Semi-supervised Adversarial Learning to Generate Photorealistic Face Images of New Identities from 3D Morphable Model

Imperial College London

Baris Gecer, Binod Bhattarai, Josef Kittler, Tae-Kyun Kim

{b.gecer, b.bhattarai, tk.kim}@imperial.ac.uk, j.kittler@surrey.ac.uk

ECCV 2018

Abstract

- We propose a novel end-to-end adversarial training framework to generate photorealistic face images of new identities constrained by synthetic 3DMM images with identity, pose, illumination and expression diversity. The resulting synthetic face images are visually plausible and can be used to boost face recognition as additional training data or any other graphical purposes.
- We propose a novel semi-supervised adversarial style transfer approach that trains an inverse mapping network as a discriminator with paired synthetic-real images.
- We employ a novel set-based loss function to preserve consistency among unknown identities during GAN training.

Project page: https://github.com/barisgecer/facegan

Related Work

1. **Pixel Level Domain Adaptation:**
   - Supervised[7]/unsuperv.[1] style transfer GANs
2. **Syntetic Training Data:**
   (a) Graphically rendered images [8]
   (b) Simulated Unsupervised GAN [4]
3. **Identity Preservation:**
   (a) Cross entropy loss for known ids [5]
   (b) Set-based loss for unkown ids [6]

References

[1] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.
[2] David Berthelot, Tom Schumm, and Luke Metz. Began: Boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717*, 2017.
[3] Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *ICML*, 2016.
[4] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Josh Susskind, Wenda Wang, and Russ Webb. Learning from simulated and unsupervised images through adversarial training. *CVPR*, 2017.
[5] Xi Yin, Xiang Yu, Kihyuk Sohn, Xiaoming Liu, and Manmohan Chandraker. Towards large-pose face frontalization in the wild. *ICCV*, 2017.
[6] Baris Gecer, Vassileios Balntas, and Tae-Kyun Kim. Learning deep convolutional embeddings for face representation using joint sample-and set-based supervision. In *ICCVW*, 2017.
[7] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NIPS*, 2017.
[8] Iacopo Masi, Anh Tuấn Trần, Tal Hassner, Jatuporn Toy Leksut, and Gérard Medioni. Do we really need to collect millions of faces for effective face recognition. In *ECCV*, 2016.

Overview of the Proposed Framework

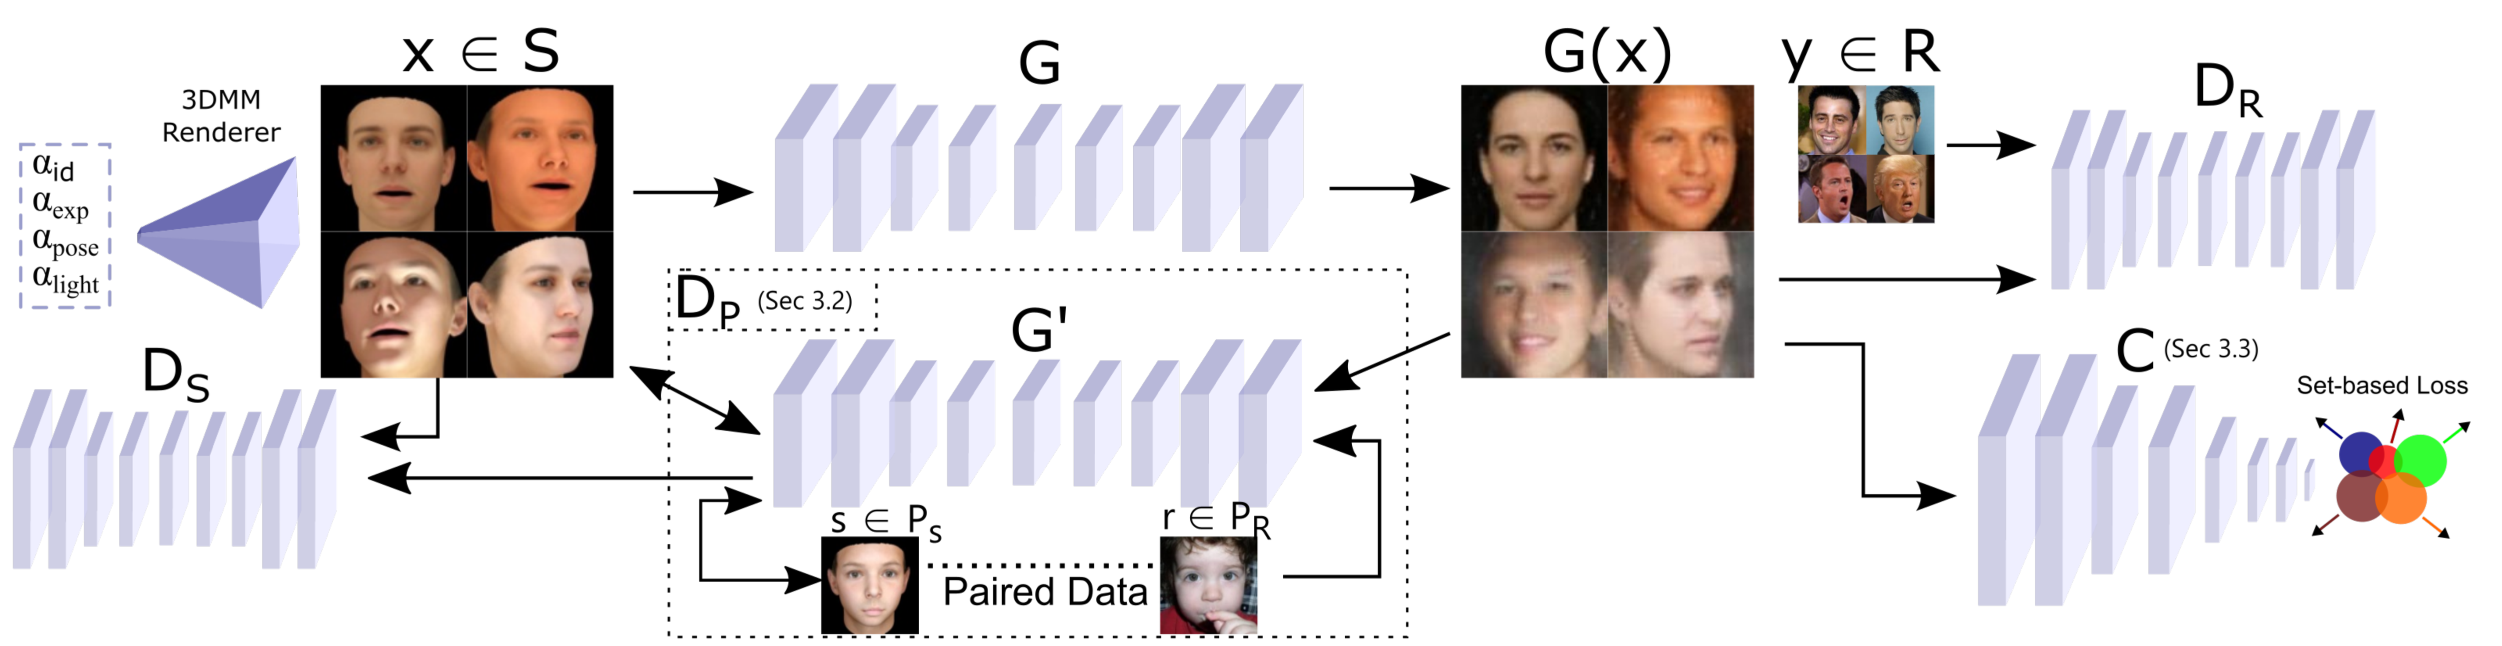

Figure 1: Our approach aims to synthesize photorealistic images ($G(x) \in \hat{\mathcal{R}}$) conditioned by a given synthetic image ($x \in \mathcal{S}$) by 3DMM. It regularizes cycle consistency [1] by introducing an additional adversarial game between the two generator networks (G, G') in an unsupervised fashion. Thus the under-constraint cycle loss is supervised to have correct matching between the two domains by a limited number of paired data ($s \in \mathcal{P_S}, r \in \mathcal{P_R}$). The generator is also encouraged to preserve face identity by a set-based supervision through a pretrained classification network (C).

Proposed Adversarial Identity Generation from 3DMM

- **Synthetic images** ($x \in \mathcal{S}$) **of new IDs** are rendered by sampling from a 3DMM with random pose, expression and lighting attributes

- **Unsupervised Domain Adaptation**
  - Cycle consistency for Syn. -> Real -> Syn.

$$\mathcal{L}_{cyc} = \mathbb{E}_{x \in \mathcal{S}} \| G'(G(x)) - x \|_1$$

  - Adversarial domain transfer by BEGAN discriminators for Real and Syn.

$$\mathcal{L}_G = \mathbb{E}_{x \in \mathcal{S}} \| G(x) - D_R(G(x)) \|_1$$
$$\mathcal{L}_{G'} = \mathbb{E}_{x \in \mathcal{S}} \| G'(G(x)) - D_S(G'(G(x))) \|_1$$
$$\mathcal{L}_{D_R} = \mathbb{E}_{x \in \mathcal{S}, y \in \mathcal{R}} \| y - D_R(y) \|_1 - k_t^{D_R} \mathcal{L}_G$$
$$\mathcal{L}_{D_S} = \mathbb{E}_{x \in \mathcal{S}} \| x - D_S(x) \|_1 - k_t^{D_S} \mathcal{L}_{G'}$$

- **Set-based Identity Preservation**

$$\mathcal{L}_C = \mathbb{E}_{x \in \mathcal{S}, i_x \in \mathbb{N}^+} \sum_x^M -log \frac{\exp(\frac{\|C(G(x)) - c_{i_x}\|_2^2}{2\sigma^2} - \eta)}{\sum_{j \neq i_x} \exp(\frac{\|C(G(x)) - c_j\|_2^2}{2\sigma^2})}$$

- **Adv. Pair Matching** with modified BEGAN
  - $G'(\hat{\mathcal{R}} \rightarrow \hat{\mathcal{S}})$ network functions as a pair-matching discriminator supervising G network to align the correlation distributions of resulting synthetic pairs ($x \in \mathcal{S}, G(x) \in \hat{\mathcal{R}}$) and paired training data ($s \in \mathcal{P_S}, r \in \mathcal{P_R}$).

$$\mathcal{L}_{D_P} = \mathbb{E}_{s \in \mathcal{P_S}, r \in \mathcal{P_R}} \| s - G'(r) \|_1 - k_t^{D_P} \mathcal{L}_{cyc}$$

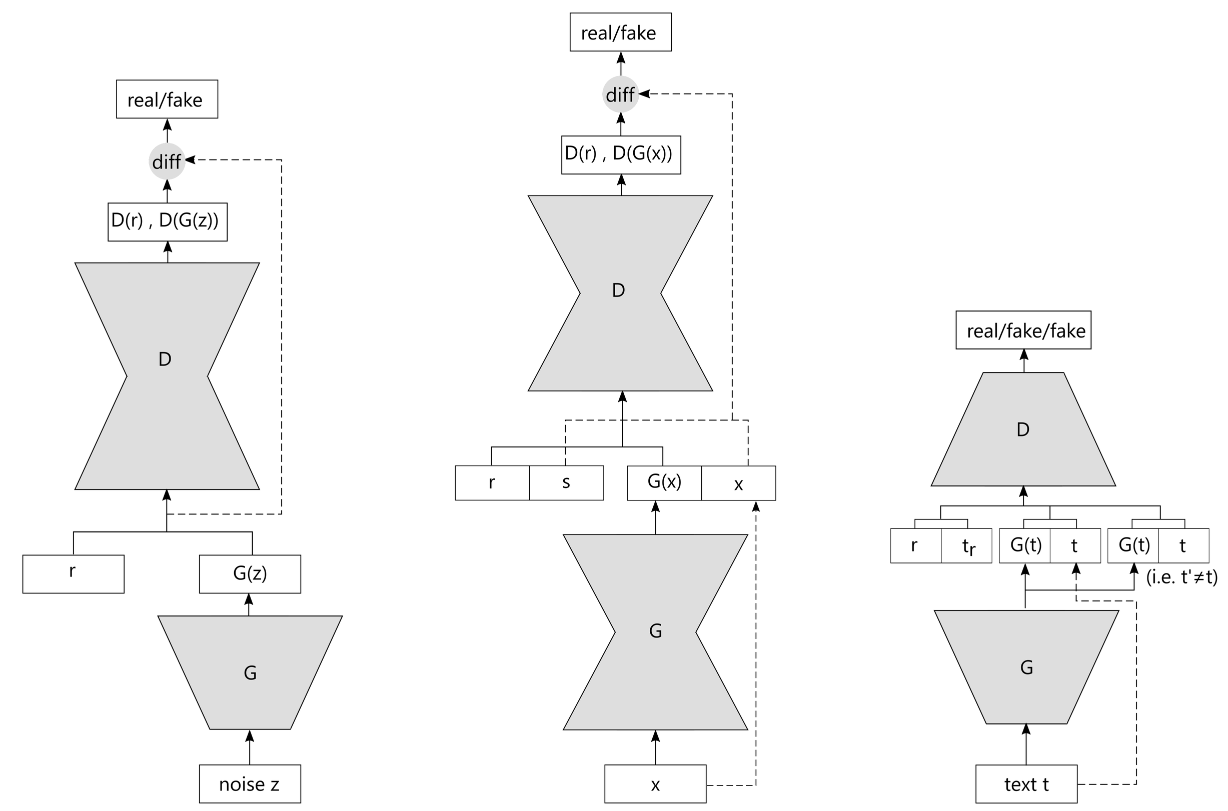

a: BEGAN [2] b: **Adv. Pair Matching** c: GAN-CLS [3]

Figure 2: Quality of 9 images of 3 identities during the training. Background plot shows the error by the proposed identity preservation layer over the iterations. Notice the changes on the level of fine-details on the faces which is the main motivation of using set-based identity preservation.

Qualitative Results - GANFaces

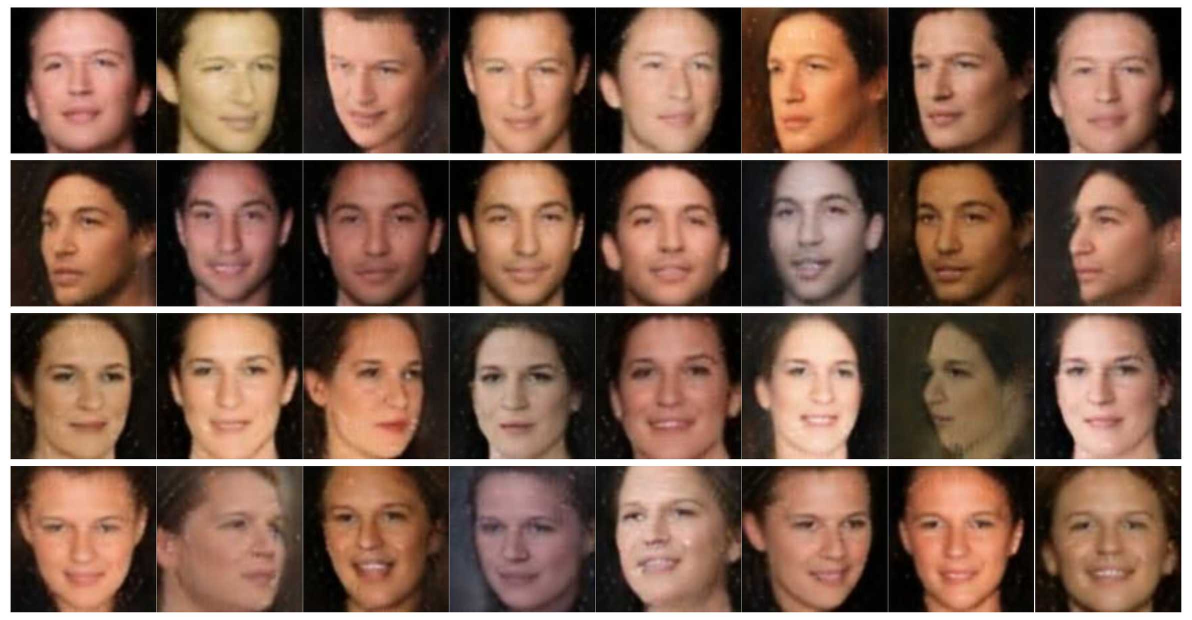

Figure 3: Random samples from GANFaces dataset. Each row belongs to same identity. Notice the variation in pose, expression and lighting.

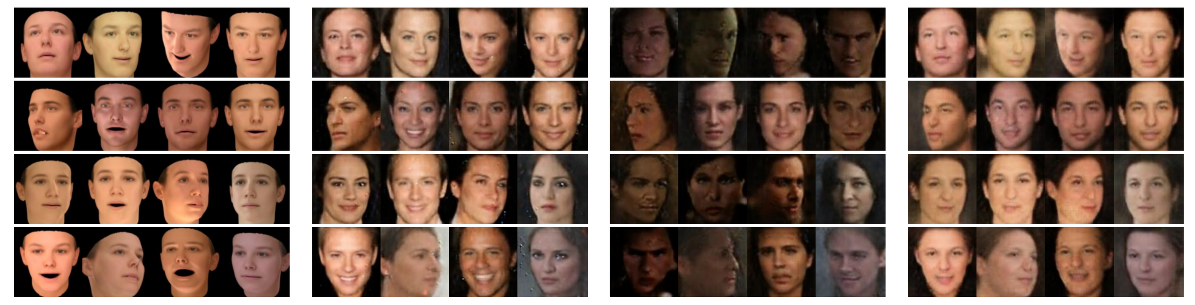

(a) 3DMM inputs (b) simGAN [4] (c) CycleGAN [1] (d) Reconstruction Err.

Figure 4: Comparison to the state-of-the-art studies

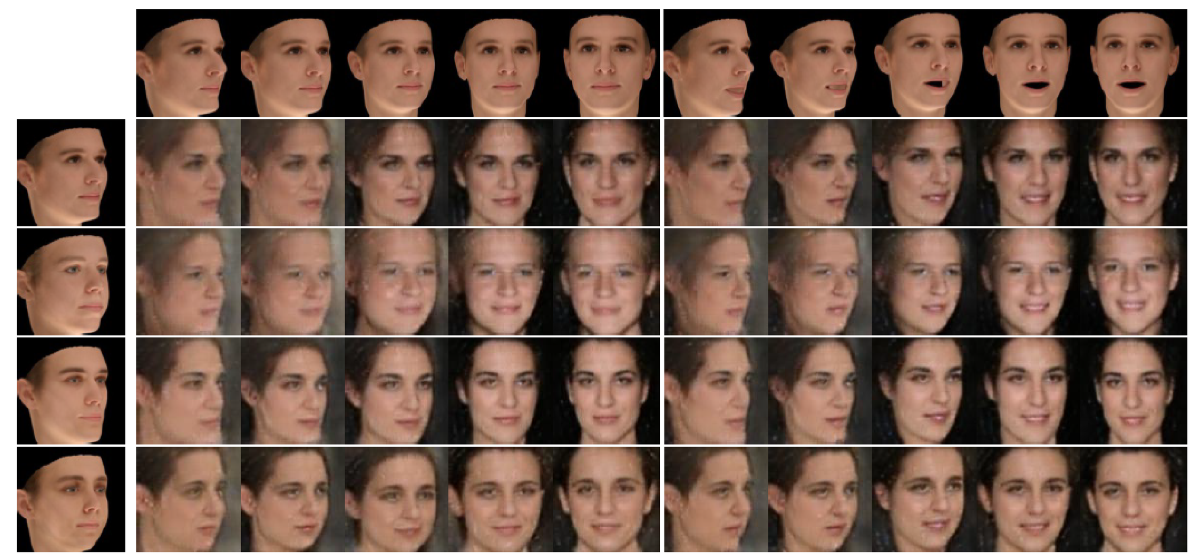

Figure 5: Our model is capable of generating photorealistic images preserving the pose and expression conditioned by the 3DMM input images. Identity variation in vertical axis, normalized and mouth open expression in left and right blocks and pose variation in horizontal axis.

Face Recognition with GANFaces

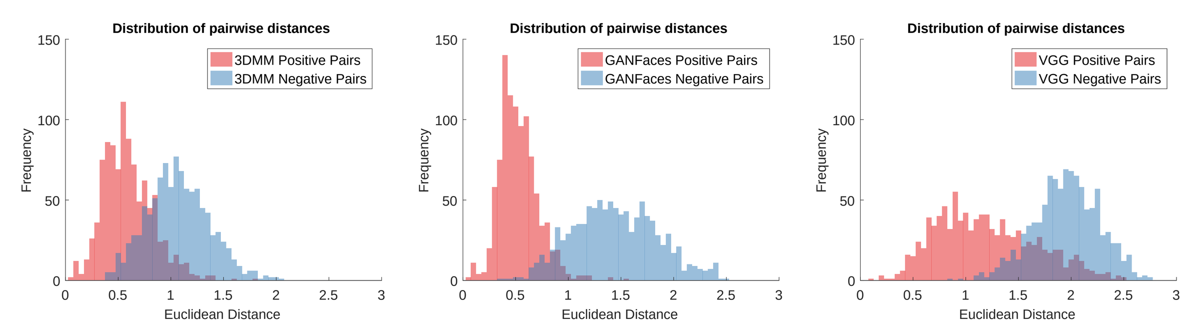

Figure 6: Distances of 1000 positive and 1000 negative pairs from three different datasets (3DMM synthetic images, GANFaces, Oxford VGG)

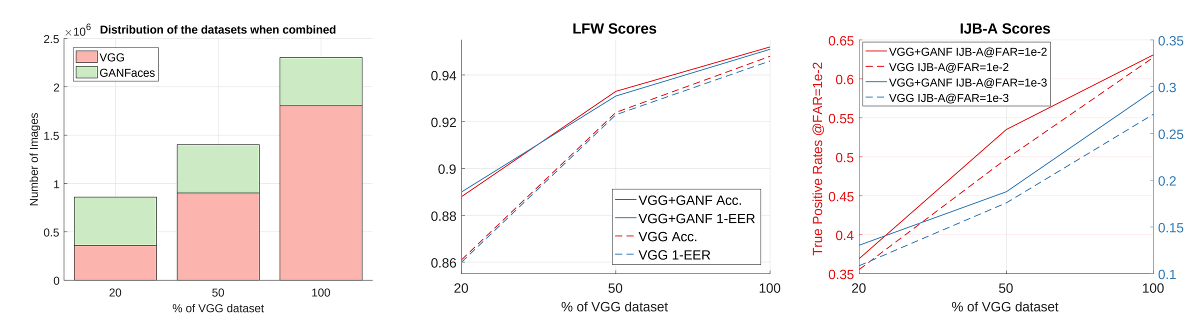

Figure 7: (Left) Number of images from VGG and GANFaces in the experiments. (Middle) Acc. on the LFW with and without GANFaces. (Right) TPRs on IJB-A verification task with and without GANFaces.