

A Hybrid Model for Identity Obfuscation by Face Replacement

Qianru Sun Ayush Tewari Weipeng Xu
Mario Fritz Christian Theobalt Bernt Schiele

Max Planck Institute for Informatics, Saarland Informatics Campus
{qsun, atewari, wxu, mfritz, theobalt, schiele}@mpi-inf.mpg.de

Abstract. As more and more personal photos are shared and tagged in social media, avoiding privacy risks such as unintended recognition, becomes increasingly challenging. We propose a new hybrid approach to obfuscate identities in photos by head replacement. Our approach combines state of the art parametric face synthesis with latest advances in Generative Adversarial Networks (GAN) for data-driven image synthesis. On the one hand, the parametric part of our method gives us control over the facial parameters and allows for explicit manipulation of the identity. On the other hand, the data-driven aspects allow for adding fine details and overall realism as well as seamless blending into the scene context. In our experiments we show highly realistic output of our system that improves over the previous state of the art in obfuscation rate while preserving a higher similarity to the original image content.

1 Introduction

Visual data is shared publicly at unprecedented scales through social media. At the same time, however, advanced image retrieval and face recognition algorithms, enabled by deep neural networks and large-scale training datasets, allow to index and recognize privacy relevant information more reliably than ever. To address this exploding privacy threat, methods for reliable identity obfuscation are crucial. Ideally, such a method should not only effectively hide the identity information but also preserve the realism of the visual data, i.e., make obfuscated people look realistic.

Existing techniques for identity obfuscation have evolved from simply covering the face with often unpleasant occluders, such as black boxes or mosaics, to more advanced methods that produce natural images [1–3]. These methods either perturb the imagery in an imperceptible way to confuse specific recognition algorithms [2, 3], or substantially modify the appearance of the people in the images, thus making them unrecognizable even for generic recognition algorithms and humans [1]. Among the latter category, recent work [1] leverages a generative adversarial network (GAN) to inpaint the head region conditioned on facial landmarks. It achieves state-of-the-art performance in terms of both recognition rate and image quality. However, due to the lack of controllability of the image generation process, the results of such a purely data-driven method

inevitably exhibit artifacts by inpainting faces of unfitting face pose, expression or implausible shape. In contrast, parametric face models [4] give us complete control of facial attributes and have demonstrated compelling results for applications such as face reconstruction, expression transfer and visual dubbing [4–6]. Importantly, using a parametric face model allows to control the identity of a person as well as to preserve attributes such as face pose and expression by rendering and blending an altered face over the original image. However, this naive face replacement yields unsatisfactory results, since (1) fine level details cannot be synthesized by the model, (2) imperfect blending leads to unnatural output images and (3) only the face region is obfuscated while the larger head and hair regions, which also contain a lot of identity information, remain untouched.

In this paper, we propose a novel approach that combines a data-driven method and a parametric face model, and therefore leverages the best of two worlds. To this end, we disentangle and solve our problem in two stages (see Fig. 1): In the first stage, we replace the face region in the image with a rendered face of a different identity. To this end we replace the identity related component of the original person in the parameter vector of the face model while preserving attributes of original facial expression. In the second stage, a GAN is trained to synthesize the complete head image given the rendered face and an obfuscated region around the head as conditional inputs. In this stage, the missing region in the input is inpainted and fine grained details are added, resulting in a photo-realistic output image. Our qualitative and quantitative evaluations show that our approach significantly outperforms the baseline methods on publicly available datasets with both lower recognition rate and higher image quality.

2 Related work

Identity obfuscation. Blurring the face region or covering it with occluders, such as a mosaic or a black bar, are still the predominant techniques for visual identity obfuscation in photos and videos. The performance of these methods in concealing identity against machine recognition systems has been studied in [7] and [8]. They show that these simple techniques not only introduce unpleasant artifacts, but also become less effective due to the improvement of CNN-based recognition methods. Hiding the identity information while preserving the photo-realism of images is still an unsolved problem. Only a few works have attempted to tackle this problem.

For *target-specific* obfuscations, Sharif et al. [3] and Oh et al. [2] used *adversarial example* based methods which perturb the imagery in an imperceptible manner aiming to confuse specific machine recognition systems. Their obfuscation patterns are invisible to humans and the obfuscation performance is strong. However, obfuscation can only be guaranteed for target-specific recognizers.

To confuse *target-generic* machine recognizers and even human recognizers, Brkic et al. [9] generated full body images that overlay with the target person masks. However, synthesized persons with uniform poses do not match scene context which leads to blending artifacts in final images. The recent work of [1]

inpaints fake head images conditioned on the context and blends generated heads with diverse poses into varied background and body poses in social media photos. While achieving state-of-the-art performance in terms of both recognition rate and image quality, the results of such a purely data-driven method inevitably exhibits artifacts like the change of attributes such as face poses and expressions.

Parametric face models. Blanz and Vetter [10] learn an affine parametric 3D Morphable Model (3DMM) of face geometry and texture from 200 high-quality scans. Higher-quality models have been constructed using more scans [11], or by using information from in-the-wild images [12, 13]. Such parametric models can act as strong regularizers for 3D face reconstruction problems, and have been widely used in optimization based [5, 14, 12, 15–17] and learning-based [18–23] settings. Recently, a model-based face autoencoder (MoFA) has been introduced [4] which combines a trainable CNN encoder with an expert-designed differentiable rendering layer as decoder, which allows for end-to-end training on real images. We use such an architecture and extend it to reconstruct faces from images where the face region is blacked out or blurred for obfuscation. We also utilize the semantics of the parameters of the 3DMM by replacing the identity-specific parameters to synthesize overlaid faces with different identities. While the reconstructions obtained using parametric face models are impressive, they are limited to the low-dimensional subspace of the models. Many high-frequency details are not captured and the face region does not blend well with the surroundings in overlaid 3DMM renderings. Some reconstruction methods go beyond the low-dimensional parametric models [24, 6, 16, 17, 21, 19] to capture more detail, but most lack parametric control of the captured high-frequency details.

Image inpainting and refinement. We propose a GAN based method in the second stage to refine the rendered 3DMM face pixels for higher realism as well as to inpaint the obfuscated head pixels around the rendered face. In [25, 26], rendered images are modified to be more realistic by means of adversarial training. The generated data works well for specific tasks such as gaze estimation and hand pose estimation, with good results on real images. Raymond et al. [27] and Pathak et al. [28] have used GANs to synthesize missing content conditioned on image context. Both of these approaches assume strong appearance similarity or connection between the missing parts and their contexts. Sun et al. [1] inpainted head pixels conditioned on facial landmarks. Our method, conditioned on parametric face model renderings, gives us control to change the identity of the generated face while also synthesizing more photo-realistic results.

3 Face replacement framework

We propose a novel face replacement approach for identity obfuscation that combines a data-driven method with a parametric face model.

Our approach consists of two stages (see Fig. 1). Experimenting on different modalities of input results in different levels of obfuscation¹. In the first stage, we

¹ Stage-I input image choices: *Original image*, *Blurred face* and *Blacked-out face*.
Stage-II input image choices: *Original hair*, *Blurred hair* and *Blacked-out hair*.

can not only render a reconstructed face on the basis of a parametric face model (3DMM), but can also replace the face region in the image with the rendered face of a different identity. In the second stage, a GAN is trained to synthesize the complete head image given the rendered face and a further obfuscated image around the face as conditional inputs. The obfuscation here protects the identity information contained in the ears, hair, etc. In this stage, the obfuscated region is inpainted with realistic content and fine grain details missing in the rendered 3DMM are added, resulting in a photo-realistic output image.

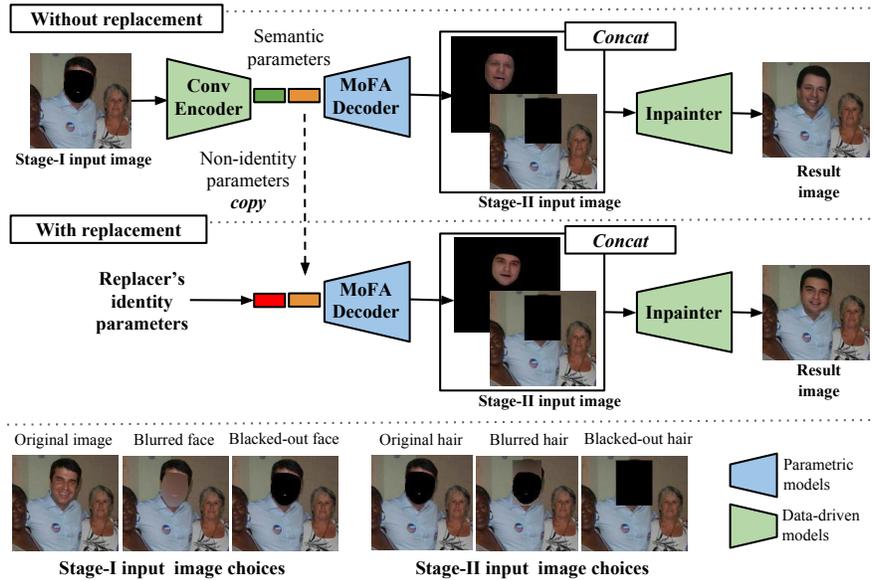


Fig. 1: Our obfuscation method based on data-driven deep models and parametric face models. The bottom row shows the input image choices for stage-I and stage-II. Different input combination results in different levels of obfuscation.

3.1 Stage-I: Face replacement

Stage-I of our approach reconstructs 3D faces from the input images using a parametric face model. We train a convolutional encoder to regress the model’s semantic parameters from the input. This allows us to render a synthetic face reconstructed from a person and also gives us the control to modify its rendered identity based on the parameter vector.

Semantic parameters. We denote the set of all semantic parameters as $p = (\alpha, \beta, \delta, \phi, \gamma)$, $|p| = 257$. These parameters describe the full appearance of the face. We use an affine parametric 3D face model to represent our reconstructions. α and β represent the shape and reflectance of the face, and correspond to the

identity of the person. These parameters are the coefficients of the PCA vectors constructed from 200 high-quality face scans [10]. δ are the coefficients of the expression basis vectors computed using PCA on selected blend shapes of [29] and [30]. We use 80 α , 80 β and 64 δ parameters. Together, they define the per-vertex position and reflectance of the face mesh represented in the topology used by [13]. In addition, we also estimate the rigid pose (ϕ) of the face and the scene illumination (γ). Rigid pose is parametrized with 6 parameters corresponding to a 3D translation vector and Euler angles for the rotation. Scene illumination is parameterized using 27 parameters corresponding to the first 3 bands of the spherical harmonic basis functions [31].

Our stage-I architecture is based on the Model-based Face Autoencoder (MoFA) [4] and consists of a convolutional encoder and a parametric face decoder. The encoder regresses the semantic parameters p given an input image².

Parametric face decoder. As shown in Fig. 1, the parametric face decoder takes the output of the convolutional encoder, p , as input and generates the reconstructed face model and its rendered image. The reconstructed face can be represented as $v_i(p) \in \mathbb{R}^3$ and $c_i(p) \in \mathbb{R}^3$, $\forall i \in [1, N]$, where $v_i(p)$ and $c_i(p)$ denote the position in camera space and the shaded color of the vertex i , and N is the total number of vertices. For each vertex i , the decoder also computes $u_i(p) \in \mathbb{R}^2$ which denotes the projected pixel location of $v_i(p)$ using a full perspective camera model.

Loss function. Our auto-encoder in stage-I is trained using a loss function that compares the input image to the output of the decoder as

$$E_{loss}(p) = E_{land}(p) + w_{photo}E_{photo}(p) + w_{reg}E_{reg}(p). \quad (1)$$

Here, $E_{land}(p)$ is a landmark alignment term which measures the distance between 66 fiducial landmarks [13] in the input image with the corresponding landmarks on the output of the parametric decoder,

$$E_{land}(p) = \sum_{i=1}^{66} \|l_i - u_x(p)\|_2^2. \quad (2)$$

l_i is the i th landmark’s image position and x is the index of the corresponding landmark vertex on the face mesh. Image landmarks are computed using the dlib toolkit [33]. $E_{photo}(p)$ is a photometric alignment term which measures the per-vertex appearance difference between the reconstruction and the input image,

$$E_{photo}(p) = \sum_{i \in V} \|I(u_i(p)) - c_i(p)\|_2. \quad (3)$$

V is the set of visible vertices and I is the image for the current training iteration. $E_{reg}(p)$ is a Tikhonov style statistical regularizer which prevents degenerate reconstructions by penalizing parameters far away from their mean,

$$E_{reg}(p) = \sum_{i=1}^{80} \frac{\alpha_i}{(\sigma_s)_i} + w_e \sum_{i=1}^{64} \frac{\delta_i}{(\sigma_e)_i} + w_r \sum_{i=1}^{80} \frac{\beta_i}{(\sigma_r)_i}. \quad (4)$$

² We use AlexNet[32] as the encoder.

σ_s , σ_e , σ_r are the standard deviations of the shape, expression and reflectance vectors respectively. Please refer to [4, 13] for more details on the face model and the loss function. Since the loss function $E_{loss}(p)$ is differentiable, we can back-propagate the gradients to the convolutional encoder, enabling self-supervised learning of the network.

Replacement of identity parameters. The controllable semantic parameters of the face model have the advantage that we can modify them after face reconstruction. Note that the shape and reflectance parameters α and β of the face model depend on the identity of the person [10, 20]. We propose to modify these parameters (referred to as identity parameters from now on) and render synthetic overlaid faces with different identities, while keeping all other dimensions fixed. While all face model dimensions could be modified we want to avoid unfitting facial attributes. For example, changing all dimensions of the reflectance parameters can lead to misaligned skin color between the rendered face and the body. To alleviate this problem, we keep the first, third and fourth dimensions of β , which control the global skin tone of the face, fixed.

After obtaining the semantic parameters on all our training set (over 2k different identities), we first cluster the identity parameters into 15 different identity clusters with the respective cluster means as representatives. We then replace the identity parameters of the current test image with the parameters of the cluster that is either closest (**Replacer1**), at middle distance (**Replacer8**) or furthest away (**Replacer15**) to evaluate different levels of obfuscation (Fig. 2). Note that each test image has its own **Replacers**.

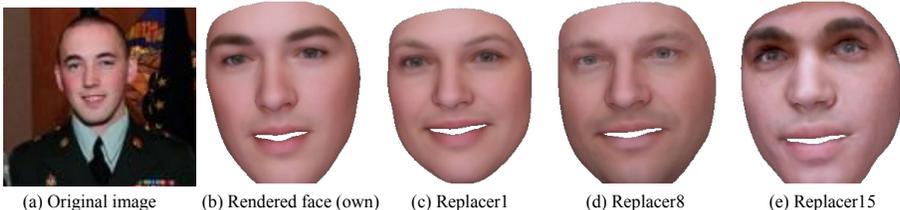


Fig. 2: Replacement of identity parameters in Stage-I allows us to generate faces with different identities.

Input image obfuscation. In addition to replacing the identity parameters, we also optionally allow additional obfuscation by blurring or blacking out the face region in the input image for Stage-I (the face region is determined by reconstructing the face from the original image). These obfuscation strategies force the Stage-I network to predict the semantic parameters only using the context information (Fig. 3), thus reducing the extent of facial identity information captured in the reconstructions. We train networks for these strategies using the full body images with the obfuscated face region as input while using the original

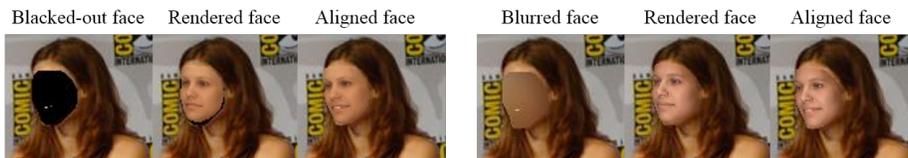


Fig. 3: Stage-I output: If the face in the input image is blacked out or blurred, our network can still predict reasonable parametric reconstructions which align to the contour of the face region. The appearance is also well estimated from the context information. The results are further aligned using an optimization-based strategy.

unmodified images in the loss function $E_{loss}(p)$ ³. This approach gives us results which preserve the boundary of the face region and the skin color of the person even for such obfuscated input images (Fig. 3). The rigid pose and appearance of the face is also nicely estimated.

In addition to reducing the identity information in the rendered face, the Stage-I network also removes the expression information when faces in the input images are blurred or blacked out. To better align our reconstructions with the input images without adding any identity-specific information, we further refine only the rigid pose and expression estimates of the reconstructions. We minimize part of the energy term in (1) after initializing all parameters with the predictions of our network.

$$p^* = \operatorname{argmin}_p E_{refine}(p) \quad (5)$$

$$E_{refine}(p) = E_{land}(p) + w_{reg} E_{reg}(p) \quad (6)$$

Note that only ϕ and δ are optimized during refinement. We use 10 non-linear iterations of a Gauss-Newton optimizer to minimize this energy. As can be seen in Fig 3, this optimization strategy significantly improves the alignment between the reconstructions and the input images. Note that input image obfuscation can be combined with identity replacement to further change the identity of the rendered face.

The output of stage-I is the shaded rendering of the face reconstruction. The synthetic face lacks high-frequency details and does not blend perfectly with the image as the expressiveness of the parametric model is limited. Stage-II enhances this result and provides further obfuscation by removing/reducing the context information from the full head region.

3.2 Stage-II: Inpainting

Stage-II is conditioned on the rendered face image from Stage-I and an obfuscated region around the head to inpaint a realistic image. There are two objectives for this inpainter: (1) inpainting the blurred/blacked-out hair pixels in

³ If the input image is not obfuscated in Stage-I, we directly use the pre-trained coarse model of [13] to get the parameters and the rendered face.

the head region; (2) modifying the rendered face pixels to add fine details and realism to match the surrounding image context. The architecture is composed of a convolutional generator G and discriminator D , and is optimized by L1 loss and adversarial loss.

Input. For the generator G , RGB channels of both the obfuscated image I and the rendered face from Stage-I F are concatenated as input. For the discriminator D , we take the inpainted image as *fake* and the original image as *real*. Then, we feed the (*fake*, *real*) pairs into the discriminator. We use the whole body image instead of just the head region in order to generate natural transitions between the head and the surrounding regions including body and background, especially for the case of obfuscated input.

Head Generator (G) and Discriminator (D). The head generator G is a ‘‘U-Net’’-based architecture [34], i.e. Convolutional Auto-encoder with skip connections between encoder and decoder⁴, following [1][35][36]. It generates a natural head image given both the surrounding context and the rendered face. The architecture of the discriminator D is the same as in DCGAN [37].

Loss function. We use L1 reconstruction loss plus adversarial loss, named \mathcal{L}^G , to optimize the generator and the adversarial loss, named \mathcal{L}^D , to optimize the discriminator. For the generator, we use the head-masked L1 loss such that the optimizer focuses more on the appearance of the targeted head region,

$$\mathcal{L}^G = \mathcal{L}_{bce}(D(G(I, F)), 1) + \lambda \| (G(I, F) - I_O) \odot M_h \|_1, \quad (7)$$

where M_h is the head mask (from the annotated bounding box), I_O denotes the original image and \mathcal{L}_{bce} is the binary cross-entropy loss. λ controls the importance of the L1 term⁵. Then, for the discriminator, we have the following losses:

$$\mathcal{L}^D = \mathcal{L}_{adv}^D = \mathcal{L}_{bce}(D(I_O), 1) + \mathcal{L}_{bce}(D(G(I, F)), 0). \quad (8)$$

We also tried to add a de-identification loss derived from verification models [38], in order to change the identity of the person in the generated image. However, this has a conflicting objective with the L1 loss and we were not able to find a good trade-off between them.

Fig. 4 shows the effect of our inpainter. In (a) when the original hair image is given, the inpainter refines the rendered face pixels to match surroundings, e.g., the face skin becomes more realistic in the bottom image. In (b)(c), the inpainter not only refines the face pixels but also generates the blurred/missing head pixels based on the context.

⁴ Network architectures and hyper-parameters are given in supplementary materials.

⁵ When λ is too small, the adversarial loss dominates the training and it is more likely to generate artifacts; when λ is too big, the generator mainly uses the L1 loss and generates blurry results.

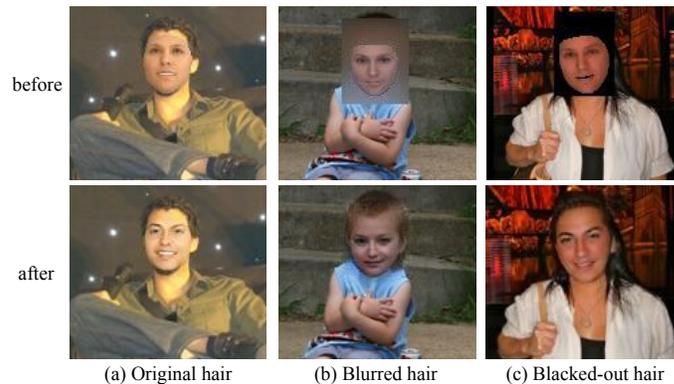


Fig. 4: Visualization results before and after inpainting. On the top row, rendered faces are overlaid onto the color images for better comparison of details.

4 Recognizers

Identity obfuscation in this paper is target-generic: it is designed to work against any recognizer, be it machine or human. In this paper, we use both recognizers to test our approach.

4.1 Machine recognizers

We use an automatic recognizer `naeil` [39], the state-of-the-art for person recognition in social media images [1][40]. In contrast to typical person recognizers, `naeil` also uses body and scene context cues for recognition. It has thus proven to be relatively immune to common obfuscation techniques like blacking-out or blurring the head region [7].

We first train feature extractors over head and body regions, and then train SVM identity classifiers on those features. We can concatenate features from multiple regions (e.g. head+body) to make use of multiple cues. In our work, we use GoogleNet features from `head` and `head+body` for evaluation. We have also verified that the obfuscation results show similar trends against AlexNet-based analogues (see supplementary materials).

4.2 Human recognizers

We also conduct human recognition experiments to evaluate the obfuscation effectiveness in a perceptual way. Given an original head image and the head images inpainted by variants of our method and results of other methods, we ask users to recognize the original person from the inpainted ones, and to also choose the farthest one in terms of identity. Users are guided to focus on identity recognition rather than the image quality. For each method, we calculate the percentage of times its results were chosen as the farthest identity (higher number implies better obfuscation performance).

5 Experiments

An obfuscation method should not only effectively hide the identity information but also produce photo-realistic results. Therefore, we evaluate our results on the basis of recognition rate and visual realism. We also study the impact of different levels of obfuscation yielded from different input modalities at two stages.

5.1 Dataset

Our obfuscation method needs to be evaluated on realistic social media photos. PIPA dataset [41] is the largest social media dataset (37,107 Flickr images with 2,356 annotated individuals), which shows people in diverse events, activities and social relations [42]. In total, 63,188 person instances are annotated with head bounding boxes from which we create head masks. We split the PIPA dataset into a training set and a test set without overlapping identities, following [1]. In the training set, there are 2,099 identities, 46,576 instances and in the test set 257 identities, 5,175 instances. We further prune images with strong profile or back of the head views from both sets following [1], resulting in 23,884 training and 1,084 test images. As our pipeline takes a fixed-size input ($256 \times 256 \times 3$), we normalize the image size of the dataset. To this end, we crop and zero-pad the images so that the face appears in the top middle block of a 3×4 grid in the entire image. Details of our crop method are given in supplementary materials.

5.2 Input modalities

Our method allows 18 different combinations of input modalities, which is a combination of 3 types of face modalities, 3 types of hair modalities and the choice of modifying the face identity parameters (default replacer is `Replacer15`). Note that, only 17 of them are valid for obfuscation, since the combination of original face and hair aims to reconstruct the original image. Due to space limitations, we compare a representative subset, as shown in Table 1. The complete results can be found in the supplementary material.

In order to blur the face and hair regions in the input images, we use the same Gaussian kernel as in [1, 7]. Note that in contrast to those methods, our reconstructed face model provides the segmentation of the face region allowing us to precisely blur the face or hair region.

5.3 Results

In this section, we evaluate the proposed hybrid approach with different input modalities in terms of the realism of images and the obfuscation performance.

Image realism. We evaluate the quality of the inpainted images compared to the ground truth (original) images using Structure Similarity Score (SSIM) [43]. During training, the body parts are not obfuscated, so we report the mask-SSIM [1, 35] for the head region only (SSIM scores are in supplementary materials). This score measures how close the inpainted head is to the original head.

Table 1: Quantitative results comparing with the state-of-the-art methods [1]. Image quality: Mask-SSIM and HPS scores (both the higher, the better). Obfuscation effectiveness: recognition rates of machine recognizers (lower is better) and confusion rates of human recognizers (higher is better). **v*** simply represents the method in that row.

Obfuscation method			Evaluation				
Stage-I	Stage-II		Image quality		Machine		Human
	Hair	Rendered face	Mask-SSIM	HPS	head	body+head	
Original	-	-	1.00	0.93	85.6%	88.3%	
[1]	Blackhead+Detect		0.41	0.19	10.1%	21.4%	-
[1]	Blackhead+PDMDec.		0.20	0.11	5.6%	17.4%	-
[1], our crop	Blackhead+Detect		0.43	0.34	12.7%	24.0%	4.1%
[1], our crop	Blackhead+PDMDec.		0.23	0.15	9.7%	19.7%	20.1%
v1 , Original	Overlay-No-Inpainting		0.75	0.58	66.9%	68.9%	
v2 , Original	Original Own		0.87	0.71	70.8%	71.5%	-
v3 , Original	Original Replacer15		-	0.49	47.6%	57.4%	-
v4 , Blurred	Original Own		0.86	0.59	59.9%	65.2%	-
v5 , Blurred	Original Replacer15		-	0.41	26.3%	41.7%	-
v6 , Blurred	Blurred Own		0.55	0.55	25.8%	38.0%	-
v7 , Blurred	Blurred Replacer15		-	0.40	12.7%	29.3%	-
v8 , Blacked	Original Own		0.85	0.60	59.3%	64.4%	-
v9 , Blacked	Blacked Own		0.47	0.41	14.2%	25.7%	2.9%
v10 , Blacked	Blacked Replacer1		-	0.45	11.8%	23.5%	6.2%
v11 , Blacked	Blacked Replacer8		-	0.39	9.3%	22.4%	31.3%
v12 , Blacked	Blacked Replacer15		-	0.33	7.1%	18.1%	35.4%

The SSIM metric is not applicable when using a **Replacer**, as ground truth images are not available. Therefore, we conduct a human perceptual study (HPS) on Amazon Mechanical Turk (AMT) following [1, 35]. For each method, we show 55 real and 55 inpainted images in a random order to 20 users, who are asked to answer whether the image looks real or fake within 1s.

Obfuscation performance. Obfuscation evaluation is to measure how well our methods can fool automatic person recognizers as well as humans. We have defined machine recognizers and human recognizers in Section 4.

For machine recognizers, we report the average recognition rates for 1,084 test images in Table 1. For human recognition, we randomly choose 45 instances then ask recognizers to verify the identity, given the original image as reference, from the obfuscated images of six representative methods: two methods in [1] and four methods indexed by **v9-v12**, see the last column of Table 1.

Comparison to the state-of-the-art. In Table 1, we report quantitative evaluation results on different input modalities and in comparison to [1]. We also implement the exact same models of [1] on our cropped data for fair comparisons. We also compare the visual quality of our results with [1], see Fig. 5.

Our best obfuscation rate is achieved by **v12**. The most comparable method in [1] is *Blackhead+PDMDec*, where the input is an image with a fully blacked-

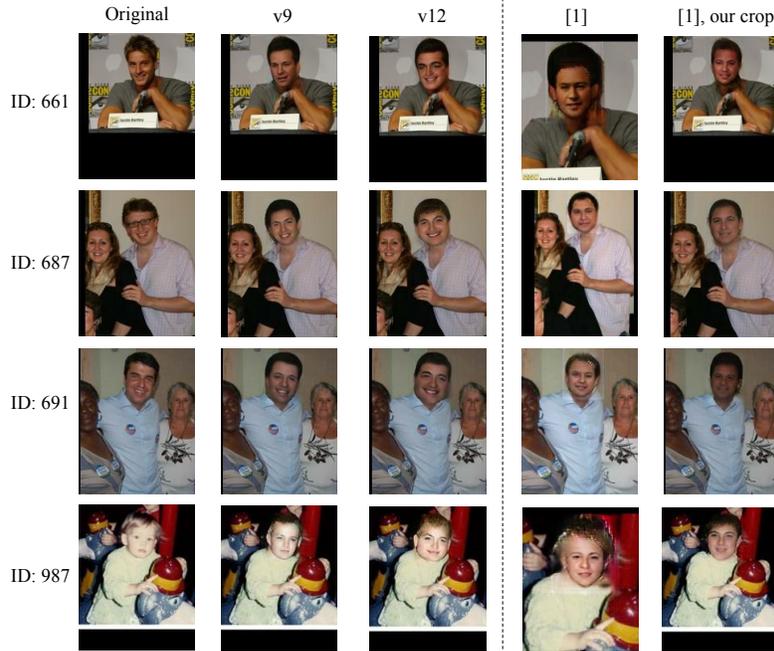


Fig. 5: Result images by methods **v9** and **v12**, compared to original images and the results of the *Blackhead* scenario using *PDMDec* landmarks in [1]. Note that the image scale difference with [1] is because of different cropping methods.

out head and the landmarks are generated by *PDMDec*. Comparing **v12** with it, we achieve 2.6% lower recognition rate (2.6% higher for confusing machine recognizers) using **head** features. Our method does even better (15.3% higher) in fooling human recognizers. In addition, our method has clearly higher image quality in terms of HPS, 0.33 vs. 0.15 [1]. Figure 5 shows that our method generates more natural images in terms of consistent skin colors, proper head poses and vivid face expressions.

Parametric model versus GAN. For the ablation study of our hybrid model, we replace the parametric model in Stage-I with a GAN. We use the same architecture as the stage-II network of [1], but without the landmark channel. This is inspired by the regional completion method using GANs [44]. We consider two comparison scenarios: when the input face is blacked-out (indexed by **v13**), we compare with **v8**; when the head is blacked-out (indexed by **v14**), we compare with **v9**. We observe that **v13** results in a lower mask-SSIM score of 0.80 (**v8** has 0.85) with the same recognition rate of 64.4%. This means that the GAN generates lower-quality images without performing better obfuscation. **v14** has a lower recognition rate of 19.7% vs. **v9**'s 25.7%, but its mask-SSIM (image quality) is only 0.23, 0.24 lower than **v9**'s 0.47. If we make use of face replacement (only applicable when using our parametric model-based approach), we are able to achieve a lower recognition rate of 18.1%, sacrificing only 0.08 in terms of image quality (see HPS of **v9** and **v12** in Table 1).

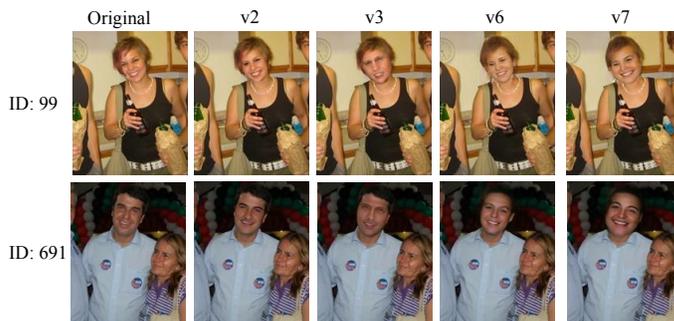


Fig. 6: Result images of methods **v2**, **v3**, **v6** and **v7**, compared to the original images.

Analysis of different face/hair modalities. Table 1 shows that different modalities of the input yield different levels of obfuscation and image quality. In general, the image quality is roughly correlated to the recognition rate. With higher level of modification to an image, the identity will be more effectively obfuscated, but the image quality will also deteriorate accordingly. However, we can observe that the recognition rate drops quicker than the image quality.

It is worth noting that when there is no inpainting on the rendered faces (**v1**), HPS score is 0.58, 0.13 lower than **v2**, verifying that rendered faces are less realistic than inpainted ones. Not surprisingly, the best image quality is achieved by **v2** which aims to reconstruct the original image without obfuscation. On top of that, when we use blurred faces in Stage-I (**v4**), the machine recognition rate (**head**) drops from 70.8% to 59.9%. This indicates that blurring the face region indeed partially conceals the identity information.

When we blur the hair region (**v6**), the recognition rate sharply drops to 25.8%, which implies that the region around the face contains a large amount of identity information. When we remove all information from the face and hair regions (**v9**), we get an even lower recognition rate of 14.2%.

Face replacement is of great effectiveness. We can see from Table 1 that replacing the face parameters with those of another identity is an effective way of hiding the identity information. Regardless of the face and hair input modalities, the obfuscation performances on both recognizers are significantly improved using **Replacer15** rendered faces than using **Own** rendered faces. Replacing faces from close to far identities also has an obvious impact on the obfuscation effectiveness. From **v10** to **v12** in Table 1, we can see using **Replacer8** yields clearly better obfuscation than the **Replacer1**, e.g., obfuscation for humans gets 25.1% improvement. This is further evidenced by the comparison between the **Replacer15** and **Replacer1**. Visually, Fig. 6 and Fig. 5 show that replacing the face parameters indeed makes the faces very different.

Trade-off between image quality and obfuscation. Fig. 7 shows the machine recognition rate vs. image quality plots for different obfuscation methods (some are not in Table 1 but in supplementary materials). Points on the curves from left to right are the results of using *Blacked-out*, *Blurred* and *Original* hair inputs for stage-II.

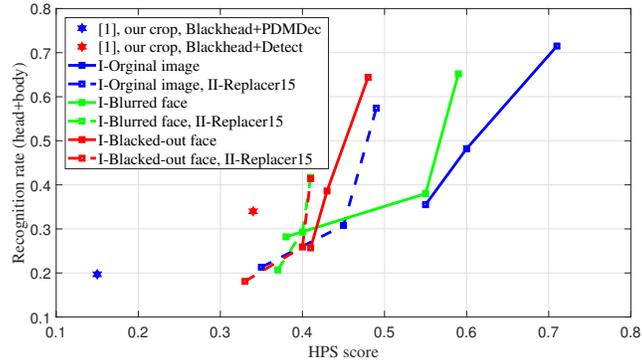


Fig. 7: Scatter curves of different obfuscation methods. HPS scores change along X-axis for different obfuscation levels (*Blacked-out*, *Blurred*, *Original*) on hair regions.

This figure allows users to select the method with the highest image quality given a specified obfuscation threshold. For example, if a user would like to take the risk of 30% recognizability at most, the highest image quality he/she can get is about 0.45, corresponding to the middle point on the blue dashed line (the method of *Original image*, *Blurred hair*, *Replacer15*). On the other hand, if a user requires the image quality to be at least 0.30, the best obfuscation possible corresponds to the first point of the red dashed line (the method of *Blacked-out face*, *Blacked-out hair*, *Replacer15*). The global coverage of these plots show the selection constrains, such as when a user strictly controls the privacy leaking rate under 20%, there are only two applicable methods: *Blackhead+PDMDec* [1] (image quality is only 0.15) and ours (*Blacked-out face*, *Blacked-out hair*, *Replacer15*) where the image quality is higher at 0.33.

6 Conclusion

We have introduced a new hybrid approach to obfuscate identities in photos by head replacement. Thanks to the combination of a parametric face model reconstruction and rendering, and the GAN-based data-driven image synthesis, our method gives us complete control over the facial parameters for explicit manipulation of the identity, and allows for photo-realistic image synthesis. The images synthesized by our method confuse not only the machine recognition systems but also humans. Our experimental results have demonstrated output of our system that improves over the previous state of the art in obfuscation rate while generating obfuscated images of much higher visual realism.

Acknowledgments

This research was supported in part by German Research Foundation (DFG CRC 1223) and the ERC Starting Grant CapReal (335545). We thank Dr. Florian Bernard for the helpful discussions.

References

1. Sun, Q., Ma, L., Oh, S.J., Gool, L.V., Schiele, B., Fritz, M.: Natural and effective obfuscation by head inpainting. In: CVPR. (2018)
2. Oh, S.J., Fritz, M., Schiele, B.: Adversarial image perturbation for privacy protection – a game theory perspective. In: ICCV. (2017)
3. Sharif, M., Bhagavatula, S., Bauer, L., Reiter, M.K.: Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. (2016)
4. Tewari, A., Zollhöfer, M., Kim, H., Garrido, P., Bernard, F., Perez, P., Theobalt, C.: Mofa: Model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. In: ICCV. Volume 2. (2017)
5. Thies, J., Zollhöfer, M., Stamminger, M., Theobalt, C., Nießner, M.: Face2Face: Real-time Face Capture and Reenactment of RGB Videos. In: CVPR. (2016)
6. Garrido, P., Zollhöfer, M., Casas, D., Valgaerts, L., Varanasi, K., Perez, P., Theobalt, C.: Reconstruction of personalized 3d face rigs from monocular video. ACM Trans. Graph. (Presented at SIGGRAPH 2016) **35**(3) (2016) 28:1–28:15
7. Oh, S.J., Benenson, R., Fritz, M., Schiele, B.: Faceless person recognition; privacy implications in social media. In: ECCV. (2016)
8. McPherson, R., Shokri, R., Shmatikov, V.: Defeating image obfuscation with deep learning. arXiv **1609.00408** (2016)
9. Brkic, K., Sikiric, I., Hrkac, T., Kalafatic, Z.: I know that person: Generative full body and face de-identification of people in images. In: CVPR Workshops. (2017) 1319–1328
10. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: SIGGRAPH, ACM Press/Addison-Wesley Publishing Co. (1999) 187–194
11. Booth, J., Roussos, A., Zafeiriou, S., Ponniah, A., Dunaway, D.: A 3d morphable model learnt from 10,000 faces. In: CVPR. (2016)
12. Booth, J., Antonakos, E., Ploumpis, S., Trigeorgis, G., Panagakis, Y., Zafeiriou, S.: 3d face morphable model “in-the-wild”. In: CVPR. (2017)
13. Tewari, A., Zollhöfer, M., Garrido, P., Bernard, F., Kim, H., Pérez, P., Theobalt, C.: Self-supervised multi-level face model learning for monocular reconstruction at over 250 hz. In: CVPR. (2018)
14. Roth, J., Tong, Y., Liu, X.: Adaptive 3d face reconstruction from unconstrained photo collections. IEEE Transactions on Pattern Analysis and Machine Intelligence **39**(11) (2017) 2127–2141
15. Romdhani, S., Vetter, T.: Estimating 3D Shape and Texture Using Pixel Intensity, Edges, Specular Highlights, Texture Constraints and a Prior. In: CVPR. (2005)
16. Garrido, P., Valgaxerts, L., Wu, C., Theobalt, C.: Reconstructing detailed dynamic face geometry from monocular video. In: ACM Trans. Graph. (Proceedings of SIGGRAPH Asia 2013). Volume 32. (2013) 158:1–158:10
17. Shi, F., Wu, H.T., Tong, X., Chai, J.: Automatic acquisition of high-fidelity facial performances using monocular videos. ACM Transactions on Graphics (TOG) **33**(6) (2014) 222
18. Richardson, E., Sela, M., Kimmel, R.: 3D face reconstruction by learning from synthetic data. In: 3DV. (2016)
19. Sela, M., Richardson, E., Kimmel, R.: Unrestricted facial geometry reconstruction using image-to-image translation. In: ICCV. (2017)

20. Tran, A.T., Hassner, T., Masi, I., Medioni, G.G.: Regressing robust and discriminative 3d morphable models with a very deep neural network. In: CVPR. (2017)
21. Richardson, E., Sela, M., Or-El, R., Kimmel, R.: Learning detailed face reconstruction from a single image. In: CVPR. (2017)
22. Dou, P., Shah, S.K., Kakadiaris, I.A.: End-to-end 3d face reconstruction with deep neural networks. In: CVPR. (2017)
23. Kim, H., Zollöfer, M., Tewari, A., Thies, J., Richardt, C., Christian, T.: InverseFaceNet: Deep Single-Shot Inverse Face Rendering From A Single Image. arXiv:1703.10956 (2017)
24. Cao, C., Bradley, D., Zhou, K., Beeler, T.: Real-time high-fidelity facial performance capture. *ACM Trans. Graph.* **34**(4) (2015) 46:1–46:9
25. Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., Webb, R.: Learning from simulated and unsupervised images through adversarial training. In: CVPR. (2017) 2242–2251
26. Mueller, F., Bernard, F., Sotnychenko, O., Mehta, D., Sridhar, S., Casas, D., Theobalt, C.: Generated hands for real-time 3d hand tracking from monocular rgb. In: CVPR. (2018)
27. Yeh, R., Chen, C., Lim, T., Hasegawa-Johnson, M., Do, M.N.: Semantic image inpainting with perceptual and contextual losses. arXiv **1607.07539** (2016)
28. Pathak, D., Krähenbühl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: CVPR. (2016)
29. Alexander, O., Rogers, M., Lambeth, W., Chiang, M., Debevec, P.: The Digital Emily Project: photoreal facial modeling and animation. In: ACM SIGGRAPH Courses, ACM (2009) 12:1–12:15
30. Cao, C., Weng, Y., Zhou, S., Tong, Y., Zhou, K.: Facewarehouse: A 3d facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics* **20**(3) (2014) 413–425
31. Ramamoorthi, R., Hanrahan, P.: A signal-processing framework for inverse rendering. In: SIGGRAPH, ACM (2001) 117–128
32. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: NIPS. (2012) 1097–1105
33. King, D.E.: Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research* **10**(Jul) (2009) 1755–1758
34. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. (2015) 234–241
35. Ma, L., Jia, X., Sun, Q., Schiele, B., Tuytelaars, T., Gool, L.V.: Pose guided person image generation. In: NIPS. (2017) 405–415
36. Ma, L., Sun, Q., Georgoulis, S., Gool, L.V., Schiele, B., Fritz, M.: Disentangled person image generation. In: CVPR. (2018)
37. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. In: ICLR. (2016)
38. Mirjalili, V., Raschka, S., Namboodiri, A.M., Ross, A.: Semi-adversarial networks: Convolutional autoencoders for imparting privacy to face images. In: 2018 International Conference on Biometrics, ICB 2018, Gold Coast, Australia, February 20-23, 2018. (2018) 82–89
39. Oh, S.J., Benenson, R., Fritz, M., Schiele, B.: Person recognition in personal photo collections. In: ICCV. (2015)
40. Oh, S.J., Benenson, R., Fritz, M., Schiele, B.: Person recognition in social media photos. arXiv **1710.03224** (2017)

41. Zhang, N., Paluri, M., Taigman, Y., Fergus, R., Bourdev, L.D.: Beyond frontal faces: Improving person recognition using multiple cues. In: CVPR. (2015)
42. Sun, Q., Schiele, B., Fritz, M.: A domain based approach to social relation recognition. In: CVPR. (2017)
43. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Processing* **13**(4) (2004) 600–612
44. Iizuka, S., Simo-Serra, E., Ishikawa, H.: Globally and locally consistent image completion. *ACM Trans. Graph.* **36**(4) (2017) 107:1–107:14