

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

Digital Object Identifier DOI

Stegano-morphing: Concealing attacks on face identification algorithms

LUIS CÁRABE¹ AND EDUARDO CERMEÑO² (Senior Member, IEEE)

¹Computer Science Department, UAM, Madrid 28049, Spain (e-mail: luiscarabe@gmail.com) ²Research Department, Vaelsys, Madrid 28043, Spain (e-mail:eduardo.cm@vaelsys.com)

Corresponding author: E. Cermeño (e-mail: eduardo.cm@vaelsys.com).

This work was partially funded by the Consejería De Ciencia, Universidad e Innovación, Comunidad de Madrid.

ABSTRACT Face identification is becoming a well-accepted technology for access control applications, both in the real or virtual world. Systems based on this technology must deal with the persistent challenges of classification algorithms and the impersonation attacks performed by people who do not want to be identified. Morphing is often selected to conduct such attacks since it allows the modification of the features of an original subject's image to make it appear as someone else. Publications focus on impersonating this other person, usually someone who is allowed to get into a restricted place, building, or software app. However, there is no list of authorized people in many other applications, just a blacklist of people no longer allowed to enter, log in, or register. In such cases, the morphing target person is not relevant, and the main objective is to minimize the probability of being detected. In this paper, we present a comparison of the identification rate and behavior of six recognizers (Eigenfaces, Fisherfaces, LBPH, SIFT, FaceNet, and ArcFace) against traditional morphing attacks, in which only two subjects are used to create the altered image: the original subject and the target. We also present a new morphing method that works as an iterative process of gradual traditional morphing, combining the original subject with all the subjects' images in a database. This method multiplies by four the chances of a successful and complete impersonation attack (from 4% to 16%), by deceiving both face identification and morphing detection algorithms simultaneously.

INDEX TERMS Access control, ArcFace, biometrics, deep learning, FaceNet, face recognition, identification, morphing, security, spoofing attack.

I. INTRODUCTION

ACE recognition is gaining momentum. Continuous improvements in this research field [9] have led to an increasing number of commercial applications. Today face recognition algorithms are implemented in a wide range of products and solutions. People counting cameras are used to estimate the number of clients getting into a store and draw some statistical conclusions on their characteristics, e.g., age or gender [2]. Most mobile phones in the market have embedded technology to unlock them with a simple look at the device [3]. More and more websites implement "Know your Customer" policies by comparing a photo ID with realtime capture of the applicant [4]. People identification is likely one of the most important uses in the field. Mobile phones or websites are just two examples of its utility in everyday life. However, this technology is implemented in a wide range of areas, such as access control in offices or airports, where it has proven particularly efficient.

Like in any other biometric technologies, people have tried to deceive face recognition systems [37]. We can find several examples in the literature. For instance, printing a photo of a subject and trying to use it for impersonation purposes [37], [38]. More sophisticated methods include creating a mask able to deceive 3D face recognition [5] or using a wearable face projector [6].

For some specific applications, like in the case of airports with Automated Border Control (ABC) systems, where it is unlikely that anyone would place an image in front of the camera without being noticed, morphing techniques have been particularly relevant. Originally, morphing techniques consisted of generating intermediate frames between two images to achieve a smooth transition between them. If used on two images of different faces, one could get frames that merge features of both faces into one. Depending on the level of morphing applied, one person would be recognized better than the other. In the ABC scenario, M. Ferrara *et al.* [7]

VOLUME , 2021

studied a way of applying morphing to successfully verify two different subjects by using only one ID picture.

This approach is interesting because it proved able to fool face verification systems. The morphing process however can be discovered, making the spoofing attempt a failure. Our work focuses on concealing the attack in such a way that humans or automated systems cannot detect that an image has been altered. In a first phase of the work, we have researched how different face identification methods behave against the morphing process. Face identification differs from face verification because in the latter we already have information about who the subject might potentially be. In some applications however this is not the case, for instance, when using a face image to check whether a user has already been registered in a web site.

This paper is divided into seven sections. In Section II, we present an overview of the state-of-the-art face recognition and morphing software, as well as a brief review of past spoofing attacks to face recognition algorithms. In Section III, we describe the method used for selecting the most robust algorithm against morphing and the proposed algorithm to defeat it. In Section IV, we explain the implementation of the method. In sections V and VI, we present the results of the experiments and their discussion. Finally, in Section VII, we make conclusions about the findings of our experiments.

subsectionContributions As far as we know, our work constitutes a novel contribution to the field since we have not found other research publications covering face identification algorithms tested with morphed images or any similar method specifically designed to deceive the face identification algorithm while passing undetected. Next, we summarize our main contributions.

- We present a study of six current techniques of face identification. Each technique is tested with morphed images to find the more robust one, considering robustness the quality of requiring a higher amount of morphing alteration to misclassify a subject.
- We propose a new method to reduce the amount of morphing alteration required to make a face identification algorithm misclassify a subject.

II. RELATED WORK

A. FACE RECOGNITION

As seen by A. Agrawal *et al.*, the face recognition process involves the location of the face in the image, followed by an analysis of the located face (for instance, extracting its features) and then a comparison of the analysis results against all the faces stored in the database, using a classifier. Face recognition methods can be divided into four main categories: holistic, local, hybrid and deep learning approaches [9],[10]. The local approach classifies according to specific facial features, whereas the holistic approach considers the whole face as a unit. The hybrid approach combines both techniques. Many recent advances have been made in the deep learning approach, using Convolutional Neural Networks (CNNs) that offer higher speed and better accuracy. The first simple and fast algorithm that worked well in a constrained environment, Eigenfaces, came in 1991 [11], based on the Principal Component Analysis (PCA) technique, which is included in the holistic approach. Fisherfaces was developed, using Linear Discriminative Analysis (LDA) and achieving a better performance over variation in lighting [12]. The third most popular technique in the holistic approach is Independent Component Analysis (ICA) [13], which has excellent efficiency. Other methods can be included in this category. Some of them, combined with the three techniques mentioned before, can obtain good recognition performance. For instance, Hafez *et al.* [14] used a Gabor filter and LDA.

In the local approach, we can find a simple method, Local Binary Pattern (LBP), used to extract features from any object. Zhang *et al.* [15] was the first to use it for face recognition. Other well-studied feature extractors utilized for face recognition are Scale-invariant Feature Transform (SIFT) [16] and Speeded-up Robust Features (SURF) [17], inspired by SIFT but with better execution time.

Hybrid techniques can offer high recognition rates. Because of their high complexity however, they are considered more challenging to implement and are therefore less popular than the others. An example can be found in A.A. Fathima *et al.* [18], where the authors combined Gabor wavelet with Linear Discriminant Analysis. Further examples can be found in [9].

The deep learning approach can be considered as a nonlinear holistic technique [9]. Nevertheless, some references [10] define it as a new category due to its novelty and great accuracy. A few examples showing very good accuracy with the verification problem in the Labeled Faces in the Wild database (LFW) [24] are: deepFace [19], developed in 2014, with an accuracy of 97.35%; DeepID3 [20] (2015, 99.43%), FaceNet [21] (2015, 99.63%), VGGFace [22] (2015, 98.95%), and Arcface [23] (2018, 99.83%).

LFW is an excellent database to test face recognition algorithms because it is an unconstrained database. Usually, algorithms struggle with lighting, location, setting, pose, or age variations, as well as occlusions or misalignment. For example, A. K. Agrawal *et al.* [25] discuss about how face recognition methods deal with pose, illumination or expression variations. Moreover X. Zhang *et al.* [26] provide more information about how the pose of the subjects affect the recognition. Lastly, G. H. Givens *et al.* [27] study a database with easy, medium and hard verification tasks. However, over time, algorithms have improved significantly in this area, so recent local and deep learning approaches can handle these problems better.

B. MORPHING

For many years the film and television industries have used morphing to obtain fluid transformations between two different frames using mesh warping methods [28] based on three stages: feature specification, warping, and blending. In the first step, correspondence between the two images is created (using a mesh). In the second stage, a geometrical alignment of the mesh is performed using warping [29]. In the latter, all warped images are aligned, so it only remains to merge each pixel's color value, using a cross-dissolve method.

A review of this morphing approach with other firstgeneration morphing methods, such as field morphing or radial basis functions, can be found in [30]. In his work [31], M. Steyvers analyzes field morphing with a greater mathematical perspective. More recently, U. Scherhag *et al.* presented an overview of the publicly available latest commercial and open-source face morphing tools [32]. Most of them are based on Delaunay triangulation [48], which we consider the principal approach to morphing until the appearance of Generative Adversarial Networks (GAN) that also showed promising results [34]. However, to this day, GAN performs worse than the more classic methods against face recognition systems like OpenFace (a face recognition implementation based on FaceNet) [34].

The steps followed by Delaunay triangulation-based methods are the same as in mesh warping but using different techniques to achieve each goal. The correspondence between the two images is made by determining face key landmarks (i.e., eyes, mouth, nose, face contour, etc) either manually or automatically (using software). Then, a Delaunay triangulation is applied using the landmarks as vertices for the nonoverlapping triangles. During warping, the corresponding triangles of both images suffer a geometrical transformation in order to be aligned. Finally, a linear blending is applied.

C. SPOOFING ATTACKS

Attacks on biometric recognition systems are not only carried out on facial recognition devices. In [35], the authors conduct spoofing attacks on fingerprint sensors, iris scanners, and facial recognizers. They conclude that the method's performance does not correlate with its vulnerability. In fact, in all of them, a satisfactory attack can be achieved. This statement is also supported by A. Hadid *et al.* [37], using 3D masks (face recognition spoofing) and fake fingerprints (fingerprintrecognition spoofing), among others. They also study how anti-spoofing methods can reduce the vulnerability of the systems.

If we focus on facial recognition attacks, morphing is not the only issue. In [37], [38], the authors explore some databases with presentation attacks. Presentation attacks consist of showing a printed image (or printed mask) to a camera with facial recognition software to fool it. In addition, A. Mohammadi *et al.* [38], prove that the higher the face verification accuracy, the higher is its vulnerability to presentation attacks. Apart from this, M. Ferrara *et al.* [39] study the effects of geometric distortions (barrel distortion, vertical contraction, and extension) and digital beautification on face recognition accuracy. Other digital manipulation techniques can be very harmful, e.g., face synthesis, attribute manipulation, and identity or expression swap [40].

M. Ferrara *et al.* [7] were the first to present a successful morphing attack in a simulation of an ABC, using

two commercial face recognition software tools. Applying GIMP+GAP, manually morphed images were created to verify the two contributing subjects with the same photo. They were able to achieve that for eleven pairs of subjects in both face verification tools. Moreover, in M. Ferrara et al. [39], the authors extend the experiment proving that human experts (border guard group) and non-experts, in most cases, do not detect morphed images. However, D. J. Robertson et al. [41] reveal that although the attack may go more unnoticed in untrained subjects, when the subjects receive morphing training, they tend to detect morphing with higher probability. Nonetheless, in their experiment, they use Psychomorph, which creates lower quality morphings (with more ghost artifacts) than GIMP+GAP. More examples of verification attacks can be found in L. Wandzik et al. [42] and U. Scherhag [43]. In the first one, they carried out the experiment using FaceNet, utilizing more than 3000 pairs with 22 morphed images between each pair, working with triplets of images (impostor-accomplice-morphing). In the second one, experiments were conducted to prove face verification's vulnerability both with printed and scanned images.

Another morphing attack perspective may be to protect the privacy of the users in video surveillance systems. P. Korshunov and T. Ebrahimi [44] study this problem along with its robustness and reversibility.

Finally, we would like to refer to some studies in which we can find out how some parameters can affect the success of a morphing attack, such as U. Scherhag *et al.* [32], M. Gómez-Barrero *et al.* [35], and U. Scherhag *et al.* [45]. In the first one the authors discuss how the morphing quality can be an important parameter. In the second one, they study the importance of the recognizer's threshold. The latter also adds the significance of the similarity between the impostor and the accomplice.

D. MORPHING DETECTION

Face anti-spoofing detectors already existed before morphing attacks [60]. H. Chen *et al.* [61], M. De Marsico *et al.* [62], and W. Sun *et al.* [63] present more recent anti-spoofing methods. In the first one they propose a two-stream convolutional neural network (TSCNN) to detect spoofing attacks even with lighting variations. In the second one, they present an efficient solution that can verify if a face is truly 3D. In the latter they show a method that detects presentation attacks and outperforms other detectors.

Regarding morphing, the first detector was presented by R. Raghavendra *et al.* [64], which successfully verified all the 450 morphed face images from a database. It belongs to a category of morphing detectors that operate in Single Image Morphing Attack Detection scenarios (S-MAD). It refers to algorithms that only analyze one photograph to verify its morphing. In contrast, Differential Morphing Attack Detection (D-MAD) group algorithms that analyze a pair of images, one of them being a trusted unaltered photograph that the algorithm uses to verify the morphing on the other image. Some state-of-the-art S-MAD algorithms can be found in [67]–[69]. In order to detect morphing successfully, the authors use different techniques, such as Fourier spectrum of sensor pattern noise, Binarised Statistical Image Features (BSIF), or Local Binary Pattern (LBP), respectively. Additional approaches can be found in [65], [66], where authors use Photo Response Non-Uniformity (PRNU) and deep learning techniques to perform the detection. Scherhag *et al.* [32] present a review of some of these methods, along with other D-MAD algorithms.

III. CONCEALING ATTACKS ON FACE IDENTIFICATION ALGORITHMS WITH MORPHING

A morphing attack is the alteration of a subject's portrait using morphing techniques leading to his misidentification. In our work, it is complete only when it meets two criteria: first, a face identification algorithm should not identify the morphed image, and second, the morphed image should appear as a genuine image to a potential auditor. Face recognition algorithms might be beaten or defeated by a morphing attack when the image resulting from a morphing process is not identified as the original subject. However, if the resulting image does not appear genuine, the attack cannot be considered complete. We call our method Stegano-morphing because we try to hide (stegano means *hidden* in Greek) the morphing process from both detectors, the face recognizer and the morphing detector.

Considering the research done in Section II-B, we have chosen a morphing method based on Delaunay triangulation, hereafter referred to as the traditional morphing method. At the warping and blending steps of the process, a parameter is taken into account. In the case of warping (α_w) , it conditions how much each position of each face's landmarks contributes to the morphed image. If $\alpha_w = 0$, only the first image's landmarks are taken into account. If $\alpha_w = 1$, only the landmarks of the second image are considered. The inbetween values achieve a linear combination of the positions of the landmarks of both contributing images. The blending step (α_b) has a similar behavior, the color of all the correlated pixels are combined using a linear transformation. $\alpha_b = 0$ only considers the first image and $\alpha_b = 1$ the second.

For simplicity, some implementations only use one parameter α , that reflects the general percentage of contribution of both faces in each step ($\alpha_w = \alpha_b = \alpha$). In our study, we use this simplification as a quantifier of the morphing process. For example, a morphing process of 5% means that $\alpha = 0.05$. The first subject of the pair will contribute to the final image by 95% in both the landmarks' position and the pixels' value. The second subject will contribute with the remaining 5%.

A. FACE IDENTIFICATION AND MORPHING DETECTION

As we have seen in Section II-A, face recognition is a very active research field, and different approaches to it are being studied. We have selected the more promising ones with care to include two from each category (excluding the hybrid):

• Eigenfaces

- Fishefaces
- LBPH
- SIFT
- FaceNet
- ArcFace

However, the experiments found in the literature do not consider morphing attacks against face identification. Our objective is to select the approach that performs better against these attacks. From our perspective, good performance means that the algorithm can correctly identify the original subject in images that have been morphed. Since morphing is an incremental process, we consider an algorithm to be more robust than another when the amount of morphing required to make it fail is higher. Therefore, the selection criteria is related to the frame at which the face recognition algorithm does not recognize the original subject but another (either the target subject or any other person).

The original image is morphed into 100 images with n% morphing $(n \in \{1, ..., 100\})$. We consider that the original image has been morphed 0%, the target image has been morphed 100%, and any other image in between has n% $(n \in \{1, ..., 99\})$ as the amount of morphing. The higher the percentage required to avoid that the recognizer correctly identifies the original subject, the more robust it will be considered.

We recommend that face recognition algorithms are trained with a database composed of N subjects, with a number of photos per subject between 5 and 20. This quantity helps to avoid imbalanced data and biased results. We have chosen pairs of similar-looking subjects. This should reduce the amount of alteration required to pass from the original image (referred to as A) to the target image (referred to as B).

Since a complete morphing attack has to pass undetected, we need to define a method to detect morphed images. The easy procedure is to invite human experts that will evaluate the resulting image. However, this method might not be the most consistent because the same person can change his evaluation of a particular image or because different people may have different opinions. Therefore, using a morphing detector algorithm seems a good idea. Since we only provide one image to the detector to get a morphing verification, we will need an S-MAD algorithm.

B. PROPOSED ALGORITHM

We have approached the problem of concealing morphing attacks as an optimization problem. We searched for the range of potential solutions while trying to minimize the amount of morphing required to beat the face recognition algorithm. This approach makes sense if we assume that the lower the amount of morphing, the higher the chances of passing undetected. Starting with Subject A's original image, potential solutions are created by an iterative process of gradual morphing that combines the original image and all the subjects' images in a database.

L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms

The problem can be represented as an *m*-ary tree, being m the number of images stored in the database. The root vertex would be the unaltered image of Subject A and the other vertices, the morphed images. Each branch would represent an n% morphing between the parent and a person in the database. The less modification the image has, the less detectable it will be, so the algorithm searches the vertex that causes misidentification with the lowest percentage of morphing (lowest depth) and lowest morphing detection, using a Breadth-first search. It starts at the root, then searches for a misidentification in all of its child nodes' images, then moves to the next depth level, and so on.

We use a morphing detector as an additional evaluator of the probability of a complete attack. The resulting solution is the combination of morphing procedures with the lowest amount of alteration and the lowest evidence reported by the morphing detector. In summary, our method requires:

- A robust face recognizer.
- A testing database.
- A morphing algorithm.
- A morphing detector (S-MAD).

It follows these steps:

- 1) The original photo is morphed 5% separately with all the photos available in the testing database.
- 2) The morphed photos are passed on to the face identifier:
 - a) If in all the images it still identifies the original subject, it goes to Step 3.
 - b) If a different person is identified in one or more morphed photos, it goes to Step 6.
- 3) The first ten images that most reduce the identification confidence are selected.
- 4) The morphing detector evaluates the ten images to get the photo with the least detectable morphing, outputting the one with the least morphing detection confidence. In case of a tie, it is resolved by the alphabetical order of the subjects used.
- 5) The algorithm goes back to Step 1, replacing the original photo of the subject with the surviving image.
- 6) The morphing detector also evaluates those images to get the resulting photo with the least detectable morphing, ending the algorithm.

Algorithm 1 shows the pseudocode of the proposed method. The traditional morphing algorithm (Delaunay) acts when the *morphing* subroutine is called, the recognizer when *identify* is called, and the morphing detector when *detector* is invoked.

When we say that the original image is morphed 5%, we mean that the first face of each morphing contributes 95% to the warping and blending process. For each iteration, the original image's contribution is reduced by the formula:

% of original image's contribution = $0.95^t, t \in \mathbb{N}$,

Algorithm 1: Proposed morphing attack

```
Function main():
   current_photo = original subject photo;
   iterations = 0;
   not_finished = TRUE;
   while not_finished do
       iterations++;
       morph list = [];
       for every image im in testing database do
           morph = morphing(current photo, im);
           add morph to morph_list;
       identified_morphs = identify(morph_list);
       if original subject is identified in all
        identified morphs then
           best_ten = first ten morphs that most
            reduce the identification confidence;
           surviving image = detector(best ten);
           current_photo = surviving_image;
       else
           valid_morphs = list with all misidentified
            images;
           not_finished = FALSE;
   return detector(valid morphs);
Function morphing (Subject A, Subject B):
   return 5% morphing of A with B;
Function identify (List of images):
```

return list of identifications of the input images;

Function detector (*List of images*): return the photo with the least morphing detection confidence among the input list;

being t the number of iterations performed. For instance, in the third iteration t = 3, three morphings have taken place, so the original subject's contribution is $95\%^3 \approx 85.74\%$.

Dealing with all the possible morphing paths has an exponential complexity over the testing database size. Suppose the database size is m (with N subjects, N < m), the database size without the original subject's pictures is m'. If the algorithm needs t iterations to finish, the complexity would be $O((m')^t)$. This is because, in each iteration, all the images of the previous iteration would be morphed with all the database. To reduce that computational cost, we have implemented a heuristic. It is reflected in steps 3–4 and manages to reduce to one the number of images that pass to the successive iteration. The heuristic chooses the photo that is closest to the goal in each iteration, and the complexity becomes linear $(O(t \cdot m'))$.

Additionally, the morphing detector is also used in the sixth step to make sure that we select the picture that gets closer to a complete attack. Fig. 1 shows a flowchart of the process.

VOLUME , 2021

L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms



FIGURE 1. Flowchart of the proposed algorithm.

IV. EXPERIMENTS

We have carried out experiments to compare the morphing robustness of face identification algorithms in order to select the best one. We have also tested our proposed attack method on the face identifier selected, and compared the results obtained with the traditional attack.

All the experiments have been performed in an HP Pavilion x360 14-cd0005ns laptop, with 8 Gbs of RAM and an Intel Core i3-8130U chip, running Ubuntu 18.04.5 LTS with bash 4.4.20(1)-release, Python 3.6.9 and Python 2.7.17. The versions of the libraries used are:

- OpenCV[33]: 3.4.2.
- Scikit-learn[53]: 0.21.3 in Python3 and 0.20.4 in Python2.
- Tensorflow[55]: 1.14.0 in Python3 and 1.7.0 in Python2.
- Numpy[71]: 1.18.2 in Python3 and 1.16.6 in Python2.
- Dlib[49]: 19.18.0.

A. SELECTION OF THE FACIAL IDENTIFICATION ALGORITHM

We have chosen different solutions for each face recognition algorithm category (excluding the hybrid). Within the holistic approach, Eigenfaces [11] and Fisherfaces [12] have been selected. As representatives of the second category (local approach), we have picked Local Binary Patterns Histogram LBPH [50] and SIFT [16]. The LBPH algorithm works by creating histograms of the binary patterns extracted by LBP [15]. As seen in [9], [12], [16], [46], these techniques have been well studied and have good performance when using frontal views of faces. FaceNet [21] and ArcFace [23] have been selected out of the deep learning category due to their excellent performance [10].

For the first three algorithms (Eigenfaces, Fisherfaces, and LBPH), we have downloaded a Python implementation of R. Raja [51] that uses the Face library of OpenCV to cover the feature extraction and classification. A Haar cascade classifier [52] is then used for face detection. Slightly

modifying the previous implementation, we have gotten a SIFT deployment, using the xfeatures2d OpenCV class to perform the SIFT feature extraction and the Scikit-learn library for classification using a Support Vector Machine (SVM) [54]. We have used a Tensorflow implementation of FaceNet [56] written in Python. It uses a pre-trained model for VGGFace2 [58] as the training dataset and the Inception-ResNet-v1 architecture [59], achieving an LFW accuracy of 99.65+-0.00252%. It also uses an SVM for classification. For ArcFace, we have used the Python implementation provided by its creators [57], adding an SVM for classification.

All the recognizers expect the testing subjects to be included in their training database, which is known as closedset identification. We have also made small changes in the code files to get similar behavior in all the implementations. Every algorithm used can output the top 5 identification matches of the face presented. The parameters of the Haar cascade classifier that worked better with our database were *scaleFactor=1.001*, *minNeighbors=2*, *min-Size=(90,90)*, *outputRejectLevels=True*. Regarding the SVM used on SIFT, we have considered the settings *kernel="poly"*, C=10, *gamma=0.0001*. We have left all the other configurations according to the original sources.

As we have seen, we need a fully automatic morphing implementation. We have used the Python code presented by S. Patel [47], based on OpenCV functions [36]. In order to find the face landmarks, it uses Dlib's facial landmark detector. Then, as we have seen, those landmarks are employed as vertices of the Delaunay triangles. Using the corresponding triangles, it performs warping and blending to obtain all the intermediate frames.

We have created a database based on LFW [24]. As seen in [10], it is a widely used database to test state-of-the-art face recognizers. The database has 5749 subjects, but, as mentioned earlier, we want only the ones that have between 5 and 20 images each (both numbers included). That filters the database to 366 people with a total number of 3062 images. The Haar cascade face detector does not correctly detect the subject face in 5 of the 3062 images because those images have more than one face present and the wrong face is detected. We deleted those images from the database. The deleted images are *Erika_Harold_0003*, *Hugh_Grant_0008*, *Igor_Ivanov_0014*, *Jean_Charest_0004*, and *Joe_Lieberman_0004*. That implies that Erika Harold now has four images instead of 5, we consider this as an exception.

To determine the pairs of subjects who look more alike, we have used the Similar-looking LFW database (SLLFW) [70], which offers 3000 pairs of similar-looking faces (using the images of LFW). We have picked 25 pairs of images from it considering two factors. First, the individuals must be included in our 366 subjects database. Second, the subjects need to have more than five photos to train once the similar-looking images selected are removed from the training database. Fig. 2 shows an example of one selected pair.

Considering all the pairs, there are 49 different images (Re-



(a) Matthew_Perry_0007.



(b) Rubens_Barrichello_0011.

FIGURE 2. Similar-looking pair.

nee_Zellweger_0009 appears twice). The training database of the morphing robustness comparison and selection experiments consists of 3062 - 5 - 49 = 3008 images of 366 subjects. In Table 1, we provide all the pairs used.

TABLE 1. Similar-looking pairs selected.

No.	Original subject	Target subject
1	Amelia_Vega_0003	Norah_Jones_0015
2	Ana_Guevara_0002	Ian_Thorpe_0006
3	Andy_Roddick_0008	Richard_Virenque_0004
4	Angelina_Jolie_0002	Britney_Spears_0004
5	Anna_Kournikova_0011	Jelena_Dokic_0007
6	Ben_Affleck_0002	Ian_Thorpe_0007
7	Bill_McBride_0010	Jon_Gruden_0002
8	Bill_Simon_0011	Ron_Dittemore_0001
9	Catherine_Zeta-Jones_0001	Salma_Hayek_0001
10	Edmund_Stoiber_0004	John_Snow_0003
11	Eduardo_Duhalde_0006	George_HW_Bush_0005
12	Fidel_Castro_0018	Mohamed_ElBaradei_0003
13	Hillary_Clinton_0010	Renee_Zellweger_0009
14	Howard_Dean_0003	Kevin_Costner_0005
15	James_Blake_0006	Mark_Philippoussis_0003
16	Jason_Kidd_0003	Leonardo_DiCaprio_0003
17	Jean-Pierre_Raffarin_0001	Joschka_Fischer_0012
18	Jimmy_Carter_0006	John_Snow_0004
19	Joan_Laporta_0007	Pierce_Brosnan_0006
20	John_Kerry_0005	Robert_Redford_0002
21	Julianne_Moore_0019	Nancy_Pelosi_0002
22	Kate_Hudson_0008	Mariah_Carey_0006
23	Matthew_Perry_0007	Rubens_Barrichello_0011
24	Mike_Martz_0005	Paul_ONeill_0003
25	Renee_Zellweger_0009	Sheryl_Crow_0001

B. PROPOSED METHOD AGAINST FACE IDENTIFICATION AND MORPHING DETECTION

Our aim is to evaluate the performance of our method as a complete morphing attack system. Therefore, we have tested it with the (already trained) most robust face recognition system selected in the previous part and a morphing detector.

The basic morphing operation required in the algorithm is implemented with the traditional morphing processing technique based on Delaunay Triangulation.

Regarding the Single Image Morphing Attack Detector, we have tried the algorithms of [67]–[69]. The one that had the best performance and integration in our scenario has been

the detector presented by R. Raghavendra *et al.* [68], which has better results than other state-of-the-art alternatives. Although it is designed to detect morphing in printed-scanned photographs, it achieves excellent detection results in our scenario (Fig. 3), and therefore, it is the morphing detector used.

To train and test the S-MAD, we have picked the LFW subjects' images not used in the training database of the recognizers (people with n images, n < 5 or n > 20). We have split the subjects randomly into two groups, one for testing and the other one for training. Due to Matlab memory limitations, we have trained the detector using 3000 bonafide (not altered) images from the training group and 3500 morphed images. The morphed images were created randomly using pairs from the subjects included in the training group, covering all percentages between 1 and 99. Analogously, we have tested the detector using 500 bonafide images and 500 morphed images. Fig. 3 represents the ROC and FAR vs. FRR curves obtained, showing the excellent performance achieved.



FIGURE 3. Performance of the morphing detector.

To create the testing database, we have randomly selected

2350 images from LFW subjects not used to train the recognizers or the S-MAD.

We have tested the proposed morphing attack using 25 individuals as original subjects based on the first image of the similar-looking pairs used previously. Table 2 presents all the participants.

TABLE 2. Subjects selected to test our proposed method.

No.	Subject	No.	Subject
1	Amelia_Vega_0003	14	Howard_Dean_0003
2	Ana_Guevara_0002	15	James_Blake_0006
3	Andy_Roddick_0008	16	Jason_Kidd_0003
4	Angelina_Jolie_0002	17	Jean-Pierre_Raffarin_0001
5	Anna_Kournikova_0011	18	Jimmy_Carter_0006
6	Ben_Affleck_0002	19	Joan_Laporta_0007
7	Bill_McBride_0010	20	John_Kerry_0005
8	Bill_Simon_0011	21	Julianne_Moore_0019
9	Catherine_Zeta-Jones_0001	22	Kate_Hudson_0008
10	Edmund_Stoiber_0004	23	Matthew_Perry_0007
11	Eduardo_Duhalde_0006	24	Mike_Martz_0005
12	Fidel_Castro_0018	25	Renee_Zellweger_0009
13	Hillary_Clinton_0010		

C. ATTACK COMPARISON

We have conducted experiments comparing the proposed method with the traditional morphing attack to have a better feeling of its concealing features.

Since we have used the same recognizer, training database, and original subjects in both traditional and proposed morphing techniques, we can compare the percentages of alteration needed to cause misidentification using both approaches. Moreover, we have used the selected S-MAD with the morphed images that cause misidentification with the traditional method, so we can also compare the the morphing detected and complete attacks achieved with both morphing methods.

V. RESULTS

A. SELECTION OF FACE RECOGNIZER

Fig. 4 shows the face identification algorithms comparison of their robustness against morphing. It is divided into three plots. Fig. 4a exhibits the face recognizers' comparison analyzing the top 1 identification matches. Fig. 4b analyzing the top 3. Fig. 4c the top 5. Their x-axes represent the level of morphing in the pairs. 0% morphing symbolizes the unaltered image of the first subject of the pair (original subject), 100% the second subject, and the rest of percentages the inbetween morphings. Their y-axes reflect the percentage of couples who still have their original subject identified within the top for each morphing level.

We can observe that the identification percentages rise as we increase the top analyzed. However, the three graphs show a similar robustness ranking:

- 1) FaceNet ArcFace
- 2) LBPH
- 3) Eigenfaces

By far, FaceNet and ArcFace are above all the other recognizers. They are the ones that take the longest to misidentify the original subject. We can see that the identification percentage of ArcFace is above FaceNet with low levels of morphing, but at some point, it drops below the percentage achieved with FaceNet. In the top 1, this point is reached with 65% morphing, in the top 3 with 47% and in the top 5 with 49%. Next is LBPH, having a misidentification distance with FaceNet of more than 50% and more than 60% with ArcFace in some cases. Then Eigenfaces, followed by Fisherfaces and SIFT, which are the last ones and have a very similar performance (especially analyzing the top 3 and 5). These positions are maintained in practically all the three graphs' morphing levels, except for some exceptions or ties, e.g., beyond 80% morphing in Fig. 4a.

Each top's best identification scores are achieved, with 0% morphing, by ArcFace, being 100% in the three cases. FaceNet is in the second position, achieving 84%, 96%, and 100%, respectively. Not even LBPH passes the 50% of identification of the original subject. However, once the 100% morphing is reached, only in the top 3 and 5 the original subject is still identified in some pairs of ArcFace, FaceNet, LBPH, and Eigenfaces.

As ArcFace is the recognizer that achieves the highest identification rate in most of the in-between morphings (in the top 1), we consider it to be the most robust recognizer. Therefore, it is the selected algorithm. These experiments' complete results can be found in Appendix A, reflecting all the percentages where Subject A (first member of the pair) or B (second member) are recognized for every pair of images.

B. PROPOSED MORPHING ATTACK

Fig. 5 presents the summary of the results of our proposed morphing attack. It contains three plots. The first one (Fig. 5a) represents the number of iterations required to make ArcFace misidentify the original subject. The second plot (Fig. 5b) presents the necessary decrease in the original subject's contribution to the morphing in order to achieve the misidentification. Fig. 5c shows the percentage of complete (undetected) attacks depending on the morphing detector's confidence threshold. The threshold is the confidence needed to classify an image as morphed.

We can see that most images needed eight iterations or less to finish, being seven iterations the most common case (28% of the images). This means that in the majority of cases, 30% of modification is enough to cause misidentification $(1 - 0.95^7)$. In addition, even with a small alteration (5– 10%) some images achieve an incorrect identification. Only 16% of the subjects needed more than eight iterations. The maximum number of iterations required has been twelve, so, if the subject contribution is down to 54% (0.95¹²), all the images obtain the original subject's misidentification.

With a 100% confidence threshold of morphing detection, 16% of the subjects achieve a complete attack. However, this percentage drops to 8% with a threshold of 99.46%. This reflects the excellent performance of the morphing detector. Nevertheless, it is not infallible, and we can conceal 8% of the attacks if the threshold is set above 1.5%.

L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms



FIGURE 4. Percentage of morphed images identified as the original subject for each level of morphing.

The experiments' complete results, including the initial and final image of each subject can be found in Appendix B.

C. ATTACK COMPARISON

Fig. 6 displays the comparison of the traditional morphing attack technique and our proposed method. It is divided into two plots. The first one (Fig. 6a) compares the percentage of subjects misidentified by ArcFace for each level of contribution of the original subject in the morphed image. In the second plot (Fig. 6b), we can see the comparison of the percentages of complete (undetected) attacks depending on the morphing detector's confidence threshold.

The new method achieves misidentification much faster than the traditional method. With the subject's contribution down to 63%, they get 92% and 20% of misidentification, respectively. In addition, with the subject's contribution down to 54%, the misidentification rises to 100% and 40%, respectively. The traditional method needs the contribution to go down to 29% so that the misidentification reaches 100%. Our method also achieves a higher percentage of complete attacks, up to 16% depending on the threshold, compared to the 4% that the traditional attack accomplishes at most.

VI. DISCUSSION

A. PERFORMANCE OF FACE RECOGNIZERS AGAINST TRADITIONAL MORPHING

The results obtained in facial identification on the LFW database are considerably worse than those obtained in verification (except with ArcFace). This might be expected since, for identification, we work 1 vs. N (N = 366 in our database), and regarding verification, we work 1 vs. 1. Thus, as mentioned in [72], the difficulty of identification is related to the number of subjects contained in the database. Some examples are Eigenfaces, in which we have obtained 16% of identification accuracy in contrast with 60.02% of verification accuracy [72], and FaceNet, with 84% and 99.6% of identification and verification accuracy, respectively [72].

FaceNet and ArcFace present a good performance identifying the in-between morphed images correctly. This means that in the case of a real attack, the attacker would need to significantly alter the image to fool the recognizers. Taking a look at Table 3, we can see that analyzing the top 1, the attacker would need a 43% morphing alteration with FaceNet





(a) Percentage of subjects that are misidentified by ArcFace.

(b) Percentage of subjects that are misidentified by ArcFace.



(c) Percentage of complete attacks depending on the morphing detector's classification threshold, i.e., the confidence needed to classify an image as morphed.

FIGURE 5. Summary of the results of the proposed morphing attack.

and 51% with ArcFace to have more than a 50% chance of the attack being successful. If we analyze the top 3, the required morphing alteration is higher than 66% and 57 % respectively. Finally, if we analyze the top 5, the alteration needed rises to 71% and 64%. FaceNet and ArcFace show such good results that some attacks will fail even with the original image wholly modified (100% morphing) if we consider top 3 or top 5 lists.

TABLE 3. Accuracy of FaceNet (F) and ArcFace (A) at different percentages of morphing using the traditional method.

					% o	f morp	hing			
		0%	43%	50%	51%	57%	64%	66%	71%	100%
F	Top 1	84%	48%	32%	32%	16%	16%	12%	12%	0%
	Top 3	96%	80%	72%	72%	64%	52%	48%	32%	8%
	Top 5	100%	80%	76%	76%	72%	68%	64%	48%	12%
	Top 1	100%	72%	52%	40%	24%	16%	8%	0%	0%
A	Top 3	100%	80%	64%	64%	48%	20%	16%	12%	4%
	Top 5	100%	84%	72%	72%	60%	48%	32%	16%	8%

Since the morphing process converts the original image progressively into the target one, we may expect to obtain identification results transitioning from the former to the latter. However, this only happens with ArcFace and FaceNet (in the latter only with some pairs). The other recognizers studied have behaviors such that they identify other subjects in some intermediate morphings, and they might even recognize the original subject intermittently. For example, Table 7 shows that in the fourth pair, Fisherfaces recognize the original subject in 0-28% and 34-36% of morphing. On the contrary, FaceNet and ArcFace have a much more regular and expected performance. For instance, Table 10 exhibits that in the case of the fourth pair, the original subject is identified in the top 1 in 0-54% morphing, then she goes to the second and third position in the top in 55-58% morphing and finally to the fourth and fifth position in 59-64% morphing. On the contrary, the target appears in the fourth and fifth position in 44-53% morphing. Then she goes up to the second and third position at 54-59% morphing. Finally, she remains in the top



FIGURE 6. Comparison of attack methods.

1 in 60-100% morphing.

The percentage considered for the results has been the first percentage at which the original subject ceases to be recognized, regardless of whether he is recognized again in later percentages or not. Another interesting approach could be to study all these intermediate percentages where the original subject is identified again.

B. RESULTS ACHIEVED BY OUR PROPOSED METHOD

Our morphing method requires a considerably lower amount of morphing process to fool ArcFace. Table 4 shows that ArcFace misidentifies 16% of the images where the original subject contributes with 90% of the information. This is especially interesting if we consider that with the traditional morphing technique, the success rate is 0%. Moreover, our method successfully beats ArcFace in all the cases when the original subject contributes with 54% or less to the morphed image. The traditional method is much less capable since it requires that only 29% of the original image remains to get all the attacks passed by.

TABLE 4. Comparison of misidentification of the original subject by ArcFace for each level of his contribution on the image using the traditional (trad) morphing and our proposal (ppsd). Higher is better.

		Contribution of the original subject											
	100%	90%	81%	73%	66%	63%	54%	42%	29%				
Trad.	0%	0%	4%	16%	20%	20%	40%	80%	100%				
Ppsd.	0%	16%	24%	48%	84%	92%	100%	100%	100%				

Moreover, the performance of the morphing detector is also remarkable. The traditional morphing is not able to reach more than 4% of complete attacks (Table 5). This means that the morphing detector can detect 96% of the attacks even if the required threshold is higher than 99.93%. If it is below this value, then it detects 100% of the attacks. Our method improves these results in a very significant way. For example, when 100% confidence is required we can achieve 16% of complete attacks, four times more than the traditional morphing technique. The improvement decreases with the demanded confidence such that when we require a certitude between 1.49 and 99.46%, we achieve 8% of complete attacks versus 0% from the other method. Only when the threshold is lowered to meaningless values (0.09%) the number of complete attacks is equal in both cases, 0%.

 TABLE 5.
 Comparison of complete attacks depending on the morphing detector's classification threshold of the morphing detector using the traditional (trad) morphing and our proposal (ppsd). Higher is better.

		Threshold											
	100%	99.46%	93.93%	1.49%	1.48%	0.09%							
Trad.	4%	4%	0%	0%	0%	0%							
Ppsd.	16%	8%	8%	8%	4%	0%							

If we consider the option of a human being as a morphing detector, the detection accuracy might be lower, and therefore the number of complete attacks could be higher. Fig. 7 shows two images that most people would consider equal, whereas the morphing detector is 100% sure that Fig. 7b has been morphed.





(b) Morphed image.

(a) Unaltered image.

FIGURE 7. Initial and final (misidentified) images of subject number 17.

VOLUME , 2021

L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms

The experiments' results are also conditioned by other parameters, such as the size of the database or the amount of morphing per iteration. The bigger the database, the more possible morphing combinations. With a smaller percentage of morphing added at each iteration, we could get closer to an optimal result in exchange for increasing the number of iterations.

VII. CONCLUSIONS

Most of the existing face recognition techniques perform poorly when trying to classify images with some morphing degree. The exceptions are Deep Learning based algorithms (ArcFace) that can robustly detect 100% of the images that have less than 18% of morphing alteration.

We have presented a new way of attacking face identification systems that minimizes the chances of being detected by a morphing detector system. We have called the new method: Stegano-Morphing. The results outperform previous morphing techniques by 300% in the best case. A soft modification of 34% of the original image is enough to make the best identification algorithm misclassify almost 85% of the subjects.

However, a morphing detector algorithm is still able to catch more than 84% of the attacks. The results from this detector are awe-inspiring compared with our own human perception, which would be unable to notice any alteration on the images in many positive cases. Using machine-based auditors to verify if an image has been morphed seems necessary when trying to avoid spoofing attacks based on morphing faces.

APPENDIX A DETAILED RESULTS OF THE ROBUSTNESS AGAINST MORPHING

In tables 6 to 11, all the results of the face recognizers' comparisons against traditional morphing are presented. For each recognizer, the tables show in what percentages of morphing alteration is the original or the target subject identified (Column A for the original subject and B for the target). Column No. reflects the pair number (to see the individuals' names check Table 1). Pos. 1, Pos. 2–3 and Pos. 4-5 represent the first, second and third, fourth and fifth positions (respectively) of each face recognition algorithm's top identification matches. For clarity, rows of couples that have not been identified within the top 5 at any percentage have been deleted.

L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms

	Pos. 1		Pos	. 2–3	Pos	s. 4–5
No.	Α	В	Α	В	Α	В
1	72–100	-	70-71	98-100	69	94
2	_	-	-	_	_	72–76
4	0-20	-	21-34	-	35–37	_
5	100	-	91-92, 95-99	95-100	93–94	92, 94
6	_	-	_	_	_	94-95, 97-100
8	0-15	72-89	16-19	52-71,90-100	20-28	48-51
10	43-92	-	9-42, 93-97, 100	_	0-8, 98-99	_
13	_	-	0-16,38	-	17-20, 27, 29-30, 35,	_
					39-41, 43-44	
15	_	95-100	-	82–94	0-12, 14-16	77-81
16	_	-	-	56, 59-69, 71-82	_	51-55, 57-58,70, 83-89,
						91–92
18	30-31, 33-38, 40-44,	-	25-29, 32, 39, 45, 48-52,	_	22-24, 75-76,78, 80,	_
	46-47, 53, 55, 60-62, 74		54, 56-59, 63-73, 77		86-89, 91-92, 98-100	
19	_	-	_	89, 91	_	71-88, 90, 92-100
20	0-12, 15	-	13-14, 16-19	-	20-25, 54, 56, 58, 63	_
22	_	-	-	_	_	40-41, 49, 51-54
24	0-54	-	55-60	14–19	61–63	0-13, 20-22
25	_	-	-	64, 67-71, 73, 77, 80-83	_	61-63, 65-66, 72, 75-76,
						78-79, 84-90, 92-93

TABLE 6. Complete results of the robustness of Eigenfaces against traditional morphing.

TABLE 7. Complete results of the robustness of Fisherfaces against traditional morphing.

	~				D 4 5			
	Pos.	.1	Pos.	. 2–3	Pos	. 4–5		
No.	Α	В	A	В	Α	В		
4	0-28, 34-36	-	29-33, 37-50	-	51-52, 59	-		
7	-	-	-	-	-	53-54, 58-59, 62-68,		
						70-75, 77-81, 83-93, 95,		
						98–99		
12	33, 69, 81	_	70–74, 77, 79, 84, 86–87	-	41, 50, 65–68, 75–76, 78,	_		
					80-81, 83, 85, 89, 91, 93			
13	-	-	-	-	-	87, 99		
15	-	_	0-1	-	_	_		
16	-	-	-	100	-	-		
17	-	_	-	50-51,94	_	52-57, 60, 62-63, 81, 83,		
						96, 99–100		
18	_	_	25-26	_	_	27		
19	0-1, 3, 6	28-29, 77-78	2, 4–5, 7–13, 16–17, 19	25-26, 31, 33, 35, 75-76	15	15, 24, 27, 30, 32, 74, 83,		
						85		
20	-	_	0-4	-	5, 11, 55, 60-62	-		
25	-	-	2	-	0	-		



L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms

TABLE 8. Complete results of the robustness of LBPH against traditional morphing.

	Pos	.1	Pos.	2-3	Pos. 4–5			
No.	Α	В	Α	В	Α	В		
1	38-39, 47-48, 50	_	40-41, 45, 49, 51, 53-55,	-	25, 30, 35–36, 42–44, 52,	90, 100		
			65, 75, 78–79, 81–90, 92,		56-60, 64, 77, 80, 91, 93			
			94					
2	-	37-38, 40-100	-	34–36, 39	28	-		
3	-	-	-	-	30	_		
4	-	95–96	6	69-81, 83-86, 88, 90-94	1-2, 44, 47, 51-52, 54-56, 66-67, 90-94, 97-100	82, 87, 89		
5	90–95, 97–98	-	67–68, 70–79, 83–85, 87–89, 96, 99–100	-	61, 63–66, 69, 81–82, 86	-		
7	-	-	17, 26–27, 29, 31, 33, 36,	-	22, 24, 30, 37, 39, 41, 43,	-		
	0 6 8 10 22 24	95 96 07 100	7 20 22 25 25 20	82 84 87 02 04 06	45	02		
ô	0-0, 8-19, 25-24	83-80, 97-100	7, 20–22, 23–33, 39	82-84, 87-92, 94-90	50-56, 40	93		
9	_	_	_	_	0, 5, 20–22, 24–52, 30, 40, 44, 51–52, 63, 78, 80–82,	_		
12	5 (46 54 95 06	4	44 45 47 52 86 05	84, 87	27 42 42		
12	3-0	40, 54–85, 90	4	44-45, 47-55, 86-95, 97-100	2-3	37, 42–43		
13	0-5, 7-15, 22, 44, 51-57	98-100	6, 16–21, 23–43, 45–50,	78, 80, 83–90, 92–97	61–62	76–77, 79, 81–82, 91		
14			58-00		0 1 2 7 14 20 22			
14	-	_			0-1, 3, 7-14, 20-23, 25-26, 28	-		
15	0-21, 26-32, 34-35, 37,	43-100	22–25, 33, 36, 38, 43–44,	41	45-48, 53, 71, 76, 82, 86,	31-40, 42		
	39–42		49–52, 70, 74, 85, 87, 90		88, 91			
16	0-13, 15	56, 63–64, 70, 77, 79,	14, 16, 20	49, 53–55, 57–62, 65–67,	17, 19, 21	50-51, 68		
		81-100		69, 71–76, 78, 80				
17	-	23–25, 27–28, 30–71, 73–100	_	20–22, 26, 29, 72	10	16–19		
18	0-80, 83-85, 87-94	-	81-82, 86, 95-100	82, 86, 89–92	-	47–48, 59, 61–81, 83–85, 87–88, 93–97		
19	0-1, 3-11	31, 42, 44, 57-58	2, 18–19, 21–23, 28	9-11, 14, 32-40, 46-48,	14-17, 20, 24-27	7, 12–13, 15–16, 19, 24,		
		- , , ,	, , - , -	59, 65–67, 71–95, 99–100		28-30, 41, 45, 49-51,		
						54-55, 61-64, 68, 70,		
						96–98		
20	0-35, 37-54, 56-67	81, 86, 91-97, 99-100	36, 55, 68–69	85, 87, 89-90, 98	70-72, 74, 76	84, 88		
22	_	_		45, 48-49, 51-63	_	46-47, 50, 64		
24	0-50	-	51, 53, 55–56	_	52, 54, 58	_		
25	2	-	0-1, 3-16, 19-22, 31, 33	94–99	17-18, 23, 32, 34, 36	86		

L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms

TABLE 9. Complete results of the robustness of SIFT against traditional morphing.

	Pos	s. 1	Pos.	2-3	Pos	. 4–5
No.	Α	В	Α	В	Α	В
1	-	55, 71, 75, 86, 95–96	-	39, 59, 62, 64–65, 67,	-	61, 84-85, 93, 98
				69–70, 72–74, 82, 76–79,		
	12.14			88, 90, 94, 97	27 12 52 57	
3	43-44	-	38, 46–47, 49, 52, 55, 97	-	37, 42, 53, 57	-
4	—	07, 81, 80-87, 89-100	_	58, 66, 69-70, 78-80,	9	55, 85
5	34 79 81	_	33.82	02-04, 00	21 20 73 83	_
6	-	_	55,82	_		_
8	68-69, 80, 82, 84-85	_	23, 72, 75, 79, 96	63	5, 24, 86	_
9	39	10. 33-37. 39-40. 50. 54.	49	31-32, 38, 41-43, 45, 51,	_	49, 95
-		57-59, 64-65, 67-68, 70,	-	53, 56, 60–63, 69, 71–73,		- ,
		86-87, 93		75, 83-84, 94, 96, 99		
10	-	32, 34, 50	-	33, 36	-	-
11	10	70	4, 11	-	-	-
12	18, 20, 22, 28–29, 38	-	1-2, 8, 12-14, 19, 25-27,	55, 60, 76, 80	0, 4 , 10	-
			30–31, 33–34, 75			
13	0–9, 11, 13–14, 16, 18,	-	10, 15, 17, 20–22, 27, 29,	43, 80, 91	12, 25	-
1.5	30-31	55 56 50 100	32	50 54 50	5 ((7 00	50.57
15	3, 11-12, 14, 16-19, 21,	55-56, 59-100	0-2, 4, 7-8, 10, 22-26, 34,	50, 54, 58	5-6, 67, 89	52, 57
	27-55, 55-45, 45-49,		44, 50, 58-59, 61-62, 69, 71, 77, 80, 81, 87, 61, 64			
	51-52, 50, 00		99_100			
16	0 10 14 17-19 22 27	_	6 9 15-16 20 23-24	_	25-26 32	_
10	0, 10, 11, 17, 17, 22, 27		28–30, 33		20 20, 02	
17	_	33-40, 42-44, 47-53, 56,	_	0, 13, 41, 46, 54, 57, 62,	_	8, 32, 45, 69
		58, 61, 66-68, 70-100		65		
18	0-10, 12, 13, 16-18, 20,	62, 73, 77, 81	11, 14, 27, 33, 35–37, 52,	56, 65, 76, 79-80	38	83
	23, 25, 28–30, 34, 39,		55, 59			
	43–51, 53–54, 57, 64					
19		-	12	-	_	45
20	6, 8, 11, 13–15, 21, 25, 27,	-	7, 10, 16, 19–20, 31, 53	-	9, 28, 38	-
21	29	72	12 15 16 18 10 21 22	72 75		0
21	0, 3, 20, 27	12	13, 15–16, 18–19, 21–22,	/3-/5	-	0
22			20-20	18 36 47		34 48
23	-	_	15 18 44	-	_	-
24	-	76	_	75, 84, 95, 99	8	-
25	10, 18, 94	_	11-12, 14-17, 19-20, 30.	_	13, 31–32, 45	-
-	- / - / -		51, 91, 95–96		-,, -	

TABLE 10. Complete results of the robustness of FaceNet against traditional morphing.

	Po	os. 1	Pos.	2–3	Pos	5. 4–5
1	-	45-100	0–7	32–44	8-19, 21	20, 22–31
2	0-24	25-100	25-43	8-24	44-49	4–7
3	0-72, 74	75-100	73, 75–100	65, 68, 71-74	_	18, 20–21, 23–25, 33–34,
						36-39, 42-64, 66-67,
						69–70
4	0–54	60-100	55–58	54–59	59-64	44–53
5	0–77 78–100		78–100	0–77	-	-
6	-	41-100	-	37, 39–40	0-4, 6	33-34, 36, 38
7	0-64	71-100	65–68	69–70	69–73	63, 65–68
8	0–36	37-39, 41, 75-100	37	30-36, 40, 42-74	38-39, 41-44	10, 17–29
9	0-14, 18	66–100	15–17, 19–70, 72–73	47, 49–65	71, 74–79	27-46, 48
10	0–47	67-100	48-62	56-66	63–66	54–55
11	0-48, 52 53, 55-100		49–51, 53–65	47–52, 54	66	41–46
12	0-46, 48	47, 49–100	47, 49–54, 56	32, 34–46, 48	55, 57–58	29-31, 33
13	0-37, 39	38, 40–100	38, 40–72, 74	38, 40–72, 74 17–37, 39		10–16
14	0-21, 35	78-100	22-34, 36-67	68–77	68-71, 73-75	66–67
15	0–54	55-100	55–96, 98–99	21-22, 26-54	97, 100	14-17, 19-20, 23-25
16	0-55	65-100	56-61, 63	55, 57–64	62, 64–70	41, 45–54, 56
17	5, 7–36	37-100	0-4, 6, 37-74, 78	0-3, 5-36	75–77, 79–94	4
18	0-53	54-100	54-71	50-53	72–75, 77–80, 82	46-49
19	0-42	49, 52–100	43-69	22-48, 50-51	70-82	16, 19–21
20	0-14, 16-19, 21-24	20, 25–54, 56–57, 59,	15, 20, 25–35	21-24, 55, 58, 60	36-38, 40, 42-48, 52,	12–19
		61-100			60-67, 69-70	
21	0-45	46-100	46-81	21-45	82-91	14, 16–20
22	-	62, 64–100	0-18	33-61, 63	19–24	25–32
23	0–26	-	27–55 –		56–67	69-70, 73-100
24	0-20, 22-25	21, 26–100	21, 26–45 0–20, 22–25		46–56, 58 –	
25	0-80	81-100	81-88	34-80	89–98	1–33



L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms

TABLE 11. Complete results of the robustness of ArcFace against traditional morphing.

	P	os. 1	Pos	. 2–3	Pos.	4-5
1	0–23	24-100	24-31	7–23	32–37	4–6
2	0-17	46, 48–100	18-23, 25	37, 39-45, 47	24,26-31	32-36, 38
3	0–68	69–100	69–73	60-61, 63-68	74-80	56-59, 62
4	0-48	57-100	49-58, 60-61	37–56	59, 62–63	35-36
5	0–64	65-100	65–78 54–64		79–100	51-53
6	0–26	63-100	27-32	58-62	33-43, 45	56–57
7	0–50	66-100	51–56	57-65	57–58	
8	0-64, 66	69–100	65, 67	68	68–71	59–67
9	0–50	51-100	51–57, 59	39-40, 42-50	58, 60-65	31, 33-38,41
10	0-51	60-100	52–59	55–59	60-63	52-53
11	0–57	0–57 61–100		57-60	63–66	50-56
12	0–50	51-100	51–55	44–50	-	-
13	0–47	48-100	48–53, 55	38–47	54, 56-64, 66	28-29, 34-37
14	0-62	63-100	63–64	59-62	65	55, 57–58
15	0–54	55-100	55-60	51–54	61–64	45-50
16	0–56	57-100	57–59	54–56	60-68, 70	45-53
17	0-31	42, 44–100	32–39	38-41, 43	41	35, 37
18	0-43, 45	50-100	44, 46	47–49	47–52	39–46
19	0-41	53-100	42-46	49, 52	47-48, 50	47, 51
20	0-43	52-54, 56-100	44-45	46-51, 55	46-48	38–45
21	0–70	71-100	71–100	51-70	-	44, 47–50
22	0–25	58-100	26-35	52–57	36–37	49–50
23	0-44	62, 64–100	45-51 50-61, 63		52-66, 68	42-51
24	0–40	63-100	41–46 55–62		47–52	53–54
25	0–52	73–74, 76–100	53-61	53-72, 75	62–68	51–52

L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms



APPENDIX B DETAILED RESULTS OF THE PROPOSED METHOD

Table 12 presents the results of our proposed morphing attack. It shows all the original subjects' portraits and their respective images produced by our method that achieve misidentification with ArcFace and reduces the amount of morphing needed. Column *No.* reflects the image number (check Table 2), *It.* the number of iterations needed to get the misidentification and *Mor. conf.* the morphing detector confidence of the misidentified image.

We can see that some images (especially those that required a greater number of iterations) present some blurred areas. They are ghost artifacts, which are generated by mismatched parts of the images used in the morphing process. As the background of the photos in the LFW database is usually very different in each of them, in some cases some shadows can be observed due to the alpha blending.

L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms

TABLE 12. Complete results of the proposed method.

No.	Initial image	Final image	It.	Mor. conf.	No	Initial image	Final image	It.	Mor. conf.	No	. Initial image	Final image	It.	Mor. conf.
1	(Q)	(Q)	2	99.46	10			6	100	18			8	100
2	9	9	2	100	11			9	100	19	-	R	7	100
3	R		7	99.95	12			12	100	20			6	100
4			5	100	13			6	100	21		Ø	11	100
5	<u>S</u>	Đ.	8	100	14			8	100	22			3	100
6			1	100	15		R	7	100	23		. E	5	1.48
7			9	100	16			8	100	24	heard Jones	want jones	4	100
8	S		7	0.09	17			1	100	25	CRE	CRE CRE	7	100
9		0	6	100										

L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms

IEEEAccess

REFERENCES

- S. F. Kak, F. M. Mustafa, P. Valente, "A review of person recognition based on face model," *Eurasian Journal of Science & Engineeting*, vol. 4, issue 1, pp. 157–168, 2018. [Online]. Available doi: 10.23918/eajse.v4i1sip157.
- [2] E. P. Ijjina, G. Kanahasabai and A. S. Joshi, "Deep learning based approach to detect customer Age, gender and expression in surveillance video," in 2020 11th International Conference on Computing, Communication and Networking Technologies (ICCCNT), Kharagpur, India, 2020, pp. 1–6. [Online]. Available doi: 10.1109/ICCCNT49239.2020.9225459.
- [3] V. Mirchandani, 17 Best Smartphones with Face Unlock Yoy Can Buy Right Now, Beebom, March 2019. Accessed on: Nov. 18, 2020. [Online]. Available: https://beebom.com/10-best-smartphones-face-unlockbuy-right-now/.
- [4] "System, method, and computer program product for verifying the identity of social network users," by S. Ufford and J.Tanis (2012, Nov. 20). U.S. Patent No 8,316,086. Accessed 18 Nov. 2020. [Online]. Available: https://patents.google.com/patent/US8316086B2/en.
- [5] N. Erdogmus and S. Marcel, "Spoofing face recognition with 3D masks," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 7, pp. 1084–1097, July 2014. [Online]. Available doi: 10.1109/TIFS.2014.2322255.
- [6] "Anonymous," HKU University of the Arts Utrecht. Accessed on: Nov. 16, 2020. [Online]. Available: https://www.hku.nl/het-werk-vanhku/studentenwerk/anonymous-meerdere-makers.
- [7] M. Ferrara, A. Franco and D. Maltoni "The magic passport," in *IEEE International Joint Conference on Biometrics*, Clearwater, FL, 2014, pp. 1–7. [Online]. Available doi: 10.1109/BTAS.2014.6996240.
- [8] A. Agrawal and S. Samson, "A review on feature extraction techniques and general approach for face recognition," *International Journal of Computer Applications Technology and Research*, vol. 5, issue 3, pp. 156–158, 2016.
- [9] Y. Kortli, M. Jridi, A. Al Falou and M. Atri, "Face recognition systems: A survey," *Sensors*, vol. 20, no. 2:342, 2020. [Online]. Available doi: 10.3390/s20020342.
- [10] M. Wang and W. Deng, "Deep face recognition: A survey," 2018, arXiv:1804.06655. [Online]. Available: https://arxiv.org/abs/1804.06655.
- [11] M. A. Turk and A. P. Pentland, "Face recognition using eigenfaces," in Proceedings of Computer Vision and Pattern Recognition IEEE Computer Society, June 1991, pp. 586–591.
- [12] P. N. Belhumeur, J. P. Hespanha and D. J. Kriegman, "Eigenfaces vs fisherfaces: Recognition using class specific linear projection," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 711–720, 1997. [Online]. Available doi: 10.1109/34.598228.
- [13] M. S. Bartlett, J. R. Movellan and T. J. Sejnowski, "Face recognition by independent component analysis," *IEEE Trans. Neural Networks*, vol. 13, no. 6, pp. 1450–1464, Nov. 2002. [Online]. Available doi: 10.1109/TNN.2002.804287.
- [14] S. F. Hafez, M. M. Selim and H. H. Zayed, "2D face recognition system based on selected gabor filters and linear discriminant analysis LDA," *IJCSI International Journal of Computer Science Issues*, vol. 12, no. 1, pp. 33–41, 2015.
- [15] G. Zhang, X. Huang, S. Z. Li, Y. Wang and X. Wu, "Boosting local binary pattern (LBP)-based face recognition," in *Proc. of the 5th Chinese* conference on Advances in Biometric Person Authentication, 2004, pp. 179–186. [Online]. Available doi: 10.1007/978-3-540-30548-4_21.
- [16] M. Bicego, A. Lagorio, E. Grosso and M. Tistarelli, "On the Use of SIFT Features for Face Authentication," in 2006 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'06), New York, NY, USA, 2006, pp. 35–35. [Online]. Available doi: 10.1109/CVPRW.2006.149.
- [17] G. Du, F. Su and A. Cai, "Face recognition using SURF features," in *MIPPR 2009: Pattern Recognition and Computer Vision*, International Society of Optics and Photonics; SPIE, vol. 7496, 2009, pp. 749628.1– 749628.7. [Online] Available doi: 10.1117/12.832636.
- [18] A. A. Fathima, S. Ajitha, V. Vaidehi, M. Hemalatha, R. Karthigaiveni and R. Kumar, "Hybrid approach for face recognition combining gabor wavelet and linear discriminant analysis," in 2015 IEEE International Conference on Computer Graphics, Vision and Information Security (CGVIS), Bhubaneswar, 2015, pp. 220–225. [Online]. Available doi: 10.1109/CGVIS.2015.7449925.
- [19] Y. Taigman, M. Yang, M. Ranzato and L. Wolf, "Deepface: Closing the gap to human-level performance in face verification," in *PProceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 1701–1708. [Online]. Available doi: 10.1109/CVPR.2014.220.
- [20] Y. Sun, D. Liang, X. Wang and X. Tang, "Deepid3: Face recognition with very deep neural networks," *arXiv preprint arXiv:1502.00873*, 2015.

- [21] F. Schroff, D. Kalenichenko and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 815–823. [Online]. Available doi: 10.1109/CVPR.2015.7298682.
- [22] O. M. Parkhi, A. Vedaldi and A. Zisserman, "Deep face recognition" in Proc. British Machine Vision Conf., Jan. 2015, pp. 41.1–41.12.
- [23] J. Deng, J. Guo, N. Xue and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 4690– 4699. [Online]. Available doi: 10.1109/CVPR.2019.00482.
- [24] G. B. Huang, M. Ramesh, T. Berg and E. Learned-Miller, "Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments," University of Massachusetts, Technical Report 07-49, Oct. 2007.
- [25] A. K. Agrawal and Y. N. Singh, "Evaluation of face recognition methods in unconstrained environments," *Proceedia Computer Science*, vol. 48, pp. 644–751, 2015.
- [26] X. Zhang and Y. Gao "Face recognition across pose: A review," *Pattern Recognition*, vol. 42, no. 11, pp. 2876–2896, 2009. [Online]. Available doi: 10.1016/j.patcog.2009.04.017.
- [27] G. H Givens, J. R. Beveridge, P. J. Phillips, B. Draper, Y. M. Lui and D. Bolme, "Introduction to face recognition and evaluation of algorithm performance," *Computational Statistics and Data Analysis*, vol. 67, pp. 236–247, 2013. [Online]. Available doi: 10.1016/j.csda.2013.05.025.
- [28] D. B. Smythe, "A Two-Pass Mesh Warping Algorithm for Object Transformation and Image Interpolation," Technical Memo 1030, Industrial Light and Magic, Computer Graphics Department, Lucasfilm Ltd., 1990.
- [29] G. Wolberg, Digital Image Warping, IEEE Computer Society Press, Los Alamitos, California, 1990.
- [30] G. Wolberg, "Image morphing: A survey," *The Visual Computer*, vol. 14, no. 8–9, pp. 360–372, 1998.
- [31] M. Steyvers, "Morphing techniques for manipulating face images," *Behavior Research Methods*, vol. 31, no. 2, pp. 359–369, 1999. [Online]. Available doi: 10.3758/BF03207733.
- [32] U. Scherhag, C. Rathgeb, J. Merkle, R. Breithaupt and C. Busch, "Face recognition systems under morphing attacks: A survey," *IEEE Access*, vol. 7, pp. 23012–23026, 2019. [Online]. Available doi: 10.1109/AC-CESS.2019.2899367.
- [33] G. Bradski and A. Kaehler, *Learning OpenCV: Computer Vision with the OpenCV Library*. Sebastopol, CA, USA: O'Reilly Media, 2008.
- [34] N. Damer, A. M. Saladié, A. Braun and A. Kuijper, "MorGAN: Recognition Vulnerability and Attack Detectability of Face Morphing Attacks Created by Generative Adversarial Network," 2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS), Redondo Beach, CA, USA, 2018, pp. 1–10. [Online]. Available doi: 10.1109/BTAS.2018.8698563.
- [35] M. Gómez-Barrero, C. Rathgeb, U. Scherhag and C. Busch, "Is your biometric system robust to morphing attacks?," in 2017 5th International Workshop on Biometrics and Forensics (IWBF), Coventry, 2017, pp. 1–6. [Online]. Available doi: 10.1109/IWBF.2017.7935079.
- [36] S. Mallick, Face Morph Using OpenCV C++/Python, Learn OpenCV, March 11, 2016. Accessed on: January 21, 2021. [Online]. Available: https://learnopencv.com/face-morph-using-opencv-cpp-python/
- [37] A. Hadid, N. Evans, S. Marcel and J. Fierrez, "Biometrics systems under spoofing attack: An evaluation methodology and lessons learned," *IEEE Signal Processing Magazine*, vol. 32, no. 5, pp. 20–30, Sept. 2015. [Online]. Available doi: 10.1109/MSP.2015.2437652.
- [38] A. Mohammadi, S. Bhattacharjee and S. Marcel, "Deeply vulnerable: A study of the robustness of face recognition to presentation attacks," *Institution of Engineering and Technology Biometrics*, vol. 7, issue 1, pp. 15–26, 2018. [Online]. Available doi: 10.1049/iet-bmt.2017.0079.
- [39] M. Ferrara, A. Franco and D. Maltoni, "On the effects of image alterations on face recognition accuracy," in *Face Recognition Across the Image Spectrum*. Springer Nature, 2016, pp. 195–222.
- [40] R. Tolosana, R. Vera-Rodríguez, J. Fiérrez, A. Morales and J. Ortega-García, "DeepFakes and beyond: A survey of face manipulation and fake detection," arXiv preprint arXiv:2001.00179, 2020.
- [41] D. J. Robertson, R. S. S. Kramer and A. M. Burton, "Fraudulent ID using face morphs: Experiments on human and automatic recognition," *PLoS ONE*, vol. 12, no. 3:e0173319. [Online]. Available doi: doi:10.1371/journal.pone.0173319.
- [42] L. Wandzik, R. V. García, G. Kaeding and X. Chen, "CNNs under attack: On the vulnerability of deep neural networks based face recognitionto



image morphing," in Proc. 16th Int. Workshop on Digital Forensics and Watermarking, 2017, pp. 121–135.

- [43] U. Scherhag, R. Raghavendra, K.B. Raja, M. Gómez-Barrero, C. Rathgeb and C. Busch, "On the vulnerability of face recognition systems towards morphed face attacks," in 2017 5th International Workshop on Biometrics and Forensics (IWBF), Coventry, 2017, pp. 1–6. [Online]. Available doi: 10.1109/IWBF.2017.7935088.
- [44] P. Korshunov and T. Ebrahimi, "Using face morphing to protect privacy," in 2013 10th IEEE International Conference on Advanced Video and Signal Based Surveillance, Krakow, 2013, pp. 208–213. [Online]. Available doi: 10.1109/AVSS.2013.6636641.
- [45] U. Scherhag et al., "Biometric Systems under Morphing Attacks: Assessment of Morphing Techniques and Vulnerability Reporting," in 2017 International Conference of the Biometrics Special Interest Group (BIOSIG), Darmstadt, 2017, pp. 1–7. [Online]. Available doi: 10.23919/BIOSIG.2017.8053499.
- [46] M. Sharif, F. Naz, M. Yasmin, M. A. Shahid, and A. Rehman, "Face recognition: A survey," *Journal of Engineering Science and Technology Review*, vol. 10, no. 2, pp. 166–177, 2017.
- [47] S. Patel, *Face Morphing*, Github, 2018. Accessed on: Nov.29, 2020. [Online]. Available: https://github.com/cirbuk/face-morphing.
- [48] B. Delaunay "Sur la sphère vide. A la mémoire de Georges Vorono," Bulletin de l'Académie des Sciences de l'URSS. Classe des sciences mathématiques et naturelles vol. 6, p. 793, 1934.
- [49] D. E. King, "Dlib-ml: A machine learning toolkit," Journal of Machine Learning Research (JMLR), vol. 10, pp. 1755–1758, 2009.
- [50] T. Ahonen, A. Hadid and M. Pietikainen, "Face Description with Local Binary Patterns: Application to Face Recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 12, pp. 2037–2041, Dec. 2006. Available doi: 10.1109/TPAMI.2006.244.
- [51] R. Raja, Face Recognition with OpenCV and Python, Github, 2017. Accessed on: Nov. 29, 2020. [Online]. Available: https://github.com/informramiz/opencv-face-recognition-python.
- [52] V. Garg and K. Garg, "Face recognition using haar cascade classifier," *Journal of Emerging Technologies and Innovative Research (JETIR)*, vol. 3, no. 12, December 2016.
- [53] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, et al. "Scikit-learn: Machine learning in Python," the Journal of machine Learning research, vol. 12, pp. 2825–2830, 2011.
- [54] W. Noble, "What is a support vector machine?," *Nature Biotechnology*, vol. 24, pp. 1565–1567, 2006. [Online]. Available doi: 10.1038/nbt1206-1565.
- [55] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis *et al.* "Tensorflow: A system for large-scale machine learning," in *12th USENIX symposium on* operating systems design and implementation (OSDI 16), 2016, pp. 265– 283.
- [56] D. Sandberg, Face Recognition using Tensorflow, Github, 2016. Accessed on: Nov. 29, 2020. [Online]. Available: https://github.com/davidsandberg/facenet.
- [57] J. Deng and J. Guo, InsightFace: 2D and 3D Face Analysis Project, Github, 2018. Accessed on: Apr. 12, 2021. [Online]. Available: https://github.com/deepinsight/insightface.
- [58] Q. Cao, L. Shen, W. Xie, O. M. Parkhi and A. Zisserman, "VGGFace2: A dataset for recognising faces across pose and age," in 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018), Xi'an, 2018, pp. 67–74. [Online]. Available doi: 10.1109/FG.2018.00020.
- [59] C. Szegedy, S. Ioffe, and V. Vanhoucke, "Inception-v4, inception-resnet and the impact of residual connections on learning," arXiv preprint arXiv:1602.07261, 2016.
- [60] J. Galbally, S. Marcel and J. Fiérrez, "Biometric antispoofing methods: A survey in face recognition," *IEEE Access*, vol. 2, pp. 1530–1552, 2014. [Online]. Available doi: 10.1109/ACCESS.2014.2381273.
- [61] H. Chen, G. Hu, Z. Lei, Y. Chen, N. M. Robertson and S. Z. Li, "Attention-Based Two-Stream Convolutional Networks for Face Spoofing Detection," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 578–593, 2020. [Online]. Available doi: 10.1109/TIFS.2019.2922241.
- [62] M. De Marsico, M. Nappi, D. Riccio and J. Dugelay, "Moving face spoofing detection via 3D projective invariants," in *5th IAPR International Conference on Biometrics (ICB)*, 2012, pp. 73–78. [Online]. Available doi: 10.1109/ICB.2012.6199761.
- [63] W. Sun, Y. Song, H. Zhao and Z. Jin, "A Face Spoofing Detection Method Based on Domain Adaptation and Lossless Size Adaptation," *IEEE Access*, vol. 8, pp. 66553–66563, 2020. [Online]. Available doi: 10.1109/ACCESS.2020.2985453.

- [64] R. Raghavendra, K. B. Raja, and C. Busch. "Detecting Morphed Face Images," in 8th IEEE International Conference on Biometrics: Theory, Applications, and Systems (BTAS), pp. 1–8, 2016.
- [65] U. Scherhag, L. Debiasi, C. Rathgeb, C. Busch and A. Uhl, "Detection of face morphing attacks based on PRNU analysis," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 1, no. 4, pp. 302–317, Oct. 2019. [Online]. Available doi: 10.1109/TBIOM.2019.2942395.
- [66] C. Seibold, W. Samek, A. Hilsmann, and P. Eisert, "Detection of face morphing attacks by deep learning," in *Digital Forensics Watermarking*. Cham, Switzerland: Springer, 2017, pp. 107–120.
- [67] L. Zhang, F. Peng and M. Long, "Face morphing detection using Fourier spectrum of sensor pattern noise," in 2018 IEEE International Conference on Multimedia and Expo (ICME), San Diego, CA, 2018, pp. 1–6. [Online]. Available doi: 10.1109/ICME.2018.8486607.
- [68] R. Raghavendra, K. Raja, S. Venkatesh and C. Busch, "Face morphing versus face averaging: Vulnerability and detection," in 2017 IEEE International Joint Conference on Biometrics (IJCB), Denver, CO, 2017, pp. 555–563. [Online]. Available doi: 10.1109/BTAS.2017.8272742.
- [69] L. Spreeuwers, M. Schils and R. Veldhuis, "Towards robust evaluation of face morphing detection," in 2018 26th European Signal Processing Conference (EUSIPCO), Rome, 2018, pp. 1027–1031. [Online]. Available doi: 10.23919/EUSIPCO.2018.8553018.
- [70] W. Deng, J. Hu, N. Zhang, B. Chen and J. Guo, "Fine-grained face verification: FGLFW database, baselines, and human-DCMN partnership," *Pattern Recognition*, vol. 66, pp.63–73, 2017.
- [71] C. R, Harris, K. J. Millman, S. J. van der Walt *et al.* "Array programming with NumPy," *Nature*, vol. 585, num. 7825, pp. 357–362, 2020. [Online]. Available doi: 10.1038/s41586-020-2649-2.
- [72] E. Learned-Miller, G. Huang, A. RoyChowdhury, H. Li, Haoxiang and G. Hua., "Labeled Faces in the Wild: A Survey," in *Advances in Face Detection and Facial Image Analysis*, 2016, pp. 189–248. [Online]. Available doi: 10.1007/978-3-319-25958-1_8.



EDUARDO CERMEÑO (M'2012–SM'2019) obtained the degrees of Superior Engineer in computer science, Master in computer science and artificial intelligence and PhD in computer science and telecommunications from the Autonomous University of Madrid (UAM).

He currently serves as General and Research Director of the company Vaelsys in Madrid and Associate Professor at the Autonomous University of Madrid in the department of Computer Engi-

neering. He has directed several projects and research activities related to video surveillance that have obtained public funding and received recognition of excellence by the European Commission. He is co-author of several patents and articles in high impact journals. Eduardo has participated as a speaker in different international conferences related to crime prevention technologies, evaluation of video analysis systems for video surveillance, as well as forensic and biometric techniques. Eduardo also regularly collaborates as a reviewer in several international scientific journals in the area of video solutions. He is a member of the European Association of Artificial Intelligence and Senior Member of the IEEE.

L. Cárabe et al.: Stegano-morphing: Concealing attacks on face identification algorithms





LUIS CÁRABE studied a degree in Mathematics and a degree in Computer Science at the Autonomous University of Madrid (UAM).

He has been researching in some mathematical fields, such as differential privacy. Interested in artificial intelligence and computer vision, he has collaborated with Vaelsys on public funding projects involving face recognition systems and morphing attacks since 2019.

...