

Neural Point-based Volumetric Avatar: Surface-guided Neural Points for Efficient and Photorealistic Volumetric Head Avatar

Cong Wang Tsinghua University, Tencent AI Lab China wangcong20@mails.tsinghua.edu.cn

> Linchao Bao Tencent AI Lab China linchaobao@gmail.com

Di Kang Tencent AI Lab China di.kang@outlook.com

Ying Shan Tencent AI Lab China yingsshan@tencent.com Yan-Pei Cao Tencent AI Lab China caoyanpei@gmail.com

Song-Hai Zhang* Tsinghua University China shz@tsinghua.edu.cn



Figure 1: Neural Point-based Volumetric Avatar explores point-based neural representation combined with volume rendering, achieving high-fidelity facial animations (images and depth maps) while maintaining efficiency comparable to mesh-based methods (Tab. 1). During training, NPVA can adaptively allocate more points to challenging facial regions, forming a thicker "shell" (i.e., a higher variance of projected distances onto the face in the normal direction) and increasing capacity as needed. In the leftmost image, we show our rendering and the ground truth (GT) side-by-side on the right and left parts, respectively. Observe the close resemblance between our rendering quality and the GT.

ABSTRACT

Rendering photorealistic and dynamically moving human heads is crucial for ensuring a pleasant and immersive experience in AR/VR and video conferencing applications. However, existing methods often struggle to model challenging facial regions (e.g., mouth interior, eyes, and beard), resulting in unrealistic and blurry results. In this paper, we propose Neural Point-based Volumetric Avatar (NPVA), a method that adopts the neural point representation as

*Song-Hai Zhang is the corresponding author (shz@tsinghua.edu.cn).



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

SA Conference Papers '23, December 12–15, 2023, Sydney, NSW, Australia © 2023 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0315-7/23/12. https://doi.org/10.1145/3610548.3618204 well as the neural volume rendering process and discards the predefined connectivity and hard correspondence imposed by meshbased approaches. Specifically, the neural points are strategically constrained around the surface of the target expression via a highresolution UV displacement map, achieving increased modeling capacity and more accurate control. We introduce three technical innovations to improve the rendering and training efficiency: a patch-wise depth-guided (shading point) sampling strategy, a lightweight radiance decoding process, and a Grid-Error-Patch (GEP) ray sampling strategy during training. By design, our NPVA is better equipped to handle topologically changing regions and thin structures while also ensuring accurate expression control when animating avatars. Experiments conducted on three subjects from the Multiface dataset demonstrate the effectiveness of our designs, outperforming previous state-of-the-art methods, especially in handling challenging facial regions.

Cong Wang, Di Kang, Yan-Pei Cao, Linchao Bao, Ying Shan, and Song-Hai Zhang

CCS CONCEPTS

• Computing methodologies \rightarrow Rendering.

KEYWORDS

Neural representation, volume rendering, high-fidelity head avatars

ACM Reference Format:

Cong Wang, Di Kang, Yan-Pei Cao, Linchao Bao, Ying Shan, and Song-Hai Zhang. 2023. Neural Point-based Volumetric Avatar: Surface-guided Neural Points for Efficient and Photorealistic Volumetric Head Avatar. In *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers '23), December 12–15, 2023, Sydney, NSW, Australia.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3610548.3618204

1 INTRODUCTION

Realizing photorealistic rendering of an animatable human head is a pivotal goal in computer graphics and vision, which has broad applications such as AR/VR communications [He et al. 2020; Lombardi et al. 2018; Orts-Escolano et al. 2016], gaming [Waggoner 2009], and remote collaboration [Wang et al. 2021a]. However, providing a satisfying and immersive experience in these applications remains immensely challenging due to our innate ability to express and perceive emotions through subtle facial cues [Ekman 1980]. Existing data-driven learning methods often generate noticeable artifacts in the mouth area and blurry beard textures [Egger et al. 2020; Grassal et al. 2022; Khakhulin et al. 2022; Lombardi et al. 2018; Ma et al. 2021; Zollhöfer et al. 2018]. This limitation is primarily attributed to their underlying mesh-based representations, since the predefined mesh has a fixed topology and limited discretization resolution.

To illustrate this limitation, consider the Deep Appearance Model (DAM) [Lombardi et al. 2018] that decodes a head mesh and the corresponding view-specific texture UV map. Despite achieving high-quality rendering results for the skin regions, DAM produces conspicuous artifacts in the mouth and hair regions due to inaccurate correspondences across frames (e.g., mouth interior) and the mesh's inability to model thin structures (e.g., beard). To alleviate these issues, Pixel Codec Avatar (PiCA) [Ma et al. 2021] proposes the use of neural textures, allowing the subsequent neural renderer to address inaccurate shape estimation and topological inconsistencies (e.g., closed/open mouth), attaining moderately improved facial geometry and renditions for the mouth region. Similarly, Mixture of Volumetric Primitives (MVP) [Lombardi et al. 2021] attaches volumetric primitives (predicted by a CNN) to mesh vertices, replacing vertex colors with more flexible volumetric primitives. However, these enhanced "texture"-like representations remain embedded in a predefined topology and tend to produce blurry results if inaccurate correspondences occur.

Therefore, we propose Neural Point-based Volumetric Avatar (NPVA) that leverages highly flexible neural points [Aliev et al. 2020] and versatile neural volume rendering, enabling sharper rendering for both topologically changing geometries (e.g., mouth interior) and translucent thin beard structures.

To create animatable head avatars, another key technical challenge lies in enhancing the controllability of neural points to generate accurate target expressions. To address this, we reconstruct an intermediate coarse geometry (represented as a UV position map) of the driving signal (i.e., target expression) and constrain the movable neural points close to its surface. Further, to maximize the benefits of neural volume rendering, we introduce an additional displacement map, which allows the points to move to more optimal positions around the surface. For instance, after training, more points are located inside the mouth, resulting in a thicker "point shell" and increased modeling capability for volume rendering.

Efficient rendering and training are also indispensable for practical applications. To this end, we propose three technical innovations. (1) We introduce a novel depth-guided [Lin et al. 2022] sampling method that incorporates local depth context information (i.e. a patch), achieving more realistic rendering while reducing the rendering time by ~ 10× compared to the vanilla NeRF. (2) We develop a lightweight radiance decoding process that eliminates unnecessary per-point processing used in [Bai et al. 2023; Xu et al. 2022], significantly improving rendering efficiency (~ 7×) and offering better generalization in our dynamic modeling task. (3) Lastly, to speed up training, we propose a novel Grid-Error-Patch (GEP) ray sampling strategy comprising three stages: a uniform grid-sampling stage for rapid initialization of a coarse result, an error-based importance sampling that delves into more challenging regions, and a patch-based stage to impose high-level perceptual image losses.

Our contributions are summarized below:

- We propose a novel volumetric representation based on neural points that are dynamically allocated around the surface (i.e., the target expression) for animatable head avatar creation. This representation is inherently capable of better handling thin geometry and topological changes.
- We introduce three technical innovations to enable efficient rendering and training, including a patch-wise depth-guided shading point sampling method, a lightweight radiance decoding process, and a Grid-Error-Patch training strategy.
- Experiments on the Multiface dataset show that our approach produces higher-quality images for novel expressions and novel views while being ~ 70× faster than NeRF.

2 RELATED WORKS

2.1 Continuous Neural Representations.

Neural implicit representations, such as ONet [Mescheder et al. 2019], DeepSDF [Park et al. 2019], and NeRF [Mildenhall et al. 2020], parameterize a continuous occupancy/SDF/radiance fields with a network (e.g., a multilayer perceptrons, MLP). The property (e.g., occupancy, color/density) of any location in the scene could be queried by feeding its coordinate (optionally along with specified viewing direction) to the MLP, which is also referred as the coordinate network.

Rapid progresses have been achieved on shape modeling [Chen and Zhang 2019; Martel et al. 2021; Tretschk et al. 2020; Wang et al. 2021b; Wu et al. 2023b] and appearance modeling [Kellnhofer et al. 2021; Schwarz et al. 2020; Sitzmann et al. 2019; Zheng et al. 2023b]. For example, PVA [Raj et al. 2021] and KeypointNeRF [Mihajlovic et al. 2022] obtain facial neural radiance fields from multiview images for high-fidelity static face reconstruction. Later, these neural implicit representations are further extended to handle dynamic scenes (e.g., a talking head) by introducing deformation techniques [Gafni et al. 2021; Tang et al. 2022; Wu et al. 2023a; Yao

Neural Point-based Volumetric Avatar

SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia



Figure 2: Overview of Neural Point-based Volumetric Avatar (NPVA). The core of our approach is a neural point-based volumetric representation (middle), with points distributed around the surface of the target expression. This surface is defined by the low-resolution position map \hat{G}_o with intermediate supervision. A high-resolution displacement map \hat{G}_d allows the points to adaptively move within a certain range, as needed to provide increased capacity in more challenging regions (e.g., mouth, hair/beard). The attached point features are obtained from the feature map \hat{F} . \hat{G}_o , \hat{G}_d , and \hat{F} are decoded from the latent code z (left), which is trained in a variational auto-encoding style (encoder omitted). In addition, we propose three technical innovations with the aim of achieving rendering efficiency on par with mesh-based methods (right).

et al. 2022; Yu et al. 2023]. For instance, NerFace [Gafni et al. 2021] conditions the NeRF on 3DMM coefficients, making a small step to animatable avatar creation. DFA-NeRF [Yao et al. 2022] further conditions NeRF with disentangled face attribute features and enables more detailed control over the talking head.

Methods using neural implicit representations assume no fixed topology and have infinite resolution (in theory), achieving impressive novel view synthesis results for dynamic scenes. However, they usually suffer from various artifacts (e.g., blur, distortion) and inaccurate control when rendering novel expressions/poses, which is critical for creating high-quality animatable avatars. In contrast, our approach employs an explicit point-based representation driven by a coarse surface (represented as a UV position map) of the novel expression, achieving better generalization on novel expression and precise expression control.

2.2 Discrete Neural Representations.

Common discrete representations in CV/CG include mesh, grid/voxel, and point cloud. Recently, these representations have been extended with neural features to not only reconstruct the shape, but also create photorealistic image renderings thanks to the rapid progress in differentiable neural rendering [Cao et al. 2022; Laine et al. 2020; Lassner and Zollhöfer 2021; Ma et al. 2023; Rakhimov et al. 2022; Rückert et al. 2022].

For example, neural volumes [Lombardi et al. 2019] learns a radiance field in the canonical space and introduce another warp field to handle dynamic scenes. However, regular grid/voxel-based representations usually suffer from the required cubic memory footprint and slow rendering speed due to the processing of empty voxels.

Another line of works uses mesh-based neural representations [Lombardi et al. 2021; Ma et al. 2021]. For example, PiCA employs a neural UV texture map to render the head avatar in different expressions from any given viewpoint. Neural head avatar [Grassal et al. 2022] introduces additional geometry and texture networks in complementary to a base FLAME [Li et al. 2017] head model, resulting in a better shape and texture modeling. However, mesh-based representations require accurate surface geometry and semantically consistent correspondence, which is usually hard to acquire. Alternatively, MVP combines mesh and grid-based representations, and uses the volume rendering technique to render images. However, the color and density of MVP's primitives are decoded from 2D CNNs and further aggregated to obtain the radiance of query points, which inevitably causes blurry renderings. As a result, we propose to explore a more flexible neural point-based representation. Concurrent with our work, PointAvatar [Zheng et al. 2023a] also uses a point-based representation. Different from ours, their points only store color information and use splatting for rendering. We utilize powerful volume rendering, which has the potential to render higher-quality hairs and beards.

2.3 Differentiable Neural Rendering.

Rapid progress in differentiable neural rendering [Jiang et al. 2020; Liu et al. 2019, 2020; Müller et al. 2022] makes analysis-by-synthesis pipeline more powerful than before and enables avatar creation from even a short monocular video [Grassal et al. 2022; Zheng et al. 2023a]. For example, deferred neural rendering [Thies et al. 2019] proposes a neural texture map, which is decoded by a neural renderer into high quality renderings from any viewpoint. Recently, neural volume rendering [Mildenhall et al. 2020] achieves impressive high-quality renderings by introducing volume rendering to neural implicit fields. What's more, volume rendering can naturally model translucent objects (e.g., hair, smoke) in very high fidelity, which is very suitable for head modeling. However, a major drawback of neural volume rendering is low rendering efficiency. Our approach employs volume rendering to ensure high-fidelity visual

results, and proposes a series of acceleration methods to achieve rendering speed matching neural texture rendering.

3 METHODS

An overview of our Neural Point-based Volumetric Avatar (NPVA) is shown in Fig. 2. Given a latent code corresponding to the target facial state, we employ three decoders to generate a position map \hat{G}_o , a displacement map \hat{G}_d , and a feature map \hat{F} , respectively. Our NPVA representation is constructed from these maps, followed by a point-based neural volume rendering to efficiently produce high-fidelity images and detailed depth maps from any viewpoint.

In Sec. 3.1, we introduce our NPVA representation that leverages flexible point clouds for improved modeling of topological changes and thin structures, and utilizes volume rendering to produce high-fidelity images. In Sec. 3.2, we present our efficient rendering and training strategies, enabling NPVA to render photorealistic images comparable to NeRF while being $\sim 70 \times$ faster. Sec. 3.3 details our training losses. Sec. 3.4 describes our implementation details.

3.1 Neural Point-based Volumetric Avatar

3.1.1 Neural Radiance Field. NeRF [Mildenhall et al. 2020] effectively encodes a static scene using a Multi-Layer Perceptron network (MLP), achieving unprecedented high-quality novel view synthesis results. During rendering, the trained MLP takes the scene coordinates $x \in \mathbb{R}^3$ and viewing direction (θ, ϕ) as input and produces the corresponding density σ_x and view-dependent color c_x . The final color of an image pixel is obtained via volume rendering (Eq. (1)), which integrates all shading points on the ray that passes through this image pixel. Mathematically, NeRF uses piece-wise constant density and color as an approximation:

$$c = \sum_{i=1}^{N} T_i (1 - \exp(-\sigma_i \delta_i)) c_i, \quad T_i = \exp(-\sum_{j=1}^{i-1} \sigma_j \delta_j)$$
(1)

where σ_i and c_i denote density and color, respectively, and δ_i is the distance between two adjacent shading points.

3.1.2 Animatable Neural Points. We employ *explicit* neural points in our NPVA to achieve more controllable deformation of the underlying implicit neural radiance field. Leveraging this geometry proxy can also significantly enhance the rendering efficiency (Sec. 3.2).

Inspired by Point-NeRF [Xu et al. 2022], our representation consists of a set of neural points, denoted as $\mathcal{A} = \{(\mathbf{p}_i, f_i) | i = 1, ..., N\},\$ where $p_i \in \mathbb{R}^3$ denotes the location of point *i*, and f_i represents its associated feature. In our network, these features are learned in a variational auto-encoding style, similar to DAM [Lombardi et al. 2018] and PiCA [Ma et al. 2021], to create an animatable avatar. Specifically, the geometry sub-network in the decoder predicts a UV position map \hat{G}_o and a UV displacement map \hat{G}_d . \hat{G}_o stores the vertex positions of a coarse head mesh. We apply additional intermediate supervision on the 256² position map \hat{G}_o to obtain better expression control. To achieve better expressiveness of the point-based neural radiance field, we upsample the position map \hat{G}_o to 1024² and incorporate the high-resolution displacement map \hat{G}_d to compensate for inaccuracy contained in the coarse geometry. By compositing the upsampled \hat{G}_o and \hat{G}_d , we determine the positions of the neural points. The final neural points can adjust

Cong Wang, Di Kang, Yan-Pei Cao, Linchao Bao, Ying Shan, and Song-Hai Zhang



Figure 3: Lightweight Radiance Decoding. Given a query point, we find its *K* nearest neighboring neural points. We weighted sum these points to obtain an "average" feature for the subsequent radiance decoding. Note that we removed the per-point processing MLP used in Point-NeRF. This lightweight decoding process runs faster and obtains better generalization on novel expressions for our dynamic modeling task.

their positions adaptively around the surface, as we only apply a regularization term to penalize unreasonably large displacements.

3.1.3 Lightweight Radiance Decoding. The radiance (i.e., color c_x and density σ_x) at position $x \in \mathbb{R}^3$ is extracted based on (up to) its *K* nearest neighboring points and a neural decoding MLP (see Fig. 3), inspired by prior methods [Xu et al. 2022; Zimny et al. 2022]. To increase efficiency and ensure better generalization on novel expressions, we design a novel lightweight radiance decoding process that directly aggregates the neural points and passes this "average" feature to a lightweight neural decoding network, achieving ~ 7× speedup (compared to Point-NeRF) and better renderings. Specifically, for a given shading point, we compute a weighted average of the features and relative positions of (up to) its *K* nearest neural points inside a sphere of radius *R* as in Point-NeRF:

$$f_{\mathbf{x}} = \sum_{i} \frac{w_{i}}{\sum w_{i}} f_{i}, \quad v_{\mathbf{x}} = \sum_{i} \frac{w_{i}}{\sum w_{i}} v_{i}$$
(2)

where $v_i \in \mathbb{R}^6$ is the position encoding of the displacement vector between neural point p_i and the shading point x, $w_i = \frac{1}{||p_i - x||_2}$ is the weight that is inversely proportional to the Euclidean distance between x and p_i . Then, two shallow MLPs are applied to decode this average feature to density σ_x and view-dependent color c_x :

$$\mathcal{F}_{d}: (f_{\mathbf{x}}, v_{\mathbf{x}}) \to \sigma_{\mathbf{x}}, \quad \mathcal{F}_{c}: (f_{\mathbf{x}}, v_{\mathbf{x}}, d) \to c_{\mathbf{x}}$$
(3)

Positional encoding [Mildenhall et al. 2020] is applied to every dimension of the input vector in \mathcal{F}_d and \mathcal{F}_c . Finally, a pixel color is obtained via the integration in Eq. (1).

The primary difference between our radiance decoding and previous methods [Bai et al. 2023; Xu et al. 2022] is that we omit the per-point feature processing MLP before the feature aggregation in Eq. (2). This modification provides two benefits: it results in $\sim 7 \times$ speedup and improved generalization to unseen expressions in our dynamic modeling task (Fig. 5).

Neural Point-based Volumetric Avatar

SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia



Figure 4: Qualitative Comparisons with State-of-the-Art Methods. Our NPVA produces more photorealistic facial renditions on held-out test expressions compared to previous state-of-the-art methods, particularly in challenging facial regions (i.e., eyes, beard, and mouth interior). The normal expression is presented in Row 1, while the extreme expressions are shown in Row 2. The bottom left corners of the leftmost images show the shell thickness of our NPVA. The thickness for a specified face is the variance of point-to-surface distance for the points corresponding to it. Red (blue) indicates larger (smaller) variance. During learning, our NPVA automatically increases capacity (i.e., thicker shell) to better model the more challenging facial regions (e.g., hair and mouth interior).

3.2 Efficient Rendering and Training

3.2.1 Patch-wise Depth-guided Volume Rendering. Since we have some prior knowledge about the scene (i.e., a head and its coarse shape), we can focus on sampling the shading points around the surface, significantly improving the rendering efficiency compared to the original NeRF.

We propose a patch-wise depth-guided sampling strategy, taking into account that the envelope of a head mesh is often not very accurate and the visible facial parts could appear on different depth levels (e.g., jaw and neck). Specifically, we define a fixed connectivity on the 256² position map \hat{G}_o so that we can easily rasterize a depth map D_{rast} . For a ray passing through $[p_x, p_y]$, we consider a local depth patch centered on it and obtain the minimum and maximum depth values D_{\min} and D_{\max} . In our implementation, we consider only nine pixels, i.e., { $D_{\text{rast}}(p_x + i, p_y + j) | i, j \in \{-s, 0, s\}$ }, where *s* is a hyper-parameter set to 16 on 1024×667 images.

If $D_{\max} - D_{\min} < \delta_d$ (smaller than a threshold), there is only one depth level (e.g., no sudden depth changes like jaw and neck), and we use $d_c = (D_{\max} + D_{\min})/2$ as our sampling center. If $D_{\max} - D_{\min} \ge \delta_d$, there are likely two depth levels, and we split our budget equally between them and sample shading points around D_{\min} and D_{\max} separately. For every depth level d_c , we uniformly randomly sample points $\mathbf{p}(t_i)$ within evenly spaced bins t_i centered at d_c as follows:

$$\boldsymbol{p}(t_i) = \boldsymbol{o} + t_i \overrightarrow{\boldsymbol{d}}, \quad t_i \sim \mathcal{U}[d_c + (\frac{2(i-1)}{N} - 1)r, d_c + (\frac{2i}{N} - 1)r] \quad (4)$$

where o is the optical center, \vec{d} indicates the view direction, \mathcal{U} represents uniform distribution, r is the sampling radius and equals 20 in our experiments, and N is the number of sampling points.

Discussion: E-NeRF [Lin et al. 2022] proposes a similar *pixel-wise* depth-guided sampling method, reducing the required shading points for volume rendering. However, this sampling method only considers the depth of the current pixel and cannot properly handle facial parts that appear on different depth levels (e.g., beard), yielding suboptimal results like mesh-based methods (see Fig. 6).

3.2.2 *GEP Training Strategy.* For head image rendering, artifacts usually appear on several difficult but small regions (i.e., mouth, eyes); therefore, uniformly sampling rays covering the entire head region is inefficient. To address this issue, we propose a three-stage ray sampling strategy that consists of a <u>Grid-based uniform</u> sampling stage to initiate the training, an <u>Error-based importance</u> sampling stage to refine the challenging regions, and a <u>Patch-based</u> sampling stage to improve perceptual quality.

Grid-Sample Stage (*G-Stage*). In this stage, we prioritize full coverage of the images. The image is split into equal-sized grids without overlap, and we randomly sample one ray per grid G-Stage ensures uniform sampling across all regions and generates an initial model that produces reasonable results for all regions of the image. Moreover, we keep track of the error for each grid, obtaining an error map *E* (8×8) for later error-based importance sampling.

Error-Sample Stage (E-Stage). During this E-Stage, we adjust the sampling probabilities of the grids based on the grid error map initialized in the previous G-Stage. Specifically, the sampling probability of a grid region is proportional to its grid error, resulting in an error-based importance sampling similar to [Sucar et al. 2021; Zhu et al. 2022]. As a result, we allocate more computing budget to difficult facial regions (e.g., mouth interior, hair, and eyes in "Error Sample" of Fig. 2), significantly improving image quality within the same number of training epochs (see Fig.8a). Note that we maintain

an error map of a smaller size in this stage and dynamically update the sampling probability.

Patch-Sample Stage (P-Stage). In this stage, we sample rays belonging to an image patch (instead of individual pixels) so that we can apply patch-based perceptual loss [Zhang et al. 2018]. Using a perceptual loss along with per-pixel losses (e.g., L_2) can help reduce image blur, resulting in sharper images and visually better results [Luo et al. 2021; Menon et al. 2020].

3.3 Training Losses

We have a set of images $\{I^{(i)}\}, i \in \{1, 2, ..., N\}$, their corresponding tracked meshes $\{\mathcal{M}^{(i)}\}$, UV position maps $\{G_o^{(i)}\}$ converted from $\{\mathcal{M}^{(i)}\}$, and reconstructed depth maps $\{D^{(i)}\}$ using Metashape software [Over et al. 2021]. Our training losses include per-pixel photometric loss \mathcal{L}_{pho} , patch-based perceptual loss \mathcal{L}_{per} [Zhang et al. 2018], coarse mesh loss \mathcal{L}_m , two depth losses \mathcal{L}_d and \mathcal{L}_{rd} , and three regularization losses \mathcal{L}_s , $\mathcal{L}_{\text{disp}}$ and \mathcal{L}_{kl} .

The appearance losses are defined as:

$$\mathcal{L}_{pho} = \sum_{p \in \mathcal{P}} ||I_p^{(i)} - \hat{I}_p^{(i)}||_2, \quad \mathcal{L}_{per} = \text{LPIPS}(I_{\mathcal{P}}^{(i)}, \hat{I}_{\mathcal{P}}^{(i)})$$
(5)

where, \hat{I} is our rendering image, and \mathcal{P} is a set of image coordinates used to train our networks. Note that \mathcal{L}_{per} is only used in the last P-stage due to its patch-based property.

The geometric losses, which help generate more controllable head avatars, are defined as:

$$\mathcal{L}_{m} = ||G_{o} - G_{o}||_{2},$$

$$\mathcal{L}_{rd} = ||(D^{(i)} - D_{rast}^{(i)}) \odot M_{D_{rast}}||_{1},$$

$$\mathcal{L}_{d} = \sum_{p \in \mathcal{P}} ||(D^{(i)}_{p} - \hat{D}^{(i)}_{p}) \odot M_{D}||_{1}$$
(6)

where \hat{G}_o is the decoded 256² position map, and D_{rast} is a coarse depth map rasterized using \hat{G}_o , \hat{D} is our fine depth map obtained using volume rendering. Note that depth masks M_D and $M_{D_{\text{rast}}}$ are used to penalize only those pixels whose depth errors are less than a depth threshold δ_D (set to 10mm) to handle outliers.

We also include three regularization losses to improve our model's generalization ability. $\mathcal{L}_{disp} = ||G_d \odot M_{G_d}||_2$ is a regularization term on the displacement map to constrain the final points close to the surface and prevent overfitting, where M_{G_d} is a mask to only penalize the points whose displacement values are larger than δ_{disp} (set to 10mm). \mathcal{L}_s is a Total Variation (TV) loss applied on the 256² UV position map G_o to encourage a smooth surface. \mathcal{L}_{kl} is the common Kullback-Leibler (KL) divergence prior applied on the latent space in VAE training.

In summary, the complete training loss is the weighted sum of these loss terms:

$$\mathcal{L} = \lambda_{\rm pho} \mathcal{L}_{\rm pho} + \lambda_{\rm per} \mathcal{L}_{\rm per} + \lambda_d \mathcal{L}_d + \lambda_{\rm rd} \mathcal{L}_{\rm rd} + \lambda_m \mathcal{L}_m + \lambda_s \mathcal{L}_s + \lambda_{\rm disp} \mathcal{L}_{\rm disp} + \lambda_{\rm kl} \mathcal{L}_{\rm kl}$$
(7)

3.4 Network Structures & Implementation Details

Our network is trained in a variational auto-encoding fashion [Kingma and Ba 2015] following DAM and PiCA [Lombardi et al. 2018, 2021;

Ma et al. 2021]. The encoder comprises 5 and 7 convolution layers (with the last 5 layers shared) respectively and encodes a UV position map, which is converted from a coarse tracked mesh (~ 5K vertices) and an average texture map into a latent code $z \in \mathbb{R}^{8 \times 8 \times 4}$ as in PiCA [Ma et al. 2021]. The average texture map is obtained from an open-mouth expression by averaging the unwrapped textures of all camera views. Note that the tracked mesh does not contain vertices for the tongue and teeth.

The decoder contains 5/7/7 convolution layers and predicts a position map, a displacement map, and a feature map, all of which are used for later radiance decoding. The position map G_p at 1024^2 represents a coarse (i.e., less detailed) surface as it is upsampled from G_o at 256^2 , which is supervised by the input coarse mesh \mathcal{M} . The displacement map G_d at 1024^2 increases details to compensate for lost geometric details of the coarse mesh \mathcal{M} during estimation. The 32-D feature map F at 1024^2 contains local appearance information around the point for radiance decoding.

We set loss weights of { λ_{pho} , λ_{per} , λ_d , λ_{rd} , λ_m , λ_s , λ_{disp} , λ_{kl} } in Eq. (7) as {5, 0.1, 0.1, 0.2, 0.2, 1, 0.1, 0.001} respectively. And the G/E/P-stages take 10/15/5 epochs, respectively.

4 EXPERIMENTS

We test on the Multiface dataset [Wuu et al. 2022]. It is an opensourced multi-view human face dataset that captures high-quality facial details from a camera array. Processed data include calibrated camera parameters, tracked meshes, and unwrapped UV texture maps (1024 × 1024). We obtain a depth map for each frame individually using Metashape [Over et al. 2021] based on the provided camera parameters.

Following MVP [Lombardi et al. 2021], we use downsampled images (1024×667) during training. Unless stated otherwise, all our experiments are trained on a subset of expressions and tested on held-out expressions (15 randomly chosen expressions and fixed), resulting in ~11K frames for training and ~1K frames for testing. Following PiCA [Ma et al. 2021], MSE and LPIPS are calculated based on image pixels under the rasterized mask for evaluation.

4.1 Comparisons with State-of-the-Art Methods

We compare with DAM, PiCA, and MVP to demonstrate the superiority of our approach via the rendered images under novel expressions. In Tab. 1, NPVA achieves the best MSE (up to 5.37 lower than the 2nd best). Fig. 4 and Fig. 10 demonstrate that NPVA produces more realistic facial renditions, especially in challenging facial regions (e.g., eyes, beard, mouth interior). Refer to our Supp. Mat. for video comparisons.

4.2 Ablation Studies

In this section, we present a series of ablation studies to verify the effectiveness of our major design choices.

Effect of using different numbers of points. We investigate the impact of using different numbers of points on the rendering quality and inference time in Tab. 2 (a.1)-(a.3). The visual results are shown in Fig. 9. When the point number is small (i.e., NPVA-256), NPVA generates poor results (46.45 MSE) and contains holes due to insufficient point resolution. With increased point number (NPVA-1k),

Table 1: Comparisons with State-of-the-Art Methods. Compared to state-of-the-art methods, NPVA achieves great improvements on LPIPS and MSE (up to 5.37 lower than the 2nd best) with slightly slower rendering speed. We bold (underline) the best (2nd best) results.

Methods	Subject 1		Subject 2		Subject 3		Subject 4		Subject 5		Inference
	MSE ↓	LPIPS \downarrow	MSE ↓	LPIPS \downarrow	MSE ↓	$ $ LPIPS \downarrow	$MSE\downarrow$	LPIPS \downarrow	MSE ↓	LPIPS \downarrow	Time (ms)
PiCA	34.50	0.232	28.83	0.108	25.49	0.239	30.22	0.285	21.32	0.224	73
DAM	28.40	0.208	23.53	0.088	23.18	0.183	24.38	0.248	20.21	0.185	<u>107</u>
MVP	48.59	0.242	28.25	0.102	38.82	0.262	36.23	0.268	25.46	0.230	144
NPVA (Ours)	23.70	0.160	18.16	0.075	21.95	0.183	21.88	0.180	17.13	0.141	482

Table 2: Ablation Studies. We demonstrate the impact of different components on the results of our NPVA method (evaluated using MSE). The first row shows our full method as reference. (a) shows the importance of using an extra displacement map, which cannot be replaced by simply increasing point numbers. (b) shows that our lightweight radiance decoding can produce better rendering results on held-out expressions with a faster rendering speed ($\sim 7 \times$). (c) shows the superiority of our patch-wise depth-guided sampling method. (d) shows that our GEP ray sampling strategy achieves great performance improvements by properly allocating the computation budget and introducing perceptual loss.

Label	Name	Point Num.	Disp. Map	KNN Search Radius - R (mm)	Radiance Decoding	GEP strategy	Sampling Method	MSE ↓	Inference Time (ms)
	NPVA-full (Ours)	1024×1024	 ✓ 	R=3	lightweight	 ✓ 	Patch-Depth	23.70	482
(a.1)	NPVA-256	256×256	X	R=3	lightweight	1	Patch-Depth	46.45	397
(a.2)	NPVA-1k	1024×1024	X	R=3	lightweight	1	Patch-Depth	26.36	423
(a.3)	NPVA-2k	2048×2048	×	R=3	lightweight	1	Patch-Depth	27.36	506
(b)	NPVA-heavy	1024×1024	1	R=3	Point-NeRF	/	Patch-Depth	24.51	3129
(c.1)	NPVA-noDepth	1024×1024	1	R=3	lightweight	1	NeRF	406.49	406
(c.2)	NPVA-PixDepth	1024×1024	1	R=3	lightweight	1	Pixel-Depth	24.92	413
(d)	NPVA-R4	1024×1024	1	R=4	lightweight	1	Patch-Depth	23.54	563
(e)	NPVA-noGEP	1024×1024	 ✓ 	R=3	lightweight	x	Patch-Depth	30.08	472

Table 3: Comparison with NeRF. Our NPVA produces high-fidelity rendering results comparable to NeRF with a rather faster rendering speed (\sim 70×), when trained both on a single frame and multiple frames.

Methods	Training Data	MSE ↓	PSNR ↑	Inference Time (ms)	
NeRF	single frame	13.63	36.79	38392	
NPVA (Ours)	single frame	17.22	35.77	524	
NPVA (Ours)	49 frames	18.75	35.54	531	

we notice obvious improvement (from 46.45 to 26.36 MSE) and do not see holes, indicating the point resolution is sufficient. Further increasing the point number (NPVA-2k) results in perceptually very similar results and slightly worse MSE (27.36 vs. 26.36).

Importance of using an extra displacement map. Introducing an extra displacement map reduces MSE from 26.36 to 23.70 (Tab. 2 (a.2)). A visual comparison is provided in Fig. 9. We can also see using a displacement map is much more effective than using more points. This is because using a displacement map enables a more flexible arrangement of the neural points so that they can not only move on the surface (i.e., along the tangent plane) but also move along the normal direction, forming a thicker shell with increased capacity (see Fig. 4 and Fig. 10).

Influence of the lightweight radiance decoding process. Using the proposed lightweight radiance decoding in Sec. 3.1.3 not only greatly



Figure 5: Effect of lightweight radiance decoding. Our lightweight radiance decoding not only gains $\sim 7 \times$ speedup over Point-NeRF's radiance decoding, but also produces better facial renditions for unseen expressions.

reduces the inference time (3129 vs. 482 ms), but also improves the quality (23.70 vs. 24.51 MSE). A visual comparison is provided in Fig. 5. A possible explanation is that using networks with too much capacity may cause overfitting and hinder the generalization of novel expressions for dynamic scene modeling tasks.

Different shading point sampling methods. We investigate the impact of different shading point sampling methods on a given ray. Tab. 2 (c.1, c.2) present results obtained using different sampling methods. Our patch-wise depth-guided sampling method

Cong Wang, Di Kang, Yan-Pei Cao, Linchao Bao, Ying Shan, and Song-Hai Zhang



Pixel-Wise Depth Sampling

Patch-Wise Depth Sampling (Ours)

Figure 6: Effect of different depth-guided shading point sampling strategies. There are two depth levels around the orange mark. With pixel-wise depth sampling [Lin et al. 2022], the model generates "mesh-like" artifacts. In contrast, our patchwise depth sampling samples around both the front and back depth levels, leading to more realistic results.

reduces MSE from 24.92 to 23.70 compared to pixel-wise depthguided proposed in [Lin et al. 2022]. Using "Pixel-Depth" easily causes inaccuracies, especially for the jaw region with two different depth levels, resulting in a mesh-like beard rendering (see Fig. 6). Training with a naive sampling strategy is slower and could not give a similar result (406.49 MSE) compared to the depth-guided sampling methods when using only 20 sample points per ray (~200 in NeRF).

Different KNN search radius. Increasing the KNN search radius from 3mm to 4mm improves rendering results (23.54 vs. 23.70 MSE) with longer processing time (563 vs. 482 ms). As shown in Fig. 8b, a larger search radius can eliminate holes in some extreme expressions.

GEP training strategy. We compare our GEP ray sampling strategy (Sec. 3.2.2) with a naive alternative that always uses uniformly sampled rays. GEP achieves lower MSE (23.70 vs. 30.08) in Tab. 2 (d). The training loss curves and visual comparisons are shown in Fig. 8a. The model trained with the naive strategy converges to a sub-optimal solution. Although achieving satisfactory renderings in smooth facial regions (e.g., skins), it struggles to handle challenging facial regions (e.g., eyes and mouth interior). In contrast, the model trained with our GEP strategy allocates more computing budget to these difficult regions and obtains more realistic facial renditions.

4.3 Analysis on Volumetric Methods

We compare with single-frame NeRF fitting [Mildenhall et al. 2020], which can be viewed as the upper limit of different volumetric avatars. The results are shown in Tab. 3 and Fig. 7. On single-frame fitting, our NPVA generates high-fidelity results comparable to NeRF while being \sim 70× faster (524 vs. 38392 ms) during inference. What's more, our NPVA can handle dynamic scenes (a 49-frame sequence) effectively with minor performance drops. In Fig. 7, we notice NPVA also generates visually better results (e.g., sharper and more realistic reflection effects) than NeRF, possibly due to the help of coarse geometry prior and the perceptual loss.

5 CONCLUSION & DISCUSSION

In this paper, we present a novel volumetric representation based on movable neural points for animatable avatar creation, focusing on both high-quality rendering and time efficiency. To ensure controllability and accurate expression control, we guide point locations with decoded coarse meshes of target expressions and constrain the points around the surface, which is supervised with the driving signal. To further enhance rendering quality, we increase the point number and incorporate an additional displacement map that adaptively adjusts after training. Moreover, our approach features three technical innovations tailored to improve training and rendering efficiency: lightweight radiance decoding, patch-wise depth-guided sampling, and a GEP training strategy.

Limitation. We rely on coarse mesh tracking for modeling and optimization, which generally works well but does not account for very long hair or diverse hairstyles, such as those not present in tested female subjects. When relaxing the regularization on the displacement map for these cases, our method tends to produce blurry results for novel expressions (Fig. 7b).

ACKNOWLEDGMENTS

This work was supported by the Natural Science Foundation of China (Project Number 62132012) and Tsinghua-Tencent Joint Laboratory for Internet Innovation Technology.

REFERENCES

- Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor S. Lempitsky. 2020. Neural Point-Based Graphics. In Computer Vision - ECCV 2020 -16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXII (Lecture Notes in Computer Science, Vol. 12367), Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, UK, 696–712. https://doi. org/10.1007/978-3-030-58542-6_42
- Ziqian Bai, Feitong Tan, Zeng Huang, Kripasindhu Sarkar, Danhang Tang, Di Qiu, Abhimitra Meka, Ruofei Du, Mingsong Dou, Sergio Orts-Escolano, Rohit Pandey, Ping Tan, Thabo Beeler, Sean Fanello, and Yinda Zhang. 2023. Learning Personalized High Quality Volumetric Head Avatars from Monocular RGB Videos. In IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023. IEEE, Canada, 16890–16900. https://doi.org/10.1109/ CVPR52729.2023.01620
- Chen Cao, Tomas Simon, Jin Kyu Kim, Gabe Schwartz, Michael Zollhöfer, Shunsuke Saito, Stephen Lombardi, Shih-En Wei, Danielle Belko, Shoou-I Yu, Yaser Sheikh, and Jason M. Saragih. 2022. Authentic volumetric avatars from a phone scan. *ACM Trans. Graph.* 41, 4 (2022), 163:1–163:19.
- Zhiqin Chen and Hao Zhang. 2019. Learning Implicit Fields for Generative Shape Modeling. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. Computer Vision Foundation / IEEE, USA, 5939–5948.
- Bernhard Egger, William A. P. Smith, Ayush Tewari, Stefanie Wuhrer, Michael Zollhöfer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, Christian Theobalt, Volker Blanz, and Thomas Vetter. 2020. 3D Morphable Face Models - Past, Present, and Future. ACM Trans. Graph. 39, 5 (2020), 157:1–157:38.
- P. Ekman. 1980. The Face of Man: Expressions of Universal Emotions in a New Guinea Village. Garland STPM Press, virtual. https://books.google.com.hk/books?id= cMO0AAAAIAAJ
- Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. 2021. Dynamic Neural Radiance Fields for Monocular 4D Facial Avatar Reconstruction. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. Computer Vision Foundation / IEEE, virtual, 8649–8658.
- Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. 2022. Neural Head Avatars from Monocular RGB Videos. In IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. IEEE, USA, 18632–18643.
- Zhenyi He, Ruofei Du, and Ken Perlin. 2020. CollaboVR: A Reconfigurable Framework for Creative Collaboration in Virtual Reality. In 2020 IEEE International Symposium on Mixed and Augmented Reality, ISMAR 2020, Recife/Porto de Galinhas, Brazil, November 9-13, 2020. IEEE, Brazil, 542–554. https://doi.org/10.1109/ISMAR50242. 2020.00082
- Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. 2020. SDFDiff: Differentiable Rendering of Signed Distance Fields for 3D Shape Optimization. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020. Computer Vision Foundation / IEEE, USA, 1248–1258.
- Petr Kellnhofer, Lars Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. 2021. Neural Lumigraph Rendering. In IEEE Conference on Computer Vision

Neural Point-based Volumetric Avatar

SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia

and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. Computer Vision Foundation / IEEE, virtual, 4287–4297.

- Taras Khakhulin, Vanessa Sklyarova, Victor Lempitsky, and Egor Zakharov. 2022. Realistic One-Shot Mesh-Based Head Avatars. In Computer Vision – ECCV 2022 – 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 13662), Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (Eds.). Springer, Israel, 345–362.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In 3rd International Conference on Learning Representations, ICLR 2015, May 7-9, 2015, Conference Track Proceedings, Yoshua Bengio and Yann LeCun (Eds.). Computer Vision Foundation / IEEE Computer Society, San Diego, CA, USA, .
- Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. 2020. Modular primitives for high-performance differentiable rendering. ACM Trans. Graph. 39, 6 (2020), 194:1–194:14.
- Christoph Lassner and Michael Zollhöfer. 2021. Pulsar: Efficient Sphere-Based Neural Rendering. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. Computer Vision Foundation / IEEE, ., 1440–1449.
- 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), 194:1–194:17.
- Haotong Lin, Sida Peng, Zhen Xu, Yunzhi Yan, Qing Shuai, Hujun Bao, and Xiaowei Zhou. 2022. Efficient Neural Radiance Fields for Interactive Free-viewpoint Video. In SIGGRAPH Asia 2022 Conference Papers, SA 2022, December 6-9, 2022, Soon Ki Jung, Jehee Lee, and Adam W. Bargteil (Eds.). ACM, Daegu, Republic of Korea, 39:1–39:9.
- Shichen Liu, Weikai Chen, Tianye Li, and Hao Li. 2019. Soft Rasterizer: A Differentiable Renderer for Image-Based 3D Reasoning. In 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, October 27 - November 2, 2019. IEEE, Seoul, Korea (South), 7707–7716.
- Shaohui Liu, Yinda Zhang, Songyou Peng, Boxin Shi, Marc Pollefeys, and Zhaopeng Cui. 2020. DIST: Rendering Deep Implicit Signed Distance Function With Differentiable Sphere Tracing. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, June 13-19, 2020. Computer Vision Foundation / IEEE, Seattle, WA, USA, 2016–2025.
- Stephen Lombardi, Jason M. Saragih, Tomas Simon, and Yaser Sheikh. 2018. Deep appearance models for face rendering. *ACM Trans. Graph.* 37, 4 (2018), 68.
- Stephen Lombardi, Tomas Simon, Jason M. Saragih, Gabriel Schwartz, Andreas M. Lehrmann, and Yaser Sheikh. 2019. Neural volumes: learning dynamic renderable volumes from images. ACM Trans. Graph. 38, 4 (2019), 65:1–65:14.
- Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhöfer, Yaser Sheikh, and Jason M. Saragih. 2021. Mixture of volumetric primitives for efficient neural rendering. ACM Trans. Graph. 40, 4 (2021), 59:1–59:13.
- Huiwen Luo, Koki Nagano, Han-Wei Kung, Qingguo Xu, Zejian Wang, Lingyu Wei, Liwen Hu, and Hao Li. 2021. Normalized Avatar Synthesis Using StyleGAN and Perceptual Refinement. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021.* Computer Vision Foundation / IEEE, ., 11662–11672.
- Shugao Ma, Tomas Simon, Jason M. Saragih, Dawei Wang, Yuecheng Li, Fernando De la Torre, and Yaser Sheikh. 2021. Pixel Codec Avatars. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021.* Computer Vision Foundation / IEEE, ., 64–73.
- Zhiyuan Ma, Xiangyu Zhu, Guojun Qi, Zhen Lei, and Lei Zhang. 2023. OTAvatar: One-Shot Talking Face Avatar with Controllable Tri-Plane Rendering. In IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023. IEEE, ., 16901–16910.
- Julien N. P. Martel, David B. Lindell, Connor Z. Lin, Eric R. Chan, Marco Monteiro, and Gordon Wetzstein. 2021. Acorn: adaptive coordinate networks for neural scene representation. ACM Trans. Graph. 40, 4 (2021), 58:1–58:13.
- Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin. 2020. PULSE: Self-Supervised Photo Upsampling via Latent Space Exploration of Generative Models. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020. Computer Vision Foundation / IEEE, ., 2434–2442.
- Lars M. Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. 2019. Occupancy Networks: Learning 3D Reconstruction in Function Space. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. Computer Vision Foundation / IEEE, ., 4460–4470.
- Marko Mihajlovic, Aayush Bansal, Michael Zollhöfer, Siyu Tang, and Shunsuke Saito. 2022. KeypointNeRF: Generalizing Image-Based Volumetric Avatars Using Relative Spatial Encoding of Keypoints. In Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XV (Lecture Notes in Computer Science, Vol. 13675), Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (Eds.). Springer, ., 179–197.
- 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 Computer Vision - ECCV 2020 - 16th European Conference,

Glasgow, UK, August 23-28, 2020, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 12346), Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, ., 405–421.

- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.* 41, 4 (2022), 102:1–102:15.
- Sergio Orts-Escolano, Christoph Rhemann, Sean Ryan Fanello, Wayne Chang, Adarsh Kowdle, Yury Degtyarev, David Kim, Philip L. Davidson, Sameh Khamis, Mingsong Dou, Vladimir Tankovich, Charles T. Loop, Qin Cai, Philip A. Chou, Sarah Mennicken, Julien P. C. Valentin, Vivek Pradeep, Shenlong Wang, Sing Bing Kang, Pushmeet Kohli, Yuliya Lutchyn, Cem Keskin, and Shahram Izadi. 2016. Holoportation: Virtual 3D Teleportation in Real-time. In Proceedings of the 29th Annual Symposium on User Interface Software and Technology, UIST 2016, Tokyo, Japan, October 16-19, 2016, Jun Rekimoto, Takeo Igarashi, Jacob O. Wobbrock, and Daniel Avrahami (Eds.). ACM, ., 741–754.
- Jin-Si R Over, Andrew C Ritchie, Christine J Kranenburg, Jenna A Brown, Daniel D Buscombe, Tom Noble, Christopher R Sherwood, Jonathan A Warrick, and Phillipe A Wernette. 2021. Processing coastal imagery with Agisoft Metashape Professional Edition, version 1.6—Structure from motion workflow documentation. Technical Report. US Geological Survey.
- Jeong Joon Park, Peter R. Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. 2019. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. Computer Vision Foundation / IEEE, ., 165–174.
- Amit Raj, Michael Zollhöfer, Tomas Simon, Jason M. Saragih, Shunsuke Saito, James Hays, and Stephen Lombardi. 2021. PVA: Pixel-aligned Volumetric Avatars. CoRR abs/2101.02697 (2021), .
- Ruslan Rakhimov, Andrei-Timotei Ardelean, Victor Lempitsky, and Evgeny Burnaev. 2022. NPBG++: Accelerating Neural Point-Based Graphics. In IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. IEEE, ., 15948–15958.
- Darius Rückert, Linus Franke, and Marc Stamminger. 2022. ADOP: approximate differentiable one-pixel point rendering. ACM Trans. Graph. 41, 4 (2022), 99:1– 99:14.
- Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. 2020. GRAF: Generative Radiance Fields for 3D-Aware Image Synthesis. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.)...,.
- Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. 2019. Scene Representation Networks: Continuous 3D-Structure-Aware Neural Scene Representations. In Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.)..,., 1119–1130.
- Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J. Davison. 2021. iMAP: Implicit Mapping and Positioning in Real-Time. In 2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021. IEEE, ., 6209–6218.
- Junshu Tang, Bo Zhang, Binxin Yang, Ting Zhang, Dong Chen, Lizhuang Ma, and Fang Wen. 2022. Explicitly controllable 3d-aware portrait generation. arXiv preprint arXiv:2209.05434., (2022), .
- Justus Thies, Michael Zollhöfer, and Matthias Nießner. 2019. Deferred neural rendering: image synthesis using neural textures. ACM Trans. Graph. 38, 4 (2019), 66:1–66:12.
- Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Carsten Stoll, and Christian Theobalt. 2020. PatchNets: Patch-Based Generalizable Deep Implicit 3D Shape Representations. In Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XVI (Lecture Notes in Computer Science, Vol. 12361). Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, ., 293–309.
- Zach Waggoner. 2009. My avatar, my self: Identity in video role-playing games. McFarland, .
- Peng Wang, Xiaoliang Bai, Mark Billinghurst, Shusheng Zhang, Xiangyu Zhang, Shuxia Wang, Weiping He, Yuxiang Yan, and Hongyu Ji. 2021a. AR/MR Remote Collaboration on Physical Tasks: A Review. *Robotics Comput. Integr. Manuf.* 72 (2021), 102071.
- 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. In Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual, Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (Eds.)..., 27171–27183.
- Sijing Wu, Yichao Yan, Yunhao Li, Yuhao Cheng, Wenhan Zhu, Ke Gao, Xiaobo Li, and Guangtao Zhai. 2023b. GANHead: Towards Generative Animatable Neural Head Avatars. In IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023. IEEE, ., 437–447.

- Xiuzhe Wu, Pengfei Hu, Yang Wu, Xiaoyang Lyu, Yan-Pei Cao, Ying Shan, Wenming Yang, Zhongqian Sun, and Xiaojuan Qi. 2023a. Speech2Lip: High-fidelity Speech to Lip Generation by Learning from a Short Video. arXiv preprint arXiv:2309.04814.,. (2023),.
- Cheng-hsin Wuu, Ningyuan Zheng, Scott Ardisson, Rohan Bali, Danielle Belko, Eric Brockmeyer, Lucas Evans, Timothy Godisart, Hyowon Ha, Alexander Hypes, Taylor Koska, Steven Krenn, Stephen Lombardi, Xiaomin Luo, Kevyn McPhail, Laura Millerschoen, Michal Perdoch, Mark Pitts, Alexander Richard, Jason M. Saragih, Junko Saragih, Takaaki Shiratori, Tomas Simon, Matt Stewart, Autumn Trimble, Xinshuo Weng, David Whitewolf, Chenglei Wu, Shoou-I Yu, and Yaser Sheikh. 2022. Multiface: A Dataset for Neural Face Rendering. CoRR abs/2207.11243 (2022), .
- Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. 2022. Point-NeRF: Point-based Neural Radiance Fields. In IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. IEEE, ., 5428–5438.
- Shunyu Yao, Ruizhe Zhong, Yichao Yan, Guangtao Zhai, and Xiaokang Yang. 2022. DFA-NeRF: Personalized Talking Head Generation via Disentangled Face Attributes Neural Rendering. CoRR abs/2201.00791 (2022), .
- Wangbo Yu, Yanbo Fan, Yong Zhang, Xuan Wang, Fei Yin, Yunpeng Bai, Yan-Pei Cao, Ying Shan, Yang Wu, Zhongqian Sun, et al. 2023. NOFA: NeRF-based One-shot Facial Avatar Reconstruction. In ACM SIGGRAPH 2023 Conference Proceedings..., 1–12.
- Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018. Computer Vision Foundation / IEEE Computer Society, ., 586–595.
- Mingwu Zheng, Haiyu Zhang, Hongyu Yang, and Di Huang. 2023b. NeuFace: Realistic 3D Neural Face Rendering from Multi-View Images. In IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023. IEEE, ., 16868–16877.
- Yufeng Zheng, Wang Yifan, Gordon Wetzstein, Michael J. Black, and Otmar Hilliges. 2023a. PointAvatar: Deformable Point-Based Head Avatars from Videos. In IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023. IEEE, ., 21057–21067.
- Runsong Zhu, Di Kang, Ka-Hei Hui, Yue Qian, Xuefei Zhe, Zhen Dong, Linchao Bao, and Chi-Wing Fu. 2022. Semi-signed neural fitting for surface reconstruction from unoriented point clouds. arXiv preprint arXiv:2206.06715. (2022).
- D. Zimny, Tomasz Trzcinski, and Przemysław Spurek. 2022. Points2NeRF: Generating Neural Radiance Fields from 3D point cloud. *CoRR* abs/2206.01290 (2022), .
- Michael Zollhöfer, Justus Thies, Pablo Garrido, Derek Bradley, Thabo Beeler, Patrick Pérez, Marc Stamminger, Matthias Nießner, and Christian Theobalt. 2018. State of the Art on Monocular 3D Face Reconstruction, Tracking, and Applications. *Comput. Graph. Forum* 37, 2 (2018), 523–550.

Neural Point-based Volumetric Avatar

SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia



Figure 7: Qualitative comparison with NeRF and our failure case. (a) shows our qualitative comparison with NeRF. When trained on one single frame, our NPVA achieves results even visually better than NeRF (i.e., sharper and more realistic reflection effects), particularly in challenging facial regions (i.e., eyes and mouth interior) possibly due to the help of coarse geometry prior and the perceptual loss. (b) is one failure example, in which the regions far outside the rasterized mask contain obvious artifacts. This result is obtained by relaxing the constraint applied on the displacement map during training. Since no prior knowledge about the hair is provided in the coarse mesh, we replace our depth-guided shading point sampling strategy with the NeRF sampling strategy.



Figure 8: Comparisons on different training strategies (a) and different KNN search radius (b). The left part of (a) shows the photometric losses and validation MSE during training with different training strategies, while the right part of (a) shows qualitative rendering results. The model using our GEP training strategy converges to a lower photometric loss and achieves lower validation MSE, thereby producing more realistic facial reditions. As expected, the model using our GEP training strategy achieves a significant reduction of validation MSE during the E-stage. (b) shows that a larger search radius can eliminate holes in some extreme expressions with longer processing time.

SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia

Cong Wang, Di Kang, Yan-Pei Cao, Linchao Bao, Ying Shan, and Song-Hai Zhang



Figure 9: Ablation Study on Different Point Number and Displacements. When the number of points is low (i.e., 256 × 256), the model produces unsatisfactory results, even with holes present. As the point number grows to 1024×1024 , the model obtains sufficient points and fills holes, generating improved visual results. Increasing the point number to 2048×2048 results in perceptually similar facial renditions. At this time, introducing our displacements can further improve the model performance, generating more realistic results, particularly in challenging facial regions (i.e., eyes and mouth interior).



NPVA (Ours)

DAM [Lombardi et al. 2018]

PiCA [Ma et al. 2021]

MVP [Lombardi et al. 2021]

Figure 10: Qualitative Comparisons with State-of-the-Art Methods.