

Single-Image 3D Human Digitization with Shape-Guided Diffusion

Badour AlBahar
Kuwait University
Kuwait City, Kuwait
badour.albahar@ku.edu.kw

Changil Kim
Meta
Seattle, Washington, USA
changil@meta.com

Shunsuke Saito
Meta
Pittsburgh, Pennsylvania, USA
shunsukesaito@meta.com

Johannes Kopf
Meta
Seattle, Washington, USA
jkopf@meta.com

Hung-Yu Tseng
Meta
Seattle, Washington, USA
hungyutseng@meta.com

Jia-Bin Huang
University of Maryland
College Park, Maryland, USA
jbhuang@umd.edu



Figure 1: 3D Human Digitization from a Single Image. For a single image as input, our approach synthesizes the 3D consistent texture of a person without relying on any 3D scans for supervised training. Our key idea is to leverage high-capacity 2D diffusion models pretrained for general image synthesis tasks as a human appearance prior. Images from Adobe Stock.

ABSTRACT

We present an approach to generate a 360-degree view of a person with a consistent, high-resolution appearance from a *single* input

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SA Conference Papers '23, December 12–15, 2023, Sydney, NSW, Australia

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0315-7/23/12...\$15.00

<https://doi.org/10.1145/3610548.3618153>

image. NeRF and its variants typically require videos or images from different viewpoints. Most existing approaches taking monocular input either rely on ground-truth 3D scans for supervision or lack 3D consistency. While recent 3D generative models show promise of 3D consistent human digitization, these approaches do not generalize well to diverse clothing appearances, and the results lack photorealism. Unlike existing work, we utilize high-capacity 2D diffusion models pretrained for general image synthesis tasks as an appearance prior of clothed humans. To achieve better 3D consistency while retaining the input identity, we progressively synthesize multiple views of the human in the input image by inpainting missing regions with shape-guided diffusion conditioned

on silhouette and surface normal. We then fuse these synthesized multi-view images via inverse rendering to obtain a fully textured high-resolution 3D mesh of the given person. Experiments show that our approach outperforms prior methods and achieves photorealistic 360-degree synthesis of a wide range of clothed humans with complex textures from a single image.

CCS CONCEPTS

• **Computing methodologies** → *Texturing*.

KEYWORDS

Digital humans, single-image 3D reconstruction, diffusion models

ACM Reference Format:

Badour AlBahar, Shunsuke Saito, Hung-Yu Tseng, Changil Kim, Johannes Kopf, and Jia-Bin Huang. 2023. Single-Image 3D Human Digitization with Shape-Guided Diffusion. In *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers '23)*, December 12–15, 2023, Sydney, NSW, Australia. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3610548.3618153>

1 INTRODUCTION

A photorealistic 3D human synthesis is indispensable for a myriad of applications in various fields, including fashion, entertainment, sports, and AR/VR. However, creating a photorealistic 3D human model typically requires multi-view images [Kwon et al. 2021; Liu et al. 2021a; Peng et al. 2021a,b] or 3D scanning systems [Bagautdinov et al. 2021; Saito et al. 2021] as input, which hinders everyone from effortlessly experiencing personalized 3D human digitization. In this work, we aim to create a photorealistic 3D human that can be rendered from arbitrary viewpoints from a *single* input image. Despite its attractive utility, reducing the input to monocular data is highly challenging because the person's backside is not observable, and 3D reconstruction from a single image inherently suffers from depth ambiguity.

To address these challenges, data-driven methods have made significant progress in recent years by incorporating prior information into various 3D representations such as meshes [Alldieck et al. 2019a], voxels [Varol et al. 2018], and neural fields [Saito et al. 2019]. While the geometric fidelity of 3D reconstruction drastically improved over the last several years [Alldieck et al. 2022a; He et al. 2021; Huang et al. 2020; Saito et al. 2020; Xiu et al. 2022; Zheng et al. 2021], its *appearance*, especially for the occluded regions, is still far from photorealistic (Figure 2). This is primarily because these approaches require 3D ground-truth data for supervision, and the available 3D scans of clothed humans are insufficient to learn the entire span of clothing appearance. The appearance of clothing is significantly more diverse than the geometry, and creating a large set of high-quality textured 3D scans of people remains infeasible.

An image collection in the wild is another source of human appearance prior. Images are easily accessible at scale and provide a high variation of clothing appearances. By leveraging large-scale image datasets and high-capacity generative models [Karras et al. 2019, 2020], 2D human synthesis approaches show impressive reposing of clothed humans from a single image [AlBahar et al. 2021; Lewis et al. 2021]. However, they often produce an incoherent appearance with the input image for large rotations because their



Figure 2: Limitations of existing methods. Existing 3D human generation approaches from a single image lack photorealism. Existing methods such as PIFu [Saito et al. 2019] suffer from blurriness; Impersonator++ [Liu et al. 2021b] tends to duplicate content from the front view, suffering from projection artifacts; TEXTure [Richardson et al. 2023] fails to preserve the appearance of the input view and results in saturated colors; Magic123 [Qian et al. 2023] fails to synthesize realistic shape and appearance. Images from Adobe Stock.

underlying representation is not in 3D. While 3D generative models have recently demonstrated 3D-consistent view synthesis of clothed humans [Bergman et al. 2022; Hong et al. 2023; Zhang et al. 2022], we observe that these approaches do not generalize well to various clothing appearances and the results are not sufficiently photorealistic.

In this paper, we argue that the suboptimal performance of existing approaches stems from the limited diversity of training data. However, expanding existing 2D-clothed human datasets also requires nontrivial curation and annotation efforts. To address this limitation, we propose a simple yet effective algorithm to create a 3D consistent textured human from a single image *without* relying on a curated 2D clothed human dataset for appearance synthesis. Our key idea is to utilize powerful 2D generative models trained on an extremely large corpus of images as a human appearance prior. In particular, we use latent diffusion models [Rombach et al. 2022], which allows us to synthesize diverse and photorealistic images. Unlike recent works that leverage 2D diffusion models for 3D object generation from text inputs [Lin et al. 2023; Poole et al. 2022; Richardson et al. 2023], we employ diffusion models to reconstruct a 360-degree view of a real person in the input image in a 3D-consistent manner.

We first reconstruct the 3D geometry of the person using an off-the-shelf tool [Saito et al. 2020] and then generate the back-view of the input image using a 2D single image human reposing approach [AlBahar et al. 2021] to ensure that the completed views are consistent with the input view. Next, we synthesize multi-view images of the person by progressively inpainting novel views utilizing a pretrained inpainting diffusion model guided by both normal and silhouette maps to constrain the synthesis to the underlying 3D structure. To generate a (partial) novel view, we aggregate all other views by blending their RGB color based on importance. Similar to previous work [Buehler et al. 2001; Rong et al. 2022; Xiang et al. 2023], we use the angular differences between the visible pixels of

those views and the current view of interest as well as their distance to the nearest missing pixel to determine the appropriate weight for each view in the blending process. This ensures that the resulting multi-view images are consistent with each other. Finally, we perform multi-view fusion by accounting for slight misalignment in the synthesized multi-view images to obtain a fully textured high-resolution 3D human mesh.

Our experiments show that the proposed approach achieves a more detailed and faithful synthesis of clothed humans than prior methods without requiring high-quality 3D scans or curated large-scale clothed human datasets.

Our contributions include:

- We demonstrate, for the first time, that a 2D diffusion model trained for general image synthesis can be utilized for 3D textured human digitization from a *single* image.
- Our approach preserves the shape and the structural details of the underlying 3D structure by using both normal maps and silhouette to guide the diffusion model.
- We enable 3D consistent texture reconstruction by fusing the synthesized multi-view images into the shared UV texture map.

2 RELATED WORK

2.1 2D human synthesis.

Generative adversarial networks (GANs) enable the photorealistic synthesis of human faces [Karras et al. 2019, 2020] and bodies [Fu et al. 2022]. While these models are unconditional, several works extend them to conditional generative models such that we can control poses while retaining the identity of an input subject. By incorporating additional conditions these works can achieve human reposing [AlBahar and Huang 2019; AlBahar et al. 2021; Liu et al. 2021b; Ma et al. 2017, 2018; Men et al. 2020; Ren et al. 2020; Sarkar et al. 2021; Siarohin et al. 2018; Zhu et al. 2019], virtual try-on [AlBahar et al. 2021; Lewis et al. 2021], motion transfer [Aberman et al. 2019; Chan et al. 2019; Liu et al. 2021b; Yoon et al. 2021]. Pose-with-style [AlBahar et al. 2021] utilizes dense pose [Güler et al. 2018] to warp input images to the target view as an initialization of the synthesis. Impersonator++ [Liu et al. 2021b] further improves the robustness to a large pose change by leveraging a parametric human body model [Loper et al. 2015] and warping blocks to better preserve the information from the input. While these methods enable the control of viewpoints by changing the input pose, the results suffer from view inconsistency. In contrast, our approach achieves 3D consistent generation of textured clothed humans.

2.2 Unconditional 3D human synthesis.

More recently, neural fields and inverse rendering techniques allow us to train 3D GANs with only 2D images [Chan et al. 2022, 2021; Niemeyer and Geiger 2021]. These 3D GANs are extended to articulated full-body humans using warping based on linear blend skinning [Bergman et al. 2022; Hong et al. 2023; Zhang et al. 2022]. By applying inversion [Roich et al. 2022], these methods can generate a 360-degree rendering of a clothed human from a single image. While these results are 3D consistent, we observe that they are plausible only for relatively simple clothing and degrade for

more complex texture patterns. Achieving photorealistic and generalizable 3D human digitization with 3D GANs remains an open problem. Our work achieves better generalization and photorealism by incorporating more general yet highly expressive image priors from diffusion models.

2.3 3D human reconstruction from a single image.

3D reconstruction of clothed humans from a single image is a long-standing problem. A parametric body model [Loper et al. 2015] provides strong prior about the underlying shape of a person, but only for minimally clothed bodies [Kanazawa et al. 2018; Kolotouros et al. 2019; Lassner et al. 2017; Pavlakos et al. 2018]. To enable clothed human reconstruction, regression-based 3D reconstruction has been extended to various shape representations such as voxels [Varol et al. 2018], mesh displacements [Alldieck et al. 2019a,b; Bhatnagar et al. 2019], silhouettes [Natsume et al. 2019], depth maps [Gabeur et al. 2019; Wang et al. 2020], and neural fields [Corona et al. 2021; He et al. 2021; Huang et al. 2020; Saito et al. 2019, 2020; Smith et al. 2019; Xie et al. 2022; Xiu et al. 2023, 2022]. Among them, several works also support texture synthesis for the occluded regions. SiCloPe [Natsume et al. 2019] shows that an image-to-image translation network in screen space can infer occluded textures. PIFu [Saito et al. 2019] infers continuous texture fields [Oechsle et al. 2019] in 3D, which is later improved by explicitly modeling reflectances [Alldieck et al. 2022b]. These approaches, however, often fail to produce photorealistic textures for the back side due to the limited 3D scan data for supervised training. Differentiable rendering based on NeRFs [Mildenhall et al. 2020] has also been applied to learn 3D human representations from images. Both person-specific models [Liu et al. 2021a; Peng et al. 2021b; Weng et al. 2022] and generalizable models across identities [Choi et al. 2022; Gao et al. 2022; Hu et al. 2023; Huang et al. 2022; Kwon et al. 2021; Mihajlovic et al. 2022] have been proposed, but the training requires multi-view images or videos. They are difficult to collect at scale such that the collected data covers a sufficient span of clothing types and textures. Our approach, on the other hand, does not require multi-view images or person-specific video capture.

2.4 Diffusion models for 3D synthesis.

Denoising diffusion models have shown impressive image synthesis results. These powerful 2D generative models are recently adopted to learn 3D scene representations. Recent methods [Chen et al. 2023; Lin et al. 2023; Metzger et al. 2023; Poole et al. 2022; Wang et al. 2022, 2023] have shown that text-to-image models can be repurposed for 3D object generation from text input with remarkable results. Unlike these methods, our method is conditioned on a human input image to create a 3D consistent texture of the person, where the results are photorealistic. Diffusion models can be customized for a specific subject, but this customization typically requires multiple images and a considerable amount of time [Gal et al. 2022; Ruiz et al. 2022]. Moreover, such methods may not consistently maintain the subject's appearance details (i.e. clothing, hairstyle, facial expression, etc.) [Rinon Gal 2023]. These customization methods can be utilized to generate 3D objects conditioned on a single image [Qian et al. 2023; Xu et al. 2022]. Unlike these customization methods, our

method can generate 3D textured human models without test-time finetuning. Moreover, current image-to-3D techniques [Qian et al. 2023; Tang et al. 2023; Xu et al. 2022] lack human-specific prior and hence struggle to synthesize realistic and detailed textured human models. The closest to our work is TEXTure [Richardson et al. 2023], which utilizes 2D diffusion models to synthesize texture of an input mesh. We observe that their shape guidance based on depth maps is insufficient for photorealistic clothed human synthesis. Instead of progressively refining the texture based on viewing angles, we improve the consistency by blending the RGB color of existing views, weighted by visibility, viewing angles, and distance to missing regions. We also improve the per-view synthesis by incorporating normal and silhouette maps as guidance signals.

3 METHOD

Our goal is to generate a 360-degree view of a person with a consistent, high-resolution appearance from a *single* input image. To this end, we first synthesize a set of multi-view images of the person $\{\hat{I}_2, \dots, \hat{I}_N\}$ that are consistent among each other and coherent with the input image I_1 (Figure 3). In particular, we use the reconstructed 3D geometry of the person to guide the inpainting with diffusion models (Figure 4). For 3D shape reconstruction, we employ an off-the-shelf method [Saito et al. 2020] to obtain a triangular mesh G of the input person using Marching cubes [Lorensen and Cline 1987].

We synthesize the multi-view images in an *auto-regressive* manner. More specifically, we start with synthesizing the back-view of the person with [AlBahar et al. 2021] (Section 3.1). The input and the synthesized back-view images form an initial *support set* V (i.e., currently available views). Using the images from the support set and the mesh G , we can render a new view of the person (Section 3.2). Here, this blended view is consistent with the previously generated images but may have missing regions (that are not covered by any of the images in the support set). We use a shape-guided diffusion model to inpaint the appearance details while respecting the estimated shape (Section 3.3). We expand the support set by adding this inpainted view and proceed to a new view until all the views are generated. We sample views at intervals of 45° , specifically in the order of $[45^\circ, -45^\circ, 90^\circ, -90^\circ, 135^\circ, -135^\circ, 180^\circ]$. Thus, our support set will have a total of 8 views ($N = 8$). When we use more viewpoints, the missing regions become very small. In such cases, we found that the inpainting performance deteriorates. On the other hand, when we use less viewpoints, the missing regions become very large. We found that the inpainting fails to preserve the input appearance.

We then fuse these multi-view images $\{I_1, \hat{I}_2, \dots, \hat{I}_N\}$ via inverse rendering robust to slight misalignment and optimize a UV texture map T (Figure 5). We finally use this UV texture map T to render the 360-degree view of the person. Note that our approach assumes weak perspective projection for simplicity, following [Saito et al. 2019, 2020; Xiu et al. 2022], but extending it to a perspective camera is also possible.

3.1 Back-view Synthesis

The input frontal and back views have strong semantics correlations (e.g., the back side of a T-shirt is likely a T-shirt with similar textures), and its silhouette contour provides structural guidance.

Thus, we first synthesize the back-view of the person for guidance *prior* to synthesizing other views. While prior works [He et al. 2021; Natsume et al. 2019] show that front-to-back synthesis is highly effective with supervised training, our approach achieves the front-to-back synthesis without relying on ground-truth paired data. More specifically, we apply the SoTA 2D human synthesis method [AlBahar et al. 2021] with the inferred dense pose prediction for the back-view. To generate a dense pose prediction that aligns precisely with the input image, we render the surface normal and depth map of the shape G from the view opposite to the input view and create a photorealistic back-view using ControlNet [Zhang and Agrawala 2023] with the text prompt of “*back view of a person wearing nice clothes in front of a solid gray background, best quality.*” We then run dense pose [Güler et al. 2018], which is finally fed into Pose-with-Style [AlBahar et al. 2021]. We empirically find that using Pose-with-Style [AlBahar et al. 2021] with the aforementioned procedure leads to a more semantically consistent back-view than just using ControlNet [Zhang and Agrawala 2023]. See Figure 7 for the impact of the back-view initialization.

3.2 Multi-view visible texture aggregation

Prior to inpainting, we aggregate all the views in the support set V to the target view V_c . However, naively averaging all views leads to a blurry image due to slight misalignment in each view. To ensure that high-resolution details are all retained, we use weighted averaging using confidence based on visibility, viewing angles, and distance.

For each view V_o in the set of synthesized views V_o , we render the normal map N_o^c as well as its color C_o^c from V_c . In addition, we set the visibility mask M_o of each view V_o by comparing its visible faces to the visible faces from V_c . We use this visibility mask M_o to compute distance transform d_o from the boundary of the visible pixels and the invisible pixels in each view V_o . We also compute the angular difference ϕ_o of each visible pixel between view V_o and the current view of interest V_c as follows:

$$\phi_o = M_o \arccos \left(\frac{N_o^c \cdot N_c}{\max(\|N_o^c\|_2 \cdot \|N_c\|_2, \epsilon)} \right), \quad (1)$$

where $\epsilon = 10^{-8}$ is a small value to avoid dividing by zero.

Finally, we compute the blending weight w_o of view V_o as follows:

$$w_o = \frac{M_o B_o e^{-\alpha \phi_o} d_o^\beta}{\sum_{i \in V} M_i B_i e^{-\alpha \phi_i} d_i^\beta + \epsilon}. \quad (2)$$

In our experiments, we set both α , which determines the strength of the angular difference, and β , which determines the strength of the Euclidean distance, to 3. Using the angular difference ϕ_o ensures a higher weight to closer views, while using the Euclidean distance d_o ensures a lower weight for pixels close to the missing region. Moreover, if only one existing view contains a specific pixel, we mark its boundary B_o as invisible. This ensures that the target view does not suffer from boundary artifacts.

We use the computed weights w_o to blend the color C_o of the previously synthesized views V_o together, where the blended image of the current view I_c and its visibility mask M_c are as follows:

$$M_c = \bigcup_{i \in V} M_i, \quad \text{and} \quad I_c = \sum_{i \in V} w_i C_i. \quad (3)$$

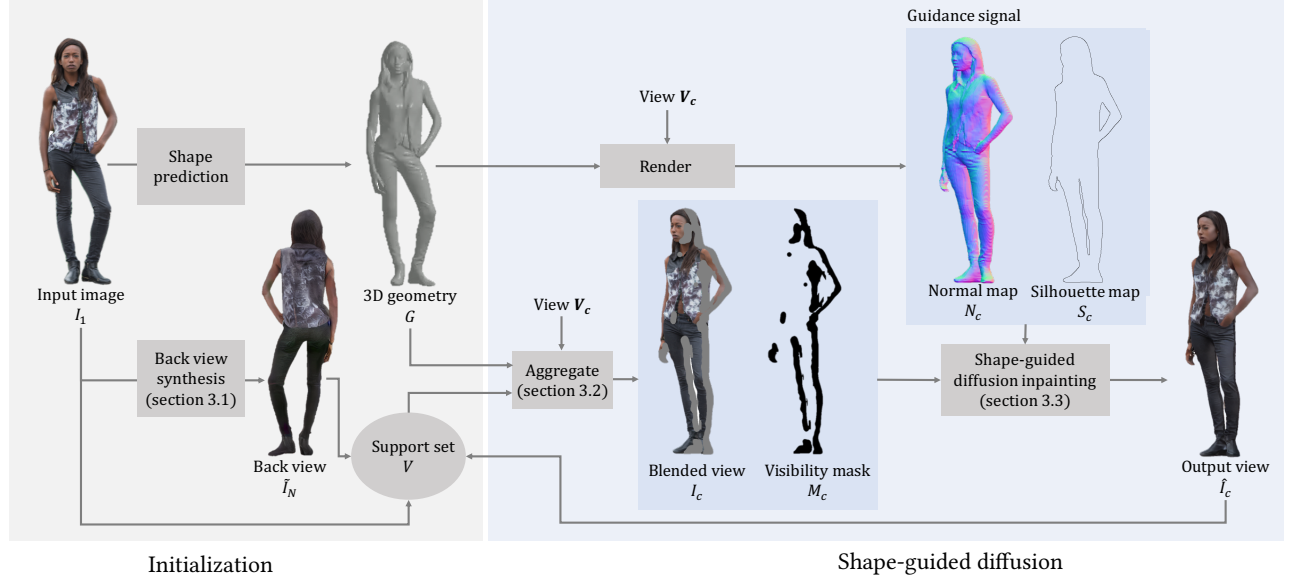


Figure 3: Person image generation with shape-guided diffusion. To generate a 360-degree view of a person from a *single* image I_1 , we first synthesize multi-view images of the person. We use an off-the-shelf method to infer the 3D geometry [2020] and synthesize an initial back-view \hat{I}_N of the person [2021] as a guidance. We add our input view I_1 and the synthesized initial back-view \hat{I}_N to our support set V . To generate a new view V_c , we aggregate all the visible pixels from our support set V by blending their RGB color, weighted by visibility, viewing angles, and the distance to missing regions. To hallucinate the unseen appearance and synthesize view \hat{I}_c , we use a pretrained inpainting diffusion model guided by shape cues (normal N_c and silhouette S_c maps). We include the generated view \hat{I}_c in our support set and repeat this process for all the remaining views. Images from Adobe Stock.

The final blended image I_c and its visibility mask M_c are then used to synthesize a complete view \hat{I}_c using our shape-guided diffusion.

3.3 Shape-guided diffusion inpainting

To synthesize the unseen appearance indicated by the visibility mask M_c in the blended image I_c , we use a 2D inpainting diffusion model [Rombach et al. 2022]. However, we observe that without any guidance, the inpainted regions often do not respect the underlying geometry G (see Figure 4(a)). To address this, we use the method of ControlNet [Zhang and Agrawala 2023] by incorporating additional structural information into the diffusion model. When we use normal maps as a control signal, we can preserve the structural details of the mesh but not the shape of the human body (Figure 4(b)). On the other hand, using the silhouette map alone preserves the shape of the human body, but not the structural details of the mesh (Figure 4(c)). To best guide the inpainting model to respect the underlying 3D geometry, we propose to use both normal map and silhouette maps, as shown in Figure 4(d). We add this generated view to our support set V and proceed to the next view until all N views are synthesized.

3.4 Multi-view fusion

Since the latent diffusion model operates inpainting in the low-resolution latent space, the final synthesized images do not form geometrically consistent multi-view images. Therefore, we consolidate these slightly misaligned multi-view images $I_1, \hat{I}_2, \dots, \hat{I}_N$ into



Figure 4: Shape-guided diffusion inpainting. To synthesize the unseen appearance in a new view, we use a pretrained inpainting diffusion model. With no guidance, the inpainted regions often do not preserve the shape (red silhouette) nor the structural details of the 3D geometry (a). If we use normal maps as a control signal for ControlNet [2023] (b), the inpainted region preserves the structural details of the mesh (e.g., fingers), but not the shape of the human body. Using the silhouette map preserves the shape of the human body, but not the structural details of the mesh (c). We propose to use both normal and silhouette maps to guide the inpainting model to respect the underlying 3D geometry (d). Images from Adobe Stock.

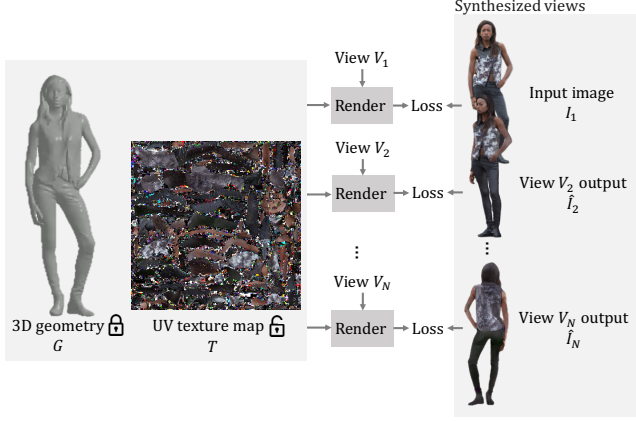


Figure 5: Multi-view fusion. We fuse the synthesized multi-view images $\{I_1, \hat{I}_2, \dots, \hat{I}_N\}$ (see Figure 3) to obtain a textured 3D human mesh. We use the computed UV parameterization [2021] to optimize a UV texture map T with the geometry G fixed. In each iteration, we differentially render the UV texture map T in every synthesized view from our set of views $\{V = V_1, V_2, \dots, V_N\}$. We minimize the reconstruction loss between the rendered view and our synthesized view using both LPIPS loss [2018] and L1 loss. The fusion results in a textured mesh that can be rendered from any view. Images from Adobe Stock.

a single consistent 3D texture map T . We show the overview of our multi-view fusion in Figure 5.

We first compute the UV parameterization of the reconstructed 3D geometry using xatlas [Young 2021]. Then, we optimize a UV texture map T via inverse rendering with loss functions that are robust to small misalignment. In every iteration, we render the UV texture map T in each view i from our set of synthesized views $\{V = V_1, V_2, \dots, V_N\}$ and minimize the reconstruction loss of this rendered view and the synthesized view using both LPIPS loss [Zhang et al. 2018] and L1 loss such that:

$$L(T) = \sum_{i \in V} L_{\text{LPIPS}}(\text{Render}(T; G, i), \hat{I}_i) + \lambda L_1(\text{Render}(T; G, i), \hat{I}_i), \quad (4)$$

where $\hat{I}_1 = I_1$ and λ is set to 10.

Once the texture map T is optimized, one can render the textured mesh from arbitrary viewpoints.

4 EXPERIMENTAL RESULTS

4.1 Experimental Setup

4.1.1 Implementation details. We implement our approach with PyTorch on a single RTX A6000 GPU. We set the guidance scale of the pretrained inpainting diffusion model to 15 and the number of inference steps per view to 25. In all our experiments, we use a generic text prompt for all subjects: “a person wearing nice clothes in front of a solid white background, <VIEW> view, best quality, extremely detailed”, where <VIEW> is set to “front” for frontal views; “left” and “right” for 45° and −45° views, respectively; “side” for ±90° views; and “back” for the rest of viewing angles (±135° and

Table 1: Quantitative comparisons with baseline methods on the THuman2.0 dataset [Yu et al. 2021].

Methods	PSNR↑	SSIM↑	FID↓	LPIPS↓	CLIP-score↑
PwS baseline	17.8003	0.8888	132.4511	0.1320	0.7733
PIFu	18.0934	0.9117	150.6622	0.1372	0.7721
Impersonator++	16.4791	<u>0.9012</u>	106.5753	0.1468	0.8168
TEXTure	16.7869	0.8740	215.7078	0.1435	0.7272
Magic123	14.5013	0.8768	137.1108	0.1880	<u>0.7996</u>
S3F	14.1212	0.8840	165.9806	0.1868	0.7475
Ours	17.3651	0.8946	<u>115.9918</u>	0.1300	0.7992

180°). We use the ADAM optimizer with a learning rate of 0.1 and with $\beta_1 = 0.9$ and $\beta_2 = 0.999$ to learn the UV texture map T . The entire process of generating a 3D textured model from a single image takes approximately 7 minutes on an RTX A6000 GPU.

4.1.2 Datasets. To evaluate our approach, we utilize the THuman2.0 dataset [Yu et al. 2021], using 30 subjects, evenly split between 15 males and 15 females. We use front-facing images as input. We also evaluate our approach on the DeepFashion dataset [Liu et al. 2016] to compare with ELICIT [Huang et al. 2022]. We additionally use in-the-wild images from Adobe Stock¹ to showcase results from images with diverse subjects, clothing, and poses.²

4.1.3 Baselines. We compare our 360-degree view synthesis approach with Pose with Style (PwS) baseline. We use Pose with Style [AlBahar et al. 2021] to generate multi-view images and then fuse them using our multi-view fusion. We also compare with PIFu [Saito et al. 2019], Impersonator++ [Liu et al. 2021b], TEXTure [Richardson et al. 2023], Magic123 [Qian et al. 2023], and S3F [Corona et al. 2023]. To make TEXTure [Richardson et al. 2023] conditional on an input image, we use the input image directly instead of generating an initial view from the depth-to-image diffusion model. We also compare our work with ELICIT [Huang et al. 2022] on a subset of the DeepFashion dataset [Liu et al. 2016] provided by its authors.

4.2 Quantitative Comparison

To quantify the quality of our results, we measure peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM), Frechet Inception Distance (FID) [Parmar et al. 2022], learned perceptual image patch similarity (LPIPS) [Zhang et al. 2018], and CLIP-score. CLIP-score measures the cosine similarity between the CLIP embeddings of an input image and each of the synthesized views. We use a total of 90 synthesized views with 4° spacing. We compare these metrics on the THuman2.0 dataset [Yu et al. 2021] with other baselines in Table 1. Quantitative results show that existing metrics are not consistent in evaluating 3D textured humans. PSNR favors blurry images as in PIFu [Saito et al. 2019], and FID does not provide accurate results for sparse view distributions. To quantitatively compare with ELICIT [Huang et al. 2022], we

¹<https://stock.adobe.com/>

²All datasets used in this research were exclusively downloaded, accessed, and utilized on UMD clusters.




Figure 6: Limitations. Our approach inherits limitations from existing methods for shape reconstruction (unusual foot shape (left)) and back-view synthesis (misaligned skirt length due to lack of geometry awareness (middle)). We also show the baked specularities on the face and garment texture, which is ideally view-dependent (right). Images from Adobe Stock.




Figure 7: The need of back-view synthesis. Having an initial back-view encourages all other views to preserve the appearance of the person in the input image especially when a target view is far from the input view. Images from Adobe Stock.

compute the CLIP-score (where higher values indicate better performance) on their provided subset of the DeepFashion dataset [Liu et al. 2016]. Our method achieved a CLIP-score of 0.7732, surpassing their score of 0.7236.

4.3 Qualitative Comparison

We show visual comparisons of our results with the baselines on in-the-wild images from Adobe Stock in Figures 1 and 8, and on the THuman2.0 dataset [Yu et al. 2021] in Figure 9. These results demonstrate that our method produces high-resolution, photorealistic 3D human models that respect the appearance of the input, for a variety of input images.

4.4 Ablation Study

4.4.1 Guidance signals. We validate our shape-guided diffusion inpainting in Table 2. We show the effect of using no guidance (B), only normal maps (C), only silhouette maps (D), and both normal and silhouette maps (E). We also show visual comparison in Figure 4. The use of both normal maps and silhouette maps leads to better preserving the synthesized person’s shape and details and thus enhancing the quality of resulting 3D human models.

4.4.2 Back-view synthesis. We validate the initial back-view synthesis using a human reposing technique [AlBahar et al. 2021] in Table 2 (A vs. E). We also show visual comparison in Figure 7. Having an initial back view encourages all other views to preserve

Table 2: Ablation study on the THuman2.0 dataset [Yu et al. 2021]. We use the ground truth mesh to evaluate the effectiveness of initializing the back-view (B), and using normal (N) and silhouette (S) maps as guidance signals.

ID	B	N	S	PSNR↑	SSIM↑	FID↓	LPIPS↓	CLIP-score↑
A		✓	✓	23.9463	0.9373	117.7447	0.0538	0.8013
B	✓			24.0494	0.9389	129.4944	0.0592	0.7896
C	✓	✓		25.8709	0.9449	108.5836	0.0506	<u>0.8041</u>
D	✓		✓	25.7199	0.9435	<u>101.3901</u>	<u>0.0480</u>	0.8013
E	✓	✓	✓	<u>25.8465</u>	0.9453	98.9282	0.0473	0.8069

the appearance of the input person, especially when clothing has nontrivial textures.

4.5 Limitations and Future Work

Our main limitation is the dependence on off-the-shelf methods [AlBahar et al. 2021; Saito et al. 2020] for the base geometry reconstruction and back-view synthesis. Figure 6 shows that our approach inherits the limitations of these methods. Another limitation is the lack of view-dependency. While clothing is mostly diffuse, human skin may exhibit view-dependent specular highlights. Extending our approach to view-dependent radiance would be an exciting direction, which can be addressed by future work. Furthermore, our work does not support human reposing and it requires per-subject UV texture optimization. For the generality of our approach, we use off-the-shelf 3D shape reconstruction methods for clothed humans [Saito et al. 2020; Xiu et al. 2022], which are trained on 3D ground-truth data. We also use off-the-shelf human reposing method [AlBahar et al. 2021] for the back-view synthesis. Future work should also enable the high-fidelity 3D shape reconstruction of clothed humans and back-view synthesis with general-purpose 2D diffusion models.

5 CONCLUSIONS

We introduced a simple yet highly effective approach to generate a fully textured 3D human mesh from a *single* image. Our experiments show that synthesizing a high-resolution and photorealistic texture for occluded views is now possible with shape-guided inpainting based on high-capacity latent diffusion models and a robust multi-view fusion method. While 3D human digitization relies on curated human-centric datasets either in 3D or 2D, our approach, for the first

time, achieves superior synthesis results by leveraging a general-purpose large-scale diffusion model. We believe our work will shed light on unifying data collection efforts for 3D human digitization and other general 2D/3D synthesis methods.

REFERENCES

- Kfir Aberman, Mingyi Shi, Jing Liao, Dani Lischinski, Baoquan Chen, and Daniel Cohen-Or. 2019. Deep video-based performance cloning. In *Computer Graphics Forum*, Vol. 38. 219–233.
- Badour AlBahar and Jia-Bin Huang. 2019. Guided image-to-image translation with bi-directional feature transformation. In *ICCV*.
- Badour AlBahar, Jingwan Lu, Jimei Yang, Zhixin Shu, Eli Shechtman, and Jia-Bin Huang. 2021. Pose with Style: Detail-Preserving Pose-Guided Image Synthesis with Conditional StyleGAN. *ACM TOG* (2021).
- Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. 2019a. Learning to reconstruct people in clothing from a single RGB camera. In *CVPR*.
- Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. 2019b. Tex2shape: Detailed full human body geometry from a single image. In *ICCV*.
- Thiemo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. 2022a. Photorealistic monocular 3d reconstruction of humans wearing clothing. In *CVPR*.
- Thiemo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. 2022b. Photorealistic Monocular 3D Reconstruction of Humans Wearing Clothing. In *CVPR*.
- Timur Bagautdinov, Chenglei Wu, Tomas Simon, Fabian Prada, Takaaki Shiratori, Shih-En Wei, Weipeng Xu, Yaser Sheikh, and Jason Saragih. 2021. Driving-signal aware full-body avatars. *ACM TOG* 40, 4 (2021), 1–17.
- Alexander W. Bergman, Petr Kellnhofer, Wang Yifan, Eric R. Chan, David B. Lindell, and Gordon Wetzstein. 2022. Generative Neural Articulated Radiance Fields. In *NeurIPS*.
- Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. 2019. Multi-garment net: Learning to dress 3d people from images. In *ICCV*.
- Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. 2001. Unstructured lumigraph rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. 425–432.
- Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. 2019. Everybody dance now. In *ICCV*.
- 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, et al. 2022. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*.
- Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. 2021. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *CVPR*.
- Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. 2023. Fantasia3D: Disentangling Geometry and Appearance for High-quality Text-to-3D Content Creation. In *ICCV*.
- Hongsuk Choi, Gyeongsik Moon, Matthieu Armando, Vincent Leroy, Kyoung Mu Lee, and Gregory Rogez. 2022. MonoNHR: Monocular Neural Human Renderer. *International Conference on 3D Vision*.
- Enric Corona, Albert Pumarola, Guillem Alenya, Gerard Pons-Moll, and Francesc Moreno-Noguer. 2021. Smplicit: Topology-aware generative model for clothed people. In *CVPR*.
- Enric Corona, Mihai Zanfir, Thiemo Alldieck, Eduard Gabriel Bazavan, Andrei Zanfir, and Cristian Sminchisescu. 2023. Structured 3D Features for Reconstructing Relightable and Animatable Avatars. In *CVPR*.
- Jianglin Fu, Shikai Li, Yuming Jiang, Kwan-Yee Lin, Chen Qian, Chen Change Loy, Wayne Wu, and Ziwei Liu. 2022. Stylegan-human: A data-centric odyssey of human generation. In *ECCV*.
- Valentin Gabeur, Jean-Sébastien Franco, Xavier Martin, Cordelia Schmid, and Gregory Rogez. 2019. Moulding humans: Non-parametric 3d human shape estimation from single images. In *ICCV*.
- Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H. Bermanto, Gal Chechik, and Daniel Cohen-Or. 2022. An Image is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion. (2022). <https://doi.org/10.48550/ARXIV.2208.01618>
- Xiangjun Gao, Jiaolong Yang, Jongyoo Kim, Sida Peng, Zicheng Liu, and Xin Tong. 2022. MPS-NeRF: Generalizable 3D Human Rendering From Multiview Images. *IEEE TPAMI* (2022), 1–12.
- Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. 2018. Densepose: Dense human pose estimation in the wild. In *CVPR*.
- Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. 2021. Arch++: Animation-ready clothed human reconstruction revisited. In *Proceedings of the IEEE/CVF international conference on computer vision*. 11046–11056.
- Fangzhou Hong, Zhaoxi Chen, Yushi Lan, Liang Pan, and Ziwei Liu. 2023. EVA3D: Compositional 3D Human Generation from 2D Image Collections. In *ICLR*.
- Shoukang Hu, Fangzhou Hong, Liang Pan, Haiyi Mei, Lei Yang, and Ziwei Liu. 2023. SHERF: Generalizable Human NeRF from a Single Image. In *ICCV*.
- Yangyi Huang, Hongwei Yi, Weiyang Liu, Haofan Wang, Boxi Wu, Wenxiao Wang, Binbin Lin, Debing Zhang, and Deng Cai. 2022. One-shot Implicit Animatable Avatars with Model-based Priors. *arXiv preprint arXiv:2212.02469* (2022).
- Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. 2020. Arch: Animatable reconstruction of clothed humans. In *CVPR*.
- Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. 2018. End-to-end recovery of human shape and pose. In *CVPR*.
- Tero Karras, Samuli Laine, and Timo Aila. 2019. A style-based generator architecture for generative adversarial networks. In *CVPR*.
- Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. Analyzing and improving the image quality of stylegan. In *CVPR*.
- Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. 2019. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *ICCV*.
- Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. 2021. Neural human performer: Learning generalizable radiance fields for human performance rendering. *Advances in Neural Information Processing Systems* 34 (2021).
- Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J Black, and Peter V Gehler. 2017. Unite the people: Closing the loop between 3d and 2d human representations. In *CVPR*.
- Kathleen M Lewis, Srivatsan Varadharajan, and Ira Kemelmacher-Shlizerman. 2021. Tryongan: Body-aware try-on via layered interpolation. *ACM TOG* 40, 4 (2021), 1–10.
- Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiao-hui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. 2023. Magic3D: High-Resolution Text-to-3D Content Creation. In *CVPR*.
- Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. 2021a. Neural Actor: Neural Free-view Synthesis of Human Actors with Pose Control. *ACM TOG* (2021).
- Wen Liu, Zhixin Piao, Zhi Tu, Wenhan Luo, Lin Ma, and Shenghua Gao. 2021b. Liquid warping GAN with attention: A unified framework for human image synthesis. *IEEE TPAMI* (2021).
- Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. 2016. DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations. In *CVPR*.
- Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. 2015. SMPL: A skinned multi-person linear model. *ACM TOG* 34, 6 (2015), 1–16.
- William E Lorensen and Harvey E Cline. 1987. Marching cubes: A high resolution 3D surface construction algorithm. *ACM TOG* 21, 4 (1987), 163–169.
- Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. 2017. Pose guided person image generation. In *NeurIPS*.
- Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc Van Gool, Bernt Schiele, and Mario Fritz. 2018. Disentangled person image generation. In *CVPR*.
- Yifan Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. 2020. Controllable person image synthesis with attribute-decomposed gan. In *CVPR*.
- Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. 2023. Latent-NeRF for Shape-Guided Generation of 3D Shapes and Textures. In *CVPR*.
- Marko Mihajlovic, Aayush Bansal, Michael Zollhoefer, Siyu Tang, and Shunsuke Saito. 2022. KeypointNeRF: Generalizing image-based volumetric avatars using relative spatial encoding of keypoints. In *ECCV*.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- Ryota Natsume, Shunsuke Saito, Zeng Huang, Weikai Chen, Chongyang Ma, Hao Li, and Shigeo Morishima. 2019. Siclope: Silhouette-based clothed people. In *CVPR*.
- Michael Niemeyer and Andreas Geiger. 2021. Giraffe: Representing scenes as compositional generative neural feature fields. In *CVPR*.
- Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. 2019. Texture fields: Learning texture representations in function space. In *ICCV*.
- Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. 2022. On Aliased Resizing and Surprising Subtleties in GAN Evaluation. In *CVPR*.
- Georgios Pavlakos, Luyang Zhu, Xiao-wei Zhou, and Kostas Daniilidis. 2018. Learning to estimate 3D human pose and shape from a single color image. In *CVPR*.
- Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiao-wei Zhou, and Hujun Bao. 2021a. Animatable Neural Radiance Fields for Modeling Dynamic Human Bodies. In *ICCV*.
- Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiao-wei Zhou. 2021b. Neural Body: Implicit Neural Representations with Structured Latent Codes for Novel View Synthesis of Dynamic Humans. In *CVPR*.
- Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. 2022. DreamFusion: Text-to-3D using 2D Diffusion. In *ICLR*.
- Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, and Bernard Ghanem. 2023. Magic123: One Image to High-Quality 3D Object Generation Using Both 2D and 3D Diffusion Priors. *arXiv preprint arXiv:2306.17843* (2023).

- Yurui Ren, Xiaoming Yu, Junming Chen, Thomas H Li, and Ge Li. 2020. Deep image spatial transformation for person image generation. In *CVPR*.
- Elad Richardson, Gal Metzger, Yuval Alaluf, Raja Giryes, and Daniel Cohen-Or. 2023. TEXTure: Text-Guided Texturing of 3D Shapes. *ACM TOG* (2023).
- Yuval Atzmon Amit H. Bermano Gal Chechik Daniel Cohen-Or Rinon Gal, Moab Arar. 2023. Encoder-based Domain Tuning for Fast Personalization of Text-to-Image Models. (2023). <https://arxiv.org/abs/2302.12228>
- Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. 2022. Pivotal tuning for latent-based editing of real images. *ACM TOG* 42, 1 (2022), 1–13.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *CVPR*.
- Xuejian Rong, Jia-Bin Huang, Ayush Saraf, Changil Kim, and Johannes Kopf. 2022. Boosting View Synthesis with Residual Transfer. In *CVPR*.
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2022. DreamBooth: Fine Tuning Text-to-image Diffusion Models for Subject-Driven Generation. (2022).
- Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. 2019. PIFu: Pixel-Aligned Implicit Function for High-Resolution Clothed Human Digitization. In *ICCV*.
- Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. 2020. PIFuHD: Multi-Level Pixel-Aligned Implicit Function for High-Resolution 3D Human Digitization. In *CVPR*.
- Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. 2021. SCANimate: Weakly supervised learning of skinned clothed avatar networks. In *CVPR*.
- Kripasindhu Sarkar, Vladislav Golyanik, Lingjie Liu, and Christian Theobalt. 2021. Style and Pose Control for Image Synthesis of Humans from a Single Monocular View. *arXiv preprint arXiv:2102.11263* (2021).
- Aliaksandr Siarohin, Enver Sangineto, Stéphane Lathuiliere, and Nicu Sebe. 2018. Deformable gans for pose-based human image generation. In *CVPR*.
- David Smith, Matthew Loper, Xiaochen Hu, Paris Mavroidis, and Javier Romero. 2019. Facsimile: Fast and accurate scans from an image in less than a second. In *ICCV*.
- Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. 2023. Make-It-3D: High-Fidelity 3D Creation from A Single Image with Diffusion Prior. *arXiv preprint arXiv:2303.14184* (2023).
- Gul Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. 2018. Bodynet: Volumetric inference of 3d human body shapes. In *ECCV*.
- Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A. Yeh, and Greg Shakhnarovich. 2022. Score Jacobian Chaining: Lifting Pretrained 2D Diffusion Models for 3D Generation. *arXiv preprint arXiv:2212.00774* (2022).
- Lizhen Wang, Xiaochen Zhao, Tao Yu, Songtao Wang, and Yebin Liu. 2020. NormalGAN: Learning Detailed 3D Human from a Single RGB-D Image. In *ECCV*.
- Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. 2023. ProlificDreamer: High-Fidelity and Diverse Text-to-3D Generation with Variational Score Distillation. *arXiv preprint arXiv:2305.16213* (2023).
- Chung-Yi Weng, Brian Curless, Pratul P. Srinivasan, Jonathan T. Barron, and Ira Kemelmacher-Shlizerman. 2022. HumanNeRF: Free-Viewpoint Rendering of Moving People From Monocular Video. In *CVPR*.
- Jianfeng Xiang, Jiaolong Yang, Binbin Huang, and Xin Tong. 2023. 3D-aware Image Generation using 2D Diffusion Models. *arXiv preprint arXiv:2303.17905* (2023).
- Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. 2022. Neural fields in visual computing and beyond. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 641–676.
- Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J. Black. 2023. ECON: Explicit Clothed humans Optimized via Normal integration. In *CVPR*.
- Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J. Black. 2022. ICON: Implicit Clothed humans Obtained from Normals. In *CVPR*.
- Dejia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. 2022. NeuralLift-360: Lifting An In-the-wild 2D Photo to A 3D Object with 360° Views. *arXiv preprint arXiv:2211.16431*.
- Jae Shin Yoon, Lingjie Liu, Vladislav Golyanik, Kripasindhu Sarkar, Hyun Soo Park, and Christian Theobalt. 2021. Pose-Guided Human Animation from a Single Image in the Wild. In *CVPR*.
- Jonathan Young. 2021. xatlas: Mesh parameterization / UV unwrapping library. <https://github.com/jpcy/xatlas>.
- Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. 2021. Function4D: Real-time Human Volumetric Capture from Very Sparse Consumer RGBD Sensors. In *CVPR*.
- Jianfeng Zhang, Zihang Jiang, Dingdong Yang, Hongyi Xu, Yichun Shi, Guoxian Song, Zhongcong Xu, Xinchao Wang, and Jiashi Feng. 2022. AvatarGen: A 3D Generative Model for Animatable Human Avatars. *Arxiv* (2022).
- Lvmin Zhang and Maneesh Agrawala. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. *arXiv:2302.05543 [cs.CV]*
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*.
- Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. 2021. PaMIR: Parametric Model-Conditioned Implicit Representation for Image-based Human Reconstruction. *IEEE TPAMI* (2021).
- Zhen Zhu, Tengfeng Huang, Baoguang Shi, Miao Yu, Bofei Wang, and Xiang Bai. 2019. Progressive Pose Attention Transfer for Person Image Generation. In *CVPR*.



Figure 8: Visual comparison on in-the-wild images from Adobe Stock. We compare our 3D human digitization approach with prior methods [Corona et al. 2023; Liu et al. 2021b; Qian et al. 2023; Richardson et al. 2023; Saito et al. 2019] on images in-the-wild to showcase the generalizability of our approach. Our approach demonstrates high-resolution photorealistic results that preserve the appearance of the input image.



Figure 9: Visual comparisons on the THuman2.0 dataset. We compare our approach with prior methods [AlBahar et al. 2021; Corona et al. 2023; Liu et al. 2021b; Qian et al. 2023; Richardson et al. 2023; Saito et al. 2019] on the THuman2.0 dataset [Yu et al. 2021]. Our results showcase photorealistic images with consistent views that are consistent with the input images.