

DEFORMTOON3D: Deformable Neural Radiance Fields for 3D Toonification

Junzhe Zhang^{1,3*} Yushi Lan^{1*} Shuai Yang¹ Fangzhou Hong¹
Quan Wang³ Chai Kiat Yeo² Ziwei Liu¹ Chen Change Loy¹

¹S-Lab, Nanyang Technological University

²Nanyang Technological University ³SenseTime Research

{shuai.yang, asckyeo, ziwei.liu, ccloy}@ntu.edu.sg

{junzhe001, yushi001, fangzhou001}@e.ntu.edu.sg {wangquan}@sensetime.com

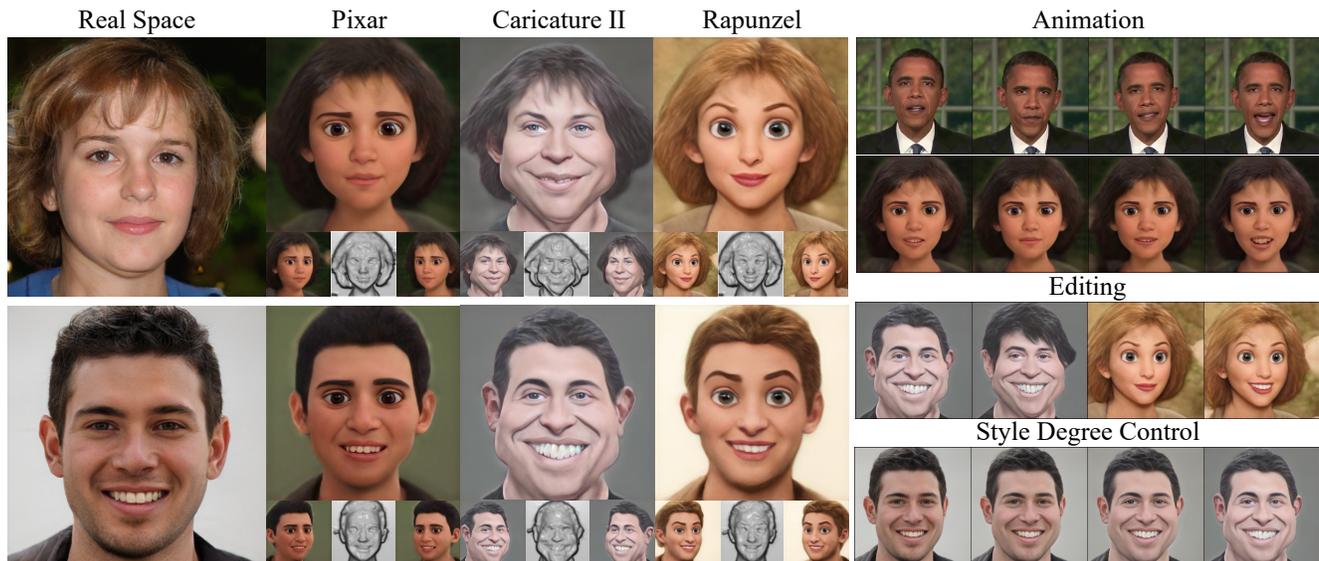


Figure 1: We propose DEFORMTOON3D, an efficient 3D toonification framework which supports geometry-texture decomposed toonification over multiple styles. DEFORMTOON3D is fine-tuning free and could be easily extended to a series of downstream applications designed for the pre-trained GAN, *e.g.*, animation given a driving video, semantic attributes editing (+ bangs and smile), and flexible style degree control.

Abstract

In this paper, we address the challenging problem of 3D toonification, which involves transferring the style of an artistic domain onto a target 3D face with stylized geometry and texture. Although fine-tuning a pre-trained 3D GAN on the artistic domain can produce reasonable performance, this strategy has limitations in the 3D domain. In particular, fine-tuning can deteriorate the original GAN latent space, which affects subsequent semantic editing, and requires independent optimization and storage for each new style, limiting flexibility and efficient deployment. To overcome these challenges, we propose DEFORMTOON3D, an effective toonification framework tailored for hierarchical 3D GAN. Our approach decomposes 3D toonification into

subproblems of geometry and texture stylization to better preserve the original latent space. Specifically, we devise a novel StyleField that predicts conditional 3D deformation to align a real-space NeRF to the style space for geometry stylization. Thanks to the StyleField formulation, which already handles geometry stylization well, texture stylization can be achieved conveniently via adaptive style mixing that injects information of the artistic domain into the decoder of the pre-trained 3D GAN. Due to the unique design, our method enables flexible style degree control and shape-texture-specific style swap. Furthermore, we achieve efficient training without any real-world 2D-3D training pairs but proxy samples synthesized from off-the-shelf 2D toonification models. Code is released at <https://github.com/junzhezhang/DeformToon3D>.

*Equal contribution.

1. Introduction

Artistic portraits are prevalent in various applications such as comics, animation, virtual reality, and augmented reality. In this work, our main objective is to propose an effective approach for 3D-aware artistic toonification, a critical problem that involves transferring the style of an artistic domain onto a target 3D face with stylized geometry and texture. The task opens up potential applications for quick high-quality 3D avatar creation based on a photograph with the style of a designated artwork, which would typically require highly professional handcraft skills.

Substantial progress has been made in automatic portrait style transfer over 2D images. Starting with image style transfer [14, 52, 34, 28] and image-to-image translation [27, 33, 10, 6, 55], recent advancements in StyleGAN-based generators [23, 24] have shown their potential in high-quality toonification via efficient transfer learning [47]. Specifically, a pre-trained StyleGAN generator on face images is fine-tuned to transfer to the artistic portrait domain. With the progress of 3D-aware GANs [8, 40], researchers have extended this pipeline to 3D with well-designed domain adaptation frameworks [21, 31, 26, 1], enabling remarkable 3D portrait toonification.

Although fine-tuning the pre-trained StyleGAN-based model for toonification achieves superior quality, it has several limitations. **First**, fine-tuning the pre-trained generator shifts its generative space from the real face domain to the artistic portrait domain at the cost of deteriorating the original GAN latent space. With tremendous off-the-shelf tools [53, 61] trained for the original GAN space, altering the well-learned style space would affect the performance of downstream applications over the toonified portrait, *e.g.*, semantic editing. **Second**, despite fine-tuning-based domain adaptations have been thoroughly investigated for 2D GANs [48, 64], applying this technique to 3D GANs fails to leverage the full potential of the architecture of 3D generator models [8, 40] for characterizing view-consistent shape and high-frequency textures in the artistic domain. **Third**, it is inevitable to fine-tune a heavy generator for each new style, which requires hours of training time and additional storage. This limitation affects scalability when deploying dozens of fine-tuned generators for real-time user interactions. Therefore, 3D toonification remains a challenging task that requires further exploration.

To better preserve the pre-trained GAN latent space and to better exploit the 3D GAN generator, we propose **DEFORMTOON3D** that decomposes geometry and texture stylization into more manageable subproblems. In particular, unlike conventional 3D toonification approaches that fine-tune the whole 3D GAN generator following existing 2D fine-tuning schemes, we carefully consider the characteristics of 3D GANs to decompose the stylization of geometry and texture domains. To achieve geometry stylization,

we introduce a novel **StyleField** on top of a pre-trained 3D generator to deform each point in the style space to the pre-trained real space guided by an instance code. This allows for easy extension to multiple styles with a single stylization field by introducing a style code to guide the deformation. Since StyleField already handles geometry stylization well, texture stylization can be easily achieved through adaptive style mixing which injects artistic domain information into the network for effective texture toonification. Notably, our unique design enables training of the method at minimal cost using synthetic paired data with realistic faces generated by a pre-trained 3D GAN and corresponding paired stylized data generated by an off-the-shelf 2D toonification model [66].

The proposed DEFORMTOON3D achieves high-quality geometry and texture toonification over a vast variety of styles, as demonstrated in Fig. 1. Additionally, our approach preserves the original GAN latent space, enabling compatibility with existing tools built on the real face space GAN, including inversion [31], editing [53], and animation [61]. Furthermore, our design significantly reduces the storage footprint by requiring only a small stylization field with a set of AdaIN parameters for artistic domain stylization. In summary, our work makes the following contributions:

- We propose a novel StyleField that separates geometry toonification from texture, providing a more efficient method for modeling 3D shapes than fine-tuning and enabling flexible style control.
- We present an approach to achieve multi-style toonification with a single model, facilitating cross-style manipulation and reducing storage footprint.
- We introduce a full synthetic data-driven training pipeline that offers an efficient and cost-effective solution to training the model without requiring real-world 2D-3D training pairs.

2. Related Work

3D Generative Models. Inspired by the success of Generative Adversarial Networks (GAN) [15] in generating photo-realistic images [23, 5, 25], researchers have been making efforts towards 3D-aware generation [37, 18, 41]. Starting with explicit intermediate shape representations, such as voxels [37, 18] and meshes [41], which lack photo-realism and are memory-inefficient, researchers have recently shifted towards using implicit functions [43, 35, 9] along with physical rendering processes [57, 36] as intrinsic 3D inductive biases. Among these approaches, 3D generative models [7, 51] extended from neural radiance fields (NeRF) [36] have demonstrated impressive view-consistency in synthesized results. While the original NeRF

is limited to modeling static scenes, recent research has introduced deformation fields to enable NeRF to model dynamic volumes [44, 45, 62, 30]. To increase the resolution of generated images, recent studies [8, 19] have resorted to voxel-based representations or adopted a hybrid design [38, 40, 8, 16]. This hybrid design involves a cascade model $G = G_1 \circ G_0$ pairing a 3D generator G_0 with a 2D super-resolution decoder G_1 . Both G_0 and G_1 follow the style-based architecture [23, 24] to accept a latent code \mathbf{w} to control the style of the generated object. By super-resolving the intermediate low-resolution 2D features produced by the G_0 with the G_1 , the hybrid design achieves view-consistent synthesis at high resolution, *e.g.*, 1024^2 .

Domain Adaptation for StyleGAN Toonification. Researchers have typically employed domain adaptation in 2D space to achieve toonification with StyleGANs [23, 25]. Typically, a portrait StyleGAN pre-trained on real images [32, 22] is fine-tuned on an artistic domain dataset to generate toonified faces. Building upon this straightforward framework [48], a series of in-depth research has been conducted to further improve style control [66, 65], choices of latent code [58], few-shot training [39], and text-guided adaptation [13].

Pre-trained 3D GANs offer high-quality generation and thus have the potential to facilitate downstream applications such as portrait stylization through 3D GAN inversion [31, 8]. To extend 2D StyleGAN domain adaptation to 3D, CIPS-3D [69] proposed fine-tuning only the super-resolution decoder module for view-consistent texture toonification. However, it is limited to texture toonification since the 3D generator is left unchanged. Dr.3D [21], E3DGE [31], and 3DAvatarGAN [1] fine-tune the entire 3D generator [8, 40] for both geometry and texture toonification, while DATID-3D [26] leverages a Stable Diffusion [50] generated corpus for text-guided domain adaptation. While these methods yield impressive results, they come with limitations, as they require costly fine-tuning and independent model storage for each new style. Furthermore, previous fine-tuning-based toonification methods suffer from limited generality due to the incompatibility with abundant editing techniques developed for the original StyleGAN latent space [53, 54, 46]. In contrast, DEFORMTOON3D fully preserves the original 3D GAN latent space, which makes it intrinsically compatible with the editing methods trained for the original StyleGAN space. It achieves comparable visual quality while being 10 times more storage-efficient than previous methods.

3. DEFORMTOON3D

We present the framework of DEFORMTOON3D in Fig. 2. Our approach begins with a typical hybrid 3D-aware design [40, 8] that generates real-domain faces, and reformulates it to a 3D toonification framework. The approach

starts with a cascade model $G = G_1 \circ G_0$, which pairs a 3D generator G_0 with a 2D super-resolution decoder G_1 . The generator G_0 captures the underlying geometry with the instance code \mathbf{w} and camera pose ξ , and produces an intermediate feature map \mathbf{F} with volume rendering [36]. Then, G_1 upsamples \mathbf{F} to obtain a high-resolution image \mathbf{I} with high-frequency details added. To adapt G from the real domain to the artistic or cartoon domain, existing methods [21, 1, 69, 31] view G_1 and G_0 as a whole and simply fine-tune the pre-trained G , failing to take advantage of the decomposed characteristics of the hybrid framework design. By comparison, DEFORMTOON3D fully exploits this cascaded synthesis process by using a novel StyleField module for geometry stylization, which in turn also benefits appearance stylization, allowing it to adopt a simple adaptive style mixing strategy. In Sec. 3.1, we elaborate the proposed StyleField along with the pre-trained G_0 to handle geometry stylization. In Sec. 3.2, we explain how the adaptive style mixing injects the style of the target domain into G_1 to achieve texture stylization. Lastly, we present our training pipeline in Section 3.3.

3.1. Geometry Toonification with StyleField

To train a 3D generator \tilde{G}_0 capable of synthesizing artistic domain geometry, previous methods [21, 26, 31, 1, 59] fine-tune the pre-trained G_0 with a target-domain dataset, which can be computationally expensive and could potentially deteriorate the original GAN latent space. To address this issue, we propose to establish a correspondence between the stylized NeRF \mathcal{N}_S and the real-space NeRF \mathcal{N}_R . More specifically, we use a stylization field (the StyleField), H_D , to bridge the correspondence, such that $\tilde{G}_0(\mathbf{x}_S) = (G_0 \circ H_D)(\mathbf{x}_S)$. As shown in Fig. 2, given a stylized NeRF $\mathcal{N}_S : \mathbb{R}^3 \mapsto \mathbb{R}^4$ of the target style domain, our goal is to estimate a 3D deformation residual, $H_D : \mathbb{R}^3 \mapsto \mathbb{R}^3$, which maps \mathcal{N}_S back to \mathcal{N}_R via:

$$\mathcal{N}_S \rightarrow \mathcal{N}_R : \mathbf{x}_R = (\mathbf{x}_S + H_D(\mathbf{x}_S)), \forall \mathbf{x}_S \in \mathcal{N}_S, \quad (1)$$

where H_D represents the residual 3D deformation $H_D(\mathbf{x}_S) = \Delta \mathbf{x}_S$ in the 3D space of the Stylized NeRF \mathcal{N}_S and maps each 3D point $\mathbf{x}_S \in \mathcal{N}_S$ in the stylized space to its corresponding position in the real space \mathcal{N}_R .

To improve expressiveness, we extend H_D as a conditional neural deformation field that outputs the offsets under the conditions of style and identity:

$$\mathcal{N}_S \rightarrow \mathcal{N}_R : \mathbf{x}_R = \mathbf{x}_S + H_D(\mathbf{x}_S, \mathbf{w}_S, \mathbf{w}_R) \quad (2)$$

where \mathbf{w}_S is the style code that specifies the artistic domain, \mathbf{w}_R is the instance code corresponding to \mathcal{N}_R that represents the identity of the 3D face in the source domain. Both \mathbf{w}_S and \mathbf{w}_R serve as the holistic geometry indicators to guide the deformation.

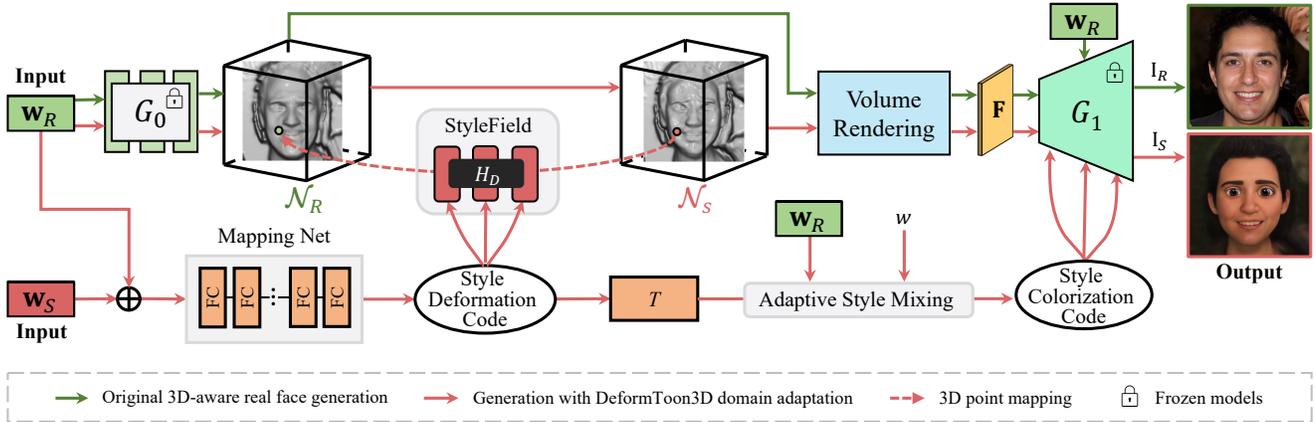


Figure 2: **DEFORMTOON3D framework.** Given sampled instance code \mathbf{w}_R and style code \mathbf{w}_S as conditions, DEFORMTOON3D first deforms a point from the style space \mathcal{N}_S to the real space \mathcal{N}_R , which achieves geometry toonification without modifying the pre-trained G_0 . Afterwards, we leverage adaptive style mixing with weight w to inject the texture information of the target domain into the pre-trained G_1 for texture toonification. Both pre-trained generators G_0 and G_1 are kept frozen.

We build H_D as an MLP consisting of four SIREN [56] layers due to its superior high-frequency modeling capacity. After all the points $\mathbf{x}_S \in \mathcal{N}_S$ are deformed to the real space, we generate the new feature map $\hat{\mathbf{F}} = \tilde{G}_0(\mathbf{w}_R, \mathbf{w}_S, \xi)$ for the input of G_1 and synthesize the high-resolution image with the geometry of the artistic domain. By introducing the 3D deformation module H_D , we no longer need to fine-tune pre-trained G_0 to achieve geometry deformation of the target domain. This greatly alleviates the parameters to optimize by 50% and fully preserves the original GAN latent space. Moreover, with \mathbf{w}_S serving as the style condition, a single H_D can support multiple styles, which further saves storage by 98.5% compared to fine-tuning the whole model per style with 10 styles.

3.2. Texture Transfer with Adaptive Style Mixing

Thanks to the StyleField H_D that handles the geometry toonification within the cascade 3D GAN model, we only need to inject the artistic domain texture information into G_1 for texture stylization. Here, G_1 is a 2D style-based architecture, where image styles are effectively adjusted by AdaIN [20]. Inspired by style mixing [23, 25, 66] that controls the AdaIN parameters, we inject the texture information of the target style \mathbf{w}_S by mixing the style parameters of G_1 , as shown in Fig. 2. To bridge the domain gap between the real space and target domain, we further add a lightweight MLP, T , for each layer of G_1 to adjust the style code \mathbf{w}_S . The adapted $T(\mathbf{w}_S)$ and the instance code \mathbf{w}_R are fused with a weight w by weighted average, and sent to the affine transformation block of G_1 to obtain the final style parameters for AdaIN. This mechanism allows us to model and control multi-domain textures with T and w , without fine-tuning the original decoder G_1 .

Let \tilde{G}_1 denote G_1 with the adaptive style mixing. The image generation process with domain adaptation given w

and \mathbf{w}_S can be formulated as $\hat{\mathbf{I}} = \tilde{G}_1(\hat{\mathbf{F}}, \mathbf{w}_R, \mathbf{w}_S, w)$.

3.3. Training

Data Preparation. We follow Sim2Real [63, 68, 31] to generate paired data for training. Specifically, to generate the training corpus for each iteration of the training process, we pre-calculate a set of real space NeRFs \mathcal{N}_R with corresponding latent codes \mathbf{w}_R and rendered images \mathbf{I}_R . To generate pair-wise stylized ground truths, we stylize the rendered image with existing 2D toonification models to obtain the target ground truth \mathbf{I}_S . Here, to validate our method’s performance on multi-style domain adaptation, we adopt the exemplar-based DualStyleGAN [66] as our 2D toonification model, since it supports hundreds of diverse styles, such as Cartoon, Pixar, and Caricature.

Finally, we define $\mathcal{X} = \{\mathbf{I}_R, \mathbf{w}_R, \mathbf{w}_S, \mathbf{I}_S\}$ as a training set for DEFORMTOON3D with $\{\mathbf{w}_R, \mathbf{I}_R, \mathbf{w}_S\}$ to serve as the training inputs and $\{\mathbf{I}_S\}$ is the set of training ground truth. $\mathbf{I}_R \in \mathcal{X}$ is drawn i.i.d from distribution $P(G(\mathbf{z}, \xi))$ where $\mathbf{z} \sim \mathcal{N}(0, 1)$ and ξ is the camera pose distribution of pre-trained 3D GAN G . Some samples are shown in Fig. 3.

Reconstruction Loss. Here we use the LPIPS loss [67] to evaluate stylized image quality:

$$\mathcal{L}_{\text{Rec}}(\mathcal{X}) = \mathbb{E}_{\mathcal{X}} \left[\|P(\hat{\mathbf{I}}) - P(\mathbf{I}_S)\|_2 \right], \quad (3)$$

where $P(\cdot)$ denotes the perceptual feature extractor.

Smoothness Regularization. To encourage the smoothness of stylization field deformation offsets and reduce spatial distortion, a smoothness regularization is included to regularize H_D . Here we penalize the norm of the Jacobian matrix $\mathbb{J}_{H_D} = \nabla H_D$ of the deformation field [44] to ensure the learned deformations are physically smooth:

$$\mathcal{L}_{\text{Elastic}} = \text{ReLU}(\|\nabla H_D(\mathbf{x}_S, \mathbf{w}_S, \mathbf{w}_R)\|_2^2 - \epsilon), \quad (4)$$

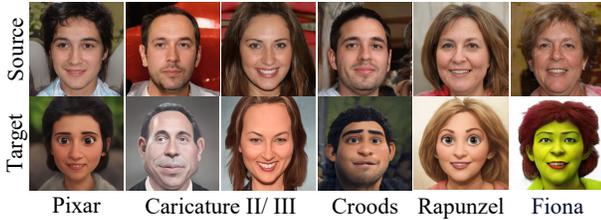


Figure 3: **Training Samples.** We show the source image (row 1) and the stylized images of the target domain (row 2) for training supervisions.

where ϵ is the slack parameter for the smoothness regularization. We set $\epsilon = 0.1$ for all the experiments.

Adversarial Training. Additionally, we apply a non-saturating adversarial loss \mathcal{L}_{Adv} [23] to bridge the domain gap of toonified results.

In summary, the overall loss function is defined as

$$\mathcal{L} = \mathcal{L}_{Rec} + \lambda_{Elastic}\mathcal{L}_{Elastic} + \lambda_{Adv}\mathcal{L}_{Adv},$$

where we set $\lambda_{Adv} = 0.05$ and $\lambda_{Elastic} = 0.01$ in all the experiments.

4. Experiments

Datasets. We mainly focus on the human face domain and use synthesized data for the whole training. We follow the data synthesis procedure mentioned in Sec. 3.3 for generating training corpus and adopt DualStyleGAN [66] as the 2D stylization model to infer paired toonified image. We generate images with the following 10 styles: Pixar, Comic, Slam Dunk, The Croods, Fiona (Shrek), Rapunzel (Disney Princess), Hiccup Horrendous Haddock III (How To Train Your Dragon), and three different caricature styles. We use StyleSDF [40] pre-trained on FFHQ [23] as our 3D generator. For evaluating toonification on real-world images, we evaluate on CelebA-HQ [22].

Implementation Details. In all the experiments, we set the learning rate to 5×10^{-4} . We adopt Adam [29] optimizer to train the toonification models. We train the DEFORMTOON3D for 100 epochs. To expedite the training process, we disable the adversarial loss for the initial 50 epochs. The training takes approximately 24 hours using 8 Tesla V100 GPUs with a batch size set to 16. For adaptive style mixing, $w = 1$ is fixed during training and could be manually selected from $[0, 1]$ to achieve texture interpolation during inference. More implementation details and experiment results are included in the supplementary material.

Baselines. Here we design three baselines for extensive evaluations. The first is CIPS-3D [69], which naively fine-tunes the super-resolution G_1 for view-consistent toonification. The second is E3DGE [31], which fine-tunes both G_0 and G_1 independently for true-3D toonification. Another prominent method is to leverage directional CLIP

loss [13, 2, 26] for adapting a pre-trained style-based generator. Here we extend StyleGAN-NADA [13] to 3D GAN and employ image CLIP directional loss for the evaluation. The baseline models are trained on the same dataset as DEFORMTOON3D, with each fine-tuning process applied to a single style. In contrast, DEFORMTOON3D is capable of accommodating all styles within a single model.

4.1. Comparisons with Baselines

Qualitative Results. We show the qualitative comparisons against the baselines in Fig. 4. DEFORMTOON3D fully captures the characteristics of the target domain with consistent identity preservations. CIPS-3D only provides texture-wise toonification and ignores geometry deformation. E3DGE suffers from mode collapse and tends to lose identity. StyleGAN-NADA fails to capture the characteristics of the target domain. Our method produces high-quality toonification with consistent identity preservations.

Table 1: **Quantitative evaluation over 10 styles.** DEFORMTOON3D achieves the best identity preservation (IP) and FID.

	CIPS-3D	E3DGE	NADA	Ours
IP \uparrow	0.681	0.707	0.535	0.781
FID \downarrow	50.6	34.0	59.3	27.6

Quantitative Results. To evaluate the fidelity and quality of toonification, we compare their identity preservation(IP) [11] and FID respectively in Tab. 1. In terms of IP, DEFORMTOON3D outperforms all baseline methods across the 10 styles provided, which underscores the benefits of retaining the 3D generator. In terms of FID, CIPS-3D [69] only fine-tunes G_1 and achieves worse performance compared to E3DGE [31], which fine-tunes the whole generator. StyleGAN-NADA achieves the worst FID performance, which we attribute to the challenges of directly adopting 2D CLIP-based supervision on 3D GANs. A detailed breakdown by individual styles is available in the supplementary material.

Table 2: **User preference study.**

	CIPS-3D	E3DGE	NADA	Ours
Shape	20.8%	17.4%	3.4%	58.4%
Appearance	16.1%	16.4%	3.4%	64.1%
Identity	23.1%	7.1	5.4	64.4%
Overall	21.8%	9.9%	2.7%	65.6%

User Study. For the user preference study, we collected 2400 votes to select the preferred rendering results in terms of shape stylization, appearance stylization, identity preservation, and overall performance. As shown in Tab. 2, the proposed method gives the most preferable results despite the fact that the baseline methods are trained on a single style.

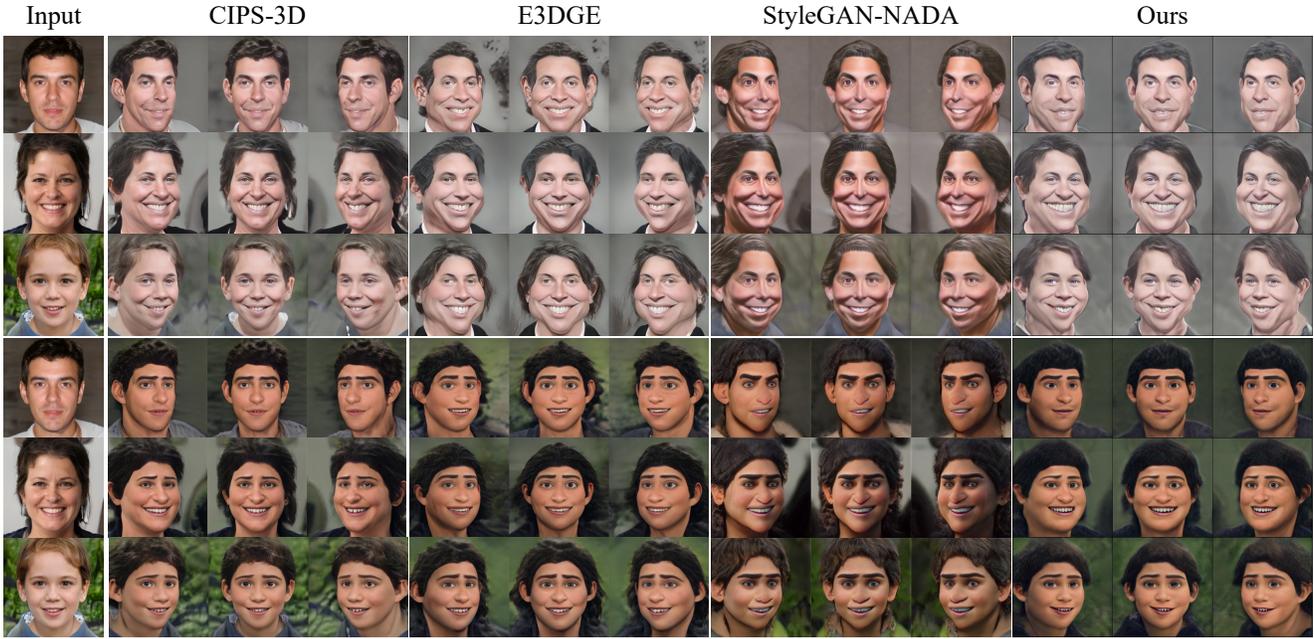


Figure 4: **Qualitative comparisons with baseline methods.** DEFORMTOON3D produces better performance against all baselines regarding toonification fidelity, diversity and identity preservation.

Table 3: **Storage cost comparison.** Values are averaged across 10 styles and shown in MB. DEFORMTOON3D achieves a considerably more efficient storage footprint against the baselines.

Methods	Trainable Params↓	Model Storage↓
CIPS-3D [69]	5.81	59.93
E3DGE [31]	7.64	76.4
StyleGAN-NADA [13]	7.64	76.4
Ours (single-style)	3.82	11.46

Storage Cost Comparison. We detail the storage costs in Table 3. Thanks to our unique design that disentangles geometry and texture, our method requires no fine-tuning and fewer parameters for training. In a single-style scenario, our method reduces storage needs by 85%. This advantage becomes even more pronounced with the multi-style version of H_D , achieving a storage saving of 98.5% in a 10-style scenario. This makes our approach particularly feasible for potential mobile applications.

4.2. Applications

To demonstrate the generality of full GAN latent space preservation, we show that DEFORMTOON3D could be easily extended to a series of downstream applications proposed for the original pre-trained GAN, including inversion, editing, and animation. To further validate DEFORMTOON3D’s unique geometry-texture decomposed toonification, we show results of flexible toonification style control. **Inversion and Editing.** With fully preserved 3D generative prior, DEFORMTOON3D could directly adopt pre-trained



Figure 5: **Editing of toonified results.** We show two attribute editing results over six styles. In row 1, we add bangs to the male identities, and in row 2, we add “Smile” to female identities. The edited results fully preserve the identity and abide by the style of the target domain.

3D GAN inversion framework and latent editing directions for semantic-aware editing over the toonified portrait. Here we adopt E3DGE [31] for 3D GAN inversion and show the inverted results of the real image in Fig. 6. With fully preserved GAN latent space, DEFORMTOON3D could be applied to real images over multiple styles with high quality. We also include the editing results of diverse styles in Fig. 5. As can be seen, our method produces consistent editing results and fully preserves the identity and the style of the target domain, regarding both the geometry and texture.

Animatable Toonification. Inspired by the success of 2D method [61] that obtains a rig-like control over StyleGAN-generated 2D faces, we propose a straightforward pipeline that aligns 3DMM [4] parameters with pre-trained 3D GAN latent space. Specifically, we train two MLP $\mathcal{F}_w : \mathbb{R}^{|\mathcal{W}|} \rightarrow \mathbb{R}^{|\mathcal{M}|}$ and $\mathcal{F}_m : \mathbb{R}^{|\mathcal{M}|} \rightarrow \mathbb{R}^{|\mathcal{W}|}$ to learn the bi-directional mapping between 3DMM parameter space \mathcal{M} and 3D GAN



Figure 6: **DEFORMTOON3D on the real images.** Our method enables multiple styles toonification with a single model, where both the texture and the geometry matches the target domain.

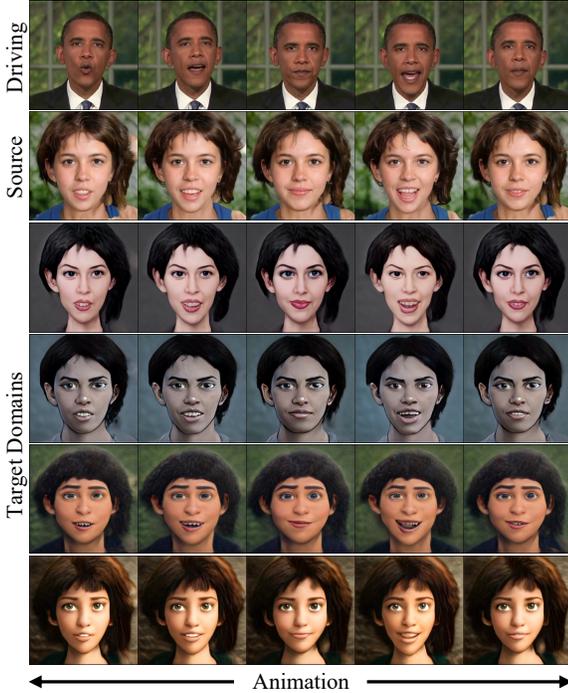


Figure 7: **Animation of stylized results.** We drive the real space images (row 2) with an input video (row 1). Since DEFORMTOON3D fully preserves the pre-trained GAN latent space, the driving direction of the real space could be directly applied to the style space (row 3 – 6). The expression on of animated toonified identities fully abide with the driving frames without affecting the toonification performance of target domain.

\mathcal{W} space. To impose 3DMM-based control, given a latent code \mathbf{w} , we first infer its 3DMM parameters $\tilde{\mathbf{m}} = \mathcal{F}_{\mathbf{w}}(\mathbf{w})$ and reconstructed 3DMM code $\tilde{\mathbf{w}} = \mathcal{F}_{\mathbf{m}}(\tilde{\mathbf{m}})$. After imposing 3DMM-based editing $\tilde{\mathbf{m}}_{\Delta} = \tilde{\mathbf{m}} + \mathbf{m}_{\Delta}$, we infer the corresponding edited \mathcal{W} code $\tilde{\mathbf{w}}_{\Delta} = \mathcal{F}_{\mathbf{m}}(\tilde{\mathbf{m}}_{\Delta})$. The final result is synthesized from $\hat{\mathbf{w}} = \mathbf{w} + (\tilde{\mathbf{w}}_{\Delta} - \tilde{\mathbf{w}})$. The whole framework is trained in a self-supervised manner with cycle consistency regularizations. Please refer to the supplementary material for more technical details and animation results.

We show the animated toonification results in Fig. 7. Here, we extract the 3DMM parameters from a driving video [17] using a pre-trained predictor [12]. The expression dimension of extracted parameters are injected to the sampled \mathbf{w}_R using the procedure described above. Four styles of animated toonified results are included. As can be seen, our designed animation pipeline could accurately drive the identity both in the real space (row 2) and the style spaces (row 3 – 6). The reenacted expressions are natural and abide with the driving input, which validates the generality of our method.

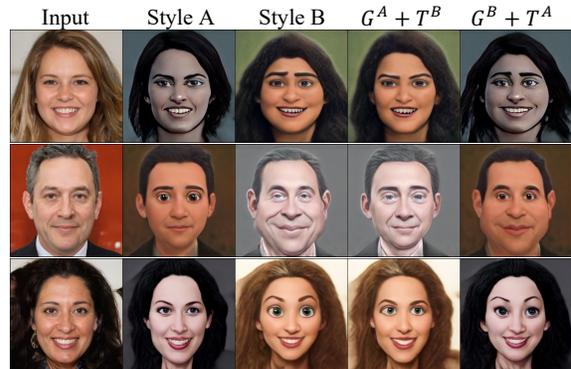


Figure 8: **Style swap.** Besides style interpolation, DEFORMTOON3D supports style swap of geometry and texture only across two styles. As shown here, given the real space input (col 1) and the toonification results of two spaces (col 2 – 3), DEFORMTOON3D could swap the geometry G and texture T of two styles independently (col 4 – 5), which cannot be achieved by previous methods.

Toonification Style Control. Due to the unique geometry-texture decomposition design, DEFORMTOON3D offers flexible style control. **First**, we achieve style degree control and show the results in Fig. 9. Geometry-wise, since H_D outputs the 3D deformation offsets $\Delta \mathbf{x}_S$, we simply interpolate the offsets with $\tau * \Delta \mathbf{x}_S$ where $\tau = 0$ represents an identical mapping of the real space. Texture-wise, we rescale the style mixing weight w of G_1 , where $w = 0$ preserves the color of the real space images. **Second**, due to the geometry-texture disentangled property, DEFORM-

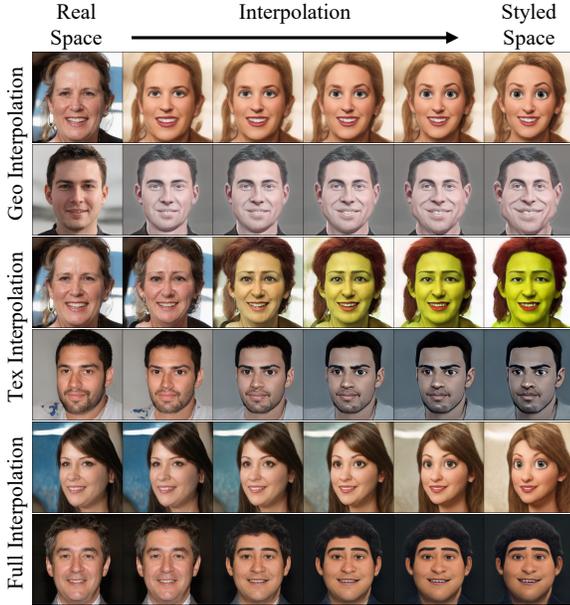


Figure 9: **Style degree control.** Thanks to the geometry-texture toonification decomposition, DEFORMTOON3D can specify geometry, texture, and full style control.

Table 4: Ablation study.

Method	IP \uparrow	FID \downarrow
w/o \mathbf{w}_R condition	0.735	33.4
MLP as H_D	0.746	31.8
w/o \mathcal{L}_{Adv}	0.769	29.9
Ours	0.781	27.6

TOON3D naturally supports the toonification of shape and texture only. As shown in Fig. 8, we achieve the geometry-texture swap between multiple styles, where the geometry of one style could be combined with the texture of another style. This could not be achieved by all previous methods and opens up broader potential downstream applications.

4.3. Ablation Study

We ablate the effectiveness of our design choices in Tab. 4. **1) Instance code condition:** By removing instance code \mathbf{w}_R , StyleField tends to learn similar offsets for different identities, whereas adopting \mathbf{w}_R as the instance deformation conditions facilitates better toonification. **2) Style-Field architecture:** We replace the SIREN architecture with an MLP network as defined by D-NeRF [49] and observe the significant performance drop. This demonstrates the representation power of SIREN deformation field. **3) Adversarial training:** We also validate the effectiveness of adversarial loss, which brings noticeable improvement regarding both FID and identity similarity, respectively.

4.4. Limitations

As shown in Fig. 10 pertaining to the style ‘‘Comic’’ and ‘‘Slam Dunk’’, though DeformToon3D produces reasonable texture-wise stylization, the corresponding geometry still has noticeable artifacts. The proposed StyleField implicitly learns the correspondence between the paired data from the real space and the style space. Such correspondence is easier to learn with information cues such as illumination from the 3D-ish styles like Pixar or noticeable keypoints from caricature styles, but harder for styles with limited information cues like Comic.

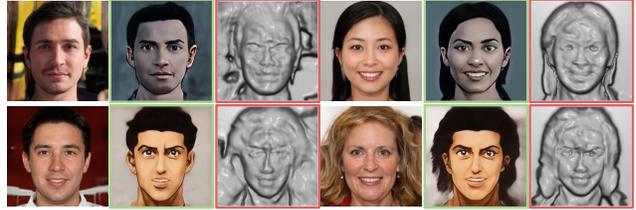


Figure 10: **Failure cases.**

5. Conclusion and Future Work

In this paper, we propose a novel 3D toonification framework DEFORMTOON3D for a fine-tuning free, geometry-texture decomposed 3D face toonification. We fully exploit the hierarchical characteristics of 3D GAN and introduce a StyleField to handle 3D geometry toonification of G_0 , along with adaptive style mixing that injects texture information into G_1 . Our method achieves high-quality toonification on both geometry and texture, outperforming existing methods. Thanks to the preservation of the 3D generative prior, DEFORMTOON3D facilitates a range of downstream applications.

As a pioneering effort in this field, we believe this work will inspire future works on free 3D toonification. First, to mitigate the geometry-texture ambiguity present in certain styles, introducing re-lighting during training could serve as a potential solution [42]. Second, a more flexible training paradigm could directly guide the 3D toonification process with a pre-trained vision-language model [50]. Third, future research could focus on integrating a comprehensive 3D animation pipeline [3, 60] into the toonification process. Moreover, the potential applicability of DeformToon3D to other 3D GANs [8] and shapes beyond human faces, such as full-body settings, is worth investigating.

Acknowledgement. This study is supported under the RIE2020 Industry Alignment Fund Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s). It is also supported by Singapore MOE AcRF Tier 2 (MOE-T2EP20221-0011, MOE-T2EP20221-0012) and NTU NAP Grant.

References

- [1] Rameen Abdal, Hsin-Ying Lee, Peihao Zhu, Menglei Chai, Aliaksandr Siarohin, Peter Wonka, and S. Tulyakov. 3DAvatarGAN: Bridging domains for personalized editable avatars. *arXiv*, abs/2301.02700, 2023. 2, 3
- [2] Aibek Alanov, Vadim Titov, and Dmitry Vetrov. HyperDomainNet: Universal domain adaptation for generative adversarial networks. 2022. 5
- [3] Alexander W. Bergman, Petr Kellnhofer, Yifan Wang, Eric Chan, David B. Lindell, and Gordon Wetzstein. Generative neural articulated radiance fields. In *NIPS*, 2022. 8
- [4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3D faces. In *SIGGRAPH*, 1999. 6
- [5] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*. OpenReview.net, 2019. 2
- [6] Kaidi Cao, Jing Liao, and Lu Yuan. CariGANs: unpaired photo-to-caricature translation. *ACM TOG*, 37(6):1–14, 2018. 2
- [7] Eric Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and G. Wetzstein. Pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis. In *CVPR*, 2021. 2
- [8] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 2022. 2, 3, 8
- [9] Wenzheng Chen, Jun Gao, Huan Ling, Edward J. Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3D objects with an interpolation-based differentiable renderer. 2
- [10] Min Jin Chong and David Forsyth. GANs N’ Roses: Stable, controllable, diverse image to image translation. *arXiv preprint arXiv:2106.06561*, 2021. 2
- [11] Jiankang Deng, Jia Guo, Xue Niannan, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 5
- [12] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3D face reconstruction with weakly-supervised learning: From single image to image set. In *CVPR*, 2019. 7
- [13] Rinon Gal, Or Patashnik, Haggai Maron, Gal Chechik, and Daniel Cohen-Or. StyleGAN-NADA: Clip-guided domain adaptation of image generators. *arXiv*, abs/2108.00946, 2021. 3, 5, 6
- [14] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, pages 2414–2423, 2016. 2
- [15] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [16] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. StyleNeRF: A style-based 3D-aware generator for high-resolution image synthesis. In *ICLR*, 2021. 3
- [17] Yudong Guo, Keyu Chen, Sen Liang, Yongjin Liu, Hujun Bao, and Juyong Zhang. AD-NeRF: Audio driven neural radiance fields for talking head synthesis. In *ICCV*, 2021. 7
- [18] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Escaping plato’s cave: 3D shape from adversarial rendering. In *ICCV*, 2019. 2
- [19] Fangzhou Hong, Zhaoxi Chen, Yushi Lan, Liang Pan, and Ziwei Liu. EVA3D: Compositional 3D human generation from 2D image collections. In *ICLR*, 2023. 3
- [20] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, pages 1501–1510. 4
- [21] Wonjoon Jin, Nuri Ryu, Geon-Yeong Kim, Seung-Hwan Baek, and Sunghyun Cho. Dr.3D: Adapting 3D GANs to artistic drawings. *SIGGRAPH Asia 2022 Conference Papers*, 2022. 2, 3
- [22] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *ICLR*, 2018. 3, 5
- [23] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 2, 3, 4, 5
- [24] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 2, 3
- [25] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 2, 3, 4
- [26] Gwanghyun Kim and Se Young Chun. DATID-3D: Diversity-preserved domain adaptation using text-to-image diffusion for 3D generative model, 2022. 2, 3, 5
- [27] Junho Kim, Minjae Kim, Hyeonwoo Kang, and Kwang Hee Lee. U-GAT-IT: Unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. In *ICLR*, 2019. 2
- [28] Sunnie SY Kim, Nicholas Kolkin, Jason Salavon, and Gregory Shakhnarovich. Deformable style transfer. In *ECCV*, 2020. 2
- [29] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, volume abs/1412.6980, 2015. 5
- [30] Yushi Lan, Chen Change Loy, and Bo Dai. DDF: Correspondence Distillation from NeRF-based GAN. *arXiv preprint arXiv:2212.09735*, 2022. 3
- [31] Yushi Lan, Xuyi Meng, Shuai Yang, Chen Change Loy, and Bo Dai. E3DGE: Self-supervised geometry-aware encoder for style-based 3D gan inversion. In *CVPR*, 2022. 2, 3, 4, 5, 6
- [32] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. MaskGAN: Towards diverse and interactive facial image manipulation. In *CVPR*, 2020. 3
- [33] Bing Li, Yuanlue Zhu, Yitong Wang, Chia-Wen Lin, Bernard Ghanem, and Linlin Shen. AniGAN: Style-guided generative adversarial networks for unsupervised anime face generation. *IEEE TMM*, 2021. 2
- [34] Jing Liao, Yuan Yao, Lu Yuan, Gang Hua, and Sing Bing Kang. Visual attribute transfer through deep image analogy. *ACM Transactions on Graphics*, 36(4):120, 2017. 2

- [35] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *CVPR*, June 2019. 2
- [36] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*. Springer, 2020. 2, 3
- [37] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yongliang Yang. HoloGAN: Unsupervised Learning of 3D Representations From Natural Images. In *ICCV*, 2019. 2
- [38] Michael Niemeyer and Andreas Geiger. GIRAFFE: Representing scenes as compositional generative neural feature fields. In *CVPR*, 2021. 3
- [39] Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A. Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. Few-shot image generation via cross-domain correspondence. *CVPR*, pages 10738–10747, 2021. 3
- [40] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. StyleSDF: High-Resolution 3D-Consistent Image and Geometry Generation. In *CVPR*, 2021. 2, 3, 5
- [41] Xingang Pan, Bo Dai, Ziwei Liu, Chen Change Loy, and Ping Luo. Do 2D GANs know 3D shape? Unsupervised 3D Shape Reconstruction from 2D Image GANs. In *ICLR*, 2021. 2
- [42] Xingang Pan, Xudong Xu, Chen Change Loy, Christian Theobalt, and Bo Dai. A Shading-Guided Generative Implicit Model for Shape-Accurate 3D-Aware Image Synthesis. In *NIPS*, 2021. 8
- [43] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*. IEEE, 2019. 2
- [44] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, 2021. 3, 4
- [45] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *TOG*, 40(6), dec 2021. 3
- [46] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. StyleCLIP: Text-driven manipulation of stylegan imagery. *ICCV*, pages 2065–2074, 2021. 3
- [47] Justin NM Pinkney and Doron Adler. Resolution dependent gan interpolation for controllable image synthesis between domains. *arXiv preprint arXiv:2010.05334*, 2020. 2
- [48] Justin N. M. Pinkney and Doron Adler. Resolution dependent gan interpolation for controllable image synthesis between domains. *arXiv*, abs/2010.05334, 2020. 2, 3
- [49] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *CVPR*, 2020. 8
- [50] Robin Rombach, A. Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *CVPR*, pages 10674–10685, 2021. 3, 8
- [51] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. GRAF: Generative radiance fields for 3D-aware image synthesis. In *NIPS*, 2020. 2
- [52] Ahmed Selim, Mohamed Elgharib, and Linda Doyle. Painting style transfer for head portraits using convolutional neural networks. *ACM TOG*, 35(4):1–18, 2016. 2
- [53] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. InterFaceGAN: Interpreting the disentangled face representation learned by GANs. *PAMI*, PP, 2020. 2, 3
- [54] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *CVPR*, 2021. 3
- [55] Yichun Shi, Debayan Deb, and Anil K Jain. WarpGAN: Automatic caricature generation. In *CVPR*, pages 10762–10771, 2019. 2
- [56] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *NIPS*, 2020. 4
- [57] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene Representation Networks: Continuous 3D-structure-aware neural scene representations. In *NIPS*, 2019. 2
- [58] Guoxian Song, Linjie Luo, Jing Liu, Wan-Chun Ma, Chun-Pong Lai, Chuanxia Zheng, and Tat-Jen Cham. AgileGAN: stylizing portraits by inversion-consistent transfer learning. *ACM Transactions on Graphics (TOG)*, 40:117:1–117:13, 2021. 3
- [59] Kunpeng Song, Ligong Han, Bingchen Liu, Dimitris Metaxas, and Ahmed Elgammal. Diffusion guided domain adaptation of image generators. *arXiv preprint https://arXiv.org/abs/2212.04473*, 2022. 3
- [60] Jingxiang Sun, Xuan Wang, Lizhen Wang, Xiaoyu Li, Yong Zhang, Hongwen Zhang, and Yebin Liu. Next3D: Generative neural texture rasterization for 3D-aware head avatars. In *CVPR*, 2023. 8
- [61] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. Stylerig: Rigging StyleGAN for 3D control over portrait images. In *CVPR*, pages 6142–6151, 2020. 2, 6
- [62] Ayush Tewari, R. MallikarjunB., Xingang Pan, Ohad Fried, Maneesh Agrawala, and Christian Theobalt. Disentangled3D: Learning a 3d generative model with disentangled geometry and appearance from monocular images. *CVPR*, pages 1506–1515, 2022. 3
- [63] Erroll Wood, Tadas Baltruvsaitis, Charlie Hewitt, Sebastian Dziadzio, Matthew Johnson, Virginia Estellers, Thomas J. Cashman, and Jamie Shotton. Fake it till you make it: face analysis in the wild using synthetic data alone. *ICCV*, pages 3661–3671, 2021. 4
- [64] Zongze Wu, Yotam Nitzan, Eli Shechtman, and Dani Lischinski. StyleAlign: Analysis and applications of aligned stylegan models. *arXiv preprint arXiv:2110.11323*, 2021. 2

- [65] Shuai Yang, Liming Jiang, Ziwei Liu, , and Chen Change Loy. VToonify: Controllable high-resolution portrait video style transfer. *ACM Transactions on Graphics (TOG)*, 41(6):1–15, 2022. [3](#)
- [66] Shuai Yang, Liming Jiang, Ziwei Liu, and Chen Change Loy. Pastiche master: Exemplar-based high-resolution portrait style transfer. In *CVPR*, 2022. [2](#), [3](#), [4](#), [5](#)
- [67] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [4](#)
- [68] Yuxuan Zhang, Huan Ling, Jun Gao, Kangxue Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, and Sanja Fidler. DatasetGAN: Efficient labeled data factory with minimal human effort. In *CVPR*, 2021. [4](#)
- [69] Peng Zhou, Lingxi Xie, Bingbing Ni, and Qi Tian. CIPS-3D: A 3D-Aware Generator of GANs Based on Conditionally-Independent Pixel Synthesis. *arXiv*, 2021. [3](#), [5](#), [6](#)