

Fast and Robust Face-to-Parameter Translation for Game Character Auto-Creation

Tianyang Shi,¹ Zhengxia Zou,² Yi Yuan,^{1*} Changjie Fan¹

¹NetEase Fuxi AI Lab

²University of Michigan, Ann Arbor

{shitianyang, yuanyi, fanchangjie}@corp.netease.com, zzhengxi@umich.edu

Abstract

With the rapid development of Role-Playing Games (RPGs), players are now allowed to edit the facial appearance of their in-game characters with their preferences rather than using default templates. This paper proposes a game character auto-creation framework that generates in-game characters according to a player’s input face photo. Different from the previous methods that are designed based on neural style transfer or monocular 3D face reconstruction, we re-formulate the character auto-creation process in a different point of view: by predicting a large set of physically meaningful facial parameters under a self-supervised learning paradigm. Instead of updating facial parameters iteratively at the input end of the renderer as suggested by previous methods, which are time-consuming, we introduce a facial parameter translator so that the creation can be done efficiently through a single forward propagation from the face embeddings to parameters, with a considerable 1000x computational speedup. Despite its high efficiency, the interactivity is preserved in our method where users are allowed to optionally fine-tune the facial parameters on our creation according to their needs. Our approach also shows better robustness than previous methods, especially for those photos with head-pose variance. Comparison results and ablation analysis on seven public face verification datasets suggest the effectiveness of our method.

Introduction

The character customization system in many role-playing games (RPGs) provides an interactive interface where players can create their in-game characters according to their preferences. To improve players’ immersion and interactivity, character customization systems in RPGs are becoming more and more sophisticated. In many modern RPGs, e.g., “Grand Theft Auto Online¹”, “Dark Souls III²”, and “Justice³”, players are now allowed to precisely manipulate specific part of their characters, e.g., the corner of the eye, the wing of the nose, hairstyles, and makeups. As a result,

*Corresponding author

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://www.rockstargames.com/GTAOnline>

²<https://www.darksouls.jp>

³<https://n.163.com>




Figure 1: We propose a new method to generate in-game character of the players according to their input face photo. Our method is three magnitudes faster than the previous method (Shi et al. 2019) and shows better robustness (especially for pose changes).

the character customization process turns out to be laborious and time-consuming. To create an in-game character with a desired facial appearance (e.g. a pop star or the players themselves), most players need to spend several hours by manually adjusting hundreds of parameters, even after considerable practice. In this paper, we study an interesting problem, i.e. how to automatically generate in-game characters according to a player’s face photo.

In computer vision, efforts have been made in generating 3D faces based on single input face photo, where some representative methods include the 3DMM (Blanz and Vetter 1999) and 3DMM-CNN (Tran et al. 2017), etc. However, these methods are difficult to be applied in-game environments. This is because most of the above methods are built based on “morphable face models” (Blanz and Vetter 1999) while the faces in most RPGs are rendered by “bone-driven face models”. There is a huge gap between the two different type of face models due to their different rendering mode,

skeletons, and physical meaning of facial parameters.

To solve the character auto-creation problem in RPGs, Shi *et al.* (Shi *et al.* 2019) proposed a method for bone-driven face model called “Face-to-Parameters” (F2P). In the F2P method, a CNN model is trained to make game engine differentiable, and the character auto-creation is formulated as a facial similarity measurement and parameter searching problem. The facial parameters are optimized under a neural style transfer framework (Gatys, Ecker, and Bethge 2016; Gu *et al.* 2018), where they iteratively adjust the parameters at the input end of the renderer by using gradient decent. Despite their promising results, the F2P method still has some limitations in both robustness and processing speed. As the local feature they used dominates the facial similarity measurement, their method is sensitive to pose changes. Besides, since the parameter generation is performed based on an iterative updating process, the F2P suffers from a low inference speed.

In this paper, we propose a new framework to solve the above problems. Different from the previous methods which are designed based on neural style transfer (Shi *et al.* 2019) or monocular 3D face reconstruction (Banz and Vetter 1999; Richardson *et al.* 2017; Tewari *et al.* 2017), we re-formulate the character auto-creation process under a *self-supervised learning* paradigm. Instead of updating facial parameters iteratively at the input end of the renderer, as suggested by Shi *et al.* (Shi *et al.* 2019), we introduce a facial parameter translator so that the optimal facial parameters can be obtained efficiently through a single time forward propagation. This improvement increases the inference speed of the previous method by three orders of magnitude. Despite its high efficiency, the interactivity is also preserved in our method where users are allowed to optionally fine-tune the generated facial parameters according to their needs. To improve the robustness of our framework, especially for pose variance of the input face, instead of directly predicting facial parameters in their original parameter space, we integrated the facial priors by learning their low-dimensional representation and making predictions in their orthogonal subspace. We shed light on how these simple modifications improvement the robustness of our framework.

Fig. 2 shows an overview of the proposed framework. Our framework consists of multiple components:

- **An imitator** G . As the rendering process of most game engines is not differentiable, we train an imitator G to imitate the behavior of a game engine, as suggested by previous methods (Shi *et al.* 2019; Genova *et al.* 2018), and make the rendering process differentiable. In this way, our framework can be smoothly trained in an end-to-end fashion.

- **A translator** T , which aims to transform the input facial image embeddings to output facial parameters. The generated parameters either can be further fine-tuned by players or can be used to render 3D faces in-game environments.

- **A face recognition network** F_{recg} , which encodes an input face image to a set of pose-irrelevant face embeddings.

- **A face segmentation network** F_{seg} , which extracts position-sensitive face representations.

We formulate the training of our framework under a multi-

Table 1: A comparison between the Morphable Face Model (MFM) and Bone-Driven Face Model (BDFM).

	MFM	BDFM
Physical Meaning	ambiguous	explicit
Degree of Freedom	normal	high
Texture Style	real	game-style
Structure	parametrized	bone-driven
Ground Truth	a few	none
Makeup	a few	many

task regression paradigm by constructing multiple self-supervised loss functions. To measure the similarity between an input face and a generated one, we define two loss functions, an “identity loss” \mathcal{L}_{idt} and a “facial content loss” \mathcal{L}_{ctt} for effective facial similarity measurement, where the former one focuses on facial identity (pose-irrelevant) and the later one compute similarity base on pixel-wise representations. To further improve the robustness and stability, we also introduce a “loopback loss” \mathcal{L}_{loop} , which ensures the translator can correctly interpret its own output. Our contributions are summarized as follows:

- 1) We propose a new framework for game character auto-creation under a self-supervised learning paradigm. The training of our method neither requires any ground truth references nor any user interactions.

- 2) Different from previous method (Shi *et al.* 2019) that the facial parameters are optimized iteratively by gradient descent, we introduce a translator T where the prediction of the facial parameter only requires single time forward-propagation. We improve the inference speed over three orders of magnitude.

- 3) Facial priors are integrated as auxiliary constraints by performing whitening transformation on a large set of facial parameters. An orthogonal space with reduced dimensions is used for improving robustness.

Related Work

3D Face Reconstruction

3D Face Reconstruction, as a classical ill-posed problem in computer vision, aims to recover 3D information from a 2D photo. A representative of the early method of this field is the well-known 3D Morphable Model (3DMM), which was originally proposed by Banz *et al.* (Banz and Vetter 1999). They use principal component analysis (PCA) to parameterize the 3D mesh data and then optimize the parameter of the face model to fit input photo (Banz and Vetter 1999). On this basis, some improved methods e.g., BFM09 (Paysan *et al.* 2009) and BFM17 (Gerig *et al.* 2018), were then proposed to improve the shape, texture, and expression of the morphable model. In recent years, deep Convolutional Neural Networks (CNNs) are also introduced to predict the parameters of the morphable model with the help of high-level image representations (Tran *et al.* 2017; Dou, Shah, and Kakadiaris 2017; Jackson *et al.* 2017). Besides, the Generative Adversarial Network (GAN) has

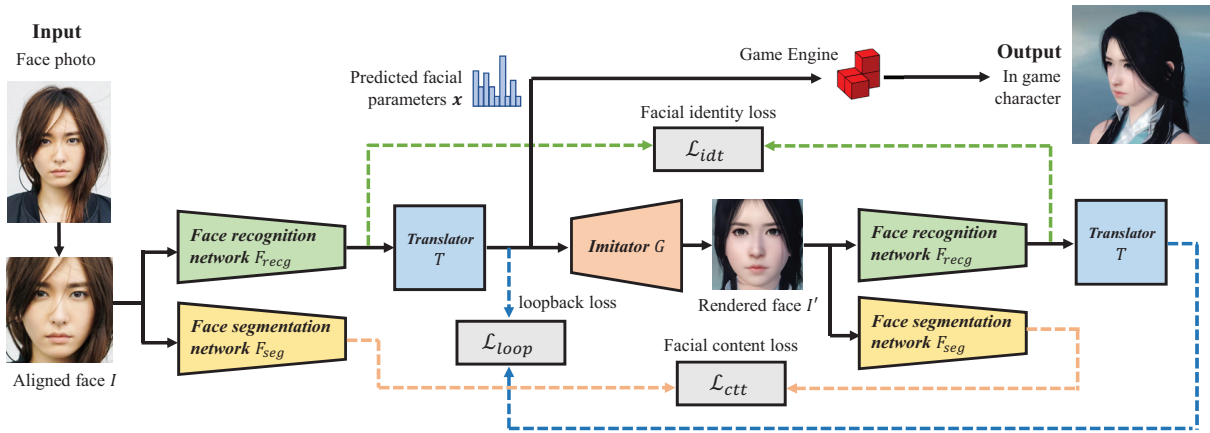


Figure 2: Our model consists of four networks: 1) an Imitator G to imitate the behavior of a game engine, 2) a Translator T to predict facial parameters, 3) a face recognition network F_{recg} to encode an input image to face embeddings, and 4) a face segmentation network F_{seg} to extract position-sensitive facial representations. The key to our method is a self-supervised learning framework where a “recursive consistency” is introduced to enforce the facial representation of the rendered image I' to be similar with that of its input I : $I' = G(T(F_{recg}(I))) \approx I$.

also been introduced to this field recently to generate high-fidelity textures (Gecer et al. 2019).

As the 3DMM-based reconstruction can be essentially considered as a parameter fitting problem between 3D face model and input facial image, making renderer differentiable is the key to solve this problem. With the development of the differentiable rendering technology, self-supervised loss functions can now be used for training neural network without using pre-scanned 3D faces (Tewari et al. 2017; Genova et al. 2018).

But unfortunately, most of the above methods cannot be directly applied to game environments. The reasons are twofold. First, these methods are designed based on morphable face models while in most RPGs, the 3D faces are rendered by the bone-driven face model. Second, the morphable face model is not friendly for user interactions as each of its parameters lacks a clear physical meaning. Table 1 gives a list of differences between these two models.

Character Auto-Creation

Character auto-creation is of great importance for computer games and augmented reality. As an emerging technology, it has drawn increasing attention in recent years (Wolf, Taigman, and Polyak 2017; Shi et al. 2019). Wolf *et al.* (Wolf, Taigman, and Polyak 2017) first proposed an adversarial training based method named Tied Output Synthesis (TOS) for creating parameterized avatars based on face photos, in which they take advantage of adversarial training to select the facial components from a set of pre-defined templates. However, this method is designed to predict discrete attributes rather than continuous facial parameters. A similar work to ours is “Face-to-Parameter (F2P)” method (Shi et al. 2019) proposed by Shi *et al.*, where they build their method based on the bone-driven face model and frame the character creation as a neural style transfer process. To make the renderer differentiable, they train a generative network to imi-

tate the input-output behavior of a game engine. They also introduce two discriminative networks to measure and maximize the similarity between the input photo and the generated character.

Methodology

Our model consists of four neural networks: a imitator G , a facial parameter translator T , a pre-trained face recognition network F_{recg} , and a face segmentation network F_{seg} , as shown in Fig. 2.

Imitator

We train a convolutional neural network as our imitator to fit the input-output mapping of a game engine so that to make the character customization system differentiable. We take the similar configuration with the generator of DCGAN (Radford, Metz, and Chintala 2015) in our imitator $G(x)$, which also has been used in F2P (Shi et al. 2019). We frame the training of our imitator as a standard regression problem, where we aim to minimize the difference between the in-game rendered image and the generated one in their pixel space. When training the imitator, we use pixel-wise l_1 loss function as it encourages less blurring effect:

$$\mathcal{L}_G(x) = E_{x \sim u(x)} \{ \|G(x) - \text{Engine}(x)\|_1 \}, \quad (1)$$

where x is the facial parameters sampled from a uniform distribution. Given x as an input, the $G(x)$ and $\text{Engine}(x)$ are the outputs of our imitator and a game engine (ground truth), respectively. For simplicity, our imitator G only fits the front view of the facial model with the corresponding facial parameters.

Translator

We train a multi-layer perceptron as our translator T to map the face embeddings f_{recg} (the output of our face recogni-

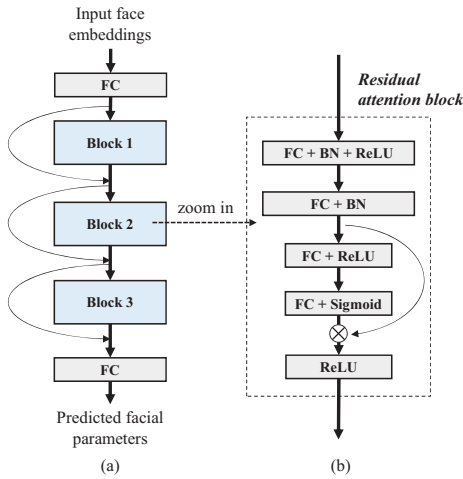


Figure 3: (a) The architecture of our translator T . Our translator consists of three residual attention blocks and two fully connected layers (FC). (b) The details of the residual attention block. \otimes denotes element-wise product between two neural activations.

tion network) to facial parameters \mathbf{x} .

$$\mathbf{x} = T(\mathbf{f}_{recg}) = T(F_{recg}(I)), \quad (2)$$

where F_{recg} is the face recognition network, and I is an input face photo. Since the face embeddings correspond to a global description of the input identity while the facial parameters depict local facial details, to learn better correspondence of the two fields, we introduce an attention-based module as the core building block of our translator. Our translator consists of three residual attention blocks for generating internal representations and two Fully Connected (FC) layers for input and output. The details of our translator are shown in Fig. 3.

The attention mechanism was originally proposed in machine translation to improve the performance of an RNN model (Bahdanau, Cho, and Bengio 2014). It has now been widely used in many computer vision tasks, such as object detection (Zhang, Yang, and Schiele 2018), image super-resolution (Zhang et al. 2018), etc. To introduce the attention mechanism in our translator, we compute the element-wise importance of the neurons and then perform feature re-calibration by multiplying the representations with them. We make a simple modification of the squeeze and excitation block in SENet (Hu, Shen, and Sun 2018) to apply it in an FC layer (the global pooling layer thus is removed). Specifically, let $\mathbf{v} \in \mathbb{R}^{c \times 1}$ represents a c dimensional internal representation of our translator. We use a gating function to learn the element-wise attention weights, which can be written as:

$$\alpha = \sigma(\mathbf{W}_2 \delta(\mathbf{W}_1 \mathbf{v})), \quad (3)$$

where $\sigma(\cdot)$ and $\delta(\cdot)$ denote the sigmoid and ReLU activation function, respectively. \mathbf{W}_1 and \mathbf{W}_2 are the weights of two FC layers. Then, internal representation \mathbf{v} is re-scaled as follows:

$$\tilde{v}_k = \alpha_k v_k, \quad (4)$$



(a) Random character creation *with facial priors*



(b) Random characters *without facial priors*

Figure 4: Randomly generated characters (a) with and (b) without integrating facial priors. The facial priors can be considered as an additional constraint on the predicted parameters and prove to help generate more stable and meaningful characters.

where $\alpha_k, k = 1, \dots, c$ are the learned attention weights.

Integrating facial priors

To further improve the robustness of our framework, instead of predicting facial parameters directly in the parameter space, we perform a whitening transform on the facial parameters and predict in their intrinsic low-dimensional subspace. This can be considered as an integration of facial priors or an additional constraint on the predicted parameters. We apply F2P (Shi et al. 2019) on CelebA (Liu et al. 2015), a large-scale face attributes dataset with over 200k facial images and obtain a large set of facial parameters⁴.

Suppose \mathbf{m} is the mean of facial parameters and \mathbf{P} is the projection matrix, which can be easily obtained by performing singular value decomposition on the parameter data. Before feeding the facial parameter \mathbf{x} into the imitator, we first project it into a subspace by multiplying with the projection matrix (whitening transformation with reduced dimensions):

$$\hat{\mathbf{x}} = \mathbf{P}^T (\mathbf{x} - \mathbf{m}). \quad (5)$$

The final predicted facial parameters can be reconstructed as follows:

$$\mathbf{x} = (\mathbf{P}\mathbf{P}^T)^{-1} \mathbf{P} \hat{\mathbf{x}} + \mathbf{m}. \quad (6)$$

Fig. 4 gives an illustration of the importance of facial prior and how it affects the rendering results. Fig. 5 plots the reconstruction energy with different number of subspace dimension. With the above operation, the dimension of original facial parameters can be significantly reduced and the model becomes much easier to train.

Facial Similarity Measurement

After we obtain a well-trained imitator G and a facial parameter translator T , the parameter generation essentially

⁴Since F2P is sensitive to head pose, we only use the frontal faces in CelebA training set ($\sim 18k$ images) to learn the whitening matrix.

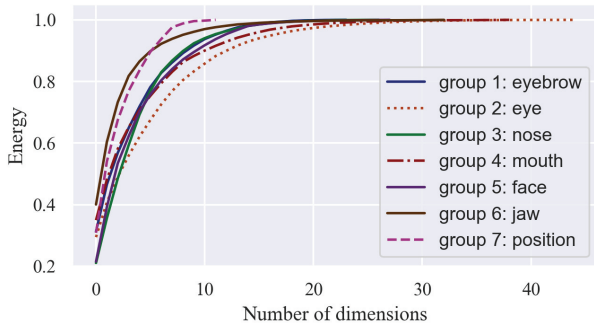


Figure 5: Reconstruction energy vs. Number of principal components in the parameter space of CelebA dataset. Regular faces are embedded in a low-dimensional subspace. We divide the facial parameters into seven groups and perform PCA dimension reduction accordingly.

becomes a face similarity measurement problem. The key to our self-supervised learning framework is to *enforce* the facial representation of the rendered image I' to be similar to that of its input face photo I :

$$I' = G(T(F_{recg}(I))) \approx I. \quad (7)$$

As the input face photo and the rendered game character belong to different image domains, we introduce additional two networks, a pre-trained face recognition network F_{recg} and a facial segmentation network F_{seg} , to learn effective facial similarity measurement in terms of both identity and local details.

Facial identity loss. We use a popular face recognition network named LightCNN-29v2 (Wu et al. 2018) as our face recognition network. We use its backbone to extract 256-d face embeddings as a representation of facial identity. We define the identity loss of two faces as the cosine distance on their embeddings:

$$\mathcal{L}_{idt} = 1 - e_1^T e_2 / \sqrt{\|e_1\|_2 \|e_2\|_2}, \quad (8)$$

where $e_i, i = 1, 2$ are the face embeddings of an input face photo I and a rendered face image I' :

$$e_1 = F_{recg}(I), \quad e_2 = F_{recg}(I'). \quad (9)$$

Facial content loss. In addition to the facial identity loss, we also define a content loss by computing pixel-wise distance based on the facial representations extracted from a pre-trained face semantic segmentation model. The motivation behind is intuitive, i.e., if two faces share the same displacement of different facial components, e.g. the contour of the face, the eyes, and nose, they will share a high similarity regardless of the image domains. We build our facial segmentation model F_{seg} based on Resnet-50 (He et al. 2016). We remove its fully connected layers and increase its output resolution from 1/32 to 1/8. We train this model on Helen face semantic segmentation dataset (Le et al. 2012) and then freeze its weights as a fixed feature extractor. We define the facial content loss on the pixel-wise l_1 distance between the semantic representations of two faces:

$$\mathcal{L}_{ctt} = E_k \{\|f_1 - f_2\|_1\}, \quad (10)$$

where k is the pixel location of the feature map. $f_i, i = 1, 2$ are the facial semantic features of the image I and I' :

$$f_1 = F_{seg}(I), \quad f_2 = F_{seg}(I'). \quad (11)$$

Loopback loss. Inspired by the unsupervised 3D face reconstruction method proposed by Genova *et al.* (Genova et al. 2018), we also introduce a “loopback loss” to further improve the robustness of our parameter output. After we obtain the face embeddings of the rendered image $F_{recg}(I')$, we further feed it into our translator T to produce a set of new parameters $x' = T(F_{recg}(I'))$ and force the generated facial parameters before and after the loop unchanged. The loopback loss can be written as follows:

$$\mathcal{L}_{loop} = \|x - T(F_{recg}(I'))\|_1. \quad (12)$$

The final loss function in our model can be written as the summary of the above three losses:

$$\mathcal{L}(G, T, F_{recg}, F_{seg}) = \lambda_1 \mathcal{L}_{idt} + \lambda_2 \mathcal{L}_{ctt} + \lambda_3 \mathcal{L}_{loop}, \quad (13)$$

where $\lambda_i > 0, i = 1, 2, 3$ control the balance between them. Since our imitator G , face recognition network F_{recg} and face segmentation networks F_{seg} are all pre-trained models. All we need to train is the translator T . We aim to solving the following optimization problem:

$$T^* = \operatorname{argmin}_T \mathcal{L}(G, T, F_{recg}, F_{seg}). \quad (14)$$

A complete implementation pipeline of our method is summarized as follows:

- **Training (I).** Train the imitator G by minimizing Eq (1). Train the face recognition network F_{recg} and the face segmentation network F_{seg} .
- **Training (II).** Fix G, F_{recg}, F_{seg} , and train the translator T according to Eq (14).
- **Inference.** Given an input photo I , predict the facial parameters based on Eq (2): $x = T(F_{recg}(I))$.

Training Details

We use the SGD optimizer with momentum to train G and F_{seg} with the learning rate of 10^{-2} and 10^{-3} , respectively. We set `batch_size = 16`, `momentum = 0.9`. The learning rate decay is set to 10% per 50 epochs and the training stops at 500 epochs. We then freeze above networks, set $\lambda_1 = 0.01$, $\lambda_2 = 1$, and $\lambda_3 = 1$, and use the Adam optimizer (Kingma and Ba 2014) to train T with the learning rate of 10^{-4} and max-iteration of 20 epochs.

We use the CelebA dataset (Liu et al. 2015) to train our translator T . We set $\lambda_2 = 0$ every 4 training steps. When the λ_2 is set to 0, we update the T by sampling from the full CelebA training set, while when the λ_2 is set > 0 , we update the T by sampling from a subset of the CelebA training set, which only contains high-quality frontal faces. In this way, our translator not only can be well-trained on frontal images with the content loss, but also can generalize to the rest images with the help of the identity loss, which is pose-invariant.

The face alignment is performed by using dlib library (King 2009) to align the input photo before it is fed into our framework.

Table 2: Performance comparison. We compare our method with two previous methods on seven face verification datasets. Higher scores indicate better. (Face embeddings are normalized by PCA as 3DMM-CNN for a fair comparison)

Method	Datasets							Speed*
	LFW	CFP_FF	CFP_FP	AgeDB	CALFW	CPLFW	Vggface2_FP	
3DMM CNN	0.9235	-	-	-	-	-	-	$\sim 10^2$ Hz
F2P	0.6977	0.7060	0.5800	0.6013	0.6547	0.6042	0.6104	~ 1 Hz
Ours	0.9402	0.9450	0.8236	0.8408	0.8463	0.7652	0.8190	$\sim 10^3$ Hz
LightCNN-29v2**	0.9958	0.9940	0.9494	0.9597	0.9433	0.8857	0.9374	$\sim 10^3$ Hz

* Inference time under GTX 1080Ti. The time cost for model loading and face alignment are not considered.

** Here we use the performance of LightCNN-29v2 on input photos as a reference. (upper-bound accuracy)

Table 3: Ablations analysis on different technical components of our method. The integration of each technique yields consistent improvements in face verification accuracy.

Ablations				Datasets							
\mathcal{L}_{ctt}	\mathcal{L}_{idt}	\mathcal{L}_{loop}	ResAtt	LFW	CFP_FF	CFP_FP	AgeDB	CALFW	CPLFW	Vggface2_FP	
✓	×	×	✓	0.7880	0.7930	0.6666	0.6868	0.6792	0.6252	0.6696	
✓	✓	×	✓	0.8843	0.8777	0.7507	0.7917	0.7675	0.7032	0.7432	
✓	✓	✓	×	0.8870	0.8901	0.7626	0.7875	0.7725	0.7042	0.7618	
✓	✓	✓	✓	0.9243	0.9200	0.7896	0.8152	0.8130	0.7400	0.7854	
LightCNN-29v2*				0.9948	0.9939	0.9476	0.9537	0.9438	0.8872	0.9326	

* Here we use the performance of LightCNN-29v2 on input photos as a reference. (upper-bound accuracy)

Experiments

We test our method on two games with East Asian faces, “Justice” and “Revelation⁵”, where the former one is a PC game and the latter one is a mobile game on Android/iOS devices. We use seven well-known face verification datasets to quantitatively evaluate effectiveness of our method, including LFW (Huang et al. 2008), CFP_FF (Sengupta et al. 2016), CFP_FP (Sengupta et al. 2016), AgeDB (Moschoglou et al. 2017), CALFW (Zheng, Deng, and Hu 2017), CPLFW (Zheng and Deng 2018), and Vggface2_FP (Cao et al. 2018).

Comparison with other methods

We compare our method with some other approaches, including the F2P (Shi et al. 2019) and 3DMM-CNN (Tran et al. 2017) on both their accuracy and speed. Fig. 6 shows a comparison of generated characters of different methods. Our method creates game characters with high similarity in terms of both global appearance and local details while the 3DMM method can only generate masks with similar facial outlines.

To quantitatively evaluate the generated results, we follow the 3DMM-CNN (Tran et al. 2017), where the face verification accuracy is used as the evaluation metric. To do this, we first generate the facial parameters for every input face, then we use the face verification benchmark toolkit “face.evoLve” (Zhao et al. 2019a; 2019b; 2018) to compute the verification accuracy based on the generated parameters. Clearly, if the two photos belong to the same person, they should have similar facial parameters. Table 2 shows

⁵<https://tym.163.com/>

Table 4: Subjective evaluation.

Method	Selection Ratio	LightCNN	evoLVE
3DMMCNN	17.4% \pm 1.2%	0.13	0.06
F2P	33.9% \pm 1.2%	0.30	0.18
Ours	48.7% \pm 1.1%	0.34	0.22

the verification scores of different methods. A higher score suggests a better performance. The performance of 3DMM-CNN is reported by Tran *et al.*’s paper (Tran et al. 2017).

We further follow the subjective evaluation in F2P (Shi et al. 2019) by inviting 15 volunteers to rank the results generated by three methods, and finally calculate the Selection Ratio for each method. Due to the large number of images, we only randomly collect 50 images from CelebA test set for evaluation. Table 4 shows the subjective evaluation results and cosine similarities measured by face recognition networks (LightCNN-29v2 and face.evoLVE-IR50 (Zhao et al. 2019a)). Our method achieves consistent improvements over previous methods.

Ablation Studies

The ablation experiments are conducted to verify the importance of each component in our framework, including:

- 1) the facial identity loss \mathcal{L}_{idt} , (Eq. 8);
- 2) the facial content loss \mathcal{L}_{ctt} , (Eq. 10);
- 3) the loopback loss \mathcal{L}_{loop} , (Eq. 12);
- 4) the residual attention block (ResAtt) in our facial parameter translator, (Fig. 3).




Figure 6: A comparison with other methods: 3DMM-CNN and F2P. The first three rows show results on the PC game “Justice” and the last row shows the results on the mobile game “Revelation”. Our method achieves better creation results.

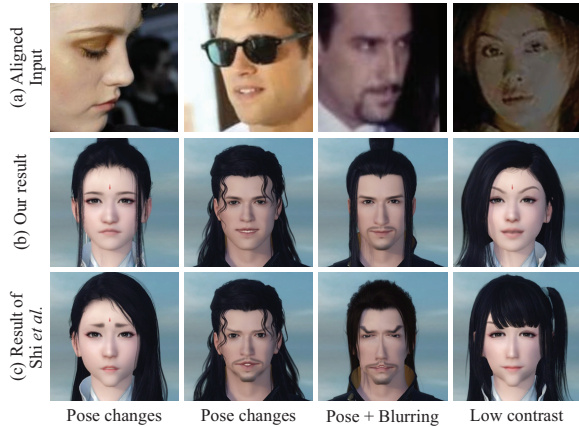


Figure 7: Comparison results on robustness.

We use the same evaluation metric as we used in our comparison experiment. We first evaluate our baseline method, where we train our model only based on the facial content loss, then we gradually add other technical components. Table 3 shows their performance on different datasets. The integration of the facial identity loss, loopback loss, and the residual attention mechanism in our translator yield consistent improvements in the face verification accuracy. Be-

sides, the identity loss brings noticeable improvement for our method. This is because the face recognition network F_{recg} and the face verification metric all focus on the facial identities. We also can observe a significant improvement by introducing the attention mechanism to our translator rather than by simply applying the vanilla multi-layer perceptron which is used in Genova’s method (Genova et al. 2018).

Experiment on robustness

We further evaluate our method on different blurring conditions, different illuminations conditions, and the faces with different poses. Fig. 7 shows the comparison results. Our method proves to be robust to these changes and has better robustness than the F2P, especially for the pose variants.

Conclusion

This paper proposes a new framework for game character auto-creation, i.e., to generate game characters of players based on their input face photos. Different from the previous methods that are designed based on neural style transfer and predicting facial parameters based on an iterative inference pipeline, we construct our method under a self-supervised learning paradigm and introduce a facial parameter translator so that the prediction can be done efficiently through a single forward propagation. Our method achieves a considerable 1000x computational speedup. Our method also

proves to be more robust than previous methods, especially for the pose variants. Comparison results and ablation analysis on seven benchmark datasets suggest the effectiveness of our method.

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Blanz, V., and Vetter, T. 1999. A morphable model for the synthesis of 3d faces. In *Siggraph*, volume 99, 187–194.
- Cao, Q.; Shen, L.; Xie, W.; Parkhi, O. M.; and Zisserman, A. 2018. Vggface2: A dataset for recognising faces across pose and age. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, 67–74. IEEE.
- Dou, P.; Shah, S. K.; and Kakadiaris, I. A. 2017. End-to-end 3d face reconstruction with deep neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Gatys, L. A.; Ecker, A. S.; and Bethge, M. 2016. Image style transfer using convolutional neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Gecer, B.; Ploumpis, S.; Kotsia, I.; and Zafeiriou, S. 2019. Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1155–1164.
- Genova, K.; Cole, F.; Maschinot, A.; Sarna, A.; Vlastic, D.; and Freeman, W. T. 2018. Unsupervised training for 3d morphable model regression. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Gerig, T.; Morel-Forster, A.; Blumer, C.; Egger, B.; Luthi, M.; Schönborn, S.; and Vetter, T. 2018. Morphable face models—an open framework. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, 75–82. IEEE.
- Gu, S.; Chen, C.; Liao, J.; and Yuan, L. 2018. Arbitrary style transfer with deep feature reshuffle. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7132–7141.
- Huang, G. B.; Mattar, M.; Berg, T.; and Learned-Miller, E. 2008. Labeled faces in the wild: A database for studying face recognition in unconstrained environments.
- Jackson, A. S.; Bulat, A.; Argyriou, V.; and Tzimiropoulos, G. 2017. Large pose 3d face reconstruction from a single image via direct volumetric cnn regression. In *Proceedings of the IEEE International Conference on Computer Vision*, 1031–1039.
- King, D. E. 2009. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research* 10:1755–1758.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Le, V.; Brandt, J.; Lin, Z.; Bourdev, L.; and Huang, T. S. 2012. Interactive facial feature localization. In *European conference on computer vision*, 679–692. Springer.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Moschoglou, S.; Papaioannou, A.; Sagonas, C.; Deng, J.; Kotsia, I.; and Zafeiriou, S. 2017. Agedb: the first manually collected, in-the-wild age database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 51–59.
- Paysan, P.; Knothe, R.; Amberg, B.; Romdhani, S.; and Vetter, T. 2009. A 3d face model for pose and illumination invariant face recognition. In *2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance*, 296–301. Ieee.
- Radford, A.; Metz, L.; and Chintala, S. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- Richardson, E.; Sela, M.; Or-El, R.; and Kimmel, R. 2017. Learning detailed face reconstruction from a single image. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Sengupta, S.; Chen, J.-C.; Castillo, C.; Patel, V. M.; Chellappa, R.; and Jacobs, D. W. 2016. Frontal to profile face verification in the wild. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 1–9. IEEE.
- Shi, T.; Yuan, Y.; Fan, C.; Zou, Z.; Shi, Z.; and Liu, Y. 2019. Face-to-parameter translation for game character auto-creation. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Tewari, A.; Zollhofer, M.; Kim, H.; Garrido, P.; Bernard, F.; Perez, P.; and Theobalt, C. 2017. Mofa: Model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Tran, A.; Hassner, T.; Masi, I.; and Medioni, G. 2017. Regressing robust and discriminative 3d morphable models with a very deep neural network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wolf, L.; Taigman, Y.; and Polyak, A. 2017. Unsupervised creation of parameterized avatars. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Wu, X.; He, R.; Sun, Z.; and Tan, T. 2018. A light cnn for deep face representation with noisy labels. *IEEE Transactions on Information Forensics and Security* 13(11):2884–2896.
- Zhang, Y.; Li, K.; Li, K.; Wang, L.; Zhong, B.; and Fu, Y. 2018. Image super-resolution using very deep residual channel attention networks. In *Proceedings of the European Conference on Computer Vision*, 286–301.
- Zhang, S.; Yang, J.; and Schiele, B. 2018. Occluded pedestrian detection through guided attention in cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6995–7003.
- Zhao, J.; Cheng, Y.; Xu, Y.; Xiong, L.; Li, J.; Zhao, F.; Jayashree, K.; Pranata, S.; Shen, S.; Xing, J.; et al. 2018. Towards pose invariant face recognition in the wild. In *CVPR*, 2207–2216.
- Zhao, J.; Cheng, Y.; Cheng, Y.; Yang, Y.; Lan, H.; Zhao, F.; Xiong, L.; Xu, Y.; Li, J.; Pranata, S.; et al. 2019a. Look across elapse: Disentangled representation learning and photorealistic cross-age face synthesis for age-invariant face recognition. In *AAAI*.
- Zhao, J.; Li, J.; Tu, X.; Zhao, F.; Xin, Y.; Xing, J.; Liu, H.; Yan, S.; and Feng, J. 2019b. Multi-prototype networks for unconstrained set-based face recognition. In *IJCAI*.
- Zheng, T., and Deng, W. 2018. Cross-pose lfw: A database for studying crosspose face recognition in unconstrained environments. *Beijing University of Posts and Telecommunications, Tech. Rep* 18–01.
- Zheng, T.; Deng, W.; and Hu, J. 2017. Cross-age lfw: A database for studying cross-age face recognition in unconstrained environments. *arXiv preprint arXiv:1708.08197*.