

Facial Identity Anonymization via Intrinsic and Extrinsic Attention Distraction

Zhenzhong Kuang¹ Xiaochen Yang¹ Yingjie Shen¹ Chao Hu¹ Jun Yu^{* 1, 2}

¹ School of Computer Science, Hangzhou Dianzi University, Hangzhou, China

² Department of Computer Science and Technology, Harbin Institute of Technology (Shenzhen), Shenzhen, China

{zzkuang, xcyang, shenyongjie, chaos, yujun}@hdu.edu.cn

Abstract

The unprecedented capture and application of face images raise increasing concerns on anonymization to fight against privacy disclosure. Most existing methods may suffer from the problem of excessive change of the identity-independent information or insufficient identity protection. In this paper, we present a new face anonymization approach by distracting the intrinsic and extrinsic identity attentions. On the one hand, we anonymize the identity information in the feature space by distracting the intrinsic identity attention. On the other, we anonymize the visual clues (i.e. appearance and geometry structure) by distracting the extrinsic identity attention. Our approach allows for flexible and intuitive manipulation of face appearance and geometry structure to produce diverse results, and it can also be used to instruct users to perform personalized anonymization. We conduct extensive experiments on multiple datasets and demonstrate that our approach outperforms state-of-the-art methods.

1. Introduction

By making full use of face images, modern AI technologies have enabled us a more convenient life [9, 24, 39]. However, this may raise a wide social concern on privacy because face images are easy to capture but cannot be easily changed. Although some strict constraints (e.g. laws) were set up in the last few years [37, 56, 63], the privacy disclosure events continue to emerge one after another.

Anonymization has attracted increasing attention, which usually has two basic requirements. The first one is to ensure identity safety by fighting against re-identification. Another is to preserve the data utility, such as image quality, face detectability, expression and user-defined attributes, which may vary under different scenarios. Besides protecting the original identity, we also take identity intrusion into consideration to reduce the risk of bringing troubles for the others. The kind of technology has multiple advantages,




Figure 1. Demonstration of our approach for face anonymization.

such as: (1) prevent unauthorized users, organizations and applications from freely collecting and using personal data; (2) help people to avoid troubles by blocking the relationship disclosure between identity and the other factors, such as location, action, event and so on; (3) maintain the data usability in various applications, like autonomous driving and remote medical system, without worrying about information leakage even if the data were attacked or misused.

Traditional methods (e.g. pixelation and blurring [34, 58]), seem simple and effective for privacy protection, but may easily damage the image content and quality, resulting in poor data reusability (e.g. face may become undetectable). Recently, the generative method (e.g. GAN) show promising performances on supporting realistic face synthesis [13, 20, 23, 44], which makes it possible to improve image quality and utility preservation. On this basis, many excellent anonymization trials were conducted from different viewpoints [1, 3, 6, 18, 19, 25, 29, 37, 43, 49, 50, 55, 65]. However, many of them may suffer from the problem of excessive change of the identity-independent information to ensure anonymity, or insufficient identity protection in order to preserve more data utility. The former may lead to performance drop on utility preservation, and the latter may lead to degraded protection against re-identification or identity intrusion, which would prevent the existing methods from achieving a good privacy-utility (PU) tradeoff.

To address the above problem, we present a new face anonymization approach by exploiting the intrinsic and extrinsic face characteristics for identity attention distraction, where a deep generative model is employed to synthesize

*Jun Yu is the corresponding author

anonymous face images. To enable flexible control on anonymization, we divide the input data into two types, including intrinsic identity feature and extrinsic visual clues. Many works intend to embed additional PU tradeoff constrains in their model, but this may increase the difficulty of model optimization. Differently, we propose to perform data anonymization in ahead and let the deep generative model only focus on synthesizing high utility images. Since attention reflects the intrinsic characteristics of the recognition process [21, 47, 53, 62, 64], we perform identity feature anonymization (IFA) by distracting the attention of the original identity to let the face recognition model make wrong prediction. Since the extrinsic face characteristics may attract human attention for re-identification, we perform visual clue anonymization (VCA) by distracting the identity attention of visual clues (i.e. visual appearance and geometry structure) that may easily lead to privacy disclosure. Figure 1 briefly demonstrates the idea of our approach.

Notice that, by proper modeling, IFA can achieve a low loss of identity independent information for utility preservation. In the meanwhile, VCA can enable users more freedom in producing diverse results without significant damage on data utility. For example, it can support fine-grained adjustment on geometry structure to support more effective anonymization, which was rarely considered previously (e.g. [3, 18, 26, 27, 29, 30, 32, 37, 59, 60]). During the anonymization process, our approach can enable users to easily spot what kind of changes were made by comparing input and output, which can be used to instruct users on how to perform personalized anonymization. In summary, the main contributions of this paper are as follows:

- We propose a new synthetic face anonymization approach from the viewpoint of identity attention distraction by exploiting the intrinsic and extrinsic face characteristics.
- We propose an intrinsic identity attention distraction method for IFA anonymization in the feature space.
- We propose an extrinsic identity attention distraction method for VCA anonymization in the visual space.
- We demonstrate that the proposed approach can achieve state-of-the-art performance with the help of extensive experiments on different public datasets.

2. Related Works

Class Activation Mapping. Interpretability is very important for deep learning based AI systems. Visualization of CNN predictions has received wide attention to interpret deep networks [5, 21, 36]. The most relevant approach is CAM [64] which can highlight class-specific discriminative regions by mapping the predicted class score back to the last convolutional layer of a face classification network. In [47], CAM was generalized to gradient CAM that exhibited excellent ability in providing faithful visual explanations. In [30], CAM was used to locate and change the identity-

independent regions and attributes which were utilized to anonymize face images to fool human instead of machine. But, in our study, the output of gradient CAM is used as an indicator to find and recast the identity feature to fool both human and machine, which can enable us to reduce the information loss during the anonymization process for achieving a better privacy-utility tradeoff.

Face Synthesis. GAN has already been used in face image synthesis by playing an adversarial game between generator and discriminator [20, 23, 41]. In [61], a landmark driven synthesis method was proposed for talking head generation. In [28], MaskGAN was proposed for interactive face image manipulation. In [65], DeepFake was used to perform face swapping to protect medical video data. In [31], FaceShifter was introduced to perform face swapping by focusing on identity transformation, which was further applied in [3] and [55] to support face anonymization.

Face Anonymization. Along with the unprecedented application of face images, face anonymization becomes increasingly important and lots of interesting methods were proposed [3, 6, 14, 18, 27, 30, 32, 33, 35, 37, 40, 42, 50, 55, 59, 60]. [18, 25, 50] relied on inpainting to synthesize anonymous face. [33, 37, 40, 42] adopted attribute editing, classifier or control vector to support face anonymization. [3, 14, 42] studied the reversible face anonymization methods based on password or attribute vector. [32, 35, 55] employed disentanglement or identity perturbation to de-identify face identity. [1, 6, 29, 60] synthesized anonymous faces in the StyleGAN latent space. [30, 59] only focused on fooling human eyes by preserving the original identity. Different from previous works, we present a new solution from the viewpoint of identity attention distraction.

3. The Proposed Approach

In this section, we elaborate on our proposed approach. To alleviate developing complicated generative models, we propose a very simple two-step based anonymization process in Figure 2 (a). Given a face image x , we rely on step1 to preprocess it by using IFA and VCA, and their outputs are used in step2 to synthesize an anonymous face \hat{x} , where step1 is only responsible for anonymization, step2 is only responsible for face image synthesis, and model training only happens in step2. Under this setting, we are able to reduce the difficulty of model design and optimization for processing a complicated privacy-utility tradeoff. Next, we present the details in several sections.

3.1. Identity Feature Anonymization

In this subsection, we present how to perform identity feature anonymization in the feature space. As shown in Figure 2 (b), a pre-trained classification network Ψ (e.g. FaceNet [46]) is employed for identity feature anonymization by conducting intrinsic identity attention distraction at the last

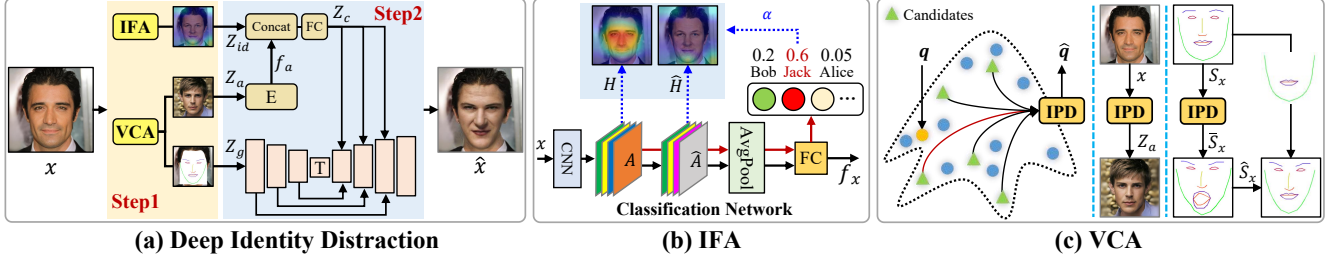


Figure 2. Overview of our approach: (a) the flowchart, (b) identity feature anonymization (IFA), and (c) visual clue anonymization (VCA).

convolutional layer (its output is denoted as A). First, the input data flow goes along the red solid lines to find the identity related feature maps in A on top of the calculated CAM heatmap H [47]. Then, the data flow switches to the black lines after performing attention distraction on A , goes through \hat{A} , and finally output the recasted identity feature f_x for anonymization. \hat{A} is calculated from A by distracting the identity attention away from H , where the visual results of H and \hat{H} in Figure 2 (b) illustrate the function of this operation. Note that the identity of x may be new for Ψ and, thus, the prediction may be incorrect. But, it does not matter. We simply employ the pre-trained classes as the codebook to interpret the identity of any input faces regardless of whether these identities were trained or not. The final softmax output is used as the indicator to determine which pre-trained identities are more related to x .

Let c_i denote the top i -th prediction result (i.e. identity class) of Ψ . The CAM heatmap of any given class c_i can be calculated as the weighted combining of the forward activation maps by following [47]

$$H = \text{ReLU}(\alpha^{c_i} A) = \text{ReLU}\left(\sum_j \alpha_j^{c_i} A^j\right), \quad (1)$$

where $\alpha_j^{c_i}$ denotes the neuron importance weight of the j -th activated feature map A^j , which is calculated by average-pooling the gradients flowing back

$$\alpha_j^{c_i} = \frac{1}{Z} \sum_k \sum_l \frac{\partial y^{c_i}}{\partial A_{kl}^j}. \quad (2)$$

We analyze the importance of facial features on top of H . Motivated by the existing studies that some attributes are critical for identification while the others not [9, 30, 47, 51], we can reasonably suppose that anonymization can be formulated as a min-max optimization problem by suppressing the identity dependent attributes and preserving the identity-independent ones. We model this as an attention distraction problem by using the identity correlated CAM heatmaps and enforcing $H_{c_i} \rightarrow 0$. Then, we have

$$\hat{A} = A + \xi, \quad \text{s.t.} \quad \alpha^{c_i} \hat{A} = 0 \quad (3)$$

where ξ is a modulation item. It is an invalid solution by directly using $\xi = -A$ because the resulting distracted feature map \hat{A} would become meaningless for identity representation. To solve it, we introduce an assistant matrix Φ^{c_i} to distract the attention of the identity correlated CAM heatmaps

and define $\xi = (\alpha^{c_i})^T \Phi^{c_i}$ by taking the importance of the activated feature map into consideration. By substituting ξ back to Equ. (3), we obtain

$$\Phi^{c_i} = -\frac{\alpha^{c_i} A}{\alpha^{c_i} (\alpha^{c_i})^T}. \quad (4)$$

To enhance anonymization, we propose to jointly distract the top K predictions of Ψ because they may be closely related to identity representation. We redefine Equ. (3) as

$$\hat{A} = A + \sum_{i=1}^K w_i (\alpha^{c_i})^T \Phi^{c_i}, \quad (5)$$

where w_i denotes the contribution of the i -th item.

3.2. Visual Clue Anonymization

In this subsection, we present how to perform visual clue anonymization by extrinsic identity attention distraction using both the visual appearance and the geometry structure in the visual space. One may think of directly replace the original data q with a completely (or predefined) different delegate \hat{q} (e.g. from white skin to black skin) so that no one can re-identify it, but this may easily damage data utility (e.g. ethnic and expression preservation [18, 29, 37]). A possible better choice is to sample some \hat{q} based on q so that they share more identity independent information than identity dependent information in a random manner. As shown in Figure 2 (c), we first introduce an instance-level probabilistic delegate (IPD) sampling method and then use it to anonymize the visual clues.

IPD Sampling. Given data q , we build a candidate set M_q by finding its top k -nearest neighbors in the feature space (e.g. Arcface [8]). For each candidate $X_i \in M_q$, we rely on the simple random sampling to obtain a delegate $\hat{q} = X_k \in M_q$ according to the probability set $\{P(X_i), 1 \leq i \leq \text{size}(M_q) | \sum_i P(X_i) = 1\}$, where $P(X_i)$ is defined by following the idea of differential privacy (DP) [10, 11, 38]

$$P(X_i) = \frac{e^{\{\epsilon u(q, X_i)/(2\Delta u)\}}}{\sum_{X_j \in M_q} e^{\{\epsilon u(q, X_j)/(2\Delta u)\}}}, \quad (6)$$

where ϵ is the privacy budget, u is the utility function and Δu is its ℓ_1 sensitivity. Notice that DP received increasing attention in face anonymization since it can provide a theoretically sound privacy protection by adding random perturbation. In previous works, DP is usually used with low-level

or middle-level data (e.g. pixels [7, 12] and identity features [6, 32, 55]) by adding Laplace noise. Differently, we generalize it to perform instance-level data sampling to reduce the disclosure risk of the identity information.

Visual Appearance Anonymization (VAA). Since the visual appearance may be correlated with some useful attributes, such as ethnic and age, the significant change of it may easily damage the data utility. To solve the problem, as shown in Figure 2 (c), we rely on IPD to sample a delegate face Z_a by using the following utility function (i.e. $u = u_a$)

$$u_a(x, X_i) = \frac{\max_j d(x, X_j) - d(x, X_i)}{\max_j d(x, X_j) - \min_j d(x, X_j)} \quad (7)$$

so that Z_a and x would have a high probability to share the same set of data utility, where $d(x, X_j)$ is the ℓ_2 distance between the features of x and X_j .

Geometry Structure Anonymization (GSA). The detected landmarks [2] are used to describe the facial geometry structure S_x of x . Instead of directly modifying the landmarks (which is a complicated task to make the result look real), we prefer to perform instance-level anonymization by replacing S_x with another realistic delegate. As shown in Figure 2 (c), the process consists of two steps. We first rely on IPD to sample a delegate \tilde{S}_x that has the same pose as S_x by using the following utility function (i.e. $u = u_g$)

$$u_g(S_x, X_i) = \frac{d(S_x, X_i) - \min_j d(S_x, X_j)}{\max_j d(S_x, X_j) - \min_j d(S_x, X_j)} \quad (8)$$

that tends to sample a more distinct geometry structure. Since this may violate the original pose and expression, we then recover them by adjusting the contour and mouth of \tilde{S}_x according to S_x and preserving the thickness of the upper and lower lips of \tilde{S}_x , resulting in \hat{S}_x . To recover the original background, we fuse \hat{S}_x and the background x_b of x as the geometry input Z_g . Note that \hat{S}_x is not the same as \tilde{S}_x , which would reduce the probability of identity intrusion.

3.3. Conditional Face Synthesis

This subsection focuses on face synthesis. As shown in the step2 of Figure 2 (a), our generator G takes appearance image Z_a , identity feature Z_{id} and geometry input Z_g as input, goes through the appearance encoder E and the conditional translator T to produce a realistic face image \hat{x} .

Appearance Encoder E is adopted to process Z_a to obtain an appearance feature f_a . We realize it by stacking six ResBlocks [15] and a SumPooling layer [61]. Semantic segmentation [57] is used to obtain the foreground face image of Z_a as input, when it is not available, we can approximately use the detected landmarks to realize this.

Conditional Translator T aims at translating Z_g and the condition input Z_c to a realistic face image \hat{x} . We employ a U-Net like structure to build T by following [20, 45, 61] via downsampling and upsampling with ResBlocks, where

adaptive instance normalization (AdaIN) [17] is employed to fuse the identity and appearance information encoded in Z_c which is defined as the fusion of Z_{id} and f_a by using a concatenation (Concat) layer and a fully connected (FC) layer: $Z_c = \text{FC}(\text{Concat}(Z_{id}, f_a))$. Note that, if the out contour were changed, T would adaptively inpaint the background so that the generated face could smoothly dissolve into the original background. Users can realize personalized anonymization by manipulating Z_{id} , Z_a and Z_g under different scenarios. For example, one can simply use the facial region of x to preserve the original appearance or attributes.

3.4. Training and Optimization

Let $\{x, y\}$ be a set of randomly sampled face images, x acts as the image to be anonymized and y acts as the identity provider, we present a 1:1 alternative reconstruction and cycle swap-reconstruction strategy for network training. For the former, \hat{x} is reconstructed from x , denoted as $Z_{id} = f_x$. For the latter, we change the identity from x to y and back to x in a loop, denoted as $Z_{id} = f_y$. The following multi-task loss function is used to optimize our generator

$$L_{All} = \lambda_1 L_1 + \lambda_2 L_2 + \lambda_3 L_3 + \lambda_4 L_4 + \lambda_5 L_5 + \lambda_6 L_6, \quad (9)$$

where L_1 is the adversarial loss, L_2 is the feature matching loss borrowed from [54] to stabilize the training process by matching the multi-layer features of discriminator D for the input and output images, L_3 is the perceptual loss, L_4 is the appearance loss, L_5 is the identity loss, $L_6 = \mathbb{E}[|x_b - g_b|_1]$ is the ℓ_1 between the background images of x and the generated image g , and $\lambda_1 \sim \lambda_6$ are parameters.

The adversarial loss L_1 is defined as

$$L_1 = \begin{cases} L_G(x, \hat{x}) + L_D(\hat{x}), & Z_{id} = f_x \\ \beta_1(L_G(x, \hat{y}) + L_D(\hat{y})) + \\ L_G(\hat{y}, \bar{x}) + L_D(\bar{x}), & Z_{id} = f_y, \end{cases} \quad (10)$$

where $L_G(x, \hat{x}) = \mathbb{E}[\max(0, 1 + D(\hat{x}, S_x)) + \max(0, 1 - D(x, S_x))]$ is used to optimize the generator G with the help of D which takes paired data (image, structure) as input, $L_D(\hat{x}) = \mathbb{E}[-D(\hat{x}, S_x)]$ is used to optimize D , $\hat{x} = G(f_x, Z_a(x), Z_g(x))$, $\hat{y} = G(f_y, Z_a(x), Z_g(x))$, $\bar{x} = G(f_x, Z_a(\hat{y}), Z_g(\hat{y}))$. The perceptual loss L_3 is defined as

$$L_3 = \begin{cases} \mathbb{E}[\sum_i \rho_i(x, \hat{x})], & Z_{id} = f_x \\ \mathbb{E}[\sum_i (\beta_3 \rho_i(x, \hat{y}) + \rho_i(\bar{x}, x))], & Z_{id} = f_y, \end{cases} \quad (11)$$

where VGG19 and FaceNet [22, 46, 48] are used to calculate ρ_1 and ρ_2 . The appearance loss L_4 is defined as

$$L_4 = \mathbb{E}[d_a(x, \hat{x})] + \mathbb{E}[d_e(x)], \quad (12)$$

where $d_a(x, \hat{x})$ is the ℓ_1 distance between $f_a(x)$ and $f_a(\hat{x})$, $d_e(x) = \max(0, \cos(f_a(x), f_x))$ discourages E from encoding the identity information. The identity loss is defined as




Figure 3. Intuitive comparison of our approach with the existing anonymization methods, where the first column presents the original faces.

$$L_5 = \begin{cases} \mathbb{E}[d_f(\hat{x}, x)], & Z_{id} = f_x \\ \mathbb{E}[\beta_2 d_f(\hat{x}, y) + d_f(\hat{x}, x)], & Z_{id} = f_y, \end{cases} \quad (13)$$

where $d_f(\hat{x}, x) = |f_{\hat{x}} - f_x|_2$ is the ℓ_2 distance between the identity features of \hat{x} and x .

4. Experiments

In this section, we show the performance of our approach by carrying out comparative study and user study on public datasets. We also rely on ablation study to show the influences of each component of our approach.

4.1. Settings

Dataset. Three popular datasets CelebA-HQ [28], VggFace2 [4] and LFW [16] are used. CelebA-HQ has 30,000 high quality facial images from 6,217 persons, where 5,000 images are used as the test set. VggFace2 has 3.31 million images from 9,131 persons, where a subset of 5,000 images from 1,000 identities are used for systematical analysis and fair comparison. LFW has 13,233 face images from 5,749 individuals, where 1,680 identities have more than one image and we use a subset of 5,000 images from them.

Implementation Details. For pre-processing, we employ [2] to detect facial landmarks and BiSeNet [57] to perform semantic segmentation. We rely on [20] and [61] to build G and D by stacking ResBlocks. We train our network to generate 256×256 images by using the Adam optimizer. We set $\lambda_1 \sim \lambda_4 = 1$, $\lambda_5 = \lambda_6 = \beta_2 = 2$, $\beta_1 = 0.6$, and $\beta_3 = 0.8$. Our approach is trained on CelebA-HQ and evaluated on all.

Evaluation Measures. Our approach is evaluated from the perspectives of privacy protection (anonymization and identity intrusion) and data reusability. For anonymization, we calculate re-identification (ReID) rate (in percentage %).

For identity intrusion, we calculate identity swapping (IDS) rate (%). The pre-trained FaceNet [46] are ArcFace [8] are used for face verification. The cosine similarity is used for ArcFace with two thresholds 0.30 and 0.35, and the ℓ_2 distance is used for FaceNet with three thresholds 0.9, 1.0 and 1.1. Face alignment [2] is used to evaluate face detection rate (%). LPIPS and SSIM [3] are used to evaluate image quality. Pre-trained classifiers [15] are used to evaluate attribute preservation (%), including expression, ethnic, gender, age and makeup.

4.2. Main Results

We mainly compare our approach with the following representative and state-of-the-art (SOTA) methods: CIAGAN [37], PIFD [3], DeepPrivacy (DP1) [18], DeepPrivacy2 (DP2) [19], LDFA [25], FALCO [1] and Riddle [29].

Qualitative Results. As shown in Figure 3, the success of Blurring and Pixelation can be contributed to the destruction of image content, which would tell the observers that the data is under protection. In contrast, the generative methods not only show excellent anonymization performance but also show higher probability of making the results imperceptible to observers. DP1 and DP2 may fail to retain some facial attributes, like expression. CIAGAN and LDFA may generate distorted faces. PIFD may bring some artifacts. FALCO and Riddle may lose the facial details. Compared with the other methods, our results can not only look realistic but also preserve more original attributes.

Quantitative Results. Since the image content of Pixelation and Blurring is significantly destroyed, we only compare with the generative methods in Table 1. Ours, Riddle [29] and PIFD [3] outperform the other methods on ReID across different face recognition backbones, especially when larger thresholds are used, where DP2 exhibits competitive results. Besides, it is very important to see if

Table 1. Results on privacy protection: (ReID, IDS). Lower ReID and IDS rates indicates better performance.

Method	ArcFace↓		FaceNet↓		
	$\cos > 0.30$	$\cos > 0.35$	$\ell_2 < 0.9$	$\ell_2 < 1.0$	$\ell_2 < 1.1$
CIAGAN	(1.7, 93.3)	(0.5, 51.4)	(0.5, 87.0)	(3.2, 99.9)	(11.3, 100)
PIFD	(0.0, 79.3)	(0.0, 19.1)	(0.0, 58.6)	(0.0, 99.7)	(0.0, 100)
DP1	(0.9, 70.9)	(0.2, 19.4)	(1.2, 73.5)	(4.8, 99.9)	(15.0, 100)
DP2	(0.5, 78.5)	(0.1, 26.4)	(0.6, 67.4)	(3.7, 99.8)	(7.4, 100)
LDFA	(4.9, 81.3)	(2.1, 33.6)	(4.3, 74.6)	(15.9, 99.8)	(24.8, 100)
FALCO	(6.7, 81.8)	(2.7, 37.0)	(1.4, 83.7)	(1.7, 99.9)	(5.4, 100)
Riddle	(0.0, 72.9)	(0.0, 26.2)	(0.6, 77.7)	(1.2, 99.8)	(2.3, 100)
Ours	(0.0, 78.1)	(0.0, 19.6)	(0.1, 34.9)	(0.3, 98.1)	(1.0, 100)

Table 2. The ReID and IDS results of using the Adaface model.

	CIAGAN	PIFD	DP1	DP2	FALCO	LDFA	Riddle	Ours
ReID	0.7	0.3	0.8	0.6	8.0	18.3	0.4	0.2
IDS	9.0	7.6	6.3	5.0	17.0	12.7	7.2	5.7

Table 3. Results on preserving facial attributes and image quality.

Method	Attribute↑					LPIPS↓	SSIM↑
	Express.	Ethnic	Gender	Age	Makeup		
CIAGAN	78.7	46.7	80.9	81.3	74.7	0.558	0.358
PIFD	82.3	48.5	84.4	82.9	63.7	0.124	0.771
DP1	54.8	52.0	84.7	84.4	66.8	0.192	0.785
DP2	59.1	52.6	84.3	85.3	78.9	0.127	0.779
LDFA	76.1	48.8	77.9	78.8	74.2	0.124	0.733
FALCO	82.6	51.8	84.8	86.3	77.6	0.307	0.475
Riddle	77.9	41.8	81.0	84.4	68.9	0.300	0.530
Ours	84.2	51.5	85.1	83.3	80.1	0.120	0.799

there happens identity intrusion after anonymization. CIAGAN suffers from the highest IDS rates. Compared with the best performed PIFD and DP1, our approach obtains competitive results (e.g. 58.6% vs. 34.9% for PIFD and Ours), which can be contributed to identity attention distraction. Also, we test another recent AdaFace model [24] for face verification. According to Table 2, it is easy to observe that our approach exhibits similar behavior to that in Table 1.

Utility Preservation. We compare the data utility of different methods in Table 3. DP1 and DP2 performs poor on expression. PIFD, DP1 and Riddle perform poor on makeup. All these methods perform poor on ethnic. Compared with the other methods, our approach achieves a much better balance on all items and performs well on preserving expression and makeup. We also find that, as with SOTA methods, our approach can also achieve 100% face detection rate, which also reveals its high data utility.

Diversity and Controllability. As shown in the first two rows of Figure 4, our approach can produce diverse results that look different from each other. Besides, we have tried to add a similar distraction item to Equ(5) to test the diversity of \hat{A} by using the bottom j -th prediction ($1 \leq j \leq 3$). According to the last row of Figure 4, although we can produce diverse results, the facial expression has changed from



Figure 4. Demonstration of our diverse anonymous results.




Figure 5. Demonstration of our results on controlling the geometry structures and visual appearances. Each of the top row shows the employed geometry structures and visual appearance images.

non-smile to smile, which indicates that adding such diversity may affect the data utility. Our approach can also support flexible anonymization according to user requirements and practical applications. For example, in Figure 5, we can produce different anonymous faces by controlling the geometry and visual appearance inputs, where the influences of changing the geometry structure is more significant.

Analysis. According to the above results, it is obvious that achieving high anonymization performance is relatively easy but it is somewhat difficult to (a) prevent identity intrusion and (b) preserve data utility. The reason why (a) is hard lies in that it is uneasy to ensure the non-existence of a synthesized identity in reality. The reason why (b) is hard lies in that identity is closely related to some critical facial attributes and the change of identity would inevitably lead to some variations on facial attributes. In Figure 3, 4 and 5, we can intuitively observe the attribute changes on the faces (e.g. eye and nose) and they may vary for different persons. Most existing methods pay much less attention on this and anonymization is usually achieved at the cost of damaging too much useful information. Although our approach can not perform the best all the time, it has achieved a better privacy-utility tradeoff. This can be mainly contributed to our anonymization strategy of minimizing the changes on identity independent attributes. But, our approach still suffers from the drawbacks of well preserving some facial attributes, such as ethnic and eye gaze direction, which are left for the follow up works.




Figure 6. The influential curves of identity attention distraction in the feature space by increasing K from 1 to 10.

Table 4. Ablation study results by removing the key components.

Method	ArcFace↓	FaceNet↓	Attribute↑	LPIPS↓	SSIM↑
w/o. IFA	(45.5, 54.0)	(67.6, 85.8)	88.0	0.085	0.860
w/o. VAA	(0.2, 20.1)	(0.2, 36.1)	77.2	0.120	0.800
w/o. GSA	(0.2, 22.7)	(0.1 , 36.6)	78.2	0.121	0.810
Full Model	(0.1, 19.6)	(0.1, 34.9)	76.8	0.120	0.799



Figure 7. Visualization of the distracted CAM heatmaps results by increasing K from 1 to 7 in Equation (5).

4.3. Ablation Study

In Figure 6, we plot the influential curves for the top K predictions of the face classification network. With top 1 distraction, our approach can remove over 75% of identity information. When increasing K , the ReID and IDS rates keep decreasing and the trend would slow down when $K \geq 2$. The utility preservation performance would decrease slightly along with K (see Attribute and LPIPS). Most of the curves would become almost flat after $K \geq 4$. Thus, we generally recommend to set K vary from 2 to 4 to reduce the computational costs and the loss on data utility.

In Figure 7, we qualitatively present some visual comparison results. With larger K , the CAM heatmap would focus much farther from the facial parts, especially the eyes. By comparing the face images before and after feature distraction, one can find some significant changes on the facial parts (e.g. the nose may vary from small to big or vice versa) but they may vary for different persons. The results also show that the joint distraction in Equation (5) would push some critical facial features or attributes heading to the opposite direction to realize identity anonymization, but this may also lead to some other unexpected changes, such as the age of the first person in Figure 7. This negative effect may come from the significant change of the identity related information in A . These observations show that some facial




Figure 8. Visualization of our results before (left) and after (right) applying IFA. The colored numbers denote different persons.

Table 5. Quantitative performance study on IFA and GSA.

Method	ArcFace↓	FaceNet↓	Attribute↑	LPIPS↓	SSIM↑
IFA-Rand	(0.5, 22.8)	(1.1, 36.1)	79.7	0.111	0.821
IFA-KFN	(0.2 , 23.1)	(0.1 , 37.3)	79.2	0.111	0.821
IFA-Ours	(0.2, 21.8)	(0.1, 36.3)	79.1	0.112	0.819
GSA-Rand	(45.5, 89.7)	(47.6, 79.1)	79.9	0.117	0.807
GSA-KFN	(30.7 , 89.9)	(56.7, 81.0)	80.2	0.117	0.808
GSA-Ours	(43.9, 90.7)	(38.5, 83.2)	86.6	0.083	0.853

parts or attributes are critical because they are correlated to identity representation, and the change of them may more or less lead to the performance drop on utility preservation.

In Figure 8, we compare the identity features before and after IFA by using the classical T-SNE embedding [52]. The intra-class differences would increase after identity feature distraction, but they still exhibit clustering characteristics regardless of some outliers, which would favor utility preservation. Note that our results can preserve the personalized facial attributes according to the status of each image. For example, as for the 7-th person, our results can still retain the makeup of each instance.

In Table 4, we present the ablation study results by removing the key components. w/o. IFA means removing the identity input and only using Z_a as the conditional input. w/o. VAA means feeding the original face appearance to Z_a . w/o. GSA means feeding the original geometry structure to Z_g . IFA can significantly help to reduce the ReID and IDS rates, but it suffers from performance drops on attribute preservation and image quality. GSA and VAA can further help to improve the ReID and IDS performance, but may lead to some drops on utility preservation. We have also verified the expression recovery in GSA by removing it and observed a significant performance drop.

In Table 5, we study the performance of IFA and GSA by using different strategies. IFA-Rand, IFA-KFN and IFA-Ours denote using random delegate, k -th farthest neighbor and our method to anonymize the identity feature, respectively. GSA-Rand, GSA-KFN and GSA-Ours denote using random delegate, k -farthest neighbor and our IPD method to anonymize the geometry structure, respectively. It is obvious that IFA-Ours and GSA-Ours have achieved a better balance between privacy protection and utility preservation.

In Equ.(7) and Equ.(8), u_a and u_g are determined by jointly considering anonymity and data utility. According

Table 6. Quantitative performance study on u_a and u_g .

Methods	ArcFace↓	FaceNet↓	AdaFace↓	Attribute↑	LPIPS↓	SSIM↑
w/o. u_a	(0.0, 19.1)	(0.0, 36.2)	(0.1, 6.6)	72.3	0.160	0.750
w/o. u_g	(0.1, 75.1)	(0.1, 36.4)	(0.7, 5.5)	77.7	0.128	0.784
Ours	(0.0, 19.6)	(0.1, 34.9)	(0.2, 5.7)	76.8	0.120	0.799




Figure 9. Our results (bottom) on the Vggface2 and LFW datasets.

Table 7. Results on Vggface2 (top) and LFW (bottom) datasets.

	ArcFace↓	FaceNet↓	Attribute↑	LPIPS↓	SSIM↑
CIAGAN	(0.4, 21.6)	(1.7, 53.2)	66.0	0.443	0.419
PIFD	(0.0, 10.2)	(0.2, 38.3)	68.0	0.520	0.363
DP1	(0.4, 39.1)	(0.7, 76.4)	74.5	0.082	0.891
DP2	(0.0, 14.1)	(0.9, 58.7)	62.2	0.479	0.520
Riddle	(0.0, 24.3)	(0.0, 77.9)	58.8	0.455	0.502
Ours	(0.0, 10.7)	(0.1, 27.8)	75.6	0.200	0.838
CIAGAN	(0.6, 3.9)	(1.5, 27.3)	72.2	0.458	0.397
PIFD	(0.0, 4.8)	(0.3, 8.0)	66.8	0.591	0.314
DP1	(0.4, 8.7)	(10.3, 61.6)	79.2	0.077	0.895
DP2	(0.5, 7.3)	(8.3, 56.9)	79.6	0.071	0.804
Riddle	(0.0, 8.5)	(0.1, 43.9)	61.7	0.460	0.428
Ours	(0.0, 4.7)	(0.7, 16.1)	80.9	0.074	0.895

to Table 6, u_a can help to improve the data utility and u_g can help to improve the protection ability.

4.4. Results on Vggface2 and LFW

We show the generalization ability of our approach on Vggface2 and LFW datasets. According to Table 7, most methods show quite high privacy protection ability (e.g. 0.0% ReID rate). Our approach outperforms the contrast methods on attribute, LPIPS and SSIM, which indicates that our approach has achieved a better privacy-utility tradeoff. According to Figure 9, one can find that the anonymized faces on both datasets look realistic and different from their original versions. These results are in consistent with the results reported in the previous subsections, which have again helped to verify the performance of our approach.

4.5. User Study

We conduct a simple user study to verify the performance of our approach from the perspectives of human. As shown in Figure 10, we asked around 30 participants to answer two kinds of questionnaires: for Q1, each anonymized face is paired with another image of the original face; for Q2, the top 5 retrieved results are presented. For each of the contrast methods, we randomly assign each participant: (1) 20 Q1 to calculate the ReID rate of choosing B and C; (2) 20 Q2



Figure 10. The employed questionnaires for user study.




Figure 11. The human evaluation results on ReID and IDS.

from dataset retrieval to calculate the IDS rate of choosing 'not appear'; (3) 20 Q2 from Google retrieval to calculate the ReID rate of choosing 'not appear'.

As shown in Figure 11, one can observe that: (1) our ReID and IDS results closely follow CIAGAN, which work better than the other methods; (2) the IDS rate of all the generative methods are high (over 45.0%), which may easily lead to identity intrusion. Since the image distortion would prevent the observers from correct recognition, it is easy for CIAGAN to achieve the best performance. This study show consistent results with that of machine recognizers.

5. Conclusion

In this paper, we present a distinct face anonymization approach from another viewpoint based on identity attention distraction. On top of ablation study, we have showed how and why our approach works. By performing comparative study and user study, we have validated our approach for improving the performance of privacy-utility tradeoff. Our approach allows for flexible manipulation of the facial appearance and geometry structure for more diverse anonymization and it has also demonstrated the generalizability in the other datasets. Future work includes exploring the correspondences between the convolutional feature maps and facial attributes for more effective anonymization, and exploring how to retain some more complex signal (e.g. psychological and physiological) hidden in visual data.

Acknowledgement

This work was supported in part by National Natural Science Foundation of China (Grant No. 62372147, 62125201, U21B2040), Zhejiang Provincial Natural Science Foundation of China (Grant No. LY22F020028).

References

- [1] Simone Barattin, Christos Tzelepis, Ioannis Patras, and Nicu Sebe. Attribute-preserving face dataset anonymization via latent code optimization. In *CVPR*, pages 8001–8010, 2023. [1](#), [2](#), [5](#)
- [2] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem?(and a dataset of 230,000 3d facial landmarks). In *ICCV*, pages 1021–1030, 2017. [4](#), [5](#)
- [3] Jingyi Cao, Bo Liu, Yunqian Wen, Rong Xie, and Li Song. Personalized and invertible face de-identification by disentangled identity information manipulation. In *ICCV*, pages 3334–3342, 2021. [1](#), [2](#), [5](#)
- [4] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: a dataset for recognising faces across pose and age. In *FG*, pages 67–74, 2018. [5](#)
- [5] Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. This looks like that: deep learning for interpretable image recognition. *NIPS*, 32, 2019. [2](#)
- [6] Jia-Wei Chen, Li-Ju Chen, Chia-Mu Yu, and Chun-Shien Lu. Perceptual indistinguishability-net (pi-net): Facial image obfuscation with manipulable semantics. In *CVPR*, pages 6478–6487, 2021. [1](#), [2](#), [4](#)
- [7] William L Croft, Jörg-Rüdiger Sack, and Wei Shi. Obfuscation of images via differential privacy: from facial images to general images. *Peer-to-Peer Networking and Applications*, 14(3):1705–1733, 2021. [4](#)
- [8] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, pages 4690–4699, 2019. [3](#), [5](#)
- [9] Matheus Alves Diniz and William Robson Schwartz. Face attributes as cues for deep face recognition understanding. In *FG*, pages 307–313, 2020. [1](#), [3](#)
- [10] Cynthia Dwork. Differential privacy: A survey of results. In *TAMC*, pages 1–19, 2008. [3](#)
- [11] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014. [3](#)
- [12] Liyue Fan. Image pixelization with differential privacy. In *DBSec*, pages 148–162, 2018. [4](#)
- [13] Oran Gafni, Lior Wolf, and Yaniv Taigman. Live face de-identification in video. In *ICCV*, pages 9378–9387, 2019. [1](#)
- [14] Xiuye Gu, Weixin Luo, Michael S Ryoo, and Yong Jae Lee. Password-conditioned anonymization and deanonymization with face identity transformers. In *ECCV*, pages 727–743, 2020. [2](#)
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [4](#), [5](#)
- [16] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: a database for studying face recognition in unconstrained environments. In *ECCV Workshop*, pages 1–14, 2008. [5](#)
- [17] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, pages 1501–1510, 2017. [4](#)
- [18] Håkon Hukkelås, Rudolf Mester, and Frank Lindseth. Deepprivacy: a generative adversarial network for face anonymization. In *ISVC*, pages 565–578, 2019. [1](#), [2](#), [3](#), [5](#)
- [19] Håkon Hukkelås and Frank Lindseth. Deepprivacy2: Towards realistic full-body anonymization. In *WACV*, pages 1329–1338, 2023. [1](#), [5](#)
- [20] Phillip Isola, Junyan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, pages 5967–5976, 2017. [1](#), [2](#), [4](#), [5](#)
- [21] Peng-Tao Jiang, Chang-Bin Zhang, Qibin Hou, Ming-Ming Cheng, and Yunchao Wei. Layercam: Exploring hierarchical class activation maps for localization. *TIP*, 30:5875–5888, 2021. [2](#)
- [22] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711, 2016. [4](#)
- [23] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410, 2019. [1](#), [2](#)
- [24] Minchul Kim, Anil K Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *CVPR*, pages 18750–18759, 2022. [1](#), [6](#)
- [25] Marvin Klemp, Kevin Rösch, Royden Wagner, Jannik Quehl, and Martin Lauer. Ldfa: Latent diffusion face anonymization for self-driving applications. In *CVPRW*, pages 3198–3204, 2023. [1](#), [2](#), [5](#)
- [26] Zhenzhong Kuang, Zhiqiang Guo, Jinglong Fang, Jun Yu, Noboru Babaguchi, and Jianping Fan. Unnoticeable synthetic face replacement for image privacy protection. *Neurocomputing*, 457:322–333, 2021. [2](#)
- [27] Zhenzhong Kuang, Huigui Liu, Jun Yu, Aikui Tian, Lei Wang, Jianping Fan, and Noboru Babaguchi. Effective de-identification generative adversarial network for face anonymization. In *ACMMM*, page 3182–3191, 2021. [2](#)
- [28] Cheng-Han Lee, Ziwei Liu, Lingyu Wu, and Ping Luo. Maskgan: towards diverse and interactive facial image manipulation. In *CVPR*, pages 5549–5558, 2020. [2](#), [5](#)
- [29] Dongze Li, Wei Wang, Kang Zhao, Jing Dong, and Tieniu Tan. Riddle: Reversible and diversified de-identification with latent encryptor. In *CVPR*, pages 8093–8102, 2023. [1](#), [2](#), [3](#), [5](#)
- [30] Jingzhi Li, Lutong Han, Ruoyu Chen, Hua Zhang, Bing Han, Lili Wang, and Xiaochun Cao. Identity-preserving face anonymization via adaptively facial attributes obfuscation. In *ACMMM*, pages 3891–3899, 2021. [2](#), [3](#)
- [31] Lingzhi Li, Jianmin Bao, Hao Yang, Dong Chen, and Fang Wen. Advancing high fidelity identity swapping for forgery detection. *CVPR*, pages 5074–5083, 2020. [2](#)
- [32] Tao Li and Chris Clifton. Differentially private imaging via latent space manipulation. *S&P*, 2021. [2](#), [4](#)
- [33] Tao Li and Lei Lin. Anonymousnet: natural face de-identification with measurable privacy. In *CVPRW*, pages 56–65, 2019. [2](#)

- [34] Yifang Li, Nishant Vishwamitra, Bart Knijnenburg, Hongxin Hu, and Kelly Caine. Blur vs. block: investigating the effectiveness of privacy-enhancing obfuscation for images. In *CVPRW*, pages 1343–1351, 2017. 1
- [35] Tianxiang Ma, Dongze Li, Wei Wang, and Jing Dong. Cfnet: Controllable face anonymization network with identity representation manipulation. *arXiv:2105.11137*, 2021. 2
- [36] Aravindh Mahendran and Andrea Vedaldi. Salient deconvolutional networks. In *ECCV*, pages 120–135, 2016. 2
- [37] Maxim Maximov, Ismail Elezi, and Laura Leal-Taixé. Ciagan: conditional identity anonymization generative adversarial networks. In *CVPR*, pages 5447–5456, 2020. 1, 2, 3, 5
- [38] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *FOCS*, pages 94–103, 2007. 3
- [39] Cade Metz. Facial recognition tech is growing stronger, thanks to your face, July 13, 2019. *The New York Times*. 1
- [40] Vahid Mirjalili, Sebastian Raschka, and Arun Ross. Privacynet: semi-adversarial networks for multi-attribute face privacy. *TIP*, 29:9400–9412, 2020. 2
- [41] Youxin Pang, Yong Zhang, Weize Quan, Yanbo Fan, Xiaodong Cun, Ying Shan, and Dong-Ming Yan. Dpe: Disentanglement of pose and expression for general video portrait editing. In *CVPR*, pages 427–436, 2023. 2
- [42] Hugo Proença. The uu-net: Reversible face de-identification for visual surveillance video footage. *TIP*, 32(2):496–509, 2021. 2
- [43] Zhongzheng Ren, Yong Jae Lee, and Michael S Ryoo. Learning to anonymize faces for privacy preserving action detection. In *ECCV*, pages 620–636, 2018. 1
- [44] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. 1
- [45] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. 4
- [46] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: a unified embedding for face recognition and clustering. In *CVPR*, pages 815–823, 2015. 2, 4, 5
- [47] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, pages 618–626, 2017. 2, 3
- [48] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 4
- [49] Abhishek Singh, Ethan Garza, Ayush Chopra, Praneeth Vepakomma, Vivek Sharma, and Ramesh Raskar. Decouple-and-sample: Protecting sensitive information in task agnostic data release. In *ECCV*, pages 499–517, 2022. 1
- [50] Qianru Sun, Liqian Ma, Seong Joon Oh, Luc Van Gool, Bernt Schiele, and Mario Fritz. Natural and effective obfuscation by head inpainting. In *CVPR*, pages 5050–5059, 2018. 1, 2
- [51] Fariborz Taherkhani, Nasser M Nasrabadi, and Jeremy Dawson. A deep face identification network enhanced by facial attributes prediction. In *CVPR*, pages 553–560, 2018. 3
- [52] Laurens Van der Maaten and Geoffrey Hinton. Visualizing non-metric similarities in multiple maps. *Machine learning*, 87(1):33–55, 2012. 7
- [53] Jiakai Wang, Aishan Liu, Zixin Yin, Shunchang Liu, Shiyu Tang, and Xianglong Liu. Dual attention suppression attack: Generate adversarial camouflage in physical world. In *CVPR*, pages 8565–8574, 2021. 2
- [54] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *CVPR*, pages 8798–8807, 2018. 4
- [55] Yunqian Wen, Bo Liu, Ming Ding, Rong Xie, and Li Song. Identitydp: Differential private identification protection for face images. *Neurocomputing*, 501:197–211, 2022. 1, 2, 4
- [56] Kaiyu Yang, Jacqueline Yau, Li Fei-Fei, Jia Deng, and Olga Russakovsky. A study of face obfuscation in imagenet. In *ICML*, 2022. 1
- [57] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *ECCV*, pages 325–341, 2018. 4, 5
- [58] Jun Yu, Zhenzhong Kuang, Baopeng Zhang, Wei Zhang, Dan Lin, and Jianping Fan. Leveraging content sensitiveness and user trustworthiness to recommend fine-grained privacy settings for social image sharing. *TIFS*, 13(5):1317–1332, 2018. 1
- [59] Lin Yuan, Linguo Liu, Xiao Pu, Zhao Li, Hongbo Li, and Xinbo Gao. Pro-face: A generic framework for privacy-preserving recognizable obfuscation of face images. In *ACMMM*, pages 1661–1669, 2022. 2
- [60] Zhuowen Yuan, Zhengxin You, Sheng Li, Zhenxing Qian, Xinpeng Zhang, and Alex Kot. On generating identifiable virtual faces. In *ACMMM*, pages 1465–1473, 2022. 2
- [61] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *ICCV*, pages 9459–9468, 2019. 2, 4, 5
- [62] Robert J Zatorre, Todd A Mondor, and Alan C Evans. Auditory attention to space and frequency activates similar cerebral systems. *Neuroimage*, 10(5):544–554, 1999. 2
- [63] Liming Zhai, Qing Guo, Xiaofei Xie, Lei Ma, Yi Estelle Wang, and Yang Liu. A3gan: Attribute-aware anonymization networks for face de-identification. In *ACMMM*, pages 5303–5313, 2022. 1
- [64] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016. 2
- [65] Bingquan Zhu, Hao Fang, Yanan Sui, and Luming Li. Deepfakes for medical video de-identification: privacy protection and diagnostic information preservation. In *AIES*, pages 414–420, 2020. 1, 2