

AHOY! Animatable Humans under Occlusion from YouTube Videos with Gaussian Splatting and Video Diffusion Priors

Aymen Mir¹, Riza Alp Guler³, Xiangjun Tang⁴,
Peter Wonka^{4†}, and Gerard Pons-Moll^{1,2†}

¹ Tübingen AI Center, University of Tübingen, Germany

² Max Planck Institute for Informatics, Saarland Informatics Campus, Germany

³ Imperial College London, UK

⁴ King Abdullah University of Science and Technology (KAUST), Saudi Arabia



Fig. 1: Given a monocular YouTube video with occlusion, AHOY reconstructs a complete, animatable 3D human avatar using video diffusion priors and 3D Gaussian Splatting, enabling photorealistic human animation within 3D scenes.

Abstract. We present AHOY, a method for reconstructing complete, animatable 3D Gaussian avatars from in-the-wild monocular video despite heavy occlusion. Existing methods assume unoccluded input: a fully visible subject, often in a canonical pose. This excludes the vast majority of real-world footage where people are routinely occluded by furniture, objects, or other people. Reconstructing from such footage poses fundamental challenges: large body regions may never be observed, and multi-view supervision per pose is unavailable. We address these challenges with four contributions: (i) a hallucination-as-supervision pipeline that uses identity-finetuned diffusion models to generate dense supervision

[†] Equal advising.

for previously unobserved body regions; (ii) a two-stage canonical-to-pose-dependent architecture that bootstraps from sparse observations to full pose-dependent Gaussian maps; (iii) a map-pose/LBS-pose decoupling that absorbs multi-view inconsistencies from the generated data; (iv) a head/body split supervision strategy that preserves facial identity. We evaluate on YouTube videos and on multi-view capture data with significant occlusion and demonstrate state-of-the-art reconstruction quality. We also demonstrate that the resulting avatars are robust enough to be animated with novel poses and composited into 3DGS scenes captured using cell-phone video. Our project page is available at <https://miraymen.github.io/ahoy/>.

1 Introduction

Reconstructing photorealistic, animatable 3D humans from monocular input is a long-standing goal in computer vision and graphics, with applications spanning gaming, virtual reality, and film production. Recent methods based on NeRFs [37, 103, 131] and 3DGS [26, 48, 60, 90, 105, 106, 156, 157] have made remarkable progress, producing high-fidelity avatars that can be re-posed in real time. Yet these methods assume unoccluded monocular input, or a fully visible subject, often in a canonical pose.

These assumptions exclude the vast majority of real-world footage. YouTube alone hosts billions of videos of people in everyday activities (cooking, exercising, conversing) where the body is occluded by furniture, other people, or the scene itself. The viewpoints are limited, and large portions of the body may never be observed or cut off in the camera shot. If we could reconstruct complete, animatable avatars from such videos, we would unlock an essentially unlimited source of diverse human appearance (every body type and clothing style) without controlled capture, multi-view rigs, or cooperation from the subject.

In this work, we present **AHOY** (**A**nimatable **H**umans under **O**cclusion from **Y**ouTube **V**ideos), a method that reconstructs complete, animatable 3D Gaussian avatars from in-the-wild monocular video despite heavy occlusion, and animates them within photorealistic 3D scenes (Fig. 1). Achieving this requires overcoming several challenges that do not arise when the subject is fully visible.

Bootstrapping a complete avatar from partial observations. With large body regions never observed, there is no direct supervision for a significant portion of the avatar. We address this in three stages. First, we build a coarse 3DGS avatar from sparse multi-view images generated from the partially observed subject (Sec. 3.1). Second, we render this coarse avatar in structured motion sequences and refine the imperfect renderings via an identity-finetuned video diffusion model and diffusion inversion producing hallucinated supervision videos that cover previously unobserved body surfaces from multiple viewpoints and poses (Sec. 3.2). Third, we train a full avatar with pose-dependent deformation using these hallucinated videos as supervision (Sec. 3.3).

Learning pose-dependent appearance without multi-view capture. High-fidelity avatars require pose-dependent deformation to capture how clothing

changes with body configuration. In the first stage, no multi-view supervision for pose-dependent deformation exists. We therefore learn an avatar that does not model pose-dependent deformations. We adapt the 2D Gaussian map design from Animatable Gaussians [71] to learn a canonical pose-independent Gaussian map that captures the texture and structure of the human. We also adopt the use of per-identity networks that use 2D maps to represent both input pose and output 3DGS properties. To obtain the pose-dependent multi-view supervision missing from the original video, we design the structured motion sequences of the second stage (full-body turns paired with target poses such as sitting and reaching) to provide both pose-dependent deformations and in-effect 360° multi-view coverage per pose. In the third stage, these hallucinated videos supply the dense multi-view supervision needed to upgrade from canonical to pose-dependent Gaussian maps that capture high-fidelity appearance changes across body configurations. This canonical-to-pose-dependent progression is the architectural insight that makes pose-dependent learning possible (Secs. 3.2.2 and 3.3.1).

Absorbing multi-view inconsistencies from diffusion-generated data. Each hallucinated video is independently sampled from the diffusion model and is therefore not geometrically consistent across viewpoints and frames. Naively training on these videos produces artifacts. To address this, we decouple per-frame variations by introducing two distinct poses: a ‘map pose’ and an ‘LBS pose.’ Specifically, the map pose, which is taken as input for our per-identity networks, is shared across frames with similar poses to ensure consistent prediction of Gaussian attributes for a given pose. In contrast, the LBS pose, which deforms the output Gaussians from our per-identity networks, is optimized per-frame to absorb the geometric inconsistencies inherent in the video diffusion model (Sec. 3.3.2).

Preserving facial identity under unreliable diffusion outputs. Video diffusion models generate identity-inconsistent faces. We partition the avatar into head and body regions via SMPL-to-FLAME correspondence, mask head pixels from the diffusion-based body supervision, and instead supervise the head through a dedicated multi-view face diffusion model that preserves identity (Sec. 3.3.3).

We evaluate AHOY on YouTube videos featuring humans under significant occlusion and on multi-view capture data with comparable occlusion, and demonstrate state-of-the-art reconstruction quality. The resulting avatars are robust enough to be animated with novel poses and composited into 3DGS scenes reconstructed from casual cell-phone video, yielding an end-to-end pipeline from commodity RGB input to photorealistic human–scene rendering.

In summary, our contributions are:

- 1) A method for reconstructing complete, animatable 3DGS avatars from in-the-wild monocular video under heavy occlusion, unlocking YouTube-scale footage as a source of 3D human assets.
- 2) A hallucination-as-supervision approach using identity-finetuned video diffusion and diffusion inversion.
- 3) A two-stage canonical-to-pose-dependent architecture that allows for learning high-fidelity pose-dependent deformations in 3DGS avatars.

- 4) A map-pose/LBS-pose decoupling that enables pose-dependent Gaussian maps for high-fidelity avatar learning from multi-view-inconsistent generated data.
- 5) A head/body split supervision for facial identity preservation.

2 Related Work

Neural Rendering. NeRF [82] and its extensions [9, 10, 91, 120, 135] enable high-quality novel-view synthesis but remain computationally expensive. 3DGS [58] addresses this with explicit Gaussian primitives [63, 132] and has been extended to dynamic scenes [64, 69, 76, 85, 116, 134], SLAM [57], mesh extraction [30, 46], sparse-view reconstruction [81], and avatar-scene composition [84]. Beyond efficiency, the explicit nature of 3DGS makes it easy to edit, deform, and composite into new scenes, properties we rely on when animating reconstructed avatars and placing them in captured environments.

Human Reconstruction and Neural Rendering. Dense body representations such as DensePose [31, 93, 94] provide pixel-level surface correspondences for human understanding. Mesh-based approaches [6, 7, 13, 36, 56, 74, 86, 100] and implicit functions [8, 23, 24, 38, 49, 80, 83, 99, 113] recover 3D humans but lack photorealism or reposability. NeRF-based methods [32, 37, 52, 67, 73, 103, 131, 138, 154] produce photorealistic renderings from video, while 3DGS-based avatars [5, 25, 33, 44, 53, 55, 60, 65, 70, 71, 89, 90, 97, 105, 140, 150, 157] offer real-time rendering. None of these methods work with YouTube videos with significant occlusion.

Humans from Monocular Input. Early implicit methods [47, 112, 137, 148, 151] and optimization-based approaches [45, 51] reconstruct humans from single images or monocular video. Feed-forward models [41, 106, 107, 118, 156] predict animatable avatars in seconds. Many methods leverage diffusion to synthesize multi-view images as an intermediate step [15, 22, 26, 39, 48, 66, 75, 96, 102, 108, 121], while others extend to 4D or unconstrained settings [72, 98, 115, 119, 144]. All these methods assume a fully visible subject as input. Text-to-3D methods [19, 61, 104] generate humans from language prompts, but cannot reconstruct from visual observations. In contrast, our method recovers an animatable avatar of a particular person from occluded in-the-wild YouTube videos.

Video Diffusion Models. Image and video diffusion models [12, 14, 16–18, 28, 40, 50, 62, 95, 101, 109, 111, 123, 133, 145] have driven rapid progress in visual generation. A large body of work specializes these for human-centric synthesis, spanning pose-driven character animation [21, 29, 35, 42, 43, 77, 79, 122, 126, 127, 141, 142, 147, 153, 155], portrait animation [34, 78, 130, 136, 139, 143], and video editing [54], building on earlier warping-based methods [117, 149]. LoRA-based personalization enables identity-specific finetuning of video diffusion models [1, 2, 27], while inversion techniques [3, 4, 88, 110, 124] allow embedding real videos into the model’s latent space for faithful reconstruction and editing. Video diffusion has also been used as supervision for 3DGS under sparse-view setups [92, 152]. Unlike these works, we leverage video diffusion as a prior for 3DGS-based avatar reconstruction. Rather than using the diffusion model only to render novel views, we synthesize




Fig. 2: Method overview. **Block 1** (Sec. 3.1): We map observed textures from partially occluded video onto a canonical pose via DensePose UV correspondences, inpaint missing regions with FLUX, and generate multi-view canonical-pose images with multi-view diffusion to supervise a coarse 3DGS avatar using canonical Gaussian maps. **Block 2** (Sec. 3.2): We finetune a video diffusion model (Wan 2.2) with LoRA to capture the subject’s identity, render the coarse avatar under structured motion sequences, and refine these renderings via RF-Inversion through the identity-finetuned latent space to produce hallucinated supervision videos. **Block 3** (Sec. 3.3): The hallucinated videos supervise a full avatar with pose-dependent Gaussian maps, where per-frame poses and cameras are jointly optimized to absorb multi-view inconsistencies; a separate FLAME-based head path preserves facial identity. **Block 4** (Sec. 3.4): At inference, the avatar is driven by novel poses and composited into 3DGS scenes.

plausible appearance for body regions the input video never reveals and turn these generations into dense supervision for fitting the avatar.

3 Method

Given a monocular video in which a human subject is partially occluded, our goal is to reconstruct a complete, animatable 3D Gaussian avatar that can be driven by arbitrary poses. Our pipeline (Fig. 2) proceeds in four stages: (1) coarse avatar construction from the partially observed subject (Sec. 3.1); (2) hallucinated supervision generation via identity-finetuned video diffusion (Sec. 3.2); (3) full avatar learning with pose-dependent Gaussian maps trained on this supervision (Sec. 3.3); and (4) animation and scene composition (Sec. 3.4).

3.1 Coarse Avatar Construction

We first construct a coarse 3DGS avatar from the partially observed subject (Block 1 in Fig. 2).

3.1.1 Texture mapping and inpainting. Given a monocular YouTube video with frames $\{\mathbf{I}_t\}_{t=1}^T$, we first estimate per-frame SMPL parameters $\{(\boldsymbol{\theta}_t, \boldsymbol{\beta})\}$ using NLF [114], and obtain visibility masks $\{\mathbf{M}_t\}$ via SAM 3 [20].

Using DensePose [31], we assign each visible pixel a UV coordinate and accumulate them frame by frame into a canonical texture atlas with a first-write-wins rule to avoid temporal blending artifacts. We then define a fixed front-facing canonical pose $\boldsymbol{\theta}_*$ and warp the atlas back to this pose via inverse DensePose correspondence [93], producing a partial canonical-pose image $\mathbf{I}_*^{\text{part}}$. Regions that were never observed remain empty. A pretrained FLUX [28] inpainting model fills in these missing regions, yielding a complete canonical-pose image $\mathbf{I}_*$ that serves as the basis for multi-view generation. (Full details in the supplementary.)

3.1.2 Multi-view generation and coarse 3DGS avatar. The inpainted canonical image $\mathbf{I}_*$ provides only a single viewpoint. To obtain broader coverage, we generate $V - 1 = 3$ additional canonical-pose views (back, left, right) using a pretrained multi-view diffusion model [22], discarding its 3D reconstruction branch. Together with the frontal view $\mathbf{I}_*^1 = \mathbf{I}_*$, this yields $V = 4$ canonical-pose images $\{\mathbf{I}_*^v\}_{v=1}^V$ at known orthogonal cameras $\{\boldsymbol{\gamma}_v\}_{v=1}^V$.

Canonical Gaussian maps. Following [71, 85], we densely sample N_H points on the SMPL mesh in canonical pose $\boldsymbol{\theta}_*$ to define a canonical Gaussian template $\mathcal{G}_*^H$. These are projected from front and back views to create canonical position maps with one-to-one pixel-to-Gaussian correspondence. Unlike the pose-dependent human maps in [85], we use time-invariant canonical maps, as our monocular occluded setting lacks the multi-view coverage needed for pose-conditioned offsets.

StyleUNet prediction. A StyleUNet [125] S processes the canonical maps and predicts per-pixel offsets for each Gaussian property from the template, yielding the canonical Gaussian set $\mathcal{G}^H = S(\boldsymbol{\theta}_*)$. Additionally, S predicts per-Gaussian skinning weights $\mathbf{w}_{k,j}^H$ for $j = 1, \dots, J$ joints.

LBS deformation. For each frame t with estimated pose $\boldsymbol{\theta}_t$, we deform the canonical Gaussians into posed space via Linear Blend Skinning (LBS) using the predicted skinning weights $\mathbf{w}_{k,j}^H$. We denote the complete deformation as $\mathcal{T}(\mathcal{G}^H; \boldsymbol{\theta}_t)$; full details are provided in the supplementary material. Colors and opacities remain as predicted by the StyleUNet.

Training objective. The coarse avatar is supervised by two sources: (i) the V canonical-pose images at known cameras, and (ii) the T observed video frames at estimated poses and observed camera $\boldsymbol{\gamma}^{\text{obs}}$ (the camera associated with the NLF reference frame). Let $\hat{\mathbf{I}}_t = \mathcal{R}(\mathcal{T}(\mathcal{G}^H; \boldsymbol{\theta}_t); \boldsymbol{\gamma}^{\text{obs}})$ denote the rendering at frame t . The training loss is:

$$\mathcal{L}_{\text{coarse}} = \sum_{v=1}^V \|\mathbf{I}_*^v - \mathcal{R}(\mathcal{G}^H; \boldsymbol{\gamma}_v)\|_1 + \sum_{t=1}^T \|(\mathbf{I}_t - \hat{\mathbf{I}}_t) \odot \mathbf{M}_t\|_1, \quad (1)$$

where $\odot$ denotes element-wise multiplication and $\mathbf{M}_t$ masks out occluded regions. We additionally employ a perceptual loss [146] with the same masking (see

supplementary). This yields a coarse avatar faithful to the visible portions of the subject but with artifacts in unobserved regions.

3.2 Hallucinated Supervision Generation

The coarse avatar (Sec. 3.1) is supervised by only $V=4$ canonical-pose images and the occluded input frames. This limited coverage means that unobserved regions (the back of the body, inner arms, or any area persistently hidden by the occluder) remain poorly reconstructed, and the avatar cannot generalize to novel poses and camera angles beyond those seen during training. We now proceed to generate hallucinated videos for further supervision (Block 2 in Fig. 2).

3.2.1 Identity-specific video diffusion finetuning. A generic video diffusion model lacks knowledge of the subject’s specific appearance, so we finetune one on the input video to build an identity-specific prior. We adapt the identity encoding stage of Dynamic Concepts [2] to Wan 2.2 [123]. Let $\mathbf{v}_\phi$ denote the velocity field of the pretrained rectified flow model. We learn a LoRA update $\mathbf{AB}$ together with a learnable text token $[v]$ that binds to the target identity in the text encoder’s embedding space. Following [2], we treat the input video as an *unordered* set of frames (removing temporal information) so that the LoRA captures appearance and identity without entangling motion dynamics. The training objective is the flow matching loss:

$$\mathcal{L}_{\text{id}} = \mathbb{E}_{\tau, \mathbf{z}_0} \left[\left\| (\mathbf{z}_1 - \mathbf{z}_0) - \mathbf{v}_\phi(\mathbf{z}_\tau, \tau; \text{“a } [v] \text{ person”}) \right\|^2 \right], \quad (2)$$

where $\mathbf{z}_0 = \text{Enc}(\mathbf{I}_t)$ is the VAE-encoded latent, $\mathbf{z}_1 \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$ is noise ($\mathbb{I}$ denotes the identity matrix), $\mathbf{z}_\tau = (1-\tau)\mathbf{z}_0 + \tau\mathbf{z}_1$ is the interpolated latent at flow time $\tau \in [0, 1]$, and the velocity field is parameterized by the LoRA update $\mathbf{W}' = \mathbf{W} + \mathbf{AB}$ applied to the pretrained weights $\mathbf{W}$. After training, the finetuned model $\mathbf{v}_\phi$ generates videos that faithfully preserve the subject’s identity, appearance, and clothing, which we exploit in the next stages to hallucinate occluded regions.

3.2.2 Supervision video generation. We now use the identity-finetuned model to generate supervision videos that cover the missing body regions. Rather than generating arbitrary motions, we curate a set of N motion sequences designed to systematically expose all body surfaces to the camera while providing the multi-view coverage needed to learn pose-dependent Gaussian maps (Sec. 3.3).

Structured Motion (turn). Each turn-based sequence pairs a target body action (*e.g.*, arms lifting or both arms raised) with a full 360° body rotation. The key observation is that all frames within such a sequence can be treated as views of the *same* pose from different angles: the body action remains approximately constant while the rotation sweeps the viewpoint around the person. A turn with T' frames thus provides in-effect T' -view supervision of a single pose, enabling the model to learn how the Gaussian representation should look from every direction under that specific body configuration.

Structured Motion (static). Certain poses, such as sitting, are incompatible with a full-body turn. For these, we generate multiple videos of the same sitting pose from different camera angles. We then synchronize frames across these videos where the pose matches and supervise each pose with the images from all corresponding videos, providing cross-video multi-view supervision analogous to the within-video coverage of the turn sequences.

Coarse rendering. For each motion sequence n , we animate the coarse avatar with the corresponding poses $\{\theta'_t\}_{t=1}^{T'}$ and render from the observed camera γ^{obs} :

$$\hat{\mathbf{V}}_n = \{\mathcal{R}(\mathcal{T}(\mathcal{G}^H; \theta'_t); \gamma^{\text{obs}})\}_{t=1}^{T'}. \quad (3)$$

These renderings capture the correct body pose and silhouette but contain artifacts and missing textures in regions that were occluded in the input video. We refine them through the diffusion latent space below.

3.2.3 Hallucination via RF-Inversion. The coarse renderings $\hat{\mathbf{V}}_n$ from above preserve the target body layout but suffer from texture artifacts in unobserved regions. We refine them by passing each video through the latent space of the identity-finetuned model $\mathbf{v}_\phi$, leveraging RF-Inversion [110] to replace artifacts with identity-consistent appearance.

Let Enc and Dec denote the VAE encoder and decoder, and let $\mathbf{z}_0 = \text{Enc}(\hat{\mathbf{V}}_n)$ be the encoded coarse video. The forward rectified flow ODE maps this to a noisy latent $\mathbf{z}_1$, and reversing the flow with the identity-finetuned model decodes a hallucinated video:

$$\mathbf{z}_1 = \mathbf{z}_0 + \int_0^1 \mathbf{v}_\phi(\mathbf{z}_\tau, \tau) d\tau, \quad \tilde{\mathbf{z}}_0 = \mathbf{z}_1 + \int_1^0 \mathbf{v}_\phi(\mathbf{z}_\tau, \tau) d\tau, \quad \tilde{\mathbf{V}}_n = \text{Dec}(\tilde{\mathbf{z}}_0). \quad (4)$$

The *forward pass encodes the spatial structure* (body pose and silhouette) of the coarse rendering into $\mathbf{z}_1$, while the *reverse pass preserves this layout but steers the appearance* toward the subject’s identity via the LoRA, filling in previously missing regions with plausible textures. Applying this process to all N motion sequences yields the hallucinated video set $\{\tilde{\mathbf{V}}_n\}_{n=1}^N$, which provides multi-view body coverage that serves as supervision for the full avatar (Sec. 3.3).

3.3 Learning the Full Avatar under Multi-View Inconsistencies

The hallucinated videos $\{\tilde{\mathbf{V}}_n\}_{n=1}^N$ from Sec. 3.2 provide dense body coverage but are not perfectly multi-view consistent, as each video is independently sampled from the diffusion model. We address this with (i) a switch from canonical to pose-dependent Gaussian maps, enabled by the in-effect multi-view supervision from the hallucinated sequences, (ii) joint optimization of Gaussian, pose, and camera parameters to absorb inconsistencies, and (iii) a separate FLAME-based head supervision path that preserves facial identity.

3.3.1 Pose-dependent Gaussian maps. The coarse stage (Sec. 3.1) used time-invariant canonical maps because multi-view supervision for multiple poses was unavailable. The hallucinated videos now provide this coverage: in a turn-based sequence, all T' frames observe approximately the same body pose from different viewpoints, yielding in-effect T' -view supervision for that pose. For static-pose sequences (*e.g.*, sitting), we synchronize frames across videos rendered from different angles where the same pose occurs, providing cross-video multi-view supervision. This dense coverage enables us to upgrade to pose-dependent Gaussian maps, following [71, 85]: we project posed SMPL vertices onto front and back UV maps for each frame, and a re-initialized StyleUNet S takes these pose-dependent maps as input and predicts per-pixel Gaussian property offsets from the canonical template, yielding $\mathcal{G}_n^H = S(\theta_n^{\text{map}})$.

A crucial distinction arises between two roles of the pose: the *map pose* θ_n^{map} , which is the input to the StyleUNet and remains *fixed* across all frames of sequence n , and the *LBS pose* $\theta_{n,t}^{\text{lbs}}$, which deforms the predicted Gaussians into posed space and varies per frame. For a turn-based sequence, the map pose encodes the target body configuration (*e.g.*, arms raised) so that S produces a single set of Gaussians representing that pose; the per-frame LBS pose then handles the body rotation, mapping these Gaussians to the correct 3D orientation for each viewpoint.

3.3.2 Joint optimization with pose and camera correction. Since the hallucinated videos are occlusion-free, we estimate SMPL poses with NLF [114] and assume cameras in the NLF reference frame. These estimates initialize both the map pose θ_n^{map} (used to create pose-dependent input maps for S) and the per-frame LBS pose $\theta_{n,t}^{\text{lbs}}$. We optimize S jointly with per-frame LBS pose corrections $\{\Delta\theta_{n,t}\}$ and camera corrections $\{\Delta\gamma_{n,t}\}$ for each hallucinated frame. Let $\tilde{\mathbf{I}}_{n,t}$ denote frame t of the n -th hallucinated video. We exclude head pixels using a segmentation mask $\mathbf{M}_{n,t}^{\text{head}}$ obtained from Sapiens [59] and SAM 3 [20], since diffusion-generated faces are unreliable for identity preservation. The body supervision loss is:

$$\mathcal{L}_{\text{body}} = \sum_{n,t} \|(\tilde{\mathbf{I}}_{n,t} - \hat{\mathbf{I}}_{n,t}) \odot (1 - \mathbf{M}_{n,t}^{\text{head}})\|_1, \quad (5)$$

where $\hat{\mathbf{I}}_{n,t} = \mathcal{R}(\mathcal{T}(\mathcal{G}_n^H; \theta_{n,t}^{\text{lbs}} + \Delta\theta_{n,t}); \gamma^{\text{obs}} + \Delta\gamma_{n,t})$ is the rendered image with corrected LBS pose and camera. The map pose θ_n^{map} is shared across all frames of sequence n , so every hallucinated view supervises the same pose-dependent Gaussian map. The LBS pose corrections $\Delta\theta_{n,t}$ absorb per-frame multi-view inconsistencies from the diffusion-generated data, introducing a deliberate divergence between the map pose and the LBS pose that allows the canonical Gaussians to remain consistent while the deformation adapts to each frame.

We retain observed-frame supervision $\mathcal{L}_{\text{obs}} = \sum_t \|(\mathbf{I}_t - \hat{\mathbf{I}}_t^{\text{obs}}) \odot \mathbf{M}_t\|_1$ from the input video, analogous to the second term of Eq. 1, where $\hat{\mathbf{I}}_t^{\text{obs}}$ is rendered with the original estimated poses and cameras.

3.3.3 Head identity preservation via FLAME. Head regions in the Wan-generated videos are unreliable for facial identity. Instead, we leverage the multi-view face diffusion model from CAP4D [121] to generate identity-consistent head images $\{\mathbf{I}_p^{\text{face}}\}_{p=1}^P$ from P viewpoints with known cameras $\{\gamma_p^{\text{face}}\}$.

Using the SMPL-to-FLAME [68] vertex correspondence, we partition the canonical Gaussians into head and body regions:

$$\mathcal{G}^{\text{head}} = \{(\mu_k, \dots) \in \mathcal{G}^H \mid k \in \mathcal{H}_{\text{head}}\}, \quad (6)$$

where $\mathcal{H}_{\text{head}} \subset \{1, \dots, N_H\}$ is the set of Gaussian indices corresponding to the SMPL head region. For head supervision, we deform only these Gaussians using FLAME expression parameters ψ_p :

$$\mathcal{L}_{\text{head}} = \sum_{p=1}^P \|\mathbf{I}_p^{\text{face}} - \mathcal{R}(\mathcal{T}_{\text{FLAME}}(\mathcal{G}^{\text{head}}; \psi_p); \gamma_p^{\text{face}})\|_1, \quad (7)$$

where $\mathcal{T}_{\text{FLAME}}$ applies FLAME-based deformation to the head Gaussians while leaving their canonical positions shared with the full body representation.

3.3.4 Total objective. The full avatar is trained with the combined loss:

$$\mathcal{L} = \lambda_{\text{body}} \mathcal{L}_{\text{body}} + \lambda_{\text{obs}} \mathcal{L}_{\text{obs}} + \lambda_{\text{head}} \mathcal{L}_{\text{head}} + \lambda_{\text{per}} \mathcal{L}_{\text{per}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}, \quad (8)$$

where $\mathcal{L}_{\text{per}}$ is a perceptual loss [146] and $\mathcal{L}_{\text{reg}}$ regularizes the skinning weights toward SMPL-derived values. The weighting coefficients λ_{body} , λ_{obs} , λ_{head} , λ_{per} , and λ_{reg} balance the respective terms. Because the hallucinated sequences span a wide range of body configurations (standing, sitting, arms raised, turning), the training pose distribution is broad enough that the learned pose-dependent Gaussian maps generalize well to novel driving poses at inference time.

3.4 Animation in 3D Scenes

The preceding stages produce a complete, animatable 3DGS avatar that can be driven by arbitrary SMPL poses via LBS with pose-dependent Gaussian maps (Sec. 3.3.1). Following AHA! [87], we animate this avatar within reconstructed 3DGS scenes. To prevent out-of-distribution artifacts, we apply PCA-based pose projection at inference, following [71], constraining novel driving poses to the training distribution (details in the supplementary). Given a scene reconstructed as a static Gaussian set $\mathcal{G}^S$ from casual cell-phone video [58] and a driving motion of T'' frames $\{\theta_t\}_{t=1}^{T''}$, we deform the avatar at each timestep and render $\mathbf{I}_{\text{out}} = \mathcal{R}(\mathcal{T}(\mathcal{G}_t^H; \theta_t) \cup \mathcal{G}^S; \gamma)$, where $\mathcal{G}_t^H = S(\theta_t)$ uses the (PCA-projected) pose-dependent Gaussian maps, producing photorealistic novel-view videos of the reconstructed human animated within the target environment.




Fig. 3: Animation comparison (Zoom In). Top: canonical-pose input. Bottom (below line): occluded input. Ours produces higher-fidelity avatars in both settings.

4 Experiments

We evaluate AHOY along four axes: *animation on YouTube* (Sec. 4.2); *static reconstruction on YouTube* (Sec. 4.3); *multi-view evaluation on BEHAVE* with ground-truth side views measuring novel-view and novel-pose quality under occlusion (Sec. 4.4); and *component analysis* (Sec. 4.5). We further demonstrate animation within photorealistic 3DGS scenes (Sec. 4.6).

4.1 Experimental Setup

Datasets. We evaluate on two benchmarks that provide naturally occluded human subjects. **BEHAVE** [11] contains multi-view captures (4 Kinect cameras) of humans interacting with objects, producing significant body occlusion. We evaluate on 8 sequences. The multi-view ground truth enables rigorous quantitative evaluation. **YouTube videos.** We collect 50 videos of humans in diverse everyday activities with significant occlusion. For these sequences, ground-truth multi-view images are unavailable, so we compare against held-out frames from the original video. **Metrics.** We report PSNR, SSIM [129], and LPIPS [146].

Baselines. Our method is video-based, while all baselines are single-image methods. We consider recent animatable avatar methods with public code: LHM [106] and IDOL [156]. We evaluate baselines under two input settings: (1) receiving the original occluded frame (first video frame) directly, and (2) receiving the unoccluded canonical-pose image generated by our pipeline, giving them the best possible input. We additionally compare with static reconstruction methods PSHuman [66] and SyncHuman [22], which do not produce animatable




Fig. 4: Static reconstruction comparison (Zoom In). Top: canonical-pose input. Bottom: occluded input. Our animatable avatar is posed to match the target view.



Fig. 5: Reconstruction under difficult fully visible poses. Static methods struggle to reconstruct humans under difficult poses, while our method allows reposing of 3DGS avatars to match the target pose.

avatars. We note that AdaHuman [48], Dream-Lift-Animate [15], MoGA [26], and MEAT [128] are plausible baselines, but none have publicly available code.

4.2 Animation on YouTube

We compare with LHM [106] and IDOL [156] (the only baselines that produce animatable avatars) under two input settings (Tab. 2, Fig. 3). We animate reconstructed avatars to novel poses and compare against held-out video frames. **Canonical-pose input.** To provide a fairer comparison, we give baselines the unoccluded canonical-pose images generated by our pipeline (Sec. 3.1.1). Even with complete input, our full pipeline outperforms these methods (Fig. 3, top). **Occluded input.** When baselines receive the original occluded video frames, the realism of their avatars drops significantly (Fig. 3, bottom).

4.3 Static Reconstruction on YouTube

We compare with PSHuman [66], SyncHuman [22], LHM [106], and IDOL [156] on YouTube videos under two input settings (Tab. 1, Figs. 4–5). Each method reconstructs from its respective input, and all reconstructions are rendered from the input viewpoint and compared against the original video frame. **Occluded**

Table 1: Quantitative evaluation. Left: static reconstruction on YouTube videos under occluded and canonical-pose input. Right: evaluation on BEHAVE with novel-view and novel-pose settings.

Method	YouTube Static						BEHAVE					
	Occluded			Canonical			Novel View			Novel Pose		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
PSHuman [66]	19.47	0.811	0.197	21.88	0.858	0.145	18.47	0.812	0.198	<i>n/a</i>		
SyncHuman [22]	19.82	0.819	0.189	22.03	0.863	0.140	19.34	0.831	0.172	<i>n/a</i>		
LHM [106]	19.12	0.803	0.207	21.18	0.845	0.160	18.21	0.807	0.203	16.93	0.774	0.238
IDOL [156]	17.31	0.771	0.241	19.78	0.819	0.189	17.38	0.783	0.231	15.87	0.748	0.267
Ours	22.01	0.881	0.109	23.05	0.892	0.098	24.12	0.905	0.090	22.81	0.889	0.110

Table 2: Animation comparison on YouTube videos. We animate reconstructed avatars to novel poses and compare against held-out video frames.

Method	Occluded input			Canonical-pose input		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
LHM [106]	19.03	0.821	0.181	20.17	0.842	0.159
IDOL [156]	16.52	0.769	0.231	18.52	0.808	0.191
AHOY (Ours)	22.81	0.887	0.107	22.83	0.888	0.106

input. Under occlusion, all baselines fail to produce complete reconstructions, while our method generates a full avatar that can be accurately posed to match the target view. **Canonical-pose input.** To provide a fairer comparison, we give all baselines the unoccluded canonical-pose images generated by our pipeline. Even with complete input, our method outperforms all baselines. In Fig. 5, we demonstrate that static methods struggle to reconstruct humans under difficult poses, while our method allows reposing of avatars to match the target pose.

4.4 BEHAVE Evaluation

We evaluate all methods on BEHAVE (Tab. 1, Fig. 6). All methods receive the occluded front-camera view as input. For *novel-view* evaluation, we render each reconstruction from held-out side cameras and compare against ground-truth side views. For *novel-pose* evaluation, we animate the reconstructed avatars to unseen ground-truth SMPL-X poses and render from held-out side cameras. Our method significantly outperforms all baselines in both settings.

4.5 Ablation Study

We ablate each key component of our pipeline on BEHAVE to isolate its contribution, evaluating both novel-view and novel-pose settings as in Sec. 4.4 (Tab. 3, Fig. 7). **(A) Coarse avatar only.** We skip the hallucination and full avatar stages, using only the coarse canonical-map avatar (Sec. 3.1). This removes pose-dependent appearance and the diffusion-based supervision for occluded regions. **(B) Without RF-Inversion.** We skip RF-Inversion and supervise the avatar without any new hallucinated videos of the person in new poses; this



Fig. 6: Novel-view evaluation on BEHAVE. We render each reconstruction from held-out BEHAVE side cameras and compare against the ground-truth side views.

Table 3: Ablation study on BEHAVE. Each row removes one component from the full pipeline. *Novel View*: render from held-out side cameras. *Novel Pose*: animate to unseen ground-truth SMPL-X poses and render from held-out side cameras.

Configuration	Novel View			Novel Pose		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
AHOY (full)	24.12	0.905	0.090	22.81	0.889	0.110
(A) Coarse avatar only	16.10	0.753	0.258	14.30	0.718	0.301
(B) w/o RF-Inversion	21.10	0.862	0.143	19.40	0.838	0.168
(C) w/o map/LBS decoupling	22.60	0.886	0.118	20.60	0.856	0.152
(D) w/o head/body split	23.60	0.898	0.098	22.20	0.882	0.120

results in significant artifacts in unobserved regions (Sec. 3.2.3). **(C) Without map-pose/LBS-pose decoupling.** We do not use pose-dependent maps for StyleUNet training and do not optimize LBS deformation, removing the decoupling that absorbs multi-view inconsistencies (Sec. 3.3.2). **(D) Without head/body split.** We supervise the entire avatar, including the head, using the hallucinated Wan videos, without the separate head path (Sec. 3.3.3).

4.6 Animation in 3D Scenes

We demonstrate that the avatars reconstructed by AHOY are robust enough to be animated with novel poses and composited into 3DGS scenes reconstructed from casual cell-phone video (Fig. 8). Please see the supplementary for videos.

5 Conclusion

We have presented AHOY, a method that reconstructs complete, animatable 3D Gaussian avatars from in-the-wild monocular video despite heavy occlusion, unlocking YouTube-scale footage as a source of 3D human assets. Our key insight is to use an identity-finetuned video diffusion model as a bridge between imperfect coarse renderings and dense, realistic supervision: by embedding coarse avatar renderings into the model’s latent space via RF-Inversion and decoding them back, the identity prior fills in previously unobserved regions while preserving body layout. The resulting avatars are robust enough to be animated with novel poses and composited into 3DGS scenes reconstructed from casual cell-phone video. Our method inherits limitations from its diffusion backbone: the



Fig. 7: Visual ablations. Each pair shows the ablated result (w/o) and the full pipeline. (A) Without hallucination, the coarse-only avatar lacks detail in unobserved regions. (B) Without RF-Inversion, hallucinated supervision retains rendering artifacts. (C) Without map/LBS decoupling, multi-view inconsistencies cause blurring. (D) Without the head/body split, facial identity is lost.



Fig. 8: Animation in 3DGS scenes. Avatars from occluded YouTube video, animated and composited into 3DGS scenes from casual cell-phone video.

identity-finetuned model may hallucinate plausible but incorrect details for body regions that are never observed, and the quality of the final avatar is bounded by the fidelity of the video diffusion model. Additionally, our pipeline requires several stages of optimization and diffusion inference, making it substantially slower than feed-forward methods. Addressing these limitations suggests several promising directions: scaling to stronger video diffusion models as they improve and incorporating geometric consistency constraints into the hallucination stage. We hope that AHOY demonstrates that occlusion, long treated as a nuisance to be avoided, can be overcome, opening the door to learning from the vast diversity of human appearance captured in everyday video.

Acknowledgements. We gratefully acknowledge support from the hessian.AI Service Center (funded by the Federal Ministry of Research, Technology and Space, BMFTR, grant no. 16IS22091) and the hessian.AI Innovation Lab (funded by the Hessian Ministry for Digital Strategy and Innovation, grant no. S-DIW04/0013/003). We especially thank Patrick Blauth at the hessian.AI Service Center for technical assistance. This work was further supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A. Gerard Pons-Moll is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645. The research reported in this publication was also supported by funding from King Abdullah University of Science and Technology (KAUST) – Center of Excellence for Generative AI, under award number 5940 and a gift from Google.

References

1. Abdal, R., Patashnik, O., Deyneka, E., Chen, H., Siarohin, A., Tulyakov, S., Cohen-Or, D., Aberman, K.: Zero-shot dynamic concept personalization with grid-based LoRA. In: ACM SIGGRAPH Asia Conference Proceedings (2025)
2. Abdal, R., Patashnik, O., Skorokhodov, I., Menapace, W., Siarohin, A., Tulyakov, S., Cohen-Or, D., Aberman, K.: Dynamic concepts personalization from single videos. In: ACM SIGGRAPH Conference Proceedings (2025)
3. Abdal, R., Qin, Y., Wonka, P.: Image2stylegan: How to embed images into the stylegan latent space? In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
4. Abdal, R., Qin, Y., Wonka, P.: Image2stylegan++: How to edit the embedded images? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
5. Abdal, R., Yifan, W., Shi, Z., Xu, Y., Po, R., Kuang, Z., Chen, Q., Yeung, D.Y., Wetzstein, G.: Gaussian shell maps for efficient 3d human generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
6. Alldieck, T., Magnor, M., Bhatnagar, B.L., Theobalt, C., Pons-Moll, G.: Learning to reconstruct people in clothing from a single RGB camera. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
7. Alldieck, T., Magnor, M., Xu, W., Theobalt, C., Pons-Moll, G.: Video based reconstruction of 3d people models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
8. Alldieck, T., Xu, H., Sminchisescu, C.: imGHUM: Implicit generative models of 3d human shape and articulated pose. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
9. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
10. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-nerf: Anti-aliased grid-based neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
11. Bhatnagar, B.L., Xie, X., Petrov, I.A., Sminchisescu, C., Theobalt, C., Pons-Moll, G.: Behave: Dataset and method for tracking human object interactions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
12. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorber, D., Levi, Y., English, Z., Voleti, V., Letts, A., Jampani, V., Rombach, R.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)
13. Bogo, F., Kanazawa, A., Lassner, C., Gehler, P., Romero, J., Black, M.J.: Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In: European Conference on Computer Vision (ECCV) (2016)
14. Brooks, T., Peebles, B., Holmes, C., DePue, W., Guo, Y., Jing, L., Schnurr, D., Taylor, J., Luhman, T., Luhman, E., Ng, C., Wang, R., Ramesh, A.: Video generation models as world simulators. <https://openai.com/research/video-generation-models-as-world-simulators> (2024)

15. Buehler, M.C., Yuan, Y., Li, X., Huang, Y., Nagano, K., Iqbal, U.: Dream, lift, animate: From single images to animatable gaussian avatars. In: International Conference on 3D Vision (3DV) (2026)
16. ByteDance Seed Team: Seedance 1.0: Exploring the boundaries of video generation models. arXiv preprint arXiv:2506.09113 (2025)
17. ByteDance Seed Team: Seedance 1.5 pro: A native audio-visual joint generation foundation model. arXiv preprint arXiv:2512.13507 (2025)
18. ByteDance Seed Team: Seedance 2.0. https://seed.bytedance.com/en/seedance2_0 (2026)
19. Cao, Y., Cao, Y.P., Han, K., Shan, Y., Wong, K.Y.K.: DreamAvatar: Text-and-shape guided 3D human avatar generation via diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
20. Carion, N., et al.: Sam 3: Segment anything with concepts. arXiv preprint arXiv:2511.16719 (2025)
21. Chang, D., Shi, Y., Gao, Q., Xu, H., Fu, J., Song, G., Yan, Q., Zhu, Y., Yang, X., Soleymani, M.: MagicPose: Realistic human poses and facial expressions retargeting with identity-aware diffusion. In: International Conference on Machine Learning (ICML) (2024)
22. 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. In: Advances in Neural Information Processing Systems (NeurIPS) (2025)
23. Chen, X., Zheng, Y., Black, M.J., Hilliges, O., Geiger, A.: Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
24. Deng, B., Lewis, J.P., Jeruzalski, T., Pons-Moll, G., Hinton, G., Norouzi, M., Tagliasacchi, A.: Nasa neural articulated shape approximation. In: European Conference on Computer Vision (ECCV) (2020)
25. Dhamo, H., Nie, Y., Moreau, A., Song, J., Shaw, R., Zhou, Y., Pérez-Pellitero, E.: Headgas: Real-time animatable head avatars via 3d gaussian splatting. In: European Conference on Computer Vision (ECCV) (2024)
26. Dong, Z., Duan, L., Song, J., Black, M.J., Geiger, A.: MoGA: 3D generative avatar prior for monocular gaussian avatar reconstruction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2025)
27. Dravid, A., Gandelsman, Y., Wang, K.C., Abdal, R., Wetzstein, G., Efros, A.A., Aberman, K.: Interpreting the weight space of customized diffusion models. In: Advances in Neural Information Processing Systems (NeurIPS) (2024)
28. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., Podell, D., Dockhorn, T., English, Z., Rombach, R.: Scaling rectified flow transformers for high-resolution image synthesis. In: International Conference on Machine Learning (ICML) (2024)
29. Gan, Q., Ren, Y., Zhang, C., Ye, Z., Xie, P., Yin, X., Yuan, Z., Peng, B., Zhu, J.: HumanDiT: Pose-guided diffusion transformer for long-form human motion video generation. arXiv preprint arXiv:2502.04847 (2025)
30. Guédon, A., Lepetit, V.: Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)

31. Güler, R.A., Neverova, N., Kokkinos, I.: DensePose: Dense human pose estimation in the wild. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
32. 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 (CVPR) (2023)
33. 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 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2025)
34. Guo, J., Zhang, D., Liu, X., Zhong, Z., Zhang, Y., Wan, P., Zhang, D.: LivePortrait: Efficient portrait animation with stitching and retargeting control. arXiv preprint arXiv:2407.03168 (2024)
35. Guo, Y., Yang, C., Rao, A., Liang, Z., Wang, Y., Qiao, Y., Agrawala, M., Lin, D., Dai, B.: AnimateDiff: Animate your personalized text-to-image diffusion models without specific tuning. In: International Conference on Learning Representations (ICLR) (2024)
36. Guzov, V., Mir, A., Sattler, T., Pons-Moll, G.: Human poseitioning system (hps): 3d human pose estimation and self-localization in large scenes from body-mounted sensors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
37. Habermann, M., Liu, L., Xu, W., Pons-Moll, G., Zollhoefer, M., Theobalt, C.: Hdhumans: A hybrid approach for high-fidelity digital humans. Proceedings of the ACM on Computer Graphics and Interactive Techniques **6**(3) (2023)
38. 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 (ICCV) (2021)
39. He, X., Li, X., Kang, D., Ye, J., Zhang, C., Chen, L., Gao, X., Zhang, H., Wu, Z., Zhuang, H.: MagicMan: Generative novel view synthesis of humans with 3D-aware diffusion and iterative refinement. arXiv preprint arXiv:2408.14211 (2024)
40. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Advances in Neural Information Processing Systems (NeurIPS) (2020)
41. 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. In: International Conference on Learning Representations (ICLR) (2024)
42. Hu, L., Gao, X., Zhang, P., Sun, K., Zhang, B., Bo, 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 (CVPR) (2024)
43. Hu, L., Wang, G., Shen, Z., Gao, X., Meng, D., Zhuo, L., Zhang, P., Zhang, B., Bo, L.: Animate anyone 2: High-fidelity character image animation with environment affordance. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2025)
44. Hu, L., Zhang, H., Zhang, Y., Zhou, B., Liu, B., Zhang, S., Nie, L.: Gaussianavatar: Towards realistic human avatar modeling from a single video via animatable 3d gaussians. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
45. Hu, S., Hu, T., Liu, Z.: GauHuman: Articulated gaussian splatting from monocular human videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)

46. Huang, B., Yu, Z., Chen, A., Geiger, A., Gao, S.: 2d gaussian splatting for geometrically accurate radiance fields. In: ACM SIGGRAPH Conference Proceedings (2024)
47. Huang, Y., Yi, H., Xiu, Y., Liao, T., Tang, J., Cai, D., Thies, J.: TeCH: Text-guided reconstruction of lifelike clothed humans. In: International Conference on 3D Vision (3DV) (2024)
48. Huang, Y., Yuan, Y., Li, X., Kautz, J., Iqbal, U.: Adahuman: Animatable detailed 3D human generation with compositional multiview diffusion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2025)
49. 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 (CVPR) (2020)
50. Imagen Team Google: Imagen 3. arXiv preprint arXiv:2408.07009 (2024)
51. Jiang, B., Hong, Y., Bao, H., Zhang, J.: SelfRecon: Self reconstruction your digital avatar from monocular video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
52. Jiang, W., Yi, K.M., Samei, G., Tuzel, O., Ranjan, A.: Neuman: Neural human radiance field from a single video. In: European Conference on Computer Vision (ECCV) (2022)
53. Jiang, Y., Shen, Z., Hong, Y., Guo, C., Wu, Y., Zhang, Y., Yu, J., Xu, L.: Robust dual gaussian splatting for immersive human-centric volumetric videos. *ACM Transactions on Graphics* **43**(6) (2024)
54. Jiang, Z., Han, Z., Mao, C., Zhang, J., Pan, Y., Liu, Y.: VACE: All-in-one video creation and editing. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2025)
55. Junkawitsch, H., Sun, G., Zhu, H., Theobalt, C., Habermann, M.: Eva: Expressive virtual avatars from multi-view videos. In: ACM SIGGRAPH Conference Proceedings (2025)
56. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
57. Keetha, N., Karhade, J., Jatavallabhula, K.M., Yang, G., Scherer, S., Ramanan, D., Luiten, J.: Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
58. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (2023)
59. Khrodar, 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 (ECCV) (2024)
60. Kocabas, M., Chang, J.H.R., Gabriel, J., Tuzel, O., Ranjan, A.: HUGS: Human gaussian splats. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
61. Kolotouros, N., Alldieck, T., Zanfir, A., Bazavan, E., Fieraru, M., Sminchisescu, C.: DreamHuman: Animatable 3D avatars from text. In: Advances in Neural Information Processing Systems (NeurIPS) (2023)
62. Kong, W., et al.: HunyuanVideo: A systematic framework for large video generative models. arXiv preprint arXiv:2412.03603 (2024)
63. Lassner, C., Zollhöfer, M.: Pulsar: Efficient sphere-based neural rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)

64. Lee, J., Won, C., Jung, H., Bae, I., Jeon, H.G.: Fully explicit dynamic gaussian splatting. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2024)
65. Lei, J., Wang, Y., Pavlakos, G., Liu, L., Daniilidis, K.: GART: Gaussian articulated template models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2024)
66. Li, P., Zheng, W., Liu, Y., Yu, T., Li, Y., Qi, X., Chi, X., Xia, S., Cao, Y.P., Xue, W., Luo, W., Guo, Y.: Pshuman: Photorealistic single-image 3D human reconstruction using cross-scale multiview diffusion and explicit remeshing. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2025)
67. 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)
68. Li, T., Bolkart, T., Black, M.J., Li, H., Romero, J.: Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics (SIGGRAPH Asia)* (2017)
69. Li, Z., Chen, Z., Li, Z., Xu, Y.: Spacetime gaussian feature splatting for real-time dynamic view synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2024)
70. Li, Z., Sun, Y., Zheng, Z., Wang, L., Zhang, S., Liu, Y.: Animatable and relightable gaussians for high-fidelity human avatar modeling. *arXiv preprint arXiv:2311.16096* (2024)
71. 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 (CVPR)* (2024)
72. Lin, G., Jiang, J., Yang, J., Zheng, Z., Liang, C.: OmniHuman-1: Rethinking the scaling-up of one-stage conditioned human animation models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2025)
73. Liu, L., Habermann, M., Rudnev, V., Sarkar, K., Gu, J., Theobalt, C.: Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Transactions on Graphics* (2021)
74. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics* **34**(6) (2015)
75. Lu, Y., Dong, J., Kwon, Y., Zhao, Q., Dai, B., De la Torre, F.: GAS: Generative avatar synthesis from a single image. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2025)
76. Luiten, J., Kopanas, G., Leibe, B., Ramanan, D.: Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In: *International Conference on 3D Vision (3DV)* (2024)
77. Luo, Y., Rong, Z., Wang, L., Zhang, L., Hu, T., Zhu, Y.: DreamActor-M1: Holistic, expressive and robust human image animation with hybrid guidance. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2025)
78. Ma, Y., Liu, H., Wang, H., Pan, H., He, Y., Yuan, J., Zeng, A., Cai, C., Shum, H.Y., Liu, W., Chen, Q.: Follow-your-emoji: Fine-controllable and expressive freestyle portrait animation. In: *ACM SIGGRAPH Asia Conference Proceedings* (2024)
79. Men, Y., Yao, Y., Cui, M., Bo, L.: MIMO: Controllable character video synthesis with spatial decomposed modeling. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2025)

80. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
81. Mihajlovic, M., Prokudin, S., Tang, S., Maier, R., Bogo, F., Tung, T., Boyer, E.: SplatFields: Neural gaussian splats for sparse 3d and 4d reconstruction. In: European Conference on Computer Vision (ECCV) (2024)
82. 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: European Conference on Computer Vision (ECCV) (2020)
83. Mir, A., Alldieck, T., Pons-Moll, G.: Learning to transfer texture from clothing images to 3d humans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
84. Mir, A., Guler, R.A., Wang, J., Wonka, P., Zhou, B., Pons-Moll, G.: Raga: Real time ray traced gaussian shadow casting for 3dgs avatar-scene interaction. In: European Conference on Computer Vision (ECCV) (2026)
85. Mir, A., Moreau, A., Dhano, H., Zhang, Z., Pons-Moll, G., Pérez-Pellitero, E.: Gaspacho: Gaussian splatting for controllable humans and objects. arXiv preprint arXiv:2503.09342 (2025)
86. Mir, A., Puig, X., Kanazawa, A., Pons-Moll, G.: Generating continual human motion in diverse 3d scenes. In: International Conference on 3D Vision (3DV) (2024)
87. Mir, A., Wang, J., Guler, R.A., Guo, C., Pons-Moll, G., Zhou, B.: Aha! animating human avatars in diverse scenes with gaussian splatting. arXiv preprint arXiv:2511.09827 (2025)
88. Mokady, R., Hertz, A., Aberman, K., Pritch, Y., Cohen-Or, D.: Null-text inversion for editing real images using guided diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
89. Moon, G., Shiratori, T., Saito, S.: Expressive whole-body 3D gaussian avatar. In: European Conference on Computer Vision (ECCV) (2024)
90. Moreau, A., Song, J., Dhano, H., Shaw, R., Zhou, Y., Pérez-Pellitero, E.: Human gaussian splatting: Real-time rendering of animatable avatars. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
91. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics* **41**(4) (2022)
92. Nazarczuk, M., Catley-Chandar, S., Tanay, T., Zhang, Z., Slabaugh, G., Pérez-Pellitero, E.: Vidar: Video diffusion-aware 4d reconstruction from monocular inputs. In: Advances in Neural Information Processing Systems (NeurIPS) (2025)
93. Neverova, N., Güler, R.A., Kokkinos, I.: Dense pose transfer. In: European Conference on Computer Vision (ECCV) (2018)
94. Neverova, N., Novotny, D., Szafraniec, M., Khalidov, V., Labatut, P., Vedaldi, A.: Continuous surface embeddings. In: Advances in Neural Information Processing Systems (NeurIPS) (2020)
95. OpenAI: Sora 2. <https://openai.com/index/sora-2/> (2025)
96. Pan, P., Su, Z., Lin, C., Fan, Z., Zhang, Y., Li, Z., Shen, T., Mu, Y., Liu, Y.: HumanSplat: Generalizable single-image human gaussian splatting with structure priors. In: Advances in Neural Information Processing Systems (NeurIPS) (2024)

97. 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) (2024)
98. Pang, H.E., Liu, S., Cai, Z., Yang, L., Zhang, T., Liu, Z.: DISCO4D: Disentangled 4D human generation and animation from a single image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2025)
99. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
100. 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 of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
101. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
102. Peng, H.Y., Zhang, J.P., Guo, M.H., Cao, Y.P., Hu, S.M.: CharacterGen: Efficient 3D character generation from single images with multi-view pose canonicalization. *ACM Transactions on Graphics* (2024)
103. 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 (ICCV) (2021)
104. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: DreamFusion: Text-to-3D using 2D diffusion. In: International Conference on Learning Representations (ICLR) (2023)
105. Qian, Z., Wang, S., Mihajlovic, M., Geiger, A., Tang, S.: 3dgs-avatar: Animatable avatars via deformable 3d gaussian splatting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
106. Qiu, L., Gu, X., Li, P., Zuo, Q., Shen, W., Zhang, J., Qiu, K., Yuan, W., Chen, G., Dong, Z., Bo, L.: LHM: Large animatable human reconstruction model for single image to 3D in seconds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2025)
107. 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)
108. Qiu, L., Zhu, S., Zuo, Q., Gu, X., Dong, Y., Zhang, J., Xu, C., Li, Z., Yuan, W., Bo, L., Chen, G., Dong, Z.: Anigs: Animatable gaussian avatar from a single image with inconsistent gaussian reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2025)
109. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
110. Rout, L., Chen, Y., Ruiz, N., Caramanis, C., Shakkottai, S., Chu, W.S.: Semantic image inversion and editing using rectified stochastic differential equations. In: International Conference on Learning Representations (ICLR) (2025)
111. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S.K.S., Gontijo-Lopes, R., Ayan, B.K., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M.: Photorealistic text-to-image diffusion models with deep language understanding. In: Advances in Neural Information Processing Systems (NeurIPS) (2022)

112. 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 (ICCV) (2019)
113. 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 (CVPR) (2020)
114. Sárándi, I., Pons-Moll, G.: Neural localizer fields for continuous 3D human pose and shape estimation. In: Advances in Neural Information Processing Systems (NeurIPS) (2024)
115. Shao, R., Pang, Y., Zheng, Z., Sun, J., Liu, Y.: Human4DiT: 360-degree human video generation with 4D diffusion transformer. *ACM Transactions on Graphics* (2024)
116. Shaw, R., Nazarczuk, M., Song, J., Moreau, A., Catley-Chandar, S., Dharmo, H., Perez-Pellitero, E.: SWinGS: Sliding windows for dynamic 3D gaussian splatting. In: European Conference on Computer Vision (ECCV) (2024)
117. Siarohin, A., Lathuilière, S., Tulyakov, S., Ricci, E., Sebe, N.: First order motion model for image animation. In: Advances in Neural Information Processing Systems (NeurIPS) (2019)
118. Sim, G., Moon, G.: PERSONA: Personalized whole-body 3D avatar with pose-driven deformations from a single image. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2025)
119. Tan, J., Xiang, D., Tulsiani, S., Ramanan, D., Yang, G.: DressRecon: Freeform 4D human reconstruction from monocular video. In: International Conference on 3D Vision (3DV) (2025)
120. Tancik, M., Weber, E., Ng, E., Li, R., Yi, B., Kerr, J., Wang, T., Kristoffersen, A., Austin, J., Salahi, K., Ahuja, A., McAllister, D., Kanazawa, A.: Nerfstudio: A modular framework for neural radiance field development. In: ACM SIGGRAPH Conference Proceedings (2023)
121. Taubner, F., Zhang, R., Tuli, M., Lindell, D.B.: CAP4D: Creating animatable 4D portrait avatars with morphable multi-view diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2025)
122. Tu, S., Xing, Z., Han, X., Cheng, Z.Q., Dai, Q., Luo, C., Wu, Z.: StableAnimator: High-quality identity-preserving human image animation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2025)
123. Wan Team: Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314* (2025)
124. Wang, J., Pu, J., Qi, Z., Guo, J., Ma, Y., Huang, N., Chen, Y., Li, X., Shan, Y.: Taming rectified flow for inversion and editing. In: International Conference on Machine Learning (ICML) (2025)
125. Wang, L., Zhao, X., Sun, J., Zhang, Y., Zhang, H., Yu, T., Liu, Y.: StyleAvatar: Real-time photo-realistic portrait avatar from a single video. In: ACM SIGGRAPH Conference Proceedings (2023)
126. Wang, X., Zhang, S., Gao, C., Wang, J., Zhou, X., Zhang, Y., Yan, L., Sang, N.: UniAnimate: Taming unified video diffusion models for consistent human image animation. *Science China Information Sciences* (2025)

127. Wang, X., Zhang, S., Tang, L., Zhang, Y., Gao, C., Wang, Y., Sang, N.: UniAnimate-DiT: Human image animation with large-scale video diffusion transformer. arXiv preprint arXiv:2504.11289 (2025)
128. Wang, Y., Hong, F., Yang, S., Jiang, L., Wu, W., Loy, C.C.: MEAT: Multiview diffusion model for human generation on megapixels with mesh attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2025)
129. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* (2004)
130. Wei, H., Yang, Z., Wang, Z.: AniPortrait: Audio-driven synthesis of photorealistic portrait animation. arXiv preprint arXiv:2403.17694 (2024)
131. 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: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
132. Westover, L.A.: Splatting: a parallel, feed-forward volume rendering algorithm. Ph.D. thesis, University of North Carolina at Chapel Hill (1992)
133. Wu, B., et al.: HunyuanVideo 1.5 technical report. arXiv preprint arXiv:2511.18870 (2025)
134. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
135. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond. *Computer Graphics Forum* (2022)
136. Xie, Y., Xu, H., Song, G., Wang, C., Shi, Y., Luo, L.: X-Portrait: Expressive portrait animation with hierarchical motion attention. In: ACM SIGGRAPH Conference Proceedings (2024)
137. 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 (CVPR) (2023)
138. Xu, H., Alldieck, T., Sminchisescu, C.: H-nerf: Neural radiance fields for rendering and temporal reconstruction of humans in motion. In: Advances in Neural Information Processing Systems (NeurIPS) (2021)
139. Xu, S., Chen, G., Guo, Y.X., Yang, J., Li, C., Zang, Z., Zhang, Y., Tong, X., Guo, B.: VASA-1: Lifelike audio-driven talking faces generated in real time. In: Advances in Neural Information Processing Systems (NeurIPS) (2024)
140. Xu, Y., Chen, B., Li, Z., Zhang, H., Wang, L., Zheng, Z., Liu, Y.: Gaussian head avatar: Ultra high-fidelity head avatar via dynamic gaussians. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
141. 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 (CVPR) (2024)
142. Xu, Z., Huang, Z., Cao, J., Zhang, Y., Cun, X., Shuai, Q., Wang, Y., Bao, L., Li, J., Tang, F.: AnchorCrafter: Animate cyber-anchors selling your products via human-object interacting video generation. arXiv preprint arXiv:2411.17383 (2024)

143. Xu, Z., Yu, Z., Zhou, Z., Zhou, J., Jin, X., Hong, F.T., Ji, X., Zhu, J., Cai, C., Tang, S., Lin, Q., Li, X., Lu, Q.: HunyuanPortrait: Implicit condition control for enhanced portrait animation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2025)
144. Yang, X., Chen, X., Gao, D., Wang, S., Han, X., Wang, B.: HAVE-FUN: Human avatar reconstruction from few-shot unconstrained images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
145. Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., Xu, J., Yang, Y., Hong, W., Zhang, X., Feng, G., Yin, D., Zhang, Y., Wang, W., Cheng, Y., Xu, B., Gu, X., Dong, Y., Tang, J.: CogVideoX: Text-to-video diffusion models with an expert transformer. In: International Conference on Learning Representations (ICLR) (2025)
146. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
147. 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. In: International Conference on Machine Learning (ICML) (2025)
148. Zhang, Z., Yang, Z., Yang, Y.: SIFU: Side-view conditioned implicit function for real-world usable clothed human reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
149. 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 (CVPR) (2022)
150. 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)
151. 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* (2021)
152. Zhong, Y., Li, Z., Chen, D.Z., Hong, L., Xu, D.: Taming video diffusion prior with scene-grounding guidance for 3d gaussian splatting from sparse inputs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2025)
153. Zhou, J., Wu, Y., Li, S., Wei, M., Fan, C., Chen, W., Jiang, W., Wang, F.: RealisDance-DiT: Simple yet strong baseline towards controllable character animation in the wild. *arXiv preprint arXiv:2504.14977* (2025)
154. 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 Transactions on Graphics* (2024)
155. Zhu, S., Chen, J.L., Dai, Z., Su, Q., Xu, Y., Cao, X., Yao, Y., Zhu, H., Zhu, S.: Champ: Controllable and consistent human image animation with 3D parametric guidance. In: European Conference on Computer Vision (ECCV) (2024)
156. 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. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2025)

157. Zielonka, W., Bagautdinov, T., Saito, S., Zollhöfer, M., Thies, J., Romero, J.: Drivable 3D gaussian avatars. In: International Conference on 3D Vision (3DV) (2025)