

DNA-Rendering: A Diverse Neural Actor Repository for High-Fidelity Human-centric Rendering

Wei Cheng¹ Ruixiang Chen^{2*} Wanqi Yin^{2*} Siming Fan^{1,2*} Keyu Chen^{1*}
 Honglin He¹ Huiwen Luo¹ Zhongang Cai³ Jingbo Wang⁴ Yang Gao²
 Zhengming Yu¹ Zhengyu Lin² Daxuan Ren³
 Lei Yang^{1,2} Ziwei Liu³ Chen Change Loy³ Chen Qian¹
 Wayne Wu¹ Dahua Lin^{1,4} Bo Dai^{1†} Kwan-Yee Lin^{1,4†}
¹ Shanghai AI Laboratory ² SenseTime Research ³ S-Lab, NTU ⁴ CUHK




Figure 1: **Overview of our dataset.** DNA-Rendering is a large-scale human-centric dataset, with high-quality multi-view images and videos for various human actors. The dataset comes with grand categories of motion, cloth, accessory, body shape, and human-object interaction. We hope it could boost the development of human-centric rendering and related tasks.

Abstract

Realistic human-centric rendering plays a key role in both computer vision and computer graphics. Rapid progress has been made in the algorithm aspect over the years, yet existing human-centric rendering datasets and benchmarks are rather impoverished in terms of diversity (e.g., outfit’s fabric/material, body’s interaction with objects, and motion sequences), which are crucial for rendering effect. Researchers are usually constrained to explore and evaluate a small set of rendering problems on cur-

rent datasets, while real-world applications require methods to be robust across different scenarios. In this work, we present **DNA-Rendering**, a large-scale, high-fidelity repository of human performance data for neural actor rendering. DNA-Rendering presents several alluring attributes. First, our dataset contains over 1500 human subjects, 5000 motion sequences, and 67.5M frames’ data volume. Upon the massive collections, we provide human subjects with grand categories of pose actions, body shapes, clothing, accessories, hairdos, and object intersection, which ranges the geometry and appearance variances from everyday life to professional occasions. Second, we provide rich assets for each subject – 2D/3D human body keypoints, foreground

*Joint-first authors with W. Cheng.

†Equal advising.

Dataset	Attribute					Scale				Realism
	Ethnicity	Age	Cloth	Motion	Interactivity	#ID × #Outfit	#Motions	#View	#Frames	HRes
Human3.6M [24]	✗	✗	✗	✓	✓	11 × 1	17	4	3.6M	1000P
CMU Panoptic [29]	✓	✓	✗	✓	✓	97 × 1	65	31 + 480*	15.3M	1080P
ZJU-MoCap [53]	✗	✗	✗	✓	✗	10 × 1	10	24	180K	1024P
HUMBI [76]	✓	✓	✓	✗	✗	772 × 1	—	107	26M	1080P
AIST++ [65, 34]	✗	✗	✗	✗	✗	30 × 1	—	9	10.1M	1080P
THuman 4.0 [60]	✗	✗	✓	✓	✗	3 × 1	—	24	10K	1150P
HuMMan [5]	✓	✓	✓	✓	✗	1000 × 1	500	10	60M	1080P
GeneBody [11]	✓	✓	✓	✓	✓	50 × 2	61	48	2.95M	2048P
DNA-Rendering(Ours)	✓	✓	✓	✓	✓	500 × 3	1187	60	67.5M	4096P

Table 1: **Dataset comparison on attributes and scales.** We compare the proposed dataset with previous human-centric multiview datasets in terms of attribute coverage, scale, and realism. ‘Ethnicity’ denotes whether the dataset contains actors from multiple ethnic groups. ‘Age’ means if there is a wide age range containing elders or infants. ‘Cloth’ separates datasets with only daily costumes or with extra diverse clothing. ‘Attribute-Motion’ denotes whether it has human motion in different scenarios. ‘Interactivity’ tells whether there contains human-object interaction. We mark these attributes with ✓ and ✗. In scale, we list the number of key factors with compared dataset, Note that ‘Scale-#Motions’ means the number of motion categories, and superscript * means low-resolution VGA cameras, we exclude them during ‘#View’ ranking and ‘#Frames’ calculation. Cells indicate the best, and second best in the specific category among all datasets. We abbreviate resolution at height as ‘HRes’.

masks, SMPLX models, cloth/accessory materials, multi-view images, and videos. These assets boost the current method’s accuracy on downstream rendering tasks. Third, we construct a professional multi-view system to capture data, which contains 60 synchronous cameras with max 4096×3000 resolution, 15 fps speed, and stern camera calibration steps, ensuring high-quality resources for task training and evaluation.

Along with the dataset, we provide a large-scale and quantitative benchmark in full-scale, with multiple tasks to evaluate the existing progress of novel view synthesis, novel pose animation synthesis, and novel identity rendering methods. In this manuscript, we describe our DNA-Rendering effort as a revealing of new observations, challenges, and future directions to human-centric rendering. The dataset, code, and benchmarks will be publicly available at <https://dna-rendering.github.io/>.

1. Introduction

Understanding humans is an everlasting problem in our research community, and extensive literature on perceiving and synthesizing humans shows great efforts toward this goal. Over the decades, multiple pioneers have constructed large-scale and diverse datasets, such as COCO [37] for human pose estimation, and ActivityNet [4] for analyzing human action. These datasets are the driving force behind flourishing developed human-centric perceiving algorithms.

However, for human-centric rendering, the vacancy for a comprehensive dataset is still there. Capturing high-quality and massive 3D/4D human avatars is difficult due to the requirements of high-end equipment as well as an efficient data processing pipeline. Existing datasets [24, 28, 53, 60, 23] partially narrow the gaps but have significant limitations on sample diversity (e.g., clothing, motion, body shape, and human-object interaction), or have insufficient realism (e.g.,

camera resolution, and capture speed). These factors are crucial to rendering effects.

To drive advance in human-centric rendering, we contribute a large-scale multi-view human performance capture dataset, named DNA-Rendering, which includes the factors that are important to rendering in great diversity and granularity. *On the hardware side*, we build up a 360-degree indoor system equipped with 60 calibrated RGB cameras and 8 synchronized depth sensors. The captured videos are under the fidelity of up to 12MP (4096×3000) resolution and recorded at 15 fps. *From the dataset’s footage design aspect*, we intend to cover most attributes (e.g., age, ethnicity, shape, motion, cloth, accessory, and interactive objects) that could reflect the rendering differences with respect to texture, materials, primary/secondary motion deformation, and category priors. In practice, we design over 1500 outfits and 1187 motion types to ensure the comprehensive coverage of real-world scenarios. We invite 500 actors to participate in the data capture process. We record each person with three different outfits and at least nine unique motion sequences. The full dataset contains 5000 video sequences with over 67.5M frames. Compared with the existing human-centric dataset like CMU Panoptic [29], ZJU-MoCap [53], THuman [60], and Human3.6M [24], DNA-Rendering comprises the most multi-view body performance samples and reaches the highest image quality. The unfold comparisons between DNA-Rendering and the others are given in Tab. 1.

Meanwhile, we provide essential annotations attached to each frame to facilitate the application of downstream tasks. The ultra-large scale of the DNA-Rendering dataset raises great challenges to the corresponding data processing steps. To this end, we develop an automatic annotation pipeline encompassing camera calibration, color correction, image matting, 2D /3D landmark estimation, and SMPLX model fitting. To ensure the labeling quality, we developed a series

of technical refinements to the annotation toolchain. With these efforts, the automatic pipeline can generate faithful data annotations both effectively and efficiently. To further assist the community in the public use of external data, we will open-source the annotation tool used in the DNA-Rendering project.

The unprecedented richness of DNA-Rendering dataset provides fertile data soil for researchers to develop, and dissect their rendering methods in depth. To set up a kickoff example, we further construct benchmarks upon the dataset with extensive experiments. We evaluate the performances of several state-of-the-art full-body rendering and animation approaches under three major tasks, *i.e.*, *novel view synthesis*, *novel pose animation*, and *novel identity rendering*. To better analyze current methods in terms of the model capacity, module necessity, and methodology generality, we set up multiple test set splits under different levels of challenging aspects. For instance, we divide the *easy*, *medium*, and *hard* subsets *w.r.t.* the cloth looseness, the texture complexity, the motion difficulty, and the human-object interactivity, respectively. We conclude a series of key observations based on the benchmarks, such as how human prior influences the robustness of rendering, how sensitive the multi-view/frame relationship module design is to data volume/distribution, and how loss design affects the performance in terms of different rendering metrics.

In summary, we contribute DNA-Rendering project to fulfill the requirement of a high-fidelity human performance capture dataset for the research community. We establish by far the largest multi-view human body performance dataset for high-fidelity human-centric rendering research, with an emphasis on image quality and data attributes. The attached benchmarks provide baseline standards for three major tasks, with rigorous evaluations and dissections on multiple state-of-the-art methods. We believe the dataset, the attached benchmarks, and the tools will boost a wide range of digital human applications and inspire future research.

2. Related Works

2.1. Human-centric Datasets

Perception Datasets. Perceiving human is a long-standing problem in the research community. Over the decades, researchers keep dedicating their efforts to building relevant datasets. Earlier efforts in the computer vision community present large-scale datasets like COCO [37] for human segmentation or keypoint detection from in-the-wild images. Later research works follow the inspiration to establish open-world datasets [8, 19], while with emphasis on parsing more precise human body parts. Some researchers focus on constructing datasets [7, 49, 59] that capture daily activities and help the perception of human action recognition by using RGB-D cameras. Despite the wild variety of data samples, these datasets are not capable of hu-

man rendering tasks, due to a lack of multiview images as groundtruth references for evaluating methods' performance. In computer graphics society, there is another parallel branch that contributes datasets [33, 30] for avatar animation, with recording human motion via maker-based motion capture systems. AMASS [44] further integrates these motion capture databases with fully rigged surface mesh representation. In the last decade, computer vision and graphics society fit in with each other in the field of perceiving 3D humans. Datasets like Human3.6M [24] and MPI-INF-3DHP [45] capture humans under in-door multiview environment, with 3D marker label or multiview segmentation which further encourage the applications in recovering human in 3D. These databases facilitate the development of numerous algorithms. However, due to the limits of the data sample, camera views, and resolution, they cannot well reflect the pros and cons of rendering methods.

Rendering Datasets. Representing 3D /4D human appearances and performances are important in both research communities and commercial applications. The progress of relevant algorithms relies on human scan datasets with accurate geometry or dense views. Such datasets are the foundational factor to close the gap between virtual avatars and real humans. THuman [82, 74, 63] and commercial scan datasets [17, 16] capture static human scan reconstructed by either depth sensors or camera array. [2, 3, 77, 60, 80] provide dynamic human scans with minimal clothing and daily costumes. These datasets are usually biased centering on standing poses due to the sophisticated capture process. With the emergence of neural rendering techniques, rendering realistic humans directly from images has become a trend. Such a setting usually requires the dataset equipped with high-quality dense view images and accurate annotations like human body keypoints and foreground segmentation. CMU Panoptic [28] uses a 30-HD-camera system and annotates the humans with 3D keypoints. HUMBI [76] focuses on local motions like gesture, facial expression, and gaze movements, rather than factors that have influences on rendering quality, such as cloth texture, object interactivity, *etc.* ZJU-MoCap [53] is a widely used dataset for human rendering algorithms. It includes ten video sequences with 24 cameras under 1K resolution and provides annotations of human segmentation and estimated SMPL model for each frame. However, ZJU-MoCap is limited in narrow motion and clothing diversity, which might lead the evaluation to great bias. AIST++ [65, 34] is a dance database with various dance motions while sticking in the one-fold scenario and lacking view density. Recently proposed HuMMan [5] and GeneBody [11] datasets, expand the motion and clothing diversity, while the effective human resolution is still below 1K. Differing with these efforts, we make a step further and contribute the largest high-fidelity multiview dataset for human-centric rendering tasks, with 5000 human perfor-

mance sequences under a large variation in ethnicity, age, cloth, motion, and interaction. Concurrent works [83, 23] also contribute datasets for human avatar tasks, while centering on detailed human geometry with long-lens cameras to film human body parts. In practice, [83] captures on low frame rate, and [23] enjoys dense multi-view capturing but limits to 16 motion sequences. Moreover, their setups are vulnerable to the capture of aboriginal motion tracks at full-body levels.

2.2. Implicit Neural Body Representation

Different from previous works that represent humans with explicit representations (such as skeleton, parametric model, or mesh), recent work models human appearance as neural implicit function, *e.g.*, neural radiance fields [47] or neural signed distant functions. PIFu [56, 57] presents the orthogonal camera space as an occupancy function conditioned by pixel-aligned features and depth. It learns to reconstruct the human body with single-view images. Neural-Body [53] learns a neural radiance field of dynamic humans conditioned by body structure and temporal latent code from sparse multi-view videos. Recently, many category-agnostic implicit representations, PixelNeRF [73], IBR-Net [69], VisionNeRF [35], *etc.*, can generalize NeRF to arbitrary unseen scenes given a set of reference views. The key insight behind these methods is that the novel views of new scenes can be recovered from visual clues of reference views and neural implicit functions that are derived by rendering equations. Since rafting these two factors together can help models learn visual consistency on certain camera distributions. The intrinsic differences among these methods are the design of feature aggregation, which varies from average [73], max [54] pooling to more adaptive weighted pooling [69] and vision transformer [35]. Given human rendering is more challenging due to the large variation in pose and appearance, recent generalizable human rendering methods [79, 46, 11, 32] condition such image feature-aligned NeRF with human priors. For example, Neural-HumanPerformer [32] uses structured latent code and KeypointNeRF [46] deploys human keypoints. These methods outperform category-agnostic methods on human rendering on existing human rendering datasets, while such improvement is not convincing due to the limited scale and coverage bias of these datasets. We believe the proposed large-scale dataset with large variation overall dimensions will help both the robustness of generalizing humans and better generalization ability assessment.

2.3. Animatable Digital Human

The challenge of creating realistic animatable human avatars from images is two folds – (1) how to reconstruct the human body from motion sequences and (2) how to disentangle non-rigid deformation. Early seminal work A-

NeRF [62] learns dynamic body from sequences, it conditions the radiance field with relative pose coordinate of the query point, which fails to model the non-rigidity of clothed humans. To reconstruct the human body from sequences, AnimatableNeRF [52] learns a static canonical radiance field together with a ray blending network from the current frame to canonical space. To further better disentangle motion deformation from pose recent works [58, 9] use a complementary forward blending network or root-finding algorithm to regularize the learned blending with cycle consistency loss. Other works [71, 25, 75] learn animatable models from more challenging monocular video, with a tighter assumption of Gaussian distributed occupancy along bone or fixed SMPL motion weights.

3. DNA-Rendering

In this section, we discuss the core features of our DNA-Rendering from the perspective of dataset construction, including the hardware systems setup, data collection, and relevant dataset statistics. By introducing these aspects, we present a way to ensure high fidelity and effectiveness under the massive amount of data capture.

3.1. Dataset Capture

System Setup. Our capture system contains a high-fidelity camera array, with 60 high-resolution RGB cameras and 16 lighting boards uniformly distributed in a sphere with a radius of three meters. The cameras are adjusted to point at the sphere’s center, where the participants perform. Concretely, the array consists of 48 high-end 2448×2048 industrial cameras, and 12 ultra-high resolution cameras with up to 4096×3000 resolution. We use these industrial cameras to capture subjects’ body and facial performances, and the ultra-high resolution cameras to capture more detailed textures on clothing and accessory. We additionally place eight Kinect cameras to capture additional depth streams as auxiliary geometric data. The high-fidelity video streams and depth streams are synchronized at 15 frames per second. To guarantee the cameras can clearly capture subjects’ performances and clothing patterns from all views, the lighting is adjusted to the setting of $5600K \pm 300K$ color temperature and 4500 Lux/m illuminance. The above designs ensure the system could record the sharp texture edges, fine-grained color changes of clothing patterns, and the reflection effects caused by different clothing materials. Please refer to Sec. A.4 for more details.

Data Collection Protocol. To enable subsequent research probing into the factors that have influences on rendering, we design a data collection protocol with both interlaced and hierarchical data attributes. Specifically, we ask each actor to wear three sets of outfits and perform at least three actions in different hallucinated scenarios for each outfit, which maximize the identity scale and diversity. Each mo-

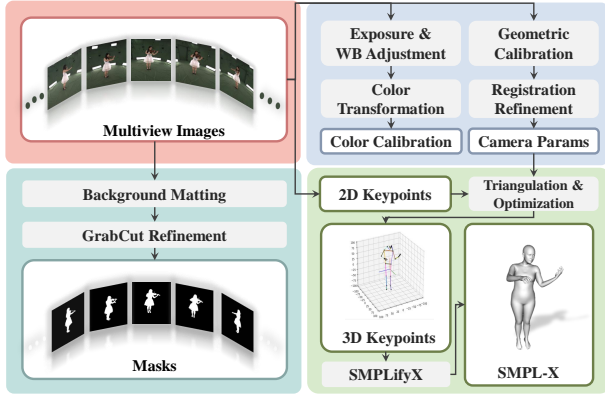


Figure 2: **Annotation pipeline.** The illustration of annotation pipeline for camera calibration, camera color calibration, masks, keypoints, and parametric model.

tion sequence is recorded under specific action category instruction with a free-style performance lasting for 15 seconds, which ensures the diversity of action performance. As an auxiliary feature, we also capture a static frame of A-pose for actors in each outfit for canonical pose recording, and a frame with only empty background for image matting. Multiview videos are inspected on-site with a quick preview generation tool, which ensures the movement fits best to the cameras’ field of view. For accurate camera pose annotation, extrinsic calibration data are collected at a daily frequency. The color data and intrinsic calibration data are collected whenever system adjustments are made. Please refer to Sec. A.4 for details.

3.2. Dataset Statistics

In order to cover diverse attributes that relate to rendering quality, we have carried out a detailed design from the selection of the actors’ gender, age, and skin color, to their actions, clothing, and makeup. Compared with previous work, to our best knowledge, our dataset contains the largest number of human subjects, covering the most diverse action categories, clothing types, and human-object interaction scenarios. The key statistics of our dataset are shown at the bottom of Fig. 1. Specifically, to preserve authenticity in action behavior, we invite 153 professional actors to perform special scenes with corresponding costumes/makeup, and 347 normal performers to act under footage of daily-life scenes. The special scenes constitute 153 sub-categories, including sports, dances, and unique event performances such as typical costumes in ancient Chinese dynasties, traditional costumes around the world, cosplay, etc. Common scenes can be divided into 269 sub-categories, covering scenes such as daily indoor activities, communication, entertainment, and new trends. We describe the comprehensive distribution of data in Sec. A.1 and the limitation of data in Sec. A.5. Please refer to the corresponding sections for more visualized figures and detailed discussion.

3.3. Data Annotation

To enable applications in human rendering and animation, DNA-Rendering provides rich annotations attached with the raw data, *i.e.*, camera calibration, camera color calibration, image matting, and parametric model fitting. The overall annotation pipeline is shown in Fig. 2.

Camera Calibration. As illustrated in the blue box of Fig. 2, we separate the multi-camera calibration into two stages. In the first stage, we calibrate the intrinsic parameters of each camera individually. Specifically, we divide the camera’s field of view into a 3×3 Sudoku, and capture images with ± 30 degree rotation in pitch, row, and yaw angle of checkerboard in all Sudoku grids, referring to Fig. S4. For extrinsic calibration, we deploy multiple ChArUco boards and spin the main board in the capture volume. We use the toolbox xrprimer [14] and multical [1] to optimize intrinsic parameters, distortion coefficients, and extrinsic parameters with the captured data. To eliminate the depth camera pose error caused by the large resolution gap between industrial cameras and Kinect depth cameras, we further adopt a point cloud registration stage to refine the depth camera extrinsic parameters in the second stage. More concretely, for each depth camera, we project the partial point cloud and estimate a full point cloud from the MVS algorithm such as [67] as a reference. We jointly optimize the pose graph of the depth camera for neighboring pointclouds with overlaps through a multi-way registration [12] with MVS pointcloud as reference.

Benefiting from the above camera calibration refinement strategy, the reprojection error of optimized corner points to detected corner points across all the camera views is optimized from 1.37 to 0.94 pixels. For detailed camera calibration estimation, please refer to Section. B.1 in the supplementary.

Color Calibration. The identical color response across different cameras could be vital for a multi-view, mixed-type camera system to provide qualified data for rendering applications, as it is an essential data basis for algorithms to render realistic view-dependent effects. Different from other multi-camera datasets, *e.g.*, ZJU-MoCap [53] that has noticeable camera brightness difference, or Multiface [72, 40] which use a network to optimize the color transformation during model training, we pay attention to ensure the color consistency of data collection across different cameras. First, we conduct careful adjustments on hardware parameters such as exposure and white balance to make the captured color of the color checkerboard under the standard light as close as possible. Then, the 2-order polynomial color correction coefficients could be optimized by least square regression of transforming the detected color to the true value on the color checkerboard. Please refer to Sec. B.1 and Fig. S5 for the detailed descriptions. We also analyze the impact of color consistency of multi-camera

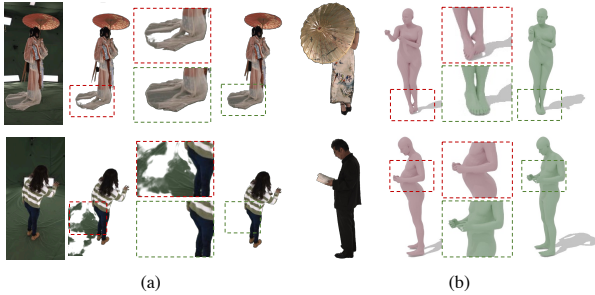


Figure 3: **Annotation quality improvements.** The zoom-in boxes with red dot lines show the annotation quality before pipeline optimization. The green ones show quality improvements over (a) mask annotation and (b) SMPLX annotation with the optimized pipeline.

datasets on generalizable neural rendering in Sec. D.4.

Matting. Considering the large quantities of the captured images, we develop an automatic matting pipeline to extract the foreground objects from the backgrounds. We first adopt an off-the-shelf background matting model [36] to eliminate most background pixels. However, due to the complicated nature of the capture settings, the learning-based model inevitably generates unsatisfying results in some challenging cases, leaving some pieces of labeled data with artifacts such as broken holes or noisy patches (See Fig. 3). Thus, we further propose a refinement strategy by applying the *HSV* filtering and the grabcut [55] algorithms to improve the matting quality. According to the manual assessment of the refined masks, the error rates of problematic cases are reduced from 11% to 2% on average. We compare matting with and without refinement, and visualize detailed manual assessment in Fig. S8.

Keypoints and Parametric Model. Inspired by existing works [5, 6], we develop an automatic pipeline to annotate keypoints and parametric model parameters. 1) First, 2D keypoints in COCO-Wholebody [26] format (including body, hand, and face keypoints) are detected for each camera view, with pretrained model HRNet-w48 [64]. 2) Then, we triangulate 3D keypoints with known camera intrinsic and extrinsic parameters from the multi-view 2D keypoints with optimization and post-processing strategies [15] including keypoint selection, bone length constraint, as well as outlier removal. 3) Finally, we register the SMPLX, a commonly used parametric model, via 3D keypoints. Body shape $\beta \in \mathbb{R}^{n \times 10}$ (or $\beta \in \mathbb{R}^{n \times 11}$ for children [22, 50]), pose parameters (body pose, hand pose, and global orientation) $\theta \in \mathbb{R}^{n \times 156}$, and translation parameters $t \in \mathbb{R}^{n \times 3}$ (n is the number of frames) are estimated via a modified SMPLify-X [51] for dynamic poses.

Our automatic annotation pipeline is proved effective and robust in getting natural SMPLX models. As shown in Fig. 3, twists in feet and body trunk caused by insufficient keypoints regularization can be fixed by introducing human priors during registration. We evaluate the fitting er-

ror between 3D keypoints and corresponding regressed SMPLX joints. The mean and median ‘Mean Per Joint Position Error’ (MPJPE) of our system is 30.20 mm and 29.80 mm. The error is on par with the oracle fitting accuracy of 29.34 mm in Human3.6M [24, 41], which includes data from an optical motion capture system. Detailed analysis is conducted in Sec. B.3 and Sec. B.4. A thorough comparison of our fitting pipeline with other fitting methods [61, 81, 11] is described in Sec. B.5.

4. Benchmarking Human-centric Rendering

Our DNA-Rendering dataset could be used to unfold the reflections and boot the developments of research on high-fidelity human body rendering tasks, due to its *large-scale volume, diverse scenarios, multi-level challenges, and high-resolution multi-view data* properties. To kick off an example of how to utilize this dataset, we set up benchmarks with exclusive experiments centered around three fundamental tasks of human body rendering, *i.e., novel view synthesis, novel pose animation, and novel identity rendering*. In this section, we introduce the benchmark settings and key observations.

4.1. Data Splits

To unfold each method in depth, and thoroughly evaluate the effectiveness of our dataset, we construct multiple training and testing data splits to conduct level tests for each method. We consider the four most influential factors of rendering quality for the benchmark test, *i.e., the looseness of clothes, the texture complexity, the pose difficulty, and the interactivity between the human body and manipulated object*.

The Cloth Looseness. We define the cloth’s challenging levels by the deformation distance between the minimal-cloth human body and the clothing outline, and the softness of cloth materials. For the *Easy* level, we collect simple cases wearing tight-fitting clothes like yoga wear and sports t-shirts. The *Medium* level includes the daily clothes such as coats, skirts, jeans, loose t-shirts, *etc.* As for the *Hard* level, we make a split containing ethical costumes, national clothing, and fancy decorations.

The Texture Complexity. The texture distribution also plays an important role in the human body rendering tasks. To examine the correlations between texture complexity and rendering performance, we build three data splits for texture evaluation. The *Texture-Easy* split is composed of single-color clothes. The *Texture-Medium* split includes most daily clothes in a few colors and plain patterns. The *Texture-Hard* split contains the most complicated texture clothes with intricate patterns like dots, stripes, plaids, *etc.*

The Pose Difficulty. In the novel pose animation task, it is vital to probe if the trained models could handle different levels of motion sequences in terms of (non-rigid) diffi-

culties and degree of out-of-distribution (OOD). Therefore, we prepare the *Motion-Easy*, *Motion-Medium*, and *Motion-Hard* splits for training and evaluation. The *Easy* data are simple motions with limited body parts involved, like shaking and waving hands. The *Medium* level refers to casual motions including full-body actions such as walking, eating, sitting, kneeling, stretching, *etc.* Moreover, the *Hard* split is designed to cover the extremely challenging motion cases that are performed by professional sports players or actors, *e.g.*, instrument playing, sports action, yoga, and dancing.

The Human-Object Interactivity. Finally, we propose to evaluate the impact of human-object interactivity by novel view synthesis and novel pose animation tasks. The data difficulty level is determined by the object size. Specifically, we define four data splits: the *Interaction-No* split contains pure human motions with no interactive objects; the *Interaction-Easy* split includes rigid small-size hand-held objects like cellphones, pencils, cigarettes, and cups. The *Interaction-Medium* split has middle-size hand-held objects, *e.g.*, handbag, volleyball, newspaper, *etc.* This split includes both rigid motions and non-rigid object motions; and the *Interaction-Hard* split consists of large-size assets such as yoga mats, desks, chairs, and sofas.

To sum up, we construct an overall train split consisting of 400 sequences with even distribution on all human factors and difficulties, and 13 test factor-difficulty splits in total with three sequences in each test split.

4.2. Task Definition

Human body rendering and animation problems have been popular research topics in the digital human area for decades. With the development of implicit scene representations like NeRF [47], many researchers design algorithms upon the methodology. In conducting the benchmark experiment on the DNA-Rendering dataset, we intend to provide a thorough comparison of the state-of-the-art methods on data splits with different difficulty levels. Depending on the generalizability of the state-of-the-art methods, we categorize the recently published works into two classes: case-specific methods and generalizable ones. We evaluate the methods under multiple problem settings according to their categories. Concretely, *we set up novel view synthesis and novel pose animation tasks for the case-specific methods, and the novel identity rendering task for the generalization approaches.* In this section, we present the key observations of the benchmarks. We provide a detailed review of the benchmark methods and necessary modifications to apply them to our dataset in Sec. C.1 of the supplementary.

Novel View Synthesis. Recent *dynamic* human rendering works like NeuralBody [53], A-NeRF [62], AnimatableNeRF [52], and NeuralVolumes [39] obtained impressive results by training on a single case with multi-view

video data. HumanNeRF [71] demonstrated the ability to render realistic novel view images of humans from monocular video sequences. In this task, we adopt the official architecture implementation of the case-specific methods and train each individual model for every single case in the DNA-Rendering test set. For a fair comparison, we unify the training setting of NeuralVolumes [39], A-NeRF [62], NeuralBody [53], AnimatableNeRF [52], and HumanNeRF [71] with 42 dense training views. We evaluate the image rendering quality of these methods on the other 18 unseen testing camera poses. Meanwhile, we also train two general scene *static* methods – Instant-NGP [48] and NeuS [68], in each testing frame with the same training views. These two methods’ performances could serve as the per-frame static reconstruction baseline reference. The rendering results are analyzed based on the difficulty level of data splits.

Novel Pose Animation. Similar to the novel view synthesis task, we conduct novel pose animation benchmark on the four case-specific methods [53, 52, 62, 71]. For each test case, we split the sequence into two parts, where images from the first 80% frames are used for training and the ones from the last 20% are used for testing. Besides, for the SMPL-guided pose animation methods [53, 52, 71], we provide the *SMPL parameters* of test images for the models to infer rendering. As for the SMPL-free method [62], the trained models take the target *pose images* as the input (*i.e.*, the underlying skeletons), and optimize the novel pose representations for inference.

Novel Identity Rendering. The other category of our benchmark methods is the generalizable algorithms that can be trained on multiple cases and infer across different unseen identities. Specifically, we probe three general scene generalizable methods – PixelNeRF [73], VisionNeRF [35], IBRNet [69], and two human-centric methods – NeuralHumanPerformer [32], and KeypointNeRF [46]. To fairly compare their performances on unseen identities, we use the same training set (all training samples of the three splits, which results in 400 sequences in total) to train the generalizable models. In the inference stage, we evaluate the image rendering quality on novel cases from each test split respectively.

4.3. Benchmark Results

As introduced in Sec. 4.1, we construct a test set with 13 sub-splits according to the four most concerned attributes (*Deformation*, *Motion*, *Texture*, and *Interaction*) in different difficulty levels (*Easy*, *Medium*, and *Hard*), and an extra *No* level for *Interaction*. This results in a data volume of 39 motion sequences for testing. For all rendered images, three metrics are computed – PSNR, SSIM [70], and LPIPS [27] (LPIPS* denotes LPIPS×1000). We evaluate more than 10 state-of-the-art methods on these splits and

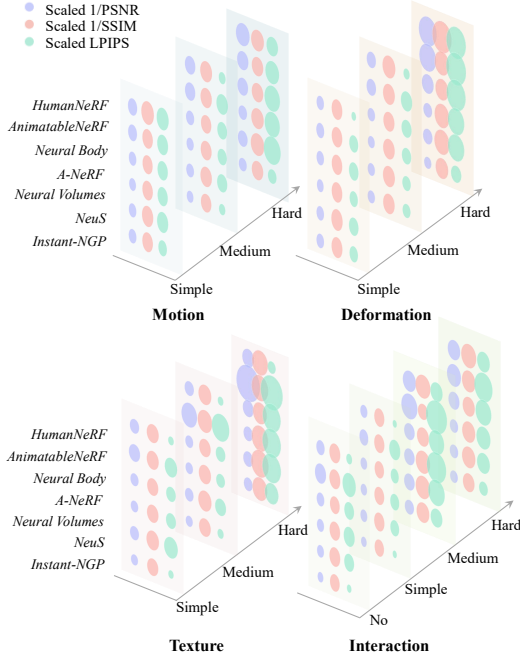


Figure 4: **Quantitative results visualization of novel view synthesis test across benchmarks splits and difficulties.** The colored circles denote different metrics, the smaller the circles indicate the better the novel view quality the method achieved. The numbers are reported in the appendix (Tab. S2).

analyze their performances under the same metrics. The experiment analysis is given below. Noted that due to limited space in the main paper, we provide the detailed setting, thorough discussions, and additional results in Sec. C in the supplementary.

Novel View Synthesis. We visualize the bubble diagram of quantitative results across all benchmark splits in Fig. 4. The precise numbers of the quantitative results are reported

in Tab. S2 in the supplementary. We conclude three key observations in the main paper: (1) Generally speaking, the rendering quality is inversely proportional to split difficulties, as reported in Fig. 4, where the circles get bigger when the difficulty grows. (2) Among all case-specific dynamic methods, A-NeRF [62] achieves the best PSNR performance, and NeuralBody [53] and HumanNeRF [71] gets the best SSIM and LPIPS respectively. Qualitative results are shown in Fig. 5, NeuralBody [53] and A-NeRF [62] could render novel view image with fewer background artifacts than other methods, while HumanNeRF [71] can better preserve high fidelity textures, especially in high-frequency texture regions. (3) When it comes to rendering novel views for trained human action frames, the difficulties increase on *Motion* and *Interaction* dimensions will not separate the performances among dynamic methods to a large extent, while hard *Texture* cases enlarge the performance gap among dynamic methods (refer to T-shirt case with stripe pattern in Fig. 5). Meanwhile, dynamic methods’ performances on *Texture* degrade the most when difficulty rises compared to the static baselines (refer to bubbles in Fig. 4). *Texture-Hard* and *Deformation-Hard* shows the high texture and non-rigidity are still challenging for simple dynamic modeling with bone-coordinates [62] or SMPL-skinning based blending [52], even in seen-poses. More qualitative results in each benchmark split are shown in Fig. S9, and we analyze the conceptual difference of these methods in Sec. C.2.1 in the appendix.

Novel Pose Animation. Similar to novel view synthesis, when split difficulty increases the rendering quality decreases as shown in Tab. 2. Among all data factors, we found that *Deformation* and *Interaction* are insurmountable factors for current methods to model in novel poses. Qualitative results are displayed in Fig. 6, none of the

Splits	PSNR \uparrow					SSIM \uparrow					LPIPS* \downarrow				
	NV	AN	NB	AnN	HN	NV	AN	NB	AnN	HN	NV	AN	NB	AnN	HN
Motion-Simple	22.05	26.65	25.84	22.78	24.65	0.947	0.965	0.974	0.958	0.953	78.30	58.04	58.33	74.33	62.76
Motion-Medium	19.30	21.73	21.84	21.41	21.14	0.941	0.951	0.969	0.957	0.952	92.80	71.81	65.46	80.97	54.46
Motion-Hard	19.17	21.49	20.43	19.64	22.48	0.938	0.952	0.965	0.949	0.964	105.46	83.58	82.85	97.98	51.18
Deformation Simple	20.42	25.44	24.57	23.62	26.15	0.939	0.957	0.968	0.958	0.967	84.65	53.30	59.04	61.12	30.18
Deformation-Medium	23.09	27.26	27.05	23.52	24.97	0.945	0.963	0.974	0.961	0.958	61.09	48.41	49.91	65.43	33.94
Deformation-Hard	20.11	20.88	20.27	19.41	19.70	0.925	0.926	0.956	0.943	0.924	117.31	108.89	102.84	103.22	102.67
Texture-Simple	20.99	26.21	25.54	23.12	25.65	0.954	0.974	0.982	0.970	0.974	77.68	49.88	50.48	67.40	28.81
Texture-Medium	25.44	27.94	25.77	23.15	27.19	0.959	0.966	0.977	0.962	0.969	56.68	43.94	48.44	67.00	24.04
Texture-Hard	20.95	23.22	22.05	18.45	23.78	0.916	0.927	0.951	0.943	0.945	117.93	98.43	96.01	101.09	41.84
Interaction-No	22.64	26.32	25.41	22.44	25.93	0.957	0.968	0.980	0.967	0.968	71.98	62.55	54.71	69.29	31.61
Interaction-Simple	24.28	27.57	26.42	23.18	27.18	0.965	0.976	0.983	0.968	0.975	55.54	48.35	45.86	65.36	23.31
Interaction-Medium	20.37	23.67	21.96	20.81	23.20	0.934	0.950	0.965	0.953	0.951	95.79	84.74	87.41	97.32	52.10
Interaction-Hard	21.14	25.00	22.10	21.29	22.29	0.931	0.949	0.961	0.953	0.940	94.04	79.40	89.63	91.82	70.54
Overall	21.53	24.88	23.79	21.76	24.18	0.942	0.956	0.970	0.957	0.957	85.33	68.56	68.54	80.18	46.73

Table 2: **Benchmark results on novel pose task.** State-of-the-art methods’ rendering performance on novel poses for each benchmark split. We abbreviate NeuralVolumes [39] as ‘NV’, A-NeRF [62] as ‘AN’, NeuralBody [53] as ‘NB’, AnimatableNeRF [52] as ‘AnN’ and HumanNeRF [71] as ‘HN’. Cell color ■ ■ ■ indicate the best, second best, and third best performance in the same split respectively.

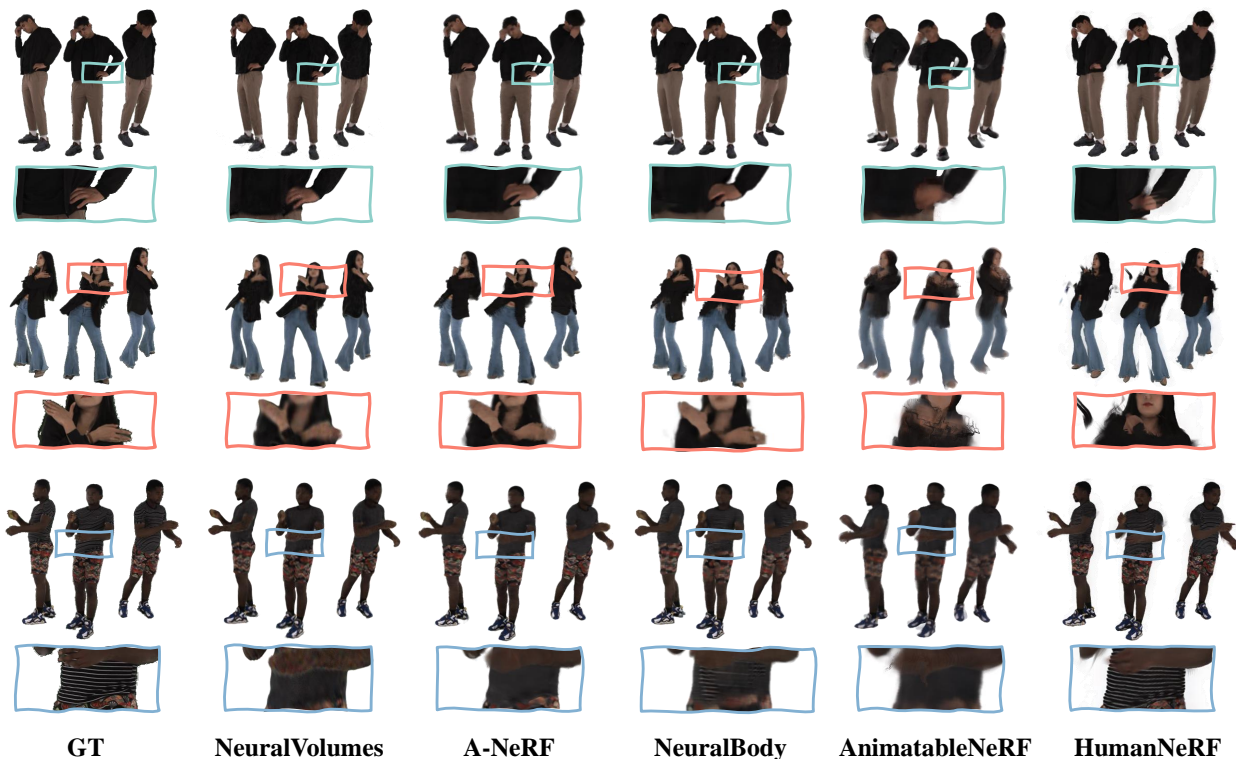


Figure 5: **Visualization of novel view synthesis result samples.** We qualitatively compare the novel views of the test frames in our test splits. More qualitative novel results in each split are shown in Fig. S9.

methods can generate reasonable deformation in the case of the Peking opera costume. NeuralBody [53] and AnimatableNeRF [52] can not model the interactive objects, and the objects are stretched when given large poses in A-NeRF [62]. Conclusively, current state-of-the-art methods can learn relatively reasonable human avatars with even hard *Motion* and *Textures*, while stuck in the imperfectness of modeling hard *Deformation* and *Interaction*. These two animation challenges should stimulate the communities for further investigation. More detailed analysis is provided in

Sec. C.2.2 in the appendix.

Novel Identity Rendering. We report the quantitative metrics of all 39 novel identities in Tab. 3. Generally, generalizable methods with human prior [46, 32] performs better with higher robustness than category-agnostic methods [73, 69, 35]. Among category-agnostic methods, IBR-Net [69] directly blends pixel color from source views, and it outperforms PixelNeRF and VisionNeRF that predict radiance color only from image features. We draw the conclusion that, in generalizable human rendering, human

Splits	PSNR \uparrow					SSIM \uparrow					LPIPS* \downarrow				
	IBR	PN	VN	NHP	KN	IBR	PN	VN	NHP	KN	IBR	PN	VN	NHP	KN
Motion-Simple	26.13	26.04	25.90	25.61	24.67	0.964	0.957	0.959	0.961	0.964	65.48	72.68	72.29	65.53	44.77
Motion-Medium	25.56	25.84	25.22	24.63	24.34	0.966	0.960	0.961	0.963	0.971	59.71	63.80	64.83	61.08	38.22
Motion-Hard	23.78	24.49	23.72	23.43	24.36	0.959	0.950	0.949	0.956	0.973	79.18	89.93	93.04	80.20	43.79
Deformation Simple	26.72	26.85	26.31	26.41	26.01	0.965	0.960	0.960	0.963	0.965	53.73	61.92	63.48	48.45	34.68
Deformation-Medium	27.46	27.55	27.90	27.28	25.83	0.965	0.961	0.965	0.963	0.972	57.30	66.82	62.82	56.02	38.00
Deformation-Hard	23.98	20.77	20.05	22.64	23.00	0.942	0.841	0.838	0.936	0.943	87.04	282.36	264.57	92.22	67.76
Texture-Simple	25.70	26.27	25.62	25.43	22.72	0.973	0.967	0.968	0.973	0.971	63.66	69.16	70.93	58.70	50.07
Texture-Medium	26.15	26.76	26.40	26.25	25.14	0.965	0.960	0.963	0.963	0.968	54.45	56.66	56.58	53.95	35.10
Texture-Hard	23.34	24.08	23.61	23.45	22.91	0.932	0.921	0.922	0.933	0.935	98.77	106.05	104.58	90.84	72.87
Interaction-No	26.08	25.91	26.39	25.46	23.43	0.968	0.958	0.963	0.966	0.968	61.99	72.67	66.91	64.35	44.95
Interaction-Simple	27.60	25.76	27.54	26.67	26.12	0.976	0.944	0.973	0.974	0.977	50.50	82.30	53.38	50.77	28.60
Interaction-Medium	24.04	24.44	24.09	23.61	22.70	0.950	0.937	0.941	0.947	0.950	84.64	96.53	93.92	86.29	63.72
Interaction-Hard	24.78	25.76	24.93	24.02	24.12	0.951	0.944	0.942	0.946	0.953	79.06	82.30	82.54	78.82	58.43
Overall	25.49	25.42	25.21	24.99	24.26	0.960	0.943	0.946	0.957	0.962	68.89	92.55	88.45	68.25	47.77

Table 3: **Benchmark results on novel identity task.** State-of-the-art methods’ performance on novel identity rendering in each benchmark split. We abbreviate IBRNet [69] as ‘IBR’, PixelNeRF [73] as ‘PN’, VisionNeRF [35] as ‘VN’, NeuralHumanPerformer [32] as ‘NHP’ and KeypointNeRF [71] as ‘KN’.



Figure 6: **Visualization of novel pose animation result samples.** From top to bottom, we illustrate the reposing results generated by (a-e): NeuralVolumes [39], A-NeRF [62], NeuralBody [53] AnimatableNeRF [52], and HumanNeRF [71].

prior and appearance references from observation could help boost the generalization ability on data with large variations of poses and appearances. We illustrate the qualitative results in Fig. S11. We provide additional results and analysis in Sec. C.2.3 in supplementary.

4.4. Cross-dataset Comparison

Apart from the benchmark experiments conducted on different data splits of DNA-Rendering dataset, we also evaluate the *data generalizability* provided by our dataset and the other competitive ones, *i.e.*, GeneBody[11], ZJU-MoCap [53] and HuMMan [5].

Setting and Implementations. To eliminate the scale and annotation differences across all datasets, we train three general scene generalizable rendering methods [73, 69, 35] on these datasets with the same pixel batch per-iteration and stop training with the same 200K global iterations. For each method, we train each individual model on each dataset mentioned above, with a fixed image resolution 512×512 and four balanced views. To thoroughly evaluate the datasets’ generalizability, we cross-verify the rendering images of *novel identities* on each dataset.

Results. The experimental results are presented in Fig. 7 in terms of the average PSNR of all three methods. From this colored error map, we conclude that training on DNA-Rendering dataset is beneficial for generalizing to the other datasets. Generally, due to the existence of domain gaps, a model would perform better in the situation of an in-domain setting, where the training set and test set follow the same distribution, see diagonal elements in Fig. 7. The off-diagonal numbers report the cross-domain performances of models trained on one dataset and tested directly on other datasets’ test sets. We observe an interesting phenomenon that, compared to datasets with limited data diversity and high data bias (like ZJU-MoCap [53] and HuMMan [5]), the proposed dataset enables generalization methods to achieve more plausible results even with large domain gaps. Moreover, opposite to DNA-Rendering, HuMMan [5] generalize poorly on other datasets even on cases with simple motions and appearances in ZJU-MoCap [53], despite the fact that both HuMMan [5] and our DNA-Rendering have large data volume. From a data engineering perspective, this demonstrates the construction of the proposed dataset benefits the community not merely with the amount of data, more im-




Figure 7: **Results of cross-dataset experiments.** We visualize the ‘affinity’ matrix of cross-dataset evaluation results.

portantly, the significant improvement in data completeness and richness. Due to space limit, we provide the detailed setup and additional results in Sec. D of the supplementary. It is worth mentioning that, we also unfold the generalization performance across testing cameras and reveal the impact of color consistency for multi-camera datasets in Sec. D.4.

5. Conclusion

In this work, we present DNA-Rendering, a large-scale, high-fidelity repository of neural actors with neural implicit representations. We contribute a multiview human body capture dataset that covers many diverse factors like ethnicity, age, body shape, clothing, motion, and interactive objects. According to these attributes, we construct several data splits under multiple difficulty levels. Besides, we provide faithful annotations attached to the raw data. Finally, we present benchmarks to evaluate the state-of-the-art research works on the DNA-Rendering dataset with in-depth discussions, and compare our dataset with the other ones via cross-dataset experiments on generalization capability provided to human avatar models. We hope our DNA-Rendering project could help the research community to boost the development of human-centric rendering and related domains with new reflections, new challenges, and new opportunities.

References

- [1] Oliver Batchelor. Multi-camera calibration using one or more calibration patterns. <https://github.com/oliver-batchelor/multical>, 2021. 5
- [2] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *CVPR*, 2014. 3
- [3] Federica Bogo, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Dynamic FAUST: Registering human bodies in motion. In *CVPR*, 2017. 3
- [4] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015. 2
- [5] Zhongang Cai, Daxuan Ren, Ailing Zeng, Zhengyu Lin, Tao Yu, Wenjia Wang, Xiangyu Fan, Yang Gao, Yifan Yu, Liang Pan, et al. Humman: Multi-modal 4d human dataset for versatile sensing and modeling. In *ECCV*, 2022. 2, 3, 6, 10, 18, 26, 27, 28, 29
- [6] Zhongang Cai, Mingyuan Zhang, Jiawei Ren, Chen Wei, Daxuan Ren, Jiatong Li, Zhengyu Lin, Haiyu Zhao, Shuai Yi, Lei Yang, et al. Playing for 3d human recovery. *arXiv preprint*, 2021. 6
- [7] Chen Chen, Roozbeh Jafari, and Nasser Kehtarnavaz. Utdmhad: A multimodal dataset for human action recognition utilizing a depth camera and a wearable inertial sensor. In *ICIP*, 2015. 3
- [8] Xianjie Chen, Roozbeh Mottaghi, Xiaobai Liu, Sanja Fidler, Raquel Urtasun, and Alan Yuille. Detect what you can: Detecting and representing objects using holistic models and body parts. In *CVPR*, 2014. 3
- [9] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, 2021. 4
- [10] Yue Chen, Xuan Wang, Xingyu Chen, Qi Zhang, Xiaoyu Li, Yu Guo, Jue Wang, and Fei Wang. Uv volumes for real-time rendering of editable free-view human performance. In *CVPR*, 2023. 26
- [11] Wei Cheng, Su Xu, Jingtian Piao, Chen Qian, Wayne Wu, Kwan-Yee Lin, and Hongsheng Li. Generalizable neural performer: Learning robust radiance fields for human novel view synthesis. *arXiv preprint*, 2022. 2, 3, 4, 6, 10, 18, 19, 25, 26, 27, 28, 29
- [12] Sungjoon Choi, Qian-Yi Zhou, and Vladlen Koltun. Robust reconstruction of indoor scenes. In *CVPR*, pages 5556–5565, 2015. 5
- [13] EasyMocap Contributors. Easymocap - make human motion capture easier. <https://github.com/zju3dv/EasyMocap>, 2021. 18, 19, 21
- [14] XRPrimer Contributors. Openxrlab foundational library for xr-related algorithms. <https://github.com/openxrlab/xrprimer>. 5
- [15] XRMoCap Contributors. Openxrlab multi-view motion capture toolbox and benchmark. <https://github.com/openxrlab/xrmocap>, 2022. 6
- [16] 3D People Cooperation. 3d people dataset. <https://3dpeople.com/>. 3
- [17] RenderPeople Cooperation. Renderpeople dataset. <http://https://renderpeople.com/>, 2017. 3
- [18] T. Derose, M. Kass, and T. Truong. Subdivision surfaces in character animation. In *SIGGRAPH*, 1998. 25
- [19] Ke Gong, Xiaodan Liang, Yicheng Li, Yimin Chen, Ming Yang, and Liang Lin. Instance-level human parsing via part grouping network. In *ECCV*, 2018. 3
- [20] Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. The lumigraph. In *SIGGRAPH*, 1996. 19

- [21] Peter Hedman and Johannes Kopf. Instant 3d photography. *TOG*, 2018. 20
- [22] Nikolas Hesse, Sergi Pujades, Javier Romero, Michael J Black, Christoph Bodensteiner, Michael Arens, Ulrich G Hofmann, Uta Tacke, Mijna Hadders-Algra, Raphael Weinberger, et al. Learning an infant body model from rgb-d data for accurate full body motion analysis. In *MICCAI*, 2018. 6, 18
- [23] Mustafa Işık, Martin Rünz, Markos Georgopoulos, Taras Khakhulin, Jonathan Starck, Lourdes Agapito, and Matthias Nießner. Humanrf: High-fidelity neural radiance fields for humans in motion. *TOG*, 2023. 2, 4
- [24] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *TPAMI*, 2013. 2, 3, 6, 18, 26
- [25] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. In *ECCV*, 2022. 4
- [26] Sheng Jin, Lumin Xu, Jin Xu, Can Wang, Wentao Liu, Chen Qian, Wanli Ouyang, and Ping Luo. Whole-body human pose estimation in the wild. In *ECCV*, 2020. 6
- [27] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 7, 23
- [28] Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *ICCV*, 2015. 2, 3, 26
- [29] Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, Sean Banerjee, Timothy Scott Godisart, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social interaction capture. *TPAMI*, 2017. 2
- [30] Sai Charan Mahadevan Karunanidhi Durai Kumar, Huang Geng. Sfu motion capture database. <https://mocap.cs.sfu.ca/>. 3
- [31] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint*, 2013. 21
- [32] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. *NeurIPS*, 2021. 4, 7, 9, 21, 25, 26
- [33] CMU Graphics Lab. Cmu graphics lab motion capture database. <http://mocap.cs.cmu.edu/>. 3
- [34] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *ICCV*, 2021. 2, 3, 26
- [35] Kai-En Lin, Lin Yen-Chen, Wei-Sheng Lai, Tsung-Yi Lin, Yi-Chang Shih, and Ravi Ramamoorthi. Vision transformer for nerf-based view synthesis from a single input image. In *WACV*, 2023. 4, 7, 9, 10, 21, 25, 27
- [36] Shanchuan Lin, Andrey Ryabtsev, Soumyadip Sengupta, Brian L Curless, Steven M Seitz, and Ira Kemelmacher-Shlizerman. Real-time high-resolution background matting. In *CVPR*, 2021. 6, 19
- [37] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 2, 3
- [38] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *CVPR*, 2016. 14
- [39] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: learning dynamic renderable volumes from images. *TOG*, 2019. 7, 8, 10, 20, 21, 22, 23, 24
- [40] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *TOG*, 2021. 5
- [41] Matthew Loper, Naureen Mahmood, and Michael J Black. Mosh: Motion and shape capture from sparse markers. *TOG*, 2014. 6, 18
- [42] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *TOG*, 2015. 18
- [43] MR Luo, G Cui, and B Rigg. The development of the cie 2000 colour-difference formula: Ciede2000. *Color Research and Application*, 2001. 17
- [44] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *ICCV*, 2019. 3
- [45] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *3DV*, 2017. 3
- [46] Marko Mihajlovic, Aayush Bansal, Michael Zollhoefer, Siyu Tang, and Shunsuke Saito. Keypointnerf: Generalizing image-based volumetric avatars using relative spatial encoding of keypoints. In *ECCV*, pages 179–197, 2022. 4, 7, 9, 21, 25, 26
- [47] 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. 4, 7, 20
- [48] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *TOG*, 2022. 7, 20, 23, 30
- [49] Ferda Ofli, Rizwan Chaudhry, Gregorij Kurillo, René Vidal, and Ruzena Bajcsy. Berkeley mhad: A comprehensive multimodal human action database. In *WACV*, 2013. 3
- [50] Priyanka Patel, Chun-Hao P Huang, Joachim Tesch, David T Hoffmann, Shashank Tripathi, and Michael J Black. Agora: Avatars in geography optimized for regression analysis. In *CVPR*, 2021. 6, 18
- [51] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 6, 18, 19
- [52] 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. 4, 7, 8, 9, 10, 21, 22, 23, 24

- [53] 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. 2, 3, 4, 5, 7, 8, 9, 10, 18, 20, 21, 22, 23, 24, 26, 27, 28, 29
- [54] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 4
- [55] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. "grabcut" interactive foreground extraction using iterated graph cuts. *TOG*, 2004. 6, 29
- [56] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 4
- [57] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, 2020. 4
- [58] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. In *CVPR*, 2021. 4
- [59] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+d: A large scale dataset for 3d human activity analysis. In *CVPR*, 2016. 3
- [60] Ruizhi Shao, Zerong Zheng, Hongwen Zhang, Jingxiang Sun, and Yebin Liu. Diffustereo: High quality human reconstruction via diffusion-based stereo using sparse cameras. In *ECCV*, 2022. 2, 3
- [61] Qing Shuai, Chen Geng, Qi Fang, Sida Peng, Wenhao Shen, Xiaowei Zhou, and Hujun Bao. Novel view synthesis of human interactions from sparse multi-view videos. In *SIGGRAPH*, 2022. 6, 18
- [62] Shih-Yang Su, Frank Yu, Michael Zollhöfer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. *NeurIPS*, 2021. 4, 7, 8, 9, 10, 20, 21, 22, 23, 24
- [63] Zhaoqi Su, Tao Yu, Yangang Wang, and Yebin Liu. Deepcloth: Neural garment representation for shape and style editing. *TPAMI*, 2023. 3
- [64] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, 2019. 6
- [65] Shuhei Tsuchida, Satoru Fukayama, Masahiro Hamasaki, and Masataka Goto. Aist dance video database: Multi-genre, multi-dancer, and multi-camera database for dance information processing. In *ISMIR*, 2019. 2, 3, 26
- [66] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 21
- [67] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, Pablo Speciale, and Marc Pollefeys. Patchmatchnet: Learned multi-view patchmatch stereo. In *CVPR*, 2021. 5
- [68] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. 7, 20, 23
- [69] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snaveley, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021. 4, 7, 9, 10, 21, 25, 27, 28
- [70] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 2004. 7
- [71] 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. 4, 7, 8, 9, 10, 21, 22, 23, 24
- [72] Cheng-hsin Wu, Ningyuan Zheng, Scott Ardisson, Rohan Bali, Danielle Belko, Eric Brockmeyer, Lucas Evans, Timothy Godisart, Hyowon Ha, Alexander Hypes, et al. Multiface: A dataset for neural face rendering. *arXiv preprint*, 2022. 5
- [73] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *CVPR*, 2021. 4, 7, 9, 10, 21, 25, 27
- [74] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4d: Real-time human volumetric capture from very sparse consumer rgbd sensors. In *CVPR*, 2021. 3
- [75] Zhengming Yu, Wei Cheng, Xian Liu, Wayne Wu, and Kwan-Yee Lin. Monohuman: Animatable human neural field from monocular video. In *CVPR*, 2023. 4
- [76] Zhixuan Yu, Jae Shin Yoon, In Kyu Lee, Prashanth Venkatesh, Jaesik Park, Jihun Yu, and Hyun Soo Park. Humbi: A large multiview dataset of human body expressions. In *CVPR*, 2020. 2, 3
- [77] Chao Zhang, Sergi Pujades, Michael J. Black, and Gerard Pons-Moll. Detailed, accurate, human shape estimation from clothed 3d scan sequences. In *CVPR*, 2017. 3
- [78] 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. 21, 23
- [79] Fuqiang Zhao, Wei Yang, Jiakai Zhang, Pei Lin, Yingliang Zhang, Jingyi Yu, and Lan Xu. Humannerf: Efficiently generated human radiance field from sparse inputs. In *CVPR*, 2022. 4
- [80] Zerong Zheng, Han Huang, Tao Yu, Hongwen Zhang, Yandong Guo, and Yebin Liu. Structured local radiance fields for human avatar modeling. In *CVPR*, 2022. 3
- [81] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *TPAMI*, 2021. 6, 18, 19
- [82] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3d human reconstruction from a single image. In *CVPR*, 2019. 3
- [83] Taotao Zhou, Kai He, Di Wu, Teng Xu, Qixuan Zhang, Kuixiang Shao, Wenzheng Chen, Lan Xu, and Jingyi Yu. Relightable neural human assets from multi-view gradient illuminations. In *CVPR*, 2023. 4

Appendix

In the appendix, we provide detailed information about the proposed DNA-Rendering dataset and the attached benchmarks. We first provide dataset statistics, hardware design, and data collection protocol in Sec. A. Then, we discuss the additional information about the annotations, as well as a comparison to other publicly released toolchains in Sec. B. Moreover, we conduct compact discussions on the benchmarks by introducing more detailed settings, additional results, and unfolded comparisons of benchmark methods’ conceptual differences in Sec. C. We provide an in-depth discussion on competing datasets and highlight our comparative contributions to society in Sec. D. Finally, we discuss our future work in Sec. E.

A. Dataset Details

A.1. Dataset Statistics

DNA-Rendering has a wide distribution over ethnicity, clothing, actions, and human-object-interaction scenarios. In this section, we present the detailed data distribution in key data aspects, namely ethnicity, age, shape, actions, clothing, and interactive objects.

Ethnicity, Age and Shape. We invite 500 actors with a uniform distribution of gender and a ratio of 4 : 3 : 2 : 1 for Asian, Caucasian, Black, and Hispanic individuals, respectively. The quota has a wide coverage of age and body shape. We visualize the distribution of actors’ age, height, and weight in Fig. S1.

Human Actions. DNA-Rendering covers both normal actions and professional actions. We maintain a library of 269 human action definitions, including daily-life activities, simple exercises, and social communication. All normal performers are asked to select 9 actions from the action library and perform the picked actions in a free-style manner. There are 153 professional actors among the total 500 performers. These professional actors are asked to dress in their special costumes and perform 6 unique professional actions with skills, including special costume performances, artistic movements, sports activities, *etc.*. Note that differ-

ent from the intuitive visualization in Fig. 1, we visualize fine-grain categories of professional and normal action in Fig. S2a. These labels are classified in terms of a standard human activity subcategory definition. The sunburst chart of distribution is visualized in the middle, and samples of specific categories of labels are visualized in the outer word cloud.

Clothing and Interactive Objects. We create a clothing repository with 527 items, which covers all 50 clothing types in DeepFashion [38] while with a random distribution of color, material, texture, and looseness for each clothing type. We ask each performer to wear three sets of outfits, where one comes from the performer’s self-prepared outfit (for both special and normal actors), and the other two are randomly coordinated from our clothing repository. The distribution of cloth statistical distribution on all action sequences and samples of cloth labels is illustrated in Fig. S2b.

A.2. Meta Attributes

We have designed an attribute system for each dimension of the collected data, including basic information about the actors, clothing, and action information for each action sequence. Fig. S3 shows an example of meta attribute information of an action sequence for a professional actor. In terms of actor information, we record the actor’s name, gender, ethnicity, age, height, and weight. For clothing information, we describe the upper and lower clothing, and shoe information. These descriptions include information on color, type, and other significant visual features. For action information, we describe the overall content of the action and, if there are any interactive objects, we also describe the type and other significant visual features of those objects.

A.3. Hardware Construction

The main structure of DNA-Rendering’s capture system is a dome with a radius of three meters. The camera array built upon the dome consists of multiple types of cameras – ultra-high-resolution 12MP cameras, 5MP industrial

https://en.wikipedia.org/wiki/Wikipedia:Contents/Human_activities






Figure S1: **The distribution of actors’ attributes.** We record the age, height, and weight of our invited actors. The statistical results reflect the wide range of the actors’ personal attributes.



(a) Action distribution and labels



(b) Cloth distribution and labels

Figure S2: **Illustration of the action and clothes label distribution.** (a). The distribution of action categories and sub-categories is visualized by a sunburst chart in the middle which is surrounded by the word cloud of normal and professional action labels. (b). The distribution of clothes categories and labels are visualized in the same form with (a).

cameras, and Azure Kinect cameras. The lighting system provides natural lighting conditions. All cameras are triggered and synchronized by hardware, and synchronized multi-view data are transferred and recorded through our data streaming system.



Figure S3: **An example of our meta attribute system.** We record the actors' basic information, costumes, and actions.

Camera System. DNA-Rendering has 68 cameras, including 12 ultra-high resolution cameras with 12MP resolution (short for 4096×3000 resolution), 48 industrial cameras at 5MP resolution (*i.e.*, 2448×2048 resolution), and 8 RGB-D Kinect cameras with depth resolution of 576×640 . Specifically, the 5MP cameras are mounted on three cycles on the dome skeleton with 1, 2, and 3 meters in height, each circle in height has 16 balanced 5MP cameras with 22.5° angle interval. 12MP cameras are placed uniformly on another two intermediate height level circles, 1.5 and 2.5 meters height

respectively. 12MP cameras are installed with a 60° angle interval and interlaced with 5MP cameras. The Kinect cameras are mounted close to the middle level of 5MP cameras, providing the best RGB texture references for depth maps. Such construction of the camera array achieves dense coverage of the human body at multiple heights and angles. 5MP cameras and 12MP cameras are equipped with lenses of 8 mm and 6 mm respectively to achieve the best trade-off between full body proportion-in-view and size of capture volume. Note that, we use data captured from 12MP and 5MP cameras to construct our rendering dataset. The data captured from depth cameras comprise the auxiliary data which provide coarse geometry of human. Noted that we abandon the Kinect RGB cameras during the entire process, due to the bad color consistency.

Lighting System. Our lighting system consists of 16 flat light sources with a color temperature of $5600K \pm 300K$ and an illuminance of 4500 Lux/m. The lighting scale of each light source is 700×500 mm. There are eight flat lights on the ground installed with a 45° tilt towards actors in the middle to provide the best lighting on actors. There are extra eight flat lights hung on the roof to strengthen the lighting of upper body parts, especially for human heads. These uniformly distributed flat lights irradiate the whole scene with strong, natural, and balanced illumination.

Data Streaming. To collect, transfer and store the multi-view camera data, we construct a data streaming system that consists of two pieces of equipment for data synchronization – a 10 Giga-byte network, and a high data throughput workstation. The camera system is synchronized by Kinect's trigger signal. First, eight Kinects are configured in a daisy chain and the out-trigger signal is converted to the TTL signal, and the other 60 cameras are triggered by synchronization equipment. The 5MP cameras are connected

to six workstations via USB-3.0 ports and four-channel USB cards with PCI-E interfaces. The 12MP cameras are connected to the other three workstations via 10 GigE networks, capture cards, and PCI-E interfaces. To reduce active light interference of Kinect depth cameras, we adopt a $160 \mu\text{s}$ time delay for each slave device on the chain. The maximum synchronization error of Kinect is 1.12 ms in theory. The maximum synchronization error among all industrial cameras is less than 2 ms, we measure this error by utilizing the image of high-speed flashing LED timer arrays and computing the displayed time differences.

A.4. Data Collection Protocol

We discuss the detailed data collection protocol from five aspects, *i.e.*, data content, system check, core data collection, auxiliary data collection, and post-processing.

Data Content. During everyday data collection, we gather a comprehensive set of data sources, including action data, background data, actors' A-pose data for each outfit, extrinsic calibration data, and the record of performance attributes. We collect intrinsic data and color calibration data only when we apply any modification to the system.

System Check. We conduct a daily system check before formal data collection. The process focuses on the verification of camera parameters and synchronization. Concretely, we will check 1) if the camera parameters remain the same with recorded optimal values (*e.g.*, white balance, gamma, focal length, the valid field of view (FOV), *etc.*). Checking these factors ensures capturing under excellent image quality and valid capture volume. 2) We monitor the network's condition and check the synchronization via a system probe using high-speed flashing LEDs. 3) Finally, we collect extrinsic camera calibration data via a standard data collection process that records the checkerboard rotating as described in Sec. 3.3 in the main paper.

Core Data Collection. We invite 4-6 actors per day to perform actions in our studio in different appointed time slots. Once the actors arrive, we will briefly introduce the collection procedure and ask them to sign the authorization agreements first. If the actors agree, they are asked to prepare their outfits, makeup, and actions. Meanwhile, we record the basic information for each actor, including the height, weight, age, ethnicity, and other appearance attributes like the type, color, and material of his/her self-prepared outfit. After the preparation, we ask each actor to perform several actions in his/her self-prepared outfit. Specifically, a normal actor will pick at least three actions from our pre-defined daily activities and perform them in a free-style manner. A professional actor will wear a special outfit and perform at least six unique sets of footage that fit with the professional skill or costume. Then, we ask each actor to change his/her outfits with another two sets that are randomly coordinated from our clothing library. For each new outfit, the actor will

perform another three different normal actions. To ensure the performed motion is rational and authentic enough, we will ask each performer to rehearse outside the studio before the formal shooting. After our staff confirms that the action is performed correctly, the actor will perform in the studio again for the formal data collection.

Auxiliary Data Collection. Aside from the core performance data collection, we also record auxiliary data, including the blank background data for the matting process, and A-pose data as a record for the canonical actor model before each round of new outfit recording. To record A-pose data, we require 1) the actor's hands tilt 45° downward the legs with clear distance; 2) the hands should slightly open without clenched fingers or put them together; 3) the farical expression should keep expressionless with the eyes open and looking straight ahead.

Post-processing. After the action is completed, the center workstation generates fast multi-view preview videos for all cameras, and we check whether the performance content or the filming on each camera view meets the requirement. Actors are asked to re-play the performance if the recorded data is invalid. After collecting all qualified data, we post-process the data in per-day shooting volume Image sequences in RAW format will be converted to the lossless BMP format, and then compressed into a video with a low constant rate factor with the x264 library. The processed data are then uploaded to the cloud server for subsequent dataset processing.

A.5. Limitations on Data Collection

Data Content. To achieve high-fidelity data collection, we set the lighting with invariant and uniform illumination and set the acquisition frame rate to 15 frames per second. We also constrain the field of view of the cameras, to ensure each one can capture the *full-body* movements of a single actor (including the interacted object if there is one) while maintaining the FOV as max as possible. This allows us to capture details such as facial makeup and clothing textures. In future work, we would update our hardware systems and upgrade our capture processes to accommodate different lighting conditions, multiple FOV ranges for multi-person scenes, high-speed capture of subtle movements, and multi-sensory (*e.g.*, auditory, and tactile data) collection.

Failure Cases. During the data collection process, various factors can lead to failure, such as large movements that exceed the field of view of the multi-camera system, or loss of frames due to large volume data transmission fluctuations. Following the standardized capture process, our operators will manually inspect the completeness and effectiveness of all camera data and actor movements after each capture cycle. If any issues are found, the hardware will be instantly checked and the data will be re-captured. Most failure cases are identified and promptly resolved at this stage.

B. Data Annotation Details

B.1. Camera Intrinsic Calibration

As this project targets capturing high-fidelity *whole-body* data, we adopt a long lens that enlarges the human proportion in the camera view. This setup requires a high-quality estimation of camera distortion since some subjects' body parts might appear on the image boundary region. Thus, we use a 3×3 Soduku data collection protocol for intrinsic calibration, as illustrated in Fig. S4. Specifically, to maximize the checkerboards' coverage across the whole image space, we separate the image into a 3×3 Soduku and capture the images of the checkerboards' movement in grids. For each grid, we rotate the checkerboard in pitch, row, and yaw direction to enlarge the angle of the boards.




Figure S4: **Intrinsic calibration.** To ensure better distortion coefficient estimation, (a) we separate the camera view into 3×3 Soduku and capture the images of a checkerboard (about $1/4$ size of the one used in estimating extrinsic parameters) in every Soduku grid. (b) For each grid, we rotate the checkerboard at pitch, row, and yaw angles. This calibration step forces the checker to appear in every corner of the camera view. (c) Zoom-in for small-size checkerboard for intrinsic calibration.

B.2. Camera Color Calibration

To ensure color consistency across multiple cameras, we inject a color calibration process into our data collection. A standard color board could be used as the criterion for color calibration, and the fixed lighting condition in the dome could be treated as a standard condition during calibration. Specifically, the calibration lies in two aspects: 1) Hardware parameter adjustment. We make a rough adjustment on the hardware parameters to make the white balance and color balance of each camera as consistent as possible by human eyes; 2) Fine adjustment. Under a standard light source, we make the standard color board face straightforwardly to the camera to be calibrated at a constant distance, and a single image under this setting is collected; the corner detection algorithm is used to automatically identify the position of the color board in the image, and the color sampling is performed with the center radius $p = 10$ pixels of each color square. The average color value is taken as the color sampling value. We carry out the polynomial projection of the color sampling value to the standard value via

least squares. Note that, we calibrate in RGB form and take $n = 2$ to prevent overfitting. The overall procedure and illustrated results are presented in Fig. S5.

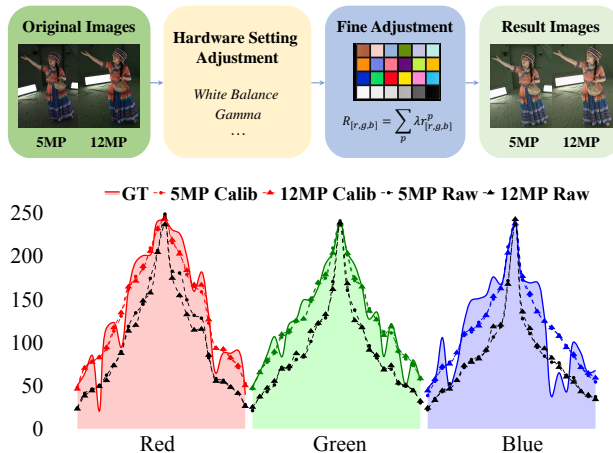


Figure S5: **Color calibration pipeline and calibrated color response.** The color-calibrated images are listed on the right of the flow chart (at the top of the figure). The color responses of two calibrated cameras (Camera 25 which is a 5MP camera, and Camera 51 which is a 12MP camera) compared with groundtruth color value of the color checkerboard are plotted below the flow chart, and smoothed spline curve is used. We show the 'Raw responses' after hardware setting adjustment for reference. With the help of the color correction process, the average RGB value consistency between these two cameras ΔE_{00} [43] is improved from 37.79 to 4.15.

B.3. Keypoints

We highlight that having a large number of camera views allows us to rectify the occasional failures of single-view 2D keypoint detection. For the more natural and stable 3D keypoints, we adopt the following optimization and post-processing strategies: 1) *Keypoint selection.* We dynamically select views for each keypoint in the data sequence, which with the most confident score to the keypoint while ensuring the keypoint can be triangulated. 2) *Bone length constraint.* The bone length is constrained with a fixed length. We use the median bone length after initial triangulation as the target in the optimization. Only the lengths of the main limbs are considered in this step. 3) *Outlier removal.* As a post-processing pipeline, filter modules are designed based on human priors, including a 3D bounding box filter, a movement filter, and a relative position filter. 3D keypoints outliers, which are too far from the body trunk, move too fast between frames, or lead to an inconsistent relative position between frames are removed. An interpolation is applied to recover the missing keypoints. Such a post-processing scheme can assure reliable and consistent face and hand keypoints, even with large-scale occlusions. As shown in Fig. S6, these optimization and post-




Figure S6: **Evaluation of keypoint quality from every camera view.** We compute the mean reprojection error of 3D keypoints with 2D detection results. Optimization effectively reduces the error to below 30 pixels. Note that camera IDs 0-47 are 5MP cameras, and camera IDs 48-59 are 12MP cameras, high-lighted with red x-ticks.

processing strategies effectively reduce reprojection error compared to triangulation with all available 2D keypoints.

B.4. Parametric Model

In our automatic parametric model annotation pipeline, body shape $\beta \in \mathbb{R}^{n \times 10}$ (or $\beta \in \mathbb{R}^{n \times 11}$ for children [22, 50]) is first estimated based on the bone length calculated from 3D keypoints with the static and less challenging A-pose sequence. We use the estimated body shape parameters as initial values and optimize the full parametric model parameters including pose parameters (body pose, hand pose, and global orientation) $\theta \in \mathbb{R}^{n \times 156}$, and translation parameters $t \in \mathbb{R}^{n \times 3}$ (n is the number of frames) via a modified SMPLify-X for other sequences with dynamic poses. The main energy terms in the optimization are keypoint energy $E_{\mathcal{P}}$, full-body joint angle prior energy E_a , bone length energy $E_{\mathcal{B}}$, and body shape prior energy E_{β} [51, 42, 5]. The main modification of SMPLify-X in our annotation pipeline is the decoupling body shape optimization and pose optimization, which we empirically find to produce more stable results.

Concretely, we employ bone length energy $E_{\mathcal{B}}$ and body shape prior energy E_{β} to fine-tune body shape parameters for each sequence of a subject, with the same shape initialization from the A-pose static sequence. Body shape values are kept consistent throughout all the frames in a sequence.

$$E_{shape}(\theta, \beta, t) = \lambda_1 E_{\mathcal{B}} + \lambda_2 E_{\beta} \quad (S1)$$

We then leverage keypoint energy $E_{\mathcal{P}}$ and full-body joint angle prior energy E_a for pose optimization with body shape fixed.

$$E_{pose}(\theta, \beta, t) = \lambda_3 E_{\mathcal{P}} + \lambda_4 E_a \quad (S2)$$

As shown in Fig. S7, we evaluate the fitting error between 3D keypoints and corresponding regressed SMPLX joints. Body-only keypoints, hand-only keypoints, and all keypoints are evaluated separately. With our multi-view capture system and annotation pipeline, the MPJPE of

body-only keypoints is on par with the optical motion capture system in Human3.6M [24, 41], MPJPE of hand-only keypoints is 15.87mm.




Figure S7: **Evaluation of parametric model registration quality.** We evaluate body-only keypoints, hand-only keypoints, and all keypoints separately. The orange line indicates the median value, the box indicates the lower to the higher quartile, and the whiskers indicate the range of data.

B.5. Comparison to Other SMPLX Fitting Methods

Baselines. To analyze the effectiveness of the proposed SMPLX fitting pipeline, we evaluate the accuracy of SMPLX fitting and compare it with three publicly available pipelines, *i.e.*, the baseline **MultiviewSMPLifyX** [81, 51], **EasyMoCap** [13, 61] used in ZJU-MoCap [53, 61] dataset, and **BodyFitting** used in GeneBody [11] dataset. Specifically, **MultiviewSMPLify** [81] and **BodyFitting** [11] directly optimize the error of reprojected 3D SMPLX keypoints to 2D detections. Such a naive strategy is straightforward but lacks outlier robustness (might stuck in absolutely wrong detections or detection flip between left and right), and it is also computationally expensive. On the contrary, both **EasyMoCap** and the proposed pipeline adopt another strategy that separates the SMPLify process by a triangulation process. This strategy optimizes 3D keypoints from 2D detection and then fits SMPLX from directly on optimized 3D keypoints. As robust designs could be adapted during the triangulation process to eject outliers caused by flipping or occlusion, such a two-step strategy is faster and more robust to outliers. Whereas, one drawback is that the final SMPLX totally rely on the results of triangulation in the first stage by hand-crafted optimization and filtering.

Methods	2D Reprojection Error (pixel)				3D MPJPE (mm)				Run Time (s)
	Body	Hand	Face	Overall	Body	Hand	Face	Overall	
MultiviewSMPLifyX [81]	42.27	33.36	23.91	28.77	55.46	23.25	22.24	27.54	81.33
BodyFitting (GeneBody) [11]	45.68	33.43	34.17	35.67	42.37	32.19	30.67	32.89	29.50
EasyMoCap (ZJU-MoCap) [13]	32.71	33.75	32.72	33.64	36.04	25.37	38.10	33.96	0.69
Ours	29.63	31.41	19.08	24.08	30.20	15.87	16.46	17.52	3.23

Table S1: **Comparison among multi-view SMPLX fitting methods.** Cell color ■ ■ ■ indicates the best, second best, and third best performance, respectively. Runtime in seconds indicates the average time required for the fitting process for one multi-view frame.

Compared to EasyMoCap, incorporate a more sophisticated designed 2D keypoints postprocessing phase, where movement filtering and relative position filtering are used when the given 2D keypoints are not accurate.

Settings. For fairness, we use the same 2D keypoints consisting of human-inspected body labels and hands and faces auto-detection results. We also force all SMPLX models to have 10 facial expression coefficient and 45 hand PCA components. We run the SMPLX fitting on our benchmark test data with their default SMPLX setting, namely with default penalty energies, their coefficient, and other settings except for the aforementioned modifications. We quantitatively evaluate the MPJPE of 3D keypoints and the reprojected 2D error across all views.

Results. The quantitative results are listed in Tab. S1, we also separately evaluate the accuracy on body, hand, face, and whole body. 3D MPJPE is computed by regressed SMPLX 3D keypoints to human-inspected 3D keypoints. 2D reprojection error is compared by reprojected SMPLX 3D keypoint to input 2D keypoints. The runtime performances are also recorded, indicating the average fitting time usage (excluding other time namely, data IO, *etc.*) for one multi-view frame. Our proposed pipeline outperforms other fitting methods in all categories in terms of both 2D and 3D metrics, and has a more acceptable runtime requirement than EasyMocap [13]. Moreover, MultiviewSMPLify [81, 51] achieves the second-best performance, while its time consumption is exploded by an order of magnitude. In a nutshell, our pipeline ensures the best trade-off between performance and efficiency.

B.6. Matting and Segmentation Refine

As described in the main text, despite the state-of-the-art background matting method [36] achieving impressive matting performance in the majority of our data, there are still several corner cases that fail to extract the foreground correctly, *e.g.*, noisy backgrounds, broken bodies, and missing bodies and objects. We demonstrate these most common corner cases in Fig. S8. To further improve the matting quality, we adopt the traditional computer vision algorithm-GraphCut [20] to refine the predicted masks, and we find such a classical method plays a good fit to the CNN-based method which generates good results on these failure cases.

In order to quantify the improvement, we introduce a manual inspection process to grade the results generated by

CNN-based method [36] only and by a subsequent refinement procedure. Noted that due to the large scale of data, generating masks with manual labeling is impractical. More specifically, we ask three annotators to conduct such grading surveys on 500 random multiview sequences. We report the error rates in terms of the type of corner cases in the bar chart in Fig. S8. From the human grading probe, we can conclude that with our proposed hybrid strategy, the error rates in all category decrease by a large margin compared with [36] only, with the overall error rate reduced from 11% to 2%.

B.7. Quality Control of Auto Annotation Results

To ensure the quality of annotation data, we conduct manual quality checks on the auto-annotated results. Specifically, we perform a human-in-the-loop quality evaluation for the SMPLX and matting results generated by the annotation pipeline.

For SMPLX quality control, we overlay each SMPLX result on the original action data to create a multi-view video and manually verify the quality with the labeling task that requires our human annotator to grade the SMPLX quality. We subdivide the process into three stages. 1) Binary filtering. If the SMPLX human body completely overlaps with the human body in the image or is within the natural shape range of the human body throughout the entire video, it is considered as a qualified SMPLX annotation; otherwise, if there is severe misalignment or distortions on the main body, it is considered as an unqualified one. 2) Quality Grading. For qualified data cases, we further evaluate their subdivision quality, dividing them into five scores based on the unnaturalness of fingers or faces, the alignment of the head and shoulders with the image, *etc.* 3) Keypoints re-annotation. For unqualified cases, we ask the annotators to re-annotate the main skeleton in views with large errors by auto-annotators. The new annotation results are used to re-run the SMPLX results. We repeat the whole process until we achieve valid SMPLX models in all cases.

For Matting quality control, we manually evaluate the quality of the annotated video after matting by grading each video’s quality. Quality is divided into three levels: A-level, where the entire human body is fully displayed without occlusion, and any interactive objects are fully shown, with no excess areas; B-level, where a small part of the human body or object is missing, or there are a few extra cutouts,




Figure S8: **Examples and statistics on matting refinement.** In the upper image, we show three kinds of common challenging cases and the comparisons among color filtering only, background matting only, and our optimized solution (*Refined Matting*). From top to bottom cases, the problems (before optimization) are noisy backgrounds, broken body areas, and missing areas like body parts and incomplete objects; In the bottom figure, we show the error rates from the sampling survey of the three categories.

with the erroneous pixel area not exceeding one-third of the human body area; and C-level, where there are serious problems and the erroneous pixel area exceeds one-third of the effective human body area. We treat A-level and B-level as acceptable mask annotations, while C-level as failure annotations. Noted that the cases in training and testing data split in our benchmark were manually selected to ensure high-quality annotations. For the sake of rigor, we will release the mask rating as confidence for mask annotation.

with the erroneous pixel area not exceeding one-third of the human body area; and C-level, where there are serious problems and the erroneous pixel area exceeds one-third of the effective human body area. We treat A-level and B-level as acceptable mask annotations, while C-level as failure annotations. Noted that the cases in training and testing data split in our benchmark were manually selected to ensure high-quality annotations. For the sake of rigor, we will release the mask rating as confidence for mask annotation.

C. Benchmark Details

C.1. Methods Overview and Modifications

In this subsection, we review the state-of-the-art methods benchmarked in this paper, and describe the major modification we made to the default implementation for adapting to the proposed dataset.

C.1.1 Static Methods

For static methods, we target to anchor the performances of novel view rendering on static test frames, which could be used as the baseline reference for dynamic methods on certain frozen times.

Instant-NGP [48] as an alter of NeRF [47], which utilizes the multi-resolution hash embedding and smaller network to accelerate the training and evaluation cost without loss of quality. Given the original implementation of Instant-NGP is under the underlying assumption of a moving single camera input or cameras sharing the intrinsic parameter across all camera positions, we modify it to suit multi-camera data with different intrinsic parameters.

NeuS [68] is a hybrid representation that combines neural radiance field with neural SDF, which produces better 3D reconstruction ability than NeRF-based methods [21, 47, 53] on existing datasets, while the rendered images of NeuSare typically not as sharp as NeRF-based methods. When adapting NeuS [68] on the proposed dataset, no special modification is required.

C.1.2 Dynamic Methods

To construct the novel view synthesis and novel pose animation benchmark, we select the most recent state-of-the-art dynamic neural human rendering methods which can learn a neural body avatar from video sequences.

NeuralVolumes [39] formulates a category-agnostic dynamic scene by a canonical voxel-grid decoder, and models the per-frame deformation as a mixture of affine warps that are parameterized by an auto-encoder with image input. Due to this property, we feed the network with 4 balanced views of images from the training views. We also center and scale the camera system by 0.3 to fit the voxel-grid system. During testing on novel pose, novel pose images of the 4 view are input to the auto-encoder.

A-NeRF [62] learns a human NeRF by conditioning the field with coordinates in each bone’s local system. Note that its default setting only trains the network in the foreground, which usually leads to artifacts on the floor, we improve this by forcing pixel sampling on non-foreground space which helps to reduce the artifacts.

NeuralBody [53] conditions a dynamic NeRF by time codes as well as structural latent features by sparsely involving parametric model’s vertices in 3D. To run Neural-

Body [53], we transform our standard definition of the parametric model to EasyMoCap [13] style, and we train the network using 42 dense views. Note that during novel pose estimation, we fed the network with novel pose SMPLs and linearly extrapolate the time step.

AnimatableNeRF [52] introduces neural blend weights with 3D human skeletons to generate observation-canonical correspondences in dynamic human NeRF. We do the same transformation like NeuralBody [53].

HumanNeRF [71] learns a dynamic neural human model from monocular video. It decouples the motion field by a corrected skeleton movement and non-rigid motion. Different from its original setting, we train the model with dense views by stacking multiview video sequences. It is *important* to point out that HumanNeRF [71] models may collapse on certain data sequences producing meaningless images like the left bottom case in Fig. 6. Such a phenomenon is consistent even with multiple trials of random initialization. Considering to deliver a more straightforward metric meaning in the benchmark, the report numbers of HumanNeRF in Tab. S2 and Tab. 2 only include the valid models.

C.1.3 Generalizable Methods

For novel identity generalization, we evaluate three category-agnostic methods [73, 69, 35] and two methods with human structure priors [32, 46].

PixelNeRF [73] is one of the first generalizable NeRFs that generalize novel objects’ color and opacity by pixel-aligned feature-conditioned NeRF. We train it on our dataset with 4 selected views and fuse the multiview image feature with average pooling.

IBRNet [69] predicts the radiance color of novel objects by blending observed color from source views, and inference the opacity from multiview feature fusion.

VisionNeRF [35] upgrades PixelNeRF’s [73] image encoder with a global transformer [66] and fuse the multi-level features with 2D CNN features. Like PixelNeRF [73], we fuse the multiview feature with average pooling.

NeuralHumanPerformer [32] combines key components of PixelNeRF [73] and NeuralBody [53], and fuse them with multiview transformer and predict the radiance of human body. Noted there is a slight contradiction between the technical paper and the released implementation on the window size of the temporal transformer. We follow the open-sourced implementation and set the window size to 1 to avoid memory explosion which means the temporal transformer is a dummy module.

KeypointNeRF [46] use IBRNet [69] as the backbone, and tailors 3D keypoints as human prior into the framework. It conditions the radiance field with relative depth to every 3D keypoint in each source camera coordinate. We train the network with 24 SMPL main skeleton keypoints. Noted

that different from other generalizable methods which allow arbitrary resolution rendering, KeypointNeRF [46] is suited to render square-sized images with 2^n width and height. We render out a minimum squared image that can cover the desired resolution then crop out the valid part.

C.2. Benchmark Details

As a supplement to the benchmark part in the main paper, we describe the detailed benchmark settings and additional analysis of results.

C.2.1 Novel View Synthesis

Detailed Settings. As we described the task in Sec. 4.2 and reviewed the methods in Sec. C.1, we evaluate the dynamic methods’ novel view synthesis ability on the benchmark test set, which consists of 13 splits with 39 performance sequences. During training, we train models on each sequence separately, with the 42 multiview (training views) image sequences of the first 80% frames. Evaluations are performed on the same seen human poses but with every 45 frame skip and only calculated on 18 unseen camera poses. For static methods, we train separate models on multi-view images of each single frame. To evaluate the high-fidelity rendering of these benchmark methods, we train and test the models on *half of the origin resolution*, namely 1024×1224 and 1500×2048 for 5MP and 12MP images respectively.

Detailed Results. As a supplement to the result analysis in the main text, we present more detailed results and analysis in this subsection. Detailed quantitative results across our testing splits are listed in Tab. S2, which correspond to the bubble charts in Fig. 4. We also illustrate the qualitative results in Fig. S9. From the ranking in Tab. S2, we can observe that A-NeRF [62], NeuralBody [53] and HumanNeRF [71] achieve the best numbers in most of the test splits in terms of PSNR, SSIM and LPIPS [78]. As the module designs of these methods might play vital roles in such distinct results, we further analyze the phenomenon by unfolding their conceptual differences as follows: *NeuralVolumes* [39] adopts a VAE [31]-style neural rendering framework that encodes and decodes both the affinity transformation field and rendering volume from sparse view references. Such a paradigm absorbs the strength of VAE that compresses input multi-view features into one compact latent representation space, which follows the Gaussian assumption. Thus it could generalize well on novel views (achieving top-three performance in PSNR) with acceptable rendering quality. Whereas, such a framework also inherits the smooth problem of VAE that leads to not sharp enough qualitative results. *A-NeRF* [62] is a conditioned NeRF that utilizes local joint ordinate information of query points. It samples a small box center at a random point in the foreground and adds a proportion of background points to reg-




Figure S9: Novel view synthesis results on each data split. From top to bottom, we illustrate the rendering results generated by (a-e). NeuralVolumes [39], A-NeRF [62], NeuralBody [53], AnimatableNeRF [52], HumanNeRF [71]. Please zoom in for better visualization.

ularize the empty space, using only mean square error loss (MSE). This strategy enforces the network to encode the dynamic NeRF with local bone coordinates, which is efficient in the human foreground region. The overall novel view synthesis ability of A-NeRF [62] is appealing, especially in PSNR. However, due to the sparsity characteristic of skeleton representation, A-NeRF tends to generate dilated artifacts (see *Motion-Medium* and *Motion-Hard* in Fig. S9), more obvious with novel pose in off-body parts of *Interaction-Hard* case in Fig. S10. *NeuralBody* [53] and *AnimatableNeRF* [52] first compute a 3D bounding box from SMPL, then train/infer on the reprojected 3D box region and fill the outer region with the background color. Thus, their SSIM scores are typically greater than other methods that infer the whole image. Whereas, the bounding box only helps the network consider the main body and ignore the object, which leads to both methods can not reconstruct large interacting objects (as illustrated in the sword part of *Interaction-Medium* case in Fig. S9) *HumanNeRF* [71] combines the strength of previous methods’ design, it samples squared boxes on reprojected 3D bounding box, and trained model with a perceptual [27] loss and MSE loss. This results in best LPIPS performance, as well as the best visual performance with sharp texture in qualitative results. However, as *HumanNeRF* [71] designs a human motion prior that is Gaussian distribution along body parts or bones. This prior may lead to training failure on loose clothing and interactive objects, as illustrated in the *Deformation-Hard* and *Interaction-Medium* cases in Fig. S10.

C.2.2 Novel Pose Animation

Detailed Settings We conduct novel pose experiments on dynamic methods on the same models in novel view synthesis. Specifically, by training on the first 80% frames of each case, we test the animatable model with input of the

pose sequences extracted from the last 20% frames with a 15-frame skip. Depending on the pose-condition scheme of different methods, the test input can be divided into two categories, the SMPL parameters and the image features. We use the same testing view and rendering resolution as the novel view synthesis experiment.

Detailed Results. We present the detailed novel pose animation results on 13 testing splits in Tab. 2 in the main paper, and show the qualitative samples of results in Fig. S10 in this subsection. Similar to novel view synthesis task, A-NeRF [62] achieves the best PSNR performance, *NeuralBody* [53] has the best SSIM score, and *HumanNeRF* [71] gets the best LPIPS [78]. Differently, *NeuralVolumes* [39]’s performance decrease by a large margin, especially in *Motion* splits. One of the underlying reasons is that the affinity field learning in the *NeuralVolumes* [39] only relies on the latent code that is learned from multiview images and regularized by KL-divergence. Such a methodology is tied in a global warping manner, which might be relatively less affected by factors like global deformation in distance within a short movement change, but is vulnerable to unseen local motion (e.g., wrong head pose of *Interaction-Medium* case in Fig. S10, and strained border of actor’s shirt in *Texture-Easy* case in Fig. S10). Moreover, the design cannot preserve global scale in unseen poses, due to the wrong prediction of global affine transform sometimes (e.g., the zoom scale of the actor in *Texture-Easy* case in Fig. S10)). The methods [62, 53, 52, 71] with the explicit human pose information as input can typically generate reasonable animation results in terms of the local first motion, as shown in Fig. S10. Whereas, we draw the other major conclusion that current methods fail to model *Deformation* and *Interaction* properly. The typical examples shown in Fig. S10 are the loose cloth case of *Deformation-Hard*, and the interactive objects in *Interaction-Medium* and *Interaction-Hard* cases. How to properly model non-rig or out-of-body motions while preserving the advantages from explicit body

Splits	PSNR [↑]								SSIM [↑]								LPIPS* [↓]							
	NGP	NS	NV	AN	NB	AnN	HN		NGP	NS	NV	AN	NB	AnN	HN		NGP	NS	NV	AN	NB	AnN	HN	
Motion-Simple	30.97	27.49	27.85	29.15	27.84	25.89	25.49	0.979	0.973	0.966	0.974	0.978	0.974	0.955	31.52	44.18	57.74	52.75	53.32	56.10	62.08			
Motion-Medium	31.40	30.04	28.16	29.07	27.47	24.93	24.80	0.980	0.980	0.970	0.975	0.981	0.971	0.966	25.12	31.33	50.03	45.28	48.12	56.43	33.66			
Motion-Hard	29.05	28.49	26.10	27.55	25.16	24.54	22.93	0.972	0.976	0.959	0.967	0.976	0.976	0.964	41.35	40.13	77.54	69.75	71.25	63.35	53.42			
Deformation-Simple	31.63	28.01	28.09	29.63	28.18	27.42	28.30	0.981	0.972	0.968	0.975	0.976	0.975	0.974	29.02	42.62	48.17	41.68	48.18	45.44	23.70			
Deformation-Medium	30.01	29.65	29.77	30.52	29.22	26.29	26.60	0.972	0.975	0.971	0.974	0.979	0.972	0.963	41.14	37.18	39.36	43.51	46.95	52.69	29.12			
Deformation-Hard	29.79	30.80	27.19	28.11	24.77	21.93	21.48	0.967	0.973	0.954	0.957	0.969	0.958	0.934	46.09	44.83	70.04	70.16	81.14	84.59	83.36			
Texture-Simple	30.53	31.39	27.85	30.45	29.13	25.36	27.39	0.978	0.988	0.974	0.984	0.988	0.979	0.979	36.02	23.53	58.78	43.09	41.53	53.69	24.72			
Texture-Medium	30.85	31.33	28.50	30.53	29.62	22.46	27.40	0.978	0.982	0.968	0.977	0.984	0.959	0.971	29.32	27.04	47.33	37.99	41.46	75.08	25.01			
Texture-Hard	29.16	28.23	26.73	27.36	25.69	19.98	24.78	0.966	0.956	0.942	0.947	0.963	0.946	0.950	36.48	58.55	79.68	79.17	77.36	94.60	34.66			
Interaction-No	31.31	31.90	29.05	28.71	27.77	23.82	26.77	0.978	0.985	0.972	0.975	0.984	0.971	0.971	34.00	23.65	50.36	56.07	49.79	67.40	30.00			
Interaction-Simple	31.55	32.13	29.09	30.12	29.56	25.93	28.59	0.982	0.987	0.976	0.981	0.987	0.977	0.978	27.58	21.79	45.55	46.48	41.69	57.18	22.00			
Interaction-Medium	28.82	27.15	25.59	25.72	25.65	22.02	23.51	0.967	0.968	0.955	0.956	0.975	0.997	0.953	43.33	52.52	73.05	85.50	68.03	92.90	54.39			
Interaction-Hard	30.03	29.29	28.09	28.25	25.00	22.92	23.87	0.972	0.977	0.962	0.964	0.972	0.961	0.951	43.64	39.71	57.97	63.42	71.00	81.49	57.38			
Overall	30.39	29.68	27.85	28.86	27.31	24.11	25.53	0.975	0.976	0.964	0.970	0.978	0.970	0.962	35.74	37.47	58.12	56.53	56.91	67.76	41.04			

Table S2: **Benchmark results on novel view synthesis task.** State-of-the-art methods’ performance of novel test views on seen poses in each benchmark split. We abbreviate Instant-NGP [48] as ‘NGP’, NeuS [68] as ‘NS’, A-NeRF [62] as ‘AN’, *NeuralVolumes* [39] as ‘NV’, *NeuralBody* as ‘NB’ [53], *AnimatableNeRF* [52] as ‘AnN’, and *HumanNeRF* [71] as ‘HN’. Cell color indicate the best, second best, and third best performance in the same split respectively. We exclude the static methods NGP and NS during ranking and separate them with dash lines.



Figure S10: **Novel pose animation results on each data split.** From top to bottom, we illustrate the reposing results generated by (a-e): NeuralVolumes [39], A-NeRF [62], NeuralBody [53] AnimatableNeRF [52], and HumanNeRF [71]. Please zoom in for better visualization.

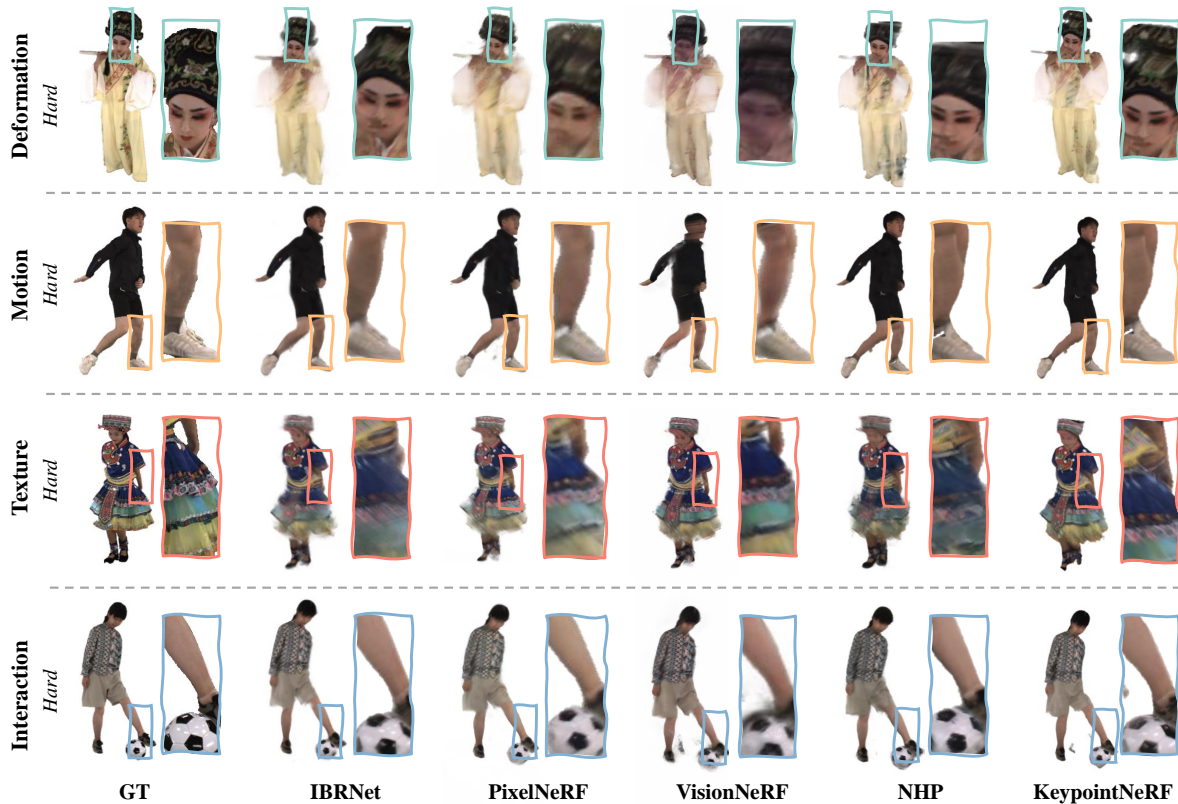


Figure S11: **Novel ID synthesis results on each data split.** Splits with hard difficulty are visualized to illustrate the robustness of generalizable methods on hard cases.

representations is worth great pondering.

C.2.3 Novel Identity Rendering

Detailed Settings. For novel identity synthesis, we review five selected state-of-the-art methods [73, 69, 46, 18, 32], and described their modification in Sec. C.1.3. The testing set is the same 39 testing sequences in the novel view and the novel pose benchmark. The training set consists of 400 sequences with full coverage of all categories and difficulties. We select four balanced views as source views. These source view images are cropped and resized into 512×512 resolution (same with the official implementation in [46, 11, 32]). For category-agnostic methods [73, 69, 35], we only provide segmentation and camera parameters during training and testing. For methods with human prior, we also input the fitted SMPLX or 3D keypoints. We train models on the full 60 views in the training identity sequences. For inference, we evaluate the unseen identities on the same 18 test views used in novel view and novel pose tasks, but on full sequences with frame skip at 45. All methods are trained under the same 8-V100 machine environment, and evaluated on a single V100.

Detailed Results. In the main text, we draw the conclusion that generalization methods use human prior [46, 32] is more robust than category-agnostic methods [73, 69, 35] according to the results reported in Tab. 3. To further illustrate

this conclusion, we compare the qualitative results in *Hard* level of each data factor in Fig. S11. Human prior methods generally render better images with more precise human shape and texture compared to category-agnostic methods, especially in the *Texture-Hard* and *Deformation-Hard* cases. Moreover, in addition to the influence of conceptual difference between image blending strategies to direct radiance prediction in the main text, we also compare the generalization of image feature extractors between PixelNeRF [73] and VisionNeRF [35]. VisionNeRF [35] uses a similar structure of PixelNeRF [73], but mainly incorporates the local CNN-based image encoder with a global vision transformer. Such a design achieves better visual quality on average with sharper texture details and higher scores in both Fig. S11 and Tab. 3, since the transformer is capable to learn more global coherence features across source views.

D. Cross-dataset Details

In this section, we provide more details as a supplement to the cross-dataset evaluation mentioned in the main paper. Specifically, we first describe the criteria for selecting the compared datasets, and review the key attributes of these datasets in Sec. D.1. Then, we introduce the implementation and setting details in Sec. D.2. We discuss the results in Sec. D.3. Finally, we analyze the impact of color consis-




Figure S12: **Cross-dataset affinity map for category-agnostic generalization methods.** We crossly evaluate models trained on each dataset and plot their performance on testing splits on each dataset. The PSNR, SSIM, and LPIPS* are plotted in separate matrices.

tency on multi-camera datasets in Sec. D.4.

D.1. Compared Datasets

To evaluate the potential of the proposed dataset on boosting algorithms’ generalizability from the data engineering aspect, we compare the proposed dataset with the most commonly used human-centric multiview datasets on the generalizable neural rendering task. For a fair comparison, we select datasets with foreground segmentation annotations and dense camera views. Thus, several well-known datasets are not suitable for this evaluation. For example, Human3.6M [24] only contains four RGB cameras, CMU Panoptic [28] and AIST++ [65, 34] lack official segmentation annotation. We select ZJU-MoCap [53], HuMMan [5] and GeneBody [11] for comparison. In this subsection, we discuss their main features and their adaption to generalizable methods.

ZJU-MoCap [53] is currently the most widely used dataset in human neural rendering domain. It contains 10 multiview performance sequences, with accurate camera calibration,

Some human rendering methods [32, 10] use their own tools to generate mask, we exclude these mask sources for fairness.

human segmentation as well as SMPL annotation. The main drawback of this dataset is the lack of diversity in clothing and motion and without human objection interaction. As mentioned in Sec. 3.3 in the main paper, color consistency design might be ignored in ZJU-MoCap [53], where obvious color differences can be observed between neighboring cameras, as shown in Fig. S15a. When training generalizable models on this dataset, we adopt the official splits and follow the implementation in KeypointNeRF [46].

HuMMan [5] is a human action dataset with data captured under 10 synchronized Kinect RGB-D cameras. It contains 400k sequences and 500 human actions which emphasize muscle-related movements. The clothing diversity is marginal where most subjects wear sports and daily costumes, and there is no human-object interaction either. As the source images come from the Kinect sensor, they might be stuck in low-quality, and obvious color differences can be found in HuMMan [5] dataset. Note that the full dataset is still unreachable, we train models on the released version, with its official list that contains a training split with 317 sequences and a testing split with 22 sequences. Noted that different from other datasets, where cameras are organized in a world coordinate near the origin that axis alignment

with the real world, HuMMan [5] uses a coordinate system relative to the first camera. Thus, we make a rigid transformation to eliminate the coordinate system difference.

GeneBody [11] is a recent multi-view human performance capture dataset, which contains relatively wide diversity coverage among clothing, motion, and interactions. It captures human performance with 48 synchronized 5MP cameras with a low proportion-in-view of the main subject, where the average height of the human bounding box is around 600 pixels. We use its official splits with 40 training sequences and 10 testing sequences.

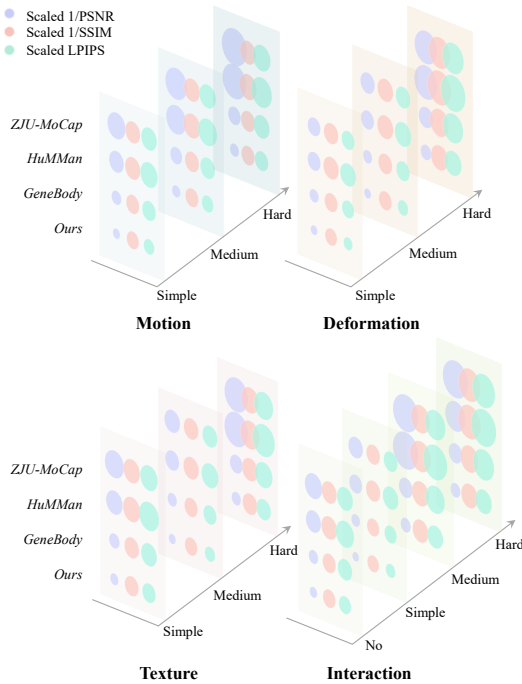


Figure S13: **Cross-dataset evaluation on our DNA-Rendering data splits.** We visualize the performance of models trained on different datasets on the proposed dataset’s splits.

D.2. Detailed Settings

Multiple factors might affect the impact assessment across different datasets on neural rendering tasks, *eg.*, the proportion of subjects in camera views, data scale, annotation accuracy, source view selection, training status, *etc.* Our main goal is to investigate where the *diversity* of the proposed dataset can benefit the generalization of human rendering. We conduct the experiments in the following settings. 1) In order to ensure the fairness of dataset comparison, we need to unify several input conditions, *e.g.*, number and resolution of source views, *etc.* Meanwhile, we only investigate the category-agnostic generalizable methods, namely PixelNeRF [73], IBRNet [69] and VisionNeRF [35], to avoid difference of input and accuracy of human prior in multiple datasets. 2) Evaluating the whole

object-centric images with background removal produces a large rendering metric gap if the center subjects’ proportions in camera views differ a lot. Thus, different from the novel pose benchmark in Sec. 4 and Sec. C.2.3 where a half resolution is used, we crop the subjects out from the original image in all datasets with square bounding boxes and resize them to 512×512 resolution for both source views and target views. 3) As the discussed datasets are all captured in dense circling camera arrays, we manually select four balanced views as the reference views at the same height that have a roughly 90-degree interval. 4) During training, we use the same learning rate over all datasets and stop the training process with the same global step. Each model is trained on one 8-V100 machine with distributed data-parallel stopping at $200k$ iterations. 5) Finally, we train and evaluate all the models based on the official splits of each dataset. For the comparable data volume magnitude of test samples, we size the data volume of test frames or test views on each dataset with $\langle 45$ frame-skip, 18 test views \rangle on GeneBody and DNA-Rendering, $\langle 45$ frame-skip, 12 uniform sampled test views \rangle on ZJU-MoCap, and $\langle 8$ frame-skip, 6 test views \rangle on HuMMan, respectively.

D.3. Additional Results

In the main paper, we present the average PSNR performance over all three general scene methods, *i.e.*, PixelNeRF [73], IBRNet [69], and VisionNeRF [35]. Here, we present the performances of these methods individually (Fig. S12), and illustrate the qualitative results (Fig. S14). We unfold the result analysis in terms of in-domain, and cross-domain in Sec. D.3.1 and Sec. D.3.2 respectively. A discussion on the impact of color consistency is discussed in Sec. D.4.

D.3.1 In-domain

In-domain refers to the problem of evaluating models with the trainset and testset sharing the same underlying data distribution. We observe two key phenomena: 1) models trained on datasets with low data diversity achieve better in-domain results. As shown in the diagonal elements of the matrices in Fig. S12, for in-domain generalization performance, methods trained on HuMMan [5] and ZJU-MoCap [53] achieve the best and second performances with relatively appealing metric values. In contrast, their in-domain performances on GeneBody and DNA-Rendering are worse than the other two datasets. Noting that test sets in ZJU-MoCap and HuMMan only contain cases with textureless clothing and easy motion illustrated in Fig. S14, which are easy cases in terms of data

Note that this setting leads to the absolute values in cross-dataset evaluation worse than the ones in novel identity task, due to the larger proportion of human foreground.

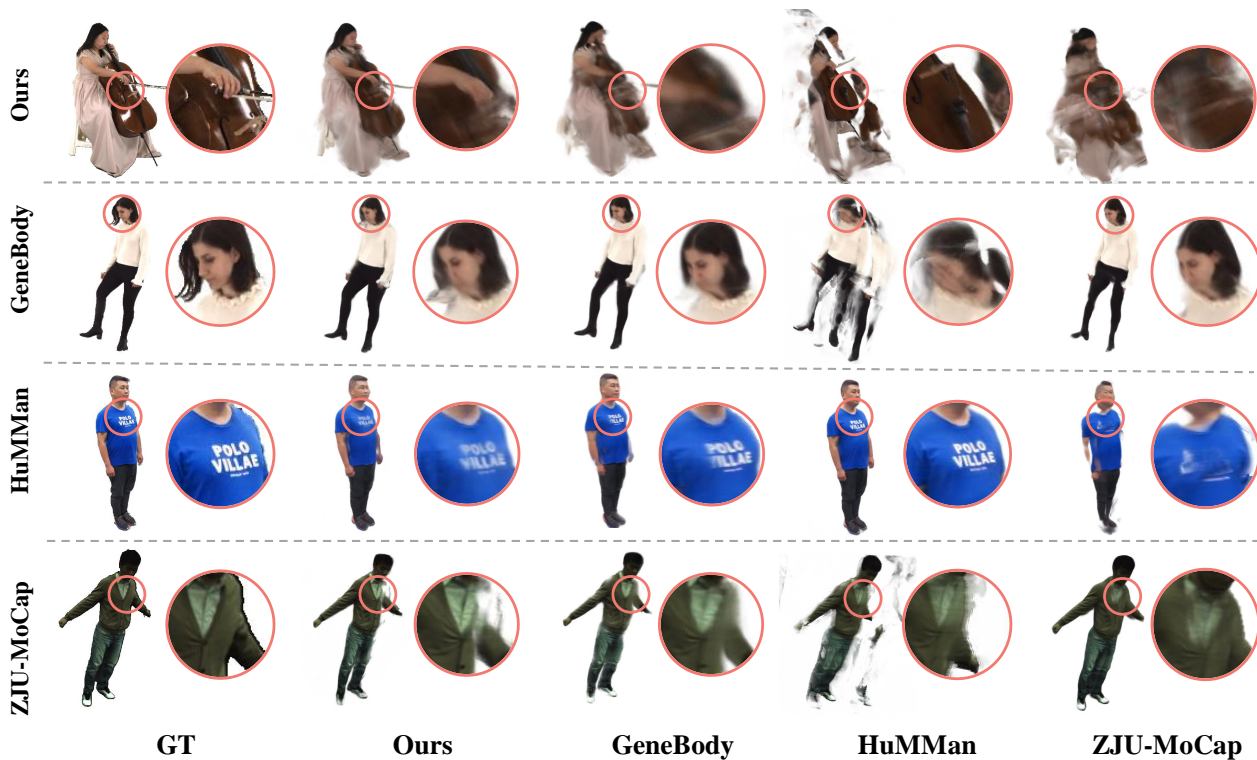


Figure S14: **Qualitative cross-dataset results.** We demonstrate samples from different datasets (left labels) generated by IBRNet [69] models trained on different datasets (bottom labels).

difficulty. Besides, the nature of textureless data and easy human geometry is relatively friendly to category-agnostic generalization methods that conduct multiview image feature aggregation in a common manner without geometry prior. 2) Larger data volume and diversity boosts in-domain performance. Concretely, like the proposed dataset, GeneBody [11] contains a train and test split with a wide distribution of clothing, accessories, and motion, while with far less data volume and diversity compared to DNA-Rendering (about 10% data volume of our dataset). Despite both test sets of GeneBody and proposed dataset containing cases with even distribution in multiple difficulties, all three methods demonstrate our boost on in-domain performance.

D.3.2 Cross-domain

Cross-domain refers to directly evaluating the pre-trained model on one dataset to the test split of another dataset, which is represented by the off-diagonal elements in Fig. S12. We expand the result analysis in two folds: 1) datasets with large variations of data attributes and difficulties boost the cross-domain generalization. Higher performance degradation can be observed in off-diagonal elements in each row in ZJU-MoCap [53] and HuMMan [5] in Fig. S12. Besides, even when ZJU-MoCap [53] trained model test on cross-domain case with easy clothing and motion, the cross-domain performance is still far from acceptable, refer to left most blue T-shirt case in Fig. S14. On the

contrary, GeneBody [11] and DNA-Rendering experience a very marginal degradation when evaluating other test sets. 2) Larger data volume and diversity can also boost cross-domain robustness. As mentioned in the in-domain part, DNA-Rendering is 10 times than GeneBody in terms of data volumes, such improvement helps increase the model’s generalization ability in considering margin. In the overwhelming majority of first-row elements, models with the proposed dataset get the best cross-domain results and even perform better than in-domain results of GeneBody.

Cross-domain on DNA-Rendering Splits. To further investigate the performance of different dataset-trained models’ performance on different human performance dimensions and difficulties, we visualize their performance on DNA-Rendering splits in Fig. S13. We plot the in-domain result of our dataset-trained model as a reference bar. GeneBody [11] has a relatively wide range of dimensions across our dimension and it achieves the best rendering quality among all splits. On the other hand HuMMan [5] and ZJU-MoCap [53], only contain easy clothing, motion, and interaction, the performance on more difficult splits drop significantly espec compared to GeneBody [11].

D.4. Impact of Color Consistency

To further analyze the impact of color consistency of training dataset in generalizable rendering, we unfold the stats across views on the cross-dataset results. Due to the different groundtruth in different views, it is hard to draw

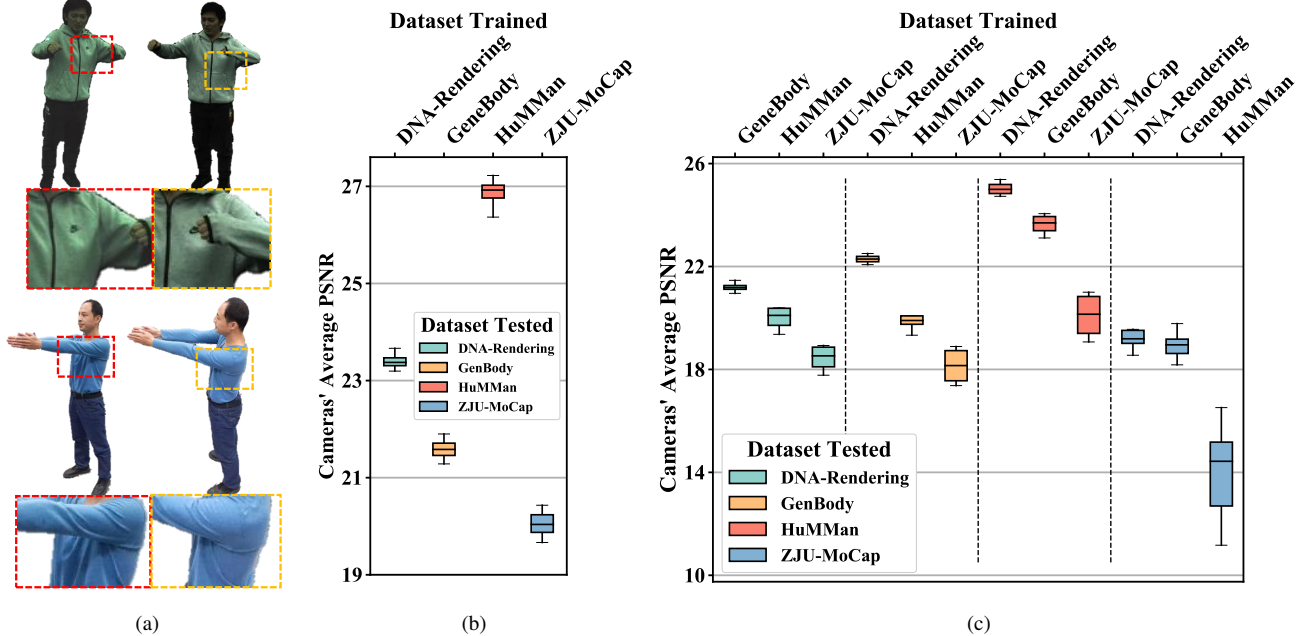


Figure S15: **Cross-dataset evaluation across views.** (a) Examples of the color differences between neighboring cameras in ZJU-MoCap [53] (top) and HuMMan [5] (bottom). (b) **In-domain** statistics on view average PSNR when trained and tested on the same dataset. (c) **Cross-domain** statistics on view average PSNR when trained and tested on different datasets. In both (b) and (c), the horizontal axis means different datasets trained, and the color of the box separates the datasets tested. The line in the middle indicates the median value, the box indicates the lower to the higher quartile, and the whiskers indicate the range of average PSNR across views.

any conclusion from any single frame. Thus, we expand the average PSNR across camera views and analyze the statistics. Noted that we only select the test views which have very close angle distances from the nearest source view, to erase the performance gap from the viewpoints. The average PSNR across testing cameras is plotted in Fig. S15. More concretely, we visualize the in-domain statistics of cameras' average PSNR in Fig. S15b. When training and testing a model on the same dataset, the camera color distinction will remain constant. Models trained on the datasets that cannot ensure color consistency across views (illustrated in Fig. S15a) might treat the color difference of different views as the view-depend effect and memorize it. The variance of cameras' average performance in such datasets is slightly higher than GeneBody [11] and the proposed dataset. Cross-domain generalization span on views is also plotted in Fig. S15c. Different from in-domain statistics, when generalizing on other datasets, models trained on datasets with color inconsistency all suffer a major average performance dropping, and the variance of camera performance becomes even larger. This phenomenon is very noticeable on cross-evaluation between ZJU-MoCap [53] and HuMMan [5]. While models trained on the proposed dataset as well as GeneBody [11], have very small view performance variations between each other. The increased view variation on ZJU-MoCap [53] and HuMMan [5] is due to the nature color difference on their groundtruth. In a nutshell, with the best color consistency, the proposed dataset

can benefit the community by providing high-quality data with faithful probing capability across views.

E. Future Work

Leaderboard. In the current human-centric rendering community, researchers from different institutions use different datasets and experimental settings to evaluate the performances of their algorithms. There is no agreement across institutions to benchmark human rendering methods under the same criterion yet. To reduce such divergence and align the standards, we proposed a large-scale diverse dataset DNA-Rendering, and construct a complete benchmark in three human-centric rendering tasks. Benchmarks are evaluated on different data splits on different factors and difficulties. Additionally, we conduct a cross-dataset evaluation which demonstrates the proposed dataset can benefit the community from its diversity and coverage. In DNA-Rendering, all actors are signed with agreements before data collection. Thus, all of the data is with a Creative Commons license and free for use under certain usage agreements. In the future, we will host a web-based leaderboard, and release the easy-to-run tools to the community to better reduce the divergence.

Robust Human-centric Matting Refinement. In the annotation pipeline, we use Grab-cut [55] to refine bad results of CNN-based methods, yet it is still not perfect. Since per-frame human labeling is impractical due to the volume of captured data, we involve human checks over segmentation

results, and only cases without major artifacts in the whole sequence across views will be released. We tried 3D methods to further refine results, *e.g.*, using Instant-NGP [48] to train with valid views and infer the bad views. However, the results are not appealing enough (blurry edges). We will further investigate more robust tools. These challenges could also benefit matting research. We believe with the development of relevant techniques, matting robustness will be improved in the near future.

New benchmarks. In this paper, we set up three task benchmarks as a kick-off of the DNA-Rendering dataset. Our datasets could potentially be used in many other tasks related to human-centric rendering, such as garment modeling/animation and human shape completion. We encourage and welcome the community to join us to unlock more downstream tasks.