# CageNeRF: Cage-based Neural Radiance Field for Generalized 3D Deformation and Animation

**Yicong Peng**[1] **Yichao Yan**[2*] **Shenqi Liu**[2] **Yuhao Cheng**[2]
**Shanyan Guan**[2] **Bowen Pan**[3] **Guangtao Zhai**[1*] **Xiaokang Yang**[2]
[1] Institute of Image Communication and Network Engineering, Shanghai Jiao Tong University
[2] MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University
[3] Alibaba Group
{jack-sparrow,yichaoyan,lsqlsq,chengyuhao,shyanguan,zhaiguangtao,xkyang}@sjtu.edu.cn
bowen.pbw@alibaba-inc.com

## Abstract

While implicit representations have achieved high-fidelity results in 3D rendering, deforming and animating the implicit field remains challenging. Existing works typically leverage data-dependent models as deformation priors, such as SMPL for human body animation. However, this dependency on category-specific priors limits them to generalize to other objects. To solve this problem, we propose a novel framework for deforming and animating the neural radiance field learned on *arbitrary* objects. The key insight is that we introduce a cage-based representation as deformation prior, which is category-agnostic. Specifically, the deformation is performed based on an enclosing polygon mesh with sparsely defined vertices called *cage* inside the rendering space, where each point is projected into a novel position based on the barycentric interpolation of the deformed cage vertices. In this way, we transform the cage into a generalized constraint, which is able to deform and animate arbitrary target objects while preserving geometry details. Based on extensive experiments, we demonstrate the effectiveness of our framework in the task of geometry editing, object animation and deformation transfer. The code and supplementary materials are available at https://pengyicong.github.io/CageNeRF/.

## 1 Introduction

Editing and manipulating 3D object is a fundamental task for generating immerse contents, which enables numerous applications in VR/AR. Although recent advances in implicit neural rendering, such as neural radiance field (NeRF) [34], are able to produce photo-realistic modeling of static objects, how to edit the 3D content remains a challenging task. In this paper, we focus on learning a general deformation field for NeRF, enabling the deformation and animation of general objects.

Editing explicit 3D representations, such as point cloud and mesh, has been well studied. Point cloud and deformable mesh [7, 3, 17] store the geometry information in 3D space, which are directly editable and are widely utilized in computer graphics for modeling dynamic objects [35]. Cage-based deformation [24, 25, 30], for example, has been used to animate polygon mesh, where per-vertex deformation can be generated by translations prescribed on a coarse cage. While mesh and point cloud representations are relatively flexible to edit and animate, it is difficult to produce geometrically accurate and visually-faithful complex objects and scenes. Additional data such as normal maps and textures are usually required to achieve high-quality rendering results.

Recent works on implicit neural rendering, such as NeRF, employ multi-layer perceptron (MLP) as a universal approximator [19, 43], allowing them to approximate a continuous volumetric scene

*Corresponding authors

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

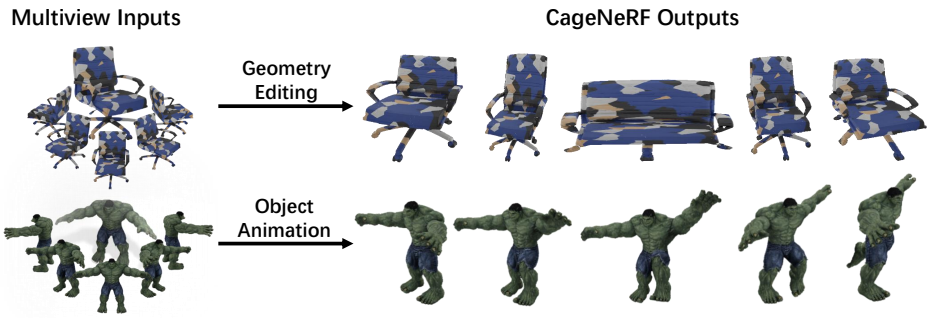


Figure 1: Given multiview images of an object, our method learns an implicit deformable 3D representation, which is used to deform and animate the object.

function to reconstruct high-fidelity results with only sparsely sampled images. Therefore, it is highly desirable if the neural rendering models can be also applied to model dynamic objects. However, optimizing a deformable radiance field end-to-end for dynamic objects like human body is a highly under-constrained problem. To address this issue, when modeling dynamic objects, most works optimize inside a deformation-invariant canonical space with data-dependent priors, such as 3DMM [4] for face deformation [1, 18], and SMPL [31] for human animation [40, 38]. While using data-dependent priors when modeling specific dynamic objects proved to be effective in providing stable correspondence between the object structure and the implicit field, it also limits the model's capability to generalize for broader applications.

To address this issue, we propose a novel framework which focuses on deforming and animating general objects represented by neural radiance fields, as shown in Figure 1. We take inspiration from the vanilla cage deforming approaches [24, 53] and extend them to work with implicit neural radiance fields. However, directly applying cage deformation into implicit field is non-trivial, due to the following issues. i) Traditional cage deformation methods work on polygon mesh with finite vertices and faces, while implicit radiance field is a continuous 3D volume. ii) Cage-based deformation transfer requires explicit geometry of the target, which is not directly accessible in the implicit field. iii) It is computationally expensive to densely sample and discretize the implicit field. To address these issues, we propose to estimate a coarse geometry of the target and generate a discrete field which supports bi-directional mapping between the canonical space and the deformed space, while keeping the radiance field continuous. Moreover, we further design a novel sampling strategy to reduce the computational cost while maintaining the render quality.

The advantages of our framework are three-fold. First, by utilizing category-agnostic cage-based deformation, our framework does not rely on data-dependent priors and can be generalized to work with a wide range of objects. Second, since cage-based deformation is a self-contained method, once the model is trained on the static scene, no additional training is required to apply deformation. Third, both cage deformation and volumetric rendering operate on the Euclidean space, we can design our framework in a decoupled manner as long as we keep the underlying coordinate system consistent between the two modules. This increases the flexibility and versatility of our framework. Moreover, as a well-developed method in computer graphics, cage-based deformation can be used in various tasks, which further expands the application scenario of our framework. We demonstrate the effectiveness of our framework in geometry editing, object animation, and deformation transfer.

## 2 Related Work

**Novel View Synthesis.** Traditional methods for novel view synthesis rely on explicit representations of 3D scenes. In these methods, 3D objects are modeled by meshes, point clouds [56, 44], etc, and then novel views are generated by shading and rendering. However, high-fidelity textures and well-designed 3D models are required for photo-realistic rendering, which makes the procedure highly expensive and inefficient. Recently, NeRF [34] has achieved impressive performance on novel view synthesis, where 3D objects are represented by implicit functions which map positions to densities and colors. A number of follow-up works have made further improvements to NeRF

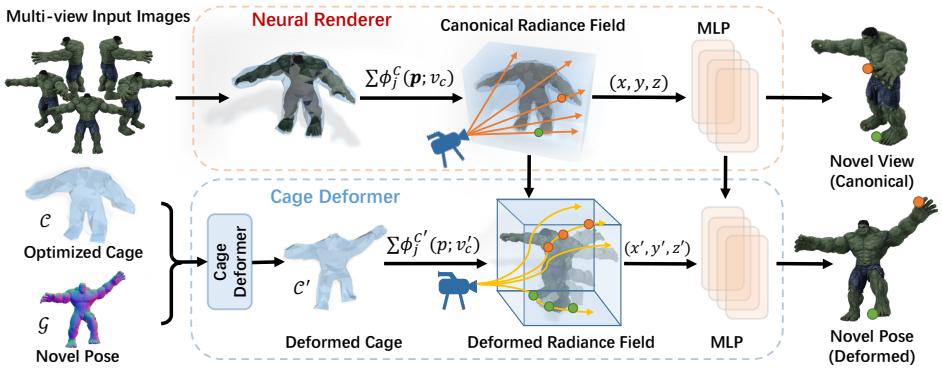


Figure 2: **Overview of CageNeRF.** CageNeRF consists of two modules: (1) a neural rendering module, which learns a radiance field from multi-view input images in the canonical space; (2) a cage deformer module that produces general cage deformation to manipulate the learned radiance field. CageNeRF deforms the implicit radiance field in an end-to-end manner.

to yield synthesis results with higher quality [59, 2, 58, 51]. Although these NeRF-based implicit methods are good at representing static scenes, it is still challenging to model dynamic scenes.

**Explicit Geometry Deformation.** Modeling deformable 3D objects is a fundamental task in computer graphics and computer vision, and triangle meshes are most commonly utilized due to their flexibility in manipulation. Traditional methods achieve this by deforming a fixed set of vertex points and then interpolating values among them [26, 52], such as skinning on skeletons [21, 33] or cages-based methods [25, 46, 62]. Along with the advances of large-scale datasets [6, 8, 57], learning-based methods are proposed by analyzing the prior principle of mesh deformation [12, 53, 13, 23]. Wang et al. [53] proposed a differentiable cage-based deformation module to synthesize shape variations and perform deformation transfer. However, these deformation methods can be only applied to explicit models, e.g., meshes. It is expensive to generate high-fidelity textures for photo-realistic rendering, and these methods cannot be directly applied to implicit models.

**Implicit Geometry Deformation.** Compared to explicit representations that only discretely control the vertex of meshes, the deformation of implicit geometry is more challenging, since the space is represented with a continuous function. Recently, various approaches build different deformation fields on implicit representations for animation [36, 22, 48, 49]. However, the generalization abilities of these methods are limited with respect to novel poses and different objects. Referring to the widely concerned digital humans, a vast range of methods are proposed for the deformation of the implicit face [16, 63, 18] and body [40, 9, 42, 39]. However, these methods heavily rely on parametric models as priors, such as 3DMM [5, 29] and SMPL [31], which are difficult to transfer to non-humanoid objects. In this paper, we propose a cage-based implicit geometry deformation method, which utilizes a sparse neural cage as guidance to build the deformation field for NeRF. The proposed CageNeRF is a general algorithm which is applicable to a wide variety of objects.

## 3 Methodology

We propose a novel framework, dubbed **CageNeRF**, for editing and animating arbitrary 3D objects represented by a neural radiance field. Radiance fields can render photo-realistic content. However, it remains challenging to perform deformation on radiance fields, since it's non-trivial to learn correspondence between implicit representations. We address this challenge by employing the cage-based deformation, which offers an elegant way of manipulating points inside a continuous space through a predefined coarse triangle mesh called cage. As shown in Figure 2, our framework consists of two modules: the neural rendering module optimizes the radiance field of a given static object inside canonical space using multiview images with known camera parameters. Then the cage deformer module deforms the cage inside this canonical space and reconstructs the implicit radiance field by translating points to novel positions.

### 3.1 Canonical Space Neural Rendering

We first train our neural renderer module to acquire the NeRF volume [34], which acts as a reference canonical space. NeRF models the scene along with its geometry, lighting, and material as a continuous volume and approximates it as a 5D radiance field using an MLP network. Given a 3D position of a point $\mathbf{p} = (x, y, z)$, the network learns a mapping function of its density $\sigma$ and color $\mathbf{c} = (r, g, b)$ variations with respect to the view direction $\mathbf{d} = (\theta, \phi)$:

$$\mathbf{F}_{\mathbf{\Theta}} : (\gamma(\mathbf{p}), \gamma(\mathbf{d})) \rightarrow (\mathbf{c}, \sigma), \tag{1}$$

where $\gamma(\cdot)$ is a positional embedding function of sinusoidal basis and $\Theta$ is the trainable parameters. Then, volumetric rendering is performed to synthesize images by marching camera rays $\mathbf{r} = \mathbf{o} + t\mathbf{d}$ from viewpoint $\mathbf{o}$, with direction $\mathbf{d}$, through the radiance field. The pixel color $\hat{C}$ is estimated using discrete integration on points sampled along each ray:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^{N} \exp(-\sum_{j=1}^{i-1} \sigma_j \delta_j)(1 - \exp(-\sigma_i \delta_i))c_i, \tag{2}$$

where $\delta_i = t_{i+1} - t_i$ is the interval between adjacent sample points. We then optimize our neural render module inside this canonical space using $L_2$ loss calculated between the estimated color $\hat{C}(\mathbf{r})$ and ground truth pixel value $C(\mathbf{r})$ of the input images:

$$\mathcal{L} = \sum_{\mathbf{r} \in \mathcal{R}} \|\hat{C}(\mathbf{r}) - C(\mathbf{r})\|_2^2, \tag{3}$$

where $\mathcal{R}$ denotes the set of rays that contributes to sampled pixel.

### 3.2 Cage Optimization

To deform the neural radiance field, we need to establish a consistent correspondence between the canonical and the deformed fields. To achieve this goal, a cage, i.e., a coarse polygon mesh inside the canonical space is first optimized to be geometrically aligned and enclose the object. Then the deformation field can be interpreted as the deformation based on the structural constraint of this optimized cage. Similar to Neural Cages, we provide our network with predefined polygon meshes as the starting point of this optimization process.

The geometry of the object inside the canonical space is represented by a polygon mesh $\mathcal{G} = \{\mathbf{v}, \mathbf{f}\}$, where $\mathbf{v}$ and $\mathbf{f}$ represent the vertex and face, respectively. When an explicit surface is not available, $\mathcal{G}$ can be obtained by applying the marching cube algorithm [32] on a radiance field or SDF [51].

The template cage $\mathcal{C}$ is a polygon mesh containing $N$ vertices, which is optimized w.r.t. the mean value coordinates (MVC) of sampled vertices of $\mathcal{G}$. Since negative MVC occurs when 1) the point lies outside the cage and 2) the cage itself is self-overlapping or highly concave. By penalizing the negative values, this optimization process automatically drives the vertices of the template cage to enclose the object inside the canonical space:

$$\min_{\mathbf{v}} |\min(\phi_i(\mathbf{v}), 0)|^2, \mathbf{v} \in \mathcal{G}, i = 1, 2, \cdots, N, \tag{4}$$

where $\phi_i(\cdot)$ is the MVC weight function w.r.t. sampled vertex $\mathbf{v}$ from polygon mesh $\mathcal{G}$ and the $i$-th vertex on cage $\mathcal{C}$.

### 3.3 Cage-based Deformation

After the cage $\mathcal{C}$ is optimized w.r.t. an object, we can perform deformation based on the target novel pose/shape and the current cage shape. While it is possible to perform the cage-based deformation by explicitly editing the cage vertices, we take inspiration from Neural Cage [53] and design a cage deformer module, which generates a deformed cage based on a novel polygon mesh $\mathcal{G}_n$ given as the deformation target. This cage deformer module makes our framework applicable to tasks like deformation transfer and object reenactment. Similar to the network proposed in Neural Cages, this module takes an encoder-decoder structure with PointNet++ [41] as its backbone to automatically calculate the displacement prescribed on each vertex of the optimized cage $\mathcal{C}$, as shown in Figure 3.

We pretrain our cage deformer in two different datasets: 1) ShapeNet [8] for deformation transfer and geometry editing on general objects; 2) an extended version of Surreal dataset [14] for better generalization ability on novel pose synthesis and animation.

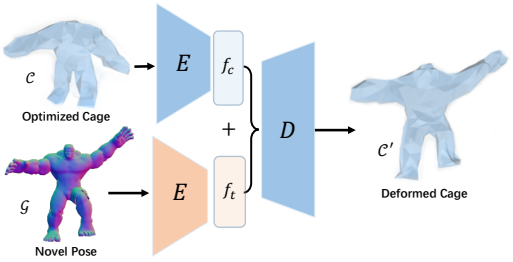

Figure 3: **Structure of the cage deformer.**

When performing the cage deformation, the encoder $\mathbf{E}$ extracts geometry features from the optimized cage $\mathcal{C}$ and the given deformation target $\mathcal{G}_t$, and outputs $f_c$ and $f_t$, respectively. The decoder $\mathbf{D}$ takes in the concatenated features and outputs the displacement $\Delta\mathbf{v}_i$ for each vertex on cage $\mathcal{C}$:

$$\Delta\mathbf{v} = \mathbf{D}(f_c \oplus f_t), \Delta\mathbf{v} \in \mathbb{R}^{N\times 3}, \tag{5}$$

where $\oplus$ refers to vector concatenation and $\Delta\mathbf{v}$ is the per-vertex displacement w.r.t cage $\mathcal{C}$. Since the cage deformation only changes vertex position while preserving its topology, the deformed cage $\mathcal{C}^{'}$ can be obtained by adding the displacement $\Delta\mathbf{v}$:

$$\mathcal{C}^{'} = \{\mathbf{v} + \Delta\mathbf{v}, \mathbf{f}\}, \mathbf{v}, \mathbf{f} \in \mathcal{C}. \tag{6}$$

The deformed cage $\mathcal{C}^{'}$ is later used to generate an MVC field for radiance field editing.

### 3.4 Radiance Field Editing

Once the deformed cage $\mathcal{C}^{'}$ is acquired, we can edit the radiance field with the following two steps. 1) Deforming the original geometry $\mathcal{G}$ acquired in Section 3.2 into a novel mesh $\mathcal{G}^{'}$ using traditional cage-based deformation. This newly generated mesh is used as the 3D mask to filter out redundant sample points and speed up the rendering process. 2) Building a dense MVC field, which acts as the consistent correspondence during deformation, based on the deformed cage $\mathcal{C}^{'}$ through grid sampling.

**Deformed mesh as 3D Mask.** The deformed mesh $\mathcal{G}^{'}$ can be obtained by performing cage-based deformation directly on the extracted mesh $\mathcal{G}$ of the object inside the canonical space. For every vertex $\mathbf{v}_i \in \mathcal{G}$, we first express its position using MVC coordinates:

$$\mathbf{v}_i = \mathcal{C}(\mathbf{v}_i) = \sum_{j=1}^{N} \phi_j^{\mathcal{C}}(\mathbf{v_i})\mathbf{v}_j, \mathbf{v}_i \in \mathcal{G}, \mathbf{v}_j \in \mathcal{C}, \tag{7}$$

where $\phi_j^{\mathcal{C}}(\cdot)$ is a series of weight functions depending on the relative position between the cage vertices $\mathbf{v}_j$ and the given mesh vertex $\mathbf{v}_i$. The detailed derivation of the 3D barycentric interpolation is given in Appendix A. Then, the vertices of the deformed mesh $\mathcal{G}^{'}$ can be obtained by:

$$\mathbf{v}_i^{'} = \sum_{j=1}^{N} \phi_j^{\mathcal{C}}(\mathbf{v}_i)\mathbf{v'}_j, \mathbf{v}_i^{'} \in \mathcal{G}^{'}, \mathbf{v'}_j \in \mathcal{C}^{'}, \tag{8}$$

where every vertex $\mathbf{v}_i^{'}$ on $\mathcal{G}^{'}$ has its novel position determined by the deformed cage vertices $\mathbf{v}_j^{'} \in \mathcal{C}^{'}$ as its new basis.

Since our radiance field is sparsely occupied by the render target, in order to increase the rendering speed while maintaining visual quality, we use the deformed mesh $\mathcal{G}^{'}$ as a 3D mask, where only the points within the vicinity threshold $\epsilon$ of the mesh surface are sampled during volumetric rendering.

Moreover, this explicit mesh enables us to robustly establish the backward deformation, i.e., from the deformed space to canonical space, when canonical information is not present in datasets like Human 3.6M [20]. See appendix for a detailed discussion.

**Anti-aliasing MVC Sampling.** As shown in Equations 7,8, while the coordinate representation of point changed in Euclidean space, the MVC representation remains invariant. We extend this

property on meshes into 3D volumes, which yields our editable radiance field. By utilizing this invariant correspondence, we are able to effectively edit the radiance field via volumetric cage-based deformation:

$$\mathcal{F}_d = \{\mathbf{p} = (\phi_1^{\mathcal{C}}, \phi_2^{\mathcal{C}}, \cdots, \phi_N^{\mathcal{C}}) | \mathbf{p} \in \mathcal{F}\}, \tag{9}$$

where every point $\mathbf{p}$ in this deformable radiance field $\mathcal{F}_d$ has its MVC coordinates calculated under a given cage $\mathcal{C}$. $\mathcal{F}$ is the canonical radiance field learned by our neural renderer. And the deformation in this neural volume can be carried out by replacing the original cage $\mathcal{C}$ with a novel one $\mathcal{C}^{'}$.

Since the MVC values remain invariant once the cage is fixed, to reduce computation cost, we build a pre-sampled MVC field based on the deformed mesh $\mathcal{G}^{'}$, as is shown in Figure 4. Specifically, we apply a grid sampling strategy to the deformed space, and the MVC of the sampled grid point $\mathbf{p}_g$ is equal to the barycentric coordinates of the closest point on the deformed mesh $\mathbf{p}_n \in \mathcal{G}^{'}$ to $\mathbf{p}_g$.

While this sampling strategy is efficient, it causes the aliasing issue as shown in Figure 4. The points $\{\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2\}$ are sampled grid points, since they all project to the same closest point $\mathbf{p}_n$ on the deformed mesh $\mathcal{G}^{'}$, i.e., they share the same MVC value. When rendering the deformed object, sampled points $\{\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2\}$ are projected to the same position inside the canonical space and thus have the same color and density value. Aliased sampled points such as $\mathbf{p}_2$ have the same color and density as $\mathbf{p}_n$ instead of the correct transparent color. This causes the "expanding" effect in the deformation space, e.g., the deformed object in the figure will be rendered as a circle with radius $r + \epsilon$, which should be $r$ instead. One naive solution would be directly reducing the vicinity threshold $\epsilon$. This approach imposes a higher demand on the accuracy of both the implicit representation and the extracted polygon mesh, which is non-trivial. To mitigate this issue while maintaining a relatively large $\epsilon$, we additionally define a shell $\mathcal{G}_s$, which is also a polygon mesh that closely bounds the surface of $\mathcal{G}^{'}$ by a small gap. This shell can be obtained by translating vertices on $\mathcal{G}^{'}$ along their surface normal while maintaining the topology. When assigning MVC value to grid points, we compare its distance to both $\mathcal{G}_s$ and $\mathcal{G}^{'}$ and determine the MVC value according to:

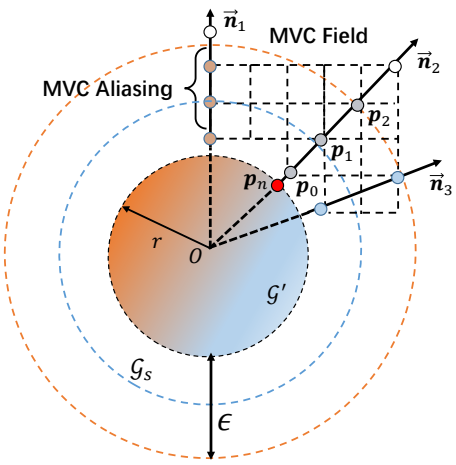


Figure 4: **Illustration of anti-aliasing MVC sample strategy in 2D case.** Grid points are projected onto the target geometry w.r.t. the surface normal $\vec{n}_i$. Given a sample range $\epsilon$, Points $\mathbf{p}_{0,1,2}$ among $\epsilon$ have the same nearest point $\mathbf{p}_n$ on deformed mesh $\mathcal{G}^{'}$ and thus share the same MVC value. This causes aliasing when rendering the deformed radiance field. Note that Grid points outside $\epsilon$ are filtered out. Considering this problem, we devise a shell $\mathcal{G}_s$ (dotted blue circle) which closely bounds the render geometry by a small margin. Then, we assign the MVC value by comparing the distance from the grid point to both the deformed mesh $\mathcal{G}^{'}$ and the shell $\mathcal{G}_s$ to negate this effect.

$$\phi_{grid} = \begin{cases} \phi_g, & \|\mathbf{p}_{grid} - \mathbf{p}_g\|_2 \leq \|\mathbf{p}_{grid} - \mathbf{p}_s\|_2 \\ \phi_s, & \|\mathbf{p}_{grid} - \mathbf{p}_g\|_2 > \|\mathbf{p}_{grid} - \mathbf{p}_s\|_2 \end{cases}, \tag{10}$$

where $\phi_{grid}$ is the MVC of the sampled grid point, $\phi_g$ is the MVC of the closest point $\mathbf{p}_g$ on mesh $\mathcal{G}^{'}$, $\phi_s$ is the MVC of the closest point $\mathbf{p}_s$ on shell $\mathcal{G}_s$. When rendering the deformed neural volume, points such as $\mathbf{p}_{1,2}$ with MVC of $\phi_s$ are projected outside the object in canonical space and thus have correct transparent color.

**Deformed Rendering.** After acquiring this MVC field, we can render the deformed radiance field through volumetric accumulation. 1) Sample points along the camera ray in render space using the deformed mesh $\mathcal{G}^{'}$ as a 3D mask. 2) Calculate the MVC value of filtered points using the MVC field via trilinear interpolation based on their coordinates. 3) Apply the cage deformation using the sampled MVC value to project points back to canonical space. 4) Infer the color and density of sampled points inside canonical space and apply volumetric rendering:

$$\mathbf{F}_\theta : (\gamma(\mathcal{C}(\mathbf{p})), \gamma(\mathbf{d})) \rightarrow (\mathbf{c}, \sigma), \tag{11}$$

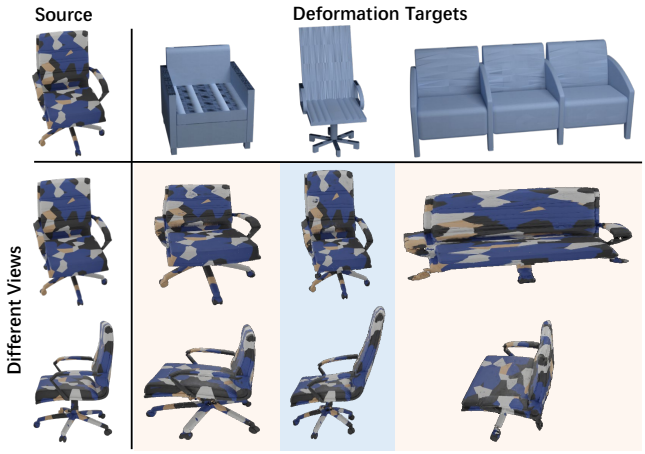


Figure 5: **Results of geometry editing via deformation transfer.** Top row shows the source object with texture and deformation target geometries. Row 2-3: deformed model under different views.

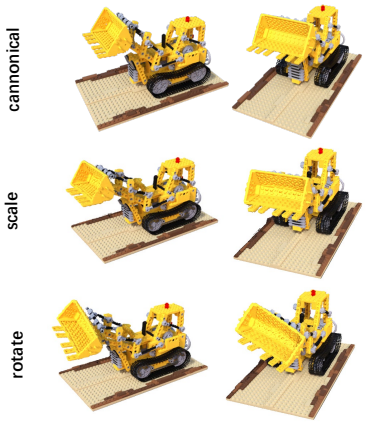


Figure 6: **Results of explicit geometry editing.** Top row shows the original object rendered in canonical space. Row 2-3: deformed model under different views.

Table 1: Animation Performance on Synthetic Datasets

| Dataset | Hulk | | | | Whale | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Frame Index | 1 | 27 | 47 | Avg. | 1 | 35 | 80 | 100 | Avg. |
| PSNR ↑ | 36.434 | 35.327 | 34.078 | 35.466 | 36.146 | 35.723 | 35.755 | 35.085 | 35.742 |
| SSIM ↑ | 0.963 | 0.957 | 0.948 | 0.959 | 0.972 | 0.972 | 0.971 | 0.966 | 0.972 |
| LPIPS ↓ | 0.0126 | 0.0156 | 0.0188 | 0.0153 | 0.0117 | 0.0121 | 0.0119 | 0.0126 | 0.0120 |

where $\mathcal{C}(\cdot)$ is the cage deformation that translates sampled points back to canonical space.

## 4 Experiments

To verify the generalization of cage deformation, we evaluate CageNeRF on three tasks:

**Geometry editing (Section 4.1).** In this task, we demonstrate the capability of CageNeRF in editing the geometry of the implicit object by deforming the radiance field. This is achieved by either explicitly editing the cage shape or using the cage deformer module to generate novel cage via deformation transfer. Deformation transfer is a task that takes two objects with different geometry properties (e.g., shape, structure and scale), and modifies these properties on the source to make it resemble the target. In this task, CageNeRF is trained on the lego dataset used in NeRF [34] and a chair dataset generated in the same manner.

**Neural animation (Section 4.2).** In this task, we evaluate the continuous deformation quality of CageNeRF by comparing the results with targets animated using Linear Blend Skinning [27]. We use both synthesized objects and human videos as our benchmarks. The synthesized 3D objects contain two animated characters: 'Hulk' and a blue whale. Furthermore, we evaluate CageNeRF on real humans using Human 3.6M [20].

**Pose reenactment (Section 4.3).** In this task, we generate a synthetic dataset using a static model of a humanoid robot, which only contains polygon mesh and texture without rigged animation. This dataset contains the rendered 360° multi-view images with paired camera parameters. We randomly sample posing SMPL models from the Surreal dataset [50] as the reenactment target.

We further provide the ablation study w.r.t. the MVC sampling strategy. Following the majority of previous literature [34, 38, 40, 51], we use PSNR/SSIM [54] and LPIPS [60] as evaluation metrics. We refer readers to Appendix C for more details and results.

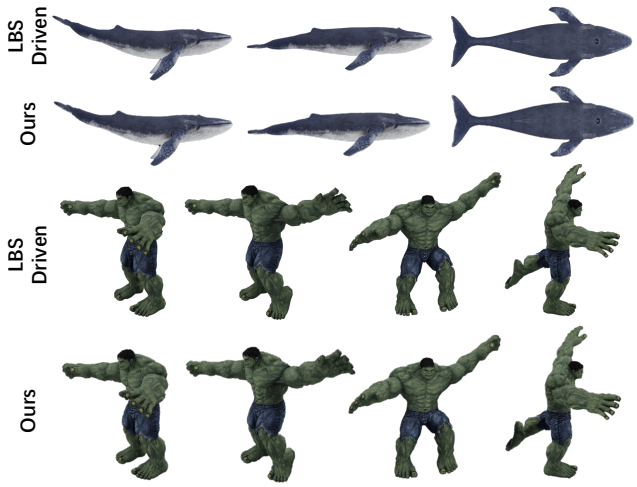


Figure 7: Qualitative results of neural animation on Whale and Hulk datasets. **Row 1, 3**: driven by linear blend skinning. **Row 2, 4**: driven by our framework.

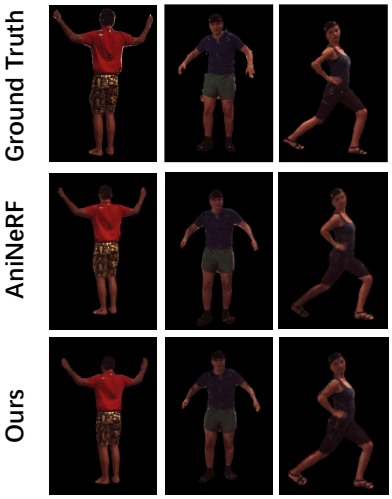


Figure 8: Qualitative comparison of CageNeRF and AniNeRF [38] for human synthesis on Human 3.6M.

### 4.1 Geometry Editing

We demonstrate the effectiveness of deforming raidance field via explicit cage editing on the lego dataset as shown in Figure 6. We edit the cage geometry to achieve the effect of stretching and rotating the bucket. In our framework, deformation transfer can be applied to edit the geometric properties of the implicit 3D objects. Our framework is first trained on an office chair in the canonical space, as a deformation source. We randomly select objects from the ShapeNet dataset [8] as the deformation targets and then apply our cage deformer to edit the implicit geometry stored in the canonical radiance field. As shown in Figure 5, our framework is capable of editing the implicit geometry of object in the canonical space by deforming the radiance field learned by our neural renderer module. The deformed results resemble the selected target in aspects of size, shape, and scale. These results demonstrate that our framework can edit the radiance field such that the deformed geometry resembles the target.

### 4.2 Neural Animation

**Synthesized Object Animation.** To quantitatively determine the animation precision, we analyze the rendering quality of an animated sequence and compare it to the traditional linear blend skinning (LBS) animation method which operates on the explicit deformable mesh. We train our neural renderer using multiview images of the object rendered at the starting frame. To animate the object, we extract the animated mesh from the animation sequence and use it as the deformed geometry $\mathcal{G}^{'}$. Our cage deformer then builds the MVC field using Equation 10. We animate our object by applying the extracted meshes $\mathcal{G}^{'}$ sequentially to our framework.

As is shown in Figure 7, our framework is capable of deforming the radiance field and generating continuous animation of objects. CageNeRF also achieves nice render quality which is comparable to that driven by explicit animation methods. As shown in Table 1, our framework achieves promising results in animating the neural radiance field when compared to explicit mesh deforming methods.

**Human Animation.** We also evaluate on a real-world dataset (i.e., Human 3.6M), and compare with the state-of-the-art method AniNeRF [38] as well as other human sythesis methods [47, 55]. The Human 3.6M dataset is captured by 4 synchronized cameras and has 2D and 3D joint annotations. We first estimate the geometry proxy of the posing human using the SMPL model by optimizing the 3D joint error and 2D re-projection joint error. Then, we utilize the geometry proxy as the "novel pose" and apply the operation in Figure 3 to generate a deformed cage with the matching pose. Then we build the MVC field using this deformed cage and posing geometry proxy. Since our model only stores the information inside the canonical space, CageNeRF can generalize to novel poses without

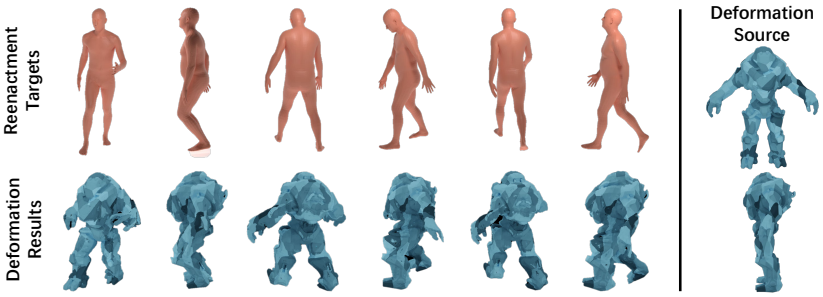


Figure 9: **Qualitative results of animation reenactment using SMPL models as the reenactment target. Top Row** shows SMPL models sampled from Surreal dataset as the reenactment target. **Bottom Row** reenactment results defromed from the souce model shows on the right.

Table 2: **Quantitative performance comparison on Human 3.6M dataset.** "NT", "NHR" and "AniNeRF" refer to Neural Textures, Neural Human Rendering and Animatable NeRF respectively.

| | PSNR | | | | SSIM | | | |
|---|---|---|---|---|---|---|---|---|
| | NT [47] | NHR [55] | AniNeRF [38] | Ours | NT | NHR | AniNeRF | Ours |
| S5 | 19.87 | 20.64 | 23.27 | 23.01 | 0.855 | 0.872 | 0.892 | 0.877 |
| S7 | 20.47 | 20.29 | 22.50 | 22.07 | 0.856 | 0.868 | 0.890 | 0.873 |
| S9 | 22.96 | 23.04 | 24.72 | 23.58 | 0.873 | 0.879 | 0.908 | 0.886 |
| S11 | 21.71 | 21.91 | 24.55 | 24.19 | 0.859 | 0.871 | 0.902 | 0.891 |
| Avg. | 21.25 | 21.47 | 23.76 | 23.21 | 0.861 | 0.873 | 0.898 | 0.882 |

additional training. This is in contrast to AniNeRF, which needs to optimize an additional blend weight network using the estimated geometry in order to synthesize novel pose results.

As shown in Figure 8 and Table 2, the performance of our framework in synthesizing dynamic human body is comparable to AniNeRF across different subjects. CageNeRF achieves slightly lower PSNR due to the inaccurate estimation of the geometry used in our cage deformer module. Moreover, we employ SMPL model as the geometry proxy during cage deformation, where SMPL only models human body, while the performer in Human 3.6M dataset wears loosely fitted clothes which has complex nonrigid deformations during motion. The inaccuracy in the geometry proxy further induces sampling displacement during volumetric rendering and thus impacts the performance of our framework. Despite the fact that our framework only takes a coarse estimation of the object, it is still able to correctly deform and synthesize complex dynamic objects such as the human body.

### 4.3 Pose Reenactment

To analyze the generalization ability of CageNeRF, we apply our framework to tasks of pose reenactment and motion retargeting from humans to humanoid objects. We keep the same training process as Section 4.1. Afterward, we randomly sample posed SMPL models from the Surreal dataset as the target for deformation transfer. Both the posed SMPL model and the optimized robot cage are fed into our cage deformer to calculate the cage deformation. Finally, the barycentric interpolation is applied to deform the static robot model using Equation 8. As is shown in Figure 9, although the structure of the source robot is notably different from the target human, the deformed pose and geometric structure correspond well to the given target, while keeping the source identity and texture.

### 4.4 Impact of the MVC Sampling Strategy

To evaluate the effectiveness of our anti-aliasing MVC sampling strategy, we generate an MVC field using different sampling strategies and qualitatively and quantitatively compare the rendering results. As shown in Figure 10, the rendering result generated with grid-sampled MVC field displays inflated geometry and blurred texture, due to the inaccurate mapping between the canonical and the deformed space. In contrast, the result rendered using our anti-aliasing sampled MVC field has notably sharper

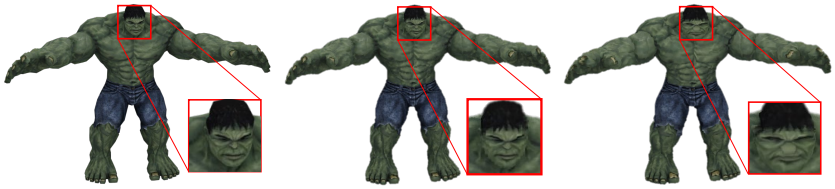


Figure 10: **Comparison between Anti-aliasing Sampling and Grid Sampling.** Our anti-aliasing sampling strategy achieves notably better rendering results w.r.t. geometry and texture.

Table 3: **Quantitative results of animated hulk using different MVC sampling strategies.** AA Sampling: the anti-aliasing sampling strategy.

| Metrics | Grid Sampling | AA Sampling |
|---|---|---|
| PSNR | 30.695 | 35.280 |
| SSIM | 0.916 | 0.956 |
| LPIPS | 0.0293 | 0.0156 |

textures and a more accurate geometry structure. As also shown in Table 3, the anti-aliasing sampling strategy significantly increases the render quality on all three evaluation metrics.

## 5 Limitations

Our approach achieves high quality results in deforming and animating implicit objects by combining cage-based deformation with neural radiance field. However, there remains a few limitations of our approach we would like to address. 1) While the cage-based deformation is category-agnostic and easy to generalize under various settings, it is not suited for modeling complex and subtle deformations such as facial expressions. Our work is a proof of the concept that explicit deformation can be applied to implicit neural representations. We believe our framework is also compatible with other explicit deformation schemes, which may provide the extra capability in adapting to complex movements. 2) Since the color and density of the canonical radiance field are relatively stable, our framework is not suited for modeling dynamic objects with changing color or rendering objects under varying lighting conditions. Recent works [45, 61] have developed novel techniques on factorizing basic color and density information into several decoupling elements such as surface normal, albedo, and light visibility, this further enables material editing and neural scene relighting. It would be interesting to further improve and augment our framework to have editable material, color, and lighting by incorporating novel neural rendering algorithms and structures.

## 6 Conclusion

We propose a versatile framework for deforming and animating implicit field. Comparing to previous methods, our framework excels in aspects of generalization, structural flexibility, and editing capability. This is achieved by adopting a deformable cage as the universal constraint for our radiance field editing which does not require any data-dependent prior knowledge. We demonstrate the effectiveness of our framework in the tasks of object animation, geometry editing, and deformation transfer. Experimental results show that our framework is capable of animating and deforming arbitrary objects represented by implicit field, and produces results comparable to other state-of-the-art methods.

## Acknowledgments and Disclosure of Funding

This work is supported by NSFC (62225112, 62101325, 61831015), the National Key Research and Development Program of China (2021YFE0206700), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), and CCF-Alibaba Innovative Research Fund For Young Scholars.

## References


[1] ShahRukh Athar, Zhixin Shu, and Dimitris Samaras. Flame-in-nerf : Neural control of radiance fields for free view face animation. *CoRR*, abs/2108.04913, 2021.

[2] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *ICCV*, pages 5835–5844, 2021.

[3] Matthew Berger, Andrea Tagliasacchi, Lee M. Seversky, Pierre Alliez, Gaël Guennebaud, Joshua A. Levine, Andrei Sharf, and Cláudio T. Silva. A survey of surface reconstruction from point clouds. *Comput. Graph. Forum*, 36(1):301–329, 2017.

[4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH*, SIGGRAPH '99, page 187–194, 1999.

[5] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH*, pages 187–194, 1999.

[6] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: dataset and evaluation for 3d mesh registration. In *CVPR*, pages 3794–3801, 2014.

[7] Mario Botsch, Leif Kobbelt, Mark Pauly, Pierre Alliez, and Bruno Lévy. *Polygon Mesh Processing*. 2010.

[8] Angel X. Chang, Thomas A. Funkhouser, Leonidas J. Guibas, Pat Hanrahan, Qi-Xing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. Shapenet: An information-rich 3d model repository. *CoRR*, abs/1512.03012, 2015.

[9] Xu Chen, Yufeng Zheng, Michael J. Black, Otmar Hilliges, and Andreas Geiger. SNARF: differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, pages 11574–11584, 2021.

[10] Michael S. Floater. Mean value coordinates. *CAGD*, 20(1):19–27, 2003.

[11] Michael S. Floater, Géza Kós, and Martin Reimers. Mean value coordinates in 3d. *CAGD*, 22(7):623–631, 2005.

[12] Lin Gao, Yu-Kun Lai, Dun Liang, Shu-Yu Chen, and Shihong Xia. Efficient and flexible deformation representation for data-driven surface modeling. *ACM TOG*, 35(5):158:1–158:17, 2016.

[13] Lin Gao, Yu-Kun Lai, Jie Yang, Ling-Xiao Zhang, Shihong Xia, and Leif Kobbelt. Sparse data driven mesh deformation. *IEEE TVCG*, 27(3):2085–2100, 2021.

[14] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. 3d-coded: 3d correspondences by deep deformation. In *ECCV*, volume 11206, pages 235–251, 2018.

[15] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. A papier-mâché approach to learning 3d surface generation. In *CVPR*, pages 216–224, 2018.

[16] Yudong Guo, Keyu Chen, Sen Liang, Yong-Jin Liu, Hujun Bao, and Juyong Zhang. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *ICCV*, pages 5784–5794, 2021.

[17] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennamoun. Deep learning for 3d point clouds: A survey. *IEEE TPAMI*, 43(12):4338–4364, 2021.

[18] Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. Headnerf: A real-time nerf-based parametric head model. *CoRR*, abs/2112.05637, 2021.

[19] Kurt Hornik, Maxwell B. Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.

[20] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE TPAMI*, 36(7):1325–1339, 2014.

[21] Alec Jacobson, Ilya Baran, Ladislav Kavan, Jovan Popovic, and Olga Sorkine. Fast automatic skinning transformations. *ACM TOG*, 31(4):77:1–77:10, 2012.

[22] Wonbong Jang and Lourdes Agapito. Codenerf: Disentangled neural radiance fields for object categories. In *ICCV*, pages 12929–12938, 2021.

[23] Timothy Jeruzalski, David I. W. Levin, Alec Jacobson, Paul Lalonde, Mohammad Norouzi, and Andrea Tagliasacchi. Nilbs: Neural inverse linear blend skinning. *CoRR*, abs/2004.05980, 2020.

[24] Pushkar Joshi, Mark Meyer, Tony DeRose, Brian Green, and Tom Sanocki. Harmonic coordinates for character articulation. *ACM TOG*, 26(3):71, 2007.

[25] Tao Ju, Scott Schaefer, and Joe D. Warren. Mean value coordinates for closed triangular meshes. *ACM TOG*, 24(3):561–566, 2005.

[26] Ladislav Kavan, Steven Collins, Jirí Zára, and Carol O'Sullivan. Geometric skinning with approximate dual quaternion blending. *ACM TOG*, 27(4):105:1–105:23, 2008.

[27] Ladislav Kavan, Steven Collins, Jiří Žára, and Carol O'Sullivan. Geometric skinning with approximate dual quaternion blending. *ACM TOG*, 27(4):1–23, 2008.

[28] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

[29] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM TOG*, 36(6):194:1–194:17, 2017.

[30] Yaron Lipman, David Levin, and Daniel Cohen-Or. Green coordinates. *ACM TOG*, 27(3):78, 2008.

[31] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: a skinned multi-person linear model. *ACM TOG*, 34(6):248:1–248:16, 2015.

[32] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *SIGGRAPH*, pages 163–169, 1987.

[33] Bruce Merry, Patrick C. Marais, and James E. Gain. Animation space: A truly linear framework for character animation. *ACM TOG*, 25(4):1400–1423, 2006.

[34] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, pages 405–421, 2020.

[35] Ramakrishnan Mukundan. *3D Mesh Processing and Character Animation - With Examples Using OpenGL, OpenMesh and Assimp*. 2022.

[36] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B. Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, pages 5845–5854, 2021.

[37] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8024–8035, 2019.

[38] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *ICCV*, pages 14294–14303, 2021.

[39] Sida Peng, Shangzhan Zhang, Zhen Xu, Chen Geng, Boyi Jiang, Hujun Bao, and Xiaowei Zhou. Animatable neural implicit surfaces for creating avatars from videos. *CoRR*, abs/2203.08133, 2022.

[40] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, pages 9054–9063, 2021.

[41] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, pages 5099–5108, 2017.

[42] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. In *CVPR*, pages 2886–2897, 2021.

[43] Franco Scarselli and Ah Chung Tsoi. Universal approximation using feedforward neural networks: A survey of some existing methods, and some new results. *Neural Networks*, 11(1):15–37, 1998.

[44] Daeyun Shin, Charless C Fowlkes, and Derek Hoiem. Pixels, voxels, and views: A study of shape representations for single view 3d object shape prediction. In *CVPR*, pages 3061–3069, 2018.

[45] Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *CVPR*, pages 7495–7504, 2021.

[46] Jean-Marc Thiery, Julien Tierny, and Tamy Boubekeur. Jacobians and hessians of mean value coordinates for closed triangular meshes. *CoRR*, abs/1109.1914, 2011.

[47] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: image synthesis using neural textures. *ACM TOG*, 38(4):66:1–66:12, 2019.

[48] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *ICCV*, pages 12959–12970, 2021.

[49] Wei-Cheng Tseng, Hung-Ju Liao, Lin Yen-Chen, and Min Sun. Cla-nerf: Category-level articulated neural radiance field. *CoRR*, abs/2202.00181, 2022.

[50] Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *CVPR*, pages 4627–4635, 2017.

[51] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *NIPS*, pages 27171–27183, 2021.

[52] Xiaohuan Corina Wang and Cary B. Phillips. Multi-weight enveloping: least-squares approximation techniques for skin animation. In *SIGGRAPH*, pages 129–138, 2002.

[53] Yifan Wang, Noam Aigerman, Vladimir G. Kim, Siddhartha Chaudhuri, and Olga Sorkine-Hornung. Neural cages for detail-preserving 3d deformations. In *CVPR*, pages 72–80, 2020.

[54] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004.

[55] Minye Wu, Yuehao Wang, Qiang Hu, and Jingyi Yu. Multi-view neural human rendering. In *CVPR*, pages 1679–1688, 2020.

[56] Kai Xu, Vladimir G. Kim, Qixing Huang, Niloy J. Mitra, and Evangelos Kalogerakis. Data-driven shape analysis and processing. In *SIGGRAPH Asia*, pages 4:1–4:38, 2016.

[57] Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. Facescape: A large-scale high quality 3d face dataset and detailed riggable 3d face prediction. In *CVPR*, pages 598–607, 2020.

[58] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *NeurIPS*, 34, 2021.

[59] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *CoRR*, abs/2010.07492, 2020.

[60] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018.

[61] Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul E. Debevec, William T. Freeman, and Jonathan T. Barron. Nerfactor: neural factorization of shape and reflectance under an unknown illumination. *ACM TOG*, 40(6):237:1–237:18, 2021.

[62] Yuzhe Zhang, Jianmin Zheng, and Yiyu Cai. Proxy-driven free-form deformation by topology-adjustable control lattice. *Comput. Graph.*, 89:167–177, 2020.

[63] Yufeng Zheng, Victoria Fernández Abrevaya, Xu Chen, Marcel C Bühler, Michael J Black, and Otmar Hilliges. IM avatar: Implicit morphable head avatars from videos. *CoRR*, abs/2112.07471, 2021.

## Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] , our work focuses on providing editing method for implicit 3D representations in a generalized manner, it is possible that someone utilizes our method to produce fake videos.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A] We do not include theoretical results.
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We describe the implementation details in the supplementary material. The data utilized is publicly available, and we will make our code publicly available.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We discuss the training details in supplementary material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See supplementary for detailed descriptions.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We use existing datasets and cited them in our paper.
   (b) Did you mention the license of the assets? [No]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] We discuss this in the supplementary material.

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] The data used in our work does not contain such information.

5. If you used crowdsourcing or conducted research with human subjects...

   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [No] Our work does not involve crowdsourcing as such, same as follows.

   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [No]

   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [No]

# A Derivation of Mean Value Coordinate

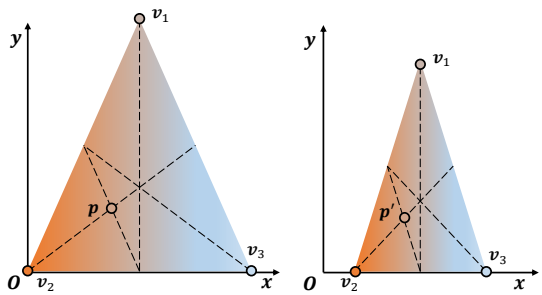


Figure 11: **MVC and cage deformation in 2D case.** MVC defined by basis vertices $\mathbf{v}_{1,2,3}$ are visualized through color. By translating the basis vertices to novel positions point $\mathbf{p}$ is shifted to $\mathbf{p}'$ while maintaining its MVC.

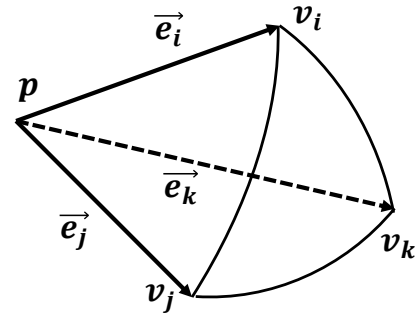


Figure 12: **MVC extended to 3D scenario via unit sphere projection.** $\mathbf{v}_{i,j,k}$ are projected points on the unit sphere with center $\mathbf{p}$ from $\hat{\mathbf{v}}_{i,j,k}$ of closed polygon mesh $\Omega$, and $\vec{\mathbf{e}}_{i,j,k}$ are unit normals with origin $\mathbf{p}$.

Mean value coordinates (MVC) were first introduced in [10] as a way of representing points using convex combinations of vertices from a polygon. It was later extended to work with 3D closed meshes [11, 25]. Mean value coordinates encode the position of a given point with respect to a set of predefined vertices. The relationship between the MVC of a given point and the predefined vertices is analogous to that between a vector and its corresponding basis in linear algebra, i.e., the predefined vertices act as the basis in MVC space. Figure 11 illustrates the cage-based deformation using MVC coordinates in the 2D scenario. Given the predefined vertices $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$, which act as the basis in the MVC space, we can express arbitrary points inside the enclosed triangle as linear combinations of these vertices. For example, the midpoint of line segment $\overline{\mathbf{v}_1\mathbf{v}_2}$ can be expressed as $\frac{1}{2} \cdot \mathbf{v}_1 + \frac{1}{2} \cdot \mathbf{v}_2$ and hence having the MVC of $(0.5, 0.5, 0)$. In Figure 11, we visualize the MVC of points inside the triangle by using different colors, and provide the Cartesian plane with origin $\Omega$ as the reference coordinate system. Compared to the left, basis vertices are translated to novel positions. All points inside the enclosed triangle are shifted and thus have different coordinates in the Cartesian coordinate system. Notice that while the position of an arbitrary point $\mathbf{p}$ is shifted to $\mathbf{p}'$ after the deformation, its MVC remain unchanged, i.e., $\mathbf{p}$ and $\mathbf{p}'$ share the same color.

In general, given an arbitrary point $\mathbf{p}$ inside the convex kernel $\mathbf{K}$ of a polygon $\Omega$ with $n$ vertices, we can calculate the MVC by constructing a series of non-negative functions $\phi_1, \phi_2, \cdots, \phi_n : \mathbf{K} \to \mathbb{R}$, such that:

$$\sum_{i=1}^{n} \phi_i(\mathbf{p}) = 1 \quad \text{and} \quad \sum_{i=1}^{n} \phi(\mathbf{p})\mathbf{v}_i = \mathbf{p}, \tag{12}$$

where $\mathbf{v}_i$ is the $i$-th ordered vertex of $\Omega$. The MVC is defined by:

$$\phi_i = \frac{w_i}{\sum_{j=1}^{n} w_j}, \quad w_i = \frac{1}{r_i}(\tan\frac{\alpha_i}{2} + \tan\frac{\alpha_{i-1}}{2}), \tag{13}$$

where $\alpha_i$ is the angle $\angle \mathbf{v}_i\mathbf{p}\mathbf{v}_{i+1}$, and $r_i(\mathbf{p}) = \|\mathbf{v}_i - \mathbf{p}\|$. To extend this case to the 3D scenario, given a closed polygon mesh $\Omega$, its vertices $\hat{\mathbf{v}}_i \in \Omega$ are first projected to a unit sphere with $\mathbf{p}$ at its center. As shown in Figure 12, $\{\vec{\mathbf{e}}_i, \vec{\mathbf{e}}_j, \vec{\mathbf{e}}_k\}$ are unit vectors from $\mathbf{p}$ to the projected vertices $\{\mathbf{v}_i, \mathbf{v}_j, \mathbf{v}_k\}$ on the unit sphere. The MVC of the point $\mathbf{p}$ takes the same form as Equation 12, except that $\phi_i$ is defined by:

$$\phi_i = \frac{w_i}{\sum_{j=1}^{n} w_j}, \quad w_i = \frac{1}{r_i} \sum_{\mathbf{v}_i \in \mathbf{T}} \mu_{i,\mathbf{T}}, \tag{14}$$

where $T$ is the spherical triangle enclosed by $\{\mathbf{v}_i, \mathbf{v}_j, \mathbf{v}_k\}$ and $\mu_{i,\mathbf{T}}$ is given by:

$$\mu_{i,\mathbf{T}} = \frac{\beta_{jk} + \beta_{ij}\vec{\mathbf{n}}_{ij} \cdot \vec{\mathbf{n}}_{jk} + \beta_{ki}\vec{\mathbf{n}}_{ki} \cdot \vec{\mathbf{n}}_{jk}}{2\vec{\mathbf{e}}_i \cdot \vec{\mathbf{n}}_{jk}}, \quad \text{and} \quad \beta_{rs} = \angle \mathbf{v}_r\mathbf{p}\mathbf{v}_s \tag{15}$$

$$\vec{\mathbf{n}}_{rs} = \frac{\vec{\mathbf{e}}_r \times \vec{\mathbf{e}}_s}{\|\vec{\mathbf{e}}_r \times \vec{\mathbf{e}}_s\|}. \tag{16}$$

# B Datasets

We train our framework on different datasets and evaluate CageNeRF on several tasks. In the following section, we provide a detailed description of the employed datasets. We will make the synthetic data publicly available.

For the task of **Geometric Editing via Deformation Transfer**, we use a synthetic dataset of an office chair, which act as the static object inside the canonical space, to train our neural renderer. We also employ the sampled objects from ShapeNet [8] as deformation targets and comparison baseline to evaluate the editing results. Our synthetic dataset of the office chair is created from a polygon mesh with colored texture using 3D render software. The dataset contains around 1.2K 360° inward-facing multiview images with corresponding masks and camera parameters. Some samples of the multiview images and masks are shown in Figure 13. As for the camera parameters, we store the camera rotation $\mathbf{R} \in \mathbb{R}^{3\times3}$, camera translation $\mathbf{T} \in \mathbb{R}^{1\times3}$, and the camera intrinsic parameters $\mathbf{K} \in \mathbb{R}^{3\times3}$. Since we use the same camera for the entire render process, the intrinsic parameters are shared across all the images, while the rotation and translation vary with the rendering view.

Figure 13: **Samples from synthetic dataset.** We render the inward-facing 360° multiview images along with corresponding masks. The masks are used during ray sampling on the training image to improve sample efficiency.

In the experiment of **Neural Animation**, we use two types of datasets: 1) a synthetic dataset of animated objects driven by linear blend skinning, and 2) a real-world dataset, i.e., Human3.6M [20], which is often utilized as a benchmark for human body synthesis. The synthetic datasets contain two characters, one contains a comic character Hulk, and the other contains a swimming blue whale. The deformable meshes of both objects are fully rigged with corresponding skeletons and blend weights. Both objects are capable of producing short animation clips using liner blend skinning as driving method. Similar to the synthetic dataset used in **Geometric Editing**, we render 1.2K 360° inward-facing multiview images of the animatable objects at the starting frame. The deformed polygon meshes of the following frames are also extracted as the animation target for our cage deformer module to generate new deformation fields. The Human3.6M dataset was captured by 4 synchronized inward-facing cameras from different viewpoints. Every section of this dataset captures a short clip of a posing performer with 2D and 3D joint annotations and per-frame masks. For each frame in this clip, a SMPL [31] human model is estimated based on the annotations and masks for its position, rotation, and joint locations. This estimated SMPL model, similar to the animated mesh in our synthetic dataset, is used in our cage deformer to generate deformations on the radiance field. We provide additional results on synthesizing human body on different subjects in Figure 14.

Figure 14: **Additional results in comparison with AniNeRF.** We render various poses of different subjects in Human 3.6M dataset and achieve comparable results with the state-of-the-art human synthesis method. Note that our framework does not require additional training or linear blend skinning information when rendering novel poses.

As for **Pose Reenactment**, we also create a synthetic dataset in the same manner as the office chair dataset used in **Geometry Editing**. We use a static robot model with colored texture, instead of an office chair, to create this dataset. We sample our reenactment targets from the Surreal [14] dataset.

# C Implementation Details

## C.1 Training Configuration

We train our network on a single Nvidia RTX3090 GPU for 500 epochs, depending on the complexity of the geometry and texture. The overall training process takes around 22 hours for a dataset with 1K multi-view images. We take ADAM optimizer [28] with an exponential learning rate scheduler. We set our initial learning rate to $5e-4$ and our learning rate decays by a factor of 10 for every 500 epochs. We implement our entire framework using PYTORCH [37] deep learning library. It roughly takes 12G GPU memory to train our framework. Our framework achieves the average test-time rendering speed of 2.071 seconds per frame with the resolution of $800 \times 800$.

## C.2 Network Structure

Different from the original NeRF [34], which optimizes two radiance fields, i.e., a coarse and a fine model, the neural render in our framework only optimizes a single multi-layer perceptron (MLP) network. We sample 128 points along each camera ray during volumetric rendering. We also provide a trainable latent code to compensate for the texture and normal variances of the object inside the canonical space. The structure of the neural renderer is shown in Figure 15. Our neural renderer has a similar structure to the original NeRF network except for the additional latent code input. We take inspiration from [53, 15] and design our cage deformer in a similar encoder-decoder manner. We use the PointNet++ [41] with multi-scale grouping strategy to extract the feature from 3D objects in our encoder. For the decoder, we adopt an MLP network to predict the vertex-wise offset for the input cage base on the extracted feature.

# D Additional Experiment Analysis

## D.1 Impact of the Network Structure

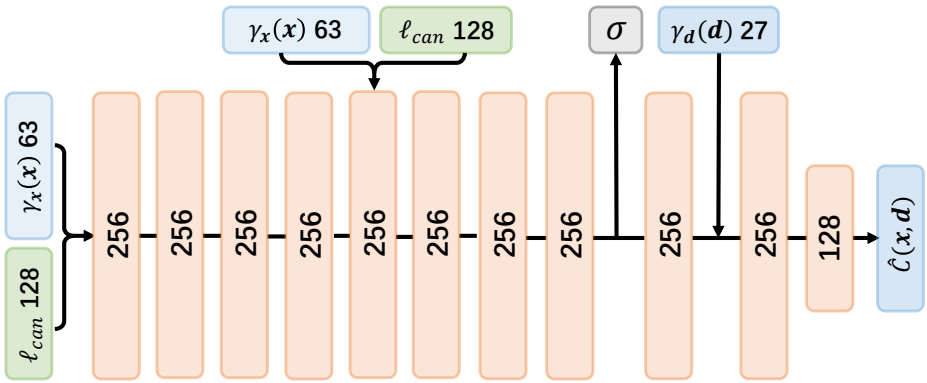


Figure 15: **Network structure of the Neural Renderer module.** Our neural renderer adopts a structure similar to the vanilla NeRF save for the additional latent code input. Numbers inside the colored block indicate the dimension of the input or the linear layer. $\gamma(\cdot)$ is the embedding function. The neural renderer takes in concatenated input of embedded position $\mathbf{x}$ and latent code $l_{can}$ and outputs density $\sigma$ and estimated color $\hat{C}$.

As mentioned in our paper, we design our framework in a decoupled manner, the underlying structure of our neural renderer module, which learns the implicit representation of a static object in canonical space, can be replaced with other designs without affecting the integrity of our framework. To validate this point, we prepare two neural renderers with different structures and carry out experiments for the task of neural animation. Our neural renderer adopts a modified NeRF structure, as described in Section C.2, with additional latent code to represent the appearance of the object inside the canonical space, while the other uses the vanilla version of NeRF. Table 4 shows the impact of different network structures on the rendering quality across three datasets. When using the vanilla NeRF structure as the neural renderer, the rendering quality slightly drops w.r.t. PSNR. The introduced additional latent code is effective in improving the rendering quality for both the synthetic and real-world datasets.

Table 4: **Impact of network structure in render quality measured in PSNR.**

| | Hulk | Whale | H36M |
|---|---|---|---|
| **NeRF** | 34.672 | 35.154 | 22.605 |
| **Ours** | 35.280 | 35.857 | 23.583 |

### D.2 Impact of Mesh Accuracy

Our cage-based deformation method utilizes a geometry proxy extracted from the learned radiance field as a 3D mask. The quality of this geometry proxy affects the point sampling process during volumetric rendering. To examine its influence on our framework, we apply deformation transfer to our canonical radiance field using the MVC field generated from the ground truth mesh and the mesh reconstructed by the signed distance function [51]. As shown in Figure 16, while the reconstructed mesh has minor artifacts in terms of smoothness and geometric detail, the overall shape and structure match the ground truth. The artifacts and dilation in certain areas such as the armrest causes blurred and distorted texture. This result shows the accuracy of geometry proxy has a noticeable effect in rendering quality. In the future, one possible direction to further improve the flexibility of CageNeRF is to learn geometry directly from multi-view images.

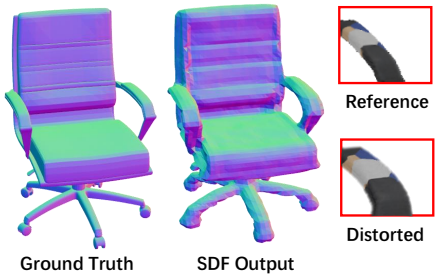


Figure 16: **Comparison between the ground truth and mesh extracted using SDF.**

### D.3 Functionality of Explicit Mesh

The explicit mesh extracted in our method has the following functionalities:

**Enhance Deformation Capability.** While our framework supports explicit deformation by manually editing the shape of the canonical cage, we want to enhance the deformation capability further. Neural Cage [53] offers impressive performance with explicit mesh in tasks like deformation transfer and pose reenactment. By incorporating an explicit mesh into our framework, we are able to transfer the capability of Neural Cage, i.e., a mesh-based method, to neural radiance field.

**Increase Sampling Efficiency.** As mentioned in our paper Section 3.4, the canonical radiance field is sparsely occupied by the render target. Restricting the sample points inside the vicinity space of the render target can effectively reduce the sample points down to $10\%$ of the original number required in volume rendering.

**Provide Robust Correspondence.** Our framework handles deformation by first learning a static radiance field inside the canonical space. However, this canonical space is not always directly accessible. For example, the Human 3.6M dataset only contains posing human subjects captured by multi-view cameras, the canonical space, i.e., "T-pose" images of the target is not given. This requires first estimating the deformed cages of the targets in various poses and performing cage-based deformation backward to transform posing targets to canonical "T-pose". An explicit mesh can be used to generate correct MVC field which enables us to robustly establish this backward deformation. Cage-based deformation can express the forward deformation correctly, i.e., from the canonical space to the deformed space, but not always vice versa. In some cases, both Neural Cages and our cage deformer module will output a deformed cage with overlapping faces or collapsed vertices. Naively taking the MVC value calculated in such deformed cages will induce artifacts, since points that reside in such regions cannot be correctly projected into the canonical space. This property can be interpreted as: CBD using a cage with a self-overlapping structure is similar to applying a matrix that is not full rank, and hence not invertible. We provide backward deformation results on SMPL model using cage with and without self-overlapping structures in. Additionally, if the cage vertices are manually edited without overlapping faces and collapsed vertices, we can generate MVC field directly based on this deformed cage, i.e., without the need of an explicit mesh, since this cage deformation is invertible from deformed space to canonical space.

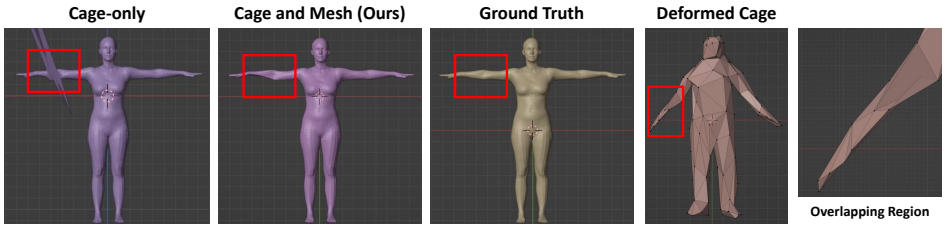


Figure 17: **Comparison of inverse cage-based deformation w/o an explicit mesh.** We perform a backward deformation a SMPL model using a self-overlapping cage. The "A-pose" cage resides in the deformed space, and the "T-pose" SMPL is the object in canonical space. As is shown in the top row, backward deformation using only the cage will produce artifacts such as spikes and outlier vertices, while using an explicit mesh can avoid this artifact. The difference between our backward deformation result and the ground truth can be compensated by the sampling vicinity $\epsilon$.

## E Broader Impacts

As shown in our experiments, our work is capable of producing high-fidelity results of edited and animated 3D objects with 2D images as its input. Our framework may find application in 3D animation, digital human creation, high fidelity 3D object editing. Especially, the deformation method used in our framework is category-agnostic, our framework can be applied in cross-domain animation retargeting and reenactment.