

NeuS2: Fast Learning of Neural Implicit Surfaces for Multi-view Reconstruction

Yiming Wang*

University of Pennsylvania, Peking University
 wym12416@pku.edu.cn

Qin Han*

Peking University
 hanqin@pku.edu.cn

Marc Habermann

Max Planck Institute for Informatics
 mhaberma@mpi-inf.mpg.de

Kostas Daniilidis

University of Pennsylvania
 kostas@cis.upenn.edu

Christian Theobalt

Max Planck Institute for Informatics
 theobalt@mpi-inf.mpg.de

Lingjie Liu

University of Pennsylvania, Max Planck Institute for Informatics
 lingjie.liu@seas.upenn.edu

Abstract

Recent methods for neural surface representation and rendering, for example NeuS [59], have demonstrated the remarkably high-quality reconstruction of static scenes. However, the training of NeuS takes an extremely long time (8 hours), which makes it almost impossible to apply them to dynamic scenes with thousands of frames. We propose a fast neural surface reconstruction approach, called NeuS2, which achieves two orders of magnitude improvement in terms of acceleration without compromising reconstruction quality. To accelerate the training process, we parameterize a neural surface representation by multi-resolution hash encodings and present a novel lightweight calculation of second-order derivatives tailored to our networks to leverage CUDA parallelism, achieving a factor two speed up. To further stabilize and expedite training, a progressive learning strategy is proposed to optimize multi-resolution hash encodings from coarse to fine. We extend our method for fast training of dynamic scenes, with a proposed incremental training strategy and a novel global transformation prediction component, which allow our method to handle challenging long sequences with large movements and deformations. Our experiments on various datasets demonstrate that NeuS2 significantly outperforms the state-of-the-arts in both surface reconstruction accuracy and training speed for both static and dynamic scenes. The code is available at our website: <https://vc.ai.mpi-inf.mpg.de/projects/NeuS2/>.

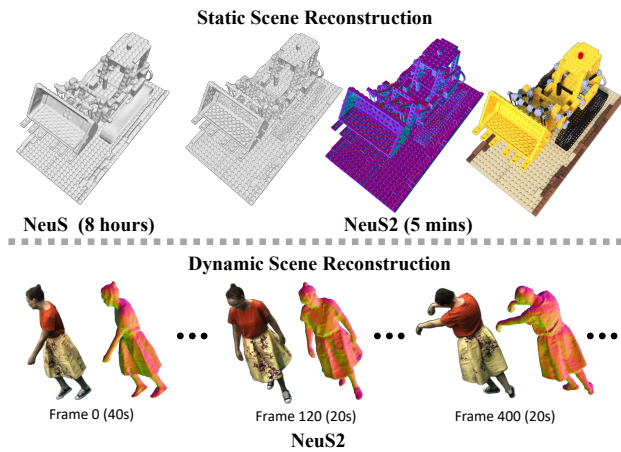


Figure 1. We present NeuS2, a fast neural scene reconstruction method. Given a set of multi-view images, NeuS2 can accurately reconstruct the scene geometry and appearance in the order of minutes. This is in stark contrast to previous work [59], which only recovers medium-scale details at significantly increased time (about 8 hours). Moreover, we demonstrate that NeuS2 can also be applied to dynamic scene reconstruction from multi-view videos, where we recover per-frame reconstruction in about 20 seconds.

1. Introduction

Reconstructing the dynamic 3D world from 2D images is crucial for many Computer Vision and Graphics applications, such as AR/VR, 3D movies, games, telepresence, and 3D printing. Classical stereo algorithms employ computer vision methods, e.g. feature matching, to capture the geometry and appearance of 3D contents from multi-view 2D images. Despite great progress, these methods are still comparably slow and struggle to reconstruct high-quality results.

* Equal contribution

Recently, 3D reconstruction with neural implicit representations has become a promising alternative to traditional methods, because of its high spatial resolution and highly detailed reconstructions outperforming classical stereo algorithms [4, 12, 13, 50, 10, 5, 51]. NeuS [59], as a representative work, models geometry surfaces as a neural network encoded Signed Distance Field (SDF), and renders an image via differentiable volume rendering. While NeuS [59] produces high-quality reconstruction results, its training process is extremely slow, i.e. about 8 hours for a static object. This makes it nearly impossible to reconstruct dynamic scenes. Instant-NGP [36] has explored the training acceleration of neural radiance fields (NeRF) [35] by utilizing multi-resolution hash tables to augment neural network-encoded radiance fields. Though Instant-NGP [36] synthesizes impressive novel view synthesis results, the extracted geometry from the learned density fields contains discernible noise since the 3D representation lacks surface constraints.

To overcome these drawbacks, we propose *NeuS2*, a new method for fast training of highly-detailed neural implicit surfaces (see Fig. 1). *NeuS2* reconstructs a static object in minutes, and a moving object sequence in up to 20 seconds per frame. To achieve this, we parameterize the neural network-encoded SDF using multi-resolution hash tables of learnable feature vectors [36]. Notably, in this design, a surface constraint and the rendering formulation require calculating the second-order derivatives. The main challenge is to have a simple and memory-efficient calculation to achieve the highest possible GPU computing performance. Therefore, we derive a simple formula of the second-order derivatives tailored to ReLU-based MLPs, which enables an efficient CUDA implementation with a small memory footprint at a significantly reduced computational cost. To further enforce and accelerate the training convergence, we introduce an efficient progressive training strategy, which updates the hash table features in a coarse-to-fine manner.

We further extend our method to multi-view dynamic scene reconstruction. Instead of training each frame in the sequence separately, we propose a new incremental learning strategy to efficiently learn a neural dynamic representation of objects with large movements and deformations. Specifically, we exploit the similarity of the shape and appearance information shared in two consecutive frames by first training the first frame and sequentially fine-tuning the subsequent frames. While generally, this strategy works well, we observe that when the movement between two consecutive frames is relatively large, the predicted SDF of the occluded regions that are not observed in most images may get stuck in the learned SDF of the previous frame. To address this, we predict a global transformation to roughly align these two frames before learning the representation of the new frame. In summary, our technical contributions are:

- We propose a new method, *NeuS2*, for fast learning of neural surface representations from multi-view RGB input for both static and dynamic scenes, which achieves a significant speed up over the state-of-the-art while achieving an unprecedented reconstruction quality.
- A simple formulation of the second-order derivatives tailored to ReLU-based MLPs is presented to enable efficient parallelization of GPU computation.
- A progressive training strategy for learning multi-resolution hash encodings from coarse to fine is proposed to enforce better and faster training convergence.
- We design an incremental learning method with a novel global transformation prediction component for reconstructing long sequences (e.g., 2000 frames) with large movements in an efficient and stable manner.

2. Related Work

Multi-view Stereo. Traditional multi-view 3D reconstruction methods can be categorized into depth-based and voxel-based methods. Depth-based methods [4, 12, 13, 50] reconstruct a point cloud by identifying point correspondences across images. However, the reconstruction quality is heavily affected by the accuracy of correspondence matching. Voxel-based methods [10, 5, 51] side-step the difficulties in explicit correspondence matching by recovering occupancy and color in a voxel grid from multi-view images with a photometric consistency criterion. However, the reconstruction of these methods is limited to low resolution due to the high memory consumption of voxel grids.

Classical Multi-view 4D Reconstruction. A large body of work [6, 57, 2, 9, 64, 16] in multi-view 4D reconstruction utilizes a precomputed deformable model, which is then fit to the multi-view images. In contrast, our method does not rely on a precomputed model, can reconstruct detailed results, and handles topology changes. The most relevant to our work is [8], which is also a model-free method. They leverage RGB and depth inputs to reconstruct high-quality point clouds for each frame, and then produce temporally coherent geometry. Instead, we only require RGB as input and can learn the high-quality geometry and appearance for each frame in an end-to-end manner in 20 seconds per frame.

Neural Implicit Representations. Neural implicit representations have made remarkable achievements in novel view synthesis [53, 33, 22, 35, 30, 52] and 3D/4D reconstruction [47, 48, 66, 37, 23, 20, 32, 39, 59, 65]. NeRF [35] has shown high-quality results in the novel view synthesis task, but it cannot extract high-quality surfaces since the geometry representation lacks surface constraints. NeuS [59]

represents the 3D surface as an SDF for high-quality geometry reconstruction. However, the training of NeuS is very slow, and it only works for static scenes. Instead, our method is 100 times faster and can be further accelerated to 20 seconds per frame when applied to dynamic scene reconstruction.

Some NeRF-based works [49, 55, 36] introduce voxel-grid features to represent 3D properties for fast training. However, these methods cannot extract high-quality surfaces as they inherit the volume density field as the geometry representation from NeRF [35]. In contrast, our method can achieve high-quality surface reconstruction as well as fast training. For dynamic scene modeling, many works [56, 40, 45, 41, 28, 62, 26] propose to disentangle a 4D scene into a shared canonical space and a deformable field per frame.

[11] represents a 4D scene with a time-aware voxel feature. [29] proposes a static-to-dynamic learning paradigm for fast dynamic scene learning. [25] presents a grid based method for efficiently reconstructing radiance fields frame by frame. [58] presents a novel Fourier PlenOctree method to compress a dynamic scene into one model. These four methods focus on the novel view synthesis and, thus, are not designed to reconstruct high-quality surfaces, which is different from our goal of achieving high-quality surface geometry and appearance models. While these works make the training for dynamic scenes more efficient, the training is still time-consuming. Furthermore, these methods are not able to handle large movements and only reconstruct medium-quality surfaces. Some works in human performance modeling [54, 31, 43, 7, 63, 38, 19, 44, 24, 60] can model large movements by introducing a deformable template as a prior. In contrast, our method can handle large movements, does not require a deformable template and, thus, is not restricted to a specific dynamic object. Moreover, we can learn high-quality surfaces of dynamic scenes for 20 seconds per frame. [68] proposes a method for human modeling and rendering. It first reconstructs a neural surface representation for each frame; then it applies non-rigid deformation to obtain a temporally coherent mesh sequence. Our work focuses on the first part, that is, fast reconstruction of dynamic scenes, where we exploit the temporal consistency between two consecutive frames to accelerate the learning of dynamic representation. Therefore, our work is orthogonal to [68], and can be integrated into [68] as the first step.

Concurrent Work. Voxurf [61] proposes a voxel-based surface representation for fast multi-view 3D reconstruction. While it enables 20x speedup over the baseline (i.e., NeuS [59]), our proposed method is over 3x faster than Voxurf and achieves better geometry quality when compared with Voxurf’s results reported in their paper, as shown in the Suppl. document. Neuralangelo [27] presents a novel




Figure 2. Qualitative comparison on DTU Scan 110, concerning the quality impact of smoothing post-process on Instant-NGP and the Eikonal Penalty on NeuS2.

method that leverages multi-resolution hash grids with numerical gradient computation for neural surface reconstruction. It can achieve dense and high-fidelity geometry reconstruction results for large-scale scenes from multi-view images with multiple delicate designs while sacrificing its training cost, 100x slower when compared to ours. Also, Voxurf and Neuralangelo are not designed for dynamic scene reconstruction. Last, Unbiased4d [21] proposes a monocular dynamic surface reconstruction approach by extending the NeuS formulation for bending rays. In stark contrast to our approach, their focus lies on proving that unbiasedness also holds in the case of ray bending and the challenging monocular setting, and less on the highest possible quality at the fastest speeds.

3. Background

NeuS. Given calibrated multi-view images of a static scene, NeuS [59] implicitly represents the surface and appearance of a scene as a signed distance field $f(\mathbf{x}) : \mathbb{R}^3 \rightarrow \mathbb{R}$ and a radiance field $c(\mathbf{x}, \mathbf{v}) : \mathbb{R}^3 \times \mathbb{S}^2 \rightarrow \mathbb{R}^3$, where \mathbf{x} denotes a 3D position and $\mathbf{v} \in \mathbb{S}^2$ is a viewing direction. The surface \mathcal{S} of the object can be obtained by extracting the zero-level set of the SDF $\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^3 | f(\mathbf{x}) = 0\}$. To render an object into an image, NeuS leverages volume rendering. Specifically, for each pixel of an image, we sample n points $\{\mathbf{p}(t_i) = \mathbf{o} + t_i \mathbf{v} | i = 0, 1, \dots, n - 1\}$ along its camera ray, where \mathbf{o} is the center of the camera and \mathbf{v} is the view direction. By accumulating the SDF-based densities and colors of the sample points, we can compute the color \hat{C} of the ray. As the rendering process is differentiable, NeuS can learn the signed distance field f and the radiance field c from the multi-view images. However, the training process is very slow, taking about 8 hours on a single GPU.

InstantNGP. To overcome the slow training time of deep coordinate-based MLPs, which is also a main reason for the slow performance of NeuS, recently, Instant-NGP [36] proposed a multi-resolution hash encoding and has proven its effectiveness. Specifically, Instant-NGP assumes that the object to be reconstructed is bounded in multi-resolution voxel grids. The voxel grids at each resolution are mapped to a hash table with a fixed-size array of learnable feature vectors. For a 3D position $\mathbf{x} \in \mathbb{R}^3$, it obtains a hash encoding at each level $h^i(\mathbf{x}) \in \mathbb{R}^d$ (d is the dimension of a feature vector, $i = 1, \dots, L$) by interpolating the feature vectors

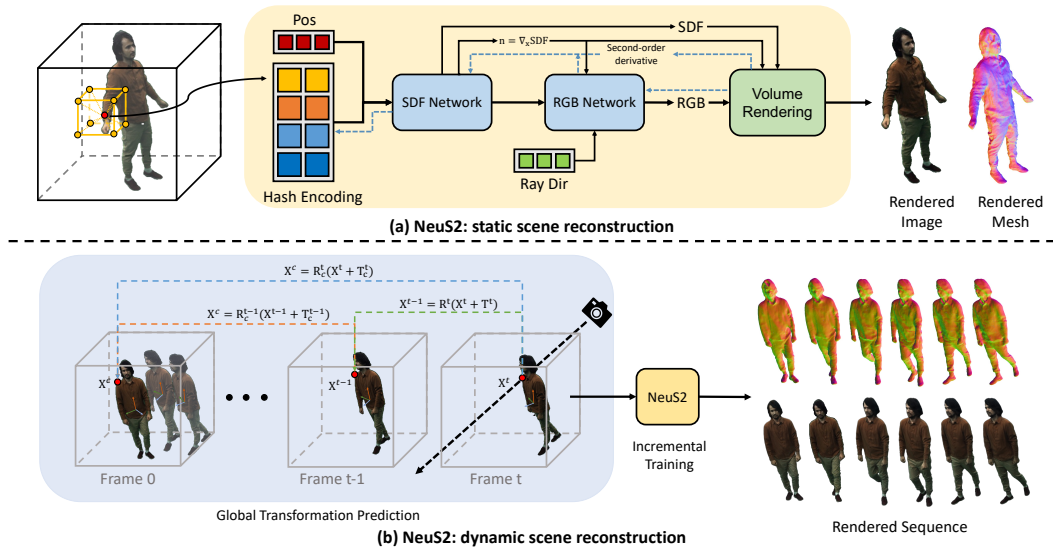


Figure 3. Method overview. (a) **Static Scene Reconstruction:** Given a 3D point \mathbf{x} , we concatenate its queried feature from the multi-resolution hash grid and its 3D position as the input to the SDF network. The SDF network outputs the SDF value and geometry features, which are combined with the viewing direction and further fed into our RGB network to generate the RGB value. Notably, during backward propagation the second-order derivatives are efficiently computed using our CUDA implementation. (b) **Dynamic Scene Reconstruction:** Given a sequence of multi-view images, we first construct the first frame following our static reconstruction method. For every subsequent frame, we predict its global transformation with respect to the previous frame and accumulate the transformation to convert it into the canonical space (i.e. the first frame). Then, we fine-tune the parameters of NeuS2 for incremental training to generate the rendered results.

assigned at the surrounding voxel grids at this level. The hash encodings at all L levels are then concatenated to be the multi-resolution hash encoding $h(\mathbf{x}) = \{h^i(\mathbf{x})\}_{i=1}^L \in \mathbb{R}^{L \times d}$. Besides the hash encoding, another key factor to the training acceleration is the CUDA implementation of the whole system, which makes use of GPU parallelism. While the runtime is significantly improved, Instant-NGP still does not reach the quality of NeuS in terms of geometry reconstruction accuracy.

Challenges. Given the above discussion, one can ask whether a naive combination of NeuS [59] and Instant-NGP [36] can unite the best of the two worlds, i.e. high 3D surface reconstruction quality and efficient computation. We highlight that *it is far from being trivial to achieve training as fast as Instant-NGP [36] and, meanwhile, a reconstruction as high-quality as NeuS [59]*. Specifically, to ensure high-quality surface learning, the Eikonal constraint used in NeuS [59] is indispensable, as shown in Fig. 2 and Suppl. materials; and the key challenge of adding the Eikonal loss to CUDA-based MLPs (a key factor to fast training in Instant-NGP [36]) is how to calculate the second-order derivatives efficiently for backpropagation. Instant-NSR [68] addresses this by approximating the second-order derivatives using finite differences, which suffer from precision problems and it can cause unstable training. Instead, we propose a simple, precise, and efficient formulation of second-order derivatives tailored to MLPs (Sec. 4.2), which leads to fast and high-quality reconstruction. The superiority of our approach over Instant-NSR is shown in Tab. 1 and

Fig. 4.

For dynamic scene reconstruction, there are two key challenges: how to exploit the temporal information for acceleration, and how to handle long sequences with large movements and deformations. To address them, we first propose an incremental training strategy to exploit the similarity in geometry and appearance information shared across two consecutive frames, which enabling faster convergence (Sec. 5.1). To handle large movements and deformations, we propose a novel Global Transformation Prediction component, which prevents the predicted SDF training from getting stuck in a local minimum. Further, it can bound a dynamic sequence in a small volume to save memory and to improve reconstruction accuracy (Sec. 5.2).

4. Static Neural Surface Reconstruction

We first present how our formulation can effectively learn the signed distance field of a *static* scene from calibrated multi-view images (see Fig. 3 a)). To accelerate the training process, we first demonstrate how to incorporate multi-resolution hash encodings [36] for representing the SDF of the scene, and how volume rendering can be applied to render the scene into an image (Sec. 4.1). Next, we derive a simplified expression of second-order derivatives tailored to ReLU-based MLPs, which can be efficiently parallelized in custom CUDA kernels (Sec. 4.2). Finally, we adopt a progressive training strategy for learning multi-resolution hash encodings, which leads to faster training convergence

and better reconstruction quality (Sec. 4.3).

4.1. Volume Rendering of a Hash-encoded SDF

For each 3D position \mathbf{x} , we map it to its multi-resolution hash encodings $h_\Omega(\mathbf{x})$ with learnable hash table entries Ω . As $h_\Omega(\mathbf{x})$ is an informative encoding of spatial position, the MLPs for mapping \mathbf{x} to its SDF d and color c can be very shallow, which results in more efficient rendering and training without compromising quality.

SDF Network. In more detail, our SDF network

$$(d, \mathbf{g}) = f_\Theta(\mathbf{e}), \quad \mathbf{e} = (\mathbf{x}, h_\Omega(\mathbf{x})). \quad (1)$$

is a shallow MLP with weights Θ , which takes the 3D position \mathbf{x} along with its hash encoding $h_\Omega(\mathbf{x})$ as input and outputs the SDF value d and a geometry feature vector $\mathbf{g} \in \mathbb{R}^{15}$. Concatenating the position serves as a geometry initialization [3] leading to a more stable learning of the geometry.

Color Network. The normal of \mathbf{x} can be computed as

$$\mathbf{n} = \nabla_{\mathbf{x}} d. \quad (2)$$

where $\nabla_{\mathbf{x}} d$ denotes the gradient of the SDF with respect to \mathbf{x} . We then combine the normal with the geometry feature \mathbf{g} , the SDF d , the point \mathbf{x} , and the ray direction \mathbf{v} serving as the input to our color network

$$\mathbf{c} = c_\Upsilon(\mathbf{x}, \mathbf{n}, \mathbf{v}, d, \mathbf{g}), \quad (3)$$

which predicts the color \mathbf{c} of \mathbf{x} .

Volume Rendering. To render an image, we apply the unbiased volume rendering of NeuS [59]. Additionally, we adopt a ray marching acceleration strategy used in Instant-NGP[36]. More details are provided in the supp. document.

Supervision. To train NeuS2, we minimize the color difference between the rendered pixels \hat{C}_i with $i \in \{1, \dots, m\}$ and the corresponding ground truth pixels C_i without any 3D supervision. Here, m denotes the batch size during training. We also employ an Eikonal term [14] to regularize the learned signed distance field leading to our final loss

$$\mathcal{L} = \mathcal{L}_{\text{color}} + \beta \mathcal{L}_{\text{eikonal}}, \quad (4)$$

where $\mathcal{L}_{\text{color}} = \frac{1}{m} \sum_i \mathcal{R}(\hat{C}_i, C_i)$, and \mathcal{R} is the Huber loss [17]. $\mathcal{L}_{\text{eikonal}} = \frac{1}{mn} \sum_{k,i} (|\|\mathbf{n}_{k,i}\| - 1|^2)$, and k indexes the k th sample along the ray with $k \in \{1, \dots, n\}$, and n is the number of sampled points. $\mathbf{n}_{k,i}$ is the normal of a sampled point (see Eq. 2).

4.2. Efficient Handling of Second-order Derivatives

To avoid the computational overhead that learning frameworks introduce, we implement our whole system in CUDA. In contrast to Instant-NGP [36], which only requires the first-order derivatives during the optimization, we

must calculate the second-order derivatives for the parameters associated with the normal term $\mathbf{n} = \nabla_{\mathbf{x}} d$ (Eq. 2), which is input to the color network c_Υ (Eq. 3) and the Eikonal loss term $\mathcal{L}_{\text{eikonal}}$.

Second-order Derivatives. To accelerate this computation, we directly calculate them using simplified formulas instead of applying the computation graph of PyTorch [42]. Specifically, we calculate the second-order derivatives of the hash table parameters Ω and the SDF network parameters Θ using the chain rule as

$$\frac{\partial \mathcal{L}}{\partial \Omega} = \frac{\partial \mathcal{L}}{\partial \mathbf{n}} \left(\frac{\partial \mathbf{e}}{\partial \mathbf{x}} \frac{\partial \frac{\partial d}{\partial \mathbf{e}}}{\partial \mathbf{e}} \frac{\partial \mathbf{e}}{\partial \Omega} + \frac{\partial d}{\partial \mathbf{e}} \frac{\partial \frac{\partial \mathbf{e}}{\partial \mathbf{x}}}{\partial \Omega} \right) \quad (5)$$

$$\frac{\partial \mathcal{L}}{\partial \Theta} = \frac{\partial \mathcal{L}}{\partial \mathbf{n}} \left(\frac{\partial \mathbf{e}}{\partial \mathbf{x}} \frac{\partial \frac{\partial d}{\partial \mathbf{e}}}{\partial \Theta} + \frac{\partial d}{\partial \mathbf{e}} \frac{\partial \frac{\partial \mathbf{e}}{\partial \mathbf{x}}}{\partial \Theta} \right) \quad (6)$$

Note that the color network c_Υ only takes \mathbf{n} as input, so we do not need to calculate its second-order gradients of the color network parameters Υ . The derivation of Eq. 5 and 6 is included in the supplementary document.

To speed up the computation of Eqs. 5 and 6, we found that ReLU-based MLPs can greatly simplify the above terms leading to less computation overhead. In the following, we discuss this in more detail and provide the proof of this proposition in the supplementary document. We first introduce some useful definitions as follows.

Definition 1 Given a ReLU based MLP f with L hidden layers taking $x \in \mathbb{R}^d$ as input, it computes the output $y = H_L g(H_{L-1} \dots g(H_1 x))$, where $H_l \in \mathbb{R}^{n_l \times \mathbb{R}^{n_{l-1}}}$, $l \in \{1, \dots, L\}$ is the layer index, and g is the ReLU function. We define $P_l^j \in \mathbb{R}^{n_{l-1} \times \mathbb{R}^1}$ and $S_l^i \in \mathbb{R}^1 \times \mathbb{R}^{n_l}$ as

$$P_l^j = G_l H_{l-1} \dots G_2 H_1^{(-j)} \quad (7)$$

$$S_l^i = H_L^{(i,-)} G_L \dots H_{l+1} G_{l+1}$$

where $H_1^{(-j)}$ is the j th column of H_1 , $H_L^{(i,-)}$ is the i th row of H_L , and $G_l = \begin{cases} 1, & H_{l-1} \dots g(H_1 x) > 0 \\ 0, & \text{otherwise} \end{cases}$.

Now the second order derivatives of a ReLU-based MLP with respect to its input and intermediate layers can be defined.

Theorem 1 (Second-order derivative of ReLU-based MLP)

Given a ReLU based MLP f with L hidden layers with the same definition in Definition 1. The second-order derivative of the MLP f is:

$$\frac{\partial \frac{\partial y}{\partial x^{(i,j)}}}{\partial H_l} = (P_l^j S_l^i)^T, \quad \frac{\partial^2 y}{\partial \mathbf{x}^2} = 0 \quad (8)$$

where $\frac{\partial y}{\partial x^{(i,j)}}$ is the matrix element (i,j) of $\frac{\partial y}{\partial \mathbf{x}}$, and S_l^i and P_l^j are defined in Definition 1.

Coming back to our original second-order derivatives (Eqs. 5 and 6), by Theorem 1, we obtain $\frac{\partial \frac{\partial d}{\partial \mathbf{e}}}{\partial \mathbf{e}} = 0$. Since $\frac{\partial \mathbf{e}}{\partial \mathbf{x}}$ is irrelevant to Θ , we have $\frac{\partial \frac{\partial \mathbf{e}}{\partial \mathbf{x}}}{\partial \Theta} = 0$. This results in the following simplified form

$$\frac{\partial \mathcal{L}}{\partial \Omega} = \frac{\partial \mathcal{L}}{\partial \mathbf{n}} \frac{\partial d}{\partial \mathbf{e}} \frac{\partial \frac{\partial \mathbf{e}}{\partial \mathbf{x}}}{\partial \Omega} \quad (9)$$

$$\frac{\partial \mathcal{L}}{\partial \Theta} = \frac{\partial \mathcal{L}}{\partial \mathbf{n}} \frac{\partial \mathbf{e}}{\partial \mathbf{x}} \frac{\partial \frac{\partial d}{\partial \mathbf{e}}}{\partial \Theta} \quad (10)$$

for the second-order derivatives, which result in improved efficiency and less computational overhead. The simplest form of $\frac{\partial \mathcal{L}}{\partial \Omega}$ and $\frac{\partial \mathcal{L}}{\partial \Theta}$ can be obtained by substituting the terms of $\frac{\partial \frac{\partial \mathbf{e}}{\partial \mathbf{x}}}{\partial \Omega}$ and $\frac{\partial \frac{\partial d}{\partial \mathbf{e}}}{\partial \Theta}$ using Eq. 8. We show in Fig. 9 that our computation using Eq. 8 is more efficient than PyTorch.

4.3. Progressive Training

Although our highly-optimized gradient calculation already improved training time, we found there is still room for improvement in terms of training convergence and speed. Additionally, we empirically observed that using grids at low resolutions results in underfitting of the data, in which the geometry reconstruction is smooth and lacks details, whereas using grids at high resolutions induces overfitting, leading to increased noise and artifacts of the results. Therefore, we introduce a progressive training method by gradually increasing the bandwidth of our spatial grid encoding denoted as:

$$h_{\Omega}(\mathbf{x}, \lambda) = (w_1(\lambda)h_{\Omega}^1(\mathbf{x}), \dots, w_L(\lambda)h_{\Omega}^L(\mathbf{x})), \quad (11)$$

where h_{Ω}^i is the hash encoding at level i and the weight w_i for each grid encoding level is defined by $w_i(\lambda) = I[i \leq \lambda]$, and the parameter λ modulates the bandwidth of the low-pass filter applied to the multi-resolution hash encodings. Smaller parameter λ leads to faster training speed, but limits the model’s capacity to model high-frequency details. Thus, we initialize λ as 2, and then gradually increase by 1 for every 2.5% of the total training steps in all experiments.

5. Dynamic Neural Surface Reconstruction

We have explained how NeuS2 can produce highly accurate and fast reconstructions of static scenes. Next, we extend NeuS2 to dynamic scene reconstruction. That is, given multi-view videos of a moving object and camera parameters of each view, our goal is to learn the neural implicit surfaces of the object in each video frame (see Fig. 3 b)).

5.1. Incremental Training

Even though our reconstruction method of static objects can achieve promising efficiency and quality, constructing dynamic scenes by training every single frame independently is still time-consuming. However, scene changes

from one frame to the other are typically small. Thus, we propose an incremental training strategy to exploit the similarity in geometry and appearance information shared between two consecutive frames, which enables faster convergence of our model. In detail, we train the first frame from scratch as presented in our static scene reconstruction, and then fine-tune the model parameters for the subsequent frames from the learned hash grid representation of the preceding frame. Using this strategy, the model is able to produce a good initialization of the neural representation of the target frame and, thus, significantly accelerates its convergence speed.

5.2. Global Transformation Prediction

As we observed during the incremental training process, the predicted SDF easily get stuck in the local minima of the learned SDF of the previous frame, especially when the object’s movement between adjacent frames is relatively large. For instance, when our model reconstructs a walking sequence from multi-view images, the reconstructed surface appears to have many holes, as shown in Fig. 10. To address this issue, we propose a global transformation prediction to roughly transform the target SDF into a canonical space before the incremental training. Specifically, we predict the rotation R and transition T of the object between two adjacent frames. For any given 3D position \mathbf{x}_i in the coordinate space of the frame i , it is transformed back to the coordinate space of the previous frame $i - 1$, denoted as \mathbf{x}_{i-1}

$$\mathbf{x}_{i-1} = R_i(\mathbf{x}_i + T_i). \quad (12)$$

The transformations can then be accumulated to transform the point \mathbf{x}_i back to \mathbf{x}_c in the first frame’s coordinate space

$$\mathbf{x}_c = R_{i-1}^c(\mathbf{x}_{i-1} + T_{i-1}^c) = R_i^c(\mathbf{x}_i + T_i^c), \quad (13)$$

where $R_i^c = R_{i-1}^c R_i$ and $T_i^c = T_i + R_i^{-1} T_{i-1}^c$.

The global transformation prediction also allows us to model dynamic sequences with large movement in a small region, rather than covering the entire scene with a large hash grid. Since the hash grid only needs to model a small portion of the entire scene, we can obtain more accurate reconstructions and reduce memory cost.

Notably, our method can handle large movements and deformations which are challenging for existing dynamic scene reconstruction approaches[56], [45], thanks to the following designs of our approach: (1) Global transformation prediction that accounts for large global movements in the sequence; (2) Incremental training that learns relatively small deformable movements between two adjacent frames rather than learns relatively large movements from each frame to a common canonical space.

We realize incremental training combined with global transformation prediction as an end-to-end learning

	COLMAP	NeuS	InstantNGP	InstantNSR	Ours
CD ↓	1.36	0.77	1.84	1.68	0.70
PSNR ↑	-	28.00	28.86	25.81	28.82
Runtime	1 h	8 h	5 min	8.5 min	5 min

Table 1. The quantitative comparison on DTU dataset. We color code the **best** and **second best** results. Our method outperforms other baselines for geometry reconstruction regarding to the Chamfer Distance (CD) and is on par with Instant-NGP of novel view synthesis in terms of PSNR.

	D-NeRF		TiNeuVox		Ours	
Dataset	PSNR↑	CD↓	PSNR↑	CD↓	PSNR↑	CD↓
Lego	24.25	59.0	28.06	19.90	29.5	17.1
Lion	31.45	-	31.86	-	33.60	-
Human	29.33	5.73	29.00	9.21	33.20	1.86
Runtime	20h		1h		1h	

Table 2. Quantitative comparisons on synthetic scenes. The Chamfer Distance of Lion sequence is omitted since the ground truth geometry is not provided. Compared to D-NeRF and TiNeuVox, our method achieves much better appearance and geometry reconstruction results.

	D-NeRF		TiNeuVox		Ours	
Dataset	PSNR↑	LPIPS↓	PSNR↑	LPIPS↓	PSNR↑	LPIPS↓
D1	21.48	0.122	18.17	0.159	26.41	0.036
D2	17.47	0.154	14.95	0.198	27.76	0.037
D3	20.80	0.155	14.16	0.222	27.25	0.042
Runtime	50h		3h		3h	

Table 3. Quantitative comparisons on real scenes. We found that our method outperforms D-NeRF and TiNeuVox in all metrics.

scheme, as illustrated in Fig. 3(b). When processing a new frame, we first predict the global transformation independently, and then fine-tune the model’s parameters and global transformation together to efficiently learn the neural representation.

6. Experiments

All experiments are conducted on a single GeForce RTX3090 GPU. Implementation details, additional results, and video results are provided in the supplementary material.

6.1. Static Scene Reconstruction

For the static scene reconstruction, we use 15 scenes of DTU dataset [18] for evaluation. There are 49 or 64 images with a resolution of 1600×1200 , and we test each scene with foreground masks provided by IDR [67]. To demonstrate the performance of *NeuS2* on the static scene reconstruction task, we compare our approach with the state-of-art methods NeuS [59], Instant-NGP [36], and Instant-

NSR [68]. We also provide a comparison with a concurrent work, Voxurf [61], in the supplemental material.

For quantitative measurement, we use the Chamfer Distance [18] to evaluate the geometry reconstruction quality, in the same way as NeuS [59] did. We also measure the novel view synthesis quality by the peak signal-to-noise ratio (PSNR) between the reference images and the synthesized images. The quantitative comparison results are reported in Tab. 1, and for per-scene breakdown results please refer to the Suppl. materials. The results show that our method outperforms the baseline methods on the geometry reconstruction task with significantly less training time compared to NeuS [59]. Meanwhile, our method achieves comparable performance in the novel view synthesis task to Instant-NGP [36] requiring the same training time (5 minutes).

A qualitative comparison of geometry reconstruction and novel view synthesis results for all methods is presented in Fig. 4. As shown in Fig. 4 for the 3D geometry reconstruction results, NeuS [59] exhibits limited performance in terms of reconstructed details with excessively smooth surfaces. The extracted meshes of Instant-NGP [36]’s results are noisy since it lacks surface constraints in the geometry representation – volume density field. The results obtained from Instant-NSR [68] exhibit many artifacts, which can be attributed to the use of a finite difference method to approximate the second derivative. This method is prone to precision problems and may cause unstable training. Furthermore, since Instant-NSR requires multiple forward calculations using the finite difference method, their approach is slower than ours. Regarding the novel view synthesis results, our method outperforms NeuS [59] and Instant-NSR [68], presenting detailed rendering results on par with Instant-NGP. To conclude, our method achieves high-quality geometry and appearance reconstruction, with additional details demonstrated in, both, the surface and rendered images without inducing noise, e.g. it can recover the complex structures of the windows and render detailed textures in scan 24.

6.2. Dynamic Scene Reconstruction

We conduct experiments on both synthetic and real *dynamic* scenes for the tasks of novel view synthesis and geometry reconstruction. We compare our method with the state-of-the-art neural-based method for dynamic scenes, D-NeRF [46] and TiNeuVox [11] quantitatively and qualitatively. D-NeRF [46] and TiNeuVox [11] models general scenes for dynamic novel view synthesis by combining NeRF [35] with deformation fields, which is learned from all the frames simultaneously. Compared with D-NeRF which purely uses MLPs, TiNeuVox further accelerates the training by utilizing explicit voxel grids. We also provide the results of Instant-NGP [36] in the Suppl. materials.



Figure 4. Qualitative comparisons on DTU dataset for static scene geometry reconstruction and novel view synthesis. Our method demonstrates high-quality rendering quality, superior to NeuS and comparable to Instant-NGP in terms of complex texture reconstruction. In addition, it outperforms all baselines regarding to 3D geometry reconstruction, with fine details without inducing noise.

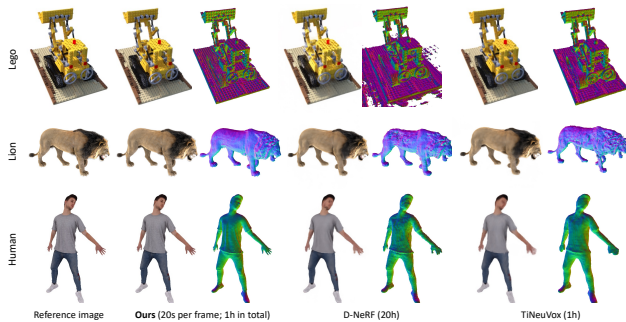


Figure 5. Qualitative comparisons of synthetic scene for geometry reconstruction and novel view synthesis. Our method produces photo-realistic rendering results and accurate geometry reconstructions with only 20 seconds of training time per frame.

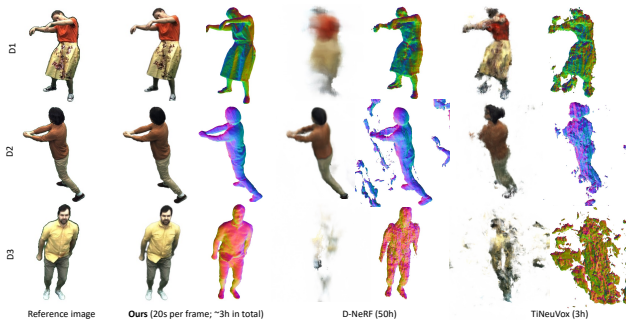


Figure 6. Qualitative comparisons of real scenes for geometry reconstruction and novel view synthesis. Our method outperforms D-NeRF and TiNeuVox for both tasks, demonstrating sharp and accurate reconstruction quality, while they fail to handle complex transformations in real scenes.

Synthetic Scenes. We choose three different types of datasets for synthetic scene reconstruction, including a Lego scene shared by NeRF [35] (150 frames), a Lion sequence provided by Artemis [34] (177 frames), and a human character in the RenderPeople [1] dataset (100 frames). For the quantitative evaluation, the Chamfer Distances and PSNR scores are calculated and averaged over all frames and all testing views of all frames, respectively. As shown in Tab. 2, our method shows significantly improved novel view synthesis and geometry reconstruction results compared to D-NeRF and TiNeuVox. Notably, our method takes less




Figure 7. Long sequence reconstruction results. Our method exhibits robustness even for extremely long sequences, without any performance degradation, e.g., over 20K frames. Video results are provided in the supplementary material.




Figure 8. Free view-point reconstruction results. Video results are provided in the supplementary material.

than 1 hour to complete the learning of a sequence, with 40 seconds (80 seconds for the Lego sequence) of training time for the first frame and 20 seconds for each subsequent frame; while the training time for each scene of D-NeRF is about 20 hours. The qualitative results are provided in Fig. 5, showing that our method outperforms D-NeRF and TiNeuVox in terms of dynamic novel view synthesis and geometry reconstruction.

Real Scenes. To further evaluate the effectiveness of our method on real scenes with large and non-rigid movement, we select three sequences from the Dynacap [15] dataset. Each sequence contains 500 frames with about 50 to 100 camera views for training and about 5 to 10 camera views for testing. More details are provided in the supp. document. Tab. 3 summarizes the quantitative comparisons of our method, D-NeRF [45] and TiNeuVox [11]. Since the real scene datasets do not have ground truth geometry, we can only calculate the PSNR and LPIPS score of novel view synthesis results to evaluate the rendering quality. For long sequences with 500 frames consisting of challenging movements, D-NeRF and TiNeuVox struggle to reconstruct the dynamic real scenes. Even when training D-NeRF for 50 hours, our method achieves significantly better scores,




Figure 9. Ablation study for second-order derivative backward computation. We compare the speedup of our 2nd-order derivative with PyTorch across various MLP layer numbers and batchsizes.

taking only 20 seconds per frame. Also, according to the qualitative evaluation results shown in Fig. 6, D-NeRF and TiNeuVox show blurred rendering results and inaccurate geometry reconstruction. In contrast, our approach produces photo-realistic renderings and detailed geometry. We also conducted experiments on long sequences of 2k and even 20k frames. Here, we present an example with 20K frames in Fig. 7. More results can be found in the Suppl. materials. We provide free view-point results in Fig. 8, showcasing the sequence in suppl. video at 2:55.

6.3. Ablation

We first compare our efficient second-order derivative backward computation (Theorem 1) implemented in CUDA with a baseline where we automatically compute the second-order derivative in PyTorch [42] using their computational graph. As shown in Fig. 9, our method achieves a faster speed for second-order derivative backpropagation than the PyTorch implementation in all settings. Moreover, we evaluated our method for calculating MLP’s second-order derivatives by eliminating the influence of other factors. The ablation study in Tab. 6 further showcases the versatility of NeuS2, demonstrating that it can be easily integrated into other methods for second-order derivatives calculation to greatly bolster training speed and performance.

Second, we evaluate the performance of the individual components of *NeuS2*: Global Transformation Prediction (GTP) and Progressive Training strategy (PT) on dynamic scenes, as shown in Fig. 10 and Tab. 4. The geometry and appearance reconstruction quality of the full model are better than other ablated models, both, quantitatively and qualitatively. The ablated model shows noisy holes on the surface and blurred renderings since it gets stuck in local minima during the incremental training, which can be alleviated by the Global Transformation Prediction and Progressive Training. We also ablated the effectiveness of the Progressive Training Strategy and Eikonal Loss, applied to static scenes within the DTU dataset, as shown in Tab. 5. We evaluate the impact of predicted transformation accuracy on reconstruction quality, as shown in Suppl. materials.

7. Conclusion

Limitations. While our method reconstructs each frame of a dynamic scene in high quality, there is no dense surface correspondences across the frames. A possible way could be to deform a mesh template to fit our learned neural surfaces for each frame, like the mesh tracking used in [8, 68].




Figure 10. Ablation study for Global Transformation Prediction (GTP) and progressive training (PT). Ablated models perform poorly on both novel view synthesis and geometry reconstruction.

Methods	w/o GTP + w/o PT	w/o PT	Full model
PSNR \uparrow	27.78	27.72	28.18
LPIPS \downarrow	0.0595	0.0559	0.0474

Table 4. Ablation study of our design choices. Note that, both, the Global Transformation Prediction (GTP) as well as the Progressive Training (PT) improve the overall results.

	W/O Progressive	W/O Eikonal	Full Model
CD \downarrow	0.75	1.48	0.70
PSNR \uparrow	28.75	28.73	28.82
Runtime	5min 40s	5min 10s	5min

Table 5. Ablation study of Progressive Training and Eikonal loss on DTU.

	PyTorch’s Autograd	NSR’s Finite difference	Ours (in PyTorch)
Runtime \downarrow	10 min	11.5 min	7 min
PSNR \uparrow	35.3	33.4	35.5

Table 6. Comparison of Pytorch Autograd, Instant-NSR finite differences, and our second-order derivatives. We tested all of them in PyTorch and on the same codebase, eliminating other influences, and found that our formulation is superior in terms of speed and accuracy.

Currently, we also need to save the network parameters of each frame (25M). As future work, a compression of such parameters encoding the dynamic scene could be explored.

We proposed a learning-based method for accurate multi-view reconstruction of both static and dynamic scenes at an unprecedented runtime performance. To achieve this, we integrated multi-resolution hash encodings into neural SDF and introduced a simple calculation of the second-order derivatives tailored to our dedicated network architecture. To enhance the training convergence, we presented a progressive training strategy to learn multi-resolution hash encodings. For dynamic scene reconstruction, we proposed an incremental training strategy with a global transformation prediction component, which leverages the shared geometry and appearance information in two consecutive frames.

Acknowledgement

Christian Theobalt and Marc Habermann were supported by ERC Consolidator Grant 4DReply (770784). Lingjie Liu was supported by Lise Meitner Postdoctoral Fellowship.

References

- [1] Renderpeople. <https://renderpeople.com/3d-people/>, 2018.
- [2] Naveed Ahmed, Christian Theobalt, Petar Dobrev, Hans-Peter Seidel, and Sebastian Thrun. Robust fusion of dynamic shape and normal capture for high-quality reconstruction of time-varying geometry. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [3] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2565–2574, 2020.
- [4] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24, 2009.
- [5] Adrian Broadhurst, Tom W Drummond, and Roberto Cipolla. A probabilistic framework for space carving. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 1, pages 388–393. IEEE, 2001.
- [6] Joel Carranza, Christian Theobalt, Marcus A. Magnor, and Hans-Peter Seidel. Free-viewpoint video of human actors. *ACM Trans. Graph.*, 22(3):569–577, jul 2003.
- [7] Jianchuan Chen, Ying Zhang, Di Kang, Xuefei Zhe, Linchao Bao, and Huchuan Lu. Animatable neural radiance fields from monocular rgb video, 2021.
- [8] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. 34(4), jul 2015.
- [9] Edilson de Aguiar, Carsten Stoll, Christian Theobalt, Naveed Ahmed, Hans-Peter Seidel, and Sebastian Thrun. Performance capture from sparse multi-view video. *ACM Trans. Graph.*, 27(3):1–10, aug 2008.
- [10] Jeremy S De Bonet and Paul Viola. Poxels: Probabilistic voxelized volume reconstruction. In *Proceedings of International Conference on Computer Vision (ICCV)*, pages 418–425, 1999.
- [11] Jiemin Fang, Taoran Yi, Xinggang Wang, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Matthias Nießner, and Qi Tian. Fast dynamic radiance fields with time-aware neural voxels. *arxiv:2205.15285*, 2022.
- [12] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE transactions on pattern analysis and machine intelligence*, 32(8):1362–1376, 2009.
- [13] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Gipuma: Massively parallel multi-view stereo reconstruction. *Publikationen der Deutschen Gesellschaft für Photogrammetrie, Fernerkundung und Geoinformation e. V.*, 25(361-369):2, 2016.
- [14] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020.
- [15] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. *ACM Transactions on Graphics (TOG)*, 40(4):1–16, 2021.
- [16] Marc Habermann, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Livecap: Real-time human performance capture from monocular video. *ACM Trans. Graph.*, 38(2):14:1–14:17, Mar. 2019.
- [17] Peter J. Huber. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73 – 101, 1964.
- [18] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014.
- [19] Zhang Jiakai, Liu Xinhang, Ye Xinyi, Zhao Fuqiang, Zhang Yanshun, Wu Minye, Zhang Yingliang, Xu Lan, and Yu Jingyi. Editable free-viewpoint video using a layered neural representation. In *ACM SIGGRAPH*, 2021.
- [20] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [21] Erik Johnson, Marc Habermann, Soshi Shimada, Vladislav Golyanik, and Christian Theobalt. Unbiased 4d: Monocular 4d reconstruction with a neural deformation model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6597–6606, 2023.
- [22] Srinivas Kaza et al. *Differentiable volume rendering using signed distance functions*. PhD thesis, Massachusetts Institute of Technology, 2019.
- [23] Petr Kellnhofer, Lars Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. Neural lumigraph rendering. *arXiv preprint arXiv:2103.11571*, 2021.
- [24] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. *NeurIPS*, 2021.
- [25] Lingzhi Li, Zhen Shen, Zhongshu Wang, Li Shen, and Ping Tan. Streaming radiance fields for 3d video synthesis. *arXiv preprint arXiv:2210.14831*, 2022.
- [26] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5521–5531, 2022.
- [27] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [28] Z. Li, Simon Niklaus, Noah Snaveley, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. *ArXiv*, abs/2011.13084, 2020.
- [29] Jia-Wei Liu, Yan-Pei Cao, Weijia Mao, Wenqiao Zhang, David Junhao Zhang, Jussi Keppo, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Devrf: Fast deformable voxel radiance fields for dynamic scenes. *arXiv preprint arXiv:2205.15723*, 2022.

- [30] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33, 2020.
- [31] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Trans. Graph.*, 40(6), dec 2021.
- [32] Shaohui Liu, Yinda Zhang, Songyou Peng, Boxin Shi, Marc Pollefeys, and Zhaopeng Cui. Dist: Rendering deep implicit signed distance function with differentiable sphere tracing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2019–2028, 2020.
- [33] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Transactions on Graphics (TOG)*, 38(4):65, 2019.
- [34] Haimin Luo, Teng Xu, Yuheng Jiang, Chenglin Zhou, Qiwei Qiu, Yingliang Zhang, Wei Yang, Lan Xu, and Jingyi Yu. Artemis: Articulated neural pets with appearance and motion synthesis. *arXiv preprint arXiv:2202.05628*, 2022.
- [35] 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 *European Conference on Computer Vision*, pages 405–421. Springer, 2020.
- [36] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022.
- [37] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020.
- [38] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *International Conference on Computer Vision*, 2021.
- [39] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *International Conference on Computer Vision (ICCV)*, 2021.
- [40] 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 *International Conference on Computer Vision (ICCV)*, 2021.
- [41] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *ACM Trans. Graph.*, 40(6), dec 2021.
- [42] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward 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 *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [43] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Animatable neural radiance fields for human body modeling. *ICCV*, 2021.
- [44] 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. *CVPR*, 1(1):9054–9063, 2021.
- [45] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [46] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021.
- [47] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2304–2314, 2019.
- [48] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 84–93, 2020.
- [49] Sara Fridovich-Keil and Alex Yu, Matthew Tancik, Qinlong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022.
- [50] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision*, pages 501–518. Springer, 2016.
- [51] Steven M Seitz and Charles R Dyer. Photorealistic scene reconstruction by voxel coloring. *International Journal of Computer Vision*, 35(2):151–173, 1999.
- [52] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019.
- [53] Vincent Sitzmann, Michael Zollhofer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, pages 1119–1130, 2019.
- [54] Shih-Yang Su, Frank Yu, Michael Zollhofer, and Helge Rhodin. A-nerf: Surface-free human 3d pose refinement via neural rendering. *arXiv preprint arXiv:2102.06199*, 2021.
- [55] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, 2022.

- [56] 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 *IEEE International Conference on Computer Vision (ICCV)*. IEEE, 2021.
- [57] Daniel Vlasic, Ilya Baran, Wojciech Matusik, and Jovan Popović. Articulated mesh animation from multi-view silhouettes. *ACM Trans. Graph.*, 27(3):1–9, aug 2008.
- [58] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yanshun Zhang, Yingliang Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Fourier plencotrees for dynamic radiance field rendering in real-time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13524–13534, 2022.
- [59] 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. *NeurIPS*, 2021.
- [60] Yiming Wang, Qingzhe Gao, Libin Liu, Lingjie Liu, Christian Theobalt, and Baoquan Chen. Neural novel actor: Learning a generalized animatable neural representation for human actors. 2022.
- [61] Tong Wu, Jiaqi Wang, Xingang Pan, Xudong Xu, Christian Theobalt, Ziwei Liu, and Dahua Lin. Voxurf: Voxel-based efficient and accurate neural surface reconstruction, 2022.
- [62] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [63] Hongyi Xu, Thiemo Alldieck, and Cristian Sminchisescu. H-nerf: Neural radiance fields for rendering and temporal reconstruction of humans in motion, 2021.
- [64] Weipeng Xu, Avishek Chatterjee, Michael Zollhöfer, Helge Rhodin, Dushyant Mehta, Hans-Peter Seidel, and Christian Theobalt. Monoperfcap: Human performance capture from monocular video. *ACM Trans. Graph.*, 37(2):27:1–27:15, May 2018.
- [65] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [66] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33, 2020.
- [67] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33:2492–2502, 2020.
- [68] Fuqiang Zhao, Yuheng Jiang, Kaixin Yao, Jiakai Zhang, Liao Wang, Haizhao Dai, Yuhui Zhong, Yingliang Zhang, Minye Wu, Lan Xu, et al. Human performance modeling and rendering via neural animated mesh. *arXiv preprint arXiv:2209.08468*, 2022.