

ContactGen: Contact-Guided Interactive 3D Human Generation for Partners

Dongjun Gu, Jaehyeok Shim, Jaehoon Jang, Changwoo Kang, Kyungdon Joo*

Artificial Intelligence Graduate School, UNIST
 {djku1020, jh.shim, erick1997, kangchangwoo, kyungdon}@unist.ac.kr

Abstract

Among various interactions between humans, such as eye contact and gestures, physical interactions by contact can act as an essential moment in understanding human behaviors. Inspired by this fact, given a 3D partner human with the desired interaction label, we introduce a new task of 3D human generation in terms of physical contact. Unlike previous works of interacting with static objects or scenes, a given partner human can have diverse poses and different contact regions according to the type of interaction. To handle this challenge, we propose a novel method of generating interactive 3D humans for a given partner human based on a guided diffusion framework (*ContactGen* in short). Specifically, we newly present a contact prediction module that adaptively estimates potential contact regions between two input humans according to the interaction label. Using the estimated potential contact regions as complementary guidances, we dynamically enforce *ContactGen* to generate interactive 3D humans for a given partner human within a guided diffusion model. We demonstrate *ContactGen* on the CHI3D dataset, where our method generates physically plausible and diverse poses compared to comparison methods. Source code is available at <https://dongjunku.github.io/contactgen>.

1 Introduction

Can you imagine a scenario where a human doing any action without interacting with others? Human actions, instead of being solely independent actions like walking and running, often stem from various interactions. For instance, we constantly interact with diverse objects in our surroundings and engage in daily communication with others (Yi et al. 2023; Hassan et al. 2023; Yi et al. 2022; Zhang et al. 2020). Particularly, interactions among humans are vital for establishing social relationships, exchanging information, sharing emotions, and fulfilling various other aspects of life (Suvilehto et al. 2015; Rudovic et al. 2017; Leclère et al. 2017).

Interactions among humans take on diverse forms, ranging from eye contact or gestures during conversations to physical contact. Among various human interactions, physical contact is common and serves as a crucial communication tool, conveying unambiguous meanings. Even a brief

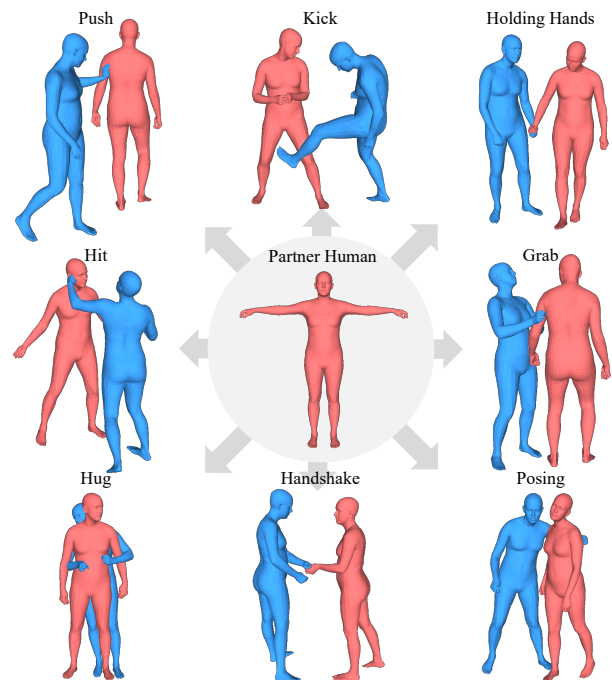


Figure 1: Examples of generated 3D interactive humans using the proposed *ContactGen*. Given a partner human (red-colored human) and interaction label as conditions, *ContactGen* generates various interactive 3D humans (blue-colored human) through a guided DDM, where we focus on physical interaction by contact.

moment of physical contact can provide valuable insights into the nature of the interaction (see Fig. 1). Furthermore, it can act as a significant anchor posture for predicting human motion, as demonstrated in (Wu et al. 2022). Despite the potential of interaction-based approaches, most research works in 3D human generation have concentrated on generating humans from sound (Zhu et al. 2023), scene (Kim et al. 2023), text (Petrovich, Black, and Varol 2022; Kim, Kim, and Choi 2023) and prior movement inputs (Lee, Gyeongsik, and Mu 2023), overlooking the aspect of the interaction between humans. Instead, physical interactions are primarily addressed in hand-object interaction (Grady et al. 2021; Karunratanakul et al. 2020) and scene-based interaction studies (Zhang et al. 2020; Hassan et al. 2021).

*Corresponding author.

In this work, we investigate a new problem of a 3D human generation that maintains natural physical contact with a given partner human based on input interaction labels. Unlike previous works on interactions between humans and objects (or scenes) (Zhao et al. 2022; Huang et al. 2023; Yi et al. 2023), generating 3D human interacting with a given partner human involves several challenges. Specifically, the target of interaction (*i.e.*, partner humans) is non-rigid and dynamic, leading to numerous possible combinations of pose and shape. In addition, compared to objects (or scenes) that have a clearly defined purpose, humans can have different physical contact regions based on the type of interactions. For instance, when two humans interact, such as a greeting, the points of contact can exhibit considerable variation. They might opt for a handshake, a hug, or even a cheek-to-cheek kiss. Each of these interactions has different physical contact regions, making the interaction more complex and dynamic than with static objects or scenes.

To overcome these challenges, we present a novel approach to generate an interactive 3D human that maintains natural physical contact with a given partner human according to the desired interaction label. We call this proposed method `ContactGen` (see Fig. 2). `ContactGen` basically adopts a guided denoising diffusion model (DDM) (Dhariwal and Nichol 2021; Huang et al. 2023), which enables to consider additional optimization as guidance during the sampling process of DDM. In particular, we aim to exploit physical interaction, such as contact, as guidance within a guided DDM. To this end, we propose a contact prediction module that adaptively estimates the probability of physical contact for each body part between two input humans according to the interaction label. Based on the predicted contact regions, we can seamlessly incorporate interactive optimization related to physical contact into DDM, which allows us to synthesize interactions that involve meaningful and accurate points of contact, enhancing the realism and coherence of the generated human interactions. We demonstrate the effectiveness and realism of our method on the modified version of the CHI3D dataset (Fieraru et al. 2020), where `ContactGen` shows a promising solution for generating rich and dynamic human interactions.

In summary, our main contributions are:

- We introduce a new and exciting task of generating interactive 3D humans for a given partner human and desired interaction label.
- We propose a novel guided DDM for generating interactive 3D humans, `ContactGen`, that considers physical interactions, such as contact during sampling, leading to physically feasible human interactions.
- We present a contact prediction module that adaptively predicts potential interaction regions between humans, which allows us to optimize valid interaction regions within a guided DDM.

2 Related Work

Our method of generating physical interactions by contact between humans has not been studied before, but related

studies do exist. To identify our work, we summarize relevant research on human action generation, physics-based prediction, and the human-to-human interaction dataset.

Human action generation. Since human action recognition and prediction are crucial for understanding human behavior, many studies have been actively conducted in the computer vision field (Tran et al. 2018; Shi et al. 2019; Wang et al. 2021; Duan et al. 2022). Inspired by the success of the generation model in the 2D domain, many works have been proposed to synthesize 3D human motion from various modalities. Given an action label, Action2Motion (Guo et al. 2020) and ACTOR (Petrovich, Black, and Varol 2021) adopt a variational autoencoder (VAE) architecture (Kingma and Welling 2013) that learns latent space representing human motions. Beyond this simple conditioning, TEMOS (Petrovich, Black, and Varol 2022) and TEACH (Athanasidou et al. 2022) take text descriptions that contain multiple actions as input, creating a natural motion sequence. However, these VAE-based approaches, which are prone to posterior collapse, have limitations of human motion diversity.

To alleviate this issue, Tevet *et al.* (Tevet et al. 2023) propose MDM, the first time to apply a diffusion model (Ho, Jain, and Abbeel 2020) to human motion synthesis. PhysDiff (Yuan et al. 2022) extends MDM to enforce the generated motion sample to be physically plausible by projecting physical constraints from the physics simulator to the diffusion process. Although the above methods show great performance, they only focus on synthesizing a single human motion, not considering interaction between humans. On the other hand, we propose a diffusion-based framework for generating a 3D human who comes into contact with various body parts of a given partner based on input interaction.

Physics-based prediction in 3D. Physical contact-based interactions are most extensively covered in scene-aware 3D human generation (Zhang et al. 2020; Zhao et al. 2022; Huang et al. 2023), object-aware 3D hand pose estimation (Taheiri et al. 2022; Wu et al. 2022), and motion re-targeting (Van Gemeren, Poppe, and Veltkamp 2016; Mehta et al. 2018; Guo et al. 2022). Several research works use proximity between human and a given scene to induce physical affordance and use negative SDF to prohibit interpenetration (Zhang et al. 2020). Other methods use object alignment by point normals (Kim et al. 2023; Grady et al. 2021), and point-wise embedded closest semantics (Kim et al. 2023). In contrast to previous studies, human-human interactions are challenging due to their dynamic and non-rigid nature. As a result, research on human-human interaction is limited to the subject of 3D motion re-targeting and is given relatively little attention. To handle this challenge, we propose a contact prediction module that can predict potential contact regions between each human body during the sampling process and provide guidance for diffusion.

Dataset of human-human interaction in 3D. Contrary to many studies (Van Gemeren, Poppe, and Veltkamp 2016; Mehta et al. 2018; Guo et al. 2022) that introduce 3D datasets addressing human interactions, they represented the 3D human solely based on joint locations, omitting con-

Algorithm 1: Training

Input: learning rate η
repeat
 $(\mathbf{x}_h^0, \mathbf{x}_p, \mathbf{l}) \sim q(\mathbf{x}_h^0, \mathbf{x}_p, \mathbf{l})$
 $t \sim \mathcal{U}(\{1, \dots, T\})$
 $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 $\mathbf{x}_h^t = \sqrt{\hat{\alpha}_t} \mathbf{x}_h^0 + \sqrt{1 - \hat{\alpha}_t} \epsilon$
 $\tilde{\epsilon} = \epsilon_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l})$ \triangleright noise prediction
 $\theta \leftarrow \eta \nabla_\theta \mathcal{L}(\tilde{\epsilon}, \epsilon)$
until ϵ_θ is converged

$$p_\theta(\mathbf{x}_h^{t-1} | \mathbf{x}_h^t, \mathbf{x}_p, \mathbf{l}) = \mathcal{N}(\mathbf{x}_h^{t-1}; \mu_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l}), \sigma_t^2 \mathbf{I}), \quad (5)$$

where $p(\mathbf{x}_h^T)$ is the standard Gaussian distribution, $\mu_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l})$ is the mean induced from the noise prediction network $\epsilon_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l})$ parameterized by θ , and σ_t is a time-step dependant variance. By repeating this process, we can eventually generate an interactive 3D human \mathbf{x}_h^0 .

Training objective. Noise prediction network ϵ_θ is trained with MSE between predicted and ground truth noises:

$$\mathcal{L}(\tilde{\epsilon}, \epsilon) = \sum_i \frac{1}{N_i} \|\tilde{\epsilon}_i - \epsilon_i\|_2^2, \quad (6)$$

where $\tilde{\epsilon}_i$ and ϵ_i indicate parts of predicted noise $\tilde{\epsilon} = \epsilon_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l})$ and ground truth noise ϵ , respectively, and N_i means the dimension of ϵ_i . We split the predicted noise $\tilde{\epsilon}$ into four parts of $\tilde{\epsilon}_i$ according to their characteristics, such as translation, rotation, body pose, and hand poses. By dividing, we aim to lead our model focusing on each human characteristic with the same weight. The detailed procedure for training conditional DDM is available in Alg. 1.

Guided sampling. Beyond the conditional sampling from $p_\theta(\mathbf{x}_h^{0:T} | \mathbf{x}_p, \mathbf{l})$ by Eq. (4), we introduce a guided sampling that allows us to sample from $p_\theta(\mathbf{x}_h^{0:T} | \mathbf{x}_p, \mathbf{l}, \mathcal{O})$, while satisfying a certain objective \mathcal{O} with the learned $p_\theta(\mathbf{x}_h^{0:T} | \mathbf{x}_p, \mathbf{l})$.

First of all, the distribution $p_\theta(\mathbf{x}_h^{0:T} | \mathbf{x}_p, \mathbf{l}, \mathcal{O})$ we desire to sample can be formulated using the form of a Markov chain:

$$p_\theta(\mathbf{x}_h^{0:T} | \mathbf{x}_p, \mathbf{l}, \mathcal{O}) = p(\mathbf{x}_h^T) \prod_{t=1}^T p_\theta(\mathbf{x}_h^{t-1} | \mathbf{x}_h^t, \mathbf{x}_p, \mathbf{l}, \mathcal{O}). \quad (7)$$

Here, we use the reverse process $q'(\mathbf{x}_h^{t-1} | \mathbf{x}_h^t, \mathbf{x}_h^0)$ defined in denoising diffusion implicit models (DDIM) (Song, Meng, and Ermon 2020). In each iteration of this reverse process, the model predicts \mathbf{x}_h^0 from a noisy input \mathbf{x}_h^t and the predicted noise:

$$f_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l}) = (\mathbf{x}_h^t - \sqrt{1 - \hat{\alpha}_t} \epsilon_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l})) / \sqrt{\hat{\alpha}_t}. \quad (8)$$

Then, the model subsequently generates a sample \mathbf{x}_h^{t-1} using the reverse process $q'(\mathbf{x}_h^{t-1} | \mathbf{x}_h^t, \mathbf{x}_h^0)$.

Concretely, $p_\theta(\mathbf{x}_h^{t-1} | \mathbf{x}_h^t, \mathbf{x}_p, \mathbf{l}, \mathcal{O})$ can be represented as:

$$p_\theta(\mathbf{x}_h^{t-1} | \mathbf{x}_h^t, \mathbf{x}_p, \mathbf{l}, \mathcal{O}) = q'(\mathbf{x}_h^{t-1} | \mathbf{x}_h^t, f_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l}, \mathcal{O})), \quad (9)$$

where $f_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l}, \mathcal{O})$ indicates objective-aware denoised human prediction function based on $f_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l})$ and \mathcal{O} . To consider a certain objective \mathcal{O} , we introduce guidance as:

$$\mathbf{g} = \nabla_{\mathbf{x}_h} \mathcal{O}(\mathbf{x}_h, \mathbf{x}_p, \mathbf{l}) |_{\mathbf{x}_h = f_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l})}. \quad (10)$$




Figure 3: Illustration of contact prediction module. We use the contact prediction module to predict the part-wise contact probability map \mathbf{C} between interacting humans. After estimating \mathbf{C} , a set of potential contact regions \mathbb{C} (right) can be obtained from a certain threshold τ .

Algorithm 2: Guided Sampling

Input: partner 3D human \mathbf{x}_p , interaction label \mathbf{l}
Output: generated interactive 3D human \mathbf{x}_h^{t-1}
 $\mathbf{x}_h^T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
for $t = T, \dots, 1$ **do**
 $\tilde{\mathbf{x}}_h^0 = f_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l})$ \triangleright denoised sample by Eq. (8)
 $\mathbf{g} = \nabla_{\mathbf{x}_h} \mathcal{O}(\mathbf{x}_h, \mathbf{x}_p, \mathbf{l}) |_{\mathbf{x}_h = \tilde{\mathbf{x}}_h^0}$ \triangleright guidance by Eq. (10)
 $\hat{\mathbf{x}}_h^0 = \tilde{\mathbf{x}}_h^0 - \lambda \odot \mathbf{g}$ \triangleright interactive optimization by Eq. (11)
 $\mathbf{x}_h^{t-1} = \sqrt{\hat{\alpha}_{t-1}} \tilde{\mathbf{x}}_h^0 + \frac{\sqrt{1 - \hat{\alpha}_{t-1}}}{\sqrt{1 - \hat{\alpha}_t}} (\mathbf{x}_h^t - \sqrt{\hat{\alpha}_t} \tilde{\mathbf{x}}_h^0)$ \triangleright reverse process of DDIM
end for

Then, we define $f_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l}, \mathcal{O})$ as:

$$f_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l}, \mathcal{O}) = f_\theta(\mathbf{x}_h^t, t, \mathbf{x}_p, \mathbf{l}) - \lambda \odot \mathbf{g}, \quad (11)$$

where λ denotes the scaling vector for the guidance and \odot indicates the element-wise product. In Eq. (11), we apply the guidance \mathbf{g} in a compact form for iterative computation.

In this work, we consider physical contact between humans as the objective \mathcal{O} to provide guidance (*cf.* Sec. 3.2 for details of guidance objective). Specifically, in each iteration of the reverse process described above, we estimate noise-less human $\tilde{\mathbf{x}}_h^0$ from the estimated noise $\tilde{\epsilon}$ and a noisy input \mathbf{x}_h^t in Eq. (8). Then, $\tilde{\mathbf{x}}_h^0$ is further optimized to satisfy the objective \mathcal{O} , leading to optimized human $\hat{\mathbf{x}}_h^0$ in Eq. (11).

Note that instead of directly optimizing \mathbf{x}_h^{t-1} , as in (Huang et al. 2023), ContactGen indirectly optimizes \mathbf{x}_h^{t-1} by optimizing $\tilde{\mathbf{x}}_h^0$ with the DDIM sampling process. As a result, ContactGen can effectively estimate contact regions from less noisy humans $\tilde{\mathbf{x}}_h^0$ instead of noisy ones \mathbf{x}_h^{t-1} . The detailed guided sampling algorithm is summarized in Alg. 2.

3.2 Interactive Optimization

We design optimization objectives that help the diffusion process by providing a guiding direction in terms of physical interaction by contact.

Contact objective. We exploit the contact objective to promote physical contact between a partner human and an interactive human so that their interaction can be physically natural. The contact objective is formulated as:

$$\mathcal{O}_{\text{contact}}(\mathbf{x}_h, \mathbf{x}_p, \mathbf{l}) = \sum_{(r_i, r_j) \in \mathbb{C}} d_{CD}(V_h^{r_i}, V_p^{r_j}), \quad (12)$$

where $d_{CD}(\cdot, \cdot)$ is the Chamfer distance between two sets of vertices, \mathbb{C} denotes the contact region set predicted by the contact prediction module (*cf.* Sec. 3.3), and $V_h^{r_i}$ and $V_p^{r_j}$ are the set of vertices of interactive human and partner human, belonging to regions r_i and r_j , respectively.

Interactive optimization. Based on the above objective, we expand $\mathcal{O}(\mathbf{x}_h, \mathbf{x}_p, \mathbf{l})$ in Eq. (10) into customized objectives for interaction, following (Huang et al. 2023):

$$\mathcal{O}(\mathbf{x}_h, \mathbf{x}_p, \mathbf{l}) = \mathcal{O}_{contact}(\mathbf{x}_h, \mathbf{x}_p, \mathbf{l}). \quad (13)$$

3.3 Contact Prediction Module

To adaptively provide potential contact regions within the guided DDM, we propose a new transformer-based contact prediction module that estimates the contact probability map between each contact region of a given partner human \mathbf{x}_p and an interactive human $\tilde{\mathbf{x}}_h^0$ (see Fig. 3).

To this end, we divide the SMPL-X mesh of humans into N_{reg} number of regions with regard to vertex indices following CI3D (Fieraru et al. 2020) and predict a part-wise contact map $\mathbf{C} \in \{0, 1\}^{N_{reg} \times N_{reg}}$ (we set N_{reg} as 75). Concretely, the contact prediction module takes as input each center location of N_{reg} regions of partner and interactive humans and passes through a sequence of transformers (Vaswani et al. 2017) composed of self-attention and cross-attention to compute each feature vector. Then, we estimate part-wise contact probability map \mathbf{C} using our contact prediction module by mutually extracting features from both SMPL-X parameters of the partner and the denoised interactive human. The contact prediction module is trained with binary-cross-entropy objective function:

$$\mathcal{L}_{CE} = -\frac{1}{N_{reg}^2} \sum_i \sum_j \left(\mathbf{C}_{i,j} \log(\tilde{\mathbf{C}}_{i,j}) + (1 - \mathbf{C}_{i,j}) \log(1 - \tilde{\mathbf{C}}_{i,j}) \right), \quad (14)$$

where $\mathbf{C}_{i,j}$ indicates the probability of contact occurring between i -th and j -th regions from interactive human and partner, respectively, and \mathbf{C} and $\tilde{\mathbf{C}}$ are the ground truth and estimated contact prediction maps, respectively.

After estimating $\tilde{\mathbf{C}}$, we compute a contact region set \mathbb{C} that contains a set of contact region pairs corresponding to probability over a certain threshold τ (we set τ as 0.5). By predicting \mathbb{C} for $\tilde{\mathbf{x}}_h^0$ and \mathbf{x}_p at each time t in an adaptive way, we can provide dynamic contact regions within the guided DDM. That is, ContactGen can sample diverse interactive humans of high quality by avoiding the conflict between diffusion direction and guiding direction.

4 Experiments

4.1 Implementation Details

The training policy of our guided DDM basically follows DDM (Ho, Jain, and Abbeel 2020). We selectively adopt SMPL-X representation (Pavlakos et al. 2019), a differentiable function that maps from a set of low-dimensional body parameters to a 3D human mesh. We use the Adam optimizer (Kingma and Ba 2015) to train our networks. For the training procedure, noise is sampled from the linear noise

scheduler initialized with $\beta_0 = 5e - 6$, $\beta_T = 5e - 3$, and diffusion step $T = 1000$. In the sampling phase, we utilize classifier-free guidance (Ho and Salimans 2022) by hiding partner and both partner and interaction label conditions in the training and sampling of our network. More details of implementation are described in supplementary materials.

4.2 Dataset

We basically use CHI3D (Fieraru et al. 2020), which is a 3D motion capture dataset of 8 close human interaction scenarios: push, posing, grab, hug, kick, handshake, holding hands and hit. For each interaction, CHI3D provides various sequences (*i.e.*, video, not contact interaction moment) containing physical contact events and corresponding 3D human models, such as SMPL-X (Pavlakos et al. 2019) and GHUM (Xu et al. 2020). To generate a 3D human in contact with a given partner based on input interaction, we proceed with several pre-processing on CHI3D to acquire a valid dataset for training interactive human generation in terms of physical contact.

Concretely, we perform frame sampling that captures the moment of contact between human subjects. We then normalize the rotation parameters so that all partners head in the same direction, and transform the rotation and translation parameters of the actor to maintain the interaction. In addition, we augment the 3D human representation in terms of the body parts of a human with symmetry to ensure the diversity of generation results. For example, we flip the two people holding their right hands so that they hold their left hands. We also utilize contact annotation between human subjects to train the contact prediction module. We will release additional code to preprocess CHI3D dataset for research purposes.

4.3 Evaluation Metrics

We evaluate and compare ContactGen in terms of sample quality and physical plausibility. To this end, we propose Fréchet human interaction distance (FHID) to assess sample quality by comparing distributional discrepancy between a set of generated samples and ground truth, motivated by FID (Heusel et al. 2017). To achieve this, we train an MLP-based classifier network with the modified CHI3D dataset to estimate interaction label from given partner and human SMPL-X parameters. The FHID metric is measured with the MLP feature of the network and this network is also used for top-1 accuracy metrics.

For physical plausibility, we adopt the contact score based on the Chamfer distance and non-collision score. In the case of contact score, we count body part-wise contact frequencies of interactive humans to calculate potential contact regions corresponding to each interaction label. Then, the contact score is computed by unidirectional Chamfer distance from interactive human to partner human. Additionally, the non-collision score, which refers to the ratio of vertices with positive SDF values for the entire vertex of the interactive human is measured following (Zhang et al. 2020; Huang et al. 2023).




Figure 4: Qualitative evaluation with comparison methods on the modified CHI3D dataset. We visualize generated interactive humans (blue-colored humans) alongside their given partners (red-colored humans). The first three rows show generated humans by *ContactGen* according to different partner poses, where *ContactGen* shows diverse and plausible generations in terms of physical contact. For comparison with SAGA and IMM, we visualize generated humans under the same given partner conditions (last three rows). Compared to our results showing visually satisfactory outcomes, humans generated by SAGA and IMM show improper pose and penetration.

4.4 Comparison Methods

Since there are no exact methods that generate an interactive 3D human at the moment of physical contact, we modify and re-implement two related methods for comparisons. We here briefly discuss how we modify them at a high level.

Interaction Mix and Match (Goel, Men, and Ho 2022).

Interaction Mix and Match (IMM) synthesizes interactive human motions, not the contact moment of two humans. In addition, IMM generates humans in the form of 3D skeletons, not 3D body meshes like SMPL-X (Pavlakos et al. 2019). For a fair comparison, we thus pick out 15 joints among the body joints of SMPL-X provided by the CHI3D dataset (Fieraru et al. 2020) for compatibility. We then train IMM on the modified CHI3D dataset and utilize SMPLify (Bogo et al. 2016) to render the predicted output into the 3D human mesh. Lastly, we select frames whose distance between two humans is closer than 2cm as interactive 3D humans from the generated motions.

SAGA (Wu et al. 2022). Originally, SAGA is designed to generate human motions conditioned on human-object contacts. To this end, SAGA consists of two-stage: 1) human generation focusing on human-object contacts and 2) motion generation conditioned on the generated human. For comparison, we adopt and train the first stage of SAGA for our task on the modified CHI3D dataset.

4.5 Qualitative Evaluation

We conduct qualitative experiments on the modified CI3D dataset to compare our method with SAGA and IMM (see Fig. 4). SAGA and IMM have difficulties generating physically plausible contact and modeling the dynamic nature of various interactions. For instance, IMM synthesizes a human that has interpenetration with the given partner without pushing. Unlike comparison methods, *ContactGen* generates physically plausible interactive humans while taking diverse poses corresponding to each type of interaction.

Furthermore, we additionally compare ours with SAGA

Interaction	IMM				SAGA				Ours			
	FHID	top-1	contact	non-col	FHID	top-1	contact	non-col	FHID↓	top-1↑	contact↓	non-col↑
Push	51.00	0.00	17.52	94.08	34.54	18.43	7.44	95.10	56.78	87.52	3.61	94.77
Posing	174.78	11.90	14.81	93.54	53.23	62.50	4.33	92.15	71.70	84.15	8.58	92.72
Grab	59.23	92.78	13.26	96.98	46.70	84.41	5.95	97.00	28.88	91.32	4.76	96.50
Hug	235.84	18.72	12.45	88.17	57.16	93.38	3.16	76.21	21.43	99.75	5.96	78.71
Kick	29.89	66.67	10.01	99.70	42.04	30.17	4.62	97.41	46.55	95.69	11.98	98.49
Handshake	146.90	11.90	16.14	99.10	34.39	66.12	4.14	98.16	12.97	92.12	7.34	99.13
Hit	43.01	61.40	10.94	98.52	44.10	50.60	6.22	96.98	20.38	67.47	5.16	96.55
Holding Hands	499.84	0.00	15.17	96.80	54.07	73.93	4.27	96.41	26.20	99.02	6.42	97.12
All	92.25	34.49	14.24	94.41	24.26	67.47	5.23	94.62	6.15	90.30	6.07	94.89

Table 1: Quantitative comparison on the modified CHI3D dataset.




Figure 5: Comparison of predicted contact regions. In our method, the same colored regions describe the estimated contact regions between two humans. We can observe plausible contact regions according to the interaction label. In SAGA, yellow indicates the estimated contactable vertices between humans, but they do not know the correspondences.

Method	FHID	top-1	contact	non-col
w/o contact obj.	6.91	90.70	15.11	96.48
w/ contact obj.	6.15	90.30	6.07	94.89

Table 2: Ablation study for the interactive optimization.

in terms of contact prediction (see Fig. 5). In the case of SAGA, it just estimates the probability of vertex-wise contact, not the corresponding contact vertex between humans. Thus, SAGA that minimizing distances between entire feasible regions limits its capacity to capture dynamic interactions. However, *ContactGen* can predict more diverse kinematic relations between interacting humans by estimating part-wise information.

4.6 Quantitative Evaluation

We quantitatively evaluate comparison methods and ours on various metrics, such as FHID, top-1 accuracy, contact, and non-collision metrics for both per-category and whole test datasets (see Table 1). *ContactGen* shows better results than comparison methods in most categories and metrics. Especially in FHID and top-1 accuracy metrics, our method outperforms comparison methods with large margins. It means that our method can generate high-quality

samples compared to comparison methods.

SAGA, in particular, shows the best contact scores due to its post-optimization step encouraging contact. However, it has limitations in extensively modeling human-to-human interaction, a crucial element for generating natural humans based on the partner. This tendency of SAGA is apparent in both top-1 accuracy and FHID. On the other hand, despite slightly lower contact scores of *ContactGen* compared to SAGA, *ContactGen* outperforms all other scores of SAGA, indicating that *ContactGen* generates more plausible humans maintaining similar contact quality to SAGA. In the case of IMM, it shows comparable non-collision scores but performs worse in all other metrics. This can be attributed to the fact that IMM does not explicitly model close interaction and contact between humans, allowing for the generation of humans who maintain distance from their partners rather than engaging in close interaction with them.

4.7 Ablation Study

We conduct an ablation study to validate the effectiveness of contact objective in the interactive optimization phase. As shown in Table 2, *ContactGen* with or without contact objective shows comparable performance in general (better for FHID, and comparable to accuracy and non-collision). However, when we measured the contact metric, *ContactGen* with contact objective outperforms one without contact objective with two times better performance. From this result, we deduce that the contact objective strongly drives the predicted contact regions to be close, resulting in high-quality samples that maintain physical contact.

5 Conclusion

In this work, we have proposed *ContactGen*, a novel guided DDM for generating interactive 3D humans in terms of physical contact for a given partner human and desired interaction label. Especially for dealing with dynamic and non-rigid partners according to desired interaction, we propose the contact prediction module that flexibly estimates potential contact regions. This contact prediction module provides valid contact regions to optimize valid interaction regions, which leads the guided DDM to consider physical interaction as guidance. Thus, *ContactGen* offers a promising solution for generating rich and dynamic human interactions, which involve meaningful and accurate points of contact, enhancing the realism and coherence of the generated human interactions.

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2022-0-00612, Geometric and Physical Commonsense Reasoning based Behavior Intelligence for Embodied AI and No.2020-0-01336, Artificial Intelligence Graduate School Program (UNIST)), and the Ministry of Science and ICT and NIPA through the HPC Support Project.

References

- Athanasios, N.; Petrovich, M.; Black, M. J.; and Varol, G. 2022. Teach: Temporal action composition for 3d humans. In *3DV*.
- Bogo, F.; Kanazawa, A.; Lassner, C.; Gehler, P.; Romero, J.; and Black, M. J. 2016. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *ECCV*.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *NeurIPS*.
- Duan, H.; Zhao, Y.; Chen, K.; Lin, D.; and Dai, B. 2022. Revisiting Skeleton-Based Action Recognition. In *CVPR*.
- Fieraru, M.; Zanfir, M.; Oneata, E.; Popa, A.-I.; Oлару, V.; and Sminchisescu, C. 2020. Three-dimensional reconstruction of human interactions. In *CVPR*.
- Goel, A.; Men, Q.; and Ho, E. S. 2022. Interaction Mix and Match: Synthesizing Close Interaction using Conditional Hierarchical GAN with Multi-Hot Class Embedding. In *Computer Graphics Forum*.
- Grady, P.; Tang, C.; Twigg, C. D.; Vo, M.; Brahmabhatt, S.; and Kemp, C. C. 2021. ContactOpt: Optimizing Contact To Improve Grasps. In *CVPR*.
- Guo, C.; Zuo, X.; Wang, S.; Zou, S.; Sun, Q.; Deng, A.; Gong, M.; and Cheng, L. 2020. Action2motion: Conditioned generation of 3d human motions. In *ACM MM*.
- Guo, W.; Bie, X.; Alameda-Pineda, X.; and Moreno-Noguer, F. 2022. Multi-person extreme motion prediction. In *CVPR*.
- Hassan, M.; Ghosh, P.; Tesch, J.; Tzionas, D.; and Black, M. J. 2021. Populating 3D scenes by learning human-scene interaction. In *CVPR*.
- Hassan, M.; Guo, Y.; Wang, T.; Black, M.; Fidler, S.; and Peng, X. B. 2023. Synthesizing Physical Character-scene Interactions. In *SIGGRAPH*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *NeurIPS*.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Huang, S.; Wang, Z.; Li, P.; Jia, B.; Liu, T.; Zhu, Y.; Liang, W.; and Zhu, S.-C. 2023. Diffusion-Based Generation, Optimization, and Planning in 3D Scenes. In *CVPR*.
- Karunratanakul, K.; Yang, J.; Zhang, Y.; Black, M. J.; Muan-det, K.; and Tang, S. 2020. Grasping Field: Learning Implicit Representations for Human Grasps. *2020 International Conference on 3D Vision (3DV)*, 333–344.
- Kim, J.; Kim, J.; and Choi, S. 2023. Flame: Free-form language-based motion synthesis & editing. In *AAAI*.
- Kim, M.; Kang, C.; Park, J.; and Joo, K. 2023. Pose-Guided 3D Human Generation in Indoor Scene. In *AAAI*.
- Kingma, D.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Leclère, C.; Avril, M.; Viaux, S.; Bodeau, N.; Achard, C.; Missonnier, S.; Keren, M.; Chetouani, M.; and Cohen, D. 2017. Interaction and behaviour imaging: a novel method to measure mother–infant interaction using video 3D reconstruction. *Translational Psychiatry*, 6: 05.
- Lee, T.; Gyeongsik, M.; and Mu, L. K. 2023. MultiAct: Long-Term 3D Human Motion Generation from Multiple Action Labels. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(1): 1231–1239.
- Mehta, D.; Sotnychenko, O.; Mueller, F.; Xu, W.; Sridhar, S.; Pons-Moll, G.; and Theobalt, C. 2018. Single-shot multi-person 3d pose estimation from monocular rgb. In *3DV*.
- Muller, L.; Ye, V.; Pavlakos, G.; Black, M.; and Kanazawa, A. 2023. Generative Proxemics: A Prior for 3D Social Interaction from Images. *arXiv preprint 2306.09337v1*.
- Pavlakos, G.; Choutas, V.; Ghorbani, N.; Bolkart, T.; Osman, A. A.; Tzionas, D.; and Black, M. J. 2019. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*.
- Petrovich, M.; Black, M. J.; and Varol, G. 2021. Action-conditioned 3D human motion synthesis with transformer VAE. In *ICCV*.
- Petrovich, M.; Black, M. J.; and Varol, G. 2022. TEMOS: Generating diverse human motions from textual descriptions. In *ECCV*.
- Rudovic, O.; Lee, J.; Mascarell Maricic, L.; Schuller, B.; and Picard, R. 2017. Measuring Engagement in Robot-Assisted Autism Therapy: A Cross-Cultural Study. *Frontiers in Robotics and AI*, 4: 36.
- Shi, L.; Zhang, Y.; Cheng, J.; and Lu, H. 2019. Skeleton-Based Action Recognition With Directed Graph Neural Networks. In *CVPR*.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*.
- Suvilehto, J. T.; Glerean, E.; Dunbar, R. I. M.; Hari, R.; and Nummenmaa, L. 2015. Topography of social touching depends on emotional bonds between humans. *Proceedings of the National Academy of Sciences*, 112(45): 13811–13816.
- Taheri, O.; Choutas, V.; Black, M. J.; and Tzionas, D. 2022. GOAL: Generating 4D whole-body motion for hand-object grasping. In *CVPR*.
- Tevet, G.; Raab, S.; Gordon, B.; Shafir, Y.; Cohen-or, D.; and Bermano, A. H. 2023. Human Motion Diffusion Model. In *ICLR*.
- Tran, D.; Wang, H.; Torresani, L.; Ray, J.; LeCun, Y.; and Paluri, M. 2018. A Closer Look at Spatiotemporal Convolutions for Action Recognition. In *CVPR*.

- Van Gemeren, C.; Poppe, R.; and Veltkamp, R. C. 2016. Spatio-temporal detection of fine-grained dyadic human interactions. In *Human Behavior Understanding: 7th International Workshop, HBU 2016, Amsterdam, The Netherlands, October 16, 2016, Proceedings 7*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *NeurIPS*, 30.
- Wang, L.; Tong, Z.; Ji, B.; and Wu, G. 2021. TDN: Temporal Difference Networks for Efficient Action Recognition. In *CVPR*.
- Wu, Y.; Wang, J.; Zhang, Y.; Zhang, S.; Hilliges, O.; Yu, F.; and Tang, S. 2022. Saga: Stochastic whole-body grasping with contact. In *ECCV*.
- Xu, H.; Bazavan, E. G.; Zafir, A.; Freeman, W. T.; Sukthankar, R.; and Sminchisescu, C. 2020. Ghum & ghuml: Generative 3d human shape and articulated pose models. In *CVPR*.
- Yi, H.; Huang, C.-H. P.; Tripathi, S.; Hering, L.; Thies, J.; and Black, M. J. 2023. MIME: Human-Aware 3D Scene Generation. In *CVPR*.
- Yi, H.; Huang, C.-H. P.; Tzionas, D.; Kocabas, M.; Hassan, M.; Tang, S.; Thies, J.; and Black, M. J. 2022. Human-Aware Object Placement for Visual Environment Reconstruction.
- Yin, Y.; Guo, C.; Kaufmann, M.; Zarate, J.; Song, J.; and Hilliges, O. 2023. Hi4D: 4D Instance Segmentation of Close Human Interaction. In *CVPR*.
- Yuan, Y.; Song, J.; Iqbal, U.; Vahdat, A.; and Kautz, J. 2022. PhysDiff: Physics-Guided Human Motion Diffusion Model. *arXiv preprint arXiv:2212.02500*.
- Zhang, Y.; Hassan, M.; Neumann, H.; Black, M. J.; and Tang, S. 2020. Generating 3D People in Scenes Without People. In *CVPR*.
- Zhao, K.; Wang, S.; Zhang, Y.; Beeler, T.; and Tang, S. 2022. Compositional human-scene interaction synthesis with semantic control. In *ECCV*.
- Zhu, L.; Liu, X.; Liu, X.; Qian, R.; Liu, Z.; and Yu, L. 2023. Taming Diffusion Models for Audio-Driven Co-Speech Gesture Generation. In *CVPR*.