SHIVANGI ANEJA, Technical University of Munich, Germany JUSTUS THIES, Max Planck Institute for Intelligent Systems, Tübingen, Germany ANGELA DAI, Technical University of Munich, Germany MATTHIAS NIESSNER, Technical University of Munich, Germany

(b) Expression Manipulation

(a) Texture Manipulation

Fig. 1. ClipFace learns a self-supervised generative model for jointly synthesizing geometry and texture leveraging 3D morphable face models, that can be guided by text prompts. For a given 3D mesh with fixed topology, we can generate arbitrary face textures as UV maps (top). The textured mesh can then be manipulated with text guidance to generate diverse set of textures and geometric expressions in 3D by altering (a) only the UV texture maps for Texture Manipulation and (b) both UV maps and mesh geometry for Expression Manipulation.

We propose ClipFace, a novel self-supervised approach for text-guided editing of textured 3D morphable model of faces. Specifically, we employ userfriendly language prompts to enable control of the expressions as well as appearance of 3D faces. We leverage the geometric expressiveness of 3D morphable models, which inherently possess limited controllability and texture expressivity, and develop a self-supervised generative model to jointly synthesize expressive, textured, and articulated faces in 3D. We enable highquality texture generation for 3D faces by adversarial self-supervised training, guided by differentiable rendering against collections of real RGB images. Controllable editing and manipulation are given by language prompts to adapt texture and expression of the 3D morphable model. To this end, we propose a neural network that predicts both texture and expression latent codes of the morphable model. Our model is trained in a self-supervised fashion by exploiting differentiable rendering and losses based on a pre-trained CLIP model. Once trained, our model jointly predicts face textures in UV-space, along with expression parameters to capture both geometry and texture changes in facial expressions in a single forward pass. We further show the

SIGGRAPH '23 Conference Proceedings, August 6–10, 2023, Los Angeles, CA, USA © 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-0159-7/23/08...\$15.00 https://doi.org/10.1145/3588432.3591566 applicability of our method to generate temporally changing textures for a given animation sequence.

ACM Reference Format:

Shivangi Aneja, Justus Thies, Angela Dai, and Matthias Niessner. 2023. ClipFace: Text-guided Editing of Textured 3D Morphable Models. In Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH '23 Conference Proceedings), August 6– 10, 2023, Los Angeles, CA, USA. ACM, New York, NY, USA, 17 pages. https: //doi.org/10.1145/3588432.3591566

1 INTRODUCTION

Modeling 3D content is central to many applications in our modern digital age, including asset creation for video games and films, as well as mixed reality. In particular, modeling 3D human face avatars is a fundamental element towards digital expression. However, current content creation processes require extensive time from highly-skilled artists in creating compelling 3D face models.

In contrast to implicit representations for human faces [Chan et al. 2021] that do not follow fixed mesh topology, 3D morphable models present a promising approach for modeling animatable avatars, with popular blendshape models used for human faces (e.g., FLAME [Li et al. 2017]) or bodies (e.g., SMPL [Loper et al. 2015]). In particular, they offer a compact, parametric representation to model an object, while maintaining a mesh representation that fits the classical graphics pipelines for editing and animation. Additionally,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Shivangi Aneja, Justus Thies, Angela Dai, and Matthias Niessner

the shared topology of the representation enables deformation and texture transfer capabilities.

Despite such morphable models' expressive capability in geometric modeling and potential practical applicability towards artist creation pipelines, they remain insufficient for augmenting artist workflows. This is due to limited controllability, as they rely on PCA models, and lack of texture expressiveness, since the models have been built from very limited quantities of 3D-captured textures; both of these aspects are crucial for content creation and visual consumption. We thus address the challenging task of creating a generative model to enable synthesis of expressive, textured, and articulated human faces in 3D.

We propose ClipFace, to enable controllable generation and editing of 3D faces. We leverage the geometric expressiveness of 3D morphable models, and introduce a self-supervised generative model to jointly synthesize textures and adapt expression parameters of the morphable model. To facilitate controllable editing and manipulation, we exploit the power of vision-language models [Radford et al. 2021] to enable user-friendly generation of diverse textures and expressions in 3D faces. This allows us to specify facial expressions as well as the appearance of the human via text while maintaining a clean 3D mesh representation that can be consumed by standard graphics applications. Such text-based editing enables intuitive control over the content creation process.

Our generative model is trained in a self-supervised fashion, leveraging the availability of large-scale face image datasets with differentiable rendering to produce a powerful texture generator that can be controlled along with the morphable model geometry by text prompts. Based on our texture generator, we learn a neural network that can edit the texture latent code as well as the expression parameters of the 3D morphable model with text prompt supervision and losses based on CLIP. Our approach further enables generating temporally varying textures for a given driving expression sequence.

The main focus of our work lies in enabling text-guided editing and control of 3D morphable face models. As recent 3D face models [Gerig et al. 2017; Li et al. 2017] are limited in their texture space, we propose a generative model to synthesize UV textures with a realistic appearance; this texture space is a pre-requisite for our text-guided editing. To summarize, we present the following contributions:

- We propose a novel approach to controllable editing of textured, parametric 3D morphable models through user-friendly text prompts, by exploiting CLIP-based supervision to jointly synthesize texture and expressions of a 3D face model.
- The controllable 3D face model is supported by our texture generator, trained in a self-supervised fashion on 2D images only.
- Our approach additionally enables generating temporally varying textures of an animated 3D face model from a driving video sequence.

2 RELATED WORK

Texture Generation: There is a large corpus of research works in the field of generative models for UV textures [Gecer et al. 2021a; Gecer et al. 2020; Gecer et al. 2019, 2021b; Lattas et al. 2020, 2021;

Lee et al. 2020; Li et al. 2020; Luo et al. 2021; Wang et al. 2022b]. These methods achieve impressive results; however, the majority is fully supervised in nature, requiring ground truth textures, which in turn necessitate collection in a controlled capture setting. Learning self-supervised texture generation is much more challenging, and only a handful of methods exist. For instance, Marriott *et al.* [2021] were among the first to leverage Progressive GANs [Karras et al. 2018] and 3D Morphable Models [Blanz and Vetter 1999] to generate textures for facial recognition; however, the textures generated remain relatively low resolution and are not suitable to perform language-driven texture and expression manipulations.

Textures generated by Slossberg *et al.* [2022] made a significant improvement in quality by using pretrained StyleGAN [Karras et al. 2020b] and StyleRig [Tewari et al. 2020a]. The closest inspiration to our texture generator is StyleUV [Lee et al. 2020] which also operates on a mesh. Both methods achieve stunning results, but do not take the head and ears into account, which limits their practical applicability to directly use them as assets in games and movies. In our work, we propose a generative model to synthesize UV textures for the full-head topology to enable text-guided editing and control of 3D morphable face models.

Semantic Manipulation of Facial Attributes: Facial manipulation has also seen significant studies following the success of StyleGAN2 [Karras et al. 2020b] image generation. In particular, its disentangled latent space facilitates texture editing as well as enables a level of control over pose and expressions of generated images. Several methods [Abdal et al. 2021b; Deng et al. 2020; Ghosh et al. 2020; Kowalski et al. 2020; Liu et al. 2022; Tewari et al. 2020a,b] have made significant progress to induce controllability to images by embedding 3D priors via conditioning StyleGAN on known facial attributes extracted from synthetic face renderings. However, these methods operate on 2D images, and although they achieve high-quality results on a per-frame basis, consistent and coherent rendering from varying poses and expressions remains challenging. Motivated by such impressive 2D generators, we propose to lift them to 3D and directly operate in UV space of a 3D mesh, producing temporally-consistent results when rendering animation sequences.

Text-Guided Image Manipulation: Recent progress in 2D language models has opened up significant opportunities for textguided image manipulation [Abdal et al. 2021a; Avrahami et al. 2022; Bau et al. 2021; Crowson 2021; Dayma et al. 2021; Ramesh et al. 2022, 2021; Ruiz et al. 2022]. For instance, contrastive languageimage pre-training (CLIP) [Radford et al. 2021] model has been used for text-guided editing for a variety of 2D/3D applications [Canfes et al. 2022; Gal et al. 2021; Hong et al. 2022; Khalid et al. 2022; Kocasari et al. 2022; Michel et al. 2022; Patashnik et al. 2021; Petrovich et al. 2022; Wang et al. 2022a; Wei et al. 2022; Youwang et al. 2022]. StyleClip [Patashnik et al. 2021] presented seminal advances in stylizing human face images by leveraging the expressive power of CLIP in combination with the generative power of StyleGAN to produce unique manipulations for faces. This was followed by StyleGAN-Nada [Gal et al. 2021], which enables adapting image generation to a remarkable diversity of styles from various domains, without requiring image examples from those domains. However, these manipulations are designed for the image space, and are not 3D consistent.



Fig. 2. **Texture Generation**: We learn self-supervised texture generation from collections of 2D images. An RGB image is encoded by pretrained encoder DECA [Feng et al. 2021] to extract shape $\vec{\beta}$, pose $\vec{\theta}$ and expression $\vec{\psi}$ coefficients in FLAME's latent space, which are decoded by FLAME [Li et al. 2017] to deformed mesh vertices. The background and mouth interior are then masked out, generating the 'Real Image' for our adversarial formulation. In parallel, latent code $z \in \mathbb{R}^{512}$ is sampled from a gaussian distribution $\mathcal{N}(0, I)$ and input to mapping network M to generate intermediate latent $w \in \mathbb{R}^{512\times 18}$, which is used by synthesis network G to generate the UV texture image. The predicted texture is differentiably rendered on a randomly deformed FLAME mesh to generate the 'Fake Image'. Two discriminators interpret the generated and masked real image, at full resolution and at patch size 64×64 . Frozen models are denoted in blue, and learnable networks in green.

Text-Guided 3D Manipulation: Following the success of textguided image manipulation, recent works have adopted powerful vision-language models to enable text guidance for 3D object manipulation. Text2Mesh [Michel et al. 2022] was one of the first pioneering methods to leverage a pre-trained 2D CLIP model as guidance to generate language-conditioned 3D mesh textures and geometric offsets. Here, edits are realized as part of a test-time optimization that aims to solve for the texture and mesh offsets in a neural field representation, such that their re-renderings minimizes a 2D CLIP loss from different viewpoints. Similar to Text2Mesh, CLIP-Mesh [Khalid et al. 2022] produces textured meshes by jointly estimating texture and deformation of a template mesh, based on text inputs using CLIP. Recently, Canfes et al. [2022] adapt TB-GAN [Gecer et al. 2020], an expression-conditioned generative model to produce UV-texture maps, with a CLIP loss to produce facial expressions in 3D, although the quality of textures and expressions is limited, due to reliance on 3D scan data for TB-GAN training. Our method is inspired by these lines of research; however, our focus lies on leveraging the parametric representations of 3D morphable face models with our StyleGAN-based texture generator, which can enable content creation for direct use in many applications such as games or movies.

3 METHOD

ClipFace targets text-guided synthesis of textured 3D face models. It consists of two fundamental components: (i) an expressive generative texture space for facial appearances (Sec 3.1), and (ii) a text-guided prediction of the latent codes for the texture generator and the expression parameters of the underlying statistical morphable model (Sec 3.2). In the following, we will detail these contributions and further demonstrate how they enable producing temporally changing textures for a given animation sequence Sec. 3.3.

3.1 Generative Synthesis of Face Appearance

Since there does not exist any large-scale datasets for UV textures, we propose a self-supervised method to learn the appearance manifold of human faces, as depicted in Fig. 2. Rather than learning from ground truth UV textures, we instead leverage large-scale RGB image datasets of faces, which we use in an adversarial formulation through differentiable rendering. For our experiments, we use the FFHQ dataset [Karras et al. 2019] which consists of 70,000 diverse, high-quality images. As we focus on the textures of the human head, we remove images that contain headwear (caps, scarfs, etc.) and eyewear (sunglasses and spectacles) using face parsing [zllrunning 2018], resulting in a subset of 45,000 images. Based on this data, we train a StyleGAN-ADA [Karras et al. 2020a] generator to produce UV textures that when rendered on top of the FLAME mesh [Li et al. 2017] results in realistic imagery.

More specifically, we use the FLAME model as our shape prior to produce different geometric shapes and facial expressions. It can be defined as:

$$\mathcal{F}(\vec{\beta}, \vec{\theta}, \vec{\psi}) : \mathbb{R}^{|\beta| \times |\dot{\theta}| \times |\dot{\psi}|} \to \mathbb{R}^{3N}, \tag{1}$$

where $\vec{\beta} \in \mathbb{R}^{100}$ are the shape parameters, $\vec{\theta} \in \mathbb{R}^6$ refers to the jaw and head pose, and $\vec{\psi} \in \mathbb{R}^{50}$ are the expression coefficients.

To recover the distribution of face shapes and expressions from the training dataset, we employ DECA [Feng et al. 2021], a pretrained encoder that takes an RGB image as input and outputs the corresponding FLAME parameters $\vec{\beta}, \vec{\theta}, \vec{\psi}$, including orthographic camera parameters c. We use the recovered parameters to remove the backgrounds from the original images, and only keep the image region that is covered by the corresponding face model. Using this distribution of face geometries and camera parameters $\mathcal{D} \sim [\vec{\beta}, \vec{\theta}, \vec{\psi}, c]$, along with the masked real samples, we train the StyleGAN network using differentiable rendering [Laine et al. 2020]. We sample a latent code $z \in \mathbb{R}^{512}$ from Gaussian distribution $\mathcal{N}(0, I)$ to generate the



Fig. 3. **Text-guided Synthesis of Textured 3D Face Models:** From a given textured mesh with texture code w_{init} , we synthesize various styles by adapting both texture and expression to the target text prompt. w_{init} is input to the texture mappers $\mathcal{T} = [\mathcal{T}^1, ..., \mathcal{T}^{18}]$ to obtain texture offsets $w_{delta} \in \mathbb{R}^{512 \times 18}$ for 18 different levels of w_{init} . The expression mapper \mathcal{E} takes mean latent code $w_{mean} = ||w_{init} + w_{delta}||_2$ as input, and predicts expression offset ψ_{delta} to obtain deformed mesh geometry \mathcal{M}_{tgt} . The generated UV map T_{tgt} and deformed mesh \mathcal{M}_{tgt} are differentiably rendered to generate styles that fit the text prompt.

intermediate latent code w $\in \mathbb{R}^{512 \times 18}$ using the a mapping network M: w = M(z). This latent code w is passed to the synthesis network G to generate UV texture map $T \in \mathbb{R}^{512 \times 512 \times 3}$: T = G(w). This predicted texture T is then rendered on a randomly sampled deformed mesh from our discrete distribution of face geometries \mathcal{D} . We use an image resolution of 512×512 .

Both the generated image and masked real image are then passed to the discriminator during training. To further improve the texture quality, we use a patch discriminator alongside a full-image discriminator, which encourages high-fidelity details in local regions of the rendered images. We apply image augmentations (e.g., color jitter, image flipping, hue/saturation changes) to both full-image and image patches before feeding them to the discriminator. The patch size is set to 64×64 for all of our experiments. Note that the patch discriminator is critical to producing high-frequency texture details; see Section 4.

3.2 Text-guided Synthesis of Textured 3D Models

For a given textured mesh with texture code $w_{init} = \{w_{init}^1, ..., w_{init}^{18}\} \in \mathbb{R}^{512 \times 18}$ in neutral pose θ_{init} and neutral expression ψ_{init} , our goal is to learn optimal offsets w_{delta} , ψ_{delta} for texture and expression respectively defined through text prompts. As a source of supervision, we use a pretrained CLIP [Radford et al. 2021] model due to its high expressiveness, and formulate the offsets as:

$$\mathbf{w}_{\text{delta}}^{*}, \boldsymbol{\psi}_{\text{delta}}^{*} = \underset{\mathbf{w}_{\text{delta}}, \boldsymbol{\psi}_{\text{delta}}}{\arg\min} \mathcal{L}_{\text{total}}, \tag{2}$$

where \mathcal{L}_{total} formulates CLIP guidance and expression regularization, as defined in Eq. 9.

In order to optimize this loss, we learn a texture mapper $\mathcal{T} = [\mathcal{T}^1, ..., \mathcal{T}^{18}]$ and an expression mapper \mathcal{E} . The texture mapper predicts the latent variable offsets across the different levels $\{1, 2, ...18\}$

of the StyleGAN generator:

$$\mathbf{w}_{\text{delta}} = \begin{cases} \mathcal{T}^{1}(\mathbf{w}_{\text{init}}^{1}) \\ \mathcal{T}^{2}(\mathbf{w}_{\text{init}}^{2}) \\ \vdots \\ \mathcal{T}^{18}(\mathbf{w}_{\text{init}}^{18}). \end{cases}$$
(3)

The expression mapper \mathcal{E} learns the expression offsets and takes as input $w_{mean} = ||w_{init} + w_{delta}||_2$, where $w_{mean} \in \mathbb{R}^{512}$ is the mean of 18-different levels of the latent space, and outputs the expression offsets $\boldsymbol{\psi}_{delta}$:

$$\boldsymbol{\psi}_{\text{delta}} = \mathcal{E}(\mathbf{w}_{\text{mean}}).$$
 (4)

We notice that it is critical to design separate texture and expression mappers to maintain disentangled texture and expression spaces. Conditioning the expression mapper on texture codes correlates them meaningfully for realistic expression generation, significantly improving generation quality. We show results for different input conditions in supplemental. We use a 4-layer MLP architecture with LeakyReLU activations for the mappers. The method is shown in Fig. 3.

Naively using a CLIP loss as in StyleClip [Patashnik et al. 2021] to train the mappers tends to result in unwanted identity and/or illumination changes in texture. Thus, we draw inspiration from [Gal et al. 2021], and leverage the CLIP-space direction between the initial style and the to-be-performed manipulation in order to perform consistent and identity-preserving manipulation. We compute the 'text-delta' direction Δt in CLIP-space between the initial text prompt t_{init} and the target text prompt t_{tgt} , indicating which attributes from the initial style should be changed:

$$\Delta t = E_T(t_{tgt}) - E_T(t_{init}), \qquad (5)$$

where E_T refers to the CLIP text encoder. We use the same initial text $t_{init} = A$ photo of a face' for all our experiments and alter the target text prompt t_{tgt} , depending on the desired style change. For example, to generate a Mona Lisa style texture, we use the text prompt $t_{tgt} = A$ photo of a face that looks like Mona Lisa'. Guided

SIGGRAPH '23 Conference Proceedings, August 6-10, 2023, Los Angeles, CA, USA



Fig. 4. **Texture Manipulation for Animation Sequences:** Given a video sequence $\mathcal{V} = [\psi^{1:T}; \theta^{1:T}]$ of T frames, an initial texture code w_{init}, and a text prompt, we synthesize a 3D textured animation to match the text. We concatenate \mathcal{V} to w_{init} across different timestamps to obtain e^t, which is input to the time-shared texture mappers \mathcal{T} to obtain time-dependent textures offsets w^{1:T}_{delta} for all frames. The new texture codes w^{1:T}_{tgt} generated using importance weighting, are then passed to our texture generator *G* to obtain time-dependent UV textures $T^{1:T}_{tgt}$, which are then differentiably rendered to generate the final animation, guided by the CLIP loss across all frames.

via the CLIP-space image direction between the initial rendered image i_{init} and the target rendered image i_{tgt} generated using our texture generator, we train the mapping networks to predict style specified by the target text prompt t_{tgt} .

$$\Delta \mathbf{i} = E_I(\mathbf{i}_{tgt}) - E_I(\mathbf{i}_{init}), \tag{6}$$

where i_{init} is the image rendered with the initial texture and initial flame parameters (w_{init} , β , θ , ψ_{init}), and i_{tgt} is the image with the target texture and target flame parameters (w_{tgt} , β , θ , ψ_{tgt}), and E_I refers to the CLIP image encoder.

Note that we do not alter the pose θ and shape code β of the FLAME model. The CLIP loss \mathcal{L}_{clip} is then computed as:

$$\mathcal{L}_{\text{clip}} = 1 - \frac{\Delta i.\Delta t}{|\Delta i| |\Delta t|}.$$
(7)

In order to prevent the mesh from taking unrealistic expressions, we further regularize the expressions using the Mahalanobis prior as:

$$\mathcal{L}_{\rm reg} = \boldsymbol{\psi}^T \boldsymbol{\Sigma}_{\boldsymbol{\psi}}^{-1} \boldsymbol{\psi}, \tag{8}$$

where Σ_{ψ}^{-1} is the diagonal expression covariance matrix of FLAME model. As we show in our results, this regularization is critical to prevent the 3D morphable model from taking an unrealistic shape.

The full training loss can then be written as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{clip}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}.$$
(9)

Note that we can also only alter the texture without changing expressions by keeping the expression mapper frozen and not fine-tuning it. We pre-train the mapper networks to predict zero-offsets (details in supplemental).

3.3 Texture Manipulation for Video Sequences

Given an expression video sequence, we propose a novel technique to manipulate the textures for every frame of the video guided by a CLIP loss (see Fig. 4). That is, for a given animation sequence $\mathcal{V} = [\theta^{1:T}; \psi^{1:T}]$ of *T* frames, with expression codes $\psi^{1:T} = [\psi^1, \psi^2, ... \psi^T]$, pose codes $\theta^{1:T} = [\theta^1, \theta^2, ... \theta^T]$, and a given texture code w_{init}, we use a multi-layer perceptron as our texture mapper $\mathcal{T} = [\mathcal{T}^1, ..., \mathcal{T}^{18}]$ to generate time-dependent texture offsets w^{1:T}_{delta} for different levels of the texture latent space. This mapper receives as input e^{1:T}, the concatenation of the initial texture code w_{init} with the time-dependent expression and pose code $[\psi^t; \theta^t]$. Mathematically, we have:

$$e^{1:T} = [e^1, \dots e^T]$$
 (10)

$$\mathbf{e}^{t} = [\mathbf{w}_{\text{init}}; \boldsymbol{\psi}^{t}; \boldsymbol{\theta}^{t}], \tag{11}$$

where $\boldsymbol{\psi}^t$ and $\boldsymbol{\theta}^t$ refer to the expression and pose code at timestamp t extracted from sequence \mathcal{V} . Next, we pass $e^{1:T}$ to the time-shared texture mapper \mathcal{T} to obtain texture offsets $w_{\text{delta}}^{1:T}$. To ensure a coherent animation and smooth transition across frames, we weight the predicted offsets $w_{\text{delta}}^{1:T}$ using importance weights $\mathcal{I} = [i_1, ..., i_T]$ extracted from video sequence \mathcal{V} , before adding them to w_{init} :

$$\mathbf{w}_{\text{tgt}}^t = \mathbf{w}_{\text{init}} + i_t \cdot \mathbf{w}_{\text{delta}}^t.$$
(12)

We compute importance weights by measuring the deviation between the neutral shape $[\theta_{neutral}; \psi_{neutral}]$ and per-frame face shape $[\theta^t; \psi^t]$, with by min-max normalization:

$$i_t = \frac{\delta^t - \min(\boldsymbol{\delta}^{1:T})}{\max(\boldsymbol{\delta}^{1:T}) - \min(\boldsymbol{\delta}^{1:T})},$$
(13)

with $\delta^t = \|[\theta_{\text{neutral}}; \psi_{\text{neutral}}] - [\theta^t; \psi^t]\|_2$. The importance weighting ensures that key frames with strong expressions are emphasized. The predicted target latent codes $w_{\text{tgt}}^{1:T}$ are then used to generate the UV maps $T_{\text{tgt}}^{1:T}$, which are differentiably rendered onto the given animation sequence \mathcal{V} .

To train texture mapper \mathcal{T} , we minimize clip loss (Eq. 7) for the given text prompt t_{tgt} and the rendered frames i_{tgt}^t aggregated over T timesteps for all the frames from the video.

$$\mathbf{w}_{\text{delta}}^{1:T} = \underset{\mathbf{w}_{\text{delta}}^{1:T}}{\operatorname{arg\,min}} \sum_{t=1}^{T} \mathcal{L}_{\text{clip}}(\mathbf{t}_{\text{tgt}}, \mathbf{i}_{\text{tgt}}^{t}).$$
(14)

4 RESULTS

We evaluate ClipFace on the tasks of texture generation, text-guided synthesis of textured 3D face models, and text-guided manipulation of animation sequences. For texture generation, we evaluate on standard GAN metrics FID and KID. We evaluate both of these metrics with respect to masked FFHQ images [Karras et al. 2019] (with background and mouth interior masked out) as ground truth, and generated textures rendered at 512×512 resolution for ≈ 45 K rendered textures and ground truth images. For text-guided manipulation, we evaluate perceptual quality using FID & KID and similarity to text prompt using CLIP score, which is evaluated as the cosine similarity to the text prompt using pre-trained CLIP models. We use two different CLIP variants, 'ViT-B/16' and 'ViT-L/14', each on 224×224 pixels as input. We report average scores for these pre-trained variants.

Implementation Details: For our texture generator, we produce 512×512 texture maps. We use an Adam optimizer with a learning rate of 2e-3, batch size 8, gradient penalty 10, and path length regularization 2 for all our experiments. We use a learning rate of 0.005 and 0.0001 for the expression and texture mappers, also using Adam. For differentiable rendering, we use NvDiffrast [Laine et al. 2020]. For the patch discriminator, we use a patch size of 64×64 . We train for 300,000 iterations until convergence. For the text-guided manipulation experiments, we use the same model architecture for expression and texture mappers, a 4-layer MLP with LeakyReLU activations. For CLIP supervision, we use the pretrained 'ViT-B/32' variant. For text manipulation tasks, we train for 5,000 iterations.

Texture Generation. We evaluate the quality of our generated textures and compare with existing unsupervised texture generation methods in Tab. 1 and Fig. 5. ClipFace outperforms other baselines in perceptual quality. Although Slossberg *et al.* [2022] can obtain good textures for the interior face region, it does not synthesize head and ears.

Texture & Expression Manipulation. We compare with CLIPbased texturing techniques for texture manipulation in Fig. 6 and Tab. 2. Note that for comparisons with Text2Mesh [Michel et al. 2022], we follow the authors' suggestion to first perform remeshing to increase vertices from 5023 to 60,000 before optimization. Our approach generates consistently high-quality textures for various

Method	$\mathrm{FID}\downarrow$	$\mathrm{KID}\downarrow$
FlameTex [Feng 2019]	76.627	0.063
Slossberg et al. [2022]	32.794	0.021
Ours (w/o Patch)	16.640	0.013
Ours (w/ Patch)	9.559	0.006

Table 1. Quantitative evaluation of texture quality. Our approach significantly outperforms baselines in both FID and KID scores.

Shivangi Aneja, Justus Thies, Angela Dai, and Matthias Niessner



Fig. 5. Comparison against unsupervised texturing methods. The left column shows the mesh geometry, followed by the textured mesh and corresponding UV map. FlameTex [Feng 2019] generates textures for the full head region, however, the texture quality is fairly limited. Slossberg et. al. [2022] generates plausible textures, but does not take the head and ears into account, limiting its practical applicability. Our approach is able to synthesize diverse texture styles ranging across different skin colors, ethnicities, and demographics. Note that the patch discriminator helps not only to improve texture quality but also in correctly aligning the UV texture with mesh geometry, especially around the mouth region.

prompts, in comparison to baselines. In particular, our texture generator enables effective editing even in small face regions (e.g., lips

Method	$\mathrm{FID}\downarrow$	$\mathrm{KID}\downarrow$	CLIP Score ↑
Latent3d [Canfes et al. 2022]	205.27	0.260	0.227 ±0.041
FlameTex [Feng 2019]	88.95	0.053	0.235 ± 0.053
ClipMatrix [Jetchev 2021]	198.34	0.180	0.243 ±0.049
Text2Mesh [Michel et al. 2022]	219.59	0.185	0.264 ±0.044
Ours	80.34	0.032	0.251 ± 0.059

Table 2. Evaluation of text manipulation. ClipFace effectively matches text prompts while maintaining high perceptual fidelity.

SIGGRAPH '23 Conference Proceedings, August 6-10, 2023, Los Angeles, CA, USA



Fig. 6. Qualitative comparison on texture manipulation: We compare our method against several 3D texturing methods: Latent3D [Canfes et al. 2022], Clip-Matrix [Jetchev 2021], FlameTex [Feng 2019], Text2Mesh [Michel et al. 2022]. Our method obtains consistently superior textures in comparison to the baselines, even capable deftly adapting identity when guided by text prompt.

and eyes). While Text2Mesh yields a high CLIP score, it produces semantically implausible results, as the specified text prompts highly match rendered colors irrespective of the global face context (i.e., which region should be edited). In contrast, our method generates perceptually high-quality face texture, evident in the perceptual KID metric.

We show additional ClipFace texturing results on a wide variety of prompts, including on fictional characters, in Fig. 7, demonstrating our expressive power.

Furthermore, we show results for expression manipulation in Fig. 9. ClipFace faithfully deforms face geometry and texture to match a variety of text prompts, where expression regularization maintains plausible geometry and directional loss enables balanced adaptation of geometry and texture. We refer to the supplemental for more visuals.

Texture Manipulation for Video Sequences. Finally, we show results for texture manipulation for given animation sequences in Fig. 8. ClipFace can produce more expressive animation compared to a constant texture that looks monotonic. We show results for

only 3 frames; however, we refer readers to the supplemental video for more detailed results.

Limitations. Although ClipFace can generate high-quality textures and expressions, it still has various limitations. For instance, our method does not capture accessories like jewelry, headwear, or eyewear, due to our use of the FLAME [Li et al. 2017] model, which does not represent accessories or complex hair. We believe that this could be further improved by augmenting parametric 3D models with artist-designed assets for 3D hair, headwear, or eyewear.

5 CONCLUSION

In this paper, we have introduced ClipFace, a novel approach to enable text-guided editing of textured 3D morphable face models. We jointly synthesize high-quality textures and adapt geometry based on the expressions of the morphable model, in a self-supervised fashion. This enables compelling 3D face generation across a variety of textures, expressions, and styles, based on user-friendly text prompts. We further demonstrate the ability of ClipFace to synthesize of animation sequences, driven by a guiding video sequence.

Shivangi Aneja, Justus Thies, Angela Dai, and Matthias Niessner



Fig. 7. Texture manipulations. 'Original' (first column) shows the initial textured mesh sampled from our ClipFace generator without any text prompt, followed by ClipFace-generated textures for various text prompts.



Fig. 8. Given a 3D face motion sequence (top row), we compare our dynamic texturing approach (bottom row) against static-only texturing (middle row). Geo + Tex shows textures overlaid on the animated mesh, and Tex Only shows texture in the neutral pose. Our proposed dynamic texture manipulation technique generates more compelling animation, particularly in articulated expressions (e.g., t=23). Here, we show results for the text prompt "Laughing". We further refer interested readers to our supplemental video.

SIGGRAPH '23 Conference Proceedings, August 6-10, 2023, Los Angeles, CA, USA



Fig. 9. ClipFace generates a large variety of expressions, faithfully deforming the mesh geometry and manipulating texture for more expressiveness. Our expression regularization is important to produce realistic geometric expressions, and directional loss for balanced manipulation of both texture and geometry.

We believe this is an important first step towards enabling controllable, realistic texture and expression modeling for 3D face models, dovetailing with conventional graphics pipelines, which will enable many new possibilities for content creation and digital avatars.

ACKNOWLEDGMENTS

This work was supported by the ERC Starting Grant Scan2CAD (804724), the Bavarian State Ministry of Science and the Arts and coordinated by the Bavarian Research Institute for Digital Transformation (bidt), the German Research Foundation (DFG) Grant "Making Machine Learning on Static and Dynamic 3D Data Practical," the German Research Foundation (DFG) Research Unit "Learning and Simulation in Visual Computing," and Sony Semiconductor Solutions Corporation. We would like to thank Yawar Siddiqui for help with implementation of differential rendering codebase, Artem Sevastopolsky and Guy Gafni for helpful discussions, and Alexey Bokhovkin, Chandan Yeshwanth, and Yueh-Cheng Liu for help during internal review.

APPENDIX

We provide additional ablation studies and results in Section 6, network architecture and training details in Section 7, and further discussion of baseline method experimental setup in Section 8.

6 ADDITIONAL RESULTS

We provide additional results for texture and expression manipulation, as well as an additional comparison to texturing baselines. For results related to video animation, we refer readers to the supplementary video.

Additional Baseline Comparisons: We compare our method with texturing baselines for additional text prompts, and show results in Figure 10. Our method outperforms these baselines and achieves high quality manipulations.

Expression Manipulation: We show additional results for expression manipulation, and analyze the effect of our expression regularization and directional clip loss in Figure 11. Our proposed technique outperforms others and achieves realistic texture manipulation.

Texture Manipulation: We show additional texture manipulation results with a large variety of text prompts in Figure 12. As can be seen, our method is able to generate a wide variety of textures, even capable of adapting identity when implied by the text prompt.

Effect of Pre-training: Here, we analyze the effect of pre-training the texture and expression mappers. We first pre-train the texture mapper \mathcal{T} and expression mapper \mathcal{E} to predict zero offsets using ℓ_2 regularization before training them with our CLIP loss. We show the effect of pre-training mapper networks to predict zero offsets in Figure 13. Without pre-training, the textures begin with unrealistic values and converge to low quality styles with visible artifacts.

Effect of input conditioning for Expression Mapper: Finally, we analyze the effect of different input conditions on the generated expressions during expression manipulation. Conditioning the expression mapper \mathcal{E} on initial expression code ψ_{init} generates uncanny textures, does not noticeably alter the geometry and attempts to encode all information into the texture. Our proposed approach conditions the \mathcal{E} on mean texture code w_{mean} providing a meaning-ful signal to modify texture and expression cohesively. Results are shown in Figure 14.

7 ARCHITECTURE & TRAINING DETAILS

ClipFace is implemented in the Pytorch Lightning framework [Falcon et al. 2019; Paszke et al. 2017]. For differentiable rendering, we use the NvDiffrast [Laine et al. 2020] library.

Texture Generation: For texture generation, we use the Style-GAN2 architecture with adaptive discriminator augmentation [Karras et al. 2020a]. For all of our experiments, we operate at a texture and image resolution of 512×512 . We apply augmentation to both the full-image discriminator as well as the patch discriminator, which operates on patches of size 64×64 . For augmentations, we apply geometric transformations such as image flipping, rotation

and scaling, as well as color transformations such as changing image brightness, contrast, hue, saturation, etc. For training, we used the Adam [Kingma and Ba 2014] optimizer with a learning rate of 0.002, batch size 8 per GPU, gradient penalty 10, and path length regularization 2. We perform multi-GPU training on 3 RTX A6000 GPUs and train for 300.000 iterations.

Texture and Expression Manipulation: For the text-guided manipulation experiments, we use a 4-layer MLP architecture with LReLU activations. For the texture mapper, we use 18 identical MLPs to predict texture offsets for different levels of the latent code $w_{init} = \{w_{init}^1, w_{init}^2, ..., w_{init}^{18}\} \in \mathbb{R}^{512 \times 18}$.

$$\mathbf{w}_{\text{delta}} = \begin{cases} \mathcal{T}^{1}(\mathbf{w}_{\text{init}}^{1}) \\ \mathcal{T}^{2}(\mathbf{w}_{\text{init}}^{2}) \\ \vdots \\ \mathcal{T}^{18}(\mathbf{w}_{\text{init}}^{18}). \end{cases}$$
(15)

Each texture MLP \mathcal{T}^i takes as input 512-dimensional latent code w_{init}^i and outputs the 512-dimensional offset w_{delta}^i . The expression mapper \mathcal{E} takes the mean latent code as input $w_{mean} \in \mathbb{R}^{512}$, and predicts the expression offset $\psi_{delta} \in \mathbb{R}^{50}$ as output. The network architecture for both these mappers is shown in Figure 15. For language supervision, we use the CLIP model [Radford et al. 2021]. For our experiments, we use the pre-trained 'ViT-B/32' variant for computing the CLIP loss. We use a learning rate of 0.005 for the expression mapper and 0.0001 for the texture mapper.

Texturing for Animation Sequences: For the task of texturing for animations, we learn only the texture mapper. We use the same architecture as shown in Figure 15(a). Since a given animation sequence consists of multiple frames, we share the texture mapper across different timestamps. For all text manipulation experiments, we train for 20,000 iterations.

8 BASELINE IMPLEMENTATIONS

Latent3d [Canfes et al. 2022]: This method builds upon the TB-GAN [Gecer et al. 2020], a generative model \mathcal{G} that takes onehot encoded facial expression vector \vec{e} and a random noise vector $\vec{z} \in \mathbb{R}^d$ as input and generates shape, shape-normal and texture images. Given a pretrained generator \mathcal{G} , the method optimizes offset Δc for the intermediate layer c which is 4×4 dense layer of TB-GAN. The offset Δc gives the direction in which the target attributes specified by text prompt t are enhanced, while other attributes stay unchanged. The authors use a Clip-loss $\mathcal{L}_{\text{CLIP}}$, supplemented with an identity loss \mathcal{L}_{ID} and L2 regularization \mathcal{L}_{L2} to perform meaningful manipulation of meshes:

$$\underset{\Delta \mathbf{c} \in C}{\arg\min} \mathcal{L}_{\text{CLIP}} + \lambda_{\text{ID}} \mathcal{L}_{\text{ID}} + \lambda_{\text{L2}} \mathcal{L}_{\text{L2}}, \tag{16}$$

where λ_{ID} and λ_{L2} are hyperparameters for the \mathcal{L}_{ID} and \mathcal{L}_{L2} respectively. The identity loss \mathcal{L}_{ID} minimizes the distance between the identity of original renders and manipulated renders:

$$\mathcal{L}_{\rm ID} = 1 - \langle R(\mathcal{G}(c)), R(\mathcal{G}(c + \Delta c)) \rangle, \tag{17}$$

where *R* is the ArcFace [Deng et al. 2019], a facial recognition network and $\langle ., . \rangle$ computes the cosine similarity between the identities

SIGGRAPH '23 Conference Proceedings, August 6-10, 2023, Los Angeles, CA, USA



Fig. 10. Additional texture manipulation results in comparison with baselines, from the text prompts shown. Our method outperforms baseline approaches in texture quality.

of the initial rendering and manipulated rendering. The L2 loss is used by authors to prevent artifact generation and can be written as:

$$\mathcal{L}_{L2} = \|c - (c + \Delta c)\|_2$$
(18)

For the Clip loss, the authors use a list of text templates like 'a photo of a ...'; 'a face of a ...', etc prefixed to the target style:

$$\mathcal{L}_{\text{CLIP}} = \frac{\sum_{j=1}^{K} \sum_{i=1}^{N} D_{\text{CLIP}}(\mathcal{I}_i, t_j)}{K \cdot N},$$
(19)

where I_i is the rendered image from a list of N rendered images, t_j is the target text t embedded in a text template from a list of K templates. D_{CLIP} minimizes the cosine distance between CLIP embeddings of the rendered image I_i and the set of text prompts t_j .

ClipMatrix [Jetchev 2021]: Given a 3D mesh and initial texture map T_{init} , ClipMatrix optimizes the texture image offset T_{delta} to match the image rendering I to the text prompt t from random

camera view c:

$$\mathcal{L}(T_{delta}) = \sum_{t} \mathop{\mathbb{E}}_{c \sim \pi_c} \mathcal{L}_{\text{CLIP}}(I, t).$$
(20)

By sampling from random camera angles $c \sim \pi_c$ during optimization, the method ensures that output mesh shows the desired properties from different viewing angles. The image rendering can be obtained as:

$$I = \mathcal{R}(\mathcal{M}, T_{tqt}, c), \tag{21}$$

where \mathcal{M} refers to the 3D mesh, $T_{tgt} = T_{init} + T_{delta}$ denotes the final UV texture map and *c* denotes the camera view. The clip loss $\mathcal{L}_{\text{CLIP}}(I, t)$ minimizes the negative cosine similarity in CLIP embedding space between image *I* and the fixed text prompt *t*.

$$\mathcal{L}_{\text{CLIP}}(I,t) = -\cos(\phi_i(I),\phi_t(t))$$
(22)

where ϕ_i and ϕ_t refer to the clip image and text encoder respectively. The texture offset T_{delta} is initialized with zero and during optimization clipping is applied to final texture image $T_{tgt} \in [-1, 1]$ to ensure that it stays in valid image range.

Shivangi Aneja, Justus Thies, Angela Dai, and Matthias Niessner



Fig. 11. Additional expression manipulation results: our proposed method achieves realistic expression manipulation. Both our expression regularization and directional clip loss contribute notably to realistic output quality.





Fig. 12. Additional texturing results: Our proposed method generates a diverse range of textures.

Shivangi Aneja, Justus Thies, Angela Dai, and Matthias Niessner



Fig. 13. We evaluate the effect of pre-training texture and expression mappers. w/o Pretrain' refers to the case when mappers are not trained to predict zero offsets before performing text manipulation. Pre-training helps to produce realistic texture changes.



Fig. 14. We evaluate the effect of different conditioning inputs on the expression mapper \mathcal{E} during expression manipulation. 'w/ initial expression cond' refers to the case when expression mapper is conditioned on initial expression code ψ_{init} , 'w/ texture code cond' refers to the case when \mathcal{E} is conditioned on mean texture code w_{mean} . Conditioning on w_{mean} provides the meaningful signal to modify texture and expression cohesively.

FlameTex [Feng 2019]: FlameTex is the PCA-based texturing model designed specifically for FLAME face model [Li et al. 2017]. The texture space for FlameTex is built using randomly selected 1500 images from the FFHQ dataset [Karras et al. 2019] and the base texture from the Basel Face Model [Paysan et al. 2009]. Given a mean texture $T_{mean} \in \mathbb{R}^{512 \times 512 \times 3}$ and texture basis $T_{basis} \in$

 $\mathbb{R}^{50\times512\times512\times3}$ from the FlameTex texture model, we optimize for the texture basis coefficients $\omega \in \mathbb{R}^{50}$ to match the target text prompt *t* to generate the desired texture map T_{tqt} :

$$\boldsymbol{\omega}^{*} = \arg\min_{\boldsymbol{\omega}} \mathcal{L}_{\text{CLIP}}(I, t) + \lambda_{\text{L2}} \mathcal{L}_{\text{L2}}(\boldsymbol{\omega}), \quad (23)$$





(a) Texture Mapper Architecture

(b) Expression Mapper Architecture

Fig. 15. Architecture Overview: Network architecture for Texture Mapper (left) and Expression mapper (right). FC(x, y) refers to a fully-connected layer, where x and y denote the input and output dimensions, respectively. LReLU refers to LeakyReLU activations.

where $\mathcal{L}_{\text{CLIP}}$ refers to the clip loss between text prompt *t* and the rendered image *I* and \mathcal{L}_{L2} refers to the L2 regularization for the texture coefficients ω with λ_{L2} controlling the strength of the regularization. The desired texture map is given by:

obtained as:

$$I = \mathcal{R}(\mathcal{M}, [f_{\theta}(\gamma(p_i))]_{i=1}^N, c), \tag{29}$$

where \mathcal{R} refers to the differentiable rendering, c refers to the random camera view and $p_i \in \mathbb{R}^3$ refers to the vertex coordinate of the mesh and N refers to the total number of mesh vertices.

$$T_{tqt} = T_{mean} + \omega * T_{basis}.$$
 (24)

The image rendering can be obtained as:

$$I = \mathcal{R}(\mathcal{M}, T_{tqt}, c), \tag{25}$$

where \mathcal{M} refers to the Flame 3D mesh, T_{tgt} denotes the final UV texture map and *c* denotes the camera view. The L2 regularization is applied to prevent the model from producing unrealistic texture and is given by:

$$\mathcal{L}_{L2} = \|\boldsymbol{\omega}\|_2. \tag{26}$$

We initialize ω with zero and start from base texture T_{mean} during optimization.

Text2Mesh [Michel et al. 2022]: Given a 3D mesh, the method uses coordinate-based MLPs to predict per-vertex color and displacement conforming to the target text prompt *t* used for stylizing the mesh. In our experiments, we first remesh the Flame 3D mesh to increase vertices from 5K to 60K as the method works reasonably well for meshes with a higher vertex count. We used the default hyperparameters used by authors for stylizing human body meshes, as the authors did not perform experiments on human face meshes.

For every vertex point p, first the positional encoding γ is applied to obtain high frequency features, before passing them to the MLP:

$$\gamma(p) = [\cos(2\pi Bp), \sin(2\pi Bp)]^T, \qquad (27)$$

where $B \in \mathbb{R}^{n \times 3}$ is the random Gaussian matrix. We first pretrain the MLP f_{θ} to predict a base texture T_{init} to begin learning from a reasonable starting texture. Since we do not wish to change the geometry, we do not perform vertex displacements in our experiments. Our sanity experiments for vertex displacements produced unrealistic geometries. We train with loss function and augmentations proposed in the main paper. The loss function can be written as:

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \mathcal{L}_{\text{CLIP}}(I, t), \tag{28}$$

where I refers to the image rendered from different viewpoints and t refers to the target text prompt. The image rendering I can be

Shivangi Aneja, Justus Thies, Angela Dai, and Matthias Niessner

REFERENCES

- Rameen Abdal, Peihao Zhu, John Femiani, Niloy J. Mitra, and Peter Wonka. 2021a. CLIP2StyleGAN: Unsupervised Extraction of StyleGAN Edit Directions. CoRR abs/2112.05219 (2021). arXiv:2112.05219 https://arxiv.org/abs/2112.05219
- Rameen Abdal, Peihao Zhu, Niloy J. Mitra, and Peter Wonka. 2021b. StyleFlow: Attribute-Conditioned Exploration of StyleGAN-Generated Images Using Conditional Continuous Normalizing Flows. ACM Trans. Graph. (May 2021). https: //doi.org/10.1145/3447648
- Omri Avrahami, Dani Lischinski, and Ohad Fried. 2022. Blended diffusion for text-driven editing of natural images. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 18208–18218.
- David Bau, Alex Andonian, Audrey Cui, YeonHwan Park, Ali Jahanian, Aude Oliva, and Antonio Torralba. 2021. Paint by Word. arXiv:arXiv:2103.10951
- Volker Blanz and Thomas Vetter. 1999. A Morphable Model for the Synthesis of 3D Faces. In Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99). ACM Press/Addison-Wesley Publishing Co., USA, 187–194. https://doi.org/10.1145/311535.311556
- Zehranaz Canfes, M. Furkan Atasoy, Alara Dirik, and Pinar Yanardag. 2022. Text and Image Guided 3D Avatar Generation and Manipulation. https://doi.org/10.48550/ ARXIV.2202.06079
- 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. 2021. Efficient Geometry-aware 3D Generative Adversarial Networks. In arXiv.
- Katherine Crowson. 2021. VQGAN-CLIP. https://github.com/nerdyrodent/VQGAN-CLIP
- Boris Dayma, Suraj Patil, Pedro Cuenca, Khalid Saifullah, Tanishq Abraham, Phuc Le Khac, Luke Melas, and Ritobrata Ghosh. 2021. DALL-E Mini. https://doi.org/10. 5281/zenodo.5146400
- Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. 2019. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 4685–4694. https://doi.org/10. 1109/CVPR.2019.00482
- Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. 2020. Disentangled and Controllable Face Image Generation via 3D Imitative-Contrastive Learning. In *IEEE Computer Vision and Pattern Recognition*.
- William Falcon et al. 2019. Pytorch lightning. GitHub. Note: https://github. com/PyTorchLightning/pytorch-lightning 3, 6 (2019).
- Haven Feng. 2019. Photometric FLAME Fitting. https://github.com/HavenFeng/ photometric_optimization.
- Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. 2021. Learning an Animatable Detailed 3D Face Model from In-the-Wild Images. ACM Transactions on Graphics (ToG), Proc. SIGGRAPH 40, 4 (Aug. 2021), 88:1–88:13.
- Rinon Gal, Or Patashnik, Haggai Maron, Gal Chechik, and Daniel Cohen-Or. 2021. StyleGAN-NADA: CLIP-Guided Domain Adaptation of Image Generators. arXiv:2108.00946 [cs.CV]
- Baris Gecer, Jiankang Deng, and Stefanos Zafeiriou. 2021a. OSTeC: One-Shot Texture Completion. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 7628–7638.
- Baris Gecer, Alexander Lattas, Stylianos Ploumpis, Jiankang Deng, Athanasios Papaioannou, Stylianos Moschoglou, and Stefanos Zafeiriou. 2020. Synthesizing Coupled 3D Face Modalities by Trunk-Branch Generative Adversarial Networks. In *Proceedings* of the European conference on computer vision (ECCV). Springer.
- Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. 2019. GANFIT: Generative Adversarial Network Fitting for High Fidelity 3D Face Reconstruction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos P Zafeiriou. 2021b. Fast-GANFIT: Generative Adversarial Network for High Fidelity 3D Face Reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021).
- Thomas Gerig, Andreas Morel-Forster, Clemens Blumer, Bernhard Egger, Marcel Lüthi, Sandro Schönborn, and Thomas Vetter. 2017. Morphable Face Models - An Open Framework. https://doi.org/10.48550/ARXIV.1709.08398
- Partha Ghosh, Pravir Singh Gupta, Roy Uziel, Anurag Ranjan, Michael J. Black, and Timo Bolkart. 2020. GIF: Generative Interpretable Faces. In International Conference on 3D Vision (3DV). 868–878. http://gif.is.tue.mpg.de/
- Fangzhou Hong, Mingyuan Zhang, Liang Pan, Zhongang Cai, Lei Yang, and Ziwei Liu. 2022. AvatarCLIP: Zero-Shot Text-Driven Generation and Animation of 3D Avatars. ACM Transactions on Graphics (TOG) 41, 4 (2022), 1–19.
- Nikolay Jetchev. 2021. ClipMatrix: Text-controlled Creation of 3D Textured Meshes. https://doi.org/10.48550/ARXIV.2109.12922
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In International Conference on Learning Representations. https://openreview.net/forum?id=Hk99zCeAb
- Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. 2020a. Training Generative Adversarial Networks with Limited Data. In Proc.

NeurIPS.

- Tero Karras, Samuli Laine, and Timo Aila. 2019. A Style-Based Generator Architecture for Generative Adversarial Networks. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 4396–4405. https://doi.org/10.1109/CVPR. 2019.00453
- Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020b. Analyzing and Improving the Image Quality of StyleGAN. In Proc. CVPR.
- Nasir Mohammad Khalid, Tianhao Xie, Eugene Belilovsky, and Popa Tiberiu. 2022. CLIP-Mesh: Generating textured meshes from text using pretrained image-text models. (December 2022).
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. https://doi.org/10.48550/ARXIV.1412.6980
- Umut Kocasari, Alara Dirik, Mert Tiftikci, and Pinar Yanardag. 2022. StyleMC: Multi-Channel Based Fast Text-Guided Image Generation and Manipulation. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). 895– 904.
- Marek Kowalski, Stephan J. Garbin, Virginia Estellers, Tadas Baltrušaitis, Matthew Johnson, and Jamie Shotton. 2020. CONFIG: Controllable Neural Face Image Generation. In European Conference on Computer Vision (ECCV).
- Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. 2020. Modular Primitives for High-Performance Differentiable Rendering. ACM Transactions on Graphics 39, 6 (2020).
- Alexandros Lattas, Stylianos Moschoglou, Baris Gecer, Stylianos Ploumpis, Vasileios Triantafyllou, Abhijeet Ghosh, and Stefanos Zafeiriou. 2020. AvatarMe: Realistically Renderable 3D Facial Reconstruction "In-the-Wild". In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Abhijeet Ghosh, and Stefanos P Zafeiriou. 2021. AvatarMe++: Facial Shape and BRDF Inference with Photorealistic Rendering-Aware GANs. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021).
- Myunggi Lee, Wonwoong Cho, Moonheum Kim, David I. Inouye, and Nojun Kwak. 2020. StyleUV: Diverse and High-fidelity UV Map Generative Model. *ArXiv* abs/2011.12893 (2020).
- Ruilong Li, Karl Bladin, Yajie Zhao, Chinmay Chinara, Owen Ingraham, Pengda Xiang, Xinglei Ren, Pratusha Prasad, Bipin Kishore, Jun Xing, and Hao Li. 2020. Learning Formation of Physically-Based Face Attributes. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Tianye Li, Timo Bolkart, Michael. J. Black, Hao Li, and Javier Romero. 2017. Learning a model of facial shape and expression from 4D scans. ACM Transactions on Graphics, (Proc. SIGGRAPH Asia) 36, 6 (2017), 194:1–194:17.
- Yuchen Liu, Zhixin Shu, Yijun Li, Zhe Lin, Richard Zhang, and S. Y. Kung. 2022. 3D-FM GAN: Towards 3D-Controllable Face Manipulation. ArXiv abs/2208.11257 (2022).
- Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. 2015. SMPL: A Skinned Multi-Person Linear Model. ACM Trans. Graphics (Proc. SIGGRAPH Asia) 34, 6 (Oct. 2015), 248:1–248:16.
- Huiwen Luo, Koki Nagano, Han-Wei Kung, Qingguo Xu, Zejian Wang, Lingyu Wei, Liwen Hu, and Hao Li. 2021. Normalized Avatar Synthesis Using StyleGAN and Perceptual Refinement. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 11662–11672.
- Richard T. Marriott, Sami Romdhani, and Liming Chen. 2021. A 3D GAN for Improved Large-pose Facial Recognition. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021), 13440–13450.
- Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. 2022. Text2Mesh: Text-Driven Neural Stylization for Meshes. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 13492– 13502.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. (2017).
- Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. 2021. StyleCLIP: Text-Driven Manipulation of StyleGAN Imagery. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 2085–2094.
- Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. 2009. A 3D Face Model for Pose and Illumination Invariant Face Recognition. 2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance (2009), 296–301.
- Mathis Petrovich, Michael J. Black, and Gül Varol. 2022. TEMOS: Generating diverse human motions from textual descriptions. In European Conference on Computer Vision (ECCV).
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139), Marina Meila and Tong Zhang (Eds.). PMLR, 8748–8763. https://proceedings.mlr.press/

SIGGRAPH '23 Conference Proceedings, August 6-10, 2023, Los Angeles, CA, USA

v139/radford21a.html

- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. https://doi. org/10.48550/ARXIV.2204.06125
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-Shot Text-to-Image Generation. https: //doi.org/10.48550/ARXIV.2102.12092
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2022. DreamBooth: Fine Tuning Text-to-image Diffusion Models for Subject-Driven Generation. (2022).
- Ron Slossberg, Ibrahim Jubran, and Ron Kimmel. 2022. Unsupervised High-Fidelity Facial Texture Generation and Reconstruction. In Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIII (Tel Aviv, Israel). Springer-Verlag, Berlin, Heidelberg, 212–229. https://doi.org/10. 1007/978-3-031-19778-9_13
- Ayush Tewari, Mohamed Egharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zöllhofer, and Christian Theobalt. 2020a. StyleRig: Rigging StyleGAN for 3D Control over Portrait Images, CVPR 2020. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE.

- Ayush Tewari, Mohamed Elgharib, Mallikarjun B R, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhöfer, and Christian Theobalt. 2020b. PIE: Portrait Image Embedding for Semantic Control. ACM Trans. Graph. (2020).
- Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. 2022a. CLIP-NeRF: Text-and-Image Driven Manipulation of Neural Radiance Fields. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 3835– 3844.
- Lizhen Wang, Zhiyuan Chen, Tao Yu, Chenguang Ma, Liang Li, and Yebin Liu. 2022b. FaceVerse: A Fine-Grained and Detail-Controllable 3D Face Morphable Model From a Hybrid Dataset. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 20333–20342.
- Tianyi Wei, Dongdong Chen, Wenbo Zhou, Jing Liao, Zhentao Tan, Lu Yuan, Weiming Zhang, and Nenghai Yu. 2022. HairCLIP: Design Your Hair by Text and Reference Image. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 18072–18081.
- Kim Youwang, Kim Ji-Yeon, and Tae-Hyun Oh. 2022. CLIP-Actor: Text-Driven Recommendation and Stylization for Animating Human Meshes. In ECCV.
- zllrunning. 2018. face-parsing.PyTorch. https://github.com/zllrunning/face-parsing. PyTorch.