barrero, R. N. J. Veldhuis, L. Spreeuwers, M. Schils, D. Maltoni, P. Grother, S. Marcel, R. Breithaupt, R. Ramachandra, and C. Busch, "Biometric systems under morphing attacks: Assessment of morphing techniques and vulnerability reporting," in Proceedings of the 2017 International Conference of the Biometrics Special Interest Group (BIOSIG). IEEE, sep 2017.

[53] M. Farias and S. Mitra, "No-reference video quality metric based on artifact measurements," in Proceedings of the 2005 International Conference on Image Processing (ICIP). IEEE, 2005.

[54] E. Ong, W. Lin, Z. Lu, X. Yang, S. Yao, F. Pan, L. Jiang, and F. Moschetti, "A no-reference quality metric for measuring image blur," in Proceedings of the 7th International Symposium on Signal Processing and Its Applications (ISSPA). IEEE, 2003.

[55] Z. Wang, H. Sheikh, and A. Bovik, "No-reference perceptual quality assessment of JPEG compressed images," in Proceedings. International Conference on Image Processing. IEEE, 2002.

[56] A. Mittal, A. K. Moorthy, and A. C. Bovik, "No-reference image quality assessment in the spatial domain," IEEE Transactions on Image Processing, vol. 21, no. 12, pp. 4695–4708, dec 2012.

[57] J. P. Vyas, M. V. Joshi, and M. S. Raval, "Automatic target image detection for morphing," Journal of Visual Communication and Image Representation, vol. 27, pp. 28–43, feb 2015.

[58] A. M. Megreya and A. M. Burton, "Unfamiliar faces are not faces: Evidence from a matching task," Memory & Cognition, vol. 34, no. 4, pp. 865–876, jun 2006.

[59] V. Bruce, Z. Henderson, K. Greenwood, P. J. B. Hancock, A. M. Burton, and P. Miller, "Verification of face identities from images captured on video," Journal of Experimental Psychology: Applied, vol. 5, no. 4, pp. 339–360, 1999.

[60] R. Kemp, N. Towell, and G. Pike, "When seeing should not be believing: Photographs, credit cards and fraud," Applied Cognitive Psychology, vol. 11, no. 3, pp. 211–222, jun 1997.

[61] A. M. Megreya and A. M. Burton, "Matching faces to photographs: Poor performance in eyewitness memory (without the memory)." Journal of Experimental Psychology: Applied, vol. 14, no. 4, pp. 364–372, 2008.

[62] J. P. Davis and T. Valentine, "CCTV on trial: Matching video images with the defendant in the dock," Applied Cognitive Psychology, vol. 23, no. 4, pp. 482–505, may 2009.

[63] A. M. Megreya, A. Sandford, and A. M. Burton, "Matching face images taken on the same day or months apart: the limitations of photo ID," Applied Cognitive Psychology, vol. 27, no. 6, pp. 700–706, oct 2013.

[64] D. White, R. I. Kemp, R. Jenkins, M. Matheson, and A. M. Burton, "Passport officers' errors in face matching," PLoS ONE, vol. 9, no. 8, p. e103510, aug 2014.

[65] D. J. Robertson, R. S. S. Kramer, and A. M. Burton, "Fraudulent ID using face morphs: Experiments on human and automatic recognition," PLOS ONE, vol. 12, no. 3, p. e0173319, mar 2017.

[66] T. Jäger, K. H. Seiler, and A. Mecklinger, "Picture database of morphed faces (mofa)," Universit"at des Saarlands, Tech. Rep., 2005.

[67] A. J. Mansfield and J. L. Wayman, "Best practices in testing and reporting performance of biometric devices," Centre for Mathematics and Scientific Computing, Tech. Rep., 2002.

[68] M. Gomez-Barrero, C. Rathgeb, U. Scherhag, and C. Busch, "Is your biometric system robust to morphing attacks?" in Proceedings of the 5th International Workshop on Biometrics and Forensics (IWBF). IEEE, apr 2017.

[69] ——, "Predicting the vulnerability of biometric systems to attacks based on morphed biometric information," IET Biometrics, vol. 7, no. 4, pp. 333–341, jul 2018.

[70] ISO/IEC JTC1 SC37 Biometrics, "Information technology – biometric performance testing and reporting – part 1: Principles and framework," International Organization for Standardization, Geneva, Switzerland, ISO ISO/IEC 19795-1:2006, 2006.

[71] M. Ngan, P. Grother, and K. Hanaoka, "Performance of automated facial morph detection and morph resistant face recognition algorithms," National Institue of Standards and Technology (NIST), Tech. Rep., 2018.

[72] R. Ramachandra, K. Raja, S. Venkatesh, and C. Busch, "Face morphing versus face averaging: Vulnerability and detection," in Proceedings of the 2017 International Joint Conference on Biometrics (IJCB). IEEE, oct 2017.

[73] P. Phillips, P. Flynn, T. Scruggs, K. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek, "Overview of the face recognition grand challenge," in 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). IEEE, 2005.

[74] A. Agarwal, R. Singh, M. Vatsa, and A. Noore, "SWAPPED! digital face presentation attack detection via weighted local magnitude pattern," in Proceedings of the 2017 International Joint Conference on Biometrics (IJCB). IEEE, oct 2017.

[75] A. Asaad and S. Jassim, "Topological data analysis for image tampering detection," in Digital Forensics and Watermarking. Springer International Publishing, 2017, pp. 136–146.

[76] S. Jassim and A. Asaad, "Automatic detection of image morphing by topology-based analysis," in Proceedings of the 26th European Signal Processing Conference (EUSIPCO), 2018.

[77] "Utrecht ECVP," European Conference on Visual Perception, 2008.

[78] 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 Proceedings of the 5th ACM Workshop on Information Hiding and Multimedia Security - IHMMSec '17. ACM Press, 2017.

[79] U. Scherhag, C. Rathgeb, and C. Busch, "Morph detection from single face images: a multi-algorithm fusion approach," in Proceedings of the 2018 International Conference on Biometrics Engineering and Application (ICBEA). ACM, 2018.

[80] P. Phillips, H. Wechsler, J. Huang, and P. J. Rauss, "The FERET database and evaluation procedure for face-recognition algorithms," Image and Vision Computing, vol. 16, no. 5, pp. 295–306, apr 1998.

[81] A. Martinez and R. Benavente, "The AR face database," Computer Vision Center (CVC), Tech. Rep. 24, Jun. 1998.

[82] L. Spreeuwers, M. Schils, and R. Veldhuis, "Towards robust evaluation of face morphing detection," in Proceedings of the 26th European Signal Processing Conference (EUSIPCO), 2018.

[83] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in 2015 IEEE International Conference on Computer Vision (ICCV). IEEE, dec 2015, celebA.

[84] L. Debiasi, U. Scherhag, C. Rathgeb, A. Uhl, and C. Busch, "PRNU-based detection of morphed face images," in Proceedings of the 6th International Workshop on Biometrics and Forensics (IWBF). IEEE, 2018.

[85] L. Debiasi, C. Rathgeb, U. Scherhag, A. Uhl, and C. Busch, "PRNU variance analysis for morphed face image detection," in Proceedings of the 9th IEEE International Conference on Biometrics: Theory, Applications, and Systems (BTAS). IEEE, 2018.

[86] L.-B. Zhang, F. Peng, and M. Long, "Face morphing detection using fourier spectrum of sensor pattern noise," in 2018 IEEE International Conference on Multimedia and Expo (ICME). IEEE, jul 2018.

[87] C. E. Thomaz and G. A. Giraldi, "A new ranking method for principal components analysis and its application to face image analysis," Image and Vision Computing, vol. 28, no. 6, pp. 902–913, jun 2010.

[88] C. Seibold, A. Hilsmann, and P. Eisert, "Reflection analysis for face morphing attack detection," in Proceedings of the 26th European Signal Processing Conference (EUSIPCO), 2018.

[89] R. Ramachandra, S. Venkatesh, K. Raja, and C. Busch, "Detecting face morphing attacks with collaborative representation of steerable features," in Proceedings of the 3rd Computer Vision and Image Processing (CVIP2018), 2018, pp. 1–11.

[90] U. Scherhag, D. Budhrani, M. Gomez-Barrero, and C. Busch, "Detecting morphed face images using facial landmarks," in Proceedings of the 2018 International Conference on Image and Signal Processing (ICISP). Springer International Publishing, 2018, pp. 444–452.

[91] M. Ferrara, A. Franco, and D. Maltoni, "Face demorphing," IEEE Transactions on Information Forensics and Security, vol. 13, no. 4, pp. 1008–1017, apr 2018.
[92] ——, "Face demorphing in the presence of facial appearance variations," in Proceedings of the 26th European Signal Processing Conference (EUSIPCO), 2018.
[93] W. Gao, B. Cao, S. Shan, D. Zhou, X. Zhang, and D. Zhao, "The CAS-PEAL large-scale chinese face database and baseline evaluations," Chinese Academy of Sciences, Tech. Rep. JDL-TR-04-FR-001, May 2004.
[94] R. Ramachandra, K. B. Raja, S. Venkatesh, and C. Busch, "Transferable deep-CNN features for detecting digital and print-scanned morphed face images," in Proceedings of the 2017 Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). IEEE, jul 2017.
[95] C. Seibold, W. Samek, A. Hilsmann, and P. Eisert, "Accurate and robust neural networks for security related applications exampled by face morphing attacks," Computer Vision and Pattern Recognition, pp. 1–16, 2018.
[96] L. Yin, X. Wei, Y. Sun, J. Wang, and M. Rosato, "A 3d facial expression database for facial behavior research," in 7th International Conference on Automatic Face and Gesture Recognition (FGR06). IEEE, 2006.
[97] D. S. Ma, J. Correll, and B. Wittenbrink, "The chicago face database: A free stimulus set of faces and norming data," Behavior Research Methods, vol. 47, no. 4, pp. 1122–1135, jan 2015.
[98] A. Kasinski, A. Florek, and A. Schmidt, "The PUT face database," Image Processing & Communications, Jan. 2008.
[99] M. Grgic, K. Delac, and S. Grgic, "SCface surveillance cameras face database," Multimedia Tools and Applications, vol. 51, no. 3, pp. 863–879, oct 2009.
[100] L. Wandzik, G. Kaeding, and R. V. Garcia, "Morphing detection using a general-purpose face recognition system," in Proceedings of the 26th European Signal Processing Conference (EUSIPCO), 2018.
[101] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, "Multi-PIE," in 2008 8th IEEE International Conference on Automatic Face & Gesture Recognition. IEEE, sep 2008.
[102] T. Ojala, M. Pietikäinen, and D. Harwood, "A comparative study of texture measures with classification based on featured distributions," Pattern Recognition, vol. 29, no. 1, pp. 51–59, jan 1996.
[103] J. Kannala and E. Rahtu, "Bsif: Binarized statistical image features," in Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012), Nov 2012, pp. 1363–1366.
[104] S. Cai, L. Zhang, W. Zuo, and X. Feng, "A probabilistic collaborative representation based approach for pattern classification," in Proceedings of the 2016 Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, jun 2016.
[105] D. Chen, X. Cao, F. Wen, and J. Sun, "Blessing of dimensionality: High-dimensional feature and its efficient compression for face verification," in Proceedings of the 2013 Conference on Computer Vision and Pattern Recognition. IEEE, jun 2013.
[106] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," International Journal of Computer Vision, vol. 60, no. 2, pp. 91–110, nov 2004.
[107] H. Bay, A. Ess, T. Tuytelaars, and L. V. Gool, "Speeded-up robust features (SURF)," Computer Vision and Image Understanding, vol. 110, no. 3, pp. 346–359, jun 2008.
[108] C. Shu, X. Ding, and C. Fang, "Histogram of the oriented gradient for face recognition," Tsinghua Science and Technology, vol. 16, no. 2, pp. 216–224, apr 2011.
[109] B. Amos, B. Ludwiczuk, and M. Satyanarayanan, "Openface: A general-purpose face recognition library with mobile applications," School of Computer Science Carnegie Mellon University, Tech. Rep., 2016.
[110] T. Neubert, "Face morphing detection: An approach based on image degradation analysis," in Digital Forensics and Watermarking. Springer International Publishing, 2017, pp. 93–106.
[111] A. Makrushin, C. Kraetzer, T. Neubert, and J. Dittmann, "Generalized benfords law for blind detection of morphed face images," in Proceedings of the 6th ACM Workshop on Information Hiding and Multimedia Security - IH&MMSec '18. ACM Press, 2018.
[112] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," Computer Vision and Pattern Recognition, pp. 1–14.
[113] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," Communications of the ACM, vol. 60, no. 6, pp. 84–90, may 2017.
[114] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in Proceedings of the 2015 Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, jun 2015.
[115] ISO/IEC JTC1 SC37 Biometrics, Information technology – Biometric data interchange formats – Part 5: Face image data, 2005.
[116] A. Othman and A. Ross, "Mixing fingerprints for generating virtual identities," in Proceedings of the 2011 International Workshop on Information Forensics and Security (WIFS). IEEE, nov 2011.
[117] K. Nandakumar and A. K. Jain, "Biometric template protection: Bridging the performance gap between theory and practice," IEEE Signal Processing Magazine, vol. 32, no. 5, pp. 88–100, sep 2015.
[118] C. Rathgeb and A. Uhl, "A survey on biometric cryptosystems and cancelable biometrics," EURASIP Journal on Information Security, vol. 2011, no. 1, sep 2011.
[119] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone, Handbook of Applied Cryptography. Taylor & Francis Inc, 2001.
[120] B. Gupta, D. P. Agrawal, and S. Yamaguchi, Eds., Handbook of Research on Modern Cryptographic Solutions for Computer and Cyber Security. IGI Global, 2016.
[121] M. A. Alsmirat, F. Al-Alem, M. Al-Ayyoub, Y. Jararweh, and B. Gupta, "Impact of digital fingerprint image quality on the fingerprint recognition accuracy," Multimedia Tools and Applications, jan 2018.
[122] Y. Chen, S. C. Dass, and A. K. Jain, "Fingerprint quality indices for predicting authentication performance," in Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2005, pp. 160–170.
[123] N. D. Kalka, J. Zuo, N. A. Schmid, and B. Cukic, "Estimating and fusing quality factors for iris biometric images," IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans, vol. 40, no. 3, pp. 509–524, may 2010.
[124] ——, "Image quality assessment for iris biometric," in Biometric Technology for Human Identification III, P. J. Flynn and S. Pankanti, Eds. SPIE, apr 2006.