

CIAGAN: Conditional Identity Anonymization Generative Adversarial Networks

Maxim Maximov*
Technical University Munich

Ismail Elezi*
University of Venice

Laura Leal-Taixé
Technical University Munich

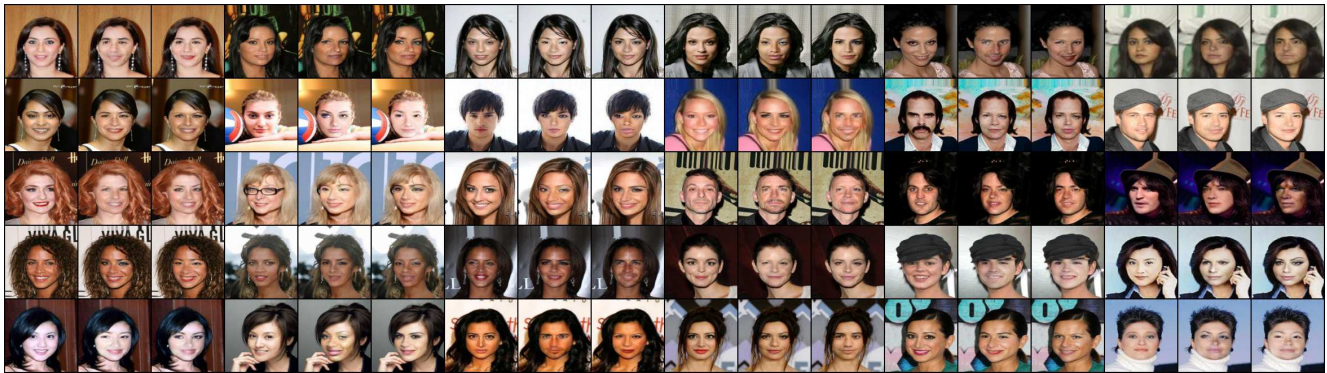


Figure 1: Given a face image, our network anonymizes the face based on the desired identity. In the figure, it can be seen the variability of the generated faces, controlled from a given label. In each triplet, the first image is the real image, while the other two images are different anonymized versions of the real image.

Abstract

The unprecedented increase in the usage of computer vision technology in society goes hand in hand with an increased concern in data privacy. In many real-world scenarios like people tracking or action recognition, it is important to be able to process the data while taking careful consideration in protecting people’s identity. We propose and develop CIAGAN, a model for image and video anonymization based on conditional generative adversarial networks. Our model is able to remove the identifying characteristics of faces and bodies while producing high-quality images and videos that can be used for any computer vision task, such as detection or tracking. Unlike previous methods, we have full control over the de-identification (anonymization) procedure, ensuring both anonymization as well as diversity. We compare our method to several baselines and achieve state-of-the-art results. To facilitate further research, we make available the code and the models at <https://github.com/dvl-tum/ciagan>.

“All human beings have three lives: public, private, and secret.”

Gabriel García Márquez, [24]

* Authors contributed equally.

1. Introduction

The ubiquitous usage of computer vision technology in society implies the automatic processing of large-scale visual data containing, more often than not, personal data. While we are eager to take advantage of technology for house monitoring, video conferencing and surveillance, we are not willing to do so by giving away our personal privacy. In fact, data privacy is an increasing concern, and entities such as the European Union have already passed laws such as the General Data Protection Regulations (GDPR)[1], to guarantee data privacy. For computer vision researchers, the creation of high-quality datasets which include people is becoming extremely challenging, as every person in the dataset needs to consent for the usage of his or her image data. Recently, a popular dataset for person re-identification, the Duke MTMC dataset [10], was taken offline for privacy reasons.

Our key observation is that many computer vision tasks such as person detection, multiple people tracking, or action recognition, do not need to *identify* the people on the videos, they just need to *detect* them. Classic anonymization techniques, such as face blurring, significantly alter the image and consequently result in a large detection performance drop.

We propose a model to anonymize (or de-identify) images and videos by removing the *identification* characteristics of people, while still keeping necessary features to

allow face and body *detectors* to work. Importantly, the images should still look realistic to a human observer, but people on them should not be identifiable. Our proposed approach can be used to anonymize computer vision datasets while retaining necessary information for tasks such as detection, recognition or tracking. We leverage the generative power of Conditional Generative Adversarial Networks (CGAN) [25, 15] to generate anonymized images and videos that look realistic. In existing GAN-based methods, the image generation process is typically controlled by a random noise vector to generate diverse outputs. Such a random process is not suitable for anonymization purposes, where we need to have guarantees that the identity has actually changed from input to output. To tackle this issue, we propose a novel identity control discriminator. Our CIA-GAN model fulfills the following important properties that an anonymization system should have:

- (i) **Anonymization:** the produced output must hide the identity of the person in the original image. Essentially, we are generating a new fake identity out of the input image.
- (ii) **Control:** the fake identity of the generated images is governed by a control vector, so we have full control over the real person-fake identity mapping.
- (iii) **New identities:** the generated images must contain only new identities not present in the training set.
- (iv) **Realistic:** output images must look realistic in order to be used by state-of-the-art detection and recognition systems.
- (v) **Temporal consistency:** temporal consistency and pose preservation in videos should be ensured for tasks like people tracking or action recognition.

By satisfying the aforementioned five properties, we ensure the anonymization of images and videos and the protection of data privacy. At the same time, our method guarantees that detectors will be able to use the anonymized data, as our experiments demonstrate.

Our **contribution** in this work is four-fold:

- We propose a general framework that is suitable for the anonymization of people in images and video streams.
- We show that images anonymized by our method can be used by existing detection and recognition systems.
- We demonstrate state-of-the-art results in several datasets, while at the same time qualitatively demonstrate diversity and control over the generated images.
- We perform a comprehensive ablation study, showing the importance of each building block of our model.

2. Related Work

Face Generation Since the advent of Generative Adversarial Networks [8, 28], generation of realistic faces has been an active research area [20, 16, 17]. Current state-of-the-art models [17] are able to generate high-resolution face images by progressively training large convolutional neural networks. Diversity in appearance, race, hair and eye color is achieved by adaptive instance normalization [13]. Despite their impressive quality, by conditioning on random noise and having no information about the original face, those methods do not have control over the pose of the generated face. Thus the blending of the face with the other parts of the body is challenging and remains an open research problem. Consequently, their usability in anonymization applications is limited.

Image-to-Image and Video-to-Video Translation Pix2Pix networks [15] and their unsupervised variant [40] have shown impressive results on cross-domain image translation, *e.g.*, from winter to summer. Nevertheless, it is not clear how suitable they are on making subtle but important changes in images coming from the same domain, such as faces or bodies. Closely related, there has been recent work on ensuring temporal consistency across videos for the task of face translation [38]. To ensure temporal consistency, [38] condition the generator on both the previous real and generated frames, in addition to the estimated optical flow between the frames. While the work demonstrates smooth temporal consistency, generated faces usually resemble closely the original identities, and are not suitable for the anonymization task.

Face Anonymization Until recently, face anonymization has been achieved by pixelization, blurring or masking the faces. Alternatively, [32] proposes to use segmentation-based methods. Being based on heuristics rather than being learned, there is no guarantee that these operations are optimal for the task of de-identification. Critically, these methods often make faces undetectable and therefore unusable in standard computer vision pipelines. We advocate the usage of machine *learning* to achieve anonymization that preserves important features that are necessary for computer vision tasks such as detection and tracking. This was previously explored in [29, 34, 14, 35, 7]. However, all these works come with important limitations. The faces generated by [14] are still, in general, identifiable by humans. [29] exhibits a similar issue and, furthermore, the method has no control over the generation process, with every identity being mapped to the same fake identity. The work of [34] focuses on altering facial landmarks which can lead to unnatural looking results. Furthermore, their method does not have explicit control over generated appearance. The results of [35] are visually appealing, but due to an optimization procedure for face alignment, the method is not

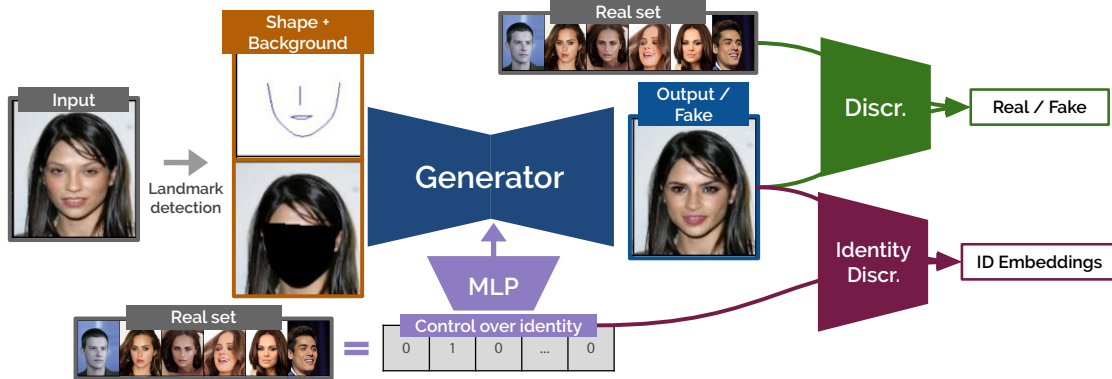


Figure 2: Our CIAGAN model takes as input the image, its landmark, the masked face and the desired identity. The generator is an encoder-decoder model where the encoder embeds the image information into a low dimensional space. The identity given as a one-hot label is encoded via a transposed convolutional neural network and is fed in the bottleneck of the generator. Then the decoder decodes the combined information of source images and the identities into a generated image. The generator plays an adversarial game with a discriminator in a standard GAN-setting. Finally, we introduce an identity discriminator network, whose goal is to provide a guiding signal to the generator about the desired identity of the generated face.

computationally efficient. Additionally, the method is designed to work only with faces since it is based on a parametric face model, so it is not straightforward to extend it to other domains such as full human bodies.

The state-of-the-art method is that of [7], where authors show good qualitative results and an unprecedented de-identification rate. However, while the generated images can fool recognition systems, in general, humans can recognize the identity of the presented faces. More critically, none of these methods except for [14] and [7] make any attempt at working with videos. [14] shows limited experiments in video processing, but the temporal consistency is not well-preserved. [7] shows very good temporal consistency, but as in the case of images, some identities are clearly not anonymized and easy to be spotted by a human eye. Additionally, they lack control and diversity to show differently anonymized outputs for the same input face. Our CIAGAN model provides a general framework in anonymizing both images and videos. By providing the labels for the desired generated identities and mixing styles of different identities, we have direct control over the de-identification process. This results in not only high-quality generated images but also with a much higher variability between the images of the same identity (see Fig. 1).

3. CIAGAN

In this section, we detail our methodology to anonymize images and videos. Our proposed Conditional Identity Anonymization Generative Adversarial Network (CIAGAN) leverages the power of generative adversarial networks to produce realistic images. In order to have control over the identity generation process and guarantee anonymization, we propose a novel identity discriminator to train CIAGAN. For the remainder of the section, we will

specifically refer to face anonymization, even though the method is directly applicable to full bodies.

3.1. Method overview

We show a full diagram of CIAGAN in Fig. 2. The main components of our method are the following:

Pose preservation and temporal consistency. We propose using a landmark-based representation of the input face (or body). This has two advantages: it ensures pose preservation which is especially useful for, *e.g.*, tracking; it provides a simple but efficient way to maintain temporal consistency when working with videos.

Conditional GAN. We exploit the generative power of GANs to produce realistic-looking results. It is important that standard detection and tracking systems can be applied to generated images without accuracy loss. Naturally, the realistic-looking generated faces are easily detectable. We achieve pose preservation by conditioning on the landmark representation. We train the conditional GAN in an adversarial game-theoretical way, where the discriminator judges the realism of the images generated by the generator.

Identity guidance discriminator. We propose a novel module that controls the identifying characteristics that the generator injects to create the new image. The identity discriminator and the generator play a collaboration game where they work together to achieve their common goal of generating realistic anonymized images. We now provide a more detailed description of the three modules of our method.

3.2. Pose preservation and temporal consistency

Several de-identification methods [29, 7] take as input the RGB image of the face to be anonymized. It is not surprising that there typically is some leakage of face infor-

mation to the generated image. Thus, while those methods produce high-quality images, the identity of the generated faces is not fully anonymized and can be often identified by a human.

Landmark image. To make sure our generated faces can not be linked to the original identity, we propose to work with an abstraction of faces. More precisely, we use face landmark images. This has two advantages: (i) the landmark image contains a sparse representation of the face with little identity information left, avoiding identity leakage, and (ii) the generator is conditioned on the face shape, allowing us to preserve the input pose in the output. This is especially important as we intend to use the generated images and videos as inputs to computer vision algorithms. In many vision applications, *e.g.* tracking, methods typically leverage the pose of the face or the body. Thereby, assuring the method will not change the pose of the anonymized face or body is extremely useful. To hide as much identity information as possible but still preserve the pose, instead of using all 68 landmarks [18], we use only the face silhouette, the mouth, and the bridge of the nose (see Fig. 2). This allows the network the freedom of choosing several facial characteristics, such as eye distance or eye shape, while at the same time, preserving expressions that are dependent on the mouth region, *e.g.* smiling or laughing, and the global pose is given by the nose position. The landmarks are represented as a binary image which is fed as input to the generator.

Masked background image. Our goal is to generate only the face region of an image and inpainting it to the original image background. This allows our algorithm to focus the learning power to the generation of faces (and not that of background), and at the same time guarantees that we do not have background changes that could interfere with the detection or tracking algorithms. To achieve this, we provide the generative model with the *masked background image* together with the *landmark image*. The masked background image still contains the forehead region of the head. Once the generator has access to this information, it can learn to match the skin appearance of the generated face to the forehead skin color. This leads to overall better quality of visual results. In cases there are multiple faces in the same image, we detect each face on the image and sequentially apply our anonymization framework.

Our pipeline can be additionally used for full-body anonymization by simply replacing the mask image with a segmentation mask representing the silhouette of the body. In our case, we do not use body joints as a replacement for the landmark image, as the silhouette of a person suffices as a pose prior.

Temporal consistency. In order to work on videos, any de-anonymization pipeline must ensure the temporal consistency of generated images on the video sequence. The

state-of-the-art video translation model [38] ensures temporal consistency by using a discriminator that is conditioned on the optical flow between corresponding frames. The optical flow is computed via an external neural network [6], which makes the framework both complex and computationally expensive. In our work, due to the nature of our input representation, we obtain temporal consistency for free. The landmarks of every frame are smoothed using a spline interpolation over neighboring frames. Therefore, we provide the same framework for both images and videos, with the only difference being the computationally cheap interpolation done at inference time.

3.3. Conditional generative adversarial networks

GAN framework. In simple terms, GANs combine two neural networks: a generator G whose goal is to generate realistic-looking samples, and a discriminator D whose goal is to differentiate between the real samples and the generated ones. The networks are trained in an adversarial manner with D being trained to maximize the probability of assigning the correct label to both training and generated examples, and G being trained to minimize the probability of D predicting the correct labels for the generated samples. In other words, D learns to separate the real samples from the generated ones, while G learns to fool D into classifying the generated samples as real ones. It is well-known that GAN training is hard and requires many tricks [23, 9, 4]. In this work, we choose to train CIAGAN with the LSGAN loss function [23]. The idea of using least-squares loss function for GAN training is simple yet powerful: the least-squares loss function is able to move the fake samples toward the decision boundary, as it penalizes also samples that are correctly classified but still lie far away from real samples. This is opposed to the cross-entropy loss that mostly penalizes wrongly classified samples. Based on this property, LSGANs are able to generate samples that are closer to real data.

The objective function of the discriminator in the LSGAN setting is defined as follows:

$$\min_D V_{LSGAN}(D) = \frac{1}{2} \mathbb{E}_{x \sim p_{data}(x)} [(D(x) - b)^2] + \frac{1}{2} \mathbb{E}_{z \sim p_z(z)} [(D(G(z)) - a)^2], \quad (1)$$

where a and b are the labels for the fake and the real data.

The loss of the generator is defined as:

$$\min_G V_{LSGAN}(G) = \frac{1}{2} \mathbb{E}_{z \sim p_z(z)} [(D(G(z)) - b)^2], \quad (2)$$

Without loss of generality, LSGAN can be replaced with any of the other common loss functions used for GAN training [9, 4].

Conditional GAN. In the classic GAN training setting, a vector of random noise is given as input to the generator in order to provide variability to the generated images. In our case, it is necessary for the generated faces to be aligned with the landmarks from the input image, for pose preservation and temporal consistency. Furthermore, we also need to blend the generated faces seamlessly with a background. To this end, we use the conditional GAN framework [15], where we condition the generator with a landmark and a masked image (background), as explained in Section 3.2. The generator uses an encoder-decoder architecture [22]. The encoder brings the landmark and the masked image to a low-dimensional representation (bottleneck), where it is combined with identity representation while the decoder takes combined representation and upsamples it to generate the anonymized RGB image.

3.4. Identity guidance

With the two modules explained above, our model learns to generate faces that look realistic and preserve the pose of original images. However, if the entire variability in image generation is provided by the landmark input, the network quickly overfits on the training set, effectively doing only image reconstruction. By doing so, it generates faces that are very similar to those in the training dataset, forfeiting the final anonymization goal. To solve this problem, we introduce a novel identity guidance discriminator. More precisely, for every given real image, we randomly choose the desired identity of its corresponding generated image. This identity – represented as on one hot-vector – is given as input to a transposed convolutional neural network. The network produces a parametrized version of the identity and feeds it into the bottleneck of the generator. In this way, the generator learns to generate a face with some of the characteristics of the desired identity while keeping the pose of the real image. In other words, the generated image is a composition of both the landmark identity and the desired identity. The identity of the generated image must not be the same as any of the real identities in order for the generated image to be not recognizable.

The identity discriminator is designed as a siamese neural network pretrained using Proxy-NCA loss [26]. Pre-training is done using real images, where the discriminator is trained to bring together features coming from the images that belong to the same identity. During the GAN training, we finetune the siamese network using contrastive loss [2]. During this fine-tuning step, we allow the siamese network to bring together the ID representation of the fake images and the real images. The identity discriminator and the generator are trained jointly in a collaborative manner. The identity discriminator’s goal is to provide a guidance signal to the generator to guide it towards creating images whose representational features are similar to those of a par-

ticular identity.

The case for multiple object tracking. Of particular interest is the *control* on the fake identity generation. We need to be able to keep the same real person-fake identity mapping within a sequence taken from a camera *e.g.*, multiple object tracking, but at the same time change the mapping for different cameras to avoid long-term tracking and potential misuse of the data. To do so, when a person moves from one camera to another, we replace its *control* vector with a new one, in turn giving the person a new identity. This is a simple yet powerful way of doing multiple object tracking within the frames taken from a camera, without the undesirable consequence of doing long-term tracking.

4. Experiments

In this section, we compare CIAGAN with several classic and learning-based methods commonly used for identity anonymization. Our method achieves state-of-the-art qualitative and quantitative results in different image and video datasets. We also present a set of comprehensive ablation studies to demonstrate the effect of our design choices. We first introduce the datasets, evaluation metrics and baselines used throughout this section.

Datasets. We perform experiments on 3 public datasets:

- **CelebA [21]** The dataset consists of 202,599 face images of 10,177 unique identities. We use the aligned version where each image is centered on a point in-between person’s eyes, and then padded and resized to have 178×218 resolution, while maintaining original face proportions. Each identity contains up to 35 photos. We use HOG [5] to construct face landmarks for each face.
- **MOTS [37]** Our method can also be adapted to work in other domains such as full-body anonymization. Instead of using face landmarks we use body segmentation masks. The dataset consist of 3,425 videos of 1,595 different people.
- **Labeled Faces in the Wild (LFW) [12]** The dataset consists of 6,000 pair images, split in 10 different splits, where half of the pairs contains images of same identity, and the remaining pairs consist of images that have different identities.

Baselines. We compare to standard anonymization methods as well as learning-based methods.

- **Simple Anonymization methods.** We use pixelization, blur and masking of faces and compare them with our method.
- **Image Translation methods.** We use the popular pix2pix [15] and CycleGAN [40] methods. We use the official code given by the authors, and we present the results in the supplementary material.

- **Face Replacement methods.** We compare the results in de-identification with the state-of-the-art results given by [7].

4.1. Implementation details

We generate the landmarks and the masks using Dlib-ml library [18]. We train our network on 128×128 resolution images, and use an encoder-decoder U-Net [31] architecture for the generator. The identity vector is parametrized by a transposed convolutional neural network containing a fully connected layer followed by multiple transposed convolutional layers. The features coming from the landmark and the identity branches are concatenated in the generator’s bottleneck. For the discriminator, we use a standard convolutional neural network that has the same architecture as the identity guidance network. We train our model for 60 epochs using the ADAM optimizer [19] with $1e-5$ learning rate. We set the beta hyperparameters β_1, β_2 to 0.5 and 0.9. The total training time for the model is one day in a single GPU. In the supplementary material, we give the remaining implementation details and the network architectures.

4.2. Evaluation metrics

We evaluate all models in face detection and re-identification metrics. We perform detection using HOG [5] and SSH Detector [27]. To evaluate the performance of the detectors, we use the percentage of detected faces. For re-identification, we train a siamese neural network using Proxy-NCA [26]. Additionally, we use a pre-trained FaceNet model [33] based on Inception-Resnet backbone [36]. We use the standard Recall@1 evaluation metric for re-identification. It measures the ratio of samples whose nearest neighbor is from the same class. The metric can take values from 0 to 100 with 0 showing perfect de-identification rate and 100 showing perfect identification rate. Note that in a balanced dataset, a random classifier will produce (on average) a Recall@1 of $1/|C|$ where C is the number of classes. Finally, we evaluate the visual quality of the generated images quantitatively using the Fréchet Inception Distance (FID) [11], a metric that compares the statistics of generated samples to those of real samples. The lower the FID, the better, corresponding to more similar real and generated samples.

4.3. Ablation Study

In this section, we perform an ablation study of our method in order to demonstrate the value of our design choices. In Table 1, we show several variants of our model. *Siamese* indicates our full model, with a siamese identity discriminator, and using landmarks as input. *Classification* indicates replacing the siamese identity discriminator by a classification network. As we can see, the results of the detection drop more than 35 percentage points (*pp*). We also

Models	Detection (\uparrow)	Recall@1 (\downarrow)	FID (\downarrow)
Siamese	99.9	1.3	2.1
Classification	64.6	0.4	63.2
Faces	98.3	1.1	6.5

Table 1: Ablation study of our model. First row presents the result of our model, second row shows the result of the model where the siamese identity guidance network is replaced with a classification network, while the third row shows the result of the model where the generator accepts full face images instead of landmarks.

Models	Detection (\uparrow)		Identification (\downarrow)	
	Dlib	SSH	PNCA	FaceNet
Original	100	100	70.7	65.1
Pixelization 16 by 16	0.0	0.0	0.3	0.3
Pixelization 8 by 8	0.0	0.0	0.4	0.3
Blur 9 by 9	90.6	38.6	16.9	57.2
Blur 17 by 17	68.4	0.3	1.9	0.5
Ours	99.9	98.7	1.3	1.0

Table 2: Results of common existing detection and recognition pre-trained methods. Lower (\downarrow) results imply a better anonymization. Upper (\uparrow) results imply a better detection.

show what happens if instead of landmarks, entire face images are provided as input. In these cases, the detection rate drops 1.6*pp*, and the FID score increases, showing that the faces are both more difficult to be detected, and have lower visual quality.

4.4. Quantitative results

Detection and Recognition. The first experiment evaluates two important capabilities that an anonymization method should have: high detection rate and low identification rate. That is, we do not want a trained system to be able to find the identity of the newly generated face, but at the same time, we still want a face detector to have a high detection rate.

In Table 2, we show the detection and identification results of our method compared to the other methods on the CelebA dataset [21]. The detection rate of classical HOG [5] and deep learning-based SSH [27] detectors in our anonymized images is at almost 100%. Blurring methods have a much lower detection rate in the images, while faces in the pixelized images are not detectable at all.

The identification performance drops from more than 70% recall on the original dataset to $1 - 1.5\%$ recall on our anonymized images. The CIAGAN generated images are almost unrecognizable by any of the two identification systems. Note, pixelization methods reach 0.3% recall, equivalent to random guessing, but that comes at the cost of deleting all content from the image, making both detection and identification impossible. In a setup where we want to further use computer vision algorithms on anonymized data, neither pixelization nor blurring is an option.

De-ID method	VGGFace2 (\downarrow)	CASIA (\downarrow)
Original	0.986 ± 0.010	0.965 ± 0.016
Gafni et al. [7]	0.038 ± 0.015	0.035 ± 0.011
Ours	0.034 ± 0.016	0.019 ± 0.008

Table 3: Comparisons with SOTA in LWF dataset. Lower (\downarrow) identification rates imply better anonymization.

Recognition based on landmarks. Given that the input to our generator is a landmark image, and not the actual image we want to de-anonymize, one might argue that an identification method that focuses on image pixels is easily fooled by our method, as we saw in Table 2. What happens if CIAGAN is attacked by an identification method trained only on landmarks? Will it still maintain its anonymization capabilities? We perform this experiment by training a similar identification method [26] as before, but with landmarks as the only input. We evaluate that by only using landmarks, we can identify up to 30.5% compared to 70.7% when using full images. However, when the same identifier is evaluated on the landmarks extracted from our anonymized faces, the performance drops to 1.9% recall. Even if the original landmarks are used as input to our generation model, CIAGAN uses them only as prior information to fuse it with the information coming from the embedder network.

Are we just doing face swapping? In Section 3.3, we introduced a novel identity guidance network that guides the generator to produce images with similar features to those of a given identity. One might argue that by doing so, the generator is learning to only do face swapping, replacing the face of the chosen identity into the landmarks of the source image. We show that this is not the case, by evaluating the identification rate of our generated images on the training set of the real images. We set the label of the generated images to the labels of their desired identity. If the generator is learning to do only face swapping, then a recognizer would be able to identify correctly all generated images. However, we show that this is not the case. Neither FaceNet [33] nor our model trained in P-NCA [26] are able to achieve a higher recognition rate than a random guesser. Additionally, in Fig. 3 we present a qualitative experiment, where the first image of each row contains the source images while the first image in each column is a randomly chosen image from the desired identities. The other images are generated. We see that the generated images take high-level characteristics of their given identities (such as race or sex), but differ greatly from the real images of those identities.

4.4.1 Comparison to SOTA in de-identification

In this section, we compare the de-identification (anonymization) power of our model compared to state-of-the-art [7]. We follow their evaluation protocol on the LFW dataset [12]. The dataset consists of 10 different




Figure 3: Generated faces of our model, where a source image is anonymized based on different identities.

splits, where each contains 600 pairs. A pair is defined as *positive* if both elements share the same identity, otherwise as *negative*. In every split, the first 300 pairs are positive, and the remaining 300 pairs are negative. As in [7], we anonymize the second image in every pair.

We use FaceNet [33] identification model, pre-trained on two public datasets: VGGFace2 [3] and CASIA-Webface [39]. The main evaluation metric is the true acceptance rate: the ratio of true positives for a maximum 0.001 ratio of false positives. We present the results in Table 3. The network evaluated in real faces reaches a score of almost 0.99, nearly perfect identification. [7] achieves an impressive anonymization performance by obtaining a score of less than 0.04 using the networks trained in both datasets. CIAGAN improves on this result and lowers the identification rate to 0.034 using the network trained in [3] and 0.019 using the network trained in [39], thus improving anonymization. On average, CIAGAN shows a 10.5% better de-identification rate on the first dataset, and a 45.7% better de-identification rate on the second dataset, while keeping a high detection rate of 99.13%. The average performance of 2.65% true positive rate shows that even a near-flawless system would completely fail to find the true identities in our CIAGAN-processed data, showing the strength of our method in achieving image anonymization.

4.5. Visual quality of the results

As can be seen in Table 1, our method reaches an FID score of 2.08. Simple baselines such as blurring and image translation methods reach a significantly higher (worse) FID score. The FID score comparison and qualitative results of the baselines can be found in the supplementary material.

We show a series of qualitative results. In Fig. 1, we show the diversity of the generated images, when the con-



Figure 4: Qualitative comparison with [7]. The image in the first column is the source image. In the first row, we show the generated images from the framework of [7], while in the second row, we show the generated images from CIAGAN.



Figure 5: Qualitative comparison with [7] on temporal consistency. From left to right: original frames; faces generated by our model using faces as input; faces generated by our model using landmarks as input, faces generated by [7].

control vector of the identity discriminator changes. We see that the generated images take high-level characteristics of the desired identity (such as eye shape, race or sex) while generating realistically looking images.

In Fig. 4, we qualitatively compare our results with those of [7]. We see that our method not only provides images which are more dissimilar to the source image, but by changing the control vector, our network is able to provide much more diverse images than those of [7] where the authors gradually change their control parameter λ (we can still recognize N. Cage).

In Fig. 5, we demonstrate the temporal consistency of our method and compare the results with those of [7]. We see that the pose is preserved in all cases, yielding excellent temporal consistency. At the same time, we see that the CIAGAN version which works with landmarks produces better looking images than the version trained on full faces.

Finally, in Fig. 6, we show an experiment on full-body anonymization. The first image in each row is the

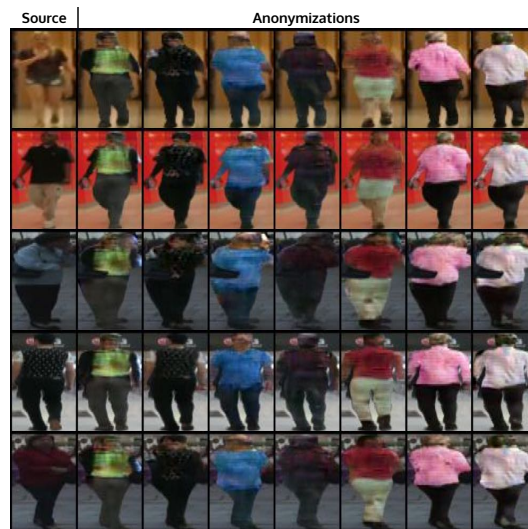


Figure 6: The full-body anonymization results of our framework in MOTs dataset.

source image, while the other images are the generated anonymized images. In each case, we see that the generated images keep the same posture as the source image counterpart, but the clothes, colors and other parts of the body are changed. To the best of our knowledge, this is the first time that the same framework has been successfully used to perform both face and body de-identification.

5. Conclusions and Future Work

Data privacy in images and videos is a serious problem. As computer vision researchers, we aim to do our part from the technological side. In this paper, we proposed a framework for face and body anonymization in images and videos. Our novel CIAGAN model is based on conditional generative adversarial networks, and faces are anonymized based on a guiding identity signal provided by a siamese network. We have shown that our method outperforms the state-of-the-art in de-identification while showing large diversity in the generated images.

A weakness of all current de-identification methods [30, 34, 7] is that they need the original faces to be initially detected before they can be anonymized. Consequently, any face that has not been detected can not be anonymized. Therefore, these methods are not deployable in systems where anonymization must be guaranteed. Our method suffers from a similar issue as it is dependent on landmark detection. As future work, we plan on working on full image anonymization and further eliminate the need for landmark detection in order to be able to handle extreme poses.

Acknowledgements. This research was partially funded by the Humboldt Foundation through the Sofja Kovalevskaja Award. We thank Aljosa Osep, Tim Meinhardt and Patrick Dendorfer for their helpful insights.

References

- [1] 2018 reform of eu data protection rules. <https://gdpr-info.eu>, 2018. 1
- [2] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a” siamese” time delay neural network. In *Advances in Neural Information Processing Systems*, pages 737–744, 1994. 5
- [3] Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *13th IEEE International Conference on Automatic Face & Gesture Recognition, FG 2018, Xi’an, China, May 15-19, 2018*, pages 67–74, 2018. 7
- [4] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2172–2180, 2016. 4
- [5] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2005), 20-26 June 2005, San Diego, CA, USA*, pages 886–893, 2005. 5, 6
- [6] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Häusser, Caner Hazirbas, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 2758–2766, 2015. 4
- [7] Oran Gafni, Lior Wolf, and Yaniv Taigman. Live face identification in video. In *IEEE International Conference on Computer Vision, ICCV 2019, Seoul, South Korea, October 27 - November 2, 2019*, 2019. 2, 3, 6, 7, 8
- [8] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2672–2680, 2014. 2
- [9] Ishaan Gulrajani, Faruk Ahmed, Martín Arjovsky, Vincent Dumoulin, and Aaron C. Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5767–5777, 2017. 4
- [10] Jules. Harvey, Adam. LaPlace. Megapixels: Origins, ethics, and privacy implications of publicly available face recognition image datasets, 2019. 1
- [11] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 6626–6637, 2017. 6
- [12] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007. 5, 7
- [13] Xun Huang and Serge J. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 1510–1519, 2017. 2
- [14] Håkon Hukkelås, Rudolf Mester, and Frank Lindseth. Deepprivacy: A generative adversarial network for face anonymization. *CoRR*, abs/1909.04538, 2019. 2, 3
- [15] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 5967–5976, 2017. 2, 5
- [16] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018. 2
- [17] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *CoRR*, abs/1812.04948, 2018. 2
- [18] Davis E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009. 4, 6
- [19] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2014. 6
- [20] Ming-Yu Liu and Oncel Tuzel. Coupled generative adversarial networks. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 469–477, 2016. 2
- [21] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 5, 6
- [22] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 3431–3440, 2015. 5
- [23] Xudong Mao, Qing Li, Haoran Xie, Raymond Y. K. Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2813–2821, 2017. 4
- [24] Gerald Martin. *Gabriel García Márquez: A Life*. Vintage Books, 2010. 1
- [25] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *CoRR*, abs/1411.1784, 2014. 2

- [26] Yair Movshovitz-Attias, Alexander Toshev, Thomas K. Leung, Sergey Ioffe, and Saurabh Singh. No fuss distance metric learning using proxies. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 360–368, 2017. [5](#), [6](#), [7](#)
- [27] Mahyar Najibi, Pouya Samangouei, Rama Chellappa, and Larry S. Davis. SSH: single stage headless face detector. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 4885–4894, 2017. [6](#)
- [28] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. [2](#)
- [29] Zhongzheng Ren, Yong Jae Lee, and Michael S. Ryoo. Learning to anonymize faces for privacy preserving action detection. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part I*, pages 639–655, 2018. [2](#), [3](#)
- [30] Zhongzheng Ren, Yong Jae Lee, and Micheal S. Ryoo. Learning to anonymize faces for privacy preserving action detection. In *European Conference on Computer Vision (ECCV)*, 2018. [8](#)
- [31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015 - 18th International Conference Munich, Germany, October 5 - 9, 2015, Proceedings, Part III*, pages 234–241, 2015. [6](#)
- [32] Michael S. Ryoo, Kiyoon Kim, and Hyun Jong Yang. Extreme low resolution activity recognition with multi-siamese embedding learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 7315–7322, 2018. [2](#)
- [33] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 815–823, 2015. [6](#), [7](#)
- [34] Qianru Sun, Liqian Ma, Seong Joon Oh, Luc Van Gool, Bernt Schiele, and Mario Fritz. Natural and effective obfuscation by head inpainting. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 5050–5059, 2018. [2](#), [8](#)
- [35] Qianru Sun, Ayush Tewari, Weipeng Xu, Mario Fritz, Christian Theobalt, and Bernt Schiele. A hybrid model for identity obfuscation by face replacement. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part I*, pages 570–586, 2018. [2](#)
- [36] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A. Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 4278–4284, 2017. [6](#)
- [37] Paul Voigtlaender, Michael Krause, Aljosa Osep, Jonathon Luiten, Berin Balachandar Gnana Sekar, Andreas Geiger, and Bastian Leibe. MOTs: multi-object tracking and segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 7942–7951. Computer Vision Foundation / IEEE, 2019. [5](#)
- [38] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Nikolai Yakovenko, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada.*, pages 1152–1164, 2018. [2](#), [4](#)
- [39] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z. Li. Learning face representation from scratch. *CoRR*, abs/1411.7923, 2014. [7](#)
- [40] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2242–2251, 2017. [2](#), [5](#)