

# Neural Parameterization for Dynamic Human Head Editing

LI MA\*, The Hong Kong University of Science and Technology, China

XIAOYU LI, Tencent AI Lab, China

JING LIAO, City University of Hong Kong, China

XUAN WANG, Tencent AI Lab, China

QI ZHANG, Tencent AI Lab, China

JUE WANG, Tencent AI Lab, China

PEDRO V. SANDER, The Hong Kong University of Science and Technology, China




Fig. 1. Given synchronized multi-view videos, our method reconstructs a volumetric representation (a,d) that enables geometry and appearance editing. We demonstrate applications of our method in geometry editing (b) appearance editing (e), and joint editing (c,f).

Implicit radiance functions emerged as a powerful scene representation for reconstructing and rendering photo-realistic views of a 3D scene. These representations, however, suffer from poor editability. On the other hand,

\* Author did this work during the internship at Tencent AI Lab.

Authors' addresses: Li Ma, The Hong Kong University of Science and Technology, Hong Kong, China, lmaag@connect.ust.hk; Xiaoyu Li, Tencent AI Lab, Shenzhen, China, lixiaoyu306@gmail.com; Jing Liao, City University of Hong Kong, Hong Kong, China, jingliao@cityu.edu.hk; Xuan Wang, Tencent AI Lab, Shenzhen, China, xwang.cv@gmail.com; Qi Zhang, Tencent AI Lab, Shenzhen, China, nwpuzhang@gmail.com; Jue Wang, Tencent AI Lab, Shenzhen, China, arphid@gmail.com; Pedro V. Sander, The Hong Kong University of Science and Technology, Hong Kong, China, psander@cse.ust.hk.

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 ACM 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.

© 2022 Association for Computing Machinery.

0730-0301/2022/12-ART236 \$15.00

<https://doi.org/10.1145/3550454.3555494>

explicit representations such as polygonal meshes allow easy editing but are not as suitable for reconstructing accurate details in dynamic human heads, such as fine facial features, hair, teeth, and eyes. In this work, we present Neural Parameterization (NeP), a hybrid representation that provides the advantages of both implicit and explicit methods. NeP is capable of photo-realistic rendering while allowing fine-grained editing of the scene geometry and appearance. We first disentangle the geometry and appearance by parameterizing the 3D geometry into 2D texture space. We enable geometric editability by introducing an explicit linear deformation blending layer. The deformation is controlled by a set of sparse key points, which can be explicitly and intuitively displaced to edit the geometry. For appearance, we develop a hybrid 2D texture consisting of an explicit texture map for easy editing and implicit view and time-dependent residuals to model temporal and view variations. We compare our method to several reconstruction and editing baselines. The results show that the NeP achieves almost the same level of rendering accuracy while maintaining high editability.

CCS Concepts: • **Computing methodologies** → **Rendering**.

Additional Key Words and Phrases: Neural Rendering, Scene Representation, Editable Neural Radiance Field, Dynamic Scenes

**ACM Reference Format:**

Li Ma, Xiaoyu Li, Jing Liao, Xuan Wang, Qi Zhang, Jue Wang, and Pedro V. Sander. 2022. Neural Parameterization for Dynamic Human Head Editing. *ACM Trans. Graph.* 41, 6, Article 236 (December 2022), 15 pages. <https://doi.org/10.1145/3550454.3555494>

**1 INTRODUCTION**

Modeling, editing and rendering photo-realistic scenes have a wide range of applications in computer games, movie industry, virtual and augmented reality. Traditionally, polygonal meshes combined with texture maps have been adopted as the standard 3D representations in rendering pipelines. Since the geometry and appearance are explicitly modeled, editing can be done naturally by deforming the meshes or modifying the texture maps. However, despite the long-lasting development of mesh reconstruction methods [Huang et al. 2018; Luo et al. 2019; Schönberger and Frahm 2016; Wei et al. 2020], it is still challenging to reconstruct accurate, well-aligned meshes for scenes like dynamic human heads from video, especially at the places of hair, teeth, and the eyes, which exhibit complex geometry and appearance.

To break through the limited representation capacity of mesh-based methods, implicit volumetric representations have gained considerable attention recently for photo-realistic rendering. In particular, neural radiance field (NeRF) [Mildenhall et al. 2020] models scenes as some continuous functions defined densely in a 3D volume. By parameterizing the function using a multilayer perceptron (MLP), scenes are implicitly encoded into the parameters of the neural networks, which is more compact than explicitly storing values in a dense 3D grid, while allowing the volume to model complex scenes at arbitrary resolution. However, encoding the entire scene into the parameters of the neural networks implicitly makes this representation more obscure for interpreting and editing, which greatly restricts the practical usage of these methods.

In this work, we address the challenge of reconstructing high-fidelity dynamic scenes with implicit representations while allowing explicit editing of both the geometry and the appearance. Existing works usually train implicit representations that is conditioned on some latent codes. Scene editing can then be implemented by propagating the editing from the rendered results back to the latent codes [Deng et al. 2021; Liu et al. 2021], or using another encoder network to directly map the editing target to the parameters [Chiang et al. 2022; Wang et al. 2021a]. Similar ideas have also been introduced into implicit human head representations [Gafni et al. 2021; Hong et al. 2022; Sun et al. 2021; Zheng et al. 2022, 2021]. However, latent-code-based editing usually does not allow fine-grained and out-of-domain editing.

NeuTex [Xiang et al. 2021] proposes to model the appearance of implicit scene representation using 2D texture maps. The appearance of the scene can then be edited by modifying the 2D texture maps. However, NeuTex only focuses on static scenes, and how to edit the scene geometry is also unclear. Neural Atlas [Kasten et al. 2021] introduces a similar idea for consistent video editing. It synthesizes each frame of a video by mapping each pixel location to texture atlas space shared by the entire video sequence. The video can then be consistently modified by editing the texture atlas. We adopt the same idea for appearance modeling and editing.

In this work, we present Neural Parameterization (NeP), a volumetric hybrid implicit-explicit 3D representation for dynamic human heads. NeP achieves photo-realistic rendering while allowing for explicit editing of both geometry and appearance. We draw inspiration from traditional texture mapping and propose to decompose a 3D dynamic human head into three components: a density volume, a UV volume and a 2D texture. The decomposition enables the disentanglement of geometry and appearance, which allows us to edit them separately. The density volume, UV volume and the texture are all modeled implicitly as MLPs in the first place to improve the reconstruction quality, and to reduce the memory requirement when modeling the dynamic view-dependent volume. To enable explicit editing of both geometry and appearance, we introduce explicit layers into the implicit MLPs for these two components. The explicit geometry layer is modeled as a temporally-varying 3D warping field controlled by semantic keypoints. This allows us to freely deform the geometry on one frame and propagate the deformation to the other frames for consistent geometry editing. For appearance, the editability is provided by using an explicit texture map that is shared by all the frames. By editing the texture map, we achieve consistent appearance editing for different views and time.

We reconstruct the NeP from multi-view videos captured by a set of synchronized and calibrated cameras. Freely optimizing the hybrid representation using the multi-view videos will lead to sub-optimal solutions that is unsuitable for the purpose of editing, so we design several regularization methods and a two-stage optimizing strategy in the training process. We also develop a user interface that enables interactive editing and preview using an extracted mesh representation. Extensive experiments show that NeP achieves photo-realistic rendering results and outperforms several baselines in terms of editing capability. To summarize, our key technical contributions are:

- A novel hybrid representation NeP that enables intuitive and consistent editing on both geometry and appearance of dynamic human heads.
- Regularization methods and a training strategy that greatly improves the editability of NeP while maintaining a high-fidelity reconstruction of human heads.
- Several applications are made possible in 3D by using our method, like virtual makeup, artistic stylization and face shape editing, with a UI for interactive editing.

**2 RELATED WORK**

Since our editing method is based on a new scene representation, we review existing works on related scene representations, from explicit mesh to the current state-of-the-art implicit neural representation. In particular, we will focus more on 3D facial reconstruction and editing methods since our main interest is 3D human heads.

*Mesh-Based Reconstruction and Editing.* The polygon mesh is the most commonly used geometry representation in real-time applications, such as in the game and movie industry. Mesh-based methods explicitly model the geometry of 3D surfaces. To reconstruct a 3D mesh from images captured from the real world, classical methods typically use structure from motion [Griwodz et al. 2021; Schönberger and Frahm 2016; Schönberger et al. 2016] to get a dense point



cloud, and then convert it to the triangle mesh using surface reconstruction algorithms [Bernardini et al. 1999; Kazhdan et al. 2006]. To improve the reconstruction quality, there are several attempts that introduce differentiable rendering to the mesh rendering, such as rasterization [Cole et al. 2021; Kato et al. 2018; Liu et al. 2019; Loper and Black 2014; Ravi et al. 2020] or ray tracing [Li et al. 2018; Loubet et al. 2019; Luan et al. 2021]. By making the rendering process differentiable, the loss can be back propagated from the rendering results to the scene parameters such as vertex positions.

The human face is a strong prior that can be used to improve the reconstruction quality. The prior is usually modeled using a 3D morphable model (3DMM) [Blanz and Vetter 1999; Booth et al. 2018; Gerig et al. 2018; Li et al. 2017; Paysan et al. 2009]. The geometry of the face is determined by morphing a pre-defined template face mesh. In a typical linear 3DMM, the morphing is parameterized as a linear combination of a fixed set of basis offsets. By learning a set of principle basis offsets from a large 3D face dataset, the 3DMM model can be used to generate arbitrary faces using the coefficient vector. The problem of mesh reconstruction is then simplified to computing the coefficient vector of the face. The reconstruction of the 3DMM can be modeled as an analysis-by-synthesis problem using differentiable rendering [Feng et al. 2021; Gecer et al. 2019; Genova et al. 2018; Wang et al. 2022], or directly regressed using deep neural networks [Guo et al. 2020].

With accurate reconstructed geometry, the mesh-based model can achieve impressive rendering results. Moreover, due to the explicitness of this representation, meshes with texture maps can be edited naturally. Many tools have been developed to edit the mesh representation, such as mesh deformation [Sorkine and Alexa 2007], smoothing [Sorkine et al. 2004], subdivision [Catmull and Clark 1978]. Unfortunately, despite the progress in generic surface and face reconstruction, an accurate mesh is still hard to acquire for human heads, especially for hair, eyes, and the mouth interior [Lombardi et al. 2021]. One option to compensate for the inaccurate reconstruction is to use neural texture representations [Ma et al. 2021; Riegler and Koltun 2020a,b; Thies et al. 2019] or deep appearance models [Bi et al. 2021; Lombardi et al. 2018]. However, the use of implicit textures corrupts the editability of the representation.

*Volumetric Representation.* Volumetric representation models the scene by densely storing parameters in a 3D regular grid, such as signed distance [Kar et al. 2017], colors [Lombardi et al. 2019] and radiance function [Yu et al. 2021]. Neural Volumes [Lombardi et al. 2019] use a decoder structure to produce a volume on the fly from a latent code. MVP [Lombardi et al. 2021] reduces the memory requirement by attaching small volumetric primitives on a coarse mesh proxy, so the empty space is not modeled. One variation is multiplane images (MPIs) [Flynn et al. 2019; Mildenhall et al. 2019; Srinivasan et al. 2019; Zhou et al. 2018], which models the scene using several fronto-parallel RGB $\alpha$  image planes in the frustum of a reference camera. Unlike meshes, volumetric representation is able to model complex geometry like thin structures. Besides, the rendering of volumetric representation is naturally differentiable, which makes it easy to optimize a volume that achieves photorealistic rendering results. One drawback of the volumetric representation

is the large memory consumption for densely storing the scene parameters, especially for dynamic scenes. In terms of editing, directly manipulating each voxel would be tedious and error-prone. There is a lack of more advanced tools to assist the editing of the volumetric representation.

*Implicit Representation.* Recently, many works propose to implicitly model a volume function by coordinate-based Multi-Layer Perceptrons (MLPs) instead of explicitly storing values in a regular grid. The implicit representation can be used to model signed distance function [Niemeyer et al. 2020; Wang et al. 2021b; Yariv et al. 2021, 2020], occupancy function [Mescheder et al. 2019], radiance fields [Mildenhall et al. 2020], BRDF parameters [Bi et al. 2020; Boss et al. 2021; Srinivasan et al. 2021]. Especially, NeRF [Mildenhall et al. 2020] proposes to model the 3D scene by mapping the spatial location and view direction to the density and color of the point. By optimizing the neural network using multi-view input images, NeRF achieves impressive novel view synthesis quality. To extend the NeRF to represent dynamic scenes, modulation-based methods [Li et al. 2021c; Xian et al. 2021] directly provide time information to the MLP, while deformation-based methods [Park et al. 2021a; Pumarola et al. 2021] introduce implicit deformation field to non-rigidly warp a template NeRF. Park et al. [2021b] proposes to combine the two methods, where the template NeRF is also conditioned on the time information. Implicit representation has also been used to help reconstruct the human head. A common approach is to condition the implicit field on some latent vectors, such as geometry and appearance code [Li et al. 2021a], expression and identity code [Zheng et al. 2022], latent code sampled from a prior distribution [Ramon et al. 2021] or feature vectors [Zhang et al. 2022a]. HeadNeRF [Hong et al. 2022] reconstruct NeRFs for human heads, which is also conditioned on the identity, expression, and appearance parameters. Gafni et al. [2021] proposes to reconstruct a dynamic facial avatar by training a NeRF that depends on pre-estimated expression parameters.

One problem of the implicit representation is that it encodes the scene into the parameters of the neural networks, which makes it extremely difficult to edit directly. To edit the implicit representation, [Liu et al. 2021] train a NeRF conditioned on latent code and then fine-tune the latent code based on the editing target. This takes a large amount of time, even for a simple editing operation. CLIP-NeRF [Wang et al. 2021a] introduces an image or text encoder that directly maps the editing target to the latent code. UV Volumes [Chen et al. 2022] adopts feature volumes and neural texture stack for editing 3D humans. Other GAN-based NeRF methods [Chan et al. 2021a,b; Gu et al. 2021; Niemeyer and Geiger 2021; Schwarz et al. 2020; Sun et al. 2022] also have the potential to be manipulated by GAN inversion [Xia et al. 2021]. However, editing the latent code could not achieve fine-grained, and out-of-domain edits [Wang et al. 2021a]. NeuTex [Xiang et al. 2021] introduces the idea of texture mapping from the mesh-based representation into the implicit representation. The appearance of the scene can be edited by manipulating the texture map. This allows fine-grained editing, but the editing is limited to the static objects, and how to edit the geometry of the scene is not explored. In this work, we propose a method that can edit both the geometry and appearance of dynamic human heads.

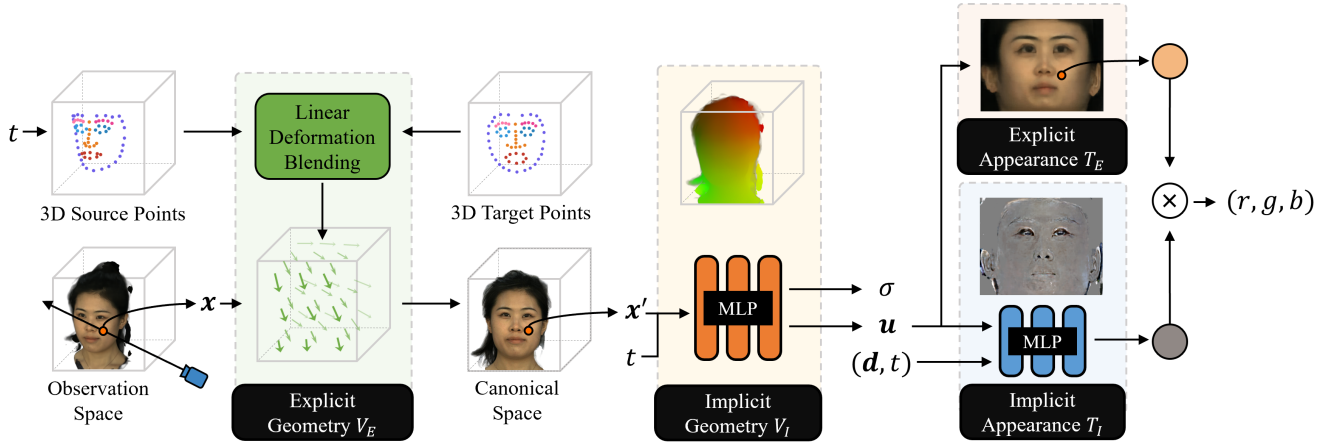


Fig. 2. Our hybrid scene representation. We disentangle a dynamic radiance field to a geometry representation and appearance representation. The geometry is composed of an explicit deformation field controlled by sparse control points, and an implicit UV and density field that maps any 3D position to a 2D texture coordinate. The appearance is made up of an explicit texture map shared by all the frames, and a MLP that models all the view and time dependent effects on the map.

### 3 OVERVIEW

We propose a hybrid implicit and explicit representation to model a dynamic 3D volume of a human head. The blueprint is that the fundamental representation is implicit, so that it can achieve photorealistic reconstruction results while maintaining a certain level of editability through the introduction of explicit layers. Our goal is to model a volume that defines a dynamic 3D radiance field  $F$  that maps the spatial 3D location  $\mathbf{x} = (x, y, z) \in \mathbb{R}^3$ , view direction  $\mathbf{d} = (\theta, \phi) \in \mathbb{R}^2$  and time index  $t \in \mathbb{Z}$  to a single RGB value  $\mathbf{c} \in \mathbb{R}^3$ , as well as the volume density  $\sigma \in \mathbb{R}$ . Formally,

$$(\mathbf{c}, \sigma) = F(\mathbf{x}, \mathbf{d}, t). \quad (1)$$

The time index  $t$  is encoded into a latent code of length 256 through a learnable embedding layer as proposed by Bojanowski et al. [2017]. We apply Fourier position encoding to  $\mathbf{x}$  and  $\mathbf{d}$  to increase the model capacity of the implicit representation as illustrated in [Tancik et al. 2020].

To be able to edit both the geometry and appearance of the human head, we disentangle the representation into the appearance representation, namely a neural texture  $T$ , as well as the geometry representation, which is essentially a UV and density field  $V$ :

$$(\mathbf{u}, \sigma) = V(\mathbf{x}, t), \mathbf{c} = T(\mathbf{u}, \mathbf{d}, t). \quad (2)$$

where  $\mathbf{u} = (u, v) \in \mathbb{R}^2$  is the UV coordinate in texture space. We further decompose the geometry representation  $V$  and the neural texture  $T$  into implicit parts  $V_I$  and  $T_I$  for implicit reconstruction and explicit modules  $V_E$  and  $T_E$  for explicit editing. Fig. 2 shows the system overview. We model the geometry using an explicit warping field cascaded with an implicit UV and density field, while the appearance is a two layered structure that is composed of an explicit texture map and a view and time dependent implicit texture.

To construct such a hybrid representation, we take synchronized multi-view videos of a person’s head as our main source of supervision. We also propose several regularizations to achieve better editability. The paper is organized as follows. In Section 4, we illustrate the data acquisition and its pre-processing. In Sections 5 and

6, we explain the geometry and appearance aspects of our representation, respectively, and how to optimize the model parameters to achieve better reconstruction quality and editability. We introduce the rendering and training details in Sections 7 and 8. Then in Sections 9 and 10, we conduct comparisons and ablations and show related applications. Finally, we discuss limitations and conclude in Section 11.

### 4 DATA ACQUISITION

The input to our system is a set of calibrated and synchronized multi-view videos of a dynamic head. We capture the videos using a camera array with 12 cameras spread over approximately 120 degrees. We use 11 views for constructing the model and 1 view for evaluation. The videos are captured at 30fps and downsampled to the resolution of  $512 \times 374$  for training. The camera poses and intrinsics are calibrated using the off-the-shelf Structure from Motion software [Epic Games 2018] using the first frame of the videos. We also compute a loose bounding box  $\mathcal{B}$  from the reconstructed point cloud as the boundary of the volume. Since our main focus is to model a dynamic head, we mask out the background of each video using the current state-of-the-art video matting algorithm [Lin et al. 2022]. Our method also requires a coarse tracked face mesh to initialize the UV mapping. To that end, we apply PRNet [Feng et al. 2018] to predict the projection of densely aligned face mesh in image space for each view. We then construct the face mesh in 3D by minimizing the reprojection error in each view. We repeat the process for each frame to get a sequence of tracked 3D face mesh. From the mesh we extract the position of each vertex  $\mathbf{p}_i \in \mathbb{R}^3$  and its corresponding UV coordinate  $\mathbf{u}_i \in \mathbb{R}^2$ .

### 5 GEOMETRY MODELING

Recall that the representation  $V$  is a composition of a UV field and a density field that maps the spatial location  $\mathbf{x}$  and time  $t$  to the UV coordinate  $\mathbf{u}$  and volume density  $\sigma$ . With  $V$  modeled merely using MLPs, it would be hard to edit the geometry. Deformation-based



Fig. 3. A visualization of the 96 control points. The control points are manually selected and have rich semantic meaning.

NeRFs [Park et al. 2021a,b; Pumarola et al. 2021] use a deformation field that maps a spatial location  $\mathbf{x}$  of one frame  $t$  in the observation space  $\mathcal{B}$  to  $\mathbf{x}'$  in canonical space  $\mathcal{B}'$ . We adopt this idea by also defining a deformation field  $V_E$  that is determined by some explicitly defined parameters. Ideally, the explicit deformation field should have the following properties for easy reconstruction and editing:

- (1) To enable volumetric rendering, the field  $V_E$  should be densely defined in the bounding box  $\mathcal{B}$  for any time index  $t$ , so that for any query point  $\mathbf{x}$ , we could find the mapped new location  $V_E(\mathbf{x}, t)$ .
- (2) The  $V_E$  should be differentiable with respect to the explicit parameters, so the parameters could be jointly optimized.
- (3) The  $V_E$  should be locally controlled by a portion of the parameters to allow for fine-grained editing.

To meet these requirements, we propose a simple but effective explicit deformation method. Specifically, since the head pose change can be modeled as a rigid transformation, we first transform the volume to a stable head space  $\mathcal{B}$  using a global rigid transformation  $\mathbf{H} \in SE(3)$ , i.e.  $\bar{\mathbf{x}} = \mathbf{H}\mathbf{x}$ . We then define the control parameters as set of sparse control points  $\{\bar{\mathbf{s}}_i\}$  and corresponding target positions  $\{\bar{\mathbf{z}}_i\}$ . The explicit deformation field  $V_E$  is then formulated as a linear deformation blending:

$$V_E(\mathbf{x}) = \bar{\mathbf{x}} + \frac{\sum_i \psi_i(\bar{\mathbf{x}})(\bar{\mathbf{z}}_i - \bar{\mathbf{s}}_i)}{\sum_i \psi_i(\bar{\mathbf{x}})}, \text{ where} \quad (3)$$

$$\psi_i(\bar{\mathbf{x}}) = \exp\left(-(\bar{\mathbf{x}} - \bar{\mathbf{s}}_i)^2 / r_i^2\right). \quad (4)$$

Intuitively,  $V_E$  is a linear interpolation of the set of displacements  $\{\bar{\mathbf{z}}_i - \bar{\mathbf{s}}_i\}$ , with the weights determined by a Gaussian radian basis function  $\psi_i(\bar{\mathbf{x}})$  that has the influence radius  $r_i$ . The bar  $\bar{\cdot}$  denotes that the control points and target positions are all defined in the stable head space  $\mathcal{B}$ . Note that  $V_E$  is differentiable with respect to the parameters  $(\mathbf{H}, \bar{\mathbf{z}}_i, \bar{\mathbf{s}}_i, r_i)$ , which means that it is possible to jointly optimize the explicit parameters using a stochastic gradient descent (SGD) algorithm. Since we are modeling a dynamic volume, we have the explicit deformation parameters  $(\mathbf{H}, \bar{\mathbf{s}}_i, r_i)$  defined in every time index. The target positions  $\{\bar{\mathbf{z}}_i\}$  are kept static for different time indexes to serve as a canonical target that is shared by all the frames. By editing the control points, we could explicitly manipulate the 3D deformation field in an intuitive way.

To edit the scene easily and consistently, we would like to make the control point semantically consistent. In other words, if each control point  $\mathbf{s}_i^{(t)}$  has the same semantic meaning for all frames  $t$ , then we could edit the entire sequence by editing only one frame, and then propagating the editing result. An example of semantically consistent points would be tracked facial landmarks. In order to construct a set of semantically consistent control points, we predefined

a subset of 96 vertices in the tracked face mesh  $\{\hat{\mathbf{s}}_i^{(t)}\}$  in frame  $t$  as our initial control points as shown in Fig. 3, and apply a semantic loss over the course of optimization:

$$\mathcal{L}_{semantic} = \sum_i \sum_t \|\mathbf{s}_i^{(t)} - \hat{\mathbf{s}}_i^{(t)}\|_2. \quad (5)$$

The reason why we only supervise the control points to be close to the vertices instead of directly using the vertices is that the tracked mesh may be inaccurate, and the deformation field constructed by the vertices may not be optimal. So we finetune the vertex locations using differentiable rendering. The global transformation  $\mathbf{H}^{(t)}$  is also initialized from a least squares estimation of all the vertices of the tracked face mesh.  $\{\mathbf{z}_i\}$  is fixed to be the same subset of vertices of a predefined canonical 3D face mesh.

After the explicit deformation field deforms each location to a canonical space, we define our final UV field as:

$$V = V_I(\mathbf{x}', t), \quad \mathbf{x}' = V_E(\mathbf{x}, t). \quad (6)$$

where the  $V_I(\mathbf{x}', t)$  is an implicit UV field that maps the canonical location  $\mathbf{x}'$  and frame index  $t$  to the UV coordinate  $\mathbf{u}$  and density  $\sigma$ . We show that the explicit deformation field not only enables editing the geometry, but also helps to align the volume to a canonical space, which eases the complexity for reconstructing the dynamic head of the implicit representation. We conditioned the  $V_I$  also on time index  $t$  since the explicit deformation  $V_E(\mathbf{x}, t)$  is a smooth and continuous deformation field that warps the geometry to the canonical space. We need  $V_I$  to model temporal varying topological changes that require a discontinuity in the deformation field such as eye blink.

## 6 APPEARANCE MODELING

Now that we have the UV coordinates, we seek to look up the RGB value given the UV coordinate, view direction  $\mathbf{d}$ , and time index  $t$ . We adopt the idea of view dependent textures and dynamic textures by allowing the appearance model  $T$  to be conditioned on  $\mathbf{d}$  and  $t$ . The view dependent textures [Lombardi et al. 2018; Xiang et al. 2021] are able to model view dependent effects such as specularly on the eyes, while compensating for inaccurate geometry and improving the realism of the rendering results. Our geometry representation  $V$  is able to model the pose and the expression change of the head, and also fit most of the skin motions, such as mouth opening and eye blinking. However, dynamic textures [Cao et al. 2021; Gotardo et al. 2018; Li et al. 2020; Yang et al. 2020a; Zhang et al. 2022b; Zheng and Xu 2021] are still needed because not all temporal variations can be modeled by a dynamic UV field, such as the sudden appearance of ambient occlusion caused by wrinkles.

The view and time dependent textures can be simply implemented using implicit representations, but this will leads to poor editability. A desired texture should be simply a single static image, so that by editing the image, every frame of the 3D volume will be consistently edited. Therefore, we use a two-layer representation that combines the explicit and implicit texture maps. The main part of the texture is stored as an explicit texture map  $T_E(\mathbf{u})$ , while the view changes and temporal variations of the texture are modulated as residuals  $T_I(\mathbf{u}, \mathbf{d}, t)$ . Thus, the appearance component is computed as:

$$T(\mathbf{u}, \mathbf{d}, t) = T_E(\mathbf{u}) * \exp(T_I(\mathbf{u}, \mathbf{d}, t)). \quad (7)$$

The lookup  $T_E$  for an arbitrary floating point UV coordinate is performed using bilinear interpolation and the  $*$  is the element-wise product that is individually applied on each RGB color channel. The use of  $\exp(\cdot)$  guarantees that the multiplier is greater than zero.  $T_I$  can be interpreted as modulating the local lighting changes of the texture map. We would like the implicit texture map to only model the residual, thus we apply a sparsity loss:

$$\mathcal{L}_{sparsity} = \sum_k |T_I(\mathbf{u}_k, \mathbf{d}_k, t)|, \quad (8)$$

where  $k$  is the index of the sampled point during rendering for a specific time  $t$ . This loss encourages the multiplier to be close to 1, so that the main contribution of the final result is from the explicit texture map  $T_E$ .

## 7 RENDERING

Given a radiance field  $F$  at frame  $t$ , we render a pixel using volume rendering [Levoy 1990]. Specifically, we first shoot rays from the camera center to the pixel locations and check whether the ray hit the volume boundary  $\mathcal{B}$  using axis-aligned bounding boxes (AABB). Then, if the ray passes the AABB test, we then sample  $N$  points along the ray between the hit near and far points. The final color is computed as:

$$\hat{\mathbf{c}} = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \text{ where } T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right) \quad (9)$$

The terms  $\sigma_i$  and  $\mathbf{c}_i$  are the density and color of  $i$ -th point, and  $\delta_i$  is the distance between adjacent samples. We adopt the same hierarchical volume sampling strategy as in [Mildenhall et al. 2020], which first uses 64 stratified sample points on a coarse volume  $V_{coarse}$ , followed by another round of importance sampling of 64 points on a fine volume  $V_{fine}$ .

## 8 TRAINING

To reconstruct the representation  $F$ , we adopt an analysis by synthesis approach and train the model using Adam [Kingma and Ba 2014] optimizer for 120k iterations, which takes about one day on three V100 GPUs. In each iteration, we randomly select a time index, and then randomly sample  $B = 6000$  rays for optimization. Next, we describe the losses and the two-stage training process that we use to train our representation.

### 8.1 Main Supervision

Our main source of supervision is the pixel reconstruction loss between the rendering result and the input images.

$$\mathcal{L}_{MSE} = \sum_i^B (\|\mathbf{c}_i - \hat{\mathbf{c}}_i\|_2 + \|\alpha_i - \hat{\alpha}_i\|_2). \quad (10)$$

Note that we also supervise the rendered foreground mask to be close to the input mask, which is computed by using the accumulated transmittance. This prevents the network from using the background to fake shading effects that should be in the texture [Xiang et al. 2021].



(a)  $\mathcal{L}_{uv}$  w/ decay (b) w/o  $\mathcal{L}_{uv}$  (c)  $\mathcal{L}_{uv}$  w/o decay

Fig. 4. An example of different configurations of  $\mathcal{L}_{uv}$ .

### 8.2 Texture Unwrap Regularization

Given only the pixel reconstruction supervision, the radiance field is able to achieve good reconstruction quality. However, we find that freely optimizing the UV field and the texture will result in a noisy UV field that leads to poor editability. Thus, we introduce several regularizations to ensure the implicit representation is optimized to be suitable for editing.

First, we would like to provide a coarse guidance of the UV field using the tracked 3D face mesh and its UV mapping, so that the model could have a big picture of where each component of the face maps to. Recall the tracked 3D face mesh has the form of a set of vertices  $\{\mathbf{p}_i\}$  and its corresponding UV coordinate  $\{\mathbf{u}_i\}$ . We apply a  $uv$  loss to the  $\mathbf{u}$  output of UV field  $V$ :

$$\mathcal{L}_{uv} = \sum_i^P \|V(\mathbf{p}_i, t) - \mathbf{u}_i\|_2, \quad (11)$$

where  $P$  is the total number of mesh vertices. Since the tracked mesh may be inaccurate, we only use it as initialization by setting the weight to exponentially decay to 0 at approximately 20,000 iterations. Fig. 4 demonstrates an example of a texture with different configurations of the  $\mathcal{L}_{uv}$ . Without  $\mathcal{L}_{uv}$ , the texture mapping is freely optimized, which fails to utilize the full texture space as shown in Fig. 4b. Without the weight decay, the inaccurate face mesh tracking leads to an incorrect mapping, and the model tries to compensate for the mapping error by generating artifacts on the texture map as shown in Fig. 4c.

Next, we would like the UV field to reasonably parameterize the surface of the head. We extend the cycle loss proposed by Xiang et al. [2021] to dynamic scenes. Specifically, we jointly optimize an inverse mapping network  $V_I^{-1}$  that maps the UV coordinate  $\mathbf{u}$  back into a 3D location:

$$\hat{\mathbf{x}}' = V_I^{-1}(\mathbf{u}, t). \quad (12)$$

Note that the output of  $V_I^{-1}$  is a 3D location in canonical space  $\mathcal{B}'$  instead of the observation space  $\mathcal{B}$ . This is because the canonical space contains fewer temporal variations, which eases the complexity of the inverse mapping network. The cycle loss is then defined as:

$$\mathcal{L}_{cycle} = \sum_i^B \|\mathbf{x}'_i - \hat{\mathbf{x}}'_i\|_2. \quad (13)$$

To reduce the computational overhead, rather than applying the cycle loss to all the sampling points, we only apply it to the points  $\mathbf{x}'_i$  that have the maximum contribution to the rendering results along each ray. Specifically, when rendering a ray in the volume,



the final color is formulated as a weighted sum of all sampled colors. We then determine the maximum contribution point by selecting the point that has the maximum weight during rendering.

The  $\mathcal{L}_{cycle}$  successfully constrains the surface of the head to be mapped to a 2D texture. However, in practice we find that the smoothness of  $(u, v)$  along the surface direction is not guaranteed using only the cycle loss. We would like the UV field to have the ability to model high frequency discontinuities, such as closing of the eyelids and lip, while also being piece-wise smooth in regions like the cheek and jaws for better editability. This discontinuity can be achieved by using the positional encoding illustrated in [Tancik et al. 2020]. To ensure piece-wise smoothness, we employ an angle preserving loss:

$$\mathcal{L}_{angle} = \sum_i^B \frac{|\nabla_{\mathbf{x}} u_{\perp} \cdot \nabla_{\mathbf{x}} v_{\perp}|}{\|\nabla_{\mathbf{x}} u_{\perp}\| \|\nabla_{\mathbf{x}} v_{\perp}\|}, \quad (14)$$

where  $\nabla_{\mathbf{x}} u_{\perp}$  and  $\nabla_{\mathbf{x}} v_{\perp}$  are the gradient of the texture coordinate  $(u, v)$ , i.e. the output of the UV field  $V$ , with respect to the input position  $\mathbf{x}$ , after being projected to the plane perpendicular to the surface normal at  $\mathbf{x}$ . The surface normal is defined as the gradient of the density  $\nabla_{\mathbf{x}} \sigma$ . This loss regularizes the UV mapping along the surface to be as conformal as possible. This term is also applied to the sampling points with maximum contribution along a ray. We will show in the experiment that this loss leads to the smoothness of the UV field.

### 8.3 Two-stage Training

The time variation of the scene can be modeled either by the dynamic UV field  $V$  or the dynamic texture  $T$ . In order to achieve consistent appearance editing, we would like the dynamic textures to be as static as possible, while leaving the majority of the time variation to be fitted using the UV field. In an ideal representation, the points of the same feature on the face for different frames should be mapped to the same texture coordinate. To improve the alignment of the dynamic texture, we propose a two-stage training procedure. In the first stage, we fix the texture  $T$  to be time independent. This is implemented by setting the embedding of the time index to be the same for all frames. Thus, the dynamic UV field tries its best to model the dynamic human head at this stage. After the training converges (at 80,000 iterations in all of our experiments), we start optimizing a dynamic texture to model the rest of the time variations that a dynamic UV field fails to model.

The final loss is the combination of all the losses introduced above:

$$\mathcal{L} = \mathcal{L}_{MSE} + \lambda_{uv} \mathcal{L}_{uv} + \lambda_{cycle} \mathcal{L}_{cycle} + \lambda_{angle} \mathcal{L}_{angle} + \lambda_{sparsity} \mathcal{L}_{sparsity} + \lambda_{semantic} \mathcal{L}_{semantic}. \quad (15)$$

where  $\lambda_{(\cdot)}$  is the corresponding weight for each loss. We set  $\lambda_{cycle} = 1$  as suggested in [Xiang et al. 2021]. We set the  $\lambda_{uv}$  to 1 and quickly decrease to 0 at 20,000 iterations as described above. In all the experiments we set  $\lambda_{angle} = \lambda_{sparsity} = \lambda_{semantic} = 0.05$ . We ablate the weighting scheme in the experiment section.

## 9 EXPERIMENTS

In this section, we perform quantitative and qualitative experiments to validate our model on the captured multi-view video dataset,

which contains 10 video sequences, 120 - 170 frames each. We implement both  $V_I$  and  $T_I$  using MLPs that have 6 layers and 256 neurons per layer, with a skip connect at the 4th layer. We first describe several metrics that we use in the experiments and then compare reconstruction and appearance editing quality with several baselines. We also conduct a series of ablations to validate the effectiveness of our model.

### 9.1 Metrics

For the reconstruction quality, we use the task of novel view synthesis, and compare the synthesized novel view and ground truth. To measure editability, we develop several metrics to compute how well the model could be used for editing. First, since we want to consistently edit the entire sequence by editing on one texture, we measure the quality of the alignment by computing the average standard deviation (*ASTD*) of each pixel color in the face region over the entire sequence of a dynamic texture. A bigger *ASTD* indicates bigger time variations of the dynamic texture, which leads to poor editability. In the experiment, we find that conformal UV mapping greatly affects the editing results. We measure the average angle preserving error (*Angle*) similar to how we compute  $\mathcal{L}_{angle}$ , but averaged over all the pixels in the test view.

### 9.2 Comparison

*Novel View Synthesis.* We first evaluate the novel view synthesis quality, which reflects the reconstruction performance of the models. We compare with two dynamic NeRF baselines, namely *HyperNeRF* [Park et al. 2021b] and *DyNeRF* [Li et al. 2021c]. We do not compare with NeRF-like human head reconstruction methods such as *NerFace* [Gafni et al. 2021] and *IM Avatar* [Zheng et al. 2021], since these methods take monocular input and exploit the head pose changes to collect multi-view information. We find that extending these methods for multi-view input is a non-trivial problem. We then compare with mesh-based methods, *Hifi3D* [Bao et al. 2020], *DFNRMVS* [Bai et al. 2020], as well as the *PRNet* [Feng et al. 2018] that our method uses as the initialization. We additionally provide results compared on the *Facescape* [Yang et al. 2020b] and *Beeler* [Beeler et al. 2011] datasets in the supplementary material. For the *DFNRMVS* method, which reconstructs facial meshes from multi-views, we find that blending multiple textures from multi-views results in blurry textures and worsens the performance. Therefore, we compute the mesh texture from a nearest neighbor view, which has roughly the same self-occlusion as the test view. For both mesh-based methods, we apply the algorithm frame by frame. We could have compared our method with other multi-view face reconstruction methods [Fyffe et al. 2017; Li et al. 2021a; Lombardi et al. 2018], but the source codes are not publicly available before the submission of this paper. We also try to compare with *Tofu* [Li et al. 2021b], but we find the reconstruction quality is poor, possibly due to the different camera, lighting, and subject settings. We report novel view synthesis results in Tab. 1. We separately report reconstruction metrics in the face region and full head. Results show that our method is slightly worse than the NeRF-based method regarding the reconstruction quality. This is expected since *HyperNeRF* and *DyNeRF* focus on reconstruction only and do not support editing,

Table 1. Reconstruction and texture temporal alignment. Our approach is slightly inferior to the NeRF-based methods due to the regularization for improving editability. Overall, our method achieves the best trade-off between reconstruction and editability. (Invalid cells are denoted as -.  $\uparrow$  means higher values are better.)

	Editable	Face region				Full head		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	ASTD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
HyperNeRF	No	31.31	0.8078	0.04270	-	29.11	0.8092	0.0979
DyNeRF	No	<b>32.61</b>	<b>0.8381</b>	0.03960	-	<b>29.53</b>	<b>0.8313</b>	0.09273
Ours	Yes	30.62	0.7998	0.03382	<b>2.713</b>	28.38	0.7964	<b>0.08892</b>
DFNRMVS	Yes	28.75	0.7778	<b>0.02633</b>	4.984	-	-	-
HiFi3D	Yes	22.46	0.5601	0.07053	5.616	-	-	-
PRNet	Yes	27.48	0.7745	0.05148	11.60	-	-	-




Fig. 5. Novel view synthesis results. Each row shows the results of one subject. Ours enables editing of the full head while achieving similar rendering results to NeRF-based methods. Some regions of ours, such as the teeth are less detailed, while other regions often outperform NeRF-based methods due to the use of the explicit deformation field.

while we apply strong regularizations to our volume to guarantee the editing quality. Our method outperforms the mesh-based methods (except for the LPIPS value in the face region) either on the reconstruction quality or the texture temporal alignment (*ASTD*). The results also show that the initial fitting (PRNet) is inaccurate and that our method is robust to the inaccurate initialization and could further improve the accuracy based on the initialization. We also show several qualitative results in Fig. 5. It can be seen that compared to mesh-based methods, our method naturally reconstructs the hair and the inside of the mouth. Our method lost some of the details in the teeth compared with the NeRF-based method. We

also find that in cases where the head moves, our method predicts sharper textures. This is probably because of the use of the explicit deformation field  $V_E$ , which aligns the head.

*Editing.* We then compare the quality of appearance editing with mesh-based methods. Although mesh-based methods focus mainly on the reconstruction, we could achieve appearance editing easily using tracked mesh by overlaying the editing on top of the input. Here we focus on the temporal consistency of the UV mapping, which reflects the temporal consistency of the editing. We first show the UV checkered pattern (UV checker) overlays on Fig. 6. For video results, please see the supplementary material. Note the

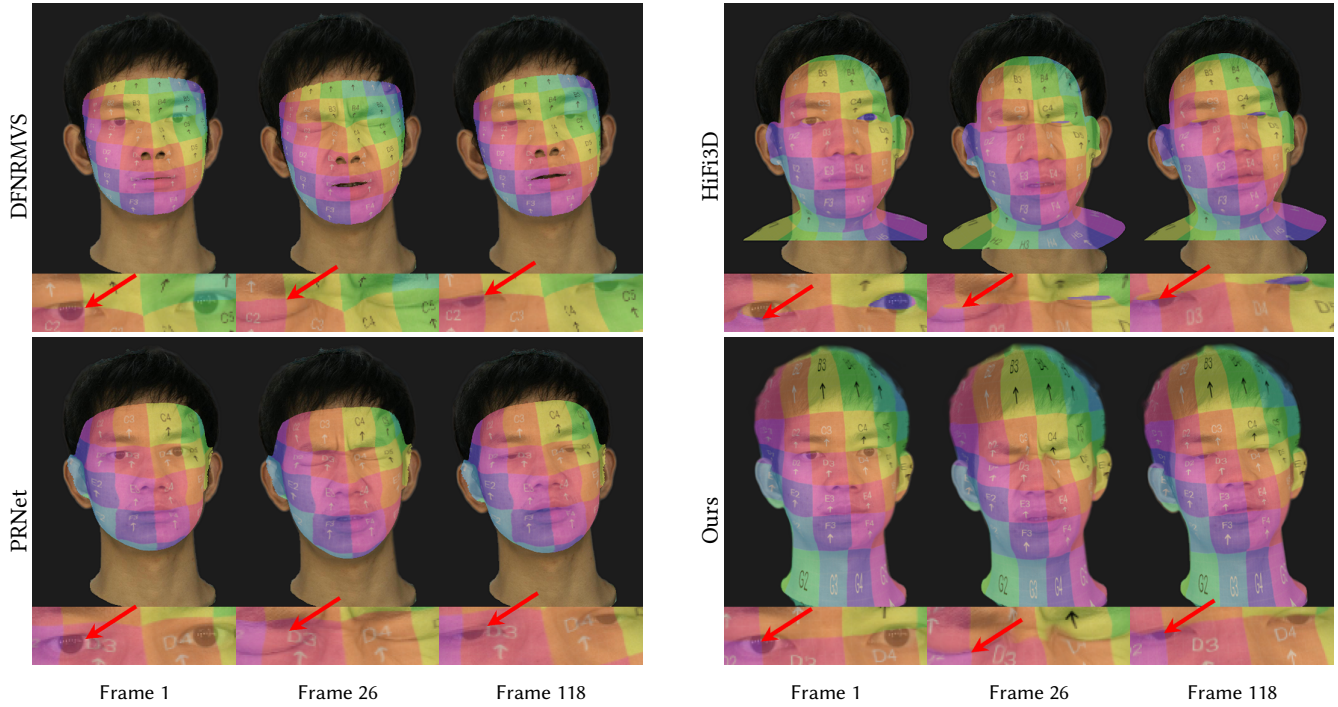


Fig. 6. Visualizations of UV checker overlay. In each column, we show a frame of the same sequence. We highlight an equivalent point on the UV checker with a red arrow. Our UV mapping achieves better temporal consistency.




Fig. 7. Visualizations of textures in different frames. In the last row, we visualize the *STD* maps of each pixel color over the entire sequence. It can be seen that ours achieves the best texture alignment.

points highlighted in the red arrow, which is the same point in the UV checker. Although mesh-based methods achieve plausible UV mapping in the face region for a single frame, temporal consistency is not guaranteed. The same point on the texture can be mapped to different locations of the face (inside the eye and on the eyelid) in the same video sequence. Besides, ours achieves mapping for the entire head, while HiFi3D fails to generate a plausible mapping in the hair, ear, and neck. We show the textures of different frames in Fig. 7. We also show the texture *STD* maps alongside the textures. Ours demonstrates smaller temporal variations of the dynamic textures, which also indicates that ours has a more consistent mapping compared to mesh-based methods.

### 9.3 Ablations

We conduct a series of ablations on the multi-view videos dataset, where we remove or modify one component at a time to validate the effectiveness of each component.

*Sparsity Loss.* We first validate the necessity of  $\mathcal{L}_{sparsity}$ . We visualize the explicit textures  $T_E$ , implicit textures  $T_I$ , the final textures  $T$  and the rendering results of different loss weights  $\lambda_{sparsity}$  in Fig. 8. Recall that we want the temporal variations to be modeled as residuals using  $T_I$  so we can edit the explicit texture  $T_E$ . Results show that a smaller  $\lambda_{sparsity}$  leaves too much information on  $T_I$  and a bigger  $\lambda_{sparsity}$  eliminates important temporal variations such as the wrinkles.

*Two-Stage Training.* Then we test the effectiveness of the proposed two-stage training strategy. We show the textures with and without two-stage training in Fig. 9. Results show that the two-stage



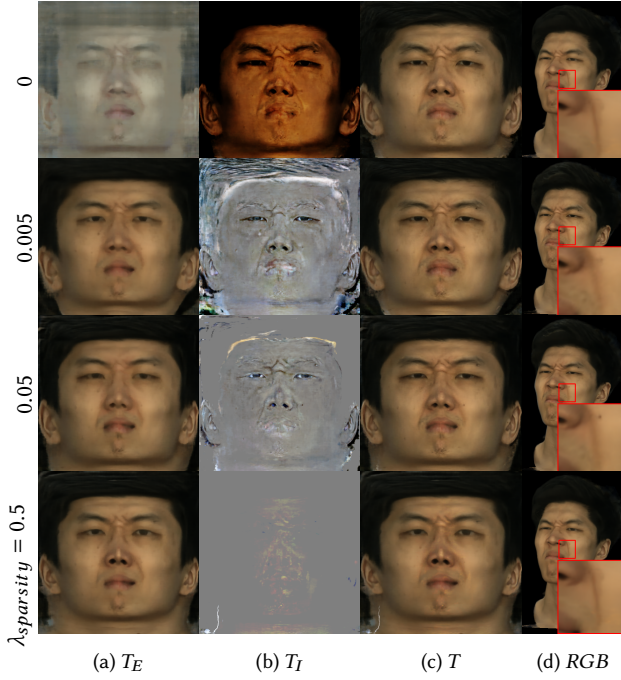


Fig. 8. Visualizations of the explicit texture  $T_E$ , the implicit texture  $T_I$ , the final texture  $T$  and the rendering results using different  $\lambda_{sparsity}$  settings. A darker color in  $T_I$  indicates a smaller value, and the gray color indicates the  $T_I = 0$ . We use  $\lambda_{sparsity} = 0.05$  for all of the experiments, which successfully decomposes the temporal variations to  $T_I$ .

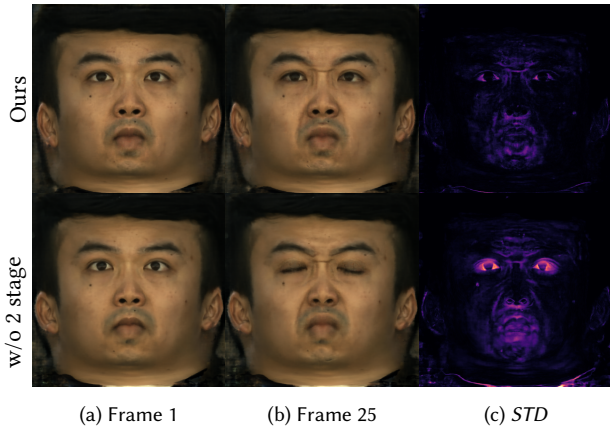


Fig. 9. Visualizations of the texture  $T$  in different frames of the same sequence and the  $STD$  maps. With the two-stage training strategy, textures tend to have better temporal alignment.

strategy reduces the temporal variations in the dynamic textures and achieves better temporal alignment, especially in the eye and mouth regions.

**Angle Loss.** We experiment with different  $\lambda_{angle}$  and visualize the rendering results, UV checker visualizations, *Angle* error maps, and editing examples in Fig. 10. We note that a small weight of  $\mathcal{L}_{angle}$  will produce better reconstruction quality since the UV mapping is less regularized. However, this will lead to a noisy UV mapping,




Fig. 10. Visualizations under different settings of  $\lambda_{angle}$ . The Angle Error Map is a color-coded map for the per-pixel angle error, the darker color indicates a smaller error. A smaller weight  $\lambda_{angle}$  results in better reconstruction quality, but leads to noisy UV mapping and poor angle and area preservation. On the other hand, larger weight leads to divergence of the model.

which is not suitable for editing since the editing results will also be noisy. It also has poor performance in terms of angle preservation. On the other hand, we find that the model does not converge to plausible results after increasing  $\lambda_{angle}$  beyond a threshold. In summary, setting  $\lambda_{angle} = 0.05$  achieves the best trade-off between reconstruction and editability. We used this setting in all of our experiments.

**Explicit Deformation Layer.** We first demonstrate the necessity of  $\mathcal{L}_{semantic}$  by training the model without this term. We show one such example in Fig. 12. It can be seen that freely optimizing the explicit deformation parameters will not converge to semantically meaningful control points, which makes the editing inconvenient.

We also find that although the design of the explicit deformation layer  $V_E$  is to enable geometry editing, it actually helps improve the reconstruction quality of the radiance field. We show one such example in Fig. 13. The reconstruction is sharper with the explicit



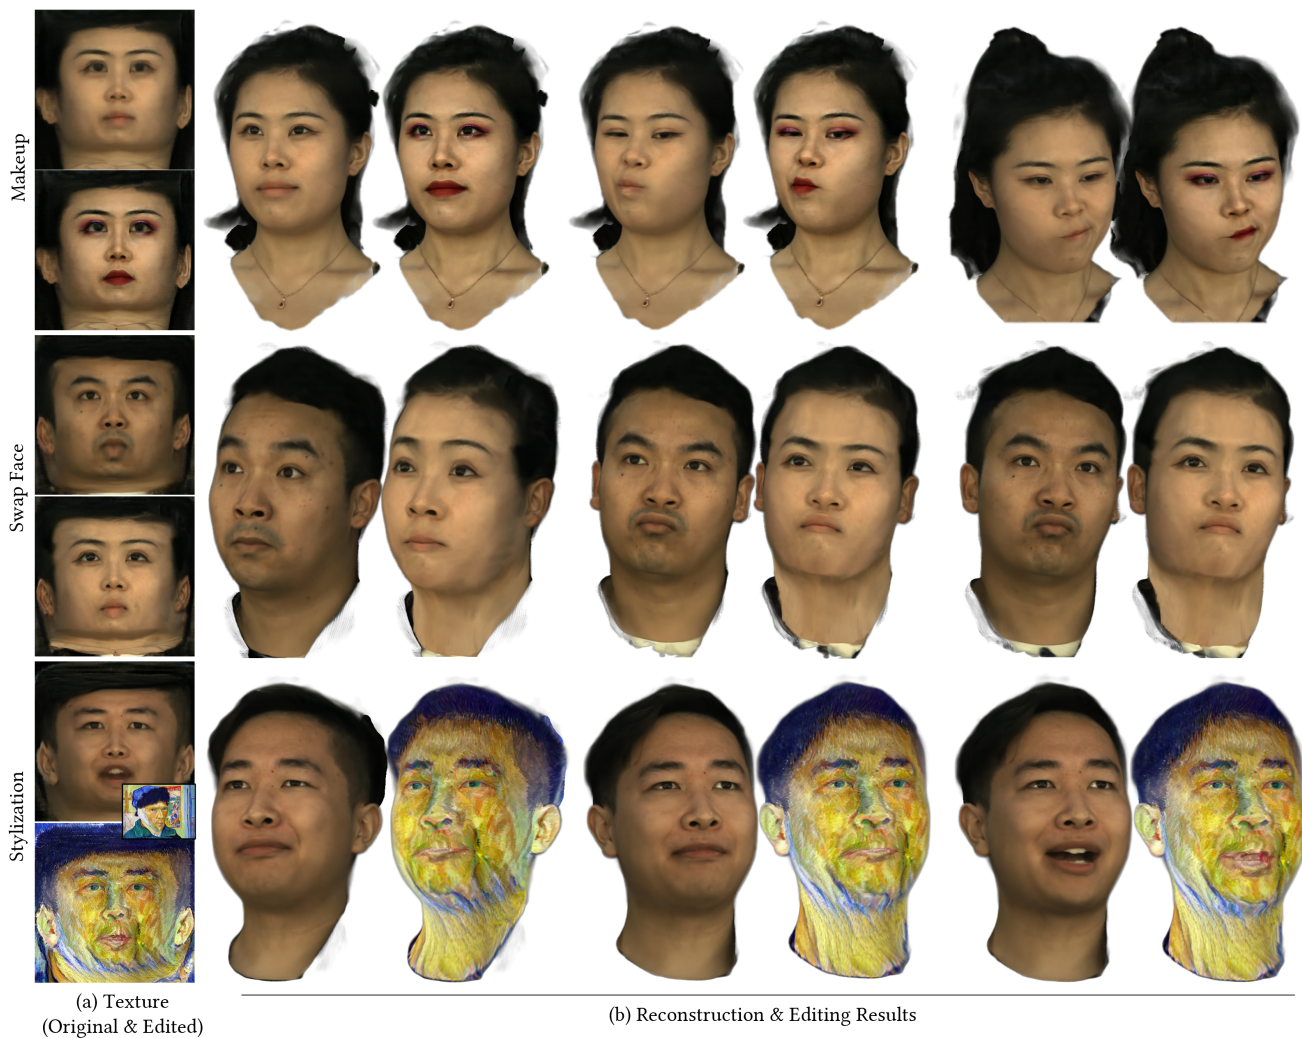



Fig. 11. Examples of appearance editing.

Fig. 12. Visualization of control points with and without  $L_{semantic}$ .

deformation layer. This is due to the fact that the explicit deformation is initialized from the mesh tracking results, which helps to make the training of the entire model easier.

## 10 APPLICATIONS

In this section, we demonstrate several applications that are made possible by using our model. These include geometry editing (Fig. 1b,


Fig. 13. Ablations with and without explicit warping module  $V_E$ . Without  $V_E$ , the reconstruction tends to be blurrier.

Fig. 14), appearance editing (Fig. 1e, Fig. 11) and joint editing (Fig. 1cf, Fig. 15). We edit the appearance by editing the explicit texture  $T_E$ . We do the face swapping by manually aligning the source texture

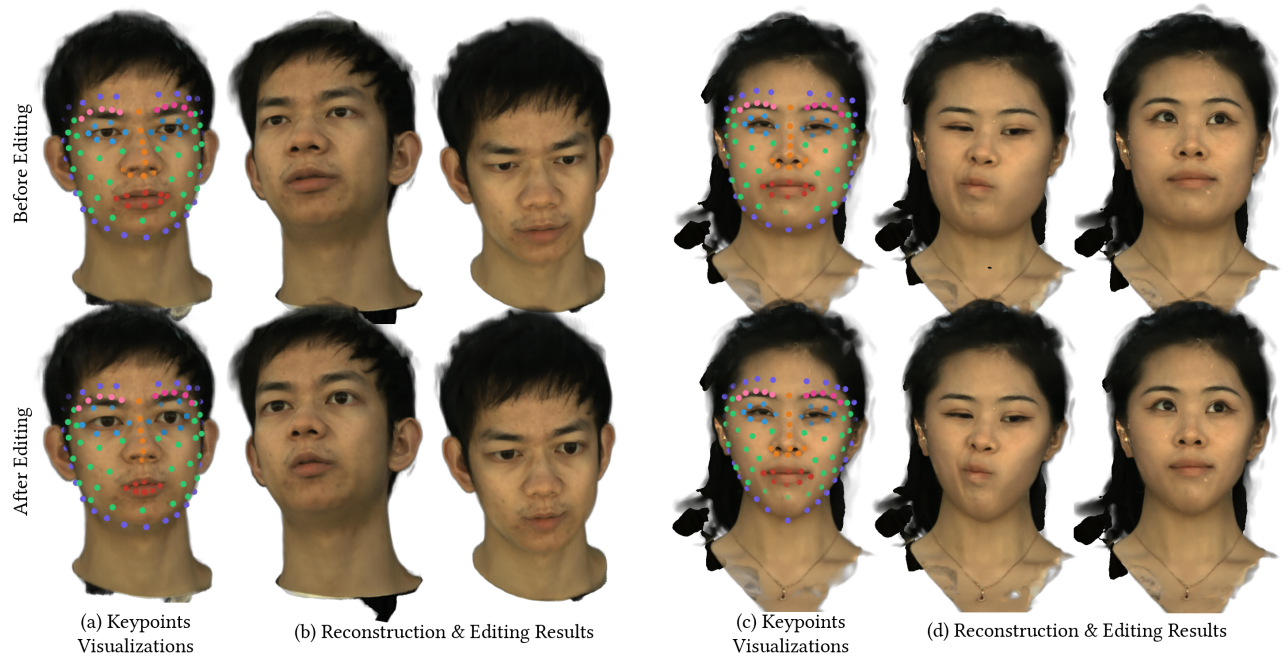


Fig. 14. Examples of geometry editing.




Fig. 15. Examples of joint appearance editing and geometry editing.



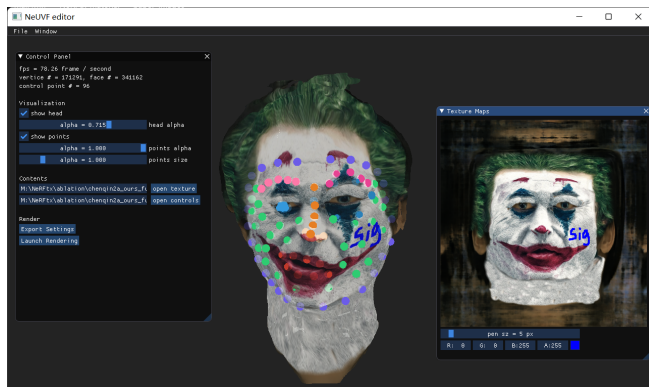


Fig. 16. A screen shot of our user interface.

map to the texture map of the editing target using photo processing software. And we use the method proposed by Kolkin et al. [2019] for stylizing the textures. For geometry editing, we manually adjust the control points in one frame using our UI, and then propagate the delta positions of each control point to the other frames. As shown in Fig. 16, we develop an interactive UI for intuitively editing the control points and the texture map. Video editing results, as well as an editing UI demo, can be found in the supplementary video. To produce a preview of the rendering results, we perform ray marching on the density field and extract a coarse mesh. We then forward our model to get the UV coordinates of each vertex, along with the texture map. In the UI, users can freely adjust the control points. We deform all the vertices of the coarse mesh using the control points based on the method described in Sec. 5. The UI is implemented in OpenGL and runs at over 70fps on a commercial laptop. After the editing, offline volume rendering is used to render the final results.

## 11 DISCUSSION AND LIMITATIONS

In this paper, we present a method for reconstructing and editing 3D dynamic human heads using multi-view videos as input. The method extends the implicit representation and improves editability by introducing explicit geometry and appearance layers. We show that our method achieves photorealistic reconstruction while enabling consistent, intuitive, and powerful editing of both appearance and geometry. Although our main focus here is the human head, it is relatively easy to extend to other objects as long as a tracked coarse mesh is available.

Our method comes with some limitations. Like most NeRF-based methods, training a model takes days, which may prevent it from practical use in some settings. We believe this could be improved with the help of existing methods for accelerating NeRF training [Müller et al. 2022; Yu et al. 2021].

Our method sometimes could not achieve the same level of reconstruction accuracy as other dynamic NeRF methods, which focus only on reconstruction, as illustrated in the experiment section. For challenging area such as mouth interior, our method tends to fit using coarse geometry and view-dependent textures. One example is shown in Fig. 17. It can be seen that the reconstructed depth map in the mouth region is over-smoothed. Furthermore, our appearance representation bakes in the lighting to the view-dependent texture.




Fig. 17. Results of a challenging mouth example. The inset visualizes the reconstructed depth map of the mouth region. Our method tends to model the mouth interior using coarse geometry and view dependent textures.




Fig. 18. An example of expression and hair editing. Extra control points are used for editing the hair.

Therefore, it is still challenging to edit the lighting and the material. An interesting future work would be to introduce a lighting model and physically based rendering techniques so that we could have more control over the shading.

To enable geometry editing, we adopt a simple yet intuitive geometric editing module driven by sparse control points. While our primary focus is to consistently edit the facial attribute of a dynamic head (e.g. size of the eye, shape of face contour), we show that it is possible to extend our method for editing one's expression. By adding more control points on the head we could perform basic hair editing. We show one example of expression and hair editing in Fig. 18. However, one limitation of our geometry editing module is that it cannot handle topology changes, such as opening the mouth. A potential solution is to develop more advanced interpolation techniques that allow modeling discontinuities while also supporting explicit editing.

## ACKNOWLEDGMENTS

The authors from HKUST were partly supported by the Hong Kong Research Grants Council (RGC). The author from CityU was supported by DON\_RMG grant from CityU, Hong Kong (Project No. 9229064).

## REFERENCES

- Ziqian Bai, Zhaopeng Cui, Jamal Ahmed Rahim, Xiaoming Liu, and Ping Tan. 2020. Deep facial non-rigid multi-view stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5850–5860.
- Linchao Bao, Xiangkai Lin, Yajing Chen, Haoxian Zhang, Sheng Wang, Xuefei Zhe, Di Kang, Haozhi Huang, Xinwei Jiang, Jue Wang, et al. 2020. High-Fidelity 3D Digital Human Head Creation from RGB-D Selfies. *arXiv preprint arXiv:2010.05562* (2020).

- Thabo Beeler, Fabian Hahn, Derek Bradley, Bernd Bickel, Paul Beardsley, Craig Gotsman, Robert W. Sumner, and Markus Gross. 2011. High-quality passive facial performance capture using anchor frames. *ACM Trans. Graph.* 30, Article 75 (August 2011), 10 pages. Issue 4. <https://doi.org/10.1145/2010324.1964970>
- Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Cláudio Silva, and Gabriel Taubin. 1999. The ball-pivoting algorithm for surface reconstruction. *IEEE transactions on visualization and computer graphics* 5, 4 (1999), 349–359.
- Sai Bi, Stephen Lombardi, Shunsuke Saito, Tomas Simon, Shih-En Wei, Keyvyn Mcphail, Ravi Ramamoorthi, Yaser Sheikh, and Jason Saragih. 2021. Deep relightable appearance models for animatable faces. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–15.
- Sai Bi, Zexiang Xu, Pratul Srinivasan, Ben Mildenhall, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Goeffroy, David Kriegman, and Ravi Ramamoorthi. 2020. Neural Reflectance Fields for Appearance Acquisition. [arXiv:2008.03824 \[cs.CV\]](https://arxiv.org/abs/2008.03824)
- 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>
- Piotr Bojanowski, Armand Joulin, David Lopez-Paz, and Arthur Szlam. 2017. Optimizing the Latent Space of Generative Networks. [arXiv:1707.05776 \[stat.ML\]](https://arxiv.org/abs/1707.05776)
- James Booth, Anastasios Roussos, Allan Ponniah, David Dunaway, and Stefanos Zafeiriou. 2018. Large scale 3d morphable models. *International Journal of Computer Vision* 126, 2 (2018), 233–254.
- Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P.A. Lensch. 2021. NeRD: Neural Reflectance Decomposition from Image Collections. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)* (Oct 2021). <https://doi.org/10.1109/iccv48922.2021.01245>
- Chen Cao, Vasu Agrawal, Fernando De La Torre, Lele Chen, Jason Saragih, Tomas Simon, and Yaser Sheikh. 2021. Real-time 3D neural facial animation from binocular video. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–17.
- Edwin Catmull and James Clark. 1978. Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-aided design* 10, 6 (1978), 350–355.
- 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. 2021a. Efficient Geometry-aware 3D Generative Adversarial Networks. [arXiv:2112.07945 \[cs.CV\]](https://arxiv.org/abs/2112.07945)
- Eric R. Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. 2021b. pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (Jun 2021). <https://doi.org/10.1109/cvpr46437.2021.00574>
- Yue Chen, Xuan Wang, Qi Zhang, Xiaoyu Li, Xingyu Chen, Yu Guo, Jue Wang, and Fei Wang. 2022. UV Volumes for Real-time Rendering of Editable Free-view Human Performance. [arXiv preprint arXiv:2203.14402](https://arxiv.org/abs/2203.14402) (2022).
- Pei-Ze Chiang, Meng-Shiun Tsai, Hung-Yu Tseng, Wei-Sheng Lai, and Wei-Chen Chiu. 2022. Stylizing 3D Scene via Implicit Representation and HyperNetwork. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 1475–1484.
- Forrester Cole, Kyle Genova, Avneesh Sud, Daniel Vlasic, and Zhoutong Zhang. 2021. Differentiable surface rendering via non-differentiable sampling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 6088–6097.
- Yu Deng, Jiaolong Yang, and Xin Tong. 2021. Deformed implicit field: Modeling 3d shapes with learned dense correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10286–10296.
- Epic Games. 2018. CapturingReality. <https://www.capturingreality.com>
- 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)* 40, 4 (2021), 1–13.
- Yao Feng, Fan Wu, Xiaohu Shao, Yanfeng Wang, and Xi Zhou. 2018. Joint 3d face reconstruction and dense alignment with position map regression network. In *Proceedings of the European conference on computer vision (ECCV)*. 534–551.
- John Flynn, Michael Broxton, Paul Debevec, Matthew DuVall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. 2019. Deepview: View synthesis with learned gradient descent. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2367–2376.
- Graham Fyffe, Koki Nagano, Loc Huynh, Shunsuke Saito, Jay Busch, Andrew Jones, Hao Li, and Paul Debevec. 2017. Multi-View Stereo on Consistent Face Topology. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 295–309.
- Guy Gafni, Justus Thies, Michael Zollhofer, and Matthias Nießner. 2021. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8649–8658.
- 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*. 1155–1164.
- Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, and William T Freeman. 2018. Unsupervised training for 3d morphable model regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8377–8386.
- Thomas Gerig, Andreas Morel-Forster, Clemens Blumer, Bernhard Egger, Marcel Luthi, Sandro Schönborn, and Thomas Vetter. 2018. Morphable face models—an open framework. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*. IEEE, 75–82.
- Paulo Gotardo, Jérémy Riviere, Derek Bradley, Abhijeet Ghosh, and Thabo Beeler. 2018. Practical dynamic facial appearance modeling and acquisition. (2018).
- Carsten Griwodz, Simone Gasparini, Lilian Calvet, Pierre Gurdjos, Fabien Castan, Benoit Maujean, Gregoire De Lillo, and Yann Lanthony. 2021. AliceVision Meshroom: An open-source 3D reconstruction pipeline. In *Proceedings of the 12th ACM Multimedia Systems Conference - MMSys '21*. ACM Press. <https://doi.org/10.1145/3458305.3478443>
- Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. 2021. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. [arXiv preprint arXiv:2110.08985](https://arxiv.org/abs/2110.08985) (2021).
- Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. 2020. Towards fast, accurate and stable 3d dense face alignment. In *European Conference on Computer Vision*. Springer, 152–168.
- Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. 2022. HeadNeRF: A Real-time NeRF-based Parametric Head Model. (2022).
- Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. 2018. DeepMVS: Learning Multi-View Stereo. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Abhishek Kar, Christian Häne, and Jitendra Malik. 2017. Learning a multi-view stereo machine. *Advances in neural information processing systems* 30 (2017).
- Yoni Kasten, Dolev Ofri, Oliver Wang, and Tali Dekel. 2021. Layered neural atlases for consistent video editing. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–12.
- Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. 2018. Neural 3d mesh renderer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3907–3916.
- Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. 2006. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, Vol. 7.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. [arXiv:1412.6980 \[cs.LG\]](https://arxiv.org/abs/1412.6980)
- Nicholas Kolkin, Jason Salavon, and Gregory Shakhnarovich. 2019. Style transfer by relaxed optimal transport and self-similarity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10051–10060.
- Marc Levoy. 1990. Efficient ray tracing of volume data. *ACM Transactions on Graphics (TOG)* 9, 3 (1990), 245–261.
- Jiaman Li, Zhengfei Kuang, Yajie Zhao, Mingming He, Karl Bladin, and Hao Li. 2020. Dynamic facial asset and rig generation from a single scan. *ACM Trans. Graph.* 39, 6 (2020), 215–1.
- Moran Li, Haibin Huang, Yi Zheng, Mengtian Li, Nong Sang, and Chongyang Ma. 2021a. Implicit Neural Deformation for Multi-View Face Reconstruction. [arXiv preprint arXiv:2112.02494](https://arxiv.org/abs/2112.02494) (2021).
- 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 Trans. Graph.* 36, 6 (2017), 194–1.
- Tianye Li, Shichen Liu, Timo Bolkart, Jiayi Liu, Hao Li, and Yajie Zhao. 2021b. Topologically Consistent Multi-View Face Inference Using Volumetric Sampling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3824–3834.
- Tianye Li, Mira Slavcheva, Michael Zollhofer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, and Zhaoyang Lv. 2021c. Neural 3d video synthesis. [arXiv preprint arXiv:2103.02597](https://arxiv.org/abs/2103.02597) (2021).
- Tzu-Mao Li, Miika Aittala, Frédéric Durand, and Jaakko Lehtinen. 2018. Differentiable monte carlo ray tracing through edge sampling. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–11.
- Shanchuan Lin, Linjie Yang, Imran Saleemi, and Soumyadip Sengupta. 2022. Robust High-Resolution Video Matting with Temporal Guidance. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 238–247.
- Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. 2019. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7708–7717.
- Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. 2021. Editing conditional radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5773–5783.
- Stephen Lombardi, Jason Saragih, Tomas Simon, and Yaser Sheikh. 2018. Deep appearance models for face rendering. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–13.
- Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. 2019. Neural volumes: Learning dynamic renderable volumes from images. [arXiv preprint arXiv:1906.07751](https://arxiv.org/abs/1906.07751) (2019).
- Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhofer, Yaser Sheikh, and Jason Saragih. 2021. Mixture of volumetric primitives for efficient neural



- rendering. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–13.
- Matthew M Loper and Michael J Black. 2014. OpenDR: An approximate differentiable renderer. In *European Conference on Computer Vision*. Springer, 154–169.
- Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. 2019. Reparameterizing discontinuous integrands for differentiable rendering. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–14.
- Fujun Luan, Shuang Zhao, Kavita Bala, and Zhao Dong. 2021. Unified shape and svbrdf recovery using differentiable monte carlo rendering. In *Computer Graphics Forum*, Vol. 40. Wiley Online Library, 101–113.
- Keyang Luo, Tao Guan, Lili Ju, Haipeng Huang, and Yawei Luo. 2019. P-mvsnet: Learning patch-wise matching confidence aggregation for multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10452–10461.
- Shugao Ma, Tomas Simon, Jason Saragih, Dawei Wang, Yuecheng Li, Fernando De La Torre, and Yaser Sheikh. 2021. Pixel codec avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 64–73.
- Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. 2019. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4460–4470.
- Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. 2019. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–14.
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2020. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*. Springer, 405–421.
- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *arXiv preprint arXiv:2201.05989* (2022).
- Michael Niemeyer and Andreas Geiger. 2021. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11453–11464.
- Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. 2020. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3504–3515.
- Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. 2021a. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5865–5874.
- Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. 2021b. HyperNeRF: A Higher-Dimensional Representation for Topologically Varying Neural Radiance Fields. *arXiv preprint arXiv:2106.13228* (2021).
- Pascal Paysan, Reinhard Kothne, Brian Amberg, Sami Romdhani, and Thomas Vetter. 2009. A 3D face model for pose and illumination invariant face recognition. In *2009 sixth IEEE international conference on advanced video and signal based surveillance*. Ieee, 296–301.
- Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. 2021. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10318–10327.
- Eduard Ramon, Gil Triginer, Janna Escur, Albert Pumarola, Jaime Garcia, Xavier Giro-i Nieto, and Francesc Moreno-Noguer. 2021. H3d-net: Few-shot high-fidelity 3d head reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5620–5629.
- Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. 2020. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501* (2020).
- Gernot Riegler and Vladlen Koltun. 2020a. Free View Synthesis. *CoRR abs/2008.05511* (2020). [arXiv:2008.05511](https://arxiv.org/abs/2008.05511) <https://arxiv.org/abs/2008.05511>
- Gernot Riegler and Vladlen Koltun. 2020b. Stable View Synthesis. *CoRR abs/2011.07233* (2020). [arXiv:2011.07233](https://arxiv.org/abs/2011.07233) <https://arxiv.org/abs/2011.07233>
- Johannes Lutz Schönberger and Jan-Michael Frahm. 2016. Structure-from-Motion Revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. 2016. Pixelwise View Selection for Unstructured Multi-View Stereo. In *European Conference on Computer Vision (ECCV)*.
- Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. 2020. Graf: Generative radiance fields for 3d-aware image synthesis. *arXiv preprint arXiv:2007.02442* (2020).
- Olga Sorkine and Marc Alexa. 2007. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, Vol. 4. 109–116.
- Olga Sorkine, Daniel Cohen-Or, Yaron Lipman, Marc Alexa, Christian Rössl, and H-P Seidel. 2004. Laplacian surface editing. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*. 175–184.
- Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. 2021. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7495–7504.
- Pratul P Srinivasan, Richard Tucker, Jonathan T Barron, Ravi Ramamoorthi, Ren Ng, and Noah Snavely. 2019. Pushing the boundaries of view extrapolation with multiplane images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 175–184.
- Jingxiang Sun, Xuan Wang, Yichun Shi, Lizhen Wang, Jue Wang, and Yebin Liu. 2022. IDE-3D: Interactive Disentangled Editing for High-Resolution 3D-aware Portrait Synthesis. *arXiv preprint arXiv:2205.15517* (2022).
- Jingxiang Sun, Xuan Wang, Yong Zhang, Xiaoyu Li, Qi Zhang, Yebin Liu, and Jue Wang. 2021. FENeRF: Face Editing in Neural Radiance Fields. *arXiv preprint arXiv:2111.15490* (2021).
- Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. 2020. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems* 33 (2020), 7537–7547.
- Justus Thies, Michael Zollhöfer, and Matthias Nießner. 2019. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–12.
- Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. 2021a. CLIP-NeRF: Text-and-Image Driven Manipulation of Neural Radiance Fields. *arXiv preprint arXiv:2112.05139* (2021).
- Lizhen Wang, Zhiyuan Chen, Tao Yu, Chenguang Ma, Liang Li, and Yebin Liu. 2022. FaceVerse: a Fine-grained and Detail-controllable 3D Face Morphable Model from a Hybrid Dataset. *arXiv preprint arXiv:2203.14057* (2022).
- Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. 2021b. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689* (2021).
- Xingkui Wei, Yinda Zhang, Zhuwen Li, Yanwei Fu, and Xiangyang Xue. 2020. DeepSfm: Structure from motion via deep bundle adjustment. In *European conference on computer vision*. Springer, 230–247.
- Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. 2021. GAN Inversion: A Survey. [arXiv:2101.05278 \[cs.CV\]](https://arxiv.org/abs/2101.05278)
- Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. 2021. Space-time neural irradiance fields for free-viewpoint video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9421–9431.
- Fanbo Xiang, Zexiang Xu, Milos Hasan, Yannick Hold-Geoffroy, Kalyan Sunkavalli, and Hao Su. 2021. Neutex: Neural texture mapping for volumetric neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7119–7128.
- Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. 2020a. Facescape: a large-scale high quality 3d face dataset and detailed riggable 3d face prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 601–610.
- Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. 2020b. FaceScape: a Large-scale High Quality 3D Face Dataset and Detailed Riggable 3D Face Prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. 2021. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems* 34 (2021).
- Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. 2020. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems* 33 (2020), 2492–2502.
- Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. 2021. Plenoxels: Radiance Fields without Neural Networks. *arXiv preprint arXiv:2112.05131* (2021).
- Jingbo Zhang, Xiaoyu Li, Ziyu Wan, Can Wang, and Jing Liao. 2022a. FDNeRF: Few-shot Dynamic Neural Radiance Fields for Face Reconstruction and Expression Editing. *arXiv preprint arXiv:2208.05751* (2022).
- Longwen Zhang, Chuxiao Zeng, Qixuan Zhang, Hongyang Lin, Ruixiang Cao, Wei Yang, Lan Xu, and Jingyi Yu. 2022b. Video-driven Neural Physically-based Facial Asset for Production. *arXiv preprint arXiv:2202.05592* (2022).
- Chengwei Zheng and Feng Xu. 2021. DTFusion: Dynamic Texture Fusion using a Consumer RGBD Sensor. *IEEE Transactions on Visualization and Computer Graphics* (2021).
- Mingwu Zheng, Hongyu Yang, Di Huang, and Liming Chen. 2022. ImFace: A Nonlinear 3D Morphable Face Model with Implicit Neural Representations. *arXiv preprint arXiv:2203.14510* (2022).
- Yufeng Zheng, Victoria Fernández Abrevaya, Xu Chen, Marcel C Bühler, Michael J Black, and Otmar Hilliges. 2021. IM Avatar: Implicit Morphable Head Avatars from Videos. *arXiv preprint arXiv:2112.07471* (2021).
- Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. 2018. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817* (2018).