PLACE: Proximity Learning of Articulation and Contact in 3D Environments

Conference Paper**Author(s):**

Zhang, Siwei; Zhang, Yan; Ma, Qianli; Black, Michael J.; Tang, Siyu

Publication date:

2020

Permanent link:

<https://doi.org/10.3929/ethz-b-000452857>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

1, <https://doi.org/10.1109/3DV50981.2020.00074>

PLACE: Proximity Learning of Articulation and Contact in 3D Environments

Siwei Zhang¹ Yan Zhang¹ Qianli Ma^{1,2} Michael J. Black² Siyu Tang¹

¹ETH Zürich ²Max Planck Institute for Intelligent Systems

{siwei.zhang, yan.zhang, siyu.tang}@inf.ethz.ch {qma, black}@tue.mpg.de

Figure 1: Given a 3D scene mesh, **PLACE** generates human-scene proximal relationships and infers a plausible human body. The selected local scene vertices are denoted by spheres, and the color denotes the generated minimum distance to the body surface. Distances are normalized to $[0, 1]$ for visualization, and red to purple means close to far.

Abstract

High fidelity digital 3D environments have been proposed in recent years, however, it remains extremely challenging to automatically equip such environment with realistic human bodies. Existing work utilizes images, depth or semantic maps to represent the scene, and parametric human models to represent 3D bodies. While being straightforward, their generated human-scene interactions are often lack of naturalness and physical plausibility. Our key observation is that humans interact with the world through body-scene contact. To synthesize realistic human-scene interactions, it is essential to effectively represent the physical contact and proximity between the body and the world.

To that end, we propose a novel interaction generation method, named **PLACE** (Proximity Learning of Articulation and Contact in 3D Environments), which explicitly models the proximity between the human body and the 3D scene around it. Specifically, given a set of basis points on a scene mesh, we leverage a conditional variational autoencoder to synthesize the minimum distances from the basis points to the human body surface. The generated proximal relationship exhibits which region of the scene is in contact with the person. Furthermore, based on such synthesized proximity, we are able to effectively obtain expressive 3D human bodies that interact with the 3D scene naturally. Our perceptual study shows that **PLACE** significantly improves the state-of-the-art method, approaching the realism of real

human-scene interaction. We believe our method makes an important step towards the fully automatic synthesis of realistic 3D human bodies in 3D scenes. The code and model are available for research at <https://sanweiliti.github.io/PLACE/PLACE.html>.

1. Introduction

Automated synthesis of realistic humans posed naturally in a 3D scene is essential for many applications, such as augmented and virtual reality, special movie effect and synthetic dataset generation for machine learning algorithms. To achieve high realism, manual animation of human-scene interactions is often required. Recent works [21, 42] have proposed methods for the fully automated generation of human bodies that interact with the 3D world, and yet the naturalness and the realism of their results are still far behind the captured real human-scene interaction such as [14].

We argue that the primal reason is that current approaches lack an explicit interaction representation. Humans interact with the 3D scene through *contact*. To explicitly model the contact relationship between the body and the scene is key to the realism of the synthesized humans in the scene. However, existing methods use images, semantic and depth maps to represent the 3D environment [21, 42]. While being easily integrated into deep convolutional networks, the 3D scene structure and the proximity between the body and the scene are not explicitly modeled, especially for the regions that are occluded from the camera view, making it hard to effectively enforce constraints in 3D, such as no inter-penetration and proper contact.

To address these problems, in this paper we propose explicit representations for modelling 3D scene geometry and person-scene contact relationships in a coherent manner. Inspired by the Basis Point Set (BPS) method [29], which efficiently represents a 3D human body by its minimum distances to a set of random points, we propose a *two-stage* BPS encoding scheme: given a 3D scene mesh and a fixed set of basis points in 3D space, we first compute the minimum distances from the basis points to the scene mesh vertices, and use this distance-based feature to represent the 3D scene. Furthermore, we use the selected scene vertices to compute their minimum distances to the human body surface, and use this distance-based feature to represent the person-scene interaction. Consequently, the scene representation and the person-scene interaction representation are modelled in a coherent manner regardless of different scene structures and body poses. Note that, our distance-based interaction representation explicitly captures the contact and the proximal relationships between the human body and the scene and can be regarded as the scene ‘affordance’. Compared with other person-scene interaction representations [7, 30, 31, 32, 34], ours is purely geometry-

based, is more efficient than dense affordance maps, and requires neither scene semantics nor action type labeling, nor affordance class annotation.

With our novel person-scene representation and 3D scene encoding scheme, we propose a new generation method for placing 3D human bodies in 3D environments naturally. We name it as **PLACE** (Proximity Learning of Articulation and Contact in 3D Environments). Specifically, we learn a conditional variational autoencoder (cVAE) to generate plausible contact and proximal relationships between a human body and a 3D scene. We further infer a plausible full body mesh that complies with the generated human-scene proximity. Following [14, 42], we refine the body mesh via optimization; going beyond prior work, we introduce a novel objective term based on our interaction representation that further improve the performance.

We evaluate the performance of our method on three datasets: **PROX** [14], **MP3D** [2, 42], and **Replica** [37], with diverse metrics, namely diversity, physical plausibility, and perceptual naturalness. Experiments show that our method outperforms state-of-the-art methods [21, 42]. Particularly, the extensive perceptual studies demonstrate that our method significantly improves the naturalness and the realism of the synthesized human-scene interactions, making an important step towards the fully automatic generation of realistic human bodies in 3D environments.

In summary, our contributions are: (1) We propose a novel human-scene interaction representation that is efficient, compact and consistent across various bodies and scene structures. (2) Based on the proposed representation, we propose a generative model (**PLACE**) that is able to synthesize full human body meshes with natural poses and plausible contact relationships in novel scenes. (3) We show through comprehensive experiments that **PLACE** achieves new state-of-the-art performance, without using additional modalities (e.g. depth and mesh semantics) that are required in prior works [21, 42].

2. Related Work

2D human-object interaction. Perceiving human-object interactions in 2D images has been studied extensively [5, 9, 10, 15, 16, 22, 23, 24, 35, 39, 40, 43]. Gkioxari et al. [10] detect (human, verb, object) triplets using human appearance as cues to localize interacted objects. Fang et al. [9] learn a pairwise body-part attention model, which focuses on crucial body parts and their corresponding interactions. Wang et al. [39] predict interaction points to localize and classify the interaction directly. Li et al. [22] learn a joint 2D-3D representation by estimating 3D human pose and interacted object position. Besides detecting interactions, generating objects with natural interactions in 2D images is also studied. Tan et al. [38] learn a human instance composition model, which selects and transforms a

human to an input image background. Lee et al. [20] insert an object or human instance mask into the semantic map of an image, and predicts the shape and pose of the object mask. Dvornik et al. [8] predict which object is suitable to be placed in a given region in the image as a data augmentation method.

3D human-scene interaction. The human-scene interaction in 3D has also been widely studied in the literature, e.g. in [1, 11, 13, 17, 33, 34, 44]. Savva et al. [34] learn a joint distribution of human poses and object structures, and generate plausible human body poses and object arrangements given the action. Gupta et al. [13] propose a human-centric scene representation to model physical interactions between the body and the given indoor scene, to predict human poses supported by the scene geometry. Kim et al. [17] predict a human pose on a given 3D object, by learning an affordance model for a class of shapes. Grabner et al. [11] learn an affordance detector to identify regions in a 3D scene that support a particular functional category such as ‘sitting’. Zhu et al. [44] infer forces or pressures applied on the body parts interacting with the scene, and learn the comfort intervals by physical simulation. Monszpart et al. [26] recover both a plausible scene arrangement and human motions to fit an input monocular video by jointly reasoning about scene objects and human motions over space-time. Chen et al. [4] jointly learn scene parsing, object bounding-boxes, camera pose, room layout and 3D human pose estimation given a single-view image. Li et al. [21] learn a 3D pose generative model to automatically put 3D body skeletons into the input scene represented by RGB, RGB-D, or depth image. Cao et al. [1] predict 3D human paths and pose sequences given a scene image and 2D pose histories. Most similar with our work, recently Zhang et al. [42] propose a cVAE to generate 3D human mesh modeled by SMPL-X [28] model given a 3D environment represented by RGB-D image.

Ours versus related work. Compared to previous work, we explicitly use the proximity to model human-scene interaction in 3D. With this compact representation, we train a cVAE from real human-scene interaction data to synthesize the proximal relations between the human body and the scene. We show significant improvement on a large-scale perceptual study. The results indicate that our method is approaching the realism of real human-scene interaction.

3. Proximity Representation for Interaction

In this section, we introduce our novel proximity representation for modeling human-scene interaction. Specifically, we use the Basis Point Sets (BPS) encoding scheme [29] to represent the scene, the human body, and their interactions. The BPS method [29] encodes a 3D object into a feature vector, which is defined on a fixed set of randomly

selected basis points. This feature vector, i.e. the distances from the basis points to their nearest neighbors on the target object, essentially captures the proximity relationship between the basis points and the object surface. Such property is well in line with our objective: by selecting basis points on the scene mesh, this distance-based feature naturally represents the body-scene contact and proximity. Furthermore, in order to synthesize novel human bodies in an arbitrary 3D environment, we need an automatic way to select and encode the basis points on a scene mesh. To that end, we propose a *two-stage* BPS scheme that encodes the scene and the body consecutively (see Fig. 2 for detailed illustration).

Stage 1: BPS encoding for 3D Scene. To represent a 3D scene, we follow [29] and create a set of $N = 10^4$ uniformly random basis points $\mathcal{V}_o \subset \mathbb{R}^{N \times 3}$ within a unit sphere. $\mathcal{V}_o$ are kept fixed throughout both the training and testing pipeline for all 3D scenes. The scene meshes are also normalized into the unit sphere. Then, for each basis point in $\mathcal{V}_o$, we find its nearest neighbor in the normalized scene mesh vertices, select the corresponding scene vertex and compute the distance. Such minimum distances form the *scene feature* $\mathbf{x}_s \in \mathbb{R}^{N \times 1}$; and the selected N scene vertices, $\mathcal{V}_s \subset \mathbb{R}^{N \times 3}$, will serve as *scene BPS* to encode the body in the second stage.

In practice, only the part of the scene that is close to the human body is most relevant for modeling the interaction. This motivates us to consider ‘local scenes’, i.e. encode only part of the scene surrounding the body. To do so, we specify a virtual 3D cage around the body, and only consider the scene vertices within the cage. Furthermore, we regularly sample points on the top and the 4 side faces of the cage, and consider them as the ‘ceiling’ and ‘walls’ of the local scene. The local scene then goes through the encoding pipeline described above. We find the inclusion of the virtual ceiling and walls useful in many cases, as otherwise the scene BPS $\mathcal{V}_s$ can fail to capture the full human body. For example, in the simplest configuration where the human stands in a 3D environment that only consists a floor, without the virtual ceiling and the walls, $\mathcal{V}_s$ only consists the vertices on the floor, then the body except the feet will hardly be encoded by the scene BPS.

Stage 2: BPS encoding for 3D Body. The obtained scene BPS $\mathcal{V}_s$, which also includes vertices on the cage surfaces, is employed to encode the body mesh, following the same nearest neighbor encoding pipeline. We use the SMPL-X [28] model to represent the 3D human body as mesh with vertices $\mathcal{V}_b \subset \mathbb{R}^{10475 \times 3}$. The BPS encoding compresses $\mathcal{V}_b$ into the *body feature*, $\mathbf{x}_b \in \mathbb{R}^{N \times 1}$, which is obtained by computing the distance between each point in scene BPS $\mathcal{V}_s$ to its nearest point of the body mesh vertices. The proposed *body feature* is scene dependent, and explicitly encodes the contact and the proximal relations between the human body and the scene. Therefore it can be used to as the interac-

Figure 2: Illustration of human-scene interaction feature extraction (Sec. 3). From left to right: (1) A pair of human-scene meshes from the **PROX** dataset [14]. (2) The cropped scene mesh as well as the 3D cage (in pink) are normalized into the unit sphere. (3) The derived scene BPS $\mathcal{V}_s$ is denoted by colored points, and the color indicates the distance-based body BPS features, i.e. the minimum distance from each scene basis point to the body mesh vertices. Small to large distance is denoted by bright yellow to dark purple. (4) The scene BPS and their minimum distances to the body only on the original scene mesh.

tion representation to model and synthesize realistic human bodies in a 3D scene. Note that the basis points in the unit sphere (used in **Stage 1**) are always fixed. Consequently, every dimension in the scene feature always corresponds to the same dimension in the body feature. Therefore the proposed scene-body representation is consistent across various human body poses and scene structures.

4. PLACE

Given the proposed proximity representation for a pair of interacting human body and 3D scene, here we introduce our generative model, PLACE, which takes the scene BPS $\mathcal{V}_s$, the scene feature x_s , and the body feature x_b as input to synthesize plausible human-scene contact and proximal relationships and to infer expressive 3D human bodies that interact with the 3D scene naturally.

Network Architecture The network architecture is illustrated in Fig. 3. It incorporates a conditional variational autoencoder (cVAE) [19, 36] to generate body features based on a scene feature auto-encoder E . The generator G is followed by a multilayer perceptron H as the regressor to produce full body meshes, based on the reconstructed body features and the scene BPS auto-encoder F . Intermediate features from the decoder of E are introduced to the decoder of G for better scene conditioning. Details of the network architecture and its ablation study are demonstrated in the Appendix (Fig. S1 and Tab. S1).

4.1. Losses and Learning

Following Sec. 3, we obtain a set of representations from each ground truth sample, namely $\{x_s, x_b, \mathcal{V}_s, \mathcal{V}_b\}$ denoting the scene feature, the body feature, the scene BPS, and the full body vertices, respectively. Note that the point coordinates are rescaled in the unit sphere, and set relative to the rescaled 3D cage center. We denote the operations of

Figure 3: Illustration of the network architecture for PLACE. The training pipeline is illustrated in (a). x_s , $\mathcal{V}_s$, x_b and $\mathcal{V}_b$ denote the scene feature, the scene BPS, the body feature, and the full body mesh vertices, respectively. $\Delta \mathcal{V}_b^{rec}$ denotes a global 3D translation for all the body vertices. Note that all point coordinates are relative to the center of the normalized 3D cage in the unit sphere. The “stronger scene condition” introduces additional intermediate features to the decoder of G . During inference stage, the scene-conditioned VAE is shown in (b), with the same body regressor as in training stage.

transforming from the rescaled 3D cage back to the original world as $\pi(\cdot)$, and its inverse operation as $\pi^{-1}(\cdot)$. The training loss is given by

$$\mathcal{L} = \mathcal{L}_{rec} + \alpha_1 \mathcal{L}_{KL} + \alpha_2 \mathcal{L}_{coll} + \alpha_3 \mathcal{L}_{contact}, \quad (1)$$

in which $\{\alpha_1, \alpha_2, \alpha_3\}$ are loss weights as hyper-parameters.

Reconstruction term: This term is given by

$$\mathcal{L}_{rec} = |x_s - x_s^{rec}| + |x_b - x_b^{rec}| + |\mathcal{V}_b - \tilde{\mathcal{V}}_b^{rec}|, \quad (2)$$

in which the last term denotes vertex-wise reconstruction.

KL-divergence term: Following [19], we specify the latent prior as the standard normal distribution, and employ the same re-parameterization approach to model the inference posterior $q(z|x)$. To effectively avoid posterior collapse,

we use the robust KL-divergence term as in [41]. Our KL-divergence term is given by

$$\mathcal{L}_{KL} = \Psi(q(z|\mathbf{x})||\mathcal{N}(0, \mathbf{I})), \quad (3)$$

in which $\Psi(\cdot)$ is the Charbonnier function $\Psi(s) = \sqrt{s^2 + 1} - 1$ [3]. According to [41], the smaller is the KL-divergence between the inference posterior and the prior, the smaller are the gradients to update this KL-divergence.

The collision term and the contact term: As in [42], these two terms are given by

$$\mathcal{L}_{coll} = \frac{1}{|\mathcal{V}_b|} \sum_{\mathbf{v}_b \in \pi(\mathcal{V}_b^{rec})} |\Phi^-(\mathbf{v}_b)|, \text{ and} \quad (4)$$

$$\mathcal{L}_{contact} = \sum_{\mathbf{v}_b \in C(\pi(\mathcal{V}_b^{rec}))} \min_{\mathbf{v}_s \in \mathcal{M}_s} \rho(|\mathbf{v}_b - \mathbf{v}_s|), \quad (5)$$

in which $C(\cdot)$ denotes the mask to select body part for contact, $|\mathcal{V}_b|$ denotes the constant number of body mesh vertices, $\Phi^-(\cdot)$ is the scene negative signed distance field (SDF), $\mathcal{M}_s$ is the provided scene mesh, and $\rho(\cdot)$ is the Geman-McClure error function to remove the influence of scene mesh vertices far away from the body [14, 42]. In contrast to [14, 42], the contact mask $C(\cdot)$ only includes feet, which we find is better to avoid body floating.

4.2. Placing People in Novel Scenes

Given a 3D scene without humans, the generation procedures are: (1) We randomly create a 3D cage within the scene and transform it into the pre-defined unit sphere. (2) Within the unit sphere, we compute the scene BPS $\mathcal{V}_s$ and the scene feature $\mathbf{x}_s$. (3) Given z sampled from $\mathcal{N}(0, \mathbf{I})$, the cVAE is employed to generate a body feature $\mathbf{x}_b$ (as illustrated in Fig. 3 (b)). (4) Based on $\mathcal{V}_s$ and $\mathbf{x}_b$, the regressor is employed to produce a set of body mesh vertices, which is then transformed back to the original world.

Interaction-based Optimization. The body vertices produced by the regressor H are not guaranteed to form a valid body shape: The raw output from the network can suffer from noisy surface geometry or inter-penetration with the scene. Therefore, we conduct an additional optimization step to convert the regressed body mesh vertices to a SMPL-X body mesh. Given the scene mesh $\mathcal{M}_s$, the body feature $\mathbf{x}_b$, and the regressed body mesh vertices $\mathcal{V}_b$, the optimization loss is given by

$$\begin{aligned} \mathcal{L}(\theta) = & |\pi(\mathcal{V}_b) - \mathcal{V}_{SMPLX}(\theta)| + |\mathbf{x}_b - f(\pi^{-1}(\mathcal{V}_{SMPLX}(\theta)))| \\ & + \lambda_1 \mathcal{L}_{coll} + \lambda_2 \mathcal{L}_{contact} + \lambda_3 \mathcal{L}_{vposer} \\ & + \lambda_4 \mathcal{L}_{hand} + \lambda_5 \mathcal{L}_{shape}, \end{aligned} \quad (6)$$

in which $\mathcal{V}_{SMPLX}(\theta)$ denotes vertices of the SMPL-X body mesh with the attribute θ (including global translation,

Figure 4: Illustration of optimization results. The left column shows two generated body mesh vertices, and the right column shows their optimized results.

orientation, body shape, pose and hand pose), $\mathcal{L}_{coll}$ and $\mathcal{L}_{contact}$ are the same in Eq. (4). In addition, $\mathcal{L}_{vposer}$, $\mathcal{L}_{hand}$, $\mathcal{L}_{shape}$ are the VPoser loss [14, 28], L2 prior loss for hand and shape parameters in SMPL-X model to ensure natural body, hand poses and shapes of the optimized body, respectively. λ 's are loss weights as hyper-parameters, and $f(\cdot)$ is the operation to extract the body feature from the body mesh, demonstrated in Sec. 3.

In the optimization loss, the first term denotes the vertex-wise reconstruction, in order to discourage dramatic deviation of the resultant SMPL-X body mesh from the regressed body vertices. Additionally, we employ the second term to encourage the resultant SMPL-X body mesh have the same contact relation with the scene. Unlike the heuristic contact term $\mathcal{L}_{contact}$, this BPS feature reconstruction term is data-driven, since the body feature can be learned from data and generated by our cVAE model.

Fig. 4 shows two results of optimization. We can see that the human-scene contact is preserved, and their inter-penetration is considerably alleviated.

5. Experiments

5.1. Datasets

PROX [14] includes 12 different room 3D scans, and captures natural actions of 20 subjects, which are represented by SMPL-X body meshes [28]. Following [42], we select ‘MPH1Library’, ‘MPH16’, ‘N0SittingBooth’ and ‘N3OpenArea’ for testing, and use the rests for training, from which we extract interaction features from 63,599 human-scene pairs. To improve the generalization ability to new scenes, we perform data augmentation 4 times for each human-scene pair, and obtain 254,396 training samples in total. During testing, we randomly set 1200 3D cages in

each test scene, and generate one random sample in each cage. This results in 4800 examples for evaluation.

MP3D [2] has 90 building-scale scene scans with semantic annotations. Following the setup in [42], we perform testing on the same 7 scenes with corresponding scene SDFs, and generate the same number of random samples for each scene as in [42], obtaining 6400 random samples in total.

Replica [37] contains reconstructed indoor scenes of clean dense geometry. It includes 18 scenes of 5 different room types, among which we select 1 room in each type, resulting in 5 rooms ('office.2', 'hotel.0', 'room.0', 'frl.apartment.0', 'apartment.1') for testing. 1200 random samples are generated on each scene for evaluation.

5.2. Our method and Baselines

In this paper, we denote our model trained without $\mathcal{L}_{\text{coll}}$ and $\mathcal{L}_{\text{contact}}$ as Ours. Its ablation study is referred to Sec. 5.6 and Appendix (Sec. B.2). We compare our method with two baselines: (1) PSI [42], which generates SMPL-X parameters given the scene depth and the semantic segmentation. Specifically, we perform comparison with their S1 model without the human-scene interaction training loss, both before and after the geometry-aware fitting, which are reported to outperform other methods. (2) Li et al. [21], which generates 3D body stick figures based on scene appearance and depth. We perform comparison with its modified version proposed in [42], which generates SMPL-X human bodies for fair comparison.

It is noted that these two baselines generate SMPL-X parameters, yet our method directly produces body vertices, which might not form a valid body shape. To fairly compare with the baselines without the geometry-aware fitting step, we perform an optimization step to derive SMPL-X parameters from the body vertices, for which $\mathcal{L}_{\text{contact}}$ and $\mathcal{L}_{\text{coll}}$ in Eq. 6 are excluded. We denote this optimization as SimOptim, and the optimization with all terms as AdvOptim in the following content. Moreover, we denote the scene geometry-aware fitting proposed in [42] as $\mathcal{L}_f$.

5.3. Quantitative Evaluation

We evaluate the performance in terms of perceptual naturalness, diversity, and physical plausibility, following Zhang et al. [42].

5.3.1 Perceptual Naturalness

We run perceptual studies on Amazon Mechanical Turk (AMT), and use two protocols to perform the comparison. For each scene and each model, we generate 100 and 400 examples for **PROX** and **MP3D** respectively, as in [42].

Unary user study. Like [42], we ask turkers to evaluate each individual generated result or ground truth, giving a score from 1 (strongly not natural) to 5 (strongly natural).

Table 1: **Unary user study.** We calculate the results from all turkers, and report the the scores w.r.t. the average $\pm$ the standard deviation. Best results except the ground truth are shown in boldface.

model	PROX	MP3D
Li et al. [21] + $\mathcal{L}_f$	3.37 ± 1.31	3.52 ± 1.34
PSI [42] + $\mathcal{L}_f$	3.47 ± 1.27	3.23 ± 1.40
Ours + $\mathcal{L}_{\text{contact}}$ + SimOptim	$3.61 \pm \mathbf{1.16}$	3.77 ± 1.15
Ours + $\mathcal{L}_{\text{contact}}$ + AdvOptim	$\mathbf{3.71} \pm 1.17$	$\mathbf{3.84} \pm \mathbf{1.11}$
Ground truth	3.94 ± 0.99	–

Table 2: **Binary user study.** Numbers show the percentage of the users that rate the corresponding method as more realistic.

Datasets	% users rated as “better”	
MP3D	Ours + $\mathcal{L}_{\text{contact}}$ + AdvOptim	PSI + $\mathcal{L}_f$
	70.1%	29.9%
PROX	Ours + $\mathcal{L}_{\text{contact}}$ + AdvOptim	PSI + $\mathcal{L}_f$
	69.0%	31.0%
PROX	Ours + $\mathcal{L}_{\text{contact}}$ + AdvOptim	GT
	48.5%	51.5%

Each example is rated by 3 turkers. The results are shown in Tab. 1. On both **PROX** and **MP3D**, our methods with SimOptim and AdvOptim outperform the baseline methods after their geometry-aware fitting.

Binary user study. To perform a direct comparison, we follow [21, 25], and show results of two methods to the turkers at the same time. The turkers are asked to pick the one that they think is more perceptually natural. For **PROX**, we compare the best instance of our method (the one trained with $\mathcal{L}_{\text{contact}}$ and AdvOptim) against [42] with the geometry-aware fitting $\mathcal{L}_f$, and against the ground truth, respectively. For **MP3D**, we compare with [42] with $\mathcal{L}_f$.

As shown in Tab. 2, Our method outperforms PSI [42] on both datasets by a large margin. More interestingly, in the direct comparison with the **PROX** ground truth, our generated results are regarded as more realistic by nearly half (48.5%) of the users, indicating that our results are hardly distinguishable from the real human-scene interaction. It suggests that our method makes an important step towards the fully automatic synthesis of realistic 3D human bodies in 3D scenes.

5.3.2 Diversity

Evaluation metrics. Like [42], we perform K-means to cluster the SMPL-X parameters of generated bodies into 20

clusters, and evaluate the model by: 1) the entropy of the cluster ID histogram for all samples, and 2) the cluster size, which denotes the average distance between the sample and the corresponding cluster center. The higher the better for both metrics. We report average scores over all test samples.

Results. Tab. 3 presents the diversity results. Comparing with baselines, Ours+ $\mathcal{L}_{\text{contact}}$ +SimOptim achieve comparable entropy and cluster size on both **PROX** and **MP3D**. When applying AdvOptim, the cluster size is further increased by a large margin, indicating that the generated bodies are more scattered in the scene. Since the cluster size is correlated with the test scene size, its values are higher on **MP3D** and **Replica** than **PROX** due to larger room sizes.

5.3.3 Physical Plausibility

Evaluation metrics. We evaluate the non-collision and contact scores between the generated body and scene mesh, following the metrics in [42]. Specifically, the non-collision score is the ratio of the number of body vertices with non-negative scene SDF values to the number of all body vertices (10475). The contact score is 1 if any body mesh vertex has a non-positive scene SDF value, and 0 otherwise. Then we report the average non-collision scores and average contact scores of all test examples. Thus, a higher non-collision score indicates fewer body-scene inter-penetration, while a higher contact score suggests more contacts between the generated body and scene mesh. Higher values for both scores are desirable.

Results. Tab. 3 shows the results on physical plausibility. On **PROX**, Ours+ $\mathcal{L}_{\text{contact}}$ +SimOptim achieves higher non-collision score and comparable contact score with the baselines, indicating that our proposed method can effectively alleviate scene-body inter-penetration, while preserving plausible contact relations. When additionally employing the interaction-based optimization, performance is further consistently improved. In addition, we find the SDFs on **MP3D** are not sufficiently reliable probably due to noisy scans¹. Thus, we further perform evaluation on the **Replica** dataset, from which high-quality scene SDFs can be derived. The results on **Replica** verify the effectiveness of our model, which is consistent with the performances on **PROX**.

5.4. Qualitative Results.

Fig. 5 shows some generated human bodies in the three datasets. To further analyze our cVAE, we perform interpolation in its latent space, and show the results in Fig. 6. It can be seen that the body changes smoothly, while preserving natural interactions with the scene. Note that the results in Fig. 6 are directly produced by the body vertex regressor H , without any optimization-based refinement.

¹Illustrations are presented in the Appendix (Fig. S3).

Figure 5: Generated results on **PROX**, **MP3D** and **Replica**.

5.5. Limitation and Failure Cases

Despite the effectiveness verified by our evaluations, our method still has limitations. For example, the generated body might inter-penetrate with thin structures in the scene mesh, e.g. table surface. In this case, AdvOptim cannot resolve this issue, since the value of the collision loss $\mathcal{L}_{\text{coll}}$ is already very small. We thus expect a novel collision representation in future could resolve this issue. Furthermore, detailed and diverse hand poses can not be well synthesized by the current model. Because the basis point set on the scene mesh are too sparse to model the hand-object interaction. We expect a coarse-to-fine encoding scheme could help. We consider these challenges as exciting future directions. Furthermore, the visualizations of the failure cases can be found in the Appendix (Fig. S7).

5.6. Ablation Study

We investigate the influence of the training loss and the optimization scheme on our model performance on **PROX** dataset. More ablation studies are referred to the Appendix.

Tab. 4 and Tab. 5 present model performances when trained with or without contact loss $\mathcal{L}_{\text{contact}}$ and collision loss $\mathcal{L}_{\text{coll}}$ in Eq. (1), respectively. We find that the model without $\mathcal{L}_{\text{contact}}$ and $\mathcal{L}_{\text{coll}}$ has inferior contact scores, indicating that it tends to produce more floating bodies. Also, in contrast to Zhang et al. [42], we find the loss term $\mathcal{L}_{\text{contact}}$ is effective to improve the contact score in our model. Additionally, the model trained without $\mathcal{L}_{\text{coll}}$ already delivers high non-collision scores, indicating the effectiveness of our BPS-based to avoid inter-penetration. However, $\mathcal{L}_{\text{coll}}$ in training reduces the cluster size. It is possibly because there exists inter-penetration cases in the ground truth data, which causes contradictions between $\mathcal{L}_{\text{coll}}$ and reconstruction loss during training. Thus, we mostly use Ours trained with $\mathcal{L}_{\text{contact}}$ to compare with baseline methods. Moreover, when comparing the results between SimOptim and AdvOptim, we can observe that the collision and contact terms in Eq. (6) significantly improve both physical plausibility and diversity metrics, which is consistent with [42].

Table 3: **Evaluation on diversity and physical plausibility:** For the entropy, the cluster size, the non-collision score and the contact score, their values are the higher the better. Row-wisely, the first part and the second part shows the comparison to baseline methods before and after geometry-aware fitting, respectively. For each part, the best results are in boldface.

Methods	Diversity						Physical plausibility					
	entropy			cluster size			non-coll			contact		
	PROX	MP3D	Replica	PROX	MP3D	Replica	PROX	MP3D	Replica	PROX	MP3D	Replica
Li et al. [21]	2.89	2.93	-	1.49	1.84	-	0.89	0.92	-	0.93	0.78	-
PSI [42]	2.96	2.99	-	2.51	2.81	-	0.93	0.94	-	0.95	0.80	-
Ours + $\mathcal{L}_{\text{contact}}$	-	-	-	-	-	-	0.95	0.97	0.92	0.90	0.57	0.90
Ours + $\mathcal{L}_{\text{contact}}$ + SimOptim	2.91	2.92	2.93	2.59	2.93	2.99	0.96	0.97	0.91	0.95	0.62	0.92
Li et al. + $\mathcal{L}_f$	2.93	2.92	-	1.52	1.94	-	0.93	0.97	-	0.99	0.89	-
PSI + $\mathcal{L}_f$	2.97	2.98	-	2.53	2.86	-	0.94	0.97	-	0.99	0.88	-
Ours + $\mathcal{L}_{\text{contact}}$ + AdvOptim	2.91	2.93	2.92	2.72	3.07	3.12	0.98	0.98	0.93	0.99	0.57	1.00

Figure 6: Latent space interpolation. In each row, the left and right images denote two different random samples, and middle images correspond to interpolated latent variables. For each row, the scene BPS $\mathcal{V}_s$ and the scene feature $\mathbf{x}_s$ are fixed.

Table 4: Ablation study w.r.t. diversity metrics.

Optimization	Method	entropy	cluster size
SimOptim	Ours	2.89	2.56
	Ours + $\mathcal{L}_{\text{contact}}$	2.91	2.59
	Ours + $\mathcal{L}_{\text{coll}}$	2.89	2.50
	Ours + $\mathcal{L}_{\text{contact}}$ + $\mathcal{L}_{\text{coll}}$	2.90	2.48
AdvOptim	Ours	2.90	2.67
	Ours + $\mathcal{L}_{\text{contact}}$	2.91	2.72
	Ours + $\mathcal{L}_{\text{coll}}$	2.88	2.63
	Ours + $\mathcal{L}_{\text{contact}}$ + $\mathcal{L}_{\text{coll}}$	2.89	2.61

6. Conclusion

This paper introduces a novel representation to model human-scene interactions by explicitly encoding the proximity between the body and the environment. It is purely geometry-based and consistent across various bodies and scene structures, thus only requiring the 3D mesh as the input. By training a generative model to synthesize such representations and body meshes, our approach is able to generate humans in natural interactions with new environments, which demonstrates the practical merits of our work.

Acknowledgements. We sincerely thank Mohamed Hassan for his scene SDF calculation protocol. This work was partially sup-

Table 5: Ablation study w.r.t. physical plausibility metrics.

Optimization	Method	non-coll	contact
w/o	Ours	0.95	0.83
	Ours + $\mathcal{L}_{\text{contact}}$	0.95	0.90
	Ours + $\mathcal{L}_{\text{coll}}$	0.95	0.83
	Ours + $\mathcal{L}_{\text{contact}}$ + $\mathcal{L}_{\text{coll}}$	0.95	0.88
SimOptim	Ours	0.96	0.84
	Ours + $\mathcal{L}_{\text{contact}}$	0.96	0.95
	Ours + $\mathcal{L}_{\text{coll}}$	0.96	0.87
	Ours + $\mathcal{L}_{\text{contact}}$ + $\mathcal{L}_{\text{coll}}$	0.95	0.94
AdvOptim	Ours	0.98	0.99
	Ours + $\mathcal{L}_{\text{contact}}$	0.98	0.99
	Ours + $\mathcal{L}_{\text{coll}}$	0.98	0.98
	Ours + $\mathcal{L}_{\text{contact}}$ + $\mathcal{L}_{\text{coll}}$	0.98	0.99

ported by the German Research Foundation (DFG): SFB 1233, Robust Vision: Inference Principles and Neural Mechanisms, TP XX, project number: 276693517. Qianli Ma acknowledges the support from the Max Planck ETH Center for Learning Systems.

Disclosure. MJB has received research gift funds from Intel, Nvidia, Adobe, Facebook, and Amazon. While MJB is a part-time employee of Amazon, his research was performed solely at MPI. He is also an investor in Meshcapde GmbH.

References

- [1] Z. Cao, H. Gao, K. Mangalam, Q. Cai, M. Vo, and J. Malik. Long-term human motion prediction with scene context. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 2020. [3](#)
- [2] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang. Matterport3D: Learning from RGB-D data in indoor environments. *International Conference on 3D Vision (3DV)*, 2017. [2](#), [6](#)
- [3] P. Charbonnier, L. Blanc-Feraud, G. Aubert, and M. Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proceedings of 1st International Conference on Image Processing*, volume 2, pages 168–172. IEEE, 1994. [5](#)
- [4] Y. Chen, S. Huang, T. Yuan, S. Qi, Y. Zhu, and S.-C. Zhu. Holistic++ scene understanding: Single-view 3d holistic scene parsing and human pose estimation with human-object interaction and physical commonsense. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8648–8657, 2019. [3](#)
- [5] V. Delaitre, J. Sivic, and I. Laptev. Learning person-object interactions for action recognition in still images. In *Advances in neural information processing systems*, pages 1503–1511, 2011. [2](#)
- [6] T. Deprelle, T. Groueix, M. Fisher, V. G. Kim, B. C. Russell, and M. Aubry. Learning elementary structures for 3D shape generation and matching. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2019. [1](#)
- [7] T.-T. Do, A. Nguyen, and I. Reid. Affordancenet: An end-to-end deep learning approach for object affordance detection. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 1–5. IEEE, 2018. [2](#)
- [8] N. Dvornik, J. Mairal, and C. Schmid. On the importance of visual context for data augmentation in scene understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. [3](#)
- [9] H.-S. Fang, J. Cao, Y.-W. Tai, and C. Lu. Pairwise body-part attention for recognizing human-object interactions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 51–67. [2](#)
- [10] G. Gkioxari, R. Girshick, P. Dollár, and K. He. Detecting and recognizing human-object interactions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8359–8367, 2018. [2](#)
- [11] H. Grabner, J. Gall, and L. Van Gool. What makes a chair a chair? In *CVPR 2011*, pages 1529–1536. IEEE, 2011. [3](#)
- [12] T. Groueix, M. Fisher, V. G. Kim, B. Russell, and M. Aubry. 3D-CODED: 3D correspondences by deep deformation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. [1](#)
- [13] A. Gupta, S. Satkin, A. A. Efros, and M. Hebert. From 3d scene geometry to human workspace. In *CVPR 2011*, pages 1961–1968. IEEE, 2011. [3](#)
- [14] 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 Conference on International Conference on Computer Vision (ICCV)*, pages 2282–2292, Oct. 2019. [2](#), [4](#), [5](#), [1](#)
- [15] K. Kato, Y. Li, and A. Gupta. Compositional learning for human object interaction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 234–251, 2018. [2](#)
- [16] D.-J. Kim, X. Sun, J. Choi, S. Lin, and I. S. Kweon. Detecting human-object interactions with action co-occurrence priors. *arXiv preprint arXiv:2007.08728*, 2020. [2](#)
- [17] V. G. Kim, S. Chaudhuri, L. Guibas, and T. Funkhouser. Shape2pose: Human-centric shape analysis. *ACM Transactions on Graphics (TOG)*, 33(4):1–12, 2014. [3](#)
- [18] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In Y. Bengio and Y. LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [1](#)
- [19] D. P. Kingma and M. Welling. Auto-encoding variational Bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014. [4](#)
- [20] D. Lee, S. Liu, J. Gu, M.-Y. Liu, M.-H. Yang, and J. Kautz. Context-aware synthesis and placement of object instances. In *Advances in neural information processing systems*, pages 10393–10403, 2018. [3](#)
- [21] X. Li, S. Liu, K. Kim, X. Wang, M.-H. Yang, and J. Kautz. Putting humans in a scene: Learning affordance in 3D indoor environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#), [3](#), [6](#), [8](#)
- [22] Y.-L. Li, X. Liu, H. Lu, S. Wang, J. Liu, J. Li, and C. Lu. Detailed 2d-3d joint representation for human-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10166–10175, 2020. [2](#)
- [23] Y.-L. Li, S. Zhou, X. Huang, L. Xu, Z. Ma, H.-S. Fang, Y. Wang, and C. Lu. Transferable interactiveness knowledge for human-object interaction detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3585–3594, 2019. [2](#)
- [24] M. Liu, S. Tang, Y. Li, and J. Rehg. Forecasting human object interaction: Joint prediction of motor attention and actions in first person video. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. [2](#)
- [25] Q. Ma, J. Yang, A. Ranjan, S. Pujades, G. Pons-Moll, S. Tang, and M. J. Black. Learning to Dress 3D People in Generative Clothing. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. [6](#)
- [26] A. Monszpart, P. Guerrero, D. Ceylan, E. Yumer, and N. J. Mitra. imapper: interaction-guided scene mapping from monocular videos. *ACM Transactions on Graphics (TOG)*, 38(4):1–15, 2019. [3](#)
- [27] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. 2017. [1](#)
- [28] 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 *Pro-*

- ceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3, 5
- [29] S. Prokudin, C. Lassner, and J. Romero. Efficient learning on point clouds with basis point sets. In *Proceedings of the IEEE/CVF Conference on International Conference on Computer Vision (ICCV)*, pages 4332–4341, Oct. 2019. 2, 3
- [30] S. Qi, Y. Zhu, S. Huang, C. Jiang, and S.-C. Zhu. Human-centric indoor scene synthesis using stochastic grammar. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5899–5908, 2018. 2
- [31] A. Roy and S. Todorovic. A multi-scale cnn for affordance segmentation in rgb images. In *European conference on computer vision*, pages 186–201. Springer, 2016. 2
- [32] E. Ruiz and W. Mayol-Cuevas. Geometric affordance perception: Leveraging deep 3d saliency with the interaction tensor. *Frontiers in Neurorobotics*, 14:45, 2020. 2
- [33] M. Savva, A. X. Chang, P. Hanrahan, M. Fisher, and M. Nießner. Scenegrok: Inferring action maps in 3d environments. *ACM transactions on graphics (TOG)*, 33(6):1–10, 2014. 3
- [34] M. Savva, A. X. Chang, P. Hanrahan, M. Fisher, and M. Nießner. Pigraphs: learning interaction snapshots from observations. *ACM Transactions on Graphics (TOG)*, 35(4):1–12, 2016. 2, 3
- [35] L. Shen, S. Yeung, J. Hoffman, G. Mori, and L. Fei-Fei. Scaling human-object interaction recognition through zero-shot learning. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1568–1576. IEEE, 2018. 2
- [36] K. Sohn, H. Lee, and X. Yan. Learning structured output representation using deep conditional generative models. In *Advances in neural information processing systems*, pages 3483–3491, 2015. 4
- [37] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijmans, S. Green, J. J. Engel, R. Mur-Artal, C. Ren, S. Verma, A. Clarkson, M. Yan, B. Budge, Y. Yan, X. Pan, J. Yon, Y. Zou, K. Leon, N. Carter, J. Briales, T. Gillingham, E. Mueggler, L. Pesqueira, M. Savva, D. Batra, H. M. Strasdat, R. D. Nardi, M. Goesle, S. Lovegrove, and R. Newcombe. The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. 2, 6
- [38] F. Tan, C. Bernier, B. Cohen, V. Ordonez, and C. Barnes. Where and who? automatic semantic-aware person composition. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1519–1528. IEEE, 2018. 2
- [39] T. Wang, T. Yang, M. Danelljan, F. S. Khan, X. Zhang, and J. Sun. Learning human-object interaction detection using interaction points. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4116–4125, 2020. 2
- [40] B. Yao and L. Fei-Fei. Modeling mutual context of object and human pose in human-object interaction activities. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 17–24. IEEE, 2010. 2
- [41] Y. Zhang, M. J. Black, and S. Tang. Perpetual motion: Generating unbounded human motion, 2020. 5
- [42] Y. Zhang, M. Hassan, H. Neumann, M. J. Black, and S. Tang. Generating 3d people in scenes without people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6194–6204, 2020. 2, 3, 5, 6, 7, 8, 1
- [43] T. Zhou, W. Wang, S. Qi, H. Ling, and J. Shen. Cascaded human-object interaction recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4263–4272, 2020. 2
- [44] Y. Zhu, C. Jiang, Y. Zhao, D. Terzopoulos, and S.-C. Zhu. Inferring forces and learning human utilities from videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3823–3833, 2016. 3