 CMC, vol.57, no.1, pp.107-121, 2018

Expression Preserved Face Privacy Protection Based on Multi-mode Discriminant Analysis

Xiang Wang[1, *], Chen Xiong[1], Qingqi Pei[1] and Youyang Qu[2]

Abstract: Most visual privacy protection methods only hide the identity information of the face images, but the expression, behavior and some other information, which are of great significant in the live broadcast and other scenarios, are also destroyed by the privacy protection process. To this end, this paper introduces a method to remove the identity information while preserving the expression information by performing multi-mode discriminant analysis on the images normalized with AAM algorithm. The face images are decomposed into mutually orthogonal subspaces corresponding to face attributes such as gender, race and expression, each of which owns related characteristic parameters. Then, the expression parameter is preserves to keep the facial expression information while others parameters, including gender and race, are modified to protect face privacy. The experiments show that this method yields well performance on both data utility and privacy protection.

Keywords: Privacy protection, multi-mode, discrimination, expression-preserving.

1 Introduction

With the proliferation of inexpensive video surveillance and Internet, an increasing number of cameras are applied to a variety of public places. These cameras are used to record, store and process large amounts of image data. And in dealing with the image data in different scenes, human's privacy captured in the images is inevitably threatened. At the same time, in the rapid development of the diversification services of Internet technology, the users' private information will be leaked. As a result, people are increasingly concerned about whether their privacy information is protected.

As the monitoring system and the Internet in people's lives have become more and more popular, for how to solve security issues exist in these scenarios, people have put forward varies effective methods. Some researchers have proposed to remove face identification information from images in a scene where there is a security risk. There are two kinds of solutions: ad-hoc methods such as blurring and pixelation [Neustaedter, Greenberg and Boyle (2006)], as well as K-Same [Newton, Sweeney and Malin (2005)] and K-Same-M [Gross, Sweeney, Fernando et al. (2006)], and other conventional methods. Neustaedter et al. [Neustaedter and Greenberg (2003)] destroyed the identity information by blurring the face, Berger [Berger (2000)] employed the pixelation to carry out de-identification

[1] Xidian University, Xi'an, 710071, China.

[2] Deakin University, Burwood, 3125, Australia.

[*] Corresponding Author: Xiang Wang. Email: wangxiang@xidian.edu.cn.

operation on the images and retain skin color; Newton et al. [Newton, Sweeney and Malin (2005)] proposed K-Same-Family, where the face image to be protected is replaced by the average image of its the most similar K face images so that the probability of identifying the protected face is no more than 1/K. However, these face de-identification methods are not to remove all the information of the images, but only to hide the identity information of the images, while the expression, behavior and other awareness information are not considered. Some researchers [Hudson and Smith (1996)] believed that there was a fundamental trade-off between data utility and privacy. When the image is modified, the effective information used to understand the image may be deleted, so the modified image may become lack of practicality, hence it is particularly important to put a balance between privacy and data utility. In the face de-identification methods, people are usually concerned about the private information than the utility of images, so in the general de-identification images, the expression of face and other effective information will be affected, and the information in the video broadcast, video surveillance and other scenes are indispensable. Therefore, considering the balance between privacy protection and data utility, the proposed method preserves the expression information of the original image while ensuring that the privacy information is protected, so that the processed face image remains utility without revealing the identity information. Applying the multi-mode discriminant analysis [Sim, Zhang, Li et al. (2009)], the images containing multiple attributes are decomposed into orthogonal subspaces, and the characteristic parameters of each attributes such as gender, race and expression are obtained. By changing the parameters corresponding to the attributes such as gender and race while maintaining the expression parameters the same, we can get the de-identification images whose known facial attributes have changed in addition to expression while the privacy is protected.

2 Related works

The rapid development of video surveillance and Internet has brought convenience for people's lives, but has also caused a hidden danger for people's privacy. More and more researchers have made a lot of important work to protect people's privacy. In the field of data mining, image and computer vision, issues related to face de-identification have also been resolved. For example, Newton and Gross et al. [Newton, Sweeney and Malin (2005); Gross, Sweeney, Fernando et al. (2006)] studied the trade-off between privacy protection and data utility, and some of these ideas are effectively applied to face de-identification. In computer vision, earlier face de-identification methods are ad-hoc methods such as masking, blurring and pixelation. However, these techniques are easy to apply to any image, but the performance in the face de-identification is not guaranteed.

Newton et al. [Newton, Sweeney and Malin (2005)] proposed a k-Same algorithm based on k-anonymity for the above problem. The k-Same algorithm achieves the purpose of face de-identification by averaging the k images in the image set to obtain the average face image. The recognition rate of the face de-identification image obtained by the k-Same algorithm does not exceed 1/k, so the privacy is protected. However, although the privacy of the face images is protected, the k-Same algorithm destroys the utility of the original data, and the de-identification images will show fuzzy and ghosting phenomenon. In view of the shortcomings in k-Same algorithm, the researchers proposed an improved

algorithm for k-Same. k-Same-Select algorithm [Gross, Airoldi, Malin et al. (2006)] was proposed to obtain the de-identification images by dividing the image set into separate subsets which contain different attribute information. Using the k-Same-Select algorithm, the privacy of face images is protected and the attribute of each subset is well preserved. The k-Same-M algorithm was proposed by Gross et al. [Gross, Sweeney, Fernando et al. (2006)], which used the AAM [Cootes, Edwards and Taylor (1998)] to get the shape and appearance parameters of the image, and applied the k-Same algorithm for the appearance parameters of k images. The AAM model in the k-Same-M algorithm effectively solved the misalignment of face images.

The initial face de-identification methods only consider the identity information of the face images, and more and more face de-identification researches work to consider the attribute information of the images. Gross et al. [Gross, Sweeney, Torre et al. (2008)] introduced a multifactorial model that can decompose the test image into identity factors and non-identity factors, then applied the de-identification method to identity factors. The de-identification image obtained by this method contains the expression information. But the method has poor performance in de-identification. Du et al. [Du, Yi, Blasch et al. (2014)] retained the race, sex, and age attributes, and k-Same was applied after measuring attributes feature and selecting the corresponding AAM parameters. Each attribute corresponds to an AAM model. Othman et al. [Othman and Ross (2014)] changed the age information while protecting the privacy information.

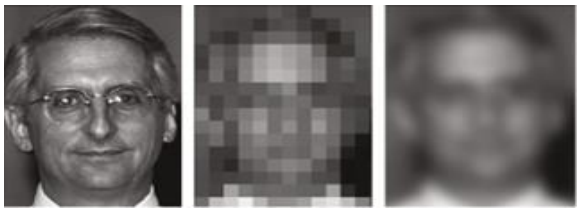

(a) Pixilation and blurring

(b) k-Same algorithm (c) k-Same-Select algorithm

(d) k-Same-M algorithm (e) Multifactor-map algorithm

Figure 1: Visual comparison of recently published de-identification approaches

In this paper, we maintain the facial expression information while protecting the privacy of face images. By establishing the multi-mode subspaces, we can get different attributes of the face. Changing the attributes other than the expression and keeping the expression features not change to produce the face de-identification images while retaining the original expression.

3 A new approach based on MMDA

In this paper, we apply the AAM algorithm to normalize the shape and appearance of the training sets, and then decompose the images in multiple modes. By decomposing the different attributes of the face into the corresponding independent subspace, we get the attribute parameters of different classes under each attribute, then the attributes feature can be changed by the related parameters. In this paper, by changing the other facial attributes of face images while maintaining the same expression, the privacy of the face images are protected and the original expression are retained.

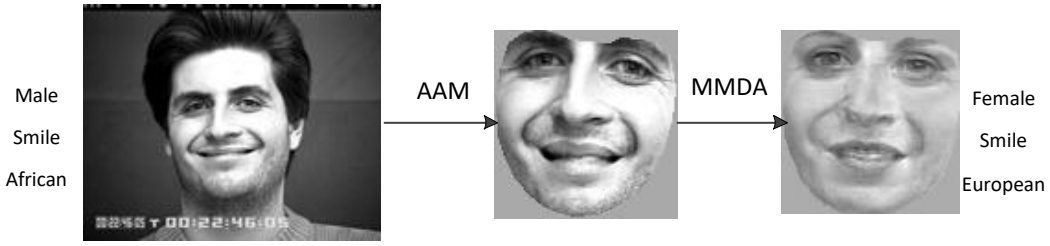

Figure 2: The proposed method can protect privacy while remaining expression attribute

Figure 3: 20 objects selected from CK+ database containing gender, race and expression

attributes

Our training set contains 20 objects with size of 640∗490 selected from the Extended Cohn-Kanade database [Lucey, Cohn, Kanade et al. (2010)]. Each image in the training set has three attributes such as gender, race and expression. The gender attribute includes male and female, and race attribute includes Africa and Europe, and the expression includes five classes such as smile, angry, sad, fear and surprise. The number of images corresponding to the different classes of each attribute is the same.

3.1 Pre-processing and normalization

Before applying multi-model discriminant analysis, we need to eliminate the effect of factors such as angle, attitude and position in the face images, also pre-process and normalize the images. In this paper, we mark the key facial points and use AAM method to get normalized images.

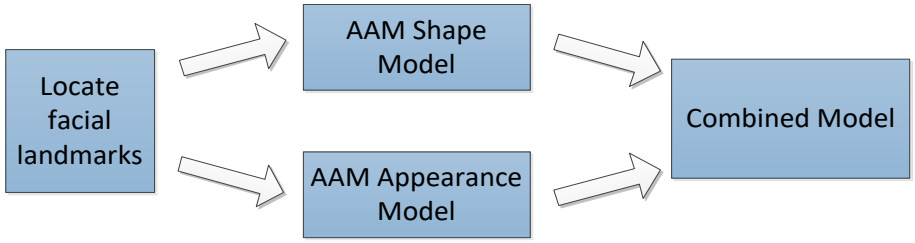

Figure 4: Overview of the pre-processing

3.1.1 Locating landmarks

Given a face image, we first locate its landmarks. By obtaining the coordinates of the landmarks, we can get the shape information. In this paper, 68 landmarks including facial contours, eyes, nose and mouth are selected. Then the shape vectors are produced by forming the coordinates into vectors.

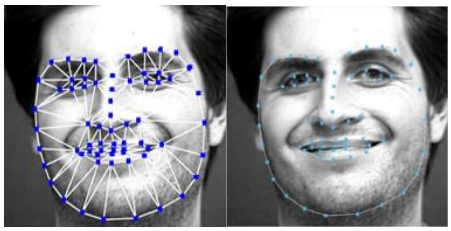

Figure 5: Locating the facial landmarks

3.1.2 Active appearance models

To normalize the image whose landmarks are located, we use the AAM to align the image and get shape-normalized appearance. By establishing the average shape model and the average appearance model respectively, and using the Procrustes algorithm [Dryden and Mardia (1998)] and PCA [Kirby and Sirovich (2002)], we can get the parameters p and λ which control the two models respectively. The shape model and appearance model are

combined to obtain a mixed model. After completing the AAM modeling, each image in the sample set is normalized to get a standardized face, so that the effects of facial attitude, angle, position and other factors can be removed.

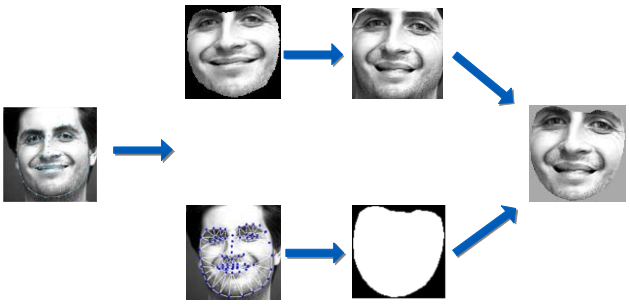

Figure 6: Normalizing a face image by shape model and appearance model

We first locate facial landmarks for the training set, and 68 landmarks are marked on the key parts of the face, such as the eyes, the nose, the mouth and the face contour. Each image is aligned with the AAM algorithm. Applying the appearance model to face images, the shape-normalized appearance can be obtained. After this, all the images in the training set are normalized, and each image is 100×100 in size.

Figure 7: Training set after processing and normalization

3.2 Decomposition and synthesis

In this paper, the face data is decomposed into multiple orthogonal subspaces by the MMDA algorithm. The face data is composed of multiple attributes such as gender, race and expression. Applying MMDA to each subspace, the orthogonal basis in each mode is calculated and the corresponding parameter is obtained. By changing these parameters, what can be affected is only the corresponding attribute. In order to protect the privacy of face image, we can selectively alter the facial attributes to change some features. At the same time, in order to ensure the integrity of the face, we retain the expression attribute, so as to protect the privacy of the face image, but left facial expression information unchanged.

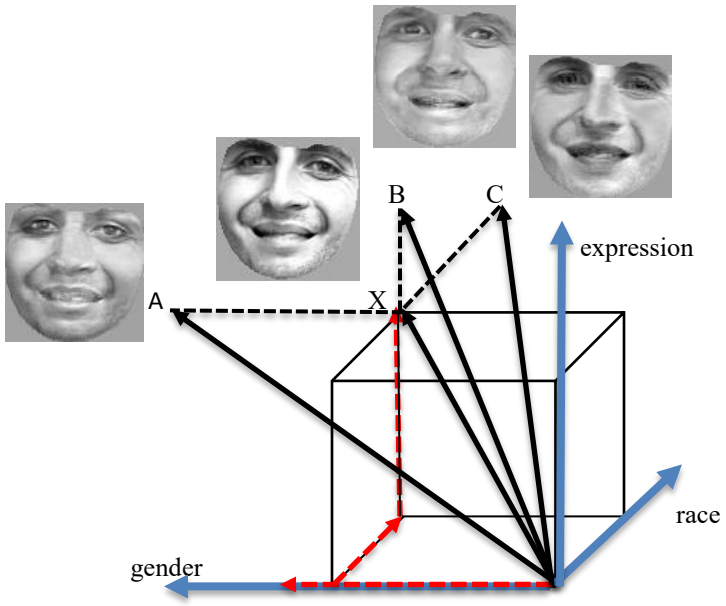

Figure 8: MMDA projects vector *X* into three mutually orthogonal subspaces (represented by the blue axis), encoding gender, race, and expression, respectively

Changing the *X* parameter in the gender subspace will only affect the gender without changing the race or expression. This corresponds to moving *X* to A. Reconstructing Vector A reveals a new face image that only shows a gender change instead of race or expression.

MMDA

Define *X* as a training matrix of $D \times N$, x_i is the training vector corresponding to any normalized image in the training set, and x_i multiply labeled has *K* attributes, $i=1, \cdots, N$, each label represents a attribute such as Male, African, Smile, etc. Each attribute *i* has C_i classes, $L_1^i, L_2^i, \cdots L_{C_i}^i$. For gender, there are two class labels, Male and Female; For race, there are two class labels, African and European; For expression, there are five class labels, Smile, Angry, Sad, Fear and Surprise. Without loss of generality, it is assumed that the overall mean of *X* is 0, if not, the overall mean is subtracted for each eigenvector. Now, the overall mean of the training samples is equal for each attribute.

Whitening

Given an initial set of training samples *X*, the sample set *X* for each attribute feature is whitened in order to reduce the correlation between the features.

We compute the global scatter matrix S_t^i for each attribute *i*,

$$S_t^i = XX^T \tag{1}$$

then Eigen-decompose it to get eigenvector matrix *U* and eigenvalue matrix *D*,

$$S_t^i = UDU^T \tag{2}$$

To reduce the redundancy of sample data, we only retain none-zero eigenvalues in D and corresponding eigenvectors in U. Then the (N-1)× D matrix P can be computed by the matrix D and U,

$$P = UD^{-1/2} \tag{3}$$

then apply it to sample data X to get whitened data $\tilde{X}$, $\tilde{X} = P^T X$, now the scatter matrix of $\tilde{X}$ is the identify matrix I. Since the sample data for each attribute comes from the same training set, only one whitening operation is required.

Building Subspace

The whitening data $\tilde{X}$ is obtained by quantizing the sample set X and removing the correlation. In order to distinguish the different attribute features as much as possible while keeping the same attribute characteristics as close as possible, the subspace for each attribute is established by using the Fish criterion, and each subspace is orthogonal to each other.

For each attribute, we firstly compute the between-class scatter matrix $\tilde{S}_b^i$ and the within-class scatter matrix $\tilde{S}_w^i$, $\tilde{m}_k$ is the whitened class mean,

$$\tilde{S}_b^i = \sum_{k=1}^{C_i} n\tilde{m}_k^i \, \tilde{m}_k^{i\,T} \tag{4}$$

$$\tilde{S}_w^i = \sum_{i=1}^{C} \sum_{\tilde{x}_i \in L_k} (\tilde{x}_i - \tilde{m}_k)(\tilde{x}_i - \tilde{m}_k)^T \tag{5}$$

then maximize the Fisher Criterion,

$$J_F(V^i) = trace\left\{\left((V^i)^T \tilde{S}_w^i (V^i)\right)^{-1}\left((V^i)^T \tilde{S}_b^i (V^i)\right)\right\} \tag{6}$$

Zhang et al. [Zhang and Sim (2006, 2007)] has proven that $J_F(V^i)$ is equal to the ratio $\frac{\lambda_b}{\lambda_w}$, where λ_b and λ_w are the eigenvalues of $\tilde{S}_b^i$, $\tilde{S}_w^i$, respectively. As is shown that, λ_b+λ_w=1, so $J_F(V^i)$ can reach the maximum value by making the λ_b equal to 1, and the subspace V^i can be obtained by keeping the eigenvectors corresponding to λ_b=1 in matrix V^i, the dimension of V^i is C_i-1 and $(V^i)^T \tilde{x}^i = (V^i)^T \tilde{m}_k^i$, $\forall \tilde{x}^i \in L_k^i$. It can be proven that the subspace obtained above is the most discriminant subspace.

Repeat the same step for all attributes to get subspaces $V^1, V^2, \cdots V^k$, the dimension of all subspaces is $\sum_{i=1}^{K} C_i$-K, but the whitened data holds a much larger subspace of dimension N-1=$\prod_{i=1}^{K} C_i - 1$, so there is a residual space V^0 which contains residual variations that do not exist in all the identify space, and residual dimension is $r^0 = N - \sum_i C_i + K - 1$.

To compute V^0, we can apply Gram-Schmidt algorithm. For each vector $\tilde{x}_i \in \tilde{X}$, subtract its projection on the basis of each subspace, and orthogonally normalize the remaining component to get V^0. The complete subspace V contains:

$$V = [V^1, V^2, \cdots V^k, V^0] \tag{7}$$

Each subspace in V is orthogonal to each other, $(V^i)^T V^j = 0, if\ i \neq j$.

Multi-mode decomposition

Assuming that the vector to be decomposed is $\tilde{x}_i \in \tilde{X}$, we can decompose it into each attribute to obtain the parameter vector y by using the trained orthogonal subspace V,

$$y = V^T P^T \tilde{x}_i \tag{8}$$

from vector y, $\{C_1\text{-}1, C_2\text{-}1, \ldots, C_i\text{-}1\}$ denote attributes respectively, that is, the first C_1-1 parameters represent the first attribute feature, the C_1 to C_2-1 parameters to be the second attribute, and the C_i to C_i-1 ones show the i-th attribute. By controlling these, we can selectively change the corresponding facial characteristics.

The training set contains K=3 facial attributes, such as gender, race and expression, where gender includes both male and female labels, race has two labels of African and European, and expression is either smile, angry, sad, fear or surprised.

Decomposing any training data $\tilde{x}_i$ to get the parameter vector y_i:

$$y_i{}^T = [g_i, r_i, e_i{}^T, s_i{}^T] \tag{9}$$

In the vector y_i, parameter of each attribute can be obtained by a certain degree:

1) The gender scalar g_i, can only take on two constant values: G_1 or G_2, representing Male and Female, respectively.

2) The race scalar r_i, can only take on two constants: R_1 or R_2, representing African or European, respectively.

3) The vector e_i, which represents expression, can only take on five constants: E_1, E_2, E_3, E_4, E_5, meaning smile, angry, sad, fear or surprised.

4) The remaining vector s_i represents residual space.

In fact, these scalars and vectors contain the average characteristics of attributes such as Male, Female, African, etc. Therefore, facial appearance can be changed by changing corresponding parameters. Meanwhile, we can control the intensity of the change, for instance, we can set the expression parameter to σE_3, and we will also get sad appearance of different degrees by varying σ.

Alteration and synthesis

Given a decomposed training data, the synthesis of face data is achieved by whitened matrix P and subspace V,

$$\tilde{x}_i = P_r V y \tag{10}$$

where P_r reverses the whitening and PCA operation of P. We decompose and synthesize the training set, where $Q = P_r V$.

$$\tilde{x} = \underbrace{\begin{bmatrix} | & | & & | \\ q_1 & q_2 & \cdots & q_{19} \\ | & | & & | \end{bmatrix}}_{Q} \underbrace{\begin{bmatrix} y_1 \\ \vdots \\ y_{19} \end{bmatrix}}_{y} \tag{11}$$

In the obtained feature parameters, the first six parameters respectively represent the characteristics of gender, race and expression attributes, so that the corresponding column vectors $q_1, \cdots, q_6$ decode these facial attribute appearance, and the decoding information remaining in the residual space contains other face information in the image.

Thus, we can define these vectors as semantic faces, which form a semantic basis for the face image, and a linear combination of any vectors results in a face that contains different gender, race, and expression appearance. For example, for $G_1 q_1$, an average male face can be obtained by the gender parameter and semantic face q_1, and $G_2 q_1$ can create an average female face. By changing scalar g, such as $3G_1$ or $3G_2$, we can get a more masculine or feminine face.

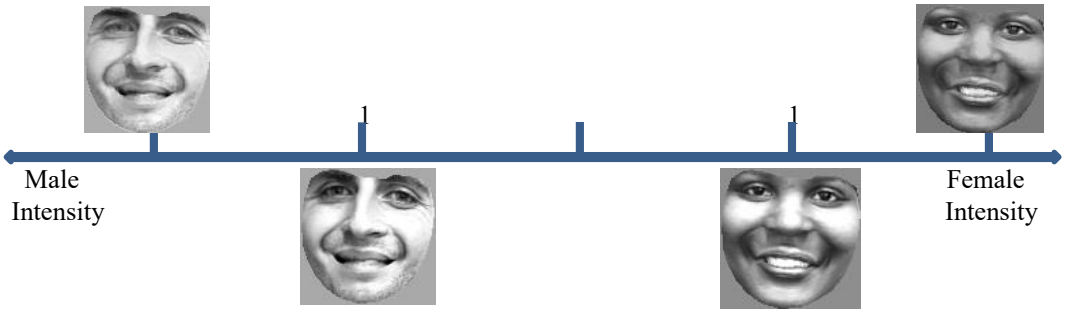

Figure 9: By changing the sole parameter of q_1 (the gender Semantic Face), Male (left) and Female (right) faces are synthesized. Varying this parameter also makes the face appear more or less Masculine (or Feminine)

Face reduction

Finally, by decomposing and synthesizing the training data, a new face image can be obtained. In order to preserve the integrity of the original image, the synthesized face is restored to the original image using AAM algorithm, so that the resulting image contains non-face information of the original image.

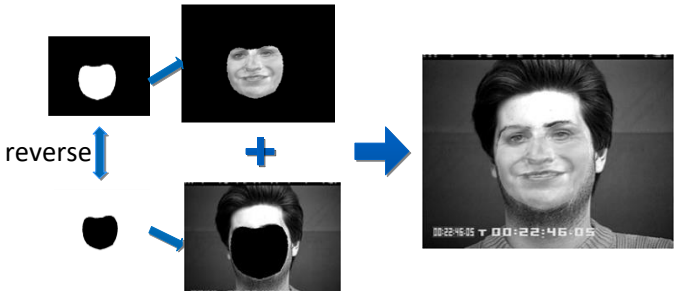

Figure 10: Reduction of the altered face image

4 Privacy protection

Applying MMDA to a face image labeled different classes, we can decompose it into gender, race and expression attributes, then a new synthetic image can be obtained by selectively changing any of the known attribute without affecting other attributes. The synthetic face image can generate new features on attributes such as gender, race and expression, which make the synthetic face image different from the original image on these attribute. Therefore, we can protect the privacy of face images through different

degrees by selectively changing the attribute characteristics.

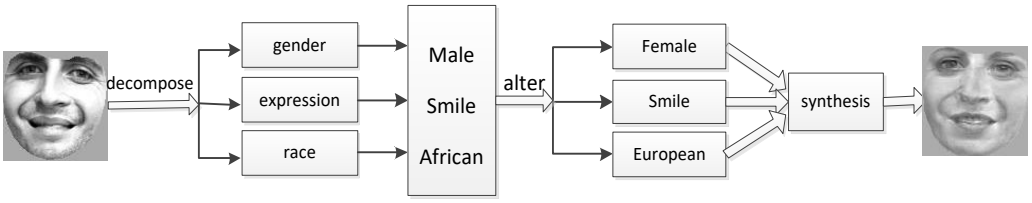

Figure 11: Overview of the proposed method containing decomposition, alteration and Synthesis

In this paper, we change the gender and race attribute to obtain the new gender and race attribute feature, and do not change the expression features, so that the privacy of face images can be guaranteed while data utility is considered.

To protect privacy of the training data, we decompose each image in the training set by MMDA algorithm to get the parameter vector *Y* related to the attribute feature,

$$y = \{y_1, y_2, \cdots, y_{20}\} \tag{12}$$

where $y_i{}^T = [g_i, r_i, e_i{}^T, s_i{}^T]$.

For a face image, facial gender and race feature will be transformed. The first and second parameter g_i, r_i in y_i represent the gender and race attribute respectively. By changing g_i and r_i in each y_i to the opposite of the current value, such as $g_i = G_1$, $r_i = R_1$ while transforming it to $g_i = \sigma G_2$, $r_i = \sigma R_2$ where σ is a constant greater than 0, the protection of facial gender and race attribute is archived. While changing the gender and the race attribute, the parameter vector e_i in y_i is kept constant so that the retention of the facial expression attribute is realized.

Hence, we can get $\bar{y}_i$ whose other face attribute has changed except that the expression attribute is invariant by changing the parameter vector y_i selectively, and applying synthesis and reduction to the reconstructed $\bar{y}_i$ to get the final face image whose expression appearance is unchanged while the face of other attributes has been protected.

$$\tilde{x} = \begin{bmatrix} | & | & & | \\ q_1 & q_2 & \cdots & q_{19} \\ | & | & & | \end{bmatrix} \tag{13}$$

5 Experiments

5.1 Evaluation criteria

Using the change detectors (CDs) proposed by Sim et al. [Sim and Li (2015)], we can determine whether the relevant attributes of the face after the synthesis have been changed. The change detector is based on the Face++ attribute classifier, and accepts two inputs, the original image and the new synthetic image. When it judges that two images differ in the relevant attribute, the output is "Changed", otherwise it outputs "Unchanged". We build three CDs, one each for gender, race and expression.

Let β be the probability that a facial attribute is correctly changed, t_p is the true positive rate which is the probability that the attribute in the input image is changed when

the CD outputs "Changed", while f_p is the false positive probability that the attribute is not changed in the input image but the CD outputs "Changed". α is the ratio of the observed changed times to the times the CD outputting "Changed". From these definitions, we can see:

$$\beta = \frac{\alpha - f_p}{t_p - f_p} \tag{14}$$

The probability β can be observed to analyze effectiveness of our method when changing the several attributes of the face image.

5.2 Experiments and analysis

Single attribute change

The gender, race and expression attributes are selected without changing other attributes feature respectively, and each attribute is changed with varying intensity σ from 0.5 to 2.5 as the interval is 0.5. All three change detectors are used to measure α, then the β is calculated.

Table 1: Actual "changed" rate ('β') values for single-attribute change. These show that our method is effective in changing an attribute (bold values) and retaining an attribute (values in normal font)

	Intensity σ	0.5	1.0	1.5	2.0	2.5
Gender Change	Gender CD	**0.554**	**0.688**	**0.743**	**0.847**	**0.996**
	Race CD	0.401	0.357	0.379	0.364	0.383
	Expression CD	0.568	0.232	0.378	0.371	0.245
Race Change	Gender CD	0.375	0.358	0.473	0.375	0.352
	Race CD	**0.589**	**0.692**	**0.751**	**0.802**	**0.924**
	Expression CD	0.218	0.225	0.475	0.694	0.274
Expression Change	Gender CD	0.362	0.374	0.352	0.383	0.213
	Race CD	0.256	0.243	0.146	0.152	0.249
	Expression CD	**0.678**	**0.867**	**0.988**	**0.995**	**0.997**

Table 2: Actual "changed" rate ('β') for multi-attribute change (G, R, E respectively represents Gender, Race and Expression)

	Gender + Race			Gender + Expression			Race + Expression			All 3 attributes		
σ	G	R	E	G	R	E	G	R	E	G	R	E
1	0.476	0.526	0.228	0.546	0.272	0.742	0.356	0.539	0.803	0.368	0.491	0.697
2	0.623	0.619	0.245	0.558	0.246	0.834	0.292	0.738	0.896	0.508	0.725	0.834

We can observe the "Gender Change" rows in the Tab. 1, for example, only when the gender attribute is changed, the gender CD holds a larger β than other CDs at each

intensity level, indicating that the gender attribute is indeed changed, while other face attribute remain unchanged. Looking at the other rows in the table, the conclusion is similar. Therefore, our algorithm shows that it is effective to change facial attributes as well as remain attributes by MMDA.

Multiple attribute change

Changing multiple facial attribute and measuring β for corresponding CDs. Similarly, when observing different β, it can be seen that when the multiple attributes are changed, the unaltered attribute is unaffected, and as the change intensity σ increases, the β of the changed attributes are getting larger, while the β of the remain attribute decreases.

Privacy based on expression unchanged

We select the gender and racial attributes to change with different intensity σ, and make no change to the expression attribute. As is shown in Tab. 2, the expression feature is not affected when other two attributes are changed. The Face++ matcher is used to authenticate the face images with unaltered expression, and the expression feature is classified by the attribute classifier. The results are as follows:

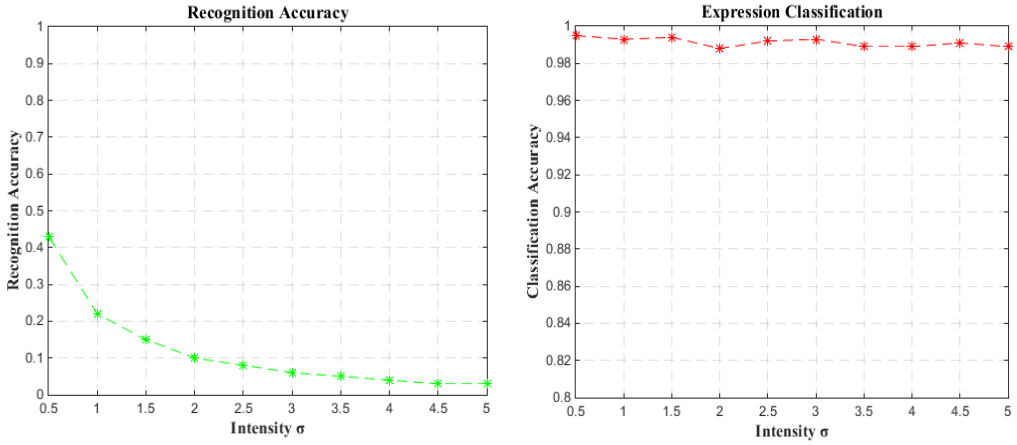

Figure 12: The result shows Recognition Accuracy and Expression Accuracy

From the experiments, it can be seen that the correct recognition rate of face is relatively low when the gender and race attributes are changed at the same time, and the result of expression classification is maintained at a high correct rate. This shows that when changing the gender and race attributes while retaining the same expression, the privacy of face image has been protected well and the expression appearance still remain unaffected, so the purpose of privacy and the face of non-privacy information retained can be achieved.

6 Conclusions

In this paper, we mainly solve the problem of protecting the facial privacy while preserving the expression information in the image. Through the orthogonal decomposition of the multi-attribute face image, we establish the independent subspace of each attribute, get

the corresponding parameter, and selectively change the parameters of other attributes in addition to the expression to retain only expression appearance the same and gender and race attributes different from the original face. Concluded from the experiment results, our method has a good effect on the decomposition of face attributes, and performs well on the privacy protection of face images by changing different attribute features while expression feature is preserved.

Acknowledgements: This work was supported by the National Key Research and Development Program of China (No. 2016YFB0800601), the Key Basic Research Plan in Shaanxi Province (Grant No. 2017ZDXM-GY-014) and the Key Program of NSFC-Tongyong Union Foundation under Grant U1636209.

References

Berger, A. M. (2000): Privacy mode for acquisition cameras and camcorders. *US 6067399 A*.

Cootes, T. F.; Edwards, G. J.; Taylor, C. J. (1998): Active appearance models. *European Conference on Computer Vision*, vol. 23, pp. 484-498.

Dryden, I. L.; Mardia, K. V. (1998): *Statistical Shape Analysis*. vol. 213, no. 6, pp. 663-669. John Wiley & Sons Ltd.

Du, L.; Yi, M.; Blasch, E.; Ling, H. (2014): GARP-face: Balancing privacy protection and utility preservation in face de-identification. *IEEE International Joint Conference on Biometrics*, pp. 1-8.

Gross, R.; Airoldi, E.; Malin, B.; Sweeney, L. (2006): Integrating utility into face de-identification. *Lecture Notes in Computer Science*, vol. 3856, no. 10, pp. 227-242.

Gross, R.; Sweeney, L.; Fernando, D. L. T.; Baker, S. (2006): Model-based face de-identificat*ion. Conference on Computer Vision and Pattern Recognition Workshop*, vol. 47, no. 10, pp. 161.

Gross, R.; Sweeney, L.; Torre, F. D. L.; Baker, S. (2008): Semi-supervised learning of multi-factor models for face de-identification. *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1-8.

Hudson, S. E.; Smith, I. (1996): Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. *ACM Conference on Computer Supported Cooperative Work*, vol. 26, pp. 248-257.

Kirby, M.; Sirovich, L. (2002): Application of the karhunen-loeve procedure for the characterization of human faces. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, vol. 12, no. 1, pp. 103-108.

Lucey, P.; Cohn, J. F.; Kanade, T.; Saragih, J. (2010): The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. *Computer Vision and Pattern Recognition Workshops*, vol. 36, pp. 94-101.

Neustaedter, C.; Greenberg, S. (2003): The design of a context-aware home media space for balancing privacy and awareness. *International Conference on Ubiquitous Computing*, pp. 297-314.

Neustaedter, C.; Greenberg, S.; Boyle, M. (2006): Blur filtration fails to preserve privacy for home-based video conferencing. *ACM Transactions on Computer-Human Interaction*, vol. 13, no. 1, pp. 1-36.

Newton, E. M.; Sweeney, L.; Malin, B. (2005): Preserving privacy by de-identifying face images. *IEEE Transactions on Knowledge & Data Engineering*, vol. 17, no. 2, pp. 232-243.

Othman, A.; Ross, A. (2014): Privacy of facial soft biometrics: Suppressing gender but retaining identity. *European Conference on Computer Vision*, vol. 8926, pp. 682-696.

Sim, T.; Li, Z. (2015): Controllable face privacy. *IEEE International Conference and Workshops on Automatic Face and Gesture Recognition*, vol. 4, pp. 1.

Sim, T.; Zhang, S.; Li, J.; Chen, Y. (2009): Simultaneous and orthogonal decomposition of data using multimodal discriminant analysis. *IEEE 12th International Conference on Computer Vision*, vol. 30, pp. 452-459.

Zhang, S.; Sim, T. (2006): When fisher meets fukunaga-koontz: A new look at linear discriminants. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 1, pp. 323-329.

Zhang, S.; Sim, T. (2007): Discriminant subspace analysis: A fukunaga-koontz approach. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, vol. 29, no. 10, pp. 1732-1745.