

SwiftAvatar: Efficient Auto-Creation of Parameterized Stylized Character on Arbitrary Avatar Engines

Shizun Wang^{1*}, Weihong Zeng^{2*}, Xu Wang², Hao Yang², Li Chen², Chuang Zhang¹,
Ming Wu^{1†}, Yi Yuan², Yunzhao Zeng², Min Zheng², Jing Liu²

¹ Beijing University of Posts and Telecommunications

² Douyin Vision

{wangshizun, zhangchuang, wuming}@bupt.edu.cn

{zengweihong, wangxu.aialab, yang.hao, chenli.phd, yuanyi.cv, zengyunzhao, zhengmin.666, jing.liu}@bytedance.com

Abstract

The creation of a parameterized stylized character involves careful selection of numerous parameters, also known as the “avatar vectors” that can be interpreted by the avatar engine. Existing unsupervised avatar vector estimation methods that auto-create avatars for users, however, often fail to work because of the domain gap between realistic faces and stylized avatar images. To this end, we propose *SwiftAvatar*, a novel avatar auto-creation framework that is evidently superior to previous works. *SwiftAvatar* introduces dual-domain generators to create pairs of realistic faces and avatar images using shared latent codes. The latent codes can then be bridged with the avatar vectors as pairs, by performing GAN inversion on the avatar images rendered from the engine using avatar vectors. Through this way, we are able to synthesize paired data in high-quality as many as possible, consisting of avatar vectors and their corresponding realistic faces. We also propose semantic augmentation to improve the diversity of synthesis. Finally, a light-weight avatar vector estimator is trained on the synthetic pairs to implement efficient auto-creation. Our experiments demonstrate the effectiveness and efficiency of *SwiftAvatar* on two different avatar engines. The superiority and advantageous flexibility of *SwiftAvatar* are also verified in both subjective and objective evaluations.

1 Introduction

The emerging of the Metaverse concept is alongside with the wide usage of virtual avatars, embodiment of Metaverse users in various styles, which are popular in modern digital lives such as socialization, e-shopping and gaming. Even though many avatar platforms, such as Zepeto, BitMoji and ReadyPlayerMe, have enabled users to create their own stylized avatars by specifying avatar vectors that can be interpreted by their avatar engines, the manual creation process is tiresome and time consuming, especially when the engine provides with a large set of options making up a long avatar vector. Besides, the avatar vectors to specify usually consists of parameters in both continuous forms (to control facial shape, eye spacing, etc.) and discrete forms (to control hair styles, wearings etc.), all requiring careful selection and

cautious adjustment to achieve a satisfactory result. Therefore, it is valuable to study *how to automatically create a stylized avatar that best matches the user’s input selfie*.

A straightforward way to realize avatar auto-creation is using supervised learning: a network is trained on labeled data to predict avatar vectors based on real face inputs. However, this requires large amount of data collection and manual labeling, which is laborious, expensive, and is not generalizable across engines. Because the definition of avatar vectors and assets vary from engine to engine, data labeled for one engine can not be used to train on other engines.

Several unsupervised learning methods have been proposed to address avatar auto-creation without using any labeled data, including Tied Output Synthesis (TOS) (Wolf, Taigman, and Polyak 2017) and the Face-to-Parameter (F2P) series (Shi et al. 2019, 2020). The main idea of these works can be abstracted by Fig. 2-a. In order to achieve an avatar vector that renders avatar image as similar to the input face as possible, these methods impose constraints on the image-level. They suffer from issues as follows: 1) The image-level similarity constraints they establish are designed for realistic avatar images, not applicable to stylized avatar images that have domain gap with real face images. 2) Leveraging the image-level supervision requires that the avatar rendering process is differentiable. So they usually introduce an imitator network that imitates the behavior of the non-differentiable avatar engine. However, the un-avoidable deviation of imitators from the original avatar engine, as illustrated by Fig. 2-a, degrades the accuracy of the similarity measure. 3) Some approaches like F2P (Shi et al. 2019) need iterative optimization to guarantee the quality of estimated avatar vectors, which is time-consuming in inference.

To address these issues, we propose *SwiftAvatar*, a novel avatar auto-creation framework shown by Figure. 2-b). Unlike previous works that use similarity constraints to find the avatar vector whose rendered image best matches a given face, the core idea of our framework is cross-domain data synthesis. *SwiftAvatar* is able to synthesize pairs of avatar vectors and corresponding realistic faces in high fidelity as many as possible. They are used to train a light-weight estimator that directly predicts avatar vectors from input selfie. Specifically, the *SwiftAvatar* framework consists of three components: *dual-domain generators*, a pipeline for *cross-*

*These authors contributed equally.

†Corresponding Author

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

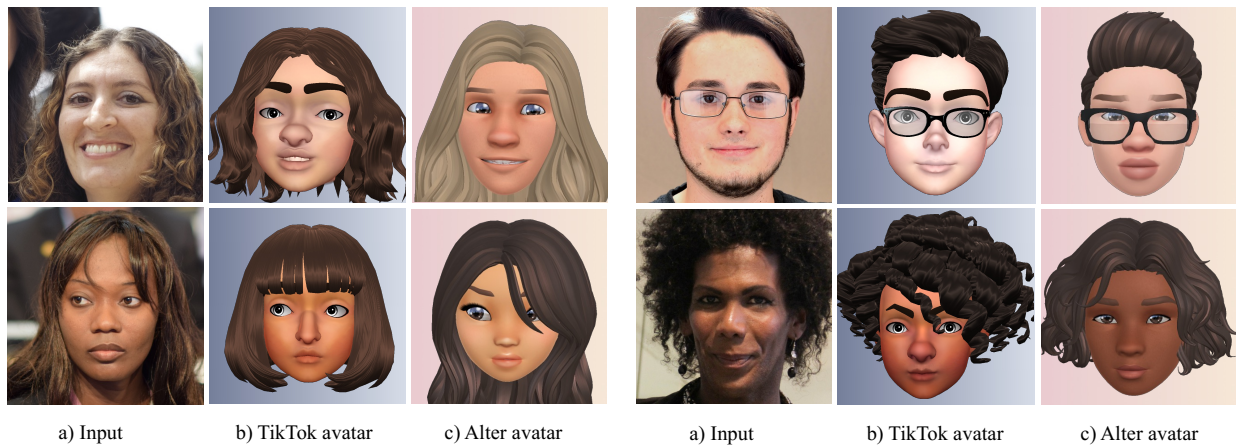


Figure 1: a) Given user’s front-facing selfie images, SwiftAvatar is able to auto-create corresponding stylized avatars on any arbitrary avatar engine, for example, b) the TikTok avatar engine ¹, or c) the Alter avatar engine ².

domain paired data production, and an *avatar estimator*.

The dual-domain generators comprise a realistic generator and an avatar generator, both adopting an architecture from SemanticStyleGAN (Shi et al. 2022). The realistic generator is pretrained on real faces, but the avatar generator is transfer-learned on engine-rendered avatar images with color consistency constraints. So that given a same latent code, the two generators could generate a realistic face image and an avatar image that naturally look similar. Data synthesis relies on the dual-domain generators. The production of synthetic data starts from randomly sampled avatar vectors. They are sent to the avatar engine to render avatar images, which are then inverted into latent codes through the avatar generator. Finally, the latent codes are fed into the realistic generator to get realistic face images corresponding to the sampled avatar vectors. Moreover, we introduce semantic augmentation to expand the diversity of produced data by adding local perturbation to latent code. With the synthetic data, we can then train an avatar estimator on them. The avatar estimator is light-weight and efficient at inference. Given user selfie image, it accomplish the auto-creation by directly predicting an avatar vector that matches the input.

Our experiments involve both objective and subjective evaluations, comparing SwiftAvatar with previous methods in different aspects. SwiftAvatar achieves advantageous results over all existing methods in creating avatars fidelity to the input images. Moreover, experiments on two diverse avatar engines verify the strong generality of SwiftAvatar. Qualitative results illustrated in Fig. 1, show that SwiftAvatar can generate reasonable avatars for both engines given input faces. In summary, our contributions are as follows:

- A novel framework, SwiftAvatar, is proposed that can automatically create a stylized avatar given a user selfie image. It can be swiftly applied to any arbitrary avatar engines without extra assumptions (e.g. differentiable ren-

dering capability) on that engine.

- SwiftAvatar presents a novel pipeline that produces paired data across domains. It involves dual-domain generators to address the domain gap between realistic faces and stylized avatars. A novel semantic augmentation is also devised to improve the diversity of data synthesis.
- Experiments show that SwiftAvatar outperforms previous methods in terms of both quality and efficiency. Results on two different avatar engines also verify the strong generalizability of SwiftAvatar.

2 Related Work

3D Face Reconstruction Many progresses on 3D face reconstruction (Tuan Tran et al. 2017; Dou, Shah, and Kakadiaris 2017) cannot be achieved without 3D morphable models (3DMM) (Banz and Vetter 1999) and its variants like BFM (Gerig et al. 2018) and FLAME (Li et al. 2017), where the geometry and texture of a photorealistic 3D face are parameterized as a vector through linear transform. Recent researches also explore representing 3D faces in other formats like dense landmarks (Feng et al. 2018b) or position maps (Feng et al. 2018a). However, stylized avatar auto-creation is not 3D face reconstruction. They differ in two folds: 1) 3D face reconstruction methods aim to recover realistic 3D faces, not stylized characters; 2) Most avatar engines utilize avatar vectors to render stylized 3D characters according to manually designed assets. Estimating avatar vectors is much more difficult than reconstructing 3DMM coefficients, since the graphics rendering of avatar engines are usually black-box and non-differentiable.

GAN and GAN Inversion The rapid growth of generative networks, such as GANs (Goodfellow et al. 2014; Karras et al. 2017) and VAEs (Kingma and Welling 2013; Razavi, Van den Oord, and Vinyals 2019), inspires methods to use *latent codes* that can implicitly parameterize face images through a pre-trained generator. Among the pre-trained generators, the most popular one is the generator of the StyleGAN (Karras, Laine, and Aila 2019; Karras et al. 2020),

¹<https://newsroom.tiktok.com/en-us/express-yourself-through-tiktok-avatars>

²<https://github.com/facemoji/alter-core>




Figure 2: Conceptual comparison of previous and our methods. a) Previous methods try to solve unknown avatar vectors under image-level constraints, which requires to train a differentiable engine imitator to participate the optimization. b) While our method can use origin avatar engine and cross-domain generation to synthesize paired data (image-label), and then a model is trained based on the synthetic dataset.

known for its impressive capability in generating high quality face images. In order to leverage the pre-trained StyleGAN generator in editing existing face images, GAN inversion (Zhu et al. 2020) is required to compute the latent code whose StyleGAN output best matches a given face image. Existing GAN inversion methods estimate latent codes using either iterative optimization (Abdal, Qin, and Wonka 2020) or training a feed-forward encoder (Richardson et al. 2021) using VGG perceptual loss (Johnson, Alahi, and Fei-Fei 2016), LPIPS loss (Zhang et al. 2018) or face identity loss (Deng et al. 2019). Though GAN inversion can not be directly applied to estimate the avatar vectors due to the hand-crafted nature of avatar engines, SwiftAvatar leverages GAN inversion in its data production pipeline to build the correspondence between avatar vectors and latent codes.

Portrait Stylization Recent literatures on portrait stylization also benefit a lot from pre-trained StyleGANs. On one hand, finetuning the pre-trained StyleGAN generator provides with an agile approach for synthesizing faces of new styles (Song et al. 2021; Back 2021). On the other hand, freezing low-level layers of the generator during finetuning helps preserving structural information (e.g. face attributes) between face images generated from the original generator (realistic domain) and the finetuned generator (stylized domain), when they are using similar latent codes (Back 2021; Huang, Liao, and Kwong 2021). Though this cross-domain latent code sharing strategy cannot be directly applied on avatar engines, we discover that it has a strong potential in cross-domain data synthesis. In specific, we design the data production pipeline of SwiftAvatar to leverage a pre-trained SemanticStyleGAN (Shi et al. 2022). It adopts a compositional generator architecture, disentangling the latent space into different semantic areas separately, and could provide more precise local controls on synthesized face images.

Stylized Avatar Auto-Creation Auto-creating avatars for individual users has become an important capability to the entertainment industry. To ensure quality, commercial solution usually involve a large amount of manual annotations, something this paper seeks to avoid. Among the published approaches that avoid manual annotations, some (Shi et al. 2019, 2020) are designed for realistic avatars only, which

share the same domain with the real face images, it is easy to define whether they match or not. Therefore, these methods can utilize straightforward in-domain supervisions to improve fidelity of creation, such as the image-level L1/L2 losses, face parsing loss, or the face identity loss (Deng et al. 2019). Creating stylized avatars from real faces, on the contrary, is more difficult than creating realistic avatars. There are only few works toward this direction. Tied Output Synthesis (TOS) (Wolf, Taigman, and Polyak 2017) devises an encoder-decoder structure, that eliminates the domain gap by sharing one encoder across two domains. The contemporaneous work AgileAvatar (Sang et al. 2022) formulates a cascaded framework which progressively bridges the domain gap. To address the non-differentiable issue of avatar engines, all these methods need to train an imitator network to imitate the engines’ rendering procedure. While the quality of the images generated by the imitator is relatively poor. Besides, they all impose explicit constraints on images from different domains, resulting in suboptimal matches because of the domain gap. By contrast, our SwiftAvatar framework employs original avatar engine to generate high-quality images and utilizes dual-domain generators to overcome the domain gap.

3 Methodology

In this section, we present our unsupervised framework for stylized avatar’s auto-creation, as shown in Fig.3. It aims at estimating avatar vector p for input real face image x , then the estimated p can be used to render the corresponding stylized avatar y by avatar engine E . Our solution is split into three parts: dual-domain generators in Sec. 3.1, paired data production in Sec. 3.2, and avatar estimator in Sec. 3.3. Dual-domain generators address the domain gap problem by generating realistic faces and stylized avatars with shared latent codes. Then, paired data production focuses on how to generate paired data consisting of avatar vectors and realistic faces. Finally, avatar estimator estimates desired avatar vectors similar to input real face images.

3.1 Dual-Domain Generators

The dual-domain generators consist of a realistic generator G_{real} and an avatar generator G_{avatar} to achieve cross-domain

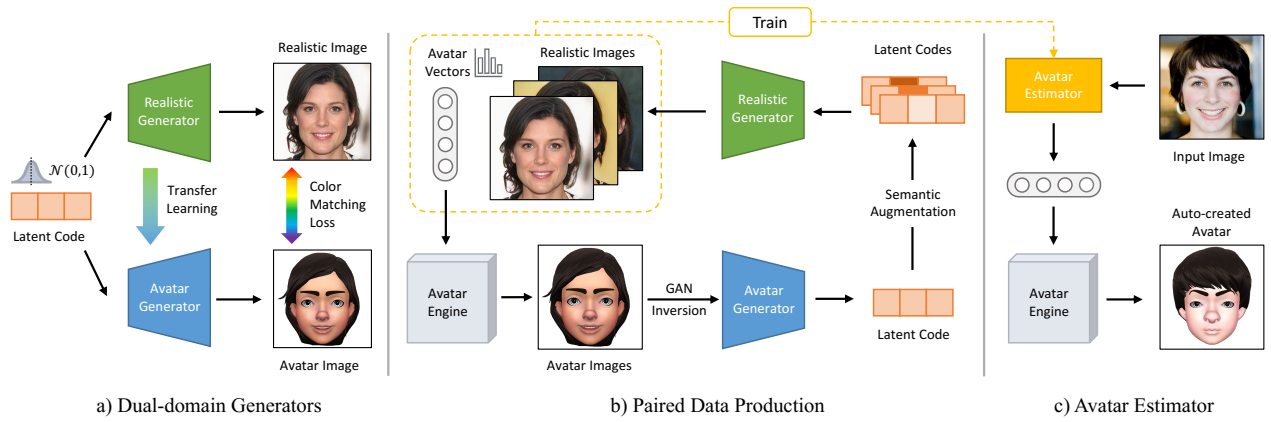


Figure 3: Overview of our method. a) Dual-domain generators consist of a fixed realistic generator and a transfer-learned avatar generator. They can generate corresponding realistic and avatar images when given the same latent codes. b) Paired data production starts from randomly sampled avatar vectors, which are sent to the avatar engine to render avatar images. These images are then inverted into latent codes through GAN inversion, and the latent codes are fed into the realistic generator to get realistic face images corresponding to the sampled avatar vectors. c) Supervised by the paired data, the avatar estimator can predict avatar vectors when given an real face input, then the avatar can be rendered by avatar engine.

image generation. Given the same latent code, they can simultaneously synthesize paired images of both realistic face and stylized avatar while preserving the same attributes (e.g. skin color, hair style). To impose an extra facial consistency constraint between the realistic and the avatar domains, we adopt SemanticStyleGAN (Shi et al. 2022) as the architecture of two generators owing to its extra semantic segmentation output.

Cross-Domain Generation Pretrained SemanticStyleGAN on CelebAMask-HQ (Lee et al. 2020) is directly used as realistic generator G_{real} , and also used to initialize the weight of G_{avatar} . We perform transfer learning on G_{avatar} : using only limited number of avatar images \mathcal{Y} to finetune avatar generator G_{avatar} . \mathcal{Y} are rendered from avatar engine using randomly sampled avatar vectors. The finetune procedure follows the settings in SemanticStyleGAN, which using the loss of StyleGAN2 (Karras et al. 2020):

$$\mathcal{L}_{StyleGAN2} = \mathcal{L}_{adv} + \lambda_{R1}\mathcal{L}_{R1} + \lambda_{path}\mathcal{L}_{path} \quad (1)$$

where λ_{R1} , λ_{path} are the constant weights of R1 regularization (Mescheder, Geiger, and Nowozin 2018) and path length regularization (Karras et al. 2020) separately. Adversarial loss \mathcal{L}_{adv} adopts non-saturating logistic loss (Goodfellow et al. 2014) and it forces G_{avatar} to generate images similar to avatar image dataset \mathcal{Y} . R1 regularization \mathcal{L}_{R1} is employed to improve the training stability and reduce the number of artifacts. And path length regularization \mathcal{L}_{path} leads to more reliable and consistently behaving models.

Facial Consistency Although directly fine-tuning SemanticStyleGAN can generate structurally similar paired data of avatar image and realistic face, the colors of each region are not well matched. Since SemanticStyleGAN learns the joint modeling of image and semantic segmentation. Such design can simultaneously synthesize face images and their

semantic segmentation results. We utilize the semantic segmentation output and introduce a color matching loss for cross-domain facial color consistency. Specifically, we extract specified pixels from same semantic areas in generated paired images, and match the mean color of them. $m^s(I)$ is the mean color of the region s in the image I , and we mainly consider matching the color in hair and skin areas: $s \in \{hair, skin\}$. The color matching loss is:

$$\mathcal{L}_{color} = \sum_s \|m^s(G_{real}(z)) - m^s(G_{avatar}(z))\|^2 \quad (2)$$

Overall, the final finetuning loss \mathcal{L}_{total} for G_{avatar} is formulated as:

$$\mathcal{L}_{total} = \mathcal{L}_{StyleGAN2} + \lambda_{color}\mathcal{L}_{color} \quad (3)$$

An example is shown in Figure. 4, when given a shared latent code, the dual-domain generators can generate a pair of images: a realistic face and a stylized avatar. They are similar in facial structure and color composition, yet belong to two different domains.

3.2 Paired Data Production

Paired data production pipeline focuses on synthesizing paired avatar vectors and realistic face images, as is illustrated in Figure. 3. We sample a number of random avatar vectors as labels \mathcal{P} , which are used by the graphics engine to generate corresponding avatar images \mathcal{Y} . Then, for every avatar image y , we pass it through GAN inversion to get its latent code w . We adopt optimization-based GAN inversion for its better performance:

$$w^* = \arg \min_w \mathcal{L}_{invert}(y, w) \quad (4)$$

where w^* represents our target latent code. For faster convergence, latent codes are initialized with mean value w_{mean} , and optimized by gradient descent. We use LPIPS




Figure 4: Generated results of our dual-domain generators. When providing randomly sampled latent codes, we can obtain a) and b), realistic faces and their semantic segmentation results. Simultaneously, we can obtain c) and d), stylized avatar images and their semantic segmentation results.

loss (Zhang et al. 2018) to measure perceptual similarity, and mean squared error (MSE) loss between original images and reconstructed avatar images to measure reconstruction similarity. Besides, MSE loss between w_{mean} and w is set as latent regularization to ensure better generation results. The loss function is formulated as:

$$\mathcal{L}_{invert} = \lambda_p LPIPS(y, G_{avatar}(w)) + \lambda_i \|y - G_{avatar}(w)\|^2 + \lambda_l \|w - w_{mean}\|^2 \quad (5)$$

where λ_p , λ_i , λ_l are constants to balance three loss terms. After obtaining desired latent code w , it is fed into the realistic generator G_{real} to generate a realistic face x which is similar to original avatar image y in identification:

$$x = G_{real}(w) \quad (6)$$

In this way, we can generate paired data (p, x) used for later avatar vector estimation training process. An example illustration of paired data is shown in Figure. 5.

Semantic Augmentation The sampled avatar vectors, as well as their rendered avatar faces suffer from the lack of diversity, due to the limited amount of assets available for avatar engines. Take the “hair style” attribute for example, in most stylized avatar engines, different hair styles are determined by selecting from different hair meshes all predefined in the engine. The limited number of predefined mesh assets restricts the capability of avatar engines to match a real-world user face, whose hair styles would be countless. To enrich the diversity of generated realistic faces in data production, we take advantage of the compositional generation ability of our dual-domain generators which are implemented as SemanticStyleGANs, and design the *semantic augmentation*. In SemanticStyleGAN, each semantic part is modulated individually with corresponding local latent codes. Such property enables us to manipulate only specific regions of the synthetic images while keeping other regions unchanged. In implementation, we add random noises to part of latent code corresponding to these ambiguous avatar

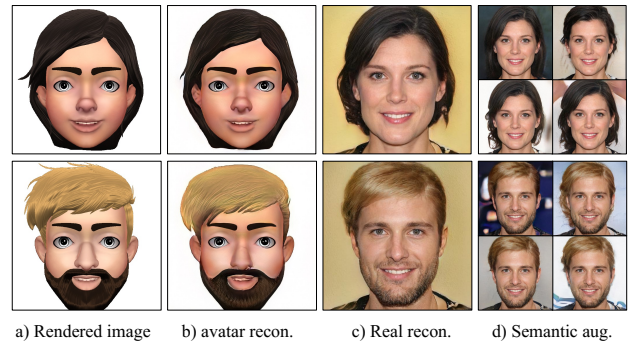


Figure 5: Examples of produced paired data. a) Engine rendered avatar images. b) Reconstructed avatar images. c) Generated realistic images. d) Semantic augmented images.

vectors (e.g. hair type). Semantic augmentation can be described as:

$$w_k \rightarrow (1 - \lambda_{aug})w_k + \lambda_{aug}n \quad (7)$$

where λ_{aug} is a hyper-parameter to adjust semantic augmentation density, w_k is the local part of latent code to be augmented, and n represents for random noise. A semantic augmentation example is shown in Figure. 5, where hair and background region are changed.

3.3 Avatar Estimator

Once the aforementioned synthetic paired dataset is produced, we can train an avatar estimator to predict avatar vectors, which contain continuous and discrete parameters. We choose ResNet-18 (He et al. 2016) pretrained on ImageNet (Deng et al. 2009) as our backbone. We remove its last fully connected layers and add multiple separate MLP heads for different parameter estimation. All continuous parameters form a target vector to be predicted in one head, supervised by L1 loss. Every discrete parameter estimation is carried out with a standalone head. Because both generation and semantic augmentation would inevitably introduce noises to discrete labels, we choose symmetric cross-entropy (SCE) loss (Wang et al. 2019), which has been proven robust to noises, for the optimization of discrete tasks. The total loss of avatar estimator is:

$$\mathcal{L}_{estimator} = \lambda_d SCE(\hat{p}_d^i, p_d^i) + \lambda_c |\hat{p}_c - p_c| \quad (8)$$

where λ_d and λ_c are hyper-parameters to balance two loss terms. \hat{p}_d^i , \hat{p}_c are the prediction results of i -th discrete head and continuous head respectively. And p_d^i , p_c are their corresponding ground-truth.

4 Experiments

Experimental Data To verify the effectiveness of our method, we conduct experiments on two stylized avatar engines: the TikTok engine and the Alter engine. The TikTok engine contains resourceful discrete and continuous avatar parameters, so we generate 50000 avatar vectors and corresponding rendered images. The Alter engine is an open-source avatar engine and only contains discrete avatar parameters which has fewer assets than the TikTok engine,




Figure 6: Visual comparison with other methods on TikTok engine and Alter engine. a) Given input images, b) our method creates stylized avatars that look similar to input images. c) Baseline method without considering the domain gap problem in stylized avatars. d) F2P (Shi et al. 2019), an image-level constraint method intending to create semi-realistic avatars, is not suitable to create stylized avatars. e) F2P v2 (Shi et al. 2020), a fast and robust version of F2P. f) Manually created avatars by professional designers, which can be regarded as ground truth.

so we just generate 10000 avatar vectors and corresponding rendered images. The detailed information of both avatar engines can be found in supplementary materials. These avatar images and avatar vectors are used for finetuning avatar generator and producing paired data. For evaluation, we choose 116 images from FFHQ dataset (Karras, Laine, and Aila 2019), which consists of diverse kinds of face shapes, hairstyles, etc. We invite designers to manually generate avatar vectors for these 116 images as ground truth.

Implementation Details We implement our methods using PyTorch 1.10 library and perform all experiments on NVIDIA V100 GPUs. When finetuning the avatar generator G_{avatar} , we use the same optimizer settings as in SemanticStyleGAN. Batch size is set to 16, style mixing probability (Karras, Laine, and Aila 2019) is set to 0.3. λ_{R1} , λ_{path} are set to 10 and 0.5 separately. Lazy regularization (Karras et al. 2020) is applied every 16 mini-batches for discriminator (R1 regularization) and every 4 mini-batches for generator (path length regularization). All the images used for generators are aligned and resized to resolution 512×512 . The optimization-based GAN inversion approach employs Adam (Kingma and Ba 2014) optimizer in the paired data production stage, and the learning rate initially follows co-

sine annealing with 0.1. We optimize 200 steps for all latent codes, and λ_i , λ_p , λ_l are set to 0.1, 1 and 1, respectively. The mean latent code w_{mean} is the average among 10^4 randomly generated latent codes from avatar generator G_{avatar} in \mathcal{W} space, and serve as the initialization. Notably, directly optimizing the latent code could be problematic since some avatar assets are transparent, e.g. glasses. Thus, we use a modified version \tilde{w}_{mean} for latent code initialization (See supplementary materials for details). For semantic augmentation, we generate 10 augmented images for each latent code, using randomly generated noise in \mathcal{W} space. We set λ_{aug} to 1 for the background to improve the model robustness of background variance, and also set λ_{aug} to 0.3, 0.06 for the hair part and glasses part to expand data diversity. In the avatar estimator training stage, the input images of avatar estimator are resized to 224×224 . We use the Adam optimizer with batch size 256 to train 100 epochs. The learning rate is set to $1e-3$, and decayed by half per 30 epochs. For the experiments on TikTok engine, there are 1 continuous head and 8 discrete heads inside the avatar estimator, so we set λ_c and λ_d to 1 and 10 separately. For the experiments on Alter engine, the avatar estimator contains 6 discrete heads and the training loss only contains discrete loss.

Method	TikTok engine ↓	Alter engine ↓	Speed ↑
Baseline	0.5033	0.3962	~ 10 ² Hz
F2P	0.3562	0.3040	~ 1 Hz
F2P v2	0.4466	0.2825	~ 10 ² Hz
Ours	0.3110	0.2405	~ 10 ² Hz

Table 1: Quantitative Evaluation. We compare our method with three other methods on TikTok engine and Alter engine. Lower distance represents results are more similar to manual-creation. We test speed on NVIDIA V100.

4.1 Comparison with Other Methods

We compare our method with other methods, including Baseline, F2P (Shi et al. 2019), F2P v2 (Shi et al. 2020) on both their auto-creation similarity and inference speed. Baseline method is setup as ignoring the domain gap problem, where the avatar estimator is trained on (y, p) paired data, that is, trained on rendered avatar images instead of real face images. Figure. 6 shows a comparison of rendered stylized avatars among different methods corresponding to their predicted avatar vectors. As can be seen, our method can better address the domain gap problem and create stylized avatars with high similarity with input real faces and approximate the quality of manual method from designers. For more results, please refer to supplementary materials.

4.2 Quantitative Evaluation

Although the avatar auto-creation has no standard answer, we generally regard the avatar manually created by professional designers as the ground truth, and quantitatively evaluate at image level. We calculate the perceptual distance (LPIPS) (Huang, Liao, and Kwong 2021) between auto-created avatar and manual-created avatar to simulate human observation. The lower distance indicates the avatar is better matched to input image. The results are presented in Table. 1, from which we can see our method significantly outperforms others. Since our method only needs one forward propagation, our inference speed is also competitive, which can be applied in real-time applications.

4.3 Human Rating

We invited 50 volunteers to subjectively evaluate all algorithms of 50 images from evaluation dataset on both graphics engine. Given the input real face images and stylized avatars generated from our and other methods, the volunteers were requested to pick the best matching avatar for each input real face. The methods were presented in random order, and the volunteers were given an unlimited time to choose. The human rating results are shown in Table. 2. 59.11% of the answers on TikTok engine and 63.68% on the Alter engine selected our method as the best matching results, showing the superiority of our method compared with others.

4.4 Ablation Study

We conduct the ablation experiments in terms of domain adaptation and semantic augmentation to verify their impor-

Method	Ours	baseline	F2P	F2P v2
TikTok	59.11%	5.66%	30.86%	4.37%
Alter	63.68%	18.94%	8.65%	8.73%

Table 2: Human subjective rating on two engines. Our method earns the most choices when asked to choose the avatar which matches the human image.

Method	TikTok engine ↓	Alter engine ↓
baseline	0.5033	0.3962
+ domain adaptation	0.3401	0.3123
+ semantic aug	0.3110	0.2405

Table 3: Ablation study on two engines. Lower distance indicates better matching results to manual-creation.

tance in our framework. We adopt the same evaluation metric described in Sec. 4.2. The ablation starts from Baseline method, where the avatar estimator is trained on rendered avatar images and avatar vectors pair. Then on the basis of baseline method, we add domain adaptation and semantic augmentation in turn. Table. 3 shows their quantitative results on different engines. The domain adaptation greatly alleviates the domain gap problem in stylized avatar auto-creation, establishes a bridge of connecting real faces and stylized avatars. The semantic augmentation brings noticeable improvement for our method due to its expansion of diversity to samples.

5 Limitations and Future Work

There are two main limitations we observed in the experiments. First, our method occasionally predicts wrong color influenced by environmental lighting. This problem might be resolved by considering lighting condition into our pipeline. Second, in paired data production, avatar vector sampling distribution directly influences the training data quality. Simply random sampling could produce some strange images and may cause long-tail problem (e.g. gender, age). In the future, we will perform attribute analysis, and introduce reasonable sampling priors to address synthetic data distribution problem.

6 Conclusion

In summary, we present a novel unsupervised framework for auto-creation of stylized avatars. We design the dual-domain generators to address the domain gap between the real images and stylized avatars. Then following the paired data production pipeline, high-quality paired data are produced, which is used for training the avatar estimator. Finally stylized avatars are created by conducting efficient avatar vector estimation. Compared with previous methods, our method is more concise in training stage and more efficient in inference stage. Results on quantitative evaluation and human rating demonstrate the superiority of our method. Also, the success of applying on two different avatar graphics engines demonstrates the generality of our method.

References

- Abdal, R.; Qin, Y.; and Wonka, P. 2020. Image2stylegan++. How to edit the embedded images? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8296–8305.
- Back, J. 2021. Fine-Tuning StyleGAN2 For Cartoon Face Generation. *arXiv preprint arXiv:2106.12445*.
- Blanz, V.; and Vetter, T. 1999. A morphable model for the synthesis of 3D faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, 187–194.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4690–4699.
- Dou, P.; Shah, S. K.; and Kakadiaris, I. A. 2017. End-to-end 3D face reconstruction with deep neural networks. In *proceedings of the IEEE conference on computer vision and pattern recognition*, 5908–5917.
- Feng, Y.; Wu, F.; Shao, X.; Wang, Y.; and Zhou, X. 2018a. Joint 3d face reconstruction and dense alignment with position map regression network. In *Proceedings of the European conference on computer vision (ECCV)*, 534–551.
- Feng, Z.-H.; Huber, P.; Kittler, J.; Hancock, P.; Wu, X.-J.; Zhao, Q.; Koppen, P.; and Rätzsch, M. 2018b. Evaluation of dense 3D reconstruction from 2D face images in the wild. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, 780–786. IEEE.
- 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.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Huang, J.; Liao, J.; and Kwong, S. 2021. Unsupervised image-to-image translation via pre-trained stylegan2 network. *IEEE Transactions on Multimedia*, 24: 1435–1448.
- Johnson, J.; Alahi, A.; and Fei-Fei, L. 2016. Perceptual Losses for Real-Time Style Transfer and Super-Resolution. In Leibe, B.; Matas, J.; Sebe, N.; and Welling, M., eds., *Computer Vision – ECCV 2016*, 694–711. Cham: Springer International Publishing. ISBN 978-3-319-46475-6.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2017. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8110–8119.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Lee, C.-H.; Liu, Z.; Wu, L.; and Luo, P. 2020. MaskGAN: Towards Diverse and Interactive Facial Image Manipulation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Li, T.; Bolkart, T.; Black, M. J.; Li, H.; and Romero, J. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6): 194:1–194:17.
- Mescheder, L.; Geiger, A.; and Nowozin, S. 2018. Which training methods for GANs do actually converge? In *International conference on machine learning*, 3481–3490. PMLR.
- Razavi, A.; Van den Oord, A.; and Vinyals, O. 2019. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32.
- Richardson, E.; Alaluf, Y.; Patashnik, O.; Nitzan, Y.; Azar, Y.; Shapiro, S.; and Cohen-Or, D. 2021. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2287–2296.
- Sang, S.; Zhi, T.; Song, G.; Liu, M.; Lai, C.; Liu, J.; Wen, X.; Davis, J.; and Luo, L. 2022. AgileAvatar: Stylized 3D Avatar Creation via Cascaded Domain Bridging. *arXiv preprint arXiv:2211.07818*.
- Shi, T.; Yuan, Y.; Fan, C.; Zou, Z.; Shi, Z.; and Liu, Y. 2019. Face-to-parameter translation for game character auto-creation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 161–170.
- Shi, T.; Zuo, Z.; Yuan, Y.; and Fan, C. 2020. Fast and robust face-to-parameter translation for game character auto-creation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 1733–1740.
- Shi, Y.; Yang, X.; Wan, Y.; and Shen, X. 2022. Semantic-StyleGAN: Learning Compositional Generative Priors for Controllable Image Synthesis and Editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11254–11264.
- Song, G.; Luo, L.; Liu, J.; Ma, W.-C.; Lai, C.; Zheng, C.; and Cham, T.-J. 2021. AgileGAN: stylizing portraits by inversion-consistent transfer learning. *ACM Transactions on Graphics (TOG)*, 40(4): 1–13.
- Tuan Tran, A.; Hassner, T.; Masi, I.; and Medioni, G. 2017. Regressing robust and discriminative 3D morphable models with a very deep neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5163–5172.
- Wang, Y.; Ma, X.; Chen, Z.; Luo, Y.; Yi, J.; and Bailey, J. 2019. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 322–330.
- Wolf, L.; Taigman, Y.; and Polyak, A. 2017. Unsupervised creation of parameterized avatars. In *Proceedings of the*

IEEE International Conference on Computer Vision, 1530–1538.

Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.

Zhu, J.; Shen, Y.; Zhao, D.; and Zhou, B. 2020. In-domain gan inversion for real image editing. In *European conference on computer vision*, 592–608. Springer.