Face unmorphing

Kevin Witlox University of Twente P.O. Box 217, 7500AE Enschede The Netherlands k.h.d.witlox@student.utwente.nl

ABSTRACT

After the recent introduction of Automated Border Control systems that rely on face recognition as a biometric method, a vulnerability has come to light that poses a security thread to international travel. Both human inspectors and face recognition software can be fooled by morphed images of two subjects, making it possible for fugitives to evade detection by hiding in an innocent face. Morph attack detection methods based on morphing artifacts have not provided satisfying results thus, in this research, a novel detection method is proposed based on subtracting a photo ID from a live capture reference and using the result to train a classifier.

Keywords

Face Morph, Morphing Attack, Automatic Face Recognition, Morph Detection

1. INTRODUCTION

Since the recent introduction of neural networks, face recognition is able to recognize people in many uncontrolled environments [24] [22]. Currently, the technique is being deployed as a biometric method for Automatic Border Control (ABC) e-gates [9]. A passenger presents his or her electronic Machine Readable Travel Document (eMRTD), e.g. a passport, after which the photo on the document is verified against a live capture of the passenger. These systems can efficiently verify the identity of all passing customers with minimal human intervention.

Recently, researchers have brought to light a vulnerability in these ABC systems [7]. The attack revolves around face morphing. A face morph is a picture in which two or more subjects can be recognized (see figure 1). As shown in [11], the attack poses a serious threat to the systems. The paper shows that morphing cannot only fool automated face recognition (AFR), but also human inspectors.

The motivation behind the attack is as follows [7]. A criminal is planning to escape a country. However, as a registered fugitive, the criminal would be caught by the Automatic Border Control system at an airport if he were to present his passport. So, the criminal requests help from his innocent colleague, the "accomplice". The accomplice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

31th Twente Student Conference on IT July. 5nd, 2019, Enschede, The Netherlands.

Copyright 2019, University of Twente, Faculty of Electrical Engineering, Mathematics and Computer Science. requests an eMRTD and provides his photo ID. Before submission this photo ID was morphed with the face of the criminal. The criminal can now use the passport of his accomplice to travel as now the photo ID on the passport resembles both the accomplice and himself.

To combat this exploitation, various techniques have been developed to detect these attacks. However, no single technique has been able to provide satisfying results. In this research, the viability of a novel approach for morph attack detection will evaluated.

2. RESEARCH QUESTIONS

- **RQ1** To what extent can face unmorphing be used to detect morph attacks?
 - RQ1.1 How should differences in lighting be accounted for while unmorphing?
 - RQ1.2 What classifier should be used on the unmorphing results to determine whether the image is a morph or not?

3. BACKGROUND

3.1 Morphing

In order to understand the approach taken in this paper for detecting morphing pictures, it is important to understand the concept of morphing itself. In the current context, the goal of face morphing is to resemble the two input pictures both in texture and in geometric features. Face morphing algorithms can take on various shapes and forms, but in general the process in this context can be divided into three steps [21], namely 1) finding correspondance, 2) warping, and 3) blending. It should be noted that face morphing does not need to produce a fifty-fifty morph between the two faces of the criminal and the accomplice. The resulting image may be set to resemble the criminal for 30% and the accomplice for 70%. This



Figure 1. Example of a manually retouched morph (middle). The (left) and (right) image are the sources for the morph.



Figure 2. Landmark detection (left) and Delaunay triangulation (right).

variance is determined by the warping factor α_w and the blending factor α_b .

3.1.1 Correspondence

Because faces can have various shapes and proportions, one needs references on both images in order to create a successful transformation. The morph should have a geometric structure that resembles that of both input images. Simply aligning and cross-dissolving the faces will produce poor results [1]. By nature of the problem at hand, all pictures contain full frontal faces with a neutral expression. For this application, the problem of automatically detecting facial features is solved [2]. The morphing process used in this project will use an automatic model-based landmark detection algorithm, as implemented in the **OpenCV** library [13].

3.1.2 Warping

The next step is to align the facial landmarks of both images, taking into account the warping factor α_w . In this project, Delaunay triangulation is used, as used in most state-of-the-art morphing algorithms, e.g. [11] [10] [18] [20]. The Delaunay triangulation process determines non overlapping triangles as depicted in figure 2. As both images have the same facial landmarks, both images have corresponding triangles. The triangles of the images are then distorted, rotated and shifted to align with with each other [21]. The warping factor determines the ratio in which the resulting landmarks resemble those of the accomplice, as shown in figure 3.

3.1.3 Blending

In the final step, the texture of both images should be blended. Now that both images are aligned, a simple linear blending can be applied. Here, the blending factor α_b determines the weight of the images.

4. RELATED WORK

Ever since the threat that is the face morphing attack came to light in 2014 [7], researchers have been searching for detection methods. There are two distinct branches in morphing attack detection. Initially, research focused on no-reference morph attack detection. This form is more general as it does not rely on the assumption that a reference image is present. A common approach in these works is to feed general purpose image descriptors into classifiers, in order to distinguish images on whether they are morphed or not. The second branch does rely on the assumption of a bona fide reference image. This situation opens up new possibilities as the photo in question can be compared with an unmodified reference of the person in question.



Figure 3. The morphing process shown with varying morphing factors. The image with factor 0.0 shows the "accomplice" and the image with factor 1.0 shows the "criminal"

4.1 No-reference morph attack detection

One of the first works that attempted to tackle the issue is [15]. In this work, Binarized Statistical Image Features (BSIF) are used to train a linear Support Vector Machine (SVM). Following this paper, various similar methods based on texture descriptors [19] [14], image signals [6] [5], and JPEG compression have been proposed [11] [12]. Most of these papers have in common that they were only tested on digital images.

Various concerns have been raised concerning the reliability of solutions based on these no-reference detection methods. A collection of methods that relies on anomalies in image descriptors was evaluated in [23] and [20]. It was shown that the results of these methods significantly deteriorate when testing on a dataset that is different from the training dataset. The results of these works indicate that this type of approach easily overfits to the training sets.

Furthermore, these methods rely on artifacts and anomalies left over by the morphing process. Because of this, the quality of the morph has significant impact on the detection rate [17]. Relying on these artifacts and anomalies is even more problematic when the presented images are printed and scanned before submission. This process of printing and scanning occurs in real-life situations where a citizen is permitted to provide his or her ID photo on paper [7]. The print/scan process has been proven to dramatically reduce the accuracy of methods within this branch of detection [18].

4.2 Differential morph attack detection

Recently, some newer works have explored reference-based detection methods. This research has shown that introducing a bone fide reference image allows for a whole new set of techniques.

The first work to utilize a reference image is [16]. In this work, the angles and distances between facial landmarks of the passport and bona fide image are compared. The angle comparison delivers the best results, but the classification error rates are not yet small enough for real-world use. Thus in future work the technique will be combined with a texture based technique.

A similar approach based on facial landmarks is used in [4]. In this paper, the relative directed distances of facial marks are fed into a Support Vector Machine (SVM). This method outperforms two existing no-reference morph attack detection methods.

4.2.1 Face Demorphing

A promising method was presented in [8] and will be elaborated on here as the method shows some similarities with method proposed in this paper. In [8], a new method is presented in which the reference image is used to demorph the input image. The idea is that demorphing (reverting the morphing process) should reveal a hidden second face if the input image was morphed. When a subject presents its passport, the AFR checks whether the passport sufficiently matches the subject. If so, the demorphing module is applied and the output is tested again against the AFR to detect whether the image was morphed. An overview of this process is depicted in figure 5.

In order to describe the demorphing algorithm, a description of the morphing process itself is needed. The morphing process, as described in section 3.1, can be viewed as a fluid transformation from one image to another [8]. The transformation is guided by the warping factor a_w and the blending factor a_b . For the sake of simplicity we will treat them as equal and denote their value with a, the morphing factor. Given two images I_0 and I_1 , the process generates a set of intermediate frames $\mathbb{M} = \{I_a, a \in \mathbb{R}, 0 < a < 1\}$.

In order to calculate the frame I_a , the warping and blending steps as described in section 3.1 need to be performed. The facial landmarks of I_0 and I_1 are assumed to be given. The process can be formally denoted as [8]:

$$I_{a}(\boldsymbol{p}) = (1-a) \cdot I_{0}(w_{P_{a} \longrightarrow P_{0}}(\boldsymbol{p})) + a \cdot I_{1}(w_{P_{a} \longrightarrow P_{1}}(\boldsymbol{p})) \quad (1)$$

where:

- **p** is a generic pixel position;
- *a* is the morphing factor;
- P₀ and P₁ are the two sets of facial landmarks in I₀ and I₁, respectively;
- P_a is the set of facial landmarks aligned according the the morphing factor a; $w_{B \longrightarrow A(p)}$ is a warping function.

The function states that the pixel value at a given location in the morphed image is equal to the blend of the pixel values at the same location in the warped frames I_0 and I_1 , where *a* determines the blending factor. The warped frames are calculated with the warping function, formally defined as:

$$P_a = \{ \boldsymbol{r} | \boldsymbol{r} = (1-a) \cdot \boldsymbol{u}_i + a \cdot \boldsymbol{v}_i, \boldsymbol{u}_i \in P_0, \boldsymbol{v}_i \in P_1 \} \quad (2)$$

where:

- **u** and **v** denote the positions of a landmark in P₀ and P₁ respectively, and
- r denotes the resulting position of the landmark in P_a .

Now, for the demorphing process, inverting and resolving equations 1 and 2 gives us:

$$I_0(\boldsymbol{q}) = \frac{I_a(w_{P_0 \longrightarrow P_a(\boldsymbol{q})}) - a \cdot I_a(w_{P_0 \longrightarrow P_1(\boldsymbol{q})})}{(1-a)} \qquad (3)$$

where:

$$P_0 = \left\{ \boldsymbol{u}_i | \boldsymbol{u}_i = \frac{\boldsymbol{r}_i - a \cdot \boldsymbol{v}_i}{1 - a}, \boldsymbol{r}_i \in P_a, \boldsymbol{v}_i \in P_1 \right\}$$
(4)

One of the problems of this approach is that in real-life scenarios the morphed images have most likely been manually retouched to hide morphing artifacts. The difference between an exact morph and a manually retouched morph can clearly be observed in (a) and (b) of figure 4. This retouching process results in loss of information about the



Figure 4. (a) shows an exact morph, (b) shows a manually retouched morph, and (c) shows a demorphed variant of b

second hidden face. Thus, the demorphing process can not produce perfect results and will show artifacts as shown in (c) of figure 4.

Another challenging aspect for differential morph detection techniques in general lies within the manner in which morphs are created. When creating a morphed image the factor α plays an important role. The factor α determines to what extends a morph resembles the criminal as opposed to the accomplice [21]. A low α makes the image resemble the accomplice more whereas a high α makes the image resemble the criminal more. In [8] it is suggested that an α of [0.2, 0.3] is the best trade-off between a morph's ability to fool a human officer as well as face recognition software. The issue with this factor is that to the defender the α is unknown.

For the face demorphing technique, this issue is particularly problematic. The paper focuses on picking a single demorphing factor $\tilde{\alpha}$ that performs best overall. However, this demorphing factor $\tilde{\alpha}$ is not the best factor for every possible morphing factor α . Furthermore, higher morphing factors deteriorate the results of the demorping process causing lower accuracy in detecting both morphs and bona fide images.

5. METHOD

The goal of this paper is to optimize and evaluate a novel morphing attack detection method based on subtraction, called "unmorphing". First, section 5.1 will provide details on the general approach of the proposed method. Next, in section 5.2, some experiments will be layed out with the purpose of finding the optimal implementation of the unmorphing method. Lastly, the dataset used will briefly be discribed.

5.1 Approach

The proposed method is applied in the same setting as the demorphing method as described in 4.2.1. However, instead of the demorphing module (shown in figure 5), an "unmorphing module" is implemented. Unmorphing is based on the subtraction of images. Given two matrices A and B representing two different gray-scale images, where $x_{ij} \in A$ corresponds to the pixel value ([0, 255]) at location (i, j) in image A, the resulting image can be calculated by A - B.

Now consider the morphing process. Assuming the morphing factor to be a, the morphing process can be described with the equation M = (1 - a)A + aC, where M represents the morph, A the accomplice (or an innocent person, if a = 0), and C the criminal. This is of course an oversimplification, as this equation does not take the geometric alignment of the faces into account.



Figure 5. Functional schema of the face demorphing procedure at ABC gates.

At an e-gate, images M and R are given. Here, M denotes the possibly morphed passport image and reference R denotes the live-captured image of the subject that has presented the passport. For the sake of simplicity, we will assume that two images of the same subject are exactly equal. As M has been determined by the face recognition software to resemble R, there are two possible situations. Either, the passport image M was not morphed and thus the person in R is the innocent person in image A, or the passport image M was morphed and the person in R is the criminal in image C.

To determine whether the passport image M is morphed, we take the reference image R and subtract it from M.

$$M - R = (1 - a)A + aC - R$$
(5)

where R is either A or C, given the assumption. If an unmodified passport image M is presented at an e-gate, subtracting the live capture R from the passport image will result in M - R = (1 - 0)A + 0C - A = 0. So, the subtraction should resolve to zero if the passport photo is unmodified. Now, consider the case in which the passport image M was morphed with the factor a. In this case, Ris actually equal to criminal C instead of the accomplice A. Subtracting R from M gives:

$$M - R = (1 - a)A + aC - C$$

= (1 - a)A + (-1 + a)C

In this case, we expect to see face-like features in both the positive and negative number-space of the resulting matrix.

5.1.1 Alignment As stated in section 5.1, the geometric alignment of the faces has not yet been taken into account. Intuitively, we would want to align two faces before subtracting them, in order to detect relevant differences between them. If for example the mouths do not line up, the difference would show noise in the areas where both mouths do not overlap, obscuring the result. Lining up the faces based on the position of the eyes is not enough to filter out this noise. Small differences in the angle at which the subject was photographed may result in notable differences. Thus, without alignment, subtracting two pictures of the same person will produce lots of noise, which will likely throw off the classifier.

To solve this problem, we warp the facial landmarks of the passport image M onto the facial landmarks of the reference image R. This process was implemented using **OpenCV** for landmark detection and **scipy** for Delaunay triangulation.

Furthermore, we should investigate whether this affects the results of the subtraction. To illustrate the problem formally, the following equation shows the unmorphing process (denoted by U) of a potentially morphed photograph without warping.

$$U(\mathbf{p}) = M(\mathbf{p}) - R(\mathbf{p})$$

$$U(\mathbf{p}) = (1 - a) \cdot A(w_{P_a \longrightarrow P_A}(\mathbf{p}))$$

$$+ a \cdot C(w_{P_a \longrightarrow P_C}(\mathbf{p})) - R(\mathbf{p})$$
(6)

Warping the passport image M to align with R gives:

$$U(\boldsymbol{p}) = M(w_{P_R \longrightarrow P_a}(\boldsymbol{p})) - R(\boldsymbol{p})$$
$$U(\boldsymbol{p}) = (1 - a) \cdot A(w_{P_R \longrightarrow P_A}(\boldsymbol{p}))$$
$$+ a \cdot C(w_{P_R \longrightarrow P_C}(\boldsymbol{p})) - \lambda R(\boldsymbol{p})$$

Now suppose that image M has not been morphed. Then a = 0 and, given that facial landmarks do not change between different images of the same subject, $w_{P_R \longrightarrow P_A}(p) = p$.

$$U(\boldsymbol{p}) = A(\boldsymbol{p}) - R(\boldsymbol{p})$$

With the assumption that R = A, we get:

$$U(\boldsymbol{p}) = 0$$

The same approach can be used to demonstrate that if M is morphed with a morphing factor a, the unmorphing process yields:

$$U(\mathbf{p}) = (1-a) \cdot A(w_{P_R \longrightarrow P_A}(\mathbf{p})) + (a-1) \cdot C(\mathbf{p})$$

The last equation shows that the remaining part of A is warped, which is to be expected. However, the conclusion that the subtraction will either produce nothing or facelike patterns, still holds.

5.2 Experiments

5.2.1 Lighting

In equation 5, we assumed that the reference image R is exactly equal to either A or C. This is clearly not the case, as the images A, C and R are captured at different moments in time. Even if all images fulfill ISO/ICAO specifications [9], differences is lighting are to be expected. This is detrimental for the detection accuracy, because if an unmorphed image produces too much noise, the classifier will not be able do distinguish the two classes. Therefore the differences in lighting need to be accounted for. In this section, possible solutions are investigated to counteract the differences in lighting. The first solution is to introduce a scaling factor into the subtraction process (denoted by λ in figure 6). The purpose of the factor is to fine-tune the subtraction process by giving more or less weight to the reference image R. As an example, if the live capture image (image R) would be overall brighter than the passport image M, then the scaling factor should be less than 1. The factor is found by minimizing the average intensity of the result of the subtraction.

However, because of natural shadowing on a face, some parts of the face will be darker than others. These shadows are of course influence by the direction of the light. To correct for the differences in lighting for certain areas of the image, the image is divided into a grid of $n = 2^x, x \in \mathbb{N}$ squares. For each square, the subtraction and minimization of the factor is done locally. The resulting optimized squares are then stitched back together for the final result. The question that remains is for which n, the classifier performance best.

The second solution is to apply histogram matching. Histogram matching can be used to normalize the global illumination of two images [3]. For each number of squares n, the subtraction will be executed both with and without histogram matching.

5.2.2 Classifier

Given the expectation that subtracting the live capture image from the passport image will produce either nothing or face-like patterns, the next step is to train a classifier. In this paper, two different classifiers will be tested, namely a linear Support Vector Machine (SVM) classifier and a linear discriminant analysis (LDA) classifier. Each classifier is trained on the training set, after which its accuracy is measured using the testing set.

When feeding the classifiers with training data, each pixel in an image is regarding as a feature. Because of this, the feature set quickly grows enormous. Because of the limited training data, it is important to find the right feature set size and thus another experiment will be performed. For both classifiers, the training set will be resized to varying dimension (d, d), where $d \in [25, 250], d/25 \in \mathbb{N}$

5.3 Dataset

This paper makes use of a subset of the Face Recognition Grand Challenge (FRGC) database. The database is divided into two sets, a training set consisting of 298 images and a testing set consisting of 266 images. Each image is a clear frontal picture of a face. For each image in each set, there is also a complementary picture of the same person, but taken at a different time with often slightly different clothing, hairstyle, facial accessories, facial hair and lighting. This second set of images is used to simulate the real-life scenario at an e-gate. The subject presents its passport and a photo is taken. These two photos will also differ slightly in the aforementioned aspects.

6. **RESULTS**

6.1 Lighting (RQ1.1)

Figure 7 shows the accuracy scores of the SVM classifier and the LDA classifier. On the x-axis, the amount of section in the grid is displayed. For example, with two sections, the subtraction phase was split in half vertically, so that both the left and the right side were subtracted with a different factor. In the case of #sections = 1, the subtraction factor was calculated over the entire image. Figure 8 shows the result of the subtraction process of the same image for different #sections.



Figure 6. An overview of the unmorphing method

For each classifier in 7 an alternative line is shown where histogram matching was applied before the subtraction, in an attempt to normalize the images. However, for both classifiers, applying histogram matching produces worse results.

As for the "section" based solution for solving the lighting issue, figure 7 shows that the optimal amount of sections for LDA without histogram matching is 2 and for SVM without histogram matching it is 4. More importantly so, it shows that locally computing the subtraction factor does increase the performance of the classifiers.

6.2 Classifier (RQ1.2)

The dimensions of the images fed into the classifiers have impact on the accuracy and speed of the classifier. The results of the experiment proposed in section 5.2.2 are shown in figure 9. The graph shows that the accuracy stabilizes for dimensions larger than 125x125. The results in the previous section 6.1 are based on the dimension 100x100, a trade-off between computation time and accuracy.

6.3 Discussion (RQ1)

From the graphs it can be concluded that the linear discriminant analysis classifier outperforms the linear support vector machine classifier. The best accuracy was achieved with an image dimensions size of 100x100 and #sections = 4, resulting in an accuracy of 75.1%. This accuracy score is far from perfect and would not be useful in real-life applications. However, it does show that this approach might be viable if further optimized. Moreover it proves that the concept might be worth further research.

The advantage of this subtraction-based method is that we do not need to know the exact morphing factor a. In a real-world scenario, the morphing factor a will be somewhere between [0.2; 0.3], so that the morph has the poten-



Figure 7. The performance of the LDA and SVM classifiers with and without histogram matching for multiple values of #sections



Figure 8. The subtraction of an unmorphed passport photo and a live capture for #sections = 1(left), #sections = 2 (middle), and #sections = 4(right).

tial to fool both a human officer and the face recognition software [8]. As long as the morphing factor a has a value significantly above 0, we expect to see some pattern in the result of the subtraction process.

The main disadvantage and the reason for the sub-optimal results is that the task of subtracting images is non-trivial. The method assumes that subtracting two similar images of the same subject will result in noise. This is however, clearly not the case, as shown in figure 8.

7. CONCLUSION

The results have shown that the "unmorphing" method for morphing attack detection based on subtraction has potential. The main obsticle is to find a subtraction method that produces close to pure nothing when two images of the same subject are subtracted. Further research should look into finding the most important areas of the images



Figure 9.

for training the classifiers and remove noise introduced by some unpredictable components of the face. Furthermore, research could look into finding better ways to normalize the input images to reduce the influence of differences in lighting.

8. REFERENCES

- B. G. Bhatt. Comparative study of triangulation based and feature based image morphing. Signal & Image Processing : An International Journal, 2(4):235-243, 2011.
- [2] O. Çeliktutan, S. Ulukaya, and B. Sankur. A comparative study of face landmarking techniques. *EURASIP Journal on Image and Video Processing*, 2013(1), 2013.
- [3] D. Coltuc, P. Bolon, and J.-M. Chassery. Exact histogram specification. *IEEE TRANSACTIONS* ON IMAGE PROCESSING, 15(5):1143âĂŞ1152, Apr 2006.
- [4] N. Damer, V. Boller, Y. Wainakh, F. Boutros, P. Terhörst, A. Braun, and A. Kuijper. Detecting face morphing attacks by analyzing the directed distances of facial landmarks shifts. *Lecture Notes in Computer Science Pattern Recognition*, pages 518–534, Jan. 2019.
- [5] L. Debiasi, C. Rathgeb, U. Scherhag, A. Uhl, and C. Busch. Prnu variance analysis for morphed face image detection. 2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS), 2018.
- [6] L. Debiasi, U. Scherhag, C. Rathgeb, A. Uhl, and C. Busch. Prnu-based detection of morphed face images. 2018 International Workshop on Biometrics and Forensics (IWBF), 2018.
- [7] M. Ferrara, A. Franco, and D. Maltoni. The magic passport. *IEEE International Joint Conference on Biometrics*, 2014.
- [8] M. Ferrara, A. Franco, and D. Maltoni. Face demorphing. *IEEE Transactions on Information Forensics and Security*, 13(4):1008–1017, Apr. 2018.
- [9] Frontex. Best practice technical guidelines for automated border control (abc) systems.
- [10] M. Hildebrandt, T. Neubert, A. Makrushin, and J. Dittmann. Benchmarking face morphing forgery detection: Application of stirtrace for impact simulation of different processing steps. 2017 5th International Workshop on Biometrics and Forensics (IWBF), Apr. 2017.
- [11] A. Makrushin, T. Neubert, and J. Dittmann. Automatic generation and detection of visually faultless facial morphs. Proceedings of the 12th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, 2017.
- [12] T. Neubert. Face morphing detection: An approach based on image degradation analysis. *Digital Forensics and Watermarking Lecture Notes in Computer Science*, pages 93–106, 2017.
- [13] OpenCV. Facemark class reference.
- [14] R. Raghavendra, K. Raja, S. Venkatesh, and C. Busch. Face morphing versus face averaging: Vulnerability and detection. 2017 IEEE International Joint Conference on Biometrics (IJCB), 2017.
- [15] R. Raghavendra, K. B. Raja, and C. Busch. Detecting morphed face images. 2016 IEEE 8th

International Conference on Biometrics Theory, Applications and Systems (BTAS), 2016.

- [16] U. Scherhag, D. Budhrani, M. Gomez-Barrero, and C. Busch. Detecting morphed face images using facial landmarks. *International Conference on Image* and Signal Processing, pages 444–452, June 2018.
- [17] U. Scherhag, A. Nautsch, C. Rathgeb, M. Gomez-Barrero, R. N. J. Veldhuis, L. Spreeuwers, M. Schils, D. Maltoni, P. Grother, S. Marcel, and et al. Biometric systems under morphing attacks: Assessment of morphing techniques and vulnerability reporting. 2017 International Conference of the Biometrics Special Interest Group (BIOSIG), 2017.
- [18] U. Scherhag, R. Raghavendra, K. B. Raja, M. Gomez-Barrero, C. Rathgeb, and C. Busch. On the vulnerability of face recognition systems towards morphed face attacks. 2017 5th International Workshop on Biometrics and Forensics (IWBF), 2017.
- [19] U. Scherhag, C. Rathgeb, and C. Busch. Detection of morphed faces from single images: a multi-algorithm fusion approach. In Proceedings of the 2018 2nd International Conference on Biometric Engineering and Applications - ICBEA 18, 2018.
- [20] U. Scherhag, C. Rathgeb, and C. Busch. Performance variation of morphed face image detection algorithms across different datasets. 2018 International Workshop on Biometrics and Forensics (IWBF), 2018.
- [21] U. Scherhag, C. Rathgeb, J. Merkle, R. Breithaupt, and C. Busch. Face recognition systems under morphing attacks: A survey. *IEEE Access*, 7:23012âĂŞ23026, Feb. 2019.
- [22] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.
- [23] L. Spreeuwers, R. Veldhuis, and M. Schils. Towards robust evaluation of face morphing detection. In Proceedings of the 26th European Signal Processing Conference (EUSIPCO), pages 1027–1031, 9 2018.
- [24] W.-Y. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld. Face recognition: A literature survey. *ACM Comput. Surv.*, 35:399–458, 12 2003.