

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

Yonwoo Choi
 SECERN AI

yonwoo.choi@secern.ai

<https://yc4ny.github.io/SVAD/>




Figure 1. **SVAD.** Our method creates high-fidelity 3D avatars from a single image through synthetic data generation. We leverage video diffusion to generate pose-conditioned animations, enhance them with identity preservation and image restoration modules, then train a 3D Gaussian Splatting avatar. The resulting avatars maintain consistent identity across novel poses and viewpoints while enabling real-time rendering, outperforming state-of-the-art approaches.

Abstract

Creating high-quality animatable 3D human avatars from a single image remains a significant challenge in computer vision due to the inherent difficulty of reconstructing complete 3D information from a single viewpoint. Current approaches face a clear limitation: 3D Gaussian Splatting (3DGS) methods produce high-quality results but require multiple views or video sequences, while video diffusion models can generate animations from single images but struggle with consistency and identity preservation. We present SVAD, a novel approach that addresses these limitations by leveraging complementary strengths of existing techniques. Our method generates synthetic training data through video diffusion, enhances it with identity preservation and image restoration modules, and utilizes this refined data to train 3DGS avatars. Comprehensive evaluations demonstrate that SVAD outperforms state-of-the-art (SOTA) single-image methods in maintaining identity consistency and fine details across novel poses and viewpoints, while enabling real-time rendering capabilities.

Through our data augmentation pipeline, we overcome the dependency on dense monocular or multi-view training data typically required by traditional 3DGS approaches. Extensive quantitative, qualitative comparisons show our method achieves superior performance across multiple metrics against baseline models. By effectively combining the generative power of diffusion models with both the high-quality results and rendering efficiency of 3DGS, our work establishes a new approach for high-fidelity avatar generation from a single image input.

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 [12, 27, 49, 50, 58, 66, 67, 73, 85, 92]. These methods rely on dense input data, typically monocular or multi-view

video [12, 27, 49, 58, 66, 73, 92], 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 adaptability to novel poses becomes a key challenge.

Recent advancements in video diffusion models offer a potential solution by enabling animation generation from a single static image [26, 39, 77, 80, 93]. These models use certain conditions in diffusion processes to create video sequences, demonstrating the powerful generative capabilities of diffusion for single-image-driven animation. However, diffusion models often struggle to maintain temporal coherence, leading to inconsistent features and identity drift across frames [4, 16, 25, 70]. Additionally, their iterative denoising process for each frame introduces significant computational overhead, limiting their feasibility for real-time or interactive applications where rapid rendering across novel views is essential.

To overcome these challenges, we propose SVAD, a novel synthetic data generation and avatar creation pipeline that synergizes the generative flexibility of diffusion models with the efficient rendering capabilities of 3DGS avatars. Our approach leverages video diffusion model [74] to generate diverse pose-conditioned synthetic training data from a single-image. This synthetic data is refined through an identity-preservation module and an image restoration module to ensure that perceptual identity consistency and structural fidelity are preserved across diverse poses and temporal sequences. The resulting high-quality synthetic dataset is then used to train a 3DGS avatar model [49], which benefits from the rapid rendering capabilities inherent to 3DGS. By combining the generative strengths of diffusion for synthetic data creation with the efficiency of 3DGS for rendering, SVAD achieves consistent, high-quality 3D avatar animations from single-image input.

In summary, our main contributions are:

- We introduce a novel pipeline that generates high-quality synthetic training data from a single-image to create detailed, animatable 3D human avatars.
- We develop a comprehensive data augmentation approach that combines identity preservation and image restoration to ensure consistent identity and fine details across diverse poses.
- We demonstrate through extensive experiments that our synthetic data-driven approach significantly outperforms SOTA single-image avatar generation methods in identity preservation and novel pose adaptation while maintaining efficient real-time rendering.

2. Related Work

Diffusion Model for Human Image Animation The use of diffusion models has led to significant advancements in human image animation, enabling the generation of real-

istic and temporally consistent animations from static images [1, 5, 8, 18, 28, 33, 56, 60, 65, 68, 69, 82, 84, 87, 90]. Early methods, such as PIDM [3] and DreamPose [37], focused on improving texture fidelity by employing texture diffusion modules to align texture patterns between reference and target images. These methods, while enhancing detail preservation, still face challenges in maintaining temporal stability across frames.

Recent works, including DisCo [77] and Animate Anyone [26], have extended diffusion models to improve temporal consistency and fine-grained control in human animation tasks. DisCo leverages dual ControlNets [86] to separately control pose and background elements, providing more robust conditioning for complex motion sequences. Similarly, Animate Anyone integrates a ReferenceNet with temporal attention layers to ensure appearance consistency and smooth transitions across frames, thereby addressing flickering issues commonly observed in earlier models.

Dynamic 3D Gaussian based Avatars The concept of Gaussian splatting for 3D avatars has emerged recently as an innovative approach to explicit scene representation [38]. This technique models a scene as a collection of 3D Gaussian elements, each containing photometric and geometric properties. During rendering, these Gaussian splats are projected onto the image plane, creating the final rendered output. The efficiency of 3DGS has been demonstrated in both static [30, 35, 42] and dynamic [17, 36, 41, 43, 48] scenes, making it a versatile tool for various applications. Recent advancements [6, 7, 11, 15, 29, 32, 44, 57, 58, 76, 94] have explored the use of 3DGS to create photorealistic human avatars across different scenarios. These methods commonly rely on multi-view data [46, 53, 91] or monocular video [27, 32, 44, 49, 58] as input to achieve high-quality, consistent results. The advantage of 3DGS lies in its ability to produce temporally stable animated avatars with superior quantitative metrics.

3. Method

To generate high-quality human avatars from a single-image, facilitating free-viewpoint rendering and realistic animation, we integrate the generative capabilities of video diffusion models with the rendering efficiency of 3D Gaussian-based avatars. We start by leveraging a pretrained video diffusion model [74] for character animation to produce initial synthetic data, as described in Sec. 3.1. Directly using these frames to train a 3DGS avatar model [49], however, often yields poor results, with challenges in preserving facial identity, clothing details, and maintaining consistent multi-view coherence across side and back views. To address these issues and enhance avatar quality, we introduce a data augmentation pipeline in Sec. 3.2 comprising identity-preservation and image-restoration modules to refine the diffusion outputs. With the augmented synthetic




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.

data, we proceed to train a 3DGS avatar model, as outlined in Sec. 3.3. The following sections detail the technical methodologies employed in our approach.

3.1. Video Diffusion Module

To generate an animated character video V from a single input image I , we leverage MusePose [74], a finetuned variant of Animate Anyone [26], which is a SOTA video diffusion model designed for realistic human animation while maintaining temporal consistency and appearance fidelity. MusePose employs a U-Net [62]-based diffusion architecture with integrated pose and temporal controls, allowing for pose-guided animation across frames. For our pipeline, we utilize a pose sequence video from a sequence from the People Snapshot [2] dataset, which depicts a subject performing a full-body rotation with arms extended horizontally. This sequence results in 189 frames that serve as pose inputs to the MusePose video diffusion model.

The model architecture incorporates several key components for effective character animation. The denoising UNet is implemented as a 3D UNet [14] with motion modules for temporal coherence. Specifically, we use Vanilla motion modules [20, 21] with temporal self-attention blocks at resolutions of [1, 2, 4, 8] and in the mid-block. Each transformer [75] block contains 8 attention heads, with temporal position encoding enabling positional awareness across

a sequence of up to 128 frames. To incorporate pose guidance, a lightweight Pose Guider encodes the motion control signal from the predefined 2D keypoints into a pose-aligned latent representation $P(p_t) \in \mathbb{R}^{H \times W \times C}$. For a pose feature $p_t \in \mathbb{R}^{J \times 2}$ at time t , where J is the number of keypoints, we align the encoding to ensure continuity between frames by adding this encoded pose signal to the noise latent z_t :

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

For the diffusion process, we adopt a v-prediction [64] formulation with zero-SNR sampling [47], using a scaled linear beta schedule with $\beta_{\text{start}} = 0.00085$ and $\beta_{\text{end}} = 0.012$. The DDIM [71] sampler is configured for efficient inference with 20 sampling steps and a classifier-free guidance [24] scale of 3.5.

A critical challenge in character animation is ensuring anatomical consistency between the reference image and the motion poses. Direct application of pose control can result in unnatural animations due to mismatches in body proportions [8]. Therefore, we employ a comprehensive pose alignment procedure that adapts the source pose to match the reference character’s physical characteristics.

Given a reference pose P_{ref} and a source pose P_{src} detected using DWpose [81], we compute scale parameters $\mathbf{S} = \{s_1, s_2, \dots, s_{10}\}$ for ten distinct body regions: neck, face, shoulders, upper arms, lower arms, hands, torso, up-



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.

per legs, and lower legs. For each body part i , we compute its scale factor s_i as the ratio between the corresponding keypoint distances. For body parts with bilateral symmetry (e.g., arms), we average the scales from both sides:

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

To apply these scales to the source pose, we use a rotation matrix transformation centered at anchor points specific to each body part:

$$p' = c_i + s_i \cdot (p - c_i) \quad (3)$$

where c_i is the anchor center for part i . This hierarchical approach ensures body proportions match the reference while maintaining the overall pose structure.

3.2. Data Augmentation Module

Training the 3DGS model using only outputs from the video diffusion model often results in low-fidelity avatars, par-

ticularly in terms of facial details and high-frequency features like hands and clothing. To address these challenges, we introduce a data augmentation module that enhances the quality of the training data. This module includes an identity preservation sub-module ensuring coherence in facial details across frames and a image restoration submodule which refines texture quality and high-frequency details, resulting in more realistic textures. This comprehensive data augmentation significantly improves the synthetic training data, enabling the 3DGS avatar model integrated in the future to generate more realistic and detailed 3D avatars.

Identity preservation sub-module. To ensure consistent and realistic facial details across frames, we implement an identity preservation module that combines 3D head reconstruction and facial fusion techniques. From a single input image, we first create a 3D Gaussian-based head avatar using a method inspired by Chu *et al.* [13], which employs a novel *dual-lifting* approach that predicts both forward and backward lifting distances.



Figure 4. **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.

Given an input image I_s , global and local features F_{local} are extracted using a frozen DINOv2 [52] backbone. These features are used to predict forward and backward lifting distances, positioning 3D Gaussians G_{pos} as follows:

$$G_{\text{pos}} = [\mathbf{p}_s + E_{\text{Conv0}}(F_{\text{local}}) \cdot \mathbf{n}_s, \mathbf{p}_s - E_{\text{Conv1}}(F_{\text{local}}) \cdot \mathbf{n}_s], \quad (4)$$

where $\mathbf{p}_s$ is the initial point plane, $\mathbf{n}_s$ is the normal vector, and E_{Conv} are convolutional layers predicting offsets. To capture expression variations, we bind 3DMM [45] features:

$$G_{\text{expr}} = \text{MLP}(F_{\text{3DMM}} + F_{\text{global}}). \quad (5)$$

To animate this 3D head avatar, we separately track FLAME [45] parameters $\Theta = \{\beta, \psi, \theta, \phi\}$ from our predefined pose sequence video (the same sequence used in the video diffusion module), where $\beta \in \mathbb{R}^{300}$ represents shape parameters, $\psi \in \mathbb{R}^{100}$ expression parameters, $\theta \in \mathbb{R}^6$ global pose parameters, and $\phi \in \mathbb{R}^6$ eye pose parameters. These tracked parameters serve as animation controls for

the reconstructed 3D head. Using these tracked FLAME parameters, we render the 3D head avatar to generate a sequence of head images that match our predefined pose sequence. These renderings provide high-quality, identity-consistent facial details across different viewpoints. Since the quality of the renderings deteriorates for back-of-head views, we selectively apply the face fusion process only to frames where the head is front-facing (front and side views).

For the face fusion process, we detect facial landmarks [40] on both the diffusion-generated frame I_{orig} and the rendered head image I_{head} , compute an affine transformation for alignment, and use Poisson image editing [55] for seamless blending:

$$\min_I \int_{\Omega} \|\nabla I - \nabla I_{\text{warp}}\|^2 dx dy, \quad \text{subject to } I|_{\partial\Omega} = I_{\text{orig}}|_{\partial\Omega}, \quad (6)$$

where Ω is defined by the facial mask. This ensures temporally consistent facial details while preserving the original identity throughout the animation sequence.

Image restoration sub-module. Finally, to preserve quality of fine detailed regions, we employ an image restoration module based on the work of Chen *et al.* [10], specifically their diffusion-based image restoration method BFRffusion. This approach leverages the generative prior encapsulated in the pretrained Stable Diffusion [61] model to enhance image details through a comprehensive architecture that effectively extracts features from low-quality images and restores realistic facial details.

For our implementation, we set the super-resolution scale factor to $s = 1.5$, which our empirical analysis showed provides an optimal balance between detail enhancement and artifact suppression. We observed that scale factors $s < 1.5$ produce insufficient detail recovery, while factors $s > 2.0$ introduce perceptual artifacts (particularly in specular regions such as eyes) and significantly increase computational demands during avatar training. The diffusion process uses 50 DDIM sampling steps with:

$$z_{t-1} = \sqrt{\alpha_{t-1}} \left(\frac{z_t - \sqrt{1 - \alpha_t} \epsilon_\theta(z_t)}{\sqrt{\alpha_t}} \right) + \sqrt{1 - \alpha_{t-1}} \epsilon_\theta(z_t) \quad (7)$$

where $\alpha_t = \prod_{i=1}^t (1 - \beta_i)$ and ϵ_θ is the denoising network. We utilize a classifier-free guidance scale of $w = 3.5$, with the guidance equation:

$$\hat{\epsilon}_\theta(z_t) = (1 + w) \epsilon_\theta(z_t) - w \epsilon_\theta(z_t, \emptyset) \quad (8)$$

where $\epsilon_\theta(z_t, \emptyset)$ represents the unconditional prediction. This achieves an optimal balance between restoration quality and processing speed. For face regions, the method employs a face restoration helper with facial landmark detection to specifically enhance facial details, ensuring identity consistency across generated frames. Restored faces are blended with Poisson image editing.

This image restoration submodule significantly improves the fidelity and realism of our synthetic training data by restoring fine facial details, enhancing texture quality in clothing and accessories, and improving overall image coherence. The refined data enables the 3DGS avatar to learn more accurate representations with consistent high-frequency details that persist across poses and viewpoints.

3.3. 3D Human Gaussian Splatting Module

We apply the architecture of a 3DGS based avatar method introduced by Moon *et al.* [49], which integrates the SMPL-X [54] model with a 3D Gaussian-based representation to produce animatable human avatars. Each 3D Gaussian acts as a vertex connected by a pre-defined mesh topology following SMPL-X. This hybrid representation combines the expressive surface modeling of SMPL-X with the flexibility of a volumetric approach, allowing for smooth interpolation across the body surface essential for realistic animations.

Each Gaussian point is associated with positional data $\mathbf{V} \in \mathbb{R}^{N \times 3}$, RGB color values $\mathbf{C} \in \mathbb{R}^{N \times 3}$, and a scale

parameter $\mathbf{S} \in \mathbb{R}^N$, where N is the number of Gaussians. The Gaussian splatting rendering equation is:

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

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

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

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

To maintain spatial coherence, a Laplacian regularizer [51, 72] minimizes the difference between the Laplacian of the canonical mesh and the deformed Gaussian points:

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

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

4. Experiments

4.1. Datasets and Metrics

People-Snapshot Dataset [2] We conduct our avatar evaluation on the People-Snapshot dataset, which features video recordings of subjects performing 360-degree rotations. Following both Anim-NeRF [79] and InstantAvatar [34], we address a known limitation in this dataset: the provided pose parameters often exhibit misalignment with the actual image content. Anim-NeRF addressed this by optimizing pose parameters for both training and test sequences. To ensure fair comparison with existing methods, we adopt these same optimized pose parameters and keep them frozen throughout our training process for fair comparison.

THuman Dataset [83] For evaluating single-image 3D human reconstruction, we employ the THuman dataset, adhering to the methodology established in Ultraman [9]. Our procedure involves randomly selecting 100 scans and generating renderings from four viewpoints (front, left, right, back). We then measure the similarity between our reconstructed outputs and the ground-truth scan renderings from these identical perspectives, facilitating objective comparison with other SOTA methods.

Evaluation Metrics Our evaluation framework uses four metrics to quantify reconstruction quality: PSNR [19], SSIM [78], LPIPS [88], and CLIP Similarity [59] (referred to as CLIP in our tables). This combination provides comprehensive assessment across different dimensions: PSNR for pixel accuracy, SSIM for structural coherence, LPIPS for perceptual alignment with human vision, and CLIP for semantic consistency at the feature level. The use of these metrics enables thorough evaluation of both fine-grained detail, and overall perceptual quality.

Method	Female-4-casual			Male-3-casual			Female-3-casual			Male-4-casual		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
HumanNeRF [79]	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 [27]	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 [49]	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 [49] (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 that use monocular input utilize approximately 200 input frames.

4.2. Quantitative Evaluation

We quantitatively evaluate the quality of single-image 3D avatars generated by our method against SOTA 3D avatar generation methods [27, 49, 79]. While current 3D avatar models generally require a monocular video as input, we assess our model’s performance using a single-image as input on ExAvatar [49]. Additionally, we report results using the original full training set of approximately 200 input frames for monocular input based avatar models for reference. As shown in Table 1, our model achieves highest scores on most of the metrics among single-image input methods. We further compare our approach with single-view 3D human reconstruction methods [9, 23, 31, 63, 89], many of which employ the SMPL model, allowing for animatability through mesh fitting and reposing techniques, such as those in Editable Humans [22]. We randomly sample 100 scans from the THuman dataset and report results. We repose our trained avatar using ground-truth SMPL-X parameters and compare with the ground-truth scan renderings from the same views. As presented in Table 2, our method surpasses all baselines, demonstrating superior quality in 3D human reconstruction tasks.

4.3. Qualitative Evaluation

Figure 4 shows the overall quality of our generated 3D avatars from single-images in the People Snapshot and the THuman dataset. Figure 5, Figure 6 shows that our method performs superior compared to current SiTH [23]. For single-image avatar generation, we evaluate on the People Snapshot dataset and compare against ExAvatar [49]. For fairness, we train ExAvatar for the same number (12,000) of iterations. Figure 7 shows that for single-image avatar generation, our method performs superior especially for the back and side views.

4.4. Ablation Study

In this section, we conduct ablation studies to validate each component of our methods. The average metrics over 4 sequences in the People Snapshot dataset are reported in Table 3. It shows that our methods modules are required




Figure 5. **Qualitative Evaluation** against SiTH [23]. Our approach better reconstructs complex contours and subtle features, resulting in a more lifelike and coherent side-view appearance.




Figure 6. **Qualitative Evaluation** against SiTH [23]. Our method reconstructs fine detail (hands), while preserving original identity in facial regions.

to reach the optimal performance reflected by all the metrics. Using the THuman dataset, we apply the same evaluation technique as in our quantitative evaluation. Results show that our method performs the best in PSNR, SSIM and CLIP similarity and performs second best in LPIPS. Figure 8 shows visual results of the effect of the image restoration module. High-detailed regions such as clothing texture,




Figure 7. **Qualitative Evaluation** against ExAvatar [49] in single-image to 3D avatar task. Our method generates more plausible back and side views with the generated synthetic dataset.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
PIFu[63]	15.62	0.8921	0.1903	0.8612
TeCH[31]	15.85	0.8892	0.1667	0.8890
Ultraman[9]	18.13	0.9019	0.1334	0.9089
SIFU[89]	18.59	0.8591	0.1402	0.8873
SiTH[23]	19.98	0.9018	0.1294	0.9084
Ours	20.92	0.9291	0.1124	0.9321

Table 2. **Quantitative Evaluation** on single-image to 3D human reconstruction tasks on 100 scan renderings of the THuman [83] Dataset. Top two results are colored as **first** **second**.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
w/o Identity Preserve	22.19	0.9419	0.0623	0.9231
w/o Image Restoration	22.61	0.9298	0.0645	0.9239
Ours (Full)	22.79	0.9502	0.0594	0.9241

Table 3. **Ablation study** on the People Snapshot dataset. Our full model consistently outperforms variants with individual components removed across all metrics.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
w/o Identity Preserve	20.12	0.9256	0.1294	0.9284
w/o Image Restoration	20.16	0.9212	0.0799	0.9201
Ours (Full)	20.92	0.9291	0.1124	0.9321

Table 4. **Ablation study** on the THuman dataset. The full model achieves superior performance in most metrics, demonstrating the importance of each component in our pipeline.

fingers, and facial details are better preserved when applying our module. Figure 9, shows the visual effect of the identity preservation module. We clearly show that original input’s facial details are more preserved our module.




Figure 8. **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 9. **Ablation study** on the identity preservation module. We show that with the module, the final avatar maintains facial details on the original input image.

5. Conclusion and Discussion

In this work, we introduced SVAD, a novel synthetic data generation approach for creating high-fidelity, animatable 3D human avatars from a single image. By combining the generative power of diffusion models with the rendering efficiency of 3D Gaussian Splatting, SVAD produces avatars that maintain consistent identity across varied poses and viewpoints. Through comprehensive experiments, we demonstrate that our method achieves SOTA performance.

Limitations and Future Work. Our method faces several limitations. First, inaccurate background segmentation of training frames produces floating artifacts. Second, our approach struggles with complex clothing textures and loose outfits due to limitations of the video diffusion model in generating detailed synthetic data. Finally, the computational requirements present practical challenges—the video diffusion step demands substantial resources, and the complete pipeline requires 5-6 hours per avatar generation. Future work will focus on improving handling of diverse clothing types and optimizing computational performance.

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 *SIGGRAPH Asia*, 2023. 2
- [2] Thimo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *CVPR*, 2018. 3, 6, 7
- [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*, 2023. 2
- [4] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *CVPR*, 2023. 2
- [5] 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*, 2024. 2
- [6] Hyunsoo Cha, Byungjun Kim, and Hanbyul Joo. Pegasus: Personalized generative 3d avatars with composable attributes. In *CVPR*, 2024. 2
- [7] Hyunsoo Cha, Inhee Lee, and Hanbyul Joo. Perse: Personalized 3d generative avatars from a single portrait. In *CVPR*, 2025. 2
- [8] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *ICCV*, 2019. 2, 3
- [9] 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, 8
- [10] Xiaoxu Chen, Jingfan Tan, Tao Wang, Kaihao Zhang, Wenhan Luo, and Xiaochun Cao. Towards real-world blind face restoration with generative diffusion prior. *IEEE TCSVT*, 2024. 6
- [11] Yufan Chen, Lizhen Wang, Qijing Li, Hongjiang Xiao, Shengping Zhang, Hongxun Yao, and Yebin Liu. Monogaussianavatar: Monocular gaussian point-based head avatar. In *SIGGRAPH*, 2024. 2
- [12] Yushuo Chen, Zerong Zheng, Zhe Li, Chao Xu, and Yebin Liu. Meshavatar: Learning high-quality triangular human avatars from multi-view videos. *ECCV*, 2024. 1, 2
- [13] Xuangeng Chu and Tatsuya Harada. Generalizable and animatable gaussian head avatar. *NeurIPS*, 2024. 4
- [14] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *MICCAI*, 2016. 3
- [15] 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*, 2025. 2
- [16] Zheng Ding, Xuaner Zhang, Zhihao Xia, Lars Jebe, Zhuowen Tu, and Xiuming Zhang. Diffusionrig: Learning personalized priors for facial appearance editing. In *CVPR*, 2023. 2
- [17] Yuanxing Duan, Fangyin Wei, Qiyu Dai, Yuhang He, Wenzheng Chen, and Baoquan Chen. 4d gaussian splatting: Towards efficient novel view synthesis for dynamic scenes. *SIGGRAPH*, 2024. 2
- [18] 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*, 2022. 2
- [19] Rafael C Gonzalez and Richard E Woods. *Digital Image Processing*. Prentice Hall, 2008. 6
- [20] Yuwei Guo, Ceyuan Yang, Anyi Rao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Sparsectrl: Adding sparse controls to text-to-video diffusion models. *ECCV*, 2024. 3
- [21] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *ICLR*, 2024. 3
- [22] Hsuan-I Ho, Lixin Xue, Jie Song, and Otmar Hilliges. Learning locally editable virtual humans. In *CVPR*, 2023. 7
- [23] I Ho, Jie Song, Otmar Hilliges, et al. Sith: Single-view textured human reconstruction with image-conditioned diffusion. In *CVPR*, 2024. 7, 8
- [24] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 3
- [25] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *NeurIPS*, 2022. 2
- [26] Li Hu. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *CVPR*, 2024. 2, 3
- [27] 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*, 2024. 1, 2, 7
- [28] Shoukang Hu, Fangzhou Hong, Liang Pan, Haiyi Mei, Lei Yang, and Ziwei Liu. Sherf: Generalizable human nerf from a single image. In *ICCV*, 2023. 2
- [29] Shoukang Hu, Tao Hu, and Ziwei Liu. Gauhuman: Articulated gaussian splatting from monocular human videos. In *CVPR*, 2024. 2
- [30] 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
- [31] 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*, 2024. 7, 8
- [32] 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
- [33] Suyi Jiang, Haoran Jiang, Ziyu Wang, Haimin Luo, Wenzheng Chen, and Lan Xu. Humangen: Generating human radiance fields with explicit priors. In *CVPR*, 2023. 2

- [34] Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Instantavatar: Learning avatars from monocular video in 60 seconds. In *CVPR*, 2023. 6
- [35] 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*, 2024. 2
- [36] Brennan Jones, Yaying Zhang, Priscilla NY Wong, and Sean Rintel. Belonging there: Vroom-ing into the uncanny valley of xr telepresence. *ACM HCI*, 2021. 2
- [37] Johanna Karras, Aleksander Holynski, Ting-Chun Wang, and Ira Kemelmacher-Shlizerman. Dreampose: Fashion image-to-video synthesis via stable diffusion. In *ICCV*, 2023. 2
- [38] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 2023. 2
- [39] Taeksoo Kim and Hanbyul Joo. Target-aware video diffusion models. *arXiv preprint arXiv:2503.18950*, 2025. 2
- [40] Davis King. Dlib-ml: A machine learning toolkit. <http://dlib.net>, 2009. 5
- [41] Agelos Kratimenos, Jiahui Lei, and Kostas Daniilidis. Dynmf: Neural motion factorization for real-time dynamic view synthesis with 3d gaussian splatting. In *ECCV*, 2025. 2
- [42] Byeonghyeon Lee, Howoong Lee, Xiangyu Sun, Usman Ali, and Eunbyung Park. Deblurring 3d gaussian splatting. *ECCV*, 2024. 2
- [43] Inhee Lee, Byungjun Kim, and Hanbyul Joo. Guess the unseen: Dynamic 3d scene reconstruction from partial 2d glimpses. In *CVPR*, 2024. 2
- [44] Jiahui Lei, Yufu Wang, Georgios Pavlakos, Lingjie Liu, and Kostas Daniilidis. Gart: Gaussian articulated template models. In *CVPR*, 2024. 2
- [45] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *TOG*, 2017. 5
- [46] Zhe Li, Zerong Zheng, Lizhen Wang, and Yebin Liu. Animatable gaussians: Learning pose-dependent gaussian maps for high-fidelity human avatar modeling. In *CVPR*, 2024. 2
- [47] Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. Common diffusion noise schedules and sample steps are flawed. In *WACV*, 2024. 3
- [48] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *3DV*, 2024. 2
- [49] Gyeongsik Moon, Takaaki Shiratori, and Shunsuke Saito. Expressive whole-body 3d gaussian avatar. *ECCV*, 2024. 1, 2, 6, 7, 8
- [50] Arthur Moreau, Jifei Song, Helisa Dharmo, Richard Shaw, Yiren Zhou, and Eduardo Pérez-Pellitero. Human gaussian splatting: Real-time rendering of animatable avatars. In *CVPR*, 2024. 1
- [51] Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. Laplacian mesh optimization. In *GRAPHITE*, 2006. 6
- [52] 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. *TMLR*, 2023. 5
- [53] Haokai Pang, Heming Zhu, Adam Kortylewski, Christian Theobalt, and Marc Habermann. Ash: Animatable gaussian splats for efficient and photoreal human rendering. In *CVPR*, 2024. 2
- [54] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 6
- [55] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. 2023. 5
- [56] Sergey Prokudin, Michael J Black, and Javier Romero. Smpix: Neural avatars from 3d human models. In *WACV*, 2021. 2
- [57] Shenhan Qian, Tobias Kirschstein, Liam Schoneveld, Davide Davoli, Simon Giebenhain, and Matthias Nießner. Gaussianavatars: Photorealistic head avatars with rigged 3d gaussians. In *CVPR*, 2024. 2
- [58] Zhiyin Qian, Shaofei Wang, Marko Mihajlovic, Andreas Geiger, and Siyu Tang. 3dgs-avatar: Animatable avatars via deformable 3d gaussian splatting. In *CVPR*, 2024. 1, 2
- [59] 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 *ICML*, 2021. 6
- [60] Yurui Ren, Ge Li, Shan Liu, and Thomas H Li. Deep spatial transformation for pose-guided person image generation and animation. *IEEE TIP*, 2020. 2
- [61] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 6
- [62] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 3
- [63] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 7, 8
- [64] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *ICLR*, 2022. 3
- [65] Kripasindhu Sarkar, Dushyant Mehta, Weipeng Xu, Vladislav Golyanik, and Christian Theobalt. Neural re-rendering of humans from a single image. In *ECCV*, 2020. 2
- [66] Zhijing Shao, Zhaolong Wang, Zhuang Li, Duotun Wang, Xiangru Lin, Yu Zhang, Mingming Fan, and Zeyu Wang. Splattingavatar: Realistic real-time human avatars with mesh-embedded gaussian splatting. In *CVPR*, 2024. 1, 2
- [67] Kaiyue Shen, Chen Guo, Manuel Kaufmann, Juan Jose Zarate, Julien Valentin, Jie Song, and Otmar Hilliges. X-avatar: Expressive human avatars. In *CVPR*, 2023. 1
- [68] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *NeurIPS*, 2019. 2

- [69] Aliaksandr Siarohin, Oliver J Woodford, Jian Ren, Menglei Chai, and Sergey Tulyakov. Motion representations for articulated animation. In *CVPR*, 2021. 2
- [70] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022. 2
- [71] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *ICLR*, 2020. 3
- [72] Olga Sorkine, Daniel Cohen-Or, Yaron Lipman, Marc Alexa, Christian Rössl, and H-P Seidel. Laplacian surface editing. In *Eurographics*, 2004. 6
- [73] David Svitov, Pietro Morerio, Lourdes Agapito, and Alessio Del Bue. Haha: Highly articulated gaussian human avatars with textured mesh prior. *ACCV*, 2024. 1, 2
- [74] Zhengyan Tong, Chao Li, Zhaokang Chen, Bin Wu, and Wenjiang Zhou. Musepose: a pose-driven image-to-video framework for virtual human generation. *arxiv*, 2024. 2, 3
- [75] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 3
- [76] Jie Wang, Xianyan Li, Jiucheng Xie, Feng Xu, and Hao Gao. Gaussianhead: Impressive 3d gaussian-based head avatars with dynamic hybrid neural field. *arXiv e-prints*, pages arXiv–2312, 2023. 2
- [77] Tan Wang, Linjie Li, Kevin Lin, Yuanhao Zhai, Chung-Ching Lin, Zhengyuan Yang, Hanwang Zhang, Zicheng Liu, and Lijuan Wang. Disco: Disentangled control for realistic human dance generation. In *CVPR*, 2024. 2
- [78] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 2004. 6
- [79] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Humanerf: Free-viewpoint rendering of moving people from monocular video. In *CVPR*, 2022. 6, 7
- [80] Zhongcong Xu, Jianfeng Zhang, Jun Hao Liew, Hanshu Yan, Jia-Wei Liu, Chenxu Zhang, Jiashi Feng, and Mike Zheng Shou. Magicanimate: Temporally consistent human image animation using diffusion model. In *CVPR*, 2024. 2
- [81] Zhendong Yang, Ailing Zeng, Chun Yuan, and Yu Li. Effective whole-body pose estimation with two-stages distillation. In *ICCV*, 2023. 3
- [82] Jae Shin Yoon, Lingjie Liu, Vladislav Golyanik, Kripasindhu Sarkar, Hyun Soo Park, and Christian Theobalt. Pose-guided human animation from a single image in the wild. In *CVPR*, 2021. 2
- [83] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4d: Real-time human volumetric capture from very sparse consumer rgb sensors. In *CVPR*, 2021. 6, 8
- [84] Wing-Yin Yu, Lai-Man Po, Ray CC Cheung, Yuzhi Zhao, Yu Xue, and Kun Li. Bidirectionally deformable motion modulation for video-based human pose transfer. In *ICCV*, 2023. 2
- [85] Yifei Zeng, Yuanxun Lu, Xinya Ji, Yao Yao, Hao Zhu, and Xun Cao. Avatarbooth: High-quality and customizable 3d human avatar generation. *arXiv preprint arXiv:2306.09864*, 2023. 1
- [86] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. 2
- [87] Pengze Zhang, Lingxiao Yang, Jian-Huang Lai, and Xiaohua Xie. Exploring dual-task correlation for pose guided person image generation. In *CVPR*, 2022. 2
- [88] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 6
- [89] Zechuan Zhang, Zongxin Yang, and Yi Yang. Sifu: Side-view conditioned implicit function for real-world usable clothed human reconstruction. In *CVPR*, 2024. 7, 8
- [90] Jian Zhao and Hui Zhang. Thin-plate spline motion model for image animation. In *CVPR*, 2022. 2
- [91] Shunyu 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 *CVPR*, 2024. 2
- [92] Yang Zheng, Qingqing Zhao, Guandao Yang, Wang Yifan, Donglai Xiang, Florian Dubost, Dmitry Lagun, Thabo Beeler, Federico Tombari, Leonidas Guibas, et al. Physavatar: Learning the physics of dressed 3d avatars from visual observations. *ECCV*, 2024. 1, 2
- [93] Shenhao Zhu, Junming Leo Chen, Zuo Zhuo Dai, Yinghui Xu, Xun Cao, Yao Yao, Hao Zhu, and Siyu Zhu. Champ: Controllable and consistent human image animation with 3d parametric guidance. *ECCV*, 2024. 2
- [94] Wojciech Zielonka, Timur Bagautdinov, Shunsuke Saito, Michael Zollhöfer, Justus Thies, and Javier Romero. Drivable 3d gaussian avatars. *3DV*, 2023. 2