

Generating Human Interaction Motions in Scenes with Text Control

Hongwei Yi^{1,2}, Justus Thies^{2,3}, Michael J. Black²,
Xue Bin Peng^{1,4}, and Davis Rempe¹

¹ NVIDIA

² Max Planck Institute for Intelligent Systems, Tübingen, Germany

³ Technical University of Darmstadt

⁴ Simon Fraser University




Fig. 1: We present TeSMo, a method for generating diverse and plausible human-scene interactions from text input. Given a 3D scene, TeSMo generates scene-aware motions, such as walking in free space and sitting on a chair. Our model can be easily controlled using textual descriptions, start positions, and goal positions.

Abstract. We present TeSMo, a method for text-controlled scene-aware motion generation based on denoising diffusion models. Previous text-to-motion methods focus on characters in isolation without considering scenes due to the limited availability of datasets that include motion, text descriptions, and interactive scenes. Our approach begins with pre-training a scene-agnostic text-to-motion diffusion model, emphasizing goal-reaching constraints on large-scale motion-capture datasets. We then enhance this model with a scene-aware component, fine-tuned using data augmented with detailed scene information, including ground plane and object shapes. To facilitate training, we embed annotated navigation and interaction motions within scenes. The proposed method produces realistic and diverse human-object interactions, such as navigation and sitting, in different scenes with various object shapes, orientations, initial body positions, and poses. Extensive experiments demonstrate that our approach surpasses prior techniques in terms of the plausibility of

human-scene interactions, as well as the realism and variety of the generated motions. Code will be released upon publication of this work at <https://research.nvidia.com/labs/toronto-ai/tesmo>.

Keywords: Scene-Aware Human Motion Generation · Text-to-Motion

1 Introduction

Generating realistic human movements that can interact with 3D scenes is crucial for many applications, ranging from gaming to embodied AI. For example, character animators for games and films need to author motions that successfully navigate through cluttered scenes and realistically interact with target objects, while still maintaining artistic control over the style of the movement. One natural way to control style is through text, e.g., “skip happily to the chair and sit down”. Recently, diffusion models have shown remarkable capabilities in generating human motion from user inputs. Text prompts [37, 48] let users control style, while methods incorporating spatial constraints enable more fine-grained control, such as specifying desired joint positions and trajectories [19, 34, 42]. However, these works have predominantly focused on characters in isolation, without considering environmental context or object interactions.

In this work, we aim to incorporate scene-awareness into user-controllable human motion generation models. However, learning to generate motions involving scene interactions is challenging, even without text prompts. Unlike large-scale motion capture datasets that depict humans in isolation [27], datasets with paired examples of 3D human motion and scene/object geometry are limited. Prior work uses small paired datasets without text annotations to train VAEs [10, 35, 50] or diffusion models [17, 31] that generate human scene interactions with limited scope and diversity. Reinforcement learning methods are able to learn interaction motions from limited supervision [12, 23, 53], and can generate behaviors that are not present in the training motion dataset. However, designing reward functions that lead to natural movements for a diverse range of interactions is difficult and tedious.

To address these challenges, we introduce a method for Text-conditioned Scene-aware Motion generation, called TeSMo. As shown in Fig. 1, our method generates realistic motions that navigate around obstacles and interact with objects, while being conditioned on a text prompt to enable stylistic diversity. Our key idea is to combine the power of general, but scene-agnostic, text-to-motion diffusion models with paired human-scene data that captures realistic interactions. First, we pre-train a text-conditioned diffusion model [37] on a diverse motion dataset with no objects (e.g., HumanML3D [8]), allowing it to learn a realistic motion prior and the correlation with text. We then fine-tune the model with an augmented scene-aware component that takes scene information as input, thereby refining motion outputs to be consistent with the environment.

Given a target object with which to interact and a text prompt describing the desired motion, we decompose the problem of generating a suitable motion in a

scene into two components, *navigation* (e.g., approaching a chair while avoiding obstacles) and *interaction* (e.g., sitting on the chair). Both stages leverage diffusion models that are pre-trained on scene-agnostic data, then fine-tuned with an added scene-aware branch. The *navigation* model generates a pelvis trajectory that reaches a goal pose near the interaction object. During fine-tuning, the scene-aware branch takes, as input, a top-down 2D floor map of the scene and is trained on our new dataset containing locomotion sequences [27] in 3D indoor rooms [7]. The generated pelvis trajectory is then lifted to a full-body motion using motion in-painting [34]. Next, the *interaction* model generates a full-body motion conditioned on a goal pelvis pose and a detailed 3D representation of the target object. To further improve generalization to novel objects, the model is fine-tuned using augmented data that re-targets interactions [10] to a variety of object shapes while maintaining realistic human-object contacts.

Experiments demonstrate that our navigation approach outperforms prior work in terms of goal reaching and obstacle avoidance, while producing full-body motions on par with scene-agnostic diffusion models [19, 42]. Meanwhile, our interaction model generates motions with fewer object penetrations than the state-of-the-art approach [53], being preferred 71.9% of the time in a perceptual study. The central contribution of this work includes: (1) a novel approach to enable scene-aware and text-conditioned motion generation by fine-tuning an augmented model on top of a pre-trained text-to-motion diffusion model, (2) a method, TeSMo, that leverages this approach for navigation and interaction components to generate high-quality motions in a scene from text, (3) data augmentation strategies for placing navigation and interaction motions with text annotations realistically in scenes to enable scene-aware fine-tuning.

2 Related Work

2.1 Scene-aware Motion Generation

Motion synthesis in computer graphics has a rich history, encompassing areas such as locomotion [1, 20, 24, 51], human-scene/object interaction [22, 36], and dynamic object interaction [4, 25, 26]. We refer readers to an extensive survey [54] for an overview and focus on scene-aware motion generation in this section.

A particular challenge in modeling scene-aware motion is the lack of paired, high-quality human-scene datasets. One line of work [38, 39] employs a two-stage method that first predicts the root path, followed by the full-body motion based on the scene and predicted path. However, these methods suffer from low-quality motion generation, attributed to the noise in the training datasets captured from monocular RGB-D videos [11]. Neural State Machine (NSM) [35] proposes the use of phase labeling [16] and local expert networks [6, 18, 46] to generate high-quality object interactions, such as sitting and carrying, after training on a small human-object mocap dataset. Nonetheless, it struggles with recognizing walkable regions in 3D scenes, often failing to avoid obstacles. Therefore, later work in this vein requires using the A* algorithm for collision-free path planning [10].

These and related approaches [49, 50] are moreover limited by the diversity of the small human-scene interaction datasets with no text annotations.

Various approaches ameliorate the data issue by creating synthetic data with captured [43, 44] or generated [21] motions placed in scenes heuristically. HUMANISE [40] does this for text-conditioned scene interactions, but rely entirely on short synthetic sequences for training, where the realism is limited by the data generation heuristics used. The reinforcement learning (RL) approach DIMOS [53] learns autoregressive policies to reach goal poses in a scene without requiring paired human-scene data for training, but still relies on A* and is constrained by the accuracy of goal pose generation [52]. RL with physical simulators [3, 12, 29, 41] has been used to produce physically plausible movements but faces challenges in generalizing across varied scenes and objects.

Unlike most prior works, our approach is text-conditioned and leverages a mix of both scene-agnostic and paired human-scene data. Pre-training is done with a diverse scene-agnostic dataset, while scene-aware fine-tuning uses motion data with scene context. For training, we adopt both synthetic data creation with real motions and data augmentation of real-world human-object interactions [10].

2.2 Diffusion-Based Motion Generation

Recently, diffusion models have demonstrated the ability to generate high-quality human motions, especially when conditioned on a text prompt [30, 37, 48]. In addition to text, several diffusion models add spatial controllability. Some works [34, 37] adopt image inpainting techniques to incorporate dense trajectories of spatial joint constraints into generated motions. OmniControl [42] and GMD [19] allow control with sparse signals and a pre-defined root path, respectively.

A few diffusion works handle interactions with objects or scenes. TRACE [33] generates 2D trajectories for pedestrians based on a rasterized street map. SceneDiffuser [17] conditions generation on a full scanned scene point cloud, but motion quality is limited due to noisy training data [11]. Another approach [31] tackles single-object interactions through hierarchical generation of milestone poses followed by dense motion, but it lacks text control. A concurrent line of work enables text conditioning for single-object interactions [5, 28], but they focus on humans manipulating dynamic objects rather than interactions in full scenes.

We leverage a pre-trained text-to-motion diffusion model [37] and a fine-tuned scene-aware branch to enable both text controllability and scene-awareness with diffusion. We break motion generation into navigation and interaction with static objects by conditioning on 2D floor maps and 3D geometry, respectively, and create specialized human-scene data to enable diversity and quality.

3 Text-Conditioned Scene-Aware Motion Generation

3.1 Overview

Given a 3D scene and a target interaction object, our goal is to generate a plausible human-scene interaction, where the motion style can be controlled by a




Fig. 2: Pipeline overview: given the start position (green arrow), goal position (red arrow), 3D scene, and text description, the navigation root trajectory is first generated and then the full-body motion is completed through in-painting. Subsequently, the interaction is generated from a start pose (i.e., the end pose from navigation), goal position, and the target object, enabling the generation of object-specific motion.

user-specified text prompt. Our approach decomposes this task into two components, *navigation* and *interaction*, as illustrated in Fig. 2. Both components are diffusion models that leverage a fine-tuning routine to enable scene-awareness without losing user controllability, as introduced in Sec. 3.2. To interact with an object, the character must first navigate to a location in the scene near the object, which is easily calculated heuristically or specified by the user, if desired. As described in Sec. 3.3, we design a hierarchical *navigation* model, which generates a root trajectory starting from an initial location that moves to the goal location while navigating around obstacles in the scene. The generated root trajectory is then lifted into a full-body motion using in-painting techniques [34, 42]. Since the navigation model gets close to the object in the first stage, for generating the actual object *interaction*, we can focus on scenarios where the character is already near the object. This allows a one-stage motion generation model that directly predicts the full-body motion from the starting pose (i.e., the last pose of navigation), a goal pelvis pose, and the object (as detailed in Section 3.4).

3.2 Background: Controllable Human Motion Diffusion Models

Motion diffusion models. Diffusion models have been successfully used to generate both top-down trajectories [33] and full-body motions [37, 48]. These models generate motions by iteratively denoising a temporal sequence of N poses (e.g., root positions or full-body joint positions/angles) $\mathbf{x} = [\mathbf{x}^1, \dots, \mathbf{x}^N]$. During training, the model learns to reverse a forward diffusion process, which starts from a clean motion $\mathbf{x}_0 \sim q(\mathbf{x}_0)$, sampled from the training data, and after T diffusion steps is approximately Gaussian $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Then at each step t of motion denoising, the reverse process is defined as:

$$p_\phi(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{c}) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\phi(\mathbf{x}_t, \mathbf{c}, t), \beta_t \mathbf{I}) \quad (1)$$

where \mathbf{c} is some conditioning signal (e.g., a text prompt), and β_t depends on a pre-defined variance schedule. The denoising model $\boldsymbol{\mu}_\phi$ with parameters ϕ predicts the denoised motion $\hat{\mathbf{x}}_0$ from a noisy input motion \mathbf{x}_t [14]. The model is trained by sampling a motion \mathbf{x}_0 from the dataset, adding random noise, and supervising the denoiser with a reconstruction loss $\|\mathbf{x}_0 - \hat{\mathbf{x}}_0\|^2$.

Augmented controllability. In the image domain, general pre-trained diffusion models are specialized for new tasks using an augmented ControlNet [47] branch, which takes in a new conditioning signal and is fine-tuned on top of the frozen base diffusion model. OmniControl [42] adopts this idea to the human motion domain. For motion diffusion models with a transformer encoder architecture, they propose an augmented transformer branch that takes in kinematic joint constraints (e.g., pelvis or other joint positions) and at each layer connects back to the base model through a linear layer that is initialized to all zeros.

As described in Secs. 3.3 and 3.4, our key insight is to use an augmented control branch to enable scene awareness. We first train a strong scene-agnostic motion diffusion model to generate realistic motion from a text prompt, and then fine-tune an augmented branch that takes scene information as input (e.g., a 2D floor map or 3D geometry). This new branch adapts generated motion to be scene-compliant, while still maintaining realism and text controllability.

Test-time guidance. At test time, diffusion models can be controlled to meet specific objectives through guidance. We directly apply guidance to the clean motion prediction from the model $\hat{\mathbf{x}}_0$ [15, 33]. At each denoising step, the predicted $\hat{\mathbf{x}}_0$ is perturbed with the gradient of an analytic objective function \mathcal{J} as $\tilde{\mathbf{x}}_0 = \hat{\mathbf{x}}_0 - \alpha \nabla_{\mathbf{x}_t} \mathcal{J}(\hat{\mathbf{x}}_0)$ where α controls the strength of the guidance and \mathbf{x}_t is the noisy input motion at step t . The predicted mean $\boldsymbol{\mu}_\phi$ is then calculated with the updated motion prediction $\tilde{\mathbf{x}}_0$ as in [15, 33]. As detailed later, we define guidance objectives for avoiding collisions and reaching goals.

3.3 Navigation Motion Generation

The goal of the navigation stage is for the character to reach a goal location near the target object using realistic locomotion behaviors that can be controlled by the user via text. We design a hierarchical method that first generates a dense root trajectory with a diffusion model, then leverages a powerful in-painting model [34] to generate a full-body motion for the predicted trajectory. This approach facilitates accurate goal-reaching with the root-only model, while allowing diverse text control through the in-painting model.

Root trajectory generation. Our root trajectory diffusion model, shown in Fig. 3(a), operates on motions where each pose is specified by $\mathbf{x}^n = [x, y, z, \cos \theta, \sin \theta]_n$, with (x, y, z) being the pelvis position and θ the pelvis rotation, both of which are represented in the coordinate frame of the *first* pose in the sequence. The model is conditioned on a text prompt along with starting and ending goal positions and orientations. In contrast to the representation from prior work [8], which uses relative pelvis velocity and rotation, our representation using absolute coordinates facilitates constraining the outputs of the model with goal poses.

Inspired by motion in-painting models [34, 37], given a start pose \mathbf{s} and end goal pose \mathbf{g} , at each denoising step, we mask out the input \mathbf{x}_t such that $\mathbf{x}_t^1 = \mathbf{s}$ and $\mathbf{x}_t^N = \mathbf{g}$, thereby providing clean goal poses directly to the model. To achieve this, a binary mask $\mathbf{m} = [\mathbf{m}^1, \dots, \mathbf{m}^N]$ with the same dimensionality as \mathbf{x}_t is defined, where \mathbf{m}^1 and \mathbf{m}^N are a vector of 1's and all other \mathbf{m}^n are 0's. During

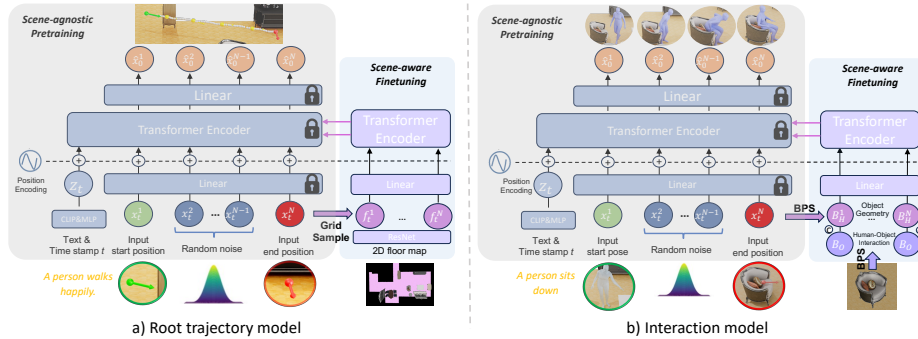


Fig. 3: Network architecture of the (a) root trajectory model and (b) interaction motion model. Initially, the base transformer encoder is trained on scene-agnostic motion data using start pose, target pose, and text as input. Subsequently, a scene-aware component is fine-tuned, which incorporates the 2D floor map (a) or 3D object (b).

training, overwriting occurs with $\tilde{\mathbf{x}}_t = \mathbf{m} * \mathbf{x}_0 + (\mathbf{1} - \mathbf{m}) * \mathbf{x}_t$ where $*$ indicates element-wise multiplication and \mathbf{x}_0 is a ground truth root trajectory. We then concatenate the mask with the overwritten motion $[\tilde{\mathbf{x}}_t; \mathbf{m}]$ and use this as input to the model to indicate which frames have been overwritten.

At test time, goal-reaching is improved using a guidance objective $\mathcal{J}_g = (\hat{\mathbf{x}}_0^N - \mathbf{g})^2$ that measures the error between the end pelvis position and orientation of the predicted clean trajectory $\hat{\mathbf{x}}_0^N$ and the final goal pose.

Incorporating scene representation. The model as described so far is trained on a locomotion subset of the HumanML3D dataset [8] to enable generating realistic, text-conditioned root trajectories. However, it will be entirely unaware of the given 3D scene. To take the scene into account and avoid degenerating the text-following and goal-reaching performance, we augment the base diffusion model with a control branch that takes a representation of the scene as input. This scene-aware branch is a separate transformer encoder that is fine-tuned on top of the frozen base model. As input, we extract the walkable regions from the 3D geometry of the scene and project them to a bird’s-eye view, yielding a 2D floor map \mathcal{M} . Following [33], a Resnet-18 [13] encodes the map \mathcal{M} as feature grid, and at denoising step t , each 2D projected pelvis position $(x, z) \in \mathbf{x}_t^n$ is queried in the feature grid \mathcal{M} to get the corresponding feature \mathbf{f}_t^n . The resulting sequence of features $\mathbf{f}_t = [\mathbf{f}_t^1, \dots, \mathbf{f}_t^N]$, along with the text prompt and noisy motion \mathbf{x}_t , become the input to the separated transformer branch.

At test time, a collision guidance objective further encourages scene compliance. This is defined as $\mathcal{J}_c = \text{SDF}(\hat{\mathbf{x}}_0, \mathcal{M})$ where SDF calculates the 2D transform distance map from the 2D floor map, then queries the 2D distance value at each time step of the root trajectory. Positive distances, indicating pelvis positions outside the walkable region, are averaged to get the final loss.

Scene-aware training and data. To train the scene-aware branch, it is important to have a dataset featuring realistic motions navigating through scenes with corresponding text prompts. For this purpose, we create the **Loco-3D-FRONT** dataset by integrating locomotion sequences from HumanML3D into

diverse 3D environments from 3D-FRONT [7]. Each motion is placed within a different scene with randomized initial translation and orientation, following the methodology outlined in [44], as depicted in Fig. 4(a). Additionally, we apply left-right mirroring to both the motion and its interactive 3D scenes to augment the dataset [8]. This results in a dataset of approximately 9,500 walking motions, each motion accompanied by textual descriptions and 10 plausible 3D scenes on average, resulting in 95k locomotion-scene training pairs.

Added control with trajectory blending. Our root trajectory diffusion model generates scene-aware motions and, unlike many prior works [10,53], does not require a navigation mesh to compute A* [9] paths to follow. However, a user may want a character to take the shortest path to an object by following the A* path, or to control the general shape of the path by drawing a 2D route themselves. To enable this, we propose to fuse an input 2D trajectory $\mathbf{p} \in \mathbb{R}^{N \times 2}$ with our model’s predicted clean trajectory at every denoising step. At step t , we extract the 2D (x, z) components $\hat{\mathbf{p}}_0$ from the predicted root trajectory $\hat{\mathbf{x}}_0$ and interpolate them with the input trajectory $\tilde{\mathbf{p}}_0 = s * \hat{\mathbf{p}}_0 + (1 - s) * \mathbf{p}$ where s is the blending scale that controls how closely the generated trajectory matches the input. We then overwrite the 2D components of $\hat{\mathbf{x}}_0$ with $\tilde{\mathbf{p}}_0$ and continue denoising. This blending procedure ensures outputs roughly follow the desired path but still maintain realism inherent to the trained diffusion model.

Lifting to full-body poses. To lift the generated pelvis trajectory to a full-body motion, we leverage the existing text-to-motion in-painting method PriorMDM [34], which takes a dense 2D root trajectory as input. By using this strong model that is pre-trained for text-to-motion, we can effectively generate natural and scene-aware full-body motion, while offering diverse stylistic control through text.

3.4 Object-Driven Interaction Motion Generation

After navigation, the character has reached a location near the target object and next should execute a desired interaction motion. Due to the fine-grained relationship between the body and object geometry during interactions, we propose a single diffusion model to directly generate full-body motion, unlike the two-stage navigation approach from Sec. 3.3.

Interaction motion generation. The interaction motion model operates on a sequence of full-body poses and is shown in Fig. 3(b). Our pose representation extends that of HumanML3D [8] to add the absolute pelvis position and heading $(x, y, z, \cos \theta, \sin \theta)$, similar to our navigation model. Each pose in the motion is $\mathbf{x}^n = [x, y, z, \sin \theta, \cos \theta, \dot{r}^a, \dot{r}^x, \dot{r}^z, r^y, \mathbf{j}^p, \mathbf{j}^v, \mathbf{j}^r, \mathbf{c}^f]_n \in \mathbb{R}^{268}$ with \dot{r}^a root angular velocity, (\dot{r}^x, \dot{r}^z) root linear velocity, r^y root height, \mathbf{c}^f foot contacts, and $\mathbf{j}^p, \mathbf{j}^v, \mathbf{j}^r$ the local joint positions, velocities, and rotations, respectively.

The model is conditioned on a text prompt along with a starting full-body pose (*i.e.*, the final pose of the navigation stage) and a final goal pelvis position and orientation. The goal pelvis pose can usually be computed heuristically, but may also be provided by the user or predicted by another network [10]. The same masking procedure described in Sec. 3.3 is used to pass the start and end goals

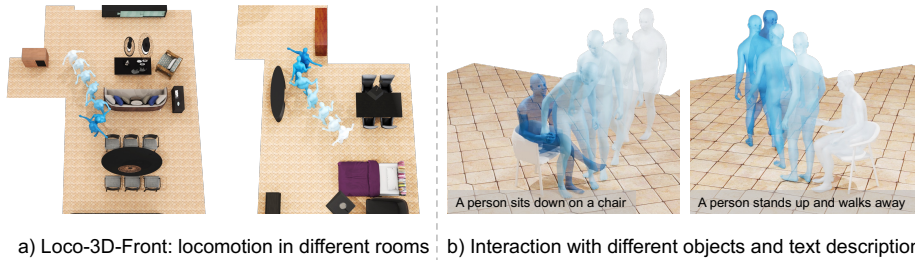


Fig. 4: (a) Loco-3D-FRONT contains locomotion placed in 3D-FRONT [7] scenes without collisions. (b) We augment SAMP [10] by randomly selecting chairs from 3D-FRONT to match the motions and annotating a text description for each sub-sequence.

as input to the model. At test time, we also use the same goal-reaching guidance to improve the accuracy of hitting the final pelvis pose.

Object representation. The base interaction diffusion model is first trained on a dataset of interaction motions from HumanML3D and SAMP [10] without any objects, which helps develop a strong prior on interaction movements driven by text prompts. Similar to navigation, we then augment the base model with a new object-aware transformer encoder and fine-tune this encoder separately.

For the input to this branch at each denoising step t , we leverage Basis Point Sets (BPS) [32] to calculate two key features: object geometry and the human-object relationship. First, a sphere with a radius of 1.0m is defined around the object’s center, and 1024 points are randomly sampled inside this sphere to form the BPS. The distance between each point in the BPS and the object’s surface is then calculated, capturing the object’s geometric features and stored as $\mathbf{B}_O \in \mathbb{R}^{1024}$. Next, for each body pose \mathbf{x}_t^n at timestep n in the noisy input sequence, we calculate the minimum distance from each BPS point to any body joint, giving $\mathbf{B}^n \in \mathbb{R}^{1024}$. The resulting sequence of features $\mathbf{B}_H = [\mathbf{B}^1, \dots, \mathbf{B}^N]$ represents the human-object relationship throughout the entire motion. Finally, the object and human-object interaction features are concatenated with the original pose representation at each timestep $[\mathbf{x}_t^n; \mathbf{B}^n; \mathbf{B}_O]$ and fed to an MLP to generate a merged representation, which serves as the input to the scene-aware branch.

At test time, a collision objective is used to discourage penetrations between human and object. This is very similar to the collision loss described in Sec. 3.3, but the SDF volume is computed for the 3D object and body vertices that are inside the object are penalized. Please see the supplementary material for details.

Scene-aware training and data. To train the scene-aware branch, we utilize the SAMP dataset [10], which captures motions and objects simultaneously. Specifically, we focus on “sitting” and “stand-up” interactions extracted from 80 sitting motion sequences in the SAMP dataset involving chairs of varying heights, as shown in Fig. 4(b). To diversify the object geometry, we randomly select objects from 3D-FRONT [7] to match the contact vertices on human poses in the original SAMP motion sequences. This matching is achieved using the contact loss and collision loss techniques outlined in MOVER [45].

The original SAMP motions are often lengthy (~ 100 sec) and lack paired textual descriptions. For instance, a “sit” motion sequence involves walking to an object, sitting down, standing up, and moving away. To effectively learn individual skills, we extract sub-sequences containing specific interactions that begin or end with a sitting pose, such as “walk then sit”, “stand up then sit”, “stand up from sitting”, and “walk from sitting.” Furthermore, we annotate textual descriptions for each sub-sequence, which often incorporate the style of sitting poses, such as “a person walks and sits down on a chair while crossing their arms.” Applying left-right data augmentation to motion and objects results in approximately 200 sub-sequences for each motion sequence, each paired with corresponding text descriptions and featuring various objects.

4 Experimental Evaluation

4.1 Implementation Details

Training. The scene-agnostic branch of our navigation model is trained on the 3D motions and text descriptions from the Loco-3D-FRONT dataset for 420k optimization steps. Subsequently, we freeze the base model weights and fine-tune the scene-aware branch, with additional 2D floor map inputs, for a further 20k steps. Similarly, the scene-agnostic base of our interaction model first trains on a mix of HumanML3D [8] and SAMP [10] data without objects for 400k steps. Then, the object-aware branch is fine-tuned on our text-annotated SAMP data with 3D object inputs for an additional 20k steps.

Test-time guidance. For the navigation model, we set the guidance weight α to 30 for goal-reaching guidance and 1000 for collision guidance. In the interaction model, we utilize weights of 1000 for goal-reaching loss and 10 for the collision SDF loss. To ensure smooth generation results, we exclude the inference guidance at the final time step of denoising. For a fair comparison with baselines, we do *not* use inference guidance unless explicitly stated in the experiment.

4.2 Evaluation Data and Metrics

Navigation. Navigation performance is assessed using the test set of Loco-3D-FRONT, comprising roughly 1000 sequences. Our metrics evaluate the generated root trajectory and the full-body motion after in-painting separately. For the root trajectory, we measure goal-reaching accuracy for the 2D (horizontal xz) root **position** (m), **orientation** (rad), and root **height** (m). The **collision ratio**, the fraction of frames within generated trajectories where a collision occurs, evaluates the consistency of root motions with the environment. For the full-body motion after in-painting, we use common metrics from prior work [8]. **FID** measures the realism of the motion, **R-precision** (top-3) evaluates consistency between the text and motion, and **diversity** is computed based on the average pairwise distance between sampled motions. Additionally, the **foot skating ratio** [19] evaluates the physical plausibility of motion-ground interaction by the proportion of frames where either foot slides a distance greater than a specified threshold (2.5 cm) while in contact with the ground (foot height < 5 cm).

Table 1: Evaluation of navigation motion generation on the Loco-3D-FRONT test set. **(Left)** For generated pelvis trajectories, our approach achieves the best goal-reaching accuracy with low collision rate. **(Right)** After in-painting the full-body motion, our method maintains diverse and realistic motion that aligns with the given text prompt, competitive with diffusion-based scene-agnostic GMD and OmniControl.

Method	Root trajectory evaluation				Full-body motion evaluation			
	Goal-reaching error ↓				FID ↓	R-precision ↑	Diversity ↑	Foot skating ↓
	Pos.	Orient.	Height	Collision ↓				
Ground Truth	-	-	-	-	0.010	0.672	7.553	0.000
GMD [19]	0.374	1.231	-	-	13.160	0.114	4.488	0.181
OmniControl [42]	1.226	1.018	1.159	-	22.930	0.458	7.128	0.094
TRACE [33]	0.205	0.152	0.010	0.055	22.669	0.144	6.501	0.058
Ours (1-stage train)	0.197	0.132	0.013	0.028	22.372	0.152	6.347	0.062
Ours	0.169	0.119	0.008	0.031	20.465	0.376	6.415	0.056

Interactions. To evaluate full-body human-object interactions, we use the established test split of the SAMP dataset [10], which contains motions related to sitting. Same as navigation, we analyze goal-reaching accuracy through position, orientation, and height errors. Furthermore, we assess physical plausibility by computing average **penetration values** and **penetration ratios** between the generated motion and interaction objects. The penetration value is the mean SDF value across all interpenetrated body vertices of the generated motions, while the ratio is the fraction of generated poses containing penetrations (*i.e.*, SDF values < -3 cm) over all generated motion frames. We also perform a **user study** to compare methods. We employ Amazon Mechanical Turk (AMT) [2] to solicit assessments from 30 individuals. Raters are presented with two side-by-side videos of generated interactions and asked which is more realistic. Please see the supplementary material for more details.

4.3 Comparisons

Navigation. We conduct a comparative analysis of our method with previous scene-aware and scene-agnostic motion generation approaches, shown in Tab. 1. Every method is conditioned on a text prompt along with a start and end goal pose, as described in Sec. 3.3. The TRACE baseline and our method TeSMo also receive the 2D floor map as input.

We first compare to GMD [19] and OmniControl [42], previous scene-agnostic text-to-motion diffusion models trained on HumanML3D to follow a diverse range of kinematic motion constraints. GMD utilizes the horizontal pelvis positions (x, z) of both the start and end goals to generate a dense root trajectory and subsequently the full-body motion. OmniControl takes as input the horizontal pelvis positions (x, z) along with the height y to directly generate full-body motion in a single stage. Our navigation model achieves better goal-reaching accuracy, e.g., 16.9 cm for root position, since it is trained specifically for the goal-reaching locomotion task. More importantly, in the right half of Tab. 1 the full-body motion from our method after in-painting is comparable in terms of realism, text-following, and diversity, while achieving the best foot skating re-




Fig. 5: Navigation generation performance. The start pose is the green arrow, and the goal pose is the red arrow. Our method more accurately reaches the goal and avoids obstacles while style is controlled by a text prompt.

Table 2: Evaluation of human-object interaction motion generation on SAMP [10] sitting test set. Compared to DIMOS, our approach reaches the goal pose more accurately and exhibits fewer object penetrations, resulting in higher human preference.

Method	Goal-reaching error ↓			Object penetration ↓		User study preference ↑
	Pos.	Height	Orient.	Value	Ratio	
DIMOS [53]	0.2020	0.1283	0.4731	0.0193	0.1076	29.1%
Ours	0.1445	0.0120	0.2410	0.0043	0.0611	71.9%

sults. This demonstrates that our approach adds scene-awareness to locomotion generation, without compromising realism or text control.

To justify our two-branch model architecture, we adapt TRACE [33], a recent root trajectory generation model designed to take a 2D map of the environment as input. The adapted TRACE architecture is very similar to our model in Fig. 3(a), but instead of using a separate scene-aware branch, the base transformer directly takes the encoded 2D floor map features as input. This results in a single-branch architecture that must be trained from scratch, as opposed to our two-branch fine-tuning approach. Tab. 1 reveals that our method generates more plausible root trajectories with fewer collisions and more accurate goal-reaching. We also see that training our full two-branch architecture from scratch (*1-stage train* in Tab. 1), instead of using pre-training then fine-tuning, degrades both goal reaching and final full-body motion after in-painting.

A qualitative comparison of generated motions in different rooms is shown in Fig. 5. GMD tends to generate simple walking-straight trajectories. OmniControl and GMD do not reach the goal pose accurately and ignore the surroundings, leading to collisions with the environment. Our method TeSMo is able to generate diverse locomotion styles controlled by text in various scenes, achieving superior goal-reaching accuracy compared to other methods.

Table 3: Test-time guidance evaluation. Adding guidance to reach goal poses and avoid collisions during inference improves performance. Lower is better for all metrics.

Guidance		Navigation		Interaction		
Goal Reach	Collision	Goal Pos.	Collision	Goal Pos.	Pen. Val.	Pen. Ratio
✗	✗	0.1568	0.0294	0.1445	0.0043	0.0611
✓	✗	0.118	0.0342	0.1453	0.0050	0.0554
✗	✓	0.1550	0.0013	0.1407	0.0040	0.0414
✓	✓	0.1241	0.0012	0.1404	0.0045	0.0494




Fig. 6: Compared with DIMOS [53], our method generates more realistic human-object interactions with reduced floating and interpenetrations.

Interaction. Tab. 2 compares our approach to DIMOS [53], a state-of-the-art method to generate interactions trained with reinforcement learning. DIMOS requires a full-body final goal pose as input to the policy, unlike our approach which uses just the pelvis pose. Despite this, DIMOS struggles to reach the goal accurately, likely due to error accumulation during autoregressive rollout. Our method showcases fewer instances of interpenetration with interaction objects and the user study reveals a distinct preference for motions generated by our approach (preferred 71.9%) over those produced by DIMOS. Fig. 6 compares the approaches qualitatively, where we see that more accurate goal-reaching reduces floating or penetrating the chair during sitting. Moreover, the interactions generated by DIMOS lack diversity, and cannot be conditioned on text.

4.4 Analysis of Capabilities

In Fig. 1, our method carries out a sequence of actions, enabling traversal and interaction with multiple objects within a scene. Fig. 7 demonstrates additional key capabilities. In the top section, our method is controlled through a variety of text prompts. For interactions in particular, diverse text descriptions disambiguate between actions like sitting or standing up, and allow stylizing the sitting motion, e.g., with crossed arms. In the middle section, we enable user control over trajectories by adhering to a predefined A* path. By adjusting the blend scale,

users can adjust how closely the generated trajectory follows A^* . At the bottom of Fig. 7, we harness guidance at test time to encourage motions to reach the goal while avoiding collisions and penetrations. As shown in Tab. 3, combining guidance losses gives improved results both for navigation and interactions.




Fig. 7: TeSMo capabilities. (**Top**) Diverse text control; (**Middle**) Following A^* path with adherence controlled by the blend scale; (**Bottom**) Test-time guidance encourages locomotion to reach the goal accurately without colliding with the environment.

5 Discussion

We introduced TeSMo, a novel method for text-controlled scene-aware motion generation. By first pre-training a scene-agnostic text-to-motion diffusion model on large-scale motion capture data and subsequently fine-tuning with a scene-aware component, our text-conditioned method enables generating realistic and diverse human-object interactions within 3D scenes. To support such training, we introduced the new Loco-3D-FRONT dataset containing realistic navigation motions placed in 3D scenes, and extended the SAMP dataset with additional objects and text annotations. Experiments demonstrate that our generated motion is on par with state-of-the-art diffusion models, while improving the plausibility and realism of interactions compared to prior work.

Limitations & Future Work. While our navigation model enables accurate goal-reaching and text-to-motion controllability, the two-stage process can sometimes lead to a disconnect between the generated pelvis trajectory and in-painted full-body poses. Exploring new one-stage models, capable of simultaneously generating pelvis trajectories and poses, would streamline the process. Additionally, our current approach, which operates on 2D floor maps, restricts the ability to handle intricate interactions, such as a person stepping over a small stool.

Our current approach is aimed at controllability to allow users to specify text prompts or goal objects and locations. However, our method may also fit into recently proposed pipelines [41] that employ LLM planners to specify a sequence of actions and contact information that could be used to guide our motion generation. Looking ahead, we also aim to broaden the spectrum of actions modeled by the system, to encompass activities such as lying down and touching. Furthermore, enabling interactions with dynamic objects will allow for more interactive and realistic scenarios.

Acknowledgments. Thanks Yangyi Huang and Yifei Liu for technical support. Thanks Mathis Petrovich and Nikos Athanasiou for the fruitful discussion about text-to-motion synthesis. Thanks Tomasz Niewiadomski, Taylor McConnell, and Tsvetelina Alexiadis for running the user study.

Disclosure. MJB has received research gift funds from Adobe, Intel, Nvidia, Meta, and Amazon. MJB has financial interests in Amazon, Datagen Technologies, and Meshcapade GmbH. While MJB is a consultant for Meshcapade, his research in this project was performed solely at, and funded solely by, the Max Planck Society.

A Ablation Study

Table 4: Ablation study comparing various full-body infilling methods and different representations of navigation motion generation using the Loco-3D-FRONT test set. **(Left)** For generated pelvis trajectories, our approach achieves the best goal-reaching accuracy with low collision rate. **(Right)** After in-painting the full-body motion, our method preserves diverse and realistic movements that align with the provided text prompt, much like the model employing an alternative OminiControl full-body in-painting technique. However, our approach distinctly outperforms the model utilizing full-body representation.

Method	Root trajectory evaluation				Full-body motion evaluation			
	Goal-reaching error ↓				FID ↓	R-precision ↑	Diversity ↑	Foot skating ↓
Pos.	Orient.	Height	Collision ↓					
Ours (OminiControl [42] in-painting)	0.459	0.999	0.090	0.073	17.927	0.396	6.288	0.0308
Ours (full-body rep)	0.844	0.016	0.110	0.124	24.642	0.189	6.967	0.169
Ours	0.169	0.119	0.008	0.031	20.465	0.376	6.415	0.056

Alternative Full-Body In-painting Approach. While our root trajectory generation approach can integrate with several motion in-painting techniques, in the

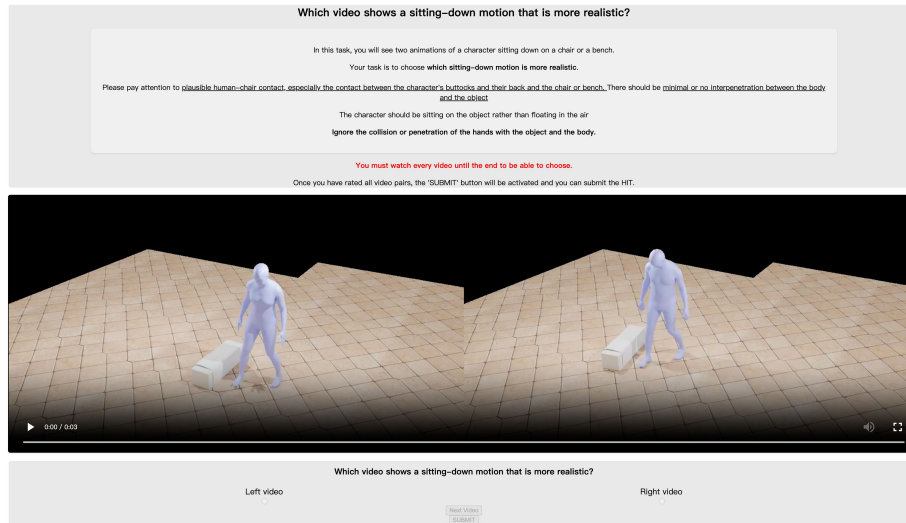


Fig. 8: The layout of our perceptual study for evaluating the plausibility of human-object interaction.

main paper we use PriorMDM [34]. As an alternative, we evaluate our method using OmniControl [42] for in-painting in Tab. 4. However, OmniControl overrides our generated dense pelvis trajectory and jointly generates full-body locomotion with a new pelvis trajectory. This severely degrades the goal-reaching ability (from 0.169 cm to 0.459 cm) as demonstrated in Table 4. Therefore, we choose to utilize PriorMDM as our body motion in-painting method. It aligns well with our generated trajectory, resulting in the generation of plausible locomotion while maintaining adherence to the goal position.

One-stage Navigation Motion Generation. To evaluate the efficacy of our two-stage navigation model design, we compare to a single-stage full-body motion generation ablation of our model. This model operates on the same input data but directly generates full-body locomotion. However, as shown in Tab. 4, this approach limits goal-reaching ability and does not produce motion styles that align with the input text. The local poses are somewhat dissociated from the global pelvis trajectories, allowing for trajectory variations while maintaining the same motion style. For instance, individuals can walk along different paths while maintaining consistency in their motion style.

B Details on User Study for Interaction Motions

To evaluate the plausibility of human-object interaction, we perform a user study to compare our method and DIMOS [53]. We employ Amazon Mechanical Turk (AMT) [2] to solicit assessments from 30 individuals. Raters are presented with

two side-by-side videos depicting generated interactions and asked to determine which appeared more realistic, particularly focusing on the contact between the character’s buttocks and their back with the chair or bench, and the presence of minimal or no interpenetration between the body and the object. We present 70 test videos with the positions of our generated videos and DIMO’s results randomly shuffled horizontally. In order to filter out poor responses, we duplicated our 5 test examples where clear preferences between two video results were evident, serving as catch trials. Ultimately, we obtained 65 useful responses out of 70 raters. The full survey page is illustrated in Fig. 8. The user study reveals a distinct preference for motions generated by our approach (preferred 71.9%) over those produced by DIMOS.

C Details on Collision Guidance Used in Interaction Motion Generation

At test time, a collision objective is used to discourage penetrations between humans and objects. Remarkably, our interaction motion generation model outputs 3D joint positions. We then link randomly sampled vertices on the SMPL mesh surfaces with the 3D skeletons in an A-pose, allowing us to obtain the posed sampled vertices for each new pose. This is defined as $\mathcal{J}_c = \text{SDF}(\hat{\mathbf{x}}_0, \mathcal{S}_O)$ where SDF calculates the SDF volume of the object O , then queries the sign distance value at each time step of the body vertices. Positive distances, indicating body vertices inside the interactive object, are averaged to get the final loss.

References

1. Agrawal, S., van de Panne, M.: Task-based locomotion. *ACM Transactions on Graphics* **35**(4), 1–11 (Jul 2016). <https://doi.org/10.1145/2897824.2925893>, <http://dx.doi.org/10.1145/2897824.2925893> 3
2. Amazon Web Services, Inc.: Amazon mechanical turk (Accessed 2024), <https://www.mturk.com/> 11, 16
3. Chao, Y.W., Yang, J., Chen, W., Deng, J.: Learning to sit: Synthesizing human-chair interactions via hierarchical control. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. pp. 5887–5895 (2021) 4
4. Corona, E., Pumarola, A., Alenyà, G., Moreno-Noguer, F.: Context-aware human motion prediction. Cornell University - arXiv, Cornell University - arXiv (Apr 2019) 3
5. Diller, C., Dai, A.: Cg-hoi: Contact-guided 3d human-object interaction generation (2024) 4
6. Eigen, D., Ranzato, M., Sutskever, I.: Learning factored representations in a deep mixture of experts. *arXiv: Learning, arXiv: Learning* (Dec 2013) 3
7. Fu, H., Cai, B., Gao, L., Zhang, L.X., Wang, J., Li, C., Zeng, Q., Sun, C., Jia, R., Zhao, B., et al.: 3d-front: 3d furnished rooms with layouts and semantics. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 10933–10942 (2021) 3, 8, 9

8. Guo, C., Zou, S., Zuo, X., Wang, S., Ji, W., Li, X., Cheng, L.: Generating diverse and natural 3d human motions from text. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5152–5161 (June 2022) [2](#), [6](#), [7](#), [8](#), [10](#)
9. Hart, P.E., Nilsson, N.J., Raphael, B.: A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics* **4**(2), 100–107 (1968) [8](#)
10. Hassan, M., Ceylan, D., Villegas, R., Saito, J., Yang, J., Zhou, Y., Black, M.J.: Stochastic scene-aware motion prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 11374–11384 (October 2021) [2](#), [3](#), [4](#), [8](#), [9](#), [10](#), [11](#), [12](#)
11. Hassan, M., Choutas, V., Tzionas, D., Black, M.J.: Resolving 3d human pose ambiguities with 3d scene constraints. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019) [3](#), [4](#)
12. Hassan, M., Guo, Y., Wang, T., Black, M., Fidler, S., Peng, X.B.: Synthesizing physical character-scene interactions. In: ACM SIGGRAPH '23 Conference Proceedings. SIGGRAPH '23, Association for Computing Machinery, New York, NY, USA (2023). <https://doi.org/10.1145/3588432.3591525>, <https://doi.org/10.1145/3588432.3591525> [2](#), [4](#)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. *CoRR* **abs/1512.03385** (2015) [7](#)
14. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 6840–6851. Curran Associates, Inc. (2020), <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf> [5](#)
15. Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. *arXiv preprint* (2022) [6](#)
16. Holden, D., Komura, T., Saito, J.: Phase-functioned neural networks for character control. *ACM Transactions on Graphics* p. 1–13 (Aug 2017). <https://doi.org/10.1145/3072959.3073663>, <http://dx.doi.org/10.1145/3072959.3073663> [3](#)
17. Huang, S., Wang, Z., Li, P., Jia, B., Liu, T., Zhu, Y., Liang, W., Zhu, S.C.: Diffusion-based Generation, Optimization, and Planning in 3D Scenes. *arXiv e-prints* arXiv:2301.06015 (Jan 2023). <https://doi.org/10.48550/arXiv.2301.06015> [2](#), [4](#)
18. Jacobs, R.A., Jordan, M.I., Nowlan, S.J., Hinton, G.E.: Adaptive mixtures of local experts. *Neural Computation* p. 79–87 (Feb 1991). <https://doi.org/10.1162/neco.1991.3.1.79>, <http://dx.doi.org/10.1162/neco.1991.3.1.79> [3](#)
19. Karunratanakul, K., Preechakul, K., Suwajanakorn, S., Tang, S.: Guided motion diffusion for controllable human motion synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2151–2162 (2023) [2](#), [3](#), [4](#), [10](#), [11](#)
20. Kovar, L., Gleicher, M., Pighin, F.: Motion graphs. In: *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. pp. 723–732 (2023) [3](#)
21. Kulkarni, N., Rempe, D., Genova, K., Kundu, A., Johnson, J., Fouhey, D., Guibas, L.: Nifty: Neural object interaction fields for guided human motion synthesis (2023) [4](#)
22. Lee, J., Chai, J., Reitsma, P.S.A., Hodgins, J.K., Pollard, N.S.: Interactive control of avatars animated with human motion data. In: Proceedings of the 29th annual conference on Computer graphics and interactive techniques (Jul 2002). <https://>

- doi.org/10.1145/566570.566607, <http://dx.doi.org/10.1145/566570.566607>
3
23. Lee, J., Joo, H.: Locomotion-action-manipulation: Synthesizing human-scene interactions in complex 3d environments. In: International Conference on Computer Vision (ICCV) (2023) 2
 24. Lee, K.H., Choi, M.G., Lee, J.: Motion patches. ACM Transactions on Graphics p. 898–906 (Jul 2006). <https://doi.org/10.1145/1141911.1141972>, <http://dx.doi.org/10.1145/1141911.1141972> 3
 25. Li, J., Clegg, A., Mottaghi, R., Wu, J., Puig, X., Liu, C.K.: Controllable human-object interaction synthesis. arXiv preprint arXiv:2312.03913 (2023) 3
 26. Li, J., Wu, J., Liu, C.K.: Object motion guided human motion synthesis. ACM Transactions on Graphics (TOG) **42**(6), 1–11 (2023) 3
 27. Mahmood, N., Ghorbani, N., Troje, N.F., Pons-Moll, G., Black, M.J.: AMASS: Archive of motion capture as surface shapes. In: International Conference on Computer Vision. pp. 5442–5451 (Oct 2019) 2, 3
 28. Peng, X., Xie, Y., Wu, Z., Jampani, V., Sun, D., Jiang, H.: Hoi-diff: Text-driven synthesis of 3d human-object interactions using diffusion models. arXiv preprint arXiv:2312.06553 (2023) 4
 29. Peng, X.B., Guo, Y., Halper, L., Levine, S., Fidler, S.: Ase: Large-scale reusable adversarial skill embeddings for physically simulated characters. ACM Trans. Graph. **41**(4) (Jul 2022) 4
 30. Petrovich, M., Litany, O., Iqbal, U., Black, M.J., Varol, G., Peng, X.B., Rempe, D.: STMC: Multi-track timeline control for text-driven 3d human motion generation. arXiv:2401.08559 (2024) 4
 31. Pi, H., Peng, S., Yang, M., Zhou, X., Bao, H.: Hierarchical generation of human-object interactions with diffusion probabilistic models. In: ICCV. pp. 15061–15073 (October 2023) 2, 4
 32. Prokudin, S., Lassner, C., Romero, J.: Efficient learning on point clouds with basis point sets. In: International Conference on Computer Vision (ICCV). pp. 4332–4341 (2019) 9
 33. Rempe, D., Luo, Z., Peng, X.B., Yuan, Y., Kitani, K., Kreis, K., Fidler, S., Litany, O.: Trace and pace: Controllable pedestrian animation via guided trajectory diffusion. In: CVPR (2023) 4, 5, 6, 7, 11, 12
 34. Shafir, Y., Tevet, G., Kapon, R., Bermano, A.H.: Human motion diffusion as a generative prior. arXiv preprint arXiv:2303.01418 (2023) 2, 3, 4, 5, 6, 8, 16
 35. Starke, S., Zhang, H., Komura, T., Saito, J.: Neural state machine for character-scene interactions. ACM Trans. Graph. **38**(6) (nov 2019). <https://doi.org/10.1145/3355089.3356505>, <https://doi.org/10.1145/3355089.3356505> 2, 3
 36. Taheri, O., Choutas, V., Black, M.J., Tzionas, D.: GOAL: Generating 4D whole-body motion for hand-object grasping. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2022), <https://goal.is.tue.mpg.de> 3
 37. Tevet, G., Raab, S., Gordon, B., Shafir, Y., Bermano, A.H., Cohen-Or, D.: Human motion diffusion model. arXiv preprint arXiv:2209.14916 (2022) 2, 4, 5, 6
 38. Wang, J., Xu, H., Xu, J., Liu, S., Wang, X.: Synthesizing long-term 3d human motion and interaction in 3d scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9401–9411 (2021) 3
 39. Wang, J., Yan, S., Dai, B., Lin, D.: Scene-aware generative network for human motion synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12206–12215 (2021) 3

40. Wang, Z., Chen, Y., Liu, T., Zhu, Y., Liang, W., Huang, S.: Humanise: Language-conditioned human motion generation in 3d scenes. In: Advances in Neural Information Processing Systems (NeurIPS) (2022) 4
41. Xiao, Z., Wang, T., Wang, J., Cao, J., Dai, B., Lin, D., Pang, J.: Unified human-scene interaction via prompted chain-of-contacts. In: International Conference on Learning Representations (ICLR) (2024) 4, 15
42. Xie, Y., Jampani, V., Zhong, L., Sun, D., Jiang, H.: Omnicontrol: Control any joint at any time for human motion generation. arXiv preprint arXiv:2310.08580 (2023) 2, 3, 4, 5, 6, 11, 15, 16
43. Ye, S., Wang, Y., Li, J., Park, D., Liu, C.K., Xu, H., Wu, J.: Scene synthesis from human motion. In: SIGGRAPH Asia 2022 Conference Papers. SA '22, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3550469.3555426>, <https://doi.org/10.1145/3550469.3555426> 4
44. Yi, H., Huang, C.H.P., Tripathi, S., Hering, L., Thies, J., Black, M.J.: MIME: Human-aware 3D scene generation. In: Computer Vision and Pattern Recognition (CVPR) (June 2023) 4, 8
45. Yi, H., Huang, C.H.P., Tzionas, D., Kocabas, M., Hassan, M., Tang, S., Thies, J., Black, M.J.: Human-aware object placement for visual environment reconstruction. In: Computer Vision and Pattern Recognition (CVPR). pp. 3959–3970 (Jun 2022) 9
46. Yuksel, S.E., Wilson, J.N., Gader, P.D.: Twenty years of mixture of experts. IEEE Transactions on Neural Networks and Learning Systems p. 1177–1193 (Aug 2012). <https://doi.org/10.1109/tnnls.2012.2200299>, <http://dx.doi.org/10.1109/tnnls.2012.2200299> 3
47. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023) 6
48. Zhang, M., Cai, Z., Pan, L., Hong, F., Guo, X., Yang, L., Liu, Z.: Motiandiffuse: Text-driven human motion generation with diffusion model. arXiv preprint arXiv:2208.15001 (2022) 2, 4, 5
49. Zhang, W., Dabral, R., Leimkühler, T., Golyanik, V., Habermann, M., Theobalt, C.: Roam: Robust and object-aware motion generation using neural pose descriptors. International Conference on 3D Vision (3DV) (2024) 4
50. Zhang, X., Bhatnagar, B.L., Starke, S., Guzov, V., Pons-Moll, G.: Couch: Towards controllable human-chair interactions. In: European Conference on Computer Vision (ECCV). Springer (October 2022) 2, 4
51. Zhang, Y., Tang, S.: The wanderings of odysseus in 3d scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20481–20491 (2022) 3
52. Zhao, K., Wang, S., Zhang, Y., Beeler, T., , Tang, S.: Compositional human-scene interaction synthesis with semantic control. In: European conference on computer vision (ECCV) (2022) 4
53. Zhao, K., Zhang, Y., Wang, S., Beeler, T., , Tang, S.: Synthesizing diverse human motions in 3d indoor scenes. In: International conference on computer vision (ICCV) (2023) 2, 3, 4, 8, 12, 13, 16
54. Zhu, W., Ma, X., Ro, D., Ci, H., Zhang, J., Shi, J., Gao, F., Tian, Q., Wang, Y.: Human motion generation: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023) 3