

NPC: Neural Point Characters from Video

Shih-Yang Su¹ Timur Bagautdinov² Helge Rhodin¹
¹The University of British Columbia ²Reality Labs Research



Figure 1: **Neural Point Characters (NPC)** is an animatable point-based body model that improves fidelity and generality. NPC can be learned from a single or multiple videos, generalizes well to novel poses and does not require a pre-built surface. **Back to front:** driving motion, estimated point cloud, and neural character model. **All faces are blurred for anonymity.**

Abstract

High-fidelity human 3D models can now be learned directly from videos, typically by combining a template-based surface model with neural representations. However, obtaining a template surface requires expensive multi-view capture systems, laser scans, or strictly controlled conditions. Previous methods avoid using a template but rely on a costly or ill-posed mapping from observation to canonical space. We propose a hybrid point-based representation for reconstructing animatable characters that does not require an explicit surface model, while being generalizable to novel poses. For a given video, our method automatically produces an explicit set of 3D points representing approximate canonical geometry, and learns an articulated deformation model that produces pose-dependent point transformations. The points serve both as a scaffold for high-frequency neural features and an anchor for efficiently mapping between observation and canonical space. We demonstrate on established benchmarks that our representation overcomes limitations of prior work operating in either canonical or in observation space. Moreover, our automatic point extraction approach enables learning models of human and animal characters alike, matching the performance of the methods using rigged surface templates despite being more general. Project website: <https://lemonatsu.github.io/npc/>.

1. Introduction

It is now possible to reconstruct photo-realistic characters from monocular videos, but reaching high-fidelity reconstructions still requires controlled conditions and dedicated capture hardware that prevents large-scale use. While static scenes can be reconstructed from multiple views recorded with a single moving camera, capturing dynamic human motion demands controlled studio conditions [2, 22, 23, 39], usually with a large number of synchronized cameras. One way to tackle the monocular case is to exploit the intuition that movement of the camera with respect to a body part is roughly equivalent to moving a body part with respect to the camera [43, 44]. However, the unconstrained setting remains difficult, as it requires establishing correspondences across frames and through dynamic deformation.

A prominent line of work follows the traditional animation pipeline—rigging a surface mesh to an underlying skeleton and equipping the mesh either with a neural texture [2, 22] or learnable vertex features [16, 34, 35]. This approach is very efficient, as the forward mapping using forward kinematics provides a robust estimate of underlying geometry in closed form, and it also allows for high-quality reconstructions, as neural textures are capable of representing high-frequency details [22, 46]. However, it does require highly accurate 3D pose and body shape, which are

typically obtained from expensive laser scans [50] or by offline fitting personalized parametric body models to multi-view captures [2, 55]. Moreover, most of the existing parametric models rely on linear blend skinning [18], which is prone to artefacts. In this work, we aim at building an animatable full-body model from a single monocular video, without relying on a pre-defined template or complex capture systems.

Another line of work reduces the problem to building a body model in canonical space, e.g., learning a neural radiance field of the person in T-pose [14, 19, 50]. In practice, this family of approaches requires finding the backward mapping from the observation space to the canonical space, which is either learned from high-fidelity scans [49], e.g. by learning a deformation on top of rigid motion [31, 52], or through root finding [19, 50]. These methods are typically limited in generalization, require high-quality training data, and are computationally heavy at test time.

Our goal is to attain the generality of surface-free canonical models and the speed of forward mappings without a need for a pre-defined template that restricts applicability. To this end, we rely on an implicit, surface-free body representation that does not depend on a precise 3D pose nor a rig [43, 44]. Namely, we use A-NeRF [44] alongside off-the-shelf pose estimators [15] to get a reasonable 3D pose estimate and DANBO [43] to learn an initial surface estimate in the canonical space.

Given such an initial estimate, our method, Neural Point Characters (NPC), reconstructs a high-quality neural character that can be animated with new poses and rendered in novel views (Figure 1). The difficulty is that the initial shape is very approximate and noisy, and is insufficient to model high-quality geometry. Central to NPC is thus a novel point representation that is designed to improve noisy shape estimates, and subsequently learns to represent fine texture details and pose-dependent deformations. Keys are our two main contributions: First, we propose to find correspondences between the canonical and observation space by inverting the forward mapping via nearest-neighbor lookup, including the non-linear pose-dependent deformation field modeled with a graph neural network (GNN). The surface points serve as anchors for inferring the backward mapping of a query point to the neural field in canonical space through an efficient amortized lookup. Our approach is more efficient than root finding and more precise than models assuming piece-wise rigidity. Second, we propose to use the non-linearly deformed points as a scaffold in observation space and represent high-frequency pose-dependent details. Point features and neural fields in canonical space are further combined with bone-relative geometry encodings to encourage better generalization. We provide a comprehensive evaluation of our method, and demonstrate state-of-the-art results on the established benchmarks in monoc-

ular and multi-view human reconstruction. We additionally demonstrate versatility by reconstructing human and animal characters alike.

2. Related Work

Neural rendering has recently emerged as a powerful family of approaches for building controllable neural representations of 3D scenes [26, 29, 45, 53].

Point-Based Representations. NPBG [1] uses a pre-computed static point cloud to represent a rough scene geometry, and learns a set of neural descriptors attached to each point, which are then rasterized at multiple resolutions and further processed with a learnable rendering network. ADOP [40] follows a similar strategy but additionally optimizes camera parameters and point cloud locations, which leads to better fidelity reconstructions. Pulsar [17] represents a scene as a collection of spheres, which have optimizable locations and radius parameters and are then rendered with a custom differentiable renderer and post-processed with neural shading ConvNet. SMPLpix [37] and Point-based human clothing [59] similarly utilize a 2D CNN to in-paint the rasterized 2D point feature maps for rendering. Power of Points [24] learns an auto-decoder to predict point displacements for modeling different garments and dynamics. Concurrent work PointAvatar [61] represents facial avatars with deformable point clouds that grow progressively to capture fine-grained details, and render with a differential point rasterizer. Our method also builds upon a point-based representation, but is equipped with an articulated model and relies on volumetric rendering.

Neural Field Representations. Neural Radiance Fields (NeRF) [29, 42] represent scenes implicitly, as a coordinate-conditioned MLPs and enable high-fidelity novel-view synthesis of static scenes. NeRF-like methods have also been applied to dynamic scenes by adding a deformation field [9, 21, 32, 38, 54]. Particularly relevant to our method is Point-NeRF [56]. Point-NeRF uses point clouds obtained from multi-view stereo as a geometry proxy, and models appearance on top with a locally-conditioned NeRF to achieve faster learning and enables high-fidelity novel-view synthesis on static scenes. Similarly, NPC uses point-based neural fields, but uses pose-driven articulated skeleton model to transform the point cloud to enable dynamic human modeling, and dynamically decodes pose-dependent per-point payload to account for pose-dependent deformations.

Neural Fields for Human Modeling. Recent work applies NeRF to reconstruct dynamic human appearance [16, 27, 35, 38, 52] and learning animatable human models [19, 22, 30, 34, 43, 44, 50, 62] from video sequences. NeRF for a *static* scene directly maps query point coordinates to opacity and color values. The key challenge for an *animatable* NeRF model is to map a per-frame 3D query to the canonical body model space.

A common approach to tackle this challenge is to learn a **backward mapping** that transports query points from observation space with the posed person to the canonical space with the person in rest pose (e.g., T-pose) [7, 14, 22, 28, 34, 41, 47]. The rendering is then done by neural fields defined in the canonical space. While learning backward mapping allows for high-quality reconstructions of dynamic human geometry and appearance, these methods typically require 3D scans for training [7, 28, 41, 47], need a texture map [22], or depend on a surface prior [14, 34]. Additionally, the backward mapping is difficult to learn as it requires solving many-to-one mapping [5, 19]. These approaches are thus prone to artefacts when rendering unseen poses.

Alternatively, some existing methods rely on a **forward mapping** that moves features in the canonical space to the observation space. Earlier approaches anchor neural features to SMPL vertices [16, 35] and deform the body mesh to the observation space via linear blend skinning (LBS) [18]. Per-vertex features are then diffused by a 3D CNN around the posed mesh, and the neural field is defined on top of the resulting voxel feature grid. These methods are still limited by the quality of the underlying LBS model and inability to capture non-linear deformations, which often results in blurry reconstructions.

An alternative backward mapping that was recently popularized is to apply a differentiable root-finding algorithm [19, 50] on learned forward mappings. TAVA [19] is a template-free approach that performs iterative root-finding to solve for a point in canonical space. ARAH [50] proposes a joint root-finding method that finds ray intersection on the NeRF body model initialized using a pre-trained hypernetwork [11, 49] for rendering. Although TAVA and ARAH are template-free and show state-of-the-art synthesis quality on unseen poses, both methods are computationally heavy at train and test time, and ARAH additionally requires a surface prior which is built using 3D scans. Compared to these approaches, NPC uses sparse point clouds to create efficient forward mapping between a canonical and observation space, which provides better query feature localization with lower computational cost, and does not rely on pre-trained surface priors.

Conceptually related to our method are SLRF [62], KeypointNeRF [27], and AutoAvatar [3]. SLRF shares a similar concept of leveraging surface points for anchoring separate radiance fields to represent local appearances. In contrast, NPC explicitly represents local structure using much denser feature point clouds and enabling dense correspondence across body poses without relying on pre-defined parametric meshes. KeypointNeRF uses sparse key points to triangulate 2D pixel-aligned features extracted from multi-view imagery for 3D face reconstructions. Unlike KeypointNeRF, NPC stores point features directly in 3D space, captures pose-dependent deformation, and is drivable. Au-

toAvatar uses K-NN-based encodings for dynamic neural body deformation fields, but does not model textures.

3. Method

NPC reconstructs a neural character that can be animated from a single video $\{\mathbf{I}_1, \dots, \mathbf{I}_N\}$ with N frames. Figure 2 provides an overview. The representation includes neural fields [29] in canonical space with the character in rest pose and a set of surface points $\{\mathbf{p}_1, \dots, \mathbf{p}_P\}$ that deform dynamically and map between canonical and observation space. In the following, we explain how we obtain these representations and how we combine them to form the geometry and appearance of the NPC.

3.1. Skeleton and Point Representation.

Pose and shape initialization. Each frame \mathbf{I}_t of the input video is processed with an off-the-shelf pose estimator¹ followed by refinement with A-NeRF [44] or if available, multi-view triangulation. The resulting pose $\theta_t = \{\mathbf{R}_{t,0}, \dots, \mathbf{R}_{t,J-1}\} \in \mathbb{R}^{J \times 3}$, the joint angle rotations of a pre-defined skeleton with J joints, is then used as input to train the existing neural body model DANBO [43]. This body model is set to a T-pose, and the iso-surface of the underlying implicit model is extracted with marching cubes. Note that we train DANBO with a much-reduced number of ray samples and small number (10k) of iterations, which corresponds to $\sim 2.5\%$ of the full training set; this process runs under 30 mins on a single NVIDIA RTX 3080 GPU.

Part-wise point cloud sampling. We obtain the location of the sparse surface points \mathbf{p} by farthest point sampling on the isosurface until we obtain 200 points $\{\mathbf{p}_{i,j}\}_{i=1}^{200}$ for every body part j of the skeleton. We drop the part subscript for simplicity when the part assignment is irrelevant. As demonstrated in Figure 2, the resulting point cloud is a very coarse approximation of the underlying geometry. In what follows, we describe our main contributions, which allow learning high-quality details on top of this rough estimate, without strict assumptions on a clean surface imposed by existing model-based approaches, and with sharper details and better texture consistency compared to DANBO and other surface-free models.

Canonical to observation space (forward mapping). The extracted points $\mathbf{p}_i \in \mathbb{R}^3$ live in canonical space and are fixed for the entire video. Deformations in frame t are modelled by a learnable mapping $\mathbf{p}_{i,j}$ to the posed observation space,

$$\mathbf{p}_{i,j,t}^o = \text{LBS}(\theta_t, \mathbf{w}, \mathbf{p}_{i,j}) + R_{t,j}(\theta_t)\Delta\mathbf{p}_{i,j}(\theta_t), \quad (1)$$

comprising a linear blend skinning (LBS) [18] term to model rigid deformations and a non-linear term $\Delta\mathbf{p}(\theta_t)$ that

¹We follow recent template-free human NeRFs [19, 43] in using SMPL-based pose estimators to acquire joint angle rotations, but other pose estimators would work equivalently.

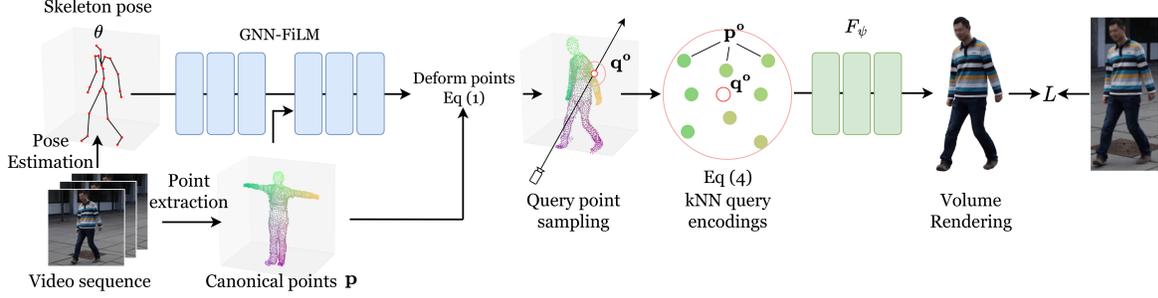


Figure 2: **Overview.** NPC produces a volume rendering of a character with a NeRF F_ψ , locally conditioned on features aggregated from a dynamically deformed point cloud. Given a raw video, we first estimate a canonical point cloud \mathbf{p} with an implicit body model (Section 3.1). GNN then deforms canonical points \mathbf{p} conditioned on skeleton pose θ , and produces a set of pose-dependent per-point features (Section 3.2, Section 3.3). Every 3D query point \mathbf{q}^o in the observation space aggregates the features from k -nearest neighbors in the posed point cloud. The aggregated feature is passed into F_ψ for the volume rendering. Our model is supervised directly with input videos (Section 3.4).

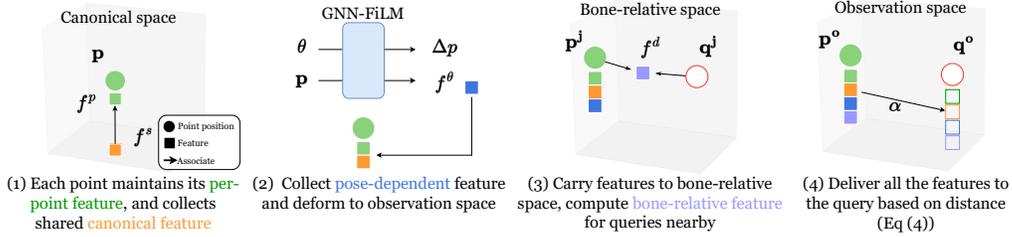


Figure 3: **Point feature encoding.** Our core idea is to employ a point cloud \mathbf{p} as an anchor to carry features from the canonical to the observation space, forming an efficient mapping between the two. (1) Each \mathbf{p} carries a learnable feature f^p and its position queries features f^s from a canonical field. (2) The GNN adds pose-dependent features f^θ and deformation $\Delta\mathbf{p}$. (3) The view direction and distance is added in bone-relative space. (4) The k -nearest neighbors of \mathbf{q}^o in $\{\mathbf{p}_i^o\}_i^N$ are used to establish forward and backward mapping from a query point to both posed and canonical points.

is learned relatively to the orientation $R_{t,j}(\theta_t)$ of the underlying skeleton.

Key to our approach is learning the pose-dependent deformation $\Delta\mathbf{p}_{i,j}(\theta_t)$ and character-specific LBS weights \mathbf{w} that define the influence of every joint on every point. This is distinct from prior work using template meshes with artist-created LBS weights [22, 34, 35, 50]. We initialize the learnable LBS weights for a point \mathbf{p} to

$$\mathbf{w}^0 = \{w_1, \dots, w_j\}, \text{ where } w_j = \frac{1}{\sqrt{d_j(\mathbf{p})}}, \quad (2)$$

with $d_j(\cdot)$ being the distance from point \mathbf{p} to the bone j . In the following, we drop the frame index t when irrelevant.

3.2. Neural Rendering with a Point Scaffold

An image $\hat{\mathbf{I}}$ of the character is created by volume rendering of an implicit radiance and signed-distance field (\mathbf{c}, s) as in [58]. The rendered image are used for optimizing NPC as detailed in Section 3.4. Specific to our method is how we combine various point features that are sparse to form a continuous field $f(\mathbf{q}^o)$. Given a body pose θ and a 3D query point \mathbf{q}^o in the observation space, we predict the corresponding color \mathbf{c} and signed distance s with a neural field

F_ψ with learnable parameter ψ from $f(\mathbf{q}^o)$, combined with established frame codes $f(t)$ [25, 35] and geometric features $g(\mathbf{q}^o)$,

$$(\mathbf{c}, s) = F_\psi(f(\mathbf{q}^o), g(\mathbf{q}^o), f(t), \mathbf{w}). \quad (3)$$

The individual features are introduced in the subsequent section and supplemental document. We continue with the sparse-to-continuous transformations and mapping between observation and canonical spaces.

Point features to observation space (forward mapping).

Our per-point features encode the local appearance and shape. They are associated with points \mathbf{p}_i , living in canonical space. To map features from canonical to observation space we first apply Eq. 1, yielding 3D locations \mathbf{p}^o . To create a continuous field, we query the $K = 8$ nearest neighbors of \mathbf{q}^o and aggregate features f_i with a Gaussian RBF kernel,

$$f(\mathbf{q}^o) = \sum_{i \in \mathcal{N}(\mathbf{q}^o)} f_i a_i, \quad a_i = \exp\left(-\frac{(\mathbf{q}^o - \mathbf{p}_i^o)^2}{\beta_i}\right), \quad (4)$$

where β_j is a learnable scale parameter that determines

the influence range of feature i among nearest neighbors $\mathcal{N}(\mathbf{q}^\circ)$.

Observation to bone space (rigid backward mapping). The kinematic chain of the skeleton in pose θ defines for every joint j a rigid transformation \mathbf{T}_j and its inverse,

$$\mathbf{q}^j = \mathbf{T}_j^{-1}(\theta)\mathbf{q}^\circ, \quad (5)$$

that maps from query \mathbf{q}° in posed to joint-relative space. It is efficient and widely used but neglects non-rigid motion.

Observation to canonical space (non-rigid backward). Inverting the non-linear pose-dependent deformation is a non-trivial task. Prior work tried to approximate it with another network [30, 34, 47], but this requires conditioning on the position in posed space and leads to poor generalization. Another approach is iterative root finding [19, 50], which is extremely slow, both at training and at inference time.

Instead, we exploit the fact that the forward mapping of surface points via Eq. 1 lets us invert the function point-wise by using points as the anchor. The forward mapping gives a finite number of pairs $(\mathbf{p}^\circ, \mathbf{p})$ in observed and canonical space that correspond exactly. Moreover, the features important for shape reconstruction live on the surface, and so we found it sufficient to query the canonical space feature f^s at the canonical positions of the nearest neighbors $\mathcal{N}(\mathbf{q}^\circ)$. Intuitively, our surface points \mathbf{p} serve as anchors to carry the appearance information from the canonical space to the point \mathbf{q}° that resides in the observation space.

Efficient nearest neighbor lookup Our forward and backward mappings both rely on nearest neighbor search, which would be slow when performed naively over the entire point cloud for each query. To make this approach feasible, we apply divide and conquer by searching only within the points associated with the three closest joints. Moreover, because our points form a surface, we found i) that the nearest neighbors to a query stays mostly unchanged in both observation and canonical space and ii) that the one nearest neighbor defines the remaining $K - 1$. This makes it possible to tabulate for each canonical point its $K - 1$ nearest neighbors. This reduces the runtime computation to finding the one nearest neighbor in posed space and looking up the remaining ones. Therefore, we can use a large number of neighbors without runtime drawbacks. Note that we allow the tabulated nearest neighbors to cross joint boundaries to prevent seam artefacts.

3.3. Neural Features and Field Representations

With the mapping between different spaces established through our point features, we now turn to the structure of the neural features in each space and their optimization.

Point features. We associate each point i with a feature vector f^p to provide an explicit representation local to the corresponding body region. It improves texture consistency across poses and captures local high-frequency ap-

pearance and geometry. Feature vector f^p and position \mathbf{p}_i are both learned, initialized with respectively $\mathcal{N}(0, 0.1)$ and the canonical point clouds from DANBO.

Pose-dependent point features. We use a pose-conditioned GNN to predict both the per-point pose-dependent feature f^θ and deformation Δp ,

$$f^\theta, \Delta p = \text{GNN-FiLM}(\mathbf{p}, \theta). \quad (6)$$

The architecture is inspired by DANBO [43], and extended with FiLM conditioning [36].

Canonical features. We represent f^s using a per-part factorized volume $\mathbf{v}_j \in \mathbf{R}^{12 \times 3V}$ [43]. To retrieve the feature for a point \mathbf{p} , we project and interpolate on the factorized vector of each axis,

$$\mathbf{v}_j(\mathbf{p}) = (v^x[x(\mathbf{p})], v^y[y(\mathbf{p})], v^z[z(\mathbf{p})]) \quad (7)$$

$$f^s = \text{MLP}_{\text{spatial}}(\mathbf{v}_j(\mathbf{p})) \in \mathbf{R}^{3V}, \quad (8)$$

where $v^{(\cdot)}[\cdot]$ is the interpolated feature, and $\text{MLP}_{\text{spatial}}$ is a 2-layer MLP combining features from all axis.

Bone-relative features. To provide the NeRF ability to reason about the approximation made in the non-linear backward mapping, we condition it on f^d - the offset of the query to each of the nearest neighbors, defined in the local bone coordinates of the closest joint j . This distance is computed in a learned feature space $\mathbf{c}_j \in \mathbf{R}^{16 \times 3C}$,

$$f_i^d = \mathbf{c}_j(\mathbf{R}_j^{-1}\mathbf{p}_{i,j}^\circ) - \mathbf{c}_j(\mathbf{R}_j^{-1}\mathbf{q}_j^\circ), \quad (9)$$

with \mathbf{c}_j a per-part factorized feature volume, like \mathbf{v} .

After the NN lookup, all of the introduced features are associated with points \mathbf{p}_i (see Figure 3). Together, they form the $f_i = (f_i^p, f_i^\theta, f_i^s, f_i^d)$ used in Eq. 4, which carries both the shared and independent traits, as well as the flexibility to adapt to deformation-induced appearance changes. Additional commonly-used features and $g(\mathbf{q})$ are explained in the supplemental document.

3.4. Training

We train NPC with the photometric loss on ground truth images \mathbf{I} ,

$$L_p = \sum_{(u,v) \in I} |I(u, v, \psi) - \hat{I}(u, v)|. \quad (10)$$

We employ L1 loss on real-world datasets for robustness and L2 otherwise. To encourage the network to predict proper level-set, we adopt the eikonal loss [10, 58]

$$L_{\text{eik}} = \sum_{\tilde{\mathbf{p}}} (|\|\nabla s|\| - 1)^2, \quad (11)$$

where $\tilde{\mathbf{p}}$ are sparse points randomly selected around the posed surface points \mathbf{p}° . We regularize the output of the deformation network with

$$L_{\Delta p} = \max(\Delta p - \delta, 0)^2, \quad (12)$$

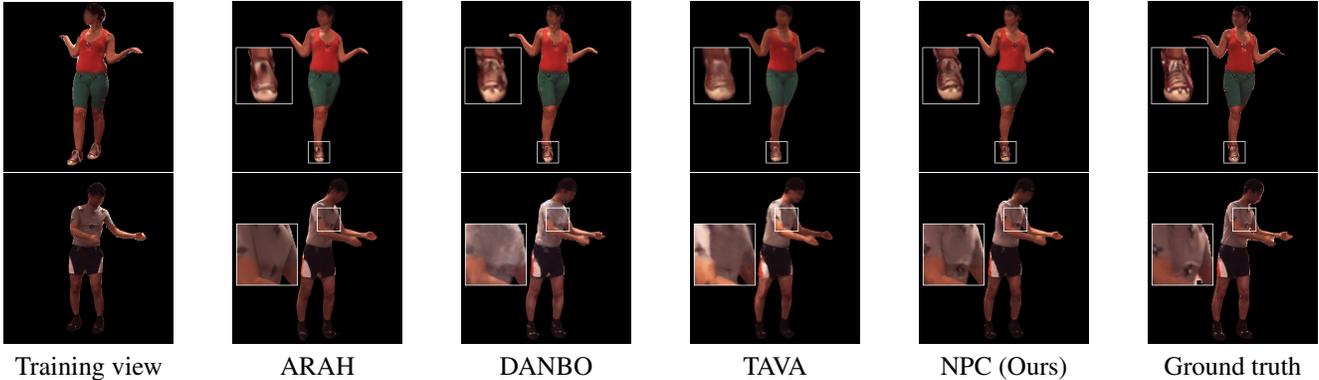


Figure 4: **Novel view synthesis on Anim-NeRF Human3.6M [12, 13, 34] test split.** Our point-based representation ensures better feature consistency across training poses, and therefore synthesizes better details for novel views.

which penalizes the non-linear deformation when Δp is over the threshold $\delta = 0.04$. A second loss discourages the network from moving points away from their neighbors,

$$L_N = \sum_{\mathbf{p}^o} \sum_{\mathbf{p}_i \in \mathcal{N}(\mathbf{p})} (\|\mathbf{p}^o - \hat{\mathbf{p}}_i^o\| - \|\mathbf{p} - \hat{\mathbf{p}}_i\|)^2. \quad (13)$$

We include a third loss to encourage the predicted signed distance on the surface points to either be on or within the body surface,

$$L_s = \sum_{\hat{\mathbf{p}}} \max(s, 0). \quad (14)$$

To summarize, our full training objective is

$$L = L_p + \lambda_{eik} L_{eik} + \lambda_{\Delta p} L_{\Delta p} + \lambda_N L_N + \lambda_S L_S, \quad (15)$$

with $\lambda_{eik} = 0.01$, $\lambda_{\Delta p} = 1.0$, $\lambda_N = 10.0$, and $\lambda_S = 0.1$. We train for 150k iterations, for 12 hours on a single NVIDIA RTX 3080 GPU including the point initialization process.

4. Experiments

We quantify the improvements our NPC brings over the most recent surface-free approach TAVA [19], DANBO [43] and A-NeRF [44], as well as most recent and established approaches that leverage template or scan-based prior, including ARAH [50], Anim-NeRF [34] and NeuralBody [35]. Moreover, we conduct ablation studies to verify how our proposed per-point encoding, pose-dependent feature, and coordinate features help improve perceptual quality. The supplemental video and document provide additional qualitative results and implementation details. The code will be made available upon publication to facilitate human body modeling research².

²All data sourcing, modeling codes, and experiments were developed at University of British Columbia. Meta did not obtain the data/codes or conduct any experiments in this work.

Metrics. We adopt standard image metrics, including pixel-wise PSNR and SSIM [51], as well as perceptual metrics LPIPS [60] and KID [4, 33] that better quantify the improvement on texture details under slight misalignment that is unavoidable when driving detailed characters with only sparse pose. All scores are on the holdout testing set. Specifically, we evaluate the performance for *novel-view synthesis* on multi-view datasets by rendering seen poses from unseen testing views and *novel-pose synthesis* to synthesize appearances of body poses unseen during training.

Datasets. We evaluate NPC on the established benchmarks for body modeling, including both indoor, outdoor, monocular video capture, and synthetic animal datasets.

- Human3.6M [12, 13]: we use the protocol from Anim-NeRF [34], evaluating on a total of 7 subjects using the foreground maps provided by the authors.
- MonoPerfCap [57]: the dataset consists of multiple monocular outdoor footages. We use the same 2 subjects as in [44]. The training poses and cameras are SPIN estimates [15] refined by A-NeRF.

We additionally evaluate NPC on one subject with loose clothing from ZJU-Mocap [8, 35], and use challenging motion sequences, including dancing and gymnastic from poses AIST++ [20] and SURREAL+CMU-Mocap [6, 48] dataset, for animating our learned characters. Finally, we use the synthetic wolf dataset from TAVA [19] as a proof-of-concept on how we can apply NPC to non-human subjects similar to other template-free methods [19, 30, 43].

4.1. Novel View Synthesis

Our point-based representation anchors high-frequency local details that are shared across all training views, enabling improved visual details even when rendering from a novel view not presented in the training data. Compared to ARAH [50] using implicit-surface prior [49] and root-finding for canonical mapping, NPC synthesizes sharper results despite using only noisy surface points extracted

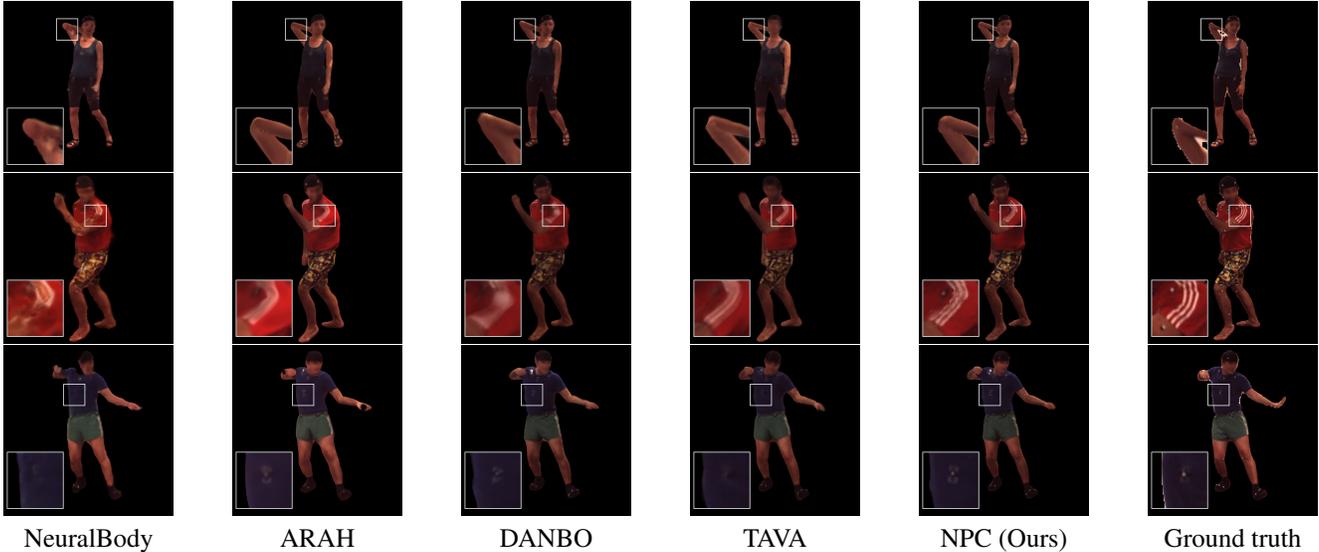


Figure 5: Unseen pose synthesis on Anim-NeRF Human3.6M [12, 13, 34] test split. Our NPC shows better appearance consistency on unseen poses, with improved articulation, sharper textures, and fine details like motion trackers and stripes.

Table 1: Novel-view synthesis comparisons on Anim-NeRF Human3.6M [12, 13, 34] test split. Our NPC benefits from the point-based representation, and achieves better overall perceptual quality.

	S1			S5			S6			S7			S8			S9			S11			Avg		
	PSNR↑	SSIM↑	LPIPS↓	PSNR	SSIM	LPIPS																		
Template/Scan-based prior																								
NeuralBody	22.88	0.897	0.139	24.61	0.917	0.128	22.83	0.888	0.155	23.17	0.915	0.132	21.72	0.894	0.151	24.29	0.911	0.122	23.70	0.896	0.168	23.36	0.905	0.140
Anim-NeRF	22.74	0.896	0.151	23.40	0.895	0.159	22.85	0.871	0.187	21.97	0.891	0.161	22.82	0.900	0.146	24.86	0.911	0.145	24.76	0.907	0.161	23.34	0.897	0.157
ARAH [†]	24.53	0.921	0.103	24.67	0.921	0.115	24.37	0.904	0.133	24.41	0.922	0.115	24.15	0.924	0.104	25.43	0.924	0.112	24.76	0.918	0.128	24.63	0.920	0.115
Template-free																								
A-NeRF	23.93	0.912	0.118	24.67	0.919	0.114	23.78	0.887	0.147	24.40	0.917	0.125	22.70	0.907	0.130	25.58	0.916	0.126	24.38	0.905	0.152	24.26	0.911	0.129
DANBO	23.95	0.916	0.108	24.86	0.924	0.108	24.54	0.903	0.129	24.45	0.920	0.113	23.36	0.917	0.116	26.15	0.925	0.108	25.58	0.917	0.127	24.73	0.918	0.115
TAVA	25.28	0.928	0.108	24.00	0.916	0.122	23.44	0.894	0.138	24.25	0.916	0.130	23.71	0.921	0.116	26.20	0.923	0.119	26.17	0.928	0.133	24.72	0.919	0.124
NPC (Ours)	24.81	0.922	0.097	24.92	0.926	0.100	24.89	0.909	0.118	24.87	0.924	0.105	24.03	0.923	0.104	26.39	0.930	0.095	25.86	0.925	0.117	25.13	0.924	0.104

[†]: we evaluate using the officially released ARAH, which has undergone refactorization, resulting in slightly different numbers to the ones in [50].

without supervision. Figure 4 shows that we recover better appearance details such as markers and shoelaces. We observe that TAVA, although using root-finding like ARAH for mapping, suffers more blurry artifacts. We conjecture that the root-finding algorithm requires good initialization for video sequences with complex motion. Table 1 quantifies the improvement over the recent template-free and template/scan-based neural body fields.

4.2. Unseen Pose and Animating NPC

Unseen pose synthesis requires the method to maintain consistent appearances in out-of-distribution pose configurations. Our NPC shows overall more fine-grained and consistent renderings as shown in Figure 5. Compared to Neural Body [35], which anchors the body representation on template body mesh and diffuses the feature using 3D CNN, our NPC produces crispier details like motion capture body trackers and wrinkles, and maintains better texture consistency like the clearer stripes. We attribute

this to our point-based feature encoding and explicit non-deformation on surface points that, in combination, enable better localization of the body features. Table 2 and Table 3 verify the consistent improvement of our NPC on the perceptual quality on not just multi-view dataset Human3.6M [12, 13], but also outdoor monocular video sequences MonoPerfCap [57]. Note that we omit PSNR and SSIM for Table 3 as these metrics are susceptible to the varying lighting conditions in outdoor sequences.

We further animate the learned NPC models with motion sequences from out-of-distribution body poses in Figure 6, to showcase how we can potentially apply our point-based characters for animation. Note that we cannot quantify the performance for these examples as there are no ground truth images available for evaluation.

Geometry Comparisons. NPC reconstructs detailed 3D geometry even from monocular video and generalizes deformations well to unseen poses, with an overall more complete body outline, as visualized in Figure 7. Different to

Table 2: **Unseen pose rendering comparisons on Anim-NeRF Human3.6M [12, 13, 34] test split.** The efficient canonical feature mapping carried out by our point-based approach enables better generalization in unseen poses.

	S1			S5			S6			S7			S8			S9			S11			Avg		
	PSNR↑	SSIM↑	LPIPS↓	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR↑	SSIM↑	LPIPS↓	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
Template/Scan-based prior																								
NeuralBody	22.10	0.878	0.143	23.52	0.897	0.144	23.42	0.892	0.146	22.59	0.893	0.163	20.94	0.876	0.172	23.05	0.885	0.150	23.72	0.884	0.179	22.81	0.888	0.157
Anim-NeRF	21.37	0.868	0.167	22.29	0.875	0.171	22.59	0.884	0.159	22.22	0.878	0.183	21.78	0.882	0.162	23.73	0.886	0.157	23.92	0.889	0.176	22.61	0.881	0.170
ARAH [†]	23.18	0.903	0.116	22.91	0.894	0.133	23.91	0.901	0.125	22.72	0.896	0.143	22.50	0.899	0.128	24.15	0.896	0.135	23.93	0.899	0.143	23.27	0.897	0.134
Template-free																								
A-NeRF	22.67	0.883	0.159	22.96	0.888	0.155	22.77	0.869	0.170	22.80	0.880	0.182	21.95	0.886	0.170	24.16	0.889	0.164	23.40	0.880	0.190	23.02	0.883	0.171
DANBO	23.03	0.895	0.121	23.66	0.903	0.124	24.57	0.906	0.118	23.08	0.897	0.139	22.60	0.904	0.132	24.79	0.904	0.130	24.57	0.901	0.146	23.74	0.901	0.131
TAVA	23.83	0.908	0.120	22.89	0.898	0.135	24.54	0.906	0.122	22.33	0.882	0.163	22.50	0.906	0.130	24.80	0.901	0.138	25.22	0.913	0.145	23.52	0.899	0.141
NPC (Ours)	23.39	0.901	0.109	23.63	0.906	0.113	24.59	0.911	0.105	23.46	0.903	0.129	22.87	0.907	0.121	24.86	0.907	0.115	25.13	0.911	0.130	23.96	0.906	0.119

[†]: we evaluate using the officially released ARAH, which has undergone refactorization, resulting in slightly different numbers to the ones in [50].

Table 3: **Unseen pose synthesis on MonoPerfCap [57].** NPC shows better overall perceptual quality over DANBO on learning generalized model from monocular videos.

	ND		WP		Avg	
	KIDx100 ↓	LPIPS ↓	KIDx100 ↓	LPIPS ↓	KIDx100 ↓	LPIPS ↓
A-NeRF	4.97	0.197	6.53	0.223	5.75	0.208
DANBO	4.83	0.194	4.66	0.214	4.74	0.202
NPC (Ours)	2.57	0.198	3.56	0.207	3.07	0.202



Figure 6: **Motion retargeting from out-of-distribution poses on various subjects.** NPC retains better appearance consistency and texture detail.

DANBO and other implicit methods, our point scaffold provides correspondences across frames and is therefore applicable to dense 3D pose estimation. Finally, we show that NPC can recover imperfect geometry, such as correcting the missing foot and trouser shape in Figure 8.

4.3. NPC Deformation on Loose Clothing

To provide a further comparison on pose-dependent deformation, we tested on subject 387 of ZJU-Mocap [8, 35], which includes loose clothing and long-range deformation dependency. We report the results in Table 4 following the established protocols from [35]. NPC match or out-

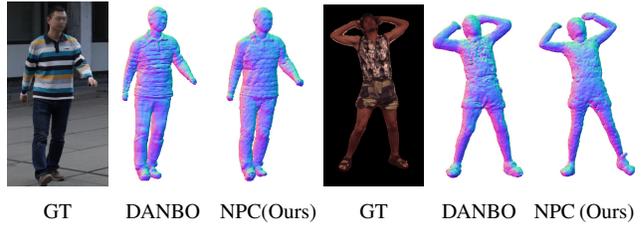


Figure 7: **Body geometry in unseen poses.** NPC matches the shape quality of DANBO.

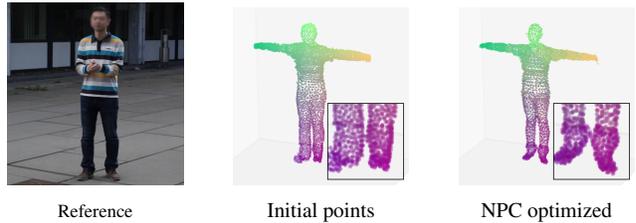


Figure 8: NPC recovers fine-grained geometry from the noisy initialization obtained by [43].

Table 4: **ZJU-Mocap [35] subject 387 novel view and pose synthesis.** NPC achieves improved FID and KID.

	ARAH				NPC (Ours)			
	PSNR	LPIPS	FID	KID×100	PSNR ↑	LPIPS ↓	FID ↓	KID×100 ↓
Novel view	25.83	0.096	38.2	1.17	25.10	0.093	36.1	0.94
Novel pose	22.93	0.127	51.4	2.74	21.88	0.134	49.2	2.59

perform ARAH [50] in perceptual metrics (LPIPS, KID, FID) on novel-view and KID & FID on novel-pose, despite ARAH using a prior pre-trained on a large-scale scan dataset [49]. A closer analysis (see Figure 9) reveals that texture details are improved, while pose-dependent wrinkle details are comparable if not better. ARAH yields overly smooth while NPC produces slightly grainy results, leading to a lower PSNR but improved FID and KID.



Figure 9: **Novel view synthesis results on ZJU-Mocap subject 387.** Compared to ARAH, NPC captures sharper but slightly grainy results.

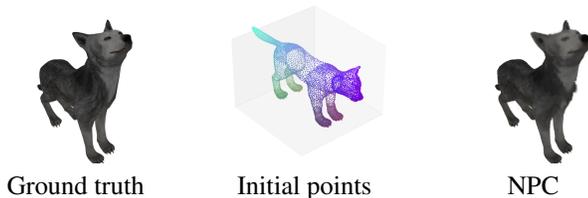


Figure 10: NPC can be used for learning non-human subjects as it does not require a pre-trained surface prior.

4.4. NPC on Animal

We verify the applicability of NPC on entities where no pretrained surface prior available on the TAVA wolf cub subject [19]. We lightly train NPC for 10k iterations, and present an out-of-distribution pose rendering in Figure 10.

4.5. Ablation Study

We use Human3.6M S9 to conduct our ablation study. In Table 5, we show that all our proposed features contribute to the final perceptual quality. Without f^θ , the rendering outcomes become noisy. Without f^p , the model produces inconsistent and blurry texture. Disabling Δp results in lower perceptual quality for poses with more extreme deformation, as the nearest neighbor mapping becomes inaccurate in this case. In Table 6, we observe that NPC is insensitive to the number of points, and using more than 200 points (our default value) yields diminished returns. We further verify that NPC learns character models with consistent quality by running our training framework 4 times, including re-training DANBO from scratch for canonical point clouds initialization. As reported in Table 7, we observe low variation in the performance across different runs on both Human3.6M S9 and MonoPerfCap weipeng.

5. Limitations and Discussion

Although NPC improves runtime over most existing implicit neural body models, it still utilizes neural fields en-

Table 5: **Ablation study on each of our proposed designs.** All of them contribute to the final perceptual quality.

	KIDx100 ↓	LPIPS ↓
No f^d	5.54	0.122
No f^p	4.51	0.122
No f^θ	4.50	0.120
No Δp	4.61	0.120
NPC (Ours)	4.43	0.115

Table 6: **Ablation study on the numbers of points per body parts.** NPC shows consistent results with different numbers of points on Human 3.6M [12].

40×19		120×19		200×19		280×19		360×19	
PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓
24.83	0.123	24.88	0.117	24.86	<u>0.115</u>	<u>24.87</u>	0.114	24.84	0.117

Table 7: **Ablation study on running with 4 different canonical point clouds initializations.** We report the standard deviations in the parenthesis. The indicates that our point initialization strategy is reliable, and NPC behaves consistently across different training runs.

S9			WP		
KIDx100 ↓	LPIPS (VGG) ↓	LPIPS (Alex) ↓	KIDx100 ↓	LPIPS (VGG) ↓	LPIPS (Alex) ↓
4.34 (± 0.12)	0.116 (± 0.000)	0.124 (± 0.000)	3.75 (± 0.29)	0.207 (± 0.001)	0.127 (± 0.001)

coded by moderately large networks, which precludes real-time applications. Our approach sometimes produces small-scale ball-shaped artifacts when the sparse points cannot fully cover the body surfaces. This can potentially be resolved via point-growing as proposed in Point-NeRF [56]. See the supplementary document for details. Finally, NPC learns a person-specific model with detailed appearance from raw video sequences, providing a more accessible way for the public to create 3D characters. On the downside, it could be exploited for DeepFakes, causing ethical concerns when used with malicious intent.

6. Conclusion

NPC demonstrated that detailed character models can be learned without having access to laser scanning technology or making restrictive assumptions on an underlying template model. As it is fully automatic, it makes digital full-body avatars available to a larger audience, including minorities that are not well-represented in existing datasets.

Acknowledgements. We thank Shaofei Wang and Rui-long Li for helpful discussions related to ARAH and TAVA. We thank Luis A. Bolaños for his help and discussions, and Frank Yu, Chunjin Song, Xingzhe He and Eric Hedlin for their insightful feedback. We also thank ARC at UBC and Compute Canada for providing computational resources.

References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *ECCV*, 2020. 2
- [2] Timur Bagautdinov, Chenglei Wu, Tomas Simon, Fabian Prada, Takaaki Shiratori, Shih-En Wei, Weipeng Xu, Yaser Sheikh, and Jason Saragih. Driving-signal aware full-body avatars. *ACM TOG (Proc. SIGGRAPH)*, 2021. 1, 2
- [3] Ziqian Bai, Timur Bagautdinov, Javier Romero, Michael Zollhöfer, Ping Tan, and Shunsuke Saito. Autoavatar: Autoregressive neural fields for dynamic avatar modeling. In *ECCV*, 2022. 3
- [4] Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. In *ICLR*, 2018. 6
- [5] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, 2021. 3
- [6] CMU Graphics Lab Motion Capture Database. <http://mocap.cs.cmu.edu>. 6
- [7] Boyang Deng, John P Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. Nasa neural articulated shape approximation. In *ECCV*, 2020. 3
- [8] 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, 8
- [9] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *ICCV*, 2021. 2
- [10] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *ICML*, 2020. 5
- [11] David Ha, Andrew M. Dai, and Quoc V. Le. Hypernetworks. In *ICLR*, 2017. 3
- [12] C. Ionescu, F. Li, and C. Sminchisescu. Latent Structured Models for Human Pose Estimation. In *ICCV*, 2011. 6, 7, 8, 9
- [13] C. Ionescu, I. Papava, V. Olaru, and C. Sminchisescu. Human3.6M: Large Scale Datasets and Predictive Methods for 3D Human Sensing in Natural Environments. *TPAMI*, 2014. 6, 7, 8
- [14] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. In *ECCV*, 2022. 2, 3
- [15] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, 2019. 2, 6
- [16] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. *NeurIPS*, 2021. 1, 2, 3
- [17] Christoph Lassner and Michael Zollhöfer. Pulsar: Efficient sphere-based neural rendering. In *CVPR*, 2021. 2
- [18] John P Lewis, Matt Corder, and Nickson Fong. Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation. In *ACM TOG (Proc. SIGGRAPH)*, 2000. 2, 3
- [19] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhofer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. In *ECCV*, 2022. 2, 3, 5, 6, 9
- [20] Ruilong Li, Shan Yang, David A. Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *ICCV*, 2021. 6
- [21] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021. 2
- [22] 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 TOG (Proc. SIGGRAPH Asia)*, 2021. 1, 2, 3, 4
- [23] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhofer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM TOG (Proc. SIGGRAPH)*, 2021. 1
- [24] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J. Black. The power of points for modeling humans in clothing. In *ICCV*, 2021. 2
- [25] 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. In *CVPR*, 2021. 4
- [26] Moustafa Meshry, Dan B Goldman, Sameh Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. Neural re-rendering in the wild. In *CVPR*, 2019. 2
- [27] Marko Mihajlovic, Aayush Bansal, Michael Zollhofer, Siyu Tang, and Shunsuke Saito. KeypointNeRF: Generalizing image-based volumetric avatars using relative spatial encoding of keypoints. In *ECCV*, 2022. 2, 3
- [28] Marko Mihajlovic, Shunsuke Saito, Aayush Bansal, Michael Zollhofer, and Siyu Tang. COAP: Compositional articulated occupancy of people. In *CVPR*, June 2022. 3
- [29] 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
- [30] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *ICCV*, 2021. 2, 5, 6
- [31] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Unsupervised learning of efficient geometry-aware neural articulated representations. In *ECCV*, 2022. 2
- [32] 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 TOG (Proc. SIGGRAPH)*, 2021. 2
- [33] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *CVPR*, 2022. 6

- [34] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *ICCV*, 2021. 1, 2, 3, 4, 5, 6, 7, 8
- [35] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. 1, 2, 3, 4, 6, 7, 8
- [36] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. Film: Visual reasoning with a general conditioning layer. In *AAAI*, 2018. 5
- [37] Sergey Prokudin, Michael J Black, and Javier Romero. Smpix: Neural avatars from 3d human models. In *WACV*, 2021. 2
- [38] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *CVPR*, 2020. 2
- [39] Edoardo Remelli, Timur Bagautdinov, Shunsuke Saito, Chenglei Wu, Tomas Simon, Shih-En Wei, Kaiwen Guo, Zhe Cao, Fabian Prada, Jason Saragih, et al. Drivable volumetric avatars using texel-aligned features. In *ACM TOG (Proc. SIGGRAPH)*, 2022. 1
- [40] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *ACM TOG (Proc. SIGGRAPH)*, 2022. 2
- [41] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J. Black. SCANimate: Weakly supervised learning of skinned clothed avatar networks. In *CVPR*, 2021. 3
- [42] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *NeurIPS*, 2020. 2
- [43] Shih-Yang Su, Timur Bagautdinov, and Helge Rhodin. Danbo: Disentangled articulated neural body representations via graph neural networks. In *ECCV*, 2022. 1, 2, 3, 5, 6, 8
- [44] Shih-Yang Su, Frank Yu, Michael Zollhöfer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. In *NeurIPS*, 2021. 1, 2, 3, 6
- [45] A. Tewari, O. Fried, J. Thies, V. Sitzmann, S. Lombardi, K. Sunkavalli, R. Martin-Brualla, T. Simon, J. Saragih, M. Nießner, R. Pandey, S. Fanello, G. Wetzstein, J.-Y. Zhu, C. Theobalt, M. Agrawala, E. Shechtman, D. B Goldman, and M. Zollhöfer. State of the Art on Neural Rendering. *Computer Graphics Forum (EG STAR 2020)*, 2020. 2
- [46] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM TOG (Proc. SIGGRAPH)*, 2019. 1
- [47] Garvita Tiwari, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Neural-gif: Neural generalized implicit functions for animating people in clothing. In *ICCV*, 2021. 3, 5
- [48] Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *CVPR*, 2017. 6
- [49] Shaofei Wang, Marko Mihajlovic, Qianli Ma, Andreas Geiger, and Siyu Tang. Metaavatar: Learning animatable clothed human models from few depth images. In *NeurIPS*, 2021. 2, 3, 6, 8
- [50] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. Arah: Animatable volume rendering of articulated human sdfs. In *European Conference on Computer Vision*, 2022. 2, 3, 4, 5, 6, 7, 8
- [51] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 2004. 6
- [52] 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. In *CVPR*, 2022. 2
- [53] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. Synsin: End-to-end view synthesis from a single image. In *CVPR*, 2020. 2
- [54] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *CVPR*, 2021. 2
- [55] Hongyi Xu, Thiemo Alldieck, and Cristian Sminchisescu. H-nerf: Neural radiance fields for rendering and temporal reconstruction of humans in motion. *NeurIPS*, 2021. 2
- [56] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *CVPR*, 2022. 2, 9
- [57] W. Xu, A. Chatterjee, M. Zollhöfer, H. Rhodin, D. Mehta, H.-P. Seidel, and C. Theobalt. Monoperfcap: Human Performance Capture from Monocular Video. *ACM TOG (Proc. SIGGRAPH)*, 2018. 6, 7, 8
- [58] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *NeurIPS*, 2021. 4, 5
- [59] Ilya Zakharkin, Kirill Mazur, Artur Grigorev, and Victor Lempitsky. Point-based modeling of human clothing. In *ICCV*, 2021. 2
- [60] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 6
- [61] Yufeng Zheng, Yifan Wang, Gordon Wetzstein, Michael J Black, and Otmar Hilliges. Pointavatar: Deformable point-based head avatars from videos. In *CVPR*, 2023. 2
- [62] Zerong Zheng, Han Huang, Tao Yu, Hongwen Zhang, Yandong Guo, and Yebin Liu. Structured local radiance fields for human avatar modeling. In *CVPR*, 2022. 2, 3