applied sciences

MDPI

Article

Face Morphing, a Modern Threat to Border Security: Recent Advances and Open Challenges

Erion-Vasilis Pikoulis [1], Zafeiria-Marina Ioannou [1], Mersini Paschou [1] and Evangelos Sakkopoulos [2,*]

[1] Computer Engineering and Informatics Department, University of Patras, 25604 Rio, Greece; pikoulis@ceid.upatras.gr (E.-V.P.); ioannouz@ceid.upatras.gr (Z.-M.I.); paschou@ceid.upatras.gr (M.P.)
[2] Department of Informatics, University of Piraeus, 18534 Piraeus, Greece
* Correspondence: sakkopul@unipi.gr; Tel.: +30-2104142312

Abstract: Face morphing poses a serious threat to Automatic Border Control (ABC) and Face Recognition Systems (FRS) in general. The aim of this paper is to present a qualitative assessment of the morphing attack issue, and the challenges it entails, highlighting both the technological and human aspects of the problem. Here, after the face morphing attack scenario is presented, the paper provides an overview of the relevant bibliography and recent advances towards two central directions. First, the morphing of face images is outlined with a particular focus on the three main steps that are involved in the process, namely, landmark detection, face alignment and blending. Second, the detection of morphing attacks is presented under the prism of the so-called on-line and off-line detection scenarios and whether the proposed techniques employ handcrafted features, using classical methods, or automatically generated features, using deep-learning-based methods. The paper, then, presents the evaluation metrics that are employed in the corresponding bibliography and concludes with a discussion on open challenges that need to be address for further advancing automatic detection of morphing attacks. Despite the progress being made, the general consensus of the research community is that significant effort and resources are needed in the near future for the mitigation of the issue, especially, towards the creation of datasets capturing the full extent of the problem at hand and the availability of reference evaluation procedures for comparing novel automatic attack detection algorithms.

Keywords: face morphing; morphing attack detection; deep learning; differential attack detection; single-image attack detection

check for updates

Citation: Pikoulis, E.-V.; Ioannou, Z.-M.; Paschou, M.; Sakkopoulos, E. Face Morphing, a Modern Threat to Border Security: Recent Advances and Open Challenges. *Appl. Sci.* **2021**, *11*, 3207. https://doi.org/10.3390/app11073207

Academic Editor: Theodoros Giannakopoulos

Received: 25 February 2021
Accepted: 29 March 2021
Published: 2 April 2021

1. Introduction

Automatic identity verification based on stored biometric features is progressively replacing traditional verification procedures based on paper documents. Since 2002, the face has been selected by the International Civil Aviation Organization (ICAO) as the primary biometric trait for machine-assisted identity confirmation in electronic Machine Readable Travel Documents (eMRTD) [1]. In parallel, recent advances in the field of automatic face recognition and especially the performance breakthrough achieved by Deep Convolutional Neural Networks (DCNNs) [2], have enabled the deployment of face recognition technologies in a variety of applications, including critical security applications such as ABC.

Face photos that are eligible for an eMRTD can be provided by the citizen in printed form to the agency issuing the document, which exposes the system to image-alteration attacks. The vulnerability of automatic FRS has been well-documented, with studies showing that, while state-of-the-art systems are able to cope with limited alterations [3], the generalizability of (deep) FRS increases their vulnerability against more relevant attacks [4–6]. One of the most severe such image-alteration attacks is based on morphed face images (also referred to as morphing attacks), a scenario that was first brought into the attention of the research community by the paper of Ferrara et al. [7].

The severity of the security threat posed by morphing attacks is nowadays prompting countries towards legislative acts with the goal of preventing them [8]. It has also turned Morphing Attack Detection (MAD) into a very active research domain with a plethora of published articles proposing numerous solutions to mitigate the problem [6].

Based on the information available at the time of detection, MAD techniques can be categorized into on-line (or differential), and off-line (single-image, or no-reference). In the on-line scenario (applicable e.g., during authentication at an ABC gate), the detector is presented with a stored (i.e., passport) test image, and also with a live capture of the individual, and has to decide whether the test image has been altered or not, by using also information from the reference live image. On the other hand, in the off-line scenario (applicable e.g., during an eMRTD application process), the detector has to assess the authenticity of the input test image, with no other reference available.

As it is the case with the majority of classification/detection problems, morphing attack detection is carried out in a feature space of the input. Based on the tools used for feature extraction, namely, whether the features are automatically learned using DCNNs, or pre-selected (e.g., texture descriptors), a further possible categorization of MAD techniques (also inherent to many classification tasks since the advent and widespread use of DCNNs) is that of "deep" vs. "non-deep" detection. In general, promising results have been reported on both sides, with the main difference being the requirement for large datasets under the "deep" scenario, which is a prerequisite for feature learning using modern DCNNs.

However, despite the reported results, the general consensus within the research community is that morphing attack detection remains still an open and extremely challenging issue [9]. This is simply a reflection of the fact that the detection of morphing attacks is in its heart an ill-posed problem. Indeed, apart from the general concept of what constitutes a morphing attack, along with the efforts of the research community to recreate such a scenario, a clear definition of the morphed image, as well as a documentation of real-life morphing attacks are still missing (and are difficult to exist, in general). As such, the true extent of the problem and accordingly, the severity of the measures necessary to address it, remain largely unknown, as it will be discussed later in the paper.

The rest of this paper is organized as follows: In Section 2 we describe the morphing attack scenario. Subsequently, in Sections 4 and 5 we present a selected bibliography on the creation and detection, respectively, of morphed images. In Section 6 we present the metrics used for evaluating the capability of the attack detection. Section 7 holds a discussion on the recent advances and open challenges posed by the morphing attack problem, and finally, Section 8 holds our conclusions.

2. The Morphing Attack Scenario

The main idea behind morphing attack is based on blending face images from two individuals, so that the resulting "morphed" image can be used for the verification of both involved participants [7]. In other words, the biometric information extracted by the morphed image resembles the true information of both participants enough to yield a successful match of both, by the FRS.

In more detail, the morphing attach scenario can be generally described as follows [6,7]: A wanted "criminal" with the goal of travelling without alerting the authorities, morphs his/her facial image with that of a lookalike accomplice. The accomplice, having no criminal record, applies for an eMRTD by presenting the morphed face photo as his/her own. If the image alteration is not detected at the time of eMRTD issuance, he/she will receive a perfectly valid eMRTD. At the time of verification at an ABC, a face image (acquired live) of the person presenting the travel document (in this case, the criminal) is matched against the face image stored in the eMRTD. If the morphing is successful, this means that the criminal can use the eMRTD issued to the accomplice, to overcome the security controls and travel freely.

3. The Posed Threat

Considering the fact that high-quality morphed images, such as the ones shown in Figure 1 can be realistic enough to fool even the trained eye [10,11], the security threat posed by the morphing attack scenario becomes readily apparent.

From a technological standpoint, as digital image manipulation techniques become more and more advanced, the resulting morphs are expected to be increasingly more difficult to detect. Fortunately, technological advancements work also in favor of mitigating the problem, with increasingly more sophisticated approaches being developed for the automatic detection of morphing attacks. Many of these techniques are going to be presented in the subsequent sections. However, despite these efforts, the proposed techniques are yet to be deployed or even tested under "real" attack scenarios, nor is there a standardized procedure in place for evaluating their performance, as it is going to be explained later in more detail.

What is more, the decision to accept or reject an ID image in face matching scenarios is usually left to a human operator, since, even in cases where automated techniques are initially employed, human users are often required to make the final selection from an automatically selected list of candidates.

Unfortunately, despite the belief that humans are "experts" at face identification, numerous studies have shown that we are basically familiar-face experts but are prone to errors when dealing with unfamiliar faces [12–14]. Surprisingly (and worryingly), these studies also point out that trained passport officers perform at similar levels to untrained university students when dealing unfamiliar faces, which is especially problematic when dealing with various types of fraudulent identification [15].

Concerning the issue of morphing attacks, this is perhaps best emphasized by the recent work in [11] (following similar studies presented in [16,17]), aiming at evaluating the performance of humans (with and without training) for the detection of morphed images, in a series of carefully designed experiments using high-quality morphs. In the first two experiments, the participants were asked to identify morphs, before and after receiving training regarding the detection of morphed images. In the first experiment, the task was to correctly classify entries in ten-image arrays containing both morphs and genuine photographs, while the participants of the second experiment were presented with a single image and had to decide on its authenticity. In their third experiment, the authors tested the participants' performance with respect to live face-matching, i.e., their ability to identify whether the person presenting a photograph is the one depicted in it. The general conclusion in all three experiments was that "naive participants were poor at detecting high-quality morph images and that providing training or tips aimed at improving performance resulted in little or no benefit". In fact, with few exceptions, the morphing detection performance was barely above chance levels. Furthermore, the authors concluded that "morphs were accepted as ID photos often enough that they may be feasible as tools for committing fraud."

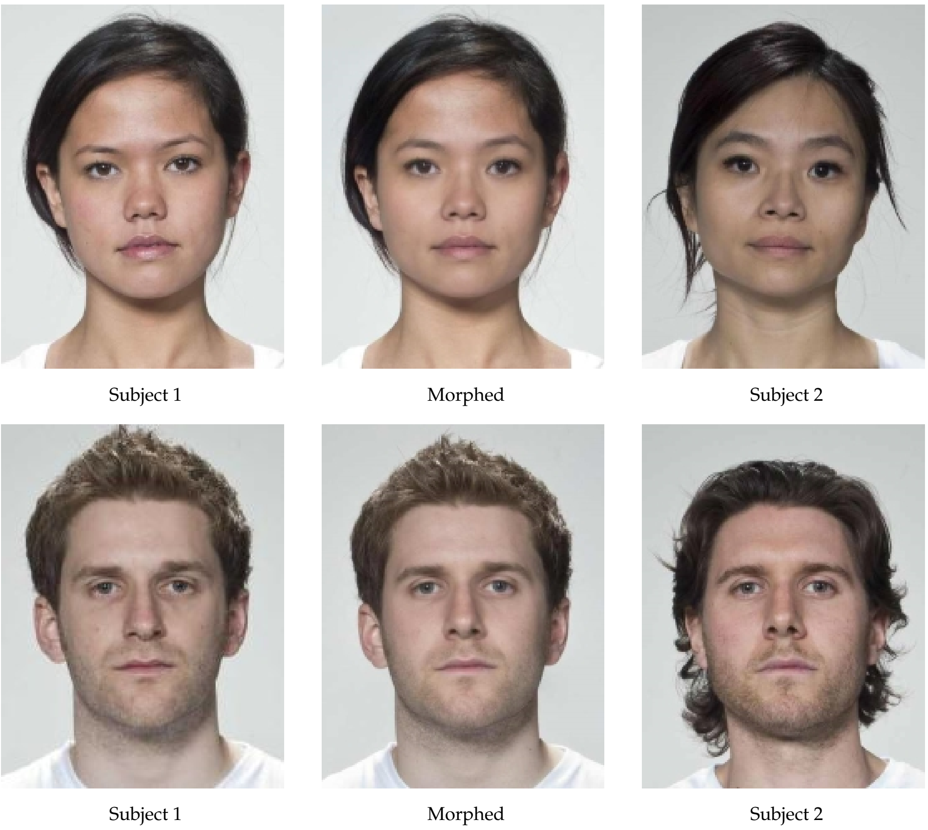

Figure 1. Morphing examples using women (**top row**) and men (**bottom row**) subjects, from the AMLS Dataset [18]. Original data are from the datasets in [19,20], while the morphing approach reported in [21] was used for the creation of morphs. It is noted that in both cases, the morphed image would yield a successful verification to both contributing subjects, by an automatic face recognition system.

4. Morphing Face Images

The original face images used for the production of morphed photographs should meet certain criteria so that the resulting morph complies with the ICAO requirements for the issuance of travel documents. These requirements entail for example that the involved originals should be frontal images exhibiting a neutral facial expression and all faces are represented equally with respect to resolution, exposure, etc.

The process of morphing face images can be generally divided into three main steps, with the likely addition of pre− and/or post− processing steps for facilitating the morphing process and removing possible morphing artefacts in order to achieve more realistic results, respectively [6,7,22]. The main morphing steps are: (a) face landmarking, (b) image alignment/warping, and (c) image blending. These steps are summarized in Figure 2.

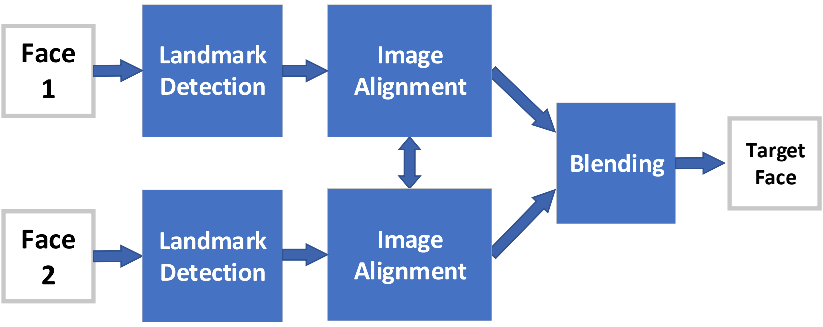

Figure 2. The main steps for creating morphed face images.

4.1. Landmark Detection

Face landmarking, refers to the detection and localization of characteristic points on the face, and constitutes an important intermediary step for many subsequent face processing operations, such as morphing. Although in the general case, this computer vision problem has proven extremely challenging due to inherent face variability as well as factors such as pose, expression, illumination and occlusions, the issue is simplified to some extent in the case of morphing applications since the involved (ICAO compliant) face photographs are taken in controlled environments, with good lighting conditions and devoid of specific poses or expressions.

With a few exceptions, face landmarks do not generally correspond to salient points in the image, meaning that low-level image processing tools are not sufficient for their detection, and higher-order face shape information is required. More specifically, the corners of the eyes, of the mouth, the nose tip, and sometimes the eyebrows are referred to as the primary or fiducial landmarks and can be detected directly using low-level image features such as gradient information, cornerness or local information extracted, e.g., with scale invariant feature transform (SIFT) [23], histogram of gradients (HOG) [24], and generic information on the face morphology. The fiducial landmarks, are primarily important in facial identity and face tracking. The detection of other landmarks, such as nostrils, chin, nasion, cheek contours, non-extremity points on lips or eyebrow midpoints, eyelids etc. is often guided by the primary landmarks.

According to the taxonomy presented in [25], landmarking methods can be split into two broad categories, namely model or shape based, which consider the ensemble of facial landmarks as a whole shape and learn "face shapes" from labeled training images, and texture based whose goal is to detect each facial landmark or local groups of landmarks independently, without the guidance of a model. Techniques based on active shape models (ASM), which seek the best-fit of predefined landmarks in an image [26], active appearance models (AAM), which constitute generalization of ASM [27,28], and elastic bunch graph-matching (EBGM), which construcut phase graphs whose nodes represent 40-dimensional Gabor jets at candidate locations [29], are well-known model-based methods. On the other hand, transform-based schemes where the image is window-scanned to produce feature vectors which are subsequently compared with learned patterns, constitute primary examples of texture based methods. Some of the transforms used by this group of methods include principal component analysis (PCA) [30], Multiresolution Wavelet Decomposition [31], Gabor wavelet transform (GWT) [32], discrete cosine transform (DCT) [33], and independent component analysis (ICA) [34].

In face morphing, detected landmarks in both images act as reference points for the alignment step. The simplest, most accurate, but also the most time-consuming way is the manual location of features such as, eye corners, eyebrows, nose tip, chin, forehead etc., a process followed in [7]. Typically, the detection of landmarks is done in an automated way using Machine Learning toolkits such as dLib [35] that offers implementations of well-known landmark detectors including e.g., the regression-trees-based technique in [36]. dLib is used for landmark detection by several morphing-related papers, including [22,37]. Other automatic landmark detection approaches range from model fitting, such as the previously mentioned ASM, AAM, and EGBM and elastic bunch graph models, to optimization techniques such as the one proposed in [38].

4.2. Face Alignment

The goal of the image alignment step is to obtain a warping of the two original images so that the locations of the corresponding landmarks (detected in the previous step) are matched. Warping is basically a geometric transformation applied to the pixel locations of both images. Early morphing techniques relied on scattered data interpolation [39], whereby the new pixel position pixels is interpolated based on the nearby control points (landmarks). More advanced approaches include the Free-Form Deformation (FFD) technique [40,41] in which the image is embedded in a grid of control points (i.e., the pixel locations are expressed as a function of the grid-point locations). This way, applying a desired deformation to the grid, results also to an analogous deformation of the image. In feature-based image metamorphosis [42], the deformation is controlled by the grid lines, which is particularly advantageous for manual morphing as the user can position lines instead of points. In [43] the optimal affine transformation based on either landmarks or lines via the minimization of the moving least squares, while in [30] a morphing process by simulating the image as a mass spring system is proposed.

The most widely used warping technique by state-of-the-art morphing algorithms is based on the Delaunay triangulation of the detected landmarks [5,9,22,44,45]; that is the partitioning of the images into a set of non-overlapping triangles whose vertices lie on the detected landmarks. The goal here is to obtain a matching between corresponding landmarks by applying geometrical transformations (rotations, translations and distortions) of the obtained Delaunay triangles.

The steps involved in the warping process can be analyzed into the following steps [45,46]:

1. Generate triangular facial mesh: Obtain Delaunay triangulation of the detected landmark locations.
2. Create interpolated facial mesh by blending original landmark locations, i.e.,

$$\mathcal{L}_w^i = \alpha_w \, \mathcal{L}_1^i + (1 - \alpha_w)\mathcal{L}_2^i, \quad i = 1, \ldots, N$$

   where $\mathcal{L}_w^i$, $\mathcal{L}_1^i$, $\mathcal{L}_2^i$, denote the blended and the (two) original landmark locations, respectively, while N is the number of used landmarks. α_w is a blending coefficient taking values in $[0, 1]$ that controls the contribution of each of the two original faces in determining the location of the landmarks on the final morphed face. For example, $\alpha_w = 0$ (or, 1) signifies that the landmark locations on the morphed image (and loosely speaking its general form and shape) will coincide with on of the involved originals.
3. Warp both original faces to the interpolated mesh: Apply affine transform to map the pixel locations of each the original triangles, to the locations of the corresponding interpolated triangle.

In addition, to avoid ghost and blur artifacts and to obtain more natural warping results, the projection of the face images into a 3D space prior to warping has been proposed [47,48].

4.3. Blending

Color interpolation can be understood as alpha-blending of intensity values of both images for each color layer separately. The parameter alpha defines the proportion of pixel intensity values obtained from source and target images.

After the alignment step, the two contributing images are blended, usually over the entire image region. The most frequent way of blending for face morph creation is linear blending, i.e., all color 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 adjusted via a blending coefficient taking values in $[0, 1]$; e.g., a value of 0.5 signifies taking the average of the corresponding pixels in each position.

4.4. Post-Processing

Post-processing steps are often incorporated in the morphing pipeline with the goal of removing artefacts and generally masking telltale signs of the morphing process. For example, the automatically created morphs can suffer from shadow or ghost artifacts, from missing or misplaced landmarks during the warping process. One way of removing such artefacts is to blend the problematic area (e.g., eyes, nostrils, hair) with corresponding content from one of the original images, an approach followed e.g., by the FaceFusion app [49], as reported in [50]. Hair artifacts can be concealed by an interpolation of the hair region [51]. Post-processing entails steps towards enhancing or reducing the image quality. For example, image processing tools such as blurring, sharpening or histogram equalization can be used to enhance color and definition, and achieve a more natural looking final image. On the other hand, unavoidable steps usually employed by issuing authorities, such as printing and/or scanning the image for its use as a passport photo, or changing the (lossy) format in which the image is stored, results in information loss that could mask morphing artifacts.

4.5. Morphing Software

In this section, a brief discussion on the software tools for face morphing, will be provided (see, also [6,52] and references therein). In general, the available solutions can be categorized into three groups, namely, (a) commercial (proprietary), (b) open-source, and (c) online. Starting with the latter, although easy-access, platform independent, online tools do exist, their use by the research community might be problematic as the procedures that are adopted for morphing faces, are not well-documented. For the other two categories, solutions that can be deployed on all well-known operating systems (either for desktop/laptop computers or smart-phones), can be found. In [6], a detailed table with various tools is provided, in which the presented solutions are characterized in terms of (a) the methods utilized for achieving the end-goal of face morphing, (b) the level of automation provided and/or the manual effort required, (c) the level of expertise desired by the operator, and (d) the achieved quality. As a final note, access to commercial software might be hindered due to their access policies, while their "black-box" procedures might be incompatible with research activities aiming to propose new morphing algorithms (either for creating or detecting morphed images). On the other hand, open-source solutions provide access to all relevant information, however, the stability of the provided software might need time for appropriate testing and verification. FaceFusion [49] and FaceMorpher [53] constitute two well known examples of commercial and open-source morphing tools, respectively.

5. Morphing Attack Detection

In this section, we present a summary of the most recently proposed techniques for the automatic detection of morphing attacks. It must be stressed that it is a problem that has sparked the interest of many research teams with great efforts being directed towards its solution. This has turned the relevant bibliography into a very active landscape, resulting in numerous published articles covering a great variety of feature extraction and classification techniques.

In general, MAD is treated as a binary classification problem in the bibliography. Given an input image, the goal is to infer the class it belongs to, namely, decide whether it represents an original image, or it is the product of morphing. As it is the case with the vast majority of classification problems, the decision is carried out in a feature space, i.e., it is performed by utilizing the values of selected features of the input. Based on the information available at the time of detection, the proposed techniques can be categorized as on- or off-line. In the first case, the system tries to decide whether an image is the product of morphing by utilizing a live capture as reference, while in the second case, the decision is based in the image itself. A further possible categorization is based on the type of the classification framework used, namely, whether the features used for classification are automatically learned, or engineered.

5.1. On-Line vs. Off-Line Detection Scenario

In the on-line (or differential) detection scenario (depicted in Figure 3), along with the input test (i.e., stored) image, the detector is also presented with a live capture, e.g., during authentication at an ABC gate. The detector then tries to decide whether the test image has been altered or not, by using also information from the live image. On the other hand, in an off-line (single-image, or no-reference) detection scenario (depicted in Figure 4), the detector is only presented with a single test image (e.g., the face image deposited during an eMRTD application process), and has to decide the authenticity of the input with no other reference available. As it is understandable, the no-reference scenario poses a more challenging problem. It should also be noted that all information extracted by no-reference morph detectors could also be used in a differential detection scenario.

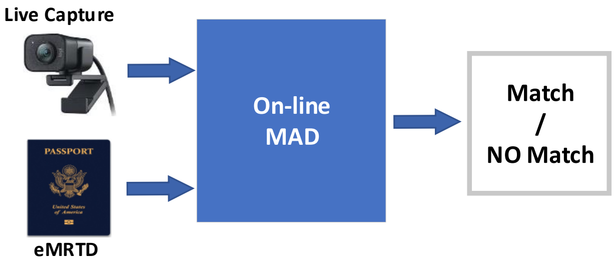

Figure 3. In the on-line scenario, MAD is based on a test image that is provided via the eMRTD and a live capture of the person's face.

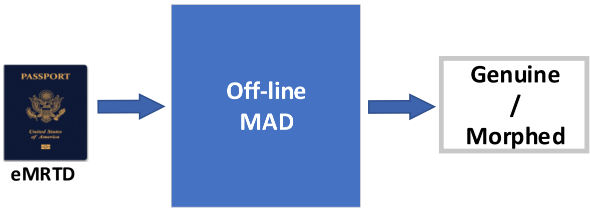

Figure 4. In the off-line scenario, MAD is based only on a test image that is provided via the eMRTD.

5.2. Deep vs. Non-Deep Classification

This categorization concerns mainly the approach used for feature extraction (i.e., whether features are automatically learned as in Figure 5 or engineered as in Figure 6) and is inherent to almost all classification tasks since the advent and widespread use of DCNNs in recent years. Nowadays, DNNs have established themselves as prominent tools for solving Machine Learning (ML) and Artificial Intelligence (AI) problems [54] (e.g., face recognition, scene analysis and driver state monitoring, autonomous driving etc.), achieving unprecedented results in challenging competitions such as the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [55], and even exceeding the accuracy of human experts in certain classification tasks [56]. As it is to be expected, DCNNs have been extensively used for morphing attack detection, with the majority of the techniques relying on the learnt features of well-known, pre-trained networks from the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [55], such as VGG [57], AlexNet [58], ResNet [59], and GoogleNet [60]. The main downside of DCNN-based approaches is the requirement of large morphing datasets, for DCNN training (i.e., feature learning).

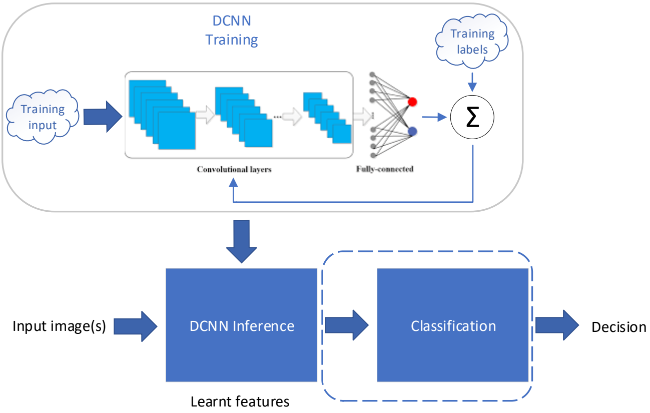

Figure 5. In deep classification, feature extraction relies on an automatic learning procedure that utilizes deep learning.

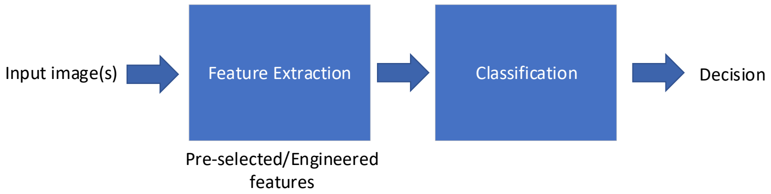

Figure 6. In non-deep classification, the features to be used, are derived via an engineering procedure with the intervention of a human designer.

On the other hand, the techniques following an "non-deep" approach, rely on engineered features for classification, with the majority of them employing general-purpose texture descriptors, widely used in Computer Vision tasks, such as the Local Binary Patterns (LBP) and the Scale-Invariant Feature Transform (SIFT).

A categorization of the MAD-related works included in the following sub-section, is presented in Table 1.

5.3. MAD-Related Literature

In the following, some indicative works treating the problem of automatic morphing attack detection will be outlined.

In [22], VGG19, AlexNet, and GoogleNet were employed for the detection of morphs, while focusing on digital morphed face images. Tests were conducted using both pre-trained (on the ILSVRC dataset) and trained-from-scratch (using morphed images) versions of the DCNNs. Unsurprisingly, the pre-trained versions, having learned rich feature sets from the extensive ILSVRC dataset, outperformed the trained-from-scratch versions in the conducted experiments.

In [44], the authors aim at overcoming a known issue related to DCNNs, which makes them vulnerable to attacks on their decision-making process and can worsen their generality. Simply put, neural networks tend to use only what they need for a task (e.g., classification), ignoring the rest of the information that is available in the training data. For example, to recognize a rooster, a network might only considers the rooster's red comb and wattle and ignores the rest of the animal. Towards mitigating this problem (in relation to morphing detection), the authors propose training schemes which are based on different alternations of the training data (morphed images), by limiting the amount and position of information available to the neural network for the decision making process. An example would be to use partial face morphs (i.e., a forged face image in which only some regions are taken from a morphed face image and the remaining part from an aligned input image), during training. A recently proposed technique called Layer-Wise Relevance Propagation (LPR) [61] was used in order to interpret the results (prediction scores) of the DCNNs, with the goal of analyzing the differences in the decision making process of the differently trained neural networks. LPR redistributes the prediction of the model backwards, using local redistribution rules, until it assigns a relevance score to each input variable (e.g., image pixel).

In [62], a face morphing attach method is proposed in which the live capture is employed in a process that aims at reverting the morphing transform, if exists, in the stored image. In the following, the so-called "demorphed" image that is produced, is decided whether is impaired or not. Ref. [63] proposes a face morphing detection method by exploiting the so-called sensor pattern noise in the Fourier domain. Using the Fourier spectrum of the two face images (live capture and stored image), facial quantification statistics are extracted that differentiate the two images due to the considered pattern noise. A linear support vector machine classifier is employed for the final decision. In [64], the face morphing detector aims to identify differences in the geometry of the two face images under consideration. This is done by analyzing the locations of facial landmarks that are biologically meaningful. The evaluation was performed for a variety of image sources and experimental settings.

The authors in [65] propose a detection technique, specifically designed for passport-scaled morphed images, comprising a pre-processing module that, if required, transforms a given face image into a passport-like image, followed by a feature extraction and a classification module (for which four statistical models are trained and evaluated). The work studies the impact of two different feature spaces, one in the frequency domain and another in the spacial domain. In [66], an "evaluation" issue is identified in the literature according to which the performance assessment of the proposed face morphing attack methods is commonly based on datasets of limited scope, namely, either a few types of morphing attacks are considered or a single dataset is employed. There, a new, richer dataset is created and an assessment of two detection techniques based on support vector machines and local binary patterns, is provided.

The authors in [67], followed the transfer-learning approach using the VGG19 and AlexNet DCNNs to detect both digital and print-scanned versions of morphed face images. The employed DCNNs were pre-trained (for the task of object classification on the ILSVRC dataset) and then further fine-tuned using morphed face images. At the inference phase, both AlexNet and VGG19 are fed with the same (test) image and the feature maps produced at their respective first fully connected layers are fused (i.e., concatenated) to form a single feature-vector which is provided to the classifier in order to decide whether the input was a genuine or a morphed face image. The goal was to utilize the full-scale feature-extraction potential of the concolutional layers of the DCNNs. The Probabilistic Collaborative Representation Classifier [68] was used for classification purposes.

In [37], the authors achieve high detection ratios (surpassing the performance of rival deep-learning-based techniques) by pairing high-dimensional Local Binary Patterns (LBP) descriptors with a linear Support Vector Machine (SVM) for classification. LBPs are also used in [69]. The authors in [70] propose training SVMs using Weighted Local Magnitude Pattern, a descriptor similar to LPB, while a detector based on Scale-Invariant Feature Transform (SIFT) is proposed in [71]. Recently, a different approach, based on Photo Response Non-Uniformity (PRNU) analysis, is proposed in [9]. PRNU, which refers to the sensor's noise pattern, can be considered as a unique fingerprint for each camera sensor. As such, it presents itself with a number of advantages (large information content, present in every sensor, robust to lossy compression, filtering etc, but sensitive to non-linear warping) that make it suitable for morphing attack detection, with greater robustness and generalization potential than descriptor-based detectors.

Treating the problem under the differential scenario, a generative adversarial network (GAN) for face de-morphing was proposed in [72] with the aim to restore the face of the accomplice that collaborated in the face morphing attack. To achieve this, the network analyses the live-capture and the morphed photo that was used in the utilized eMRTD. A similar approach for de-morphing has been proposed in [73], where the general idea is to "subtract" the information captured in-vivo (i.e., the live image) from the passport image and compare the residual of the de-morphing process, with the live capture. In case the passport image was the product of morphing between a "criminal" and an "accomplice", the residual image would resemble that of the accomplice (not present in the scene) and not live capture of the criminal. An autoencoder architecture based on VGG-face [74] was used for de-morphing. The authors in [50] propose a high-performance differential detector by using the feature vectors produced by ArcFace, a ResNet-based FRS. Detection is achieved by subtracting the feature vectors of the suspected image and the live capture. In [75] the authors use convolutional networks for image denoising and pyramid LPB for feature extraction while the same authors in [76], apart from using a DCNN-based denoising approach, utilize also AlexNet for feature extraction. Finally, Ref. [77] studied the problem of morphing attack detection when the morphing faces are produced using GANs instead of landmark-based approaches.

Table 1. Categorization of selected works on Morphing Attack Detection.

On-Line		Off-Line	
Deep	**Non-Deep**	**Deep**	**Non-Deep**
Peng et al. (2019) [72] Ortega et al (2020) [73] Scherhag et al. (2020) [50]	Ferrara et al. (2017) [62] Autherith et al. (2020) [64]	Seibold et al. (2017) [22] Seibold et al. (2020) [44] Raghavendra (2017) [67] Venkatesh et al (2019) [75] Venkatesh et al (2020) [76]	Neubert et al. (2019) [65] Spreeuwers et al. (2018) [66] Wandzik et al. (2018) [37] Ramachandra (2016) [69] Agarwal et al. (2017) [70] Kraetzer et al. (2017) [71] Scherhag et al. (2019) [9]

6. Evaluation Metrics

Recently, the community has achieved a common standard ISO (IEC 30107-3:2016) [78] to evaluate presentation attack detection (PAD), and especially morphing attack detection (MAD) systems. In this standard, the capability of the attack detection is measured with the following errors: attack presentation classification error rate (APCER) and bonafide presentation classification error rate (BPCER). This measure can be defined as follows:

- Attack presentation classification error rate (APCER) is defined as the proportion of presentation attacks that have been classified incorrectly (as bonafide) [78]:

$$\mathrm{APCER}_{\mathrm{PAI_s}} = 1 - \frac{1}{|\mathrm{PAI}|} \sum_{\omega=1}^{|\mathrm{PAI}|} \mathrm{RES}_{\omega},$$

  where $|\mathrm{PAI}|$ is the number of presentation attack instruments (PAI) and RES_{ω} takes the value 1 if the presentation ω is assessed as attack and 0 if it is evaluated as bonafide. A PAI is defined as a used object or biometric trait in a presentation attack.
- Bonafide presentation classification error rate (BPCER) is defined as the proportion of bonafide presentation incorrectly classified as presentation attacks [78]:

$$\mathrm{BPCER}_{\mathrm{PAI_s}} = 1 - \frac{1}{|\mathrm{BF}|} \sum_{\omega=1}^{|\mathrm{BF}|} \mathrm{RES}_{\omega},$$

  where $|\mathrm{BF}|$ is the cardinality of bonafide presentations and RES_{ω} takes the value 1 if the presentation ω is allocated as an attack and 0 if it is analyzed as bonafide.

An APCER vs. BPCER detection error trade-off (DET) curve and the equal error rate (EER) where both errors are identical, provides a comparison among MAD systems.

7. Discussion

Morphing attacks constitute a well-established and pertinent security issue. The bibliography regarding both morphing algorithms and morphing attack detection, is already in a mature stage, as it is clear from the works presented in this paper. However, despite the progress being made, the detection of morphing attacks remains still an open and challenging issue, both because of the difficulty posed by the problem itself - morphed images can become so realistic that even trained experts perform poorly in their detection [10,11]—but also, due to the fact that the attempts to solve it are hindered by the lack of standardized procedures for creating morphed images, and of benchmarks data, in general.

Mostly, the proposed detection techniques are trained and tested using self-created datasets, meaning that the presented results are hard to reproduce and difficult to interpret. This issue is reported widely in the bibliography, with [9,52] being the most recent examples. Due to this reason, explicit performance comparisons between rival MAD approaches are left out of this paper, and the interested reader is referred to relevant works included in the bibliography, such as [6,52].

Constructing a benchmark dataset of morphed images is a necessary step, albeit a time-consuming and a painstaking one. It should be also noted that automatically produced face morphs may not represent a realistic attack scenario. After all, a morphing attack represents a scenario where the attacker has ample time and resources to produce a "single" good image that will "fool" the FRS. The dataset needs to be as true-to-life as possible, and as extensive/representative as possible.

The previous remark is especially important for utilizing the full potential of deep networks for the problem of automatic morphing detection. The true power of deep architectures lies in the automatic learning of features, which is only made possible by the availability of extensive training datasets. To avoid over-fitting, apart from training data, new, specialized architectures and/or training protocols, tailored to the specific needs of the morphing detection problem, are required, as noted in the relevant bibliography [44]. Along the same lines, regarding the non-deep classification approach, new specialized descriptors/features must be devised and used instead of (or in conjunction with) general-purpose descriptors such as LBP or BSIF that are devised for Computer Vision tasks, which are very different in nature to morphing detection.

Finally, the need for independent evaluation of the algorithms, using dedicated platforms and testing datasets is also an issue raised in the bibliography with attempts being already made in this direction [52,79].

8. Conclusions

In this paper, the problem of morphing attacks and the tools used for its mitigation are discussed. The paper, after the description of the face morphing process, provides a thorough description of the morphing attack detection bibliography. The particular focus is on the so-called on-line (differential) and off-line (single-image) scenarios and whether the features used during the detection of a possible morphing attack are hand-crafted or automatically generated via deep learning methods. Despite recent advances showing considerable progress in the fields of face morphing and morphing attack detection, many open challenges still remain concerning, among others, the availability of representative, large-scale datasets and reference evaluation procedures of novel detection algorithms. Addressing these challenges requires considerable time, effort, and resources, but it is a necessary step towards solving the problem and ensuring border security.

Author Contributions: Methodology & writing—original draft preparation, E.-V.P., Z.-M.I. and M.P.; Supervision & writing—review and editing, E.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research has been co-financed by the European Regional Development Fund of the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH-CREATE-INNOVATE (project Happy-Cruise, project code: T1EDK-03715).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. ICAO. Biometric Deployment of Machine Readable Travel Documents. In Proceedings of the TAG MRTD/NTWG, Montreal, QC, Canada, 17–21 May 2004.
2. Schroff, F.; Kalenichenko, D.; Philbin, J. FaceNet: A unified embedding for face recognition and clustering. In Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 7–12 June 2015; pp. 815–823.
3. Ferrara, M.; Franco, A.; Maltoni, D.; Sun, Y. On the impact of alterations on face photo recognition accuracy. In Proceedings of the International Conference on Image Analysis and Processing (ICIAP), Naples, Italy, 9–13 September 2013; pp. 7743–7751.

4. Mohammadi, A.; Bhattacharjee, S.; Marcel, S. Deeply vulnerable: A study of the robustness of face recognition to presentation attacks. *IET Biom.* **2018**, *7*, 15–26. [CrossRef]
5. Wandzik, L.; Garcia, R.V.; Kaeding, G.; Chen, X. CNNs Under Attack: On the Vulnerability of Deep Neural Networks Based Face Recognition to Image Morphing. In Proceedings of the International Workshop on Digital Watermarking, Magdeburg, Germany, 23–25 August 2017; pp. 121–135.
6. Scherhag, U.; Rathgeb, C.; Merkle, J.; Breithaupt, R.; Christoph, B. Face Recognition Systems Under Morphing Attacks: A Survey. *IEEE Access* **2019**, *7*, 23012–23026. [CrossRef]
7. Ferrara, M.; Franco, A.; Maltoni, D. The Magic Passport. In Proceedings of the International Joint Conference on Biometrics (IJCB), Clearwater, FL, USA, 29 September–2 October 2014; pp. 1–7.
8. Busvine, D. *Germany Bans Digital Doppelganger Passport Photos*; Reuters: Toronto, ON, Canada, 2020.
9. Scherhag, U.; Debiasi, L.; Rathgeb, C.; Busch, C.; Uhl, A. Detection of Face Morphing Attacks Based on PRNU Analysis. *IEEE Trans. Biometr. Behav. Ident. Sci.* **2019**, *1*, 302–317. [CrossRef]
10. Ferrara, M.; Franco, A.; Maltoni, D.; Sun, Y. On the Effects of Image Alterations on Face Recognition Accuracy. In *Face Recognition Across the Electromagnetic Spectrum*; Bourlai, T., Ed.; Springer: Cham, Switzerland, 2016, pp. 195–222.
11. Kramer, R.S.S.; Mireku, M.O.; Flack, T.R.; Ritchie, K.L. Face morphing attacks: Investigating detection with humans and computers. *Cogn. Res.* **2019**, *4*. [CrossRef] [PubMed]
12. Burton, A.M.; White, D.; McNeill, A. The Glasgow face matching test. *Behav. Res. Methods* **2010**, *42*, 286–291. [CrossRef] [PubMed]
13. Jenkins, R.; White, D.; Van Montfort, X.; Burton, A.M. Variability in photos of the same face. *Cognition* **2011**, *121*, 313–323. [CrossRef] [PubMed]
14. Young, A.W.; Burton, A.M. Are we face experts? *Trends Cogn. Sci.* **2018**, *22*, 100–110. [CrossRef] [PubMed]
15. White, D.; Kemp, R.I.; Jenkins, R.; Matheson, M.; Burton, A.M. Passport officers' errors in face matching. *PLoS ONE* **2014**, *9*, e103510. [CrossRef] [PubMed]
16. Robertson, D.J.; Kramer, R.S.S.; Burton, A.M. Fraudulent ID using face morphs: Experiments on human and automatic recognition. *PLoS ONE* **2017**, *12*, e0173319. [CrossRef]
17. Robertson, D.J.; Mungall, A.; Watson, D.G.; Wade, K.A.; Nightingale, S.J.; Butler, S. Detecting morphed passport photos: A training and individual differences approach. *Cogn. Res. Princ. Implic.* **2018**, *3*, 1–11. [CrossRef]
18. AMSL Face Morph Image Data Set. Available online: https://omen.cs.uni-magdeburg.de/disclaimer/index.php (accessed on 31 March 2021).
19. Face Research Lab London Set. Available online: https://figshare.com/articles/dataset/Face_Research_Lab_London_Set/5047666 (accessed on 31 March 2021).
20. Utrecht ECVP Dataset. Available online: http://pics.stir.ac.uk/ (accessed on 31 March 2021).
21. Neubert, T.; Makrushin, A.; Hildebrandt, M.; Kraetzer, C.; Dittmann, J. Extended StirTrace Benchmarking of Biometric and Forensic Qualities of Morphed Face Images. *IET Biom.* **2018**, *7*, 325–332. [CrossRef]
22. Seibold, C.; Samek, W.; Hilsmann, A.; Eisert, P. Detection of Face Morphing Attacks by Deep Learning. In *Digital Forensics and Watermarking, Proceedings of the IWDW 2017, Magdeburg, Germany, 23–25 August 2017*; Kraetzer, C., Shi, Y.Q., Dittmann, J., Kim, H., Eds.; Springer: Cham, Switzerland, 2017; Volume 10431.
23. Lowe, D.G. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vis.* **2004**, *60*, 91–110. [CrossRef]
24. Dalal, N.; Triggs, B. Histograms of oriented gradients for human detection. In Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 7–12 June 2015; pp. 886–893.
25. Çeliktutan, O.; Ulukaya, S.; Sankur, B. A comparative study of face landmarking techniques. *J. Image Video Proc.* **2013**, *2013*, 13. [CrossRef]
26. Cootes, T.F.; Taylor, C.J.; Cooper, D.H.; Graham, J. Active shape models-their training and application. *Comput. Vis. Image Understand.* **1995**, *61*, 38–59. [CrossRef]
27. Cootes, T.F.; Edwards, G.J.; Taylor, C.J. Active appearance models. *IEEE Trans. Pattern Anal. Mach. Intell.* **2001**, *23*, 681–685. [CrossRef]
28. Seshadri, K.; Savvides, M. An analysis of the sensitivity of active shape models to initialization when applied to automatic facial landmarking. *IEEE Trans. Inf. Forensics Secur.* **2012**, *7*, 1255–1269. [CrossRef]
29. Wiskott, L.; Fellous, J.; Kruger, N.; von der Malsburg, C. Face recognition by elastic bunch graph. *IEEE Trans. Pattern Anal. Mach. Intell.* **1997**, *7*, 775–779. [CrossRef]
30. Pentland, A.; Moghaddam, B.; Starner, T. View-based and Modular Eigenspaces for Face Recognition. In Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 21–23 June 1994; pp. 84–91.
31. Çelik, T.; Özkaramanli, H.; Demirel, H. Facial feature extraction using complex dual-tree wavelet transforms. *Comput. Vis. Image Understand.* **2008**, *111*, 229–246. [CrossRef]
32. Smeraldi, F.; Bigun, J. Retinal vision applied to facial features detection and face authentication. *Pattern Recogn. Lett.* **2002**, *23*, 463–475. [CrossRef]
33. Akakin, H.Ç.; Sankur, B. Robust 2D/3D face landmarkingon. In Proceedings of the 2007 3DTV Conference, Kos, Greece, 7–9 May 2007; pp. 1–4.

34. Antonini, G.; Popovici, V.; Thiran, J.-T. Independent component analysis and support vector machine for face feature extraction. In Proceedings of the International Conference on Audio- and Video-Based Biometric Person Authentication, Guildford, UK, 9–11 June 2003; pp. 111–118.
35. King, D.E. Dlib-ml: A Machine Learning Toolkit. *J. Mach. Learn. Res.* **2009**, *4*, 1755–1758.
36. Kazemi, V.; Sullivan, J. One Millisecond Face Alignment with an Ensemble of Regression Trees. In Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR), Columbus, OH, USA, 23–28 June 2014; pp. 1867–1874.
37. Wandzik, L.; Kaeding, G.; Garcia, R.V. Morphing Detection Using a General-Purpose Face Recognition System. In Proceedings of the European Signal Processing Conference (EUSIPCO), Rome, Italy, 3–7 September 2018; pp. 1017–1021.
38. Saragih, J.M.; Lucey, S.; Cohn, J.F. Face alignment through sub-space constrained mean-shifts. In Proceedings of the ICCV, Kyoto, Japan, 29 September–2 October 2009; pp. 1034–1041.
39. Ruprecht, D.; Muller, H. Image warping with scattered data interpolation. *IEEE Comput. Graph. Appl.* **1995**, *15*, 37–43. [CrossRef]
40. Sederberg, T.; Parry, S. Free-form deformation of solid geometric models. In Proceedings of the Special Interest Group on Comp. Graphics and Interactive Techniques (SIGGRAPH), Dallas, TX, USA, 18–22 August 1986; pp. 151–160.
41. Lee, S.; Chwa, K.; Shin, S. Image metamorphosis using snakes and free-form deformations. In Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques (SIGGRAPH), Los Angeles, CA, USA, 6–11 August 1995; pp. 439–448.
42. Beier, T.; Neely, S. Feature-based image metamorphosis. In Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques (SIGGRAPH), Chicago, IL, USA, 26–31 July 1992; pp. 35–42.
43. Schaefer, S.; McPhail, T.; Warren, J. Image deformation using moving least squares. In Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques (SIGGRAPH), Boston, MA, USA, 30 July–3 August 2006; pp. 533–540.
44. Seibold, C.; Samek, W.; Hilsmann, A.; Eisert, P. Accurate and robust neural networks for face morphing attack detection. *J. Inf. Secur. Appl.* **2020**, *53*, 102526. [CrossRef]
45. Makrushin, A.; Neubert, T.; Dittmann, J. Automatic Generation and Detection of Visually Faultless Facial Morphs. In Proceedings of the International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP), Porto, Portugal, 27 February–1 March 2017, pp. 39–50.
46. Wu, J. *Face Recognition Jammer using Image Morphing*; Boston University: Boston, MA, USA, 2011.
47. Wu, E.; Liu, F. Robust image metamorphosis immune from ghost and blur. *Vis. Comput* **2012**, *29*, 311–321. [CrossRef]
48. Seitz, S.M.; Dyer, C.R. View morphing. In Proceedings of the Special Interest Group on Comp. Graphics and Interactive Techniques (SIGGRAPH), New Orleans, LA, USA, 4–9 August 1996; pp. 21–30.
49. FaceFusion. Available online: http://www.wearemoment.com/FaceFusion/ (accessed on 31 March 2021).
50. Scherhag, U.; Rathgeb, C.; Merkle, J.; Busch, C. Deep face representations for differential morphing attack detection. *arXiv* **2020**, arXiv:2001.01202.
51. Weng, Y.; Wang, L.; Li, X.; Chai, M.; Zhou, K. Hair interpolation for portrait morphing. *Comput. Graph. Forum* **2013**, *32*, 79–84. [CrossRef]
52. Raja, K.; Ferrara, M.; Franco, A.; Spreeuwers, L.; Batskos, I.; Gomez-Barrero, F.D.M.; Scherhag, U.; Fischer, D.; Venkatesh, S.; Singh, J.M.; et al. Morphing Attack Detection—Database, Evaluation Platform and Benchmarking. *arXiv* **2020**, arXiv:2006.06458.
53. Face Morpher. Available online: https://github.com/alyssaq/face_morpher (accessed on 31 March 2021).
54. Hatcher, G.; Yu, W. A survey of deep learning: Platforms, applications and emerging research trends. *IEEE Access* **2018**, *6*, 24411–24432. [CrossRef]
55. Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. ImageNet Large Scale Visual Recognition Challenge. *Int. J. Comput. Vis.* **2015**, *115*, 211—252. [CrossRef]
56. Sze, V.; Chen, Y.-H.; Yang, T.-J.; Emer, J.S. Efficient Processing of Deep Neural Networks: A Tutorial and Survey. *Proc. IEEE* **2017**, *105*, 2295–2329. [CrossRef]
57. Simonyan, K.; Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv* **2014**, arXiv:1409.1556.
58. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. Imagenet classification with deep convolutional neural networks. In Proceedings of the Advances in Neural Information Processing Systems, Lake Tahoe, NV, USA, 3–8 December 2012; pp. 1097–1105.
59. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
60. Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; Rabinovich, A. Going deeper with convolutions. In Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 7–12 June 2015; pp. 1–9.
61. Samek, W.; Wiegand, T.; Müller, K.R. Explainable Artificial Intelligence: Understanding,Visualizing and Interpreting Deep Learning Models. *arXiv* **2020**, arXiv: 1708.08296.
62. Ferrara, M.; Franco, A.; Maltoni, D. Face demorphing. *IEEE Trans. Inf. Forensics Secur.* **2017**, *13*, 1008–1017. [CrossRef]
63. Zhang, L.B.; Peng, F.; Long, M. Face morphing detection using Fourier spectrum of sensor pattern noise. In Proceedings of the 2018 IEEE International Conference on Multimedia and Expo (ICME), San Diego, CA, USA, 23–27 July 2018; pp. 1–6.
64. Autherith, S.; Pasquini, C. Detecting Morphing Attacks through Face Geometry Features. *J. Imaging* **2020**, *6*, 115. [CrossRef]
65. Neubert, T.; Kraetzer, C.; Dittmann, J. A face morphing detection concept with a frequency and a spatial domain feature space for images on eMRTD. In Proceedings of the ACM Workshop on Information Hiding and Multimedia Security, Paris, France, 3–5 July 2019; pp. 95–100.

66. Spreeuwers, L.; Schils, M.; Veldhuis, R. Towards robust evaluation of face morphing detection. In Proceedings of the 2018 26th European Signal Processing Conference (EUSIPCO), Rome, Italy, 3–7 September 2018; pp. 1027–1031.
67. Raghavendra, R.; Raja, K.B.; Venkatesh, S.; Busch, C. Transferable deep-cnn features for detecting digital and print-scanned morphed face images. In Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017; pp. 10–18.
68. Cai, S.; Zhang, L.; Zuo, W.; Feng, X. A Probabilistic Collaborative Representation Based Approach for Pattern Classification. In Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016; pp. 2950–2959.
69. Ramachandra, R.; Raja, K.B.; Busch, C. Detecting morphed face images. In Proceedings of the International Conference on Biometrics: Theory Applications and Systems (BTAS), Niagara Falls, NY, USA, 6–9 September 2016; pp. 1–7.
70. Agarwal, A.; Singh, R.; Vatsa, M.; Noore, A. SWAPPED! Digital face presentation attack detection via weighted local magnitude pattern. In Proceedings of the International Joint Conference on Biometrics (IJCB), Denver, CO, USA, 1–4 October 2017; pp. 659–665.
71. Kraetzer, C.; Makrushin, A.; Neubert, T.; Hildebrandt, M.; Dittmann, J. Modeling attacks on photo-ID documents and applying media forensics for the detection of facial morphing. In Proceedings of the Workshop on Information Hiding and Multimedia Security (IHMMSec), Philadelphia, PA, USA, 20–22 June 2017; pp. 21–32.
72. Peng, F.; Zhang, L.B.; Long, M. Fd-gan: Face de-morphing generative adversarial network for restoring accomplice's facial image. *IEEE Access* **2019**, *7*, 75122–75131. [CrossRef]
73. Ortega-Delcampo, D.; Conde, C.; Palacios-Alonso, D.; Cabello, E. Border Control Morphing Attack Detection With a Convolutional Neural Network De-Morphing Approach. *IEEE Access* **2020**, *8*, 92301–92313.
74. Parkhi, O.M.; Vedaldi, A.; Zisserman, A. Deep face recognition. In Proceedings of the British Machine Vision Conference (BMVC), Swansea, UK, 7–10 September 2015; Volume 1, p. 6.
75. Venkatesh, S.; Ramachandra, R.; Raja, K.; Spreeuwers, L.; Veldhuis, R.; Busch, C. Morphed face detection based on deep color residual noise. In Proceedings of the International Conference on Image Proceedings of the Theory, Tools and Applications (IPTA), Istanbul, Turkey, 6–9 November 2019; pp. 1–6.
76. Venkatesh, S.; Ramachandra, R.; Raja, K.; Spreeuwers, L.; Veldhuis, R.; Busch, C. Detecting morphed face attacks using residual noise from deep multi-scale context aggregation network. In Proceedings of the Winter Conference on Applications of Computer Vision (WACV), Snowmass Village, CO, USA, 1–5 March 2020; pp. 280–289.
77. Venkatesh, S.; Zhang, H.; Ramachandra, R.; Raja, K.; Damer, N.; Busch, C. Can GAN Generated Morphs Threaten Face Recognition Systems Equally as Landmark Based Morphs? Vulnerability and Detection. In Proceedings of the Int. Workshop on Biometrics and Forensics (IWBF), Porto, Portugal, 29–30 April 2020; pp. 1–6.
78. *Information Technology Biometric Presentation Attack Detection—Part 3: Testing and Reporting, Standard 30107-3:2017*; International Organization for Standardization: Geneva, Switzerland, 2017.
79. NIST. Frvt Morph Web Site. 2020. Available online: https://pages.nist.gov/frvt/html/frvt_morph.html (accessed on 31 March 2021).