

A Unified Diffusion Framework for Scene-aware Human Motion Estimation from Sparse Signals

Jiangnan Tang^{1,2} Jingya Wang^{1,2,*} Kaiyang Ji^{1,2} Lan Xu^{1,2} Jingyi Yu^{1,2} Ye Shi^{1,2,*}

¹ShanghaiTech University, China

²MoE Key Laboratory of Intelligent Perception and Human-Machine Collaboration, China

{tangjn2022, wangjingya, jiky, xulan1, yujingyi, shiye}@shanghaitech.edu.cn

Abstract

Estimating full-body human motion via sparse tracking signals from head-mounted displays and hand controllers in 3D scenes is crucial to applications in AR/VR. One of the biggest challenges to this task is the one-to-many mapping from sparse observations to dense full-body motions, which endowed inherent ambiguities. To help resolve this ambiguous problem, we introduce a new framework to combine rich contextual information provided by scenes to benefit full-body motion tracking from sparse observations. To estimate plausible human motions given sparse tracking signals and 3D scenes, we develop S^2 Fusion, a unified framework fusing Scene and sparse Signals with a conditional diffFusion model. S^2 Fusion first extracts the spatial-temporal relations residing in the sparse signals via a periodic autoencoder, and then produces time-alignment feature embedding as additional inputs. Subsequently, by drawing initial noisy motion from a pre-trained prior, S^2 Fusion utilizes conditional diffusion to fuse scene geometry and sparse tracking signals to generate full-body scene-aware motions. The sampling procedure of S^2 Fusion is further guided by a specially designed scene-penetration loss and phase-matching loss, which effectively regularizes the motion of the lower body even in the absence of any tracking signals, making the generated motion much more plausible and coherent. Extensive experimental results have demonstrated that our S^2 Fusion outperforms the state-of-the-art in terms of estimation quality and smoothness. Code is available at <https://github.com/jn-tang/S2Fusion>.

1. Introduction

With the emergence of advanced AR/VR technologies, there is an increasing demand for generating realistic human avatars in applications such as virtual conferencing and gaming. However, common AR/VR devices, e.g., HTC Vive and

Meta Quest Pro, provide only sparse tracking signals from inertial measurement units (IMU) embedded in single head-mounted displays (HMD) and hand controllers. Using sparse tracking signals to generate dense full-body motions is a one-to-many mapping with inherent ambiguities, making this problem a challenging task.

A natural method to generate full-body motion is leveraging data-driven methods that utilize both large-scale motion capture data [40] and sparse tracking signals by AR/VR devices. Recently, various data-driven models – ranging from simple regression-based methods [2, 14, 28, 63, 76] to probabilistic generative models such as VAE [16], normalizing flow [1] and diffusion model [17] – are deployed. Although these methods show various ways of finding the most probable human motion estimation, they still fail to narrow down the distribution of possible motion space, leaving the crux of one-to-many ambiguity unresolved. Existing methods [1, 14, 16, 17, 28, 60, 64] also failed to pose constraints on the pose of the lower body, making the generated motion disastrous as the legs are free to penetrate through scene geometries; the generated motion of legs may also uncorrelated with hands motion, resulting in the moving of hands and feet on the same side. The incurring of implausible leg motions is due to the lack of observations on the lower body.

Since human motion is highly related to the surrounding environments, introducing the scene modality can greatly reduce the ambiguities in estimating full-body motion from sparse tracking signals. The rich contextual information provides valuable cues to infer the lower body motion, even though no direct observations are available. Therefore, we propose to combine the scene information with sparse tracking signals for human motion estimation. To handle the uncorrelated issues between the generated motion of legs and hands, we use the temporal movement pattern presented in the upper body observations to control the generation of lower body motions.

To estimate plausible human motions given sparse tracking signals and 3D scenes, we develop S^2 Fusion, a unified framework fusing Scene and sparse Signals with a condi-

*Corresponding author




Figure 1. Given sparse tracking signals from only the head and left/right hands, our method accurately estimates full-body motion in the 3D scene.

tional *diffFusion* model. S^2 Fusion models this specific task by a conditional diffusion model, in light of the recent success of diffusion models in motion generation [9, 51]. Diffusion models support flexible control on generating samples, by incorporating various loss functions in the loss-guided sampling process [11, 46]. To tackle the lack of paired motion-scene datasets and to generate more diverse motion, the reverse diffusion process starts from a non-Gaussian motion distribution, by adopting a pre-trained motion prior on a large-scale motion dataset [40]. To coordinate the human motion in 3D scenes, we extract the periodic motion feature of the sparse tracking signals using a periodic autoencoder [49]; the periodic motion feature represents the alignment of full-body motions in time and space, resulting in more effective positional embeddings. To facilitate the mitigating of unrealistic lower body motions, we guide the diffusion sampling process by the gradient of our specially designed loss functions, under the framework of loss-guided sampling [46]. Our loss functions include scene-penetration loss and phase-matching loss, which regularizes the lower body motion *even in the absence of any tracking signals*.

In summary, our paper makes the following contributions:

- We introduce a new framework that combines scene information and sparse tracking signals for human motion estimation to greatly reduce the inherent ambiguities in estimating full-body motion. To further enhance human motion estimation, we propose to extract periodic motion features from sparse tracking signals to improve the coordination between the upper and lower body.
- We develop a unified diffusion method, i.e., S^2 Fusion, tailored for the scene-aware human motion estimation with sparse signals. S^2 Fusion integrates three main components: 1) a motion prior that provides initial value for the reverse diffusion process, which significantly improves generation quality and inference speed; 2) a periodic motion feature extractor that learns the spatial-temporal alignment of input signals varies in unit and scale; the extracted feature is then combined with sparse

tracking signals and scene geometry as conditional inputs, fed to motion generation process; 3) a set of specially designed loss functions, which effectively introduce regularization on the motion of the lower body during the loss-guided diffusion sampling process.

- We conduct extensive experiments for the scene-aware human motion estimation and verify S^2 Fusion can successfully reduce the inherent ambiguities posed by the sparse-to-dense problem in contrast to other pure sensor-based methods and achieve a notable performance boost. Moreover, comprehensive ablation studies also demonstrate the effectiveness of the developed three main components of S^2 Fusion.

2. Related Work

2.1. Motion Estimation from Sparse Sensors

There is a surging interest in studying reconstructing human motion from sparse sensor inputs, as it is not only economical compared to traditional marker-based optical solutions or vision-based solutions [69, 71], but also suits the needs of AR/VR applications. A major problem in reconstructing full-body motions where only sparse observations are available is the ambiguities inherent in this one-to-many mapping problem. One line of work reconstructs full-body motion from six IMUs located on the head, left/right wrists, left/right knees, and torso, pioneered by von Marcard et al., [56]. Yi et al. proposed a real-time RNN-based method *TransPose* [65] to predict full-body motions, and their follow-up works [66, 67] refines the predicted motions by physics-based optimization and egocentric SLAM system, respectively. Other backbones are also explored, *DIP* [26] is powered by biRNN, and *TIP* [29] deployed Transformer encoder architecture. Another line of work reconstructs full-body motion from three 6D trackers located on the head and left/right hands, providing accurate translation and rotation measurements. [1, 16, 17] tried various generative probabilistic models to solve the unconstrained problem, while [2, 28, 76] used Transformer backbone to regress full-body motion directly. Our method fits this line of research of estimating full-body motion from three 6D trackers. In addition, our approach brings scene information to resolve the inherent ambiguities induced by sparse-to-dense mapping, resulting in notable estimation improvement.

2.2. Diffusion-based Probabilistic Models

Diffusion-based probabilistic models (DPMs) [22, 47] have recently shown promising results in image generation [13], video generation [18, 23], and speech synthesis [33], demonstrating their powerful probabilistic modeling capabilities. The success of DPMs is mainly attributed to their ability to support versatile conditioning and are highly controllable. Increasing interest grows to further control the genera-

tion process of DPMs by adjusting the denoising trajectory, such as classifier guidance [13], imputation and inpainting [10, 11, 48, 79], which can be collectively termed as *loss-guided diffusion* (LGD) [46]. The loss-guided diffusion enables flexible conditioning without retraining existing models, effectively rendering DPMs a powerful generation tool. Our S²Fusion is a specifically designed diffusion model for scene-aware human motion estimation from sparse signals. It differentiates from standard diffusion models mainly in three components: 1) a motion prior module for diffusion initialization; 2) a comprehensive condition acquisition module; and 3) a set of specifically designed loss functions in the diffusion sampling stage.

2.3. Human Motion Generation

Human motion can be generated by any signal that describes the motion, including text [4, 9, 20, 31, 44, 51, 52, 61, 68, 72], audio [35, 54], action labels [9, 19, 42, 51], or unconditioned [44, 74, 75]. With increasingly available scene-motion datasets [6, 21, 24, 50, 70], there are shifts in trends to generate human motions in 3D scenes, ranging from holistically generating human motion with scene-awareness [12, 25, 27, 57–59, 73] to fine-grained reconstructing accurate human-object interactions [8, 15, 30, 34, 36, 41, 53, 62]. Our method is the first attempt to fuse scene information with sparse tracking signals to estimate full-body physical-plausible human motion with a sophisticated diffusion model.

2.4. Motion Frequency Analysis

The characteristic of motion in the frequency domain has been utilized for motion synthesis [37], editing [7, 32], stylization [55], and compression [5]. The features in the frequency domain present a holistic view of the underlying motion and are consistent over a period of time, hence are a great medium for summarizing the movement pattern. Starke et al. [49] have recently achieved success in using a deep neural network to learn the periodic feature from motion datasets and synthesizing high-quality motions, they demonstrate the power of a new network architecture called periodic autoencoder (PAE) in extracting the periodic behavior of the motions. Shi et al. [45] built upon PAE a conditional VAE, which also leverages phase features in robustly generating motions. Instead of only leveraging the phase feature residing on the frequency domain, our method explores using both the temporal and frequency domain features; the temporal domain features are useful in aligning motions time axis, and the frequency domain features enabled us to correlate the upper and lower body motions.

3. Human Motion Estimation with S²Fusion

It is not trivial to estimate full-body motion from scene geometry and sparse signals, as it is an ill-conditioned one-

to-many problem, and care must be taken to deal with *spatial* scene geometry and *temporal* sparse signals. In light of the recent success of the diffusion model in text-to-motion generation [51], we design a conditional diffusion model to generate full-body poses given scene geometry and sparse tracking signals. However, we observed several issues that hinder the motion generation quality from directly applying the conditional diffusion model. First, due to the limited data volume of the existing motion-scene datasets, the generated motion is less diverse and unrealistic. Second, the tracking signals only come from the upper body, so it is hard to generate correlated upper and lower body motions. To tackle these challenges, we designed a novel diffusion method, i.e., S²Fusion that leverages a pre-trained motion prior and augments scene and sparse tracking signals with periodic alignment features to generate more diverse and realistic motions.

Given a sequence of sparse tracking signals $\mathbf{p}^{1:N} \in \mathbb{R}^{N \times c}$ and scene geometry $\mathbf{S} \in \mathbb{R}^{P \times 3}$, we aim to predict full-body motion $\mathbf{x}^{1:N} \in \mathbb{R}^{N \times n}$. Here c and n respectively denote the dimension of input and output, and P denotes the number of points in the scene point cloud. The popular SMPL [38] model is used to represent human poses and only the first 22 joints of the SMPL model [28] are considered. We use the 6D representation [78] to represent rotations for effectively learning rotation quantities. Leveraging the scene geometry \mathcal{S} , sparse tracking signals $\mathbf{p}^{1:N}$ and the periodic alignment features $\mathbf{f}^{1:N}$, we obtain the final clean motion $\mathbf{x}_0^{1:N}$ by a conditional diffusion model, given initial noisy sample $\tilde{\mathbf{x}}^{1:N}$. We depict our full-body motion generation pipeline in Figure 2.

3.1. Motion Prior for S²Fusion Initialization

To circumvent the limited volume of scene-motion datasets and generate more diverse and realistic motions, we propose to draw initial motion distribution from a pre-trained motion prior for the reverse diffusion process [39]. We build a VAE-based generative model that resembles an encoder-decoder architecture and trained on large-scale motion dataset AMASS [40].

Once the VAE-based generative model is trained, we can sample initial motion unaware of the scene by conditioning on the tracking signals,

$$\tilde{\mathbf{x}}^{1:N} = f_\phi(z, \mathbf{p}^{1:N}), \quad z \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (1)$$

the sampled motion $\tilde{\mathbf{x}}^{1:N}$ is then refined by our conditional diffusion process Sec. 3.3.

Compared to Gaussian noise, the pre-trained generative prior captures the complex structure of the motion manifold, therefore improving the sample quality. Meanwhile, by jump-starting to a noisy motion (as compared to standard DDPM settings which require $T = 1000$ denoising steps [39]), we can also achieve accelerated inference speed.

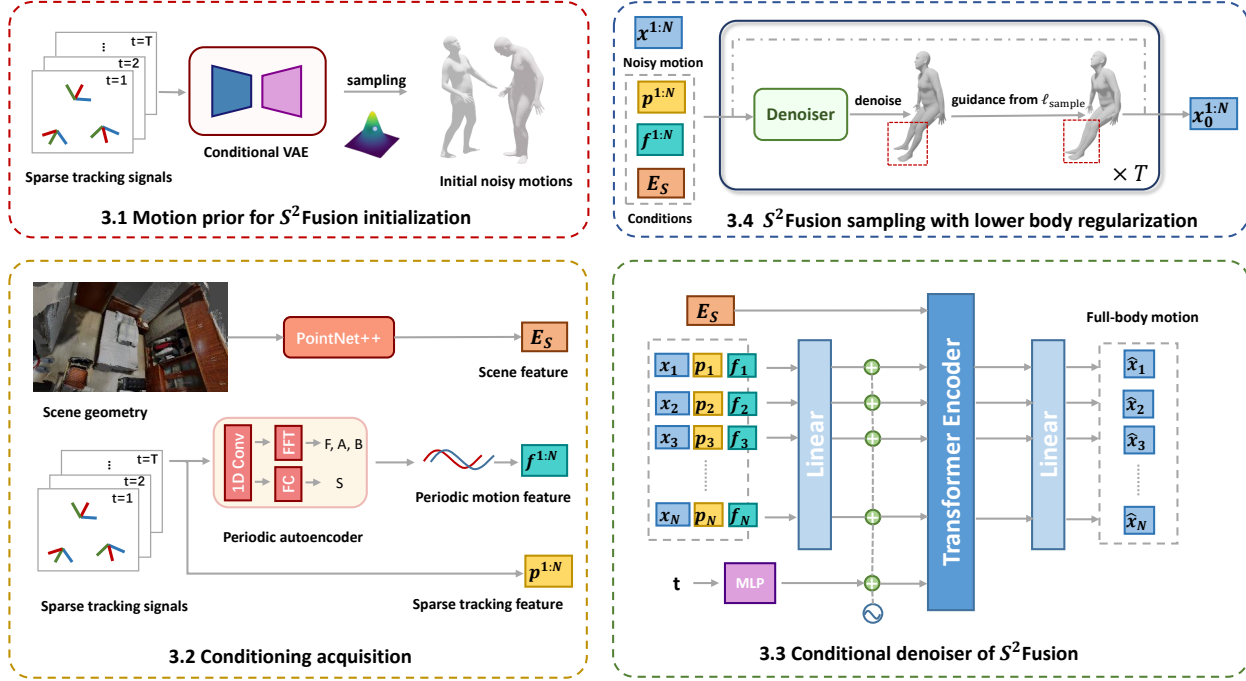


Figure 2. Illustration of S^2 Fusion pipeline. Given the sparse tracking signals $\mathbf{p}^{1:N}$ and scene geometry \mathcal{S} , S^2 Fusion generates full-body motion with scene awareness and coherent upper and lower body movements. (1) The pre-trained motion prior f_ϕ first samples the initial noisy motion $\tilde{\mathbf{x}}^{1:N}$ for the reverse diffusion process; (2) then the periodic motion features $\mathbf{f}^{1:N}$ are extracted by a periodic autoencoder, and combined with encoded scene feature \mathbf{E}_S and the sparse tracking signals $\mathbf{p}^{1:N}$ to form the final conditioning input \mathbf{c} to the reverse diffusion process; (3) the conditional diffusion model predicts the clean motion $\mathbf{x}_0^{1:N}$ from noisy motion $\tilde{\mathbf{x}}^{1:N}$ conditioned on \mathbf{c} ; (4) the diffusion sampling process is further guided by the gradient of $\ell_{\text{penetration}}$ and ℓ_{phase} to generate scene-aware and physically plausible motions.

3.2. Condition Acquisition for S^2 Fusion

The overall conditioning input to our diffusion model is

$$\mathbf{c} = (\mathbf{p}^{1:N}, \mathbf{f}^{1:N}, \mathbf{E}_S), \quad (2)$$

consists of raw tracking signals, periodic alignment features, and encoded scene features. We introduce the acquisition of the conditioning inputs as follows.

Scene conditioning. During human-scene interaction, only a small region around the human provides meaningful contextual information. Hence we crop a $2m \times 2m \times 2m$ bounding box \mathcal{B}_S around the human, given the global translation measured by HMD. The cropped point cloud is then fed into a scene encoder [43] h_θ , obtaining the scene feature vector \mathbf{E}_S .

Periodic motion feature. Though the observed tracking signals consist of varying sources with different units and scales, i.e., rotation from the gyroscope, and position from the infrared optical sensor, the changes in these signals reflect the temporal movement pattern of the underlying motions. As suggested by [49], full-body movements can happen as a composition of multiple local periodic movements, and these movements can be decomposed into multiple latent channels that capture the non-linear periodicity of different body

segments, which are aligned in time. To effectively extract the temporal alignment feature of these signals, we resort to the frequency domain and use a periodic autoencoder (PAE) [49] g_τ to compute the frequency domain parameters frequency \mathbf{F} , amplitude \mathbf{A} , offset \mathbf{B} and phase shift \mathbf{S} . Intuitively, the phase shift indicates the time-alignment of the motions, while the amplitude resembles the momentum.

The PAE g_τ consists of a combination of 1D convolution and a fast Fourier transformation (FFT) layer to compute the amplitudes \mathbf{A} , offsets \mathbf{B} , and frequencies \mathbf{F} of a temporal signal $\mathbf{p}^{1:N}$,

$$[\mathbf{A}, \mathbf{B}, \mathbf{F}] = \text{FFT}(\text{Conv}(\mathbf{p}^{1:N})). \quad (3)$$

The phase-shifts \mathbf{S} are obtained by a separate fully-connected network,

$$(s_x, s_y) = \text{FC}(\text{Conv}(\mathbf{p}^{1:N})), \mathbf{S} = \arctan(s_y, s_x). \quad (4)$$

After extracting parameters in the frequency domain, we can reconstruct a smoothed periodic alignment feature in the temporal domain,

$$\mathbf{f}_t = \mathbf{A} \cdot \sin(2\pi \cdot (\mathbf{F} \cdot t - \mathbf{S})) + \mathbf{B}, \quad (5)$$

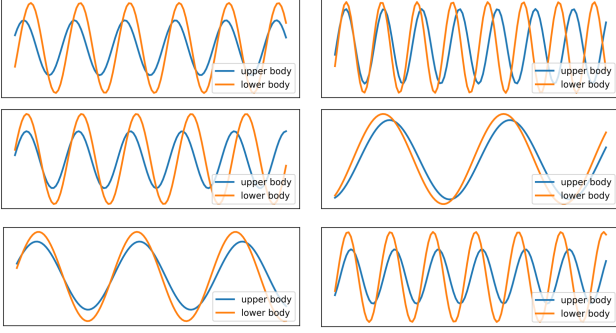


Figure 3. A visualization of the periodic motion features of the upper and lower body extracted from randomly selected motion sequences in AMASS[40]. The phase shift of the sinusoidal functions indicates the time-alignment of the upper and lower body motions, while the amplitude resembles the momentum. It can be shown that the periodic motion features of the upper body are correlated with that of the lower body.

where $t \in \{1, \dots, N\}$. The reconstructed feature in essence is a multi-resolution sinusoidal function, which summarizes the alignment of full-body motion in both time and space.

Sparse tracking signals. We take the position and rotation of the head and left/right hands as input signals, and compute angular and linear velocities as extra input following previous work [28].

3.3. Conditional Denoiser of S²Fusion

Diffusion models consist of a forward diffusion process and a reverse diffusion process. To model a distribution $\mathbf{x}_0 \sim q(\mathbf{x}_0)$, the forward process follows a Markov chain of T steps which progressively adds Gaussian noise and produces a series of time-dependent distributions $q(\mathbf{x}_t|\mathbf{x}_{t-1})$. Formally,

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}), \quad (6)$$

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}), \quad (7)$$

where $\beta_t \in (0, 1)$ is the variance schedule at timestep t .

To generate full-body motion, we model the conditioned motion generation problem by the reverse diffusion process, which gradually cleans corrupted signal \mathbf{x}_T . Instead of predicting residual Gaussian noise ϵ_t added through each step t , we predict the signal $\hat{\mathbf{x}}$ itself following [51] with the *simple objective*,

$$\mathcal{L}_{\text{simple}} = \mathbb{E}_{t \sim [1, T]} \|G(\mathbf{x}_t^{1:N}, t, \mathbf{c}) - \mathbf{x}_0\| \quad (8)$$

This iterative process at timestep t can be formulated as

$$\mathbf{x}_{t-1}^{1:N} = \sqrt{\bar{\alpha}_{t-1}}G(\mathbf{x}_t^{1:N}, t, \mathbf{c}) + \sqrt{1 - \bar{\alpha}_{t-1}}\epsilon, \quad (9)$$

where G is a network that learns to generate clean motion $\hat{\mathbf{x}}^{1:N}$ at timestep t , $\bar{\alpha}_t = \prod_{i=1}^t (1 - \beta_i)$ and $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is the injected Gaussian noise. The denoising network G is trained with $\mathcal{L}_{\text{simple}}$ and geometric loss [51] that regulates the generated motion to lie on motion manifold,

$$\mathcal{L}_{\text{train}} = \mathcal{L}_{\text{simple}} + \mathcal{L}_{\text{geometric}}, \quad (10)$$

$$\mathcal{L}_{\text{geometric}} = \|\text{FK}(\hat{\mathbf{x}}) - \text{FK}(\mathbf{x}_0)\|, \quad (11)$$

where the $\text{FK}(\cdot)$ denotes the forward kinematic process that converts joint rotations into positions.

More details about the training details of the VAE-based motion prior and periodic autoencoder are provided in the supplementary.

3.4. S²Fusion Sampling with Lower Body Motion Regularization

Given the sparse tracking signals of the *upper body*, it is already possible to reconstruct the upper body motion faithfully [17, 28, 76]. However, the lack of lower body tracking signals makes the generated leg motions often possess unrealistic behavior and are inconsistent with the upper body. Special treatment to the lower body is required to generate more realistic motions. To tackle the challenge presented by the absence of observations, we designed two loss functions that regularize the leg motions which *do not require any tracking signals from the lower body*. We incorporate the designed loss functions in the sampling procedure of our S²Fusion, enabling flexible control of the lower body motions.

Scene-penetration loss resolves the incurring of implausible human-scene penetration,

$$\ell_{\text{penetration}}(\mathbf{x}_0) = \sum_{i \in \mathcal{C}} \sum_{b \in \text{KNN}(\mathbf{x}_{0,i}, k)} \max(r - \|\mathbf{x}_{0,i} - b\|, 0), \quad (12)$$

where \mathcal{C} is the set of joints that we want to avoid contact with the scene, k is the number of neighbors to the KNN query, and points in the scene within radius r are considered to be in contact. We empirically set \mathcal{C} as the left/right ankles and the left/right knees; we set $r = 0.02\text{m}$ and $k = 4$ in our evaluation.

Phase-matching loss forces the upper body and the lower body to move in a coherent manner,

$$\ell_{\text{phase}}(\mathbf{x}_0) = \|P_{\text{upper}} - P_{\text{lower}}\|, \quad (13)$$

where P_{upper} is the phase feature computed using the amplitude and phase shift as defined in Sec. 3.2,

$$P_{\text{upper}} = [\sin(2\pi \cdot \mathbf{S}), \cos(2\pi \cdot \mathbf{S}), \mathbf{A}]. \quad (14)$$

Algorithm 1: Sampling procedure of S²Fusion for full-body human motion estimation.

Input: sparse tracking signals $\mathbf{p}^{1:N}$, scene point cloud \mathcal{S}
Output: full-body motion $\mathbf{x}_0^{1:N}$

```

/* preprocessing */
1  $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
2 Draw initial noisy motion from pre-trained prior
 $\hat{\mathbf{x}}^{1:N} \sim f_\phi(z, \mathbf{p}^{1:N})$ 
3 Extract periodic motion feature  $\mathbf{f}^{1:N} := g_\tau(\mathbf{p}^{1:N})$ 
4 Extract scene feature  $\mathbf{E}_S := h_\theta(\mathcal{B}_S)$ 
5 Form conditioning input  $\mathbf{c} := (\mathbf{p}^{1:N}, \mathbf{f}^{1:N}, \mathbf{E}_S)$ 
/* reverse diffusion process */
6  $\mathbf{x}_T^{1:N} := \hat{\mathbf{x}}^{1:N}$ 
7 for  $t = T, \dots, 1$  do
8    $\hat{\mathbf{x}}_0^{1:N} := G(\mathbf{x}_t^{1:N}, t, \mathbf{c})$  // predict clean
   signal
9    $\bar{\mathbf{x}}_0^{1:N} := \hat{\mathbf{x}}_0^{1:N} - \eta \nabla_{\hat{\mathbf{x}}_0^{1:N}} \ell(\hat{\mathbf{x}}_0^{1:N})$  // loss
   guidance
10   $\mathbf{x}_{t-1}^{1:N} := \sqrt{\bar{\alpha}_{t-1}} \bar{\mathbf{x}}_0^{1:N} + \sqrt{1 - \bar{\alpha}_{t-1}} \epsilon$ 
11 return  $\mathbf{x}_0^{1:N}$ 

```

The effectiveness of the phase-matching loss can be justified by the observation that the upper and lower bodies always move in a coordinated way: humans tend to synchronize the upper and lower body movements to achieve balance during everyday activities such as walking, running, dancing, and so on. The movement of the *upper body* provides a valuable cue to the movement of the *lower body*. We plot the periodic motion features of the upper and lower body in Figure 3. We observe there is a clear correlation between the upper and lower body movements.

To compute the phase feature P_{lower} of the lower body, we select the anchor joints that can represent the motion of the lower body: the pelvis and left/right ankles, resembling the tracking signals sent from the head and left/right hands. Then we collect the position, rotation, linear velocity, and angular velocity of the generated lower body motions, following the same procedure in Sec. 3.2. We then compute the phase shift and amplitude parameter and obtain P_{lower} .

Sampling with guidance from loss functions. The overall loss guidance in the sampling stage of S²Fusion is as follows,

$$\ell_{\text{sample}} = \alpha \cdot \ell_{\text{penetration}} + \beta \cdot \ell_{\text{phase}}, \quad (15)$$

where α and β are scaling factors controlling the strength of applying guidance.

After obtaining a clean motion sample $\hat{\mathbf{x}}_0^{1:N}$ as in Sec. 3.3, we guide the sampling process by injecting the gradient of ℓ_{sample} to regularize the motion of lower body,

$$\bar{\mathbf{x}}_0^{1:N} \leftarrow \hat{\mathbf{x}}_0^{1:N} - \eta \nabla_{\hat{\mathbf{x}}_0^{1:N}} \ell_{\text{sample}}(\hat{\mathbf{x}}_0^{1:N}). \quad (16)$$

The full pipeline of the sampling procedure is summarized in Algorithm 1. We chose $\alpha = 0.1, \beta = 0.01, \eta = 1$ empirically.

4. Experiments

We first evaluate and compare our method to others [17, 28, 76] in commonly used metrics for motion reconstruction. Then we conduct ablation studies to showcase the effectiveness of the design choices of our model.

4.1. Experiment Setup

4.1.1 Datasets

We consider two motion-scene datasets to train and evaluate our method. Both datasets consist of rich human-scene interactions and provide fine-grained scene meshes.

CIRCLE [3] contains 10 hours of full-body motion in 9 diverse scenes paired with egocentric information of the environment. We randomly split the training and testing data by a 70:30 ratio, and discarded sequences shorter than 1.0s.

GIMO [77] consists of body pose sequences, scene scans, and eye gaze information. The dataset is collected by using Hololens and IMU-based motion capture suits for motion acquisition; and iPhone 12 for scene scanning.

4.1.2 Evaluation Metrics

We evaluate the performance of our method against others using the commonly used evaluation metrics, which measure motion estimation accuracy, motion smoothness, and scene awareness, respectively.

MPJRE (mean per joint rotation error) measures the average relative rotation error of all body joints in deg. **MPJPE** (mean per joint position error) measures the pose accuracy of each frame in *mm*. **MPJVE** (mean per joint velocity error) measures the average velocity error of all body joints in *mm/s*. **Jitter** measures the average jerk, which is the time derivative of acceleration, of all body joints and reflects the smoothness of motion. **FS** represents the accumulated drift of foot joints during contact, it is computed by the average horizontal displacement between the grounding feet in adjacent frames.

We also conduct per-body part evaluations, as **Hand PE**, **Upper PE** and **Lower PE** denote the hand position error, upper body position error and lower body position error, respectively.

4.2. Comparison

We compare our method with the state-of-the-art methods, [17, 28, 76] on CIRCLE [3] and GIMO [77] datasets. For a fair comparison, we re-train these methods on CIRCLE and GIMO until convergence.

Method	GIMO [77]					CIRCLE [3]				
	MPJRE ↓	MPJPE ↓	MPJVE ↓	Jitter ↓	FS ↓	MPJRE ↓	MPJPE ↓	MPJVE ↓	Jitter ↓	FS ↓
AvatarPoser [28]	7.02	91.3	324.0	16.4	2.04	2.68	30.5	184.6	11.5	2.11
AGRoL [17]	6.58	88.6	269.4	12.5	1.70	2.62	28.7	141.1	8.2	1.72
AvatarJLM [76]	4.95	70.7	258.1	10.7	1.41	2.49	24.6	128.5	7.0	1.53
Ours	4.65	57.8	235.7	10.1	1.39	2.32	19.2	117.6	5.8	1.48

Table 1. Full-body motion estimation results evaluated on GIMO [77] and CIRCLE [3]

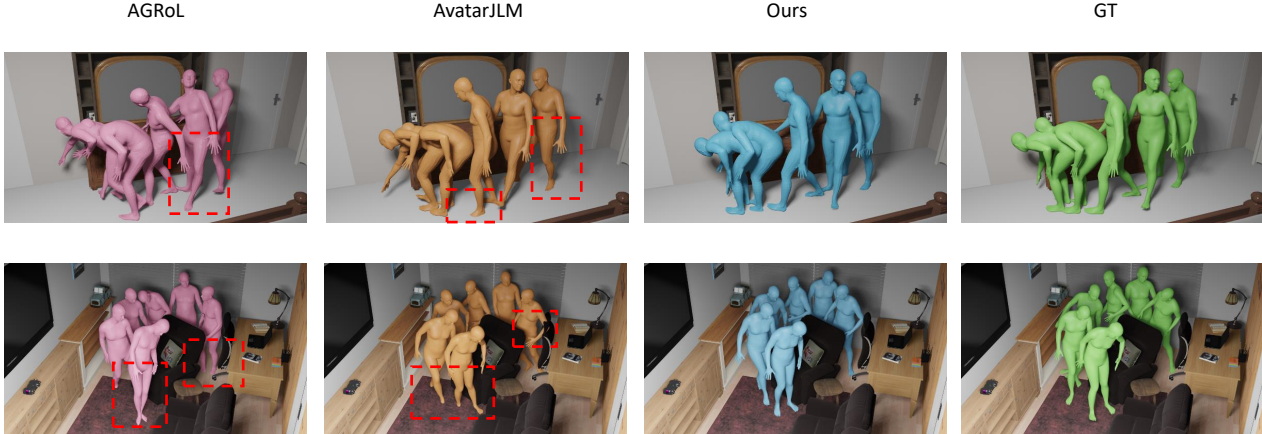


Figure 4. Qualitative results on the CIRCLE [3] dataset. We show the results of two motion sequences in different scenes and highlight the implausible motions in the red box. It can be shown that our method generates more correlated leg motions and avoids scene penetration as much as possible.

We show our quantitative results in Table 1 and Table 2. Compared to other methods, our method achieves the best MPJPE and FS among others in both datasets, demonstrating the effectiveness of incorporating scene information to resolve the sparse-to-dense ambiguities and to generate much more accurate motions. Moreover, our S^2 Fusion model improves the smoothness in motion generation, as showcased by the *MPJVE*, *Jitter* and *FS* metrics. It could be also observed that the estimation quality of the lower body motion is significantly better than other methods by a large margin. The superiority of our method is also evident in the qualitative results shown in Figure 4. By visualizing the pose trajectory in different scenes, we demonstrate that our method generates accurate and smooth motions; being able to incorporate scene information as inputs, the motion generated by our method incurred fewer scene penetrations.

4.3. Ablation Study

To further investigate the effectiveness of design choices in S^2 Fusion, we ablate the effectiveness of different components. We first study the effectiveness of our model components, including the benefit of introducing the scene as an extra modality, and the use of pre-trained motion prior and periodic autoencoder. Then we study the effectiveness of our specially designed loss function during the loss-guided sampling process.

4.3.1 Effectiveness of S^2 Fusion Modules Design

To assess our method’s effectiveness, we incorporated a pre-trained motion prior, a periodic autoencoder, and an additional scene modality. Comparing our approach against four alternatives in Tables 3 and 4, the results show that adding scene information substantially enhances motion estimation quality with limited upper body tracking data. This confirms our claim that scene data can significantly resolve ambiguities in sparse-to-dense scenarios. Moreover, initializing with a motion prior addresses data scarcity issues and improves motion smoothness and accuracy, while the periodic autoencoder enhances estimation by aligning motion temporally and spatially.

4.3.2 Effectiveness of Loss-guided Sampling

We ablate the importance of incorporating our designed loss function in the loss-guided sampling process. We show the results compared against three alternatives in Table 5 and Table 6. We notice that although the scene-penetration loss and phase-matching loss can improve the motion estimation quality, they may introduce jittering motions into the generated motions. Compared to scene-penetration loss, incorporating phase-matching loss is more effective in producing accurate motions, coinciding with our observation that the phase fea-

Method	GIMO [77]			CIRCLE [3]		
	Hand PE	Upper PE	Lower PE	Hand PE	Upper PE	Lower PE
AvatarPoser [28]	36.8	40.1	198.6	11.4	15.3	83.3
AGRoL [17]	23.7	29.8	163.9	9.1	14.5	77.4
AvatarJLM [76]	25.0	33.4	132.6	9.3	15.8	69.1
Ours	23.1	28.7	107.9	8.8	12.2	57.3

Table 2. More metrics comparison with AvatarPoser [28], AGRoL [17], and AvatarJLM [76] on GIMO [77] and CIRCLE [3]. We show the results of comparing the hand, upper body, and lower body reconstruction quality.

Components						
MP	Scene	PAE	MPJPE	MPJVE	Jit.	FS
×	×	×	26.2	135.9	7.9	1.65
×	✓	×	22.7	128.1	7.1	1.58
✓	✓	×	20.9	120.5	6.0	1.51
×	✓	✓	21.8	125.3	6.6	1.52
✓	✓	✓	19.2	117.6	5.8	1.48

Table 3. Ablation on various components of our model on CIRCLE. **MP** denotes the pre-trained motion prior, **Scene** indicates whether the model receives the scene information as an extra input, and **PAE** denotes the periodic autoencoder.

Components						
MP	Scene	PAE	MPJPE	MPJVE	Jit.	FS
×	×	×	76.6	264.9	11.7	1.68
×	✓	×	68.1	257.4	11.2	1.67
✓	✓	×	60.3	243.6	10.3	1.46
×	✓	✓	65.9	249.3	10.4	1.52
✓	✓	✓	57.8	235.7	10.1	1.39

Table 4. Ablation on various components of our model on GIMO. **MP** denotes the pre-trained motion prior, **Scene** indicates whether the model receives the scene information as an extra input, and **PAE** denotes the periodic autoencoder.

Loss fn.						
$\ell_{\text{penetration}}$	ℓ_{phase}	MPJPE	MPJVE	Jit.	FS	
×	×	20.6	119.1	5.6	1.43	
✓	×	20.1	117.9	5.8	1.40	
×	✓	19.8	118.5	6.1	1.50	
✓	✓	19.2	117.6	5.8	1.48	

Table 5. Ablation on the effect of our designed loss function during loss-guided sampling on CIRCLE.

tures are composed of time-alignment features and feature vector resembles momentum. The scene-penetration loss has merit on its own as it can avoid falsely contacting the ground, especially in the CIRCLE dataset.

5. Conclusion

Our work addresses the pivotal challenge of estimating full-body human motion in 3D scenes from sparse tracking sig-

Loss fn.					
$\ell_{\text{penetration}}$	ℓ_{phase}	MPJPE	MPJVE	Jit.	FS
×	×	59.9	240.5	10.5	1.52
✓	×	58.3	239.1	10.4	1.58
×	✓	57.6	236.8	10.1	1.42
✓	✓	57.8	235.7	10.1	1.39

Table 6. Ablation on the effect of our designed loss function during loss-guided sampling on GIMO.

nals, a critical aspect for advancing AR/VR applications. To tackle this problem, we propose S²Fusion that integrates Scene and sparse Signals through a conditional Diffusion model. S²Fusion first captures spatial-temporal relations in sparse signals using a periodic autoencoder, generating time-alignment feature embeddings as additional inputs. Employing conditional diffusion and drawing initial motion from a pre-trained prior, S²Fusion effectively fuses scene geometry and sparse tracking signals to generate full-body scene-aware motions. To improve plausibility and coherence, S²Fusion incorporates a specially designed scene-penetration loss and phase-matching loss, providing guidance for the sampling procedure. These losses effectively regularize lower-body motion, even in the absence of tracking signals. Extensive experiments demonstrate that S²Fusion outperforms state-of-the-art methods by a large margin. Extending S²Fusion for human motion estimation in more complex scenarios by systematically integrating comprehensive physically plausible constraints is under consideration for our future work.

Acknowledgement

This work was supported by Shanghai Local College Capacity Building Program (23010503100), NSFC (No.62303319), Shanghai Sailing Program (22YF1428800, 21YF1429400), Shanghai Frontiers Science Center of Human-centered Artificial Intelligence (ShangHAI), MoE Key Laboratory of Intelligent Perception and Human-Machine Collaboration (ShanghaiTech University), and Shanghai Clinical Research and Trial Center.

References

- [1] Sadegh Aliakbarian, Pashmina Cameron, Federica Bogo, Andrew Fitzgibbon, and Thomas J. Cashman. Flag: Flow-based 3d avatar generation from sparse observations, 2022. [1](#), [2](#)
- [2] Sadegh Aliakbarian, Fatemeh Saleh, David Collier, Pashmina Cameron, and Darren Cosker. Hmd-nemo: Online 3d avatar motion generation from sparse observations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9622–9631, 2023. [1](#), [2](#)
- [3] João Pedro Araújo, Jiaman Li, Karthik Vetrivel, Rishi Agarwal, Jiajun Wu, Deepak Gopinath, Alexander William Clegg, and Karen Liu. Circle: Capture in rich contextual environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21211–21221, 2023. [6](#), [7](#), [8](#)
- [4] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022. [3](#)
- [5] Philippe Beaudoin, Pierre Poulin, and Michiel van de Panne. Adapting wavelet compression to human motion capture clips. In *Proceedings of Graphics Interface 2007*, pages 313–318, 2007. [3](#)
- [6] Bharat Lal Bhatnagar, Xianghui Xie, Ilya A Petrov, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Behave: Dataset and method for tracking human object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15935–15946, 2022. [3](#)
- [7] Armin Bruderlin and Lance Williams. Motion signal processing. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 97–104, 1995. [3](#)
- [8] Zhe Cao, Ilija Radosavovic, Angjoo Kanazawa, and Jitendra Malik. Reconstructing hand-object interactions in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12417–12426, 2021. [3](#)
- [9] 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](#), [3](#)
- [10] Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. *Advances in Neural Information Processing Systems*, 35:25683–25696, 2022. [3](#)
- [11] Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023. [2](#), [3](#)
- [12] Peishan Cong, Yiteng Xu, Yiming Ren, Juze Zhang, Lan Xu, Jingya Wang, Jingyi Yu, and Yuexin Ma. Weakly supervised 3d multi-person pose estimation for large-scale scenes based on monocular camera and single lidar. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 461–469, 2023. [3](#)
- [13] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. [2](#), [3](#)
- [14] Xinhan Di, Xiaokun Dai, Xinkang Zhang, and Xinrong Chen. Dual attention poser: Dual path body tracking based on attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2794–2803, 2023. [1](#)
- [15] Christian Diller and Angela Dai. Cg-hoi: Contact-guided 3d human-object interaction generation. *arXiv preprint arXiv:2311.16097*, 2023. [3](#)
- [16] Andrea Dittadi, Sebastian Dziadzio, Darren Cosker, Ben Lundell, Thomas J. Cashman, and Jamie Shotton. Full-body motion from a single head-mounted device: Generating smpl poses from partial observations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11687–11697, 2021. [1](#), [2](#)
- [17] Yuming Du, Robin Kips, Albert Pumarola, Sebastian Starke, Ali Thabet, and Artsiom Sanakoyeu. Avatars grow legs: Generating smooth human motion from sparse tracking inputs with diffusion model. In *CVPR*, 2023. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [18] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7346–7356, 2023. [2](#)
- [19] 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. [3](#)
- [20] 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*, pages 5152–5161, 2022. [3](#)
- [21] Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J Black. Resolving 3d human pose ambiguities with 3d scene constraints. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2282–2292, 2019. [3](#)
- [22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. [2](#)
- [23] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. [2](#)
- [24] Chun-Hao P Huang, Hongwei Yi, Markus Höschle, Matvey Safroshkin, Tsvetelina Alexiadis, Senya Polikovskiy, Daniel Scharstein, and Michael J Black. Capturing and inferring dense full-body human-scene contact. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13274–13285, 2022. [3](#)

- [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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [26] Yinghao Huang, Manuel Kaufmann, Emre Aksan, Michael J Black, Otmar Hilliges, and Gerard Pons-Moll. Deep inertial poser: Learning to reconstruct human pose from sparse inertial measurements in real time. *ACM Transactions on Graphics (TOG)*, 37(6):1–15, 2018. 2
- [27] Chaofan Huo, Ye Shi, Yuexin Ma, Lan Xu, Jingyi Yu, and Jingya Wang. Stackflow: monocular human-object reconstruction by stacked normalizing flow with offset. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 902–910, 2023. 3
- [28] Jiayi Jiang, Paul Streli, Huajian Qiu, Andreas Fender, Larissa Laich, Patrick Snape, and Christian Holz. Avatarposer: Articulated full-body pose tracking from sparse motion sensing. In *Proceedings of European Conference on Computer Vision*. Springer, 2022. 1, 2, 3, 5, 6, 7, 8
- [29] Yifeng Jiang, Yuting Ye, Deepak Gopinath, Jungdam Won, Alexander W Winkler, and C Karen Liu. Transformer inertial poser: Attention-based real-time human motion reconstruction from sparse imus. *arXiv e-prints*, pages arXiv–2203, 2022. 2
- [30] Korrawe Karunratanakul, Jinlong Yang, Yan Zhang, Michael J Black, Krikamol Muandet, and Siyu Tang. Grasping field: Learning implicit representations for human grasps. In *2020 International Conference on 3D Vision (3DV)*, pages 333–344. IEEE, 2020. 3
- [31] Korrawe Karunratanakul, Konpat Preechakul, Supasorn Suwanajakorn, 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. 3
- [32] Ben Kenwright. Quaternion fourier transform for character motions. 2015. 3
- [33] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020. 2
- [34] 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. *arXiv preprint arXiv:2307.07511*, 2023. 3
- [35] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13401–13412, 2021. 3
- [36] Han Liang, Wenqian Zhang, Wenxuan Li, Jingyi Yu, and Lan Xu. Intergen: Diffusion-based multi-human motion generation under complex interactions. *International Journal of Computer Vision*, pages 1–21, 2024. 3
- [37] Zicheng Liu, Steven J Gortler, and Michael F Cohen. Hierarchical spacetime control. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 35–42, 1994. 3
- [38] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 851–866. 2023. 3
- [39] Zhaoyang Lyu, Xudong XU, Ceyuan Yang, Dahua Lin, and Bo Dai. Accelerating diffusion models via early stop of the diffusion process, 2022. 3
- [40] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. 1, 2, 3, 5
- [41] Xiaogang Peng, Yiming Xie, Zizhao Wu, Varun Jampani, Deqing Sun, and Huaizu Jiang. Hoi-diff: Text-driven synthesis of 3d human-object interactions using diffusion models. *arXiv preprint arXiv:2312.06553*, 2023. 3
- [42] 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. 3
- [43] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. 4
- [44] Yonatan Shafir, Guy Tevet, Roy Kapon, and Amit H Bermano. Human motion diffusion as a generative prior. *arXiv preprint arXiv:2303.01418*, 2023. 3
- [45] Mingyi Shi, Sebastian Starke, Yuting Ye, Taku Komura, and Jungdam Won. Phasemp: Robust 3d pose estimation via phase-conditioned human motion prior. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14725–14737, 2023. 3
- [46] Jiaming Song, Qinheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin Chen, and Arash Vahdat. Loss-guided diffusion models for plug-and-play controllable generation. 2023. 2, 3
- [47] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 2
- [48] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. *arXiv preprint arXiv:2111.08005*, 2021. 3
- [49] Sebastian Starke, Ian Mason, and Taku Komura. Deepphase: Periodic autoencoders for learning motion phase manifolds. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022. 2, 3, 4
- [50] 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. 3
- [51] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. 2, 3, 5
- [52] Jie Tian, Lingxiao Yang, Ran Ji, Yuexin Ma, Lan Xu, Jingyi Yu, Ye Shi, and Jingya Wang. Gaze-guided hand-object in-

- teraction synthesis: Benchmark and method. *arXiv preprint arXiv:2403.16169*, 2024. 3
- [53] Shashank Tripathi, Lea Müller, Chun-Hao P Huang, Omid Taheri, Michael J Black, and Dimitrios Tzionas. 3d human pose estimation via intuitive physics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4713–4725, 2023. 3
- [54] Jonathan Tseng, Rodrigo Castellon, and Karen Liu. Edge: Editable dance generation from music. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 448–458, 2023. 3
- [55] Munetoshi Unuma, Ken Anjyo, and Ryozo Takeuchi. Fourier principles for emotion-based human figure animation. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 91–96, 1995. 3
- [56] Timo Von Marcard, Bodo Rosenhahn, Michael J Black, and Gerard Pons-Moll. Sparse inertial poser: Automatic 3d human pose estimation from sparse imus. In *Computer graphics forum*, pages 349–360. Wiley Online Library, 2017. 2
- [57] 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. 3
- [58] Jingbo Wang, Yu Rong, Jingyuan Liu, Sijie Yan, Dahua Lin, and Bo Dai. Towards diverse and natural scene-aware 3d human motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20460–20469, 2022.
- [59] Zhenzhen Weng and Serena Yeung. Holistic 3d human and scene mesh estimation from single view images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 334–343, 2021. 3
- [60] Alexander Winkler, Jungdam Won, and Yuting Ye. Questsim: Human motion tracking from sparse sensors with simulated avatars. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–8, 2022. 1
- [61] Qianyang Wu, Ye Shi, Xiaoshui Huang, Jingyi Yu, Lan Xu, and Jingya Wang. Thor: Text to human-object interaction diffusion via relation intervention. *arXiv preprint arXiv:2403.11208*, 2024. 3
- [62] Sirui Xu, Zhengyuan Li, Yu-Xiong Wang, and Liang-Yan Gui. Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14928–14940, 2023. 3
- [63] Dongseok Yang, Doyeon Kim, and Sung-Hee Lee. Lobstr: Real-time lower-body pose prediction from sparse upper-body tracking signals. In *Computer Graphics Forum*, pages 265–275. Wiley Online Library, 2021. 1
- [64] Yongjing Ye, Libin Liu, Lei Hu, and Shihong Xia. Neural3points: Learning to generate physically realistic full-body motion for virtual reality users. In *Computer Graphics Forum*, pages 183–194. Wiley Online Library, 2022. 1
- [65] Xinyu Yi, Yuxiao Zhou, and Feng Xu. Transpose: Real-time 3d human translation and pose estimation with six inertial sensors. *ACM Transactions on Graphics*, 40(4), 2021. 2
- [66] Xinyu Yi, Yuxiao Zhou, Marc Habermann, Soshi Shimada, Vladislav Golyanik, Christian Theobalt, and Feng Xu. Physical inertial poser (pip): Physics-aware real-time human motion tracking from sparse inertial sensors. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [67] Xinyu Yi, Yuxiao Zhou, Marc Habermann, Vladislav Golyanik, Shaohua Pan, Christian Theobalt, and Feng Xu. EgoLocate: Real-time motion capture, localization, and mapping with sparse body-mounted sensors. *ACM Transactions on Graphics (TOG)*, 42(4), 2023. 2
- [68] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 3
- [69] Juze Zhang, Jingya Wang, Ye Shi, Fei Gao, Lan Xu, and Jingyi Yu. Mutual adaptive reasoning for monocular 3d multi-person pose estimation. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1788–1796, 2022. 2
- [70] Juze Zhang, Haimin Luo, Hongdi Yang, Xinru Xu, Qianyang Wu, Ye Shi, Jingyi Yu, Lan Xu, and Jingya Wang. Neuraldome: A neural modeling pipeline on multi-view human-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8834–8845, 2023. 3
- [71] Juze Zhang, Ye Shi, Yuexin Ma, Lan Xu, Jingyi Yu, and Jingya Wang. Ikol: Inverse kinematics optimization layer for 3d human pose and shape estimation via gauss-newton differentiation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3454–3462, 2023. 2
- [72] 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. 3
- [73] Siwei Zhang, Yan Zhang, Federica Bogo, Marc Pollefeys, and Siyu Tang. Learning motion priors for 4d human body capture in 3d scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11343–11353, 2021. 3
- [74] Yan Zhang, Michael J Black, and Siyu Tang. Perpetual motion: Generating unbounded human motion. *arXiv preprint arXiv:2007.13886*, 2020. 3
- [75] Rui Zhao, Hui Su, and Qiang Ji. Bayesian adversarial human motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6225–6234, 2020. 3
- [76] Xiaozheng Zheng, Zhuo Su, Chao Wen, Zhou Xue, and Xiaojie Jin. Realistic full-body tracking from sparse observations via joint-level modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14678–14688, 2023. 1, 2, 5, 6, 7, 8
- [77] 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. 6, 7, 8

- [78] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2019. [3](#)
- [79] Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhong Cao, Bihan Wen, Radu Timofte, and Luc Van Gool. Denoising diffusion models for plug-and-play image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1219–1229, 2023. [3](#)