

SIGMAN: Scaling 3D Human Gaussian Generation with Millions of Assets

Yuhang Yang^{1,2,*}, Fengqi Liu^{2,3,*}, Yixing Lu⁴, Qin Zhao², Pingyu Wu¹, Wei Zhai^{1,†}, Ran Yi³
 Yang Cao¹, Lizhuang Ma³, Zheng-Jun Zha¹, Junting Dong^{2,†}

¹ USTC ² Shanghai AI Lab ³ SJTU ⁴ CMU

*Equal Contribution †Corresponding Author

https://yyvhang.github.io/SIGMAN_3D/

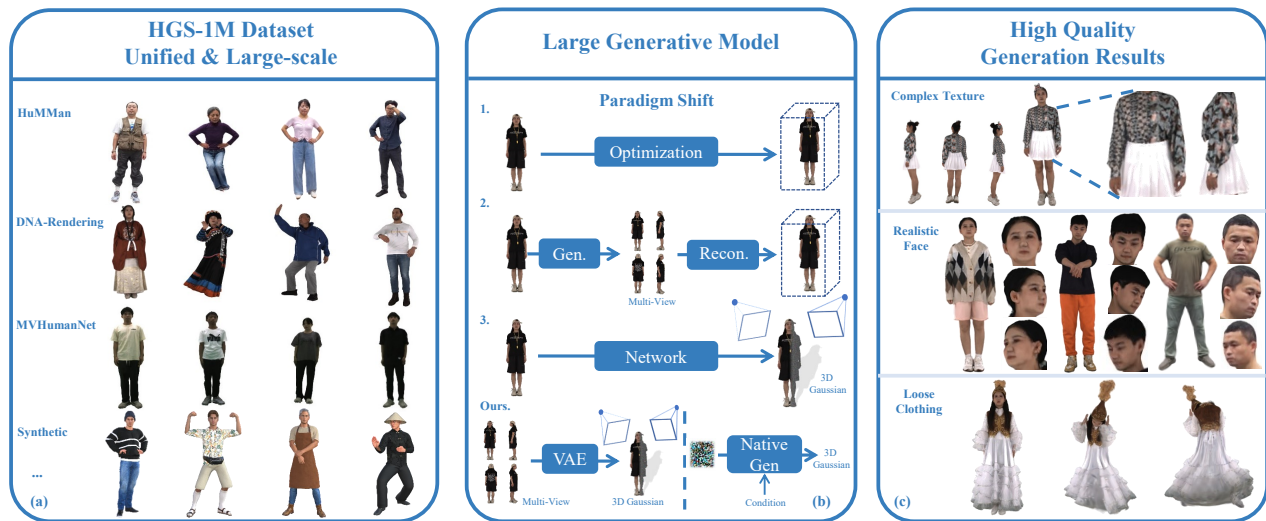


Figure 1. (a). In this work, we construct a large-scale, unified 3D Human Gaussian Dataset, called HGS-1M, to support (b) the large-scale generation model for 3D human Gaussian generation. (c) This paradigm, with large-scale data, produces high-quality 3D human Gaussians that exhibit complex textures, facial details, and realistic deformation of loose clothing.

Abstract

3D human digitization has long been a highly pursued yet challenging task. Existing methods aim to generate high-quality 3D digital humans from single or multiple views, but remain primarily constrained by current paradigms and the scarcity of 3D human assets. Specifically, recent approaches fall into several paradigms: optimization-based and feed-forward (both single-view regression and multi-view generation with reconstruction). However, they are limited by slow speed, low quality, cascade reasoning, and ambiguity in mapping low-dimensional planes to high-dimensional space due to occlusion and invisibility, respectively. Furthermore, existing 3D human assets remain small-scale, insufficient for large-scale training. To address these challenges, we propose a latent space generation paradigm for 3D human digitization, which involves compressing multi-view images into Gaussians via a UV-structured VAE, along with DiT-based conditional gener-

ation, we transform the ill-posed low-to-high-dimensional mapping problem into a learnable distribution shift, which also supports end-to-end inference. In addition, we employ the multi-view optimization approach combined with synthetic data to construct the HGS-1M dataset, which contains 1 million 3D Gaussian assets to support the large-scale training. Experimental results demonstrate that our paradigm, powered by large-scale training, produces high-quality 3D human Gaussians with intricate textures, facial details, and loose clothing deformation.

1. Introduction

Human digitization has always been a central topic in the generation and reconstruction field, it supports various applications such as gaming, AR/VR, etc. For 3D human digitization, the task usually involves mapping the input (e.g., single image, multi images, text) to 3D human assets, which

then could be animated through rigging. Compared to pixel-level character digitalization [1, 4, 5, 12, 19, 28, 35, 45, 65, 77, 78, 83], 3D character digitalization offers multi-view consistency and more stable controllability.

Recently, studies on 3D human digitization primarily focus on image-to-3D human, and these methods can be summarized into the following paradigms. 1) The optimization-based approach [74, 75, 88, 89]: which typically utilizes geometric constraints and human priors to optimize each character, the limitations of this type of method are slow speed and limited quality; 2) Feed-forward: this includes two manners, one of which is the two-stage cascade manner [36, 69, 70], this kind of method first generates multiple views that correspond to the input view, and then takes these views to reconstruct the 3D human. Its cascade structure requires both stages to work well, however, humans possess high diversity, generating multiple views from a single perspective while maintaining consistency is highly unstable, affecting the final quality. The other manner [3, 20, 29, 32, 54, 94] involves taking a single view or sparse views as the input and employing a network to directly predict the 3D human (represented in 3D Gaussian Splatting [25]). Although these methods are faster and further improve the quality of digitalization, it requires the network to learn the mapping from 2D-pixel planes to 3D geometric space through limited views, resulting in occlusion and invisibility problems, which makes it hard for the network to learn the complete mapping and low quality shows up on non-visible sides.

To address the above limitations, in this paper, we use 3D Gaussian Splatting as the representation, pushing 3D human digitization towards the paradigm of 3D native generation in latent space. Specifically, our paradigm involves training a VAE model [53] to compress multi-view consistent images into latent representations and then decode them into Gaussian attributes. Afterward, we train a conditional DiT-based [49] generation model, which takes a single image or text as the condition to generate the 3D human Gaussian. In this paradigm, the VAE stage has multiple consistent and comprehensive views as the inputs, providing sufficient geometric contexts and enabling the Gaussian latent to represent high-quality geometry, appearance details, and pose-dependent deformation. Turning to the generation, the learning objective shifts to capture the transition from conditional features to latents, it avoids learning the hard mapping from low-dimension (2D plane) to high-dimension (3D space). Moreover, in the inference stage, only a certain number of inference steps are needed to generate the corresponding latent, refrain from the cascade structure, and does not need per-case optimization.

The pose and appearance of humans have strong diversity, and there are many high-frequency details, *e.g.*, complex textures of clothes and faces. Directly compressing

multi-view images to 3D human Gaussians is non-trivial. To achieve this, we introduce a UV-structured latent representation, which takes parametric model SMPL-X [48] as human priors and leverages multi-view geometric contexts as references, employing learnable tokens defined in the SMPL-X UV coordinates to query from geometric contexts, thus modeling the latent, it can then be decoded to Gaussian attributes on the UV coordinate, by back-projecting onto the SMPLX template, we obtain the canonical human Gaussian, and the posed 3D human Gaussian could be obtained through LBS. This manner avoids directly learning the explicit geometric structure, which is conducive to the learning of appearance details. In addition, the LBS-driven method enables the latent in the VAE stage to represent the corresponding Gaussian properties under different poses and capture the implicit pose-related deformation. Furthermore, we also introduce UV-initialization and Conv-Attention dual branch mechanisms in the encoding stage to accelerate convergence and enhance the detail learning. In the generation stage, we adopt the MM-DiT architecture and a human-specific encoder [26] to better learn the transformation from conditional features to Gaussian latents.

In addition to the learning paradigm, a particularly challenging issue is that 3D human assets are very scarce, which can not support large-scale training. However, the multi-view loop optimization method can reconstruct high-quality 3D human Gaussians from dense views. In light of this, we use AnimatableGaussians [38] to build a 3D human Gaussian dataset from multi-view human datasets [2, 6, 16, 73, 81, 92], to expand the data diversity, we also combine synthetic data. Eventually, we construct the HGS-1M dataset, which contains 1 million 3D human Gaussians, including diverse races, genders, age groups, appearances, clothing, and postures. This large-scale dataset could support the training of large-scale models and serve as the new test bed for 3D human digitization.

The contributions are summarized as follows:

- 1) We introduce a paradigm of using large-scale latent generation for 3D human digitization. Through two-stage learning of compression and generation, we transform the challenging low-to-high-dimensional mapping problem into a conditional feature-to-latent transfer task, enabling high-quality modeling of 3D humans.
- 2) We propose the UV-structured latent representation, which leverages the parametric SMPL-X as priors to enhance VAE’s learning of appearance details and pose-related deformations. With the MMDiT-based generation, we perform native 3D human generation.
- 3) We construct the HGS-1M dataset, which contains 1 million 3D human Gaussians, encompassing diverse races, genders, ages, appearances, clothing, and postures. It supports large-scale training and serves as the new test bed for 3D human digitization.

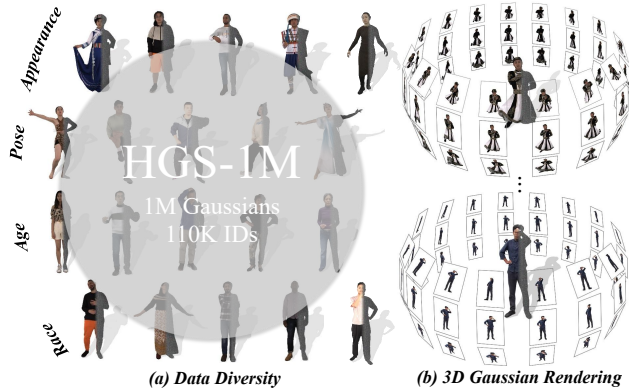


Figure 2. **HGS-1M Dataset**. The constructed HGS-1M dataset, it contains 1 million 3D Gaussian human assets of different ages, races, appearances, and poses. It supports free-view rendering.

- 4) Experiments show that large-scale data and model improve the quality of 3D human digitization, highlighting the superiority over previous methods and paradigms.

2. Related Work

2.1. 3D Human Digitization

For 3D human creation, a branch adopts implicit representation [21, 22, 29, 30, 44, 51, 56, 57, 90, 93] and explicit human prior for optimization [8, 88, 89], calculating each instance individually. While these methods enhance the robustness to a certain extent, they take relatively longer to get results for each input, lack generalizability, and deliver limited quality. To enable rapid digitization and obtain high-quality, photo-realistic 3D humans, some approaches gradually shift to the feed-forward manner, including using a regression network [3, 31, 32, 54, 91, 94] and cascaded generation with reconstruction [9, 36, 50, 62, 70]. For the regression-based method, part of it needs multiple views as the input, facing significant constraints during inference, whereas single-view approaches relax this limitation but necessitate the network to learn mappings from low-dimensional planes to high-dimensional space, encountering occlusion and invisibility issues. The cascade type that generates multi-view images at first and then reconstructs the 3D assets, it requires view-consistent generation in the first stage, which is unstable and remains challenging, affecting the subsequent reconstruction. Recent work E³Gen [87] explores native 3D generation for human digitization but focuses solely on unconditional generation. For the animation, some methods [54, 87, 94] adopt the LBS manner, using SMPL(X)-based motion sequences [13, 14, 68, 72, 80, 85] to drive, other method like DRIVE [61] employs rigging for the generated 3D human Gaussian, driving the character. In addition to obtaining 3D humans from images, a few works also explore generating 3D hu-

Dataset	#Frames	#ID	# View	# 3D Asset
Human3.6M [23]	3.6M	11	4	-
HUMBI [82]	26M	772	107	-
HuMMan [2]	60M	1000	10	-
ActorsHQ [24]	40K	8	160	-
DNA-Rendering [6]	67.5M	500	60	-
MVHumanNet [73]	645.1M	4500	48	-
HuGe100K [94]	2.4M	100K	24	-
THuman2.1 [81]	-	2500	Free	2500
2K2K [16]	-	2050	Free	2050
HGS-1M	> 90M	110K	Free	1M

Table 1. Comparison of various attributes of multiple datasets. HGS1-M contains the most IDs and large-scale 3D assets that support rendering from any perspective.

mans from text [10, 15, 27, 42, 64]. Nonetheless, prior research predominantly focuses on small-scale datasets, lacking large-scale unified datasets and benchmarks to advance the 3D human digitization. In this work, we define a new paradigm for obtaining 3D humans and construct a large-scale, unified human Gaussian dataset to promote the field forward.

2.2. Large Model for Generation

The LLM [47] opens the large-scale generation era, and the visual field also gradually shifts to this paradigm, however, the number of visual tokens is huge, and the generation speed and computational complexity are greatly challenging. Stable diffusion [52] starts the latent generation route, which compresses the raw visual data to latent and then employs diffusion models to generate the latent. It has achieved promising results in fields like image generation [33, 40, 66], video generation [1, 28, 65], and 3D generation [67, 71, 86]. For this paradigm, there are several key aspects: large-scale and high-quality data; an appropriate training objective to learn the transfer between conditional features and latents; and a large-scale model to learn the target distribution and ensure the generated quality.

3. HGS-1M Dataset

Currently, the public 3D whole human assets are very limited [16, 81, 92], which cannot meet the requirements for large-scale training. Although multi-view human datasets [2, 6, 73] can provide a large amount of human data, the quantity and pose of the camera are different in each dataset, and the views of some datasets are not dense enough to serve as supervision. To train a unified Gaussian generative model, unifying these data and getting as many perspectives as possible is necessary. Fortunately, the multi-view optimization method can reconstruct high-quality Gaussian data from multi-view human datasets, thus, we employ AnimatableGaussians [38] to merge multiple multi-view hu-

man datasets and reconstruct sequences of human Gaussians, each sequence taking around 20 4090 GPU hours for optimization. While the total computational overhead is huge, this approach allows us to obtain high-quality human Gaussian data and could render it at free-view for supervision. Next, we place each human Gaussian at the origin and unify the Gaussian position based on the corresponding SMPL-X root rotation and translation. For each human Gaussian, we select 90 views for rendering, including 30 horizontal, 30 upward, and 30 downward perspectives, shown in Fig. 2. For the small-scale data *e.g.*, Thuman2.1 [81] and 2K2K [16], we also move the asset to the origin and render it using the same cameras. Besides, we collect over 100k synthetic data and convert them into the same representation to further expand the diversity. Eventually, we construct the dataset HGS-1M, which contains 1 million 3D human Gaussians with diverse races, genders, age groups, appearances, clothing, and postures, shown in Fig. 2. Compared with previous datasets, it contains the largest amount of 3D human assets and has rich IDs and poses (Tab. 1). Moreover, it can be rendered from free-view, which supports large-scale training and evaluation.

4. Method

Our goal is to generate 3D human Gaussians from the instruction (single image or text), the method consists of a UV-structured VAE, and an MMDiT-based generation model, shown in Fig. 3. The VAE compresses multi-view images into latent Gaussian attributes in the UV structural space, and the DiT aims to capture the transformation from conditional features to the compressed Gaussian latents. Through this way, the challenging mapping from sparse 2D image planes to 3D Gaussian space is thereby transformed into a distribution shift process.

4.1. Preliminary

SMPL-X [48] is a parametric human model that represents human body with a deformable mesh, formulated as:

$$M(\beta, \theta, \psi) = \text{LBS}(T_P(\beta, \theta, \psi)), J(\beta), \theta, \mathcal{W}), \quad (1)$$

where β, θ, ψ represent shape, pose and expression, respectively. *LBS* denotes the transformation from the canonical T-pose T_p to the pose θ , based on the pre-trained skinning weights \mathcal{W} and joints $J(\beta)$.

3D Gaussian Splatting [25] explicitly represents static 3D scenes with point primitives, each primitive $G(\mathcal{X})$ could be parameterized with a 3D covariance matrix Σ and mean μ , formulated as: $G(\mathcal{X}) = e^{-\frac{1}{2}(\mathcal{X}-\mu)^T \Sigma^{-1}(\mathcal{X}-\mu)}$. The covariance matrix could be further decomposed into scaling matrix S and rotation matrix R for optimization, where $\Sigma = RSS^T R^T$. In addition, the view-dependent color is represented by coefficients of spherical harmonics.

In this paper, we utilize diagonal vector $s \in \mathbb{R}^3$ and axis angle $r \in \mathbb{R}^3$ to represent the scaling and rotation matrix, and take RGB color $c \in \mathbb{R}^3$ to replace the spherical harmonic coefficients, the position of Gaussians is denoted by $\mu \in \mathbb{R}^3$. With the SMPL-X prior, the learning objectives of μ, s, r are defined as offsets relative to the T-pose $(\hat{\mu}, \hat{s}, \hat{r})$ template, formulated as: $\bar{\mu} = \hat{\mu} + \mu, \bar{s} = \hat{s} \cdot s, \bar{r} = \hat{r} \cdot r$. Along with the c and opacity $\alpha \in \mathbb{R}^3$, the canonical human Gaussian is represented as $\mathcal{G} : (\bar{\mu}, \bar{s}, \bar{r}, c, \alpha)$, which then undergo part-aware deformation via *LBS* and output the posed human Gaussians [87]. During rendering, 3D Gaussians are projected onto the pixel plane, the pixel color C is calculated as: $C = \sum_{i=1}^N \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j) c_i$, where c_i denotes the color of the i -th Gaussian, and α_i is the blending weight calculated from the opacity and probability density.

4.2. UV-Structured VAE

The VAE leverages a set of learnable tokens defined in the SMPL-X UV space to compress multi-view renderings into latent and decode the latent into 3D human Gaussians. We emphasize that an efficient and effective VAE capable of compressing 3D human Gaussians should meet the following aspects: **1)** Quickly converge under the nature of strong diversity of humans; **2)** Decoupled appearance and pose for convenient animation during subsequent generation; **3)** Sufficiently compact and low redundancy latent space, suitable for the generative model learning.

Encoder. VAE requires numerous steps to learn geometric structures when directly compressing multi-view renderings into 3D Gaussians [17]. Some works [34, 76] introduce explicit geometric information to accelerate convergence, but the injection of explicit geometries requires the network to capture the geometric structure of different posed humans, which needs more optimization steps and may affect the convenience of subsequent pose-driven. To achieve the best of both worlds, we utilize UV-based representation [32, 87, 94] to introduce human priors and compress Gaussian attributes into learnable latent tokens, which are then decoded into Gaussian attributes and reprojected onto the canonical human Gaussians through the SMPL-X UV coordinates, and eventually obtain posed 3D human Gaussians through differentiable *LBS*.

Specifically, the VAE encoder takes multi-view renderings and their corresponding Plücker [59] ray embeddings as the input $\mathcal{V} \in \mathbb{R}^{V \times (3+6) \times H \times W}$, where V is the number of views, H, W are height and width. Then \mathcal{V} is fed into 3D convolutional blocks, obtaining multiview contexts $\mathbf{F}_v \in \mathbb{R}^{C \times V \times H' \times W'}$, where $H' = H/8, W' = W/8$. Next, we define learnable tokens $\tau \in \mathbb{R}^{N' \times C}$ ($N' = H'W'$), to accelerate the learning of appearance details (*e.g.*, face), we select several views and project the RGB image back onto SMPL-X mesh to obtain an initial UV map [63], which is then embedded and combined with τ as the initialization

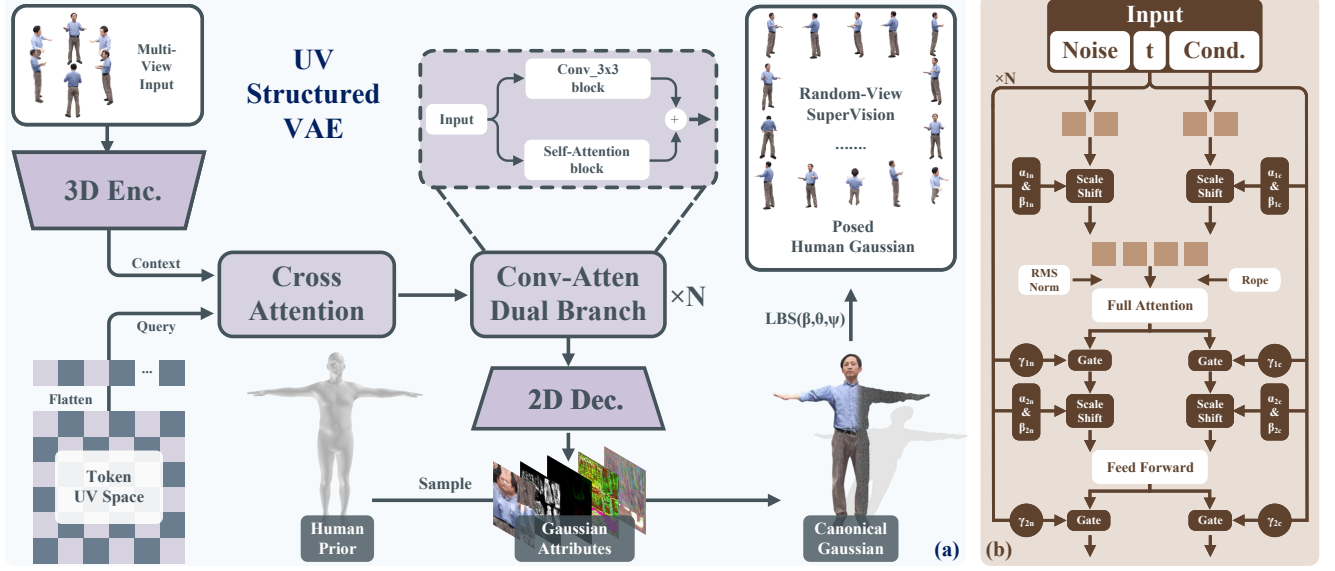


Figure 3. **Method.** The pipeline of our method. (a). The UV-structured VAE, which uses human priors to define learnable tokens in UV space and takes them to query multi-view contexts to model the Gaussian latent, then, the latent is decoded into human Gaussians in canonical space and could be driven by differentiable LBS to obtain the final posed human Gaussian. (b). The MM-DiT architecture, it treats the conditional sequence and noise as a whole sequence to complete controllable 3D human Gaussian generation.

$\hat{\tau} \in \mathbb{R}^{N' \times 2C}$. Subsequently, a cross-attention is employed among the $\hat{\tau}$ and flattened \mathbf{F}_v to convert multi-view contexts into the latent. With the following multiple Conv-Atten dual branch blocks, we model the Gaussian latent τ_L , the process could be expressed as:

$$\tau_L = \Theta(\text{Cross}(\hat{\tau} + \text{pos}, \text{Conv}(\mathcal{V}))), \quad (2)$$

where Θ represents the Conv-Atten dual branch blocks, this design captures global features and learns fine-grained local details through the inductive bias of convolution. $\text{Cross}(\text{query}, \text{key}/\text{value})$ means the cross-attention, pos is the position encoding, and Conv denotes the 3D convolutional blocks. In this way, we introduce the human prior through a simple UV projection, avoiding the learning of explicit geometries and decoupling the pose and appearance. Besides, introducing learnable tokens in the UV space effectively reduces the redundancy caused by background pixels in multi-view renderings, avoiding the latent tokens being occupied to represent white pixels, which may result in the loss of representation of appearance details.

Decoder. The latent τ_L is decoded and upsampled through several 2D convolution blocks, with flowing projection heads, we obtain UV-space Gaussian attributes \mathcal{G}_{uv} , including the position offset $\mu \in \mathbb{R}^{N \times 3}$, scale offset $s \in \mathbb{R}^{N \times 3}$, rotation $r \in \mathbb{R}^{N \times 3}$ (equal to obtain $\bar{\mu}, \bar{s}, \bar{r}$, refer to Sec. 4.1), opacity $\alpha \in \mathbb{R}^{N \times 1}$ and RGBs $c \in \mathbb{R}^{N \times 3}$, where $N = HW$. Then, we use the T-pose template to sample Gaussian attributes of the corresponding vertices from \mathcal{G}_{uv} , obtaining the canonical human Gaussians \mathcal{G}_c . Eventually,

the posed human Gaussians are calculated through a forward skinning based on LBS . In specific, the position of each Gaussian is calculated as: $\bar{\mu}' = \sum_{i=1}^{n_b} w_i \mathbf{B}_i \bar{\mu}$, where n is the number of joints, w_i represents skinning weights, and B_i denotes the transformation matrix of the i -th joint from the canonical pose space to the posed space. Plus, the rotation matrix is calculated as:

$$\mathbf{R}' = \mathbf{T}_{1:3,1:3} \mathbf{R}, \text{ where } \mathbf{T} = \sum_{i=1}^n w_i(\bar{\mu}) \mathbf{B}_i \quad (3)$$

where \mathbf{R} is the rotation matrix derived from the axis angle r . T is the transformation matrix computed as the weighted sum of the transformation B_i , $w_i(\bar{\mu})$ corresponds to the skinning weights associated with the position $\bar{\mu}$.

Optimization. The VAE is optimized in an end-to-end manner, which crosses both input views and randomly chosen views. The whole loss includes the reconstruction loss \mathcal{L}_{L1} , perception loss $\mathcal{L}_{\text{lpips}}$ (VGG-based), KL constraints \mathcal{L}_{KL} , and adversarial loss \mathcal{L}_{GAN} :

$$\mathcal{L} = \mathcal{L}_{L1} + \lambda_{\text{lpips}} \mathcal{L}_{\text{lpips}} + \lambda_{\text{kl}} \mathcal{L}_{\text{KL}} + \lambda_{\text{GAN}} \mathcal{L}_{\text{GAN}} \quad (4)$$

During the first stage of the training, we use \mathcal{L}_{L1} , $\mathcal{L}_{\text{lpips}}$, and \mathcal{L}_{KL} , after number of steps, we add the adversarial loss \mathcal{L}_{GAN} for the subsequent training.

4.3. Structured Latents Generation

Model Design. We employ the MM-DiT architecture [11, 12, 28, 79] as the generative backbone, which treats the

encoding of input conditions (*e.g.*, text or image) and noise as a whole sequence and performs full-attention, shown in Fig. 3 (b). In each Transformer block, we use Rotary Position Embedding [60] (RoPE) to enhance the model’s ability to capture both absolute and relative positional relationships. In addition, we adopt RMSNorm [84] in each block to stabilize the training. Plus, for the image-conditioned encoder, we employ the Sapiens [26], which provides more fine-grained human features. For the text-conditioned generation, we take T5-XXL [7] as the text encoder.

Training Objective. We employ DDPM v-prediction [58] to learn the transition from the conditional feature to the structured latent, the diffusion loss is formulated as:

$$L_{\text{gen}}(\theta) := \mathbf{E}_{t, x_0, \epsilon} \|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t)\|^2 \quad (5)$$

where t is uniformly distributed between 1 and T (1000 in training), $\bar{\alpha}$ is a hyperparameter related to t , controlling the noise schedule, x_0 is the input latent, ϵ denotes the noise and ϵ_θ is the network which predicts the noise. During image-to-3D training, we first use the front view as input, enabling the model to quickly learn the distribution of facial details from a large amount of data, and then use a random view as input to fine-tune the model, making it more robust to the perspective of the input image.

5. Experiment

5.1. Implementation details

For the VAE training, we render 90 views per human Gaussian, there are 6 input views with size 512×512 and the model is optimized by randomly choose 10 views, the latent channel is 16. We set $\lambda_{\text{lips}} = 1$, $\lambda_{\text{kl}} = 1e - 6$ at the first stage, and train $300k$ steps with 192 batchsize. Then, we add \mathcal{L}_{GAN} and set $\lambda_{\text{lips}} = 0.1$, using the adversarial loss dominate the decrease of perception loss and train subsequent $200k$ steps. The optimizer is AdamW [43] with parameters $\beta_1 = 0.9$, $\beta_2 = 0.95$, the weight decay is 0.05, and the learning rate is set to $2e - 5$. The DiT training uses the same optimizer and parameter, with the learning rate $1e - 4$, we train a total of $300k$ steps (256 batchsize). The largest model has 2B parameters, trained with 48 A800 GPUs. We use Classifier-Free Guidance (CFG) [18] at inference, the scale is 3.5 and the sampling step is 30.

We extract 330 samples (distinct ID) (each 20 views) as the test set for quantitative comparison. For LGM [62] and GHG [32], we train them using a subset of HGS-1M (around 300k Gaussians). All quantitative experimental results are reported on this subset (including our method). For the VAE comparison, we employ three types of methods: 1) φ_1 : directly compresses multi-view images and decodes them into multi-view Gaussians, using the video VAE [79] architecture but removing its causal padding; 2)

Type	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	Method	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow
φ_1	24.38	0.067	0.936	SIFU [89]	14.32	0.125	0.874
φ_2	27.74	0.031	0.941	LGM [62]	20.87	0.086	0.925
φ_3	28.61	0.028	0.946	GHG [32]	21.03	0.074	0.932
Ours	29.32	0.028	0.953	Ours	25.45	0.044	0.939
Ours*	30.53	0.025	0.958	Ours*	26.14	0.040	0.943

(a)

(b)

Table 2. (a). Metrics of VAE comparison, where $\varphi_1, \varphi_2, \varphi_3$ represent 3 types of methods (Sec. 5.1). (b). Metrics of image-to-3D human Gaussians comparison, SIFU performs per-instance optimization, LGM and GHG are trained with our data. * denotes the training at the whole dataset.

φ_2 : this manner structures the multi-view inputs into pixel-aligned Gaussians at first, then fine-tuning an image VAE [39]; 3) φ_3 : introducing explicit human geometry as the prior, which takes posed downsampled SMPL-X points as latent points and upsamples them in the decoder to get the final human Gaussians, similar to [34].

For the image-to-3D, the comparison baselines encompass several paradigms, *e.g.*, per-optimization: SIFU [89], feedforward: regression-based GHG [32], and cascade manner LGM [62]. For the comparison of text-to-3D human Gaussian, we conduct a quantitative comparison with Diffspat [39], please refer to the results in the Supp.

5.2. VAE Reconstruction Results

For a fair comparison, all comparison methods maintain the same latent size. The quantitative results of VAE comparison are shown in Tab. 2 (a), our method outperforms other types across all metrics. φ_1 directly compresses multi-view renderings to corresponding Gaussians without introducing any priors, the limited number of latent tokens makes it hard to represent both 3D structures and appearance; φ_2 is a cascaded structure, if the pixel-aligned Gaussians in the first stage has low quality, it affects the final reconstruction effect. Additionally, in the pixel-aligned Gaussian, some pixels representing the rendered background are redundant, which occupy certain latent tokens, resulting in no extra tokens to represent appearance details (*e.g.*, face). φ_3 achieves suboptimal results, but converges slowly and requires a large number of steps to grasp the structure, which affects the learning of appearance details.

5.3. Generation Results

The quantitative metrics of generation are reported in Tab. 2 (b). Our method outperforms other methods across all metrics. As shown in the Fig. 4, the quantitative results include different skin colors, genders, ages, poses, input views, complex textures, and loose clothing. The output includes horizontal, top-down, and bottom-up perspectives, which fully display the results of each method. As can be



Figure 4. Qualitative comparison with baselines, including inputs of different ages, races, viewpoints, and complex textures, loose cloth. Our results demonstrate more details of faces and clothes. Note: We keep the original best inference setting of the comparison method. SIFU uses one view, LGM uses four views, and GHG uses three views.

seen, our method can better express fine-grained appearance details *e.g.*, face, diverse poses, and deformable clothing. Plus, the last row of Tab. 2 (a) and (b) show that the increase of data tokens can further enhance the model capabilities, reflecting that HGS-1M can support large-scale training and bring improvements in model performance.

5.4. Ablation Study

VAE. All ablation experiments are conducted on the sub-data (300k), and ablation experiments of the VAE are reported in the Tab. 3 (a), it could be divided into three aspects: **1) Model Design.** We validate the effectiveness of the dual branch (Θ in Sec. 4) and the initialization of UV maps. **2) Latent Channels.** We verify the effect of latent size on compression quality. It can be seen that the representation ability of latent space increases with the size

of the latent channel. However, when the latent channel reaches 32, although the compression quality is high, the subsequent generation model is more difficult to learn. Considering all the above, we finally use a latent size of 16. **3) Model Scale.** Similar to the video VAE [37], we explore the impact of VAE model size on model performance, which is also shown in Tab. 3 (a).

DiT. The DiT ablations are reported in Tab. 3 (b), encompassing 4 aspects: **1) Model Architecture.** We conduct ablations between the Cross-DiT and MM-DiT, as can be seen from the results, for long sequence conditional feature *e.g.*, image, MM-DiT can better model the relation between conditional features and latents. **2) Condition Encoder.** We find that in image-to-3D Human Gaussians, the condition encoder matters, we report the results when taking CLIP [55] and DINOv2 [46] as the condition encoder, com-

	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow		PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow
$\chi \Theta$	29.04	0.030	0.950	Cross.	24.72	0.052	0.933
χ ini.	28.92	0.034	0.947	CLIP	24.07	0.071	0.927
LC-4	25.14	0.042	0.936	DINOv2	24.95	0.048	0.933
LC-8	28.07	0.037	0.944	Flow	25.38	0.046	0.937
LC-32	29.78	0.026	0.954	200M	23.73	0.070	0.931
146M	29.14	0.028	0.949	1.1B	24.70	0.060	0.936
Ours	29.32	0.028	0.953	Ours	25.45	0.044	0.939

(a)

(b)

Table 3. (a). Metrics of VAE ablations, where Θ is the Conv-Atten Dual branch, ini. denotes the UV-map initialization (Sec. 4.2). LC is the latent channel. (b). Metrics of image-to-3D human Gaussians ablations, where Cross. is the Cross-DiT, and Flow denotes training with rectified flow.



Figure 5. The loose clothing, e.g., skirts, demonstrate nature deformation under novel poses.

pared to Sapiens [26], they all lead to lower model performance. **3) Training Objective.** In our experiments, the difference of the final results between the DDPM v-prediction and rectified flow [41] is quite small. The difference in metrics may be due to the sampling strategy. However, we observe that DDPM converges faster in the early stage of scratch training (please refer to the Supp.). Finally, we use DDPM v-prediction for scaling up. **4) Model Scale.** Referring to the scaling-law curve of the MM-DiT architecture given in hunyuan video [28], we roughly calculate the scale of the model based on the amount of data token and train models of three scales, shown in Tab. 3 (b), as the model scale increases, the performance consistently improves.

5.5. Analysis.

Loose clothing deformation. Surprisingly, the model learns the correlation between pose and loose clothing deformation from a large number of diverse poses. As can be seen from Fig. 5, given an image of a person wearing loose clothing with a novel pose, we can observe that the fluffy clothing deforms naturally with the pose. This suggests that data-driven is a potential way to address this problem, and demonstrate the quality of large-scale generative models supported by large amounts of data.



Figure 6. The same 3D human Gaussian could be driven by different poses and maintain consistency.

Pose-Driven Our method decouples the learning of appearance and human pose, which allows the generated human Gaussian to be driven by different poses, and the original person ID can still maintain consistency under different poses. As shown in Fig. 6, multiple people can maintain appearance consistency under multiple poses, including translation, rotation, and movement of distinct body parts.

6. Conclusion

We propose the large-scale generation paradigm for 3D human digitization. To achieve this, we convert multi-view human datasets and synthetic data into a unified 3D Gaussian representation, constructing the HGS-1M dataset that contains 1 million 3D human assets. Building upon this, we employ a UV-structured latent representation to model high-quality Gaussian latents. Subsequently, we adopt an MM-DiT architecture to learn the transfer from conditional features to structured latents, enabling the digitization of high-quality 3D humans through generation.

Limitation and future work. Currently, we train the model with a single frame, which enables us to control the human Gaussian pose, but we have not yet learned the temporal 3D Gaussians directly from the data. One of the future works could be to further generate temporal 3D Gaussians based on our HGS-1M. Additionally, multimodal control generation could be explored to advance 3D human digitization towards multi-modality controllable synthesis based on our pre-trained model.

Acknowledgement This work is supported by the National Natural Science Foundation of China (NSFC) under Grants 62306295, 62225207 and 62436008 and Shanghai Artificial Intelligence Laboratory.

References

- [1] Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025. 2, 3
- [2] Zhongang Cai, Daxuan Ren, Ailing Zeng, Zhengyu Lin, Tao Yu, Wenjia Wang, Xiangyu Fan, Yang Gao, Yifan Yu, Liang Pan, et al. Humman: Multi-modal 4d human dataset for versatile sensing and modeling. In *European Conference on Computer Vision*, pages 557–577. Springer, 2022. 2, 3
- [3] Jinnan Chen, Chen Li, Jianfeng Zhang, Lingting Zhu, Buzhen Huang, Hanlin Chen, and Gim Hee Lee. Generalizable human gaussians from single-view image. *arXiv preprint arXiv:2406.06050*, 2024. 2, 3
- [4] Junhao Chen, Mingjin Chen, Jianjin Xu, Xiang Li, Junting Dong, Mingze Sun, Puhua Jiang, Hongxiang Li, Yuhang Yang, Hao Zhao, et al. Dancetogether! identity-preserving multi-person interactive video generation. *arXiv preprint arXiv:2505.18078*, 2025. 2
- [5] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Januspro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025. 2
- [6] Wei Cheng, Ruixiang Chen, Siming Fan, Wanqi Yin, Keyu Chen, Zhongang Cai, Jingbo Wang, Yang Gao, Zhengming Yu, Zhengyu Lin, et al. Dna-rendering: A diverse neural actor repository for high-fidelity human-centric rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19982–19993, 2023. 2, 3
- [7] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53, 2024. 6
- [8] Junting Dong, Qi Fang, Yudong Guo, Sida Peng, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Totalselfscan: Learning full-body avatars from self-portrait videos of faces, hands, and bodies. In *Advances in Neural Information Processing Systems*, 2022. 3
- [9] Junting Dong, Qi Fang, Tianshuo Yang, Qing Shuai, Chengyu Qiao, and Sida Peng. ivs-net: learning human view synthesis from internet videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22942–22951, 2023. 3
- [10] Junting Dong, Qi Fang, Zehuan Huang, Xudong Xu, Jingbo Wang, Sida Peng, and Bo Dai. Tela: Text to layer-wise 3d clothed human generation. In *European Conference on Computer Vision*, pages 19–36. Springer, 2024. 3
- [11] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis, 2024. URL <https://arxiv.org/abs/2403.03206>, 2. 5
- [12] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. 2, 5
- [13] Ke Fan, Junshu Tang, Weijian Cao, Ran Yi, Moran Li, Jingyu Gong, Jiangning Zhang, Yabiao Wang, Chengjie Wang, and Lizhuang Ma. Freemotion: A unified framework for number-free text-to-motion synthesis. In *European Conference on Computer Vision*, pages 93–109. Springer, 2024. 3
- [14] Ke Fan, Shunlin Lu, Minyu Dai, Runyi Yu, Lixing Xiao, Zhiyang Dou, Junting Dong, Lizhuang Ma, and Jingbo Wang. Go to zero: Towards zero-shot motion generation with million-scale data. *arXiv preprint arXiv:2507.07095*, 2025. 3
- [15] Chaoqun Gong, Yuqin Dai, Ronghui Li, Achun Bao, Jun Li, Jian Yang, Yachao Zhang, and Xiu Li. Text2avatar: text to 3d human avatar generation with codebook-driven body controllable attribute. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 16–20. IEEE, 2024. 3
- [16] Sang-Hun Han, Min-Gyu Park, Ju Hong Yoon, Ju-Mi Kang, Young-Jae Park, and Hae-Gon Jeon. High-fidelity 3d human digitization from single 2k resolution images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12869–12879, 2023. 2, 3, 4
- [17] Paul Henderson, Melonie de Almeida, Daniela Ivanova, and Titas Anciukevičius. Sampling 3d gaussian scenes in seconds with latent diffusion models. *arXiv preprint arXiv:2406.13099*, 2024. 4
- [18] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 6
- [19] Li Hu, Xin Gao, Peng Zhang, Ke Sun, Bang Zhang, and Liefeng Bo. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. *arXiv preprint arXiv:2311.17117*, 2023. 2
- [20] Liangxiao Hu, Hongwen Zhang, Yuxiang Zhang, Boyao Zhou, Boning Liu, Shengping Zhang, and Liqiang Nie. 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*, pages 634–644, 2024. 2
- [21] Shoukang Hu, Fangzhou Hong, Liang Pan, Haiyi Mei, Lei Yang, and Ziwei Liu. Sherf: Generalizable human nerf from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9352–9364, 2023. 3
- [22] Yangyi Huang, Hongwei Yi, Weiyang Liu, Haofan Wang, Boxi Wu, Wenxiao Wang, Binbin Lin, Debing Zhang, and Deng Cai. One-shot implicit animatable avatars with model-based priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8974–8985, 2023. 3
- [23] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. 3
- [24] Mustafa İşık, Martin Rünz, Markos Georgopoulos, Taras Khakhulin, Jonathan Starck, Lourdes Agapito, and Matthias

- Nießner. Humanrf: High-fidelity neural radiance fields for humans in motion. *ACM Transactions on Graphics (TOG)*, 42(4):1–12, 2023. 3
- [25] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 2, 4
- [26] Rawal Khirodkar, Timur Bagautdinov, Julieta Martinez, Su Zhaoen, Austin James, Peter Selednik, Stuart Anderson, and Shunsuke Saito. Sapiens: Foundation for human vision models. In *European Conference on Computer Vision*, pages 206–228. Springer, 2024. 2, 6, 8
- [27] Nikos Kolotouros, Thiemo Alldieck, Andrei Zanfir, Eduard Gabriel Bazavan, Mihai Fieraru, and Cristian Sminchisescu. Dreamhuman: Animatable 3d avatars from text. 2023. 3
- [28] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhou Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, Kathrina Wu, Qin Lin, Aladdin Wang, Andong Wang, Changlin Li, Duojuan Huang, Fang Yang, Hao Tan, Hongmei Wang, Jacob Song, Jiawang Bai, Jianbing Wu, Jinbao Xue, Joey Wang, Junkun Yuan, Kai Wang, Mengyang Liu, Pengyu Li, Shuai Li, Weiyan Wang, Wenqing Yu, Xincheng Deng, Yang Li, Yanxin Long, Yi Chen, Yutao Cui, Yuanbo Peng, Zhentao Yu, Zhiyu He, Zhiyong Xu, Zixiang Zhou, Zunnan Xu, Yangyu Tao, Qinglin Lu, Songtao Liu, Dax Zhou, Hongfa Wang, Yong Yang, Di Wang, Yuhong Liu, Jie Jiang, and Caesar Zhong. Hunyuanvideo: A systematic framework for large video generative models, 2024. 2, 3, 5, 8
- [29] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. *Advances in Neural Information Processing Systems*, 34:24741–24752, 2021. 2, 3
- [30] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural image-based avatars: Generalizable radiance fields for human avatar modeling. *arXiv preprint arXiv:2304.04897*, 2023. 3
- [31] Youngjoong Kwon, Lingjie Liu, Henry Fuchs, Marc Habermann, and Christian Theobalt. Deliffas: Deformable light fields for fast avatar synthesis. *Advances in neural information processing systems*, 36:40944–40962, 2023. 3
- [32] Youngjoong Kwon, Baole Fang, Yixing Lu, Haoye Dong, Cheng Zhang, Francisco Vicente Carrasco, Albert Mosella-Montoro, Jianjin Xu, Shingo Takagi, Daeil Kim, et al. Generalizable human gaussians for sparse view synthesis. In *European Conference on Computer Vision*, pages 451–468. Springer, 2024. 2, 3, 4, 6
- [33] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 3
- [34] Yushi Lan, Shangchen Zhou, Zhaoyang Lyu, Fangzhou Hong, Shuai Yang, Bo Dai, Xingang Pan, and Chen Change Loy. Gaussiananything: Interactive point cloud latent diffusion for 3d generation. In *ICLR*, 2025. 4, 6
- [35] Hongxiang Li, Yaowei Li, Yuhang Yang, Junjie Cao, Zhihong Zhu, Xuxin Cheng, and Long Chen. Dispose: Disentangling pose guidance for controllable human image animation. *arXiv preprint arXiv:2412.09349*, 2024. 2
- [36] Peng Li, Wangguandong Zheng, Yuan Liu, Tao Yu, Yangguang Li, Xingqun Qi, Mengfei Li, Xiaowei Chi, Siyu Xia, Wei Xue, et al. Pshuman: Photorealistic single-view human reconstruction using cross-scale diffusion. *arXiv preprint arXiv:2409.10141*, 2024. 2, 3
- [37] Zongjian Li, Bin Lin, Yang Ye, Liuhan Chen, Xinhua Cheng, Shenghai Yuan, and Li Yuan. Wf-vae: Enhancing video vae by wavelet-driven energy flow for latent video diffusion model. *arXiv preprint arXiv:2411.17459*, 2024. 7
- [38] Zhe Li, Zerong Zheng, Lizhen Wang, and Yebin Liu. 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. 2, 3
- [39] Chenguo Lin, Panwang Pan, Bangbang Yang, Zeming Li, and Yadong Mu. Diffsplat: Repurposing image diffusion models for scalable gaussian splat generation. *arXiv preprint arXiv:2501.16764*, 2025. 6
- [40] Bingchen Liu, Ehsan Akhgari, Alexander Visheratin, Aleks Kamko, Linmiao Xu, Shivam Shrirao, Chase Lambert, Joao Souza, Suhail Doshi, and Daiqing Li. Playground v3: Improving text-to-image alignment with deep-fusion large language models. *arXiv preprint arXiv:2409.10695*, 2024. 3
- [41] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 8
- [42] Xian Liu, Xiaohang Zhan, Jiaxiang Tang, Ying Shan, Gang Zeng, Dahua Lin, Xihui Liu, and Ziwei Liu. Humangaussian: Text-driven 3d human generation with gaussian splatting. *arXiv preprint arXiv:2311.17061*, 2023. 3
- [43] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [44] Yixing Lu, Junting Dong, Youngjoong Kwon, Qin Zhao, Bo Dai, and Fernando De la Torre. Gas: Generative avatar synthesis from a single image. *arXiv preprint arXiv:2502.06957*, 2025. 3
- [45] Yifang Men, Yuan Yao, Miaomiao Cui, and Liefeng Bo. Mimo: Controllable character video synthesis with spatial decomposed modeling. *arXiv preprint arXiv:2409.16160*, 2024. 2
- [46] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 7
- [47] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022. 3
- [48] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings IEEE Conf.*

- on *Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019. 2, 4
- [49] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023. 2
- [50] Hao-Yang Peng, Jia-Peng Zhang, Meng-Hao Guo, Yan-Pei Cao, and Shi-Min Hu. Charactergen: Efficient 3d character generation from single images with multi-view pose canonicalization. *ACM Transactions on Graphics (TOG)*, 43(4): 1–13, 2024. 3
- [51] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14314–14323, 2021. 3
- [52] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 3
- [53] Yunchen Pu, Zhe Gan, Ricardo Henao, Xin Yuan, Chunyuan Li, Andrew Stevens, and Lawrence Carin. Variational autoencoder for deep learning of images, labels and captions. *Advances in neural information processing systems*, 29, 2016. 2
- [54] Lingteng Qiu, Xiaodong Gu, Peihao Li, Qi Zuo, Weichao Shen, Junfei Zhang, Kejie Qiu, Weihao Yuan, Guanying Chen, Zilong Dong, and Liefeng Bo. Lhm: Large animatable human reconstruction model for single image to 3d in one second. 2025. 2, 3
- [55] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 7
- [56] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2304–2314, 2019. 3
- [57] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. 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*, pages 84–93, 2020. 3
- [58] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022. 6
- [59] Vincent Sitzmann, Semon Rezchikov, Bill Freeman, Josh Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. *Advances in Neural Information Processing Systems*, 34: 19313–19325, 2021. 4
- [60] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 6
- [61] Mingze Sun, Junhao Chen, Junting Dong, Yurun Chen, Xinyu Jiang, Shiwei Mao, Puhua Jiang, Jingbo Wang, Bo Dai, and Ruqi Huang. Drive: Diffusion-based rigging empowers generation of versatile and expressive characters. *arXiv preprint arXiv:2411.17423*, 2024. 3
- [62] Jiayang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. In *European Conference on Computer Vision*, pages 1–18. Springer, 2024. 3, 6
- [63] Jiayang Tang, Ruijie Lu, Xiaokang Chen, Xiang Wen, Gang Zeng, and Ziwei Liu. Intex: Interactive text-to-texture synthesis via unified depth-aware inpainting. *arXiv preprint arXiv:2403.11878*, 2024. 4
- [64] Junshu Tang, Yanhong Zeng, Ke Fan, Xuheng Wang, Bo Dai, Kai Chen, and Lizhuang Ma. Make-it-vivid: Dressing your animatable biped cartoon characters from text. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6243–6253, 2024. 3
- [65] Genmo Team. Mochi 1. <https://github.com/genmoai/models>, 2024. 2, 3
- [66] Kolors Team. Kolors: Effective training of diffusion model for photorealistic text-to-image synthesis. *arXiv preprint*, 2024. 3
- [67] Tencent Hunyuan3D Team. Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation, 2025. 3
- [68] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. 3
- [69] Zilong Wang, Zhiyang Dou, Yuan Liu, Cheng Lin, Xiao Dong, Yunhui Guo, Chenxu Zhang, Xin Li, Wenping Wang, and Xiaohu Guo. Wonderhuman: Hallucinating unseen parts in dynamic 3d human reconstruction, 2025. 2
- [70] Zhenzhen Weng, Jingyuan Liu, Hao Tan, Zhan Xu, Yang Zhou, Serena Yeung-Levy, and Jimei Yang. Template-free single-view 3d human digitalization with diffusion-guided lrm. *arXiv preprint arXiv:2401.12175*, 2024. 2, 3
- [71] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. *arXiv preprint arXiv:2412.01506*, 2024. 3
- [72] Lixing Xiao, Shunlin Lu, Huaijin Pi, Ke Fan, Liang Pan, Yueer Zhou, Ziyong Feng, Xiaowei Zhou, Sida Peng, and Jingbo Wang. Motionstreamer: Streaming motion generation via diffusion-based autoregressive model in causal latent space. *arXiv preprint arXiv:2503.15451*, 2025. 3
- [73] Zhangyang Xiong, Chenghong Li, Kenkun Liu, Hongjie Liao, Jianqiao Hu, Junyi Zhu, Shuliang Ning, Lingteng Qiu, Chongjie Wang, Shijie Wang, et al. Mvhumannet: A large-scale dataset of multi-view daily dressing human captures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19801–19811, 2024. 2, 3

- [74] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J Black. Icon: Implicit clothed humans obtained from normals. in 2022 *iecc*. In *CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13286–13296, 2022. 2
- [75] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J Black. Econ: Explicit clothed humans optimized via normal integration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 512–523, 2023. 2
- [76] Haitao Yang, Yuan Dong, Hanwen Jiang, Dejie Xu, Georgios Pavlakos, and Qixing Huang. Atlas gaussians diffusion for 3d generation. *arXiv preprint arXiv:2408.13055*, 2024. 4
- [77] Yuhang Yang, Wei Zhai, Hongchen Luo, Yang Cao, and Zheng-Jun Zha. Lemon: Learning 3d human-object interaction relation from 2d images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16284–16295, 2024. 2
- [78] Yuhang Yang, Wei Zhai, Chengfeng Wang, Chengjun Yu, Yang Cao, and Zheng-Jun Zha. Egochoir: Capturing 3d human-object interaction regions from egocentric views. *arXiv preprint arXiv:2405.13659*, 2024. 2
- [79] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 5, 6
- [80] Chengjun Yu, Wei Zhai, Yuhang Yang, Yang Cao, and Zheng-Jun Zha. Hero: Human reaction generation from videos. *arXiv preprint arXiv:2503.08270*, 2025. 3
- [81] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4d: Real-time human volumetric capture from very sparse consumer rgbd sensors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5746–5756, 2021. 2, 3, 4
- [82] Zhixuan Yu, Jae Shin Yoon, In Kyu Lee, Prashanth Venkatesh, Jaesik Park, Jihun Yu, and Hyun Soo Park. Humbi: A large multiview dataset of human body expressions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2990–3000, 2020. 3
- [83] Ailing Zeng, Yuhang Yang, Weidong Chen, and Wei Liu. The dawn of video generation: Preliminary explorations with sora-like models. *arXiv preprint arXiv:2410.05227*, 2024. 2
- [84] Biao Zhang and Rico Sennrich. Root mean square layer normalization. *Advances in Neural Information Processing Systems*, 32, 2019. 6
- [85] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Yong Zhang, Hongwei Zhao, Hongtao Lu, Xi Shen, and Ying Shan. Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14730–14740, 2023. 3
- [86] Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan Xu, and Jingyi Yu. Clay: A controllable large-scale generative model for creating high-quality 3d assets. *ACM Transactions on Graphics (TOG)*, 43(4):1–20, 2024. 3
- [87] Weitian Zhang, Yichao Yan, Yunhui Liu, Xingdong Sheng, and Xiaokang Yang. e³gen: Efficient, expressive and editable avatars generation. *arXiv preprint arXiv:2405.19203*, 2024. 3, 4
- [88] Zechuan Zhang, Li Sun, Zongxin Yang, Ling Chen, and Yi Yang. Global-correlated 3d-decoupling transformer for clothed avatar reconstruction. *Advances in Neural Information Processing Systems*, 36:7818–7830, 2023. 2, 3
- [89] Zechuan Zhang, Zongxin Yang, and Yi Yang. 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*, pages 9936–9947, 2024. 2, 3, 6
- [90] Fuqiang Zhao, Wei Yang, Jiakai Zhang, Pei Lin, Yingliang Zhang, Jingyi Yu, and Lan Xu. Humannerf: Efficiently generated human radiance field from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7743–7753, 2022. 3
- [91] Shunyuan Zheng, Boyao Zhou, Ruizhi Shao, Boning Liu, Shengping Zhang, Liqiang Nie, and Yebin Liu. 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*, pages 19680–19690, 2024. 3
- [92] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3d human reconstruction from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7739–7749, 2019. 2, 3
- [93] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *IEEE transactions on pattern analysis and machine intelligence*, 44(6): 3170–3184, 2021. 3
- [94] Yiyu Zhuang, Jiayi Lv, Hao Wen, Qing Shuai, Ailing Zeng, Hao Zhu, Shifeng Chen, Yujiu Yang, Xun Cao, and Wei Liu. Idol: Instant photorealistic 3d human creation from a single image. *arXiv preprint arXiv:2412.14963*, 2024. 2, 3, 4