

QUALITY PRESERVING FACE DE-IDENTIFICATION AGAINST DEEP CNNs

Panteleimon Chriskos^{}, Rosen Zhelev[†], Vasileios Mygdalis^{*†}, and Ioannis Pitas^{*†}*

^{*}Department of Informatics, Aristotle University of Thessaloniki, Thessaloniki, 54124, Greece

[†]Department of Electrical and Electronic Engineering, University of Bristol, United Kingdom

Email: {tefas,pitas}@aia.csd.auth.gr

ABSTRACT

In this paper, two face de-identification methods are proposed regarding face identification hindering against a deep neural network. Our work focuses on achieving a delicate balance, so that the facial images are miss-classified by the deep network, while the human observer can still identify the persons depicted in a scene. The proposed methods are based on achieving face de-identification by partly degrading image quality in order to hinder face recognition from deep neural networks, while maintaining the highest possible image quality, at the same time. To this end, we employ de-identification methods based on singular value decomposition and image hypersphere projections, respectively. From the conducted experiments, it can be concluded that these methods are capable of reducing correct face identification rates of the VGG-face network by over 90 %. Moreover, it is shown that these error rates preserve adequate image quality as is demonstrated through the values of the complex wavelet structural similarity index, allowing face recognition by humans contrary to most face de-identification methods.

Index Terms— De-identification, SVD-DID, PDID-M

1. INTRODUCTION

In the growing online community an ever more increasing amount of visual data is shared, viewed and stored on-line. This trend provides immediate communication and free idea distribution but also raises privacy concerns to all those involved. Through video sharing sites and social media, it is possible for a malicious user to collect facial images of specific individuals in order to train face classifiers and then monitor the activities of these individuals. Apart from video sharing, systems monitoring the World Wide Web may use face detection [1], tracking [2] and face recognition algorithms [3] [4] [5] in shared videos or images and can subsequently be

used to violate user privacy. Combined with the wide use of video surveillance in public places, and other services like Google Street View and EverySpace [6], any person can potentially be identified, infringing his privacy in most cases unintentionally. Due to the large threats posed to privacy, face de-identification methods must be developed [7] that preferably maintain a level of facial image quality rendering quality of the de-identified medium acceptable.

A large number of methods have been devised that provide face de-identification against common classification methods. However the vast majority of these methods achieve de-identification against both automatic face recognition methods and human viewers, typically destroying a significant portion of the facial image data. Perhaps the most advanced approach is to exploit generative adversarial networks [8], attempting to generate synthetic samples from the distribution of all possible images that generated query segmentations. Besides face de-identification, this method can be extended for full body de-identification in person images by also removing soft biometric and non-biometric identifiers. Another example are methods [9] [10] which de-identify not only the facial image but the entire person Region of Interest (ROI). Simple yet effective facial image de-identification methods apply black masks on parts of the face such as black bars to cover the eyes or T-shaped masks covering both eyes and nose. Other mask shapes and size can be used that cover larger face areas or the whole of the face ROI thus destroying all visual face information. Other ad-hoc methods apply a low-pass filter on the facial image ROI, add random noise, swap facial image sub regions from different individuals [11], spatially subsample a facial image, or threshold the facial image pixels. In cases that some facial information is retained, such as facial expression, the authors in [12] use variational adaptive filtering along with face key point detection, and in [13] Active Appearance Models are utilized.

A prominent family of face de-identification algorithms are based on the k -anonymity model proposed in [14]. In this model, any of the de-identified images is misclassified as at least k of original facial images. The de-identified image is calculated by averaging the k facial images that are

This work has received funding from the European Union's European Union Seventh Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE). This publication reflects only the authors' views. The European Commission is not responsible for any use that may be made of the information it contains.

closer to the input image. In [15] an objective function is formulated in order to calculate the optimal weights for fusing the k most similar images through gradient descent. In similar fashion in [16] the k least similar images to compose the de-identified image a model known as the k -Same-furthest model. An extension of this method is proposed in [17] where the de-identified facial image is unique for each one of the k original faces. Another approach in face de-identification includes replacing faces with 3D morphable facial models [18] or replacing the face depicted with a face from another person [19]. In this scope the GARP-Face framework [20] aims to balance the utility of an image and face de-identification by preserving facial attributes such as gender, age and race. Finally, in [21] face de-identification is achieved by reducing the number of eigenfaces used for reconstructing the facial images from basis facial images.

Most of the methods mentioned above the aim to completely hide the identity the individual in an image, hindering identification from both human viewers and face identification methods. The two methods evaluated in this paper however reduce the face identification accuracy of a deep neural network, while retaining adequate visual face information rendering the de-identified facial images acceptable and recognizable by humans. The first method is based on singular value decomposition (SVD) achieving face de-identification by manipulating the SVD coefficients of the initial image and reconstructing the de-identified one. The second method calculated the de-identified image by projecting the original on a hypersphere centered on a mean image.

The rest of this paper is organized in four sections, with Section 2 providing a description of the two methods, Section presenting the image quality measured use to quantify image quality reduction and Section 3 where we describe the experimental setup and results. Final conclusions are drawn in Section 4.

2. DE-IDENTIFICATION AND IMAGE QUALITY

The proposed de-identification methods [22] based on hypersphere projection and SVD for hindering the classification accuracy of deep neural networks, are described in Subsections 2.1 and 2.2, respectively. Finally,

2.1. Hypersphere projections

A hypersphere is a generalization of the ordinary circle and sphere to dimensions $n \geq 3$. For any number n of dimensions a hypersphere $S^{(n-1)}$ is the set of points which are at distance R from a center point in n -dimensional space. Such a hypersphere has a radius of R . The definition of a hypersphere S^{n-1} with a center at some origin is:

$$S^{n-1} = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| = R\} \quad (1)$$

where \mathbf{x} is a point in the n -dimensional space. The projection of a point $\mathbf{x} \in \mathbb{R}^n$ onto S^{n-1} is defined as follows:

$$P_{S^{n-1}}(\mathbf{x}) = \frac{R}{\|\mathbf{x}\|} \mathbf{x} \quad (2)$$

where $P_{S^{n-1}}(\mathbf{x})$ is the projection of the point \mathbf{x} onto the hypersphere S^{n-1} .

Projection De-Identification on Mean Image (PDID-M) uses a projection on a hypersphere centered on the mean image $\bar{\mathbf{I}}$ given by:

$$\bar{\mathbf{I}} = \frac{1}{N} \sum_{i=1}^N \mathbf{I}_i \quad (3)$$

where N is the number of facial images in a given dataset. Given an initial image \mathbf{I} the de-identified image through PDID-M is calculated as:

$$\mathbf{I}_{\text{PDID-M}} = \left(\frac{R * (\mathbf{I} - \bar{\mathbf{I}})}{\|\mathbf{I} - \bar{\mathbf{I}}\|} + \bar{\mathbf{I}} \right) \quad (4)$$

where $\bar{\mathbf{I}}$ is the mean image, R is the radius of the hypersphere and $\|\cdot\|$ denotes measure.

2.2. Singular Value Decomposition

Singular Value Decomposition factorizes a matrix $\mathbf{I} \in \mathbb{R}^{N \times M}$, in this case \mathbf{I} is a facial image, as a product of three matrices, namely matrix $\mathbf{S} \in \mathbb{R}^{N \times M}$, which contains the sorted, from largest to lowest, singular values which equal to the square roots of the eigenvalues of matrix $\mathbf{I}\mathbf{I}^T$. The other two matrices $\mathbf{U} \in \mathbb{R}^{N \times M}$ and $\mathbf{V} \in \mathbb{R}^{M \times M}$ contain the singular vectors, which are the eigenvectors of matrices $\mathbf{I}\mathbf{I}^T$ and $\mathbf{I}^T\mathbf{I}$ respectively. Through SVD image \mathbf{I} can be decomposed as:

$$\mathbf{I} = \mathbf{U}\mathbf{S}\mathbf{V}^T. \quad (5)$$

In all above cases \mathbf{A}^T denotes the transpose of matrix \mathbf{A} . Based on the above decomposition, the SVD face de-identification method (SVD-DID), modifies the input image by altering the entries of matrices \mathbf{U} , \mathbf{S} and \mathbf{V} and reconstructing the de-identified image through the altered matrices. The SVD-DID method can be broken down to three distinct steps described below.

2.2.1. SVD Coefficient Zeroing (SVD-CZ)

The first step of the SVD-DID method the first N_Z singular values of \mathbf{S} ($N_Z \leq N \leq M$) are zeroed resulting in a new \mathbf{S} matrix \mathbf{S}_{CZ} . This step tends to darken the output image, compared to the initial one and to counterbalance this effect the facial image pixel luminance values are increased at the end of the de-identification process by adding a fixed luminance value to the output facial image pixels. In the rest of the paper we assume an added luminosity value equal to 100.

2.2.2. SVD Coefficient Averaging (SVD-CA)

In the second step of this method, the values in eigenvector matrices \mathbf{U} , \mathbf{V} are low-pass filtered. This is achieved with the use of a $m \times m$ circular averaging filter, where $m = 2R + 1$ [23], and R is the radius of the circular filter. Through this filtering process matrices \mathbf{U}_{AV} and \mathbf{V}_{AV} are produced. However, reconstructing the de-identified facial image solely through the above matrices, leads to poor image quality. In order to preserve adequate image quality matrices \mathbf{U}_{AV} and \mathbf{V}_{AV} are blended with the original \mathbf{U} and \mathbf{V} matrices through weighted averaging as:

$$\mathbf{U}_{CA} = \frac{\alpha * \mathbf{U}_{AV} + \mathbf{U}}{1 + \alpha} \text{ and } \mathbf{V}_{CA} = \frac{\alpha * \mathbf{V}_{AV} + \mathbf{V}}{1 + \alpha} \quad (6)$$

where parameter α adjusts the trade-off between visual quality and face de-identification.

2.2.3. SVD Modified Sobel Filtering (SVD-MSF)

In the last step of the SVD-DID method final step matrices \mathbf{U}_{CA} and \mathbf{V}_{CA} are high pass filtered [23] using a modified Sobel filter. This filter has a 3×3 matrix form:

$$\mathbf{G} = \begin{bmatrix} d & 2d & d \\ 0 & 0 & 0 \\ -d & -2d & -d \end{bmatrix} \quad (7)$$

where parameter d specifies the intensity of the high pass filtering. As in the previous step the resulting matrices are blended with the original ones producing matrices \mathbf{U}_F and \mathbf{V}_F .

After the three steps of the SVD-DID method the output facial image $\mathbf{I}_{\text{SVD-DID}}$ is calculated as:

$$\mathbf{I}_{\text{SVD-DID}} = \mathbf{U}_F \mathbf{S}_{CZ} \mathbf{V}_F^T. \quad (8)$$

2.3. Image Quality Measure

Both de-identification methods introduce artifacts and noise in the de-identified image. This is necessary to fool automatic face recognition methods, but quality preservation is also essential to render the final image acceptable for human viewers. As a result, the need arises to quantify the effect of each of the two methods on image quality. In order to quantify the degree of image quality loss due to the application of the above face de-identification methods, the complex wavelet structural similarity index, CW-SSIM [24] was employed. This index is an extension of the structural similarity index, SSIM, proposed in [25] in the complex wavelet domain. The continuous wavelet transform of a signal $\mathbf{s}(t)$ with scaling $a \in \mathbb{R}^{+*}$ and translation $b \in \mathbb{R}$ is defined as:

$$S_w(a, b) = \frac{1}{\sqrt{|a|}} \int_{-\infty}^{\infty} \mathbf{s}(t) \bar{\psi} \left(\frac{t-b}{a} \right) dt \quad (9)$$

Table 1: Failure percentages after applying Gaussian filtering

σ	F_p	c_s
1	3.3 %	0.9971
2	3.26 %	0.9644
3	3.36 %	0.8618
4	3.56 %	0.6912
5	4.86 %	0.5019
6	14.64%	0.3397
7	47.38%	0.2225
8	84.56 %	0.1478
9	96.96 %	0.1048
10	99.26 %	0.0824

where $\psi(t)$ is a continuous function in the time and frequency domain commonly known as the mother wavelet and $\bar{\cdot}$ denotes the conjugate complex number. Through the wavelet transform it is possible to calculate the wavelet coefficients \mathbf{c} that are used to calculate the CW-SSIM index which is defined as:

$$\tilde{S}(\mathbf{c}_x, \mathbf{c}_y) = \frac{2|\sum_{i=1}^N \mathbf{c}_{x,i} \bar{\mathbf{c}}_{y,i}| + K}{\sum_{i=1}^N |\mathbf{c}_{x,i}|^2 + \sum_{i=1}^N |\mathbf{c}_{y,i}|^2 + K} \quad (10)$$

where $\mathbf{c}_x = \{c_{x,i} | i = 1, \dots, N\}$ and $\mathbf{c}_y = \{c_{y,i} | i = 1, \dots, N\}$ are the wavelet coefficients calculated for two given image regions \mathbf{x} and \mathbf{y} at the same spatial location in the two images and K is a small positive constant that improves robustness in cases of low signal to noise ratios. The CW-SSIM index takes values in the range of $[0, 1]$, equal to 1 when one image is compared to itself. As such it is preferable to have CW-SSIM values close to one after the application of the face de-identification methods.

3. EXPERIMENTS AND RESULTS

Experiments were conducted to assess the effectiveness of the SVD-DID and PDID-M face de-identification methods against the VGG deep neural network classifier [26]. To this end, a dataset of 5000 images depicting 1000 randomly selected different individuals, 5 images per individual, from the database was employed [27]. In RGB images, the two methods were applied separately on each color channel. The experiments were conducted using the MatConvNet [28] MATLAB toolbox.

To quantify the ability of the methods to hinder automatic face identification the false face identification percentage F_p was used, which is the number of missclassified images over the total number of images. In order to quantify the quality reduction after the de-identification process the mean CW-SSIM index, c_s , is used, averaged over the total number of images. As mentioned above it is preferable for CW-SSIM to have a value close to 1 so that little visual information is lost.

For comparison reasons, we have also employed a naive de-identification method. That is, performing Gaussian filter




Fig. 1: Facial images after Gaussian filtering

on the facial images for different values of σ . Experimental results for the naive de-identification method are show in Table 1. As can be seen, values of $\sigma > 9$ are required in order to achieve de-identification rates higher than 90%, resulting in images that have low quality, and make it difficult to be recognized by humans, as well. This property is also denoted by the corresponding CW-SSIM indices, which are as low as 0.10 for such values of σ .

Identification failure rates and CW-SSIM index values for applying the PDID-M method are displayed in Table 2, for different values of radius R . For large radius values the failure percentages are quite low the lowest being 13.78% for $R = 30$ and similarity index values are quite high the highest being for the aforementioned radius value equal to 0.8565. Reducing R to 20 increases the failure rate more than three times the previous value up to 43.98% and at the same time c_s drops to 0.6802. It can be easily observed that increasing the value of R leads to higher failure rates and a drop in mean image similarity. For a $R = 13$ an F_p equal to 80.02% is achieved with CW-SSIM being 0.4859. A failure rate over 90% is achieved for $R = 10$ and a corresponding similarity index value equal to 0.9379. Further reducing the radius value leads to even higher failure percentages with a maximum $F_p = 99.17\%$ and a minimum $c_s = 0.2639$ for $R = 6$.

Results after applying the SVD-DID method are tabulated in Table 3. From this table it is evident that the main parameter that affects the failure percentages is the number of zeroed singular values N_Z . For parameters $\alpha = 0.2$, $d = 0.1$ and $R = 5$ varying N_Z leads to significantly higher error rates as for $N_Z = 1$ the failure rate is 43.46%, the lowest one observed, for $N_Z = 2$ its value rises to 64.32% and for $N_Z = 5$ a failure rate of 92.36% is achieved. The similarity indexes also follow this trend, although not to such an extent, beginning from 0.8653, the highest value observed, for $N_Z = 1$ falling to 0.8227 for $N_Z = 2$ and finally reaching 0.6835 for $N_Z = 5$. It can also be observed that parameter α mainly affects image quality, the value of c_s , as can be derived from the

Table 2: Failure percentages after applying PDID-M

R	F_p	c_s
30	13.78 %	0.8565
20	43.98 %	0.6802
19	48.36 %	0.6495
18	53.06 %	0.6248
17	58.92 %	0.5990
16	64.64 %	0.5722
15	69.72 %	0.5515
14	74.96 %	0.5157
13	80.02 %	0.4859
12	84.14 %	0.4552
11	88.28 %	0.4238
10	91.64 %	0.3979
9	94.44 %	0.3649
8	96.44 %	0.3315
7	98.26 %	0.2977
6	99.14 %	0.2639




Fig. 2: Facial images after performing the PDID-M method

values observed in the table. For parameter values $N_Z = 1$, $d = 0.1$ and $R = 5$ varying α from 0.2 to 0.5 leads to a decrease of c_s from 0.8635 to 0.6875 respectively. The same can be observed for different values of N_Z , where for $N_Z = 2$ and $\alpha = 0.2$ and 0.5, c_s varies from 0.8226 to 0.6515 and similar observations can be made for $N_Z = 5$. Despite the significant drop in image quality, a significant increase is not observed for failure percentages where for $N_Z = 1$ $F_p = 43.46\%$ for $\alpha = 0.2$ and $F_p = 42.68\%$ for $\alpha = 0.5$. Similar results are obtained for different values of N_Z . Parameters d and R do not play a significant role in failure percentages and image quality. The highest failure percentage achieved with the SVD-DID method is equal to 92.64% for parameter values $N_Z = 5$, $\alpha = 0.2$, $d = 0.5$ and $R = 5$ with a CW-SSIM index equal to 0.6830. The quality of the resulting facial images from applying the Gaussian filtering, PDID-M and SVD-DID

methods is shown in Figures 1, Figure 2 and Figure 3, respectively.




Fig. 3: Facial images after performing the SVD-DID method. The same parameters were applied to all images i.e., $a = 0.2$, $d = 0.5$, $R = 5$, having variable N_z , except (c) where $R = 10$.

4. CONCLUSIONS

From the experimental results presented in the previous section it can be concluded that both the PDID-M and SVD-DID face de-identification methods provide adequate failure percentages against the VGG deep convolutional neural network. In both cases high failure percentages are attained, over 90%, while preserving acceptable image quality. The highest failure percentage observed through the PDID-M method is equal to 99.14% with $c_z = 0.2639$ and for the SVD-DID method these values are $F_p = 92.64\%$ and $c_z = 0.6830$. It must be noted that while the PDID-M methods achieves very high failure percentages this comes with heavy image quality reduction. For similar F_p values 94.44% for PDID-M and 92.64% for SVD-DID the respective c_s values are equal to 0.3649 and 0.6830. It is evident that the SVD-DID method preserves more visual information compared to the PDID-M method, making more appropriate in cases that require the de-identified image to be of higher quality.

Future work in this area will focus on developing reversible and less visible face de-identification methods that will provide privacy against deep classifiers.

5. REFERENCES

[1] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun, "Faster r-cnn: towards real-time object detection with region proposal networks," *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 6, pp. 1137–1149, 2017.

Table 3: Failure percentages after applying SVD-DID

N_z	α	d	R	F_p	c_s
1	0.2	0.1	5	43.46 %	0.8653
1	0.2	0.1	10	43.42 %	0.8654
1	0.2	0.5	5	43.94 %	0.8638
1	0.2	0.5	10	43.50 %	0.8652
1	0.5	0.1	5	42.68 %	0.6875
1	0.5	0.1	10	42.82 %	0.6876
1	0.5	0.5	5	44.84 %	0.6849
1	0.5	0.5	10	43.06 %	0.6871
2	0.2	0.1	5	64.32 %	0.8226
2	0.2	0.1	10	64.26 %	0.8227
2	0.2	0.5	5	64.34 %	0.8214
2	0.2	0.5	10	64.30 %	0.8225
2	0.5	0.1	5	63.50 %	0.6514
2	0.5	0.1	10	63.26 %	0.6515
2	0.5	0.5	5	65.18 %	0.6495
2	0.5	0.5	10	63.66 %	0.6512
5	0.2	0.1	5	92.34 %	0.6835
5	0.2	0.1	10	92.36 %	0.6836
5	0.2	0.5	5	92.64 %	0.6830
5	0.2	0.5	10	92.40 %	0.6835
5	0.5	0.1	5	91.76 %	0.5339
5	0.5	0.1	10	91.78 %	0.5339
5	0.5	0.5	5	92.38 %	0.5335
5	0.5	0.5	10	91.92 %	0.5338

[2] João F Henriques, Rui Caseiro, Pedro Martins, and Jorge Batista, "High-speed tracking with kernelized correlation filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 3, pp. 583–596, 2015.

[3] Vasileios Mygdalis, Alexandros Iosifidis, Anastasios Tefas, and Ioannis Pitas, "Graph embedded one-class classifiers for media data classification," *Pattern Recognition*, vol. 60, pp. 585–595, 2016.

[4] Alexandros Iosifidis, Anastasios Tefas, and Ioannis Pitas, "Person identification from actions based on dynemes and discriminant learning," in *Biometrics and Forensics (IWBF), 2013 International Workshop on*. IEEE, 2013, pp. 1–4.

[5] Vasileios Mygdalis, Iosifidis Alexandros, Anastasios Tefas, and Ioannis Pitas, "Large-scale classification by an approximate least squares one-class support vector machine ensemble," in *Trustcom/BigDataSE/ISPA, 2015 IEEE*. IEEE, 2015, vol. 2, pp. 6–10.

[6] Ralph Gross, Latanya Sweeney, Jeffrey Cohn, Fernando De la Torre, and Simon Baker, "Face de-identification," in *Protecting Privacy in Video Surveillance*, pp. 129–146. Springer, 2009.

- [7] Ralph Gross, Edoardo Airoldi, Bradley Malin, and Latanya Sweeney, "Integrating utility into face de-identification," in *International Workshop on Privacy Enhancing Technologies*. Springer, 2005, pp. 227–242.
- [8] Ivan Sikiric, Tomislav Hrkac, Karla Zoran Kalafatic, et al., "I know that person: Generative full body and face de-identification of people in images," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 15–24.
- [9] Prachi Agrawal and PJ Narayanan, "Person de-identification in videos," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 21, no. 3, pp. 299–310, 2011.
- [10] Suriyon Tansuriyavong and Shin-ichi Hanaki, "Privacy protection by concealing persons in circumstantial video image," in *Proceedings of the 2001 workshop on Perceptive user interfaces*. ACM, 2001, pp. 1–4.
- [11] Saleh Mosaddegh, Loic Simon, and Frédéric Jurie, "Photorealistic face de-identification by aggregating donors face components," in *Asian Conference on Computer Vision*. Springer, 2014, pp. 159–174.
- [12] Geoffrey Letournel, Aurélie Bugeau, V-T Ta, and J-P Domenger, "Face de-identification with expressions preservation," in *Image Processing (ICIP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 4366–4370.
- [13] Lily Meng, Zongji Sun, Aladdin Ariyaeinia, and Ken L Bennett, "Retaining expressions on de-identified faces," in *Information and Communication Technology, Electronics and Microelectronics (MIPRO), 2014 37th International Convention on*. IEEE, 2014, pp. 1252–1257.
- [14] Latanya Sweeney, "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, 2002.
- [15] Amin Jourabloo, Xi Yin, and Xiaoming Liu, "Attribute preserved face de-identification," in *2015 International Conference on Biometrics (ICB)*. IEEE, 2015, pp. 278–285.
- [16] Lily Meng and Zongji Sun, "Face de-identification with perfect privacy protection," in *Information and Communication Technology, Electronics and Microelectronics (MIPRO), 2014 37th International Convention on*. IEEE, 2014, pp. 1234–1239.
- [17] Zongji Sun, Li Meng, and Aladdin Ariyaeinia, "Distinguishable de-identified faces," in *Automatic Face and Gesture Recognition (FG), 2015 11th IEEE International Conference and Workshops on*. IEEE, 2015, vol. 4, pp. 1–6.
- [18] Volker Blanz, Sami Romdhani, and Thomas Vetter, "Face identification across different poses and illuminations with a 3d morphable model," in *Automatic Face and Gesture Recognition, 2002. Proceedings. Fifth IEEE International Conference on*. IEEE, 2002, pp. 192–197.
- [19] Dmitri Bitouk, Neeraj Kumar, Samreen Dhillon, Peter Belhumeur, and Shree K Nayar, "Face swapping: automatically replacing faces in photographs," *ACM Transactions on Graphics (TOG)*, vol. 27, no. 3, pp. 39, 2008.
- [20] Liang Du, Meng Yi, Erik Blasch, and Haibin Ling, "Garp-face: Balancing privacy protection and utility preservation in face de-identification," in *Biometrics (IJCB), 2014 IEEE International Joint Conference on*. IEEE, 2014, pp. 1–8.
- [21] P Jonathon Phillips, "Privacy operating characteristic for privacy protection in surveillance applications," in *International Conference on Audio-and Video-Based Biometric Person Authentication*. Springer, 2005, pp. 869–878.
- [22] Panteleimon Chriskos, Olga Zoidi, Anastasios Tefas, and Ioannis Pitas, "De-identifying facial images using singular value decomposition and projections," *Multimedia Tools and Applications*, vol. 76, no. 3, pp. 3435–3468, 2017.
- [23] Ioannis Pitas, *Digital image processing algorithms and applications*, John Wiley & Sons, 2000.
- [24] Mehul P Sampat, Zhou Wang, Shalini Gupta, Alan Conrad Bovik, and Mia K Markey, "Complex wavelet structural similarity: A new image similarity index," *IEEE transactions on image processing*, vol. 18, no. 11, pp. 2385–2401, 2009.
- [25] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [26] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [27] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, et al., "Deep face recognition.," in *BMVC*, 2015, vol. 1, p. 6.
- [28] Andrea Vedaldi and Karel Lenc, "Matconvnet: Convolutional neural networks for matlab," in *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 2015, pp. 689–692.