

Controllable Dynamic Appearance for Neural 3D Portraits

ShahRukh Athar
Stony Brook University
sathar@cs.stonybrook.edu

Zhixin Shu
Adobe Research
zshu@adobe.com

Zexiang Xu
Adobe Research
zexu@adobe.com

Fujun Luan
Adobe Research
fluan@adobe.com

Sai Bi
Adobe Research
sbi@adobe.com

Kalyan Sunkavalli
Adobe Research
sunkaval@adobe.com

Dimitris Samaras
Stony Brook University
samaras@cs.stonybrook.edu

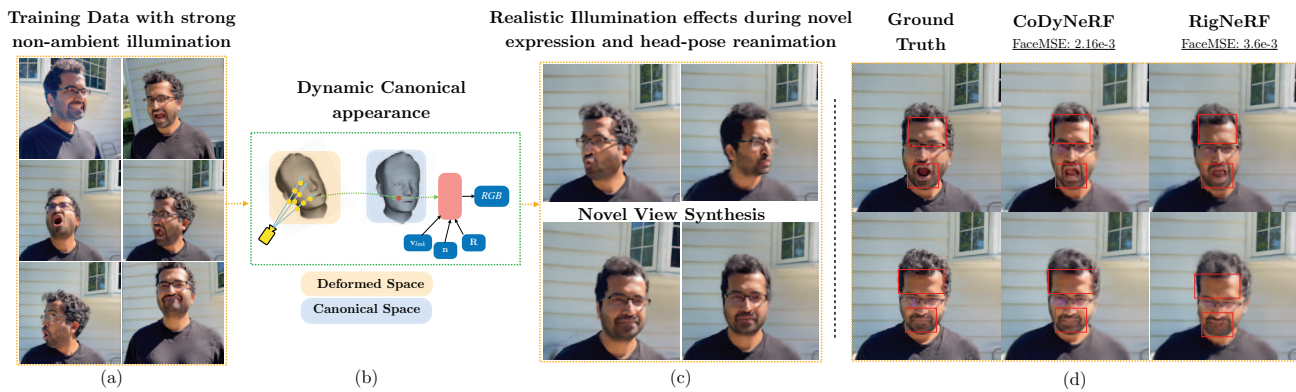


Figure 1. Our method, CoDyNeRF, enables the creation of reanimatable Neural 3D portraits from a video captured under challenging lighting conditions with a smartphone (column (a)). CoDyNeRF uses a novel dynamic canonical appearance model (column (b)), discussed in Sect 3.2) to predict illumination effects during reanimation and novel view synthesis (column (c)). Compared to prior work [3], CoDyNeRF’s prediction of illumination effects is more accurate and its renders of articulated facial expression and head-pose are of higher quality (column (d)).

Abstract

Recent advances in Neural Radiance Fields (NeRFs) have made it possible to reconstruct and reanimate dynamic portrait scenes with control over head-pose, facial expressions and viewing direction. However, training such models assumes photometric consistency over the deformed region e.g. the face must be evenly lit as it deforms with changing head-pose and facial expression. Such photometric consistency across frames of a video is hard to maintain, even in studio environments, thus making the created reanimatable neural portraits prone to artefacts during reanimation. In this work, we propose CoDyNeRF, a system

that enables the creation of fully controllable 3D portraits in real-world capture conditions. CoDyNeRF learns to approximate illumination dependent effects via a dynamic appearance model in the canonical space that is conditioned on predicted surface normals and the facial expressions and head-pose deformations. The surface normals prediction is guided using 3DMM normals that act as a coarse prior for the normals of the human head, where direct prediction of normals is hard due to rigid and non-rigid deformations induced by head-pose and facial expression changes. Using only a smartphone-captured short video of a subject for training, we demonstrate the effectiveness of our method on free view synthesis of a portrait scene with explicit head

pose and expression controls, and realistic lighting effects.

1. Introduction

The creation of photo-realistic human portraits with explicit control of head-pose and facial expressions remains a topic of active research in the computer graphics and computer vision communities. Fully controllable Neural 3D portraits are paramount in AR/VR applications where an immersive 3D experience is important. Recent advances in neural rendering and novel view synthesis [4, 5, 10, 11, 26, 27, 29, 30, 32, 35, 45, 47, 48] have demonstrated impressive image-based rendering of complex scenes and objects. Recently, these methods have also been extended to model and reanimate the human head as shown in [3, 13, 49]. These works typically employ a learnable deformation to map the deforming head to a canonical space, where the texture and geometry is predicted and rendered. The canonical space represents a mostly static appearance of the face, akin to a UV texture map, with an optional dependence on expressions [3, 13] in order to capture texture changes induced by them. The deformation, geometry, and texture are learnt via back-propagation using the photometric error with respect to the ground truth. In order for such a setup to successfully learn the appearance, deformation, and geometry, the training data must be photometrically consistent. More specifically, the color of a particular position on the human head must remain constant once mapped to the canonical space, regardless of the articulated head-pose and facial expression. However, in realistic lighting conditions, this is rarely the case.

In the real world, there is self-shadowing of the face, and the head casts its shadow on other parts of the scene as it rotates and facial expressions change. Similarly, spatially varying skin reflectance induces specularities that change with viewing angles, head poses, and facial expressions. The assumption of a static appearance in the canonical space is no longer true. As a result, deformable NeRFs trained on such data with a *static canonical appearance* assumption suffer from registration error that leads to blurriness in the renderings and inaccurate reproduction of specularities, shading and shadows. Further, if these models use a canonical space dependent on expression parameters, such as [3, 13], the aforementioned illumination-dependent effects become entangled with them. This entanglement leads to artefacts in articulated expression and inaccurate reproduction of the illumination effects during reanimation.

Accurately reproducing illumination-dependent effects due to strongly non-ambient lighting requires a modification of the canonical space. It cannot be static, it must be dynamic i.e. the canonical space must vary in appearance as illumination effects change in the deformed space. More specifically, in a dynamic portrait scene that is

captured under constant but unknown non-ambient lighting, the following lighting effects must be reproduced in the canonical space 1) Specularities, which depend on viewing directions and surface normals 2) Shading, that is dependent on head-pose and facial expression deformations as they determine the relative orientation of the surface normals to the lighting and 3) Cast shadows, which are dependent on whether or not a strong light source is occluded w.r.t a point. Capturing each of these effects, especially the cast shadows, using a physically based model is untenable in a deforming NeRF framework due to the exponentially increasing MLP evaluations that need to be performed. Our key insight is that, given enough training data, the implicit and explicit dependencies of the aforementioned illumination effects on the surface geometry, the head-pose, and the facial expression deformations can be approximated by an appropriately conditioned MLP. Based on this, we present CoDyNeRF, a method that uses a dynamic canonical appearance space, modelled by an MLP, to enable the creation of reanimatable and photorealistic neural 3D portraits using data captured in realistic lighting conditions. CoDyNeRF predicts illumination dependent effects directly in the canonical space by conditioning an MLP on the dynamic surface normals and 3DMM keypoints as well as expression and pose deformations. This conditioning makes it easier for the MLP to interpolate illumination dependent effects for novel facial expressions, head-poses and views without sacrificing the quality of facial expression and head-pose articulation.

However, a challenge remains. While 3DMM keypoints, facial expression and head-pose deformations are given by the 3DMM, the surface normals are not. Due to the lack of available ground truth geometry, it is hard to estimate accurate surface normals for each point in the scene, especially on the head and face, which are dynamic and undergo strong deformations with changing facial expression and head-pose. One possible solution is to use the normals given by the density field of the NeRF. However, as shown in Fig 4 and observed in prior work [43], these are often noisy and inaccurate. Instead, CoDyNeRF uses a carefully designed MLP, that is able to leverage both 3DMM and scene normal priors, to predict the surface normals. The MLP-predicted normals capture more detail than 3DMM-based normals and, due to the 3DMM-prior, transform correctly as the head undergoes rigid and non-rigid deformations. These normals are then used to supervise the gradient density normals of the NeRF in order to ensure accurate reconstruction of the dynamic head geometry and, consequently, the accurate reproduction of illumination effects. Once trained, CoDyNeRF realistically reproduces shadowing, shading and specular effects during reanimation with explicit control of head pose, facial expression and camera viewpoint.

In summary, our contributions in this paper are as follows:

- 1) Using a dynamic canonical appearance, we are capable of creating a fully reanimatable 3D neural portrait from data captured in strong, non-ambient lighting conditions.
- 2) We propose a method to predict accurate and detailed surface normals of the deforming human head, in addition to the static scene, which is critical for the dynamic appearance learning.
- 3) We enable a realistic re-animation of lighting and specular effects on the human face with changing head-pose and facial expressions.

2. Related works

CoDyNeRF enables the realistic rendering of illumination effects for controllable neural portraits captured in challenging lighting conditions. It is closely related to recent work on neural rendering, novel view synthesis, 3D face modeling, and controllable face generation.

Neural Scene Representations and Novel View Synthesis.

CoDyNeRF is related to recent work in neural rendering and novel view synthesis [3–5, 10, 11, 13, 23, 26–33, 35, 36, 40, 43–49]. Neural Radiance Fields (NeRF) use a Multi-Layer Perceptron (MLP), F , to learn a volumetric representation of a scene. For every 3D point on a ray, NeRFs predict an associated color and density which is volume rendered to give the final color. While NeRFs are able to reproduce specular effects, they do so at the cost of geometric fidelity of the scene [43]. RefNeRF [43] extends NeRFs to explicitly handle specularities of static scenes by improving the learnt surface normals, and consequently the scene geometry. While NeRFs such as RefNeRF are able to generate photo-realistic images for novel view synthesis, it is only designed for a static scene and is unable to represent scene dynamics. Our approach ensures the accurate reproduction of illumination effects when reanimating neural portraits since, it is specifically designed to model illumination effects of a dynamic scene.

Dynamic Neural Scene Representations. There has been an active effort to extend NeRFs to dynamic scenes. Most works do so by imposing temporal constraints either explicitly using scene flow [25, 26, 35, 45] or implicitly by using a canonical frame [32, 35]. Authors of [33] build upon [32] and model topologically changing deformations by lifting the canonical space to high dimensions. The deformation fields in these approaches are conditioned on learnt latent codes without specific physical or semantic meaning, and therefore not controllable in an intuitive manner. Further, such models only work with limited illumination change across frames, since learning a common registration becomes significantly difficult without any photometric consistency.

Controllable Face Generation. Generative Adversarial Networks (GANs) [12, 16–19, 51] have enabled high-quality image generation and have inspired a large collection of

work [2, 6–8, 22, 34, 38, 39, 41, 42] capable of face manipulation. However, it is challenging to enable high-quality view synthesis and 3D controls of the portraits as most of these works lack any 3D understanding and are purely image-based. Works such as [1, 9, 20, 21] fix this by using an intermediate 3D face representation, via a 3D Morphable Model, to reanimate face images/videos. While head poses and facial expressions are modelled with good detail in these models, thanks to the 3DMM, they are often unable to perform novel view synthesis as they focus on face region but neglect the geometry or appearance of the scene.

In a similar vein, NerFACE [10], uses NeRF to model a 4D face avatar and allows pose/expression control on the head. However, no view synthesis can be performed on the scene and the subject is assumed to be uniformly lit throughout the capture process. Neural Head Avatars [13], IMAvatar [49] and RigNeRF [3] improve upon the results of NerFACE further by using a 3DMM prior more explicitly. Neural Head Avatars learns a per-vertex pose and expression conditioned deformation for the FLAME mesh along with a detailed texture while IMAvatar [49] learns a per-point FLAME basis, used for registering different head-poses and facial expressions. RigNeRF uses the 3DMM deformation as a prior on its deformation field that maps points from the articulated space to the canonical space. Unlike our method, all three aforementioned methods require an ambiently lit face throughout the capture process and are unable to render expression and head-pose dependent illumination effects.

3. CoDyNeRF

In this section, we describe our method, CoDyNeRF, that enables the creation of reanimatable 3D portrait scenes from videos captured in the real-world with non-ambient lighting conditions. A deformable Neural Radiance Field (NeRF) [30], with a per-point 3DMM guided deformation field models facial expression and head-pose. It maps points from the deformed space to a dynamic canonical space of the model where the volume density and the appearance is predicted. The dynamic canonical space is conditioned on the surface normals, head-pose and facial expression deformations along with other shading and shadowing based cues (Sec. 3.2). The surface normals, defined in the scene world co-ordinates, are dynamic and vary with head pose and facial expression. CoDyNeRF predicts these normals using an MLP that is trained with 3DMM normals and scene normals as a prior (Sec. 3.2.1). An overview of the full architecture is shown in Fig 2. Once trained, CoDyNeRF is not only able to control facial expression and head pose of the subject but is also able to faithfully capture the varying illumination effects, such as specularities and shadows.

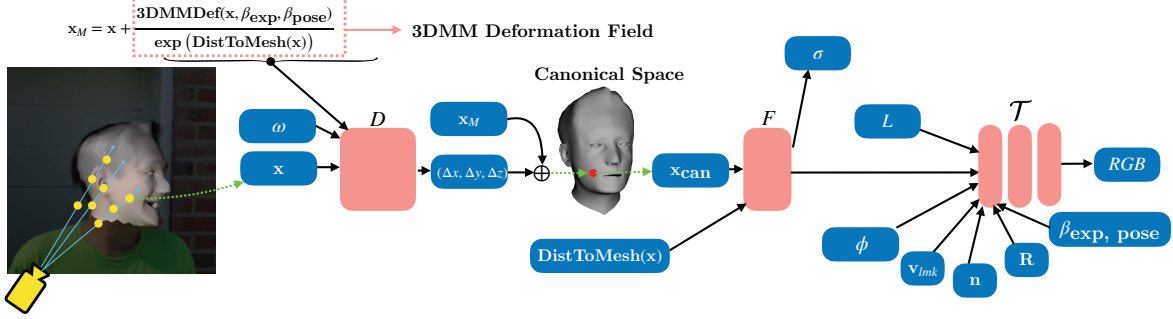


Figure 2. Overview of CoDyNeRF. CoDyNeRF is a deformable NeRF architecture that consists of three learnable MLPs: a deformation MLP D and density MLP F and a dynamic appearance MLP \mathcal{T} . Given an image, we shoot rays through each of its pixels. For every ray, we deform each point on it according to a 3DMM-guided deformation field similar to prior work [3]. Next, the deformed point is given as input to the color MLP, F , which predicts the density and neural features that are passed onto the dynamic appearance MLP \mathcal{T} . \mathcal{T} then takes as input normals, the reflection vector about the normal, the pose and expression deformations along with spherical harmonics shading and head landmark positions to predict dynamic RGB of the point. The final color of the pixel is calculated via volume rendering.

3.1. A 3DMM-guided Deformable Neural Radiance Field

A neural radiance field (NeRF) is defined as a continuous function $F : (\gamma_m(\mathbf{x}(t_i)), \gamma_n(\mathbf{d})) \rightarrow (c(\mathbf{x}(t_i), \mathbf{d}), \sigma(\mathbf{x}(t_i)))$, that, given the position of a point in the scene $\mathbf{x}(t_i) = \mathbf{o} + t_i \mathbf{d}$ that lies on a ray originating at \mathbf{o} with direction \mathbf{d} , outputs the color $c = (r, g, b)$ and the density σ . F is usually represented as a multi-layer perceptron (MLP) and $\gamma_m : \mathbb{R}^3 \rightarrow \mathbb{R}^{3+6m}$ is the positional encoding [30] defined as $\gamma_m(\mathbf{x}) = (\mathbf{x}, \dots, \sin(2^k \mathbf{x}(t_i)), \cos(2^k \mathbf{x}(t_i)), \dots)$ where m is the total number of frequency bands and $k \in \{0, \dots, m-1\}$. The expected color of the pixel through which a camera ray passes is calculated via volume rendering as follows:

$$C = \sum_t \omega_t c(x(t));$$

$$\text{where } \omega_i = \exp(-\sum_{j<i} \sigma_j (t_{j+1} - t_j)) (1 - \exp(-\sigma_i (t_{i+1} - t_i))) \quad (1)$$

The parameters of F are trained to minimize the L2 distance between the expected color and the ground-truth.

NeRFs can be extended to model dynamic scenes by using a deformation field to map each 3D point of the scene to a canonical space, where the volumetric rendering takes place [3, 32, 33, 35]. The deformation field is also represented by an MLP $D_i : \mathbf{x} \rightarrow \mathbf{x}_{\text{can}}$ where D_i is defined as $D(\mathbf{x}, \omega_i) = \mathbf{x}_{\text{can}}$ and ω_i is a per-frame latent deformation code. Following prior work [3], we use a 3DMM prior on the deformation field as follows:

$$\hat{D}(\mathbf{x}) = 3\text{DMMDef}(\mathbf{x}, \beta_{i,\text{exp}}, \beta_{i,\text{pose}})$$

$$+ D(\gamma_a(\mathbf{x}), \gamma_b(3\text{DMMDef}(\mathbf{x}, \beta_{i,\text{exp}}, \beta_{i,\text{pose}})), \omega_i) \quad (2)$$

$$\mathbf{x}_{\text{can}} = \mathbf{x} + \hat{D}(\mathbf{x})$$

where, $3\text{DMMDef}(\mathbf{x}, \beta_{i,\text{exp}}, \beta_{i,\text{pose}})$ is the deformation prior given by the 3DMM, $\beta_{i,\text{exp}}, \beta_{i,\text{pose}}$ are the articulated facial

expression and head-pose of the frame i , and γ_a, γ_b are the positional encoding functions with frequencies a and b respectively. The deformation prior, can be written as follows:

$$3\text{DMMDef}(\mathbf{x}, \beta_{i,\text{exp}}, \beta_{i,\text{pose}}) = \frac{3\text{DMMDef}(\hat{\mathbf{x}}, \beta_{i,\text{exp}}, \beta_{i,\text{pose}})}{\exp(\text{DistToMesh}(\mathbf{x}))} \quad (3)$$

where, $\hat{\mathbf{x}}$ is the closest point on the mesh to \mathbf{x} , $\text{DistToMesh}(\mathbf{x}) = \|\mathbf{x} - \hat{\mathbf{x}}\|$ is the distance between \mathbf{x} and $\hat{\mathbf{x}}$ and $3\text{DMMDef}(\hat{\mathbf{x}}, \beta_{i,\text{exp}}, \beta_{i,\text{pose}})$ is the deformation of the vertex $\hat{\mathbf{x}}$ as follows:

$$3\text{DMMDef}(\hat{\mathbf{x}}, \beta_{i,\text{exp}}, \beta_{i,\text{pose}}) = \hat{\mathbf{x}}_{\text{FLAME}(\beta_{i,\text{exp}}, \beta_{i,\text{pose}})} - \hat{\mathbf{x}}_{\text{FLAME}(\beta_{i,\text{exp}}, \beta_{i,\text{pose}})} \quad (4)$$

where, $\hat{\mathbf{x}}_{\text{FLAME}(\beta_{i,\text{exp}}, \beta_{i,\text{pose}})}$ is the position of \mathbf{x} in the canonical space and $\hat{\mathbf{x}}_{\text{FLAME}(\beta_{i,\text{exp}}, \beta_{i,\text{pose}})}$ is its position with head pose and facial expression parameters $\{\beta_{i,\text{exp}}, \beta_{i,\text{pose}}\}$.

3.2. An Illumination aware dynamic canonical appearance model

In the canonical space, CoDyNeRF predicts the density and a dynamic RGB appearance. The dynamic RGB is conditioned on the surface normals, head-pose, and expression deformations along with other shading and shadowing cues such, as the reflection vector and global location of the head. Below, we describe each aspect of the appearance model.

A spatially conditioned density prediction model. First, CoDyNeRF predicts the density at any point by conditioning on its position in the canonical space and its distance to the mesh:

$$\sigma(\mathbf{x}), \tau = F(\gamma_c(\mathbf{x}_{\text{can}}), \text{DistToMesh}(\mathbf{x})) \quad (5)$$

where, F is an MLP, τ is a feature vector, $\text{DistToMesh}(\mathbf{x}) = \|\mathbf{x} - \hat{\mathbf{x}}\|$ is the distance of \mathbf{x} to the closest mesh vertex $\hat{\mathbf{x}}$ and γ_c is the positional encoding function with c frequencies.

Additional conditioning on $\text{DistToMesh}(\mathbf{x})$ modestly boosts both the training speed and quality of results by allowing F to distinguish between points in the canonical space that have never been deformed and points that have been deformed to the canonical space. We provide experiments supporting this in the supplementary section.

An illumination aware dynamic canonical appearance model. Next, CoDyNeRF predicts dynamic RGB conditioned on inputs that capture local geometry, surface properties and viewing direction. Since the captured neural portrait is a dynamic scene, the outgoing radiance at any point \mathbf{x} , is implicitly dependent on facial expression and head-pose, $\{\beta_{\text{exp, pose}}\}$ (or $\{\beta_{\text{e,p}}\}$) due surface properties and incoming radiance being dependent on them. More specifically, at any point \mathbf{x} for a particular articulation of facial expression and head-pose, $\beta_{\text{e,p}}$, the outgoing radiance is given by the rendering equation as follows:

$$L_r(\mathbf{x}, \omega_o, \beta_{\text{e,p}}) = \int_{\omega_i} \rho(\mathbf{x}, \omega_i, \omega_o, \beta_{\text{e,p}}) (\mathbf{n} \cdot \omega_i) L_i(\mathbf{x}, \omega_i, \beta_{\text{e,p}}) d\omega_i \quad (6)$$

where, ρ is the articulation dependent BRDF, \mathbf{n} is the normal at \mathbf{x} and ω_i, ω_o are the incoming and outgoing ray directions respectively. We approximate this integral in the canonical space by an MLP, \mathcal{T} , as follows:

$$(R, G, B) = \mathcal{T}(\tau, \mathbf{n}, \mathbf{R}, \mathbf{v}_{lmk}, 3\text{DMMDef}_{\text{exp}}, 3\text{DMMDef}_{\text{pose}}, \phi_i) \quad (7)$$

where, τ are features from the density prediction network from Eq. (5), \mathbf{n} is the surface normal, \mathbf{v}_{lmk} are the facial landmarks, $\mathbf{R} = 2(\mathbf{d} \cdot \mathbf{n})\mathbf{n} - \mathbf{d}$ is the reflection vector, $3\text{DMMDef}_{\text{exp}} := 3\text{DMMDef}(\mathbf{x}, \beta_{\text{exp}}, \beta_{\text{pose, can}})$ is the expression-only deformation given by the 3DMM, $3\text{DMMDef}_{\text{pose}} := 3\text{DMMDef}(\mathbf{x}, \beta_{\text{exp, can}}, \beta_{\text{pose}})$ is the head-pose only deformation given by the 3DMM and ϕ_i is a per-frame latent vector that is learnt through optimization. Each input in Eq. (7) contains information that is essential to the prediction of accurate illumination effects. First, surface reflectance and absorption properties are captured by τ which is predicted in the canonical space and thus is forced to only model deformation-independent properties of the surface. The surface normal \mathbf{n} is used to model shading effects and, along with the reflection vector \mathbf{R} , specular effects. The face landmarks, \mathbf{v}_{lmk} , along with expression and head-pose deformations, $3\text{DMMDef}_{\text{exp}}$ and $3\text{DMMDef}_{\text{pose}}$ are used to model cast shadows, inter-reflections and any other illumination effects that depend on the global orientation of the head and deformations due to facial expressions and head-pose. The latent-code captures any appearance changes due to the camera.

3.2.1 Prediction of Dynamic Surface Normals

Prediction of shading and specular effects (through reflection vector \mathbf{R}) requires accurate surface normals. Within a NeRF,




Figure 3. Normals Prediction Architecture. The Normals prediction network takes as input the mesh normals of a given point \mathbf{x} , its distance to the mesh and the normals given by gradient of the NeRF’s density field. The details of the architecture are discussed in Section 3.2.1.

one straightforward way to calculate the normals at any point \mathbf{x} is to define it as the negative of the derivative of the density field, $\sigma(\mathbf{x})$, w.r.t \mathbf{x} i.e $\text{Grad}_{\mathbf{n}}(\mathbf{x}) = -\frac{\nabla_{\mathbf{x}}\sigma(\mathbf{x})}{|\nabla_{\mathbf{x}}\sigma(\mathbf{x})|}$. However, as shown in Fig 4 and observed in prior work [43], unless these normals are regularized [43], they are incredibly noisy and rather unusable. To get around this, we use an MLP, \mathcal{N} , to predict the normals. We design \mathcal{N} in a manner that allows it to exploit local priors from the 3D head mesh and the scene to predict the normals at each point of the dynamic neural portrait. More specifically, the normal at any point is given as follows:

$$\mathbf{n} = \mathcal{N}(\text{Mesh}_{\mathbf{n}}(\mathbf{x}), \text{Grad}_{\mathbf{n}}(\mathbf{x}), \text{DistToMesh}(\mathbf{x})) \quad (8)$$

where, $\text{Mesh}_{\mathbf{n}}(\mathbf{x})$ is normal vector of the mesh vertex closest to \mathbf{x} , $\text{Grad}_{\mathbf{n}}(\mathbf{x})$ is the normal calculated by the negative gradient of the density w.r.t the input point and $\text{DistToMesh}(\mathbf{x})$ is the distance of \mathbf{x} to the mesh. With these three inputs, \mathcal{N} , is able to rely on the 3DMM mesh normals for points close to the head, while relying on gradient normals $\text{Grad}_{\mathbf{n}}(\mathbf{x})$ everywhere else. In supplementary, we demonstrate the utility of each input to \mathcal{N} by ablating each one of them.

We train \mathcal{N} through a combination of weak supervision on mesh and scene normals, and regularization losses. The prediction of \mathcal{N} is forced to be weakly consistent with the 3DMM on its vertices as follows:

$$\mathcal{R}_{\text{Mesh}, \mathbf{n}} = \lambda_{\text{Mesh}, \mathbf{n}} \sum_v \|\mathcal{N}(v) - \text{Mesh}_{\mathbf{n}}(v)\| \quad (9)$$

where, v are the vertices of the mesh and $\lambda_{\text{Mesh}, \mathbf{n}}$ is the regularization constant. The normals predicted by \mathcal{N} are also forced to be forward facing by using the following regularization [43]:

$$\mathcal{R}_{\text{dir}, \mathbf{n}} = \sum_i w_i(\mathbf{x}_i) \max(0, \mathcal{N}(\mathbf{x}_i) \cdot \mathbf{d}_i)^2 \quad (10)$$

where, \mathbf{x}_i are points along the ray passing through pixel i with direction \mathbf{d}_i and $w_i(\mathbf{x}_i)$ is the weight of \mathbf{x}_i per Eq. (1). In order to ensure the gradient density normals $\text{Grad}_{\mathbf{n}}(\mathbf{x})$ are themselves accurate, we regularize both the normals predicted by \mathcal{N} and the gradient density normals to be consistent with each other:

$$\mathcal{R}_{\mathbf{n}} = \sum_i w_i(\mathbf{x}_i) \|\mathcal{N}(\mathbf{x}_i) - \text{Grad}_{\mathbf{n}}(\mathbf{x})\| \quad (11)$$




Figure 4. Dynamic Surface Normals. Surface normals directly influence the specularities and shading of the human head. Here we visualise the surface normals given by the density field of CoDyNeRF. As can be seen, the normals are able to capture the dynamic foreground with fine detail without sacrificing accuracy. In comparison, the normals of RigNeRF [3] are noisy and fail to capture the facial geometry accurately.

As observed in [43], this regularization ensures that highly specular surfaces of the scene are not explained away as subsurface emissive lobes. The full loss on \mathcal{N} is:

$$\mathcal{L}_{\mathcal{N}} = \underbrace{\mathcal{R}_{Mesh,n}}_{\text{Ensures consistency of normals on vertices of the head}} + \underbrace{\mathcal{R}_{dir,n}}_{\text{Ensures normals are forward facing}} + \underbrace{\mathcal{R}_n}_{\text{Ensures consistency between } \mathcal{N} \text{ and Grad}_n \text{ normals}} \quad (12)$$

Sparsity regularization for fast calculation of $\mathcal{R}_{dir,n}$ and \mathcal{R}_n . Calculating both $\mathcal{R}_{dir,n}$ and \mathcal{R}_n is very computationally expensive as it requires a second derivative calculation at each point along the ray (usually ~ 100 points for most standard NeRF architectures) for each sampled ray in the batch (typically around 1000 rays). A simple way to reduce the computational burden is to evaluate the above sum only on a subset of the points on a ray as follows

$$\mathcal{R}_n = \sum_i w_i(\mathbf{x}'_i) \|\mathcal{N}(\mathbf{x}'_i) - \text{Grad}_n(\mathbf{x}'_i)\|; \quad (13)$$

where $\mathbf{x}'_i \in \mathcal{S}_{i,k}$

where, $\mathcal{S}_{i,k}$ is the set of top k points, sorted by weight $w_i(\mathbf{x}'_i)$, of the ray passing through pixel i . However, as the weights predicted by the NeRF are relatively broadly distributed, such regularization does not minimize Eq. (13) over the whole scene consistently. To ensure the predicted weights are more tightly distributed around the surface, we enforce a Cauchy regularization to enforce sparsity [15]

$$\mathcal{R}_{cauchy} = \lambda_c \sum_i \log \left(1 + \frac{\sigma(\mathbf{x}_i)^2}{c} \right) \quad (14)$$

similar to [15], we only apply this on the coarse MLP. Eq. (13) and Eq. (14) improve the underlying dynamic scene geometry and significantly improve the quality of the gradient density normals.

4. Experimental Results

In this section, we show results of head-pose control, facial expression control, and novel view synthesis using CoDyNeRF.

For each scene, the model is trained on a short portrait video captured using a consumer smartphone.

Baseline approaches. We compare CoDyNeRF quantitatively and qualitatively to the following prior work for human head animation 1) RigNeRF [3] is a method for reanimating neural portraits with full control of facial expressions, head pose and viewing direction. Similar to us, the authors use a volumetric representation to model the dynamic scene. 2) Neural Head Avatars [13] creates a neural head model by deforming a mesh with pose and expression dependent offsets along with an MLP dependent texture space. 3) PointAvatar [50] uses a point cloud and an SDF as the geometric representation of the avatar in the canonical space. Additionally, PointAvatar separates the appearance of the avatar into an albedo and a normals-conditioned RGB shading. RigNeRF [3] and Neural Head Avatars [13] use an expression and pose dependent canonical space while PointAvatar uses a normals conditioned shading network to handle appearance variation during the capture. Unlike in [3], we *do not optimize* the appearance or deformation latent code during testing. Since we want to evaluate the fidelity of reproduction of illumination dependent effects for novel head-pose, facial expressions and views we cannot assume access to testing frames. We always use the deformation and appearance code of the first frame during testing.

Training and evaluation Data. The training and validation data was captured using an iPhone 13 Pro Max for all the experiments in the paper. In the first half of the capture, we ask the subject to enact a wide range of expressions and speech while trying to keep their head still as the camera is panned around them. In the next half, the camera is fixed at head-level and the subject is asked to rotate their head as they enact a wide range of expressions. Camera parameters are calculated using COLMAP [37]. FLAME [24] parameters and spherical harmonics coefficients L_{lm} are obtained via standard photometric and landmark fitting obtained by [14]. All videos are between 50-80 seconds long; we use the first ~ 1200 -1500 frames for training and the remaining 120-150 held out frames, with novel expressions and head-poses, for validation. Please find full details of each experiment in the supplementary document.

4.1. Evaluation on Test Data

We evaluate CoDyNeRF, along with three state-of-the-art baselines, Neural Head Avatars [13], RigNeRF [3], and PointAvatar [50] on held out testing frames. These frames contain a variety of facial expressions and head-poses. In Fig 5, we show a qualitative comparison between CoDyNeRF and the baselines. We observe that outdoor-captured videos, where pose and expression deformations under sunlight creates large appearance variations, pose a significant challenge to existing methods[3, 13, 50]. Neural Head



Figure 5. Qualitative comparison by reanimation with novel facial expression, head-pose, and camera view parameters. Here we reanimate CoDyNeRF, RigNeRF [3] and Neural Head Avatars [13] with novel facial expressions and head-pose extracted from the ground-truth frame in Row 1. As can be observed in the highlighted red boxes, while RigNeRF [3] is able to generate some shadowing and specular effects on the face, they are incorrect. Additionally, RigNeRF’s results have significant artefacts around the mouth and face regions. Similarly, due to the use of an explicit mesh and the entanglement of illumination effects with expressions, Neural Head Avatars [13] is unable to recover accurate geometries or predict accurate illuminations effects. Since PointAvatar’s design does not take into account cast shadows and specularities, it struggles to reproduce them accurately as can be seen in columns (1) - (5). In contrast, our approach, CoDyNeRF, is able to accurately reproduce cast shadows and specularities without sacrificing the quality of facial expressions and head-pose articulation.

Models	Subject 1			Subject 2			Subject 3			Subject 4		
	PSNR \uparrow	LPIPS \downarrow	FaceMSE \downarrow	PSNR \uparrow	LPIPS \downarrow	FaceMSE \downarrow	PSNR \uparrow	LPIPS \downarrow	FaceMSE \downarrow	PSNR \uparrow	LPIPS \downarrow	FaceMSE \downarrow
CoDyNeRF (Ours)	23.46	0.40	$1.92e-3$	21.24	0.33	$2.16e-3$	22.3	0.53	$1.0e-3$	23.56	0.40	$1.5e-3$
RigNeRF [3]	21.0	0.41	$3.64e-3$	20.86	0.35	$3.60e-3$	21.67	0.54	$1.7e-3$	22.26	0.45	$2.1e-3$
NHA [13]	-	-	0.058	-	-	0.0382	-	-	0.070	-	-	0.068
PointAvatar [50]	-	-	$6.3e-3$	-	-	$8.38e-3$	-	-	$1.05e-2$	-	-	$7.74e-3$

Table 1. Quantitative results of Subject 1,2,3 and 4 on test data. Here we calculate PSNR and LPIPS over the full image while FaceMSE is only restricted to the MSE calculated over the face region. Our results are better than RigNeRF [3], Neural Head Avatars [13] and IMAvatar [49] across all subjects.

Avatars [13] and RigNeRF [3], entangle illumination dependent effects with expression parameters in their canonical

space and are not able to faithfully reconstruct the target appearance, creating artefacts in the overall appearance in-

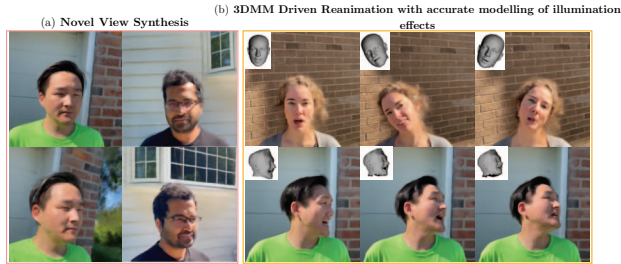


Figure 6. Applications of CoDyNeRF. CoDyNeRF enables the creation of Neural 3D Portraits captured in challenging lighting conditions with the accurate reproduction of illumination effects during novel expression and pose reanimation. In (a), we demonstrate that specularities, especially on the forehead and the nose, change realistically with changing viewing directions. In (b) we demonstrate the realistic reproduction of shading, shadows and specularities when the neural portrait is reanimated with a driving 3DMM (shown as white insets in the top left of each image). We see from results in the top row of (b) that the shading on the forehead and the specularities change realistically. Similarly, in the bottom row of (b) we see the nose being unshaded as the head rotates and the shadow cast by the ear changes realistically.

Models	Subject 1			Subject 2		
	PSNR \uparrow	LPIPS \downarrow	FaceMSE \downarrow	PSNR \uparrow	LPIPS \downarrow	FaceMSE \downarrow
CoDyNeRF $-\beta_{exp,pose}, -normals, -v_{lmk}$	19.8	0.47	4.03e-3	19.1	0.51	4.26e-3
CoDyNeRF $-v_{lmk}$	22.1	0.45	2.64e-3	20.93	0.33	3.30e-3
CoDyNeRF (Full Model)	23.46	0.40	1.92e-3	21.24	0.37	2.46e-3

Table 2. Ablation of dynamic appearance conditioning.

cluding in shadows, specularities, and in the mouth and eye region. Neural Head Avatars [13] is able to generate some shadows and specularities but they often lack of details and do not match the ground truth well. Further, due to their heavy reliance on an explicit face mesh, results from [13] are unable to accurately capture the head geometry in detail, resulting in an unnaturally deformed shape. RigNeRF [3], similar to NHA [13], uses an expression and pose (through deformation features) conditioned canonical space and is able to generate some shadows and specularities. However, they are often inaccurate and visually unnatural. For example, the specularity columns 1 and 2 of Fig 5 are incorrect, as well as the specularity around the eyes of columns 3 and 4. Similarly, the shadow on the mouth of column 5 is incorrect. Additionally, RigNeRF [3] results have significant artefacts, especially around the mouth and the eyes. Artefacts around the mouth can be seen in columns (1), (2), (3) and (5) and there are artefacts around the eyes in columns (2), (3) and (4) of Fig 5 respectively. PointAvatar [50] uses separate shading and texture networks to disentangle the albedo from the shading. However, its design does not take into account cast shadows and specularities, thus it is unable to learn and predict them*. PointAvatar is unable to reproduce the shadows and specularities on the forehead in columns (1) - (4) of Fig 5. It is also unable to predict the shadow cast by the nose

*We analyse this further in the supplementary section.

on the mouth in column (5) of Fig 5. In contrast, CoDyNeRF can faithfully reproduce the shadows and specularities on the forehead in columns (1) - (4) and the shadow cast by the nose on the mouth in column (5) of Fig 5.

In contrast to prior work, CoDyNeRF is able to faithfully reproduce illumination effects and generate high quality renders. As can be seen in Fig 5, our method captures dynamic appearance details, such as shadow patterns and specularities, under varying pose and expression more accurately *without* sacrificing the quality of expression and head-pose articulation. In Table 1, we provide quantitative evaluation of these four methods. We report image similarity measurement (PSNR), perceptual quality (LPIPS) and quality of facial appearance reconstruction (FaceMSE i.e MSE measured over the face region only) on novel expressions and head-poses. As we can see, CoDyNeRF outperforms previous approaches across all subjects.

Ablation of the Dynamic Canonical appearance Model.

In Table 2, we have ablated the conditioning of the canonical appearance model on 3DMM-based inputs and normals. First, we measured the performance of model with a static canonical space. This model does not use any surface normals or 3DMM-based information as input, and as can be seen in Table 2, it performs quite poorly. When the surface normals, expression and pose parameters are added, the results are much better, but the model still lacks global head-position conditioning via the landmarks i.e v_{lmk} , and thus is unable to reason about self-shadowing of the face. Our full model (third row), with conditioning from the surface normals, 3DMM expression, pose and the landmarks performs the best.

5. Conclusion

In conclusion, we present CoDyNeRF, a method that enables the creation of reanimatable Neural 3D Portraits captured in challenging lighting conditions. Using a dynamic canonical appearance representation, CoDyNeRF is able to accurately reproduce illumination effects with varying facial expression and head-pose. Additionally, we also propose a dynamic normals prediction module that utilizes 3DMM priors, along with a importance sampling based regularizer, to predict accurate dynamic surface normals. We believe this work takes us a step closer to creating in-the-wild Neural 3D portraits that are captured with casual smartphone devices.

6. Acknowledgements

This work was supported in part by the CDC/NIOSH through grant U01 OH012476 and a gift from Adobe.

References

- [1] ShahRukh Athar, Albert Pumarola, Francesc Moreno-Noguer, and Dimitris Samaras. Faceted3d: Facial expressions with 3d geometric detail prediction. *arXiv preprint arXiv:2012.07999*, 2020. 3
- [2] ShahRukh Athar, Zhixin Shu, and Dimitris Samaras. Self-supervised deformation modeling for facial expression editing. In *Face and Gesture*, 2020. 3
- [3] ShahRukh Athar, Zexiang Xu, Kalyan Sunkavalli, Eli Shechtman, and Zhixin Shu. Rignerf: Fully controllable neural 3d portraits. In *CVPR*, June 2022. 1, 2, 3, 4, 6, 7, 8
- [4] Mojtaba Bemana, Karol Myszkowski, Hans-Peter Seidel, and Tobias Ritschel. X-fields: Implicit neural view-, light-and time-image interpolation. *ACM TOG*, 2020. 2
- [5] Julian Chibane, Aayush Bansal, Verica Lazova, and Gerard Pons-Moll. Stereo radiance fields (srf): Learning view synthesis for sparse views of novel scenes. In *CVPR*, 2021. 2, 3
- [6] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, 2018. 3
- [7] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *CVPR*, 2020.
- [8] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and controllable face image generation via 3d imitative-contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5154–5163, 2020. 3
- [9] M. Doukas, Mohammad Rami Koujan, V. Sharmanska, A. Roussos, and S. Zafeiriou. Head2head++: Deep facial attributes re-targeting. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3:31–43, 2021. 3
- [10] Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction, 2020. 2, 3
- [11] Chen Gao, Yichang Shih, Wei-Sheng Lai, Chia-Kai Liang, and Jia-Bin Huang. Portrait neural radiance fields from a single image. *arXiv preprint arXiv:2012.05903*, 2020. 2, 3
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 3
- [13] Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. Neural head avatars from monocular rgb videos. In *CVPR*, pages 18653–18664, June 2022. 2, 3, 6, 7, 8
- [14] Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. Towards fast, accurate and stable 3d dense face alignment. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 6
- [15] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *ICCV*, pages 5875–5884, October 2021. 6
- [16] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 3
- [17] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.
- [18] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.
- [19] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 3
- [20] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. Deep video portraits. *ACM TOG*, 2018. 3
- [21] M. Koujan, M. Doukas, A. Roussos, and S. Zafeiriou. Head2head: Video-based neural head synthesis. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020) (FG)*, pages 319–326, Los Alamitos, CA, USA, may 2020. IEEE Computer Society. 3
- [22] Marek Kowalski, Stephan J Garbin, Virginia Estellers, Tadas Baltrušaitis, Matthew Johnson, and Jamie Shotton. Config: Controllable neural face image generation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 299–315. Springer, 2020. 3
- [23] Christoph Lassner and Michael Zollhöfer. Pulsar: Efficient sphere-based neural rendering. In *CVPR*, 2021. 3
- [24] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6), 2017. 6
- [25] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, and Zhaoyang Lv. Neural 3d video synthesis, 2021. 3
- [26] Zhengqi Li, Simon Niklaus, Noah Snively, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021. 2, 3
- [27] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Adv. Neural Inform. Process. Syst.*, 2020. 2
- [28] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering, 2021.
- [29] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. *arXiv:2008.02268*, 2020. 2
- [30] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. 2020. 2, 3, 4
- [31] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. *arXiv preprint arXiv:2104.10078*, 2021.
- [32] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *ICCV*, 2021. 2, 3, 4

- [33] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*, 2021. 3, 4
- [34] Albert Pumarola, Antonio Agudo, Aleix M Martinez, Alberto Sanfeliu, and Francesc Moreno-Noguer. Ganimation: One-shot anatomically consistent facial animation. *International Journal of Computer Vision*, 128(3):698–713, 2020. 3
- [35] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *CVPR*, 2021. 2, 3, 4
- [36] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *CVPR*, 2021. 3
- [37] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 6
- [38] Zhixin Shu, Mihir Sahasrabudhe, Riza Alp Guler, Dimitris Samaras, Nikos Paragios, and Iasonas Kokkinos. Deforming autoencoders: Unsupervised disentangling of shape and appearance. In *ECCV*, 2018. 3
- [39] Z. Shu, E. Yumer, S. Hadap, K. Sunkavalli, E. Shechtman, and D. Samaras. Neural face editing with intrinsic image disentangling. In *CVPR*, 2017. 3
- [40] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Adv. Neural Inform. Process. Syst.*, 2019. 3
- [41] Ayush Tewari, Mohamed Elgharib, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhöfer, and Christian Theobalt. Pie: Portrait image embedding for semantic control. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020. 3
- [42] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zöllhofer, and Christian Theobalt. Stylerig: Rigging stylegan for 3d control over portrait images, cvpr 2020. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, june 2020. 3
- [43] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T. Barron, and Pratul P. Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *CVPR*, pages 5491–5500, June 2022. 2, 3, 5, 6
- [44] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion. In *CVPR*, 2021.
- [45] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *CVPR*, 2021. 2, 3
- [46] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *NIPS*, 33, 2020.
- [47] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv:2010.07492*, 2020. 2
- [48] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv:2010.07492*, 2020. 2
- [49] Yufeng Zheng, Victoria Fernández Abrevaya, Marcel C. Bühler, Xu Chen, Michael J. Black, and Otmar Hilliges. I m avatar: Implicit morphable head avatars from videos. In *CVPR*, pages 13545–13555, June 2022. 2, 3, 7
- [50] Yufeng Zheng, Wang Yifan, Gordon Wetzstein, Michael J Black, and Otmar Hilliges. Pointavatar: Deformable point-based head avatars from videos. *arXiv preprint arXiv:2212.08377*, 2022. 6, 7, 8
- [51] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 3

Proceedings

**2024 International Conference in 3D
Vision
3DV 2024**

**18-21 March 2024
Davos, Switzerland**

Copyright © 2024 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved.

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries may photocopy beyond the limits of US copyright law, for private use of patrons, those articles in this volume that carry a code at the bottom of the first page, provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Other copying, reprint, or republication requests should be addressed to: IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 133, Piscataway, NJ 08855-1331.

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and, in the interests of timely dissemination, are published as presented and without change. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the IEEE Computer Society, or the Institute of Electrical and Electronics Engineers, Inc.

ISBN-13: 979-8-3503-6245-9
BMS Part # CFP24203-ART

Additional copies may be ordered from:

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: + 1 800 272 6657
Fax: + 1 714 821 4641
<http://computer.org/cspress>
csbooks@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: + 1 732 981 0060
Fax: + 1 732 981 9667
[http://shop.ieee.org/store/
customer-service@ieee.org](http://shop.ieee.org/store/customer-service@ieee.org)

IEEE Computer Society
Asia/Pacific Office
Watanabe Bldg., 1-4-2
Minami-Aoyama
Minato-ku, Tokyo 107-0062
JAPAN
Tel: + 81 3 3408 3118
Fax: + 81 3 3408 3553
tokyo.ofc@computer.org

Individual paper REPRINTS may be ordered at: <reprints@computer.org>

Editorial production and cover art by Javier Gurrola



*IEEE Computer Society
Conference Publishing Services (CPS)*

<http://www.computer.org/cps>