

# SVAD: From Single Image to 3D Avatar via Synthetic Data Generation with Video Diffusion and Data Augmentation

Anonymous CVPR Workshop SyntaGen submission

Paper ID 16




Figure 1. **SVAD**. We present a novel pipeline that leverages video diffusion models and data augmentation methods to generate high-quality synthetic training data from a single human image. This synthetic data generation approach enables us to train 3D Gaussian Splatting avatars with significantly improved fidelity, outperforming state-of-the-art single-image avatar creation methods while preserving identity and fine details across novel poses and viewpoints.

## Abstract

001 *Creating animatable 3D human avatars from a single image*  
002 *remains a significant challenge with applications in virtual*  
003 *reality and human-centered AI. Traditional 3D Gaussian*  
004 *Splatting (3DGS) methods produce high-quality avatars but*  
005 *require monocular video sequences or multi-view inputs,*  
006 *while video diffusion models can animate from static images*  
007 *but struggle with temporal coherence and identity preser-*  
008 *vation. We present SVAD, a novel framework for synthetic*  
009 *data generation and avatar creation that addresses these*  
010 *limitations. SVAD leverages video diffusion models to gen-*  
011 *erate an initial set of synthetic pose-conditioned anima-*  
012 *tions from a single image, then enhances this synthetic data*  
013 *through identity preservation and image restoration mod-*  
014 *ules. This high-quality synthetic dataset enables training of*  
015 *3DGS avatar models that maintain subject fidelity and fine*  
016 *details across diverse poses and viewpoints. Our approach*  
017 *combines the generative capabilities of diffusion models*  
018 *with the rendering efficiency of 3DGS, resulting in state-of-*

*the-art performance in single-image avatar creation. Ex-*  
*periments demonstrate that SVAD’s synthetic data gener-*  
*ation pipeline significantly improves temporal stability and*  
*identity consistency compared to existing methods, while*  
*enabling real-time rendering for interactive applications.*

019  
020  
021  
022  
023

## 1. Introduction

The ability to generate animatable 3D human avatars from minimal input data, such as a single image, has significant potential across a range of applications. Traditional methods, particularly those based on 3DGS, have demonstrated considerable success in producing high-quality avatars [9, 19, 37, 38, 44, 51, 52, 56, 67, 74]. These methods rely on dense input data, typically monocular or multi-view video [9, 19, 37, 44, 51, 56, 74], to achieve high fidelity across varied viewpoints and poses. This reliance on extensive video input complicates deployment in single-image scenarios, where ensuring viewpoint consistency and adapt-

024  
025  
026  
027  
028  
029  
030  
031  
032  
033  
034  
035

036 ability to novel poses becomes a key challenge.

037 Recent advancements in video diffusion models offer a  
038 potential solution by enabling animation generation from a  
039 single static image [18, 59, 62, 75]. These models use pose-  
040 conditioned diffusion processes to create video sequences,  
041 demonstrating the powerful generative capabilities of dif-  
042 fusion for single-image-driven animation. However, diffu-  
043 sion models often struggle to maintain temporal coherence,  
044 leading to inconsistent features and identity drift across  
045 frames. Additionally, their iterative denoising process for  
046 each frame introduces significant computational overhead,  
047 limiting their feasibility for real-time or interactive applica-  
048 tions where rapid rendering across novel views is essential.

049 To overcome these challenges, we propose SVAD, a  
050 novel synthetic data generation and avatar creation pipeline  
051 that synergizes the generative flexibility of diffusion mod-  
052 els with the efficient rendering capabilities of 3DGS avatars.  
053 Our approach leverages video diffusion models to generate  
054 diverse pose-conditioned synthetic training data from a single  
055 image. This synthetic data is refined through an identity-  
056 preservation module and an image restoration module to en-  
057 sure that perceptual identity consistency and structural fi-  
058 delity are preserved across diverse poses and temporal se-  
059 quences. The resulting high-quality synthetic dataset is then  
060 used to train a 3DGS-based avatar model [37], which bene-  
061 fits from the rapid rendering capabilities inherent to 3DGS.  
062 By combining the generative strengths of diffusion for syn-  
063 thetic data creation with the efficiency of 3DGS for ren-  
064 dering, SVAD achieves consistent, high-quality 3D avatar  
065 animations from single image input.

066 Our contributions can be summarized as follows:

- 067 • We introduce SVAD, a novel pipeline for generating high-  
068 quality synthetic training data from a single image, en-  
069 abling the creation of detailed and animatable 3D human  
070 avatars.
- 071 • We incorporate an identity-preservation module and an  
072 image restoration technique to refine diffusion-generated  
073 synthetic data, ensuring consistency in identity and fine  
074 details across diverse poses and viewpoints.
- 075 • We demonstrate that our synthetic data generation ap-  
076 proach significantly improves the quality of 3DGS avatars  
077 compared to state-of-the-art single-image methods, while  
078 maintaining efficient real-time rendering capabilities.
- 079 • We provide extensive experiments and evaluations show-  
080 ing that SVAD’s synthetic data-driven approach achieves  
081 superior performance in novel pose adaptation and iden-  
082 tity preservation for single-image avatar creation.

## 083 2. Related Work

084 **Diffusion Model for Human Image Animation** The use  
085 of diffusion models has led to significant advancements in  
086 human image animation, enabling the generation of real-  
087 istic and temporally consistent animations from static im-

088 ages [1, 4, 5, 14, 20, 25, 42, 46, 50, 53, 54, 64, 66, 69, 72].  
089 Early methods, such as PIDM [3] and DreamPose [29], fo-  
090 cused on improving texture fidelity by employing texture  
091 diffusion modules to align texture patterns between refer-  
092 ence and target images. These methods, while enhancing  
093 detail preservation, still face challenges in maintaining tem-  
094 poral stability across frames.

095 Recent works, including DisCo [59] and Animate Any-  
096 one [18], have extended diffusion models to improve tem-  
097 poral consistency and fine-grained control in human anima-  
098 tion tasks. DisCo leverages dual ControlNets [68] to sep-  
099 arately control pose and background elements, providing  
100 more robust conditioning for complex motion sequences.  
101 Similarly, Animate Anyone integrates a ReferenceNet with  
102 temporal attention layers to ensure appearance consistency  
103 and smooth transitions across frames, thereby addressing  
104 flickering issues commonly observed in earlier models.

105 **Dynamic 3D Gaussian based Avatars** The concept of  
106 Gaussian splatting for 3D avatars has emerged recently  
107 as an innovative approach to explicit scene representa-  
108 tion [30]. This technique models a scene as a collec-  
109 tion of 3D Gaussian elements, each containing photometric  
110 and geometric properties. During rendering, these Gaus-  
111 sian splats are projected onto the image plane, creating  
112 the final rendered output. The efficiency of 3D Gaussian  
113 splatting has been demonstrated in both static [22, 27, 32]  
114 and dynamic [13, 28, 31, 36] scenes, making it a versa-  
115 tile tool for various applications. Recent advancements [8,  
116 12, 21, 24, 33, 43, 44, 58, 76] have explored the use of  
117 3DGS to create photorealistic human avatars across differ-  
118 ent scenarios. These methods commonly rely on multi-view  
119 data [34, 40, 73] or monocular video [19, 24, 33, 37, 44] as  
120 input to achieve high-quality, consistent results. The advan-  
121 tage of 3DGS lies in its ability to produce temporally stable  
122 animated avatars with superior quantitative metrics.

## 123 3. Method

124 To generate high-quality human avatars from just a single  
125 image, facilitating free-viewpoint rendering and real-  
126 istic animation, we integrate the generative capabilities of  
127 video diffusion models with the rendering efficiency of 3D  
128 Gaussian-based avatars. We start by leveraging a pretrained  
129 video diffusion model for character animation to produce  
130 initial synthetic data, as described in Sec. 3.1. Directly  
131 using these frames to train a 3DGS avatar model, how-  
132 ever, often yields poor results, with challenges in preserv-  
133 ing facial identity, clothing details, and maintaining consis-  
134 tent multiview coherence across side and back views. To  
135 address these issues and enhance avatar quality, we intro-  
136 duce a data augmentation pipeline in Sec. 3.2 comprising  
137 identity-preservation and image-restoration modules to re-  
138 fine the diffusion outputs. With the augmented synthetic  
139 data, we proceed to train a 3DGS avatar model, as out-

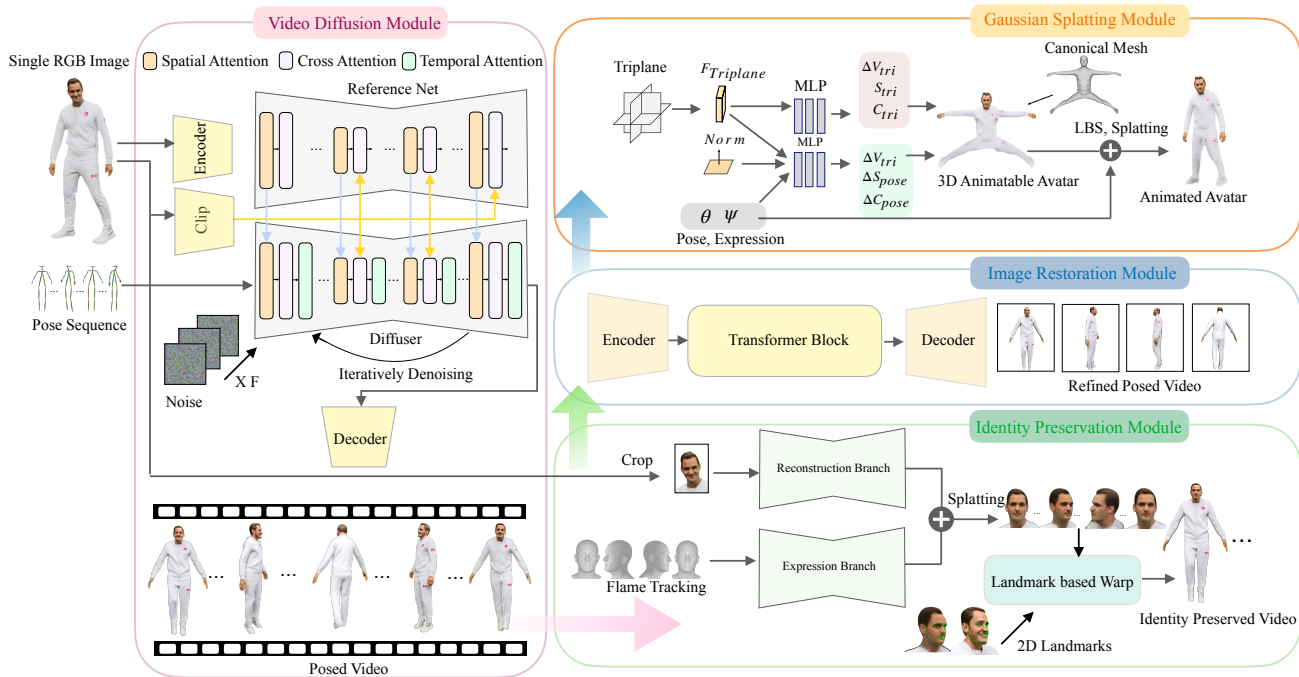


Figure 2. **Overall Pipeline of SVAD.** Starting from a single input image, the diffusion model generates pose-conditioned animations, which are refined using an identity preservation module and an image restoration module. The refined outputs are then used to train the 3DGS avatar, enabling high-fidelity, animatable 3D avatars with consistent details across poses and viewpoints.

140 lined in Sec. 3.3. The following sections detail the technical  
141 methodologies employed in our approach.

### 142 3.1. Video Diffusion Module

143 To generate an animated character video  $V$  from a single  
144 input image  $I$ , we leverage MusePose [57], a finetuned  
145 variant of Animate Anyone [18], which is a state-of-the-  
146 art video diffusion model designed for realistic human  
147 animation while maintaining temporal consistency and appear-  
148 ance fidelity. MusePose employs a U-Net-based diffusion  
149 architecture with integrated pose and temporal controls, al-  
150 lowing for pose-guided animation across frames.

151 The model architecture incorporates several key compo-  
152 nents for effective character animation. The denoising UNet  
153 is implemented as a 3D UNet [11] with motion modules for  
154 temporal coherence. Specifically, we use Vanilla motion  
155 modules with temporal self-attention blocks at resolutions  
156 of [1, 2, 4, 8] and in the mid-block. Each transformer block  
157 contains 8 attention heads, with temporal position encoding  
158 enabling positional awareness across a sequence of up to  
159 128 frames.

160 To incorporate pose guidance, a lightweight Pose Guider  
161 encodes the motion control signal from the predefined 2D  
162 keypoints into a pose-aligned latent representation  $P(p_t) \in$   
163  $\mathbb{R}^{H \times W \times C}$ . For a pose feature  $p_t \in \mathbb{R}^{J \times 2}$  at time  $t$ , where  
164  $J$  is the number of keypoints, we align the encoding to en-  
165 sure continuity between frames by adding this encoded pose

signal to the noise latent  $z_t$ :

$$z_t = z_t + P(p_t) \quad (1)$$

166 For the diffusion process, we adopt a v-prediction [49] for-  
167 mulation with zero-SNR sampling [35], using a scaled lin-  
168 ear beta schedule with  $\beta_{\text{start}} = 0.00085$  and  $\beta_{\text{end}} = 0.012$ .  
169 The DDIM sampler [55] is configured for efficient inference  
170 with 20 sampling steps and a classifier-free guidance scale  
171 of 3.5. The temporal consistency loss  $L_{\text{temp}}$  minimizes dis-  
172 crepancies across successive frames by enforcing coherence  
173 in appearance and pose:  
174  
175

$$L_{\text{temp}} = \sum_{t=1}^{T-1} \|F_t - F_{t+1}\|^2 \quad (2)$$

176 where  $F_t \in \mathbb{R}^{H \times W \times 3}$  represents the RGB frame output at  
177 time  $t$ .  
178

179 A critical challenge in character animation is ensuring  
180 anatomical consistency between the reference image and  
181 the motion poses. Direct application of pose control can  
182 result in unnatural animations due to mismatches in body  
183 proportions. Therefore, we employ a comprehensive pose  
184 alignment procedure that adapts the source pose to match  
185 the reference character’s physical characteristics.  
186

187 Given a reference pose  $P_{\text{ref}}$  and a source pose  $P_{\text{src}}$  de-  
188 tected using DWpose [63], we compute scale parameters  
189  $\mathbf{S} = \{s_1, s_2, \dots, s_{10}\}$  for ten distinct body regions: neck,  
190



Figure 3. **3D Avatars trained by SVAD.** SVAD generates high quality 3D avatars with just a single image. The trained avatars can be rendered from any view point, in any pose.

189 face, shoulders, upper arms, lower arms, hands, torso, upper  
190 legs, and lower legs. For each body part  $i$ , we compute its  
191 scale factor  $s_i$  as the ratio between the corresponding key-  
192 point distances:

$$193 \quad s_i = \frac{d_i^{ref}}{d_i^{src}} \quad (3)$$

194 where  $d_i^{ref}$  and  $d_i^{src}$  represent the Euclidean distances be-  
195 tween keypoints. For body parts with bilateral symmetry  
196 (e.g., arms), we average the scales from both sides:

$$197 \quad s_{arm\_upper} = \frac{1}{2} \left( \frac{\|p_{ref}^2 - p_{ref}^3\|}{\|p_{src}^2 - p_{src}^3\|} + \frac{\|p_{ref}^5 - p_{ref}^6\|}{\|p_{src}^5 - p_{src}^6\|} \right) \quad (4)$$

198 To apply these scales to the source pose, we use a ro-  
199 tation matrix transformation centered at anchor points spe-  
200 cific to each body part:

$$201 \quad p' = c_i + s_i \cdot (p - c_i) \quad (5)$$

202 where  $c_i$  is the anchor center for part  $i$ . This hierarchical ap-  
203 proach ensures body proportions match the reference while  
204 maintaining the overall pose structure.

### 3.2. Data Augmentation Module

205  
206 Training the 3DGS model using only outputs from the video  
207 diffusion model often results in low-fidelity avatars, par-  
208 ticularly in terms of facial details and high-frequency fea-  
209 tures like hands and clothing. To address these challenges,  
210 we introduce a data augmentation module that enhances the  
211 quality of the training data. This module includes an iden-  
212 tity preservation sub-module ensuring coherence in facial  
213 details across frames and a image restoration submodule  
214 which refines texture quality and high-frequency details, re-  
215 sulting in more realistic textures. This comprehensive data  
216 augmentation significantly improves the synthetic training  
217 data, enabling the 3DGS Avatar model integrated in the fu-  
218 ture to generate more realistic and detailed 3D avatars.

**Identity preservation sub-module.** To ensure consistent  
219 and realistic facial details across frames, we implement  
220 an identity preservation module that combines head recon-  
221 struction and facial fusion techniques. This module lever-  
222 ages a 3DGS-based head reconstruction method inspired by  
223 Chu *et al.* [10] to create a 3D Gaussian-based head avatar  
224

225 from a single input image using a novel *dual-lifting* ap-  
226 proach that predicts both forward and backward lifting dis-  
227 tances.

228 Given an input image  $I_s$ , global and local features  $F_{\text{local}}$   
229 are extracted using a frozen DINOv2 [39] backbone. These  
230 features are used to predict forward and backward lifting  
231 distances, positioning 3D Gaussians  $G_{\text{pos}}$  as follows:

$$232 \quad G_{\text{pos}} = [\mathbf{p}_s + E_{\text{Conv0}}(F_{\text{local}}) \cdot \mathbf{n}_s, \mathbf{p}_s - E_{\text{Conv1}}(F_{\text{local}}) \cdot \mathbf{n}_s], \quad (6)$$

233 where  $\mathbf{p}_s$  is the initial point plane,  $\mathbf{n}_s$  is the normal vector,  
234 and  $E_{\text{Conv}}$  are convolutional layers predicting offsets. To  
235 capture expression variations, we bind 3DMM features:

$$236 \quad G_{\text{expr}} = \text{MLP}(F_{\text{3DMM}} + F_{\text{global}}). \quad (7)$$

237 After generating the head avatar renderings, we detect  
238 facial landmarks on both the original frame  $I_{\text{orig}}$  and the  
239 generated head image  $I_{\text{head}}$ , compute an affine transforma-  
240 tion for alignment, and use Poisson image editing [41] for  
241 seamless fusion:

$$242 \quad \min_I \int_{\Omega} \|\nabla I - \nabla I_{\text{warp}}\|^2 dx dy, \quad \text{subject to } I|_{\partial\Omega} = I_{\text{orig}}|_{\partial\Omega}, \quad (8)$$

243 where  $\Omega$  is defined by the facial mask, ensuring temporally  
244 consistent facial details throughout the animation.

245 **Image restoration sub-module.** Finally, to preserve qual-  
246 ity of fine detailed regions, we employ an image restoration  
247 module based on the work of Chen *et al.* [7], specifically  
248 their diffusion-based image restoration method. BFRffu-  
249 sion leverages the generative prior encapsulated in the pre-  
250 trained Stable Diffusion model [47] to enhance image de-  
251 tails through a comprehensive architecture that effectively  
252 extracts features from low-quality images and restores real-  
253 istic facial details.

254 In our method, we set the super-resolution scale factor to  
255 2, enhancing input frames while maintaining computational  
256 efficiency. The process uses 50 DDIM sampling steps with  
257 a classifier-free guidance scale of 3.5, achieving a balance  
258 between restoration quality and processing speed. For face  
259 regions, the method employs a face restoration helper with  
260 facial landmark detection to specifically enhance facial de-  
261 tails, ensuring identity consistency across generated frames.

262 This image restoration submodule significantly improves  
263 the fidelity and realism of our synthetic training data by  
264 restoring fine facial details, enhancing texture quality in  
265 clothing and accessories, and improving overall image co-  
266 herence. The refined synthetic data enables the 3DGS  
267 Avatar to learn more accurate representations with consis-  
268 tent high-frequency details that persist across poses and  
269 viewpoints.

### 3.3. 3D Human Gaussian Splatting Module 270

271 We apply the architecture of a 3DGS based avatar method  
272 introduced by Moon *et al.* [37], which integrates the SMPL-  
273 X model with a 3D Gaussian-based representation to pro-  
274 duce animatable human avatars. Each 3D Gaussian acts as  
275 a vertex connected by a pre-defined mesh topology follow-  
276 ing SMPL-X. This hybrid representation combines the ex-  
277 pressive surface modeling of SMPL-X with the flexibility  
278 of a volumetric approach, allowing for smooth interpolation  
279 across the body surface essential for realistic animations.

280 Each Gaussian point is associated with positional data  
281  $\mathbf{V} \in \mathbb{R}^{N \times 3}$ , RGB color values  $\mathbf{C} \in \mathbb{R}^{N \times 3}$ , and a scale  
282 parameter  $\mathbf{S} \in \mathbb{R}^N$ , where  $N$  is the number of Gaussians.  
283 The Gaussian splatting rendering equation is:

$$284 \quad I = f(V, \exp(S), C, K, E), \quad (9)$$

285 where  $V$  represents positions,  $S$  denotes scale,  $C$  colors,  
286 and  $K$  and  $E$  camera parameters.

287 Pose-dependent deformations are applied through an  
288 MLP network, predicting offsets for each Gaussian based  
289 on SMPL-X pose parameters:

$$290 \quad \mathbf{V}_{\text{pose}} = \mathbf{V} + \Delta\mathbf{V}_{\text{pose}} + \Delta\mathbf{V}_{\text{expr}}. \quad (10)$$

291 To maintain spatial coherence, a Laplacian regularizer mini-  
292 mizes the difference between the Laplacian of the canonical  
293 mesh and the deformed Gaussian points:

$$294 \quad L_{\text{Lap}} = \|\Delta\mathbf{V}_{\text{canonical}} - \Delta\mathbf{V}_{\text{deformed}}\|^2. \quad (11)$$

295 This approach combined with our augmented synthetic data  
296 achieves highly realistic, animatable avatars capable of real-  
297 time rendering with smooth deformations across facial ex-  
298 pressions, body movements, and hand gestures.

## 4. Experiments 299

### 4.1. Datasets and Metrics 300

301 **People-Snapshot** We use the People-Snapshot dataset[2],  
302 which contains videos of individuals rotating in front of  
303 a stationary camera. For consistency and fair compari-  
304 son, we adhere to the evaluation protocol established by  
305 InstantAvatar[26].

306 **THuman** The PeopleSnapshot dataset has limited pose  
307 variations. To evaluate the performance on more challeng-  
308 ing test poses, we utilize the THuman dataset[65]. This  
309 dataset includes a diverse set of poses that require higher  
310 flexibility and adaptability.

311 **Evaluation Metrics.** We consider four metrics, PSNR[15],  
312 SSIM[60], LPIPS[70], and Clip Similarity[45](denoted as  
313 CLIP in tables) to assess the reconstruction quality on three  
314 datasets.

Method	Female-4-casual			Male-3-casual			Female-3-casual			Male-4-casual		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
HumanNeRF [61]	27.07	0.9615	0.0151	26.90	0.9605	0.0181	24.46	0.9516	0.0269	25.50	0.9397	0.0357
GaussianAvatar [19]	30.84	0.9771	0.0140	30.98	0.9790	0.0145	29.55	0.9762	0.0220	28.78	0.9755	0.0230
ExAvatar [37]	30.98	0.9789	0.0333	29.75	0.9628	0.0402	29.74	0.9678	0.0458	28.89	0.9666	0.0500
ExAvatar [37] (Single Image)	20.42	0.9427	0.0656	23.24	0.9448	0.0562	20.12	0.9492	0.0543	23.74	0.9497	0.0610
Ours (Single Image)	21.51	0.9442	0.0528	22.54	0.9467	0.0484	21.96	0.9609	0.0541	23.71	0.9570	0.0592

Table 1. **Quantitative Evaluation on the People-Snapshot [2] Dataset.** Our approach demonstrates superior performance on *single-image* input, outperforming the baseline on most of the metrics. The top two results for *single-image* input are highlighted in **first** and **second**, with the overall best result highlighted in **first**. Note that methods without the *Single Image* utilize approximately 200 input frames.



Figure 4. **Synthetic Data Generated by SVAD.** Comparison of the original input image (left), unprocessed synthetic data from video diffusion (middle), and our enhanced synthetic data after applying identity preservation and image restoration modules (right). Note how our data augmentation approach preserves facial identity and significantly improves detail quality in the synthetic training data.

## 315 4.2. Quantitative Evaluation

316 We quantitatively evaluate the quality of single-image 3D  
317 avatars generated by our method against SOTA 3D avatar  
318 generation methods [19, 37, 61]. While current 3D avatar  
319 models generally require a monocular video as input, we  
320 assess our model’s performance using a single image as in-  
321 put on ExAvatar [37]. Additionally, we report results us-  
322 ing the original full training set of approximately 200 input  
323 frames for monocular input based avatar models for refer-  
324 ence. As shown in Table 1, our model achieves highest

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
PIFu[48]	15.62	0.8921	0.1903	0.8612
TeCH[23]	15.85	0.8892	0.1667	0.8890
Ultraman[6]	18.13	0.9019	0.1334	0.9089
SIFU[71]	18.59	0.8591	0.1402	0.8873
SITH[17]	19.98	0.9018	0.1294	0.9084
<b>Ours</b>	<b>20.92</b>	<b>0.9291</b>	<b>0.1124</b>	<b>0.9321</b>

Table 2. **Quantitative Evaluation.** SVAD compared to recent single-image based 3D human generation SOTAs. Top two results are colored as **first** **second**.

325 scores on most of the metrics among single-image input  
326 methods. We further compare our approach with single-  
327 view 3D human reconstruction methods [6, 17, 23, 48, 71],  
328 many of which employ the SMPL model, allowing for an-  
329 imatability through mesh fitting and reposing techniques,  
330 such as those in Editable Humans [16]. We randomly sam-  
331 ple 100 scans from the THuman dataset and report results.  
332 We repose our trained avatar using ground-truth SMPL-X  
333 parameters, then render four viewpoints (front, left, right,  
334 back) for each model and compare with the ground-truth  
335 scan renderings from the same views. As presented in Ta-  
336 ble 4, our method surpasses all other baselines, demonstrat-  
337 ing superior quality in single-image 3D human reconstruc-  
338 tion tasks.

## 4.3. Qualitative Evaluation

340 Figure 5 shows the overall quality of our generated 3D  
341 avatars from single images in the People Snapshot and the  
342 THuman dataset. Figure 6, Figure 7 shows that our method  
343 performs superior compared to current single-view human  
344 reconstruction SOTA [17]. For single image avatar genera-  
345 tion, we evaluate on the People Snapshot dataset and com-  
346 pare against the 3D-GS based avatar SOTA [37]. For fair-  
347 ness, we train the SOTA with the single input image for the  
348 same amount of iterations (12k iter) as our method. Figure  
349 8 shows that for single image avatar generation, our method  
350 performs superior especially for the back and side views.



Figure 5. **Qualitative Evaluation** on the People Snapshot dataset and of THuman dataset scan renderings. From a single image input, SVAD generates high-quality, animatable 3D Avatars.




Figure 6. **Qualitative Comparison with SITH [17]**. Our approach better reconstructs complex contours and subtle features, resulting in a more lifelike and coherent side-view appearance.

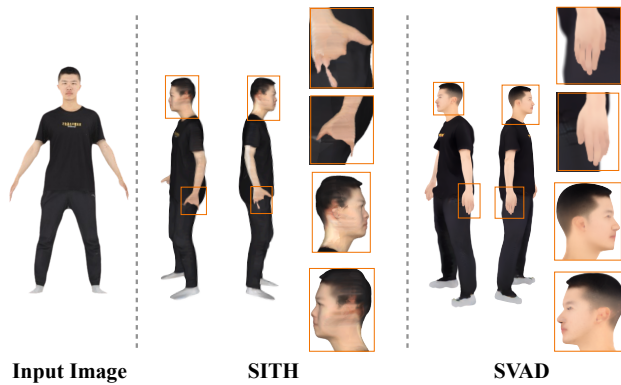


Figure 7. **Qualitative Evaluation against SITH [17]**. Our method reconstructs fine detail(hands), while preserving original identity in facial regions.

351 **4.4. Ablation Study**

352 In this section, we conduct ablation studies to validate each  
 353 component of our methods. The average metrics over 4

sequences in the People Snapshot dataset are reported in 354  
 Table 3. It shows that our methods modules are required 355  
 to reach the optimal performance reflected by all the met- 356  
 rics. Using the THuman dataset, we apply the same eval- 357




Figure 8. **Qualitative Evaluation against ExAvatar in single image input.**

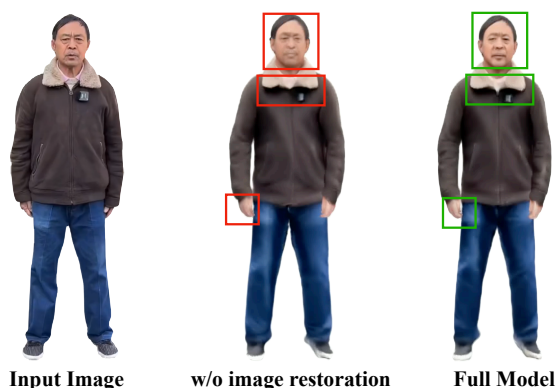


Figure 9. **Ablation study on the image restoration module.** We show that applying the module into our pipeline recover fine details on the final avatar output.




Figure 10. **Ablation study on the identity preservation module.** We show that with the module, the final avatar maintains facial details on the original input image.

358 uation technique as in our quantitative evaluation. Results  
359 show that our method performs the best in PSNR, SSIM and  
360 CLIP similarity and performs second best in LPIPS. Fig-  
361 ure 9 shows visual results of the effect of the image restora-  
362 tion module. High-detailed regions such as clothing texture,

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
w/o Identity Preserve	25.19	0.9419	0.0623	0.9231
w/o Image Restoration	25.61	0.9298	0.0645	0.9239
<b>Ours (Full)</b>	<b>25.79</b>	<b>0.9502</b>	<b>0.0594</b>	<b>0.9241</b>

Table 3. **Ablation study on the People Snapshot dataset.**

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
w/o Identity Preserve	21.12	0.9256	0.0898	0.9284
w/o Image Restoration	21.16	0.9212	0.0799	0.9201
<b>Ours (Full)</b>	<b>21.95</b>	<b>0.9291</b>	<b>0.0824</b>	<b>0.9321</b>

Table 4. **Ablation study on the THuman dataset.**

363 fingers, and facial details are better preserved when apply-  
364 ing our module. Figure 10, shows the visual effect of the  
365 identity preservation module. We clearly show that original  
366 input’s facial details are more preserved with the presence  
367 of our module.

## 368 5. Conclusion and Future Work

369 In this work, we introduced SVAD, a novel synthetic data  
370 generation approach for creating high-fidelity, animatable  
371 3D human avatars from a single image. By leveraging  
372 video diffusion models to generate pose-conditioned syn-  
373 thetic training data, and enhancing this data through iden-  
374 tity preservation and image restoration, SVAD successfully  
375 addresses key challenges in single-image avatar creation.  
376 Our method demonstrates how carefully refined synthetic  
377 data can overcome limitations in traditional approaches, en-  
378 abling stable, visually consistent avatars that retain the ori-  
379 ginal subject’s identity and details across varied poses and  
380 viewpoints. Through comprehensive evaluations, SVAD  
381 achieves state-of-the-art performance compared to exist-  
382 ing methods while maintaining the rendering efficiency of  
383 3DGS.

384 **Limitations and Future Work.** Our method has three  
385 primary limitations. First, the synthetic data generation  
386 process relies on computationally intensive video diffusion  
387 models, creating a bottleneck in the pipeline. Second, while  
388 effective for diverse poses represented in the synthetic train-  
389 ing data, SVAD struggles with extreme or unconventional  
390 poses that fall outside this distribution. Third, the method  
391 does not handle object interactions, limiting its applicability  
392 in dynamic environments.

393 In future work, we plan to explore more efficient syn-  
394 thetic data generation techniques to reduce computational  
395 requirements while maintaining quality. We aim to expand  
396 the diversity of synthetic training data to include extreme  
397 poses and object interactions, and investigate how synthetic  
398 data generation could benefit multi-subject scenarios and  
399 clothing variation to extend SVAD’s capabilities beyond  
400 single-subject avatars.



401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431  
432  
433  
434  
435  
436  
437  
438  
439  
440  
441  
442  
443  
444  
445  
446  
447  
448  
449  
450  
451  
452  
453  
454  
455  
456  
457

## References

- [1] Badour AlBahar, Shunsuke Saito, Hung-Yu Tseng, Changil Kim, Johannes Kopf, and Jia-Bin Huang. Single-image 3d human digitization with shape-guided diffusion. In *SIG-GRAPH Asia 2023 Conference Papers*, pages 1–11, 2023. 2
- [2] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *CVPR*, pages 8387–8397, 2018. 5, 6
- [3] Ankan Kumar Bhunia, Salman Khan, Hisham Cholakkal, Rao Muhammad Anwer, Jorma Laaksonen, Mubarak Shah, and Fahad Shahbaz Khan. Person image synthesis via denoising diffusion model. In *CVPR*, pages 5968–5976, 2023. 2
- [4] Yukang Cao, Yan-Pei Cao, Kai Han, Ying Shan, and Kwan-Yee K Wong. Dreamavatar: Text-and-shape guided 3d human avatar generation via diffusion models. In *CVPR*, pages 958–968, 2024. 2
- [5] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *ICCV*, pages 5933–5942, 2019. 2
- [6] Mingjin Chen, Junhao Chen, Xiaojun Ye, Huan-ang Gao, Xiaoxue Chen, Zhaoxin Fan, and Hao Zhao. Ultraman: Single image 3d human reconstruction with ultra speed and detail. *arXiv preprint arXiv:2403.12028*, 2024. 6
- [7] Xiaoxu Chen, Jingfan Tan, Tao Wang, Kaihao Zhang, Wenhao Luo, and Xiaochun Cao. Towards real-world blind face restoration with generative diffusion prior. *IEEE TCSVT*, 2024. 5
- [8] Yufan Chen, Lizhen Wang, Qijing Li, Hongjiang Xiao, Shengping Zhang, Hongxun Yao, and Yebin Liu. Monogaussianavatar: Monocular gaussian point-based head avatar. In *SIGGRAPH*, pages 1–9, 2024. 2
- [9] Yushuo Chen, Zerong Zheng, Zhe Li, Chao Xu, and Yebin Liu. Meshavatar: Learning high-quality triangular human avatars from multi-view videos. *arXiv preprint arXiv:2407.08414*, 2024. 1
- [10] Xuangeng Chu and Tatsuya Harada. Generalizable and animatable gaussian head avatar. *NeurIPS*, 2024. 4
- [11] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II 19*, pages 424–432. Springer, 2016. 3
- [12] Helisa Dharmo, Yinyu Nie, Arthur Moreau, Jifei Song, Richard Shaw, Yiren Zhou, and Eduardo Pérez-Pellitero. Headgas: Real-time animatable head avatars via 3d gaussian splatting. In *ECCV*, pages 459–476, 2025. 2
- [13] Yuanxing Duan, Fangyin Wei, Qiyu Dai, Yuhang He, Wenzheng Chen, and Baoquan Chen. 4d gaussian splatting: Towards efficient novel view synthesis for dynamic scenes. *arXiv preprint arXiv:2402.03307*, 2024. 2
- [14] Jianglin Fu, Shikai Li, Yuming Jiang, Kwan-Yee Lin, Chen Qian, Chen Change Loy, Wayne Wu, and Ziwei Liu. Stylegan-human: A data-centric odyssey of human generation. In *ECCV*, pages 1–19. Springer, 2022. 2
- [15] Rafael C Gonzalez and Richard E Woods. *Digital Image Processing*. Prentice Hall, 2008. 5
- [16] Hsuan-I Ho, Lixin Xue, Jie Song, and Otmar Hilliges. Learning locally editable virtual humans. In *CVPR*, pages 21024–21035, 2023. 6
- [17] I Ho, Jie Song, Otmar Hilliges, et al. Sith: Single-view textured human reconstruction with image-conditioned diffusion. In *CVPR*, pages 538–549, 2024. 6, 7
- [18] Li Hu. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *CVPR*, pages 8153–8163, 2024. 2, 3
- [19] 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 *CVPR*, pages 634–644, 2024. 1, 2, 6
- [20] Shoukang Hu, Fangzhou Hong, Liang Pan, Haiyi Mei, Lei Yang, and Ziwei Liu. Sherf: Generalizable human nerf from a single image. In *ICCV*, pages 9352–9364, 2023. 2
- [21] Shoukang Hu, Tao Hu, and Ziwei Liu. Gauhuman: Articulated gaussian splatting from monocular human videos. In *CVPR*, pages 20418–20431, 2024. 2
- [22] Letian Huang, Jiayang Bai, Jie Guo, Yuanqi Li, and Yanwen Guo. On the error analysis of 3d gaussian splatting and an optimal projection strategy. *CoRR*, 2024. 2
- [23] Yangyi Huang, Hongwei Yi, Yuliang Xiu, Tingting Liao, Jiaxiang Tang, Deng Cai, and Justus Thies. Tech: Text-guided reconstruction of lifelike clothed humans. In *3DV*, pages 1531–1542, 2024. 6
- [24] Rohit Jena, Ganesh Subramanian Iyer, Siddharth Choudhary, Brandon Smith, Pratik Chaudhari, and James Gee. Splatarmor: Articulated gaussian splatting for animatable humans from monocular rgb videos. *arXiv preprint arXiv:2311.10812*, 2023. 2
- [25] Suyi Jiang, Haoran Jiang, Ziyu Wang, Haimin Luo, Wenzheng Chen, and Lan Xu. Humangen: Generating human radiance fields with explicit priors. In *CVPR*, pages 12543–12554, 2023. 2
- [26] Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Instantavatar: Learning avatars from monocular video in 60 seconds. In *CVPR*, pages 16922–16932, 2023. 5
- [27] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. Gaussianshader: 3d gaussian splatting with shading functions for reflective surfaces. In *CVPR*, pages 5322–5332, 2024. 2
- [28] Brennan Jones, Yaying Zhang, Priscilla NY Wong, and Sean Rintel. Belonging there: Vroom-ing into the uncanny valley of xr telepresence. *ACM HCI*, 5(CSCW1):1–31, 2021. 2
- [29] Johanna Karras, Aleksander Holynski, Ting-Chun Wang, and Ira Kemelmacher-Shlizerman. Dreampose: Fashion image-to-video synthesis via stable diffusion. In *ICCV*, pages 22623–22633, 2023. 2
- [30] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 42(4):139–1, 2023. 2

- 515 [31] Agelos Kratimenos, Jiahui Lei, and Kostas Daniilidis. 516 Dynmf: Neural motion factorization for real-time dynamic 517 view synthesis with 3d gaussian splatting. In *ECCV*, pages 518 252–269, 2025. 2
- 519 [32] Byeonghyeon Lee, Howoong Lee, Xiangyu Sun, Usman Ali, 520 and Eunbyung Park. Deblurring 3d gaussian splatting. *arXiv 521 preprint arXiv:2401.00834*, 2024. 2
- 522 [33] Jiahui Lei, Yufu Wang, Georgios Pavlakos, Lingjie Liu, and 523 Kostas Daniilidis. Gart: Gaussian articulated template mod- 524 els. In *CVPR*, pages 19876–19887, 2024. 2
- 525 [34] Zhe Li, Zerong Zheng, Lizhen Wang, and Yebin Liu. Ani- 526 matable gaussians: Learning pose-dependent gaussian maps 527 for high-fidelity human avatar modeling. In *CVPR*, pages 528 19711–19722, 2024. 2
- 529 [35] Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. 530 Common diffusion noise schedules and sample steps are 531 flawed. In *Proceedings of the IEEE/CVF winter conference 532 on applications of computer vision*, pages 5404–5411, 2024. 533 3
- 534 [36] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and 535 Deva Ramanan. Dynamic 3d gaussians: Tracking 536 by persistent dynamic view synthesis. *arXiv preprint 537 arXiv:2308.09713*, 2023. 2
- 538 [37] Gyeongsik Moon, Takaaki Shiratori, and Shunsuke Saito. 539 Expressive whole-body 3d gaussian avatar. *ECCV*, 2024. 1, 540 2, 5, 6
- 541 [38] Arthur Moreau, Jifei Song, Helisa Dharmo, Richard Shaw, 542 Yiren Zhou, and Eduardo Pérez-Pellitero. Human gaussian 543 splatting: Real-time rendering of animatable avatars. In 544 *CVPR*, pages 788–798, 2024. 1
- 545 [39] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy 546 Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, 547 Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. 548 Dinov2: Learning robust visual features without supervision. 549 *arXiv preprint arXiv:2304.07193*, 2023. 5
- 550 [40] Haokai Pang, Heming Zhu, Adam Kortylewski, Christian 551 Theobalt, and Marc Habermann. Ash: Animatable gaussian 552 splats for efficient and photoreal human rendering. In *CVPR*, 553 pages 1165–1175, 2024. 2
- 554 [41] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson 555 image editing. In *Seminal Graphics Papers: Pushing the 556 Boundaries, Volume 2*, pages 577–582. 2023. 5
- 557 [42] Sergey Prokudin, Michael J Black, and Javier Romero. Sm- 558 plpix: Neural avatars from 3d human models. In *WACV*, 559 pages 1810–1819, 2021. 2
- 560 [43] Shenhan Qian, Tobias Kirschstein, Liam Schoneveld, Davide 561 Davoli, Simon Giebenhain, and Matthias Nießner. Gaus- 562 sianavatars: Photorealistic head avatars with rigged 3d gaus- 563 sians. In *CVPR*, pages 20299–20309, 2024. 2
- 564 [44] Zhiyin Qian, Shaofei Wang, Marko Mihajlovic, Andreas 565 Geiger, and Siyu Tang. 3dgs-avatar: Animatable avatars via 566 deformable 3d gaussian splatting. In *CVPR*, pages 5020– 567 5030, 2024. 1, 2
- 568 [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya 569 Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, 570 Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learn- 571 ing transferable visual models from natural language super- 572 vision. In *ICML*, 2021. 5
- [46] Yurui Ren, Ge Li, Shan Liu, and Thomas H Li. Deep spatial 573 transformation for pose-guided person image generation and 574 animation. *IEEE TIP*, 29:8622–8635, 2020. 2 575
- [47] Robin Rombach, Andreas Blattmann, Dominik Lorenz, 576 Patrick Esser, and Björn Ommer. High-resolution image syn- 577 thesis with latent diffusion models. In *CVPR*, pages 10684– 578 10695, 2022. 5 579
- [48] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Mor- 580 ishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned 581 implicit function for high-resolution clothed human digitiza- 582 tion. In *ICCV*, pages 2304–2314, 2019. 6 583
- [49] Tim Salimans and Jonathan Ho. Progressive distillation 584 for fast sampling of diffusion models. *arXiv preprint 585 arXiv:2202.00512*, 2022. 3 586
- [50] Kripasindhu Sarkar, Dushyant Mehta, Weipeng Xu, 587 Vladislav Golyanik, and Christian Theobalt. Neural re- 588 rendering of humans from a single image. In *ECCV*, pages 589 596–613, 2020. 2 590
- [51] Zhijing Shao, Zhaolong Wang, Zhuang Li, Duotun Wang, 591 Xiangru Lin, Yu Zhang, Mingming Fan, and Zeyu Wang. 592 Splattingavatar: Realistic real-time human avatars with 593 mesh-embedded gaussian splatting. In *CVPR*, pages 1606– 594 1616, 2024. 1 595
- [52] Kaiyue Shen, Chen Guo, Manuel Kaufmann, Juan Jose 596 Zarate, Julien Valentin, Jie Song, and Otmar Hilliges. X- 597 avatar: Expressive human avatars. In *CVPR*, pages 16911– 598 16921, 2023. 1 599
- [53] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, 600 Elisa Ricci, and Nicu Sebe. First order motion model for 601 image animation. *NeurIPS*, 32, 2019. 2 602
- [54] Aliaksandr Siarohin, Oliver J Woodford, Jian Ren, Menglei 603 Chai, and Sergey Tulyakov. Motion representations for arti- 604 culated animation. In *CVPR*, pages 13653–13662, 2021. 605 2 606
- [55] Jiaming Song, Chenlin Meng, and Stefano Ermon. 607 Denoising diffusion implicit models. *arXiv preprint 608 arXiv:2010.02502*, 2020. 3 609
- [56] David Svitov, Pietro Morerio, Lourdes Agapito, and Alessio 610 Del Bue. Haha: Highly articulated gaussian human avatars 611 with textured mesh prior. *arXiv preprint arXiv:2404.01053*, 612 2024. 1 613
- [57] Zhengyan Tong, Chao Li, Zhaokang Chen, Bin Wu, and 614 Wenjiang Zhou. Musepose: a pose-driven image-to-video 615 framework for virtual human generation. *arxiv*, 2024. 3 616
- [58] Jie Wang, Xianyan Li, Jiucheng Xie, Feng Xu, and Hao Gao. 617 Gaussianhead: Impressive 3d gaussian-based head avatars 618 with dynamic hybrid neural field. *arXiv e-prints*, pages 619 arXiv–2312, 2023. 2 620
- [59] Tan Wang, Linjie Li, Kevin Lin, Yuanhao Zhai, Chung- 621 Ching Lin, Zhengyuan Yang, Hanwang Zhang, Zicheng Liu, 622 and Lijuan Wang. Disco: Disentangled control for realistic 623 human dance generation. In *CVPR*, pages 9326–9336, 2024. 624 2 625
- [60] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P 626 Simoncelli. Image quality assessment: from error visibility 627 to structural similarity. *IEEE TIP*, 13(4):600–612, 2004. 5 628

- 629 [61] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan,  
630 Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Hu-  
631 mannerf: Free-viewpoint rendering of moving people from  
632 monocular video. In *CVPR*, pages 16210–16220, 2022. 6
- 633 [62] Zhongcong Xu, Jianfeng Zhang, Jun Hao Liew, Hanshu Yan,  
634 Jia-Wei Liu, Chenxu Zhang, Jiashi Feng, and Mike Zheng  
635 Shou. Magicanimate: Temporally consistent human image  
636 animation using diffusion model. In *CVPR*, pages 1481–  
637 1490, 2024. 2
- 638 [63] Zhendong Yang, Ailing Zeng, Chun Yuan, and Yu Li. Effec-  
639 tive whole-body pose estimation with two-stages distillation.  
640 In *Proceedings of the IEEE/CVF International Conference*  
641 *on Computer Vision*, pages 4210–4220, 2023. 3
- 642 [64] Jae Shin Yoon, Lingjie Liu, Vladislav Golyanik, Kripasindhu  
643 Sarkar, Hyun Soo Park, and Christian Theobalt. Pose-guided  
644 human animation from a single image in the wild. In *CVPR*,  
645 pages 15039–15048, 2021. 2
- 646 [65] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qiong-  
647 hai Dai, and Yebin Liu. Function4d: Real-time human vol-  
648 umetric capture from very sparse consumer rgbd sensors. In  
649 *CVPR*, 2021. 5
- 650 [66] Wing-Yin Yu, Lai-Man Po, Ray CC Cheung, Yuzhi Zhao, Yu  
651 Xue, and Kun Li. Bidirectionally deformable motion modu-  
652 lation for video-based human pose transfer. In *ICCV*, pages  
653 7502–7512, 2023. 2
- 654 [67] Yifei Zeng, Yuanxun Lu, Xinya Ji, Yao Yao, Hao Zhu, and  
655 Xun Cao. Avatarbooth: High-quality and customizable 3d  
656 human avatar generation. *arXiv preprint arXiv:2306.09864*,  
657 2023. 1
- 658 [68] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding  
659 conditional control to text-to-image diffusion models. In  
660 *ICCV*, pages 3836–3847, 2023. 2
- 661 [69] Pengze Zhang, Lingxiao Yang, Jian-Huang Lai, and Xiaohua  
662 Xie. Exploring dual-task correlation for pose guided person  
663 image generation. In *CVPR*, pages 7713–7722, 2022. 2
- 664 [70] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman,  
665 and Oliver Wang. The unreasonable effectiveness of deep  
666 features as a perceptual metric. In *CVPR*, pages 586–595,  
667 2018. 5
- 668 [71] Zechuan Zhang, Zongxin Yang, and Yi Yang. Sifu: Side-  
669 view conditioned implicit function for real-world usable  
670 clothed human reconstruction. In *CVPR*, pages 9936–9947,  
671 2024. 6
- 672 [72] Jian Zhao and Hui Zhang. Thin-plate spline motion model  
673 for image animation. In *CVPR*, pages 3657–3666, 2022. 2
- 674 [73] Shunyuang Zheng, Boyao Zhou, Ruizhi Shao, Boning Liu,  
675 Shengping Zhang, Liqiang Nie, and Yebin Liu. Gps-  
676 gaussian: Generalizable pixel-wise 3d gaussian splatting for  
677 real-time human novel view synthesis. In *CVPR*, pages  
678 19680–19690, 2024. 2
- 679 [74] Yang Zheng, Qingqing Zhao, Guandao Yang, Wang Yi-  
680 fan, Donglai Xiang, Florian Dubost, Dmitry Lagun, Thabo  
681 Beeler, Federico Tombari, Leonidas Guibas, et al. Physa-  
682 vatar: Learning the physics of dressed 3d avatars from visual  
683 observations. *arXiv preprint arXiv:2404.04421*, 2024. 1
- 684 [75] Shenhao Zhu, Junming Leo Chen, Zuozhuo Dai, Yinghui  
685 Xu, Xun Cao, Yao Yao, Hao Zhu, and Siyu Zhu. Champ:  
Controllable and consistent human image animation with 3d  
parametric guidance. *ECCV*, 2024. 2
- [76] Wojciech Zielonka, Timur Bagautdinov, Shunsuke Saito,  
Michael Zollhöfer, Justus Thies, and Javier Romero. Driv-  
able 3d gaussian avatars. *arXiv preprint arXiv:2311.08581*,  
2023. 2