

Generative Modeling of Shape-Dependent Self-Contact Human Poses

Takehiko Ohkawa^{1,2*}, Jihyun Lee^{1,3*}, Shunsuke Saito¹, Jason Saragih¹, Fabian Prada¹,
 Yichen Xu¹, Shoou-I Yu¹, Ryosuke Furuta², Yoichi Sato², and Takaaki Shiratori¹

¹Codec Avatars Lab, Meta ²The University of Tokyo ³KAIST

Project page: <https://tkhkaeio.github.io/projects/25-scgen>

Abstract

One can hardly model self-contact of human poses without considering underlying body shapes. For example, the pose of rubbing a belly for a person with a low BMI leads to penetration of the hand into the belly for a person with a high BMI. Despite its relevance, existing self-contact datasets lack the variety of self-contact poses and precise body shapes, limiting conclusive analysis between self-contact poses and shapes. To address this, we begin by introducing the first extensive self-contact dataset with precise body shape registration, **Goliath-SC**, consisting of 383K self-contact poses across 130 subjects. Using this dataset, we propose generative modeling of self-contact prior conditioned by body shape parameters, based on a body-part-wise latent diffusion with self-attention. We further incorporate this prior into single-view human pose estimation while refining estimated poses to be in contact. Our experiments suggest that shape conditioning is vital to the successful modeling of self-contact pose distribution, hence improving single-view pose estimation in self-contact.

1. Introduction

Human poses in our daily life often involve *self-contact*, such as face touching, arm crossing, or hand placement, where body parts come into contact with the body surface. These interactions with our own body are not only unconsciously displayed but also carry profound meaning across various disciplines, including psychology and social communication. Observing self-contact gestures, the areas touched can signal emotional states (*e.g.*, anxiety or tension) [13, 14, 47], and express linguistic symbols and contexts in sign language [6, 59, 65]. Notably, these self-contact poses are inherently constrained by the underlying body shapes. As shown in Fig. 1, two subjects performing “rubbing belly” gestures exhibit different poses and contacts due to variations in body shapes and proportions re-




Figure 1. **Body shape dependency in self-contact poses.** We observe that self-contact poses (*e.g.*, “rubbing belly”) are influenced by the subject’s body shape; for example, a person with a slimmer body (top) engages in different self-contact poses over one with a larger torso (bottom). Indeed, the contact maps on the template mesh (right) differ noticeably. Examples are sampled from the **Goliath-SC** dataset we captured.

lated to Body Mass Index (BMI). Despite their significance, accurately modeling self-contact poses remains a considerable challenge [11, 41]; particularly, the dependency of human poses on body shapes is underexplored.

The challenge of self-contact modeling stems from the lack of datasets containing large self-contact poses with precise body shape registration. Existing 3D body self-contact datasets, HumanSC3D [11] and MTP [41], contain small self-contact poses (1-4K poses) and suffer from inaccurate registration due to the lack of paired RGB images [41]. Other studies have highlighted specific part interactions, such as hand-hand [40] or hand-face contact [53]. However, their scopes are limited to isolated body parts and fail to capture the holistic nature of self-contact, overlooking how the full-body pose and shape influence the contact.

*Work done during the internship.

Given the limitations of the existing datasets, we begin by offering the first extensive self-contact dataset with varying full-body poses and precise body shape registration, dubbed **Goliath-SC**. Our self-contact dataset contains the largest amount of self-contact poses, comprising 383K poses from 130 subjects. Additionally, it provides accurate full-body mesh registration based on 3D scans in a multi-camera dome (Goliath [37]), which are converted to SMPL-X [46] to access body shape parameters.

Using this dataset, we model the dependency of self-contact poses on body shapes via generative models. The generative modeling is designed to learn self-contact pose distribution for the given body shapes, independent of image input. Pose-based generative training has an advantage over direct pose regression from images [11, 40, 41, 53] due to its generalizability to unknown environments and subjects. Removing image input helps debias from the captured environments. Furthermore, it facilitates interpolation in the learned pose–shape space, enabling the modeling of plausible self-contact poses for novel body shapes or contact locations not explicitly seen during training.

In more detail, our approach involves a new insight of *shape-dependency* in generative modeling with denoising diffusion. Unlike joint distribution modeling between pose and shape [27, 42] of 3D human models [36, 51], we explicitly model the shape-dependent manifold of self-contact poses using diffusion models [19, 55]. Specifically, we develop a latent diffusion model with self-attention, termed **PAPoseDiff**, which considers the relationship among highly interacting body parts (*e.g.*, hands, body, and face).

Finally, we leverage the learned diffusion prior to refine 3D poses in self-contact. Given the initial SMPL-X estimation, we refine the poses to have a smaller error to the 2D keypoint observation, while maintaining the plausibility in contact acquired by the former generative training. Our experiments demonstrate that our refinement with the shape-conditional diffusion prior surpasses a recent diffusion prior for human contact (BUDDI [42]) and the state-of-the-art foundation model with direct regression (SMPLer-X [4]) in the Goliath-SC *eval* set with unseen subjects.

Our contributions are summarized as follows:

- We introduce a new self-contact dataset **Goliath-SC** with extensive poses and precise body shape registration.
- We propose generative learning of the shape-dependent manifold of self-contact poses, along with a latent diffusion with part-aware self-attention, **PAPoseDiff**.
- We propose an efficient single-view pose refinement that fits initial SMPL-X predictions to the observed 2D keypoints using the learned diffusion model.

2. Related Work

Self-contact datasets: Human contact is taken into ac-

count in 3D human reconstruction [25, 41, 64]. These studies include self-contact of a single person (*e.g.* crossing arms) [11, 41], multi-person interactions like hugging [10, 18, 25, 42, 64], or contact with external environments, such as scene [15, 22, 31] and handheld objects [5, 8, 12, 33, 35, 44, 45, 57]. However, constructing self-contact datasets is particularly difficult due to higher self-occlusion. HumanSC3D [11] and MTP [41] datasets contain a limited number of poses (1-4K poses) and inaccurate annotations due to the absence of paired RGB images with the captured poses [41]. Other studies focus on specific part interactions; InterHand2.6M captures hand-hand interactions [40], while Decaf highlights hand-face contact [53]. Despite allowing fine interaction analysis, capturing only isolated parts disregards the holistic perspective of self-contact, *i.e.*, how the body influences hands and face in contact. In contrast, our captured Goliath-SC dataset provides extensive self-contact poses (383K) with a dense camera setup, including full-body registration with precise shapes. Our dataset also includes continuous pose variations, unlike frame-independent pose registration in MTP [41].

Self-contact estimation: Previous self-contact works follow a *regressive* approach, aiming to estimate contact states from a single image. Early attempts [11, 41] rely on the annotation of discrete 2D contact labels, representing which body parts are in contact, though the annotation process is labor-intensive and difficult to scale. Fieraru *et al.* formulate the tasks of segmenting in-contact parts and predicting interacting part pairs (contact signature) [11]. Muller *et al.* estimate human poses in self-contact [41], with two distinct training setups: (1) supervised training of a regressor on 3D GTs (*i.e.*, MTP [41]) and (2) additional optimization when 3D GTs are unavailable, instead relying on discrete contact labels (*i.e.*, in-the-wild data like DSC [41]). Without using such contact labels, recent human foundation models (*e.g.*, SMPLer-X [4]) extend (1)’s approach to train a ViT network [7] across various 3D human datasets, including self-contact scenarios (*i.e.*, HumanSC3D [11] and MTP [41]). While this simple regression strategy generalizes across domains, it still struggles to capture the nuanced self-contact. To address this, we investigate a *generative* approach to refine the regressor’s estimates, without relying on image input and manual annotations for 2D contact parts.

Diffusion models: Denoising diffusion [19, 55] is becoming a popular choice for generative prior modeling due to its higher capability compared to handcrafted methods [1, 46] or VAEs [26, 46]. Diffusion models are trained to iteratively denoise a Gaussian noise to sample from the learned data distribution [19, 55]. While they have been widely adopted, *e.g.*, for motion synthesis [58, 63], only a few works have modeled human contact using diffusion models. BUDDI [42] and InterHandGen [27] are concurrently proposed to model the contact between two bodies (either




Figure 2. **Examples of our Goliath-SC dataset and contact heatmap.** We capture self-contact poses from 130 subjects with scripted action instructions (e.g., “hand hitting forehead”). Examples from the subjects of Goliath-4 [37] are shown in the left figure. We compute vertex-wise binary contact maps to find contact frames, and the averaged heatmap is shown in the right figure.

Dataset	#SCPose	#Subj.	Params	Annot.
HumanSC3D [11]	1.0K	6 (3/3/0)	SMPL-X [46] / GHUM [62]	Mocap
FlickrSC3D [11]	<3.9K	-	SMPL-X	Pseudo-3D-GTs
MTP [41]	1.6K	148 (52/96/0)	SMPL-X	Pseudo-3D-GTs
Goliath-SC (Ours)	383K	130 (70/56/4)	SMPL-X	MV RGB scan [37] (cam: 220)

Table 1. **Comparison of full-body self-contact datasets.** We compare the number of self-contact poses, captured subjects, body parametrization, and annotation methods. The subject data include the gender ratio (female/male/non-binary).

human bodies or hands) with the DDPM formulation [19].

While these previous methods learn the joint distribution of the pose and shape parameters of SMPL [36] or MANO [51], our diffusion modeling relies on a new assumption that *pose should depend on body shape*, thus generating poses conditioned on the given body shapes. Furthermore, when adapting the diffusion prior to single-view pose estimation, our proposed refinement-based method does not require additional fine-tuning as in [27].

3. Self-Contact Analysis and Dataset

We introduce a new self-contact dataset with varying full-body poses and shapes, termed **Goliath-SC**. Our capture is based on a multi-camera dome setup of Goliath [37] with 3D full-body scans from 220 RGB cameras. The scope of captured activities lies in natural self-contacts occurring in daily life like touching the face, body, hands, etc. Tab. 1 and Fig. 2 show data statistics and examples.

Our dataset has the following advantages. (1) Our captures contain substantial self-contact data with 383K poses from 130 subjects, exceeding the existing self-contact datasets [11, 41] by two orders of magnitude. (2) Owing to numerous cameras with increased resolutions for the face and hands areas [37], it provides high-quality mesh registration with fine details for hands and face. This enables capturing fine self-contacts like “rubbing eyes” and “massaging hands”, which is distinguished from the existing self-contact scenarios [10, 11, 41, 64]. (3) Instead of collecting in-contact poses independently [41], we capture the sequence of natural self-contact poses at 30 Hz with scripted

action instructions (e.g., “rubbing belly”), leading to diverse and continuous self-contact poses. (4) Unlike targeting specific body parts (e.g., hand-hand in InterHand2.6M [40], hand-face in Decaf [53]), our dataset provides complete 3D full-body poses including hand-hand and hand-face interactions. This enables holistic behavior modeling in which hands and face are constrained by the body’s kinematics. (5) To model the shape-dependent manifold, we convert the registered meshes to SMPL-X [46], which gives the latent shape parameters. Details are found in the supplement.

Contact maps and data screening: To comprehend self-contact patterns, we compute vertex-wise contact maps; see Fig. 2. We first create binary contact maps by discriminating if the hand vertices are close ($< 3\text{mm}$) to the rest of the body vertices, and then select contact frames with positive contact maps of the captured sequence. This indicates that each sample represents a unique self-contact pose in which the hand is touching somewhere on the body.

We then calculate the contact heatmap from the binary maps, indicating contact likelihood as [57]. We observe that it includes various interactions across hands, face, neck, arms, and torso. While the studies on hand-object grasping have a high contact likelihood in the finger areas [8, 9, 45, 57], our self-contact data include frequent interactions in the palm of hands as well. This suggests that self-touching gestures are expressed by using hands widely from the palm to the fingertips.

4. Method

We present our proposed generative diffusion model for self-contact pose modeling. Here, we aim to model the manifold of self-contact poses, particularly depending on the subject’s body shape. We detail our task and model setup in Secs. 4.1 and 4.2 and training objectives in Sec. 4.3. We then provide the inference process in Sec. 4.4. Fig. 3 shows the overview of our proposed diffusion model, dubbed **PAPoseDiff**, and our refinement scheme to obtain refined pose \mathbf{X}_0^{ref} given initial 3D pose estimate \mathbf{X}_0^{init} .




Figure 3. **Shape-conditional denoising diffusion model for self-contact poses (left) and single-view refinement (right).** Our proposed diffusion model (left), **PAPoseDiff**, follows latent diffusion with part-aware attention. The model is trained to generate part-wise pose parameters conditioned on the shape information while considering their interactions with self-attention (SA). We also add small perturbations for the shapes to generalize to unseen subjects. The training losses are described in Sec. 4.3. Our refinement (right) is based on the observations of 2D keypoints and initial 3D pose estimation. We diffuse the initial 3D pose X_0^{init} and then denoise it to obtain a refined pose X_0^{ref} while fitting to the 2D observation.

4.1. Diffusion process

We follow the DDPM formulation [19], where the diffusion process consists of forward and reverse paths spanned with diffusion time steps $t \in [1, T]$. The forward process ($1 \rightarrow T$) takes an input data X_0 and gradually adds standard Gaussian noise ϵ_t to the data. We denote the process of diffusing X_0 at step t as $X_t = \text{noise}(X_0, t)$, formulated as

$$\text{noise}(X_0, t) = \sqrt{\bar{\alpha}_t} X_0 + \sqrt{(1 - \bar{\alpha}_t)} \epsilon_t. \quad (1)$$

The noisiness of X_t is controlled by noise variances β_t , *e.g.*, given by a cosine scheduler [43]. The coefficients that balance the noise and data terms are determined by $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$. T is set to 1000.

The reverse process ($T \rightarrow 1$) denoises the data in every step and finally generates a clean sample \hat{X}_0 with a learnable model f . Following [55, 58], given the noised data X_t in step t , the model directly approximates the original data X_0 as $\hat{X}_0 = f(X_t, t)$. For conditional generation, the model can take an additional conditional input c as $f(X_t, t, c)$.

4.2. Shape-dependent pose modeling

With the diffusion formulation, we construct data representations for poses and network architectures that enforce the shape-dependent constraint. We explicitly model the interactions between different body parts including hands, body, and face, which distinguishes our approach from the previous studies that only use body parameters [41, 42].

Data representation: To learn self-contact poses, we use part-wise pose parameters as the target to denoise and shape parameters as conditional input, which are obtained by the differentiable SMPL-X model [46]. The input pose data

are constructed as $X = [\theta_f, \theta_{rh}, \theta_{lh}, \theta_b]$, where $\theta_f \in \mathbb{R}^{3+10}$, $\theta_{rh}, \theta_{lh} \in \mathbb{R}^{15 \times 3}$, $\theta_b \in \mathbb{R}^{21 \times 3}$ indicate pose parameters for face, right hand, left hand, and body, respectively. To clarify, the target face parameters, jaw pose and expression, are combined in θ_f for convenience.

The shape parameters are represented as $I \in \mathbb{R}^{N_s}$ of the SMPL-X where N_s is the shape dimension (≤ 300). These encode the subject’s physical identity such as bone length and body size. In contrast to part-isolated input such as hand-hand [40] (MANO’s) or hand-face interactions (MANO’s + FLAME) [53], the whole body parametrization provides additional constraints about the location of hands and face restricted through the kinematic chain of the body. Owing to this, our representation only relies on local pose for simplicity while global orientation and translation of SMPL-X are disregarded.

Latent diffusion with part-aware self-attention: Regarding part-wise interaction modeling, we propose a part-aware self-attention transformer with latent embedding as f . Unlike learning on pose parameters directly in the diffusion process [27, 42], we train the diffusion model in the latent space, inspired by latent diffusion for image synthesis [50]. As joint movements in human motion are highly coordinated, the DOFs of whole-body joints can be represented in a lower-dimensional space [29, 30, 51]. Similarly, self-contact poses are intrinsically embedded in a latent manifold, as restricted to move along the body surface. To enforce these constraints, we utilize auto-encoders for part-wise poses, enabling the discovery of plausible and semantically meaningful latent embeddings in training.

With the embeddings of the pose, diffusion time, and shape, we then utilize a self-attention transformer [60] as the denoising module. Specifically, the query, key, and

value of the attention are given by the concatenation of the embeddings across face, right/left hand, body, time, and shape. This enables the model to consider the interactions across part-wise poses, shape, and diffusion time (*i.e.*, noisiness). Part-wise learnable embeddings are also added to facilitate part-aware relational learning as [42].

We jointly train all the modules including the latent embedding and the transformer by following recent works [34, 61] that jointly train the parameters (or adapter parameters) of both networks to improve the output quality.

Shape-conditional perturbation: Instead of naive conditioning of shape parameters, enriching the diversity of subjects in training is critical in learning the shape-dependent manifold and generalizing to unseen subjects in testing. Prior works use conditional dropout $c = \emptyset$ to emulate unconditional generation [20, 27, 42], thereby increasing the diversity of generated samples. However, in our context, a zero shape value corresponds to a plain body shape that lacks identity-specific signals.

We instead propose to perturb shape parameters slightly to augment the subject’s identity, assuming that individuals with similar identities are likely to perform similar self-contact poses. As seen in Fig. 1, two persons performing the “rubbing belly” pose differently, particularly the right arm angles are non-identical. However, we observe that people with similar body shapes can come into contact with identical pose parameters. We therefore replace a normal shape conditioning $c = \mathbf{I}$ with the perturbed shapes with a certain probability (*e.g.*, 30%) as

$$c = \mathbf{I} + s_I \epsilon \quad (2)$$

where ϵ is standard Gaussian noise and s_I is a scaling factor to control the scale of the perturbation.

4.3. Training objectives

The training loss is computed by taking the difference between the original data \mathbf{X}_0 and the generated data $\hat{\mathbf{X}}_0$. We use the L1 loss between \mathbf{X}_0 and $\hat{\mathbf{X}}_0$ for the pose space, denoted as \mathcal{L}_θ . We also compute losses on the mesh space after constructing meshes with the SMPL-X layer. We adopt the L1 loss between original and generated meshes for vertices \mathcal{L}_v , and an L1-based collision loss that penalizes vertices in collision on the generated mesh \mathcal{L}_{col} . The overall loss is formulated as

$$\mathcal{L}_D = \lambda_\theta \mathcal{L}_\theta + \lambda_v \mathcal{L}_v + \lambda_{col} \mathcal{L}_{col}. \quad (3)$$

λ_θ , λ_v , and λ_{col} are the weights for each loss. We use 6D rotation representation [66] for pose parameters.

Collision detection: The collision loss \mathcal{L}_{col} is designed to avoid heavy penetration on the mesh, which is essential to maintain plausible self-contact. For detection, [41] requires

Algorithm 1 Single-view pose refinement: given initial 3D pose estimate \mathbf{X}_0^{init} , shape parameters \mathbf{I} , 2D keypoints to fit \mathbf{P}_{2d} , projection from pose to 2d keypoints M_{p2d} , a weight for 2d keypoint fitting λ_{2d} , start diffusion time n_r , mask for poses of interest \mathbf{m}_p .

```

▷ Initialization: diffuse 3D pose at step  $n_r$ 
1:  $\mathbf{X}_{n_r} \leftarrow \text{noise}(\mathbf{X}_0^{init}, n_r)$ 
2: for  $n = n_r$  to 1 do
3:    $\hat{\mathbf{X}}_0 \leftarrow f(\mathbf{X}_n, n, \mathbf{I})$ 
   ▷ Optional: Blended pose denoising
4:    $\hat{\mathbf{X}}_0 \leftarrow \hat{\mathbf{X}}_0 \odot \mathbf{m}_p + \mathbf{X}_0^{init} \odot (1 - \mathbf{m}_p)$ 
5:    $\epsilon_n \leftarrow \frac{1}{\sqrt{1 - \bar{\alpha}_n}} (\mathbf{X}_n - \sqrt{\bar{\alpha}_n} \hat{\mathbf{X}}_0)$ 
6:    $\mathbf{X}'_{n-1} \leftarrow \sqrt{\bar{\alpha}_{n-1}} \hat{\mathbf{X}}_0 + \sqrt{1 - \bar{\alpha}_{n-1}} \epsilon_n$ 
   ▷ 2D keypoint fitting
7:    $\mathbf{X}_{n-1} \leftarrow \mathbf{X}'_{n-1} - \lambda_{2d} \nabla_{\mathbf{X}_n} \mathcal{L}_2(M_{p2d}(\hat{\mathbf{X}}_0), \mathbf{P}_{2d})$ 
8: end for
9: return  $\hat{\mathbf{X}}_0$ 

```

calculating pair-wise distances on vertices, but it is expensive to use on the fly in training. This necessitates an efficient collision detector to consider collision in training.

We implement a fast ray-tracing-based collision detector following [54]. It casts rays in the normal direction from each vertex and computes the intersections with the mesh faces. Counting the number of intersections serves to find the vertices inside the mesh. Notably, vertices in the armpit region are often detected as collisions, leading to suboptimal solutions (*e.g.*, forcing the arms to move far from the torso). To address this, we restrict collision detection to areas relevant to hands, which are explicit targets in self-contact. Similarly to the data screening of Sec. 3, we apply the loss only to penetrating hand vertices and their corresponding vertices, focusing on hand-specific collisions such as hand-over-hand, hand-over-belly, and hand-over-face.

4.4. Inference

We describe data sampling in the following tasks. We use the DDIM sampling [55] for efficiency.

Random sampling: The trained diffusion model allows random data sampling via the reverse process from random noise ($T: 1000 \rightarrow 1$). Inspired by [27], we reuse the collision loss in Sec. 4.3 as the additional guidance term, *anti-collision guidance*, which avoids the collision during the sampling phase as well. We set the sampling interval to 10. This is used to produce pose generation results of Sec. 5.2.

Single-view pose refinement: Observing self-contact poses inferred from single-view estimators [4, 28, 29, 38], the outputs often include incorrect contact states (*e.g.*, hands not in touch) due to the lack of contact prior, while detected 2D keypoints are aligned well with the given image. To address this, we develop single-view pose refinement, fitting the diffusion prior to the observed 2D keypoints with

Method	FID↓	KID↓ ($\times 10^{-3}$)	Div.↑	Prec.↑	Recall↑	Col. ratio↓
<i>Unconditional generation</i>						
VPoser* [46]	9.43	<u>0.930</u>	3.34	<u>1.0</u>	0.006	1.72
BUDDI* [42]	<u>3.19</u>	1.16	<u>6.36</u>	0.957	<u>0.528</u>	<u>1.32</u>
<i>Shape-conditional generation</i>						
VPoser* [46]	9.16	0.882	3.20	1.0	0.005	1.37
BUDDI* [42]	2.66	1.12	5.59	0.995	0.488	1.47
Ours	1.25	0.430	5.98	0.985	0.708	1.52

Table 2. **Results of self-contact pose generation.** We study sample quality and diversity in generation without (unconditional) or with shape conditioning, evaluated on the *train* split. The notation * indicates the methods adapted to our task.

the estimated initial 3D poses. This does not require additional training compared to score distillation sampling [48] of [27], which is applicable to any 2D/3D estimates.

Our refinement is efficient in sampling with fewer sampling steps (*e.g.*, only use the last 10% steps); see Algorithm 1. We assume given a single-view image, initial SMPL-X pose X_0^{init} , shape I , and 2D keypoints P_{2d} of the COCO-WholeBody format [23] can be estimated by off-the-shelf models, *e.g.*, recent vision foundation models like SMPLer-X [4] and Sapiens [24]. Then we sample data starting from the middle of the steps with diffused X_0^{init} at step n_r (*e.g.*, 100), reducing the number of sampling steps. In each iteration, we use the guidance of 2D keypoint fitting to P_{2d} by computing the gradient of the 2D keypoint error (L2 loss). We set the sampling interval to 1. This is used to produce the refinement results of Sec. 5.3.

We also provide an additional option of *blended pose denoising* in refinement. We find the 2D observation is likely to be partially unavailable or unreliable from in-the-wild videos, *e.g.*, upper-body videos in video conferences do not provide 2D keypoint cues for the lower body. Thus, inspired by image in-painting with diffusion models (*e.g.*, Blended Latent Diffusion [2]), we can only refine poses of interest (*e.g.*, upper body poses) during the reverse process, while the rest of the poses are unchanged. Specifically, given the mask for blending m_p , we replace poses not to be refined with X_0^{init} in each iteration (Line 4 of Algorithm 1), ensuring the convergence to the initial poses. This simple trick helps control the inference process flexibly and enhances the applicability of the diffusion prior.

5. Experiments

We first present our dataset and implementation details in Sec. 5.1, and then provide results for pose generation with random sampling and single-view pose estimation with our refinement method in Secs. 5.2 and 5.3. We also show qualitative results in our proposed dataset and additional results in the supplement.

Method	FID↓	Div.↑	Col. ratio↓
w/o Shape cond.	2.18	5.52	1.92
w/o PASA	1.42	5.74	1.62
w/o Shape rand.	1.27	5.89	1.41
w/o Anti-col.	1.28	6.01	1.85
Ours	1.25	5.98	1.52

Table 3. **Ablation study in our generation.** We compare methods without shape conditioning (Shape cond.), part-aware self-attention (PASA) including a single latent space for pose, shape perturbation (Shape rand.), and anti-collision guidance (Anti-col.).

5.1. Experiment setup

Datasets: We create *train/eval* sets in the Goliath-SC dataset. The *train* set is constructed with the captures with action instructions (*e.g.*, rubbing neck), which comprises 313K poses. Additionally, the *eval* set is designed for single-view pose estimation, featuring *unseen subjects*, which contains 9.7K samples. This is used to test generalizability to unseen subjects where the same action instructions are given as the *train* set.

Implementation details: For generation, we use two self-attention layers with $\text{latent_size}=256$, $\text{depth}=4$, $\text{num_heads}=4$, and set $\lambda_\theta=1$, $\lambda_v=1e-3$, $\lambda_{col}=1e-4$. We set N_s to the full size of 300 to incorporate as much detailed shape information as possible. We set the shape perturbation probability and s_I to 0.3 and $1e-4$. Following [27, 49, 58], we report Fréchet Inception Distance (FID) [17], Kernel Inception Distance (KID) [3], diversity, and precision-recall [52] for the evaluation. We also show the collision ratio of collided vertices over all SMPL-X vertices, using the detector of \mathcal{L}_{col} in Sec. 4.3.

For single-view pose estimation, we prepare different SMPL-X regressors, namely HybrIK-X [28, 29], Hand4Whole [38], and SMPLer-X [4], and use Sapiens [24] for 2D keypoint detection. We report the MPJPE for 3D keypoints of the COCO-WholeBody format [23] in the body-root aligned coordinates (disregarding global rotation and translation). We set λ_{2d} and n_r to 0.01 and 100.

5.2. Pose generation

Tab. 2 shows pose generation results with or without shape conditioning in Goliath-SC. We compare our diffusion method with VAE-based VPoser [46] and diffusion-based BUDDI [42]. We modify these baselines to our task to take the whole body pose parameters as input (aligned to X of Sec. 4.2), denoted as VPoser* and BUDDI*. Details are found in our supplement.

The results show the superiority of our shape-conditional method, remarked by the improvement over the VAE and




Figure 4. **Qualitative results of our generation with shape interpolation.** We interpolate between two shape parameters with the fixed latent code (*i.e.*, starting with the same noise at $t = T$). Our model generates plausible self-contact poses under varying shapes.

diffusion baselines. While VPoser* easily overfits to higher precision, our diffusion method (Ours) has significantly improved recall with a smaller FID score. Our method further surpasses BUDDI*, a method without latent diffusion and part-wise attention, achieving a 53% reduction in FID. Since random outputs exhibit higher diversity and samples lacking contact yield lower col. ratio, maintaining higher diversity and lower col. ratio scores with improved FID and KID (smaller distribution distance) is vital; our method shows a better balance between quality and variety.

Ablation study: Tab. 3 shows the ablation study of our proposed method. Shape-conditional generation surpasses unconditional baselines (w/o Shape cond.), which is also observed in the VPoser* and BUDDI* of Tab. 2. This indicates that the body shapes help to learn self-contact pose distribution effectively. The ablation study shows that part-aware self-attention (PASA), shape perturbation (Shape rand.), and anti-collision guidance (Anti-col.) reduce FID scores consistently over those without each module. Specifically, anti-collision guidance helps reduce col. ratio in test time with improved FID.

Qualitative results: Fig. 4 shows our qualitative results with shape interpolation. When changing the body shapes (from a large to a slim body), the generated poses continuously move on the hand surface while preserving plausible self-contact poses. This indicates that our diffusion model can learn a smooth manifold of self-contact poses with respect to body shape changes.

5.3. Single-view pose estimation and refinement

Analysis on SMPL-X regressors: Tab. 4 shows single-view pose estimation results in Goliath-SC. We first evaluate existing SMPL-X regressors [4, 28, 38] in self-contact scenarios. Hand4Whole and HybriK-X adopt CNN-based backbones (*i.e.*, ResNet [16] and HRNet [56]). SMPLer-X [4] is a foundation model trained on self-contact datasets (*i.e.*, MTP [41] and HumanSC3D [11]). We further fine-tune the model on Goliath-SC, denoted as SMPLer-X[†].

In the predictions from Hand4Whole and HybriK-X, we find frequent failures in handling to place hands in contact, *i.e.*, 2D pose is aligned in the image view but higher

Method	Avg.	Hands	Body	Face
Hand4Whole [38]	126.3	225.8	89.6	78.0
+ 2D fitting	89.5	179.9	64.8	47.1
+ BUDDI* [42]	74.5	109.2	37.8	65.9
+ Ours (w/o Shape cond.)	37.9	74.6	29.4	18.2
+ Ours	35.3	66.6	26.5	18.3
HybriK-X [28]	82.3	99.2	62.8	76.2
+ 2D fitting	51.8	63.4	38.4	50.5
+ BUDDI* [42]	65.0	90.5	36.0	58.8
+ Ours (w/o Shape cond.)	45.9	85.5	32.5	26.3
+ Ours	32.4	58.7	26.1	17.5
SMPLer-X [4]	58.0	98.7	41.6	38.9
SMPLer-X [†]	42.0	56.7	31.9	34.1
+ 2D fitting	41.7	65.7	30.6	31.6
+ BUDDI* [42]	71.7	99.9	36.3	66.4
+ Ours (w/o Shape cond.)	33.7	63.6	26.1	17.4
+ Ours	31.8	54.6	24.7	19.2

Table 4. **Results of single-view pose regression in Goliath-SC.** We evaluate our diffusion-based pose refinement in the *eval* set given initial pose estimation from SMPL-X regressors. We report MPJPE in millimeter on the body-root aligned coordinates. The notation [†] shows fine-tuned results for the dataset.

depth errors are present for hands. Owing to the higher diversity in the training data, SMPLer-X facilitates tracking better poses in our Goliath-SC data, with 58.0 mm error (see Fig. 5). We confirm the state-of-the-art performance in image-based regression with the fine-tuned SMPLer-X[†], exhibiting an overall error reduction to 42.0 mm.

Analysis on refinement: We compare our proposed refinement (Algorithm 1) with conventional refinement. This setup assumes 2D/3D observations are given, namely initial SMPL-X estimates and 2D keypoints. A naive baseline is a simple 2D keypoint fitting with optimization, which is widely used for pseudo-mesh registration of SMPL-X on in-the-wild videos [21, 32, 39]. Since the body and face can be well-constrained with 2D keypoints, their gains are better than the fine-tuned results in the SMPLer-X setting. In contrast, fitting to hands causes implausible 3D hand poses




Figure 5. **Qualitative results of our single-view refinement on Goliath-SC.** Our method successfully refines the initial poses to be valid self-contact for fine-grained poses, such as face touching and two-hand overlap.

though the 2D projection error is minimized. This underscores the need for the model prior in self-contact, particularly to correct hand placement and its local pose.

Next, we evaluate our refinement with diffusion-based priors, namely, our PAPoseDiff, and BUDDI* trained in Tab. 2. BUDDI*, a method without latent diffusion and part-wise self-attention, shows effectiveness when the initial estimation is noisy (Hand4Whole and HybRIK-X), while it has limited refinement capacity when the initialization is reasonable (SMPLer-X). This suggests that additional refinement of well-estimated poses requires more precise modeling of self-contact poses, which is essential to achieve improvements beyond the initial quality.

Our final diffusion prior achieves significantly improved results by reducing overall and part-wise errors across the three settings with different regressors. Our refinement demonstrates less dependency on the initialization and stable performance, as post-refinement converges to lower errors with the varied initializations. Compared to the unconditional prior of our method (w/o Shape cond.), the proposed shape-conditional prior achieves better results, particularly for hands and body. This indicates that the shape-dependent constraint is effective in capturing the relationship between the body and hands, as they have a higher correlation with body shapes than the face.

Qualitative results: Fig. 5 shows qualitative results of our refinement. We find that the initial predictions include ambiguities in contact and depth estimation, *e.g.*, interacting parts are not in contact, especially for fine details, and high-depth errors remain for hands. Our method can correct such failures with the generative prior derived from the contact data only, without requiring knowledge of where to contact.

Discussion: regression vs. generative prior: While recent 3D pose estimators are trained on extensive human data, the state-of-the-art baseline with regression still struggles to estimate self-contact poses of unseen subjects. In contrast, our

approach introduces a novel *generative prior modeling* of self-contact pose distribution, with the body-shape dependent assumption. Our model not only generalizes better to new subjects but also enhances robustness in handling fine-grained self-contact poses. This suggests that our generative prior offers a flexible and scalable solution for self-contact modeling over the regression approach.

6. Conclusion

To highlight challenging self-contact scenarios, we offer a comprehensive self-contact analysis, along with the newly captured **Goliath-SC** dataset with 383K poses and precise body shape registration. We then model the self-contact pose manifold depending on body shapes with the generative diffusion model. Specifically, the latent diffusion with part-aware self-attention learns pose distribution effectively and achieves the best results in pose generation. We further propose single-view pose refinement using the diffusion prior, while fitting to the observed 2D keypoints. Our experiments confirm the successful refinement of self-contact poses and show our superiority over the state-of-the-art diffusion method and the regressive foundation model.

Limitation and future work: We observe that hand-hand interaction is still a challenging scenario in the generation, in which minor interpenetration persists due to higher articulation, as studied in [27]. In addition, not only addressing in-contact scenarios only, but also generalizable modeling to non-contact cases like [41] is an interesting extension. Our success in self-contact modeling opens new avenues to include additional self-contact scenarios, *e.g.*, without scripted action instructions, with various global body poses (*e.g.*, sitting), or in multi-person conversation. Extending the diffusion prior into the temporal dimension or with linguistic contents is promising future work.

Acknowledgment: This work is partially supported by JST ASPIRE Grant Number JPMJAP2303 and JSPS KAKENHI Grant Numbers JP24K02956.

References

- [1] I. Akhter and M. J. Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1446–1455, 2015. 2
- [2] O. Avrahami, O. Fried, and D. Lischinski. Blended latent diffusion. *ACM Transactions on Graphics (ToG)*, 42(4):149:1–149:11, 2023. 6
- [3] M. Binkowski, D. J. Sutherland, M. Arbel, and A. Gretton. Demystifying MMD GANs. In *Proceedings of the International Conference on Learning Representations (ICLR)*. OpenReview.net, 2018. 6
- [4] Z. Cai, W. Yin, A. Zeng, C. Wei, Q. Sun, W. Yanjun, H. E. Pang, H. Mei, M. Zhang, L. Zhang, C. C. Loy, L. Yang, and Z. Liu. Smpler-x: Scaling up expressive human pose and shape estimation. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 2, 5, 6, 7
- [5] Y. Chen, S. K. Dwivedi, M. J. Black, and D. Tzionas. Detecting human-object contact in images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17100–17110, 2023. 2
- [6] P. Chua, C. M. Fang, T. Ohkawa, R. Kushalnagar, S. Nanayakkara, and P. Maes. Emosign: A multimodal dataset for understanding emotions in american sign language. *CoRR*, abs/2505.17090, 2025. 1
- [7] 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 *Proceedings of the International Conference on Learning Representations (ICLR)*. OpenReview.net, 2021. 2
- [8] Z. Fan, O. Taheri, D. Tzionas, M. Kocabas, M. Kaufmann, M. J. Black, and O. Hilliges. ARCTIC: A dataset for dexterous bimanual hand-object manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12943–12954, 2023. 2, 3
- [9] Z. Fan, T. Ohkawa, L. Yang, N. Lin, Z. Zhou, S. Zhou, J. Liang, Z. Gao, X. Zhang, X. Zhang, F. Li, L. Zheng, F. Lu, K. A. Zeid, B. Leibe, J. On, S. Baek, A. Prakash, S. Gupta, K. He, Y. Sato, O. Hilliges, H. J. Chang, and A. Yao. Benchmarks and challenges in pose estimation for egocentric hand interactions with objects. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 15083, pages 428–448, 2024. 3
- [10] M. Fieraru, M. Zanfir, E. Oneata, A. Popa, V. Oлару, and C. Sminchisescu. Three-dimensional reconstruction of human interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7212–7221, 2020. 2, 3
- [11] M. Fieraru, M. Zanfir, E. Oneata, A. Popa, V. Oлару, and C. Sminchisescu. Learning complex 3d human self-contact. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1343–1351, 2021. 1, 2, 3, 7
- [12] P. Grady, C. Tang, C. D. Twigg, M. Vo, S. Brahmabhatt, and C. C. Kemp. ContactOpt: Optimizing contact to improve grasps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1471–1481, 2021. 2
- [13] J. A. Harrigan. Self-touching as an indicator of underlying affect and language processes. *Social science & medicine*, 20(11):1161–1168, 1985. 1
- [14] J. A. Harrigan, K. S. Lucic, D. Kay, A. McLaney, and R. Rosenthal. Effect of expresser role and type of self-touching on observers’ perceptions 1. *Journal of Applied Social Psychology*, 21(7):585–609, 1991. 1
- [15] M. Hassan, V. Choutas, D. Tzionas, and M. J. Black. Resolving 3d human pose ambiguities with 3d scene constraints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2282–2292, 2019. 2
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 7
- [17] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, pages 6626–6637, 2017. 6
- [18] E. S. L. Ho, H. Wang, and T. Komura. A multi-resolution approach for adapting close character interaction. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST)*, pages 97–106, 2014. 2
- [19] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2, 3, 4
- [20] J. Ho and T. Salimans. Classifier-free diffusion guidance. *CoRR*, abs/2207.12598, 2022. 5
- [21] H. Hu, Z. Fan, T. Wu, Y. Xi, S. Lee, G. Pavlakos, and Z. Wang. Expressive gaussian human avatars from monocular RGB video. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 7
- [22] C.-H. P. Huang, H. Yi, M. Höschle, M. Safroshkin, T. Alexiadis, S. Polikovsky, D. Scharstein, and M. J. Black. Capturing and inferring dense full-body human-scene contact. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13264–13275, 2022. 2
- [23] S. Jin, L. Xu, J. Xu, C. Wang, W. Liu, C. Qian, W. Ouyang, and P. Luo. Whole-body human pose estimation in the wild. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 12354, pages 196–214, 2020. 6
- [24] R. Khirodkar, T. M. Bagautdinov, J. Martinez, S. Zhaoen, A. James, P. Selednik, S. Anderson, and S. Saito. Sapiens: Foundation for human vision models. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 15062, pages 206–228, 2024. 6
- [25] R. Khirodkar, J.-T. Song, J. Cao, Z. Luo, and K. Kitani. Harmony4D: A video dataset for in-the-wild close human interactions. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 2
- [26] D. P. Kingma and M. Welling. Auto-encoding variational

- bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014. 2
- [27] J. Lee, S. Saito, G. Nam, M. Sung, and T. Kim. Interhandgen: Two-hand interaction generation via cascaded reverse diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 527–537. IEEE, 2024. 2, 3, 4, 5, 6, 8
- [28] J. Li, S. Bian, C. Xu, Z. Chen, L. Yang, and C. Lu. Hybrik-x: Hybrid analytical-neural inverse kinematics for whole-body mesh recovery. *CoRR*, abs/2304.05690, 2023. 5, 6, 7
- [29] J. Li, C. Xu, Z. Chen, S. Bian, L. Yang, and C. Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3383–3393, 2021. 4, 5, 6
- [30] J. Li, S. Bian, Q. Liu, J. Tang, F. Wang, and C. Lu. NIKI: neural inverse kinematics with invertible neural networks for 3d human pose and shape estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12933–12942. IEEE, 2023. 4
- [31] Z. Li, S. Shimada, B. Schiele, C. Theobalt, and V. Golyanik. MoCapDeform: Monocular 3d human motion capture in deformable scenes. In *Proceedings of the International Conference on 3D Vision (3DV)*, pages 1–11, 2022. 2
- [32] J. Lin, A. Zeng, H. Wang, L. Zhang, and Y. Li. One-stage 3d whole-body mesh recovery with component aware transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21159–21168, 2023. 7
- [33] N. Lin, T. Ohkawa, Y. Huang, M. Zhang, M. Cai, M. Li, R. Furuta, and Y. Sato. Simhand: Mining similar hands for large-scale 3d hand pose pre-training. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. 2
- [34] G. Liu, Y. Wang, Z. Feng, Q. Wu, L. Tang, Y. Gao, Z. Li, S. Cui, J. J. McAuley, Z. Yang, E. P. Xing, and Z. Hu. Unified generation, reconstruction, and representation: Generalized diffusion with adaptive latent encoding-decoding. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2024. 5
- [35] R. Liu, T. Ohkawa, M. Zhang, and Y. Sato. Single-to-dual-view adaptation for egocentric 3d hand pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 677–686, 2024. 2
- [36] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. SMPL: a skinned multi-person linear model. *ACM Transactions on Graphics (ToG)*, 34(6):248:1–248:16, 2015. 2, 3
- [37] J. Martinez, E. Kim, J. Romero, T. Bagautdinov, S. Saito, S.-I. Yu, S. Anderson, M. Zollhöfer, T.-L. Wang, S. Bai, S.-E. Wei, R. Joshi, W. Borsos, T. Simon, J. Saragih, P. Theodosis, A. Greene, A. Josyula, S. M. Maeta, A. I. Jewett, S. Venshtain, C. Heilman, Y.-T. Chen, S. Fu, M. E. A. Elshaer, T. Du, L. Wu, S.-C. Chen, K. Kang, M. Wu, Y. Emad, S. Longay, A. Brewer, H. Shah, J. Booth, T. Koska, K. Haidle, J. C.-H. Hsu, T. Dauer, P. Selednik, T. Godisart, S. Ardisson, M. Cipperly, B. Humberston, L. Farr, B. Hansen, P. Guo, D. Braun, S. Krenn, H. Wen, L. Evans, N. Fadeeva, M. Stewart, G. Schwartz, D. Gupta, G. Moon, K. Guo, Y. Dong, Y. Xu, T. Shiratori, F. A. Prada Nino, B. R. Pires, B. Peng, J. Buffalini, A. Trimble, K. A. A. McPhail, M. R. Schoeller, and Y. Sheikh. Codec avatar studio: Paired human captures for complete, driveable, and generalizable avatars. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*. IEEE, 2024. 2, 3
- [38] G. Moon, H. Choi, and K. M. Lee. Accurate 3d hand pose estimation for whole-body 3d human mesh estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2307–2316, 2022. 5, 6, 7
- [39] G. Moon, T. Shiratori, and S. Saito. Expressive whole-body 3d gaussian avatar. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024. 7
- [40] G. Moon, S.-I. Yu, H. Wen, T. Shiratori, and K. M. Lee. InterHand2.6M: A dataset and baseline for 3D interacting hand pose estimation from a single RGB image. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 548–564, 2020. 1, 2, 3, 4
- [41] L. Müller, A. A. A. Osman, S. Tang, C. P. Huang, and M. J. Black. On self-contact and human pose. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9990–9999. Computer Vision Foundation / IEEE, 2021. 1, 2, 3, 4, 5, 7, 8
- [42] L. Müller, V. Ye, G. Pavlakos, M. J. Black, and A. Kanazawa. Generative proxemics: A prior for 3d social interaction from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9687–9697. IEEE, 2024. 2, 4, 5, 6, 7
- [43] A. Q. Nichol and P. Dhariwal. Improved denoising diffusion probabilistic models. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 139, pages 8162–8171, 2021. 4
- [44] T. Ohkawa, R. Furuta, and Y. Sato. Efficient annotation and learning for 3D hand pose estimation: A survey. *International Journal on Computer Vision (IJCV)*, 131:3193–3206, 2023. 2
- [45] T. Ohkawa, K. He, F. Sener, T. Hodan, L. Tran, and C. Keskin. AssemblyHands: Towards egocentric activity understanding via 3D hand pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12999–13008, 2023. 2, 3
- [46] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. A. Osman, D. Tzionas, and M. J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 2, 3, 4, 6
- [47] B. Pease and A. Pease. *The definitive book of body language: The hidden meaning behind people's gestures and expressions*. Bantam, 2008. 1
- [48] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *Proceedings of the International Conference on Learning Representations (ICLR)*. OpenReview.net, 2023. 6
- [49] S. Raab, I. Leibovitch, P. Li, K. Aberman, O. Sorkine-Hornung, and D. Cohen-Or. MoDi: Unconditional motion synthesis from diverse data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*

- (CVPR), pages 13873–13883. IEEE, 2023. [6](#)
- [50] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. [4](#)
- [51] J. Romero, D. Tzionas, and M. J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics (ToG)*, 36(6):245:1–245:17, 2017. [2](#), [3](#), [4](#)
- [52] M. S. M. Sajjadi, O. Bachem, M. Lucic, O. Bousquet, and S. Gelly. Assessing generative models via precision and recall. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, pages 5234–5243, 2018. [6](#)
- [53] S. Shimada, V. Golyanik, P. Pérez, and C. Theobalt. Decaf: monocular deformation capture for face and hand interactions. *ACM Transactions on Graphics (ToG)*, 42(d), 2023. [1](#), [2](#), [3](#), [4](#)
- [54] B. Smith, C. Wu, H. Wen, P. Peluse, Y. Sheikh, J. K. Hodgins, and T. Shiratori. Constraining dense hand surface tracking with elasticity. *ACM Transactions on Graphics (ToG)*, 39(6):219:1–219:14, 2020. [5](#)
- [55] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. [2](#), [4](#), [5](#)
- [56] K. Sun, B. Xiao, D. Liu, and J. Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5693–5703, 2019. [7](#)
- [57] O. Taheri, N. Ghorbani, M. J. Black, and D. Tzionas. GRAB: A dataset of whole-body human grasping of objects. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 581–600, 2020. [2](#), [3](#)
- [58] G. Tevet, S. Raab, B. Gordon, Y. Shafir, D. Cohen-Or, and A. H. Bermano. Human motion diffusion model. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023. [2](#), [4](#), [6](#)
- [59] D. Uthus, G. Tanzer, and M. Georg. YouTube-ASL: A large-scale, open-domain american sign language-english parallel corpus. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2023. [1](#)
- [60] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, pages 5998–6008, 2017. [4](#)
- [61] X. Wu, T. Liu, C. Li, Y. Ma, Y. Shi, and X. He. Fastgrasp: Efficient grasp synthesis with diffusion. In *Proceedings of the International Conference on 3D Vision (3DV)*, 2025. [5](#)
- [62] H. Xu, E. G. Bazavan, A. Zanfır, W. T. Freeman, R. Sukthankar, and C. Sminchisescu. GHUM & GHUML: generative 3d human shape and articulated pose models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6183–6192, 2020. [3](#)
- [63] S. Xu, Z. Li, Y. Wang, and L. Gui. Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14882–14894. IEEE, 2023. [2](#)
- [64] Y. Yin, C. Guo, M. Kaufmann, J. J. Zarate, J. Song, and O. Hilliges. Hi4D: 4D instance segmentation of close human interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17016–17027, 2023. [2](#), [3](#)
- [65] Z. Yu, S. Huang, Y. Cheng, and T. Birdal. SignAvatars: A large-scale 3d sign language holistic motion dataset and benchmark. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024. [1](#)
- [66] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5745–5753, 2019. [5](#)