Privacy of Facial Soft Biometrics: Suppressing Gender But Retaining Identity

Asem Othman[(✉)] and Arun Ross

Michigan State University, East Lansing, MI, USA
{aothman,rossarun}@cse.msu.edu

Abstract. We consider the problem of perturbing a face image in such a way that it cannot be used to ascertain soft biometric attributes such as age, gender and race, but can be used for automatic face recognition. Such an exercise is useful for extending different levels of privacy to a face image in a central database. In this work, we focus on masking the gender information in a face image with respect to an automated gender estimation scheme, while retaining its ability to be used by a face matcher. To facilitate this privacy-enhancing technique, the input face image is combined with another face image via a morphing scheme resulting in a mixed image. The mixing process can be used to progressively modify the input image such that its gender information is progressively suppressed; however, the modified images can still be used for recognition purposes if necessary. Preliminary experiments on the MUCT database suggest the potential of the scheme in imparting "differential privacy" to face images.

1 Introduction

Most operational face recognition systems store the original face image of a subject in the database along with the extracted feature set (template). Storing the original image would allow the system to extract new feature sets and recompute templates if the feature extractor and matching modules are changed. However, face images offer additional information about an individual which can be automatically deduced. For instance, it has been shown that *automated* schemes can be used to extract soft biometric attributes such as age [4], gender [10], and race [8] from a face image. This can be viewed as privacy leakage since an entity can learn additional information about a person (or population) from the stored data, without receiving authorization from the person for such a disclosure. Therefore, it is necessary to ensure that face images stored in a system are used *only* for the intended purpose and not for purposes that may result in a "function creep" [21].

In this work, we investigate the possibility of suppressing the soft biometric attribute of a face (e.g., gender) while simultaneously preserving the ability of the face matcher to recognize the individual (see Figure 1). Such a capability will ensure that the stored biometric data is not used for purposes beyond what was expressed during the time of data collection. However, at the same time, it is necessary that any such perturbation does not drastically impact the recognition accuracy of the automated face matcher.
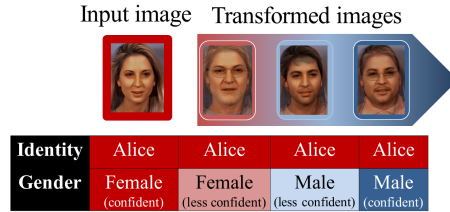
Input image   Transformed images

Identity	Alice	Alice	Alice	Alice
Gender	Female	Female	Male	Male
	(confident)	(less confident)	(less confident)	(confident)

Fig. 1. An illustration of progressively suppressing the gender of an input image while retaining identity with respect to an automated face matcher.

In this paper, a mixing approach is used to transform an input face image into a look-alike face image (i.e., similar facial features and may be appearance) while suppressing a specific soft biometric attribute, viz., gender. The **degree of suppression** is assessed by an *automated* gender classifier since the goal of this work is to disallow automated algorithms from extracting information beyond what was intended at the time of data collection[1]. A typical gender classifier outputs a label indicating the gender of a face image along with a confidence value of this determination. If the proposed method is successful, then either the confidence values associated with the transformed images will decrease (i.e., gender suppression) or their output label will change (i.e., gender conversion).

1.1 Related Work

Biometrics Privacy: In the context of biometrics, privacy refers to the assurance that the biometric data collected from an individual is not used to deduce any type of information about the individual. i.e., it should be used only for matching purposes.

Extensive work on preserving personal information has been done in the data mining community [1]. Their goal is to enable researchers and organizations to learn statistical properties of an underlying population[2] as a whole, while protecting sensitive information of the individuals in the population against "linkage attacks" [12]. Approaches such as k-anonymity [19], l-diversity [9], t-closeness [7], and differential privacy [2] have been proposed to preserve privacy of the personal data in statistical databases. These approaches employ techniques such as data perturbation and sub-sampling [19] (i.e., non-interactive privacy model), or provide an interface through which users may query about the data and get possibly noisy answers (i.e., interactive privacy model) [2]. In the context of biometrics, Newton et al. [13] and Gross et al. [5] introduced a face anonymization algorithm that minimized the chances of performing automatic face recognition on surveillance images while preserving details of the face such as expression,

---

[1] It is also possible to suppress soft biometric information from a human vision perspective - however, the work here does not explore the cognitive-psychological aspects of the transformed image.

[2] This population can be represented as statistical database.

gender and age. However, the identities of original face images are irrevocably lost, thereby undermining the use of such techniques in biometric databases.

In related literature [16], to protect the privacy of stored biometric data (e.g., face images) in a central database, template protection approaches have been proposed. Most of these template protection approaches replace the stored feature set in the central database with a transformed feature set or a cryptographic key that has been generated from the feature set or bound with it. These approaches, such as fuzzy vault cryptosystems, invariably result in loss of matching accuracy as demonstrated in the literature [16]. Further, when the feature extraction scheme is changed, the cryptosystem has to be changed. Some researchers have addressed the challenge of protecting biometric data at the image-level [15][22][6][3][14], *but their goal was to perturb identity.* Our methodology, on the other hand, only perturbs soft biometric attributes at the image level, while retaining identity, to prevent any gender profiling on an individual[3].

Face fusion for gender conversion: Fusing face images in order to change a perceived soft biometric (such as age, gender, and/or race) has been researched in both computer vision and graphics literature due to its many interesting applications. Regarding gender conversion while preserving face identity, there are two methods: a prototype-based approach [17] and a component-based approach [18]. In the prototype-based approach [17], prototypes for the two gender groups (male and female) are computed to describe the typical characteristics of males and females, respectively, and the difference between these two prototypes is used to modify the gender appearance of an input face image. A component-based approach was proposed by Suo et al. [18] as an alternative approach to gender conversion. Their approach starts by decomposing a source face image into several facial components. Next, these facial components are replaced with templates taken from the opposite gender group and the resulting mosaic is assembled using seamless image editing techniques. The identity of the source image is preserved by selecting replacement templates that are similar to that of the source components and penalizing large alterations in the image editing step. However, the goal of the gender conversion approaches described above is to generate a *single* face image that preserves the identity but modifies the gender. In this paper, our main objective is different. The input face image has to be transformed to *multiple* images that are similar to it but with the gender information suppressed at different levels. In other words, some of the generated images will be perceived to be of the same gender but with less confidence values, while other images will be perceived to be of the opposite gender with different confidence values.

1.2 Proposed Method

To generate a face image with aforementioned properties, the principle of face morphing is used. Consider two face images F_1 and F_2. The morphing algorithm generates intermediate images along the continuum from F_1 to F_2, and

---

[3] Population privacy is enhanced because an adversary cannot draw any conclusion about the gender of face database users.

their positions on this continuum are specified by the morphing parameters. The parameters, described later, are used to determine the rate of warping and color blending. So, as the morphing proceeds along the continuum from F_1 to F_2, the first image (F_1) is gradually distorted and is faded out, while the second image (F_2) is faded in (see Figure 4).

The key *contributions* of this paper are summarized as follows.

– Progressively suppressing the gender attribute of a face while preserving its identity from the face matcher's perspective. To the best of our knowledge, this paper is the first to present the *potential* of imparting differential privacy to face images via a simple face morphing technique.
– The degree of suppression has been systematically quantified by utilizing an *automated* gender classifier. The proposed method is expected to be applicable across different gender classifiers since it has not been particularly tuned to a specific one.

The rest of the paper is organized as follows. Section 2 discusses the face morphing technique to perturb gender attributes. Section 3 reports the experimental results and Section 4 concludes the paper.

2   Face Morphing

Figure 2 shows the three distinct phases in the generation of a mixed face image (MF) : facial feature extraction, image warping and cross-dissolving.
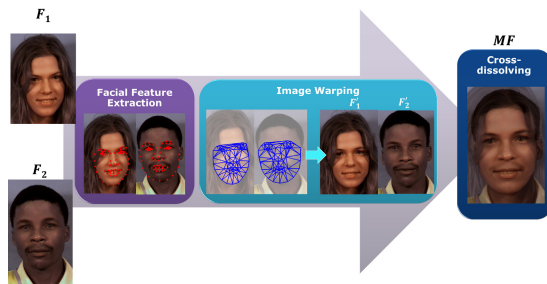

Fig. 2. Proposed approach for suppressing gender while retaining identity. Here, the gender of F_1 is perturbed by mixing it with F_2 resulting in image MF. However, an automated face matcher can successfully match MF with F_1.

2.1   Facial Feature Extraction

Morphing two face images to generate an intermediate face image involves the nontrivial task of locating facial features. For both face images, F_1 and F_2, the prominent facial features are characterized by a pre-defined set of control points.

Both sets of control points, X_1 and X_2, associated with the two face images (see Figure 2), are stored in a vector format. This representation does not include any information about the connection between the control points:

$$X_j = [x_{1j}, x_{2j}, x_{3j}, \ldots, x_{nj}, y_{1j}, y_{2j}, y_{3j}, \ldots y_{nj}]^T, \tag{1}$$

where $j \in \{1, 2\}$ and $n = 76$ is the number of control points. Errors in landmark annotation can cause a ghost-like effect on the subsequently generated image. Since extracting control points automatically is not the focus of this work, a pre-annotated face image database was used (see Section 3). This minimizes the ghost-like effect.

2.2   Image Warping

Once the corresponding control points between the two face images are known, the next step is to perform image warping by mapping each facial feature (e.g., mouth, nose and eyes) in the individual face images to its corresponding feature in the mixed image. A triangulation-based warping scheme is used to deform the face images [20]. First, the intermediate control points set (which defines the shape of the facial features of the mixed face image) is determined. From the control point sets X_1 and X_2 of the face images F_1 and F_2, respectively, the intermediate control point set (X_m) is obtained by linear interpolation as follows:

$$X_m = (1 - \alpha) \cdot X_1 + \alpha \cdot X_2, \tag{2}$$

where $\alpha \in [0, 1]$ is the **warping factor** that determines how the individual shapes of the two face images are integrated into the shape of the mixed face. Next, the face region of each face image is dissected into a suitable set of triangles by utilizing the control points as the vertices of the triangles. Generating an optimal triangulation has to be guaranteed in order to avoid skinny triangles and, therefore, Delaunay triangulation was utilized to construct a triangular mesh for each face image. An example of face images tessellated into triangular regions according to the annotated control points is shown in Figure 2.

Finally, the affine transformation that relates each triangular region in the original face image (F_1 or F_2) to the corresponding triangle in the intermediate image is computed. Suppose that $T_1 = [P_1, P_2, P_3]^T$ $(T_2 = [R_1, R_2, R_3]^T)$ is a triangular region in X_1 (X_2) and $T_m = [Q_1, Q_2, Q_3]^T$ is the corresponding triangular region in X_m (see Figure 3). A_1 (A_2) is the affine transformation that maps all points in T_1 (T_2) onto T_m.

$$T_m = A_j T_j, \tag{3}$$

where $j \in \{1, 2\}$. Together, T_1's $(T_2$'s) vertices and T_m's vertices are used in (3) to compute the parameters of the affine transformation A_1 (A_2).

As shown in Figure 2, this results in two warped face images F_1' and F_2' such that F_1' and F_2' have similar shapes.
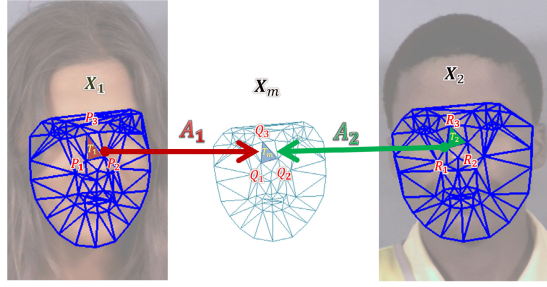
Fig. 3. Generating the corresponding triangle, T_m, in the intermediate image based on the triangles T_1 and T_2 in the original images

2.3 Image Cross-dissolving

The final step to obtain the mixed face image, is simply a cross-dissolving process of the two warped images. If F_1' and F_2' are the warped images, the mixed face image is obtained by linearly interpolating their pixel intensities, such that

$$MF = (1 - \beta) \cdot F_1' + \beta \cdot F_2', \tag{4}$$

where $\beta \in [0, 1]$ is the **color-dissolving factor** that determines the relative influence of the appearance of the two face images on the mixed face image MF. Figure 4 shows different examples of mixed face images along the continuum from F_1 to F_2 by varying the warping factor (α) and the cross-dissolving factor (β).

Fig. 4. Mixed face images along the continuum from F_1 to F_2 where $\alpha = \beta =$ (a) 0.1, (b) 0.2, (c) 0.3, (d) 0.5, (e) 0.7, (f) 0.8 and (g) 0.9

3 Experiments and Discussion

The purpose of the following experiments is to systematically investigate if mixing an input face image with different face images from the opposite gender group will (1) suppress the gender attribute of the input face image to different degrees, and (2) preserve the identity of the input image with respect to a face

matcher. To generate a mixed face image from two face images; i.e., a male face image F_m and a female face image F_f, the morphing technique described earlier is utilized and the mixed face image can be anywhere along the continuum from F_m to F_f. But where on this continuous continuum should the mixed face image be? The position of the mixed face on this continuum is specified by the morphing parameters, i.e., α and β. Although the two parameters can be different, the best visually appealing mixed face image along this continuum is observed when $\alpha = \beta$. Also, if $\alpha = \beta < 0.5$, the gender of the source image will not be suppressed effectively. Contrarily, if $\alpha = \beta > 0.5$, the identity of the source will be suppressed and the mixed image will not be similar to it. Thus, we select $\alpha = \beta = 0.5$.

3.1   Performance Metrics

The notion of similarity/dissimilarity between face images is assessed using the match scores generated by a Verilook[4] face matcher. In the context of identification, a higher rank-1 accuracy would imply a higher similarity; in the context of verification, a lower Equal Error Rate (EER) would imply higher similarity. So we use rank-1 accuracy and EER to characterize notions of similarity and dissimilarity.

The gender of a face image (male or female) is assessed using a VeriLook gender classifier [5], which also outputs classification confidence values (C') along with the gender label. These confidence values are in the [0, 100] interval. A confidence value of 0 indicates that the image is in the boundary of the male and female class[6]. However, the software labels the image as female when the confidence value is 0. Here, we mapped the resultant confidence values as follows:

$$C = \begin{cases} C'/100 & \text{if class} = \text{male}; \\ -C'/100 & \text{if class} = \text{female}, \end{cases}$$

where male and female are the labels computed by the gender classifier. This mapping results in a gender axis with two ends: 1 (i.e., male with a confidence value = 100%) and -1 (i.e., female with a confidence value = 100%). This gender axis will be used to quantify as well as visualize the degree to which gender is suppressed in the forthcoming experiments.

3.2   Database and Baseline Performance

The performance of the proposed approach was tested using a dataset from the MUCT database [11]. MUCT database consists of 3755 face images of 276 subjects. We selected the first 2 samples captured by camera "a" (usually the

---

[4] http://www.neurotechnology.com

[5] Since the proposed method is not particularly tuned to the specific gender classifier used, it is expected to work as well on other types of gender classifiers.

[6] This assertion has been confirmed by consulting the technical support at Neurotechnology.

frontal face) of each subject [7]. For each subject, one sample was added to the probe set and the other sample was added to the gallery set resulting in a probe set P and gallery set G each containing 276 face images. The images in P were matched against those in G. This resulted in a rank-1 accuracy of 95% and an Equal Error Rate (EER) of 3.5%. This dataset was used since the facial landmarks (control points) of individual images were annotated and available online, and also because it contains a comparable number of males and females (i.e., 131 males and 145 females). The ground truth for gender was obtained from the filename ("m" for male and "f" for female). The Verilook gender classifier was used to classify the face images in the gallery set G. Figure 5 shows examples of face images from G along the gender axis based on the predicted gender and confidence values. There are only 5 images from G that were misclassified (see Figure 5). Therefore, from the perspective of this automated gender classifier, the gallery set will be divided into a male dataset G_m consisting of 132 males and a female dataset G_f consisting of 144 females.
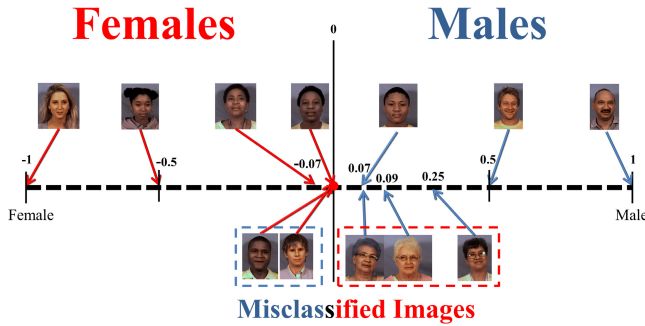

Fig. 5. Examples of frontal face images from the gallery set G shown on the gender axis. The gender axis is based on the gender as estimated by the classifier along with the confidence value, C. Misclassified images are depicted below the axis.

3.3   Degrees of Gender Suppression

In this experiment, the possibility of generating images with different gender suppression levels is tested. Every face image in the male gallery set G_m is mixed with every face image in the female gallery set G_f. This results in 19,008 mixed face images. For every male face image, there are 144 corresponding mixed face images. For every female face image, there are 132 corresponding mixed face images. Figure 6 shows the distribution of confidence values (C) of the mixed

---

[7] Camera "a" was the only camera that was directly in front of the subject's face. Images captured by other cameras exhibited some pose variations. In this paper, we used only frontal images to examine viability of the proposed approach. There are two or three images per subject captured by camera "a" and we selected two samples in order to have the same number of samples for all subjects.
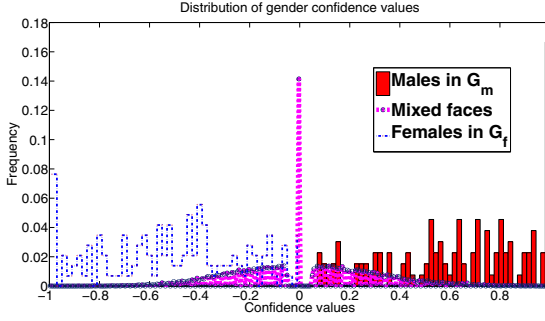
Fig. 6. Distribution of confidence values of the male, female and mixed images.

images as well as the original male and female images. This graph clearly suggests the potential of suppressing gender attributes using the proposed method. However, the objective is not just to suppress the gender attribute. We are looking to generate images that reveal different levels of gender. Another way of looking at this is as follows: if a male face image F_m is fused with different female face images F_fs, the resulting face images (MFs) should have confidence values (C_{mf}s) that vary from -1 to the confidence value of the original male image, C_m. To determine the *degree* of gender suppression, the **male suppression-level** of a mixed image is computed as follows:

$$S_m = \frac{C_m - C_{mf}}{C_m + 1}.$$ (5)

If $S_m = 0$, this indicates that the mixed image has the same confidence value as the original male image (F_m) and the gender is not suppressed. If $S_m = 1$, this indicates that the mixed image has been classified as female with $C_{mf} = -1$ and the gender of the original male image is completely suppressed. On the other hand, if the source is a female image which has been mixed with a set of male images, the goal would be to generate mixed images with suppression-levels that start from $C_{mf} = C_f$ and end at $C_{mf} = 1$. Therefore, if the input image is female, the **female suppression-level** can be computed as follows:

$$S_f = \frac{|C_f| + C_{mf}}{|C_f| + 1}.$$ (6)

Note that in this particular database S_m and $S_f \in [0, 1]$ because, as shown in Figure 6, the confidence value of a mixed face image (C_{mf}) is always less than the confidence values of the original male subjects (C_m) and greater than the confidence values of the original female subjects (C_f), i.e., $C_f \leq C_{mf} \leq C_m$. Figures 7 and 8 show S_m and S_f, respectively, for all mixed images (i.e., the 19,008 face images). These graphs suggest the possibility of having different levels of gender suppression. This can be observed by viewing the range of colors in each row or each column.
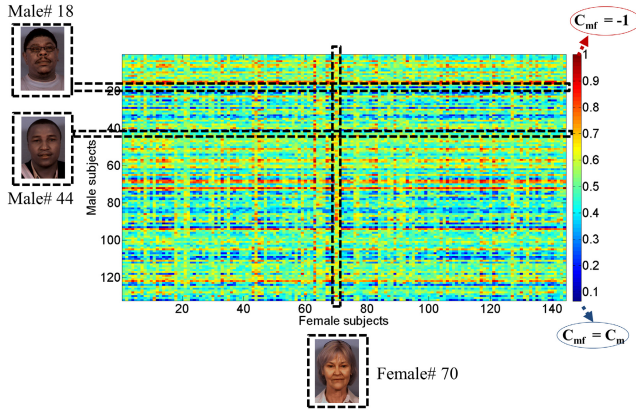
Fig. 7. Plot of male suppression-levels of the mixed images where points along the horizontal axis is the id # of female subjects and points along the vertical axis is the id # of male subjects. Male suppression-levels after mixing male subjects #18 and #44 are highlighted. Additionally, the male-suppression level when female subject #70 is used is also highlighted. See text for explanation as to why these three subjects are highlighted.
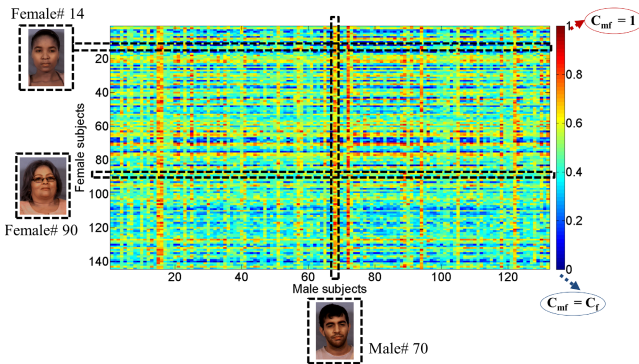

Fig. 8. Plot of female suppression-levels of the mixed images where points along the horizontal axis is the id # of male subjects and points along the vertical axis is the id # of female subjects. Female suppression-levels after mixing female subjects #14 and #90 are highlighted. Additionally, the female-suppression level when male subject #70 is used is also highlighted. See text for explanation as to why these three subjects are highlighted.

Figure 9 show two male images and the results of mixing them with different females images along with gender confidence values. Figure 11(a) show S_m for all mixed images generated by these two male subjects. Note that, the male suppression-levels (S_m) of mixed images generated by male subject #18 tend to be closer to the original gender confidence value (i.e., S_m tends to be closer to 0). On the other hand, the mixed images of male subject #44 tend to be

classified as females and are closer to the target (i.e., $C_{mf} \simeq -1$). A similar effect on female subject #18 and #44, can be seen in Figure 10. Figure 11(b) shows the female suppression-levels (S_f) for all mixed images generated by these two female subjects. We also observed that some female images when mixed with input male images cause most of the mixed images to have male suppression-levels that are closer to 1 (i.e., C_{mf} are closer to -1). For example, as shown in Figures 7 and 9, female subject #70 strongly suppressed the gender of most male images. Similarly, as shown in Figures 8 and 10, male subject #70 has strongly affected most of the female suppression-levels of the mixed face images.
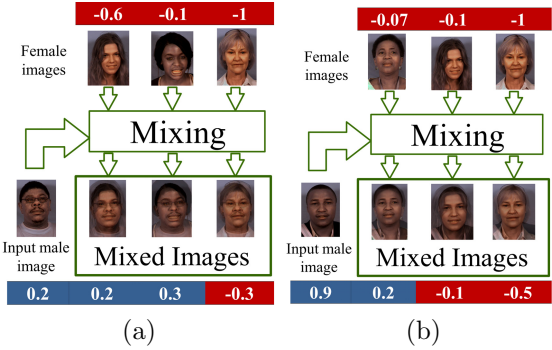

Fig. 9. Examples of mixed face images after mixing the face images of male subjects (a)# 18 and (b)#44 with different female face images, along with the confidence value (C) of each image. The blue (red) color indicates that the image is labeled as a male (female).
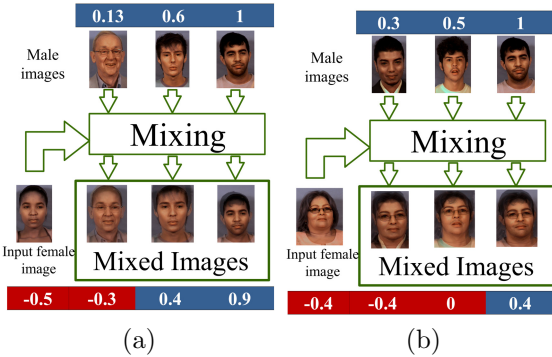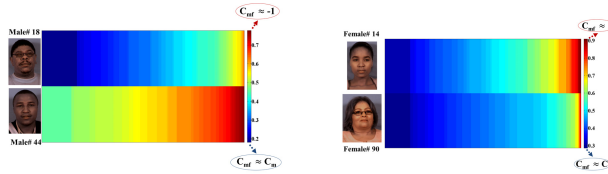

Fig. 10. Examples of mixed face images after mixing the face images of female subjects (a) # 14 and (b) #90 with different male face images, along with the confidence value (C) of each image

These results prove that we can generate different versions of an input face image with different levels of gender suppression. Note that the confidence values associated with the original images play an important role. As shown in

(a) Male suppression-levels (b) Female suppression-levels

Fig. 11. Plot of sorted suppression-levels corresponding to the mixed images generated for (a) male subjects #18 & #44, and (b) female subjects #14 & #90

Figure 9(a), the facial appearance of the input face image has also a role. Hence, the facial hair of the male subject #18, i.e., the mustache, results in mixed images with more maleness, although the original male image has a low gender confidence value.

3.4   Similarity to the Original Face Images

After suppressing or modifying the gender of the face image by mixing, the identity information should be preserved effectively in the resultant images. Therefore, in this experiment, the similarity between the mixed face image and original face images (i.e., male and female face images) was evaluated. To this end, the mixed face images generated in Experiment 1 (i.e., 19,008 face images) were matched against the original images in the probe set P (see Figure 12). Here, a genuine score is generated when the mixed face image is matched with either of the original face images (i.e., the images that were mixed) and the rest are impostor scores.

The resultant rank-1 accuracy of matching mixed images against original images in P was 95% (and the EER was 5%). These results indicate that the original images are reasonably similar to the mixed images. Hence, the identities of the originals have been preserved in the mixed faces, which is our objective.

4   Summary and Future Work

In this work, we explored the possibility of generating mixed face images that suppress the gender of a face image to *different degrees*. In this regard, we mixed an input face image with different face images from the opposite gender and determined if the mixed images suppress the gender while bearing sufficient similarity to the input face image in terms of a face matcher. We utilized a gender classifier along with the resultant confidence value to assess the gender information. To mix two face images, a face morphing technique was adopted in this work. Experiments on the MUCT dataset indicate that (a) the mixed face suppresses the gender of original face images to different degrees, and (b) the mixed face exhibits similarity with the original image and so identity with respect to a face matcher is retained. Figure 13 shows that the distribution of male and female suppression-levels are not uniform distributions. While it is
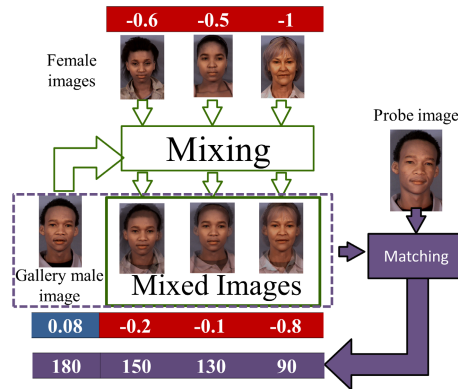
Fig. 12. Examples of mixing the face images of a male subject with different female face images, along with the confidence value (C) of each image. Match scores generated by matching the input probe against the mixed images and the gallery image of the male subject are shown below the confidence values of the gallery and mixed images. These scores are similarity scores and in the [0,180] interval.
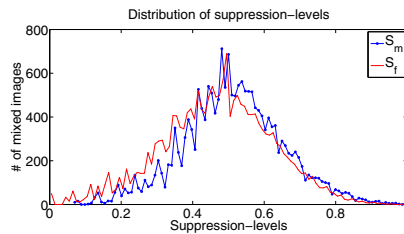

Fig. 13. Distributions of male (S_m) and female (S_f) suppression-levels of all mixed images (i.e., 19,800 face images). These distributions suggest that it is possible to suppress gender to different degrees.

possible to suppress a face image to different degrees, these degrees may not form a continuous or complete continuum[8]. Therefore, further work is needed to test this approach on a larger database having subjects with large variation in their gender confidence values. Other morphing approaches based on radial basis functions and multi-level free-form deformation [20] could be explored. The technique could potentially be extended to suppress different soft biometric attributes simultaneously (to different degrees) thereby supporting a differential privacy framework. Note that mixing more than two images is possible, but this may suppress individual identities and the mixed image is likely to be less similar to the originals. Future work will also investigate the possibility of utilizing the proposed approach as a privacy-enhancing technique by mixing faces of different subjects in order to *hide* the original identities.

---

[8] This continuum should start from the gender confidence value of the input face image and end at the maximum confidence value of the opposite gender (i.e., +1 or -1).

References

1. Agrawal, R., Srikant, R.: Privacy-preserving data mining. ACM Sigmod. Record **29**(2), 439–450 (2000)
2. Dwork, C.: Differential Privacy. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4052, pp. 1–12. Springer, Heidelberg (2006)
3. Färberböck, P., Hämmerle-Uhl, J., Kaaser, D., Pschernig, E., Uhl, A.: Transforming rectangular and polar iris images to enable cancelable biometrics. In: Image Analysis and Recognition, pp. 276–286. Springer (2010)
4. Fu, Y., Guo, G., Huang, T.S.: Age synthesis and estimation via faces: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence **32**(11), 1955–1976 (2010)
5. Gross, R., Sweeney, L., De la Torre, F., Baker, S.: Model-based face de-identification. In: Computer Vision and Pattern Recognition Workshop (CVPRW), pp. 161–168. IEEE Computer Society, Los Alamitos (2006)
6. Hämmerle-Uhl, J., Pschernig, E., Uhl, A.: Cancelable Iris Biometrics Using Block Re-mapping and Image Warping. In: Samarati, P., Yung, M., Martinelli, F., Ardagna, C.A. (eds.) ISC 2009. LNCS, vol. 5735, pp. 135–142. Springer, Heidelberg (2009)
7. Li, N., Li, T., Venkatasubramanian, S.: t-closeness: Privacy beyond k-anonymity and l-diversity. In: IEEE 23rd International Conference on Data Engineering (ICDE), vol. 7, pp. 106–115 (2007)
8. Lu, X., Jain, A.K.: Ethnicity identification from face images. In: SPIE Defense and Security Symposium, pp. 114–123 (2004)
9. Machanavajjhala, A., Kifer, D., Gehrke, J., Venkitasubramaniam, M.: l-diversity: Privacy beyond k-anonymity. In: ACM Transactions on Knowledge Discovery from Data (TKDD) (2007)
10. Makinen, E., Raisamo, R.: Evaluation of gender classification methods with automatically detected and aligned faces. IEEE Transactions on Pattern Analysis and Machine Intelligence **30**(3), 541–547 (2008)
11. Milborrow, S., Morkel, J., Nicolls, F.: The MUCT Landmarked Face Database. Pattern Recognition Association of South Africa (2010). http://www.milbo.org/muct
12. Narayanan, A., Shmatikov, V.: Robust de-anonymization of large sparse datasets. In: IEEE Symposium on Security and Privacy, pp. 111–125 (2008)
13. Newton, E., Sweeney, L., Malin, B.: Preserving privacy by de-identifying face images. IEEE Transactions on Knowledge and Data Engineering **17**, 232–243 (2005)
14. Othman, A., Ross, A.: On mixing fingerprints. IEEE Transactions on Information Forensics and Security **8**(1), 260–267 (2013)
15. Ratha, N., Connell, J., Bolle, R.: Enhancing security and privacy in biometrics-based authentication systems. IBM Systems Journal **40**(3), 614–634 (2001)
16. Rathgeb, C., Uhl, A.: A survey on biometric cryptosystems and cancelable biometrics. EURASIP Journal on Information Security **1**, 1–25 (2011)
17. Rowland, D., Perrett, D.: Manipulating facial appearance through shape and color. Computer Graphics and Applications **15**(5), 70–76 (1995)

18. Suo, J., Lin, L., Shan, S., Chen, X., Gao, W.: High-resolution face fusion for gender conversion. In: IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans, pp. 1–12 (2011)
19. Sweeney, L.: k-anonymity: A model for protecting privacy. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems **10**(05), 557–570 (2002)
20. Wolberg, G.: Image morphing: a survey. The Visual Computer **14**(8), 360–372 (1998)
21. Woodward, J., Orlans, N., Higgins, P.: Biometrics:identity assurance in the information age. McGraw-Hill/Osborne, New York (2003)
22. Zuo, J., Ratha, N., Connell, J.: Cancelable iris biometric. In: IEEE 19th International Conference on Pattern Recognition (ICPR), pp. 1–4 (2008)