

Generating 3D People in Scenes without People

Yan Zhang^{*1,3}, Mohamed Hassan², Heiko Neumann³, Michael J. Black², Siyu Tang^{*1}
¹ ETH Zürich, Switzerland

²Max Planck Institute for Intelligent Systems, Tübingen, Germany

³Institute of Neural Information Processing, Ulm University, Germany



Figure 1: 3D human bodies with various shapes and poses are automatically generated to interact with the scene. Appropriate human-scene contact is encouraged, and human-scene surface interpenetration is discouraged.

Abstract

We present a fully automatic system that takes a 3D scene and generates plausible 3D human bodies that are posed naturally in that 3D scene. Given a 3D scene without people, humans can easily imagine how people could interact with the scene and the objects in it. However, this is a challenging task for a computer as solving it requires that (1) the generated human bodies to be semantically plausible within the 3D environment (e.g. people sitting on the sofa or cooking near the stove), and (2) the generated human-scene interaction to be physically feasible such that the human body and scene do not interpenetrate while, at the same time, body-scene contact supports physical interactions. To that end, we make use of the surface-based 3D human model SMPL-X. We first train a conditional variational autoencoder to predict semantically plausible 3D human poses

conditioned on latent scene representations, then we further refine the generated 3D bodies using scene constraints to enforce feasible physical interaction. We show that our approach is able to synthesize realistic and expressive 3D human bodies that naturally interact with 3D environment. We perform extensive experiments demonstrating that our generative framework compares favorably with existing methods, both qualitatively and quantitatively. We believe that our scene-conditioned 3D human generation pipeline will be useful for numerous applications; e.g. to generate training data for human pose estimation, in video games and in VR/AR. Our project page for data and code can be seen at: <https://vlg.inf.ethz.ch/projects/PSI/>.

* This work was performed when Y. Z. and S. T. were at MPI-IS and University of Tübingen.

1. Introduction

In recent years, many high-quality datasets of 3D indoor scenes have emerged such as Matterport3D [3], Replica [42], and Gibson [47], which employ 3D scanning and reconstruction technologies to create digital 3D environments. Also, virtual robotic agents exist inside of 3D environments such as Gibson [47] and the Habitat simulator [32]. These are used to develop scene understanding methods from embodied views, thus providing platforms for indoor robot navigation, AR/VR, computer games and many other applications. Despite this progress, a significant limitation of these environments is that they do not contain people. The reason such worlds contain no people is that there are no automated tools to generate realistic people interacting realistically with 3D scenes, and manually doing this requires significant artist effort. Consequently, our goal is to automatically generate natural and realistic 3D human bodies in the scene. The generated human bodies are expected to be physically plausible (e.g. neither floating nor interpenetrating), diverse, and posed naturally within the scene. This is a step towards equipping high-quality 3D scenes and simulators (e.g. Matterport3D [3] and Habitat [32]) with semantically and physically plausible 3D humans, and is essential for numerous applications such as creating synthetic datasets, VR/AR, computer games, etc.

Our solution is inspired by how humans infer plausible interactions with the environment. According to the studies of [49], humans tend to propose interaction plans depending on the structure and the semantics of objects. Afterwards, to realize the interaction plan, physical rules will apply to determine the detailed human-object configuration, while guaranteeing that the human body can neither float in the air nor collide into the objects. Therefore, our method has two steps: (1) We propose a generative model of human-scene interaction using a conditional variational autoencoder (CVAE) [39] framework. Given scene depth and semantics, we can sample from the CVAE to obtain various human bodies. (2) Next, we transform the generated 3D human body to the world coordinates and perform scene geometry-aware fitting, so as to refine the human-scene interaction and eliminate physically implausible configurations (e.g. floating and collision).

We argue that realistically modeling human-scene interactions requires a realistic model of the body. Previous studies on scene affordance inference and human body synthesis in the literature, like [28, 46, 57], represent the body as a 3D stick figure or coarse volume. This prevents detailed reasoning about contact such as how the leg surface contacts the sofa surface. Without a model of body shape, it is not clear whether the estimated body poses correspond to plausible human poses. To overcome these issues, we use the SMPL-X model [36], which takes a set of low-dimensional body pose and shape parameters and outputs a 3D body

mesh with important details like the fingers. Since SMPL-X is differentiable, it enables straightforward optimization of human-scene contact and collision prevention [18]. In addition, we incorporate the body shape variation in our approach, so that our generated human bodies have various poses and shapes.

To train our method we exploit the PROX-Qualitative dataset [18], which includes 3D people captured moving in 3D scenes. We extend this by rendering images, scene depth, and semantic segmentation of the scene from many virtual cameras. We conduct extensive experiments to evaluate the performance of different models for scene-aware 3D body mesh generation. For testing, we extract 7 different rooms from the Matterport3D [3] dataset and use a virtual agent in the Habitat Simulator [32] to capture scene depth and semantics from different views. Based on prior work, e.g. [28, 46], we propose three metrics to evaluate the diversity, the physical plausibility, and the semantic plausibility of our results. The experimental results show that our solution effectively generates 3D body meshes in the scene, and outperforms the modified version of a state-of-the-art body generation method [28]. We will make our datasets and evaluation metrics available to establish a benchmark.

Our trained model learns about the ways in which 3D people interact with 3D scenes. We show how to leverage this in the form of a scene-dependent body pose prior and show how to use this to improve 3D body pose estimation from RGB images. In summary, our contributions are as follows: (1) We present a solution to generate 3D human bodies in scenes, using a CVAE to generate a body mesh with semantically plausible poses. We follow this with scene-geometry-aware fitting to refine the human-scene interaction. (2) We extend and modify two datasets, and propose three evaluation metrics for scene-aware human body generation. We also modify the method of [28] to generate body meshes as the baseline (see Sec. 4.1.2). The experimental results show that our method outperforms the baseline. (3) We show that our human-scene interaction prior is able to improve 3D pose estimation from RGB images.

2. Related work

Multiple studies focus on placing objects in an image so that they appear natural [11, 27, 29, 34]. For example, [11, 43, 45] use contextual information to predict which objects are likely to appear at a given location in the image. Lin et al. [29] apply homography transformations to 2D objects to approximate the perspectives of the object and background. Tan et al. [44] predict likely positions for people in an input image and retrieve a person that semantically fits into the scene from a database. Ouyang et al. [34] use a GAN framework to synthesize pedestrians in urban scenes. Lee et al. [27] learn where to place objects or people in a semantic map and then determine the pose and shape of the

respective object. However, all these methods are limited to 2D image compositing or inpainting. Furthermore, the methods that add synthetic humans do not take interactions between the humans and world into account.

To model human-object or human-scene interactions it is beneficial to know which interactions are possible with a given object. Such opportunities for interactions are referred to as affordances [14] and numerous works in computer vision have made use of this concept [7, 8, 15, 17, 21, 25, 24, 28, 38, 46, 56, 57]. Object affordance is often represented by a human pose when interacting with a given object [8, 15, 17, 21, 28, 38, 46, 56, 57]. For example, [15, 17, 57] search for valid positions of human poses in 3D scenes. Delataire et al. [8] learn associations between objects and human poses in order to improve object recognition. Given a 3D model of an object Kim et al. [21] predict human poses interacting with the given object. Given an image of an object Zhu et al. [56] learn a knowledge base to predict a likely human pose and a rough relative location of the object with respect to the pose. Savva et al. [38] learn a model connecting human poses and arrangement of objects in a 3D scene that can generate snapshots of object interaction given a corpus of 3D objects and a verb-noun pair. Monzpart et al. [33] use captured human motion to infer the objects in the scene and their arrangement. Sava et al. [37] predict action heat maps that highlight the likelihood of an action in the scene. Recently, Chen et al. [5] propose to tackle scene parsing and 3D pose estimation jointly and to leverage their coupled nature to improve scene understanding. Chao et al. [4] propose to train multiple controllers to imitate simple motions from mocap, and then use hierarchical reinforcement learning (RL) to achieve higher-level interactive tasks. The work of Zanfir et al. [50] first estimates the ground plane in the image, and requires a foreground person image as input. The above methods do not use a realistic body model to represent natural and detailed human-environment interactions.

Recently, Wang et al. [46] published an affordance dataset large enough to obtain reliable estimates of the probabilities of poses and to train neural networks on affordance prediction. The data is collected from multiple sitcoms and contains images of scenes with and without humans. The images with humans contain rich behavior of humans interacting with various objects. Given an image and a location as input, Wang et al. first predict the most likely pose from a set of 30 poses. This pose is deformed and scaled by a second network to fit it into the scene. Li et al. [28] extend this work to automatically estimate where to put people and to predict 3D poses. To acquire 3D training data they map 2D poses to 3D poses and place them in 3D scenes from the SUNCG dataset [40, 51]. This synthesized dataset is cleaned by removal of all predictions intersecting with the 3D scene or without sufficient support for the

body. The methods of [28, 46] are limited in their generalization, since they require a large amount of paired data and manual cleaning of the pose detections. Such a large amount of data might be hard to acquire for scenes that are less frequently covered by sitcoms, or in the case of [28] in 3D scene datasets. Furthermore, both methods only predict poses represented as *stick figures*. Such a representation is hard to validate visually, lacks details, and can not directly be used to generate realistic synthetic data of humans interacting with an environment.

3. Methods

3.1. Preliminaries

3D scene representation. We represent the scene from the view of an embodied agent, as in the Habitat simulator [32]. According to [52], which indicates that the depth and semantic segmentation are the most valuable modalities for scene understanding, we capture scene depth and semantics as our scene representation. For each view, we denote the stack of depth and semantics as x_s , the camera perspective projection from 3D to 2D as $\pi(\cdot)$, and its inverse operation as $\pi^{-1}(\cdot)$ for 3D recovery. Our training data, x_s , are generated from Habitat and we resize this to 128×128 for compatibility with our network; we retain the aspect ratio and pad with zeros where needed. The 3D-2D projection $\pi(\cdot)$ normalizes the 3D coordinates to the range of $[-1, 1]$, using the camera intrinsics and the maximal depth value. Note that each individual x_s is from a single camera view. We do not use multi-view data in our work.

3D human body representation. We use SMPL-X [36] to represent the 3D human body. SMPL-X can be regarded as a function $\mathcal{M}(\cdot)$, mapping a group of low-dimensional body features to a 3D body mesh. The 3D body mesh has 10475 vertices and a fixed topology. In our study, we use the body shape feature $\beta \in \mathbb{R}^{10}$, the body pose feature $\theta_b \in \mathbb{R}^{32}$, and the hand pose feature $\theta_h \in \mathbb{R}^{24}$. The body pose feature θ_b is represented in the latent space of VPoser [36], which is a variational autoencoder trained on a large-scale motion capture dataset, AMASS [31]. The global rotation R , i.e., the rotation of the pelvis, is represented by a 6D continuous rotation feature [55], which facilitates back-propagation in our trials. The global translation t is represented by a 3D vector in meters. The global rotation and translation is with respect to the camera coordinates. Based on the camera extrinsics, T_c^w , one can transform the 3D body mesh to the world coordinates.

We denote the joint body representation as $x_h := (t, R, \beta, \theta_b, \theta_h)^T \in \mathbb{R}^{75}$; i.e., the concatenation of individual body features. When processing the global and the local features separately as in [28], we denote the global translation as x_h^g , and the other body features as x_h^l .

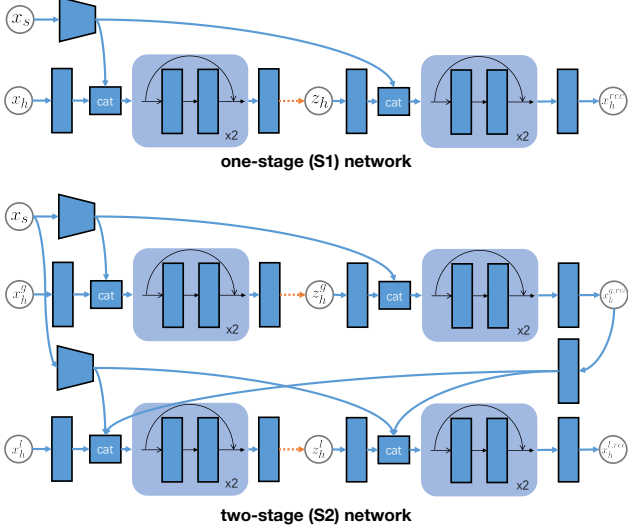


Figure 2: Network diagrams of our models. The trapezoids denote the scene encoders, which are fine-tuned from a pre-trained ResNet18 network. The blue rectangles denote fully-connected (fc) layers. In the residual blocks, Leaky-ReLU [30] is employed between fc layers. The orange dashed arrows denote the sampling operation in the VAE re-parameterization trick [23]. The blocks with “cat” denote the feature concatenation operation.

3.2. Scene context-aware Human Body Generator

3.2.1 Network architecture

We employ a conditional variational autoencoder (CVAE) [39] framework to model the probability $p(x_h|x_s)$. When inferring all body features jointly, we propose a one-stage (S1) network. When inferring x_h^g and x_h^l successively, we factorize the probability as $p(x_h^l|x_h^g, x_s)p(x_h^g|x_s)$ and use a two-stage (S2) network. The network architectures are illustrated in Fig. 2. Referring to [28], our scene encoder is fine-tuned from the first 6 convolutional layers in ResNet18 [19], which is pre-trained on ImageNet [9]. The human feature x_h is first lifted to a high dimension (256 in our study) via a fully-connected layer, and then concatenated with the encoded scene feature. In the two-stage model, the two scene encoders are both fine-tuned from ResNet18, but do not share parameters. After the first stage, the reconstructed body global feature $x_h^{g,rec}$ is further encoded, and is used in the second stage to infer the body local features.

3.2.2 Training loss

The entire training loss can be formulated as

$$\mathcal{L} = \mathcal{L}_{rec} + \alpha_{kl}\mathcal{L}_{KL} + \alpha_{vp}\mathcal{L}_{VPoser} + \alpha_{cont}\mathcal{L}_{contact} + \alpha_{coll}\mathcal{L}_{collision}, \quad (1)$$

where the terms denote the reconstruction loss, the Kullback–Leibler divergence, the VPoser loss, the human-scene contact loss and the human-scene collision loss, respectively. The set of α 's denotes the loss weights. For simplicity, we denote $\alpha_{cont}\mathcal{L}_{contact} + \alpha_{coll}\mathcal{L}_{collision}$ as \mathcal{L}_{HS} , implying the loss for human-scene interaction.

Reconstruction loss \mathcal{L}_{rec} : It is given by $\mathcal{L}_{rec} =$

$$\frac{|x_h^g - x_h^{g,rec}| + |\pi(x_h^g) - \pi(x_h^{g,rec})|}{2} + |x_h^l - x_h^{l,rec}|, \quad (2)$$

in which the global translation, the projected and normalized global translation, and the other body features are considered separately. We apply this reconstruction loss in both our S1 model and our S2 model.

KL-Divergence \mathcal{L}_{KL} : Denoting our VAE encoder as $q(z_h|x_h)$, the KL-divergence loss is given by

$$\mathcal{L}_{KL} = D_{KL}(q(z_h|x_h) || \mathcal{N}(0, \mathbf{I})). \quad (3)$$

Correspondingly, in our S2 model the KL-divergence loss is given by $\mathcal{L}_{KL} =$

$$D_{KL}(q(z_h^g|x_h^g) || \mathcal{N}(0, \mathbf{I})) + D_{KL}(q(z_h^l|x_h^l) || \mathcal{N}(0, \mathbf{I})). \quad (4)$$

We use the re-parameterization trick in [23] so that the KL divergence is closed form.

VPoser loss \mathcal{L}_{VPoser} : Since VPoser [36] attempts to encode natural poses with a normal distribution in its latent space, like in [36] and [18], we employ the VPoser loss, i.e.

$$\mathcal{L}_{VPoser} = |\theta_b^{rec}|^2, \quad (5)$$

to encourage the generated bodies to have natural poses.

Collision loss $\mathcal{L}_{collision}$: Based on the model output x_h^{rec} , we generate the body mesh and transform it to world coordinates. Then, we compute the negative signed-distance values at the body mesh vertices given the negative signed distance field (SDF) $\Psi_s^-(\cdot)$, and minimize

$$\mathcal{L}_{coll} = \mathbb{E} [|\Psi_s^-(T_c^w \mathcal{M}(x_h^{rec}))|]. \quad (6)$$

indicating the average of absolute values of negative SDFs on the body.

Contact loss $\mathcal{L}_{contact}$: Following [18], we encourage contact between the body mesh and the scene mesh. Hence, the contact loss is written as

$$\mathcal{L}_{contact} = \sum_{v_c \in C(T_c^w \mathcal{M}(x_h^{rec}))} \min_{v_s \in \mathcal{M}_s} \rho(|v_c - v_s|), \quad (7)$$

in which $C(\cdot)$ denotes selecting the body mesh vertices for contact according to the annotation in [18], \mathcal{M}_s denotes the scene mesh, and $\rho(\cdot)$ denotes the Geman-McClure error function [13] for down-weighting the influence of scene vertices far away from the body mesh.

3.3. Scene geometry-aware Fitting

We refine the body meshes with an optimization step similar to [18]. It encourages contact and helps to avoid inter-penetration between the body and the scene surfaces, while not deviating much from the generated pose. Let the generated human body configuration be x_h^0 . To refine this, we minimize a fitting loss taking into account the scene geometry, i.e.

$$\mathcal{L}_f(x_h) = |x_h - x_h^0| + \alpha_1 \mathcal{L}_{contact} + \alpha_2 \mathcal{L}_{collision} + \alpha_3 \mathcal{L}_{VPoser}, \quad (8)$$

in which the α 's denote the loss weights; the loss terms are defined above.

3.4. Implementation

Our implementation is based on PyTorch v1.2.0 [35]. For the Chamfer distance in the contact loss we use the same implementation as [10, 16]. For training, we set $\{\alpha_{kl}, \alpha_{vp}\} = \{0.1, 0.001\}$ in Eq. 1 for both our S1 and S2 models, in which α_{kl} increases linearly in an annealing scheme [1]. When additionally using \mathcal{L}_{HS} , we set $\{\alpha_{cont}, \alpha_{coll}\} = \{0.001, 0.01\}$, and enable it after 75% training epochs to improve the interaction modeling. We use the Adam optimizer [22] with the learning rate $3e^{-4}$, and terminate training after 30 epochs. For the scene geometry-aware fitting, we set $\{\alpha_1, \alpha_2, \alpha_3\} = \{0.1, 0.5, 0.01\}$ in all cases. Our data, code and models will be available for research purposes.

4. Experiments ¹

4.1. Scene-aware 3D Body Mesh Generation

4.1.1 Datasets

PROX-E: The **PROX-E** dataset (pronounced “proxy”) is extended from the PROX-Qualitative (PROX-Q) dataset [18], which records how people interact with various indoor environments. In PROX-Q, 3D human body meshes in individual frames are estimated by fitting the SMPL-X body model to the RGB-D data subject to scene constraints [18]. We use this data as pseudo-ground truth in our study, and extend PROX-Q in three ways: (1) We manually build up virtual walls, floors and ceilings to enclose the original open scans and simulate real indoor environments. (2) We manually annotate the mesh semantics following the object categorization of Matterport3D [3]. (3) We down-sample the original recordings and extract frames every 0.5 seconds. In each frame, we set up virtual cameras with various poses to capture scene depth and semantics. The optical axis of each virtual camera points towards the human body,

¹Please see appendix for more details.




Figure 3: Illustration of how we extend the PROX-Qualitative dataset [18] to **PROX-E**. In the rows of PROX-Qualitative, a video frame, a body-scene mesh and a depth map are shown from left to right. In the rows of **PROX-E**, the virtual camera setting, the mesh with semantics, and the completed depth map are shown from left to right.

and then Gaussian noise is applied on the camera translation. To avoid severe occlusion, all virtual cameras are located above half of the room height and below the virtual ceiling. As a result, we obtain about 70K frames in total. We use ‘MPH16’, ‘MPH1Library’, ‘NOSittingBooth’ and ‘N3OpenArea’ as test scenes, and use samples from other scenes for training. See Fig. 3.

MP3D-R: This name denotes “rooms in Matterport3D [3]”. From the architecture scans of Matterport3D, we extract 7 different rooms according to the annotated bounding boxes. In addition, we create a virtual agent using the Habitat simulator [32], and manipulate it to capture snapshots from various views in each room. We employ the RGB, the depth and the semantics sensor on the agent. These sensors are of height 1.8m from the ground, and look down at the scene; these are in a similar range as the virtual cameras in **PROX-E**. For each snapshot, we also record the extrinsic and intrinsic parameters of the sensors. As a result, we obtain 32 snapshots in all 7 rooms. Moreover, we follow the same procedure as in PROX-Qualitative [18] to calculate the SDF of the scene mesh. Our **MP3D-R** is illustrated in Fig. 4.

4.1.2 Baseline

To our knowledge, the most related work is Li et al. [28], which proposes a generative model to put 3D body stick figures into images². For fair comparison, we modify their method to use SMPL-X to generate body meshes in 3D

²The data and the pre-trained model in [28] are based on SUNCG [41], and not publicly available.




Figure 4: The left column shows two rooms in **MP3D-R**. The right column shows snapshots captured by the Habitat virtual agent [32] from different views, which contain RGB images, depth maps and scene semantics.

scenes. Specifically, we make the following modifications: (1) We change the scene representation from RGB (or RGB-D) to depth and semantics like ours to improve generalization. (2) During training, we perform K-means to cluster the VPoser pose features of training samples to generate the pose class. (3) The *where* module is used to infer the global translation, and the *what* module infers other SMPL-X parameters. (4) For training the geometry-aware discriminator, we project the body mesh vertices, rather than the stick figures, to the scene depth maps. We train the modified baseline model using **PROX-E** with the default architecture and loss weights in [28]. Moreover, we combine the modified baseline method with our scene geometry-aware fitting in our experiments.

4.1.3 Evaluation: representation power

Here we use **PROX-E** to investigate how well the proposed network architectures represent human-scene interaction. We train all models using samples from virtual cameras in training scenes, validate them using samples from real cameras in training scenes, and test them using samples from real cameras in test scenes. For quantitative evaluation, we feed individual test samples to our models, and report the mean of the reconstruction errors, and the negative evidenced lower bound (ELBO), i.e. $-\log P(X)$, which is the sum of the reconstruction error and the KL divergence. For fair comparison, the reconstruction error of all models is based on \mathcal{L}_{rec} in Eq. 2. As shown in Tab. 1, our models outperform the baseline model on both validation and test set by large margins. The metrics on the validation and the test sets are comparable, indicating that our virtual camera approach is effective in preventing severe over-fitting on the seen environments.

Table 1: Comparison between models, in which “ $+\mathcal{L}_{HS}$ ” denotes the model is trained with that human-scene interaction loss (Sec. 3.2.2). The best results are in boldface.

model	rec. err.		$-\log P(x)$	
	val	test	val	test
baseline [28]	0.52	0.48	0.98	0.72
S1	0.22	0.25	0.23	0.41
S1 + \mathcal{L}_{HS}	0.16	0.24	0.27	0.36
S2	0.24	0.70	0.25	0.49
S2 + \mathcal{L}_{HS}	0.20	0.23	0.30	0.39

4.1.4 Evaluation: 3D body mesh generation

Given a 3D scene, our goal is to generate diverse, physically and semantically plausible 3D human bodies. Based on [28, 46], we propose to quantitatively evaluate our method using a diversity metric and a physical metric. Also, we perform a user perceptual study to measure the semantic plausibility of the generated human bodies.

The quantitative evaluation is based on the **PROX-E** and the **MP3D-R** dataset. When testing on **PROX-E**, we train our models using all samples in the training scenes, and generate body meshes using the real camera snapshots in the testing scenes. For each individual model and each test scene, we randomly generate 1200 samples, and hence obtain 4800 samples. When testing on **MP3D-R**, we use all samples from **PROX-E** to train the models. For each snapshot and each individual model, we randomly generate 200 samples, and hence obtain 6400 samples.

(1) Diversity metric: This metric aims to evaluate how diverse the generated human bodies are. Specifically, we empirically perform K-means to cluster the SMPL-X parameters of all the generated human bodies to 20 clusters. Then, we compute the entropy (a.k.a Shannon index, a type of diversity index) of the cluster ID histogram of all the samples. We also compute the average size of all the clusters.

A higher value indicates that the generated human bodies are more diverse in terms of their global positions, their body shapes and poses. We argue that this metric is essential for evaluating the quality of the generated bodies and should always be considered together with other metrics. For instance, a posterior-collapsed VAE, which always generates an identical body mesh, could lead to a low diversity score but superior performance according to the physical metric and the semantic metric.

The results are shown in Tab. 2. Overall, our methods consistently outperform the baseline. Notably, our methods increase the average cluster size of the generated samples by large margins, indicating that the generated human bodies are much more diverse than those from the baseline.

(2) Physical metric: From the physical perspective, we evaluate the collision and the contact between the body mesh and the scene mesh. Given a scene SDF and a SMPL-

Table 2: Comparison between different models according to the diversity metric. The best results for each metric are in boldface. “S1” and “S2” denote our stage-1 and stage-2 architecture, respectively. “+ \mathcal{L}_{HS} ” denotes that the model is trained with the human-scene interaction loss (see Sec. 3.2.2). “+ \mathcal{L}_f ” denotes the results are after the scene-aware fitting process (see Sec. 3.3).

model	cluster ID entropy		cluster size average	
	PROX-E	MP3D-R	PROX-E	MP3D-R
baseline [28]	2.89	2.93	1.49	1.84
S1	2.96	2.99	2.51	2.81
S1 + \mathcal{L}_{HS}	2.93	2.99	2.40	2.73
S2	2.97	2.91	2.46	2.85
S2 + \mathcal{L}_{HS}	2.96	2.89	2.22	2.90
baseline + \mathcal{L}_f	2.93	2.92	1.52	1.94
S1 + \mathcal{L}_f	2.97	2.98	2.53	2.86
S1 + \mathcal{L}_{HS} + \mathcal{L}_f	2.94	2.96	2.43	2.79
S2 + \mathcal{L}_f	2.94	2.87	2.48	2.91
S2 + \mathcal{L}_{HS} + \mathcal{L}_f	2.91	2.90	2.26	2.95

X body mesh, we propose a non-collision score, which is calculated as the number of body mesh vertices with positive SDF values divided by the number of all body mesh vertices (10475 for SMPL-X). Simultaneously, if any body mesh vertex has a non-positive SDF value, then the body has contact with the scene. Then, for all generated body meshes, the non-collision score is the ratio of all body vertices in the free space, and the contact ratio is calculated as the number of body meshes with contact divided by all generated body meshes. Therefore, due to the physical constraints, a higher non-collision score and contact ratio indicate a better generation, in analogy with precision and recall in an object detection task.

The results are presented in Tab. 3. First, one can see that our proposed methods consistently outperform the baseline for the physical metric. The influence of the \mathcal{L}_{HS} loss on 3D body generation is not as obvious as on the interaction modeling task (see Tab. 1). Additionally, one can see that the scene geometry-aware fitting consistently improves the physical metric, since the fitting process aims to improve the physical plausibility. Fig. 7 shows some generated examples before and after the fitting.

(3) User study: In our study, we render our generated results as images, and upload them to Amazon Mechanical Turk (AMT) for a user study. Due to the superior performance of our S1 model without \mathcal{L}_{HS} , we compare it with the baseline, as well as ground truth if it exists. For each scene and each model, we generate 100 and 400 bodies in **PROX-E** and **MP3D-R**, respectively, and ask Turkers to give a score between 1 (strongly not natural) and 5 (strongly natural) to each individual result. The user study details are in the appendix. Also, for each scene in the **PROX-E** dataset, we randomly select 100 frames from the ground

Table 3: Comparison between different models according to the physical metric. The best results are in boldface.

model	non-collision score		contact score	
	PROX-E	MP3D-R	PROX-E	MP3D-R
baseline [28]	0.89	0.92	0.93	0.78
S1	0.93	0.94	0.95	0.80
S1 + \mathcal{L}_{HS}	0.89	0.95	0.88	0.65
S2	0.91	0.93	0.88	0.79
S2 + \mathcal{L}_{HS}	0.89	0.95	0.88	0.56
baseline + \mathcal{L}_f	0.93	0.97	0.99	0.89
S1 + \mathcal{L}_f	0.94	0.97	0.99	0.88
S1 + \mathcal{L}_{HS} + \mathcal{L}_f	0.92	0.98	0.99	0.81
S2 + \mathcal{L}_f	0.94	0.97	0.99	0.88
S2 + \mathcal{L}_{HS} + \mathcal{L}_f	0.93	0.97	0.99	0.81

Table 4: Comparison between models in the user study score (1-5). The best results for each metric are in boldface.

model	use study score w.r.t. mean \pm std	
	PROX-E	MP3D-R
baseline [28]	3.31 \pm 1.39	3.14 \pm 1.41
baseline + \mathcal{L}_f	3.32 \pm 1.35	3.35 \pm 1.38
S1	3.29 \pm 1.36	3.15 \pm 1.40
S1 + \mathcal{L}_f	3.49 \pm 1.26	3.30 \pm 1.30
ground truth	4.04 \pm 1.03	n/a

truth [18], and ask Turkers to evaluate them as well.

The results are presented in Tab. 4. Not surprisingly, the ground-truth samples achieve the best score from the user study. We observe that the geometry-aware fitting improves the performance both for the baseline and our model, most likely due to the improvement of the physical plausibility. Note that, although the baseline and our model achieve similar average scores, the diversity of our generated samples is much higher (Tab. 2). This indicates that, compared to the baseline, our method generates more diverse 3D human bodies, while being equally good in terms of semantic plausibility given a 3D scene.

Qualitative results are shown in Fig. 5, Fig. 6 and Fig. 8. More results are in the appendix.

4.2. Scene-aware 3D Body Pose Estimation

Here we perform a down-stream application and show our model improves 3D human pose estimation from monocular images. Given a RGB image of a scene without people, we estimate the depth map using the pre-trained model [26], and perform semantic segmentation using the model of [6] pre-trained on the ADE20K [53] dataset. To unify the semantics, we create a look-up table to convert the object IDs from ADE20K to Matterport3D. Next, we feed the estimated depth and semantics to our S1 model with \mathcal{L}_{HS} and randomly generate 100 bodies. We compute the average of the pose features in the VPoser latent space, and denote it as θ_b^s .




Figure 5: Generated human bodies in two test scenes of **PROX-E**. Results are visualized in two views.



Figure 6: Generated results in three scenes of **MP3D-R**.




Figure 7: Results before and after the scene geometry-aware fitting.



Figure 8: Two typical failure cases in our results.

When performing 3D pose estimation in the same scene, we follow the optimization framework of SMPLify-X [36] and PROX [18]. In contrast to these two methods, we use our derived θ_b^s to initialize the optimization, and change the

Table 5: Results of 3D pose estimation from RGB frames in **PROX-Quantitative**, in which “PJE”/“p.PJE” denote the mean per-joint error without/with Procrustes alignment, and “V2V”/“p.V2V” denote the mean vertex-to-vertex error without/with Procrustes alignment, respectively.

method	Error (in millimeters)			
	PJE	V2V	p.PJE	p.V2V
Simplify-X [36]	223.83	225.60	73.28	62.93
PROX [18]	171.78	173.97	73.20	64.76
Ours	174.10	171.75	71.73	62.64

VPoser term in [18, Eq. 7] from $|\theta_b|^2$ to $|\theta_b - \theta_b^s|^2$. We evaluate the performance using the **PROX-Quantitative** dataset [18]. We derive the 2D keypoints from the frames via AlphaPose [12, 48], and obtain a θ_b^s from a background image without people. Then, we use the same optimization methods and the evaluation metric in [18] for fair comparison. The results are shown in Tab. 5. We find that our method improves 3D pose estimation on the **PROX-Quantitative** dataset. This suggests that our model learns about the ways in which 3D people interact with 3D scenes. Leveraging it as a scene-dependent body pose prior can improve 3D body pose estimation from RGB images.

5. Conclusion

In this work, we introduce a generative framework to produce 3D human bodies that are posed naturally in the 3D environment. Our method consists of two steps: (1) A scene context-aware human body generator is proposed to learn a distribution of 3D human pose and shape, conditioned on the scene depth and semantics; (2) geometry-aware fitting is employed to impose physical plausibility of the human-scene interaction. Our experiments demonstrate that the automatically synthesized 3D human bodies are realistic and expressive, and interact with 3D environment in a semantic and physical plausible way.

Acknowledgments. We sincerely acknowledge: Joachim Tesch for all his work on graphics supports. Xueting Li for implementation advice about the work [28]. David Hoffmann, Jinlong Yang, Vasileios Choutas, Ahmed Osman, Nima Ghorbani and Dimitrios Tzionas for insightful discussions. Daniel Scharstein and Cornelia Köhler for proof reading. Benjamin Pellkofer and Mason Landry for IT/hardware supports. Y. Z. and S. T. acknowledge funding by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Projektnummer 276693517 SFB 1233.

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] Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. In *Proceedings of the 20th Conference on Computational Natural Language Learning (SIGNLL)*, 2016. 5
- [2] Christopher P Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. Understanding disentangling in β -VAE. *arXiv preprint arXiv:1804.03599*, 2018. 15
- [3] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D data in indoor environments. In *Proceedings of the International Conference on 3D Vision (3DV)*, 2017. 2, 5, 12, 13
- [4] Yu-Wei Chao, Jimei Yang, Weifeng Chen, and Jia Deng. Learning to sit: Synthesizing human-chair interactions via hierarchical control. *arXiv preprint arXiv:1908.07423*, 2019. 3
- [5] Yixin Chen, Siyuan Huang, Tao Yuan, Siyuan Qi, Yixin Zhu, and Song-Chun Zhu. Holistic++ scene understanding: Single-view 3D holistic scene parsing and human pose estimation with human-object interaction and physical common-sense. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 3
- [6] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 7
- [7] Ching-Yao Chuang, Jiaman Li, Antonio Torralba, and Sanja Fidler. Learning to act properly: Predicting and explaining affordances from images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [8] Vincent Delaitre, David F Fouhey, Ivan Laptev, Josef Sivic, Abhinav Gupta, and Alexei A Efros. Scene semantics from long-term observation of people. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2012. 3
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 4
- [10] Theo Deprelle, Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. Learning elementary structures for 3D shape generation and matching. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 5
- [11] Nikita Dvornik, Julien Mairal, and Cordelia Schmid. On the importance of visual context for data augmentation in scene understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, PP:1–1, 12 2019. 2
- [12] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. RMPE: Regional multi-person pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 8
- [13] Stuart Geman and Donald E. McClure. Statistical methods for tomographic image reconstruction. In *Proceedings of the 46th Session of the International Statistical Institute, Bulletin of the ISI*, volume 52, 1987. 4
- [14] James J Gibson. *The ecological approach to visual perception: classic edition*. Psychology Press, 2014. 3
- [15] Helmut Grabner, Juergen Gall, and Luc Van Gool. What makes a chair a chair? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011. 3
- [16] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan Russell, and Mathieu Aubry. 3D-CODED: 3D correspondences by deep deformation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 5
- [17] Abhinav Gupta, Scott Satkin, Alexei A Efros, and Martial Hebert. From 3D scene geometry to human workspace. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011. 3
- [18] Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J. Black. Resolving 3D human pose ambiguities with 3D scene constraints. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 2, 4, 5, 7, 8, 12, 13, 15, 16
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4
- [20] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. β -VAE: Learning basic visual concepts with a constrained variational framework. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017. 15
- [21] Vladimir G Kim, Siddhartha Chaudhuri, Leonidas Guibas, and Thomas Funkhouser. Shape2Pose: Human-centric shape analysis. *ACM Transactions on Graphics (TOG)*, 33(4):120, 2014. 3
- [22] Diederick P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015. 5
- [23] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014. 4
- [24] Hema Swetha Koppula, Rudhir Gupta, and Ashutosh Saxena. Learning human activities and object affordances from RGBD videos. *International Journal of Robotics Research*, 32(8):951–970, 2013. 3
- [25] Hema S Koppula and Ashutosh Saxena. Physically grounded spatio-temporal object affordances. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014. 3
- [26] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *Proceedings of the International Conference on 3D Vision (3DV)*, 2016. 7
- [27] Donghoon Lee, Sifei Liu, Jinwei Gu, Ming-Yu Liu, Ming-Hsuan Yang, and Jan Kautz. Context-aware synthesis and

- placement of object instances. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 2
- [28] Xueting Li, Sifei Liu, Kihwan Kim, Xiaolong Wang, Ming-Hsuan Yang, and Jan 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, 4, 5, 6, 7, 8, 14, 15, 19
- [29] Chen-Hsuan Lin, Ersin Yumer, Oliver Wang, Eli Shechtman, and Simon Lucey. ST-GAN: Spatial transformer generative adversarial networks for image compositing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [30] Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2013. 4
- [31] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 3
- [32] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A Platform for Embodied AI Research. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 2, 3, 5, 6, 13
- [33] Aron Monzspart, Paul Guerrero, Duygu Ceylan, Ersin Yumer, and Niloy J. Mitra. iMapper: Interaction-guided scene mapping from monocular videos. *ACM Transactions on Graphics (TOG)*, 38(4):15, 2019. 3
- [34] Xi Ouyang, Yu Cheng, Yifan Jiang, Chun-Liang Li, and Pan Zhou. Pedestrian-Synthesis-GAN: Generating pedestrian data in real scene and beyond. *arXiv preprint arXiv:1804.02047*, 2018. 2
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*. 2019. 5
- [36] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 3, 4, 8, 14, 15
- [37] Manolis Savva, Angel X Chang, Pat Hanrahan, Matthew Fisher, and Matthias Nießner. SceneGrok: Inferring action maps in 3D environments. *ACM Transactions on Graphics (TOG)*, 33(6):212, 2014. 3
- [38] Manolis Savva, Angel X Chang, Pat Hanrahan, Matthew Fisher, and Matthias Nießner. PiGraphs: Learning inter-action snapshots from observations. *ACM Transactions on Graphics (TOG)*, 35(4):139, 2016. 3
- [39] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2015. 2, 4
- [40] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [41] Shuran Song, Fisher Yu, Andy Zeng, Angel X. Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5
- [42] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. 2
- [43] Jin Sun and David W Jacobs. Seeing what is not there: Learning context to determine where objects are missing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [44] Fuwen Tan, Crispin Bernier, Benjamin Cohen, Vicente Ordonez, and Connelly Barnes. Where and who? Automatic semantic-aware person composition. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2018. 2
- [45] Antonio Torralba. Contextual priming for object detection. *International Journal of Computer Vision (IJCV)*, 53(2):169–191, 2003. 2
- [46] Xiaolong Wang, Rohit Girdhar, and Abhinav Gupta. Binge Watching: Scaling affordance learning from sitcoms. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 3, 6, 14, 15
- [47] Fei Xia, Amir R. Zamir, Zhi-Yang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson Env: Real-world perception for embodied agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [48] Yuliang Xiu, Jiefeng Li, Haoyu Wang, Yinghong Fang, and Cewu Lu. Pose Flow: Efficient online pose tracking. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2018. 8
- [49] Hsiao-chen You and Kuohsiang Chen. Applications of affordance and semantics in product design. *Design Studies*, 28(1):23–38, 2007. 2
- [50] Mihai Zanfir, Elisabeta Oneata, Alin-Ionut Popa, Andrei Zanfir, and Cristian Sminchisescu. Human synthesis and scene compositing. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2020. 3

- [51] Yinda Zhang, Shuran Song, Ersin Yumer, Manolis Savva, Joon-Young Lee, Hailin Jin, and Thomas Funkhouser. Physically-based rendering for indoor scene understanding using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [52] Brady Zhou, Philipp Krähenbühl, and Vladlen Koltun. Does computer vision matter for action? *Science Robotics*, 4(30), 2019. 3, 15
- [53] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ADE20K dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 7
- [54] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3D: A modern library for 3D data processing. *arXiv:1801.09847*, 2018. 12
- [55] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [56] Yuke Zhu, Alireza Fathi, and Li Fei-Fei. Reasoning about object affordances in a knowledge base representation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014. 3
- [57] Yixin Zhu, Chenfanfu Jiang, Yibiao Zhao, Demetri Terzopoulos, and Song-Chun Zhu. Inferring forces and learning human utilities from videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 3