Gender Obfuscation through Face Morphing

Shunxin Wang
Faculty of Electrical Engineering, Mathematics and Computer Science
University of Twente
Enschede, the Netherlands
Email: s.wang-2@utwente.nl

Una M. Kelly
Faculty of Electrical Engineering, Mathematics and Computer Science
University of Twente
Enschede, the Netherlands
Email: u.m.kelly@utwente.nl

Raymond N.J. Veldhuis
Faculty of Electrical Engineering, Mathematics and Computer Science
University of Twente
Enschede, the Netherlands
Email: r.n.j.veldhuis@utwente.nl

2021 IEEE International Workshop on Biometrics and Forensics (IWBF) | 978-1-7281-9556-8/20/$31.00 ©2021 IEEE | DOI: 10.1109/IWBF50991.2021.9465088

***Abstract*—With the advancement of machine learning, facial biometric data has been widely adopted for person recognition. Soft biometrics such as gender, age and ethnicity can be extracted automatically from the facial photographs without permission. This raises privacy concerns since such auxiliary information might be utilized improperly. In this work, we apply face morphing to obfuscate gender information in face images, so that gender classifiers can no longer accurately predict gender, but the resulting images can still be used for identity verification. We further explore the reversibility of our approach and the results show that gender obfuscated through face morphing cannot be recovered or retrieved easily. Our approach is especially useful for an identity verification system which is sensitive to morphing attacks.**

***Index Terms*—Image forensics, Gender obfuscation, Facial attribute manipulation**

I. Introduction

Physical and behavioral human characteristics such as fingerprint, face, iris, retina, palm-print and gait, can be used for identity verification, which can then control access to systems, devices or data [1]. Such characteristics are defined as biometrics. A typical face verification system is shown in Fig. 1. The input of the system is a face image while the output is the decision made by the system. The features extracted from the input image are compared to the templates extracted from the reference database. Through the decision module, it can either accept or reject the validity of the claimed identity. However, a face also contains some ancillary information such as gender, age and ethnicity. Such characteristics that carry some information but lack the capability to distinguish between two different persons, are defined as soft biometrics [2].

For a typical face verification system (Fig. 1), attributes such as gender, age and ethnicity can be automatically detected from the face images stored in the reference database by different facial attribute classifiers. This raises legitimate concerns about users' privacy if biometric data stored in the database were leaked, since users may not have agreed to share information other than identity. Therefore, protection of biometric data becomes critical for several reasons:

- Legitimate concern due to extracting other information without permission from users.
- Leakage of biometric data might cause identity theft.
- The extracted soft biometrics might be misused for the profiling of users.

According to General Data Protection Regulation (GDPR), it is essential to ensure the stored biometric data cannot be used for other purposes that are beyond the users' expectation [3]. For example, for the face verification system (Fig.1), the stored biometric data should only be utilized for identity verification, while it should not be possible to deduce other ancillary information..

In this paper, face morphing is applied to face images in such a way the images can still be used for identity verification while the ancillary information is suppressed. Face morphing mixes face images through image warping and cross-dissolving. It was shown in [4] that the morphed face can be successfully used to verify the identity of the two contributing faces. Therefore, the idea is that morphing with an average face of the opposite gender would help to confuse a gender classifier, but the identity of the obfuscated face remains. Specifically, gender obfuscation ensures that a gender classifier is not able to distinguish between male and female faces.

To achieve this, we morph a face (Non-obfuscated face) image with an average face of the opposite gender (Gender obfuscator), resulting in a new face (Obfuscated face) image that can confuse an arbitrary gender classifier. We investigate the utility and reversibility of face morphing in gender obfuscation. First, we determine roughly the morphing parameters suitable for gender obfuscation by comparing the degree of gender obfuscation and identity preservation. Second, we explore the influence of average faces on the performance of obfuscating gender from several aspects, such as the number of faces needed to create an average face, gender characteristics of the average face, similarity between the average face and the non-obfuscated face. Third, we investigate two possible gender retrieval approaches, which are face demorphing and gender classifier retraining. We demonstrate that face morphing is useful for obfuscating gender in face images while the gender information suppressed is hard to recover or retrieve.

978-1-7281-9556-8/21/$31.00 ©2021 IEEE

Authorized licensed use limited to: UNIVERSITY OF TWENTE.. Downloaded on September 13,2021 at 13:57:06 UTC from IEEE Xplore. Restrictions apply.

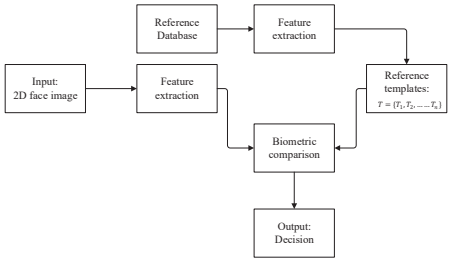

Fig. 1. A typical face verification system

II. Related work

A. Morphing

Morphing or metamorphosis, can be considered as a combination of multiple images through image warping and cross-dissolving [5]. Among those morphing techniques, feature-based approaches are used most frequently [6]–[8]. Early morphing approaches rely on human assistance to locate feature points while modern morphing techniques are able to create morphs automatically with the advancement in feature detection.

Traditional face morphing derives a morphed face from only two face images. Extending from it, polymorph, introduced in [5], create a morphing framework to morph multiple faces uniformly or non-uniformly.

Instead of morphing faces in image level, Generative Adversarial Networks (GANs) achieve face morphing in latent space and the resulted images are more realistic [9]. However, though the resulted images have better quality, training such a model can be tricky.

B. Gender Conversion and Suppression

Gender information in face images can be flipped or suppressed on image level or in feature space. A gender-conversion approach proposed in [10] decomposes a face into several components (eyes, nose, mouth) which are replaced by templates with highest similarity from the opposite gender group. This results in unnatural images since templates from different images are spliced together. Instead of replacement, Othman and Ross [11] perform gender suppression through face morphing. Different from our approach, a face other than an average face from the opposite gender group is utilized. Thus, the gender suppression level differs based on the gender characteristic of the face used for morphing. Additionally, the performance of their gender suppression approach is evaluated by one gender classifier. Thus, it is unknown whether other gender classifiers would perform similarly or not. Additionally, the potential of decompressing or retrieving the suppressed gender information is not explored. With the prevalence of adversarial networks, Mirjalili et al. [12] develop Semi-Adversarial Networks (SANs) that can generate adversarial images which are robustly misclassified by gender classifiers. In another direction, He et al. [13] propose an AttGAN method for high-quality facial attribute editing. The complexity of [12] and [13] is much higher than using face morphing although they perform well in gender conversion.

III. Methods

In this section, we specify the gender obfuscation scheme and the evaluation methods for the degree of gender obfuscation and identity preservation.

A. Gender obfuscation

The average face can be considered as prototypes with specific gender characteristics. Generally, when a face is morphed with an average face of the opposite gender, its original gender characteristics are weakened. Therefore, the gender obfuscation scheme is as follows:

1) Male face $\oplus$ Female obfuscator (Average female face)
2) Female face $\oplus$ Male obfuscator (Average male face)

The '$\oplus$' indicates the process of face morphing. An obfuscator is an average face with opposite gender (relative to the non-obfuscated face). Examples of obfuscators are shown in Fig. 2.



(a) (b) (c)

Fig. 2. Obfuscators and gender neutral background (a: Male obfuscator, b: Female obfuscator, c: Gender neutral background)

The contour and the region outside of the face such as hair style, wearing clothes etc., affect gender classification. To avoid the influence of the region outside of the face on gender classification, we generate a gender neutral background, as shown in Fig. 2c. It is created by morphing 10 female faces and 10 male faces, so the background is the average background of the 20 face images. After morphing with an obfuscator, the resulted face is pasted into the gender neutral background in the face region using Poisson blending [14]. The schema of our proposed gender obfuscation procedure is shown in Fig. 3.

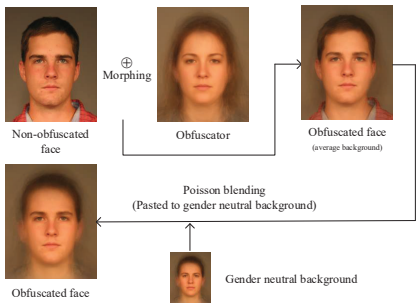

Fig. 3. Gender obfuscation scheme

Authorized licensed use limited to: UNIVERSITY OF TWENTE.. Downloaded on September 13,2021 at 13:57:06 UTC from IEEE Xplore. Restrictions apply.

1) Face Morphing: Given two face images, the corresponding facial landmarks need to be located first. In total, 71 landmarks are used. 68 landmark points are localized by *shape_predictor_68_face_landmarks* from the *dlib* library (http://dlib.net) at first. The faces we use have a neutral expression, which means that the mouth should be closed. The three landmarks depicting the upper contour of the lower lip are identical to the landmarks on the lower contour of the upper lip. We remove these three duplicates and add six landmarks on the forehead. The positions of the six landmarks are determined by the width w and height h of the bounding rectangle of the detected facial landmark points. Suppose that (x, y) is the position of the upper-left corner of the bounding rectangle. The positions of the extra three landmarks above the left (from the view of observers) eyebrow are: $(x + 0.1 \cdot w, y - 0.02 \cdot h)$, $(x + 0.15 \cdot w, y - 0.05 \cdot h)$, $(x + 0.33 \cdot w, y - 0.05 \cdot h)$ (pixel indices). The other three landmarks above the right eyebrow are in the symmetric locations.

After locating the landmark points in image I_0 and I_1, the positions of corresponding points in intermediate frame I_α are calculated using following equation:

$$P_\alpha = \{r_i | r_i = \alpha_w \cdot u_i + (1 - \alpha_w) \cdot v_i, u_i \in P_0, v_i \in P_1\} \quad (1)$$

where P_0 is the set of landmarks in I_0, P_1 is the set of landmarks in I_1 and P_α is the set of landmarks in the intermediate frame I_α. The positions of the ith point in P_0, P_1, and P_α are noted as u_i, v_i and r_i respectively. Then, Delaunay Triangulation [15] is applied, and for each pair of corresponding triangles, there is a warping function. Warping changes the facial structure, to what extent is controlled by the parameter α_w (warping parameter). Warping functions from I_α to I_0 ($w_{P_\alpha \longrightarrow P_0}$) and from I_α to I_1 ($w_{P_\alpha \longrightarrow P_1}$) are further used in the following equation:

$$I_\alpha(p) = \alpha_b \cdot I_0(\omega_{P_\alpha \to P_0}(p)) + (1 - \alpha_b) \cdot I_1(\omega_{P_\alpha \to P_1}(p)) \quad (2)$$

where p is the position of a point in I_α and $\omega_{P_\alpha \to P_0}(p)$ demonstrates its corresponding position in I_0. $I_\alpha(p)$ is the pixel value in position p in the intermediate frame I_α, computed from the pixel values in the corresponding positions in I_0 and I_1 with weights α_b (blending parameter) and $1 - \alpha_b$ through bilinear interpolation.

Morphing parameters α_w and α_b can be different. Research has shown that the influence of α_b is larger than that of α_w in face recognition [16] [17].

In our gender obfuscation scheme, average faces are needed, which are created by morphing multiple face images. The locations of corresponding points in I_α are computed by the following equation:

$$P_\alpha = \{r_i | r_i = \frac{u_{1,i} + u_{2,i} + ... + u_{N,i}}{N}, 1 \le n \le N\} \quad (3)$$

where $u_{n,i}$ is the position of the ith point in image I_n and N is the number of faces for morphing.

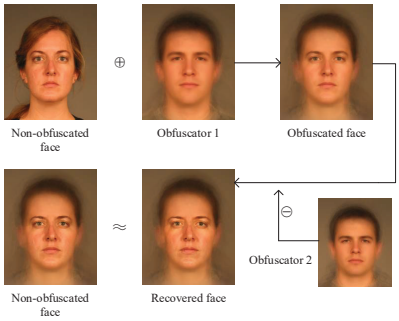

Fig. 4. An example of gender information recovery by face demorphing

Deducing from equation (2), the pixel values in an average face are computed as follows:

$$I_\alpha(p) = \frac{I_0(\omega_{P_\alpha \to P_0}(p)) + ... + I_N(\omega_{P_\alpha \to P_N}(p))}{N} \quad (4)$$

B. Gender retrieval

From the aspect of privacy protection, it would be undesirable if it were possible to recover or retrieve the protected information. In this section, potential retrieval methods are explored. One method is face demorphing. Another is to train a gender classifier which can directly distinguish between obfuscated female and male faces.

1) Face demorphing: It is the reverse process of face morphing. Suppose that we have subjects A and B, and their combined face through face morphing is noted as C, and $C = A \oplus B$. In real scenario, we have only the combined face C and a live captured face image A' which belongs to subject A. Then the demorphed image $D = C \ominus A'$, and it should be matched to subject B. In the case of recovering gender information in obfuscated faces, there is no such live captured image, but an average face. So the obfuscated face is demorphed with an average face, and the gender score of the demorphed face is utilized further. An example is shown in Fig. 4, where the '$\ominus$' indicates the process of face demorphing.

2) Gender classifier retraining: Training a gender classifier to distinguish between obfuscated female and male faces is a potential way as well. Due to the computational cost of training, a pre-trained model in [18] is used, with ShuffleNet V2 architecture [19]. To distinguish from the pre-trained gender classifier used in evaluation, the re-trained gender classifier is called Re-trained ShuffleNet while the pre-trained gender classifier is called ShuffleNet.

C. Evaluation

The goal is to obfuscate gender characteristics in face images while maintaining identity information. Therefore, two aspects need to be evaluated. The first is the degree of gender obfuscation, and the second is the degree of identity preservation.

Authorized licensed use limited to: UNIVERSITY OF TWENTE.. Downloaded on September 13,2021 at 13:57:06 UTC from IEEE Xplore. Restrictions apply.

1) Degree of gender obfuscation: For a confused gender classifier, its corresponding ROC curve should lie near the diagonal and the Area Under the Curve (AUC), which measures discrimination, should be approximately equal to 0.5, indicating the gender classifier loses its ability to distinguish between male and female faces. Therefore, the degree of gender obfuscation is computed as:

$$O_{\text{gender}} = \int_0^1 |R(t) - t| dt \tag{5}$$

where $R(t)$ is the value of the ROC curve in False positive rate t, and O_{gender} indicates the area between the ROC curve and the diagonal. The lower the O_{gender} is, the better the gender classifier is confused.

To observe intuitively the distributions of gender scores, the Empirical Cumulative Distribution Function (ECDF) is needed. Given a set of *N* ordered scores $(S_1, S_2, S_3, ..., S_N)$, ECDF is computed using Equation 6, where *n(i)* is the number of scores that are smaller than t and $t \in \{S_1, S_2, S_3, ..., S_N\}$.

$$E_N(t) = \frac{n(i)}{N} \tag{6}$$

Two gender classifiers are utilized and their performance are compared. One is from *FaceVacs SDK* (https://www.cognitec.com/) and another is provided by [18] with ShuffleNet architecture.

2) Degree of identity preservation: The identity information should be preserved mostly after obfuscating gender. Two face comparison modules are used. One is from *FaceVacs SDK*, and another is from python library *face_recognition* (https://github.com/ageitgey/face_recognition). Note that the python *face_recognition* module performs well in White faces, but its performance in other ethnicity groups cannot be guaranteed [20]. Meanwhile, *FaceVacs* face comparison module outputs a similarity score and python *face_recognition* module outputs a distance score.

Unlike confusing a gender classifier, the ROC curves for the face comparison modules are supposed to be ideal, indicating the corresponding AUC should be closed to 1. We define the degree of identity preservation as:

$$P_{\text{identity}} = AUC \cdot 100\% \tag{7}$$

IV. Experiments and Results

The purpose of the following experiments is to investigate the utility and versatility of face morphing in gender obfuscation and to explore the reversibility of our gender obfuscation scheme by two possible approaches.

A. Databases

Two databases are used for different tasks. These two databases are: FRGC [21] and CMU Multi-PIE [22]. FRGC is used for investigating the usability of our gender obfuscation approach while CMU Multi-PIE is used for exploring the reversibility of our approach.

FRGC: This database consists of around 50,000 recordings captured from different sessions. Face images captured in a controlled illumination condition (studio setting) with neutral facial expression are used to generate high-quality obfuscated faces.

CMU Multi-PIE: There are two kinds of images (Multi-view and High-resolution). The high-resolution images are utilized since they are captured frontally and have higher quality than those captured from multi-view. Subjects wearing glasses are removed manually since they affects the resulted morphed images.

Due to the limitation of python *face_recognition* module in other ethnicity groups, only White faces are used in the following experiments.

B. Obfuscate with different morphing parameters

As shown in [16], α_b influences the performance of face recognition more than α_w does. But the influence of these parameters on gender classification is unknown. Therefore, faces are morphed with different combinations of α_w and α_b. The gender classifiers and face verification modules introduced in Section III-C are used to evaluate the degree of gender obfuscation and identity preservation.

In the experiment, a face is morphed with an obfuscator. Examples of obfuscated faces are shown in Fig. 5.



Fig. 5. Example of non-obfuscated, obfuscated and average faces (When $\alpha_w = \alpha_b = 0$, face morphing is not applied to the subject face. When $\alpha_w = \alpha_b = 1$, the face is completely the obfuscator. To better compare the difference among faces, the non-obfuscated faces are pasted into the gender neutral background.)

The degree of identity preservation under different combinations of α_w and α_b is shown in Table I. It can be observed that *FaceVacs* face comparison module is vulnerable to morphing attacks [23] since the P_{identity} only decreases by at most 0.48%. The python *face_recognition module* is severely influenced when $\alpha_b = 0.4$. For the non-obfuscated faces, P_{identity} equals to 100% for *FaceVacs* face comparison module and 99.90% for python *face_recognition* module.

In terms of gender obfuscation, the degree of obfuscation with different combinations of α_w and α_b is shown in Table II. It shows that α_w has a slight influence on gender obfuscation while α_b matters more. However, the two gender classifiers are confused differently. For the *FaceVacs* gender classifier, α_w = 0.4 and α_b = 0.5 leads to the best obfuscation, while for the ShuffleNet gender classifier, α_w = α_b = 0.6 leads to the best obfuscation. Additionally, when $\alpha_w = 0.5$ and $\alpha_b = 0.6$, the difference between the two gender classifiers is smallest.

Authorized licensed use limited to: UNIVERSITY OF TWENTE.. Downloaded on September 13,2021 at 13:57:06 UTC from IEEE Xplore. Restrictions apply.

TABLE I
DEGREE OF IDENTITY PRESERVATION UNDER DIFFERENT COMBINATIONS OF α_w AND α_b

Face comparison module		P_{identity}		
		$\alpha_w = 0.4$	$\alpha_w = 0.5$	$\alpha_w = 0.6$
FaceVacs	$\alpha_b = 0.4$	99.56%	99.62%	99.49%
	$\alpha_b = 0.5$	99.91%	99.92%	99.86%
	$\alpha_b = 0.6$	99.97%	99.95%	99.96%
face_recognition	$\alpha_b = 0.4$	85.77%	89.67%	90.78%
	$\alpha_b = 0.5$	96.48%	96.48%	94.87%
	$\alpha_b = 0.6$	98.39%	98.24%	98.14%

TABLE II
DEGREE OF GENDER OBFUSCATION UNDER DIFFERENT COMBINATIONS OF α_w AND α_b

Gender classifier		O_{gender}		
		$\alpha_w = 0.4$	$\alpha_w = 0.5$	$\alpha_w = 0.6$
FaceVacs	$\alpha_b = 0.4$	0.429	0.407	0.386
	$\alpha_b = 0.5$	0.031	0.077	0.066
	$\alpha_b = 0.6$	0.261	0.153	0.173
ShuffleNet	$\alpha_b = 0.4$	0.397	0.396	0.347
	$\alpha_b = 0.5$	0.193	0.179	0.157
	$\alpha_b = 0.6$	0.154	0.165	0.106

Further decreasing α_b to 0.5 can improve gender obfuscation, especially for *FaceVacs* gender classifier.

Considering the trade-off between P_{identity} and O_{gender}, as well as the different performance of the two gender classifiers, the morphing parameters are set to: $\alpha_w = \alpha_b = 0.5$. Another reason to set $\alpha_w = 0.5$ is that some of the faces are not fully frontal, and setting $\alpha_w = 0.5$ can reduce pose variations without influence on identity verification.

Before generating obfuscators for the following experiments, we explored how the gender characteristics of the obfuscators and the number of faces used to create the obfuscators influenced P_{identity} and O_{gender}. Based on our experiments, we generated obfuscators using 20 randomly chosen faces from the Average-face gallery (FRGC FN Fall 03 & MN Spring 04) for the following experiments.

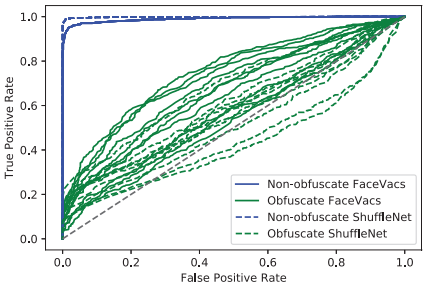

Fig. 6. ROC curves for 10-test gender obfuscation

C. Investigate versatility of gender obfuscation scheme

Due to the randomness of obfuscators, the experiment for gender obfuscation is repeated 10 times, each time with newly created male and female obfuscators. It can be observed from Fig. 6 that ROC curves lie near the diagonal but there are some deviations based on the obfuscators we used. In general, the performance of gender classifiers are poor. Selecting gender scores from one of the tests, the ECDFs of male's and female's gender scores are shown in Fig.7. After applying face morphing, the gender score distributions for males and females overlap. It demonstrates that the gender classifiers are confused. O_{gender} of the 10-test gender obfuscation is measured by Equation 5, which ranges from 0.077 to 0.298 (*FaceVacs*), and from 0.009 to 0.170 (*ShuffleNet*).

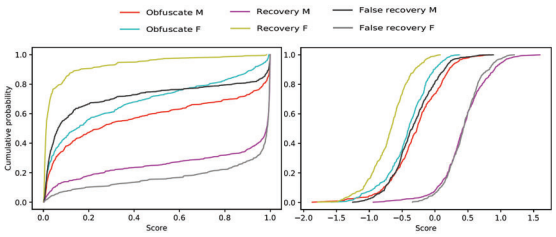

Fig. 7. Comparison of ECDFs (left: ShuffleNet gender classifier, right: *FaceVacs* gender classifier)

D. Retrieve gender in obfuscated face

It is critical that the hidden gender information cannot be recovered or retrieved easily. Two possible approaches are explored, and they are: face demorphing and gender classifier retraining. Since the obfuscators utilized in obfuscation should not be publicly available (nor the database used), a new gallery–CMU MultiPIE for obfuscator generation is utilized.

1) Retrieve by face demorphing: We consider a scenario in which someone has been presented with an obfuscated face and the gender of the obfuscator is unknown, indicating that the obfuscated face should be demorphed with both female obfuscator and male obfuscator.

In the experiments, we have the pre-known gender information of the obfuscated faces, and it enables us to analyze whether the gender information can be retrieved in this way. Before the experiments, some essential terms are defined as follows:

- Recovered face: an obfuscated face is demorphed with a correct obfuscator;
- False recovered face: an obfuscated face is demorphed with an incorrect obfuscator.

Because the morphing parameters are set to $\alpha_w = \alpha_b = 0.5$, the demorphing parameters are also set to be the same. De-obfuscation using face demorphing becomes harder if the morphing parameters are unknown.

It can be observed from Fig. 7 that face demorphing is not a feasible method to retrieve gender in obfuscated faces since the gender score distributions are still overlapped after demorphing with obfuscators with the same gender.

2) Retrain gender classifier: First, we prepare a dataset that contains obfuscated faces only. The faces are generated from 71 White males and 26 White females in CMU Multi-PIE database session 1. α_w and α_b are set to be the same and range from 0.4 to 0.6 with step = 0.003. Due to the imbalanced data distribution, under-sampling is utilized. Face images of males are removed randomly. Later, Gaussian noise is added to each

Authorized licensed use limited to: UNIVERSITY OF TWENTE.. Downloaded on September 13,2021 at 13:57:06 UTC from IEEE Xplore. Restrictions apply.

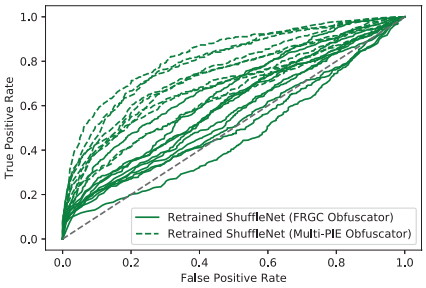

Fig. 8. ROC curves for re-trained gender obfuscation (solid lines: obfuscators generated from FRGC , dashed lines: obfuscators generated from Multi-PIE)

face with mean = 0 and variance = 0.02 for each color channel to avoid overfitting. Finally, there are 1542 obfuscated male and female faces. To be noted, the obfuscator for every non-obfuscated face is newly created from CMU Multi-PIE session 2.

The generated dataset is randomly shuffled and split into 80% for training and 20% for validation. After training, the gender classifier obtains an accuracy of 98.95% in training set and 89.97% in validation set. However, when applying the model to FRGC database (Non-obfuscated faces and obfuscators are from FRGC), the performance is not consistent with the results obtained during training (Fig. 8). Hence, we use obfuscators generated from CMU Multi-PIE session 2 instead (Fig. 8). The performance of gender classification improves slightly but the gender classifier is still confused. We further explore the similarity between obfuscated faces and randomly generated average faces, and no correlation is found. Thus, we conclude that it is complicated to train a gender classifier which can distinguish between obfuscated male and female faces.

V. Conclusion

This paper introduces a gender obfuscation scheme utilizing face morphing and investigates the reversibility of the scheme. Compared with other related methods [9] [24] [25], our approach is much simpler and requires less computation. Experimental results indicate that face morphing is a feasible method to confuse gender classifiers while identity information is well preserved, especially when the identity verification system is sensitive to morphing attacks. It has also been proven that recovering or retrieving the gender information from obfuscated face images is a tricky task, which demonstrates that the gender information is well secured. Future experiments can explore the functionality of face morphing in multiple facial attributes obfuscation, where it is assumed that there exists a trade-off between identity preservation and the number of facial attributes that are obfuscated simultaneously.

References

[1] H. T. F. Rhodes, *Alphonse Bertillon, father of scientific detection*. Greenwood Press, 1968.

[2] A. Dantcheva, C. Velardo, A. D'Angelo, and J.-L. Dugelay, "Bag of soft biometrics for person identification," *Multimedia Tools and Applications*, vol. 51, no. 2, p. 739–777, 2010.

[3] "General data protection regulation," *Official Journal of the European Union*, vol. 59, p. 35–36, May 2016.

[4] 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, 2017.

[5] S. Lee, G. Wolberg, and S. Y. Shin, "Polymorph: morphing among multiple images," *IEEE Computer Graphics and Applications*, vol. 18, no. 1, p. 58–71, 1998.

[6] N. Arad, N. Dyn, D. Reisfeld, and Y. Yeshurun, "Image warping by radial basis functions: Application to facial expressions," *CVGIP: Graphical Models and Image Processing*, vol. 56, no. 2, p. 161–172, 1994.

[7] T. Beier and S. Neely, "Feature-based image metamorphosis," *Proceedings of the 19th annual conference on Computer graphics and interactive techniques - SIGGRAPH '92*, 1992.

[8] S. Lee, G. Wolberg, K.-Y. Chwa, and S. Y. Shin, "Image metamorphosis with scattered feature constraints," *IEEE Transactions on Visualization and Computer Graphics*, vol. 2, no. 4, p. 337–354, 1996.

[9] S. Venkatesh, H. Zhang, R. Ramachandra, K. Raja, N. Damer, and C. Busch, "Can GAN generated morphs threaten face recognition systems equally as landmark based morphs? - vulnerability and detection," *2020 8th International Workshop on Biometrics and Forensics (IWBF)*, 2020.

[10] J. Suo, L. Lin, S. Shan, X. Chen, and W. Gao, "High-resolution face fusion for gender conversion," *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, vol. 41, no. 2, p. 226–237, 2011.

[11] A. Othman and A. Ross, "Privacy of facial soft biometrics: Suppressing gender but retaining identity," *Computer Vision - ECCV 2014 Workshops Lecture Notes in Computer Science*, p. 682–696, 2015.

[12] V. Mirjalili, S. Raschka, A. Namboodiri, and A. Ross, "Semi-adversarial Networks: Convolutional autoencoders for imparting privacy to face images," *2018 International Conference on Biometrics (ICB)*, 2018.

[13] Y. Viazovetskyi, V. Ivashkin, and E. Kashin, "Stylegan2 distillation for feed-forward image manipulation," 2020.

[14] P. Pérez, M. Gangnet, and A. Blake, "Poisson image editing," *ACM SIGGRAPH 2003 Papers on - SIGGRAPH '03*, 2003.

[15] F. P. Preparata and M. I. Shamos, *Proximity: Fundamental Algorithms*. Springer-Verlag, 1985, p. 185–223.

[16] M. Ferrara, A. Franco, and D. Maltoni, "Decoupling texture blending and shape warping in face morphing," in *2019 International Conference of the Biometrics Special Interest Group (BIOSIG)*, 2019, pp. 1–5.

[17] P. Korshunov and T. Ebrahimi, "Using face morphing to protect privacy," *2013 10th IEEE International Conference on Advanced Video and Signal Based Surveillance*, 2013.

[18] A. Greco, A. Saggese, M. Vento, and V. Vigilante, "A convolutional neural network for gender recognition optimizing the accuracy/speed tradeoff," *IEEE Access*, pp. 1–1, 2020.

[19] N. Ma, X. Zhang, H.-T. Zheng, and J. Sun, "ShuffleNet V2: Practical guidelines for efficient cnn architecture design," *Computer Vision – ECCV 2018 Lecture Notes in Computer Science*, p. 122–138, 2018.

[20] Ageitgey, "ageitgey/face_recognition." [Online]. Available: https://github.com/ageitgey/face_recognition/wiki/Face-Recognition-Accuracy-Problems

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

[22] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, "Multi-PIE," *2008 8th IEEE International Conference on Automatic Face and amp; Gesture Recognition*, 2008.

[23] U. Scherhag, C. Rathgeb, J. Merkle, R. Breithaupt, and C. Busch, "Face recognition systems under morphing attacks: A survey," *IEEE Access*, vol. 7, p. 23012–23026, 2019.

[24] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[25] V. Mirjalili, S. Raschka, and A. Ross, "Gender privacy: An ensemble of Semi Adversarial Networks for confounding arbitrary gender classifiers," *2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, 2018.

Authorized licensed use limited to: UNIVERSITY OF TWENTE.. Downloaded on September 13,2021 at 13:57:06 UTC from IEEE Xplore. Restrictions apply.