

# Real-Time Neural Character Rendering with Pose-Guided Multiplane Images

Hao Ouyang<sup>1</sup>, Bo Zhang<sup>2</sup>, Pan Zhang<sup>2</sup>, Hao Yang<sup>2</sup>, Jiaolong Yang<sup>2</sup>, Dong Chen<sup>2</sup>, Qifeng Chen<sup>1</sup>, and Fang Wen<sup>2</sup>

<sup>1</sup> Hong Kong University of Science and Technology

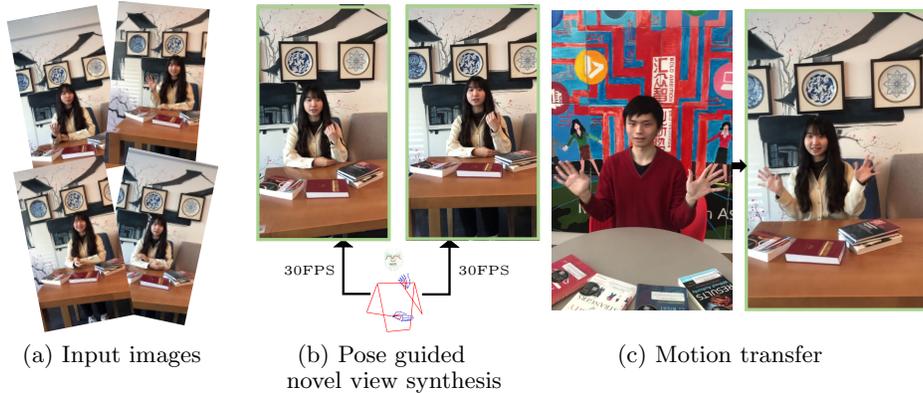
<sup>2</sup> Microsoft Research Asia

**Abstract.** We propose pose-guided multiplane image (MPI) synthesis which can render an animatable character in real scenes with photorealistic quality. We use a portable camera rig to capture the multi-view images along with the driving signal for the moving subject. Our method generalizes the image-to-image translation paradigm, which translates the human pose to a 3D scene representation — MPIs that can be rendered in free viewpoints, using the multi-views captures as supervision. To fully cultivate the potential of MPI, we propose depth-adaptive MPI which can be learned using variable exposure images while being robust to inaccurate camera registration. Our method demonstrates advantageous novel-view synthesis quality over the state-of-the-art approaches for characters with challenging motions. Moreover, the proposed method is generalizable to novel combinations of training poses and can be explicitly controlled. Our method achieves such expressive and animatable character rendering all in real time, serving as a promising solution for practical applications.

## 1 Introduction

Using a *handy camera rig* for data capturing, can we synthesize a photorealistic character with *controllable viewpoints and body poses in real-time*? Such a technique would democratize personalized, photorealistic avatars and enable various intriguing applications such as telepresence, where people will feel the virtual character of a remote person as a real person talking in real time.

Traditionally, free-viewpoint rendering of a moving person is approached by capturing a high-fidelity 3D human model in a specialized studio [5,6,8,11], which is a costly and brittle process and is not accessible to common users. Recently, researchers use data-driven methods [19,12,18,44,3,20,46] to expedite this process. These methods focus on rendering and animating human actors but do not address the interactions between human and real scenes (*e.g.*, a person sitting on a couch with arms on a table). Moreover, they cannot handle challenging motions such as finger movements. Recently, deformable NeRF methods [27,28,32,40] have been proposed to model the person and the scene by learning an implicit deformation field along with a canonical radiance field that serves as the template. However, these methods can only handle small deformations as it is hard to



**Fig. 1.** Using images from a handy capture device, our method renders a photo-realistic character in real time. The character is animatable through motion transfer.

model complex human motions with a single canonical representation. Besides, the expensive volumetric rendering process prohibits real-time video synthesis.

For character modeling in unconstrained environments, we propose a novel capture setup that comprises a portable capture rig and a fixed driving camera. Synchronized mobile phone cameras are mounted on the rig, and the videographer slightly moves the camera rig to capture the light field of the scene. The fixed camera is used to extract the driving keypoints that users can manipulate for character retargeting. With this handy setup, we strike the desired balance between 3D sensing accuracy and hardware affordability since the rig movement greatly reduces the number of cameras needed.

We also introduce *pose-guided multiplane images* for fast and controllable character rendering with high fidelity. The multiplane image (MPI) representation [49] uses a set of parallel semi-transparent planes to approximate the light field and has shown compelling quality for complex scene modeling. Rather than performing optimization directly upon MPI [49,9,25,45], we propose to use a neural network to produce the planes with pose conditioning, and the whole framework is essentially a pose-to-MPI translation network. We use 2D keypoints on the image plane of the driving camera to define the pose, and predict the multiplane images in the frustum of the same camera. Hence, the two representations are spatially aligned and compatible for network processing. During inference, the keypoints serve as the driving signal, and one can obtain the corresponding character in novel views by rendering the predicted MPI from the target viewpoint.

Our method brings several benefits. 1) Since we do not assume a human model or a canonical template but learn character modeling purely in a data-driven manner, our method offers improved flexibility to characterize the subject as well as complex interactions with the scene. In particular, we observe substantial improvement for gesture modeling, which is challenging for prior arts. 2) Taking advantage of the inductive bias of convolutional neural networks, our

method demonstrates improved generalization ability. Instead of memorizing the scene, the network trained with large data can hallucinate plausible outputs for diverse poses. 3) The MPI rendering is blazingly fast, and our network can synthesize videos of  $640 \times 360$  resolution in real time.

We further propose several techniques to achieve better MPI synthesis. As opposed to evenly placing the planes in disparity space, we propose adaptive planes whose positions are jointly optimized during training, which considerably improves the modeling quality because the planes are now densely placed near the real surfaces of the scene. To compensate for the exposure mismatch among different cameras, we also introduce a learnable exposure code for each camera. Moreover, when dealing with long video sequences (*e.g.*,  $> 4k$  frames), we observe unsatisfactory camera pose estimation using conventional structure-from-motion (SfM) pipeline [34], which leads to blurry results. To solve this, we refine the camera poses during training using the gradients of the static background pixels.

We demonstrate that our pose-guided MPI quantitatively and qualitatively outperforms state-of-the-art approaches including Video-NeRF [17], Nerfies [27], HyperNeRF [28], and NeX [45] on novel view synthesis for characters with complex motion. Moreover, the character rendered by our method can be explicitly controlled, and we achieve photorealistic results of character reenactment, as illustrated in Fig. 1. The whole rendering framework runs in real time and is scalable to high resolutions.

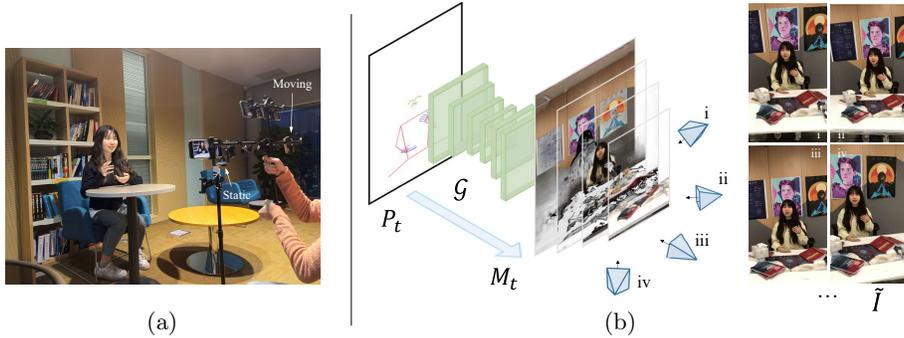
## 2 Related Work

**Neural character rendering** Traditional graphics pipelines [5,6,8,11] require a well-orchestrated studio with a dense array of cameras to build the mesh for the characters. In the past years, deep generative neural networks have been introduced to synthesize photorealistic characters [38,39]. A popular synthesis paradigm is image-to-image translation [14,43], which learns the mapping from certain representations, such as joint heatmap [23,2,4,48,50], rendered skeleton [7,31,36,35,42], and depth map [24], to real images of the character. These works have certain generalization ability to novel poses and have shown compelling rendering quality even for complex clothing and in-the-wild scenes. However, they cannot guarantee view consistency as they learn the generation in 2D screen space. More recent methods attempt to solve this by leveraging a human model, *e.g.*, SMPL model [21]. One line of works [19,33,3] unwraps the body mesh to 2D UV space where the network learns texture synthesis and then translates the rendered textured mesh to images. Meanwhile, another line of works improves view consistency by learning the deformation to a canonical 3D space enforced by the SMPL model [12]. In comparison, our method does not assume an explicit human model and hence can model complicated finger motions as well as the character’s interactions with the scene. Our method generalizes the image translation but the output representation we adopt ensures multi-view consistency.

**Neural scene representation** Instead of rendering with a black-box image translation process, recent works turn to using neural networks to model some intrinsic aspects of the scene followed by a physics-based differentiable renderer. One recent notable work is to model the scene as a neural radiance field (NeRF) whose color and volumetric density across the continuous coordinates are parameterized by the multilayer perceptron (MLP). The NeRF representation ensures that the images rendered at different views have coherent geometry, and the volumetric rendering is able to produce realistic outputs with stunning details.

Several NeRF variants [32,27,28,40,10,17] have been proposed to handle dynamic scenes. One category of these works is deformation-based [32,27,28,40] which optimizes a deformation field that warps each observed point to a canonical NeRF. These approaches have shown impressive quality even for in-the-wild scenes, but only simple deformation can be modeled as it is difficult to find a template that accounts for all the observations. To render characters with NeRF [18,44,46,29], there are works that use the SMPL model to enforce the canonical NeRF, but they may suffer from the limited modeling capability of the parametric human model. Some methods modulate NeRF with additional conditioning [17,47] and achieve enhanced expressivity to model dramatic topology change. Nonetheless, the majority of these methods are designed for replaying the dynamic scene, and it is hard to generalize to novel character poses. This is because the point-wise MLPs in NeRF models do not leverage a larger context for generative modeling. In contrast, we employ a convolutional neural network which leverages a large image context for hallucination, and thus our method is more amenable to animation and generalizes better. Moreover, our method renders characters much faster than NeRF-based approaches.

**Multiplane images representation** The MPI representation uses a stack of  $RGB\alpha$  layers arranged at various depths to approximate the light field. Initially, the MPI is used for the stereo magnification problem [49] where an MPI is predicted from a CNN given a stereo pair input. Later, a few works extend the MPI for view synthesis where images from a variable number of viewpoints can be accessed [9,25,45,37,41,16]. The multiplane images can be optimized in three ways: direct optimization [25,45,37], using a CNN to learn gradient updates upon MPI [9] or directly predicting the MPI using CNN [41,16]. The main drawback of MPI methods is that the rendered view range is limited by the number of planes. To solve this, [25] proposes to use multiple MPIs to account for local light fields, which are further blended for the final output. MPI-based approaches have limited power for handling non-Lambertian surfaces and wide viewing angles, and they are outshined by the recent advances of NeRF-based methods. Very recently, NeX [45] tackles view-dependent modeling by representing the pixel color of MPI as a combination of spherical basis functions, and it excels over NeRF in visual quality with much faster rendering speed. However, no existing method ever studies the controllability of such a representation. Our approach is motivated by the recent success of MPI, and we demonstrate the great potential of this representation for animatable character rendering.



**Fig. 2.** Overview of our method. (a) The device we use to capture data for building an animatable character in a real scene. (b) The built character is controllable. Given any pose  $\mathbf{P}_t$  at time step  $t$ , we feed its pose image into a pre-trained network  $\mathcal{G}_\theta$  to acquire a 3D character represented in MPI ( $\mathbf{M}_t$ ). The character has the same pose with  $\mathbf{P}_t$ , and can also be efficiently rendered in free views.

### 3 Approach

We aim to render an animatable character that can be controlled by a driving input. To this end, we devise a portable data capture setup to ease the multi-view capture in open scenes (Section 3.1). Once we finish the data capture for the moving character, we train a neural network conditioned on character poses to predict the multiplane images that explain the multi-view observations (Section 3.2). During inference, we can render realistic characters given a driving input (Section 3.3). Next, we elaborate on these three parts respectively.

#### 3.1 Data capture

In our design, data capturing should meet the following requirements. First, the capture device should be portable, lightweight, and low-cost to benefit as many users as possible for their character creation. Second, the captured images for the moving subject should resolve most motion ambiguity. Otherwise, the character reconstruction from sparse views is highly ill-posed. Third, we need to use some driving signal handy for user control, and such signal should be readily obtained from the same capture setup.

Taking these into consideration, we propose a novel character capture setup as shown in Fig. 2 (a), which consists of a moving camera rig along with a static camera fixed on a tripod. We mount four smartphone cameras on the capture rig, which the videographer can hold and move to capture the subject. Such capturing manner is motivated by Nerfies [27] that uses a single moving camera for selfies. In our early attempts, we find a single camera does not suffice to resolve the ambiguity for complex non-rigid motion, whereas adding a few more cameras significantly improves the modeling quality. Also, the rig motion makes the capture to cover various combinations of viewpoints and character poses,

which helps to reduce the number of multi-view cameras by requiring longer sequence capturing. On the other hand, we propose to use body keypoints as the driving signal for character animation. We believe manipulating 2D keypoints is easy for most users since people can extract such keypoints from their monocular videos to drive the virtual character. Hence, we also require a fixed camera, or “driving camera”, which is used to extract the driving keypoints. Compared to a specialized lab, it is much cheaper to build this capture setup. Once we capture the data, we use audio to synchronize the multi-view videos with the driving camera and run COLMAP [34] to estimate the camera poses for each frame.

### 3.2 Conditional MPI representation

Our method builds on the multiplane image (MPI) scene representation, which consists of  $D$  fronto-parallel planes, each with an associated  $H \times W \times 4$  RGB $\alpha$  image. As illustrated in Fig. 2 (b), the multiplane images are scaled and positioned at different depths  $d_1, \dots, d_D$  within the view frustum of the driving camera. Typically, the planes are placed equally in the depth space (for the bounded scene) or in the disparity space (for the unbounded scene).

While MPI-based methods have shown impressive quality in modeling static scenes from sparse views, its ability to model a moving character, especially those with complex motions, remains underexplored. To accomplish this, we formulate the character rendering by pose-guided MPI synthesis framework, which is illustrated in Fig. 2 (b). Given the pose image  $\mathbf{P}_t$  extracted from the driving camera at time  $t$ , we train a convolutional neural network  $\mathcal{G}_\theta$  to translate this input to multiplane images  $\mathbf{M}_t = \mathcal{G}_\theta(\mathbf{P}_t)$ , using the supervision of the multi-view observations  $\{\mathbf{I}_t^n\}_{n=1}^N$ . Here,  $N$  denotes the number of cameras we use for data capture. During training, we extract rich character pose information using [22], which includes facial landmarks, body keypoints, and finger keypoints that are extracted from the driving frames. Note that both  $\mathbf{P}_t$  and  $\mathbf{M}_t$  are spatially aligned as they are viewed from the same driving camera; hence their mapping is naturally suitable for the CNN learning.

The synthesized MPI can be rendered in the target view by compositing the colors along the rays. The implementation is efficient: the image planes are first warped towards the target view and then alpha-blended. Specifically, we refer to the RGB channels of the MPI as  $\mathbf{C} = \{c_1, \dots, c_D\}$  and the corresponding alpha channels as  $\mathbf{A} = \{\alpha_1, \dots, \alpha_D\}$ . The MPI rendering can be formulated as

$$\tilde{\mathbf{I}} = \mathcal{O}(\mathcal{W}(\mathbf{C}), \mathcal{W}(\mathbf{A})), \quad (1)$$

where the warping operator  $\mathcal{W}$  warps the images via a homography function [49] depending on the relative rotation  $\mathbf{R}$  and translation  $\mathbf{t}$  from the target to the source view and the layer depth  $d_i$ . Formally, the warping matrix is

$$\mathbf{K}_s(\mathbf{R} - \frac{\mathbf{t}\mathbf{n}^T}{d_i})\mathbf{K}_t^{-1}, \quad (2)$$

where  $\mathbf{n}$  is the normal vector;  $\mathbf{K}_s$  and  $\mathbf{K}_t$  respectively are intrinsic matrix of source camera and target camera. Besides,  $\mathcal{O}$  in Equation 1 denotes the composition [30] of the warped images ( $c'_i$  and  $\alpha'_i$ ) from back to front, *i.e.*,

$$\mathcal{O}(C, A) = \sum_{i=1}^D \left( c'_i \alpha'_i \prod_{j=i+1}^D (1 - \alpha'_j) \right). \quad (3)$$

The above rendering process is fully differentiable, so the MPI synthesis network can be trained end-to-end using 2D supervision.

This pose-guided MPI enjoys many features of CNNs and shows various advantages over implicit approaches. First of all, we do not explicitly model the geometry of the scene or assume a template for all the character motions, so our model is more expressive and can better fit challenging motions and delicate details, as proved in our view synthesis experiments. Second, compared to implicit neural representation, our method can generalize better thanks to the inductive bias of CNNs. Intuitively, it seems that our approach is good at hallucinating plausible outputs, even for unseen poses. Third, manipulating the keypoints is more straightforward for character animation instead of using latent code or body parameters. Finally, the MPI is inferred with a single feed-forward pass of the network, and its rendering is also fast.

To cultivate the full potential of the above framework and achieve state-of-the-art quality, we introduce the following key components.

**Depth-adaptive MPI** We argue that manually placing the multiplane images at fixed depth may not be optimal. Ideally, the planes should be distributed more densely around the real scene surfaces. Otherwise, some planes are wasted for modeling the vacant space. More importantly, in our scenario the bound of the scene cannot be reliably estimated because we have to mask out the moving foreground when computing the COLMAP. As a result, the MPIs initialized with a mistaken depth range may lead to scene clipping modeling.

In view of this, we propose to make the MPI depth as learnable parameters so that MPIs positions can be adaptive to the scene content. Formally, we refer to the initialized depth as  $d_i^{init}$ . During training, we learn the residual  $\delta_i$  which is initialized with zeros, so the refined depth becomes  $d_i = d_i^{init} + \delta_i$ . As we know, the homography warping  $\mathcal{W}$  in Equation 1 is the function of  $d_i$ ; hence the gradient can be back-propagated to update the depth as the training proceeds.

However, one may notice that the depth refinement may alter the plane orderings, which will mislead the alpha composition (Equation 3) that renders the planes from back to front. Therefore, we need to enforce the depth order to be unchanged in the depth refinement. To achieve this, we clamp the value of  $\delta_i$  once we find the plane shift causes the crossing over the adjacent planes. During training, the depth refinement uses 0.1x learning rate compared to the network.

**Learning with variable exposure images** There always exist exposure and color differences among cameras even if we employ cameras of the same type and

manually choose the ISO and exposure time. To use different exposed images for our training, we assume that the MPI rendering models the true light intensity of the scene but we introduce learnable exposure coefficients to account for the exposure variance among cameras. Specifically, we adopt a linear exposure model [26,1] which outputs the image as

$$\hat{\mathbf{I}} = \text{clamp}((\tilde{\mathbf{I}} + \boldsymbol{\beta}) \circ \boldsymbol{\gamma}), \quad (4)$$

where  $\text{clamp}(\cdot) = \min(\max(\cdot, 0), 1)$  whereas  $\boldsymbol{\beta} \in \mathbb{R}^3$  and  $\boldsymbol{\gamma} \in \mathbb{R}^3$  are learnable exposure coefficients associated with each camera. Here,  $\boldsymbol{\gamma}$  accounts for the exposure time whereas  $\boldsymbol{\beta}$  is for compensating the shift of black level.

**Learnable camera poses** In our experiments, the model sometimes fails to reconstruct the static background because of the inaccurate camera pose estimation from SfM. The problem becomes even worse when dealing with long video sequences to capture more diverse character poses.

To improve the robustness of inaccurate camera registration, we jointly refine the camera poses during training. The gradient through the homography warping can be used to update the camera poses for each frame. Note that the pose refinement can only leverage the gradient of a static background, whereas the MPI synthesis is updated using the whole image. Therefore, the optimization follows an alternative manner: we take two consecutive training steps to alternatively optimize the MPI synthesis network and the camera pose, with the loss computing over the whole image and the background, respectively. Since this may slow down the training speed, we only apply this strategy for sequences with blurry background reconstruction.

**Sharing textures for compact MPI** It is known that a large number of RGBA layers are helpful for high-fidelity modeling, but this brings huge memory costs when directly using the network for the MPI synthesis. To make the MPI more compact, we follow the strategy of NeX [45] and share the RGB textures for every  $K$  layers. In this way, we reduce the output channels from  $4D$  to  $(D + 3D/K)$  without obvious degradation in visual quality.

**Losses** To optimize our model, we feed input pose to generate the MPI and render an output image  $\hat{\mathbf{I}}$  using the camera pose of the ground-truth image  $\mathbf{I}$ . We use three losses for reconstruction: mean square error between  $\mathbf{I}$  and  $\hat{\mathbf{I}}$  as  $\mathcal{L}_2: \|\mathbf{I} - \hat{\mathbf{I}}\|^2$ , gradient difference along the width and height dimension as  $\mathcal{L}_{grad}: \|\nabla \mathbf{I} - \nabla \hat{\mathbf{I}}\|^2$  and the perceptual loss of the difference between VGG features:  $\mathcal{L}_{perceptual}: \|VGG_F(\mathbf{I}) - VGG_F(\hat{\mathbf{I}})\|_1$ . In total we optimize:

$$\min_{\theta, d, R, t, \beta, \gamma} \mathcal{L}_2 + \lambda_1 \mathcal{L}_{grad} + \lambda_2 \mathcal{L}_{perceptual}. \quad (5)$$

Note that for  $\mathcal{L}_2$  loss and the  $\mathcal{L}_{grad}$ , we apply a  $10\times$  weight on the foreground person using the object mask detected by [13].

### 3.3 Motion transfer

Our pose-to-MPI translation network learns to generate 3D representation with a conditional pose. Given a trained model of a character, we can transfer the motion from the driving character and generate novel views. Since characters differ in height, limb length and body shape, it is not suitable to transfer the absolute pose directly from the driving character to the source character built by our method. Relative motion transfer, which keeps the physical characteristics of the source character, is desired. Our input pose comprises face keypoints, body keypoints, and finger keypoints. Denote the face landmarks of the driving character by  $t$  and that of the source character by  $s$ . To transfer the relative motion, we find the landmarks  $t'$  most similar to  $s$  in the driving video. Thus, the transferred pose for the source character becomes  $s + t - t'$ . We treat the body landmarks as a tree structure for body motion transfer. We generate the transferred body keypoints by keeping the same limb length as the source body while utilizing the limb direction of the driving body. The tree root is the mid-point of the left shoulder and right shoulder. Finger motion transfer is similar to body transfer, except that the tree root is changed to the wrist. Please refer to the supplementary material for more details.

## 4 Experiments

### 4.1 Implementation details

The output MPI is composed of 192 alpha layers, with every 12 alpha layers sharing the same RGB texture layer, which leads to  $192/12 = 16$  RGB texture layers. The total number of channels to output thus becomes  $192 + 16 \times 3 = 240$ . Our video sequences are all captured in 1080p resolution. The temporal length of each captured sequence lies between 1 to 3 minutes. For each sequence, 1/16 frames are selected as a validation set with the rest for training. We down-sample each video frame to  $360 \times 640$  for fast inference. However, the resolution of each output MPI layer is larger than this to support rendering with wider view angles: it is equal to padding 180 pixels to all four sides of the  $360 \times 640$  frame. We train the model using Adam[15] optimizer and decay the learning rate from  $1e-3$  to  $1e-4$  in 500 epochs. The training takes 12 hours with four Tesla V100 GPUs.

### 4.2 Evaluation

Both quantitative and qualitative analysis is conducted to evaluate our methods thoroughly w.r.t the following two aspects: (i) the ability to synthesize novel views (ii) the ability to generalize to novel poses. We highly recommend the readers to watch our supplementary video for a more comprehensive evaluation. **Novel view synthesis** We compare our approach with four other scene representation methods. Among them, NeX [45] is the state-of-the-art MPI-based method. It is proposed to represent only static scenes. The other three are NeRF-variants, which comprise different extensions to handle dynamic scenes.

	Sequence 1 (755 images)			Sequence 2 (755 images)			Sequence 3 (755 images)			MEAN		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
NeX[45]	24.05	0.919	0.168	22.63	0.921	0.197	24.26	0.908	0.225	23.65	0.916	0.197
Video-Nerf[17]	27.74	0.940	0.189	26.94	0.932	0.230	27.20	0.931	0.221	27.30	0.934	0.213
Nerfies[27]	27.91	0.938	0.180	26.90	0.922	0.180	26.84	0.932	0.191	27.22	0.931	0.187
HyperNerf[28]	28.10	0.946	0.162	<b>27.13</b>	0.942	0.178	<b>27.23</b>	0.939	0.193	<b>27.48</b>	0.942	0.178
Ours	<b>28.20</b>	<b>0.954</b>	<b>0.062</b>	26.72	<b>0.957</b>	<b>0.082</b>	26.95	<b>0.945</b>	<b>0.092</b>	27.29	<b>0.952</b>	<b>0.079</b>

**Table 1.** Quantitative comparisons on different datasets in terms of PSNR $\uparrow$ , SSIM $\uparrow$ , and LPIPS $\downarrow$ . The best results are highlighted in **bold**.

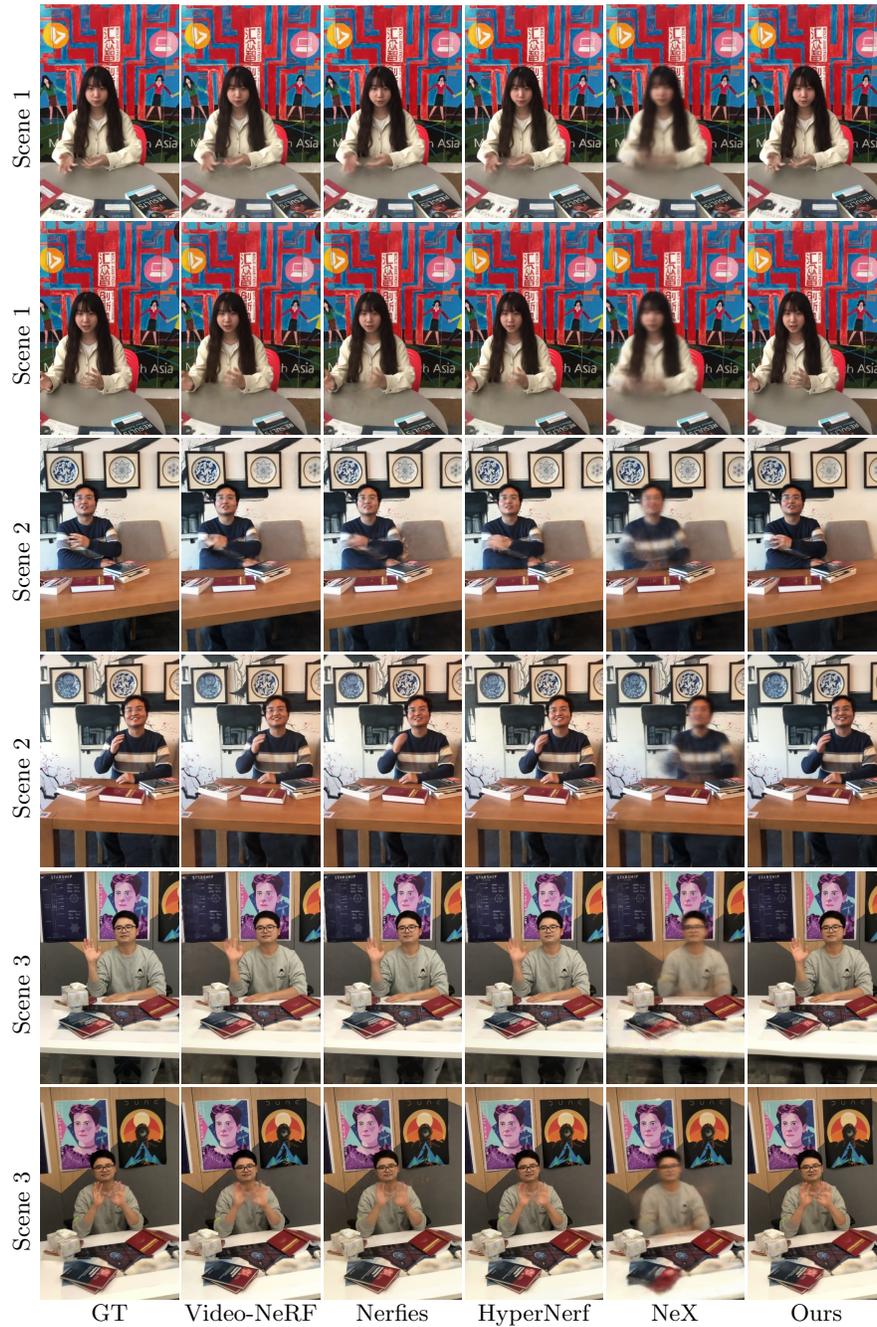
	Sequence 3 (755 images)			Sequence 4 (755 images)			Sequence 5 (5000 images)			MEAN		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
Ours w/o AD	18.00	0.793	0.314	23.27	0.918	0.110	20.34	0.812	0.205	20.54	0.841	0.210
Ours w/o EC	26.70	0.940	0.092	26.53	0.922	0.078	21.82	0.822	0.187	25.02	0.895	0.119
Ours-s	26.88	0.940	0.103	26.21	0.928	0.090	21.22	0.817	0.190	24.77	0.895	0.128
Ours	26.95	<b>0.945</b>	0.092	<b>26.63</b>	<b>0.932</b>	<b>0.077</b>	21.92	0.823	0.186	25.17	0.900	0.118
Ours w/ LCP	<b>27.01</b>	0.943	<b>0.091</b>	26.48	0.920	0.084	<b>23.12</b>	<b>0.845</b>	<b>0.162</b>	<b>25.54</b>	<b>0.902</b>	<b>0.112</b>

**Table 2.** Quantitative ablation results. The best results are highlighted in **bold**.

In specific, Video-NeRF [17] conditions NeRF functions on an extra time-variant latent code; Nerfies [27] introduces dynamic deformations into NeRF; HyperNeRF [28], on the other hand, combines both the latent code and the deformation into NeRF representation and achieves better performance. Moreover, we also introduce weighted sampling to emphasize the foreground training. Some dataset-specific settings are adopted for the baselines to ensure a fair comparison. We also make sure that images with the same timestamp share the same learned latent code and images from the same camera share the same appearance code. During training, we raise the sampling probability of foreground rays to be 10 $\times$  of the sampling probability corresponding to background. We find that employing these settings generally improves the results of baseline methods.

We illustrate visual comparisons in Fig. 3. Without a dynamic modeling capability, NeX fails to capture motions and generates only blurry bodies and faces. As for Video-NeRF, Nerfies, and Hyper-NeRF, they hardly capture the complex motion of hands since their conditioning input is latent vectors, which are too compressed to represent a detailed pose of a character. In comparison, our method well handles the challenging motions with delicate details.

We also conduct a quantitative comparison using three metrics: LPIPS, MS-SSIM, and PSNR. We report scores on three testing sequences in Table 1. Our method achieves better SSIM and LPIPS on all sequences but relatively inferior PSNR on Sequences 2 and 3. We agree with [27] that PSNR is very sensitive to small misalignments between prediction and ground truth hence it might not well reflect real perceptual quality under these cases. Moreover, we report the inference time of all these methods in Table 3. NeRF-base methods rely on a



**Fig. 3.** Visual comparisons of different methods on validation novel views.

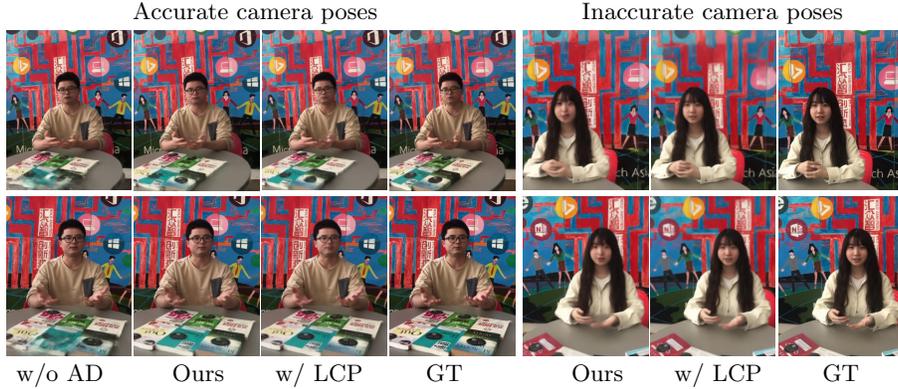


Fig. 4. Qualitative results for the ablation study.

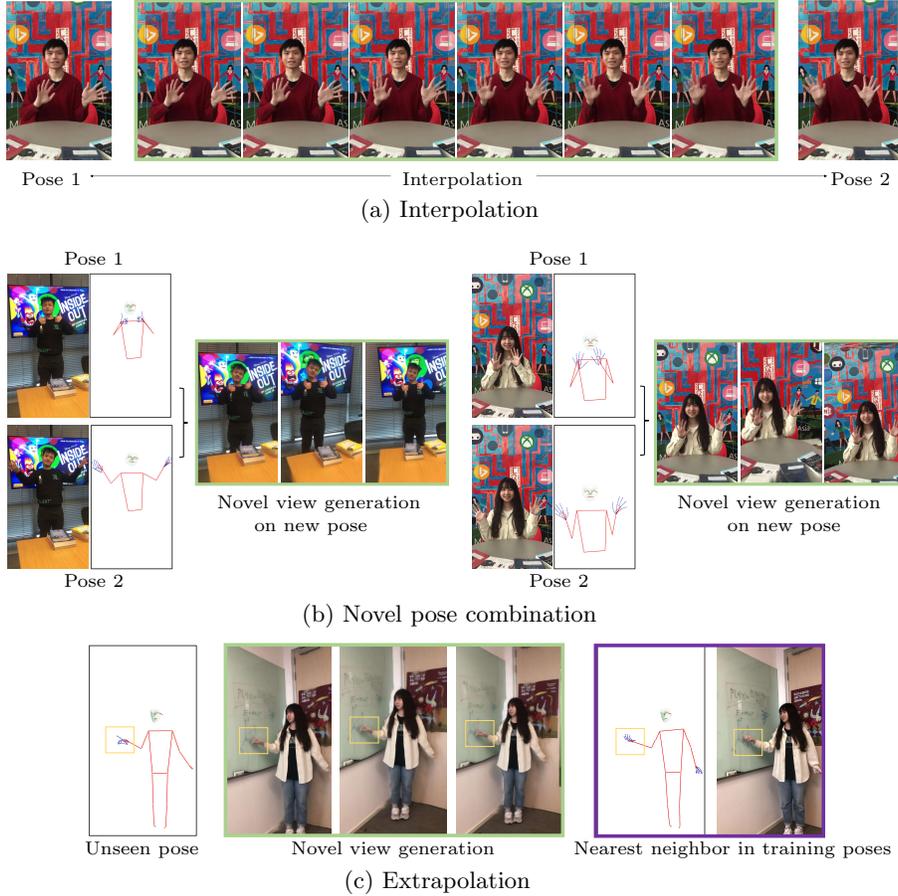
Method	NeX	VideoNerf	Nerfies	HyperNerf	Ours	Ours-s
Inference Time (s)	3.232	12.315	18.661	19.676	0.049	<b>0.031</b>
FPS	0.309	0.081	0.053	0.051	20.41	<b>32.26</b>

Table 3. Inference speed comparison. The performance is measured on images with  $360 \times 640$  resolution using a single Tesla V100 GPU.

densely sampling procedure, both time and memory-consuming. Instead, our results is inferred with a single forward pass of a CNN, which leads to hundreds of times of acceleration over all the other baseline methods.

**Novel poses** As shown in Fig. 5, we introduce three settings to illustrate the capability of our method in handling novel poses. *Interpolation*: unseen poses are smoothly interpolated between two given poses. As shown by Fig. 5 (a), our method can generate a character that smoothly acts following the interpolated poses. *Novel combination*: unseen poses are generated by combining different parts of given poses. As shown by Fig. 5 (b), the 3D character generated from our method faithfully combines the two poses, *e.g.* the tilted-head man with thumb-up gestures, the girl with both hands waving to the same side. These poses are never observed in the training dataset. *Small extrapolation*: As shown by Fig. 5 (c), our framework can also handle small extrapolation as the trained network is essentially a generative model for a single person. Please refer to our supplemental material for more results on novel poses.

**Ablation study** We conduct an ablation study to evaluate the effectiveness of each proposed design. We report scores in Table 2, with results visualized in Fig. 4. AD stands for the “adaptive depth”, EC means to use the “exposure coefficients”, while LCP stands for the “learnable camera poses”. We have the following findings: (i) Without the adaptive plane depth, the learning process is sensitive to the initial plane depth assignment and fails to represent the details of some objects, *e.g.* the book. (ii) In cases where the initial camera registration



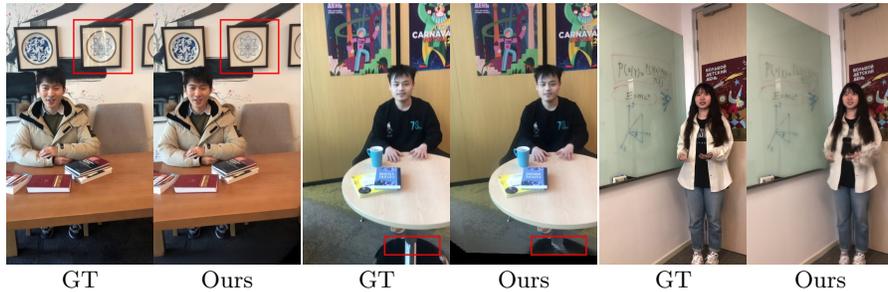
**Fig. 5.** Generalization study. Our approach performs well on motion interpolation and novel pose combination, and allows pose extrapolation to some extent.

is inaccurate, the learnable pose adjusts the initial camera poses and leads to obvious improvement in generating clearer faces. (iii) Besides, the exposure code helps adjust the global luminance level, which results in slight improvements in quantitative metrics. (iv) We also test our method with smaller U-Nets (Ours-s in Table 2) with fewer channels and only observe minor quality degradation.

**Applications** As a straightforward application, our framework can synchronize different 3D characters with the same pose, as shown by Fig. 6. With the learned representation of a 3D character, we extract a 2D pose motion from the driving character and transfer it to the character. The 3D character can be controllable by any given 2D poses and rendered to novel views in real time.



**Fig. 6.** The generated character (rightmost results) can be driven by the subject with photorealistic quality.



**Fig. 7.** Failure cases. Our model fails to generate the reflection effects and may present undesired layered textures. The quality also degrades for large extrapolated motions.

## 5 Limitations and Conclusion

This work studies a controllable representation of a real-scene 3D character based on a pose-guided MPI. Extensive experiments show that our approach is able to represent a 3D character in real scenes with unprecedented quality, which can be not only rendered into novel views in real-time speed but also be controllably guided by different human poses. Still, our method suffers from some artifacts as shown in Fig. 7. Our MPI representation only outputs diffuse colors, and the rendered images thus cannot handle view-dependent effects such as the reflections. Our method also suffers from typical MPI artifacts, such as the layered section exposed near the boundary. Even though our framework can generalize to some new unseen poses fairly well, the quality of synthesis would still degrade, especially when the new pose deviates too far away from the training data. The proposed method is promising to serve as a practical solution for character rendering and we expect further works to solve the remaining issues.

## References

1. Abdelhamed, A., Lin, S., Brown, M.S.: A high-quality denoising dataset for smart-phone cameras. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1692–1700 (2018)
2. Aberman, K., Shi, M., Liao, J., Lischinski, D., Chen, B., Cohen-Or, D.: Deep video-based performance cloning. In: Computer Graphics Forum. vol. 38, pp. 219–233. Wiley Online Library (2019)
3. Bagautdinov, T., Wu, C., Simon, T., Prada, F., Shiratori, T., Wei, S.E., Xu, W., Sheikh, Y., Saragih, J.: Driving-signal aware full-body avatars. *ACM Transactions on Graphics (TOG)* **40**(4), 1–17 (2021)
4. Balakrishnan, G., Zhao, A., Dalca, A.V., Durand, F., Gutttag, J.: Synthesizing images of humans in unseen poses. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8340–8348 (2018)
5. Carranza, J., Theobalt, C., Magnor, M.A., Seidel, H.P.: Free-viewpoint video of human actors. *ACM transactions on graphics (TOG)* **22**(3), 569–577 (2003)
6. Casas, D., Volino, M., Collomosse, J., Hilton, A.: 4d video textures for interactive character appearance. In: Computer Graphics Forum. vol. 33, pp. 371–380. Wiley Online Library (2014)
7. Chan, C., Ginosar, S., Zhou, T., Efros, A.A.: Everybody dance now. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5933–5942 (2019)
8. Collet, A., Chuang, M., Sweeney, P., Gillett, D., Evseev, D., Calabrese, D., Hoppe, H., Kirk, A., Sullivan, S.: High-quality streamable free-viewpoint video. *ACM Transactions on Graphics (ToG)* **34**(4), 1–13 (2015)
9. Flynn, J., Broxton, M., Debevec, P., DuVall, M., Fyffe, G., Overbeck, R., Snavely, N., Tucker, R.: Deepview: View synthesis with learned gradient descent. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2367–2376 (2019)
10. Gao, C., Saraf, A., Kopf, J., Huang, J.B.: Dynamic view synthesis from dynamic monocular video. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5712–5721 (2021)
11. Guo, K., Lincoln, P., Davidson, P., Busch, J., Yu, X., Whalen, M., Harvey, G., Orts-Escolano, S., Pandey, R., Dourgarian, J., et al.: The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Transactions on Graphics (ToG)* **38**(6), 1–19 (2019)
12. Habermann, M., Liu, L., Xu, W., Zollhoefer, M., Pons-Moll, G., Theobalt, C.: Real-time deep dynamic characters. *ACM Transactions on Graphics (TOG)* **40**(4), 1–16 (2021)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
14. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017)
15. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
16. Li, J., Feng, Z., She, Q., Ding, H., Wang, C., Lee, G.H.: Mine: Towards continuous depth mpi with nerf for novel view synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12578–12588 (2021)

17. Li, T., Slavcheva, M., Zollhoefer, M., Green, S., Lassner, C., Kim, C., Schmidt, T., Lovegrove, S., Goesele, M., Lv, Z.: Neural 3d video synthesis. arXiv preprint arXiv:2103.02597 (2021)
18. Liu, L., Habermann, M., Rudnev, V., Sarkar, K., Gu, J., Theobalt, C.: Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Transactions on Graphics (TOG)* **40**(6), 1–16 (2021)
19. Liu, L., Xu, W., Habermann, M., Zollhöfer, M., Bernard, F., Kim, H., Wang, W., Theobalt, C.: Neural human video rendering by learning dynamic textures and rendering-to-video translation. arXiv preprint arXiv:2001.04947 (2020)
20. Lombardi, S., Saragih, J., Simon, T., Sheikh, Y.: Deep appearance models for face rendering. *ACM Transactions on Graphics (ToG)* **37**(4), 1–13 (2018)
21. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)* **34**(6), 1–16 (2015)
22. Lugaresi, C., Tang, J., Nash, H., McClanahan, C., Uboweja, E., Hays, M., Zhang, F., Chang, C.L., Yong, M.G., Lee, J., et al.: Mediapipe: A framework for building perception pipelines. arXiv preprint arXiv:1906.08172 (2019)
23. Ma, L., Jia, X., Sun, Q., Schiele, B., Tuytelaars, T., Van Gool, L.: Pose guided person image generation. *Advances in neural information processing systems* **30** (2017)
24. Martin-Brualla, R., Pandey, R., Yang, S., Pidlypenskyi, P., Taylor, J., Valentin, J., Khamis, S., Davidson, P., Tkach, A., Lincoln, P., et al.: Lookingood: Enhancing performance capture with real-time neural re-rendering. arXiv preprint arXiv:1811.05029 (2018)
25. Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)* **38**(4), 1–14 (2019)
26. Ouyang, H., Shi, Z., Lei, C., Law, K.L., Chen, Q.: Neural camera simulators. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7700–7709 (2021)
27. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5865–5874 (2021)
28. Park, K., Sinha, U., Hedman, P., Barron, J.T., Bouaziz, S., Goldman, D.B., Martin-Brualla, R., Seitz, S.M.: Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. arXiv preprint arXiv:2106.13228 (2021)
29. Peng, S., Dong, J., Wang, Q., Zhang, S., Shuai, Q., Bao, H., Zhou, X.: Animatable neural radiance fields for human body modeling. arXiv e-prints pp. arXiv–2105 (2021)
30. Porter, T., Duff, T.: Compositing digital images. In: *Proceedings of the 11th annual conference on Computer graphics and interactive techniques*. pp. 253–259 (1984)
31. Pumarola, A., Agudo, A., Sanfeliu, A., Moreno-Noguer, F.: Unsupervised person image synthesis in arbitrary poses. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 8620–8628 (2018)
32. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10318–10327 (2021)
33. Sarkar, K., Mehta, D., Xu, W., Golyanik, V., Theobalt, C.: Neural re-rendering of humans from a single image. In: *European Conference on Computer Vision*. pp. 596–613. Springer (2020)

34. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4104–4113 (2016)
35. Shysheya, A., Zakharov, E., Aliev, K.A., Bashirov, R., Burkov, E., Iskakov, K., Ivakhnenko, A., Malkov, Y., Pasechnik, I., Ulyanov, D., et al.: Textured neural avatars. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2387–2397 (2019)
36. Si, C., Wang, W., Wang, L., Tan, T.: Multistage adversarial losses for pose-based human image synthesis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 118–126 (2018)
37. Srinivasan, P.P., Tucker, R., Barron, J.T., Ramamoorthi, R., Ng, R., Snavely, N.: Pushing the boundaries of view extrapolation with multiplane images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 175–184 (2019)
38. Tewari, A., Fried, O., Thies, J., Sitzmann, V., Lombardi, S., Sunkavalli, K., Martin-Brualla, R., Simon, T., Saragih, J., Nießner, M., et al.: State of the art on neural rendering. In: Computer Graphics Forum. vol. 39, pp. 701–727. Wiley Online Library (2020)
39. Tewari, A., Thies, J., Mildenhall, B., Srinivasan, P., Tretschk, E., Wang, Y., Lassner, C., Sitzmann, V., Martin-Brualla, R., Lombardi, S., et al.: Advances in neural rendering. arXiv preprint arXiv:2111.05849 (2021)
40. Tretschk, E., Tewari, A., Golyanik, V., Zollhöfer, M., Lassner, C., Theobalt, C.: Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12959–12970 (2021)
41. Tucker, R., Snavely, N.: Single-view view synthesis with multiplane images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 551–560 (2020)
42. Wang, T.C., Liu, M.Y., Zhu, J.Y., Liu, G., Tao, A., Kautz, J., Catanzaro, B.: Video-to-video synthesis. arXiv preprint arXiv:1808.06601 (2018)
43. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8798–8807 (2018)
44. Weng, C.Y., Curless, B., Srinivasan, P.P., Barron, J.T., Kemelmacher-Shlizerman, I.: Humannerf: Free-viewpoint rendering of moving people from monocular video. arXiv preprint arXiv:2201.04127 (2022)
45. Wizaradwongsa, S., Phongthawee, P., Yenphraphai, J., Suwajanakorn, S.: Nex: Real-time view synthesis with neural basis expansion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8534–8543 (2021)
46. Xu, H., Alldieck, T., Sminchisescu, C.: H-nerf: Neural radiance fields for rendering and temporal reconstruction of humans in motion. *Advances in Neural Information Processing Systems* **34** (2021)
47. Zhang, J., Liu, X., Ye, X., Zhao, F., Zhang, Y., Wu, M., Zhang, Y., Xu, L., Yu, J.: Editable free-viewpoint video using a layered neural representation. *ACM Transactions on Graphics (TOG)* **40**(4), 1–18 (2021)
48. Zhang, P., Zhang, B., Chen, D., Yuan, L., Wen, F.: Cross-domain correspondence learning for exemplar-based image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5143–5153 (2020)
49. Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: Learning view synthesis using multiplane images. arXiv preprint arXiv:1805.09817 (2018)

50. Zhou, X., Zhang, B., Zhang, T., Zhang, P., Bao, J., Chen, D., Zhang, Z., Wen, F.: Cocosnet v2: Full-resolution correspondence learning for image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11465–11475 (2021)