

LUNA: Learning Universal 3D Human Animation Beyond Skinning

Peng Li^{1*}, Rawal Khirodkar², Junxuan Li², Yuan Dong², Chen Cao², Yuan Liu¹, Wenhan Luo¹, Yike Guo¹, and Shunsuke Saito²

¹ The Hong Kong University of Science and Technology

² Codec Avatars Lab, Meta




Fig. 1: Given a handful of human images, LUNA reconstructs a high-fidelity animatable 3D avatar, supporting versatile 2D control signals - including RGB images, 2D keypoints, hand-drawn sketches, and other unseen characters - without any additional preprocessing. Project page: <https://penghtyx.github.io/LUNA/>.

Abstract. Creating photorealistic, animatable 3D human avatars from monocular images still largely depends on Linear Blend Skinning (LBS) and parametric body models which constrain expressivity and often introduce artifacts due to imperfect fitting. We propose LUNA, an LBS-free universal neural animation model that directly maps multiple 2D controls like images, keypoints, sketch and unseen characters into 3D Gaussian deformations, bypassing explicit body fitting. At its core, a transformer-based motion regressor disentangles global rigid motion from fine-grained local dynamics to capture both coherent movement and subtle non-rigid effects. To resolve the inherent ambiguity of 2D-to-3D lifting while scaling beyond fitted datasets, we introduce hybrid supervision that distills soft structural priors from an LBS teacher and a loss that supports training on both limited fitted data and large in-the-wild unlabeled videos. Extensive experiments show LUNA achieves competitive visual fidelity compared to LBS-based approaches, while delivering realistic human motion and zero-shot cross-identity generalization across diverse driving modalities. To the best of our knowledge, LUNA is the first end-to-end 3D animatable model that supports implicit 2D driving.

* Work was done during an internship at Meta.

Keywords: 3D Gaussian Avatars, Implicit Control, Photorealism, Animation, 3D Reconstruction, Large-Scale Supervision

1 Introduction

Creating photorealistic, fully animatable 3D human avatars from monocular inputs has long stood as a central ambition at the intersection of computer vision and graphics, driven by immense demand in film production, interactive entertainment, and VR/AR. Traditional methods [2, 36] can deliver high fidelity, but they typically depend on costly multi-view capture and iterative optimization, confining scalability to controlled studio settings. In contrast, most monocular animation systems heavily rely on predefined parametric body models [13, 40, 48], a convenient prior that nonetheless constrains expressiveness and often compromises photorealism under complex, highly dynamic motion.

Large Reconstruction Models (LRMs) [19] have recently demonstrated striking generalization in 3D reconstruction, enabled by scalable transformer architectures and large datasets. Following this, several methods [15, 51, 81] predict canonical 3D Gaussians in a feed-forward manner while incorporating parametric priors for animatability. Despite their efficiency and generalization, the continued dependence on parametric models creates a persistent representation bottleneck. First, during training, these methods commonly use Linear Blend Skinning (LBS) to map a learnable canonical representation into posed space for rendering supervision. Consequently, inevitable body fitting errors lead to suboptimal canonical learning. Second, at inference, monocular 3D pose estimation is inherently ill-posed, and small joint errors can readily amplify into temporal jitter or geometric artifacts in the rendered avatar.

In this paper, we introduce **LUNA**, an LBS-free Neural Animation model that enables direct 3D animation from flexible 2D driving signals, including RGB images, sketches, and 2D keypoints. The central idea is to remove the parametric-body-bottleneck by learning a direct 2D-to-3D deformation mapping. Concretely, we propose a transformer based *disentangled neural animator* that predicts deformations in two parts: a global rigid motion that accounts for coarse pose and camera-aligned movement, and a local non-rigid field that captured pose-dependent dynamics at the level of individual Gaussians. By leveraging the long-range context of our model, LUNA extracts motion cues from the 2D input and translates them into per-Gaussian deformations. Our end-to-end setup allows the model to handle large articulated motion while preserving fine-grained, non-rigid effects that are difficult to express with LBS.

Directly learning 3D deformation from 2D inputs using only rendering supervision, however, is severely underconstrained: without an explicit notion of human shape, reconstructions can collapse into flattened configurations, see Fig. 6. We therefore introduce a hybrid supervision strategy. While LUNA is designed to surpass the limitations of LBS, we retain its utility as a *soft* structural prior. Concretely, we utilize an LBS-based teacher model not as a hard driver, but as

a regularizer that distills structural cues into the student. Combined with photometric rendering losses, this distillation preserves geometric integrity without inheriting the teacher’s representational limits. Finally, to scale beyond scarce body-tracked data, we use a loss reweighting scheme that dynamically balances rendering and distillation objectives based on the proportion of labeled data within each batch. This enables our model to scale efficiently from limited annotated samples to massive, unlabeled in-the-wild video collections.

A notable outcome of this formulation is the emergence of **cross-identity driving**. Although trained with matched source and driving identities, LUNA learns a deformation field that captures universal humanoid kinematics. At inference, it can be driven across identities implicitly by heterogeneous modalities - keypoints, sketches, or out-of-domain character images—without further tuning. Our results demonstrate that implicit driving using LUNA drastically reduces pose-induced jitter and better captures non-rigid dynamics than parametric baselines. Our contributions are summarized as follows:

- To the best of our knowledge, LUNA is the first model to enable universal 3D human animation directly driven by versatile 2D signals, establishing a new paradigm that bypasses explicit 3D control.
- We propose a body-fitting-free animation setup that extracts motion semantics directly from 2D inputs without any preprocessing (e.g., body fitting or foreground segmentation), thereby reducing error accumulation and significantly mitigating temporal jitter.
- We introduce a disentangled neural animator with a fully LBS-free decoder, enabling direct modeling of pose-dependent deformations from coarse articulations to complex non-rigid dynamics such as loose clothing, beyond the rigid constraints of template-based approaches.

2 Related Work

3D Avatar Reconstruction. High-fidelity 3D avatar reconstruction has traditionally relied on multi-view calibrated systems to capture subjects in a canonical space [2, 5, 9, 16, 21, 23, 34, 35, 38, 45, 49, 50, 54, 56, 59, 61, 66, 78, 79]. Early efforts utilized neural implicit representations such as NeRF [42, 69] to model geometry and appearance, but required time-consuming per-avatar optimization [23, 49, 61, 66]. Recently, the integration of 3D Gaussian Splatting (3DGS) [27] has significantly accelerated this process, enabling real-time rendering and faster convergence for person-specific avatars [25, 37, 44, 46, 82]. While these methods can achieve high fidelity by integrating body priors [14, 43, 48, 63], they remain inherently subject-specific and fail to generalize to unseen identities.

To overcome these limitations, feed-forward methods are designed to reconstruct avatars instantly. For static reconstruction, early regression-based efforts utilized pixel-aligned features to infer 3D surfaces or Gaussians from a single view [1, 17, 22, 55, 57, 64, 65, 77]. However, they often suffer from texture artifacts and self-occlusions. Another line of research leverages sparse-view [6, 7, 29, 70, 73,

75, 76] or novel-view synthesis from generative models [8, 18, 33] to achieve high-fidelity results. While impressive, these static reconstructions often lack dynamic motion priors, leading to unrealistic artifacts during animation. Driven by large-scale datasets, recent animatable feed-forward works [14, 51, 81] leverage scalable transformers to encode subjects into structured latent spaces, showing promising generalization capability. Our work builds upon these generalizable backbones but significantly extends their capability from static reconstruction to LBS-free, universal implicit animation.

3D Avatar Animation. Conventional 3D animation predominantly rely on LBS tied to parametric models like SMPL [40, 47]. To model non-rigid dynamics such as loose garments, recent works augment the LBS prior with pose-dependent Gaussian mapping or neural deformation fields [5, 32, 37, 51]. AniGS [53] further studies animatable Gaussian avatars from a single image, but still relies on a human-prior-driven animation pipeline. However, this tight coupling with the SMPL topology makes these methods highly sensitive to pose estimation errors and restricts them to standard humanoid structures. Although some attempts, such as TAVA [34] and GART [30], relax these constraints via skeleton-based deformations, they still necessitate person-specific multi-view fitting. Recent feed-forward reconstruction models such as PF-LHM [52] improve pose-free avatar reconstruction from casually captured images. Most recently, HumanRAM [71] introduced an implicit animator for feed-forward 2D reposing, but it lacks a cohesive 3D representation and remains bottlenecked by human template priors.

2D Human Animation. 2D Human animation aims to generate novel-pose images from one or more reference images. Early works [4, 39, 74] primarily formulated this task as signal-driven image-to-image translation. Recently, diffusion-based methods [20, 58, 62, 67, 72, 80] have achieved remarkable success by leveraging the powerful generative priors of 2D diffusion models. While these methods exhibit impressive versatility in handling diverse driving signals (e.g., poses, depth, or edges), they are inherently restricted on 2D image plane, often suffering from temporal flickering, view inconsistency, and a lack of true 3D geometric control, while the iterative denoising process remains computationally expensive.

In contrast to the existing 2D or 3D works, we are the first to bridge the gap between the versatility of 2D driving signal and the structural integrity of 3D modeling. Our proposed LUNA is a fully LBS-free framework that completely discards template dependency and efficiently maps 2D driving signals to 3D avatar motion. Our method not only captures complex non-rigid dynamics with superior flexibility but also achieves robust cross-modality generalization that previous 3D animators could not reach.

3 Method

Given a sparse set of unposed identity images $I_{id} = \{I_i\}_{i=1}^N$ and a single 2D driving image I_d , our goal is to synthesize a high-fidelity 3D human avatar that can be animated by I_d end-to-end manner. As shown in Fig. 2, LUNA has two




Fig. 2: Overview. Given N unposed multi-view identity images and a 2D driving signal, LUNA first reconstructs canonical 3D Gaussians with an *Identity Encoder*. A transformer-based *Implicit Neural Animator* then maps them to posed space conditioned on the driving signal. During training, the driving image is randomly sampled across modalities (RGB, keypoints or sketches).

components: an *Identity Encoder* (Sec. 3.1) that lifts the input into a canonical set of 3D Gaussian, and an *Implicit Neural Animator* (Sec. 3.2) that predicts disentangled deformations and attribute residuals to pose the avatar. The model is trained with a hybrid supervision strategy (Sec. 3.3). Although I_{id} and I_d share the same identity during training, our model generalizes to cross-identity driving at inference.

3.1 Canonical Identity Encoder

Inspired by feed-forward human reconstruction model [51], we design an Identity Encoder that lifts the reference images I_{id} into canonical 3D Gaussians, together with semantic features that condition the subsequent driving.

Semantic Query Tokens. To inject human-centric structural priors, we anchor our canonical avatar to a template topology. Specifically, we define K learnable query tokens, each associated with a vertex of the template mesh. To encode anatomical semantics, each token is augmented with a learnable positional embedding $Q \in \mathbb{R}^{K \times C}$ and a semantic label embedding $E_{sem} \in \mathbb{R}^{K \times C}$. The resulting 3D query tokens are formulated as $T_{gs} = Q + E_{sem}$, where C denotes the feature dimension. This design encourages a consistent correspondence between semantic parts and Gaussian primitives, improving identity stability under large pose and viewpoint changes.

Image Tokenization. In parallel, we extract fine-grained identity features from the multi-view inputs I_{id} . We use a pre-trained Sapiens [28] encoder to tokenize each view into patch features, and then aggregate them across views to obtain identity image tokens $T_{img} \in \mathbb{R}^{M \times C}$, where M is the number of image tokens. Following [51], we include both body tokens and face tokens in T_{img} .

Joint Feature Fusion and Canonical Decoding. To lift the 2D observations into a unified 3D canonical representation, we fuse the 3D query tokens T_{gs} with

the image tokens T_{img} using Multimodal Transformer (MM-Transformer) [12] blocks. Following [51], each block contains intra-modal self-attention and cross-modal attention to enable cross-modal feature exchange. After L blocks, we obtain canonical identity tokens T_{can} , which are decoded by a lightweight MLP into canonical Gaussians G_{can} . Each Gaussian primitive is parameterized as:

$$g = \{\mu, q, s, \alpha, c\} \quad (1)$$

where $\mu \in \mathbb{R}^3$ is the center, $q \in \mathbb{R}^4$ is the rotation quaternion, $s \in \mathbb{R}^3$ is the scale, $\alpha \in \mathbb{R}$ is the opacity, and c is the color. Gaussians G_{can} and tokens T_{can} are then passed to the Implicit Neural Animator.

3.2 Implicit Neural Animator

The Implicit Neural Animator maps the canonical avatar to posed space conditioned on a 2D driving signal I_d . To stabilize learning and improve expressiveness, LUNA disentangles motion into a global rigid transformation and local, per-Gaussian deformations, both predicted in a purely feed-forward manner.

Motion Tokenization. We extract motion cues from I_d using a pre-trained DINOv3 [60] encoder, producing motion tokens T_{motion} . We choose DINOv3 over Sapiens for two reasons: (i) animation primarily depends on kinematic layout rather than fine-grained identity or texture, and (ii) DINOv3 generalizes better to out-of-domain, non-photorealistic driving signals such as 2D keypoints and abstract sketches.

Global Transformation Prediction. We model the avatar’s coarse spatial alignment with a global rotation R and translation T . We first aggregate the DINOv3 motion tokens into a global descriptor and feed it to two lightweight MLP heads. Directly regressing Euler angles is unstable due to discontinuities at $\pm\pi$, so we predict rotation in a continuous trigonometric form. Specifically, for each axis $i \in \{x, y, z\}$, the rotation head outputs a 2D vector,

$$(s_i, c_i) = \text{Tanh}(\text{MLP}_{rot}(T_{motion})) \quad (2)$$

which is normalized and then composed into a valid rotation matrix $R \in \mathbb{R}^{3 \times 3}$. For translation, regressing absolute coordinates can exhibit scale drift. Instead, we predict a normalized offset $\hat{t}$ and recover translation via dataset denormalization:

$$T = \sigma_T \cdot \hat{t} + \mu_T, \hat{t} = \text{Tanh}(\text{MLP}_{trans}(T_{motion})) \quad (3)$$

where μ_T and σ_T are the mean and standard deviation computed from the training data distribution.

Deformed Avatar Gaussian. Global R and T provide coarse alignment, while local deformations capture the coupling of articulated motion and non-rigid dynamics. We first project the canonical tokens T_{can} to a lower-dimensional latent and concatenate them with learnable motion queries $Q_{motion} \in \mathbb{R}^{K \times (C/2)}$ to form identity queries:

$$T_{id} = \text{MLP}_{proj}(T_{can}) \parallel Q_{motion} \quad (4)$$

where $\parallel$ denotes channel-wise concatenation. We fuse T_{id} with the motion tokens T_{motion} using the MM-Transformer, and decode per-Gaussian residuals with a lightweight MLP: position offset $\Delta\mu$, rotation residual Δq , and color residual Δc . For stability, we keep scale s and opacity α fixed during animation. The posed Gaussians are:

$$g_{posed} = \{\mu_p, q_p, s, \alpha, c_p\} \quad (5)$$

where $\mu_p = R(\mu + \Delta\mu) + T$, $q_p = q_R \otimes \Delta q \otimes q$, and $c_p = c + \Delta c$. Here $\otimes$ denotes quaternion multiplication and q_R is the quaternion corresponding to the global rotation R .

3.3 Hybrid Supervision Strategy

Our training corpus mixes samples with parametric body-fitting annotations and large-scale unannotated videos. We first train the Identity Encoder on labeled data by extending LHM to a multi-view setting (MV-LHM), and then train the Neural Animator using both labeled and unlabeled data. Concretely, we sample the two data types within each mini-batch and optimize a balanced hybrid objective for stable training.

Photometric Rendering Loss. For all samples, we rasterize the posed 3D Gaussians into an image I_{render} and enforce photometric consistency with the driving frame I_d :

$$\mathcal{L}_{render} = \mathcal{L}_1 + \mathcal{L}_{mask} + \mathcal{L}_{LPIPS} \quad (6)$$

where $\mathcal{L}_1$ is the pixel-wise L_1 loss, $\mathcal{L}_{mask}$ is the foreground mask loss, and $\mathcal{L}_{LPIPS}$ is the perceptual loss.

LBS Distillation Regularization. For samples with parametric pose labels, we use an LBS-based teacher to provide soft structural targets. The teacher produces pseudo attributes for posed Gaussians, including positions $\hat{\mu}$, rotations $\hat{q}$, and colors $\hat{c}$. We guide the animation transformer with the distillation loss:

$$\mathcal{L}_{distill} = \|\mu_p - \hat{\mu}\|_1 + \lambda_q \mathcal{L}_{rot}(q_p, \hat{q}) + \lambda_c \|c_p - \hat{c}\|_1 \quad (7)$$

where $\mathcal{L}_{rot}$ denotes the cosine distance between quaternions.

Projection Loss for Transformation. A natural strategy is to supervise the predicted global rotation and translation on the annotated subset and rely on photometric rendering to generalize to unlabeled data. In practice, direct supervision is stable for rotation, so we supervise R via its representation:

$$\mathcal{L}_R = \|(s, c) - (\hat{s}, \hat{c})\|_1 \quad (8)$$

However, for translation T , even marginal estimation errors (especially along the depth axis) can induce severe spatial misalignment, thus dominating the loss gradients and destabilizing the training. We therefore avoid direct 3D supervision on T and instead impose a 2D reprojection loss on posed Gaussian centers:

$$\mathcal{L}_{proj} = \frac{1}{K} \sum_{i=1}^K \|\Pi(\mu_p^i) - \hat{u}^i\|_1 \quad (9)$$

where $\Pi(\cdot)$ represents the camera projection.

Balanced Hybrid Optimization. When mixing labeled and unlabeled data, photometric gradients from the unlabeled majority can overwhelm the structural signal, leading to geometric collapse (see Fig. 6). To maintain gradient balance, we form mini-batches with a fixed labeled-to-unlabeled ratio (e.g., 1:5) and upweight the distillation term:

$$\mathcal{L}_{total} = \mathcal{L}_{render} + \mathcal{L}_R + \mathcal{L}_{proj} + \lambda_{distill} \cdot \mathcal{L}_{distill} \quad (10)$$

We set $\lambda_{distill}$ inversely proportional to the labeled fraction (e.g., $\lambda_{distill} = 5$ for 1:5), preventing the structural gradients from being diluted. This simple reweighting enables learning fine non-rigid details from large-scale unlabeled videos while retaining LBS-anchored geometric stability.

4 Experiments

4.1 Implementation Details

Data. Our training corpus integrates four distinct data sources to ensure robust generalization across diverse identities and poses.

- **Video35K:** 35,000 in-the-wild full-body video clips, uniformly sampling 60 frames per clip. Video35K was not collected from a third-party data provider under a licensing agreement that permits machine learning research and publication, and is curated from the same source as the Sapiens pre-training data [28].
- **iPhone1K:** 1,100 monocular videos captured by iPhone, featuring subjects performing basic turn-around motions. It is a proprietary dataset collected onsite by our team and collaborators, following the phone-scan capture setting used in URAvatar [31]. All participants provided written informed consent covering the use of their data for machine learning research and publication. We select 1,000 identities and uniformly sample 32 frames per identity for training.
- **Cloth10K:** To explicitly model complex non-rigid dynamics, we curate 10,000 identities with complex garments and reserve 100 distinct clips for evaluation. Cloth10K is a derivative dataset constructed by us from Video35K using a conditional video generation model, and is used under the same applicable licensing terms.
- **Dome:** A multi-view dataset of 900 identities captured in a 3D studio with 200 synchronized cameras. It is a proprietary onsite capture dataset collected by our team and collaborators, and follows the same type of studio capture source as Ava-256 [41]. All participants provided written informed consent covering the use of their data for machine learning research and publication. We uniformly sample 200 frames per identity from the original $\sim 2,000$ -frame casual motion sequences for training.



Fig. 3: Our results with diverse 2D driving signals, including sketch image, color image, 2D keypoint and other out-of-domain images.

In total, the training dataset includes over 36,000 identities and 3.7M frames.

Evaluation. We evaluate self-reenactment and animation quality on the 100 Cloth10K clips and 100 iPhone sequences. We also benchmark on the widely used NeuMan [24] dataset, which was obtained from its official public release and used following the protocol described by its authors. For optimization-based baselines (Vid2Avatar [14], ExAvatar [43]), we strictly follow the official frame splits for training. For feed-forward methods, we uniformly select one or four views from the training splits as reference input. All methods are evaluated on the official test splits. We further quantitatively evaluate cross-identity driving on NeuMan.



Fig. 4: Qualitative comparisons on Cloth10K. LBS-dependent approach (IDOL and LHM) suffer from severe ID shift or structural tearing. In contrast, LUNA preserves the structural integrity and continuous topology of the fabric.

Baselines and Metrics. We compare against monocular optimization methods (Vid2Avatar [14], ExAvatar [43]) and recent feed-forward models (IDOL [81], LHM [51], UP2YOU [3]). We also include our extended MV-LHM to benchmark multi-view inputs. Photometric quality is measured via PSNR, L1 error, and LPIPS. To quantify motion stability and temporal plausibility, we employ Mean Acceleration Error (MAE) [26] and Mean Squared Jerk (MSJ) [11]. MAE evaluates the second-order temporal derivative of 3D Gaussian trajectories, reflecting overall motion stability:

$$E_{acc} = \frac{1}{N(T-2)} \sum_{i=1}^N \sum_{t=2}^{T-1} \left\| \frac{\mathbf{x}_{i,t+1} - 2\mathbf{x}_{i,t} + \mathbf{x}_{i,t-1}}{\Delta t^2} \right\| \quad (11)$$

where $\mathbf{x}_{i,t}$ denotes the 3D position of the i -th Gaussian at frame t , and Δt is the time interval. MSJ measures the third-order derivative, serving as a highly sensitive indicator for high-frequency "jittering" artifacts,

Table 1: Quantitative comparisons of reconstruction quality. We evaluate LUNA on our collected Cloth10K (focusing on non-rigid garments) and the public NeuMan [24] dataset. Bold indicates the best performance, and * denotes our implementation for multiview inputs.

Type	Method	Cloth10K (Non-rigid)			NeuMan		
		PSNR↑	L1↓	LPIPS↓	PSNR↑	L1↓	LPIPS↓
Optim. based	Vid2Avatar [14]	18.971	0.042	0.157	26.853	0.012	0.017
	ExAvatar [43]	19.533	0.041	0.153	31.270	0.009	0.009
Feed-forward	IDOL [81]	17.430	0.067	0.213	24.700	0.037	0.051
	LHM [51]	19.201	0.047	0.167	25.310	0.029	0.039
	UP2YOU [3]	19.745	0.047	0.149	25.492	0.026	0.041
	MV-LHM*	20.124	0.038	0.158	26.832	0.017	0.023
	LUNA (Ours)	22.072	0.027	0.131	26.819	0.015	0.023

Table 2: Quantitative evaluation of cross-identity driving on NeuMan. For Cross-ID, we animate 10 curated in-the-wild identity images with NeuMan motions and use the generated videos as driving signals. The right subtable reports 2D keypoint error extracted by Sapiens and normalized by the person bounding box.

Method	PSNR↑	L1↓	LPIPS↓					
LHM++ [51]	25.744	0.023	0.032	Method	IDOL	LHM	MV-LHM	Ours
LUNA (Cr-ID)	26.800	0.017	0.027	Err.(%)↓	28.7	26.7	26.7	17.3
LUNA (Self)	26.819	0.015	0.023					

$$E_{jerk} = \frac{1}{N(T-3)} \sum_{i=1}^N \sum_{t=3}^{T-1} \left\| \frac{\mathbf{x}_{i,t+1} - 3\mathbf{x}_{i,t} + 3\mathbf{x}_{i,t-1} - \mathbf{x}_{i,t-2}}{\Delta t^3} \right\|^2 \quad (12)$$

Lower values for both metrics indicate smoother, physically coherent animations.

Hyperparameters. The Identity Encoder is initialized from our MV-LHM, pre-trained on $\sim 1\text{M}$ pose-annotated videos (details in Appendix). The Neural Animator is trained in two progressive stages. At Stage 1 (monocular pre-training), the model is trained on Video34K, iPhone1K, and the Cloth10K training split using 64 A100 GPUs (batch size 64) for 30K iterations with a learning rate of 4×10^{-4} . The training begins with a 1,000-iteration warmup supervised only by the global rotation and projection loss ($\mathcal{L}_R + \mathcal{L}_{proj}$), after which the full objective $\mathcal{L}_{total}$ is applied. At stage 2, we fine-tune the model on the Dome dataset using 16 A100 GPUs for 30K iterations with a decayed learning rate of 1×10^{-4} . The hyper-parameters are set to $\lambda_q = \lambda_c = 0.5$, $K = 8192$, $C = 1024$, $N = 4$. We employ MHR [13] as body template in our experiments.

4.2 Results

Reconstruction and Animation. By directly extracting motion cues from 2D signals, LUNA bypasses explicit 3D pose estimation and rigid SMPL fitting. This end-to-end paradigm inherently avoids the error accumulation and domain shift vulnerabilities of LBS based animation. Consequently, our approach enables flexible 3D animation driven by highly diverse, cross-modality sources, including hand-drawn sketches, 2D skeletons, and out-of-domain humanoid images (Fig. 1, 3). We refer readers to the supplementary video for dynamic visualizations.

Cross-identity driving and pose accuracy. Although LUNA is trained with matched source and driving identities, it generalizes to cross-identity driving without any explicit transfer module. To evaluate this setting on NeuMan, we curate 10 in-the-wild human images for each video and use Wan-Animate [10] to animate them with the corresponding NeuMan motions. These generated videos are then used as driving signals for LUNA. As shown in Tab. 2, cross-identity driving remains close to self-reenactment and outperforms LHM++ [51], indicating that the animator primarily extracts motion rather than appearance from the driving signal. We also evaluate normalized 2D keypoint error using Sapiens detections; LUNA achieves the lowest error, further confirming that direct implicit driving improves motion correctness.

Comparison of reconstruction quality with LBS-based methods. We conduct qualitative comparisons between state-of-the-art optimization-based and feed-forward works with ours on 100 Cloth10K samples. As shown in Fig. 4, LBS-based methods like IDOL and LHM exhibit severe ID shift or catastrophic geometric collapse on loose garments due to the skeletal binding constraints. By discarding the rigid LBS, our implicit animator can naturally capture the non-linear sliding of fabric and significantly mitigate garment tearing artifacts. Quantitatively, as reported in Tab. 1, our approach significantly outperforms all baselines across photometric metrics on the Cloth10K dataset, demonstrating a clear advantage in modeling complex motion of non-rigid garments. On the NeuMan dataset, which primarily features standard walking motions and tighter clothing, LUNA achieves comparable performance with feed-forward baselines, proving that our design maintains high fidelity on simpler, highly constrained scenarios as well.

Comparisons of animation quality. Since frame-wise photometric metrics fail to reflect temporal artifacts, we quantitatively evaluate animation stability using MAE and MSJ. As reported in Tab. 3, LUNA significantly outperforms all pure LBS and hybrid baselines, effectively mitigating the frame-to-frame jitter caused by noisy pose estimation and rigid skinning constraints and achieving a $4.5\times$ reduction over the strongest baseline (MV-LHM) on MSJ. Notably, the gains are consistent across identities and motions, indicating robust temporal behavior rather than isolated improvements.

Qualitatively, we visualize the per-Gaussian motion trajectories as streamlines (Fig. 5). Trajectories produced by LBS-based methods (Fig. 5a) exhibit




Fig. 5: Qualitative comparisons of animation smoothness. We visualize the per-gaussian trajectory of (a) LBS-based animation baseline (b) our neural approach. The LBS-based methods suffer from erratic trajectories due to pose estimation errors and rigid skinning constraints. For all methods, we use SAM-3D-Body [68] to extract body parameters from driving videos and transform them to SMPL(X) for avatar animation.

Table 3: Quantitative comparison of motion quality. We compare our neural animator and various baselines. Due to the error and ambiguity of monocular pose estimation, LBS-based methods exhibit higher motion jitter, whereas our method achieves superior smoothness.

Methods	IDOL	UP2YOU	Vid2Avatar	ExAvatar	LHM	MV-LHM*	Ours
Animation	LBS		LBS + Neural				Neural
MSJ ($\downarrow$)	0.0315	0.0321	0.0237	0.0231	0.0214	0.0144	0.0032
MAE ($\downarrow$)	0.0591	0.0591	0.0493	0.0532	0.0522	0.0477	0.0225

severe "zig-zag" artifacts and erratic spatial shifts, particularly around high-dynamic regions. In contrast, our approach (Fig. 5b) yields smooth, continuous streamlines that better preserve local coherence over time. This advantage is especially evident in the nearly stationary pose (Identity 3, rightmost column), where motion trajectories should ideally degenerate into scatter points. The baseline instead produces spurious, chaotic streamlines induced by pose estimation noise, while LUNA remains near-static and visually stable.

4.3 Ablation

Effect of distillation regularization. As shown in Fig. 6, removing LBS-guided distillation (*w/o distill reg.*) causes severe depth collapse and flattening artifacts, degrading both shape coherence and motion consistency. This confirms

that while LBS is too rigid for final non-rigid animations, its structural prior is indispensable for resolving early-stage depth ambiguity and establishing global geometric stability.

Effect of multiview finetuning. Monocular supervision (*w/o multiview finetuning*) inherently suffers from perspective distortions and tends to produce floating Gaussian splats, leading to implausible renderings (Fig. 6). Incorporating multi-view finetuning directly eliminates these perspective ambiguities and enforces the implicitly animated 3D Gaussians to be more spatially compact, significantly elevating the overall visual quality.

Effect of global rotation decoupling. We further investigate the design of the decoupled deformation field by training a variant that regresses all Gaussian offsets in a single unified model. As shown in Fig. 7, this setting (*w/o Rot*) forces the model to compensate for large articulations, leading to severe drifting artifacts and ragged boundaries. In contrast, our model focuses exclusively on local non-rigid dynamics, producing sharp results that closely align with LBS reference, as further corroborated by the quantitative results reported in Tab. 4.

5 Conclusion

We present LUNA, a novel LBS-free framework for high-fidelity 3D human animation. Bypassing explicit pose estimation entirely, LUNA leverages a Transformer-based animator to directly extract motion cues from diverse 2D signals. To prevent the geometric collapse inherent to LBS-free approaches, we introduce a hybrid supervision strategy utilizing LBS-guided distillation. Finally, the decoupled global-local formulation empowers LUNA to capture complex, uncon-





Fig. 6: Ablations of distillation regularization and multiview finetuning.

Table 4: Ablation study of each components in LUNA training. We evaluate the contribution of distillation regularization, multiview finetuning, and decoupling rotation from gaussian deformations on 100 samples of iPhone1K and 20 dome identities.

Setting	iPhone data			dome data		
	PSNR $\uparrow$	L1 $\downarrow$	LPIPS $\downarrow$	PSNR $\uparrow$	L1 $\downarrow$	LPIPS $\downarrow$
w/o Structural Distillation	21.185	0.032	0.145	21.712	0.029	0.138
w/o Global Rotation	23.102	0.025	0.118	23.294	0.022	0.115
w/o Multiview Finetuning	23.931	0.021	0.102	23.864	0.019	0.099
Full Model (Ours)	24.136	0.019	0.094	24.374	0.017	0.093



Driver w/o Rot Ours LBS

Fig. 7: Ablation of decoupled global rotation prediction. Without global rotation, the deformation model fails to handle large motion (e.g. turning round).

Fig. 8: Challenging pose-driven cases.

LUNA handles challenging poses and loose-cloth mismatch in many cases, while more extreme pose and body-shape mismatch remain difficult.

strained non-rigid dynamics. Extensive experiments demonstrate LUNA’s robust zero-shot generalization across heterogeneous identities, loose clothing, and cross-modality driving sources, establishing a highly flexible paradigm for 3D digital human animation.

Limitations and Future Work. While LUNA is highly flexible, several directions remain to further strengthen its robustness. First, severe occlusions in the driving video can degrade the extracted 2D semantic cues and occasionally introduce temporal instability. Second, although LUNA handles many challenging pose-driven and loose-cloth cases (Fig. 8), very extreme poses or large body-shape mismatch can still degrade cross-identity animation because motion and shape are not explicitly disentangled. Moving forward, incorporating explicit temporal sequence modeling, improving occlusion-aware cue extraction, and scaling training with more diverse motions, apparel, and viewpoints are promising avenues to address these cases.

References

1. Alldieck, T., Zanfir, M., Sminchisescu, C.: Photorealistic monocular 3d reconstruction of humans wearing clothing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
2. Bagautdinov, T., Wu, C., Simon, T., Prada, F., Shiratori, T., Wei, S.E., Xu, W., Sheikh, Y., Saragih, J.: Driving-signal aware full-body avatars. *ACM Transactions on Graphics (TOG)* **40**(4), 1–17 (2021)
3. Cai, Z., Li, Z., Li, X., Li, B., Wang, Z., Zhang, Z., Xiu, Y.: Up2you: Fast reconstruction of yourself from unconstrained photo collections. *arXiv preprint arXiv:2509.24817* (2025)
4. Chan, C., Ginosar, S., Zhou, T., Efros, A.A.: Everybody dance now. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5933–5942 (2019)
5. Chen, J., Hu, J., Wang, G., Jiang, Z., Zhou, T., Chen, Z., Lv, C.: Taoavatar: Real-time lifelike full-body talking avatars for augmented reality via 3d gaussian splatting. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 10723–10734 (2025)
6. Chen, J., Yi, W., Ma, L., Jia, X., Lu, H.: Gm-nerf: Learning generalizable model-based neural radiance fields from multi-view images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20648–20658 (June 2023)
7. Chen, M., Zhang, J., Xu, X., Liu, L., Cai, Y., Feng, J., Yan, S.: Geometry-guided progressive nerf for generalizable and efficient neural human rendering. In: ECCV (2022)
8. Chen, W., Li, P., Zheng, W., Zhao, C., Li, M., Zhu, Y., Dou, Z., Wang, R., Liu, Y.: Synchuman: Synchronizing 2d and 3d generative models for single-view human reconstruction. *arXiv preprint arXiv:2510.07723* (2025)
9. Chen, Y., Zheng, Z., Li, Z., Xu, C., Liu, Y.: Meshavatar: Learning high-quality triangular human avatars from multi-view videos (2024), <https://arxiv.org/abs/2407.08414>
10. Cheng, G., Gao, X., Hu, L., Hu, S., Huang, M., Ji, C., Li, J., Meng, D., Qi, J., Qiao, P., et al.: Wan-animate: Unified character animation and replacement with holistic replication. *arXiv preprint arXiv:2509.14055* (2025)
11. Du, Y., Kips, R., Pumarola, A., Starke, S., Thabet, A., Sanakoyeu, A.: Avatars grow legs: Generating smooth human motion from sparse tracking inputs with diffusion model. In: CVPR (2023)
12. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling rectified flow transformers for high-resolution image synthesis. In: ICML (2024)
13. Ferguson, A., Osman, A.A.A., Bescos, B., Stoll, C., Twigg, C., Lassner, C., Otte, D., Vignola, E., Prada, F., Bogo, F., Santesteban, I., Romero, J., Zarate, J., Lee, J., Park, J., Yang, J., Doublestein, J., Venkateshan, K., Kitani, K., Kavan, L., Farra, M.D., Hu, M., Cioffi, M., Fabris, M., Ranieri, M., Modarres, M., Kadlec, P., Khiroudkar, R., Abdrashitov, R., Prévost, R., Rajbhandari, R., Mallet, R., Pearsall, R., Kao, S., Kumar, S., Parrish, S., Yu, S.I., Saito, S., Shiratori, T., Wang, T.L., Tung, T., Xu, Y., Dong, Y., Chen, Y., Xu, Y., Ye, Y., Jiang, Z.: Mhr: Momentum human rig (2025), <https://arxiv.org/abs/2511.15586>
14. Guo, C., Jiang, T., Chen, X., Song, J., Hilliges, O.: Vid2avatar: 3d avatar reconstruction from videos in the wild via self-supervised scene decomposition. In:

- Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12858–12868 (2023)
15. Guo, C., Li, J., Kant, Y., Sheikh, Y., Saito, S., Cao, C.: Vid2avatar-pro: Authentic avatar from videos in the wild via universal prior. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 5559–5570 (2025)
 16. Habermann, M., Liu, L., Xu, W., Zollhoefer, M., Pons-Moll, G., Theobalt, C.: Real-time deep dynamic characters. *ACM Transactions on Graphics* **40**(4) (aug 2021)
 17. He, T., Xu, Y., Saito, S., Soatto, S., Tung, T.: Arch++: Animation-ready clothed human reconstruction revisited. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 11046–11056 (2021)
 18. Ho, I., Song, J., Hilliges, O., et al.: Sith: Single-view textured human reconstruction with image-conditioned diffusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 538–549 (2024)
 19. Hong, Y., Zhang, K., Gu, J., Bi, S., Zhou, Y., Liu, D., Liu, F., Sunkavalli, K., Bui, T., Tan, H.: Lrm: Large reconstruction model for single image to 3d. *arXiv preprint arXiv:2311.04400* (2023)
 20. Hu, L.: Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8153–8163 (2024)
 21. Hu, T., Yu, T., Zheng, Z., Zhang, H., Liu, Y., Zwicker, M.: Hvtr: Hybrid volumetric-textural rendering for human avatars. In: 2022 International Conference on 3D Vision (3DV) (2022)
 22. Huang, Z., Xu, Y., Lassner, C., Li, H., Tung, T.: Arch: Animatable reconstruction of clothed humans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3093–3102 (2020)
 23. Jiakai, Z., Xinhang, L., Xinyi, Y., Fuqiang, Z., Yanshun, Z., Minye, W., Yingliang, Z., Lan, X., Jingyi, Y.: Editable free-viewpoint video using a layered neural representation. In: *ACM SIGGRAPH* (2021)
 24. Jiang, W., Yi, K.M., Samei, G., Tuzel, O., Ranjan, A.: Neuman: Neural human radiance field from a single video. In: Proceedings of the European conference on computer vision (ECCV) (2022)
 25. Jung, H., Brasch, N., Song, J., Perez-Pellitero, E., Zhou, Y., Li, Z., Navab, N., Busam, B.: Deformable 3d gaussian splatting for animatable human avatars (2023), <https://arxiv.org/abs/2312.15059>
 26. Kanazawa, A., Zhang, J.Y., Felsen, P., Malik, J.: Learning 3d human dynamics from video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
 27. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *TOG* (2023), <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
 28. Khirodkar, R., Bagautdinov, T., Martinez, J., Zhaoen, S., James, A., Selednik, P., Anderson, S., Saito, S.: Sapiens: Foundation for human vision models. In: *European Conference on Computer Vision*. pp. 206–228. Springer (2024)
 29. Kwon, Y., Fang, B., Lu, Y., Dong, H., Zhang, C., Carrasco, F.V., Mosella-Montoro, A., Xu, J., Takagi, S., Kim, D., Prakash, A., la Torre, F.D.: Generalizable human gaussians for sparse view synthesis. In: *European Conference on Computer Vision* (2024)
 30. Lei, J., Wang, Y., Pavlakos, G., Liu, L., Daniilidis, K.: Gart: Gaussian articulated template models (2023)

31. Li, J., Cao, C., Schwartz, G., Khirodkar, R., Richardt, C., Simon, T., Sheikh, Y., Saito, S.: Uravatar: Universal relightable gaussian codec avatars. In: ACM SIGGRAPH 2024 Conference Papers (2024)
32. Li, M., Yao, S., Xie, Z., Chen, K.: Gaussianbody: Clothed human reconstruction via 3d gaussian splatting (2024), <https://arxiv.org/abs/2401.09720>
33. Li, P., Zheng, W., Liu, Y., Yu, T., Li, Y., Qi, X., Li, M., Chi, X., Xia, S., Xue, W., et al.: Pshuman: Photorealistic single-view human reconstruction using cross-scale diffusion. arXiv preprint arXiv:2409.10141 (2024)
34. Li, R., Tanke, J., Vo, M., Zollhofer, M., Gall, J., Kanazawa, A., Lassner, C.: Tava: Template-free animatable volumetric actors. In: European Conference on Computer Vision (ECCV) (2022)
35. Li, Z., Zheng, Z., Liu, Y., Zhou, B., Liu, Y.: Posevocab: Learning joint-structured pose embeddings for human avatar modeling. ACM SIGGRAPH Conference Proceedings (2023)
36. Li, Z., Zheng, Z., Wang, L., Liu, Y.: Animatable gaussians: Learning pose-dependent gaussian maps for high-fidelity human avatar modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19711–19722 (2024)
37. Li, Z., Zheng, Z., Wang, L., Liu, Y.: Animatable gaussians: Learning pose-dependent gaussian maps for high-fidelity human avatar modeling. In: CVPR (2024)
38. Liu, L., Habermann, M., Rudnev, V., Sarkar, K., Gu, J., Theobalt, C.: Neural actor: Neural free-view synthesis of human actors with pose control. ACM Trans. Graph.(ACM SIGGRAPH Asia) (2021)
39. Liu, W., Piao, Z., Min, J., Luo, W., Ma, L., Gao, S.: Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5904–5913 (2019)
40. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: a skinned multi-person linear model. TOG **34**(6), 1–16 (2015)
41. Martinez, J., Kim, E., Romero, J., Bagautdinov, T., Saito, S., Yu, S.I., Anderson, S., Zollhöfer, M., Wang, T.L., Bai, S., Li, C., Wei, S.E., Joshi, R., Borsos, W., Simon, T., Saragih, J., Theodosios, P., Greene, A., Josyula, A., Maeta, S.M., Jewett, A.I., Venshtain, S., Heilman, C., Chen, Y.T., Fu, S., Elshaer, M.E.A., Du, T., Wu, L., Chen, S.C., Kang, K., Wu, M., Emad, Y., Longay, S., Brewer, A., Shah, H., Booth, J., Koska, T., Haidle, K., Andromalos, M., Hsu, J., Dauer, T., Selednik, P., Godisart, T., Ardisson, S., Cipperly, M., Humberston, B., Farr, L., Hansen, B., Guo, P., Braun, D., Krenn, S., Wen, H., Evans, L., Fadeeva, N., Stewart, M., Schwartz, G., Gupta, D., Moon, G., Guo, K., Dong, Y., Xu, Y., Shiratori, T., Prada, F., Pires, B.R., Peng, B., Buffalini, J., Trimble, A., McPhail, K., Schoeller, M., Sheikh, Y.: Codec Avatar Studio: Paired Human Captures for Complete, Driveable, and Generalizable Avatars. NeurIPS Track on Datasets and Benchmarks (2024)
42. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing scenes as neural radiance fields for view synthesis. In: ECCV. pp. 405–421 (2020)
43. Moon, G., Shiratori, T., Saito, S.: Expressive whole-body 3d gaussian avatar. In: ECCV (2024)
44. Moreau, A., Song, J., Dharmo, H., Shaw, R., Zhou, Y., Pérez-Pellitero, E.: Human gaussian splatting: Real-time rendering of animatable avatars. In: CVPR (2024)
45. Noguchi, A., Sun, X., Lin, S., Harada, T.: Neural articulated radiance field. In: International Conference on Computer Vision (2021)

46. Pang, H., Zhu, H., Kortylewski, A., Theobalt, C., Habermann, M.: Ash: Animatable gaussian splats for efficient and photoreal human rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1165–1175 (June 2024)
47. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3D hands, face, and body from a single image. In: Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 10975–10985 (2019)
48. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: CVPR (2019)
49. Peng, S., Dong, J., Wang, Q., Zhang, S., Shuai, Q., Zhou, X., Bao, H.: Animatable neural radiance fields for modeling dynamic human bodies. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14314–14323 (2021)
50. Peng, S., Zhang, Y., Xu, Y., Wang, Q., Shuai, Q., Bao, H., Zhou, X.: Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9054–9063 (2021)
51. Qiu, L., Gu, X., Li, P., Zuo, Q., Shen, W., Zhang, J., Qiu, K., Yuan, W., Chen, G., Dong, Z., et al.: Lhm: Large animatable human reconstruction model for single image to 3d in seconds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14184–14194 (2025)
52. Qiu, L., Li, P., Zuo, Q., Gu, X., Dong, Y., Yuan, W., Zhu, S., Han, X., Chen, G., Dong, Z.: Pf-lhm: 3d animatable avatar reconstruction from pose-free articulated human images. arXiv preprint arXiv:2506.13766 (2025)
53. Qiu, L., Zhu, S., Zuo, Q., Gu, X., Dong, Y., Zhang, J., Xu, C., Li, Z., Yuan, W., Bo, L., et al.: Anigs: Animatable gaussian avatar from a single image with inconsistent gaussian reconstruction. In: CVPR (2025)
54. Remelli, E., Bagautdinov, T., Saito, S., Wu, C., Simon, T., Wei, S.E., Guo, K., Cao, Z., Prada, F., Saragih, J., et al.: Drivable volumetric avatars using texel-aligned features. In: ACM SIGGRAPH 2022 Conference Proceedings. pp. 1–9 (2022)
55. Saito, S., Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., Li, H.: Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 2304–2314 (2019)
56. Saito, S., Schwartz, G., Simon, T., Li, J., Nam, G.: Relightable gaussian codec avatars. In: CVPR (2024)
57. Saito, S., Simon, T., Saragih, J., Joo, H.: Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 84–93 (2020)
58. Shao, R., Pang, Y., Zheng, Z., Sun, J., Liu, Y.: Human4dit: 360-degree human video generation with 4d diffusion transformer. TOG **43**(6) (2024)
59. Shen, K., Guo, C., Kaufmann, M., Zarate, J., Valentin, J., Song, J., Hilliges, O.: X-avatar: Expressive human avatars. In: Computer Vision and Pattern Recognition (CVPR) (2023)
60. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., Massa, F., Haziza, D., Wehrstedt, L., Wang, J., Darcet, T., Moutakanni, T., Sentana, L., Roberts, C., Vedaldi, A., Tolan, J., Brandt, J., Couprie, C., Mairal, J., Jégou, H., Labatut, P., Bojanowski, P.: DINOv3 (2025), <https://arxiv.org/abs/2508.10104>

61. Wang, S., Schwarz, K., Geiger, A., Tang, S.: Arah: Animatable volume rendering of articulated human sdf. In: European Conference on Computer Vision (ECCV) (2022)
62. Wang, T., Li, L., Lin, K., Zhai, Y., Lin, C.C., Yang, Z., Zhang, H., Liu, Z., Wang, L.: Disco: Disentangled control for realistic human dance generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9326–9336 (2024)
63. Weng, C.Y., Curless, B., Srinivasan, P.P., Barron, J.T., Kemelmacher-Shlizerman, I.: Humannerf: Free-viewpoint rendering of moving people from monocular video. In: CVPR (2022)
64. Xiu, Y., Yang, J., Cao, X., Tzionas, D., Black, M.J.: Econ: Explicit clothed humans optimized via normal integration. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 512–523 (2023)
65. Xiu, Y., Yang, J., Tzionas, D., Black, M.J.: Icon: Implicit clothed humans obtained from normals. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13286–13296. IEEE (2022)
66. Xu, T., Fujita, Y., Matsumoto, E.: Surface-aligned neural radiance fields for controllable 3d human synthesis. In: CVPR (2022)
67. Xu, Z., Zhang, J., Liew, J.H., Yan, H., Liu, J.W., Zhang, C., Feng, J., Shou, M.Z.: Magicanimate: Temporally consistent human image animation using diffusion model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1481–1490 (2024)
68. Yang, X., Kukreja, D., Pinkus, D., Sagar, A., Fan, T., Park, J., Shin, S., Cao, J., Liu, J., Ugrinovic, N., Feiszli, M., Malik, J., Dollar, P., Kitani, K.: Sam 3d body: Robust full-body human mesh recovery. arXiv preprint arXiv:2602.15989 (2026)
69. Yariv, L., Gu, J., Kasten, Y., Lipman, Y.: Volume rendering of neural implicit surfaces. In: Advances in Neural Information Processing Systems (2021)
70. Yu, T., Zheng, Z., Guo, K., Liu, P., Dai, Q., Liu, Y.: Function4d: Real-time human volumetric capture from very sparse consumer rgb-d sensors. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (June 2021)
71. Yu, Z., Li, Z., Bao, H., Yang, C., Zhou, X.: Humanram: Feed-forward human reconstruction and animation model using transformers. In: Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers. pp. 1–13 (2025)
72. Zhang, Y., Gu, J., Wang, L.W., Wang, H., Cheng, J., Zhu, Y., Zou, F.: Mimic-motion: High-quality human motion video generation with confidence-aware pose guidance. arXiv preprint arXiv:2406.19680 (2024)
73. Zhao, F., Yang, W., Zhang, J., Lin, P., Zhang, Y., Yu, J., Xu, L.: Humannerf: Efficiently generated human radiance field from sparse inputs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7743–7753 (June 2022)
74. Zhao, J., Zhang, H.: Thin-plate spline motion model for image animation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3657–3666 (2022)
75. Zheng, R., Li, P., Wang, H., Yu, T.: Learning visibility field for detailed 3d human reconstruction and relighting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 216–226 (2023)
76. Zheng, S., Zhou, B., Shao, R., Liu, B., Zhang, S., Nie, L., Liu, Y.: Gps-gaussian: Generalizable pixel-wise 3d gaussian splatting for real-time human novel view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)

77. Zheng, Z., Yu, T., Liu, Y., Dai, Q.: Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *IEEE transactions on pattern analysis and machine intelligence* **44**(6), 3170–3184 (2021)
78. Zheng, Z., Zhao, X., Zhang, H., Liu, B., Liu, Y.: Avatarrex: Real-time expressive full-body avatars. *ACM Transactions on Graphics (TOG)* **42**(4) (2023)
79. Zhu, H., Zhan, F., Theobalt, C., Habermann, M.: Trihuman: A real-time and controllable tri-plane representation for detailed human geometry and appearance synthesis. *ACM Trans. Graph.* (Sep 2024). <https://doi.org/10.1145/3697140>, <https://doi.org/10.1145/3697140>
80. Zhu, S., Chen, J.L., Dai, Z., Xu, Y., Cao, X., Yao, Y., Zhu, H., Zhu, S.: Champ: Controllable and consistent human image animation with 3d parametric guidance. *arXiv preprint arXiv:2403.14781* (2024)
81. Zhuang, Y., Lv, J., Wen, H., Shuai, Q., Zeng, A., Zhu, H., Chen, S., Yang, Y., Cao, X., Liu, W.: Idol: Instant photorealistic 3d human creation from a single image. *arXiv preprint arXiv:2412.14963* (2024)
82. Zielonka, W., Bagautdinov, T., Saito, S., Zollhöfer, M., Thies, J., Romero, J.: Drivable 3d gaussian avatars (2023), <https://arxiv.org/abs/2311.08581>