

FlexNeRF: Photorealistic Free-viewpoint Rendering of Moving Humans from Sparse Views

Vinoj Jayasundara^{1*}, Amit Agrawal², Nicolas Heron², Abhinav Shrivastava¹, Larry S. Davis^{1,2}

¹University of Maryland, College Park ²Amazon.com, Inc.

{vinoj, lsdavis}@umd.edu, {aaagrawa, heron}@amazon.com, abhinav@cs.umd.edu

Abstract

We present *FlexNeRF*, a method for photorealistic free-viewpoint rendering of humans in motion from monocular videos. Our approach works well with sparse views, which is a challenging scenario when the subject is exhibiting fast/complex motions. We propose a novel approach which jointly optimizes a canonical time and pose configuration, with a pose-dependent motion field and pose-independent temporal deformations complementing each other. Thanks to our novel temporal and cyclic consistency constraints along with additional losses on intermediate representation such as segmentation, our approach provides high quality outputs as the observed views become sparser. We empirically demonstrate that our method significantly outperforms the state-of-the-art on public benchmark datasets as well as a self-captured fashion dataset. The project page is available at: <https://flex-nerf.github.io/>.

1. Introduction

Free-viewpoint rendering of a scene is an important problem often attempted under constrained settings: on subjects demonstrating simple motion carefully captured with multiple cameras [17, 19, 20]. However, photorealistic *free-viewpoint* rendering of moving humans captured from a *monocular* video still remains an unsolved challenging problem, especially with sparse views.

Neural radiance fields (NeRF) have emerged as a popular tool to learn radiance fields from images/videos for novel view-point rendering. Previous approaches assume multiple view-points and often fail on non-rigid human motions. Human-specific NeRFs have recently become popular for learning models using input videos [27, 40]. The current state-of-art approaches such as HumanNeRF [40] have shown impressive progress in this domain. However, there remain several challenges. Firstly, approaches such as HumanNeRF [40] utilize a pose prior and use a canoni-

cal configuration (*e.g.* T-pose) for optimization, which may be well outside the set of observed poses. The underlying optimization becomes challenging especially as the number of observed views become sparse. In contrast, we select a pose from the available set of poses as the canonical pose-configuration, similar to previous pose-free approaches such as D-NeRF [29]. This enables best of both worlds; it becomes easier to learn a motion field mapping due to smaller deformations while using a pose prior. In addition, having the canonical view in the training data provides a strong prior for the optimization of the canonical pose itself. Finally, it allows us to optimize the canonical configuration with our novel *pose-independent temporal deformation*. We demonstrate that this architectural change provides significantly better results compared to existing approaches [24, 40].

In addition, approaches such as HumanNeRF [40] depend on the estimated pose for the canonical configuration optimization. Errors in the initial pose estimation, for example, due to strong motion blur cause challenges in pose correction. The underlying assumption that the non-rigid motion is pose-dependent often fails in scenarios with complex clothing and accessories, hair styles, and large limb movements. Our proposed *pose-independent temporal deformation* helps to supplement the missing information in its pose-dependent counterpart.

To this end, we introduce FlexNeRF, a novel approach for jointly learning a *pose-dependent motion field* and *pose-independent temporal deformation* within the NeRF framework for modeling human motions. Moreover, we introduce a novel cycle consistency loss in our framework, further capitalizing on the fact that our canonical pose corresponds to one of the captured frames. The consistency regularizes the estimated deformation fields by mapping back and forth between each view and the canonical pose. Moreover, the information content of any frame in a motion sequence has a strong similarity to that of its neighbours. Hence, we propose to utilize this contextual information present in a consecutive set of frames to aid learning by

*Part of the work was done while the author was an intern at Amazon.

imposing a temporal consistency loss. We additionally regularize the training by adding a supplementary loss based on the segmentation masks. Our approach allows photo-realistic rendering of a moving human even when sparse views are available, by supplementing the pose-dependent motion field with additional information during learning: (i) pose-independent temporal deformation with ample pixel-wise correspondences beyond the (typically 24) pose point-correspondences, and (ii) consistency constraints/losses. In summary, our paper makes the following contributions:

- We propose a novel approach to learn pose-independent temporal deformation to complement the pose-dependent motion for modeling humans in video, using one of the views as the canonical view.
- We propose a novel cyclic-consistency loss to regularize the learned deformations.
- We propose a temporal-consistency loss to aid learning with contextual information present in neighbouring frames, as well as to maintain consistency across consecutive rendered frames.
- Our approach outperforms the state-of-the-art approaches, with significant improvement in case of sparse views.

2. Related Work

2.1. Neural Radiance Fields (NeRFs)

NeRFs attempt to learn a scene representation for novel-view synthesis by modeling the radiance field with learnable functions. A variety of approaches have been proposed recently for neural rendering, exploiting voxel grids [5, 34], neural textures [33, 37], point-clouds [1, 21], and neural implicit functions [4, 25].

The landscape of neural rendering changed with NeRF [22], which proposed a simple, yet revolutionary approach for photorealistic novel-view synthesis of static scenes. NeRF attempts to map from the 5-d light fields to 4-d space consisting of color c (RGB) and density σ : likelihood that the light ray at this 5-d co-ordinate is terminated by occlusion. Since the introduction of original NeRF formulation, several variations and improvements [15, 30, 36, 44] have been proposed.

2.2. Neural Rendering of Dynamic Scenes

While originally proposed for static scenes, NeRF based approaches have been recently extended to dynamic scenes, for both rigid and non-rigid objects. These approaches can be divided into two main categories: a) optimizing a canonical configuration, and b) directly optimizing the 4-D spatio-temporal scenes. D-NeRF [29] is an example of the first category, which attempts to map each observed frame to a

given canonical frame. Once the canonical scene has been optimized for all available views, the novel-view can be rendered from the canonical space, and mapped back to the observed space. The same approach can be seen applied to videos with simple motion and other settings [3, 26, 38]. In contrast, the approaches that directly estimate spatio-temporal scene representations [14, 42] takes positional-encoded or latent-coded time t as an input in-addition to the spatial inputs, and attempts to predict the color and the density along each ray.

2.3. Neural Rendering of Human Subjects

Compared to general rendering of dynamic scenes, human subject-specific rendering has additional challenges in terms of complex non-rigid human motions. Priors such as human pose that can provide additional information for successful scene representation. Hence, most methods [27, 28, 41] begin with assuming *SMPL* template as a prior [18]. Furthermore, most methods use multi-view videos [16, 23, 39, 43]. A few recent methods including HumanNeRF [40] and others [3, 8, 38] use monocular videos, whereas only the former attempts free-viewpoint rendering. However, these approaches have challenges rendering photorealistic outputs with sparse input views. We consider HumanNeRF [40] as the closet work to ours and address the aforementioned challenges.

3. Method

Given a sequence of frames of a monocular video with a human manifesting complex motions, our goal is to achieve photo-realistic free-viewpoint rendering and reposing. We choose a frame as the canonical configuration (*e.g.* the mid-point of the motion sequence) and learn it via: a) pose-dependent (rigid and non-rigid) motion fields and b) pose-independent temporal deformations.

3.1. Pose-Dependent Motion Fields

Given a canonical pose-configuration $p^c = (J^c, \Omega^c)$ and the observed pose $p = (J, \Omega)$, where Ω represents the local joint rotations and J represents the joint locations in 3D, we define a pose-guided motion field mapping between the observed and canonical spaces. We first compute the transformation $M_k(p^c, p)$, and hence the translation t_k and rotation R_k matrices between the joint coordinates in observed and canonical spaces, for a given body part k . $Y(w_i, j_i)$ computes the exponent of the local joint rotation w_i of the joint location j_i using the Rodrigues’s formula [35],

$$Y(w_i, j_i) = \prod_{i \in \tau(k)} \begin{bmatrix} \exp(w_i) & j_i \\ 0 & 1 \end{bmatrix}, \quad (1)$$

where $\tau(k)$ denotes the ordered set of parents of the k^{th} local joint. Subsequently, we compute the corresponding

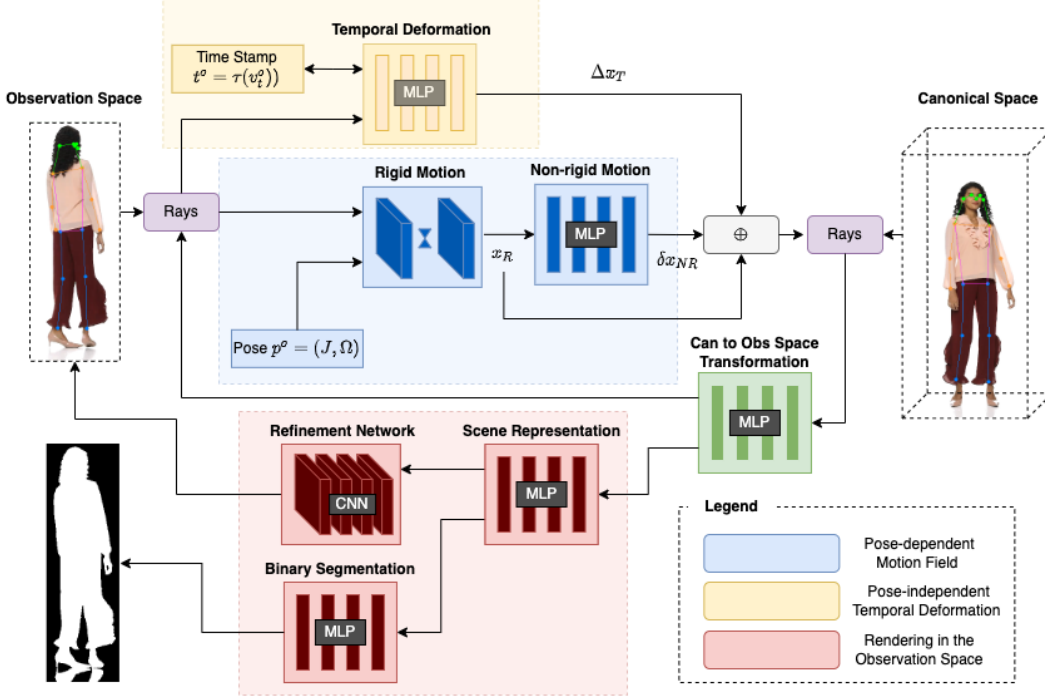


Figure 1. Overview of our approach. Pose-independent temporal deformation is used in conjunction with pose-dependent motion fields (rigid and non-rigid). We choose one of the input frames as the canonical view, allowing us to use cyclic-consistency for regularization.

translation t_k and rotation R_k matrices,

$$M_k(p^c, p) = Y(\omega_i^c, j_i^c) \{Y(\omega_i, j_i)\}^{-1} = \begin{bmatrix} R_k & t_k \\ 0 & 1 \end{bmatrix}. \quad (2)$$

Given the translation and rotation matrices, we compute the rigid deformation x_R between the observed and canonical spaces by defining

$$\mathcal{L}(x) = \sum_{k=1}^K w_k^c (R_k x + t_k), \quad (3)$$

which represents the likelihood that the position x is a part of the subject. We obtain the set of blend weight volumes in the canonical space $\{w_k^c\}_{k=1}^K$, where K is the total number of 3D joint locations. To this end, starting from a constant random latent vector z , we generate the motion weight volume $W^c(x) = CNN_{\theta_R}(x; z) \in \mathbb{R}^4$ by optimizing the parameters θ_R of the CNN_{θ_R} [40]. We add a computed approximate Gaussian bone volume as a motion weight volume prior to the output of the last transposed convolution layer before activation. Subsequently, we compute the rigid deformation x_R with the obtained $\mathcal{L}(x)$ and $W^c(x)$,

$$x_R = \frac{\sum_{k=1}^K w_k^c (R_k x^o + t_k)^2}{\mathcal{L}(x)}. \quad (4)$$

The non-rigid deformation between the observed and canonical spaces is then computed as a pose-guided offset

δx_{NR} to the rigid deformation x_R . We feed the positional encoding $\tau(x_R)$ to the non-rigid motion MLP as,

$$\delta x_{NR} = MLP_{\theta_{NR}}(\gamma(x_R); \Omega). \quad (5)$$

We follow the approach defined in [22] to obtain the positional encoding $\tau(x)$ of the position x . The non-rigid motion MLP consists of six fully-connected layers with the positional encoding $\tau(x_R)$ and the local joint rotations Ω (without global rotation) as the inputs with $\tau(x_R)$ skip-connected to the fifth layer to generate the offset. Since the initial pose estimate p obtained from off-the-shelf techniques such as SPIN [12] or VIBE [11] can be erroneous, we perform a pose correction following [40].

3.2. Pose-Independent Temporal Deformation

We strategically set the canonical configuration to an observed frame in the training set, allowing us to access observed (x^o) and canonical (x^c) positions as a source of information. Furthermore, having a common canonical anchor when learning a dynamic setting ensures that the scene is inter-connected across frames and no longer independent between time instances, which is intuitive and known to provide quality performance [29]. This grounding aids the model to learn preserving temporal consistency of the dynamic scene. Nevertheless, such an approach which optimizes a canonical time configuration does not work well alone for free-viewpoint rendering where we are required to render a 360° camera path with complex motion. Hence, we utilize a combined approach of pose-guided

and pose-independent (time-guided) canonical configuration optimization. Results in Sec. 5 show that the combined approach allows high quality photorealistic rendering with sparse views.

We compute the pose-independent temporal deformation between a point position in the observation space x^o to the canonical space x^c with a temporal deformation MLP, similar to D-NeRF [29]. This temporal deformation Δx_T is defined by,

$$\Delta x_T = MLP_{\theta_{TD}}(\gamma(x^o), \gamma(x^c); (t^o, t^c)), \quad (6)$$

where t^o is the observed time stamp defined by ($t^o = \tau(v_t^o$)), and t^c is the canonical time stamp defined by ($t^c = \tau(v_t^c$)). $v_t \in R^5$ is a learnable vector representation initialised proportional to the frame sequence index of the monocular video.

In contrast to D-NeRF [29], we set the temporal vectors to be learnable due to several reasons. Even though the progression of frame sequence indices are linear, the progression of temporal information throughout a video is highly non-linear. For instance, there can be rapid motion between two consecutive frames in a video, whereas there can be no motion between another two consecutive frames in the same video. Hence, it is not intuitive to allocate a linear representation to the temporal vectors $\{v_t\}$. Furthermore, albeit being a contrasting approach to ours, DyNeRF [13] presents strong evidence that trainable (latent) codes can better handle complex scene dynamics such as large deformations and topological changes. We heavily regularize the training of these temporal vectors in order to ensure that they are contained within practical limits.

The temporal deformation MLP, $MLP_{\theta_{TD}}$ consists of 8 fully connected layers with the positional encoded temporal vectors ($\tau(v_t^o$) and $\tau(v_t^c$)), and the positional encoded point position vectors $\tau(x^o)$ and $\tau(x^c)$ as inputs. The observed encodings $\tau(v_t^o$) and $\tau(x^o)$ are skip connected to the fifth layer to generate the deformation Δx_T . Finally, we aggregate the pose-guided rigid motion x_R , pose-guided non-rigid motion δx_{NR} (as an offset to x_R), and the pose-independent temporal deformation Δx_T to produce the predicted canonical configuration \hat{x}^c ,

$$\hat{x}^c = \underbrace{(x_R + \delta x_{NR})}_{\text{pose-guided motion field}} + \underbrace{\Delta x_T}_{\text{pose-independent temporal deformation}} \quad (7)$$

3.3. Cyclic Consistency

Having obtained the predicted canonical configuration \hat{x}^c , we predict the RGB color c and the density σ at a given spatial location. Rather than directly predicting (c, σ) from the canonical space similar to the existing approaches [29, 40, 43], we propose to break the prediction process in two steps: a) transformation from canonical (\hat{x}^c) to observed (\hat{x}^o) space, and b) (c, σ) prediction from \hat{x}^o .

The proposed approach yields the opportunity to enforce a cyclic consistency constraint (observed $x^o \rightarrow$ canonical $\hat{x}^c \rightarrow$ observed \hat{x}^o) on the output of the canonical to observed transformation MLP, $\hat{x}^o = MLP_{\theta_{CO}}(\gamma(\hat{x}^c))$. Furthermore, having two separate specialized networks rather than one network to map from the rays in the canonical space to (c, σ) in the observation space is more flexible and is empirically more effective as shown in Sec. 5.

The $MLP_{\theta_{CO}}$ has a similar architecture to $MLP_{\theta_{TD}}$, without the temporal vectors as inputs. The subsequent scene representation MLP, $(c, \sigma) = MLP_{\theta_c}(\gamma(\hat{x}^o))$ has a similar architecture to the network proposed in [22]. Prior to feeding \hat{x}^c and \hat{x}^o to the corresponding networks, each vector is positional encoded.

3.4. Volume Rendering and Refinement Network

We follow the volume rendering approach described in NeRF [22] by defining the expected alpha (density) mask $\mathcal{A}(r)$ and the expected color $\mathbf{C}(r)$ for a give ray r ,

$$\mathcal{A}(r) = \sum_{i=1}^D \left\{ \prod_{j=1}^{i-1} (1 - \alpha_j) \right\} \alpha_i \quad (8)$$

$$\mathbf{C}(r) = \sum_{i=1}^D \left\{ \prod_{j=1}^{i-1} (1 - \alpha_j) \right\} \alpha_i c(x_i) \quad (9)$$

$$\alpha_i = \mathcal{L}(x_i) \{1 - \exp(-\sigma(x_i) \Delta z_i)\}, \quad (10)$$

where D is the number of samples, and Δz_i is the interval between consecutive samples. We employ the same stratified sampling approach described in [22].

To further enhance the photorealism of the rendered images, we use a refinement network $\hat{I}_o = CNN_{\theta_{FT}}(\mathbf{C}(r), \mathcal{A}(r))$ to add fine-grained details to the rendered image, similar to latent diffusion approaches [31]. The refinement network $CNN_{\theta_{FT}}$ consists three transposed convolution layers and outputs the final rendered image \hat{I}_o .

3.5. Rendering Segmentation Mask

The segmentation masks for the input frames can be obtained using an off-the-shelf segmentation network [9]. We use them to apply an additional loss to improve the density estimation. Note that rendering $\mathcal{A}(r)$ results in the predicted segmentation $\hat{M} = \mathcal{A}(r)$, which is compared against the real segmentation mask M . This helps to eliminate the halo effects [28, 40] and provide sharper boundaries. Empirically, we observed that thresholding the predicted segmentation mask, $\hat{M} = \mathcal{A}(r)H(\mathcal{A}(r), b)$ works better, where b is a threshold value and

$$H(\mathcal{A}(r), b) = \begin{cases} 1 & \text{if } \mathcal{A}(r) > b \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

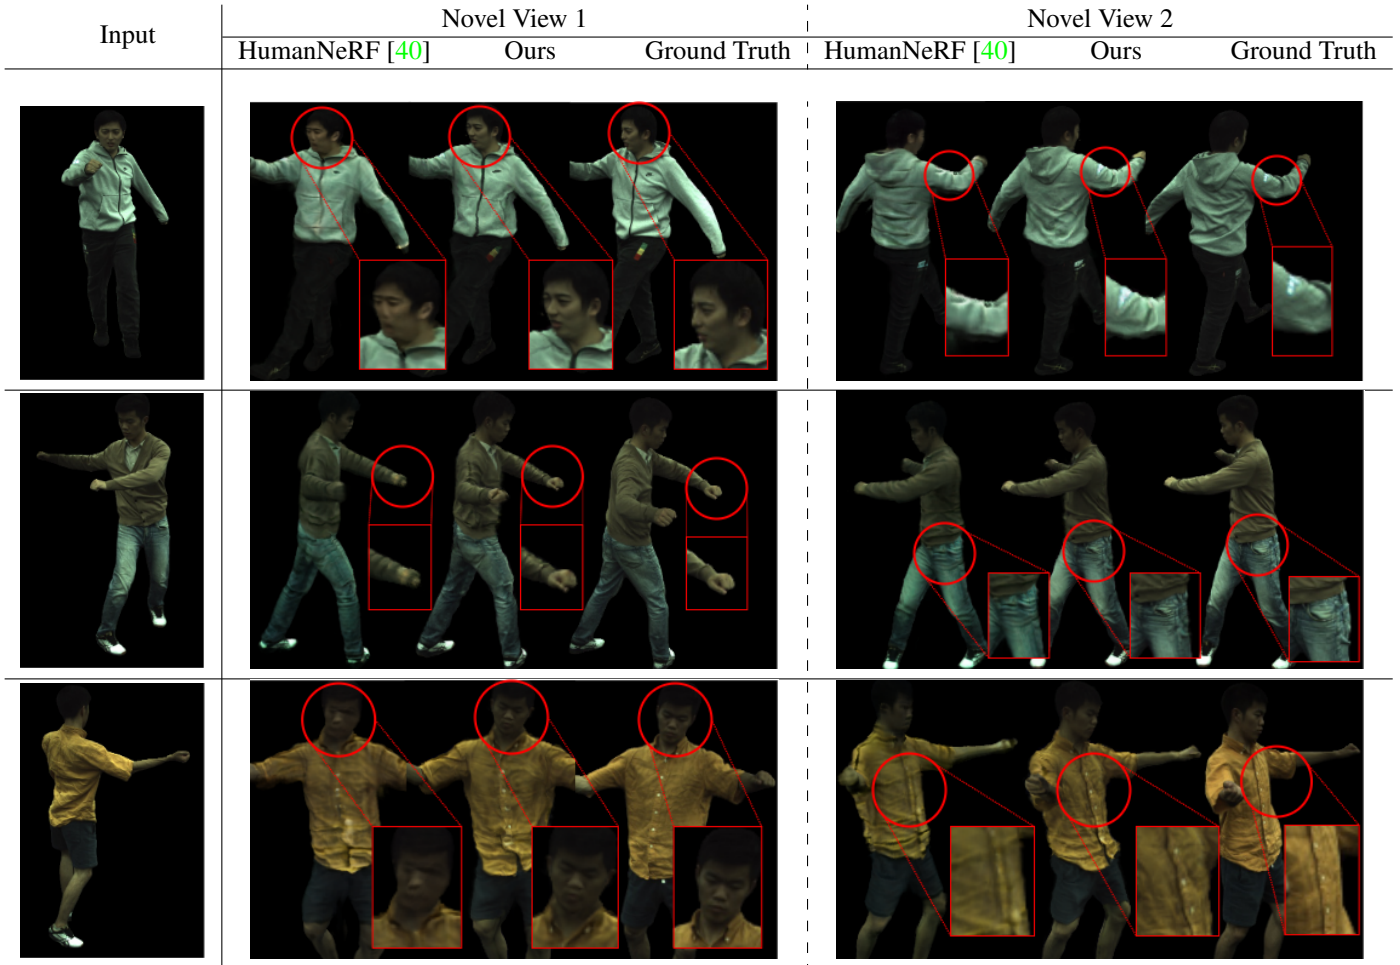


Figure 2. Qualitative comparison of rendered novel views on the ZJU-MoCap dataset. Notice the higher quality of rendered images from our method on details such as faces, buttons on shirt, etc.

However, using a fixed threshold b makes learning difficult at the start of training. To ease learning, we make b a learnable parameter and re-define $\hat{M} = (\mathcal{A}(r) + b)H(\mathcal{A}(r), b)$, so that it is differentiable with respect to b . We observe that b goes to 0 as training progresses, as in the ideal case. Compared to previous approaches such as [10], this does not require us to depend on estimated depths, which could themselves be erroneous due to complex non-rigid motions.

4. Learning the FlexNeRF model

In this section we describe the loss functions used to learn the FlexNeRF model and discuss details with respect to optimization and ray-sampling.

4.1. Loss Functions

NeRFs are typically trained with a combination of losses between the rendered and observed frames. In addition, FlexNeRF also uses a combination of segmentation loss,

cyclic consistency loss and temporal consistency loss as defined below.

Segmentation Loss: We apply the BCE-Dice loss between the predicted and ground truth binary segmentation masks

$$\mathbb{L}_S = \frac{1}{N} \sum \left[M \log \hat{M} + (1 - M) \log (1 - \hat{M}) \right] + \frac{2|M \cap \hat{M}|}{|M| + |\hat{M}|}, \quad (12)$$

where N is the number of pixels in the segmentation mask.

Cyclic Consistency Loss: We introduce a cyclic consistency constraint on the canonical to observation space transformation, using Mean Squared Error (MSE) between \hat{x}^o and x^o defined by,

$$\mathbb{L}_{CCL} = \frac{1}{L} \sum_{i=1}^L (\hat{x}_i^o - x_i^o)^2, \quad (13)$$

where L is the number of positional samples.

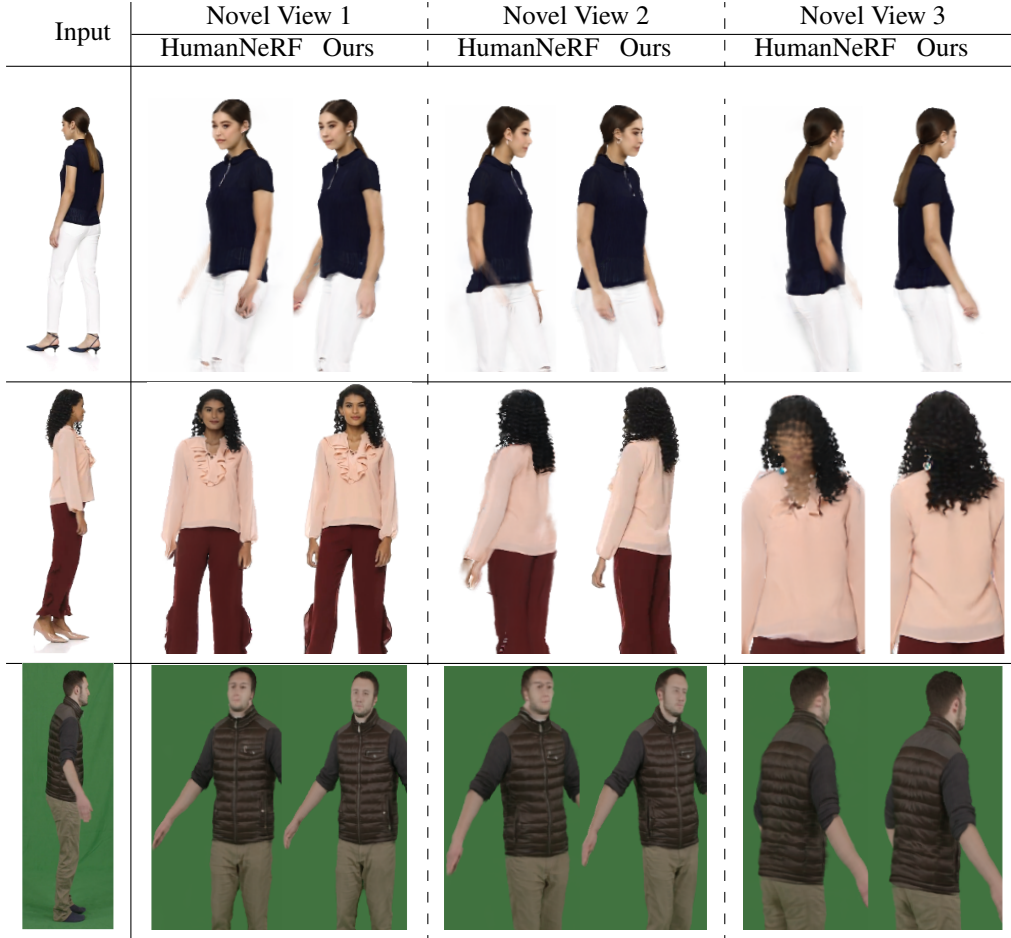


Figure 3. Qualitative comparison of novel rendered views on SCF dataset (top two rows) and the People Snapshot dataset (bottom row) using sparse views. Our approach significantly improves the results.

Dataset	Views	Method	LPIPS $\times 10^3$ \downarrow	PSNR \uparrow	SSIM \uparrow
PeopleSnapshot [2]	Sparse ‡	HumanNeRF [40]	39.27	27.65	0.8816
		Ours	37.11	28.09	0.9003
	Full	Neural Body [28]	57.67*	24.62	0.8490
		HumanNeRF [40]	36.79	28.05	0.8984
		Ours	35.63	28.77	0.9043
ZJU-MoCap [7, 28]	Sparse ‡	HumanNeRF [40]	36.02	29.82	0.9597
		Ours	31.68	30.18	0.9685
	Full	Neural Body [28]	52.28	29.07	0.9615
		HumanNeRF [40]	31.72	30.24	0.9679
		Ours	29.01	31.73	0.9765
SCF Dataset †	Sparse ‡	Neural Body [28]	48.62	25.07	0.9131
		HumanNeRF [40]	39.71	26.12	0.9366
		Ours	34.26	29.55	0.9627

Table 1. Comparison of performance across benchmark datasets. * refers to adjusted LPIPS from the values reported in [43] to fit the same scale as our experiments. \dagger refers to the Self-Captured Fashion (SCF) dataset. \ddagger indicates the model trained with sparse (~ 40) views.




Figure 4. LPIPS metric comparison on ZJU-MoCap between HumanNeRF [40] and our method with decreasing number of views.

Temporal Consistency Loss (TCL): We identify that imposing temporal consistency constraints can be valuable at two instances: a) while rendering consecutive training frames $\{\hat{I}^o\}_{t=-k}^k$ and b) while applying temporal deformation from consecutive training frames to the canonical frame $\{\Delta x_T\}_{t=-k}^k$. To this end, we employ the cycle-back regression consistency loss proposed in [6]. The cycle-back regression attempts to determine the temporal proximity of rendered frames or deformation vectors, and penalize the model if they are not in close temporal proximity. Given a rendered frame or a deformation vector u , and neighbors $\{v_k\}$, we compute the similarity vector β_k ,

$$\beta_k = \frac{\exp(-\|u - v_k\|^2)}{\sum_j \exp(-\|u - v_j\|^2)}, \quad (14)$$

where $u, v_k \in \{v_k\}$ and β is a discrete distribution of similarities over time. We impose a Gaussian prior on β by minimizing the normalized square distance,

$$\mu = \sum_k k \beta_k \quad \sigma^2 = \beta_k (k - \mu)^2 \quad (15)$$

$$\mathbb{L}_{TCL} = \frac{|i - \mu|^2}{\sigma^2} + \lambda \log(\sigma), \quad (16)$$

where λ is a regularization parameter. Finally, the rendering loss \mathbb{L}_{rend} , the canonical loss \mathbb{L}_{can} , and the overall loss \mathbb{L} are defined as

$$\mathbb{L}_{rend} = \mathbb{L}_{LPIPS}(\hat{I}^o, I^o) + \mathbb{L}_{MSE}(\hat{I}^o, I^o) + \mathbb{L}_{TCL}(\{\hat{I}^o\}_{t=-k}^k) \quad (17)$$

$$\mathbb{L}_{can} = \mathbb{L}_{MSE}(\hat{x}^c, x^c) + \mathbb{L}_{TCL}(\{\Delta x_T\}_{t=-k}^k) \quad (18)$$

$$\mathbb{L} = \mathbb{L}_{rend} + \mathbb{L}_{can} + \mathbb{L}_{CCL} + \mathbb{L}_S \quad (19)$$

4.2. Optimization Details

Delayed Modular Optimization: We follow a delayed-optimization approach similar to [40] to optimize the non-rigid motion, binary segmentation, and the refinement modules of our method. Optimizing these modules from the beginning yields lower performance as they rely on adequate inputs from the rest of the system. Hence, we freeze these modules initially, and unfreeze them gradually during the course of training.

HumanNeRF Ours HumanNeRF Ours



Figure 5. Comparing the rendering of the canonical view for SCF (left) and ZJU-MoCap (right) datasets. Our approach is able to learn a higher quality canonical view.

Ray Sampling: Since LPIPS use a convolution-based approach to extract features, we use patch-based ray sampling following [32,40] instead of random ray sampling [22] from the whole image.

5. Experiments and Results

5.1. Benchmark Datasets and Metrics

We evaluate the proposed method on two public datasets: ZJU-MoCap [7,28] and People Snapshot [2], and one Self-Captured Fashion (SCF) dataset. The People Snapshot dataset has 7 sequences of monocular videos of human subjects displaying rotating motions in front of a static camera. For the ZJU-MoCap dataset, we use 6 sequences to be compatible with [40] for comparison purposes. We use only the views from the first camera to simulate monocular video settings for training, and use the views from rest of the cameras for evaluation. Videos from both ZJU-MoCap and People Snapshot datasets are carefully captured under lab settings.

For the SCF dataset, we captured 7 sequences of monocular videos freely where the movements are solely up to the discretion of the consenting subjects. In contrast to the public datasets: a) the subjects are wearing complex clothing and accessories, and the movements are fast, b) videos are captured without any controlled settings, and c) the captured videos are brief with only one full rotation. Due to the absence of ground truths for the SCF and People Snapshot datasets, we evaluate the rendered frames by holding random frames from training. We use LPIPS, PSNR, and SSIM [44] as evaluation metrics.

5.2. Results and Analysis

Quantitative Results: Table 1 compares our method against HumanNeRF [40] and Neural-Body [28] across the three datasets. We consider two settings: *Full*, using all the frames and *Sparse*, using a sparse number of views. To generate the *Sparse* setting, we remove stationary frames from the video and any subsequent frames after the first complete rotation of the subject. Subsequently, we re-sample every k^{th} frame, where k is chosen such that $\sim 8 - 10\%$ of the original number of frames remain. Our approach outperforms HumanNeRF and NeuralBody across all metrics

Ablation	LPIPS $\times 10^3$ \downarrow	Δ	PSNR \uparrow	Δ	SSIM \uparrow	$\Delta \times 10^{-2}$
HumanNeRF	39.71	0.00	26.12	0.00	0.9366	0.00
Ours (full)	34.26	-5.45	29.55	+3.43	0.9627	+2.61
Ours (w/o PID)	37.12	-2.59	27.42	+1.30	0.9469	+1.03
Ours (w/o CO-MLP + CCL)	36.86	-2.85	28.50	+2.38	0.9509	+1.43
Ours (w/o TCL)	36.17	-3.54	27.94	+1.82	0.9521	+1.55
Ours (w/o RF)	35.04	-4.67	29.17	+3.05	0.9593	+2.27
Ours (w/o BS)	35.78	-3.93	29.07	+2.95	0.9601	+2.35

Table 2. Ablation: Effect of removing various modules and losses from our full approach on the SCF dataset. *PID*: Pose-Independent temporal Deformation module, *CO-MLP*: Canonical-to-Observed transformation MLP, *CCL*: Cyclic Consistency Loss, *RF*: Refinement MLP, *BS*: Binary Segmentation Loss, *TCL*: Temporal Consistency Loss.




Figure 6. Each cell shows the original training frame (left) with rendered frame using the *same* viewpoint after training. Note that for difficult/challenging poses, HumanNeRF (right) fails to minimize the training loss compared to ours (middle).

in the *Full* setting. Since HumanNeRF is significantly better than Neural Body using all the frames, we only compare with HumanNeRF for sparse view setting. Our approach also outperforms HumanNeRF in the *Sparse* setting for all three datasets as shown in Table 1.

Analysis of Number of Views: To further analyze the effect of the number of views, we train various models using different number of input views. Figure 4 compares the performance of LPIPS metric for our method against HumanNeRF [40] with varying number of views. Our approach is better than HumanNeRF for all settings, with significant reduction in LPIPS metric as the number of views decreases.

Quality of Canonical View: An interesting analysis is to visualize the quality of the canonical view rendering itself. This indicates how well the model can learn deformations and fuse information from various frames. As shown in Figure 5, our approach can produce significantly higher quality of rendering for canonical view, thanks to our proposed pose-independent temporal deformation.

Quality of Rendered Training View: We can infer how well the NeRF model is minimizing the training loss by rendering the trained model using the same viewpoint as

input training frames. Figure 6 shows the best and worst frames for training loss minimization. Notice that for easy poses (frontal), the rendered training frame from both HumanNeRF and our method are visually similar to the original training frame. However, for challenging poses (side view), the quality and accuracy of rendered training frame is significantly higher for our method compared to HumanNeRF. In fact, in the example shown in Figure 6 (top row), HumanNeRF totally fails to render the correct pose of the training frame itself, hindering overall learning. This again highlights the pitfalls of pose-dependent learning, and indicates that our method can better transform input frames to canonical frame, thanks to the proposed cyclic consistency constraint and pose-independent temporal deformation. Moreover, our method can render neighbouring frames that are not used for training in the *sparse* setting by passing the interpolated frame index. However, both our method and HumanNeRF share similar challenges generating completely unseen poses *beyond* the range in the input video.

Qualitative Results: Figs. 2 and 3 show qualitative results for rendering from novel view-points for the ZJU-MoCap and SCF datasets respectively. Our approach renders higher quality novel-views compared to HumanNeRF (on faces, buttons on t-shirts, etc.).

Ablations: Table 2 shows the effect of removing various modules from our full approach. We observe that all the proposed losses and constraints contribute to the performance improvement.

6. Conclusion

In this work, we presented FlexNeRF: a novel method for photorealistic free-viewpoint rendering of moving humans starting with monocular videos. Our proposed framework utilizes pose-independent temporal deformation along with cycle consistency to better model the complex human motions with sparse input views. Experiments demonstrated that our method outperforms the state-of-the-art approaches in rendering novel viewpoints. We hope that our work will spark future research in the challenging problem of free-viewpoint rendering from in-the-wild videos.

Acknowledgments: This project was partially funded by the DARPA SemaFor (HR001119S0085) program.

References

- [1] Kara-Ali Aliev, Dmitry Ulyanov, and Victor S. Lempitsky. Neural point-based graphics. *ArXiv*, abs/1906.08240, 2020. [2](#)
- [2] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8387–8397, Jun 2018. CVPR Spotlight Paper. [6](#), [7](#)
- [3] Jianchuan Chen, Ying Zhang, Di Kang, Xuefei Zhe, Linchao Bao, and Huchuan Lu. Animatable neural radiance fields from monocular rgb video. *ArXiv*, abs/2106.13629, 2021. [2](#)
- [4] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5932–5941, 2019. [2](#)
- [5] Frank Dellaert and Yen-Chen Lin. Neural volume rendering: Nerf and beyond. *ArXiv*, abs/2101.05204, 2021. [2](#)
- [6] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. Temporal cycle-consistency learning. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1801–1810, 2019. [7](#)
- [7] Qi Fang, Qing Shuai, Junting Dong, Hujun Bao, and Xiaowei Zhou. Reconstructing 3d human pose by watching humans in the mirror. In *CVPR*, 2021. [6](#), [7](#)
- [8] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5692–5701, 2021. [2](#)
- [9] Haoyu He, Jing Zhang, Qiming Zhang, and Dacheng Tao. Grapy-ml: Graph pyramid mutual learning for cross-dataset human parsing. In *AAAI*, 2020. [4](#)
- [10] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. *ArXiv*, abs/2203.12575, 2022. [5](#)
- [11] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. Vibe: Video inference for human body pose and shape estimation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5252–5262, 2020. [3](#)
- [12] Nikos Kolotouros, Georgios Pavlakos, Michael J. Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2252–2261, 2019. [3](#)
- [13] Tianye Li, Miroslava Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, S. Lovegrove, Michael Goesele, Richard A. Newcombe, and Zhaoyang Lv. Neural 3d video synthesis from multi-view video. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5511–5521, 2022. [4](#)
- [14] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6494–6504, 2021. [2](#)
- [15] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *ArXiv*, abs/2007.11571, 2020. [2](#)
- [16] Lingjie Liu, Marc Habermann, V. Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ArXiv*, abs/2106.02019, 2021. [2](#)
- [17] Lingjie Liu, Weipeng Xu, Marc Habermann, Michael Zollhöfer, Florian Bernard, Hyeonwoo Kim, Wenping Wang, and Christian Theobalt. Neural human video rendering by learning dynamic textures and rendering-to-video translation. *IEEE transactions on visualization and computer graphics*, PP, 2020. [1](#)
- [18] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Smpl: a skinned multi-person linear model. *ACM Trans. Graph.*, 34:248:1–248:16, 2015. [2](#)
- [19] Ricardo Martin-Brualla, Rohit Pandey, Shuoran Yang, Pavel Pidlypenskiy, Jonathan Taylor, Julien P. C. Valentin, S. Khamis, Philip L. Davidson, Anastasia Tkach, Peter Lincoln, Adarsh Kowdle, Christoph Rhemann, Dan B. Goldman, Cem Keskin, Steven M. Seitz, Shahram Izadi, and S. Fanello. Lookingood: Enhancing performance capture with real-time neural re-rendering. *ACM Trans. Graph.*, 37:255, 2018. [1](#)
- [20] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7206–7215, 2021. [1](#)
- [21] Moustafa Meshry, Dan B. Goldman, S. Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. Neural re-rendering in the wild. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6871–6880, 2019. [2](#)
- [22] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. [2](#), [3](#), [4](#), [7](#)
- [23] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5742–5752, 2021. [2](#)
- [24] Hao Ouyang, Bo Zhang, Pan Zhang, Hao Yang, Jiaolong Yang, Dong Chen, Qifeng Chen, and Fang Wen. Real-time neural character rendering with pose-guided multiplane images. *ArXiv*, abs/2204.11820, 2022. [1](#)
- [25] Jeong Joon Park, Peter R. Florence, Julian Straub, Richard A. Newcombe, and S. Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, 2019. [2](#)
- [26] Keunhong Park, U. Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B. Goldman, Steven M. Seitz, and Ricardo Martín Brualla. Deformable neural radiance fields. *ArXiv*, abs/2011.12948, 2020. [2](#)

- [27] Sida Peng, Junting Dong, Qianqian Wang, Shang-Wei Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14294–14303, 2021. [1](#), [2](#)
- [28] 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. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9050–9059, 2021. [2](#), [4](#), [6](#), [7](#)
- [29] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10313–10322, 2021. [1](#), [2](#), [3](#), [4](#)
- [30] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. Derf: Decomposed radiance fields. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14148–14156, 2021. [2](#)
- [31] Robin Rombach, A. Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, 2022. [4](#)
- [32] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *ArXiv*, abs/2007.02442, 2020. [7](#)
- [33] Aliaksandra Shysheya, Egor Zakharov, Kara-Ali Aliev, Renat Bashirov, Egor Burkov, Karim Isakov, Aleksei Ivakhnenko, Yury Malkov, I. Pasechnik, Dmitry Ulyanov, Alexander Vakhitov, and Victor S. Lempitsky. Textured neural avatars. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2382–2392, 2019. [2](#)
- [34] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. Deepvoxels: Learning persistent 3d feature embeddings. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2432–2441, 2019. [2](#)
- [35] Lorenzo Sorgi. Two-view geometry estimation using the rodrigues rotation formula. *2011 18th IEEE International Conference on Image Processing*, pages 1009–1012, 2011. [2](#)
- [36] Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7491–7500, 2021. [2](#)
- [37] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *arXiv: Computer Vision and Pattern Recognition*, 2019. [2](#)
- [38] 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. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12939–12950, 2021. [2](#)
- [39] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. Vid2actor: Free-viewpoint animatable person synthesis from video in the wild. *ArXiv*, abs/2012.12884, 2020. [2](#)
- [40] Chung-Yi Weng, Brian Curless, Pratul P. Srinivasan, Jonathan T. Barron, and Ira Kemelmacher-Shlizerman. Humannerf: Free-viewpoint rendering of moving people from monocular video. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16189–16199, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [41] Minye Wu, Yuehao Wang, Qiang Hu, and Jingyi Yu. Multi-view neural human rendering. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1679–1688, 2020. [2](#)
- [42] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9416–9426, 2021. [2](#)
- [43] Hongyi Xu, Thiemo Alldieck, and Cristian Sminchisescu. H-nerf: Neural radiance fields for rendering and temporal reconstruction of humans in motion. In *NeurIPS*, 2021. [2](#), [4](#), [6](#)
- [44] Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul E. Debevec, William T. Freeman, and Jonathan T. Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Trans. Graph.*, 40:237:1–237:18, 2021. [2](#), [7](#)