

SceneMI: Motion In-betweening for Modeling Human-Scene Interactions

Inwoo Hwang^{1*} Bing Zhou^{2†} Young Min Kim¹ Jian Wang² Chuan Guo^{2†}

¹Dept. of Electrical and Computer Engineering, Seoul National University, ²Snap Inc.




Figure 1. **SceneMI** synthesizes physically plausible transitions (colored in purple) that simultaneously satisfy keyframe constraints (colored in blue) and environmental affordances in challenging scenarios. The model exhibits robust generalization with *noisy* keyframes in phone-scanned scenes from real-world data, GIMO [90] (third figure). SceneMI can further assist realistic 3D human-scene interaction reconstruction only from monocular video inputs (rightmost figure).

Abstract

Modeling human-scene interactions (HSI) is essential for understanding and simulating everyday human behaviors. Recent approaches utilizing generative modeling have made progress in this domain; however, they are limited in controllability and flexibility for real-world applications. To address these challenges, we propose reformulating the HSI modeling problem as Scene-aware Motion In-betweening—a more tractable and practical task. We introduce SceneMI, a framework that supports several practical applications, including keyframe-guided character animation in 3D scenes and enhancing the motion quality of imperfect HSI data. SceneMI employs dual scene descriptors to comprehensively encode global and local scene context. Furthermore, our framework leverages the inherent denoising nature of diffusion models to generalize on noisy keyframes. Experimental results demonstrate SceneMI’s effectiveness in scene-aware keyframe in-betweening and generalization to the real-world GIMO dataset, where motions and scenes are acquired by noisy IMU sensors and smartphones. We further showcase SceneMI’s applicability in HSI reconstruction from monocular videos. Project page: <http://inwoohwang.me/SceneMI>

*Most work done during an internship at Snap Research NYC, Snap Inc.

†Co-corresponding author

1. Introduction

Modeling dynamic human motions in everyday environments is challenging due to the inherent complexity of understanding and replicating human-scene interactions (HSI). Various studies have approached this challenge by learning the data distribution of human motions in scenes through generative modeling based on textual descriptions [29, 71, 79], action categories [30, 43, 70], or past motion sequences [23, 45, 67, 90]. Despite these advances, these methods commonly lack controllability in motions and find limited flexibility in real applications. This motivates our task in this work—scene-aware motion in-betweening, which aims to synthesize natural motion transitions between specific keyframes in 3D scenes while adapting to environmental constraints, such as avoiding obstacles. The availability of keyframe context effectively reduces task complexity while supporting several applications in real scenarios. For instance, animators can directly incorporate scene information to create 3D character animations from keyframes, as shown in Figure 1. Additionally, this technique reliably decreases motion artifacts in existing real-world HSI data (Tab. 4) and enhances HSI reconstruction from monocular videos (Fig. 5).

Motion in-betweening has been extensively studied for *isolated* human motions [10, 20, 21, 21, 58, 85]. However,

trivially applying these approaches within 3D scenes often leads to undesirable body-scene penetrations. While some existing scene-aware approaches model goal pose reaching [67, 70] as a subsystem component, they typically employ conditional VAEs to capture scene and pose constraints, resulting in limited expressivity and scalability. Another challenge lies in the imperfection of keyframes and scenes in real-world scenarios—such as keyframes captured from noisy motion capture sensors and variations of scenes acquired by different devices.

In this work, we propose SceneMI, a conditional diffusion model tailored for scene-aware motion in-betweening. We extract hierarchical scene features: the global scene is represented by a coarse occupancy voxel grid, while local scenes are encoded as Basis Point Set (BPS) features computed at keyframe time steps based on selected surface body markers. This scene encoding effectively captures the high-level scene layout and spatial context anchored on immediate keyframes, preventing overfitting on specific local scene geometries. We apply imputation to keyframes to effectively encode the temporal location of keyframes within the motion sequence. We particularly leverage the inherent denoising capabilities of diffusion models to handle noisy keyframes. During inference, *noisy* keyframes guide the denoising sampling from steps T until T^* (a hyperparameter). For the remaining steps (from $T^* - 1$ to 1), the full motion sequences, including *noisy* key poses, are iteratively denoised.

In our experiment, SceneMI is trained exclusively on the TRUMANS dataset [31], a large-scale HSI dataset with motion capture data and hand-crafted indoor scenes. For evaluation, beyond analysis on TRUMANS, we test the pre-trained model on out-of-domain human-scene interactions from the GIMO dataset [90], where motions are captured using inaccurate IMU sensors and scenes are scanned by smartphones in real-world environments. Empirical results demonstrate that SceneMI generalizes well across different scene types and key pose qualities, excelling at synthesizing smooth transitions. Additionally, SceneMI reliably reduces foot skating by 37.5% (from 0.261 to 0.163) and jittering by 56.5% (from 0.573 to 0.249), which are artifacts prevalent in the GIMO dataset. To further explore SceneMI’s applicability, we develop a framework to reconstruct human-scene interactions from monocular videos using image-to-3D techniques [75, 78] and human pose estimation [13]. Our results show that SceneMI significantly enhances interaction naturalness and reduces penetration when applied to synthetic 3D scenes and keyframes.

In summary, this work formally studies the problem of scene-aware motion in-betweening for the first time and explores its application in various HSI scenarios. This is realized by a simple diffusion-based model, SceneMI, which efficiently encodes scene contexts at global and local scales, with specialized denoising procedures to accommodate *noisy*

keyframes. Our comprehensive analysis across different scene resources and motion quality, including real-world data, demonstrates SceneMI’s motion in-betweening capability and scalability. Additionally, we highlight SceneMI’s applicability in reconstructing HSI from monocular video.

2. Related works

2.1. Human Motion Synthesis

Advancements in data-driven approaches have led to appealing breakthroughs in human motion synthesis. Beyond achieving high-quality and realistic motion outputs, many applications are designed to generate motions adaptable to various conditions, such as action categories [8, 15, 24, 52, 64], text [4, 16–19, 28, 44, 53, 54, 64, 80, 81, 83], audio [1, 14, 37, 42, 57, 61, 65], interacting objects [3, 6, 12, 36, 40, 41, 63, 76, 87], and scenes [7, 25, 27, 29, 30, 38, 43, 47, 48, 62, 71, 72, 79, 88]. Recently, there has been growing interest in trajectory-based controllable motion generation. For instance, PriorMDM [60] fine-tuned an existing motion diffusion model MDM [64] to support end-effector trajectory control. GMD [60] incorporated classifier-free guidance into a diffusion model to enable root joint-based controllable generation. OmniControl [74] expanded this capability to arbitrary joint trajectories using a combination of ControlNet [82] and inference-time guidance. The recent ControlMM [55] approach, based on masked motion modeling, further enables multi-joint trajectory-based controllable motion generation.

Human-Scene Interaction Modeling. Despite promising progress in motion synthesis, modeling realistic human movements within environmental scenes remains challenging. Reinforcement learning-based methods [23, 25, 86, 89] have demonstrated the ability to produce goal-oriented locomotion by learning policies with carefully designed reward functions. UniHSI [73] leverages large language models for detailed motion planning to capture complex human-scene interactions. Another line of research adopts data-driven deep learning models. The complexity of this task often necessitates a hierarchical approach: first generating paths and start/end poses, followed by full motion synthesis [68–70]. The latest advances include end-to-end generative models. HUMANISE [71] and LINGO [29] attempt to generate 3D human motions in scenes from text prompts using diffusion models, while TRUMANS [30] experiments with an autoregressive diffusion model conditioned on action labels. SceneDiffuser [25] combines diffusion models with reinforcement learning, extending motion dynamics from a start pose within scenes. Nevertheless, existing works lack motion controllability and practical applicability, such as recognizing human-scene interactions from real-world demonstrations or aiding in their reconstruction.

2.2. Motion In-betweening

Motion in-betweening synthesizes complete motion sequences from keyframes positioned at specific time steps. While traditional approaches use spline interpolation such as Bézier curves [39, 49, 59], these methods often require extensive post-processing adjustments. Earlier research framed this task as motion planning [2, 5], employing search and optimization over motion graphs [35]. However, these approaches typically demand intensive computational resources to maintain and search large databases during inference. Deep learning has transformed this field. Given the temporal nature of motion data, RNNs have been widely adopted for in-betweening [20, 21, 85]. Transformer architectures have further improved the modeling of long-term dependencies [11, 34, 50, 58], exemplified by Qin et al. [58], which developed a two-stage Transformer generating both rough and refined transitions. While most approaches treat motion in-betweening as deterministic, recent work has explored generative methods. CondMDI [10] leverages diffusion models to produce natural, diverse transitions that align with sparse keyframes and optional text instructions. Despite these advancements, a key limitation remains: most existing methods overlook environmental factors that influence motion behavior and assume precise keyframes, which is impractical for real-world inputs.

3. Methodology

Given a 3D scene \mathcal{G} , a sparse set of key poses $\mathbf{s} \in \mathbb{R}^{N \times D}$ with key pose indicator $\mathbf{m} \in \{0, 1\}^N$, our goal is to synthesize a complete motion sequence $\mathbf{x} = \{x^n\} \in \mathbb{R}^{N \times D}$ that satisfies both the key poses and the environmental constraints of the 3D scene. Here, N represents the total number of poses in the full sequence, with each pose represented by a D -dimensional feature vector. The mask $m_n \in \mathbf{m}$ indicates whether n^{th} frame contains a key pose ($m_n = 1$) or non-key pose ($m_n = 0$), with a total of $k = \sum_n m_n \ll N$ key poses.

Motion and Shape Representation Each pose feature vector comprises global joint positions $J \in \mathbb{R}^{22 \times 3}$ (include root translation $\gamma \in \mathbb{R}^3$), 6D root orientation $\phi \in \mathbb{R}^6$, and local SMPL [46] pose parameters $\psi \in \mathbb{R}^{21 \times 6}$. These components collectively form a 201-dimensional pose feature representation. To account for body volume variations, we extract human shape features $\mathbf{b} \in \mathbb{R}^7$ as distances between a set of representative joint pairs, for example [root, head], [left.shoulder, right.shoulder].

3.1. Scene Encoding

We need to provide a compact yet descriptive scene representation as the condition of our motion in-betweening module. As shown in Fig. 2 (middle block), we adopt distinct strate-

gies to encode scenes at global and local scales. The global scene feature \mathbf{c}_g encapsulates the expansive scene layout, providing crucial information for navigating the overall motion trajectory. The local scene features \mathbf{c}_l encode spatial information centered around individual keyframes, enabling fine-grained interactions with the environment. Together, these features $\mathbf{c} = [\mathbf{c}_g, \mathbf{c}_l]$ provide a comprehensive environmental context that guides motion in-betweening, which also presents strong generalization capability to out-of-domain 3D scenes.

Global Scene Features The global scene features represent the overall spatial context of the entire scene as an occupancy voxel grid at a coarse resolution of $\mathbf{c}_g \in \{0, 1\}^{d_x \times d_y \times d_z}$ with 0.1m per voxel, where a value of 1 indicates an occupied voxel and 0 represents an unoccupied space. The center and orientation of this grid are initialized by the root position and orientation in the first frame. To effectively encode these global scene features, we employ a Vision Transformer (ViT) model, following the approach in [30]. This global scene encoder processes inputs with dimensions $48 \times 24 \times 48$ and results in a 512-dimensional feature vector. The global scene encoder is jointly trained with the rest of the framework.

Keyframe-centered Local Scene Features We extract local scene features using the Basis Point Set (BPS) [56] approach surrounding the keyframe poses. First, we determine 64 anchor points on the T-posed SMPL mesh surface via farthest-point sampling, creating a structured abstraction of the body shape geometry. These anchor vertex indices remain fixed across all keyframe poses throughout our experiments. For each keyframe, we calculate the nearest point from the 3D scene to each indexed anchor point on the body mesh. The resulting ordered offset vectors—relative positions to these nearby points—serve as the BPS features $\mathbf{c}_l^n \in \mathbb{R}^{64 \times 3}$ for the immediate local scene surrounding each keyframe. These features are embedded through an MLP and then concatenated with key pose features at their corresponding frames. Compared to traditional geometry representation, BPS features specifically target scene context close to the human body while remaining agnostic to point order, mesh topology, and resolution, thereby enhancing generalizability on different scene sources.

3.2. Scene-aware Motion In-betweening

We employ diffusion models for motion in-betweening, which consist of a forward diffuse process and a reverse denoising process. Beginning with a conditional data distribution $p(\mathbf{x}_0|\tau)$, where the conditions are the set of scene context \mathbf{c} , and body shape \mathbf{b} , $\tau = \{\mathbf{c}, \mathbf{b}\}$, and the denoised variable is the clean pose $\mathbf{x}_0 = \mathbf{x}$, the forward process progressively corrupts original data samples until reaching diffusion step T . Noisy samples at timestep t are defined




Figure 2. **Overview of SceneMI.** Given the input 3D scene, we extract global voxelized features \mathbf{c}_g , and local BPS features \mathbf{c}_l^n based on key pose meshes. During training, key poses within the motion sequence are imputed based on the indicator mask \mathbf{m} . The model integrates these scene features (\mathbf{c}_l and \mathbf{c}_g), body shape features \mathbf{b} , and the diffusion timestep t to synthesize the full motion sequence \mathbf{x} that satisfy both the keyframe and scene constraints.

as $\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\boldsymbol{\epsilon}$, where $\bar{\alpha}_t$ is determined by the diffusion noise schedule and $\boldsymbol{\epsilon}$ is sampled from an i.i.d. Gaussian distribution. In the reverse process, the denoising model \mathcal{D}_θ synthesizes clean data samples by recursively refining from pure Gaussian noise $q(\mathbf{x}_T) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. This process is conditioned on both the timestep t and the corresponding conditional features τ . The denoiser \mathcal{D}_θ is trained to reconstruct the original data \mathbf{x}_0 :

$$\mathcal{L}_{\text{simple}} = \mathbb{E}_{\mathbf{x}_0 \sim p(\mathbf{x}_0 | \tau), t \sim [1, T]} [\|\mathbf{x}_0 - \mathcal{D}_\theta(\mathbf{x}_t, t, \tau)\|_2^2] \quad (1)$$

Model Architecture. Our model employs a U-Net based architecture with Adaptive Group Normalization (AdaGN) and 1D convolution, which has demonstrated strong performance in global motion representation [33]. AdaGN dynamically adapts normalization, enhancing the model’s ability to capture motion dynamics, while 1D convolution facilitates the learning of sequential motion data. The details for the architecture are provided in the supplementary material.

Training Procedure During training, we first randomly sample a diffusion timestep $t \sim \mathcal{U}(\{1, \dots, T\})$. Then, $k \sim \mathcal{U}(\{2, \dots, N\})$ keyframes are randomly selected from the motion sequence, including the start and end frames, forming a binary mask $\mathbf{m} \in \{0, 1\}^N$ where $m_n = 1$ indicates a keyframe location. Using this mask \mathbf{m} , we perform imputation on the noisy sample \mathbf{x}_t by substituting the feature

values at keyframe locations with their corresponding clean values from \mathbf{x}_0 , as follows:

$$\mathbf{x}'_t = \mathbf{m} \odot \mathbf{x}_0 + (1 - \mathbf{m}) \odot \mathbf{x}_t$$

where \odot denotes element-wise multiplication. And we retain local scene features \mathbf{c}_l , which are visible from selected keyframes:

$$\mathbf{c}'_l = \mathbf{m} \odot \mathbf{c}_l$$

The motion-related features are obtained by concatenating the features of $\tilde{\mathbf{x}}_t = \text{spatialconcat}(\mathbf{x}'_t, \mathbf{c}'_l, \mathbf{m})$. These features are further concatenated with body shape features \mathbf{b} and global scene features \mathbf{c}_g along the temporal dimension. Subsequently, the embeddings of diffusion steps (i.e., t) are added to all input features, as illustrated in Fig. 2. During training, we randomly mask out \mathbf{c}_g with a probability of 10% to enable classifier-free guidance during inference.

In addition to the reconstruction loss $\mathcal{L}_{\text{simple}}$ defined in Eq. 1, our denoiser \mathcal{D}_θ is trained with additional losses on the global joint positions and joint velocities to enhance realism, which are obtained through forward kinematics (FK) of the predicted SMPL parameters, as following:

$$\mathcal{L}_{\text{joints}} = \|\text{FK}(\mathbf{x}_0) - \text{FK}(\mathcal{D}_\theta(\mathbf{x}_t, t, \tau))\|^2, \quad (2)$$

$$\mathcal{L}_{\text{vel}} = \|\text{diff}(\text{FK}(\mathbf{x}_0)) - \text{diff}(\text{FK}(\mathcal{D}_\theta(\mathbf{x}_t, t, \tau)))\|^2. \quad (3)$$

Overall, the final learning objective is:

$$\mathcal{L} = \mathcal{L}_{\text{simple}} + \lambda_{\text{joints}}\mathcal{L}_{\text{joints}} + \lambda_{\text{vel}}\mathcal{L}_{\text{vel}}. \quad (4)$$

Inference At inference, the goal is to synthesize the dense motion sequence \mathbf{x}_0 from the keyframes \mathbf{s} , along with the given keyframe indices \mathbf{m} , where $\mathbf{s}_n = \mathbf{x}_0^n$ if $\mathbf{m}_n = 1$. We directly calculate the global feature \mathbf{c}_g and the keyframe-centered local features \mathbf{c}'_t . At each sampling step, we impute the keyframe features following the same procedure used during training.

$$\mathbf{x}'_t = \mathbf{m} \odot \mathbf{s} + (1 - \mathbf{m}) \odot \mathbf{x}_t.$$

Then we concatenate all motion-related features $\tilde{\mathbf{x}}_t = \text{spatialconcat}(\mathbf{x}'_t, \mathbf{c}'_t, \mathbf{m})$ and apply classifier-free guidance with scale \mathbf{w} to the denoiser output:

$$\hat{\mathbf{x}}_0 = \mathbf{w} \cdot \mathcal{D}_\theta(\tilde{\mathbf{x}}_t, t, \mathbf{b}, \mathbf{c}_g) + (1 - \mathbf{w}) \cdot \mathcal{D}_\theta(\tilde{\mathbf{x}}_t, t, \mathbf{b}, \emptyset)$$

3.2.1. Handling Noisy Keyframes

In real-world scenarios, when the keyframe poses are captured by less accurate mocap sensors or extracted from videos, directly using *noisy* keyposes for motion inbetweening would compromise the output motion quality. Fortunately, diffusion models inherently add noise to motion data during the forward process, and the models learn to remove this noise during the reverse process. As suggested by MotionMix [24] and Rohm [84], the noise in imperfect motions could be analogous to the added noise in the diffusion process at certain timesteps. Motivated by this insight, we divide the diffusion/sampling timesteps into two ranges: $[T, T^* + 1]$ and $[T^*, 1]$. From early diffusion timestep T to $T^* + 1$, the prediction of \mathbf{x}_0 is guided by noisy keyposes $\mathbf{s}^{\text{noisy}}$ for keyframe alignment, while for the remaining timesteps (T^* to 1), both keyframes and in-betweening frames are iteratively denoised by \mathcal{D}_θ .

To learn motion inbetweening with noisy keyframes, we first create a noisy motion dataset $\mathbf{x}_0^{\text{noisy}}$ from clean motions \mathbf{x}_0 by adding unitary Gaussian noise scaled by a random factor $l \sim \mathcal{U}(0, 1.0)$. During training, we restrict the imputation operation to only occur until timestep T^* :

$$\mathbf{x}'_t = \begin{cases} \mathbf{m} \odot \mathbf{x}_0^{\text{noisy}} + (1 - \mathbf{m}) \odot \mathbf{x}_t, & t \in [T, T^* + 1] \\ \mathbf{x}_t, & t \in [T^*, 1] \end{cases}$$

The training objective is to predict the clean motion \mathbf{x}_0 .

During inference, we apply the same imputation strategy to noisy keyframes $\mathbf{s}^{\text{noisy}}$ for sampling:

$$\mathbf{x}'_t = \begin{cases} \mathbf{m} \odot \mathbf{s}^{\text{noisy}} + (1 - \mathbf{m}) \odot \mathbf{x}_t, & t \in [T, T^* + 1] \\ \mathbf{x}_t, & t \in [T^*, 1] \end{cases}$$

Empirically, we experiment with different T^* values (see Table 3) and find $T^* = 20$ performs the best, where the total diffusion step $T = 1000$.

4. Experiments

We provide a comprehensive evaluation of our approach in three settings: motion in-betweening with noise-free

keyframes and in-domain 3D scenes in Sec. 4.1, out-of-domain noisy keyframes and real-world scenes in Sec. 4.2, and applications for refining noisy human-scene interactions in Sec. 4.3.

Dataset. For training, we utilized TRUMANS [30], the largest high-quality motion capture dataset (30 FPS) with hand-crafted scene geometry. Motion sequences were segmented into $N = 121$ frames, with 10% randomly selected for the test set. To train for noisy keyframe handling, we added zero-mean noise with a standard deviation determined by noise level l , defined as $(l^\circ, l^\circ, l \text{ cm})$ for the ground-truth SMPL parameters (ϕ, ψ, γ) except the start and end frames. For evaluation, we tested our approach on both noise-free and synthetic noisy TRUMANS test sets. To further assess out-of-domain generalization robustness, we evaluated performance on the real-world GIMO dataset [90] and demonstrated video-based HSI reconstruction using the PROX dataset [22].

Implementation Details. We implement our pipeline using PyTorch [51] and train our model on a single NVIDIA RTX 3090 GPU. During training, we use $T = 1000$ diffusion timesteps and set $\lambda_{\text{joints}} = 2.0$ and $\lambda_{\text{vel}} = 10.0$, with $T^* = 20$ for the noisy setting. During inference, we use a guidance scale of $\mathbf{w} = 2.5$ with $T = 1000$ diffusion sampling steps.

Evaluation Metrics. We evaluate motion in-betweening results using three sets of metrics that assess motion naturalness, motion alignment, and human-scene interaction plausibility. For motion naturalness, we calculate the *Frechet Inception Distance (FID)*, which measures the distributional distance between high-level features of completed motions and ground truth data. We also quantify *foot skating* and *jerk* (the rate of change of acceleration) as indicators of motion naturalness. For motion alignment, we compute the *mean joint position error* at keyframe locations (MPJE Key) and across the full sequence (MPJE All). To evaluate HSI plausibility, we report the *collision frame ratio*—proportion of frames containing human-scene collisions relative to the total sequence length, and the *penetration max*—average maximum collision distance across each interaction sequence.

Baselines. Since scene-aware motion in-betweening has not been formally investigated in previous works, we adapt several state-of-the-art approaches from related fields as our baselines. These include diffusion models for scene-agnostic motion synthesis (MDM [64] and StableMoFusion [26]), motion in-betweening methods (OmniControl [74] and CondMDI [10]), and scene-aware motion synthesis approaches (SceneDiffuser [25] and Wang et al. [66]). Base-

| Method | FID ↓ | Foot Skating ↓ | Jerk (m/s^3) ↓ | MJPE Key (m) ↓ | MJPE All (m) ↓ | Collision Frame Ratio ↓ | Pene Max (m) ↓ |
|--------------------------------|--------------|----------------|--------------------|--------------------|--------------------|-------------------------|--------------------|
| MDM [64] | 1.422 | 0.316 | 0.972 | 0.568 | 0.576 | 0.317 | 0.112 |
| StableMoFusion [26] | 0.732 | 0.264 | 0.272 | 0.412 | 0.471 | 0.275 | 0.098 |
| SceneDiffuser [25] | 1.397 | 0.432 | 0.583 | 0.349 | 0.391 | 0.292 | 0.128 |
| Wang et al. [66] | 3.243 | 0.528 | 17.243 | 0.091 | 0.096 | 0.203 | 0.082 |
| OmniControl [74] | 0.371 | 0.294 | 0.274 | 0.217 | 0.294 | 0.211 | 0.081 |
| CondMDI [10] | 0.943 | 0.281 | 0.305 | 0.452 | 0.457 | 0.262 | 0.087 |
| Ours | 0.123 | 0.248 | 0.194 | 0.006 | 0.023 | 0.113 | 0.043 |
| w/o scene-awareness c_g, c_l | 0.136 | 0.251 | 0.103 | 0.012 | 0.059 | 0.131 | 0.049 |
| w/o global feature c_g | 0.138 | 0.254 | <u>0.131</u> | 0.011 | 0.051 | 0.128 | 0.048 |
| w/o local feature c_l | <u>0.125</u> | 0.245 | 0.196 | <u>0.008</u> | <u>0.036</u> | <u>0.119</u> | <u>0.045</u> |

Table 1. Quantitative scene-aware motion in-betweening results on TRUMANS dataset [30] with **noise-free** keyframes. Our method excels in in-betweening within scene constraints across various metrics. The keyframe interval is set to $r = 60$ frames. **Bold** represents the best




Figure 3. **Visual comparisons** with baseline models on **noise-free** TRUMANS dataset.

lines that originally did not incorporate scene information [10, 26, 64, 74] were re-trained with additional *global scene conditions*. Detailed descriptions of all baseline implementations are provided in the supplementary material.

4.1. Evaluation on Noise-free TRUMANS

We first evaluate SceneMI on the TRUMANS [30] dataset using the classical motion in-betweening setting with surrounding environments, where keyframes are clean, and 3D scenes belong to the same domain as the training data. In our experiments, we employ a sparse keyframe interval of 60 frames, corresponding to 2-second motion sequences between consecutive keyframes.

Comparison with Baselines. Table 1 presents the quantitative results. Our model demonstrates superior performance across all evaluation metrics, achieving the highest motion quality (FID of 0.123 compared to CondMDI), precise keyframe pose alignment (MJPE of 0.023 mm), and enhanced environmental awareness as evidenced by the lowest collision frame ratio (0.113). In contrast, general motion synthesis models such as MDM [64], StableMoFusion [26], and SceneDiffuser [25] lack effective in-betweening capabilities, particularly in keyframe alignment, as they primarily focus on modeling motion distribution rather than in-betweening functionality. Wang et al. [66] produce lower

| Method | Collision Frame Ratio ↓ | Collision Vertex Ratio ↓ | Pene. Max (m) ↓ |
|----------------------|-------------------------|--------------------------|---------------------|
| Ours w/o scene-aware | 0.237 | 0.035 | 0.056 |
| Ours | 0.162 | 0.020 | 0.047 |

Table 2. Quantitative evaluation on the *close-proximity* human-scene interaction frames from the TRUMANS [30].

motion quality (FID of 3.243) due to their reliance on CVAE and global point cloud features, which have limited expressivity. OmniControl[74] struggles to effectively incorporate scene information, resulting in observable body-scene collisions, as shown in Figure 3. Similarly, CondMDI [10] exhibits inferior performance in both motion quality (FID of 0.943) and keyframe alignment (see Fig. 3), attributable to its dependence on motion velocity features as model inputs, which are rarely available in practical applications.

Ablation Analysis. We compared our model to ablated variants that exclude global scene features c_g , local scene features c_l , or both (i.e., “w/o scene awareness”). We also evaluated performance on close-proximity interaction frames where the minimal distance between human and scene is less than 8 cm, with results shown in Tab. 2. From Tab. 1 and Tab. 2, we observed that scene-awareness effectively reduces the collision ratio from 0.237 to 0.162 in interaction frames and enhances alignment with ground-truth motions by over

| Method | FID ↓ | Foot Skating ↓ | Jerk (m/s^3) ↓ | MJPE Key (m) ↓ | MJPE All (m) ↓ | Collision Frame Ratio ↓ | Pene. Max (m) ↓ |
|---------------------------------|--------------|----------------|--------------------|--------------------|--------------------|-------------------------|---------------------|
| MDM [64] | 5.149 | 0.761 | 22.169 | 0.285 | 0.279 | 0.151 | 0.056 |
| OmniControl [74] | 2.981 | 0.381 | 2.198 | 0.302 | 0.308 | 0.169 | 0.058 |
| CondMDI [10] | 3.136 | 0.317 | 0.296 | 0.354 | 0.349 | 0.187 | 0.059 |
| Ours ($T^* = 20$) | 0.118 | 0.247 | 0.198 | 0.013 | 0.012 | 0.108 | 0.042 |
| $T^* = 0$ (w/o noise-awareness) | 0.157 | 0.265 | 0.230 | 0.015 | 0.014 | 0.119 | 0.046 |
| $T^* = 10$ | 0.123 | 0.253 | 0.199 | <u>0.013</u> | <u>0.013</u> | <u>0.110</u> | <u>0.044</u> |
| $T^* = 40$ | <u>0.121</u> | <u>0.249</u> | 0.187 | 0.014 | 0.014 | 0.112 | 0.045 |
| $T^* = 60$ | 0.122 | 0.250 | <u>0.189</u> | 0.015 | 0.015 | 0.114 | 0.045 |

Table 3. Quantitative scene-aware motion in-betweening results TRUMANS dataset [30] with **synthetic noise**. keyframes are provided, using an interval of $r = 3$ and a noise level of $l = 1$. We intentionally used dense keyframes, which presents a more challenging scenario than dealing with sparse, noisy keyframes. **Bold** represents the best value, and underlined represents the second-best.

50% (MJPE from 0.059 to 0.023 m). This observation is consistent with the visualized generation results in Fig. 3. Both global and local scene features enhance motion quality and alignment, with global features playing more prominent roles. Refer to the supplementary video for more results.

4.2. Robust Generalization to Noisy Data

Next, we demonstrated SceneMI’s ability to generalize from synthetic to real-world noisy data. To enhance our model’s robustness to noise, we trained exclusively on TRUMANS by introducing synthetic noise and applying a diffusion process divided into two ranges: $[T, T^* + 1]$ and $[T^*, 1]$, with $T^* = 20$, as described in Section 3.2.1.

We first evaluated SceneMI’s performance on the noisy TRUMANS test set, where synthetic noise was added to keyframes. Then, we extended our evaluation to the real-world GIMO dataset [90], which contains natural noise characteristics arising from inaccurate motion capture sensors. Additionally, GIMO’s scenes are scanned by smartphone, which significantly differs from the synthetic scenes in TRUMANS that our model was trained on.

Evaluation on TRUMANS with Synthetic Noise. We employed a dense keyframe interval of 3 frames, with a maximum noise level $l = 1$ for evaluation. Dense, noisy keyframes were intentionally used for motion in-betweening, which typically presents a more challenging scenario than in-betweening with sparse, noisy keyframes. Table 3 presents the quantitative results. Our model consistently outperforms baseline models across various metrics. Baseline models that directly use noisy keyframes for motion in-betweening face degraded motion quality (e.g., MDM with an FID of 5.149), even with denser keyframes compared to the noise-free setting in Table 1. In comparison, our noise-aware approach steadily produces high-quality motions (e.g., FID of 0.118) while respecting both keyframes (MJPE of 0.012 m) and environmental constraints.

In addition, we validate our noise-aware design in Section 3.2.1, and experiment with different T^* values, as shown in Table 3. Noise awareness ($T^* \neq 0$) consistently improves



Figure 4. SceneMI can perform motion in-betweening on the real-world GIMO dataset [90], while maintaining the original semantics in a scene-aware manner.

motion in-betweening performance on noisy keyframes. When setting $T^* = 20$, our approach produces motions with FID improved from 0.157 to 0.118, and jerk reduced from 0.230 to 0.198. Generally, a larger T^* yields smoother motion with reduced jerk; however, it also increases keyframe error due to the randomness introduced by the diffusion sampling process in the range $[T^*, 1]$. We find $T^* = 20$ as a sweet spot value, synthesizing high-quality motion while preserving keyframe accuracy.

Generalization on Real-World HSI. To emphasize the robust and generalizable performance of SceneMI under noisy conditions, we evaluate our algorithm on the real-world GIMO dataset. Figure 4 presents qualitative visualizations, showing SceneMI generating in-betweening walking motions around indoor tables and chairs. Table 4 reports quantitative results for both baseline models and our approach, including ablated versions. We omit keyframe MJPE and FID scores due to the unavailability of ground-truth motion. Though exclusively trained on the TRUMANS dataset, SceneMI demonstrates strong generalizability to the noisy real-world keyframes in GIMO, outperforming baselines by a large margin. Compared to the original GIMO dataset, the in-betweening motions from SceneMI exhibit significantly reduced *foot skating* (from 0.261 to 0.163) and *jittering* (from 0.573 to 0.249), demonstrating SceneMI’s capability

| Method | Foot Skating ↓ | Accel (m/s^2) ↓ | Jerk (m/s^3) ↓ | Collision Frame Ratio ↓ | Pene Max (m) ↓ |
|-----------------------------|----------------|---------------------|--------------------|-------------------------|--------------------|
| Real-World Data (GIMO [90]) | 0.261 | 0.347 | 0.573 | 0.057 | 0.048 |
| MDM [64] | 0.329 | 0.491 | 0.828 | 0.132 | 0.114 |
| OmniControl [74] | 0.301 | 0.413 | 0.624 | 0.102 | 0.092 |
| CondMDI [10] | 0.312 | 0.359 | 0.498 | 0.091 | 0.083 |
| Ours (full) | 0.163 | <u>0.165</u> | <u>0.249</u> | 0.060 | 0.047 |
| w/o scene awareness | 0.192 | 0.163 | 0.245 | 0.082 | 0.051 |
| w/o noise awareness | 0.391 | 0.215 | 0.301 | 0.072 | 0.049 |

Table 4. Quantitative evaluation on **real-world** GIMO [90], which naturally contains noise arising from acquisition equipment, using an interval of $r = 15$. Through motion in-betweening, our method demonstrates the ability to reduce *foot skating* and *jerk* that are prevalent in the original motion data. *Noise awareness* plays a key role in improving motion quality while *scene-awareness* effectively reduces collisions. **Bold** represents the best value, and underlined represents the second-best.

to improve imperfect HSI data.

Ablation analysis provides further insights into how scene-awareness and noise-awareness contribute to performance. Our results confirm our design motivation: *noise awareness* primarily enhances motion quality, while *scene awareness* effectively reduces the incidence of collision.

4.3. Application: Video-based HSI Reconstruction

We highlight SceneMI’s applicability and generalizability by applying it to reconstructed geometry and estimated key poses from real-world monocular videos. Given one frame from a monocular video, we first segment each instance [9] and use off-the-shelf image-to-3D techniques [75, 78] to reconstruct individual objects. This is followed by estimating depths [77] and camera parameters [32] for coordinate calibration to properly place each object according to the layout in the video frame. We then recover the human mesh from the video stream [13] and place the motion sequence in the reconstructed scene according to camera parameters.

Initially, only a partial scene can be reconstructed. Motions also contain severe jitters and foot skating due to occlusion and depth ambiguity. Moreover, since scene geometry and motions are obtained independently, they exhibit serious collisions, as shown in Fig. 5 (top). To refine the results, we uniformly select keyframes with 15-frame intervals and apply SceneMI to re-generate the motion sequence within synthetic scenes through keyframe in-betweening. The post-processed motions are much smoother, containing fewer motion artifacts and maintaining more plausible interactions with scenes, as shown in Fig. 5 (bottom). Additionally, we employ an autoregressive sampling strategy to synthesize *long-duration* sequences. By iteratively using the final $v = 60$ frames of a prior episode as initial keyframes for the subsequent segment, we synthesize continuous motion sequences across a 23-second video (30 FPS, 690 frames), as demonstrated in the supplemental videos.

To our knowledge, we are the first to build a complete pipeline for 3D HSI reconstruction from monocular videos. Since this is not the focus of the current work, we elaborate the detailed procedure in the appendix.




Figure 5. SceneMI can be applied to reconstructed scenes and keyframes from video, facilitating realistic and physically plausible human-scene interaction reconstruction from monocular video.

5. Conclusions

We propose SceneMI, a scene-aware motion in-betweening framework for modeling human-scene interactions (HSI). SceneMI synthesizes smooth and natural human motions by processing both clean and noisy keyframes within surrounding geometric contexts. This capability enables practical applications including keyframe-guided character animation in 3D scenes and enhances motion quality from imperfect HSI data captured via noisy sensors or video reconstructions. Our experiments demonstrate that SceneMI not only excels in classical in-betweening scenarios with scene constraints but also robustly generalizes to real-world GIMO data. Additionally, we showcase SceneMI’s versatility through a new HSI reconstruction pipeline for monocular video.

Limitations. SceneMI can be improved in several aspects: (i) it relies on full-pose keyframes, which limits flexibility when only partial poses are available; (ii) scene-aware motion in-betweening conditioned on text input could enhance controllability; and (iii) SceneMI models human-scene interactions primarily through feature-level fusion, potentially restricting the expressivity. Future work should consider model-level fusion modeling to address these limitations.

Acknowledgements We sincerely thank the Computational Imaging team at Snap Research NYC, Snap Inc., for their generous support and valuable collaboration. This work was partly supported by IITP grant funded by the Korea government (MSIT) [NO.RS-2021-II211343, Artificial Intelligence Graduate School Program (Seoul National University)] and BK21 FOUR program of the Education and Research Program for Future ICT Pioneers, Seoul National University in 2025.

References

- [1] Simon Alexanderson, Rajmund Nagy, Jonas Beskow, and Gustav Eje Henter. Listen, denoise, action! audio-driven motion synthesis with diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–20, 2023. 2
- [2] Okan Arikan and David A Forsyth. Interactive motion generation from examples. *ACM Transactions on Graphics (TOG)*, 21(3):483–490, 2002. 3
- [3] Jinseok Bae, Jungdam Won, Donggeun Lim, Cheol-Hui Min, and Young Min Kim. Pmp: Learning to physically interact with environments using part-wise motion priors, 2023. 2
- [4] Jinseok Bae, Inwoo Hwang, Young Yoon Lee, Ziyu Guo, Joseph Liu, Yizhak Ben-Shabat, Young Min Kim, and Mubbasir Kapadia. Less is more: Improving motion diffusion models with sparse keyframes, 2025. 2
- [5] Philippe Beaudoin, Stelian Coros, Michiel Van de Panne, and Pierre Poulin. Motion-motif graphs. In *Proceedings of the 2008 ACM SIGGRAPH/Eurographics symposium on computer animation*, pages 117–126, 2008. 3
- [6] Bharat Lal Bhatnagar, Xianghui Xie, Ilya Petrov, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Behave: Dataset and method for tracking human object interactions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2022. 2
- [7] Zhi Cen, Huaijin Pi, Sida Peng, Zehong Shen, Minghui Yang, Zhu Shuai, Hujun Bao, and Xiaowei Zhou. Generating human motion in 3d scenes from text descriptions. In *CVPR*, 2024. 2
- [8] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18000–18010, 2023. 2
- [9] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. 2022. 8
- [10] Setareh Cohan, Guy Tevet, Daniele Reda, Xue Bin Peng, and Michiel van de Panne. Flexible motion in-betweening with diffusion models. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–9, 2024. 1, 3, 5, 6, 7, 8
- [11] Yinglin Duan, Tianyang Shi, Zhengxia Zou, Yanan Lin, Zhehui Qian, Bohan Zhang, and Yi Yuan. Single-shot motion completion with transformer. *arXiv preprint arXiv:2103.00776*, 2021. 3
- [12] Anindita Ghosh, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Philipp Slusallek. Imos: Intent-driven full-body motion synthesis for human-object interactions. In *Eurographics*, 2023. 2
- [13] Shubham Goel, Georgios Pavlakos, Jathushan Rajasegaran, Angjoo Kanazawa*, and Jitendra Malik*. Humans in 4D: Reconstructing and tracking humans with transformers. In *International Conference on Computer Vision (ICCV)*, 2023. 2, 8
- [14] Kehong Gong, Dongze Lian, Heng Chang, Chuan Guo, Zihang Jiang, Xinxin Zuo, Michael Bi Mi, and Xinchao Wang. Tm2d: Bimodality driven 3d dance generation via music-text integration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9942–9952, 2023. 2
- [15] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2021–2029, 2020. 2
- [16] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5152–5161, 2022. 2
- [17] Chuan Guo, Xinxin Zuo, Sen Wang, and Li Cheng. Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *ECCV*, 2022.
- [18] Chuan Guo, Yuxuan Mu, Muhammad Gohar Javed, Sen Wang, and Li Cheng. Momask: Generative masked modeling of 3d human motions. 2023.
- [19] Chuan Guo, Inwoo Hwang, Jian Wang, and Bing Zhou. Snapmogen: Human motion generation from expressive texts, 2025. 2
- [20] Félix G Harvey and Christopher Pal. Recurrent transition networks for character locomotion. In *SIGGRAPH Asia 2018 Technical Briefs*, pages 1–4. 2018. 1, 3
- [21] Félix G Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion in-betweening. *ACM Transactions on Graphics (TOG)*, 39(4):60–1, 2020. 1, 3
- [22] Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J. Black. Resolving 3D human pose ambiguities with 3D scene constraints. In *International Conference on Computer Vision*, 2019. 5
- [23] Mohamed Hassan, Duygu Ceylan, Ruben Villegas, Jun Saito, Jimei Yang, Yi Zhou, and Michael J Black. Stochastic scene-aware motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11374–11384, 2021. 1, 2
- [24] Nhat M. Hoang, Kehong Gong, Chuan Guo, and Michael Bi Mi. Motionmix: Weakly-supervised diffusion for controllable motion generation, 2024. 2, 5
- [25] Siyuan Huang, Zan Wang, Puhao Li, Baoxiong Jia, Tengyu Liu, Yixin Zhu, Wei Liang, and Song-Chun Zhu. Diffusion-based generation, optimization, and planning in 3d scenes. *arXiv preprint arXiv:2301.06015*, 2023. 2, 5, 6
- [26] Yiheng Huang, Yang Hui, Chuanchen Luo, Yuxi Wang, Shibiao Xu, Zhaoxiang Zhang, Man Zhang, and Junran Peng. Stablemofusion: Towards robust and efficient diffusion-based motion generation framework. *arXiv preprint arXiv:2405.05691*, 2024. 5, 6

- [27] Inwoo Hwang, Jinseok Bae, Donggeun Lim, and Young Min Kim. Goal-driven human motion synthesis in diverse task. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR) Workshops*, pages 2920–2930, 2025. 2
- [28] Inwoo Hwang, Jinseok Bae, Donggeun Lim, and Young Min Kim. Motion synthesis with sparse and flexible keyjoint control, 2025. 2
- [29] Nan Jiang, Zimo He, Zi Wang, Hongjie Li, Yixin Chen, Siyuan Huang, and Yixin Zhu. Autonomous character-scene interaction synthesis from text instruction, 2024. 1, 2
- [30] Nan Jiang, Zhiyuan Zhang, Hongjie Li, Xiaoxuan Ma, Zan Wang, Yixin Chen, Tengyu Liu, Yixin Zhu, and Siyuan Huang. Scaling up dynamic human-scene interaction modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1737–1747, 2024. 1, 2, 3, 5, 6, 7
- [31] Nan Jiang, Zhiyuan Zhang, Hongjie Li, Xiaoxuan Ma, Zan Wang, Yixin Chen, Tengyu Liu, Yixin Zhu, and Siyuan Huang. Scaling up dynamic human-scene interaction modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1737–1747, 2024. 2
- [32] Linyi Jin, Jianming Zhang, Yannick Hold-Geoffroy, Oliver Wang, Kevin Matzen, Matthew Sticha, and David F. Fouhey. Perspective fields for single image camera calibration. In *CVPR*, 2023. 8
- [33] Korrawe Karunratanakul, Konpat Preechakul, Supasorn Suwanakorn, and Siyu Tang. Guided motion diffusion for controllable human motion synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2151–2162, 2023. 4
- [34] Jihoon Kim, Taehyun Byun, Seungyou Shin, Jungdam Won, and Sungjoon Choi. Conditional motion in-betweening. *Pattern Recognition*, 132:108894, 2022. 3
- [35] Lucas Kovar, Michael Gleicher, and Frédéric Pighin. Motion graphs. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 723–732. 2023. 3
- [36] Nilesh Kulkarni, Davis Rempe, Kyle Genova, Abhijit Kundu, Justin Johnson, David Fouhey, and Leonidas Guibas. Nifty: Neural object interaction fields for guided human motion synthesis, 2023. 2
- [37] Hsin-Ying Lee, Xiaodong Yang, Ming-Yu Liu, Ting-Chun Wang, Yu-Ding Lu, Ming-Hsuan Yang, and Jan Kautz. Dancing to music. *Advances in neural information processing systems*, 32, 2019. 2
- [38] Jiye Lee and Hanbyul Joo. Locomotion-action-manipulation: Synthesizing human-scene interactions in complex 3d environments. *arXiv preprint arXiv:2301.02667*, 2023. 2
- [39] Andreas M. Lehrmann, Peter V. Gehler, and Sebastian Nowozin. Efficient nonlinear markov models for human motion. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1314–1321, 2014. 3
- [40] Jiaman Li, Alexander Clegg, Roozbeh Mottaghi, Jiajun Wu, Xavier Puig, and C. Karen Liu. Controllable human-object interaction synthesis, 2023. 2
- [41] Jiaman Li, Jiajun Wu, and C Karen Liu. Object motion guided human motion synthesis. *ACM Trans. Graph.*, 42(6), 2023. 2
- [42] Ruilong Li, Shan Yang, David A. Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++, 2021. 2
- [43] Donggeun Lim, Cheongi Jeong, and Young Min Kim. Mamos: Mapping multiple human motion with scene understanding and natural interactions. In *2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 4280–4289, 2023. 1, 2
- [44] J. Lin, J. Chang, L. Liu, G. Li, L. Lin, Q. Tian, and C. Chen. Being comes from not-being: Open-vocabulary text-to-motion generation with wordless training. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [45] Xinpeng Liu, Haowen Hou, Yanchao Yang, Yong-Lu Li, and Cewu Lu. Revisit human-scene interaction via space occupancy. *arXiv preprint arXiv:2312.02700*, 2023. 1
- [46] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015. 3
- [47] Zhenyu Lou, Qiongjie Cui, Haofan Wang, Xu Tang, and Hong Zhou. Multimodal sense-informed prediction of 3d human motions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [48] Zhenyu Lou, Qiongjie Cui, Tuo Wang, Zhenbo Song, Luoming Zhang, Cheng Cheng, Haofan Wang, Xu Tang, Huaxia Li, and Hong Zhou. Harmonizing stochasticity and determinism: Scene-responsive diverse human motion prediction. In *Advances in Neural Information Processing Systems*, pages 39784–39811. Curran Associates, Inc., 2024. 2
- [49] Tomohiko Mukai and Shigeru Kuriyama. Geostatistical motion interpolation. *ACM Trans. Graph.*, 24(3):1062–1070, 2005. 3
- [50] Boris N Oreshkin, Antonios Valkanas, Félix G Harvey, Louis-Simon Ménard, Florent Bouquet, and Mark J Coates. Motion in-betweening via deep δ -interpolator. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 3
- [51] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. 5
- [52] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10985–10995, 2021. 2
- [53] Mathis Petrovich, Michael J. Black, and Gül Varol. TEMOS: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision (ECCV)*, 2022. 2
- [54] Mathis Petrovich, Or Litany, Umar Iqbal, Michael J. Black, Gül Varol, Xue Bin Peng, and Davis Rempe. STMC: Multi-track timeline control for text-driven 3d human motion generation. *arXiv:2401.08559*, 2024. 2

- [55] Ekkasit Pinyoanuntapong, Muhammad Usama Saleem, Korrave Karunratanakul, Pu Wang, Hongfei Xue, Chen Chen, Chuan Guo, Junli Cao, Jian Ren, and Sergey Tulyakov. Controlmm: Controllable masked motion generation. *arXiv preprint arXiv:2410.10780*, 2024. 2
- [56] Sergey Prokudin, Christoph Lassner, and Javier Romero. Efficient learning on point clouds with basis point sets. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4332–4341, 2019. 3
- [57] Qiaosong Qi, Le Zhuo, Aixi Zhang, Yue Liao, Fei Fang, Si Liu, and Shuicheng Yan. Diffdance: Cascaded human motion diffusion model for dance generation. In *Proceedings of the 31st ACM International Conference on Multimedia*, page 1374–1382, New York, NY, USA, 2023. Association for Computing Machinery. 2
- [58] Jia Qin, Youyi Zheng, and Kun Zhou. Motion in-betweening via two-stage transformers. *ACM Trans. Graph.*, 41(6):184–1, 2022. 1, 3
- [59] C. Rose, M.F. Cohen, and B. Bodenheimer. Verbs and adverbs: multidimensional motion interpolation. *IEEE Computer Graphics and Applications*, 18(5):32–40, 1998. 3
- [60] Yonatan Shafir, Guy Tevet, Roy Kapon, and Amit H Bermano. Human motion diffusion as a generative prior. *arXiv preprint arXiv:2303.01418*, 2023. 2
- [61] Li Siyao, Weijiang Yu, Tianpei Gu, Chunze Lin, Quan Wang, Chen Qian, Chen Change Loy, and Ziwei Liu. Bailando: 3d dance generation by actor-critic gpt with choreographic memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11050–11059, 2022. 2
- [62] Kewei Sui, Anindita Ghosh, Inwoo Hwang, Bing Zhou, Jian Wang, and Chuan Guo. A survey on human interaction motion generation, 2025. 2
- [63] Omid Taheri, Nima Ghorbani, Michael J Black, and Dimitrios Tzionas. Grab: A dataset of whole-body human grasping of objects. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 581–600. Springer, 2020. 2
- [64] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 5, 6, 7, 8
- [65] Jonathan Tseng, Rodrigo Castellon, and C Karen Liu. Edge: Editable dance generation from music. *arXiv preprint arXiv:2211.10658*, 2022. 2
- [66] Jiashun Wang, Huazhe Xu, Jingwei Xu, Sifei Liu, and Xiaolong Wang. Synthesizing long-term 3d human motion and interaction in 3d scenes. *arXiv preprint arXiv:2012.05522*, 2020. 5, 6
- [67] Jiashun Wang, Huazhe Xu, Jingwei Xu, Sifei Liu, and Xiaolong Wang. Synthesizing long-term 3d human motion and interaction in 3d scenes, 2020. 1, 2
- [68] Jiashun Wang, Huazhe Xu, Jingwei Xu, Sifei Liu, and Xiaolong Wang. Synthesizing long-term 3d human motion and interaction in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9401–9411, 2021. 2
- [69] Jingbo Wang, Sijie Yan, Bo Dai, and Dahua Lin. Scene-aware generative network for human motion synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12206–12215, 2021.
- [70] Jingbo Wang, Yu Rong, Jingyuan Liu, Sijie Yan, Dahua Lin, and Bo Dai. Towards diverse and natural scene-aware 3d human motion synthesis, 2022. 1, 2
- [71] Zan Wang, Yixin Chen, Tengyu Liu, Yixin Zhu, Wei Liang, and Siyuan Huang. Humanise: Language-conditioned human motion generation in 3d scenes. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 1, 2
- [72] Zan Wang, Yixin Chen, Baoxiong Jia, Puhao Li, Jinlu Zhang, Jingze Zhang, Tengyu Liu, Yixin Zhu, Wei Liang, and Siyuan Huang. Move as you say, interact as you can: Language-guided human motion generation with scene affordance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [73] Zeqi Xiao, Tai Wang, Jingbo Wang, Jinkun Cao, Wenwei Zhang, Bo Dai, Dahua Lin, and Jiangmiao Pang. Unified human-scene interaction via prompted chain-of-contacts. *arXiv preprint arXiv:2309.07918*, 2023. 2
- [74] Yiming Xie, Varun Jampani, Lei Zhong, Deqing Sun, and Huaizu Jiang. Omnicontrol: Control any joint at any time for human motion generation. *arXiv preprint arXiv:2310.08580*, 2023. 2, 5, 6, 7, 8
- [75] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*, 2024. 2, 8
- [76] Sirui Xu, Zhengyuan Li, Yu-Xiong Wang, and Liang-Yan Gui. Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In *ICCV*, 2023. 2
- [77] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *CVPR*, 2024. 8
- [78] Xianghui Yang, Huiwen Shi, Bowen Zhang, Fan Yang, Jiacheng Wang, Hongxu Zhao, Xinhai Liu, Xinzhou Wang, Qingxiang Lin, Jiaao Yu, Lifu Wang, Zhuo Chen, Sicong Liu, Yuhong Liu, Yong Yang, Di Wang, Jie Jiang, and Chunchao Guo. Tencent hunyuan3d-1.0: A unified framework for text-to-3d and image-to-3d generation, 2024. 2, 8
- [79] Hongwei Yi, Justus Thies, Michael J. Black, Xue Bin Peng, and Davis Rempel. Generating human interaction motions in scenes with text control. *arXiv:2404.10685*, 2024. 1, 2
- [80] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. PhysDiff: Physics-guided human motion diffusion model. In *IEEE International Conference on Computer Vision (ICCV)*, 2023. 2
- [81] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Shaoli Huang, Yong Zhang, Hongwei Zhao, Hongtao Lu, and Xi Shen. T2m-gpt: Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [82] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In

- Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. [2](#)
- [83] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. [2](#)
- [84] Siwei Zhang, Bharat Lal Bhatnagar, Yuanlu Xu, Alexander Winkler, Petr Kadlecek, Siyu Tang, and Federica Bogo. Rohm: Robust human motion reconstruction via diffusion. In *CVPR*, 2024. [5](#)
- [85] Xinyi Zhang and Michiel Van De Panne. Data-driven auto-completion for keyframe animation. In *Proceedings of the 11th ACM SIGGRAPH Conference on Motion, Interaction and Games*, pages 1–11, 2018. [1](#), [3](#)
- [86] Yan Zhang and Siyu Tang. The wanderings of odysseus in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20481–20491, 2022. [2](#)
- [87] Kaifeng Zhao, Shaofei Wang, Yan Zhang, Thabo Beeler, , and Siyu Tang. Compositional human-scene interaction synthesis with semantic control. In *European conference on computer vision (ECCV)*, 2022. [2](#)
- [88] Kaifeng Zhao, Yan Zhang, Shaofei Wang, Thabo Beeler, , and Siyu Tang. Synthesizing diverse human motions in 3d indoor scenes. In *International conference on computer vision (ICCV)*, 2023. [2](#)
- [89] Kaifeng Zhao, Yan Zhang, Shaofei Wang, Thabo Beeler, and Siyu Tang. Synthesizing diverse human motions in 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14738–14749, 2023. [2](#)
- [90] Yang Zheng, Yanchao Yang, Kaichun Mo, Jiaman Li, Tao Yu, Yebin Liu, C Karen Liu, and Leonidas J Guibas. Gimo: Gaze-informed human motion prediction in context. In *European Conference on Computer Vision*, pages 676–694. Springer, 2022. [1](#), [2](#), [5](#), [7](#), [8](#)