Chapter 2
Digital Face Manipulation in Biometric Systems

Mathias Ibsen, Christian Rathgeb, Daniel Fischer, Pawel Drozdowski, and Christoph Busch

Abstract Biometric technologies, in particular face recognition, are employed in many personal, commercial, and governmental identity management systems around the world. The processing of digitally manipulated face images within a face recognition system may lead to false decisions and thus decrease the reliability of the decision system. This necessitates the development of manipulation detection modules which can be seamlessly integrated into the processing chain of face recognition systems. This chapter discusses the impact of face image manipulation on face recognition technologies. To this end, the basic processes and key components of biometric systems are briefly introduced with particular emphasis on facial recognition. Additionally, face manipulation detection scenarios and concepts of how to integrate detection methods to face recognition systems are discussed. In an experimental evaluation, it is shown that different types of face manipulation, i.e. retouching, face morphing, and swapping, can significantly affect the biometric performance of face recognition systems and hence impair their security. Eventually, this chapter provides an outlook on issues and challenges that face manipulation poses to face recognition technologies.

M. Ibsen · C. Rathgeb (✉) · D. Fischer · P. Drozdowski · C. Busch
Hochschule Darmstadt, Darmstadt, Germany
e-mail: christian.rathgeb@h-da.de

M. Ibsen
e-mail: mathias.ibsen@h-da.de

D. Fischer
e-mail: daniel.fischer@h-da.de

P. Drozdowski
e-mail: pawel.drozdowski@h-da.de

C. Busch
e-mail: christoph.busch@h-da.de

C. Rathgeb et al. (eds.), *Handbook of Digital Face Manipulation and Detection*,
Advances in Computer Vision and Pattern Recognition,
https://doi.org/10.1007/978-3-030-87664-7_2

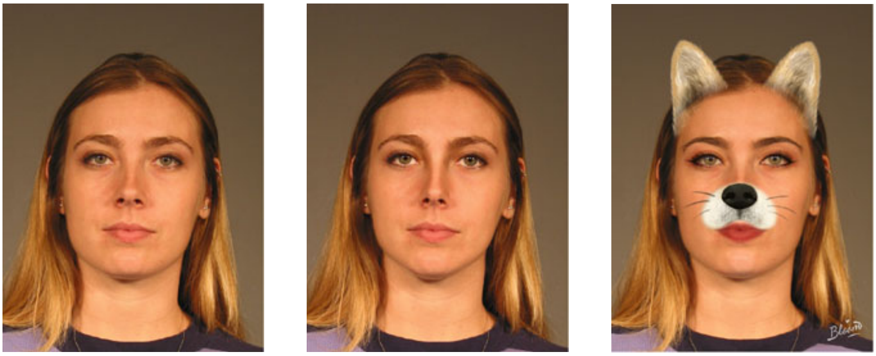

Fig. 2.1 Examples of digital face manipulation: original face image (left), a slightly retouched face image with increased eye size, slimmed nose, and cheeks (middle), and face image with a cat filter (right)

2.1 Introduction

The facial image of a subject can be altered, i.e. manipulated, in the digital domain such that the resulting digitally manipulated face image contains altered (biometric) features of the subject in a manipulated form. Digital face manipulation algorithms have advanced rapidly in recent years [50, 53]. In the scientific literature, numerous methods which can be used to alter facial images, e.g. swapping [40], morphing [43], or retouching [35], have been proposed for various application scenarios, e.g. in the film industry. Due to their popularity, face manipulation algorithms are already available in free web and mobile applications (apps). They typically allow their users to easily manipulate facial images or videos. Existing apps provide a huge variety of face manipulations ranging from funny filters to alterations in facial shape and texture; see Fig. 2.1 for examples.

Manipulated facial images which look realistic may lead to a loss of trust in digital content and can cause further harm by spreading false information [50]. Moreover, the automated processing of manipulated facial images may lead to false decisions, e.g. in a biometric system. For instance, face recognition performance might be impacted by the aforementioned manipulations. Face recognition technologies are employed for identity management in numerous application areas, e.g. mobile devices, access control, forensics, or surveillance [24, 54].

Face image manipulations might be applied for different reasons, e.g. beautification, by innocent users who have no intention of manipulating an image to impair the security of a face recognition system. However, they may also be applied by malicious users with the goal of interfering with the operation of a face recognition system. Such attacks are referred to as presentation attacks [19, 27]. Digital face image manipulation can be seen as presentation attacks in the digital domain.[1] Face recognition systems have been shown to be particularly vulnerable to presentation attacks [32], e.g. printouts of facial images or 3D masks. Presentation attacks are either performed with the aim of identity concealment, i.e. an attacker tries not to be recognized, or impersonation, i.e. an attacker tries to be recognized as somebody else (target subject). Researchers have already shown that both types of attacks are feasible with the help of digital face image manipulation [50]. In many cases, only slight alterations of original facial images are necessary to achieve alarmingly high attack success rates. This poses serious security risks to face recognition systems.

Recently, numerous methods for detecting facial image manipulations have been proposed, see [50, 53] for comprehensive surveys. Said manipulation detection methods can be applied in face recognition systems in order to protect against attacks based on manipulated face images. Moreover, detection methods may be specifically designed for integration into the processing chain of face recognition systems for different application scenarios.

This introductory chapter provides a brief overview of biometric face recognition. The potential impacts of the digital face image manipulation on facial recognition technologies are discussed, along with an empirical evaluation on a database comprising common digital face alterations using state-of-the-art face recognition systems. In addition, different face image manipulation detection scenarios and the integration of detection modules into biometric systems are described.

This chapter is organized as follows: Sect. 2.2 briefly introduces the key processes of generic biometric systems, in particular, face recognition. Subsequently, Sect. 2.3 discusses potential impacts of face image manipulation on the biometric performance of face recognition as well as detection scenarios. Experimental case studies are presented in Sect. 2.4. Finally, a summary and outlook are given in Sect. 2.5.

2.2 Biometric Systems

Biometric systems aim at establishing or verifying the identity or demographic attributes of individuals. In the international standard ISO/IEC 2382-37 [18], "biometrics" is defined as: "automated recognition of individuals based on their biological and behavioural characteristics." Humans possess, nearly universally, physiological characteristics which are highly distinctive and can, therefore, be used to distinguish between different individuals with a high degree of confidence. Prominent biometric characteristics are fingerprint, face, or iris. For a comprehensive introduction to biometrics, the interested reader is referred to [20] and the handbook series [4, 24, 26, 39, 51].

[1] In certain scenarios, digital face image manipulations can also be applied to perform presentation attacks at enrolment which may be referred to as backdoor attacks.

2.2.1 Processes

Generally, an automated biometric recognition system consists of: (1) a capture device (e.g. a camera), with which the biometric samples (e.g. facial images) are acquired; (2) a database which stores the biometric information and other personal data; (3) signal processing algorithms, which estimate the quality of the acquired sample, pre-process and extract the distinguishing features from it; and (4) comparison and decision algorithms, which enable ascertaining of similarity of two biometric samples by comparing the extracted feature vectors and establishing whether or not the two biometric samples belong to the same source.

During enrolment, a biometric capture device generates a reference sample of an individual, proceeds to pre-process it, and extracts a feature vector which is stored as a reference template. At the time of authentication, a probe sample is captured, processed in the same way, and the resulting probe template is compared against a reference template of a claimed identity (verification) or up to all stored reference templates (identification).[2] As a result, a (set of) biometric comparison score(s) is compared against a pre-defined threshold yielding acceptance or rejection decision. These processes are illustrated in Fig. 2.2.

In a biometric authentication attempt, two algorithmic errors may occur [17]:

- **False Match**: The comparison decision of "match" for a biometric probe and a biometric reference that belong to different biometric capture subjects.
- **False Non-Match**: The comparison decision of "non-match" for a biometric probe and a biometric reference that belong to the same biometric capture subject and of the same biometric characteristic.

The probabilities of each of these erroneous decision outcomes are defined as:

- **False Match Rate** (*FMR*): The proportion of the completed biometric non-mated comparison trials that result in a false match.
- **False Non-Match Rate** (*FNMR*): The proportion of the completed biometric mated comparison trials that result in a false non-match.

The *FMR* and the *FNMR* are measured at a certain decision threshold of the system. A change of the decision threshold usually results in a decrease of one of the error

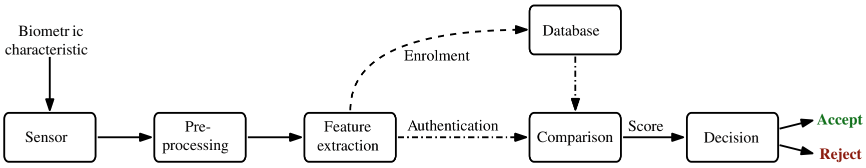

Fig. 2.2 Overview of a biometric recognition system

[2] This chapter focuses on biometric verification systems performing one-to-one comparisons.

rates at a cost of increasing the other. In other words, there exists a fundamental trade-off between system security (*FMR*) and convenience (*FNMR*) which is commonly illustrated by a detection error trade-off (DET) plot. The operation point where the *FMR* is equal to the *FNMR* is commonly referred to as Equal Error Rate (EER), which is often used as a single representative value to compare the biometric performance obtained by different algorithms.

2.2.2 Face Recognition

Face recognition systems are typically designed to process facial images captured with visible imaging sensor,[3] i.e. RGB colour cameras. In the pre-processing stage, face detection and face alignment is performed. Subsequently, face sample quality is estimated [46] and feature extraction is performed. For a long period of time, hand-crafted feature extractors, e.g. Local Binary Patterns [1] and Gabor filters [47], were predominately used. Said methods apply texture descriptors locally and aggregate extracted features into an overall face descriptor. A large variety of such systems has been proposed in the scientific literature, see [24, 25]. In contrast, current face recognition technologies utilize deep learning and massive training datasets to learn rich and compact representations of faces [15, 29]. The recent developments in Deep Convolutional Neural Networks (DCNNs) have led to breakthrough advances in facial recognition accuracy. DCNNs are usually trained using differentiable loss functions. A face embedding in the latent space is represented as a fixed-length real-valued vector. The dissimilarity of such feature vectors can be effectively estimated through simple distance measures, e.g. Euclidean distance. State-of-the-art face recognition systems have already surpassed human-level performance [34, 49] even on unconstrained (captured “in-the-wild” or with low image quality) face databases, e.g. the well-known Labeled Faces in the Wild (LFW) dataset [23].

2.3 Digital Face Manipulation in Biometric Systems

Advances in image manipulation software and machine learning technologies have made it easier to realistically manipulate face images. Some digital face manipulations are expected to impact the biometric performance of a face recognition system as they e.g. can cause severe changes in facial appearance or obscure parts of a face. Hence, methods capable of accurately detecting such manipulations are needed in order to mitigate their negative impacts on biometric systems. Development of such detection methods remains an open challenge.

[3] There are also face recognition systems which utilize other sensors, e.g. depth sensors for 3D face recognition.

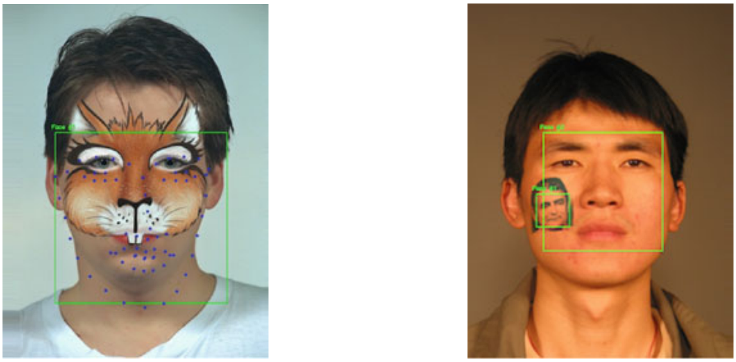

Fig. 2.3 Examples where the face of a digitally manipulated image is inaccurately detected

2.3.1 Impact on Biometric Performance

Digitally manipulated images may be introduced into a biometric system during enrolment or authentication and in systems where images are not captured live by a biometric capture device. It has been demonstrated that some manipulations (e.g. morphed images) [10, 41] can be used by attackers to circumvent the security of the system, whereas other manipulations usually carried out by bona fide users like slight retouching has little to no security implications [36]. While manipulated images can be a problem from a security-point-of-view, it can be of interest from a usability perspective, and in some applications of face recognition systems, that face recognition systems are robust to common manipulations in the digital domain. This can for instance, be relevant if images from social media are used in a face recognition system, as it is likely that users have manipulated the images without any intention of interfering with the operation of a face recognition system. Despite the intentions of digitally manipulating a face image, such images can impact different modules of a face recognition system if processed:

Face detection Digital face manipulations which occlude parts of a face or add additional texture information (e.g. synthetic tattoos) are likely to affect a face recognition system's ability to detect a face accurately. Facial manipulations can cause detection schemes to detect multiple faces or inaccurately determine the region of interest, i.e. the face (examples are given in Fig. 2.3). If a face cannot be properly detected, reliable recognition cannot be guaranteed.

Quality estimation It is expected that manipulations where part of a face is occluded, in general, will obtain a lower estimated face quality score than faces without occlusions. For other manipulations which aim at impersonation, it is not expected that the manipulation will have a significant effect on the quality score. If some types of digital manipulations receive a significantly lower quality score than bona fide images, quality estimation might be used to prevent the enrolment of such manipulated samples into the face recognition system database.

Comparison and feature extraction Digital facial manipulations are expected to impact the features extracted from facial images and affect the comparison scores of mated and non-mated comparison trials. The expected behaviour depends on the type of manipulation applied. For beautification and identity concealment, it is expected that the performance significantly drops when a face is severely manipulated or when occlusions occur over key areas such as the periocular region. For manipulations that aim at impersonation, it is expected that the manipulated image becomes more similar to the target identity than the source identity. Similarly, for manipulations that aim at merging multiple identities into a single image, it is expected that the similarity score is high for all individuals contributing to the merged image.

The impact of digital face manipulations on face recognition systems depends on the type and severity of the manipulation applied. For manipulations that only alter few aspects of a facial image e.g. lighting condition and slight beautification, it is not expected that the manipulation has a big impact as modern face recognition systems are robust to such minor changes. For manipulations where a large part of a face is occluded, the *FNMR* of the system is expected to be affected significantly. For manipulations where a face is swapped with another individual's face, it is expected that the swapped face image becomes less similar to the source identity and more similar to the target identity. For high-quality morphed images, it is expected that the system will falsely accept multiple individuals.

2.3.2 Manipulation Detection Scenarios

Several detection algorithms have been proposed to improve the robustness of face recognition systems to facial manipulations and to prevent image forgery. These algorithms can be integrated into the existing face recognition systems and be used to check the authenticity and integrity of imagery. In face recognition systems, detection algorithms can be used to prevent that facial images, which have been manipulated, are stored during enrolment or used during authentication.

For detecting manipulated face images, there are two different detection scenarios:

1. No-reference detection
2. Differential detection

In no-reference detection, a single suspected image is given as input to the detection algorithm and analyzed. Thereupon, a detection score is produced and used to determine whether the image is bona fide or manipulated. In differential detection, both the suspected image and a trusted live capture are used to determine if the suspected image has been manipulated. No-reference detection is considered a

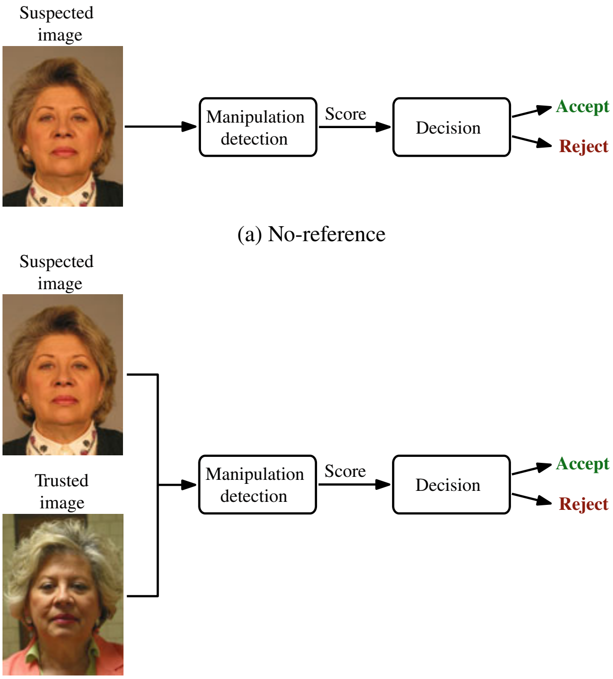

(a) No-reference

(b) Differential

Fig. 2.4 Categorisation of face manipulation detection schemes

more difficult problem and usually less accurate (see e.g. [28]). The possibility to use differential detection in face recognition systems arises due to the availability of pairs of images (reference and probe) during authentication, which is often not the case in traditional image forensics where usually only a single image or video is available. Despite the often superior performance of differential detection algorithms, no-reference detection is still important in forensic scenarios when a trusted live capture is not available. A conceptual overview of the two detection scenarios is shown in Fig. 2.4.

Several algorithms for detecting digital manipulated face images have been proposed e.g. [21, 33, 37, 40, 42, 45, 55]. In general, the existing manipulation detection schemes use (1) texture analysis, (2) digital forensics, or (3) deep-learning techniques to detect manipulated images. The use of texture descriptors has shown promising results e.g. for no-reference morphing attack detection as reported in [44]. Similarly, forensics-based detection methods, e.g. methods which analyze Photo Response Non-Uniformity (PRNU), have been shown to be useful for detecting some types of digital manipulations and have, for instance, been applied to detect retouched [36] and morphed images [7]. The features used in detection schemes based on digital forensics or texture analysis are often highly dependent on the training scenario and struggle to generalize well to unseen conditions and variations in post-processing.

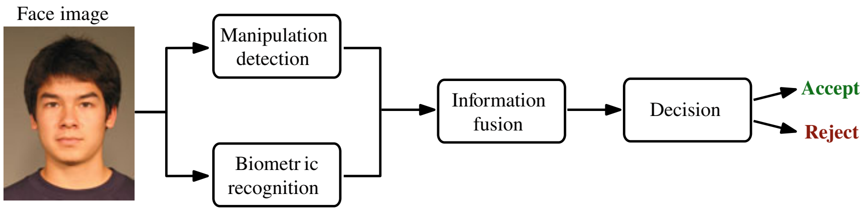

Fig. 2.5 Integration of manipulation detection and biometric recognition

Therefore, many state-of-the-art approaches utilize deep learning-based models, or features extracted from such models to detect manipulated face images. As indicated in [36] information fusion, like combining detection scores from multiple algorithms, can lead to a more robust detection system. For a more comprehensive overview of the current state-of-the-art in detecting manipulated face images, the reader is referred to [2, 35, 43, 50, 52] and to the third part of this handbook.

Figure 2.5 shows one possible integration of a manipulation detection algorithm into a face recognition system. As illustrated, the output of the manipulation detection and biometric recognition system can be fused together and used to make the final decision. Information fusion is usually based on either decision-level or score-level fusion. For decision-level fusion, the binary decision outputs of the manipulation detection and biometric system are used to determine the authentication output. For instance, a successful authentication output could be given only if neither of the systems rejects the input image. For score-level fusion, the scores produced by the two systems are combined and used together with a threshold value to determine the final authentication output. Chingovska et al. [5] investigated the impact of applying different score-level and decision-level fusion techniques for integrating a presentation attack detection algorithm with a biometric recognition system and concluded that there almost always was a trade-off between recognition and detection performance.

Another approach for making face recognition systems robust is to create algorithms capable of inverting the facial manipulations, i.e. remove the manipulation. Some authors have proposed algorithms capable of inverting specific manipulations, e.g. [11, 55].

2.4 Experiments

In this section, the vulnerability of two state-of-the-art face recognition systems towards three types of digital manipulations (retouching, morphing, and face swapping[4]) is evaluated.

[4] Face swapping is also some times referred to as identity swapping in the literature.

2.4.1 Experimental Setup

For the evaluation, one open-source and one commercial face recognition system are used. As the open-source system *ArcFace* [8] is used with the *MTCNN* face detector [56] for pre-processing. Due to terms of use, the used commercial system cannot be named explicitly and will henceforth be referred to as *COTS*.

To create an appropriate database, a subset of constrained facial images from the FERET [30] and FRGCv2 [31] database are manipulated using six different tools.

The changes imposed by the tools for the different manipulations are described below:

Retouching For the generation of retouched images, *InstaBeauty* [16] and *Fotorus* [14] are used. Both are proprietary software that offer features for beautifying facial images. Although the beautification operations performed by these and similar apps vary, common manipulations are smoothing of the skin, slimming of the nose, and enlargement of the eyes. Additionally, other manipulations might occur when beautifying an image e.g. enlargement of the mouth and slimming of the chin.

Morphing For the generation of morphs *FaceFusion* [9] and *UBO Morpher* [3, 12] are used. For *FaceFusion* a version which uses the landmarks of *dlib* [22] and Delaunay triangles is applied. Certain regions (e.g. eyes, nostrils, and hair) of the first face image are blended over the morph to hide potential artefacts. The *UBO Morpher* tool generates a morphed image by triangulation, warping, and blending. For finding landmarks for this tool, *dlib* are used. To avoid artefacts in the area outside of the face region, the morphed image generated by *UBO Morpher* is copied to the background of one of the original face images. Images generated by *UBO Morpher* might show artefacts at the border lines of the blended areas. In this evaluation and for both morphing tools, a single image is generated from the facial images of two different subjects and an equal weighting factor [12] of 0.5 is used for both blending and warping.

Face swap For the generation of face swapped images, *fewshot-face* [13] and *simple_faceswap* [48] are used. *fewshot-face* is a GAN-based approach capable of swapping a face using only a few target images; in the database used in this chapter a maximum of two target images was used to generate each face swapped image. *simple_faceswap* is a simple landmark-based approach which uses the landmarks detected by *dlib* to perform face swapping.

Example images generated using the above tools are shown in Fig. 2.6. For the generation of the swapped and morphed face images it was ensured that both individuals used to create the manipulated image were of the same gender. Additionally, to avoid artefacts, it was ensured that for the generation of the morphed images only one of the facial images contained glasses.

An overview of the total number of biometric comparisons in the generated database is given in Table 2.1. Note that for morphing and swapping where the manipulated image has been created from the facial images of two subjects, we only make

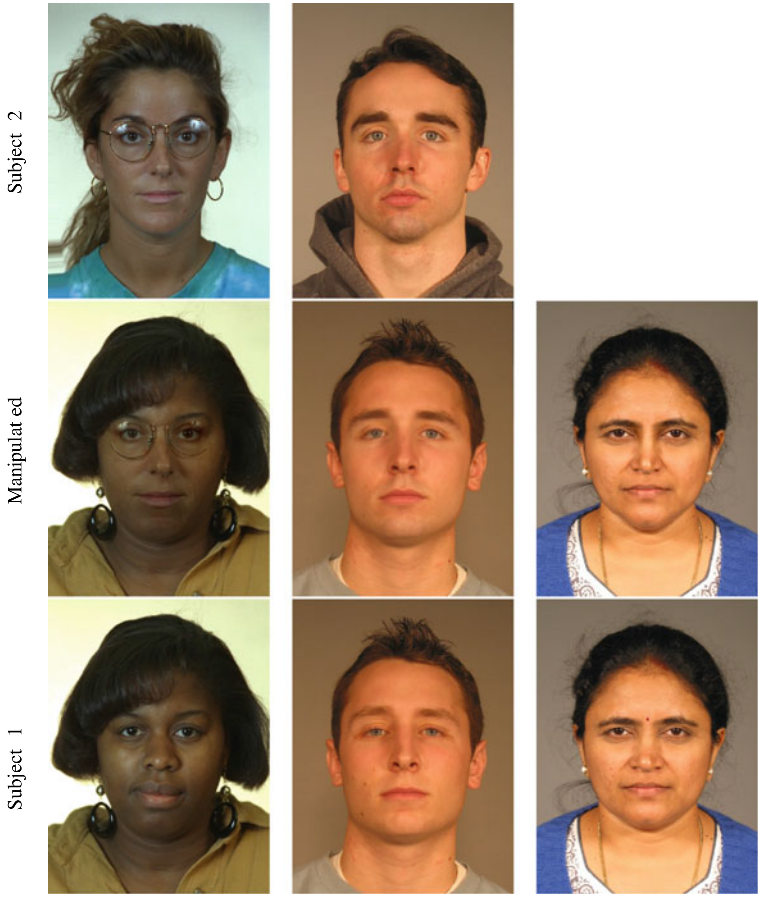

(a) Face swap (b) Morphing (c) Retouching

Fig. 2.6 Example images from the generated database

comparisons to the probe image(s) of the subject from which the area outside the face region is from. For instance, to create a mated comparison for the swapped and morphed face images in Fig. 2.6, a probe image from subject 1 is used. For the mated comparisons for retouching, morphing, and swapping, the used probe images have not been manipulated.

To evaluate the impact of the manipulations, standardized and other well-known metrics and visualizations are used. For visualizing the distributions of comparisons scores, probability density functions (PDFs) are used, and the scores produced by the algorithms are converted to similarity scores and normalized to the range [0, 1]. The degree of separation between two distributions is quantified using the decidability measure d' [6] which is calculated as follows:

Table 2.1 Number of biometric comparisons for the generated database

Scenario	Number of comparisons
Bona fide mated	2251
Retouching mated	4502
Morphing mated	4502
Face swap mated	4502
Bona fide non-mated	497,838

$$d' = \frac{|\mu_{\text{mated}} - \mu_{\text{non-mated}}|}{\sqrt{\frac{1}{2}(\sigma^2_{\text{mated}} + \sigma^2_{\text{non-mated}})}}$$

Biometric recognition performance is visualized using DET-curves which plot the *FNMR* versus the *FMR* at different decision thresholds. Furthermore, the *FNMR* at a fixed operational threshold corresponding to 0.1% *FMR* is highlighted; this security level is relevant for numerous real deployments of biometric recognition systems [38]. Finally, the equal error rate (*EER*), i.e. the point at which the *FMR* and *FNMR* are equal, is reported.

2.4.2 Performance Evaluation

This section investigates the effect of the different types of digital manipulations in the generated database (Sect. 2.4.1) on two state-of-the-art face recognition systems.

The PDFs and estimated decision thresholds at a fixed *FMR* of 0.1% for the manipulated and bona fide images in the generated database are shown in Fig. 2.7. It can be observed that the comparison scores of the manipulated images, for both *ArcFace* and *COTS*, are situated in-between the bona fide mated and non-mated score distributions. Furthermore, it can be observed that the score distribution for the retouched images is closest to the bona fide mated distribution, whereas face swapping is closest to the bona fide non-mated distribution. These observations are expected since retouching only moderately alters a face whereas for face swapping the original face identity has been replaced with the identity of another individual. From the plot, it can be observed that morphing, in general, decreases the comparison score more than retouching, but less than face swapping. Interestingly, from Fig. 2.7, it can be observed that the comparison scores for *COTS* on the swapped face images are significantly higher than the bona fide non-mated scores. When looking at the bona fide mated score distributions in Fig. 2.7 (most notable for *COTS*) two separate peaks can be observed, which is caused by using two different databases in the evaluation—the FRGCv2 database contains more unconstrained probe images than the FERET data.

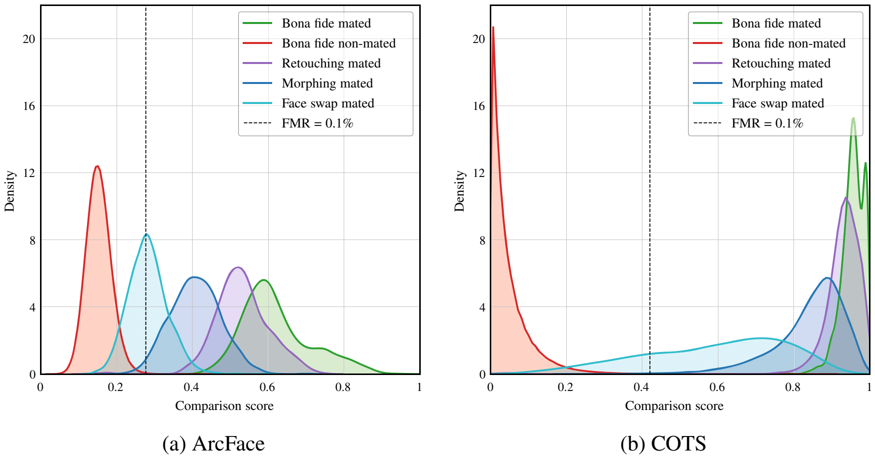

(a) ArcFace

(b) COTS

Fig. 2.7 Score distributions for manipulated and bona fide comparisons

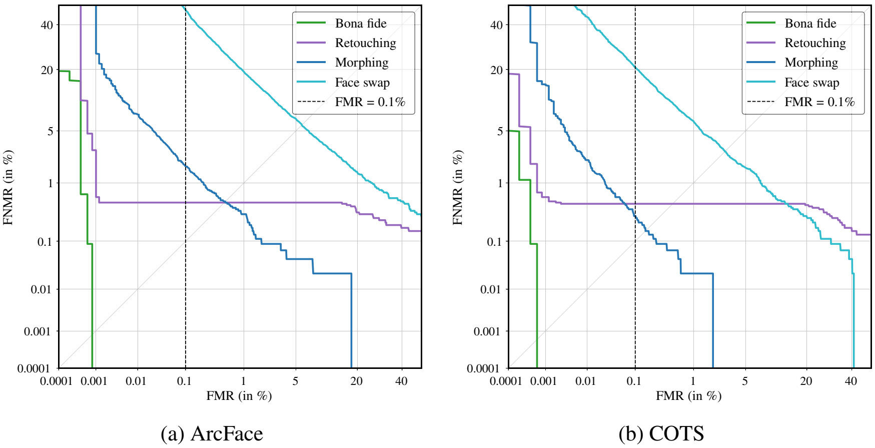

(a) ArcFace

(b) COTS

Fig. 2.8 DET-curves for performance scores of the manipulated and bona fide images

From the DET-curves in Fig. 2.8, it can be observed that face swapping has a big impact on the classification errors of both tested face recognition systems which is expected as face swapping changes the original face identity and as such makes the resulting identity less similar to the original identity. In contrast, retouching and morphing only have a moderate impact on the classification errors. For morphed images, this is a potential issue since the identity of multiple individuals contributes to a morphed image. As shown in other works, e.g. [10], morphed images can pose a security threat if accepted into a face recognition system as it is likely that the different individuals contributing to a morphed image can use the morph for authentication.

Table 2.2 Biometric performance results for ArcFace and COTS. *FNMR* is calculated at *FMR* = 0.1%. Values for *FNMR* and *EER* in the table are in %

Type	ArcFace			COTS		
	EER	FNMR	d'	EER	FNMR	d'
Bona fide	0.0004	0.0000	6.8386	0.0003	0.0000	20.4373
Retouching	0.4886	0.4887	6.7672	0.4664	0.4665	13.7155
Morphing	0.4970	1.7548	4.8792	0.1555	0.2888	10.2725
Face swap	5.8418	47.4678	2.9846	2.7765	20.7685	3.8780

Therefore, the system should ideally reject all mated comparisons where one of the images contains either a swapped or a morphed face.

In Table 2.2, the biometric performance scores are shown for *ArcFace* and *COTS*. The table shows good separability (high d') between the bona fide non-mated scores and the mated scores obtained for both the bona fide and retouched images. Least separation (lowest d') is achieved between the bona fide non-mated score distribution and the mated score distribution for the face swapped images. The performance metrics reported in the table indicate that both systems are robust to bona fide images and that face swapping has the biggest impact on the comparison scores of the systems. More specifically, it can be observed that at an operational threshold where *FMR* = 0.1%, approximately 47.5 and 20.8% of the mated comparisons for the face swapped images are rejected for *ArcFace* and *COTS*, respectively. The results show that, at best, less than half of the face swapped images are rejected, which suggest a need for algorithms capable of detecting swapped face images. For the retouched and morphed images, only moderate performance degradation can be observed. For morphed images, this means that state-of-the-art face recognition systems cannot reliably detect and reject morphing attacks. Therefore, several authors have proposed dedicated algorithms for detecting morphed images, although this remains a challenging problem [28].

2.5 Summary and Outlook

This chapter addressed the impact of digital face image manipulations on face recognition technologies. Considering the wide prevalence of face manipulation software, in particular for mobile devices and border control, face recognition systems have to cope with manipulated images. It was shown that face recognition systems can be robust to certain types of manipulations, i.e. biometric performance is maintained, while others may seriously reduce recognition capabilities. Therefore, the usability (related to the false non-match rate) and the security (related to the false match rate) of face recognition systems are impaired by digitally manipulated face images. Besides face manipulation techniques considered in this chapter, numerous face image manipulation techniques have been proposed for different fields of application. The forthcoming chapters of this book will describe many of those in detail.

Additionally, this chapter emphasized the need for reliable face manipulation detection methods to be integrated into face recognition systems. To this end, an overview of different concepts for the integration of face manipulation detection into the processing chain of a face recognition system was provided. Many of the subsequent chapters propose methods for reliable detection of different face manipulations which represent a current research challenge. Beyond that, some of the forthcoming chapters will provide more details on how to combine face recognition and face image manipulation detection effectively.

Acknowledgements This research work has been partially funded by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE and the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 860813—TReSPAsS-ETN.

References

1. Ahonen T, Hadid A, Pietikainen M (2006) Face description with local binary patterns: application to face recognition. IEEE Trans Pattern Anal Mach Intell 28(12):2037–2041
2. Akhtar Z, Dasgupta D, Banerjee B (2019) Face authenticity: an overview of face manipulation generation, detection and recognition. In: Proceedings of international conference on communication and information processing (ICCIP), pp 1–8
3. Biometric System Lab—University of Bologna. http://biolab.csr.unibo.it. Accessed 12 Mar 2021
4. Bowyer KW, Burge MJ (2016) Handbook of iris recognition. Springer International Publishing
5. Chingovska I, Anjos A, Marcel S (2013) Anti-spoofing in action: joint operation with a verification system. In: Conference on computer vision and pattern recognition workshops, pp 98–104
6. Daugman J (2000) Biometric decision landscapes. Technical Report UCAM-CL-TR-482, University of Cambridge—Computer Laboratory, January 2000
7. Debiasi L, Scherhag U, Rathgeb C, Uhl A, Busch C (2018) PRNU-based detection of morphed face images. In: 6th International workshop on biometrics and forensics, pp 1–7
8. Deng J, Guo J, Xue N, Zafeiriou S (2019) ArcFace: additive angular margin loss for deep face recognition. In: IEEE/CVF conference on computer vision and pattern recognition (CVPR), pp 4685–4694
9. FaceFusion. www.wearemoment.com/FaceFusion. Accessed 14 June 2021
10. Ferrara M, Franco A, Maltoni D (2014) The magic passport. In: IEEE International joint conference on biometrics (IJCB), pp 1–7
11. Ferrara M, Franco A, Maltoni D (2018) Face demorphing. IEEE Trans Inform Forensics Secur 13(4):1008–1017
12. Ferrara M, Franco A, Maltoni D (2019) Decoupling texture blending and shape warping in face morphing. In: International conference of the biometrics special interest group (BIOSIG), September 2019. IEEE
13. Few-shot face translation. https://github.com/shaoanlu/fewshot-face-translation-GAN. Accessed 14 June 2021
14. FotoRus (2018). https://www.apkshub.com/app/com.wantu.activity. Accessed 7 Mar 2021

15. Guo G, Zhang N (2019) A survey on deep learning based face recognition. Comput Vis Image Understanding, vol 189, pp 102805. http://dx.doi.org/10.1016/j.cviu.2019.102805
16. InstaBeauty (2017). https://www.apkshub.com/app/com.fotoable.fotobeauty. Accessed 7 Mar 2021
17. ISO/IEC JTC1 SC37 Biometrics. ISO/IEC 19795-1:2006. Information Technology—Biometric Performance Testing and Reporting—Part 1: Principles and Framework. International Organization for Standardization and International Electrotechnical Committee, March 2006
18. ISO/IEC JTC1 SC37 Biometrics. ISO/IEC 2382-37:2012 Information Technology—Vocabulary—Part 37: Biometrics. International Organization for Standardization, 2012
19. ISO/IEC JTC1 SC37 Biometrics. ISO/IEC 30107-1. Information Technology—Biometric presentation attack detection—Part 1: Framework. International Organization for Standardization, 2016
20. Jain AK, Flynn P, Ross AA (2007) Handbook of biometrics, July 2007. Springer
21. Jain A, Singh R, Vatsa M (2018) On detecting GANs and retouching based synthetic alterations. In: IEEE 9th International conference on biometrics theory, applications and systems (BTAS), pp 1–7
22. King D (2009) Dlib-ml: a machine learning toolkit. J Mach Learn Res 10(60):1755–1758
23. Learned-Miller E, Huang GB, Roy Chowdhury A, Li H, Hua G (2016) Labeled faces in the wild: a survey. In: Advances in face detection and facial image analysis. Springer, pp 189–248
24. Li SZ, Jain AK (eds) (2011) Handbook of face recognition. Springer, London
25. Liu L, Chen J, Fieguth P, Zhao G et al (2019) From BoW to CNN: two decades of texture representation for texture classification. Int J Comput Vis 127(1):74–109
26. Maltoni D, Maio D, Jain AK, Prabhakar S (2009) Handbook of fingerprint recognition, 1st edn. Springer
27. Marcel S, Nixon MS, Fierrez J, Evans N (2019) Handbook of biometric anti-spoofing: presentation attack detection, 2nd edn. Springer
28. Ngan M, Grother P, Hanaoka K, Kuo J (2021) Face recognition vendor test (FRVT) Part 4: MORPH–Performance of automated face morph detection. Technical report, National Institute of Standards and Technology, April
29. Parkhi OM, Vedaldi A, Zisserman A (2015) Deep face recognition. In: British machine vision conference (BMVC), pp 41.1–41.12
30. Phillips PJ, Wechsler H, Huang J, Rauss PJ (1998) The FERET database and evaluation procedure for face-recognition algorithms. Image Vis Comput 16(5):295–306
31. Phillips PJ, Flynn PJ, Scruggs T, Bowyer KW et al (2005) Overview of the face recognition grand challenge. In: IEEE computer society conference on computer vision and pattern recognition (CVPR), vol 1, pp 947–954. IEEE
32. Raghavendra R, Busch C (2017) Presentation attack detection methods for face recognition systems: a comprehensive survey. ACM Comput Surv 50(1):1–37
33. Raghavendra R, Raja KB, Busch C (2016) Detecting morphed face images. In: IEEE 8th international conference on biometrics theory, applications and systems (BTAS), pp 1–7
34. Ranjan R, Sankaranarayanan S, Bansal A, Bodla N et al (2018) Deep learning for understanding faces: machines may be just as good, or better, than humans. IEEE Signal Process Mag 35(1):66–83
35. Rathgeb C, Dantcheva A, Busch C (2019) Impact and detection of facial beautification in face recognition: an overview. IEEE Access 7:152667–152678
36. Rathgeb C, Botaljov A, Stockhardt F, Isadskiy S et al (2020) PRNU-based detection of facial retouching. IET Biomet 9(4):154–164
37. Rathgeb C, Satnoianu C-I, Haryanto NE, Bernardo K, Busch C (2020) Differential detection of facial retouching: a multi-biometric approach. IEEE Access 8:106373–106385
38. Research and Development Unit (2015) Best practice technical guidelines for automated border control (ABC) systems. Technical report, FRONTEX
39. Ross A, Nandakumar K, Jain A (2006) Handbook of multibiometrics. Springer

40. Rössler A, Cozzolino D, Verdoliva L, Riess C et al (2019) Faceforensics++: learning to detect manipulated facial images. In: International conference of computer vision (ICCV'19) pp 1–11
41. Sarkar E, Korshunov P, Colbois L, Marcel S (2020) Vulnerability analysis of face morphing attacks from landmarks and generative adversarial networks. arXiv e-prints, December 2020, pp 1–5
42. Scherhag U, Budhrani D, Gomez-Barrero M, Busch C (2018) Detecting morphed face images using facial landmarks. In: International conference on image and signal processing (ICISP), pp 444–452
43. Scherhag U, Rathgeb C, Merkle J, Breithaupt R, Busch C (2019) Face recognition systems under morphing attacks: a survey. IEEE Access 7:23012–23026
44. Scherhag U, Kunze J, Rathgeb C, Busch C (2020a) Face morph detection for unknown morphing algorithms and image sources: a multi-scale block local binary pattern fusion approach. IET Biomet 9(6):278–289
45. Scherhag U, Rathgeb C, Merkle J, Busch C (2020b) Deep face representations for differential morphing attack detection. IEEE Trans Inform Forensics Secur 15:3625–3639
46. Schlett T, Rathgeb C, Henniger O, Galbally J et al (2021) Face image quality assessment: a literature survey. arXiv e-prints, pp 1–29
47. Shen L, Bai L, Fairhurst M (2007) Gabor wavelets and general discriminant analysis for face identification and verification. Image Vis Comput 25(5):553–563
48. simple_faceswap. https://github.com/Jacen789/simple_faceswap. Accessed 14 June 2021
49. Taigman Y, Yang M, Ranzato M, Wolf L (2014) DeepFace: closing the gap to human-level performance in face verification. In: Conference on computer vision and pattern recognition (CVPR), pp 1701–1708
50. Tolosana R, Vera-Rodriguez R, Fierrez J, Morales A, Ortega-Garcia J (2020) Deepfakes and beyond: a survey of face manipulation and fake detection. Inform Fusion 64:131–148
51. Uhl A, Busch C, Marcel S, Veldhuis R (2020) Handbook of vascular biometrics. Springer International Publishing
52. Venkatesh S, Raghavendra R, Raja K, Busch C (2021) Face morphing attack generation & detection: a comprehensive survey. IEEE Trans Technol Soc (TTS) 2(3):128–145
53. Verdoliva L (2020) Media forensics and deepfakes: an overview. IEEE J Sel Top Signal Process 910–932
54. Wang M, Deng W (2021) Deep face recognition: a survey. Neurocomputing 429:215–244
55. Wang S, Wang O, Zhang R, Owens A, Efros A (2019) Detecting photoshopped faces by scripting photoshop. In: IEEE/CVF conference on computer vision and pattern recognition (CVPR), pp 10071–10080
56. Zhang K, Zhang Z, Li Z, Qiao Y (2016) Joint face detection and alignment using multitask cascaded convolutional networks. IEEE Signal Process Lett 23(10):1499–1503