

Chupa: Carving 3D Clothed Humans from Skinned Shape Priors using 2D Diffusion Probabilistic Models

Byungjun Kim^{1*} Patrick Kwon^{2*} Kwangho Lee² Myunggi Lee²

Sookwan Han¹ Daesik Kim² Hanbyul Joo¹

¹Seoul National University ²Naver Webtoon AI

<https://snuvclab.github.io/chupa/>




Figure 1. **Generative Human Digital Avatars.** We propose *Chupa*, a 3D human generation pipeline that combines the generative power of diffusion models [61] and neural rendering techniques [39] to create diverse, and realistic 3D humans. Our pipeline can easily generalize to unseen human poses and display realistic qualities.

Abstract

We propose a 3D generation pipeline that uses diffusion models to generate realistic human digital avatars. Due to the wide variety of human identities, poses, and stochastic details, the generation of 3D human meshes has been a challenging problem. To address this, we decompose the problem into 2D normal map generation and normal map-based 3D reconstruction. Specifically, we first simultaneously generate realistic normal maps for the front and backside of a clothed human, dubbed *dual normal maps*, using a pose-conditional diffusion model. For 3D reconstruction, we “carve” the prior SMPL-X mesh to a detailed 3D mesh according to the normal maps through mesh optimization. To further enhance the

high-frequency details, we present a diffusion resampling scheme on both body and facial regions, thus encouraging the generation of realistic digital avatars. We also seamlessly incorporate a recent text-to-image diffusion model to support text-based human identity control. Our method, namely, *Chupa*, is capable of generating realistic 3D clothed humans with better perceptual quality and identity variety.

1. Introduction

The creation of clothed 3D human characters, which we refer to as “digital avatars”, has become an essential part of many fields including gaming, animation, virtual/mixed reality, and the 3D industry in general. These digital avatars allow users to use their virtual representation for a range

*Equal contribution

of purposes, thus enhancing user immersion within such services. However, creating high-quality digital avatars often requires specialized 3D artists using a sophisticated creation pipeline [23, 42], making it a laborious process.

The recent advances in deep generative models [21, 25, 37] have enabled the creation of high-quality images that accurately reflect the textual input semantics [49, 61]. However, the usage of such generative models in creating 3D has mainly focused on object generation [56, 72, 89, 92, 94] and shown rather limited performance in generating full-body, realistic 3D human avatars due to the difficulty of collecting a large-scale ground truth dataset. Many previous 3D generative models [2, 6, 7, 22, 27, 52, 93] focus on training generative models on large-scale image datasets along with implicit 3D shape representations and differentiable volume rendering [51, 82]. However, those approaches are rather limited in generating full-body humans with realistic details and rely on computationally expensive volume rendering. Other approach [9] directly uses high-quality 3D datasets [60, 91] to train generative models based on auto-decoding frameworks [53], but the resulting stochastic details tend to be unrealistic, due to the usage of an adversarial loss [21].

In this paper, we decompose the problem of 3D generation into 2D normal map generation and 3D reconstruction, bridging the power of generative models in the image domain toward 3D generation. Following the intuition of “sandwich-like” approaches for single image-based 3D human reconstruction [19, 73, 87], we generate normal maps for frontal and backside regions of human mesh to get rich details mitigating the computational cost of 3D representations. We adopt a diffusion model [25, 61] to simultaneously create consistent normal maps for both frontal and backside regions, which we call *dual* normal maps, conditioned on a posed SMPL-X [43, 54]. Since diffusion models are well known for their mode coverage [85], we find it suitable to generate diverse 3D digital avatars. The dual normal maps are then used as input for our 3D reconstruction pipeline, in which we *carve* the initial posed SMPL-X mesh to a clothed, realistic human mesh with normal map-based mesh optimization inspired by NDS [84]. During optimization, the initial mesh is gradually deformed to match the generated normal maps through a differentiable rasterization pipeline [39] and geometric regularization including a loss function for plausible side-view. Our dual normal map-based 3D generation pipeline alleviates the difficulty of generating consistent multi-views, which is the fundamental reason that diffusion-based 3D generative models [56, 79, 89] suffer from slow convergence or fail to generate multi-view consistent results. We show that the diffusion model can generate consistent dual normal maps and they are sufficient to generate plausible 3D humans along with SMPL-X prior. Then, we can further improve the generated mesh by using a resampling scheme motivated by SDEdit [47], in which we

use separate diffusion models for the body and facial regions to refine the perceptual quality of the rendered normals in different viewpoints while preserving the view and identity consistency. The refined normal maps are subsequently used as inputs for the mesh optimization, thus creating a realistic 3D digital avatar with high-frequency details.

As shown in Fig. 1, our pipeline, which we dub it *Chupa*, can be extended to text-based generation for further controllability on the human identity (*e.g.*, gender, clothing, hair, etc.), by leveraging the power of a pre-trained text-to-image diffusion model, *e.g.*, Stable Diffusion [61]. Specifically, we modify and fine-tune the text-to-image model [3, 90] to enable conditioning on posed SMPL-X, such that the model creates detailed normal maps according to both the pose information and textual descriptions. Afterward, we pass the generated frontal normal map as guidance to the dual normal map generator to complete dual normal maps, seamlessly connecting text-based generation to our original pipeline.

Trained from posed 3D scans only, *Chupa* is capable of generating various digital avatars from pose and textual information, with realistic, high-fidelity features such as wrinkles and large varieties in human identity and clothing. We evaluate our method through established benchmarks along with a perceptual study and show that our method outperforms the previous baseline. In summary, our contributions are:

- A 3D generation pipeline that directly leverages the 2D image generation capability of diffusion models towards 3D reconstruction.
- A diffusion-based normal map generation and refinement strategy for view-consistent normal maps, targeted for 3D generation.
- A method to effectively allow text-based 3D full-body digital avatar creation, providing an intuitive scenario for digital avatar creation.

2. Related Work

3D Generative Models. Leveraging the success of generative models in producing realistic 2D images [13, 16, 17, 21, 32–34], several efforts have been made to build 3D generative models from 2D datasets while ensuring view consistency [6, 7, 22, 52]. To achieve this, 3D neural implicit representation [48, 53, 82] is employed to represent 3D targets, along with volume rendering to project the 3D scenes into 2D images [6, 7, 22, 52]. While early methods in this direction were mainly focused on rigid objects [6, 50, 69] or human faces [7, 22, 52], recent work has extended to human bodies by using LBS-based canonicalization [8] with SMPL to handle articulated pose changes [2, 27, 93]. However, these approaches suffer from low-quality 3D outputs and high computational costs due to the volume rendering.

Other methods [11, 45] utilized SMPL models with latent codes to represent clothing information. However, these

methods tend to be limited in geometric detail. gDNA [9] was the first generative model-based approach along with a neural implicit representation [55] to create diverse 3D humans with varying identities, poses, and clothing. gDNA further leverages the adversarial loss [21] to generate detailed surface normals. However, the adversarial loss made the model susceptible to mode collapse, which leads to unnatural stochastic details. In contrast, our approach is based on diffusion probabilistic models, which alleviates the mode collapsing issue while producing state-of-the-art quality.

3D Human Reconstruction. The reconstruction of 3D humans has been a long-standing problem in the field of 3D computer vision. Traditional multi-view approaches tended to rely on calibrated multi-camera systems [4, 12, 18, 20, 28, 30, 46, 76, 78]. Several 3D parametric human body models [1, 29, 43, 88] have been presented to represent the shape and pose variation of humans through parametric control, and they are widely used in human pose estimation [31, 38, 62]. Building upon such parametric models, single image-based 3D clothed human reconstruction methods with implicit 3D representation [67, 68] show outstanding results with high-frequency details. Such models, however, tend to show disembodied or broken limbs for unseen poses due to the lack of topological prior. To address the problem, recent works [86, 95] combine implicit representation [48] and parametric models [43, 54]. Inspired by sandwich-like approaches [19, 73], ECON [87] exploits front and back normal maps to build partial surfaces through normal integration [5] and stitches them with a mesh from IF-Net [10] and SMPL mesh through poisson surface reconstruction [35, 36]. Our approach achieves realistic 3D human generation via normal map-based mesh optimization with SMPL-X mesh as a prior. Rather than using the parametric model as an implicit guidance [86, 95] or stitching it with separate surfaces [87], we directly deform the SMPL-X mesh to be consistent with the input normal maps, using a differentiable rasterizer [39].

Diffusion Models. Diffusion Probabilistic Models [74] are a group of generative models that have achieved state-of-the-art results in perceptual image quality and mode coverage [13, 26, 44, 64, 66, 75]. Recent diffusion models for text-to-image generation [49, 59, 61, 65] have demonstrated the ability to produce high-quality images based on textual input. Among them, Rombach et al. [61] enhances the efficiency of diffusion models by operating in a latent space that has a lower dimension than the image space while being perceptually equivalent. We list details of the inner workings of the diffusion models in the supplementary material.

Previous methods [56, 79, 83, 89] focused on text-to-shape tasks, where the output is a small 3D object lacking photorealistic quality. Among such methods, 3DiM [83] presents view-consistent generation through stochastic con-

ditioning but is limited to expressing 3D objects in a 128 resolution. DiffuStereo [70] was one of the first methods to achieve high-quality 3D human reconstruction through diffusion models, but the usage of diffusion models was limited to refining details, while ours better utilizes the generation capability and mode coverage in generating diverse 3D models. Other work such as Rodin [80] also uses textual conditions to generate human 3D models, but are limited to the upper body, being unable to represent various human poses.

3. Method

Our model is capable of generating 3D full body human models by conditioning on a front normal map rendered from a SMPL-X [43, 54] mesh \mathcal{M} , which provides pose information, and an optional textual description that includes other identity-related information. The resulting 3D clothed human models display realistic details, while maintaining consistency to the input pose and textual description.

Conditioned on the normal map rendered from SMPL-X mesh, we first utilize a diffusion-based generative model to create full body normal maps for both frontal (observed) and backside (occluded) regions (Sec. 3.1). We then employ a normal map-based mesh optimization method inspired by NDS [84] to deform the posed SMPL-X mesh into a detailed human mesh (Sec. 3.2). To enhance the quality of our mesh, we render the normal maps from the resulting human mesh at multiple viewpoints and refine them through a diffusion-based resampling strategy [47], where we use separate diffusion models for the full body and facial regions (Sec. 3.3). The refined normal maps are subsequently used as inputs to our mesh optimization method, creating a high-quality 3D clothed digital avatar. Our pipeline also accepts additional text information to further control the identity of the digital avatar using a text-to-image diffusion model [61] (Sec. 3.4). Fig. 2 shows the overall pipeline of our method.

3.1. Dual Normal Map Generation

Following the intuition of “sandwich-like” approaches for single image-based 3D human reconstruction [19, 73, 87], we generate both the frontal and backside normal map ($\mathbf{x}^F, \mathbf{x}^B$) of clothed humans, dubbed *dual* normal maps, with the front-view SMPL-X normal map $c_N(\beta, \theta)$ as a pose condition, where β, θ are the shape parameters and pose parameters of SMPL-X, respectively. We demonstrate that dual normal maps have sufficient information to generate plausible 3D humans with our normal map-based mesh reconstruction method. By generating dual normal maps, we can mitigate the difficulty and computational cost of directly generating 3D representation (*e.g.*, voxels, point clouds, etc.) or multi-view consistent 2D representation (*e.g.*, RGB images, normal maps, etc.). Since dual normal maps can be represented as images, we can exploit a diffusion model renowned for its image generation capability. We employ a

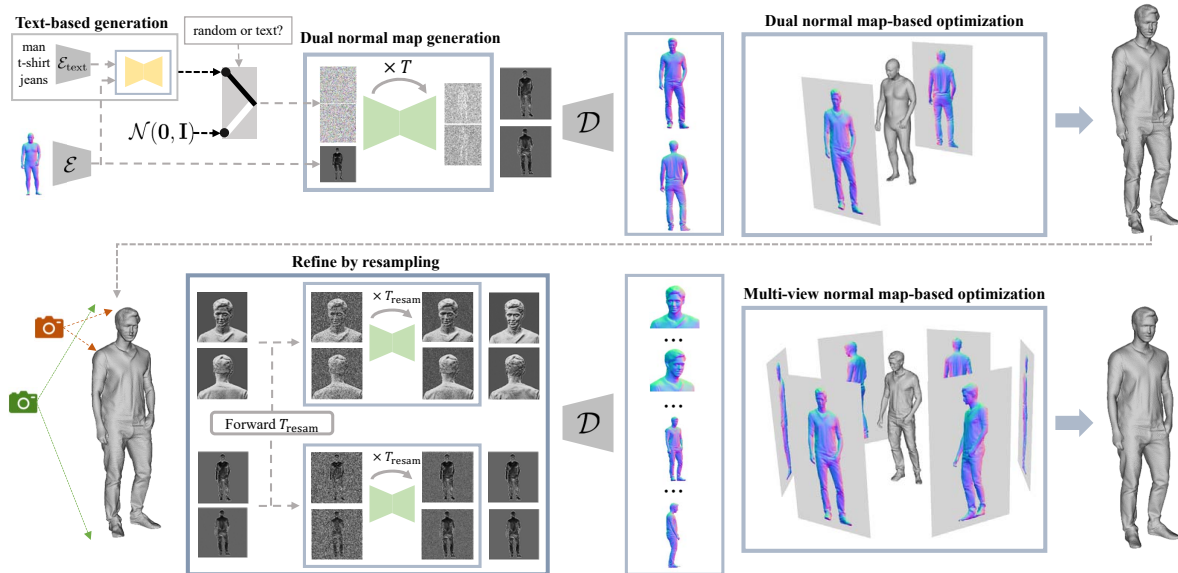


Figure 2. **Overview.** Chupa takes a posed SMPL-X mesh \mathcal{M} and its front normal map \mathbf{c}_N as input. At the first stage, Chupa generates frontal and backside clothed normal maps, $\mathbf{x}^F, \mathbf{x}^B$, conditioned on \mathbf{c}_N . These normals are then used as a reference to “carve” \mathcal{M} through our normal map-based mesh optimization process. To further increase the quality, we separately refine the multi-view normal maps rendered from the full body and facial regions through a resampling procedure and perform the second optimization to create $\mathcal{M}_{\text{final}}$. Our pipeline can also support identity control through a text description by leveraging the power of a text-to-image generation model.

latent diffusion model [61] and adapt it to generate the dual normal maps. Note that we can control the body shape and pose of the generated dual normal maps by changing β, θ with the SMPL-X normal map $\mathbf{c}_N(\beta, \theta)$ as a condition.

Following the latent diffusion model [61], we first train a vector-quantized autoencoder (\mathcal{E}, \mathcal{D}) [15, 77] to support normal maps with alpha channels which enable getting foreground mask of generated normal maps easily. Specifically, given a normal map (color-coded as RGB) with alpha channel $\mathbf{x} \in \mathbb{R}^{H \times W \times 4}$, the encoder \mathcal{E} encodes \mathbf{x} into the latent representation $\mathbf{z} \in \mathbb{R}^{h \times w \times 4}$, and the decoder \mathcal{D} reconstructs a normal map back from the latent \mathbf{z} . We train our autoencoder based on rendered normal maps from views with different yaw angles so that the autoencoder efficiently encodes these normal maps into a perceptually equivalent latent space, i.e., $\mathbf{z}^F = \mathcal{E}(\mathbf{x}^F)$ and $\mathbf{z}^B = \mathcal{E}(\mathbf{x}^B)$. For simultaneous generation, we concatenate the two latent codes \mathbf{z}^F and \mathbf{z}^B into a latent code \mathbf{z} and treat it as an 8-channel image.

During training, the latent code \mathbf{z} is perturbed by the forward diffusion process according to a timestep t , producing a noisy latent code \mathbf{z}_t . The diffusion model ϵ_θ then learns to predict the perturbed noise ϵ of \mathbf{z}_t , given the SMPL-X normal map condition $\mathbf{c}_N(\beta, \theta) \in \mathbb{R}^{H \times W \times 4}$, which is also encoded into $\mathcal{E}(\mathbf{c}_N) \in \mathbb{R}^{h \times w \times 4}$ and concatenated with \mathbf{z}_t channelwise. The corresponding objective becomes

$$L_{\text{dual}} = \mathbb{E}_{\mathbf{x}^F, \mathbf{x}^B, \mathbf{c}_N, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), t} [\|\epsilon - \epsilon_\theta(\mathbf{z}_t^F, \mathbf{z}_t^B, t, \mathcal{E}(\mathbf{c}_N))\|_2^2]. \quad (1)$$

At inference time, we start from the Gaussian noise $\mathbf{z}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and iteratively sample from the previous step until \mathbf{z}_0 , then we decode \mathbf{z}_0 to get the final frontal and backside normal maps. We use classifier-free guidance [24] to boost the sample quality during conditional generation. To enable classifier-free guidance, we randomly assign blank latent embeddings to the conditional image \mathbf{c}_N with 10% probability during training. Then, for each inference step, we use the following modification to predict the denoised latent code:

$$\hat{\epsilon}_\theta(\mathbf{z}_t, t, \mathcal{E}(\mathbf{c}_N)) = \lambda \epsilon_\theta(\mathbf{z}_t, t, \mathcal{E}(\mathbf{c}_N)) + (1 - \lambda) \epsilon_\theta(\mathbf{z}_t, t), \quad (2)$$

where λ specifies the guidance strength that can be controlled during inference, and $\epsilon_\theta(\mathbf{z}_t, t, \mathcal{E}(\mathbf{c}_N))$ and $\epsilon_\theta(\mathbf{z}_t, t)$ each corresponds to the conditional and unconditional predictions. In Fig. 3, our simultaneous dual generation scheme shows that the generated frontal and backside normal maps are more consistent, compared to separate generation.

3.2. Mesh Reconstruction with Front/Back Normals

Given the initial posed SMPL-X mesh $\mathcal{M}(\beta, \theta)$ and the generated clothed normal maps ($\mathbf{x}^F, \mathbf{x}^B$), we deform the initial mesh into a detailed 3D human mesh through iterative optimization. Our mesh reconstruction method is motivated by Neural Deferred Shading (NDS) [84], which reconstructs geometry from multi-view RGB images using a differentiable rasterizer and neural shader. Unlike NDS, we remove the neural shader as the generated normal maps provide su-

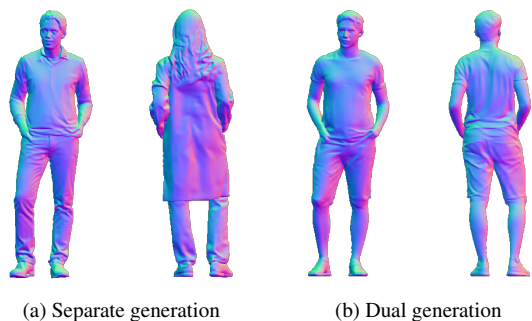


Figure 3. **Separate generation vs. Dual generation.** Comparison between (a) separate sampling for frontal/backside normal maps and (b) our dual sampling. When generated separately, attributes of two normal maps likely differ. However, generating the dual normal maps at once ensures the maps share the same semantics.

pervision for geometry, and directly optimize the 3D geometry by comparing the normal maps with the geometry buffers rendered from a differentiable rasterizer [39]. In general, mesh reconstruction via the two normal maps is an ill-posed problem due to the depth ambiguity. Using SMPL-X mesh as an initial mesh, which is a strong geometric prior, and introducing a novel side loss L_{sides} for regularizing side-views, we can reconstruct plausible 3D geometry of humans while mitigating the difficulty of generating multi-view consistent images at once. Our total objective is defined as

$$L = \lambda_{\text{normal}} L_{\text{normal}} + \lambda_{\text{mask}} L_{\text{mask}} + \lambda_{\text{sides}} L_{\text{sides}} + \lambda_{\text{laplacian}} L_{\text{laplacian}} + \lambda_{\text{normal}}^{\text{reg}} L_{\text{normal}}^{\text{reg}}. \quad (3)$$

Normal map loss. We minimize the difference between the input normal maps ($\mathbf{x}^F, \mathbf{x}^B$) and the normal maps rendered from the front/back views of the human mesh ($\mathbf{N}^F, \mathbf{N}^B$) through a L_1 loss, denoted as L_{normal} . We also minimize the discrepancy between the mask of the normal maps through a L_2 loss, L_{mask} , to match the silhouette of the mesh. Note that we can acquire the masks of the generated normal maps by a simple thresholding on the alpha channel.

Side loss. Since our initial 3D reconstruction is based on frontal/backside normal maps, the left/rightside regions of the human body tend to contain depth ambiguity [57]. We therefore introduce a novel side loss, which ensures that the body masks rendered from the side views ($\hat{\mathbf{M}}_{\text{left}}, \hat{\mathbf{M}}_{\text{right}}$) are not shrunk into the side views of the initial SMPL-X mesh ($\mathbf{M}_{\text{left}}^{\text{smpl}}, \mathbf{M}_{\text{right}}^{\text{smpl}}$). The loss function becomes

$$L_{\text{sides}} = \sum_{\mathbf{M}_{\text{view}}^{\text{smpl}}[h,w]=1} \|\mathbf{M}_{\text{view}}^{\text{smpl}}[h,w] - \hat{\mathbf{M}}_{\text{view}}[h,w]\|_2^2, \quad (4)$$

where $[h,w]$ denotes indexing with the pixel (h,w) of the mask $\mathbf{M} \in \mathbb{R}^{H \times W}$ and $\text{view} \in \{\text{left}, \text{right}\}$. Even though




Figure 4. **Body Resampling.** The initial 3D mesh displays undesired visual artifacts, such as unnatural cloth wrinkles and depth misprediction. By resampling, those artifacts are moderated to produce more natural results.

we can mitigate the problem to some extent with the 3D prior from initial SMPL-X, we further prevent the optimized mesh from having unrealistic side-views.

Geometric regularization. As noted by NDS [84], optimizing the mesh based on only the aforementioned loss terms can lead to degenerated mesh due to unconstrained vertex movement. To overcome this issue, we use geometric regularization terms following NDS [84]. Given a matrix $\mathbf{V} \in \mathbb{R}^{n \times 3}$ with vertex positions of mesh \mathcal{M} as rows, the Laplacian term is defined as $L_{\text{laplacian}} = \frac{1}{n} \sum_{i=1}^n \|\delta_i\|_2^2$, where $\delta_i = (LV)_i \in \mathbb{R}^3$ are the differential coordinates of vertex i with the graph Laplacian L . Since the differential coordinates are the sum of positional difference between its neighbors, minimizing this loss leads to a smoother mesh. We also introduce a normal consistency term, defined as $L_{\text{normal}}^{\text{reg}} = \frac{1}{|\mathcal{F}|} \sum_{(i,j) \in \mathcal{F}} (1 - \mathbf{n}_i \cdot \mathbf{n}_j)^2$, where \mathcal{F} is the set of mesh face pairs with a shared edge and $\mathbf{n}_i \in \mathbb{R}^3$ is the normal of triangle i . Minimizing the cosine similarity between face normals of neighbors encourages further smoothness.

3.3. Refine by Resampling

Resampling multi-view normal maps. After the initial mesh reconstruction, we can further improve the mesh while we already have plausible one. We refine the 3D human mesh by refining the rendered multi-view normal maps of the reconstructed mesh without losing view consistency. The refined maps are then used as inputs to the 3D reconstruction pipeline, creating an improved, realistic 3D human mesh.

Our pipeline is inspired by SDEdit [47], which proposes an image translation method by progressively denoising a noise-perturbed image. The amount of noise perturbation is decided by timestep $0 < t_0 < 1$, and as t_0 gets closer to 0, the operation focuses on editing the finer details. We repeat




Figure 5. **Face close-up resampling.** Both images are aligned according to the SMPL-X vertices for the facial region. We can observe that the perceptibility of the faces is improved.

this process by K times to improve fidelity without harming the original information. To preserve the original structure while adjusting any unrealistic information, we set $t_0 = 0.02$ and $K = 2$, which we empirically found to be sufficient.

In practice, we first render a collection of n -view normal maps $\{\mathbf{I}^1, \mathbf{I}^2, \dots, \mathbf{I}^n\}$ by evenly rotating the yaw camera angle around the 3D mesh. For refinement, we use the same dual normal map generation model in Sec. 3.1, which uses the normal map of posed SMPL-X as spatial guidance. We pair the rendered normal maps so that each pair is rendered from the backside of one another, and use the SMPL-X normal map corresponding to the frontal normal map as the condition to the diffusion model. This perturb-and-denoise process, which we call *resampling*, drives the normal maps rendered from the optimized mesh into the distribution of normal maps rendered from training 3D scans on which our diffusion model is trained, thus the normal maps become more realistic without losing overall semantics. Once the resampling is complete, we pass the refined normal maps as inputs to the 3D reconstruction stage (Sec. 3.2) to produce a refined 3D human model. Fig. 4 shows that our resampling-based refinement produces more natural details.

Facial resampling. We enhance the facial details of the optimized mesh by refining the normal maps rendered from the facial regions of the mesh. We train a latent diffusion model which shares the same architecture of the dual normal map generation model in Sec. 3.1, but trained on normal maps with face close-up. The close-up is done for the head vertices of SMPL-X based on the pre-defined part segmentation [54]. With the face close-up views, we can render facial regions of 3D scans and aligned SMPL-X mesh.

Given the aligned facial normal maps, we can train the diffusion model which generates the frontal and backside facial normal maps with facial normal maps of SMPL-X as a condition. We then apply the same resampling technique used for the full body to refine the multi-view facial normal maps rendered from the optimized mesh. Fig. 5 shows how




Figure 6. **Text-based normal map generation.** Note that our model is capable of generating a normal map consistent in gender, clothing, and hair style¹. Moreover, our guided generation method can create a view-consistent back normal map from the initial frontal map, making it possible to use it for our original pipeline.

the facial region is perceptually refined without harming the original structure. Unlike the method of Frühstück et al. [16], which performs offline optimization to blend a full body image and face image, we just do the normal map-based optimization (Sec. 3.2) with refined normal maps of both body and face, which aggregates the refined normal maps directly in 3D to generate a 3D human mesh with better details.

3.4. Text-guided Normal Map Generation

In addition to the main, pose-conditional 3D generation pipeline, we also include an optional pose-and-text conditional pipeline to further control the identity of the resulting human mesh. To generate 3D human mesh based on a textual description, we adopt a powerful text-to-image diffusion model, *e.g.*, Stable Diffusion [61], and fine-tune its weights to generate normal maps that are consistent to the text description and the posed SMPL-X normal map.

As the method of Wang et al. [81] displayed the effectiveness of fine-tuning large diffusion models for image translation tasks, we initialize the weights of our model based on a pre-trained Stable Diffusion checkpoint, leveraging its renowned generation capabilities. Following previous works [3, 90], we add additional input channels to the first layer of the U-Net [63] and initialize their weights to zero. We also use the same text conditioning based on a pre-trained CLIP model [58].

As shown in Fig. 6, our model supports the generation of detailed normal maps based on the textual description and the posed SMPL-X. Our method is the first method to support text-based full-body normal map generation by basing on Stable Diffusion.

Frontal normal map-guided generation. To get *dual* normal maps based on the frontal normal map generated from

¹For text, we used "girl, long hair, dress"

the text-based normal map generation model, we follow the intuitions of Repaint [44]. Since we already know and want to preserve the frontal shape, the goal here is to predict the unknown backside normal map, based on the frontal normal map. For each inference step, we sample the intermediate frontal latent code \mathbf{z}_t^F from the original latent \mathbf{z}^F at any timestep t , since the diffusion process is defined by a Gaussian Markov chain. In contrast, we sample the unknown, intermediate backside latent code \mathbf{z}_t^B through reverse diffusion, which is concatenated channel-wise to \mathbf{z}_t^F . Since we consider both \mathbf{z}_t^F and \mathbf{z}_t^B as a single, 8-channel latent code, the diffusion model leverages the context of the known frontal normal map while generating the unknown backside normal map, making this a *channel-wise inpainting approach*. Fig. 6 shows that our approach helps to generate backside normal maps that match the original frontal map. Through frontal normal map-guided dual normal map generation, we can seamlessly connect the generative powers of a text-to-image model with our main pipeline.

4. Experiments

In this section, we validate Chupa’s effectiveness in generating realistic 3D humans. We first compare Chupa with the previous state-of-the-art through an image quality metric and a perceptual user study. We also conduct ablation studies to illustrate the effectiveness of each part of our pipeline. Fig. 7 shows comparison of generated results from our method and the baseline [9].

Datasets. We train and test our model with Renderpeople [60] and THuman 2.0 [91] dataset, which consists of 500, 526 scans with various identities and clothing. We split both datasets with a 9:1 ratio for train/test split. For training, we render 36 multi-view normal maps of the train split scans with rotation of 10° yaw interval. We follow ICON [86] for rendering pipeline, originally from MonoPort [41], both for body and face. For rendering normal maps of facial regions, we use the pre-defined part segmentation label of SMPL-X [54] to find head vertices of fitted SMPL-X. Then, we render the facial region of 3D scans and fitted SMPL-X mesh with a weak perspective camera for rendering the head vertices of SMPL-X mesh with close-up. To create text pairs from normal maps for Stable Diffusion fine-tuning, we adopt an off-the-shelf image tagger model [40] based on ViT [14].

Baseline. We compare our method with gDNA [9] as a baseline. gDNA is the state-of-the-art method to generate 3D human mesh with given SMPL-X parameter β , Θ and randomly sampled shape latent code $\mathbf{z}_{\text{shape}}$ and detail latent code $\mathbf{z}_{\text{detail}}$ from its learned latent space.

4.1. Implementation Details

Autoencoder model training. Before training the full-body dual generation model, we trained the autoencoder model $(\mathcal{E}, \mathcal{D})$ for 1,000 epochs on $4 \times$ NVIDIA A100 GPUs following the original implementation [15]. We used a VQ-regularized autoencoder with downsampling factor $f = 4$ and channel dimension $c = 4$ such that, given a full-body normal map image with alpha transparency ($\mathbf{c}_N \in \mathbb{R}^{512 \times 512 \times 4}$), the encoder transforms the image to a latent code with 4 channels ($\mathcal{E}(\mathbf{c}_N) \in \mathbb{R}^{128 \times 128 \times 4}$), and the decoder reconstructs the image from the latent code. For training, we used the full-body normal map datasets, following the same preprocessing listed in the main paper. We used the pretrained weights for the autoencoders of facial generation models (Sec. 3.3) and text-based generation models (Sec. 3.4) provided by the original paper [61]. For the facial generation model, we used a VQ-regularized autoencoder with downsampling factor $f = 4$ and channel dimension $c = 3$. For textual generation models, we used a KL-regularized autoencoder with downsampling factor $f = 8$ and channel dimension $c = 4$. All autoencoders were frozen during diffusion training.

U-net. We adapt the U-Net [63] architecture for our diffusion models to support our dual-generation scheme. Specifically, we follow the approach of Dhariwal and Nichol [13] to further improve the sampling quality and set the input channels from 6 to 12, and the output channels from 3 to 8. By utilizing the concatenation of two input images (front and back) with the SMPL latent code $\mathcal{E}(\mathbf{c}_N)$ for conditioning, we can treat them as a single input. As a result, we can obtain two spatially aligned images for both views at the same time. For the facial generation models, we set the input channels to 9 and output channels to 6, since we used 3-channel for facial normal maps.

Dual normal map generator training. We train our full-body dual normal map generation model for 500 epochs with batch size 16 on $4 \times$ NVIDIA A100 GPUs. We set the total timesteps $T = 100$ with a linear variance schedule. During inference, we use the same 512×512 resolution and generate results with the same denoising steps used during training. We trained the facial generation model for 300 epochs with the same training settings.

Text-guided normal map generator training. We train our text-based normal map generator for 1,000 epochs on $4 \times$ NVIDIA A100 GPUs. We train at a 512×512 resolution with a total batch size of 64. We initialize our model from the EMA weights of the Stable Diffusion [61] checkpoints and adopt other training settings from the public Stable Diffusion code base. After inference, we used a thresholding operation




Figure 7. **Generation Comparison.** We display the visual comparisons between gDNA [9] and Chupa with the same SMPL input. Note that gDNA tends to amplify the unnatural artifacts from its coarse stage to the fine stage, while our results produce more natural results.

Table 1. **Quantitative Evaluation.** We report two types of FID scores for the test split of Renderpeople and Thuman 2.0.

Method	FID _{normal} ↓	FID _{shade} ↓
gDNA _{coarse} [9]	53.74	68.14
gDNA _{fine} [9]	36.43	45.57
Ours	21.90	36.58

on the 3rd channel of the image to create a transparency map before the dual generation stage.

4.2. Quantitative Results

We conduct a quantitative evaluation of the quality of generated meshes, based on given SMPL-X parameters. We generated 3D human meshes with SMPL-X parameters fitted to 103 test scans, *i.e.* 50 from Renderpeople and 53 from THuman 2.0, for both our method and gDNA [9]. Following the previous work [9, 71, 94], we render normal maps [9] and shading-images [72, 94] of groundtruth scans and generated meshes into 18 views with 20° yaw interval, and compute FID score with them, which denoted as FID_{normal}

and FID_{shade} respectively. Tab. 1 shows that our method achieves lower FID for both images than the baseline.

4.3. User Preference

We carry out a perceptual study over 78 subjects asking about their preference between the meshes from our method and gDNA. We randomly select 40 from a set of SMPL-X parameters fitted to 103 test scans. We randomly generate meshes based on them with our method and gDNA, and render shading-images in 3 views, 0°, 120°, 240° for full body images and 0°, 40°, -40° for face images. Note that we use the narrower field-of-view for better comparing facial details. Tab. 1 shows that the users preferred meshes from our method both for full-body and face images. We present more details in the supplementary material.

4.4. Ablation Study

We validate the building blocks of our pipeline through an ablation study. The evaluation is based on the same test split. The results are summarized in Tab. 3.

Table 2. **User preference.** We carry out a perceptual study asking 78 subjects to choose a more realistic one between ours and gDNA_{fine}.

Method	Body	Face	Total
gDNA _{fine}	20.89%	18.7%	19.78%
Ours	79.11%	81.3%	80.22%

Table 3. **Ablation study.** We do ablation study over our key components. We report FID_{normal} score.

dual.	L_{sides}	refine _{body}	refine _{face}	FID _{normal} ↓
				30.55
✓				26.31
✓	✓			25.50
✓	✓	✓		22.61
✓	✓	✓	✓	21.90

Front/Back normal map generation. To validate the effectiveness of our dual normal map generation method, we separately generate frontal and backside normal maps with the SMPL normal map in the corresponding view. Due to the randomness of the diffusion model, we cannot guarantee the separately generated frontal and backward normal maps are consistent (Fig. 3), which leads to performance loss.

Side loss. With the sidewise loss L_{sides} from Eq. (4), we enforce our mesh to keep better alignment with the SMPL-X prior during mesh optimization (Sec. 3.2). Fig. 8 shows the effect of utilizing L_{sides} . The first column shows the side-view normal map rendered from the mesh optimized with dual normal maps. The second column shows the same side-view normal map but overlapped with the side-view of the corresponding SMPL-X. The third column shows the normal maps after resampling (Sec. 3.3). Fig. 8a shows that the optimized mesh without L_{sides} has worse alignment with SMPL-X mesh, which leads to the artifacts on resampling results. Tab. 3 demonstrates the inclusion of L_{sides} leads to lower FID scores, indicating its effectiveness.

Refinement. To validate the effectiveness of our refinement method (Sec. 3.3), we compare 3D generation results only optimized by front/back normal maps and the results refined by body refinement and additional face refinement. Fig. 4 and Fig. 5 show that our refinement methods lead to more realistic generation results. As expected, Tab. 3 shows that our face refinement method further reduces FID.

5. Discussion

We propose Chupa, a powerful 3D generation pipeline for a large variety of dressed 3D high-quality digital avatars. By combining diffusion models for normal map generation




Figure 8. **Side loss.** We present the side-view normal maps of the optimized mesh (left), the normal maps overlapped on the SMPL-X normal maps (middle), and the normal maps after resampling (right). Without L_{sides} , the alignment between the SMPL-X mesh and the optimized mesh becomes worse, leading to artifacts on the resampling result. (Note that the blue channel of the overlapped SMPL-X normal map is flipped for visualization purposes.)

with a normal map-based mesh reconstruction method, our pipeline enables the creation of realistic 3D avatars with high levels of stochastic details. We also allow the creation of 3D humans from both pose and textual information, providing an intuitive method of digital avatar creation.

We note that while our pipeline can support text conditioning without losing visual quality, several elements that can be generated from the initial text-to-image model (e.g., bracelet, necklace, glasses) tend to be lost during the later stage of the pipeline and cannot be expressed at the final 3D model. For future work, we look forward to creating digital avatars with photorealistic textures and devising novel strategies for creating animations from our digital avatars.

Acknowledgements. This work was supported by Naver Webtoon. The work of SNU members was also supported by SNU Creative-Pioneering Researchers Program, NRF grant funded by the Korean government (MSIT) (No. 2022R1A2C2092724), and IITP grant funded by the Korean government (MSIT) (No.2022-0-00156 and No.2021-0-01343). H. Joo is the corresponding author.

References

- [1] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. Scape: shape completion and animation of people. In *ACM TOG*, pages 408–416. 2005. 3
- [2] A. W. Bergman, P. Kellnhöfer, W. Yifan, E. R. Chan, D. B. Lindell, and G. Wetzstein. Generative neural articulated radiance fields. In *NeurIPS*, 2022. 2
- [3] T. Brooks, A. Holynski, and A. A. Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proc. CVPR*, 2023. 2, 6
- [4] T. Brox, B. Rosenhahn, J. Gall, and D. Cremers. Combined region and motion-based 3D tracking of rigid and articulated objects. *IEEE TPAMI*, 2010. 3
- [5] X. Cao, H. Santo, B. Shi, F. Okura, and Y. Matsushita. Bilateral normal integration. In *Proc. ECCV*, 2022. 3
- [6] E. R. Chan, M. Monteiro, P. Kellnhöfer, J. Wu, and G. Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proc. CVPR*, 2021. 2
- [7] E. R. Chan, C. Z. Lin, M. A. Chan, K. Nagano, B. Pan, S. De Mello, O. Gallo, L. J. Guibas, J. Tremblay, S. Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *Proc. CVPR*, 2022. 2
- [8] X. Chen, Y. Zheng, M. J. Black, O. Hilliges, and A. Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *Proc. ICCV*, 2021. 2
- [9] X. Chen, T. Jiang, J. Song, J. Yang, M. J. Black, A. Geiger, and O. Hilliges. gdn: Towards generative detailed neural avatars. In *Proc. CVPR*, 2022. 2, 3, 7, 8
- [10] J. Chibane, T. Alldieck, and G. Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion. In *Proc. CVPR*, 2020. 3
- [11] E. Corona, A. Pumarola, G. Alenya, G. Pons-Moll, and F. Moreno-Noguer. Smplicit: Topology-aware generative model for clothed people. In *Proc. CVPR*, 2021. 2
- [12] E. de Aguiar, C. Stoll, C. Theobalt, N. Ahmed, H.-P. Seidel, and S. Thrun. Performance capture from sparse multi-view video. *SIGGRAPH*, 2008. 3
- [13] P. Dhariwal and A. Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, 2021. 2, 3, 7
- [14] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proc. ICLR*, 2021. 7
- [15] P. Esser, R. Rombach, and B. Ommer. Taming transformers for high-resolution image synthesis. In *Proc. CVPR*, 2020. 4, 7
- [16] A. Frühstück, K. K. Singh, E. Shechtman, N. J. Mitra, P. Wonka, and J. Lu. Insetgan for full-body image generation. In *Proc. CVPR*, 2022. 2, 6
- [17] J. Fu, S. Li, Y. Jiang, K.-Y. Lin, C. Qian, C. C. Loy, W. Wu, and Z. Liu. Stylegan-human: A data-centric odyssey of human generation. In *Proc. ECCV*, 2022. 2
- [18] Y. Furukawa and J. Ponce. Dense 3d motion capture from synchronized video streams. *CVPR*, 2008. 3
- [19] V. Gabeur, J.-S. Franco, X. Martin, C. Schmid, and G. Rogez. Moulding humans: Non-parametric 3d human shape estimation from single images. In *Proc. ICCV*, 2019. 2, 3
- [20] J. Gall, C. Stoll, E. De Aguiar, C. Theobalt, B. Rosenhahn, and H.-P. Seidel. Motion capture using joint skeleton tracking and surface estimation. In *Proc. CVPR*, 2009. 3
- [21] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 2, 3
- [22] J. Gu, L. Liu, P. Wang, and C. Theobalt. StyleNeRF: A style-based 3d aware generator for high-resolution image synthesis. In *Proc. ICLR*, 2022. 2
- [23] K. Guo, P. Lincoln, P. Davidson, J. Busch, X. Yu, M. Whalen, G. Harvey, S. Orts-Escolano, R. Pandey, J. Dourgarian, et al. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM TOG*, 38(6):1–19, 2019. 2
- [24] J. Ho. Classifier-free diffusion guidance. *ArXiv*, abs/2207.12598, 2022. 4
- [25] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2
- [26] J. Ho, C. Saharia, W. Chan, D. J. Fleet, M. Norouzi, and T. Salimans. Cascaded diffusion models for high fidelity image generation. *JMLR*, 23(47):1–33, 2022. 3
- [27] F. Hong, Z. Chen, Y. LAN, L. Pan, and Z. Liu. EVA3d: Compositional 3d human generation from 2d image collections. In *Proc. ICLR*, 2023. 2
- [28] H. Joo, T. Simon, X. Li, H. Liu, L. Tan, L. Gui, S. Banerjee, T. Godisart, B. Nabbe, I. Matthews, et al. Panoptic studio: A massively multiview system for social interaction capture. *IEEE TPAMI*, 41(1):190–204, 2017. 3
- [29] H. Joo, T. Simon, and Y. Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *Proc. CVPR*, 2018. 3
- [30] T. Kanade, P. Rander, and P. Narayanan. Virtualized reality: Constructing virtual worlds from real scenes. *IEEE Multimedia*, 1997. 3
- [31] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik. End-to-end recovery of human shape and pose. In *Proc. CVPR*, 2018. 3
- [32] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In *Proc. CVPR*, 2019. 2
- [33] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020.
- [34] T. Karras, M. Aittala, S. Laine, E. Härkönen, J. Hellsten, J. Lehtinen, and T. Aila. Alias-free generative adversarial networks. In *NeurIPS*, 2021. 2
- [35] M. Kazhdan and H. Hoppe. Screened poisson surface reconstruction. *ACM TOG*, 32(3):1–13, 2013. 3
- [36] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Proc. Eurographics*, 2006. 3
- [37] D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *Proc. ICLR*, 2014. 2
- [38] N. Kolotouros, G. Pavlakos, M. J. Black, and K. Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proc. ICCV*, 2019. 3
- [39] S. Laine, J. Hellsten, T. Karras, Y. Seol, J. Lehtinen, and T. Aila. Modular primitives for high-performance differen-

- tiabile rendering. *ACM TOG*, 39(6):1–14, 2020. 1, 2, 3, 5
- [40] S. Lee. Tagger for automatic1111’s webui. <https://github.com/toriato/stable-diffusion-webui-wd14-tagger>, 2022. 7
- [41] R. Li, Y. Xiu, S. Saito, Z. Huang, K. Olszewski, and H. Li. Monocular real-time volumetric performance capture. In *Proc. ECCV*, 2020. 7
- [42] S. Lombardi, J. Saragih, T. Simon, and Y. Sheikh. Deep appearance models for face rendering. *ACM TOG*, 37(4):68:1–68:13, 2018. 2
- [43] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. Smpl: A skinned multi-person linear model. *ACM TOG*, 34(6):1–16, 2015. 2, 3
- [44] A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and L. Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proc. CVPR*, 2022. 3, 7
- [45] Q. Ma, J. Yang, A. Ranjan, S. Pujades, G. Pons-Moll, S. Tang, and M. J. Black. Learning to dress 3d people in generative clothing. In *Proc. CVPR*, 2020. 2
- [46] T. Matsuyama and T. Takai. Generation, visualization, and editing of 3d video. *3DPVT*, 2002. 3
- [47] C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.-Y. Zhu, and S. Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *Proc. ICLR*, 2021. 2, 3, 5
- [48] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proc. CVPR*, 2019. 2, 3
- [49] A. Q. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *Proc. ICML*, 2022. 2, 3
- [50] M. Niemeyer and A. Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proc. CVPR*, 2021. 2
- [51] M. Niemeyer, L. Mescheder, M. Oechsle, and A. Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proc. CVPR*, 2020. 2
- [52] R. Or-El, X. Luo, M. Shan, E. Shechtman, J. J. Park, and I. Kemelmacher-Shlizerman. Stylesdf: High-resolution 3d-consistent image and geometry generation. In *Proc. CVPR*, 2022. 2
- [53] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proc. CVPR*, 2019. 2
- [54] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. Osman, D. Tzionas, and M. J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proc. CVPR*, 2019. 2, 3, 6, 7
- [55] S. Peng, M. Niemeyer, L. Mescheder, M. Pollefeys, and A. Geiger. Convolutional occupancy networks. In *Proc. ECCV*, 2020. 3
- [56] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *Proc. ICLR*, 2023. 2, 3
- [57] Y. Quéau, J.-D. Durou, and J.-F. Aujol. Normal integration: a survey. *Journal of Mathematical Imaging and Vision*, 60:576–593, 2018. 5
- [58] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *Proc. ICML*, 2021. 6
- [59] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 3
- [60] Renderpeople, 2018. <https://renderpeople.com/3d-people>. 2, 7
- [61] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proc. CVPR*, 2022. 1, 2, 3, 4, 6, 7
- [62] Y. Rong, T. Shiratori, and H. Joo. Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration. In *Proc. CVPR*, 2021. 3
- [63] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Proc. MICCAI*, 2015. 6, 7
- [64] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, and M. Norouzi. Palette: Image-to-image diffusion models. In *Proc. ACM SIGGRAPH*, 2022. 3
- [65] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, R. Gontijo-Lopes, B. K. Ayan, T. Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*, 2022. 3
- [66] C. Saharia, J. Ho, W. Chan, T. Salimans, D. J. Fleet, and M. Norouzi. Image super-resolution via iterative refinement. *IEEE TPAMI*, 45(4):4713–4726, 2023. 3
- [67] S. Saito, Z. Huang, R. Natsume, S. Morishima, A. Kanazawa, and H. Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proc. CVPR*, 2019. 3
- [68] S. Saito, T. Simon, J. Saragih, and H. Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proc. CVPR*, 2020. 3
- [69] K. Schwarz, Y. Liao, M. Niemeyer, and A. Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *NeurIPS*, 2020. 2
- [70] R. Shao, Z. Zheng, H. Zhang, J. Sun, and Y. Liu. Diffustereo: High quality human reconstruction via diffusion-based stereo using sparse cameras. In *Proc. ECCV*, 2022. 3
- [71] J. R. Shue, E. R. Chan, R. Po, Z. Ankner, J. Wu, and G. Wetzstein. 3d neural field generation using triplane diffusion. *arXiv preprint arXiv:2211.16677*, 2022. 8
- [72] J. R. Shue, E. R. Chan, R. Po, Z. Ankner, J. Wu, and G. Wetzstein. 3d neural field generation using triplane diffusion. In *Proc. CVPR*, 2023. 2, 8
- [73] D. Smith, M. Loper, X. Hu, P. Mavroidis, and J. Romero. Facsimile: Fast and accurate scans from an image in less than a second. In *Proc. ICCV*, 2019. 2, 3
- [74] J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proc. ICML*, 2015. 3
- [75] Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. In *NeurIPS*, 2019. 3
- [76] C. Stoll, N. Hasler, J. Gall, H.-P. Seidel, and C. Theobalt. Fast articulated motion tracking using a sums of gaussians body model. *ICCV*, 2011. 3

- [77] A. van den Oord, O. Vinyals, and K. Kavukcuoglu. Neural discrete representation learning. In *NeurIPS*, 2017. 4
- [78] D. Vlastic, I. Baran, W. Matusik, and J. Popović. Articulated mesh animation from multi-view silhouettes. *SIGGRAPH*, 2008. 3
- [79] H. Wang, X. Du, J. Li, R. A. Yeh, and G. Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *Proc. CVPR*, 2023. 2, 3
- [80] T. Wang, B. Zhang, T. Zhang, S. Gu, J. Bao, T. Baltrusaitis, J. Shen, D. Chen, F. Wen, Q. Chen, et al. Rodin: A generative model for sculpting 3d digital avatars using diffusion. *arXiv preprint arXiv:2212.06135*, 2022. 3
- [81] T. Wang, T. Zhang, B. Zhang, H. Ouyang, D. Chen, Q. Chen, and F. Wen. Pretraining is all you need for image-to-image translation. *arXiv preprint arXiv:2205.12952*, 2022. 6
- [82] Z. Wang, S. Wu, W. Xie, M. Chen, and V. A. Prisacariu. Nerf-: Neural radiance fields without known camera parameters. In *arXiv preprint arXiv:2102.07064*, 2021. 2
- [83] D. Watson, W. Chan, R. M. Brualla, J. Ho, A. Tagliasacchi, and M. Norouzi. Novel view synthesis with diffusion models. In *Proc. ICLR*, 2023. 3
- [84] M. Worchel, R. Diaz, W. Hu, O. Schreer, I. Feldmann, and P. Eisert. Multi-view mesh reconstruction with neural deferred shading. In *Proc. CVPR*, 2022. 2, 3, 4, 5
- [85] Z. Xiao, K. Kreis, and A. Vahdat. Tackling the generative learning trilemma with denoising diffusion gans. In *Proc. ICLR*, 2022. 2
- [86] Y. Xiu, J. Yang, D. Tzionas, and M. J. Black. Icon: implicit clothed humans obtained from normals. In *Proc. CVPR*, 2022. 3, 7
- [87] Y. Xiu, J. Yang, X. Cao, D. Tzionas, and M. J. Black. ECON: Explicit Clothed humans Optimized via Normal integration. In *Proc. CVPR*, 2023. 2, 3
- [88] H. Xu, E. G. Bazavan, A. Zanfir, W. T. Freeman, R. Suktankar, and C. Sminchisescu. Ghum & ghuml: Generative 3d human shape and articulated pose models. In *Proc. CVPR*, 2020. 3
- [89] J. Xu, X. Wang, W. Cheng, Y.-P. Cao, Y. Shan, X. Qie, and S. Gao. Dream3d: Zero-shot text-to-3d synthesis using 3d shape prior and text-to-image diffusion models. In *Proc. CVPR*, 2023. 2, 3
- [90] B. Yang, S. Gu, B. Zhang, T. Zhang, X. Chen, X. Sun, D. Chen, and F. Wen. Paint by example: Exemplar-based image editing with diffusion models. In *Proc. CVPR*, 2023. 2, 6
- [91] T. Yu, Z. Zheng, K. Guo, P. Liu, Q. Dai, and Y. Liu. Function4d: Real-time human volumetric capture from very sparse consumer rgbd sensors. In *Proc. CVPR*, 2021. 2, 7
- [92] X. Zeng, A. Vahdat, F. Williams, Z. Gojcic, O. Litany, S. Fidler, and K. Kreis. LION: Latent point diffusion models for 3d shape generation. In *NeurIPS*, 2022. 2
- [93] J. Zhang, Z. Jiang, D. Yang, H. Xu, Y. Shi, G. Song, Z. Xu, X. Wang, and J. Feng. Avatargen: a 3d generative model for animatable human avatars. In *Proc. ECCV*, 2022. 2
- [94] X. Zheng, Y. Liu, P. Wang, and X. Tong. Sdf-stylegan: Implicit sdf-based stylegan for 3d shape generation. In *Comput. Graph. Forum*, 2022. 2, 8
- [95] Z. Zheng, T. Yu, Y. Liu, and Q. Dai. Pamir: Parametric model-

conditioned implicit representation for image-based human reconstruction. *IEEE TPAMI*, 44(6):3170–3184, 2021. 3