

Stratified Avatar Generation from Sparse Observations

Han Feng^{1*,*} Wenchao Ma^{2*} Quankai Gao³ Xianwei Zheng¹ Nan Xue^{1,4†,*} Huijuan Xu²

¹Wuhan University ²Pennsylvania State University

³University of Southern California ⁴Ant Group

Project website: <https://fhan235.github.io/SAGENet/>




Figure 1. **Stratified avatar generation from sparse observations.** Given the sensory sparse observation of the body motion: 6-DoF poses of the head and hand marked by RGB axes in (a), our method leverages a disentangled body representation in (b) to reconstruct the upper-body conditioned on the sparse observation in (c), and lower-body conditioned on the upper-body reconstruction in (d) to accomplish the full-body reconstruction in (e).

Abstract

Estimating 3D full-body avatars from AR/VR devices is essential for creating immersive experiences in AR/VR applications. This task is challenging due to the limited input from Head Mounted Devices, which capture only sparse observations from the head and hands. Predicting the full-body avatars, particularly the lower body, from these sparse observations presents significant difficulties. In this paper, we are inspired by the inherent property of the kinematic tree defined in the Skinned Multi-Person Linear (SMPL) model, where the upper body and lower body share only one common ancestor node, bringing the potential of decoupled reconstruction. We propose a stratified approach to decouple the conventional full-body avatar reconstruction pipeline into two stages, with the reconstruction of the upper body first and a subsequent reconstruction of the lower body conditioned on the previous stage. To implement this straightforward idea, we leverage the latent diffusion model as a powerful probabilistic generator, and train it to follow the latent distribution of decoupled motions explored by a VQ-VAE encoder-decoder model. Extensive experiments on AMASS mocap dataset demonstrate our state-of-the-art performance in the reconstruction of full-body motions.

1. Introduction

Generating 3D full-body avatars from observations of Head Mounted Devices (HMDs) is crucial for enhancing immersive AR/VR experiences. HMDs primarily track the head and hands, while leaving the rest of the body unmonitored. This limited motion tracking poses a challenging scenario for accurately reconstructing full-body 3D avatars, particularly in representing the lower body. The high degree of freedom in body movements compounds this difficulty, making the task of reasoning human motion from such sparse observations significantly complex.

Tremendous efforts have been made to obtain more tracking signals by adding sensors at Pelvis [5, 10, 33] or both Pelvis and Legs [16, 19, 45, 49, 50]. While these approaches provide more data points for avatar construction, they can diminish the user’s experience. Wearing extra devices can be cumbersome, potentially interfering with the user’s comfort and immersion in the virtual environment. This trade-off highlights the need for innovative solutions that can deliver detailed body tracking without compromising the user’s comfort and immersion in AR/VR settings. Accordingly, we are interested in the problem of generating 3D full-body avatars from sparse observations of HMDs that track the motion of the head and two hands, by developing a neural solution that learns the distribution of full-body poses given the sparse observations as the input condition.

* indicates the School of Computer Science, Wuhan University

* indicates equal contributions.

† Corresponding author: xuenan@ieee.org

Recent studies have attempted to address the challenge of sparse observations in HMD-based full-body avatar generation by employing regression-based techniques, as seen in [18, 53], or by adopting generation-based approaches like [7, 11]. These methods typically use deep neural networks to predict human motion within a single, expansive motion space. However, due to the limited data provided by sparse observations, these networks often struggle to fully capture the complexities of human kinematics across such a broad and unified motion space. This limitation frequently results in reconstructions that are unrealistic and lack physical plausibility.

We introduce a new method for reconstructing full-body human motions from sparse observations, called Stratified Avatar Generation (SAGE). Instead of the upper-body motion prediction that has tracking signals of certain upper joints from sparse observations, predicting lower-body motion is not straightforward as no direct tracking signals about any lower-body joint is given. It is noteworthy that SMPL model [22] connects the upper and lower half-body by a single root joint, as shown in Fig. 1 (b), which motivates us to split the full-body motions into upper and lower half-body parts. The benefits are two-fold: 1) the smaller search space achieved by disentanglement facilitates learning and prediction; 2) our stratified design makes the modeling and inferring for lower-body motions more accurate and visually appealing by explicitly modeling the correlation and constraint between two half-body motions. To this end, we use VQ-VAE [43] to encode and reconstruct upper and lower body motions separately.

With the disentangled latent representation of the upper and lower body motions, we aim to recover the accurate full-body motions from sparse observations with a body-customized latent diffusion model (LDM) [35] in a stratified manner. Specifically, as shown in Fig. 1 (c), Fig. 1(d), and Fig. 1(e), we first find the latent of upper-body motion condition on the sparse observations (i.e., tracking signals of the head and hands in Fig. 1(a)). Then, the latent of lower-body motion is inferred condition on both the predicted upper-body latent and sparse observations. Finally, a full-body decoder takes the two half-body latents as input and outputs the full-body motion.

In the experiments, we comprehensively justified our intuitive design of disentangling the upper and lower body motion in a stratified manner. On the large-scale motion capture benchmark AMASS [25], our proposed SAGE is exhibiting superior performance in different evaluation settings and particularly performs well in terms of the evaluation metrics for lower-body motion estimation compared to previous state-of-the-art methods.

2. Related Work

2.1. Motion Reconstruction from Sparse Input

The task of reconstructing full human body motion from sparse observations has gained significant attention in recent decades within the research community [1, 3, 5, 7, 10, 11, 16, 18, 19, 45, 46, 48–50, 53]. For instance, recent works [16, 19, 45, 49, 50] focus on reconstructing full body motion from six inertial measurement units (IMUs). SIP [45] employs heuristic methods, while DIP [16] pioneers the use of deep neural networks for this task. PIP [50] and TIP [19] further enhance performance by incorporating physics constraints. With the rise of VR/AR applications, researchers turn their attention toward reconstructing full body motion from VR/AR devices, such as head-mounted devices (HMDs), which only provide information about the user’s head and hands, posing additional challenges. LoB-Str [48], AvatarPoser [18], and AvatarJLM [53] approach this task as a regression problem, utilizing GRU [48] and Transformer Network [18, 53] to predict the full body pose from sparse observations of HMDs. Another line of methods employs generative models [5, 7, 10, 11]. For example, VAEHMD [10] and FLAG [5] utilize Variational AutoEncoder (VAE) [20] and Normalizing flow [34], respectively. Recent works [7, 11] leverage more powerful diffusion models [15, 37] for motion generation, yielding promising results due to the powerful ability of diffusion models in modeling the conditional probabilistic distribution of full-body motion.

Contrasting with previous methods that model full-body motion in a comprehensive, unified framework, our approach acknowledges the complexities such methods impose on deep learning models, particularly in capturing the intricate kinematics of human motion. Hence, we propose a stratified approach that decouples the conventional full-body avatar reconstruction pipeline, first for the upper body and then for the lower body under the condition of the upper-body.

2.2. Human Motion Generation

Human motion generation is explored under various input conditions, including text [17, 30, 38, 39, 51], action labels [13, 29], 3D scenes and objects [14, 31, 47], and motion itself [8, 24, 27]. Our work shares similarities with two primary research streams. The first involves diffusion-based motion generation. For instance, [39] is the first to utilize a diffusion model [15, 37] for text-to-motion generation with a transformer network. [31] develops a hierarchical generation pipeline for human-object interaction by generating initial keyframes in the motion sequence and then interpolating between them. Secondly, our approach parallels works like PoseGPT [24], MotionGPT [17], and T2M-GPT [51] in terms of representing human motions with discrete la-

tent. These studies also utilize a VQ-VAE [43] to encode human motion into a discrete latent space, facilitating the subsequent generation process.

The task we aim to address significantly deviates from traditional motion generation tasks like text-to-motion generation, which typically aims to create motions that align with textual descriptions. Our goal is distinctly different: we focus on accurately reconstructing human motion using solely sparse observations.

3. SAGE: Stratified Avatar Generation

This section introduces the proposed SAGE Network, following our observation about the connection relationship between upper-body and lower-body motions. The overall architecture of our SAGE Net is shown in Fig. 2. Disentangled latent representations for upper-body and lower-body motions are learned in Fig. 2 (a). Subsequently, as illustrated in Fig. 2 (b), we employ a stratified latent diffusion process for full body motion reconstruction.

3.1. Problem Statement and Notation

Input Signals. Our paper follows the common setting of Head Mounted Devices (HMDs) inputs for motion generation, in which three sensors mounted on the head and left and right hands are employed to perceive the corresponding joint motions. Formally, the raw input signals are denoted by a time-dependent vector function $\mathbf{X}(t) = (\mathbf{x}_h(t), \mathbf{x}_l(t), \mathbf{x}_r(t))$, where the subscripts h, l and r indicate the head, left hand, and right hand, and all these functions are with six degree of freedom for 3D rotation and translation under the global coordinate system. Joint rotations are represented by a six-axis representation, which has been demonstrated to be more suitable for network learning in previous works [54]. Given a time interval with T sampling points, the raw input signals can be denoted in a matrix $\mathbf{X}_{\text{raw}} \in \mathbb{R}^{T \times (3 \times (3+6))}$. To enhance the input signals, we follow [18] to compute the positional velocities and angular velocities. This augmentation process adds a 9D input signal for each observed joint, resulting in an 18D input signal per joint at every timestamp. By combining these signals for all joints over all timestamps, we form the complete sparse input signal, represented as $\mathbf{X} \in \mathbb{R}^{T \times 54}$.

Kinematic Tree and SMPL Representation. As shown in Fig. 1 (b), SMPL [22] represents a human pose by a standard skeletal rig, which is widely adopted by current motion generation works. A pose $\theta^j(t)$ represents the relative rotation of joint j at t th frame with respect to its parent in the kinematic tree. The global rotations $G(\theta^j(t))$ can be calculated by:

$$G(\theta^j(t)) = \prod_{k \in A(j)} \theta^k(t) \quad (1)$$

where $A(j)$ denotes the ordered set of joint ancestors of joint j .

Articulated motion representation based on the human kinematic tree, is key for realistically simulating human motions and enables efficient control over joint parameters. However, such an intricate articulated motion representation poses a significant challenge for models to learn effectively. In this work, we seek to disentangle this complex representation to enable the model to focus on a limited set of motions and interactions, thereby simplifying the learning process.

Nevertheless, separating full-body human motions into distinct parts is nontrivial due to the complex correlations among joints. We revisit the human kinematic tree defined in SMPL model, where the upper and lower half-body is connected solely via a root joint. This insight from SMPL model provides a natural solution to separate the articulated full-body motion into two distinct parts: upper-body motion and lower-body motion. Notably, the root joint is included in both two parts as a central element since the parameters of all other joints in each half-body are defined in the local coordinate system of the root joint.

The Outputs. As discussed in the last paragraph for SMPL representation [22], the problem of 3D body avatar generation comes down to the full-body motion estimation of 22 joints (including the root joint), denoted in the set function $\Theta(t) = \{\theta^i(t) \in \text{SE}(3) | t \in \{t_1, \dots, t_T\}\}$ as the expected output of our problem. Based on the discussion of SMPL model with the disentangling nature of upper and lower body, we redefine the set function $\Theta(t)$ in the disentangled way by $\Theta(t) = \Theta_{\text{upper}}(t) \cup \Theta_{\text{lower}}(t)$, where $\Theta_{\text{upper}}(t) = \{\theta^0(t), \dots, \theta_{b_u}^u(t)\}$, $\Theta_{\text{lower}} = \{\theta^0(t), \dots, \theta_{b_l}^l(t)\}$. These two subsets have only one intersected joint: root joint θ^0 , and $b_u = 13$ and $b_l = 8$ denote the number of rest joints in the upper and lower body, respectively. For the final output of our method, the dimension of the underlying motion variables is $22 \times 6 = 132$ at every timestamp.

3.2. Disentangled Motion Representation

In this section, our objective is to disentangle full-body human motions into upper-body and lower-body parts and encode them to discrete latent spaces. This can effectively reduce the complexity and burden of encoding since each encoding takes care of only half-body motions.

We employ two autoencoders, i.e., VQ-VAE [43, 52], with identical architecture to learn the discrete latent spaces for upper-body and lower-body motions, respectively. As shown in Fig. 2 (a), our VQ-VAE model consists of an encoder and a decoder. The encoder E takes the motion sequence $\Theta = \{\theta_i\}_{i=1}^T$ as input and encodes it into a series of

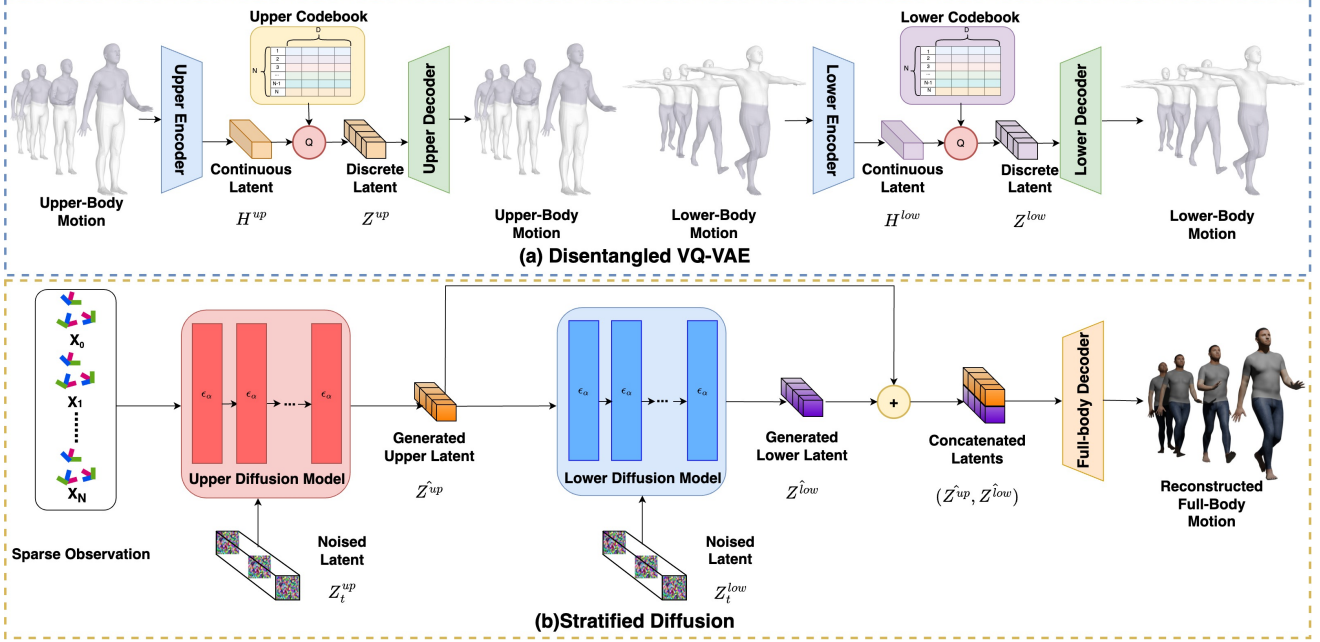


Figure 2. The overall architecture of our SAGE Net. It mainly contains two components: (a) Disentangled VQ-VAE for discrete human motion latent learning. To facilitate visualization, we incorporate zero rotations as padding for the lower body in the Upper VQ-VAE, and vice versa for the Lower VQ-VAE. Consequently, in the visualizations of the Upper VQ-VAE, the lower body remains in a stationary pose, whereas in the visualizations of the Lower VQ-VAE, the upper body is maintained in a T-pose. (b) The stratified diffusion model, which models the conditional distribution of the latent space for upper and lower motion. This model sequentially infers the upper and lower body latents, capturing the correlation between upper and lower motions. By employing a dedicated full-body decoder on the concatenated upper and lower latents, we can obtain full-body motion.

continuous latent $E(\Theta) = H$, where $H = \{h_i\}_{i=1}^{T/l}$, and l is the temporal down-sampling rate of input motion sequence.

To quantize the continuous latent, we define the discrete motion latent space by a codebook $C = \{c_j\}_{j=1}^N \in \mathbb{R}^{N \times D}$, where $N = 512$ is the number of entries in the codebook and $D = 384$ is the dimension of each element c_j . The operation Q quantizes the continuous latent h_i into discrete latent z_i by finding its most similar element in C :

$$z_i = Q(h_i) = \arg \min_{c_j \in C} \|h_i - c_j\|_2 \quad (2)$$

Since continuous latent from all data samples share the same codebook C , all the real motions in the training set could be expressed by a finite number of bases in latent space.

Subsequently, the quantized latents Z are fed into the decoder to reconstruct the original motions, given by $\hat{\Theta} = D(Z)$. The training process involves the joint optimization of the encoder and decoder by minimizing the following loss over the training dataset:

$$Loss_{vq} = \text{Smooth}_{L_1}(\hat{\Theta}, \Theta) + \|\text{sg}[Z] - H\|_2 + \beta \|Z - \text{sg}[H]\|_2 \quad (3)$$

Here, sg denotes the stop gradient operator, and β is a hyperparameter. We have two independent VQ-VAEs for

upper-body and lower-body motion encoding, which we refer to as VQ-VAE_{up} and VQ-VAE_{low} .

3.3. Stratified Motion Diffusion

After encoding and expressing different human motions as latents, we aim to properly sample from the latent space for full-body motion reconstructions and match the sparse observations.

Although disentangling the full-body motions into upper and lower parts enhances effectiveness and efficiency for motion representation learning, it's crucial to include the correlation between two body parts during generation. Otherwise, severe inconsistency would be witnessed in reconstructed full-body motions. To this end, we propose Stratified Motion Diffusion to sample upper-body and lower-body latent in a cascaded manner with explicit considerations of the correlations mentioned above.

Since the sparse observations are all from the upper body (e.g., head and hand sensors), we first generate upper-body latent z^{up} by upper diffusion model conditioning on the sparse observations X . Thus the training objective of the upper diffusion model is:

$$L_{up} := \mathbb{E}_{z^{up}, \epsilon \sim \mathcal{N}(0,1), k} \left[\|\epsilon - \epsilon_\alpha(z_k^{up}, X, k)\|_2^2 \right] \quad (4)$$

where ϵ is the random noise from the normal distribution, ϵ_α is the noise predictor of the diffusion model with network parameters α , and k is the diffusion time step.

Compared with upper-body latent prediction, directly predicting the lower-body latent $z^{\hat{low}}$ from the same sparse observations is more challenging due to the absence of direct tracking or supervision for any of the lower-body joints. To make the prediction more physically meaningful, as shown in Fig. 2 (b), we take both the sparse observations X and the generated upper-body latent $z^{\hat{up}}$ as conditions for lower-body latent prediction by lower diffusion model. This design considers the correlation between two half-body parts and allows more information to be involved for the lower-body inference. The objective for lower diffusion model training is as follows:

$$L_{low} := \mathbb{E}_{z^{low}, \epsilon \sim \mathcal{N}(0,1), k} \left[\left\| \epsilon - \epsilon_\alpha \left(z_k^{low}, (X, z^{\hat{up}}), k \right) \right\|_2^2 \right] \quad (5)$$

Once two half-body latent $z^{\hat{up}}$ and $z^{\hat{low}}$ are obtained, the full-body motions can be recovered with a decoder $\hat{\Theta} = D_{full}(z^{\hat{up}}, z^{\hat{low}})$. Instead of directly using pre-trained upper and lower decoders in Fig. 2 (a) to recover the corresponding half-body motions, we train this full-body decoder E_{full} from scratch together with our stratified motion diffusion, which is further optimized to capture the correlations between half-body motions.

3.4. Implementation Details

Since both sparse observations and human motion occur sequentially, we utilize the widely adopted sequential network, i.e., transformer [44], as the backbone network for the encoder and decoder in the disentangled VQ-VAE [43], and the denoise network in the stratified diffusion model. We set temporal down-sampling rate $l = 2$ to balance the computational cost and the performance. During the training of the latent diffusion model, instead of predicting noise ϵ_k as formulated by the standard latent diffusion model [35], we follow [32, 39] and directly predict the latent z itself, as we find that this operation can significantly reduce the sampled time steps during inference stage. For training decoders, i.e., E_{up} , E_{low} and E_{full} , in addition to the rotation-level reconstruction loss, we incorporate the forward kinematic loss proposed in [18] and the hand loss described in [53].

For the inference stage, we evaluate our model in an online manner. Specifically, we fix the sequence length at 20 for both the input and the output of our model, and only the last pose in the output motion sequence is retained. Given a sparse observation sequence, we apply our model using a sliding window approach. For the first 20 poses in the motion sequence, we predict by padding the sparse observation sequence x at the beginning with the first available observation. We make this choice considering the practicality and relevance of online inference in real-world application sce-

narios. This allows the motion sequences to be predicted in a frame-by-frame manner.

In addition, we employ a simple two-layer GRU [9] on the top of the full body decoder as a temporal memory to smooth the prediction of the output sequence with minimal computational expense, and we term it as a Refiner. To train this Refiner, we use the same velocity loss as [53]. Our model takes 0.74ms to infer 1 frame on a single NVIDIA RTX3090 GPU.

4. Experiments and Evaluation Metrics

4.1. Dataset and Evaluation Metrics

We train and evaluate our method on AMASS [25], which unifies multiple motion capture datasets [2, 4, 6, 12, 23, 26, 28, 36, 40–42] as SMPL [22] representations.

We report several metrics for evaluations and comparisons: mean per joint rotation error (MPJRE) and mean per joint position error (MPJPE) for measuring the average relative rotation and position error across all joints respectively, as well as the average position error of the root joints (Root PE), hand joints (Hand PE), upper-body joints (Upper PE), and lower-body joints (Lower PE).

Besides the above reconstruction accuracy, we also evaluate the spatial and temporal consistency of the generated sequences, as it significantly contributes to the visual quality. Specifically, we calculate the mean per joint velocity error (MPJVE) and Jitter, where MPJVE measures the average velocity error of all body joints, and Jitter quantifies the average jerk (time derivative of acceleration) of all body joints. In both cases, lower values indicate better results.

4.2. Quantitative and Qualitative Results

For fair comparison, we follow two settings used in previous works [5, 10, 11, 18, 33, 53] for quantitative and qualitative assessment. Moreover, we propose a new setting in this paper for a more comprehensive evaluation on current methods.

In the first setting, as previous works [7, 11, 18, 53], subsets CMU [6], BMLrub [40], and HDM05 [28] datasets are randomly divided into 90% for training and 10% for testing. Besides sparse observations of three joints, we also evaluate the performance of all compared methods by using four joints as input, including the root joint as an additional input, the same as in [18]. We term this setting as S1 in the following.

Tabs. 1 and 2 show that our method outperforms existing methods on most evaluation metrics, confirming its effectiveness. For the MPJVE metric, only AGRoL [11] surpasses our method when employing an offline strategy. In this scenario, specifically, AGRoL processes the entire sparse observation sequence in one pass and outputs the predicted full-body motions simultaneously. This enables

Method	MPJRE	MPJPE	MPJVE	Hand PE	Upper PE	Lower PE	Root PE	Jitter
Final IK [1]	16.77	18.09	59.24	-	-	-	-	-
LoBStTr [48]	10.69	9.02	44.97	-	-	-	-	-
VAR-HMD [10]	4.11	6.83	37.99	-	-	-	-	-
Avatarposer [18]	3.08	4.18	27.70	2.12	1.81	7.59	3.34	14.49
AvatarJLM [53]	2.90	3.35	20.79	1.24	1.42	6.14	2.94	8.39
AGRoL (Online) [11]	2.96	4.26	79.12	1.51	1.73	7.91	3.78	84.79
AGRoL (Offline) [11]	2.66	3.71	18.59	1.31	1.55	6.84	3.36	7.26
Ours	2.53	3.28	20.62	1.18	1.39	6.01	2.95	6.55

Table 1. Evaluation results under setting S1.




Figure 3. Visualization results compared with other methods. All models are trained under setting S1.

each position in the sequence to utilize the information from both preceding and subsequent time steps, offering an advantage in this particular metric. However, it’s important to note that, despite being competitive in metric numbers, offline inference has limited practical applicability in real-world scenarios where online processing capability is most important.

The second setting follows [5, 10, 11, 33, 53], where we evaluate the methods on a larger benchmark from AMASS [25]. The subsets [2, 4, 6, 12, 21, 23, 26, 26, 28,

40–42, 42] are for training, and Transition [25] and HumanEva [36] subsets are for testing. We term this setting as S2 in the following.

As shown in Tab. 3, our method achieves comparable performance with previous works on S2. However, we observe that the testing set of S2 is disproportionately small (i.e., only 1% of the training set). Such a small fraction cannot represent the overall data distribution of the large dataset and may not include sufficiently diverse motions to evaluate the models’ scalability, causing unconvincing eval-

Method	MPJRE	MPJPE	MPJVE
Final IK [1]	12.39	9.54	36.73
CoolMoves [3]	4.58	5.55	65.28
LoBStr [48]	8.09	5.56	30.12
VAE-HMD [10]	3.12	3.51	28.23
AvatarPoser [18]	2.59	2.61	22.16
AvatarJLM [53]	2.40	2.09	17.82
AGRoL [11]	2.25	2.17	16.26
Ours	2.10	1.88	14.79

Table 2. Evaluation results under setting S1 with the root joint as an additional input.

Method	MPJRE	MPJPE	MPJVE	Jitter
VAEHMD [†] [10]	-	7.45	-	-
Humor [†] [33]	-	5.50	-	-
FLAG [†] [5]	-	4.96	-	-
AvatarPoser [18]	4.70	6.38	34.05	10.21
AGRoL [11]	4.30	6.17	24.40	8.32
AvatarJLM [53]	4.30	4.93	26.17	7.19
Ours	4.62	5.86	33.54	7.13

Table 3. Evaluation result under setting S2. † indicates that these methods use additional inputs of pelvis location and rotation for training and inference, which are not directly comparable methods. The results of AvatarPoser [18] is provided by [11].




Figure 4. Visualization results on real data.

uation results. We introduce a new setting, S3, which adopts the same training and testing splitting ratio used in S1. In this setting, we randomly select 90% of the samples from the 15 subsets of S2 for training, while the remaining 10% are for testing. We train and evaluate the compared methods with this new setting. Table 4 reveals that under S3, the performance differences between the compared methods are more significant than S1 and S2. Since the test set has more




Figure 5. The visualization comparison for disentanglement. The darker the red color, the greater the deviation is between the predicted result and the ground truth.

diverse motions in S3, this benchmark evaluates the models’ scalability in a more objective way. In this context, our method outperforms existing methods in most metrics, especially in the critical metric of Lower PE, highlighting the superiority of our stratified design for lower-body modeling and inference.

Fig. 3 presents a visual comparison between our SAGE Net and baseline methods, all trained under the S1 protocol, which is commonly used by baselines for releasing their trained checkpoints. These visualizations demonstrate the significant improvements that our model offers in reconstructing the lower body. For example, in the first row of samples, baseline methods typically reconstruct the feet too close to the ground, restricting the avatar’s leg movements. Our model, however, overcomes this limitation, enabling more flexible leg movements. In the third row, for a subject climbing a ladder, the baseline methods often result in avatars with floating feet, failing to capture the detailed motion of climbing. In contrast, our SAGE Net accurately replicates complex foot movements, resulting in more realistic and precise climbing animations. We also evaluate our model on the real data, and for fair comparison, we directly use the real data release by [53]. As shown in Fig. 4, our method also achieves better reconstruction results on the real data.

4.3. Ablation Study

We perform ablation study under S1 to justify the design choice of each component in our SAGE Net.

Disentangled Codebook: We establish a baseline using a unified motion representation to evaluate the disentangle strategy. Specifically, we developed a full-body VQ-VAE model that encodes full-body motion into a single, unified discrete codebook. Other components are the same as the original model. Results shown in the first and the last rows in Table 5, demonstrate that our approach employing dis-

Method	MPJRE	MPJPE	MPJVE	Hand PE	Upper PE	Lower PE	Root PE	Jitter
AGRoL(Online) [11]	3.09	4.31	109.29	1.79	1.80	7.95	3.86	121.78
AGRoL(Offline) [11]	2.83	3.80	17.76	1.62	1.66	6.90	3.53	10.08
AvatarJLM [53]	3.14	3.39	15.75	0.69	1.48	6.13	3.04	5.33
Avatarposer [18]	2.72	3.37	21.00	2.12	1.63	5.87	2.90	10.24
Ours	2.41	2.95	16.94	1.15	1.28	5.37	2.74	5.27

Table 4. Evaluation results under setting S3.

Method	MPJRE	MPJPE	MPJVE	Jitter
w/o Disentangle	2.64	3.62	33.18	25.07
w/o Full-Body Decoder	2.71	3.69	26.07	10.80
w/o Refiner	2.54	3.26	21.99	9.29
5 Part Disentanglement	2.63	3.48	20.32	7.16
Ours	2.53	3.28	20.62	6.55

Table 5. Ablation results of different components in SAGE Net under setting S1.

Method	Lower PE	Jitter
Parallel Diffusion	6.73	14.71
Stratified Diffusion	6.46	10.83

Table 6. Evaluation results on the conditional strategy of the diffusion model under setting S1.

entangled latents significantly outperforms the baseline on all evaluation metrics. This demonstrates that the disentanglement can simplify the learning process by allowing the model to focus on a more limited set of movements and interactions. Additionally, Fig. 5 shows the visualization comparison between our model and baseline model, verifying that the disentangle can significantly improve the reconstruction results for the most challenging lower motions.

Full-Body Decoder and Refiner: The second and third rows of Tab. 5 demonstrate the impact of the full-body decoder and the refiner, respectively. Compared with utilizing the upper and lower decoder from VQ-VAE_{up} and VQ-VAE_{low}, the full-body decoder facilitates the integration of features from both the upper and lower body, improving the overall accuracy of full-body motion reconstruction. On the other hand, the refiner acts as a temporal memory, smoothing out the motion sequence to yield better visualization results.

Disentanglement Strategy: To investigate the optimal disentanglement strategy, we explore an extreme disentanglement configuration by following the path from the root (Pelvis) node to each leaf node along the kinematic tree. Specifically, we break down the body into five segments: the paths from the root to the left hand (a), right hand (b), head (c), left foot (d), and right foot (e). As reported in the last two rows of Tab. 5, the natural joint interconnections within the upper (or lower) body were disrupted when further disentangling the human body, resulting in perfor-

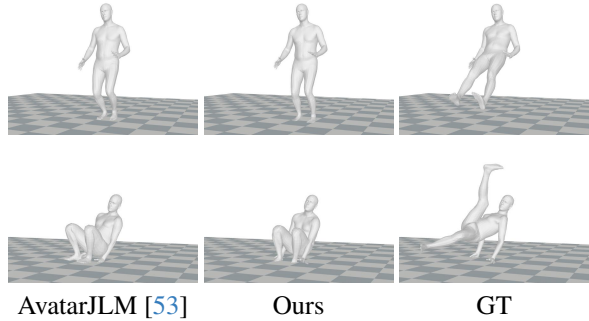


Figure 6. Failure cases. All models are trained under setting S1.

mance drops and complicating the model design.

Stratified Inference: Tab. 6 highlights the influence of our stratified design on the accuracy of lower body predictions. For comparison, we design a baseline that only uses the sparse observation for lower body latent generation without predicted upper body latent (term as Parallel Diffusion in the table). As we focus solely on the reconstruction quality of the lower body here, we use the decoding results on the generated lower latents from VQ-VAE_{low} to isolate the impact of other modules such as the full-body decoder and refiner. We report Lower PE and Jitter of the lower body for comparison. Results show that our stratified design markedly improves the accuracy of lower body predictions.

Limitation: In Fig. 6, both the previous state-of-the-art method and our model encounter difficulties in two main situations: (1) External Force-Induced Movements (the top row). (2) Unconventional Poses (the bottom row). The addition of more varied samples to the training dataset can potentially enhance the model’s performance in these areas.

5. Conclusion

We study the problem of human avatar generation from sparse observations. Our key finding is that the upper and lower body motions should be disentangled with respect to the input signals from the upper-body joints. Based on this, we propose a novel stratified solution where the upper-body motion is reconstructed first, and the lower-body motion is reconstructed next and conditioned on the upper-body motion. Our proposed stratified solution achieves superior performance on public available benchmarks.

References

- [1] *RootMotion Final IK.*, 2018. 2, 6, 7
- [2] Advanced Computing Center for the Arts and Design. ACCAD MoCap Dataset. 5, 6
- [3] Karan Ahuja, Eyal Ofek, Mar González-Franco, Christian Holz, and Andrew D. Wilson. Coolmoves: User motion accentuation in virtual reality. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 5(2): 52:1–52:23, 2021. 2, 7
- [4] Ijaz Akhter and Michael J. Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1446–1455, 2015. 5, 6
- [5] Sadeqh Aliakbarian, Pashmina Cameron, Federica Bogo, Andrew W. Fitzgibbon, and Thomas J. Cashman. FLAG: flow-based 3d avatar generation from sparse observations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13243–13252, 2022. 1, 2, 5, 6, 7
- [6] Carnegie Mellon University. CMU MoCap Dataset. 5, 6
- [7] Angela Castillo, Maria Escobar, Guillaume Jeanneret, Albert Pumarola, Pablo Arbeláez, Ali Thabet, and Arsiom Sanakoyeu. Bodiffusion: Diffusing sparse observations for full-body human motion synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 4221–4231, 2023. 2, 5
- [8] Ling-Hao Chen, Jiawei Zhang, Yewen Li, Yiren Pang, Xiaobo Xia, and Tongliang Liu. Humanmac: Masked motion completion for human motion prediction. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9510–9521, 2023. 2
- [9] Junyoung Chung, Çağlar Gülçehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *CoRR*, abs/1412.3555, 2014. 5
- [10] 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 *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11667–11677, 2021. 1, 2, 5, 6, 7
- [11] Yuming Du, Robin Kips, Albert Pumarola, Sebastian Starke, Ali K. Thabet, and Arsiom Sanakoyeu. Avatars grow legs: Generating smooth human motion from sparse tracking inputs with diffusion model. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 481–490, 2023. 2, 5, 6, 7, 8
- [12] Saeed Ghorbani, Kimia Mahdavian, Anne Thaler, Konrad P. Körding, Douglas James Cook, Gunnar Blohm, and Nikolaus F. Troje. Movi: A large multipurpose motion and video dataset. *CoRR*, abs/2003.01888, 2020. 5, 6
- [13] 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 *ACM International Conference on Multimedia (MM)*, pages 2021–2029, 2020. 2
- [14] Mohamed Hassan, Duygu Ceylan, Ruben Villegas, Jun Saito, Jimei Yang, Yi Zhou, and Michael J. Black. Stochastic scene-aware motion prediction. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11354–11364, 2021. 2
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [16] 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 Trans. Graph.*, 37(6):185, 2018. 1, 2
- [17] Biao Jiang, Xin Chen, Wen Liu, Jingyi Yu, Gang Yu, and Tao Chen. Motiongpt: Human motion as a foreign language. *Advances in Neural Information Processing Systems (NeurIPS)*, 36, 2024. 2
- [18] 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 *European Conference on Computer Vision (ECCV)*, pages 443–460, 2022. 2, 3, 5, 6, 7, 8
- [19] Yifeng Jiang, Yuting Ye, Deepak Gopinath, Jungdam Won, Alexander W. Winkler, and C. Karen Liu. Transformer inertial poser: Real-time human motion reconstruction from sparse imus with simultaneous terrain generation. In *SIGGRAPH Asia 2022 Conference Papers*, pages 3:1–3:9. ACM, 2022. 1, 2
- [20] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations (ICLR)*, 2014. 2
- [21] Matthew Loper, Naureen Mahmood, and Michael J. Black. Mosh: motion and shape capture from sparse markers. *ACM Trans. Graph.*, 33(6):220:1–220:13, 2014. 6
- [22] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: a skinned multi-person linear model. *ACM Trans. Graph.*, pages 248:1–248:16, 2015. 2, 3, 5
- [23] Eyes JAPAN Co. Ltd. Eyes Japan MoCap Dataset. 5, 6

- [24] Thomas Lucas, Fabien Baradel, Philippe Weinzaepfel, and Grégory Rogez. Posegpt: Quantization-based 3d human motion generation and forecasting. In *European Conference on Computer Vision (ECCV)*, pages 417–435, 2022. 2
- [25] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: archive of motion capture as surface shapes. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5441–5450, 2019. 2, 5, 6
- [26] Christian Mandery, Ömer Terlemez, Martin Do, Nikolaus Vahrenkamp, and Tamim Asfour. The KIT whole-body human motion database. In *IEEE International Conference on Advanced Robotics (ICAR)*, pages 329–336, 2015. 5, 6
- [27] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. Generating smooth pose sequences for diverse human motion prediction. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13289–13298, 2021. 2
- [28] M. Müller, T. Röder, M. Clausen, B. Eberhardt, B. Krüger, and A. Weber. Documentation mocap database HDM05. Technical Report CG-2007-2, Universität Bonn, 2007. 5, 6
- [29] Mathis Petrovich, Michael J. Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer VAE. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10965–10975, 2021. 2
- [30] Mathis Petrovich, Michael J. Black, and Gül Varol. TEMOS: generating diverse human motions from textual descriptions. In *European Conference on Computer Vision (ECCV)*, pages 480–497, 2022. 2
- [31] Huaijin Pi, Sida Peng, Minghui Yang, Xiaowei Zhou, and Hujun Bao. Hierarchical generation of human-object interactions with diffusion probabilistic models. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15061–15073, 2023. 2
- [32] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *CoRR*, 2022. 5
- [33] Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J. Guibas. Humor: 3d human motion model for robust pose estimation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11468–11479, 2021. 1, 5, 6, 7
- [34] Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning (ICML)*, pages 1530–1538, 2015. 2
- [35] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, 2022. 2, 5
- [36] Leonid Sigal, Alexandru O. Balan, and Michael J. Black. Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *Int. J. Comput. Vis.*, 87(1-2):4–27, 2010. 5, 6
- [37] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)*, 2021. 2
- [38] Guy Tevet, Brian Gordon, Amir Hertz, Amit H. Bermano, and Daniel Cohen-Or. Motionclip: Exposing human motion generation to CLIP space. In *European Conference on Computer Vision (ECCV)*, pages 358–374, 2022. 2
- [39] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit Haim Bermano. Human motion diffusion model. In *International Conference on Learning Representations (ICLR)*, 2023. 2, 5
- [40] Nikolaus F. Troje. Decomposing biological motion: A framework for analysis and synthesis of human gait patterns. *Journal of Vision*, 2(5):2–2, 2002. 5, 6
- [41] Matthew Trumble, Andrew Gilbert, Charles Malleson, Adrian Hilton, and John P. Collomosse. Total capture: 3d human pose estimation fusing video and inertial sensors. In *British Machine Vision Conference (BMVC)*, 2017.
- [42] Simon Fraser University and National University of Singapore. SFU Motion Capture Database. 5, 6
- [43] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 6306–6315, 2017. 2, 3, 5
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 5998–6008, 2017. 5
- [45] Timo von Marcard, Bodo Rosenhahn, Michael J. Black, and Gerard Pons-Moll. Sparse inertial poser: Automatic 3d human pose estimation from sparse imus. *Comput. Graph. Forum*, 36(2):349–360, 2017. 1, 2
- [46] Alexander W. Winkler, Jungdam Won, and Yuting Ye. Questsim: Human motion tracking from sparse sensors with simulated avatars. In *SIGGRAPH Asia 2022 Conference Papers*, pages 2:1–2:8. ACM, 2022. 2

- [47] Sirui Xu, Zhengyuan Li, Yu-Xiong Wang, and Liang-Yan Gui. Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2
- [48] Dongseok Yang, Doyeon Kim, and Sung-Hee Lee. Lobstr: Real-time lower-body pose prediction from sparse upper-body tracking signals. *Comput. Graph. Forum*, 40(2):265–275, 2021. 2, 6, 7
- [49] Xinyu Yi, Yuxiao Zhou, and Feng Xu. Transpose: real-time 3d human translation and pose estimation with six inertial sensors. *ACM Trans. Graph.*, 40(4): 86:1–86:13, 2021. 1, 2
- [50] 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)*, pages 13157–13168, 2022. 1, 2
- [51] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [52] Chuanxia Zheng and Andrea Vedaldi. Online clustered codebook. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22741–22750, 2023. 3
- [53] Xiaozheng Zheng, Zhuo Su, Chao Wen, Zhou Xue, and Xiaojie Jin. Realistic full-body tracking from sparse observations via joint-level modeling. In *IEEE/CVF international conference on computer vision (ICCV)*, 2023. 2, 5, 6, 7, 8
- [54] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5745–5753, 2019. 3