

Compositional Human-Scene Interaction Synthesis with Semantic Control

Kaifeng Zhao¹, Shaofei Wang¹, Yan Zhang¹, Thabo Beeler², and Siyu Tang¹

¹ ETH Zürich

{kaifeng.zhao, shaofei.wang, yan.zhang, siyu.tang}@inf.ethz.ch

² Google

thabo.beeler@gmail.com

Abstract. Synthesizing natural interactions between virtual humans and their 3D environments is critical for numerous applications, such as computer games and AR/VR experiences. Recent methods mainly focus on modeling geometric relations between 3D environments and humans, where the high-level semantics of the human-scene interaction has frequently been ignored. Our goal is to synthesize humans interacting with a given 3D scene controlled by high-level semantic specifications as pairs of action categories and object instances, e.g., “sit on the chair”. The key challenge of incorporating interaction semantics into the generation framework is to learn a joint representation that effectively captures heterogeneous information, including human body articulation, 3D object geometry, and the intent of the interaction. To address this challenge, we design a novel transformer-based generative model, in which the articulated 3D human body surface points and 3D objects are jointly encoded in a unified latent space, and the semantics of the interaction between the human and objects are embedded via positional encoding. Furthermore, inspired by the compositional nature of interactions that humans can simultaneously interact with multiple objects, we define interaction semantics as the composition of varying numbers of atomic action-object pairs. Our proposed generative model can naturally incorporate varying numbers of atomic interactions, which enables synthesizing compositional human-scene interactions without requiring composite interaction data. We extend the PROX dataset with interaction semantic labels and scene instance segmentation to evaluate our method and demonstrate that our method can generate realistic human-scene interactions with semantic control. Our perceptual study shows that our synthesized virtual humans can naturally interact with 3D scenes, considerably outperforming existing methods. We name our method **COINS**, for **CO**mpositional **I**nteraction **S**ynthesis with **S**emantic **C**ontrol. Code and data are available at <https://github.com/zkf1997/COINS>.

Keywords: human-scene interaction synthesis, semantic composition, virtual humans.



Fig. 1: Given a pair of action and object instance as the semantic specification, our method generates virtual humans naturally interacting with the object (*first row*). Furthermore, our method retargets interactions on unseen action-object combinations (*second row*) and synthesizes composite interactions without requiring any corresponding composite training data (*third row*).

1 Introduction

People constantly interact with their surroundings, and such interactions have semantics, specifically as combinations of actions and object instances, e.g., “sit on the chair” or “lie on the bed”. Incorporating and controlling such interaction semantics is critical for creating virtual humans with realistic behavior. An effective solution for semantic-aware virtual human generation could advance existing technologies in AR/VR, computer games, and synthetic data generation to train machine learning perception algorithms. However, existing methods focus on modeling geometric relationships between virtual humans and their 3D environments [11, 45, 46], and thus are not able to synthesize body-scene interactions with semantic control of human actions and specific interaction objects. This heavily limits their use in practice.

In this paper, we aim to incorporate semantic control into the synthesis of human-scene interactions. This is a challenging task due to the following reasons. First, semantic-aware human-scene interactions are characterized by body artic-

ulation, 3D object geometry, and the intent of the interaction. It is challenging to learn a generalizable representation to capture such heterogeneous information. Second, given a 3D scene, the space for plausible human-scene interactions is enormous, and obtaining sufficiently diverse training data remains challenging.

To address these challenges, we introduce a novel interaction synthesis approach that leverages transformer-based generative models and can synthesize realistic 3D human bodies, given a 3D scene and a varying number of action-object pairs. These pairs can specify the intent of the interactions and the interacting objects. The key advantages of our models are two-fold: First, we represent human body articulation and 3D object geometry as tokens in the proposed transformer network, and the interaction semantics are embedded via positional encoding. Combined with a conditional variational auto-encoder (VAE) [16, 33], we learn a unified latent space that captures the distribution of human articulation conditioned on the given 3D objects and the interaction semantics. Second, our models can be used to synthesize composite interactions, as they can naturally incorporate a varying number of 3D objects and interaction semantics. By training only on irreducible atomic interactions (e.g., “sit on the sofa”, “touch the table”), our model can generate novel composite interactions (e.g., “sit on the sofa *and* touch the table”) in unseen 3D scenes without requiring composite training data (Fig. 1). Furthermore, we decompose the interaction synthesis task into three stages. First, we infer plausible global body locations and orientations given pairs of actions and objects. Second, we generate detailed body articulation and body-scene contact that aligns with the object geometry and interaction semantics. Lastly, we further refine the local body-scene contact by leveraging the inferred contact map.

To train and evaluate our method, we extend the PROX dataset [10] to the PROX-S dataset, which includes 3D instance segmentation, SMPL-X body estimation, and per-frame annotation of interaction semantics. Since there are no existing methods working on the same task of populating scenes with semantic control, we adapt PiGraph [31] and POSA [11] to this new task, and train their modified versions on the PROX-S dataset as baselines. Our perceptual study shows that our proposed method can generate realistic 3D human bodies in a 3D scene with semantic control, significantly outperforming those baselines.

In summary, our contributions are: (1) a novel compositional representation for interaction with semantics; (2) a generative model for synthesizing diverse and realistic human-scene interactions from semantic specifications; (3) a framework for composing atomic interactions to generate composite interactions; and (4) the extended human-scene interaction dataset PROX-S containing scene instance segmentation and per-frame interaction semantic annotation.

2 Related Work

Human-Scene Interaction Synthesis. Synthesizing human-scene interactions has been a challenging problem in the computer vision and graphics community. Various methods have been proposed to analyze object affordances and gener-

ate static human-scene interactions [6, 8, 12, 15, 17, 30, 31]. However, the realism of synthesized interactions from these methods is often limited by their oversimplified human body representations. The recently published PSI [46] uses the parametric body model SMPL-X [23] to put 3D human bodies into scenes, conditioned on scene semantic segmentation and depth map. In addition, geometric constraints are applied to resolve human-scene contact and penetration problems. PLACE [45] learns the human-scene proximity and contact in interactions to infer physically plausible humans in 3D scenes. POSA [11] proposes an egocentric contact feature map that encodes per-vertex contact distance and semantics information. A stochastic model is trained to predict contact feature maps given body meshes, with the predicted contact features guiding the body placement optimization in scenes. Wang et al. [40] generate human bodies at specified locations and orientations, conditioning on the point cloud of the scene and implicitly modeling the scene affordance. Concurrent work [41] first generates scene-agnostic body poses from actions and then places body into scene using POSA, which is similar to the POSA-I baseline we implement.

Apart from being realistic, semantic control is another important goal in synthesizing interactions. PiGraph [31] learns a probabilistic distribution for each annotated interaction category and synthesizes interactions by selecting the highest scoring samples of fitted categorical distribution. Conditional action synthesis methods [1, 2, 24] generate human motion conditioned on action categories or text descriptions while ignoring human-scene interaction. NSM [34] uses a mixture of experts modulated by the action category to generate human motion of seven predefined categories. SAMP [9] extends NSM by incorporating stochastic motion modeling and a GoalNet predicting possible interaction goals given scene objects and action categories. NSM and SAMP are limited in interaction semantics modeling since they can only model interactions with one single object or the surrounding environment as a whole. Our work, in contrast, models the compositional nature of interaction semantics and can generate composite interactions involving multiple objects simultaneously.

Human Behavior Semantics Modelling. Understanding human behavior is an important problem that has long received a significant amount of attention. Action recognition [13, 14, 32, 37] classifies human activities into predefined action categories. Language driven pose and motion generation [1, 2, 35] incorporates natural languages that can represent more expressive and fine-grained action semantics. Human-object interaction detection [5, 7, 42] models interactions as triplets of (human, action, object), detecting humans and objects from images along with the relations between them.

One key observation is that humans can perform multiple actions simultaneously and such behaviors have composite semantics. Composite semantics are often modeled as the combination of basic semantic units, like how phrases are composed by words in natural languages. Semantic composition is termed as compound in linguistics [18, 25] and studied in natural language processing [21, 22, 43]. PiGraph [31] describes composite human-scene interaction semantics with a set of (verb, noun) pairs. BABEL [26] uses natural language

labels to annotate human motion capture data [3, 19, 20, 36, 38] and allow multiple descriptions existing in a single frame to model composite semantics. These works model semantics on the object category level and ignore the differences among instances of the same category. Our work goes beyond this limitation to model interaction semantics with specific object instances, which is critical for solving ambiguity in scenes with multiple objects of the same category.

3 Method

3.1 Preliminaries

Body Representation. We represent the 3D human body using a simplified SMPL-X model, consisting of vertex locations, triangle faces, and auxiliary per-vertex binary contact features, denoted as $\mathbf{B} = (\mathbf{V}, \mathbf{F}, \mathbf{C})$. For computational efficiency, we downsample the original SMPL-X mesh to have 655 vertices $\mathbf{V} \in \mathbb{R}^{655 \times 3}$ and 1296 faces $\mathbf{F} \in \{1 \dots 655\}^{1296 \times 3}$ using the sampling method of [28]. The binary contact features $\mathbf{C} \in \{0, 1\}^{655}$ indicate whether each body vertex is in contact with scene objects [11]. Specifically, we consider a body vertex is in contact if its distance to scene objects is below a threshold value. We introduce the auxiliary contact features because human-object contact is a strong cue for interaction semantics, e.g., “touch” implies contact on hands. We use the contact features to further refine human-scene interaction in post-optimization.

3D Scene Representation. We represent a 3D scene as a set of object instances, denoted as $\mathbf{S} = \{o^i\}_{i=1}^O$ where O is the number of object instances. Each object instance $o^i \in \mathbb{R}^{8192 \times 9}$ is an oriented point cloud of 8192 points with attributes of location, color, and normal. We represent objects as point clouds to incorporate the geometry information and enable generalization to open-set objects that do not belong to predefined object categories, unlike previous works [11, 31, 46] that only use object category information. Moreover, the effect of intra-class object variance on interactions can be learned from the object geometry to improve fine-grained human-scene contacts.

Compositional Interaction Representation. Human-scene interactions involve three types of components: human bodies, actions, and objects. Inspired by the observation that humans can perform multiple interactions at the same time, we propose a compositional interaction representation $\mathbf{I} = (\mathbf{B}, \{(a^i, o^i)\}_{i=1}^M)$, where \mathbf{B} denotes the human body, and $\{(a^i, o^i)\}_{i=1}^M$ is a set of atomic interactions. We use the term atomic interaction to refer to the basic, irreducible building block for interaction semantics, and each atomic interaction (a^i, o^i) is a pair of a one-hot encoded action and an object instance, e.g., “sit on a chair”. In addition, interactions can comprise more than one such atomic interaction unit, since one can simultaneously interact with multiple objects, e.g., “sit on a chair and type on a keyboard”. We refer to such interactions consisting of more than one atomic interaction as composite interactions.

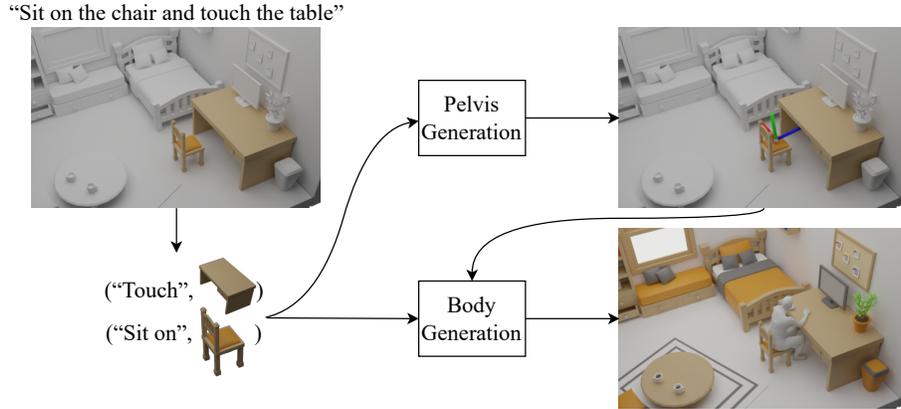


Fig. 2: Illustration of our COINS method for human-scene interaction synthesis with semantic control. Given a 3D scene and semantic specifications of actions (e.g., “sit on”, “touch”) paired with object instances (highlighted with color), COINS first generates a plausible human pelvis location and orientation and then generates the detailed body.

3.2 Interaction Synthesis with Semantic Control

Synthesizing interactions with semantic control requires generating human bodies interacting with given scene objects in a natural and semantically correct way. Specifically, given a scene S and interaction semantics specified as $\{(a^i, o^i)\}_{i=1}^M$, the task is to generate a human body \mathbf{B} that perceptually satisfies the constraints of performing action a^i with object o^i .

We propose a multi-stage method for interaction synthesis with semantic control as illustrated in Fig. 2. Given interaction semantics specified as action-object pairs, we first infer possible pelvis locations and orientations. Then we transform objects to the generated pelvis coordinate frame and sample body and contact map in the local space. We separately train two transformer-based conditional VAE (cVAE) [16, 33] models for pelvis and body generation, and name them as PelvisVAE and BodyVAE respectively. Lastly, we apply an interaction-based optimization that is guided by contact features to refine the generated interactions.

Architecture. We propose a transformer-based conditional variational auto-encoder (cVAE) architecture that can handle the heterogeneous inputs of our compositional interaction representation containing body mesh vertices, object points without predefined topology, and interaction semantics. Specifically, object points and body vertices are represented as token inputs to the transformer, and learnable action embeddings are added to paired object tokens as positional encoding.

The BodyVAE architecture is illustrated in Fig. 3. The PelvisVAE shares a similar architecture and we refer readers to the Supp. Mat. for illustration.

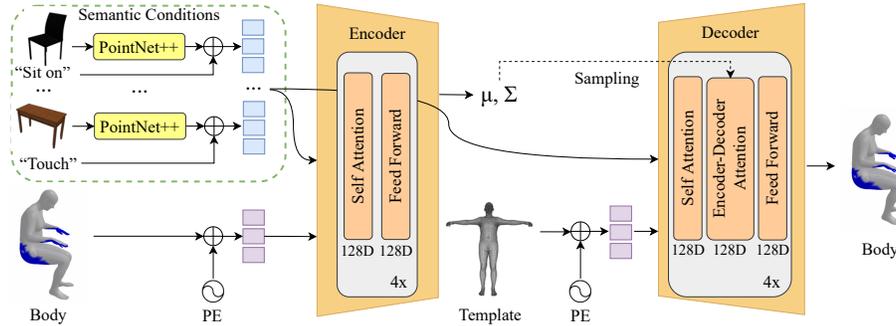


Fig. 3: Illustration of BodyVAE architecture. Human body and objects are represented as tokens of the transformers while each action category is mapped to a learnable embedding and added to the corresponding object tokens. Modules within the dashed box are the semantic conditions of a varying number of action-object pairs. We denote addition with \oplus and positional encoding with “PE”. We visualize vertices with positive contact features as blue on the body mesh.

Our cVAE networks take varying numbers of action-object pairs as conditional inputs. Each action category is mapped to a learnable embedding, and each object point cloud is first converted to 256 local keypoints using PointNet++ [27]. The action embedding is then added to the tokens of the paired object as positional encoding. Compared to treating actions as separate tokens [24], encoding actions into the positional encoding for object tokens can regulate each action to only affect the paired object, e.g., the action “sit on” is only related to the sofa, but not the table, in the example of Fig. 3. The human body input comprises mesh vertices with location and auxiliary contact features. We apply a linear layer to map vertices to tokens and vertex index based positional encoding to encode body topology.

The transformer encoder maps input interactions into a latent space and the decoder predicts human bodies with contact features from conditional action-object pairs and the sampled latent code. On top of the encoder transformer network, we apply average pooling on the output human tokens followed by a linear layer to predict the latent normal distribution. At the decoder, we use a template body with person-dependent shape parameters for body tokens. Using the personalized template body makes our model generalize to different body shapes and enables explicit control of the body shape. The sampled latent code is used for keys and values of the encoder-decoder attention layers [39].

BodyVAE shares the same transformer-based architecture and the same conditional input of action-object pairs as PelvisVAE. There are mainly two differences between BodyVAE and PelvisVAE. First, BodyVAE represents a human as a body mesh with contact features, while PelvisVAE represents a human as the pelvis joint location and orientation. Second, all input objects are in the

scene coordinate frame in PelvisVAE while input objects are transformed to the local pelvis coordinate frame in BodyVAE.

To make our sampled bodies compatible with the SMPL-X body model and ensure a valid body shape, we jointly train a Multi-Layer Perceptron (MLP) [29] with the BodyVAE to regress the corresponding SMPL-X body parameters from the sampled body meshes. We refer to the Supp. Mat. for more details on architecture and the training losses.

Interaction-Based Optimization. Our model is trained with the pseudo ground truth body estimation [10, 44], which does not guarantee natural human-scene contact. Therefore, the generated interactions can also be unnatural. We follow [11, 45, 46] to apply interaction-based optimization to refine the interactions. We use generated SMPL-X parameters as initialization and optimize body translation t , global orientation R and pose θ . The optimization objective is given by:

$$E(t, R, \theta) = \mathcal{L}_{interaction} + \mathcal{L}_{coll} + \mathcal{L}_{reg}, \quad (1)$$

where $\mathcal{L}_{interaction}$ encourages body vertices with positive contact features to have zero distance to objects, \mathcal{L}_{coll} is the scene SDF based collision term defined in [46], and the \mathcal{L}_{reg} term penalizes t, R and θ deviating from the initialization. We refer to the Supp. Mat. for detailed definitions.

3.3 Compositional Interaction Generation

Compositional interaction generation aims to generate composite interactions using models trained only on atomic interactions. Compositional interaction generation reflects the compositional nature of interaction semantics and can save the need to collect composite interaction data. Although our transformer-based generative model can synthesize bodies given composite interaction semantics, e.g. “sit on the sofa and touch the wall”, it requires collecting data and training for all composite interactions we want to synthesize. However, collecting composite interaction data is exponentially more expensive than atomic interactions considering the vast possible combinations of atomic interactions. Therefore, it is desirable to have an automatic way to generate composite interactions requiring only atomic interaction data.

To that end, we propose a simple yet effective two-step approach to generating composite interactions using models trained only on atomic interactions: (1) Compositional pelvis generation, where we formulate the task of generating pelvis frames given composite interaction semantics as finding the intersection of pelvis distributions of atomic interactions. (2) Compositional body generation, where we leverage attention masks derived from contact statistics in transformer inference to compose atomic interactions at the body parts level.

Compositional Pelvis Generation. We formulate compositional pelvis generation as finding the intersection of pelvis distributions of component atomic interactions. Using our PelvisVAE trained for atomic interactions, we can decode the

prior standard Gaussian distribution to distributions of plausible pelvis frames given atomic interaction semantics. Each component atomic interaction of a composite interaction determines a pelvis frame distribution in scenes, and intuitively the plausible pelvis of the composite interaction should lie in the intersection of all pelvis distributions. Therefore, we propose a compositional pelvis generation method where we sample pelvis frames for each atomic interaction, optimize all pelvis frames to be close in space, and average the optimized pelvis frames as the generated composite pelvis frame.

Specifically, we denote the decoder of the PelvisVAE as a function $\mathcal{G}(z^i|a^i, o^i)$ mapping a latent code z^i to a pelvis frame conditioned on the action-object pair of (a^i, o^i) , we obtain the final pelvis frame from the composition of M atomic interactions from the following optimization problem:

$$\tilde{Z} = \arg \min_{Z=\{z^i\}_{i=1}^M} \sum_{1 \leq i < M} \sum_{i < j \leq M} |\mathcal{G}(z^i|a^i, o^i) - \mathcal{G}(z^j|a^j, o^j)| - \lambda \sum_{1 \leq i \leq M} \log(p(z^i)), \quad (2)$$

Z denotes the set of M randomly sampled latent codes for M component atomic interactions, and λ is a tunable weight. The initial values of the latent codes are randomly sampled from the prior distribution. The motivation is to optimize the latent codes of M component atomic interactions to make their decoded pelvis frames as close as possible in the scene while maintaining high probabilities in each component pelvis distribution. We take the average of the M optimized component pelvis frames as the sampled composite pelvis frame.

Compositional Body Generation. For compositional body generation, we use the attention mechanism of transformer networks to compose atomic interactions at body part level. Although we use BodyVAE trained only on atomic interaction data, we can input more than one atomic interaction and achieve compositional body generation with additional attention masks derived from contact statistics.

Our method is inspired by the observation from contact probability statistics that interaction semantics can often be determined by contact on certain body parts instead of the whole body, e.g., sitting is characterized by contact around the hip area. Specifically, we calculate the per-vertex contact probability for each action using training data. The vertex contact probability for one action is defined as the ratio of the number of positive contact features of the vertex and the number of all contact features of the vertex in the interaction data of the specified action. Heatmap visualization of such categorical contact statistics is shown in Fig. 4. High contact probability indicates a high correlation between the body vertex area and action semantics.

Inspired by the observed action-part correlation, we propose leveraging the attention mechanism to compose atomic interactions according to three body parts: upper limbs, lower limbs, and torso. Given composite interaction semantics for body generation, we calculate the sum of contact probability at each part for each action. If the sum of contact probability of one action at one body part exceeds those of other actions by a threshold, we mask the attention between the



Fig. 4: Visualization of per-vertex contact probability for two actions from the side and bottom views. Brighter colors indicate higher contact probability with objects. We randomly sample a pose within the action category for visualization.

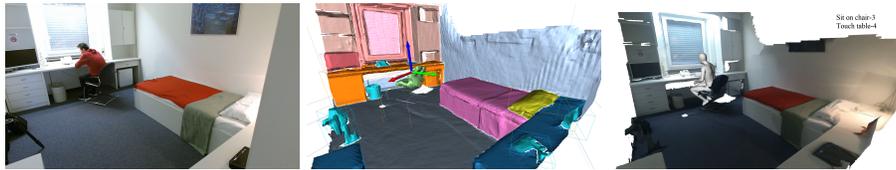


Fig. 5: Illustration of the PROX-S dataset. From left to right, it shows the RGB image, 3D scene reconstruction with instance segmentation, and the fitted SMPL-X human body with per-frame interaction semantic labels.

part and all other atomic interactions to be zero. Taking the composite interaction “sit on chair and touch table” as an example, touching has a significantly greater sum of contact probability on upper limbs so the attention between upper limb tokens and chair (objects not paired with touching) tokens will be masked as zero.

We input such attention masks and component atomic interactions to the transformer for compositional body generation. The derived attention masks can naturally compose the component atomic interactions via BodyVAE.

4 Experiments

4.1 Dataset

We annotate a new dataset called PROX-S on top of the PROX [10] dataset to evaluate interaction synthesis with semantic control since there is no existing dataset suitable for this task. The PROX-S dataset contains (1) 3D instance segmentation of all 12 PROX scenes; (2) 3D human body reconstructions within the scenes; and (3) per-frame interaction semantic labels in the form of action-object pairs, as shown in Fig. 5. We refer to the Supp. Mat. for the detailed process description and statistics.

4.2 Evaluation Metrics

Semantic Contact. Compared to the semantics-agnostic body-scene contact score used in [46], we propose a semantic contact score to evaluate whether the generated interactions adhere to the semantic specifications characterized by proper body-scene contact. Using the action-specific body vertex contact probability introduced in Sec. 3.3 as weights, we calculate a weighted sum of binary contact features for each body in the training data and name the sum as the semantic contact score. This action-dependent semantic contact score reflects how body contact features align with the semantics of one action. We use the 20 percentile semantic contact score of training data bodies as the threshold score for each action to filter low-quality training interaction data with inaccurate contact. When evaluating a generated interaction sample, we calculate the weighted sum of contact features and compare it to the action-specific threshold score. The semantic contact score is one if the sum is higher than the threshold and zero otherwise.

Physical Plausibility. We evaluate the physical plausibility of generated interactions using the non-collision metric from [4, 11, 45, 46]. It is calculated as the ratio of the number of body vertices with non-negative scene SDF values and the number of all body vertices.

Diversity. Following [46], we perform K-Means clustering to cluster the SMPL-X parameters of generated bodies into 50 clusters and evaluate the diversity using the entropy of the cluster ID histogram and the mean distance to cluster centers.

Perceptual Study. We evaluate the perceptual realism and semantic accuracy of generated human-scene interactions by conducting perceptual studies on Amazon Mechanical Turk (AMT). We perform a binary-choice perceptual study where samples of the same interaction semantics from different sources are shown and the turkers are asked to choose the more natural one. For semantic accuracy, we instruct the turkers to rate the consistency between shown interaction samples and semantic labels from 1 (strongly disagree) to 5 (strongly agree). For more details regarding the perceptual studies, please refer to the Supp. Mat.

4.3 Baselines

There is no existing method that directly applies to the problem of generating 3D human bodies given a 3D scene and interaction semantics. Therefore, we modify two state-of-the-art interaction synthesis methods for our tasks and train their models using the same PROX-S dataset.

PiGraph-X. PiGraph [31] is the most related work to our knowledges, which simultaneously generates scene placement and human skeletons from interaction category specifications. We remove the scene placement step of the original paper and replace the skeleton body representation with the more realistic SMPL-X

Table 1: Binary Realism Perceptual Study.

(a) Interaction Synthesis with Semantic Control.

% users rated as "more realistic"	
Ours	POSA-I
63.8 %	36.2 %
Ours	PiGraph-X
79.7 %	20.3 %
Ours	PseudoGT
45.3 %	54.7 %

(b) Interaction Synthesis via Semantic Composition.

% users rated as "more realistic"	
Ours-C	PiGraph-X-C
81.2 %	18.2 %
Ours-C	PseudoGT
42.4 %	57.6 %

body model. We refer to this modified PiGraph variant as PiGraph-X. For compositional interaction generation, we extend PiGraph by sampling basic bodies from all comprised atomic interactions and using the average similarity score to those atomic interactions. We denote this PiGraph extension as PiGraph-X-C.

POSA-I. POSA [11] generates human-scene interactions using optimization guided by per-vertex contact features. However, POSA lacks global control of the generated contact features and cannot generate human bodies conditioned on interaction semantics. To incorporate global semantics of human-scene interaction, we modify POSA to generate human bodies with vertex-level contact features from interaction semantics in the format of action-noun pairs. We denote the semantic-controllable POSA variant as POSA-I.

4.4 Interaction Synthesis with Semantic Control

We evaluate the synthesized interactions in terms of perceptual realism, perceptual semantic accuracy, semantic contact, physical plausibility, and diversity. Table 1 and 2 show the results of the perceptual studies on binary realism and other quantitative metrics, respectively. Our method significantly outperforms the baselines in the binary realism study and even comes close to the body poses reconstructed from the real human-scene interactions, which we consider as the pseudo ground truth. Our method also achieves higher perceptual semantic ratings than all the baselines. Quantitatively, our method achieves the highest semantic contact score (0.93) which significantly outperforms the baselines. Our method performs comparably on physical plausibility as POSA-I and significantly better than PiGraph-X. For diversity metrics, our method has a larger cluster entropy and smaller cluster size compared to the baselines. We observe more frequent unnatural interaction samples from PiGraph-X and POSA-I, which accounts for their larger cluster sizes. Note that our method is a generative pipeline where we evaluate every randomly sampled result, while both baselines need to search thousands of samples of global placement in scenes.

We show the qualitative comparison of interactions generated from ours and two baselines in Fig. 6. PiGraph-X naively models the human body as independent joints and therefore suffers from unnatural body poses, poor human-scene

Table 2: Evaluation of semantics, physical plausibility and diversity.

	Semantic Accuracy	Semantic Contact	Non-Collision	Entropy	Cluster Size	Samples
PiGraph-X	3.90 ± 0.93	0.84	0.85	3.67	1.07	25K
POSA-I	4.03 ± 0.85	0.67	0.98	3.75	1.08	> 1K
Ours	4.17 ± 0.71	0.93	0.97	3.77	0.83	1

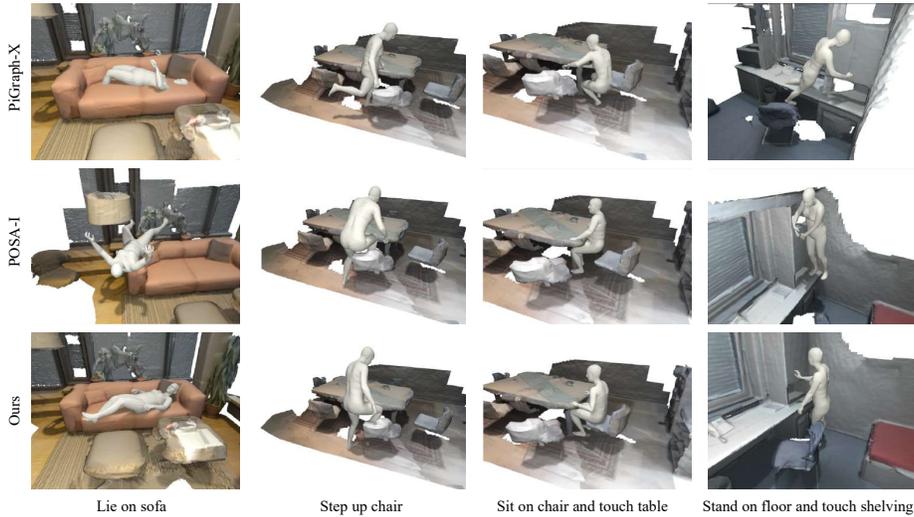


Fig. 6: Qualitative comparison of interactions generated with our method and two baselines. Each row shows interactions generated from one method. From left to right, the columns correspond to different interaction semantics of “lying on the sofa”, “stepping up on the chair”, “sitting on the chair and touching the table”, and “standing on the floor and touching the shelving”.

contact, and penetration. POSA-I generates bodies with physically plausible interactions, but since it needs to find plausible body placement in the scene, its optimization can easily get stuck at local minimums that are not semantically accurate and have inferior human-scene contact. In contrast, our method learns how humans interact with objects given different interaction semantics and generates realistic interactions that are both physically plausible and human-like.

4.5 Compositional Interaction Generation

We conduct the same evaluation for novel composite interaction generation as introduced in Sec. 4.4. We compare our method to the PiGraph-X-C extension for compositional interaction generation. POSA-I is not included in this part since it can not be easily extended for semantic composition. In the forced-alternative-choice (binary-choice) perceptual study (Tab. 1 (b)), our method significantly outperforms the baseline. The semantic accuracy, semantic contact, non-collision and diversity comparisons are shown in Tab. 3. Our method achieves significantly

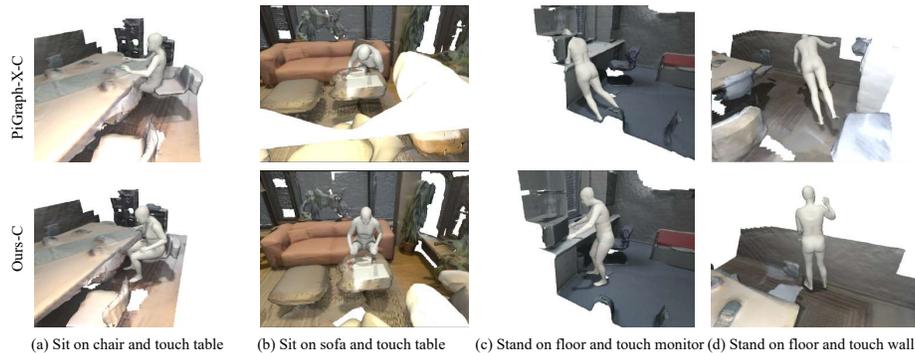


Fig. 7: Qualitative comparison of composite interactions generated using ours-C and PiGraph-X-C. Each row shows interactions generated from one method.

Table 3: Evaluation of semantics, physical plausibility and diversity for novel compositional interaction generation.

	Semantic Accuracy	Semantic Contact	Non-Collision	Entropy	Cluster Size	Samples
PiGraph-X-C	3.88 ± 1.15	0.80	0.84	3.72	1.09	25K
Ours-C	3.95 ± 1.18	0.76	0.98	3.80	0.68	1

better semantic accuracy, physical plausibility, and sample efficiency. The slight inferior semantic contact score is because the combinations of actions and objects in the test scenes may not always afford natural contact, like touching an incomplete wall in Fig. 7 (d). Touching the wall and floor simultaneously generates higher contact scores but perceptually unnatural poses to humans.

We refer to the Supp. Mat. for **more qualitative results, failures, ablation study, and limitations.**

5 Conclusion

In this paper, we introduce a method to generate humans interacting with 3D scenes given interaction semantics as action-object pairs. We propose a compositional interaction representation and transformer-based VAE models to stochastically generate human-scene interactions. We further explored generating composite interactions by composing atomic interactions, which does not demand training on composite interaction data. Our method is a step towards creating virtual avatars interacting with 3D scenes.

Acknowledgements: We sincerely acknowledge the anonymous reviewers for their insightful suggestions. We thank Francis Engelmann for help with scene segmentation and proofreading, and Siwei Zhang for providing body fitting results. This work was supported by the SNF grant 200021 204840.

References

1. Ahn, H., Ha, T., Choi, Y., Yoo, H., Oh, S.: Text2action: Generative adversarial synthesis from language to action. In: Proc. of ICRA. IEEE (2018)
2. Ahuja, C., Morency, L.P.: Language2pose: Natural language grounded pose forecasting. In: Proc. of 3DV. IEEE (2019)
3. Akhter, I., Black, M.J.: Pose-conditioned joint angle limits for 3d human pose reconstruction. In: Proc. of CVPR (2015)
4. Engelmann, F., Rematas, K., Leibe, B., Ferrari, V.: From Points to Multi-Object 3D Reconstruction. In: Proc. of CVPR (2021)
5. Gkioxari, G., Girshick, R., Dollár, P., He, K.: Detecting and recognizing human-object interactions. In: Proc. of CVPR (2018)
6. Grabner, H., Gall, J., Van Gool, L.: What makes a chair a chair? In: Proc. of CVPR. IEEE (2011)
7. Gupta, A., Kembhavi, A., Davis, L.S.: Observing human-object interactions: Using spatial and functional compatibility for recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **31**, 1775–1789 (2009)
8. Gupta, A., Satkin, S., Efros, A.A., Hebert, M.: From 3d scene geometry to human workspace. In: Proc. of CVPR. IEEE (2011)
9. Hassan, M., Ceylan, D., Villegas, R., Saito, J., Yang, J., Zhou, Y., Black, M.J.: Stochastic scene-aware motion prediction. In: Proc. of ICCV (2021)
10. Hassan, M., Choutas, V., Tzionas, D., Black, M.J.: Resolving 3D human pose ambiguities with 3D scene constraints. In: Proc. of ICCV (2019)
11. Hassan, M., Ghosh, P., Tesch, J., Tzionas, D., Black, M.J.: Populating 3D scenes by learning human-scene interaction. In: Proc. of CVPR (2021)
12. Hu, R., Yan, Z., Zhang, J., Van Kaick, O., Shamir, A., Zhang, H., Huang, H.: Predictive and generative neural networks for object functionality. arXiv preprint arXiv:2006.15520 (2020)
13. Ji, S., Xu, W., Yang, M., Yu, K.: 3d convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence* **35**(1), 221–231 (2012)
14. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017)
15. Kim, V.G., Chaudhuri, S., Guibas, L.J., Funkhouser, T.: Shape2pose: human-centric shape analysis. In: *ACM Transactions on Graphics, (Proc. SIGGRAPH)* (2014)
16. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
17. Li, X., Liu, S., Kim, K., Wang, X., Yang, M.H., Kautz, J.: Putting humans in a scene: Learning affordance in 3d indoor environments. In: Proc. of CVPR (2019)
18. Lieber, R., Stekauer, P.: *The oxford handbook of compounding* (2011)
19. Mahmood, N., Ghorbani, N., Troje, N.F., Pons-Moll, G., Black, M.J.: AMASS: Archive of motion capture as surface shapes. In: Proc. of ICCV (2019)
20. Mandery, C., Terlemez, O., Do, M., Vahrenkamp, N., Asfour, T.: The kit whole-body human motion database. In: Proc. of ICAR (2015)
21. Mineshima, K., Martínez-Gómez, P., Miyao, Y., Bekki, D.: Higher-order logical inference with compositional semantics. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing* (2015)

22. Mitchell, J., Lapata, M.: Vector-based models of semantic composition. In: Proc. of ACL (2008)
23. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: Proc. of CVPR (2019)
24. Petrovich, M., Black, M.J., Varol, G.: Action-conditioned 3d human motion synthesis with transformer vae. In: Proc. of ICCV (2021)
25. Plag, I.: Word-formation in English. Cambridge University Press (2018)
26. Punnakkal, A.R., Chandrasekaran, A., Athanasiou, N., Quiros-Ramirez, A., Black, M.J.: BABEL: Bodies, action and behavior with english labels. In: Proc. of CVPR (2021)
27. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. arXiv preprint arXiv:1706.02413 (2017)
28. Ranjan, A., Bolkart, T., Sanyal, S., Black, M.J.: Generating 3D faces using convolutional mesh autoencoders. In: Proc. of ECCV (2018)
29. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning internal representations by error propagation. Tech. rep., California Univ San Diego La Jolla Inst for Cognitive Science (1985)
30. Savva, M., Chang, A.X., Hanrahan, P., Fisher, M., Nießner, M.: Scenegrok: Inferring action maps in 3d environments. ACM transactions on graphics (TOG), (Proc. SIGGRAPH) **33**(6), 1–10 (2014)
31. Savva, M., Chang, A.X., Hanrahan, P., Fisher, M., Nießner, M.: PiGraphs: Learning Interaction Snapshots from Observations. ACM Transactions on Graphics, (Proc. SIGGRAPH) **35**(4) (2016)
32. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. Proc. of NeurIPS (2014)
33. Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. Proc. of NeurIPS (2015)
34. Starke, S., Zhang, H., Komura, T., Saito, J.: Neural state machine for character-scene interactions. In: ACM Trans. Graph.(ACM SIGGRAPH Asia) (2019)
35. Tevet, G., Gordon, B., Hertz, A., Bermano, A.H., Cohen-Or, D.: Motionclip: Exposing human motion generation to clip space. arXiv preprint arXiv:2203.08063 (2022)
36. De la Torre, F., Hodgins, J., Bargteil, A., Martin, X., Macey, J., Collado, A., Beltran, P.: Guide to the carnegie mellon university multimodal activity (cmu-mmac) database (2009)
37. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: Proc. of CVPR (2018)
38. Troje, N.F.: Decomposing biological motion: A framework for analysis and synthesis of human gait patterns. Journal of Vision **2**, 371–387 (2002)
39. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Proc. of NeurIPS (2017)
40. Wang, J., Xu, H., Xu, J., Liu, S., Wang, X.: Synthesizing long-term 3d human motion and interaction in 3d scenes. In: Proc. of CVPR (2021)
41. Wang, J., Rong, Y., Liu, J., Yan, S., Lin, D., Dai, B.: Towards diverse and natural scene-aware 3d human motion synthesis. In: Proc. of CVPR (2022)
42. Yao, B., Fei-Fei, L.: Modeling mutual context of object and human pose in human-object interaction activities. In: Proc. of CVPR (2010)
43. Yin, D., Meng, T., Chang, K.W.: Sentibert: A transferable transformer-based architecture for compositional sentiment semantics. arXiv preprint arXiv:2005.04114 (2020)

44. Zhang, S., Zhang, Y., Bogo, F., Marc, P., Tang, S.: Learning motion priors for 4d human body capture in 3d scenes. In: Proc. of ICCV (2021)
45. Zhang, S., Zhang, Y., Ma, Q., Black, M.J., Tang, S.: PLACE: Proximity learning of articulation and contact in 3D environments. In: Proc. of 3DV (2020)
46. Zhang, Y., Hassan, M., Neumann, H., Black, M.J., Tang, S.: Generating 3d people in scenes without people. In: Proc. of CVPR (2020)