
Multi-hypotheses Conditioned Point Cloud Diffusion for 3D Human Reconstruction from Occluded Images

Donghwan Kim Tae-Kyun Kim
KAIST
{kdoh2522, kimtaekyun}@kaist.ac.kr

Abstract

3D human shape reconstruction under severe occlusion due to human-object or human-human interaction is a challenging problem. Parametric models i.e. SMPL(-X), which are based on the statistics across human shapes, can represent whole human body shapes but are limited to minimally-clothed human shapes. Implicit-function-based methods extract features from the parametric models to employ prior knowledge of human bodies and can capture geometric details such as clothing and hair. However, they often struggle to handle misaligned parametric models and inpaint occluded regions given a single RGB image. In this work, we propose a novel pipeline, MHCDIFF, Multi-hypotheses Conditioned Point Cloud **Diffusion**, composed of point cloud diffusion conditioned on probabilistic distributions for pixel-aligned detailed 3D human reconstruction under occlusion. Compared to previous implicit-function-based methods, the point cloud diffusion model can capture the global consistent features to generate the occluded regions, and the denoising process corrects the misaligned SMPL meshes. The core of MHCDIFF is extracting local features from multiple hypothesized SMPL(-X) meshes and aggregating the set of features to condition the diffusion model. In the experiments on CAPE and MultiHuman datasets, the proposed method outperforms various SOTA methods based on SMPL, implicit functions, point cloud diffusion, and their combined, under synthetic and real occlusions. Our code is publicly available at <https://donghwankim0101.github.io/projects/mhcdiff>.

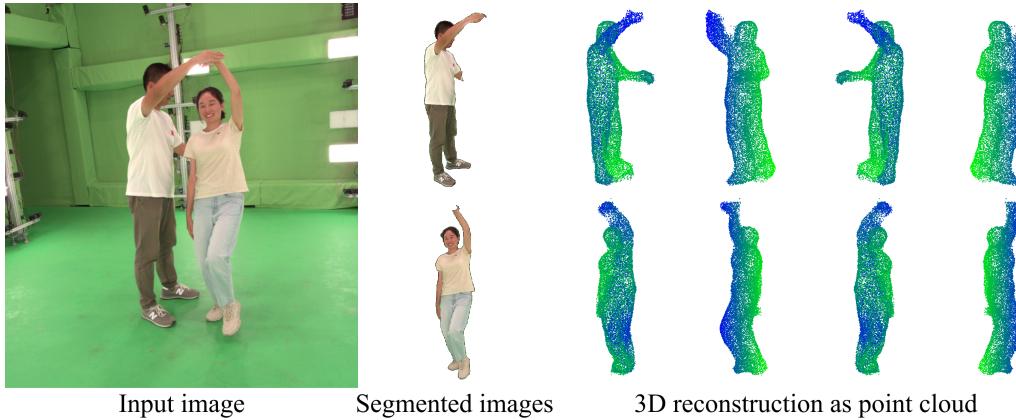


Figure 1: **Image to 3D shape.** From the segmented images, containing occlusion due to interaction, MHCDIFF reconstructs 3D human shapes as point clouds.

1 Introduction

Realistic virtual humans play a significant role in various industries, such as metaverse, tele-presence, and game modeling. However, conventional methods require expensive artist efforts and complex scanning equipments, so they are not readily applicable. A more practical approach is to reconstruct high-fidelity 3D humans from 2D images taken in the wild. This is still an ongoing research task due to its challenges; people wear a wide variety of clothing styles and adopt diverse poses. Furthermore, human-object and human-human interaction, fundamental aspects of daily social life, make it more challenging due to severe occlusions.

Existing 3D human reconstruction methods cannot predict the *pixel-aligned* 3D shapes of humans *robustly* from occluded images. The parametric body models [27, 48, 67, 93, 75] have been widely used to reconstruct 3D human shapes. Several methods [11, 28, 34, 83, 26, 15, 10, 44, 82] predict the parameters of the statistical models and are robust to occlusion because they can be trained on large scale datasets [23, 56] and parametric models are well regularized with human body priors. However, the parametric models lack geometric details like clothing and hair, so these approaches cannot align the results to the subjects with loose clothing. More recently, 3D clothed human reconstruction methods [76, 77, 106, 92, 7, 91, 90, 95, 96], which are based on implicit functions and integrate the human body prior from the 3D body models, *i.e* SMPL [48, 67], present pixel-aligned detail shapes. Despite the impressive advances of the previous methods, they are not robust to occlusion because (1) small misalignment of estimated parametric models ruins the final shapes, (2) the implicit function takes features independently and cannot inpaint the invisible regions with missing image features, and (3) datasets [74, 87, 66] usually consists of segmented full-body images.

To address the aforementioned limitations, we propose MHCDIFF (Multi-hypotheses Conditioned Point Cloud Diffusion). (1) Several existing methods [6, 60, 73, 78, 79, 35, 8, 61, 81, 14] predict multiple SMPL meshes to model uncertainty due to occlusions. The sampled distribution is also important prior knowledge of human motions, but none of the existing work utilizes the distribution for pixel-aligned 3D human reconstruction. We leverage the multi-hypotheses to be robust on the misalignment of each sample. (2) We adopt denoising diffusion probabilistic models (DDPMs) [20] to take global consistent features and generate the invisible regions. Diffusion based methods generate 3D shapes by denoising point clouds [50, 108, 63, 22], latent [101, 62, 36], neural fields [69], 3D Gaussian [84] or meshes [46]. We adopt the unstructured point clouds to project pixel-aligned image features at each diffusion step. (3) Additionally, we synthesize partial body images by random masking [107], augmenting the limited datasets.

Specifically, our goal is *pixel-aligned* and *detailed* 3D human reconstruction in a *robust* manner to occlusion in images. Given a single occluded RGB image, we extract 2D features and generate multiple plausible SMPL hypotheses using an off-the-shelf method [4, 14]. The proposed method, MHCDIFF, Multi-hypotheses Conditioned Point Cloud Diffusion, performs the diffusion process to denoise a randomly-sampled point cloud into a target human shape. To reconstruct a pixel-aligned 3D shape and leverage the human body prior, the diffusion process is conditioned on the projected image feature (Sec. 3) and local features extracted from SMPL (Sec. 4.2). The key of MHCDIFF is a novel conditional diffusion process with multiple hypotheses (Sec. 4.3), which is not sensitive to misaligned SMPL estimation. Given global 2D features and the distribution of hypotheses, the denoising diffusion model can generate the occluded parts (Sec. 4.4)

We train MHCDIFF on randomly masked THuman2.0 dataset [85]. Our experiments on CAPE dataset [53, 68] with synthesized occlusion and MultiHuman dataset [105] with real-world interaction demonstrate that MHCDIFF reconstructs pixel-aligned 3D human shapes robustly to various occlusion ratios and achieves state-of-the-art performance. Our main contributions are as follows:

- We introduce a novel multi-hypotheses conditioning mechanism that effectively captures the distribution of multiple plausible SMPL meshes. It is robust to the noise of each SMPL estimation due to the occlusion of given images. To the best of our knowledge, MHCDIFF is the first work that extends the multi-hypotheses SMPL estimation to pixel-aligned 3D human reconstruction.
- We adopt point cloud diffusion model to capture the global consistent features and inpaint the invisible parts. Unlike the previous implicit function, the misaligned SMPL estimation can be corrected during the denoising process. The point cloud diffusion model also offers detailed human meshes.

- MHCDIFF, trained on synthesized partial body images, outperforms previous methods on occluded and even full-body images.

2 Related Work

2.1 Diffusion models for point clouds

Over the past years, denoising diffusion probabilistic models (DDPMs) [20] have been applied to point clouds. For unconditional generation, Luo et al. [50], Zhou et al. [108] and LION [101] use PointNet [70], Point-Voxel-CNN [47] and latent space, respectively. PointInfinity [22] tackles the quadratic complexity of transformer [89], and generates high-resolution point clouds with a fixed-size latent vector. Otherwise, Point-E [63] is a text-conditioned generation model using CLIP [72] and PDR [51] is a point cloud completion method from partial point clouds. PC² [57], which is the baseline of MHCDIFF, reconstructs the point cloud conditioned on projected image features (please refer to Sec. 3 for more details).

2.2 Explicit-shape-based human reconstruction

Parametric models [27, 48, 67, 93, 75] have been primary representations for 3D human reconstruction. Due to the strength that they capture the statistics across a large corpus of human shapes, a lot of work [11, 28, 34, 83, 26, 15, 10, 44, 82] reconstructs 3D body meshes from an RGB image. To reduce the gaps between the image and parameter space of the statistical models and improve image alignment, they propose intermediate representations or additional supervisions, such as semantic segmentation [64, 94, 33, 100] and keypoints [9, 41]. To model the uncertainty due to occlusions or depth ambiguities, some work proposes multi-hypotheses [6], heatmaps [60], probability density functions [73, 78, 79, 35] or diffusion models [8, 61, 102, 81, 39]. ProPose [14] adopts the matrix Fisher distribution [13, 30] over $\mathcal{SO}(3)$ for the joint rotation conditioned on the von Mises-Fisher distribution [55] for the unit directions of bones, which is not only mathematically correct but also learning friendly (please refer to Sec. 3 for more details). However, these methods are limited to recovering minimally-clothed humans and lack the ability to capture geometric details such as clothing and hair.

Several works aim at modeling geometric details in explicit shapes such as meshes, voxels, depth maps and point clouds. Mesh-based methods [1, 2, 3, 37, 109, 5, 25] model 3D offsets on the vertices of SMPL [48], but they do not generalize on loose clothing such as skirts and dresses. Voxel-based methods [24, 88, 17, 86] reconstruct 3D human shapes in fine-grained voxel representations. However, free-form 3D reconstruction is challenging without prior, and they need high computation costs to output high-resolution 3D shapes. Point-cloud-based methods [52, 99, 54, 19, 85] model point clouds of clothing humans. Han et al. [19] estimate depth maps based on different body parts, and convert the depth maps into point clouds. Tang et al. [85], the most related work, reconstruct 3D humans with point cloud diffusion from an RGB image. First, they convert the estimated SMPL mesh and depth map from the RGB image to point clouds. Conditioned on this point cloud, the conditional diffusion model refines the point cloud. However, they only handle complete images without occlusion and are not robust to misaligned SMPL estimation.

2.3 Implicit-function-based human reconstruction

Implicit-function-based methods regress occupancy fields [58] or signed distance fields (SDF) [65] utilizing Multi-Layer Perceptron (MLP) decoders as implicit functions (IF). PIFu [76] and PIFuHD [77], which are pioneering works, extract pixel-aligned image features for clothed 3D human reconstruction. Later works [106, 92, 7, 91, 90, 40, 95, 103, 104, 96] leverage parametric models or body keypoints as prior information on the human body. They extract global features from voxelized SMPL meshes with a 3D encoder [106, 90] or local features such as signed distances and normals from SMPL meshes [92, 95, 103, 104] or both [7, 96]. The use of global features helps regularize global shapes and ensure consistency and local features help reconstruct local details. However, the global encoder is sensitive to global pose changes of SMPL and decreases the performance given misaligned SMPL estimation due to occlusion. The local features do not contain the global consistent features and cannot inpaint the occluded parts. Wang et al. [90] aim to reconstruct complete 3D

shapes from occluded images by primarily using the generative global encoder with a discriminator, but only assuming the accurate SMPL meshes.

3 Preliminary

PC² [57]. The projection-conditioned point cloud diffusion model is proposed for single-view 3D shape reconstruction. Denoising diffusion probabilistic model [20], which is the foundation of this framework, learns to recurrently transform noise $X_T \sim \mathcal{N}(0, \mathbf{I})$ into a sample from the target data distribution $X_0 \sim q(X_0)$ over a series of steps. In order to learn this denoising process, a neural network is trained $\mathcal{F}_\theta(X_{t-1}|X_t) \approx q(X_{t-1}|X_t)$. To reconstruct geometrically consistent 3D point clouds from single RGB images $I \in \mathbb{R}^{H \times W \times 3}$, 2D feature map $\mathcal{E}(I) \in \mathbb{R}^{h \times w \times c}$ is projected onto the partially denoised points at each step in the diffusion process. Therefore, $\mathcal{F}_\theta(\cdot) : \mathbb{R}^{(3+c)N} \rightarrow \mathbb{R}^{3N}$ is a function that predicts the noise $\epsilon \in \mathbb{R}^{3N}$ from the point cloud $X_t \in \mathbb{R}^{3N}$ and the projected features $X_t^{proj} \in \mathbb{R}^{cN}$, where c is the number of feature channels.

ProPose [14]. Recovering accurate body meshes and 3D joint rotations from single images remains a challenging problem, particularly in cases of severe occlusion, including self-occlusion and occlusion from other subjects or objects. ProPose [14] addresses this limitation by modeling the probability distributions for human mesh recovery. Since the pose parameters $\theta \in \mathbb{R}^{72}$ of SMPL [48] represent the 3D rotation of each joint and the root orientation, they adopt the matrix Fisher distribution [13, 30] over $\mathcal{SO}(3)$. Due to the gaps between the RGB images and the rotation representations, the neural network cannot easily model the distribution. ProPose [14] also introduces 3D unit vectors for bone directions as the corresponding observation on the previous matrix Fisher distribution as the prior. Leveraging Bayesian inference, they model the posterior distribution of the joint rotations from the prior distribution and observation.

4 MHCDIFF: Multi-hypotheses Conditioned Point Cloud Diffusion

4.1 Overview

Our work aims at reconstructing pixel-aligned 3D human shape as a point cloud given a single occluded RGB image via conditional point cloud diffusion, as shown in Fig. 2. Formally, the diffusion model $\mathcal{F}_\theta(\cdot)$ learns the conditional distribution $q(X_0|I)$ of 3D human shapes given the RGB images $I \in \mathbb{R}^{H \times W \times 3}$. Following PC², we extract the 2D feature map $\mathcal{E}(I) \in \mathbb{R}^{h \times w \times c}$ using ViT [12], to capture the details in the images. The image features are projected onto the partially denoised points: $X_t^{proj} = \Pi(\mathcal{E}(I), X_t)$, where Π is the projection function. This helps obtain pixel-aligned detailed body shapes. Additionally, the diffusion model is conditioned on the local features X_t^{SMPL} from SMPL mesh S to exploit statistical human body priors to complete 3D shapes from occluded body parts (Sec. 4.2). However, the SMPL estimation from single occluded RGB images has a high probability of large errors. To tackle this, we propose a novel multi-hypotheses conditioned diffusion model that considers the distribution of multiple plausible SMPL meshes $\{S_i\}_{i \in \{1, \dots, s\}}$ (Sec. 4.3). Given the partially denoised point cloud X_t , the projected image features X_t^{proj} , and the local features from SMPL X_t^{SMPL} , MHCDIFF predicts the noise ϵ :

$$\mathcal{F}_\theta(X_t, X_t^{proj}, X_t^{SMPL}) = \epsilon. \quad (1)$$

We also discuss how MHCDIFF takes the generative property and the global consistent features to reconstruct occluded parts (Sec. 4.4).

4.2 Local features from SMPL

Given the SMPL (or SMPL-X) mesh S and the partially denoised point cloud X_t at t -th diffusion step, we extract the local features X_t^{SMPL} as:

$$X_t^{SMPL} = [\gamma(d(X_t|S)), \mathbf{n}(X_t|S)], \quad (2)$$

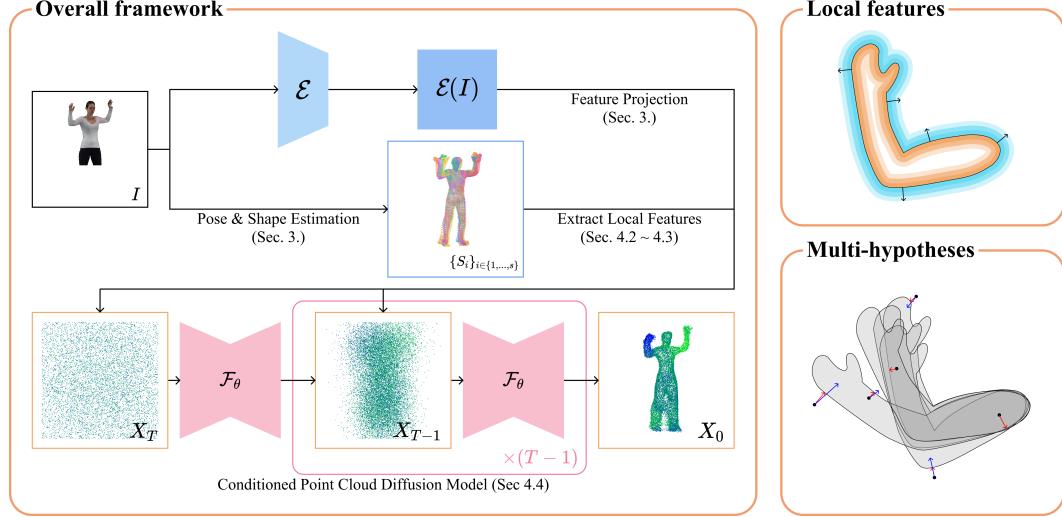


Figure 2: **(Left)** Overview of MHCDIFF. Given an occluded image I , MHCDIFF reconstructs 3D human shape as a point cloud. First, we extract the 2D feature map $\mathcal{E}(I)$ and hypothesize pose and shape parameters of multiple plausible SMPL meshes $\{S_i\}_{i \in \{1, \dots, s\}}$. Our method consists of the conditioned point cloud diffusion model (Sec. 4.4). We project the 2D image features to capture details of the image (Sec. 3) and extract local features from multiple hypothesized SMPL meshes to leverage human body priors (Sec. 4.3) **(Upper Right)** The details of local features (Sec. 4.2). The signed distance field is visualized in **positive** and **negative** regions. The arrows indicate normal vectors \mathbf{n} . **(Lower Right)** The details of multi-hypotheses (Sec. 4.3). We can consider the whole distribution during denoising process with the argmin \bar{i} , and the denoising can be approximated by **red arrows**. However, it is sensitive to extreme samples of the distribution, so we condition the mean of occupancy values, which is visualized by transparency, and the denoising can be approximated by **blue arrows**.

where $d(\cdot) : \mathbb{R}^3 \rightarrow \mathbb{R}$ and $\mathbf{n}(\cdot) : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ are the signed distance and normal obtained from the closest surface of SMPL mesh respectively. In order to map scalar values to a higher dimensional space, we adopt an encoding inspired by the positional encoding in NeRF [59]:

$$\gamma(d) = (\sin(2^0 \pi d), \cos(2^0 \pi d), \dots, \sin(s^{L-1} \pi d), \cos(s^{L-1} \pi d)). \quad (3)$$

The local features $X_t^{SMPL} \in \mathbb{R}^{(2L+3)N}$, which contain the signed distance and normal vector from SMPL, are used to predict the noise ϵ of the point cloud X_t . The local property, which is independent of global pose, helps MHCDIFF to generalize well in diverse SMPL estimation due to occlusion and capture local details.

4.3 Multi-hypotheses condition

The local features are robust to noisy SMPL estimation, but cannot correct the SMPL estimation errors. Following previous multi-hypotheses human pose estimation [6, 43, 18, 21, 8], MHCDIFF takes multi-hypotheses SMPL meshes from estimated distributions and predicts the most plausible outputs. We modify Eq. 2 to handle multiple sampled SMPL meshes $\{S_i\}_{i \in \{1, \dots, s\}}$ using ProPose [14] as an off-the-shelf method:

$$X_t^{SMPL} = [\gamma(d(X_t|S_i)), \mathbf{n}(X_t|S_i)], \quad (4)$$

where $\bar{i} = \operatorname{argmin}_{i \in \{1, \dots, s\}} |d(X_t|S_i)|$, which semantically means that each point follows the closest SMPL mesh $S_{\bar{i}}$ to consider all plausible samples in denoising steps. However, each point gets conditions from only one sample and cannot leverage off-the-shelf probability distributions. In addition to the local features, we also adopt occupancy values:

$$X_t^{SMPL} = \left[\frac{1}{s} \sum_{i=1}^s \gamma(o(X_t|S_i)), \gamma(d(X_t|S_i)), \mathbf{n}(X_t|S_i) \right], \quad (5)$$

where $o(\cdot) : \mathbb{R}^3 \rightarrow \{0, 1\}$ is the occupancy function of the given SMPL mesh, which is a binary signal while the signed distance is continuous. With the mean occupancy and closest signed distance, MHCDIFF can assume all distributions with their respective probabilities. The proposed multi-hypotheses conditioning can take an arbitrary number of SMPL, SMPL-X, and their combined.

4.4 Conditioned point cloud diffusion model

Finally, $\mathcal{F}_\theta(\cdot) : \mathbb{R}^{(3+c+4L+3)N} \rightarrow \mathbb{R}^{3N}$ predicts the noise $\epsilon \in \mathbb{R}^{3N}$ given the concatenation of partially denoised point cloud $X_t \in \mathbb{R}^{3N}$, projected image features $X_t^{proj} \in \mathbb{R}^{cN}$, and local features from SMPL $X_t^{SMPL} \in \mathbb{R}^{(4L+3)N}$ (Eq. 1). Notably, we do not need any learnable parameters to extract the local features from SMPL and aggregate the features of multiple SMPL meshes. We freeze the pre-trained 2D image encoder, so it is straightforward to train the diffusion model without additional training strategies.

The point cloud diffusion model of MHCDIFF takes the role of the decoder of previous implicit-function-based methods. Given the encoded features from RGB images or SMPL meshes, the decoder predicts 3D shapes such as point clouds, occupancy fields, or signed distance fields. The implicit-function-based methods need to sample the query points randomly, so the decoder has been primarily Multi-Layer Perceptron (MLP), which takes the input points independently. MHCDIFF consists of the point cloud diffusion model instead of MLP because (1) the point cloud model considers the global consistent features, (2) the diffusion model has the generative properties, and (3) the denoising process approximates correcting the misaligned SMPL estimation. Given the globally encoded image features X_t^{proj} and the local features from SMPL X_t^{SMPL} , MHCDIFF can inpaint or restore invisible body parts and is robust to noisy SMPL estimation due to occlusion.

Algorithm 1 Pseudocode of learning pipeline of MHCDIFF

Require: $\alpha_{1:T}$: diffusion noise scheduling

```

1: repeat
2:   Sample  $X_0$  from  $q(X_0)$ 
3:   Load the corresponding image  $I$  and ground truth SMPL-X  $S$ 
4:    $t \sim \text{Uniform}(\{1, \dots, T\})$ 
5:    $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ 
6:    $X_t = \sqrt{\bar{\alpha}_t} X_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ 
7:    $X_t^{proj} = \Pi(\mathcal{E}(I), X_t)$  ▷ Project image features (Sec. 3)
8:    $X_t^{SMPL} = [\gamma(o(X_t|S)), \gamma(d(X_t|S)), \mathbf{n}(X_t|S)]$  ▷ Extract local features from SMPL (Sec. 4.2)
9:   Take gradient descent step on
   
$$\nabla_\theta \left\| \epsilon - \mathcal{F}_\theta(X_t, X_t^{proj}, X_t^{SMPL}) \right\|^2$$
 ▷ Point cloud diffusion model (Sec. 4.4)
10:  until converged

```

5 Experiments

Implementation. We use the Pytorch3D library [71] for image feature projection (Sec. 3) and the kaolin library [16] to extract local features from SMPL (Sec. 4.2). MHCDIFF is trained with batch size 8 in 100,000 steps. We use MSN [4] as the image feature encoder. We use AdamW [31] with $\beta = (0.9, 0.999)$ and a learning rate which is decayed linearly from 0.0002 to 0. For diffusion noise schedule, we use linear scheduling from $1 \cdot 10^{-5}$ to $8 \cdot 10^{-3}$ with warmup. For inference, we denoise the point cloud for 1,000 steps. The training process takes approximately 1 day on a single 24GB NVIDIA RTX 4090 GPU with 28M learnable parameters.

Learning. We synthesize the THuman2.0 dataset [98], which contains 526 high-fidelity textured scans with corresponding SMPL-X fits. We use 500 subjects for training and the others for validation.

Algorithm 2 Pseudocode of inference pipeline of MHCDIFF

Require: Input image I

```

1: Sample  $X_T$  from  $\mathcal{N}(0, \mathbf{I})$ 
2: Estimate single or multi SMPL(-X) meshes  $\{S_i\}_{i \in \{1, \dots, s\}}$ 
3: for all  $t$  from  $T$  to 1 do
4:    $z \sim \mathcal{N}(0, \mathbf{I})$  if  $t > 1$  else  $z = 0$ 
5:    $X_t^{proj} = \Pi(\mathcal{E}(I), X_t)$                                  $\triangleright$  Project image features (Sec. 3)
6:   for all  $i$  from 1 to  $s$  do
7:     Compute  $o(X_t|S_i)$ ,  $d(X_t|S_i)$ , and  $\mathbf{n}(X_t|S_i)$        $\triangleright$  Can be accelerated by kaolin [16]
8:   end for
9:    $\bar{i} \leftarrow \text{argmin}_{i \in \{1, \dots, s\}} |d(X_t|S_i)|$ 
10:   $X_t^{SMPL} = [\frac{1}{s} \sum_{i=1}^s \gamma(o(X_t|S_i)), \gamma(d(X_t|S_{\bar{i}})), \mathbf{n}(X_t|S_{\bar{i}})]$   $\triangleright$  Multi-hypotheses conditioning (Sec. 4.3)
11:   $\hat{e} \leftarrow \mathcal{F}_\theta(X_t, X_t^{proj}, X_t^{SMPL})$ 
12:   $X_{t-1} \leftarrow \frac{1}{\sqrt{\alpha_t}}(X_t - \frac{1-\alpha_t}{\sqrt{1-\alpha_t}}\hat{e}) + \sigma_t z$                                  $\triangleright$  DDPM [20] sampling
13: end for
14: return  $X_0$ 

```

We render each human subject from 36 multiple viewpoints and randomly mask the images, resulting in partially occluded body images. We use the farthest point sampling operation to sample 16,384 points from each GT scan. During the training, local features X_t^{SMPL} are extracted from a single corresponding GT SMPL-X. The learning pipeline is presented in Algorithm 1.

Inference. First, we use the CAPE dataset [53, 68] with 150 textured scans. Similar to the training stage, we render each subject from 3 multiple viewpoints and randomly mask the images. During the inference, local features X_t^{SMPL} are extracted from multiple sampled SMPL or single estimated SMPL-X. We sample 10 SMPL meshes for our experiments. To further show the generalizability on the real-world interaction, we also evaluate MHCDIFF on the MultiHuman [105] and Hi4D [97] dataset. MultiHuman, which includes the diverse interaction with objects and people, provides 3D textured scans, so we render each subject from 3 multiple viewpoints. We evaluate the performance of MHCDIFF qualitatively on Hi4D, which includes close human-human interaction with high-fidelity meshes. The inference pipeline is presented in Algorithm 2.

Baseline models. We compare MHCDIFF with parametric models and pixel-aligned reconstruction methods. For parametric models, which are robust for occlusion, we select ProPose [14] as SMPL estimator and PIXIE [15] as SMPL-X estimator. For pixel-aligned reconstruction methods, which can capture geometric details, we select PaMIR [106] for global features, ICON [92] for local features, and HiLo [96] and SIFU [104] for both. For the fair comparison, we primarily condition with the mean of SMPL distribution estimated via ProPose, and PIXIE is also used for ICON, which supports SMPL-X. We use pre-trained weights and evaluate under our test setting.

Evaluation metrics. We employ Point-to-Surface distance and Chamfer Distance as evaluation metrics. MHCDIFF outputs a point cloud, so Chamfer Distance is the average L2 distance from the reconstructed point cloud to vertices of ground-truth scans and vice versa, and Point-to-Surface distance is the average point-to-surface from the reconstructed point cloud to ground-truth scans. The outputs of implicit-function-based methods can be converted meshes via the Marching Cubes algorithm [49]. For fair comparison, we sample the same number of points from the reconstructed meshes uniformly.

5.1 Comparison with state-of-the-art methods

MHCDIFF outperforms prior implicit-function-based methods and SMPL estimation methods on occluded and even full-body images. Fig. 3 presents the robustness of 3D human reconstruction to the occlusion ratio. PaMIR and HiLo cannot handle the occlusions because the global feature encoder is sensitive to misaligned SMPL estimation. SIFU does not use the 3D encoder, but the cross-attention from the normal map of SMPL takes global features and is sensitive to occlusion

	Methods	Chamfer Distance (cm)	Point-to-Surface (cm)
A	PaMIR [106]	12.912	12.619
	ICON [92]	2.896	2.789
	ICON (PIXIE estimation)	3.329	3.212
	SIFU [104]	14.397	14.087
	HiLo [96]	13.711	13.405
B	PIXIE (SMPL-X) [15]	2.705	2.662
	ProPose (SMPL) [14]	2.370	2.307
Ours	MHCDIFF	1.872	1.810

Table 1: **Quantitative evaluation on CAPE dataset.** We report the average Chamfer Distance (cm) and Point-to-Surface distance (cm) on CAPE dataset. We randomly mask the images about 40% in average. We compare the performance with respect to **(A)** implicit-function-based methods; and **(B)** SMPL estimation methods used to condition MHCDIFF and **(A)**. Best in bold, second-best underlined.

	Methods	single	occluded single	two natural-inter	two closely-inter	three
A	PaMIR [106]	0.690	2.349	5.154	3.752	4.714
	ICON [92]	0.555	<u>0.549</u>	<u>0.563</u>	0.786	0.669
	SIFU [104]	0.644	3.335	4.796	3.503	3.264
	HiLo [96]	0.606	2.808	4.139	3.346	4.398
B	PIXIE (SMPL-X) [15]	0.868	0.813	0.755	0.951	0.809
	ProPose (SMPL) [14]	0.675	0.567	0.574	0.766	0.688
Ours	MHCDIFF	0.591	0.491	0.536	0.703	0.673

Table 2: **Quantitative evaluation on MultiHuman dataset.** We report the average Chamfer Distance (cm) for each category. We compare the performance similar to Tab. 1.

		Chamfer Distance (cm)	Point-to-Surface (cm)
full	MHCDIFF	1.872	1.810
A	w/o occupancy	1.893	1.831
	w/o signed distance	2.016	1.949
	w/o normal	1.888	1.827
	w/o encoding	1.928	1.863
	PC ² [57]	3.640	3.533
B	conditioned on PIXIE estimation	2.314	2.237
	conditioned on single ProPose estimation	1.939	1.869
C	trained with ProPose estimation	2.708	2.624
	w/o random masking	1.940	1.868

Table 3: **Ablation study on CAPE dataset.** We validate the effectiveness of **(A)** each component; **(B)** conditioning strategies; and **(C)** training strategies.

The number of SMPL sampled	Chamfer Distance (cm)	Point-to-Surface (cm)	Evaluation time on CAPE dataset (hours)
1	1.939	1.869	4
5	1.882	1.817	8
10	1.872	1.810	12
15	1.833	1.773	16
20	1.836	1.777	20

Table 4: **The correlation between the number of SMPL sampled and the reconstruction quality.** We report the average Chamfer Distance (cm), Point-to-Surface distance (cm) and evaluation time of the various number of SMPL sampled.

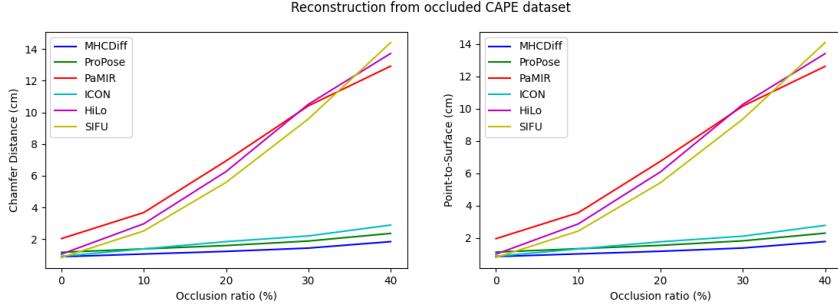


Figure 3: **A cumulative occlusion-to-reconstruction test.** This figure shows the performance of different models from the images of various occlusion ratios. From the whole-body images, which is 0% occlusion, we randomly mask the images from 10% to 40%. MHCDIFF is robust to the occlusion ratio, showing the best performance.

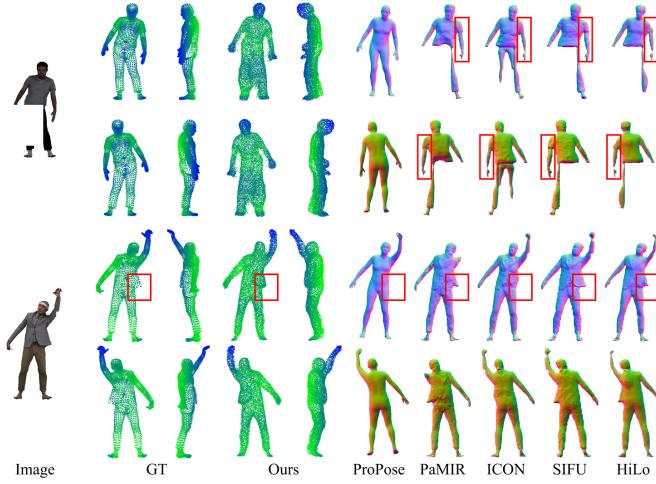


Figure 4: **Qualitative results on CAPE dataset.** We evaluate our method with SMPL estimation method and implicit-function-based methods. Given the upper image, PaMIR, ICON, and HiLo cannot generate the occluded regions. They cannot also handle the misaligned SMPL mesh on the arms, creating incomplete bodies. ProPose predicts the full-body shape, but cannot capture the details like the blazer of the lower image. However, MHCDIFF is robust to the occlusion and misalignment, and can capture pixel-aligned details.

and misaligned SMPL estimation. ICON shows comparable robustness due to its locality, but worse quality than ProPose estimation used to condition as the occlusion ratio increases. On the contrary, MHCDIFF is as robust as the statistical models, showing the most accurate results for all occlusion ratios. The results of 40% occlusion ratio are also displayed in numbers in Tab. 1. Tab. 2 presents the performance on real-world interaction scenarios with MultiHuman dataset. The dataset is divided into 5 categories by the level of occlusions: "occluded single" and "two closely-inter" show the most severe occlusion, and "single" and "three" show the least occlusion. We compare the performance in each category and similar to randomly masked settings, MHCDIFF achieves state-of-the-art on severe occluded images, and comparable performance on full-body images. The major improvements of MHCDIFF are (1) correcting the misaligned SMPL estimation as shown in Tab. 1, and (2) inpainting the invisible regions as shown in Fig. 3. The qualitative results on CAPE dataset are shown in Fig. 4, and MultiHuman and Hi4D datasets are shown in the appendices Sec. E.

5.2 Ablation study

We conduct an ablation on MHCDIFF to validate the effectiveness of each component. In Tab. 3-B, we condition the diffusion model with single SMPL-X (PIXIE) or SMPL (ProPose) estimation.

We improve the performance with multi-hypotheses condition (Sec. 4.3). In Tab. 4, we show the correlation between the number of SMPL sampled and the reconstruction quality. More SMPL hypotheses may include more accurate samples and improve the quality (15 samples), as well as extreme samples and decrease the quality (20 samples). From PC² [57], which only takes image condition, we also validate the local features from SMPL in Tab. 3-A. All of these features improve the performance, especially the signed distance. In Tab. 3-C, MHCDIFF is trained without random masking or by conditioning the distribution estimated by ProPose [14] instead of GT SMPL-X.

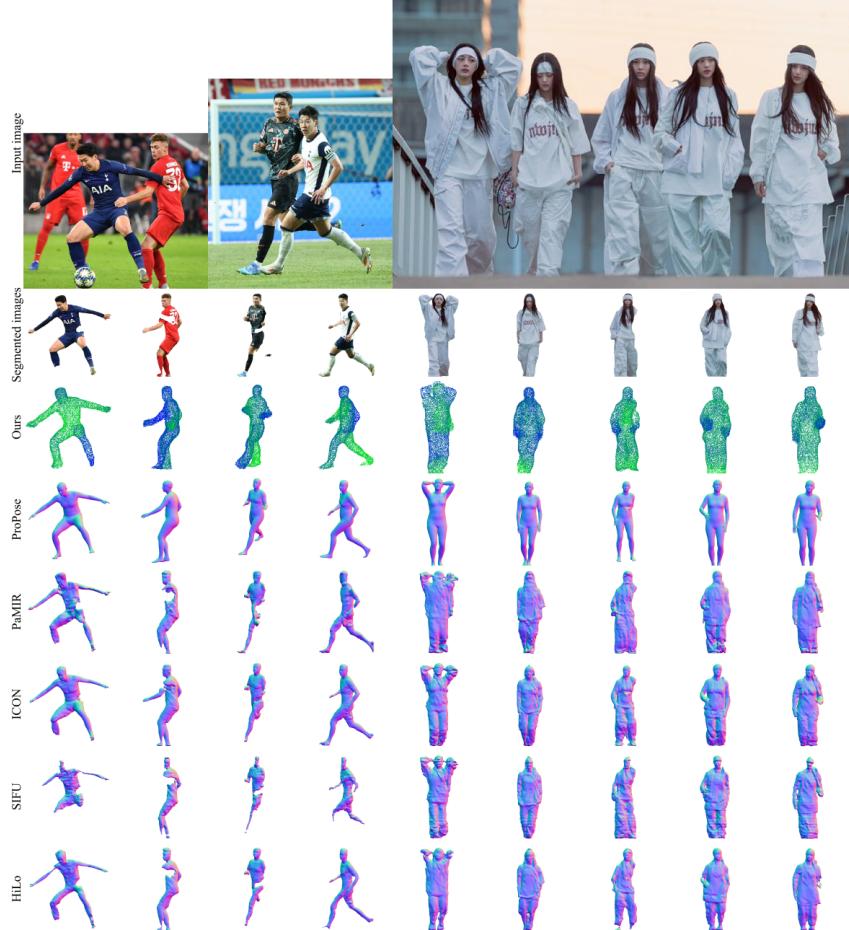


Figure 5: **Qualitative results on in-the-wild images.** Two images on the left show occlusions due to interactions, and the rightmost image shows loose clothes. From internet photos, we use [32] to segment images.

6 Conclusion

In this paper, We present MHCDIFF, which robustly reconstructs pixel-aligned and detailed 3D humans from single occluded images. Rather than implicit-function-based methods, we choose the point cloud diffusion model to generate invisible regions capturing the features globally. Our multi-hypotheses conditioning mechanism extracts local features from multiple SMPL estimations and integrates them without learnable parameters, so MHCDIFF is robust to a single erroneous SMPL due to occlusion. We augment the limited training data by random masking to synthesize occlusion by diverse interaction. The experiments demonstrate that our proposed method outperforms state-of-the-art methods from various levels of occlusion and interaction. In the future, the point cloud of human shapes can be applied to intermediate stages for implicit function [58] and human body deformation [45] or motion flow [38].

7 Acknowledgements

This work was supported by NST grant (CRC 21011, MSIT), IITP grant (RS-2023-00228996, RS-2024-00459749, MSIT) and KOCCA grant (RS-2024-00442308, MCST).

References

- [1] Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single RGB camera. In *CVPR*, 2019.
- [2] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Detailed human avatars from monocular video. In *3DV*, 2018.
- [3] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *ICCV*, 2019.
- [4] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Michael Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *ECCV*, 2022.
- [5] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *ICCV*, 2019.
- [6] Benjamin Biggs, Sébastien Ehrhart, Hanbyul Joo, Benjamin Graham, Andrea Vedaldi, and David Novotny. 3D multibodies: Fitting sets of plausible 3D models to ambiguous image data. In *NeurIPS*, 2020.
- [7] Yukang Cao, Kai Han, and Kwan-Yee K. Wong. Sesdf: Self-evolved signed distance field for implicit 3d clothed human reconstruction. In *CVPR*, 2023.
- [8] Hanbyel Cho and Junmo Kim. Generative approach for probabilistic human mesh recovery using diffusion models. In *ICCV Workshop*, 2023.
- [9] Hongsuk Choi, Gyeongsik Moon, and Kyoung Mu Lee. Pose2mesh: Graph convolutional network for 3d human pose and mesh recovery from a 2d human pose. In *ECCV*, 2020.
- [10] Hongsuk Choi, Gyeongsik Moon, JoonKyu Park, and Kyoung Mu Lee. Learning to estimate robust 3d human mesh from in-the-wild crowded scenes. In *CVPR*, 2022.
- [11] Vasileios Choutas, Lea Muller, Chun-Hao P. Huang, Siyu Tang, Dimitris Tzionas, and Michael J. Black. Accurate 3d body shape regression using metric and semantic attribute. In *CVPR*, 2022.
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- [13] Thomas D Downs. Orientation statistics. In *Biometrika*, volume 59, pages 665–676, 1972.
- [14] Qi Fang, Kang Chen, Yinghui Fan, Qing Shuai, Jiefeng Li, and Weidong Zhang. Learning analytical posterior probability for human mesh recovery. In *CVPR*, 2023.
- [15] Yao Feng, Vasileios Choutas, Timo Bolkart, Dimitrios Tzionas, and Michael J. Black. Collaborative regression of expressive bodies using moderation. In *3DV*, 2021.
- [16] Clement Fuji Tsang, Maria Shugrina, Jean Francois Lafleche, Towaki Takikawa, Jiehan Wang, Charles Loop, Wenzheng Chen, Krishna Murthy Jatavallabhula, Edward Smith, Artem Rozantsev, Or Perel, Tianchang Shen, Jun Gao, Sanja Fidler, Gavriel State, Jason Gorski, Tommy Xiang, Jianing Li, Michael Li, and Rev Lebaredian. Kaolin: A pytorch library for accelerating 3d deep learning research. <https://github.com/NVIDIAAGameWorks/kaolin>, 2022.
- [17] Andrew Gilbert, Marco Volino, John Collomosse, and Adrian Hilton. Volumetric performance capture from minimal camera viewpoints. In *ECCV*, 2018.
- [18] Jia Gong, Lin Geng Foo, Zhipeng Fan, QiuHong Ke, Hossein Rahmani, and Jun Liu. Diffpose: Toward more reliable 3d pose estimation. In *CVPR*, 2023.

- [19] Sang-Hun Han, Min-Gyu Park, Ju Hong Yoon, Ju-Mi Kang, Young-Jae Park, and Hae-Gon Jeon. High-fidelity 3d human digitization from single 2k resolution images. In *CVPR*, 2023.
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.
- [21] Karl Holmquist and Bastian Wandt. Diffpose: Multi-hypothesis human pose estimation using diffusion models. In *ICCV*, 2023.
- [22] Zixuan Huang, Justin Johnson, Shoubhik Debnath, James M Rehg, and Chao-Yuan Wu. Pointinfinity: Resolution-invariant point diffusion models. *CVPR*, 2024.
- [23] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *PAMI*, 2014.
- [24] Aaron S Jackson, Chris Manafas, and Georgios Tzimiropoulos. 3d human body reconstruction from a single image via volumetric regression. In *ECCV Workshop*, 2018.
- [25] Boyi Jiang, Juyong Zhang, Yang Hong, Jinhao Luo, Ligang Liu, and Hujun Bao. Bcnet: Learning body and cloth shape from a single image. In *ECCV*, 2020.
- [26] Hanbyul Joo, Natalia Neverova, and Andrea Vedaldi. Exemplar fine-tuning for 3d human pose fitting towards in-the-wild 3d human pose estimation. In *3DV*, 2020.
- [27] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces. In *CVPR*, 2018.
- [28] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018.
- [29] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM TOG*, 2013.
- [30] CG Khatri and Kanti V Mardia. The von mises-fisher matrix distribution in orientation statistics. In *Journal of the Royal Statistical Society: Series B (Methodological)*, volume 39, pages 95–106, 1977.
- [31] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2017.
- [32] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023.
- [33] Muhammed Kocabas, Chun-Hao P. Huang, Otmar Hilliges, and Michael J. Black. PARE: Part attention regressor for 3D human body estimation. In *ICCV*, 2021.
- [34] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, 2019.
- [35] Nikos Kolotouros, Georgios Pavlakos, Dinesh Jayaraman, and Kostas Daniilidis. Probabilistic modeling for human mesh recovery. In *ICCV*, 2021.
- [36] Juil Koo, Seungwoo Yoo, Minh Hieu Nguyen, and Minhyuk Sung. SALAD: Part-level latent diffusion for 3d shape generation and manipulation. In *ICCV*, 2023.
- [37] Verica Lazova, Eldar Insafutdinov, and Gerard Pons-Moll. 360-degree textures of people in clothing from a single image. In *3DV*, 2019.
- [38] Jihyun Lee, Junbong Jang, Donghwan Kim, Minhyuk Sung, and Tae-Kyun Kim. Fourierhandflow: Neural 4d hand representation using fourier query flow. In *NeurIPS*, 2023.
- [39] Jihyun Lee, Shunsuke Saito, Giljoo Nam, Minhyuk Sung, and Tae-Kyun Kim. Interhandgen: Two-hand interaction generation via cascaded reverse diffusion. In *CVPR*, 2024.
- [40] Jihyun Lee, Minhyuk Sung, Honggyu Choi, and Tae-Kyun Kim. Im2hands: Learning attentive implicit representation of interacting two-hand shapes. In *CVPR*, 2023.
- [41] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *CVPR*, pages 3383–3393, 2021.

[42] Peike Li, Yunqiu Xu, Yunchao Wei, and Yi Yang. Self-correction for human parsing. *PAMI*, 2020.

[43] Wenhao Li, Hong Liu, Hao Tang, Pichao Wang, and Luc Van Gool. Mhformer: Multi-hypothesis transformer for 3d human pose estimation. In *CVPR*, 2022.

[44] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. Cliff: Carrying location information in full frames into human pose and shape estimation. In *ECCV*, 2022.

[45] Minghua Liu, Minhyuk Sung, Radomír Měch, and Hao Su. Deepmetahandles: Learning deformation meta-handles of 3d meshes with biharmonic coordinates. In *CVPR*, 2021.

[46] Zhen Liu, Yao Feng, Michael J. Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu. Meshdiffusion: Score-based generative 3d mesh modeling. In *ICLR*, 2023.

[47] Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-voxel cnn for efficient 3d deep learning. In *NeurIPS*, 2019.

[48] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM TOG*, 2015.

[49] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *SIGGRAPH Comput. Graph.*, 1987.

[50] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *CVPR*, 2021.

[51] Zhaoyang Lyu, Zhifeng Kong, Xudong Xu, Liang Pan, and Dahu Lin. A conditional point diffusion-refinement paradigm for 3d point cloud completion. In *ICLR*, 2022.

[52] Qianli Ma, Shunsuke Saito, Jinlong Yang, Siyu Tang, and Michael J. Black. SCALE: Modeling clothed humans with a surface codec of articulated local elements. In *CVPR*, 2021.

[53] Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, and Michael J. Black. Learning to Dress 3D People in Generative Clothing. In *CVPR*, 2020.

[54] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J. Black. The power of points for modeling humans in clothing. In *ICCV*, 2021.

[55] Kanti V Mardia, Peter E Jupp, and KV Mardia. Directional statistics. In *Whiley Online Library*, volume 2, 2000.

[56] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *3DV*, 2017.

[57] Luke Melas-Kyriazi, Christian Rupprecht, and Andrea Vedaldi. Pc2: Projection-conditioned point cloud diffusion for single-image 3d reconstruction. In *CVPR*, 2023.

[58] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019.

[59] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.

[60] Gyeongsik Moon and Kyoung Mu Lee. I2l-meshnet: Image-to-lixel prediction network for accurate 3d human pose and mesh estimation from a single rgb image. In *ECCV*, 2020.

[61] Lea Muller, Vickie Ye, Georgios Pavlakos, Michael J. Black, and Angjoo Kanazawa. Generative proxemics: A prior for 3D social interaction from images. *CVPR*, 2023.

[62] Gimin Nam, Mariem Khlifi, Andrew Rodriguez, Alberto Tono, Linqi Zhou, and Paul Guerrero. 3d-ldm: Neural implicit 3d shape generation with latent diffusion models. *arXiv preprint arXiv:2212.00842*, 2022.

[63] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022.

[64] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter V. Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model-based human pose and shape estimation. In *3DV*, 2018.

[65] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019.

[66] Priyanka Patel, Chun-Hao P. Huang, Joachim Tesch, David T. Hoffmann, Shashank Tripathi, and Michael J. Black. AGORA: Avatars in geography optimized for regression analysis. In *CVPR*, 2021.

[67] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *CVPR*, 2019.

[68] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael Black. Clothcap: Seamless 4d clothing capture and retargeting. *ACM TOG*, 2017.

[69] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICLR*, 2023.

[70] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017.

[71] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501*, 2020.

[72] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021.

[73] Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J. Guibas. Humor: 3d human motion model for robust pose estimation. In *ICCV*, 2021.

[74] RenderPeople. Renderpeople. renderpeople.com, 2018.

[75] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM TOG*, 2017.

[76] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019.

[77] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, 2020.

[78] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Hierarchical Kinematic Probability Distributions for 3D Human Shape and Pose Estimation from Images in the Wild. In *ICCV*, 2021.

[79] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Probabilistic 3D Human Shape and Pose Estimation from Multiple Unconstrained Images in the Wild. In *ICCV*, 2021.

[80] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021.

[81] Anastasis Stathopoulos, Ligong Han, and Dimitris Metaxas. Score-guided diffusion for 3d human recovery. In *CVPR*, 2024.

[82] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Black Michael J., and Tao Mei. Monocular, One-stage, Regression of Multiple 3D People. In *ICCV*, 2021.

[83] Yu Sun, Wu Liu, Qian Bao, Yili Fu, Tao Mei, and Michael J. Black. Putting people in their place: Monocular regression of 3D people in depth. In *CVPR*, 2022.

[84] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. In *ICLR*, 2024.

[85] Yingzhi Tang, Qijian Zhang, Junhui Hou, and Yebin Liu. Human as points: Explicit point-based 3d human reconstruction from single-view rgb images. *arXiv preprint arXiv:2311.02892*, 2023.

[86] Matthew Trumble, Andrew Gilbert, Adrian Hilton, and John Collomosse. Deep autoencoder for combined human pose estimation and body model upscaling. In *ECCV*, 2018.

[87] Twindom. Twindom. twindom.com, 2018.

[88] G  l Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. BodyNet: Volumetric inference of 3D human body shapes. In *ECCV*, 2018.

[89] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.

[90] Junying Wang, Jae Shin Yoon, Tuanfeng Y. Wang, Krishna Kumar Singh, and Ulrich Neumann. Complete 3d human reconstruction from a single incomplete image. In *CVPR*, 2023.

[91] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J. Black. ECON: Explicit Clothed humans Optimized via Normal integration. In *CVPR*, 2023.

[92] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J. Black. ICON: Implicit Clothed humans Obtained from Normals. In *CVPR*, 2022.

[93] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, Bill Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Ghum & ghuml: Generative 3d human shape and articulated pose models. In *CVPR*, 2020.

[94] Yuanlu Xu, Song-Chun Zhu, and Tony Tung. Denserac: Joint 3d pose and shape estimation by dense render-andcompare. In *ICCV*, 2019.

[95] Xuetong Yang, Yihao Luo, Yuliang Xiu, Wei Wang, Hao Xu, and Zhaoxin Fan. D-IF: Uncertainty-aware Human Digitization via Implicit Distribution Field. In *ICCV*, 2023.

[96] Yifan Yang, Dong Liu, Shuhai Zhang, Zeshuai Deng, Zixiong Huang, and Mingkui Tan. Hilo: Detailed and robust 3d clothed human reconstruction with high-and low-frequency information of parametric models. In *CVPR*, 2024.

[97] Yifei Yin, Chen Guo, Manuel Kaufmann, Juan Zarate, Jie Song, and Otmar Hilliges. Hi4d: 4d instance segmentation of close human interaction. In *CVPR*, 2023.

[98] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4d: Real-time human volumetric capture from very sparse consumer rgbd sensors. In *CVPR*, 2021.

[99] Ilya Zakharkin, Kirill Mazur, Artur Grigorev, and Victor Lempitsky. Point-based modeling of human clothing. In *ICCV*, 2021.

[100] Andrei Zanfir, Eduard Gabriel Bazavan, Mihai Zanfir, Bill Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Neural descent for visual 3d human pose and shape. In *CVPR*, 2021.

[101] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. Lion: Latent point diffusion models for 3d shape generation. In *NeurIPS*, 2022.

[102] Siwei Zhang, Qianli Ma, Yan Zhang, Sadegh Aliakbarian, Darren Cosker, and Siyu Tang. Probabilistic human mesh recovery in 3d scenes from egocentric views. In *ICCV*, 2023.

[103] Zechuan Zhang, Li Sun, Zongxin Yang, Ling Chen, and Yi Yang. Global-correlated 3d-decoupling transformer for clothed avatar reconstruction. In *NeurIPS*, 2023.

[104] Zechuan Zhang, Zongxin Yang, and Yi Yang. Sifu: Side-view conditioned implicit function for real-world usable clothed human reconstruction. In *CVPR*, 2024.

[105] Yang Zheng, Ruizhi Shao, Yuxiang Zhang, Tao Yu, Zerong Zheng, Qionghai Dai, and Yebin Liu. Deepmulticap: Performance capture of multiple characters using sparse multiview cameras. In *ICCV*, 2021.

[106] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *PAMI*, 2021.

[107] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *AAAI*, 2020.

[108] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *ICCV*, 2021.

[109] Hao Zhu, Xinxin Zuo, Sen Wang, Xun Cao, and Ruigang Yang. Detailed human shape estimation from a single image by hierarchical mesh deformation. In *CVPR*, 2019.

A Broader impact

Our method can be potentially used for AR/VR applications. The real-world interaction can be captured and modeled in virtual scenes, which can be extended to reinforcement learning. However, there are potential risks associated with falsifying human avatars, which could inadvertently compromise personal privacy. Consequently, there is a pressing need to establish regulations that clarify the fair use of such technology.

B Limitations

Our method, based on DDPM [20] sampling with 1,000 steps, has limitation on efficiency. The training time is reasonable because we do not need query point sampling, which yields CPU bottleneck to learn implicit-function. However, evaluation on CAPE dataset takes about 12 hours, while other implicit-function-based methods take about 30 minutes. We can apply DDIM [80] sampling with fewer steps to shorten the inference time.

C Pointcloud to mesh

Following previous work [85, 19], we try to convert our reconstructed point cloud to mesh with Screened Poisson surface reconstruction [29]. However, the process takes about 10 hours per sample with 16,384 points. The implicit function [58] converts the point clouds to occupancy fields by encoding features with a PointNet [70]. This two-stage pipeline can generate occluded regions and capture details. We will try this pipeline in our future work.

D Statistical significance

We evaluate MHCDIFF on CAPE dataset [53, 68] with 10 different random seeds. The random seeds effect on random noise in the diffusion process and SMPL sampling from the estimated distribution vis ProPose [14]. The Chamfer Distance and Point-to-Surface are $1.872(\pm 0.008)$ and $1.810(\pm 0.008)$ with 1-sigma error bars, respectively.

E Qualitative results

For the real-world interaction, we evaluate MHCDIFF on MultiHuman [105] and Hi4D [97] datasets. We render the textured scans with Pytorch3D library [71] for MultiHuman dataset, and segment each subject with pre-trained network [42] for Hi4D dataset. Our proposed method is robust not only to the occlusion but also to noise in full images or segmentation process.



Figure 6: **Qualitative results on MultiHuman dataset.**

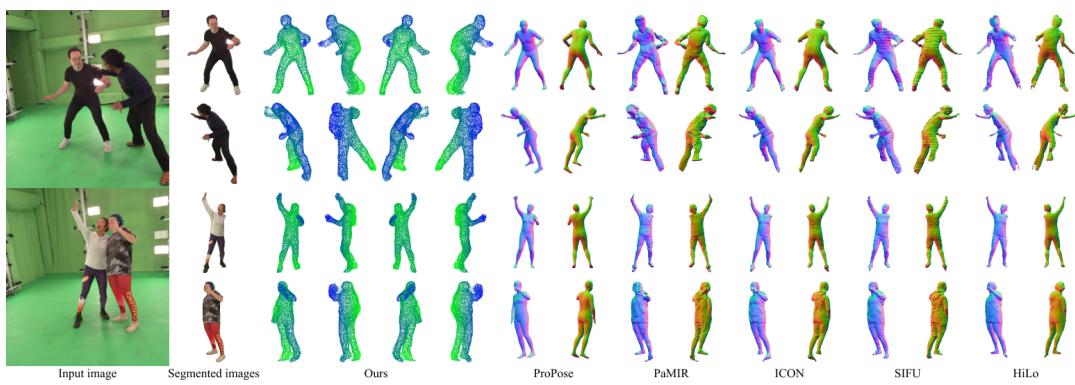


Figure 7: **Qualitative results on Hi4D dataset.**

NeurIPS Paper Checklist

1. Claims

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

Answer: [\[Yes\]](#)

Justification: Please refer to Abstract and Introduction Sections.

Guidelines:

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

2. Limitations

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

Answer: [\[Yes\]](#)

Justification: Please refer to Sec. B.

Guidelines:

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

3. Theory Assumptions and Proofs

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

Answer: [\[NA\]](#)

Justification: The paper is based on experiments.

Guidelines:

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

4. Experimental Result Reproducibility

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

Answer: [Yes]

Justification: Please refer to Implementation paragraph in Sec. 5.

Guidelines:

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

5. Open access to data and code

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

Answer: [\[Yes\]](#)

Justification: The code will be public after cleansing.

Guidelines:

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

6. Experimental Setting/Details

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

Answer: [\[Yes\]](#)

Justification: Please refer to Sec. 5.

Guidelines:

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

7. Experiment Statistical Significance

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

Answer: [\[Yes\]](#)

Justification: Please refer to Sec. D.

Guidelines:

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

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

8. Experiments Compute Resources

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

Answer: [\[Yes\]](#)

Justification: Please refer to Implementation paragraph in Sec. 5.

Guidelines:

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

9. Code Of Ethics

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

Answer: [\[Yes\]](#)

Justification: The paper conform the NeurIPS Code of Ethics.

Guidelines:

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

10. Broader Impacts

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

Answer: [\[Yes\]](#)

Justification: Please refer to Sec. A.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to

generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

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

Answer: [NA]

Justification: Our method poses no such risks.

Guidelines:

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

12. Licenses for existing assets

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

Answer: [Yes]

Justification: We cite the creators of the assets.

Guidelines:

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

13. New Assets

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

Answer: [\[Yes\]](#)

Justification: We submit the code as a zip file.

Guidelines:

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

14. **Crowdsourcing and Research with Human Subjects**

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

Answer: [\[NA\]](#)

Justification: The paper does not involve crowdsourcing, and we only use public datasets for human subjects.

Guidelines:

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

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

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

Answer: [\[NA\]](#)

Justification: The paper does not involve crowdsourcing, and we only use public datasets for human subjects.

Guidelines:

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