Human-3Diffusion: Realistic Avatar Creation via Explicit 3D Consistent Diffusion Models

Yuxuan Xue^{1,2} Xianghui Xie^{1,2,3} Riccardo Marin^{1,2} Gerard Pons-Moll^{1,2,3}

¹University of Tübingen ² Tübingen AI Center ³Max Planck Institute for Informatics, Saarland Informatics Campus https://yuxuan-xue.com/human-3diffusion/



Figure 1: Given a single image of a person (top), our method Human-3Diffusion creates 3D Gaussian Splats of realistic avatars with cloth and interacting objects with high-fidelity geometry and texture.

Abstract

Creating realistic avatars from a single RGB image is an attractive yet challenging problem. To deal with challenging loose clothing or occlusion by interaction objects, we leverage powerful shape prior from 2D diffusion models pretrained on large datasets. Although 2D diffusion models demonstrate strong generalization capability, they cannot provide multi-view shape priors with guaranteed 3D consistency. We propose **Human-3Diffusion**: Realistic Avatar Creation via Explicit **3D** Consistent **Diffusion**. Our key insight is that 2D multi-view diffusion and 3D reconstruction models provide complementary information for each other. By coupling them in a tight manner, we can fully leverage the potential of both models. We introduce a novel image-conditioned generative 3D Gaussian Splats reconstruction model that leverages the prior from 2D multi-view diffusion models, and provides an explicit 3D representation, which further guides the 2D reverse sampling process to have better 3D consistency. Experiments show that our proposed framework outperforms state-of-the-art methods and enables the creation of realistic avatars from a single RGB image, achieving high-fidelity in both geometry and appearance. Extensive ablations also validate the efficacy of our design, (1) multi-view 2D priors conditioning in generative 3D reconstruction and (2) consistency refinement

99601

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

of sampling trajectory via the explicit 3D representation. Our code and models will be released here.

1 Introduction

Realistic human avatar creation is crucial for various applications such as AR/VR, as well as the movie and gaming industry. Methods for creating a 3D avatar from a single RGB image are especially important to scale up avatar creation and make it more consumer-friendly compared to traditional studio-based capture methods. This task is, however, very challenging due to the vast diversity of human bodies and poses, further complicated by the wide variety of clothing and accessories. These challenges are exacerbated by the lack of large-scale 3D human data and ambiguities inherent in a monocular 2D view setting.

Recent image-to-3D approaches can be categorized into reconstruction-based and multi-view diffusion-based methods. Reconstruction-based approaches directly predict a 3D representation that can be rendered from any viewpoint. Due to the explicit 3D representation, these methods produce an arbitrary number of consistent viewpoint renderings. They obtain the 3D reconstruction either based on common template [23, 87, 88, 108] which utilize the SMPL [46] body model as the shape prior, or a flexible implicit function to represent loose clothing [5, 56, 57]. These methods, either template-based [23, 87, 88, 108] or template-free [5, 26, 56, 57, 74, 110], are typically deterministic which produce blurry textures and geometry in the occluded regions. More importantly, they are trained on small-scale datasets due to the limited amount of high-quality 3D data, which further restricts their ability to generalize to diverse shapes and textures.

Multi-view diffusion methods [42, 60, 77] distill the inherent 3D structure present in 2D diffusion models [55]. Typically, they fine-tune a large-scale 2D foundation model [25, 64] on a large 3D dataset of objects [13, 82, 99], to produce a *fixed* number of viewpoints. However, since these models diffuse images purely in 2D without explicit 3D constraints or representation, the resulting multi-views often lack 3D consistency [54, 41], which restricts downstream applications [67].

To address these challenges, we propose **Human-3Diffusion**: realistic avatar creation via **3D** consistent **Diffusion** models. We design our method based on two key insights: 1) 2D multi-view diffusion models provide large-scale shape priors that can help 3D reconstruction; 2) A reconstructed 3D representation ensures 3D consistency across multi-views in 2D diffusion. Specifically, we propose a novel diffusion method, which bridges 3D Gaussian Splatting (3D-GS) [36] generation with a 2D multi-view diffusion model. At every iteration, multi-view images are denoised and reconstructed to 3D-GS to be re-rendered to continue the diffusion process. This 3D lifting during iterative sampling ensures the 3D consistency of the 2D diffusion model while leveraging a large-scale foundation model trained on billions of images. Our framework elegantly combines reconstruction methods with multi-view diffusion models. In summary, our contributions are:

- We propose a novel image-conditioned 3D-GS generation model for 3D reconstruction that bridges large-scale priors from 2D multi-view diffusion models and the efficient and explicit 3D-GS representation.
- A sophisticated diffusion process that incorporates reconstructed 3D-GS to improve the 3D consistency of 2D diffusion models by refining the reverse sampling trajectory.
- Our proposed formulation enables us to jointly train 2D diffusion and our 3D model on ~ 6000 high-quality human scans and our method shows superior performance and generalization capability than prior works. Our code and pretrained models will be publicly released on our project page.

2 Related Work

Image to 3D. Creating realistic human avatar from consumer grade sensors [33, 96, 73, 92–94] is essential for downstream tasks such as human behaviour understanding [9, 51, 84, 86, 85] and gaming application [38, 44, 20, 105, 106]. Researchers have explored avatar creation from monocular RGB [31, 80], Depth [16, 93] video or single image [56, 57, 59, 87, 88]. Avatar from single image is particularly interesting and existing methods can be roughly categorized as template-based [23, 87, 88, 108] and template-free [56, 57, 59, 95]. Despite the impressive performance,

template-based approaches rely on the naked body model [46, 50] and fail to reason extremely loose clothing, while template-free methods produce blurry back side textures. Instead of SMPL shape prior, our method is template-free and leverages strong 2D image priors to create high-quality avatars. Orthogonal to humans, object reconstruction methods typically adopt template-free paradigms and early works [6, 71, 81, 83, 107] focus mainly on geometry. With the advance of 2D diffusion models [55] and efficient 3D representation [12], recent works can reconstruct 3D objects with detailed textures [26, 41, 45, 60, 67, 74, 89, 90, 110]. One popular paradigm is first using strong 2D models [42, 61, 77] to produce multi-view images and then train another model to reconstruct 3D from multi-view images [41, 40, 45, 67, 91]. In practice, their performance is limited by the accuracy of the multi-view images generated by 2D diffusion modes. Our method tightly couples 2D and 3D models and yields better performance by guiding 2D sampling with 3D reconstruction.

Shape Prior from 2D Diffusion Model. Being trained on billions of images [58], 2D image diffusion models [55] have been shown to have 3D awareness and some works tried to use score distillation sampling [53, 79] to distil 3D knowledge of 2D models [39, 48, 109]. Other works propose to further enhance the 3D reasoning ability by fine-tuning the model on large-scale datasets [13, 82, 99] to generate multi-view images [34, 37, 42, 43, 60, 61, 68, 76, 77]. Dense self-attention [75, 77], depth-aware attention [28] or epipolar attention [30, 66] are introduced to enhance the 3D consistency of multi-views. However, these methods do not have explicit 3D while our method incorporates explicit 3D consistency into the reverse sampling process and obtains better results.

3 Preliminaries

Denoising Diffusion Probabilistic Models. DDPM [25] is a generative model which learns a data distribution by iteratively adding (forward process) and removing (reverse process) the noise. Formally, the forward process iteratively adds noise to a sample \mathbf{x}_0 drawn from a distribution $p_{\text{data}}(\mathbf{x})$:

$$\mathbf{x}_{t} \sim \mathcal{N}(\mathbf{x}_{t}; \sqrt{\alpha_{t}} \mathbf{x}_{t-1}, (1 - \alpha_{t}) \mathbf{I}) := \sqrt{\bar{\alpha}_{t}} \mathbf{x}_{0} + \sqrt{1 - \bar{\alpha}_{t}} \epsilon, \text{ where } \epsilon \sim \mathcal{N}(0, \mathbf{I}),$$
 (1)

where $\alpha_t, \bar{\alpha}_t$ schedules the amount of noise added at each step t [25]. To sample data from the learned distribution, the reverse process starts from $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$ and iteratively denoises it until t = 0:

$$\mathbf{x}_{t-1} \sim \mathcal{N}(\mathbf{x}_{t-1}; \mu_{\theta}(\mathbf{x}_t, t), \tilde{\beta}_{t-1}\mathbf{I}), \text{ where } \tilde{\beta}_{t-1} = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} (1 - \alpha_t)$$
 (2)

A network parametrized by θ is trained to estimate the posterior mean μ_{θ} at each step t. One can also model conditional distribution with DDPM by adding the condition to the network input [15, 24].

2D Multi-View Diffusion Models. Many recent works [42, 43, 45, 60, 68, 77] propose to leverage strong 2D image diffusion prior [55] pre-trained on billions images [58] to generate multi-view images from a single image. Among them, ImageDream [77] demonstrated a superior generalization capability to unseen objects [67]. Given a single condition image \mathbf{x}^c and an optional text description y, ImageDream generate 4 orthogonal target views \mathbf{x}^{tgt} with a model ϵ_{θ} , which is trained to estimate the noise added at each step t. With the estimated noise ϵ_{θ} , one can compute the "clear" target views $\tilde{\mathbf{x}}_0^{\text{tgt}}$ with close-form solution in Eq. (1):

$$\tilde{\mathbf{x}}_0^{\text{tgt}} = \frac{1}{\sqrt{\bar{\alpha}_t}} (\mathbf{x}_t^{\text{tgt}} - \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_{\theta} (\mathbf{x}_t^{\text{tgt}}, \mathbf{x}^c, y, t)). \tag{3}$$

This *one-step* estimation of $\tilde{\mathbf{x}}_0^{\text{tgt}}$ can be noisy, especially when t is large and $\mathbf{x}_t^{\text{tgt}}$ is extremely noisy. Thus, the iterative sampling of $\mathbf{x}_t^{\text{tgt}}$ is required until t=0. To sample next step $\mathbf{x}_{t-1}^{\text{tgt}}$, standard DDPM [25] computes the posterior mean μ_{θ} from current $\mathbf{x}_t^{\text{tgt}}$ and estimated $\tilde{\mathbf{x}}_0^{\text{tgt}}$ at step t with:

$$\mu_{\theta}(\mathbf{x}_{t}^{\text{tgt}}, t) := \mu_{t-1}(\mathbf{x}_{t}^{\text{tgt}}, \tilde{\mathbf{x}}_{0}^{\text{tgt}}) = \frac{\sqrt{\alpha_{t}} \left(1 - \bar{\alpha}_{t-1}\right)}{1 - \bar{\alpha}_{t}} \mathbf{x}_{t}^{\text{tgt}} + \frac{\sqrt{\bar{\alpha}_{t-1}} \beta_{t}}{1 - \bar{\alpha}_{t}} \tilde{\mathbf{x}}_{0}^{\text{tgt}}, \text{ where } \beta_{t} = 1 - \alpha_{t}.$$
 (4)

Afterwards, $\mathbf{x}_{t-1}^{\text{tgt}}$ can be sampled from Gaussian distribution with mean μ_{t-1} and variance $\tilde{\beta}_{t-1}\mathbf{I}$ (Eq. (2)) and used as the input for the next iteration. The reverse sampling is repeated until t=0 where 4 clear target views are generated.

Although multi-view diffusion models [43, 60, 77] generate multiple views together, the 3D consistency across these views is not guaranteed due to the lack of an explicit 3D representation. Thus, we propose a novel 3D consistent diffusion model, which ensures the multi-view consistency at each step of the reverse process by diffusing 2D images using reconstructed 3D Gaussian Splats [36].

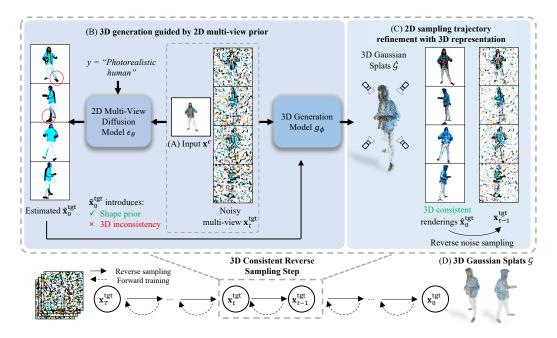


Figure 2: **Method Overview.** Given a single RGB image (A), we sample a realistic 3D avatar represented as 3D Gaussian Splats (D). At each reverse step, our 3D generation model g_{ϕ} leverages 2D multi-view diffusion prior from ϵ_{θ} which provides a strong shape prior but is not 3D consistent (B, cf. Sec. 4.1). We then refine the 2D reverse sampling trajectory with generated 3D renderings that are guaranteed to be 3D consistent (C, cf. Sec. 4.2). Our tight coupling ensures 3D consistency at each sampling step and obtains a high-quality 3D avatar (D).

4 3Diffusion

Overview. Given a single RGB image, we aim to create a realistic 3D avatar consistent with the input. We adopt an image-conditioned 3D generation paradigm due to inherent ambiguities in the monocular view. We introduce a novel 3D Gaussian Splatting (3D-GS [36]) generative model that combines shape priors from 2D multi-view diffusion models with the explicit 3D-GS representation. This allows us to jointly train our 3D generative model and a 2D multi-view diffusion model end-to-end and improves the 3D consistency of 2D multi-view generation at inference time.

In this section, we first introduce our novel generative 3D-GS reconstruction model in Sec. 4.1. We then describe how we leverage the 3D reconstruction to generate 3D consistent multi-view results by refining the reverse sampling trajectory (Sec. 4.2). An overview of our method can be found in Fig. 2.

4.1 Generative 3D-GS Reconstruction with Diffusion Priors

Given a context image \mathbf{x}^c , we use a conditional diffusion model to learn and sample from a plausible 3D distribution. Previous works demonstrated that 3D generation can be done implicitly via diffusing rendered images of a differentiable 3D representation [7, 35, 70] such as NeRF [49, 97]. In this work, we introduce a novel generative model for 3D Gaussian Splatings [36], which diffuses rendered images of 3D-GS and enables sampling of 3D-GS at inference time. Single image to 3D generation is however very challenging, we hence propose to leverage 2D multi-view diffusion models in a tightly coupled manner which allows us to train it end-to-end with our novel 3D generative model.

Generative 3D-GS Reconstruction. In this work, we propose a 3D-GS generative model g_{ϕ} , which is conditioned on input context image \mathbf{x}^c to perform reconstruction of 3D Gaussian Splats \mathcal{G} . Diffusing directly in the space of \mathcal{G} parameters requires pre-computing Gaussian Splats from scans, which is exorbitant. Instead, we diffuse the multi-view renderings of \mathcal{G} using a differentiable rendering function renderer.

We denote $\mathbf{x}_0^{\text{tgt}}$ as the ground truth images at target views to be diffused and $\mathbf{x}_0^{\text{novel}}$ as the additional

novel views for supervision. At training time, we uniformly sample a timestep $t \sim \mathcal{U}(0,T)$ and add noise to $\mathbf{x}_0^{\text{tgt}}$ using Eq. (1) to obtain noisy target views $\mathbf{x}_t^{\text{tgt}}$. Our generative model g_ϕ takes $\mathbf{x}_t^{\text{tgt}}$, diffusion timestep t, and the conditional image \mathbf{x}^c as input, and estimates 3D Gaussians $\hat{\mathcal{G}}$:

$$\hat{\mathcal{G}} = g_{\phi}(\mathbf{x}_{t}^{\text{tgt}}, t, \mathbf{x}^{c}), \text{ where } \mathbf{x}_{t}^{\text{tgt}} = \sqrt{\bar{\alpha}_{t}} \mathbf{x}_{0}^{\text{tgt}} + \sqrt{1 - \bar{\alpha}_{t}} \epsilon, \text{ and } \epsilon \sim \mathcal{N}(0, \mathbf{I})$$
 (5)

We adopt an asymmetric U-Net Transformer proposed by [67] for g_{ϕ} to directly predict 3D-GS parameters from per-pixel features of the last U-Net layer. To supervise the generative model g_{ϕ} , we use a differentiable rendering function renderer : $\{\mathcal{G}, \pi^p\} \mapsto \mathbf{x}^p$ to render images at target views π^{tgt} and additional novel views π^{novel} . Denoting $\mathbf{x}_0 := \{\mathbf{x}_0^{\text{tgt}}, \mathbf{x}_0^{\text{novel}}\}$ as ground truth and $\hat{\mathbf{x}}_0 := \{\hat{\mathbf{x}}_0^{\text{tgt}}, \hat{\mathbf{x}}_0^{\text{novel}}\}$ as rendered images, we compute the loss on images and generated 3D-GS:

$$\mathcal{L}_{gs} = \lambda_1 \cdot \mathcal{L}_{\text{MSE}}(\mathbf{x}_0, \hat{\mathbf{x}}_0) + \lambda_2 \cdot \mathcal{L}_{\text{Percep}}(\mathbf{x}_0, \hat{\mathbf{x}}_0) + \lambda_3 \cdot \mathcal{L}_{\text{reg}}(g_{\phi}(\mathbf{x}_t^{\text{tgt}}, t, \mathbf{x}^c)),$$
where $\hat{\mathbf{x}}_0 := \{\hat{\mathbf{x}}_0^{\text{tgt}}, \hat{\mathbf{x}}_0^{\text{novel}}\} = \text{renderer}(g_{\phi}(\mathbf{x}_t^{\text{tgt}}, t, \mathbf{x}^c), \{\pi^{\text{tgt}}, \pi^{\text{novel}}\}),$
(6)

here \mathcal{L}_{MSE} denotes the Mean Square Error (MSE) and \mathcal{L}_{Percep} is the perceptual loss based on VGG-19 [62]. We also apply \mathcal{L}_{reg} , a geometry regularizer [29, 100] to stabilize the generation of $\hat{\mathcal{G}}$.

With this, we can train a generative model that diffuses 3D-GS implicitly by diffusing 2D images $\mathbf{x}_t^{\text{tgt}}$. At inference time, we can generate 3D-GS given the input image by denoising 2D multi-views sampled from Gaussian distribution. We initialize $\mathbf{x}_T^{\text{tgt}}$ from $\mathcal{N}(0,\mathbf{I})$, and iteratively denoise the rendered images of predicted $\hat{\mathcal{G}}$ from our model g_{ϕ} . At each reverse step, our model g_{ϕ} estimates a clean state $\hat{\mathcal{G}}$ and render target images $\hat{\mathbf{x}}_0^{\text{tgt}}$. We then calculate target images $\mathbf{x}_{t-1}^{\text{tgt}}$ for the next step via Eq. (4) and repeat the process until t=0. For more details, please refer to Appendix A.3

Our generative 3D-GS reconstruction model archives superior performance on in-distribution human reconstruction yet generalizes poorly to unseen categories such as general objects (Sec. 5.3 Fig. 5). Our key insight for better generalization is leveraging strong priors from pretrained 2D multi-view diffusion models for 3D-GS generation.

3D-GS Generation with 2D Multi-view Diffusion. Pretrained 2D multi-view diffusion models (MVD) [43, 61, 77] have seen billions of real images [58] and millions of 3D data [13], which provide strong prior information and can generalize to unseen objects [67, 89]. Here, we propose a simple yet elegant idea for incorporating this multi-view prior into our generative 3D-GS model g_{ϕ} . We can also leverage generated 3D-GS to guide 2D MVD sampling process which we discuss in Sec. 4.2. Our key observation is that both 2D MVD and our proposed 3D-GS generative model are diffusion-based and share the same sampling state $\mathbf{x}_t^{\text{tgt}}$ at timestep t. Thus, they are *synchronized*. This enables us to couple and facilitate information exchange between 2D MVD ϵ_{θ} and 3D-GS generative model g_{ϕ} at the same diffusion timestep t. To inject the 2D diffusion priors into 3D generation, we first compute *one-step* estimation of $\tilde{\mathbf{x}}_0^{\text{tgt}}$ (Eq. (3)) using 2D MVD ϵ_{θ} , and condition our 3D-GS generative mode g_{ϕ} additionally on it. Formally, our 3D-GS generative model enhanced with 2D multi-view diffusion priors is written as:

$$\hat{\mathcal{G}} = g_{\phi}(\mathbf{x}_{t}^{\text{tgt}}, t, \mathbf{x}^{c}, \tilde{\mathbf{x}}_{0}^{\text{tgt}}), \text{ where } \tilde{\mathbf{x}}_{0}^{\text{tgt}} = \frac{1}{\sqrt{\bar{\alpha}_{t}}} (\mathbf{x}_{t}^{\text{tgt}} - \sqrt{1 - \bar{\alpha}_{t}} \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_{t}^{\text{tgt}}, \mathbf{x}^{c}, y, t))$$
(7)

The visualization of $\tilde{\mathbf{x}}_0^{\text{tgt}}$ along the whole sampling trajectory in Fig. 7 shows that the pretrained 2D diffusion model ϵ_{θ} can already provide useful multi-view shape prior even in large timestep t=1000. This is further validated in our experiments where the additional 2D diffusion prior $\tilde{\mathbf{x}}_0^{\text{tgt}}$ leads to better avatar reconstruction (Tab. 4) as well as more robust generalization to general objects (Fig. 5). By utilizing the timewise iterative manner of 2D and 3D diffusion models, we can not only leverage 2D priors for 3D-GS generation but also train both models jointly end to end, which we discuss next.

Joint Training with 2D Model. We adopt pretrained ImageDream [77] as our 2D multi-view diffusion model ϵ_{θ} and jointly train it with our 3D-GS generative model g_{ϕ} . We observe that our joint training is important for coherent 3D generation, as opposed to prior works that frozen pretrained 2D multi-view models [67, 74]. We summarize our training algorithm in Algorithm 1. We combine the loss of 2D diffusion and our 3D-GS generation loss \mathcal{L}_{qs} (Eq. (6)):

$$\mathcal{L}_{total} = \mathcal{L}_{MSE}(\epsilon, \epsilon_{\theta}) + \mathcal{L}_{qs}$$
(8)

Algorithm 1 Training

Input: Dataset of posed multi-view images $\mathbf{x}_0^{\text{tgt}}$, π^{tgt} , $\mathbf{x}_0^{\text{novel}}, \pi^{\text{novel}}$, a context image \mathbf{x}^c , text description y **Output:** Optimized 2D multi-view diffusion model ϵ_{θ} and 3D-GS generative model g_{ϕ}

1: **repeat**
2:
$$\{\mathbf{x}_0^{\text{tgt}}, \mathbf{x}_0^{\text{novel}}, \mathbf{x}^c, y\} \sim q(\{\mathbf{x}_0^{\text{tgt}}, \mathbf{x}_0^{\text{novel}}, \mathbf{x}^c, y\})$$
3: $t \sim \text{Uniform}(\{1, \dots, T\}); \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
4: $\mathbf{x}_t^{\text{tgt}} = \sqrt{\overline{\alpha}_t} \mathbf{x}_0^{\text{tgt}} + \sqrt{1 - \overline{\alpha}_t} \epsilon$
5: $\tilde{\mathbf{x}}_0^{\text{tgt}} = \frac{1}{\sqrt{\overline{\alpha}_t}} (\mathbf{x}_t^{\text{tgt}} - \sqrt{1 - \overline{\alpha}_t} \epsilon_{\theta}(\mathbf{x}_t^{\text{tgt}}, \mathbf{x}^c, y, t))$

4:
$$\mathbf{x}_t^{\text{tgt}} = \sqrt{\bar{\alpha}_t} \mathbf{x}_0^{\text{tgt}} + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}$$

5: $\tilde{\mathbf{x}}_t^{\text{tgt}} = \frac{1}{2} (\mathbf{x}_t^{\text{tgt}} - \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}) (\mathbf{x}_t^{\text{tgt}} - \mathbf{x}_t^{\text{c}})$

6:
$$\hat{\mathcal{G}} = g_{\phi}\left(\mathbf{x}_{t}^{\text{tgt}}, t, \mathbf{x}^{\text{c}}, \tilde{\mathbf{x}}_{0}^{\text{ggt}}\right)$$
 // Enhance conditional 3D generation with 2D diffusion prior $\tilde{\mathbf{x}}_{0}^{\text{tgt}}$ from ϵ_{θ}

7:
$$\{\hat{\mathbf{x}}_0^{\text{tgt}}, \hat{\mathbf{x}}_0^{\text{novel}}\} = \text{renderer}\left(\hat{\mathcal{G}}, \{\pi^{\text{tgt}}, \pi^{\text{novel}}\}\right)$$

8: Compute loss
$$\mathcal{L}_{total}$$
 (Eq. (8))

- Gradient step to update ϵ_{θ} , g_{ϕ}
- 10: until converged

Algorithm 2 3D Consistent Sampling

Input: A context image \mathbf{x}^c and text y; Converged 2D diffusion model ϵ_{θ} and 3D generative model g_{ϕ} **Output:** A 3D Gaussian Avatar \mathcal{G} of the 2D image \mathbf{x}^c

1:
$$\mathbf{x}_{T}^{\text{tgt}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

2: $\mathbf{for} \ t = T, \dots, 1 \ \mathbf{do}$
3: $\tilde{\mathbf{x}}_{0}^{\text{tgt}} = \frac{1}{\sqrt{\bar{\alpha}_{t}}} (\mathbf{x}_{t}^{\text{tgt}} - \sqrt{1 - \bar{\alpha}_{t}} \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_{t}^{\text{tgt}}, \mathbf{x}^{c}, y, t))$
4: $\hat{\mathcal{G}} = g_{\phi} \left(\mathbf{x}_{t}^{\text{tgt}}, t, \mathbf{x}^{c}, \tilde{\mathbf{x}}_{0}^{\text{tgt}} \right)$

5:
$$\hat{\mathbf{x}}_{t}^{\text{tgt}} = \text{renderer} \left(\hat{\mathcal{C}}, \pi^{\text{tgt}}\right)$$

5:
$$\hat{\mathbf{x}}_{0}^{\text{tgt}} = \text{renderer}\left(\hat{\mathcal{G}}, \pi^{\text{tgt}}\right)$$
6: $\mu_{t-1}(\mathbf{x}_{t}^{\text{tgt}}, \hat{\mathbf{x}}_{0}^{\text{tgt}}) = \frac{\sqrt{\alpha_{t}}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_{t}} \mathbf{x}_{t}^{\text{tgt}} + \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_{t}}{1-\bar{\alpha}_{t}} \hat{\mathbf{x}}_{0}^{\text{tgt}} / \mathcal{A}_{t}^{\text{tgt}}$
Guide 2D sampling with 3D consistent renderings

7:
$$\mathbf{x}_{t-1}^{\text{tgt}} \sim \mathcal{N}\left(\mathbf{x}_{t-1}^{\text{tgt}}; \tilde{\boldsymbol{\mu}}_{t}\left(\mathbf{x}_{t}^{\text{tgt}}, \hat{\mathbf{x}}_{0}^{\text{tgt}}\right), \tilde{\beta}_{t-1}\mathbf{I}\right)$$
8: end for

9: **return**
$$\mathcal{G} = g_{\phi} \left(\mathbf{x}_{0}^{\text{tgt}}, \tilde{\mathbf{x}}_{0}^{\text{tgt}}, \mathbf{x}^{\text{c}}, t = 0 \right)$$

Once trained, one can sample a plausible 3D-GS avatar \mathcal{G} conditioned on the input image from the learned 3D distributions. However, we observe that the multi-view diffusion model ϵ_{θ} can still output inconsistent multi-views along the sampling trajectory (see Fig. 2). On the other hand, our 3D generator produces explicit 3D-GS which can be rendered as 3D consistent multi-views. Our second key idea is to use the 3D consistent renderings to guide 2D sampling process for more 3D consistent mulit-view generation. We discuss this in Sec. 4.2.

Guide 2D Multi-view Sampling with Reconstructed 3D-GS

With the shared and synchronized sampling state $\mathbf{x}_t^{\text{tgt}}$ of 2D multi-view diffusion model ϵ_{θ} and 3D-GS reconstruction model g_{ϕ} , we couple both models at arbitrary t during training. Similarly, they are also connected by both using estimated clean multi-views \mathbf{x}_0^{tgt} at sampling time. To leverage the full potential of both models, we carefully design a joint sampling process that utilizes the reconstructed 3D-GS $\hat{\mathcal{G}}$ at each timestep t to guide 2D multi-view sampling, which is summarized in Algorithm 2. We observe that the key difference between the clean multi-views estimated ${f x}_0^{
m tgt}$ from 2D diffusion model and our 3D-GS generation lies in 3D consistency: 2D MVD computes multi-view $\tilde{\mathbf{x}}_0^{\text{tgt}}$ from 2D network prediction which can be 3D inconsistent while our $\hat{\mathbf{x}}_0^{\text{tgt}}$ are rendered from explicit 3D-GS representation which are guaranteed to be 3D consistent. Our idea is to guide the 2D multi-view reverse sampling process with our 3D consistent renderings $\hat{\mathbf{x}}_0^{\text{tgt}}$ such that the 2D sampling trajectory is more 3D consistent. Specifically, we leverage 3D consistent multi-view renderings $\hat{\mathbf{x}}_0^{tgt}$ to refine the posterior mean $\mu_{\theta}(\mathbf{x}_{t}^{\text{tgt}}, t)$ at each reverse step:

Original:
$$\mu_{\theta}(\mathbf{x}_{t}^{\text{tgt}}, t) := \mu_{t-1}(\mathbf{x}_{t}^{\text{tgt}}, \hat{\mathbf{x}}_{0}^{\text{tgt}}) \rightarrow \text{Ours: } \mu_{\theta}(\mathbf{x}_{t}^{\text{tgt}}, t) := \mu_{t-1}(\mathbf{x}_{t}^{\text{tgt}}, \hat{\mathbf{x}}_{0}^{\text{tgt}}),$$
where $\hat{\mathbf{x}}_{0}^{\text{tgt}} = \text{renderer}(\hat{\mathcal{G}}, \pi^{\text{tgt}}), \text{ and } \mu_{t-1}(\mathbf{x}_{t}^{\text{tgt}}, \hat{\mathbf{x}}_{0}^{\text{tgt}}) = \frac{\sqrt{\alpha_{t}} (1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_{t}} \mathbf{x}_{t}^{\text{tgt}} + \frac{\sqrt{\bar{\alpha}_{t-1}} \beta_{t}}{1 - \bar{\alpha}_{t}} \hat{\mathbf{x}}_{0}^{\text{tgt}}$ (9)

With this refinement, we guarantee the 3D consistency at each reverse step t and avoid 3D inconsistency accumulation in original multi-view sampling [77]. In Fig. 7, we visualize the evolution of originally generated multi-views $\tilde{\mathbf{x}}_0^{\text{tgt}}$ and multi-views rendering $\hat{\mathbf{x}}_0^{\text{tgt}}$ from generated 3D-GS $\hat{\mathcal{G}}$ along the whole reverse sampling process. It intuitively shows how effective the sampling trajectory refinement is. We perform extensive ablation in Sec. 5.3 showing the importance of the consistent refinement for sampling trajectory.

5 **Experiments**

In this section, we first compare against baseline methods for human reconstruction in Sec. 5.2 and then ablate our design choices in Sec. 5.3.

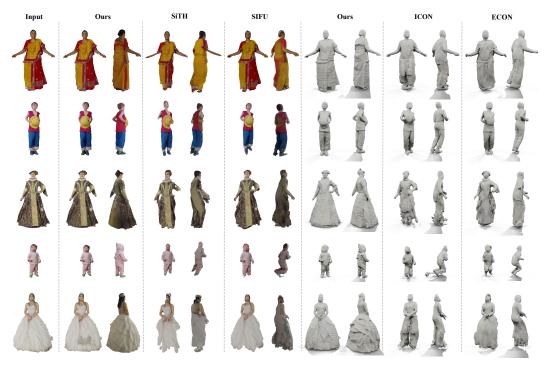


Figure 3: **Qualitative comparison with baselines**. Recent avatar reconstruction works ICON [87], ECON [88], SiTH [23] and SIFU [108]) cannot reconstruct loose clothing coherently. Additionally, SiTH and SIFU generate blurry texture in unseen regions due to their deterministic formulation of regressing 3D avatar directly from single RGB imagse. In contract, our method is able to reconstruct avatars with realistic textures and plausible 3D geometry in both seen and unseen region.

5.1 Experimental Setup

Datasets. We train our model on a combined 3D human dataset [1, 3, 4, 2, 21, 27, 65, 98] compromising ~ 6000 high quality scans. We evaluate qualitatively and quantitatively on CAPE [47, 52, 103], Sizer [8, 72] and IIIT [32] dataset. Please refer to Appendix D.1 and Appendix D.2 for more details.

Implementation Details. We trained our model on 8 NVIDIA A100 GPUs over approximately 5 days. Each GPU was configured with a batch size 2 and gradient accumulations of 16 steps to achieve an effective batch size of 256. For more training details regarding hyperparameters, diffusion schedulers, etc., please refer to Appendix A.1 for more details. Our model creates 3D Avatar from single images in 22.6 seconds on a A100 GPU and only consumes 11.7 GB VRAM, which allows the efficient large-scale avatar generation.

Evaluation Metrics. We evaluate the geometry quality using Chamfer Distance (CD in cm), Point-to-Surface Distance (P2S in cm), F-score [69] (w/ threshold of 0.01m), and Normal Consistency (NC) between the extracted mesh (Appendix A.4) and the groundtruth scan. Appearance quality is assessed by rendering the reconstructed avatar from 32 novel views with uniform azimuth and 0 elevation angle. The metrics for appearance reported include multi-scale Structure Similarity (SSIM) [78], Learned Perceptual Image Patch Similarity (LPIPS) [104], and Peak Signal to Noise Ratio (PSNR) between rendered and ground-truth views. Moreover, we report the Fréchet inception distance (FID) [22] which reflects the quality and realism of the unseen regions.

5.2 Realistic Avatar from Image

We compare our approach against prior methods for image-to-avatar reconstruction, including template-based [18, 23, 87, 88, 56, 108], template-free [56] human reconstruction methods, as well as general image-to-3D methods [67, 74, 89]. To further assess performance, we also fine-tuned the state-of-the-art object reconstruction method LGM [67] and its deployed multi-view diffusion model [77] on our training data, denoted as LGM_{human} . Quantitative evaluations reported in Tab. 1 demonstrate

Method	CD↓	P2S↓	F-score ↑	NC ↑	SSIM ↑	LPIPS ↓	PSNR ↑	FID↓
PIFu [56]	2.75	2.68	0.359	0.778	0.909	0.077	21.06	29.57
SiTH [23]	4.00	4.00	0.257	0.749	0.907	0.073	20.00	22.33
SIFU [108]	3.50	3.50	0.273	0.760	0.0.900	0.081	20.73	40.75
LGM _{human}	3.44	3.44	0.272	0.560	0.893	0.088	20.56	14.22
FoF [18]	5.43	5.34	0.183	0.683	-	-	-	-
ICON [87]	3.88	3.94	0.244	0.749	-	-	-	-
ECON [88]	2.83	3.61	0.291	0.767	-	-	-	-
Ours	1.41	1.37	0.557	0.791	0.916	0.058	21.61	8.45

Table 1: **Quantitative evaluation** on CAPE [47], SIZER [72], and IIIT [32] dataset. Our method can perform better reconstruction in terms of more accurate geometry (CD, P2S, F-score, NC) and realistic textures (SSIM, LPIPS, PSNR, FID).

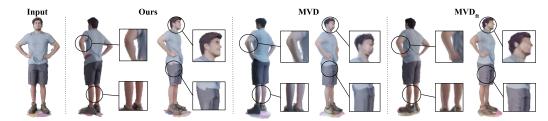


Figure 4: **3D reconstruction conditioned on different multi-view priors.** Without our 3D-consistent sampling, the 2D diffusion model cannot generate 3D consistent multi-views (MVD, MVD_{ft}), leading to artifacts like floating 3D Gaussians splats.

that our proposed method excels in reconstructing realistic avatars with accurate geometry (CD, NC, F-score) and realistic texture (SSIM, LPIPS, PSNR, FID) from a single RGB image.

We present qualitative comparison examples in Fig. 3 and Appendix B.1, highlighting the strengths and weaknesses of competing methods. Template-based methods such as SiTH [23] and SIFU [108] struggle to accurately reconstruct the geometry of loose clothing (as shown in row 4) due to their reliance on the naked SMPL body shape. In contrast, template-free methods like PIFu [56] and TripoSR [74] offer greater flexibility and better performance on loose clothing. However, they are not generative models and their deterministic formulations lead to blurry textures in unseen regions, as they tend to produce average textures rather than distinct details. Similar to our approach, LGM [67] and InstantMesh [89] utilize 2D diffusion models to generate multi-view images and perform sparse-view 3D reconstruction. Nonetheless, their separation of 2D and 3D models cannot correct the 3D inconsistencies that may arise from the 2D models. Even further fine-tuning of LGM on human scans (Fig. 4) does not adequately address these challenges due to the complex and sensitive nature of human geometry and textures. In contrast, our conditional generative formulation and inherent 3D consistency by tightly coupling 2D-3D models allow us to obtain accurate reconstruction in front view and realistic generation in unseen regions. We also show the generative power of our method in Appendix C.5: by sampling with different seed, we obtain diverse yet plausible reconstruction.

Please also refer to Fig. 6, Appendix C and our project page for additional reconstruction results on challenging subjects not previously observed, encompassing a diverse range of appearances such as loose skirts and custom suits, as well as accessories like bags and gloves.

5.3 Ablation Studies

Importance of Trajectory Refinement. One of our key ideas is leveraging our explicit 3D model to refine the 2D multi-view reverse sampling trajectory, ensuring 3D consistency in Multi-View Diffusion (MVD) generation (see Sec. 4.2 and Eq. (9)). To evaluate this, we compare the multi-view images generated by pretrained MVD, fine-tuned MVD on our data (MVD_{ft}) and MVD with our 3D consistent sampling

Method	LPIPS ↓	SSIM ↑	PSNR ↑
MVD	0.078	0.911	22.32
$\mathrm{MVD}_{\mathrm{ft}}$	0.061	0.926	24.14
Ours	0.048	0.934	24.69

Table 2: **Evaluating trajectory refinement** for 2D multi-view diffusion. Our proposed refinement improves multi-view image quality.

(ours), as shown in Tab. 2. The results demonstrate that our proposed method effectively enhances the quality of generated multi-view images by leveraging the explicit 3D model to refine sampling trajectory. Additionally, we analyze the 3D reconstruction results with the multi-view images generated by these models in Fig. 4. MVD and MVD_{ft} produce inconsistent multi-view images, which typically lead to floating Gaussian and hence blurry boundaries. In contrast, our method can generate more consistent multi-views, result in better 3D Gaussians Splats and sharper renderings.

We further quantitatively evaluate the impact of our proposed sampling trajectory refinement on

Method	CD _(cm) ↓	F-score↑	NC ↑	LPIPS↓	SSIM↑	PSNR↑
Our w/o Traj. Ref.	1.57	0.498	0.794	0.064	0.908	21.09
Ours	1.35	0.550	0.798	0.060	0.918	21.49

Table 3: **Evaluating trajectory refinement** for final 3D reconstruction. Our sampling trajectory refinement ensures multi-view consistency and hence yields better 3D results.

final 3D reconstruction in Tab. 3. We compare the reconstruction results of methods with and without our trajectory refinement while using the same 2D MVD and 3D reconstruction models. It can be clearly seen that our trajectory refinement improves the quality of 3D reconstruction.

Importance of 2D Multi-view Prior. Another key idea of our work is the use of multi-view priors $\tilde{\mathbf{x}}_0^{\text{tgt}}$ from 2D diffusion model pretrained on massive data [13, 55, 58] to enhance our 3D generative model. This additional prior information is pivotal for ensuring accurate reconstruction of both in-distribution human dataset and generalizing to out-of-distribution objects.

We evaluate the performance of our 3D model g_{ϕ} by comparing generation results with and without the 2D diffusion prior $\tilde{\mathbf{x}}_0^{\text{tgt}}$ (refer to Eq. (7) and Eq. (5)). Notably, without the 2D multi-view conditioning, the alignment of the generated 3D model in the front view is not guaranteed due to the relative camera pose settings in our 3D generative model g_{ϕ} . Therefore, we evaluate the overall quality solely through the Fréchet Inception Distance (FID).

Method	PSNR ↑
Ours w/o $\tilde{\mathbf{x}}_0^{\text{tgt}}$	20.98
Our full model	21.49

Table 4: **2D multi-view priors** $\tilde{\mathbf{x}}_0^{\text{tgt}}$ improve human reconstruction quality.

For avatars reconstruction, our powerful 3D reconstruction model can already achieve state-of-the-art performance. Moreover, our full model with multi-view prior $\tilde{\mathbf{x}}_0^{\text{tgt}}$ generates avatars with higher quality as demonstrated in Tab. 4. We further evaluate it on the GSO [17] dataset which consists of unseen general objects to our model. The improvements are even more pronounced in this setting, highlighting the challenges of generating coherent 3D structures from a single 2D image, particularly with unseen objects. For additional examples, please see Fig. 15 in Supp..

Ours 16.12	Input		Ours		generative w/o	•
w/o $\tilde{\mathbf{x}}_0^{\text{tgt}}$ 14.45		4		76		1
Method PSNR ↑			1	DU	105	
	1000	10 C C C C C C C C C C C C C C C C C C C		Will do	500	

Figure 5: 2D multi-view priors $\tilde{\mathbf{x}}_0^{\text{tgt}}$ enhances generalization to general objects in GSO [17] dataset.

6 Limitations and Future Work

Currently, our method is constrained by the 256×256 resolution of the multi-view diffusion model, which restricts the sharpness of texture details (see Appendix E). Upgrading to more powerful high-resolution (512×512) multi-view diffusion models [19, 68] could potentially resolve these issues. Moreover, our approach may struggle in reconstructing subjects with challenging poses, as we further discussed in Appendix E. Synthesizing training data with challenging poses [10, 86] could be a potential solution.

Our method is a general framework for image-to-3D reconstruction, which is applicable to various objects and compositional shapes like human-object interactions. We leave these to future works.

7 Conclusion

In this paper, we introduce **Human-3Diffusion**, a 3D consistent diffusion model for creating realistic avatars from single RGB images. Our key ideas are two folds: 1) Leveraging strong multi-view priors

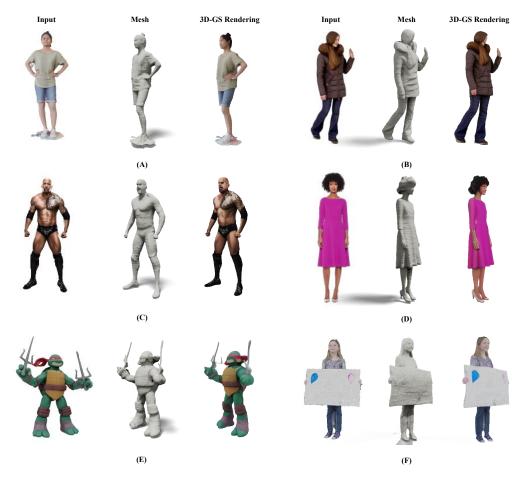


Figure 6: Visualization of reconstructed mesh and synthesized novel view of generated 3D-GS on subjects from Sizer [72], RenderPeople [2], Twindom [4], UBC Fashion [101], GSO [17] and online image. More results are presented in Appendix C and our project page.

from pretrained 2D diffusion models to generate 3D Gaussian Splats, and 2) Using the reconstructed explicit 3D Gaussian Splats to refine the sampling trajectory of the 2D diffusion model which enhances 3D consistency. We carefully designed a diffusion process that synergistically combines the strengths of both 2D and 3D models. Our experiments show that our approach outperforms all previous reconstruction works in both appearance and geometry. We also extensively ablate our method which proves the effectiveness of our proposed ideas. Our code and pretrained models will be released on our Project Page to foster future research.

Acknowledgements We appreciate G.Tiwari, Y.He, Y. Xiu, Z.Liu, Z.Qiu, S.Li and others for their feedback to improve the work. This work is made possible by funding from the Carl Zeiss Foundation. This work is also funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 409792180 (EmmyNoether Programme, project: Real Virtual Humans) and the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Y.Xue. R. Marin has been supported by innovation program under the Marie Skłodowska-Curie grant agreement No 101109330. G. Pons-Moll is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645.

References

- [1] Axyz, Nov 2023. URL https://secure.axyz-design.com. 7, 34, 36
- [2] Renderpeople, Nov 2023. URL https://renderpeople.com/. 7, 10, 34
- [3] Treedy, Nov 2023. URL https://treedys.com/. 7, 34
- [4] Twindom, Nov 2023. URL https://web.twindom.com/. 7, 10, 34, 36
- [5] Thiemo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic Monocular 3D Reconstruction of Humans Wearing Clothing. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 1496–1505, New Orleans, LA, USA, June 2022. IEEE. ISBN 978-1-66546-946-3. doi: 10.1109/CVPR52688.2022.00156. URL https://ieeexplore.ieee.org/document/9878998/. 2
- [6] Kalyan Vasudev Alwala, Abhinav Gupta, and Shubham Tulsiani. Pre-train, self-train, distill: A simple recipe for supersizing 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3773–3782, June 2022. 3
- [7] Titas Anciukevicius, Zexiang Xu, Matthew Fisher, Paul Henderson, Hakan Bilen, Niloy J. Mitra, and Paul Guerrero. Renderdiffusion: Image diffusion for 3d reconstruction, inpainting and generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 12608–12618. IEEE, 2023. doi: 10.1109/CVPR52729.2023.01213. URL https://doi.org/10.1109/CVPR52729.2023.01213. 4, 21
- [8] Dimitrije Antic, Garvita Tiwari, Batuhan Ozcomlekci, Riccardo Marin, and Gerard Pons-Moll. Close: A 3d clothing segmentation dataset and model. In *International Conference on 3D Vision*, 3DV 2024, Davos, Switzerland, March 18-21, 2024, pages 591–601. IEEE, 2024. doi: 10.1109/3DV62453.2024.00020. URL https://doi.org/10.1109/3DV62453.2024.00020. 7
- [9] Bharat Lal Bhatnagar, Xianghui Xie, Ilya Petrov, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Behave: Dataset and method for tracking human object interactions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, jun 2022. 2
- [10] Michael J. Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. BEDLAM: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 8726–8737, June 2023. 9
- [11] Xu Cao, Hiroaki Santo, Boxin Shi, Fumio Okura, and Yasuyuki Matsushita. Bilateral normal integration. 2022. 23
- [12] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J. Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3d generative adversarial networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 16102–16112. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01565. URL https://doi.org/10.1109/CVPR52688.2022.01565. 3
- [13] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 13142–13153. IEEE, 2023. doi: 10.1109/CVPR52729.2023.01263. URL https://doi.org/10.1109/CVPR52729.2023.01263. 2, 3, 5, 9, 34
- [14] Maximilian Denninger, Dominik Winkelbauer, Martin Sundermeyer, Wout Boerdijk, Markus Knauer, Klaus H. Strobl, Matthias Humt, and Rudolph Triebel. Blenderproc2: A procedural pipeline for photorealistic rendering. *Journal of Open Source Software*, 8(82):4901, 2023. doi: 10.21105/joss.04901. URL https://doi.org/10.21105/joss.04901. 34

- [15] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, Advances in Neural Information Processing Systems, 2021. URL https://openreview. net/forum?id=AAWuCvzaVt. 3
- [16] Zijian Dong, Chen Guo, Jie Song, Xu Chen, Andreas Geiger, and Otmar Hilliges. Pina: Learning a personalized implicit neural avatar from a single rgb-d video sequence. arXiv, 2022.
- [17] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas Barlow McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In 2022 International Conference on Robotics and Automation, ICRA 2022, Philadelphia, PA, USA, May 23-27, 2022, pages 2553–2560. IEEE, 2022. doi: 10.1109/ICRA46639.2022.9811809. URL https://doi.org/10.1109/ICRA46639.2022.9811809. 9, 10, 26, 31
- [18] Qiao Feng, Yebin Liu, Yu-Kun Lai, ingyu Yang, and Kun Li. Fof: Learning fourier occupancy field for monocular real-time human reconstruction. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 December 9, 2022, 2022.* 7, 8
- [19] Ruiqi Gao*, Aleksander Holynski*, Philipp Henzler, Arthur Brussee, Ricardo Martin-Brualla, Pratul P. Srinivasan, Jonathan T. Barron, and Ben Poole*. Cat3d: Create anything in 3d with multi-view diffusion models. *arXiv*, 2024. 9, 38
- [20] Vladimir Guzov, Aymen Mir, Torsten Sattler, and Gerard Pons-Moll. Human poseitioning system (HPS): 3d human pose estimation and self-localization in large scenes from bodymounted sensors. *CoRR*, abs/2103.17265, 2021. URL https://arxiv.org/abs/2103. 17265. 2
- [21] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR2023)*, June 2023. 7, 34, 36
- [22] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems* 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA, pages 6626–6637, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/8a1d694707eb0fefe65871369074926d-Abstract.html. 7
- [23] Hsuan-I Ho, Jie Song, and Otmar Hilliges. Sith: Single-view textured human reconstruction with image-conditioned diffusion. *CoRR*, abs/2311.15855, 2023. doi: 10.48550/ARXIV.2311. 15855. URL https://doi.org/10.48550/arXiv.2311.15855. 2, 7, 8, 34, 35
- [24] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. URL https://openreview.net/forum?id=qw8AKxfYbI. 3
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/ 4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html. 2, 3, 21
- [26] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. LRM: large reconstruction model for single image to 3d. *CoRR*, abs/2311.04400, 2023. doi: 10.48550/ARXIV.2311.04400. URL https://doi.org/10.48550/arXiv.2311.04400. 2, 3

- [27] Ho Hsuan-I, Xue Lixin, Song Jie, and Hilliges Otmar. Learning locally editable virtual humans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 7, 34, 36
- [28] Hanzhe Hu, Zhizhuo Zhou, Varun Jampani, and Shubham Tulsiani. MVD-Fusion: Single-view 3D via Depth-consistent Multi-view Generation, April 2024. 3
- [29] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. CoRR, abs/2403.17888, 2024. doi: 10.48550/ARXIV.2403.17888. URL https://doi.org/10.48550/arXiv.2403.17888. 5
- [30] Zehuan Huang, Hao Wen, Junting Dong, Yaohui Wang, Yangguang Li, Xinyuan Chen, Yan-Pei Cao, Ding Liang, Yu Qiao, Bo Dai, and Lu Sheng. EpiDiff: Enhancing Multi-View Synthesis via Localized Epipolar-Constrained Diffusion, April 2024. 3
- [31] Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Instantavatar: Learning avatars from monocular video in 60 seconds. *arXiv*, 2022. 2
- [32] Sai Sagar Jinka, Astitva Srivastava, Chandradeep Pokhariya, Avinash Sharma, and P. J. Narayanan. SHARP: shape-aware reconstruction of people in loose clothing. *Int. J. Comput. Vis.*, 131(4):918–937, 2023. doi: 10.1007/S11263-022-01736-Z. URL https://doi.org/10.1007/s11263-022-01736-z. 7, 8, 24, 25, 34, 35, 37
- [33] Berna Kabadayi, Wojciech Zielonka, Bharat Lal Bhatnagar, Gerard Pons-Moll, and Justus Thies. Gan-avatar: Controllable personalized gan-based human head avatar. In *International Conference on 3D Vision (3DV)*, March 2024. 2
- [34] Yash Kant, Ziyi Wu, Michael Vasilkovsky, Guocheng Qian, Jian Ren, Riza Alp Güler, Bernard Ghanem, Sergey Tulyakov, Igor Gilitschenski, and Aliaksandr Siarohin. SPAD: Spatially aware multiview diffusers. *CoRR*, abs/2402.05235, 2024. doi: 10.48550/ARXIV.2402.05235. URL https://doi.org/10.48550/arXiv.2402.05235. 3
- [35] Animesh Karnewar, Andrea Vedaldi, David Novotny, and Niloy Mitra. Holodiffusion: Training a 3D diffusion model using 2D images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023. 4, 21
- [36] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139:1–139:14, 2023. doi: 10.1145/3592433. URL https://doi.org/10.1145/3592433. 2, 3, 4, 21
- [37] Xin Kong, Shikun Liu, Xiaoyang Lyu, Marwan Taher, Xiaojuan Qi, and Andrew J. Davison. Eschernet: A generative model for scalable view synthesis. *CoRR*, abs/2402.03908, 2024. doi: 10.48550/ARXIV.2402.03908. URL https://doi.org/10.48550/arXiv.2402.03908. 3
- [38] Yifei Li, Hsiao-yu Chen, Egor Larionov, Nikolaos Sarafianos, Wojciech Matusik, and Tuur Stuyck. DiffAvatar: Simulation-ready garment optimization with differentiable simulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024. 2
- [39] Tingting Liao, Hongwei Yi, Yuliang Xiu, Jiaxiang Tang, Yangyi Huang, Justus Thies, and Michael J. Black. Tada! text to animatable digital avatars. *CoRR*, abs/2308.10899, 2023. doi: 10.48550/ARXIV.2308.10899. URL https://doi.org/10.48550/arXiv.2308.10899. 3
- [40] Minghua Liu, Ruoxi Shi, Linghao Chen, Zhuoyang Zhang, Chao Xu, Xinyue Wei, Hansheng Chen, Chong Zeng, Jiayuan Gu, and Hao Su. One-2-3-45++: Fast single image to 3d objects with consistent multi-view generation and 3d diffusion. *CoRR*, abs/2311.07885, 2023. doi: 10.48550/ARXIV.2311.07885. URL https://doi.org/10.48550/arXiv.2311.07885. 3
- [41] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 16, 2023, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/4683beb6bab325650db13afd05d1a14a-Abstract-Conference.html. 2, 3

- [42] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 9264–9275. IEEE, 2023. doi: 10.1109/ICCV51070.2023.00853. URL https://doi.org/10.1109/ICCV51070.2023.00853. 2, 3
- [43] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. *CoRR*, abs/2309.03453, 2023. doi: 10.48550/ARXIV.2309.03453. URL https://doi.org/10.48550/arXiv.2309.03453. 3, 5
- [44] Zhen Liu, Yao Feng, Yuliang Xiu, Weiyang Liu, Liam Paull, Michael J. Black, and Bernhard Schölkopf. Ghost on the shell: An expressive representation of general 3d shapes. *arXiv* preprint arXiv:2310.15168, 2023. 2
- [45] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, and Wenping Wang. Wonder3d: Single image to 3d using cross-domain diffusion. *CoRR*, abs/2310.15008, 2023. doi: 10. 48550/ARXIV.2310.15008. URL https://doi.org/10.48550/arXiv.2310.15008. 3
- [46] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: a skinned multi-person linear model. ACM Trans. Graph., 34(6):248:1–248:16, 2015. doi: 10.1145/2816795.2818013. URL https://doi.org/10.1145/2816795.2818013. 2, 3
- [47] Qianli Ma, Siyu Tang, Sergi Pujades, Gerard Pons-Moll, Anurag Ranjan, and Michael J. Black. Dressing 3d humans using a conditional mesh-vae-gan. *CoRR*, abs/1907.13615, 2019. URL http://arxiv.org/abs/1907.13615. 7, 8, 34, 35, 37
- [48] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Realfusion: 360° reconstruction of any object from a single image. In *Arxiv*, 2023. 3
- [49] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: representing scenes as neural radiance fields for view synthesis. *Commun. ACM*, 65(1):99–106, 2022. doi: 10.1145/3503250. URL https://doi.org/10.1145/3503250. 4, 21
- [50] 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. *CoRR*, abs/1904.05866, 2019. URL http://arxiv.org/abs/1904. 05866. 3
- [51] Ilya A. Petrov, Riccardo Marin, Julian Chibane, and Gerard Pons-Moll. Object pop-up: Can we infer 3d objects and their poses from human interactions alone? In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 4726–4736. IEEE, 2023. doi: 10.1109/CVPR52729.2023.00458. URL https://doi.org/10.1109/CVPR52729.2023.00458. 2
- [52] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J. Black. Clothcap: seamless 4d clothing capture and retargeting. ACM Trans. Graph., 36(4):73:1–73:15, 2017. doi: 10.1145/3072959.3073711. URL https://doi.org/10.1145/3072959.3073711. 7, 34
- [53] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. 3
- [54] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, and Bernard Ghanem. Magic 123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *CoRR*, abs/2306.17843, 2023. doi: 10.48550/ARXIV.2306.17843. URL https://doi.org/10.48550/arXiv.2306.17843. 2

- [55] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 10674–10685. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01042. URL https://doi.org/10.1109/CVPR52688.2022.01042. 2, 3, 9, 21
- [56] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Hao Li, and Angjoo Kanazawa. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 November 2, 2019, pages 2304–2314. IEEE, 2019. doi: 10.1109/ICCV. 2019.00239. URL https://doi.org/10.1109/ICCV.2019.00239. 2, 7, 8, 33
- [57] Shunsuke Saito, Tomas Simon, Jason M. Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020, pages 81-90. Computer Vision Foundation / IEEE, 2020. doi: 10.1109/CVPR42600.2020.00016. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Saito_PIFuHD_Multi-Level_Pixel-Aligned_Implicit_Function_for_High-Resolution_3D_Human_Digitization_CVPR_2020_paper.html. 2, 23, 33
- [58] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5B: an open large-scale dataset for training next generation image-text models. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 December 9, 2022, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/a1859debfb3b59d094f3504d5ebb6c25-Abstract-Datasets_and_Benchmarks.html. 3, 5, 9
- [59] Akash Sengupta, Thiemo Alldieck, Nikos Kolotouros, Enric Corona, Andrei Zanfir, and Cristian Sminchisescu. DiffHuman: Probabilistic Photorealistic 3D Reconstruction of Humans, March 2024. URL http://arxiv.org/abs/2404.00485. arXiv:2404.00485 [cs]. 2
- [60] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base model. *CoRR*, abs/2310.15110, 2023. doi: 10.48550/ARXIV.2310.15110. URL https://doi.org/10.48550/arXiv.2310.15110. 2, 3
- [61] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multiview diffusion for 3d generation. *CoRR*, abs/2308.16512, 2023. doi: 10.48550/ARXIV.2308. 16512. URL https://doi.org/10.48550/arXiv.2308.16512. 3, 5
- [62] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun, editors, 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, 2015. URL http://arxiv.org/abs/1409.1556. 5
- [63] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net, 2021. URL https://openreview.net/forum? id=St1giarCHLP. 21
- [64] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations, February 2021. URL http://arxiv.org/abs/2011.13456. arXiv:2011.13456 [cs, stat]. 2
- [65] Zhaoqi Su, Tao Yu, Yangang Wang, and Yebin Liu. Deepcloth: Neural garment representation for shape and style editing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1581–1593, 2023. doi: 10.1109/TPAMI.2022.3168569. 7, 34, 36

- [66] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Generalizable patch-based neural rendering. In Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision ECCV 2022 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXII*, volume 13692 of *Lecture Notes in Computer Science*, pages 156–174. Springer, 2022. doi: 10.1007/978-3-031-19824-3_10. URL https://doi.org/10.1007/978-3-031-19824-3_10. 3
- [67] Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. LGM: large multi-view gaussian model for high-resolution 3d content creation. *CoRR*, abs/2402.05054, 2024. doi: 10.48550/ARXIV.2402.05054. URL https://doi.org/10.48550/arXiv.2402.05054. 2, 3, 5, 7, 8, 21
- [68] Shitao Tang, Jiacheng Chen, Dilin Wang, Chengzhou Tang, Fuyang Zhang, Yuchen Fan, Vikas Chandra, Yasutaka Furukawa, and Rakesh Ranjan. Mvdiffusion++: A dense high-resolution multi-view diffusion model for single or sparse-view 3d object reconstruction. *CoRR*, abs/2402.12712, 2024. doi: 10.48550/ARXIV.2402.12712. URL https://doi.org/10.48550/arXiv.2402.12712. 3, 9, 38
- [69] Maxim Tatarchenko, Stephan R. Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 3405–3414. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00352. URL http://openaccess.thecvf.com/content_CVPR_2019/html/Tatarchenko_What_Do_Single-View_3D_Reconstruction_Networks_Learn_CVPR_2019_paper.html. 7
- [70] Ayush Tewari, Tianwei Yin, George Cazenavette, Semon Rezchikov, Josh Tenenbaum, Frédo Durand, Bill Freeman, and Vincent Sitzmann. Diffusion with forward models: Solving stochastic inverse problems without direct supervision. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 16, 2023, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/28e4ee96c94e31b2d040b4521d2b299e-Abstract-Conference.html. 4, 21
- [71] Anh Thai, Stefan Stojanov, Vijay Upadhya, and James M. Rehg. 3d reconstruction of novel object shapes from single images, 2020. 3
- [72] Garvita Tiwari, Bharat Lal Bhatnagar, Tony Tung, and Gerard Pons-Moll. SIZER: A dataset and model for parsing 3d clothing and learning size sensitive 3d clothing. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, Computer Vision ECCV 2020 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part III, volume 12348 of Lecture Notes in Computer Science, pages 1–18. Springer, 2020. doi: 10.1007/978-3-030-58580-8_1. URL https://doi.org/10.1007/978-3-030-58580-8_1. 7, 8, 10, 24, 25, 34, 35, 37
- [73] Garvita Tiwari, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Neural-gif: Neural generalized implicit functions for animating people in clothing. In 2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021, pages 11688–11698. IEEE, 2021. doi: 10.1109/ICCV48922.2021.01150. URL https://doi.org/10.1109/ICCV48922.2021.01150. 2
- [74] Dmitry Tochilkin, David Pankratz, ZeXiang Liu, Zixuan Huang, Adam Letts, Yangguang Li, Ding Liang, Christian Laforte, Varun Jampani, and Yan-Pei Cao. Triposr: Fast 3d object reconstruction from a single image. *CoRR*, abs/2403.02151, 2024. doi: 10.48550/ARXIV. 2403.02151. URL https://doi.org/10.48550/arXiv.2403.02151. 2, 3, 5, 7, 8, 33
- [75] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA,

- pages 5998-6008, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/ 3f5ee243547dee91fbd053c1c4a845aa-Abstract.html. 3
- [76] Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitry Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. SV3D: Novel Multi-view Synthesis and 3D Generation from a Single Image using Latent Video Diffusion, March 2024. 3
- [77] Peng Wang and Yichun Shi. Imagedream: Image-prompt multi-view diffusion for 3d generation. *CoRR*, abs/2312.02201, 2023. doi: 10.48550/ARXIV.2312.02201. URL https://doi.org/10.48550/arXiv.2312.02201. 2, 3, 5, 6, 7, 21, 38
- [78] Z. Wang, E.P. Simoncelli, and A.C. Bovik. Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers*, 2003, volume 2, pages 1398–1402 Vol.2, 2003. doi: 10.1109/ACSSC.2003.1292216. 7
- [79] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *arXiv preprint arXiv:2305.16213*, 2023. 3
- [80] Chung-Yi Weng, Brian Curless, Pratul P. Srinivasan, Jonathan T. Barron, and Ira Kemelmacher-Shlizerman. HumanNeRF: Free-viewpoint rendering of moving people from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16210–16220, June 2022. 2
- [81] Jiajun Wu, Chengkai Zhang, Xiuming Zhang, Zhoutong Zhang, William T Freeman, and Joshua B Tenenbaum. Learning 3D Shape Priors for Shape Completion and Reconstruction. In *European Conference on Computer Vision (ECCV)*, 2018. 3
- [82] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Jiawei Ren, Liang Pan, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, Dahua Lin, and Ziwei Liu. Omniobject3d: Large-vocabulary 3d object dataset for realistic perception, reconstruction and generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 803–814. IEEE, 2023. doi: 10.1109/CVPR52729.2023.00084. URL https://doi.org/10.1109/CVPR52729.2023.00084. 2, 3, 34
- [83] Yongqin Xian, Julian Chibane, Bharat Lal Bhatnagar, Bernt Schiele, Zeynep Akata, and Gerard Pons-Moll. Any-shot gin: Generalizing implicit networks for reconstructing novel classes. In 2022 International Conference on 3D Vision (3DV). IEEE, 2022. 3
- [84] Xianghui Xie, Bharat Lal Bhatnagar, and Gerard Pons-Moll. CHORE: contact, human and object reconstruction from a single RGB image. *CoRR*, abs/2204.02445, 2022. doi: 10.48550/ARXIV.2204.02445. URL https://doi.org/10.48550/arXiv.2204.02445. 2
- [85] Xianghui Xie, Bharat Lal Bhatnagar, and Gerard Pons-Moll. Visibility aware human-object interaction tracking from single RGB camera. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 4757–4768. IEEE, 2023. doi: 10.1109/CVPR52729.2023.00461. URL https://doi.org/10.1109/CVPR52729.2023.00461.
- [86] Xianghui Xie, Bharat Lal Bhatnagar, Jan Eric Lenssen, and Gerard Pons-Moll. Template free reconstruction of human-object interaction with procedural interaction generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024. 2, 9, 26, 32
- [87] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J. Black. ICON: implicit clothed humans obtained from normals. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022, pages 13286–13296. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01294. URL https://doi.org/10.1109/CVPR52688.2022.01294. 2, 7, 8, 34, 35, 37
- [88] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J. Black. ECON: explicit clothed humans optimized via normal integration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023, pages 512–523. IEEE, 2023. doi: 10.1109/CVPR52729.2023.00057. URL https://doi.org/10.1109/CVPR52729.2023.00057. 2, 7, 8, 34, 35

- [89] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv* preprint arXiv:2404.07191, 2024. 3, 5, 7, 8
- [90] Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli, Gordon Wetzstein, Zexiang Xu, and Kai Zhang. DMV3D: denoising multi-view diffusion using 3d large reconstruction model. *CoRR*, abs/2311.09217, 2023. doi: 10.48550/ARXIV. 2311.09217. URL https://doi.org/10.48550/arXiv.2311.09217. 3, 21
- [91] Yinghao Xu, Zifan Shi, Wang Yifan, Sida Peng, Ceyuan Yang, Yujun Shen, and Wetzstein Gordon. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. *arxiv*: 2403.14621, 2024. 3
- [92] Yuxuan Xue, Haolong Li, Stefan Leutenegger, and Joerg Stueckler. Event-based non-rigid reconstruction from contours. In 33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022, page 78. BMVA Press, 2022. URL https://bmvc2022.mpi-inf.mpg.de/78/. 2
- [93] Yuxuan Xue, Bharat Lal Bhatnagar, Riccardo Marin, Nikolaos Sarafianos, Yuanlu Xu, Gerard Pons-Moll, and Tony Tung. NSF: neural surface fields for human modeling from monocular depth. In *IEEE/CVF International Conference on Computer Vision*, *ICCV* 2023, *Paris, France, October 1-6*, 2023, pages 15004–15014. IEEE, 2023. doi: 10.1109/ICCV51070.2023.01382. URL https://doi.org/10.1109/ICCV51070.2023.01382. 2
- [94] Yuxuan Xue, Haolong Li, Stefan Leutenegger, and Jörg Stückler. Event-based non-rigid reconstruction of low-rank parametrized deformations from contours. In *International Journal of Computer Vision (IJCV)*. Springer Science and Business Media LLC, February 2024. doi: 10. 1007/s11263-024-02011-z. URL http://dx.doi.org/10.1007/s11263-024-02011-z. 2.
- [95] Xueting Yang, Yihao Luo, Yuliang Xiu, Wei Wang, Hao Xu, and Zhaoxin Fan. D-IF: Uncertainty-aware Human Digitization via Implicit Distribution Field. In 2023 IEEE/CVF International Conference on Computer Vision (ICCV), pages 9088–9098, Paris, France, October 2023. IEEE. ISBN 9798350307184. doi: 10.1109/ICCV51070.2023.00837. URL https://ieeexplore.ieee.org/document/10377954/. 2
- [96] Kim Youwang, Tae-Hyun Oh, and Gerard Pons-Moll. Paint-it: Text-to-texture synthesis via deep convolutional texture map optimization and physically-based rendering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [97] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021, pages 4578-4587. Computer Vision Foundation / IEEE, 2021. doi: 10.1109/CVPR46437.2021.00455. URL https://openaccess.thecvf.com/content/CVPR2021/html/Yu_pixelNeRF_Neural_Radiance_Fields_From_One_or_Few_Images_CVPR_2021_paper.html. 4, 21
- [98] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR2021)*, June 2021. 7, 34, 36
- [99] Xianggang Yu, Mutian Xu, Yidan Zhang, Haolin Liu, Chongjie Ye, Yushuang Wu, Zizheng Yan, Chenming Zhu, Zhangyang Xiong, Tianyou Liang, Guanying Chen, Shuguang Cui, and Xiaoguang Han. Mvimgnet: A large-scale dataset of multi-view images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 9150–9161. IEEE, 2023. doi: 10.1109/CVPR52729.2023.00883. URL https://doi.org/10.1109/CVPR52729.2023.00883. 2, 3, 34
- [100] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient high-quality compact surface reconstruction in unbounded scenes. *arXiv:2404.10772*, 2024. 5, 21, 23

- [101] Polina Zablotskaia, Aliaksandr Siarohin, Bo Zhao, and Leonid Sigal. Dwnet: Dense warp-based network for pose-guided human video generation. In 30th British Machine Vision Conference 2019, BMVC 2019, Cardiff, UK, September 9-12, 2019, page 51. BMVA Press, 2019. URL https://bmvc2019.org/wp-content/uploads/papers/1039-paper.pdf. 10, 26, 30
- [102] Andy Zeng, Shuran Song, Matthias Nießner, Matthew Fisher, Jianxiong Xiao, and Thomas Funkhouser. 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. In *CVPR*, 2017. 21, 23
- [103] Chao Zhang, Sergi Pujades, Michael J. Black, and Gerard Pons-Moll. Detailed, accurate, human shape estimation from clothed 3d scan sequences. In 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017, pages 5484–5493. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.582. URL https://doi.org/10.1109/CVPR.2017.582. 7, 34
- [104] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, pages 586-595. Computer Vision Foundation / IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00068. URL http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_The_Unreasonable_Effectiveness_CVPR_2018_paper.html. 7
- [105] Xiaohan Zhang, Bharat Lal Bhatnagar, Sebastian Starke, Vladimir Guzov, and Gerard Pons-Moll. COUCH: towards controllable human-chair interactions. In Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, Computer Vision ECCV 2022 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part V, volume 13665 of Lecture Notes in Computer Science, pages 518–535. Springer, 2022. doi: 10.1007/978-3-031-20065-6_30. URL https://doi.org/10.1007/978-3-031-20065-6_30. 2
- [106] Xiaohan Zhang, Bharat Lal Bhatnagar, Sebastian Starke, Ilya Petrov, Vladimir Guzov, Helisa Dhamo, Eduardo Pérez-Pellitero, and Gerard Pons-Moll. FORCE: dataset and method for intuitive physics guided human-object interaction. *CoRR*, abs/2403.11237, 2024. doi: 10. 48550/ARXIV.2403.11237. URL https://doi.org/10.48550/arXiv.2403.11237. 2
- [107] Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Joshua B Tenenbaum, William T Freeman, and Jiajun Wu. Learning to Reconstruct Shapes From Unseen Classes. In Advances in Neural Information Processing Systems (NeurIPS), 2018. 3
- [108] Zechuan Zhang, Zongxin Yang, and Yi Yang. SIFU: side-view conditioned implicit function for real-world usable clothed human reconstruction. *CoRR*, abs/2312.06704, 2023. doi: 10.48550/ARXIV.2312.06704. URL https://doi.org/10.48550/arXiv.2312.06704. 2, 7, 8, 33, 34, 35
- [109] Zhizhuo Zhou and Shubham Tulsiani. Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In CVPR, 2023. 3
- [110] Zi-Xin Zou, Zhipeng Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Yan-Pei Cao, and Song-Hai Zhang. Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. *CoRR*, abs/2312.09147, 2023. doi: 10.48550/ARXIV.2312.09147. URL https://doi.org/10.48550/arXiv.2312.09147. 2, 3

Appendix

Table of Contents

A	Implementation Details	21
	A.1 Training Details	 21
	A.2 Joint Framework	 21
	A.3 Generative 3D-GS Reconstruction Model	 21
	A.4 Textured Mesh Extraction	 23
В	Comparison	24
	B.1 Qualitative Comparison	 24
C	More Qualitative Results	26
	C.1 In-the-wild Data	 26
	C.2 UBC Fashion Dataset	 30
	C.3 Google Scan Objects (GSO)	 3
	C.4 Human Object Interaction	 32
	C.5 Generative Power in Reconstruction	 33
D	Dataset Overview	34
	D.1 Training Dataset	 34
	D.2 Evaluation Dataset	 34
E	Failure Cases	38
F	Broader Impacts	38

Implementation Details

Training Details

As described in Sec. 5.1, we use an effective batch size of 256. Each batch involved sampling 4 orthogonal images with zero elevation angle as target views $\mathbf{x}_0^{\text{tgt}}$, and 12 additional images as novel views $\mathbf{x}_0^{\text{novel}}$ to supervise the 3D generative model Eq. (8). The hyperparameters for training Eq. (8) were set as follows: $\lambda_1 = 1.0$, $\lambda_2 = 1.0$, and $\lambda_3 = 100.0$.

During training, we employed the standard DDPM scheduler [25] to construct noisy target images $\mathbf{x}_{t}^{\text{tg}}$. The maximum diffusion step T is set to 1000. At inference time, we use DDIM scheduler [63] to perform faster reverse sampling. The reverse steps is set to 50 in following experiments. The text prompt y used in our multi-view diffusion model(Eq. (3)) is set to "Photorealistic 3D human" for both training and inference across all subjects.

A.2 Joint Framework

Implementation. Our 2D multi-view diffusion model ϵ_{θ} is a latent diffusion model [77]. Thus, we use the frozen VAE in [55] to obtain input $\mathbf{x}_t^{\text{tgt}}$ in image space for the 3D generative model g_{ϕ} and encode refined $\tilde{\mathbf{x}}_t^{\text{tgt}}$ back to latent space for ϵ_{θ} . We extract triangle mesh from predicted Gaussian splats using Gaussian Opacity field [100] and TSDF [102]. Please refer to Appendix A.4 for more details.

Generative 3D-GS Reconstruction Model A.3

In this section, we provide details about our 3D generative model g_{ϕ} in Eq. (5) and Eq. (7) as well as the renderer(o). Following [7, 35, 70], we learn the 3D generative model by adding and removing noise on the rendered 2D images from a 3D representation. A pseudo algorithm of the training and sampling process of our 3D generative model can be found in Algorithm 3.

Since we integrate both function into the reverse sampling process iteratively (eq. 2), we expect them be efficient and fast to execute. Tewari et al. [70] base their model on pixelNeRF [97], which is a generalizable NeRF [49] conditioned on a context view image. We adopt 3D Gaussian Splats [36] as our 3D state representation \mathcal{G} due to its efficiency and simplicity. Our renderer(\circ) is the differentiable rasteraizer accelerated and implemented in CUDA, which achieves around 2700 times faster rendering than volume-rendering-based renderer(o) in [49, 70, 97].

For sampling the 3D State S from $\mathbf{x}_t^{\text{tgt}}$, $\tilde{\mathbf{x}}_0^{\text{tgt}}$, \mathbf{x}^{c} , and t (eq. 7), we adopt the time-conditioned UNet-Transformer architecture [55] due to the efficiency of convolutional layers and the scalability of transformers. For enabling the awareness of camera poses in the encoding process, we concatenate the Plücker Camera Ray Embedding $\{\mathbf{o}_i \times \mathbf{d}_i, \mathbf{d}_i\}$ [67, 90] with the image $\mathbf{x}_t^{\text{tgt}}$ and $\tilde{\mathbf{x}}_0^{\text{tgt}}$. To enhance the control ability of context view in the 3D generation process, we additionally concatenate the clear context view \mathbf{x}^c with target images \mathbf{x}_t^{tgt} following [77]. This operation enables 3D dense self-attention process between the input multi-view target images and the clear context view image, provides pixel-level local conditional signal. Since the camera pose of context view \mathbf{x}^c is unknown, we use the 0-vector as its embedding.

Algorithm 3 Learn 3D distribution

 $\mathbf{x}_0^{\text{novel}}$, π^{novel} , a context image \mathbf{x}^c

Output: Optimized 3D State diffusion network g_{ϕ}

1. Teperat.
2.
$$\{\mathbf{x}_0^{\text{tgt}}, \mathbf{x}_0^{\text{novel}}, \mathbf{x}^c, y\} \sim q(\{\mathbf{x}_0^{\text{tgt}}, \mathbf{x}_0^{\text{novel}}, \mathbf{x}^c, y\})$$
3. $t \sim \text{Uniform}(\{1, \dots, T\}); \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

4:
$$\mathbf{x}_{t}^{\text{tgt}} = \sqrt{\bar{\alpha}_{t}}\mathbf{x}_{0}^{\text{tgt}} + \sqrt{1 - \bar{\alpha}_{t}}\epsilon$$

5: $\hat{\mathcal{G}} = g_{\phi}\left(\mathbf{x}^{c}, \mathbf{x}_{t}^{\text{tgt}}, t\right)$

5.
$$\hat{G} = a_{\perp} (\mathbf{x}^{c} \ \mathbf{x}^{tgt} \ t)$$

6:
$$\{\hat{\mathbf{x}}_0^{\text{tgt}}, \hat{\mathbf{x}}_0^{\text{novel}}\} = \text{renderer}\left(\hat{\mathcal{G}}, \{\pi^{\text{tgt}}, \pi^{\text{novel}}\}\right)$$

- 7: Compute loss \mathcal{L}_{qs} (Eq. (6))
- 8: Gradient step to update q_{ϕ}
- 9: until converged

Algorithm 4 Sample from 3D distribution

Input: Dataset of posed multi-view images $\mathbf{x}_0^{\text{tgt}}$, π^{tgt} , **Input:** A context image \mathbf{x}^c ; Converged 3D diffusion

Output: A 3D Gaussian Avatar \mathcal{G} of the 2D image \mathbf{x}^c

1:
$$\mathbf{x}_{T}^{\text{tgt}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

2: **for**
$$t = T, ..., 1$$
 do

3:
$$\hat{\mathcal{G}} = q_{\phi} \left(\mathbf{x}^{c}, \mathbf{x}_{t}^{\text{tgt}}, t \right)$$

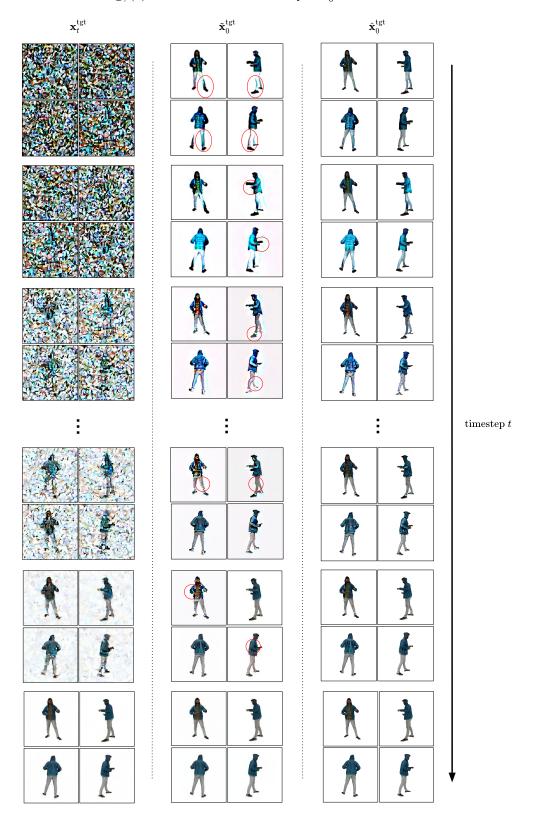
4:
$$\hat{\mathbf{x}}_0^{\text{tgt}} = \text{renderer} \left(\hat{\mathcal{G}}, \pi^{\text{tgt}}\right)$$

Output: A 3D Gaussian Avatar
$$\mathcal{G}$$
 of the 2D image \mathbf{x}^c
1: $\mathbf{x}_T^{\text{tgt}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
2: $\mathbf{for} \ t = T, \dots, 1 \ \mathbf{do}$
3: $\hat{\mathcal{G}} = g_{\phi} \left(\mathbf{x}^c, \mathbf{x}_t^{\text{tgt}}, t\right)$
4: $\hat{\mathbf{x}}_0^{\text{tgt}} = \mathbf{renderer} \left(\hat{\mathcal{G}}, \pi^{\text{tgt}}\right)$
5: $\mu_{t-1}(\mathbf{x}_t^{\text{tgt}}, \hat{\mathbf{x}}_0^{\text{tgt}}) = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t} \mathbf{x}_t^{\text{tgt}} + \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1-\bar{\alpha}_t} \hat{\mathbf{x}}_0^{\text{tgt}}$
6: $\mathbf{x}_{t-1}^{\text{tgt}} \sim \mathcal{N}\left(\mathbf{x}_{t-1}^{\text{tgt}}; \tilde{\mu}_t \left(\mathbf{x}_t^{\text{tgt}}, \hat{\mathbf{x}}_0^{\text{tgt}}\right), \tilde{\beta}_{t-1}\mathbf{I}\right)$
7: $\mathbf{end} \ \mathbf{for}$

6:
$$\mathbf{x}_{t-1}^{\text{tgt}} \sim \mathcal{N}\left(\mathbf{x}_{t-1}^{\text{tgt}}; \tilde{\boldsymbol{\mu}}_{t}\left(\mathbf{x}_{t}^{\text{tgt}}, \hat{\mathbf{x}}_{0}^{\text{tgt}}\right), \tilde{\beta}_{t-1}\mathbf{I}\right)$$

- 8: **return** $\mathcal{G} = g_{\phi} \left(\mathbf{x}_{0}^{\text{tgt}}, \mathbf{x}^{\text{c}}, t = 0 \right)$

Figure 7: Visualization of intermediate sampling steps from a Gaussian Noise (t=1000) to the last denoising step (t=0). From top to bottom: current state $\mathbf{x}_t^{\mathrm{tgt}}$, estimated clear view by 2D diffusion models $\tilde{\mathbf{x}}_0^{\mathrm{tgt}}$, and corrected clear view by generated 3D Gaussian Splatting $\hat{\mathbf{x}}_0^{\mathrm{tgt}}$. Our 2D diffusion model $\epsilon_{\phi}(\circ)$ already provides strong multi-view prior at an early stage with large t. Our 3D reconstruction model $\mathbf{g}_{\phi}(\circ)$ can correct the inconsistency in $\tilde{\mathbf{x}}_0^{\mathrm{tgt}}$ illustrated in red circle.



A.4 Textured Mesh Extraction

Gaussian Opacity Fields [100] enables extraction of triangle meshes from an existing 3D Gaussian Splatting. However, because the location of 3D-Gaussian Splats is not necessary to be on the real surface, we observe that the extracted meshes as well as the rendered depth maps are noisy. Since our method generate realistic RGB images, we use PiFU-HD [57] to estimate the normals and use Bilateral Normal Integration (BiNI) [11] to refine the noisy rendered depth with the estimated normal. As we only want the estimated normal to denoise the rendered depth map instead of modifying geometry, we set up the hyperparameter in BiNI with $\lambda=1\times10^4$. Such a large number ensures that the normal map is not used to modify the geometry but just regularize the depth map.

Assuming we have a generated 3D-Gaussian Splats \mathcal{G} from $g_{\phi}(\circ)$ and n camera views $\pi^1, \pi^2, ... \pi^n$, we obtain n paris of posed RGB-D images by Gaussian Splatting, Normal Estimation, and Bilateral Normal Integration. Finally, we perform volumetric TSDF fusion [102] to obtain high quality textured mesh from n pairs RGB-D images. Given generated 3D-GS \mathcal{G} , we set up 36 views to obtain the refined RGB-D image pairs. The rendering view of each camera i can be calculated as:

elevation_i =
$$-\frac{1}{4}\pi + \frac{1}{4}\pi * \frac{i}{36}$$
, (10)

$$\operatorname{azimuth}_{i} = 0 + 3\pi * \frac{i}{36}. \tag{11}$$

99623

B Comparison



Figure 8: Qualitative comparison on Sizer [72] and IIIT [32].



Figure 9: Qualitative comparison on Sizer [72] and IIIT [32].

C More Qualitative Results

In this section, we show more qualitative results on in-the-wild data, UBC fashion dataset [101], GSO dataset [17], and human-object interaction data [86].

C.1 In-the-wild Data

Figure 10: Qualitative results on unseen data during training. Input image is in left column. Our method successfully reconstructs different degree of loose clothing.



Figure 11: Qualitative results on more unseen data during training. Input image is in left column. Our method successfully reconstructs different types of clothing, including casual, sport, suits, custom, etc., in both appearance and geometry.



Figure 12: Qualitative results on more unseen data during training. Input image is in left column. Our method successfully reconstructs clothing and interacting objects (racket and bag here) in both appearance and geometry.



Figure 13: Qualitative results on more unseen data during training. Input image is in left column. Our method successfully reconstructs rarely seen suits and objects, in both appearance and geometry.

C.2 UBC Fashion Dataset

In this section, we show qualitative result of our model on UBC fashion [101] dataset. The input images are the first frame extracted from each video in the dataset.



Figure 14: Qualitative results on UBC fashion [101] dataset. Results demonstrate that our model generalizes well to real world images in both geometry and appearance.

C.3 Google Scan Objects (GSO)



Figure 15: Ablation study: benefit of 2D multi-view prior $\tilde{\mathbf{x}}_0^{\text{tgt}}$ in 3D generation. The 2D prior from 2D diffusion model is essntial for generalization on general objects dataset GSO [17].

C.4 Human Object Interaction



Figure 16: Qualitative results of Human-Object Interaction reconstruction on online stock images. Results show that our model is able to generalize to casual human-object-interactions.



Figure 17: Qualitative results of Human-Object Interaction reconstruction on ProciGen [86] dataset. Results demonstrate that our model can reconstruct some simple Human-Object-Interaction images with large objects.

C.5 Generative Power in Reconstruction

Our model learns a conditional distribution of the 3D representation given 2D context image. Thus, by sampling from the distribution with different seed, we obtain diverse yet plausible 3D representation. As illustrated in Fig. 18, the appearance of the occluded region (back side of subject) is different with different sampling in hair styple, texture, and cloth wrinkles.

The generative power of our approach is the key to generate clear self-occluded regions, which is impossible by non-generative reconstruction methods [56, 57, 74, 108]. As shown in Fig. 3 and Fig. 8, non-generative approaches tend to generate blurry self-occluded results because they cannot sample from distribution but only regress to a mean value of the training datasets.

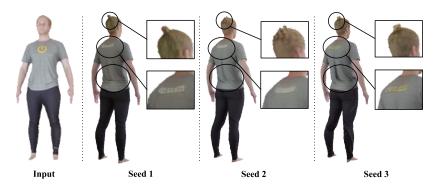


Figure 18: Our model learns 3D distribution. By different sampling from the learned distribution, we obtain diverse yet plausible 3D representations. The generative power is a key to generate clear self-occluded region, which is impossible in non-generative reconstruction approaches [56, 57, 74, 108].

99633

D Dataset Overview

To ensure robust performance and generalization, we train our model on a combined dataset comprising 3520 scans from publicly available datasets [21, 27, 65, 98] and 2320 scans from commercial 3D human datasets [1, 3, 4, 2]. These datasets encompass a diverse range of body shapes, genders, ages, clothing, accessories, and interacting objects. For further details and examples of our training datasets, please refer to section D.1. All 3D scans are rendered into RGB-A images using BlenderProc [14] along a spiral path as described in Eq. (13).

For evaluation, we note that the commonly-used CAPE dataset [47, 52, 103] in previous works [23, 87, 88, 108] often contains artifacts in scans, such as holes, and not all 3D scans are fully publicly available. To effectively and fairly evaluate performance, we propose using Sizer [72] and IIIT-Human [32] datasets, from which we randomly sample 150 scans each for evaluation. While Sizer [72] provides scans with normal human appearance similar to our training datasets, IIIT-Human [32] can be considered as out-of-distribution (o.o.d.) evaluation dataset due to its inclusion of unseen clothing types, such as traditional Indian suits. For more additional examples and analysis, please see Fig. 3 and Appendix D.2.

D.1 Training Dataset

Datasets To prevent the overfitting of our large neural network, namely the 2D multi-view diffusion models $\epsilon_{\theta}(\circ)$ and 3D generative models $\mathbf{g}_{\phi}(\circ)$, we train on as much data as we can. Unlike general objects community which has massive dataset such as Objaverse (800K) and Objaverse-XL (10M) [13]) OmniObject3D (6K) [82], MVImageNet (87K) [99], we don't have a single 3D human dataset available at such a scale. To collect data as much as possible, we collect both following public datasets and commercial human scans.

We collect several publicly available datasets inluding 2k2k (2K) [21], CustomHuman (640) [27], Thuman2.0 (520) [98], and Thuman3.0 (360) [65]). Among them, CustomHuman, Thuman2.0, and Thuman3.0 have more repeating subjects with different poses, which have less diverse subject appearance compared to 2k2k. It is worth mentioning that 2k2k [21] is a high quality dataset which contains human with diverse clothing (such as skirt) and accessories (such as cap, hat, scarf).

We also utilize in total 2320 high quality commercial scans from AXYZ [1], Treedy [3], Twindom [4], and RenderPeople [2]. All of these scans are with casual clothing and without interaction with objects.

Rendering For each scan, we render 100 views following a spiral trajectory with each view i:

elevation_i =
$$-\frac{1}{4}\pi + \frac{7}{8}\pi * \frac{i}{100}$$
, (12)

$$\operatorname{azimuth}_{i} = 0 + 5\pi * \frac{i}{100}.$$
 (13)

Additionally, we render 32 views uniformly around z-axis with each view j:

$$elevation_{j} = 0, (14)$$

$$\operatorname{azimuth}_{j} = 0 + \pi * \frac{j}{32}. \tag{15}$$

To protect the privacy of subjects in the training dataset, we only use the frontal view (with azimuth_j $\in [-\frac{\pi}{2}, \frac{\pi}{2}]$) as the input context view during training. Thus, we expect the model will not learn faces of subjects when it takes the back view as input.

D.2 Evaluation Dataset

For quantitative evaluation, we use Sizer [72] and IIIT 3D human dataset [32]. In this section, we start with introducing the two used evaluation datasets, and explain why we omit the commonly used CAPE dataset [47] in our experiments. Finally, we provide a summary of the evaluation datasets.

CAPE Cape [47] is a popular evaluation dataset which is widely used by previous methods [23, 87, 88, 108]. We observe that CAPE has limitation in appearance variety, geometrical artifacts as well as the publicly unavailability. CAPE only contains simple clothing such as T-shirts and jeans, but no garments of loose clothing. As illustrated in Fig. 22, CAPE contains several artifacts, such as holes on the head, missing hands, and wrong mesh geometry. We manually removed the noise and artifacts in scan to serve as our evaluation dataset. Moreover, CAPE doesn't have most original scans publicly available, but only the SMPL+D fitting. Due to this reason, we cannot render the CAPE scan to RGB images at desired camera view to evaluate the appearance performance such as PSNR, SSIM, and LPIPS. For above mentioned reasons, we also propose to use Sizer [72] and IIIT [32] which are publicly available.

Sizer Sizer [72] is a high quality 3D human scan dataset which contains 100 different subjects wearing casual clothing items in various sizes. We randomly sample 150 scans from Sizer [72] as one of our evaluation dataset.

IIIT 3D Humans [32] is a high quality dataset from IIIT Hyderabad in India. Different from the casual clothing setup in Sizer [72], IIIT dataset mainly focuses on subjects wearing traditional India custom suits, including ethnicity, diverse color pattern and extremely loose clothing (Fig. 21). It brings the huge variety of the subject appearance which can be considered as o.o.d. evaluation set to our model and baselines. We randomly sample 150 scans from IIIT dataset [32] for evaluation.

Summary We observe that the high quality datasets Sizer [72] and IIIT 3D Human [32] are unexplored for the community of 3D avatars reconstruction. In fact, Sizer [72] contains casual clothing which is suitable to evaluate performance, and IIIT [32] contains challenging texture and loose clothing which is suitable to evaluate robustness. All high quality scans in [32, 72] have no severe artifacts and are fully publicly available, which are the benefits unprovided in CAPE dataset [47]. By evaluating on these datasets [32, 72] and release our randomly sampled subjects which are used in our experiments, we hope the 3D avatars community can discover and benefit from

99635



Figure 19: Example scans in training datasets [1-4, 21, 27, 65, 98].



Figure 20: Example scans in Sizer [72] dataset. Sizer contains human in casual clothing.



Figure 21: Example scans in IIIT [32] dataset. IIIT contains subjects with diverse color pattern and loose garments, which rarely appear in training datasets.



Figure 22: Example artifacts in CAPE [47] dataset. Images shown here are rendered by ICON [87] due to the inaccessibility of original scans.

E Failure Cases

Limited by low resolution (256×256) of our multi-view diffusion model [77], our model can often fail in reconstructing fine details such as text on the cloth as illustrated in Fig. 23. One potential solution is to switch to a recent powerful high-resolution multi-view diffusion models [19, 68].

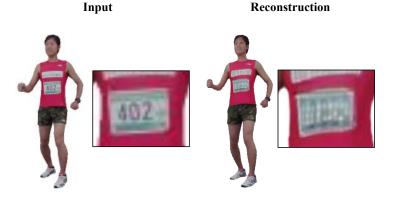


Figure 23: Failure Case: our model cannot reconstruct the numbers on the cloth.

In addition, we observe that our model can fail when reconstructing human extremely challenging poses. As shown in Fig. 24, our model cannot infer head geometry and appearance accurately due to the challenging pose in input image.

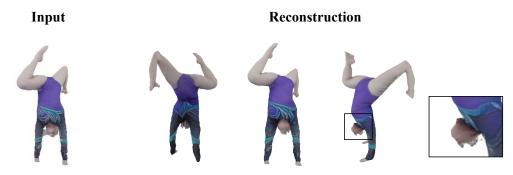


Figure 24: Failure Case: our model fails in infer appearance of human with challenging pose.

F Broader Impacts

Our work shows generality across different ethnicities and humans, providing a useful tool for a fair representation of different cultures. Having a robust method to synthesize realistic 3D geometry from a single RGB image may be used in surveillance and inappropriate content generation.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Contributions and scope are included in the abstract and the introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Please refer to Appendix E

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

99639

Answer: [NA]

Justification: The paper does not contain theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Please refer to our Appendix for Implementation details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
- (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: All code and scripts are publicly available here.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how
 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We include details about datasets and implementation in Sec. 5.1, Appendix A.1, Appendix D.1, Appendix D.2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We follow standard evaluation protocols reporting average errors on an extensive test set.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The paper includes the computational resources used in our experiments in Section 5.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: We read and adhere to the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a
 deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The paper include the Broader Impacts statement in Appendix F.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not see such high risk posed by our paper.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite and explicitly refer to all the legitimate sources of code, data, and models.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

 If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not include new assets, but we will include appropriate documentation upon code release after acceptance.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Although the method involves 3D human models, we rely on datasets collected before this work; we refer to them for their specifics.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]
Justification: NA

Justification: Although the method involves 3D human models, we rely on datasets collected before this work; we refer to them for their IRB approvals.

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.

 For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.