

RoHM: Robust Human Motion Reconstruction via Diffusion

Siwei Zhang^{1,2*} Bharat Lal Bhatnagar² Yuanlu Xu² Alexander Winkler² Petr Kadlecek²
 Siyu Tang¹ Federica Bogo²
¹ETH Zürich ²Meta Reality Labs Research
 {siwei.zhang, siyu.tang}@inf.ethz.ch
 {bharatbhatnagar, yuanluxu, winklera, petr.kadlecek, fbogo}@meta.com




Figure 1. Our method robustly reconstructs smooth and complete 3D human motion from different inputs, such as incomplete and noisy motion estimates (left), RGB-D (middle) and RGB (right) monocular videos. We learn diffusion-based models to denoise and infill both **root trajectory** in global space and local motion in body-root space for **visible** and **occluded** joints, predicting whether feet are **in contact** or **not** with the ground for improved physical plausibility. Compared with baselines such as HuMoR [67], our method reconstructs more plausible motions that faithfully match image evidence, especially under heavy occlusions.

Abstract

We propose RoHM, an approach for robust 3D human motion reconstruction from monocular RGB(-D) videos in the presence of noise and occlusions. Most previous approaches either train neural networks to directly regress motion in 3D or learn data-driven motion priors and combine them with optimization at test time. The former do not recover globally coherent motion and fail under occlusions; the latter are time-consuming, prone to local minima, and require manual tuning. To overcome these shortcomings, we exploit the iterative, denoising nature of diffusion models. RoHM is a novel diffusion-based motion model that, conditioned on noisy and occluded input data, reconstructs complete, plausible motions in consistent global coordinates. Given the complexity of the problem – requiring one to address different tasks (denoising and infilling) in different solution spaces (local and global motion) – we decompose it into two sub-tasks and learn two models, one for global trajectory and one for local motion. To capture the correlations between the two, we then introduce a novel conditioning module, combining it with an iterative inference scheme. We apply RoHM to a variety of tasks

– from motion reconstruction and denoising to spatial and temporal infilling. Extensive experiments on three popular datasets show that our method outperforms state-of-the-art approaches qualitatively and quantitatively, while being faster at test time. The code is available at <https://sanweiliti.github.io/ROHM/ROHM.html>.

1. Introduction

In this paper, we tackle the problem of 3D human motion reconstruction from monocular RGB(-D) videos in real scenarios – i.e., in the presence of noise and occlusions. Reconstructing 3D human motions is crucial for many applications, ranging from augmented and virtual reality to robotics. Many methods in the literature tackle the problem by training deep neural networks to directly regress 3D body pose and shape from monocular input [10, 19, 34, 39, 59]. However, these approaches commonly suffer from two major shortcomings: 1) they estimate only *local* motion, representing pose in body-root relative coordinates without plausible *global* motion, in world coordinates consistent over time; 2) they lack robustness when the body undergoes occlusions, in either the spatial or temporal dimension.

In such scenarios, optimization-based methods [67, 72, 100] have shown better performance. For instance, Hu-

*The work was done during an internship at Meta.

MoR [67] and PhaseMP [72] explicitly address scenarios with noisy and incomplete input by combining data-driven motion priors with application-agnostic optimizers such as L-BFGS [60]. Still, these methods may fail under heavy occlusions (Fig. 1). Furthermore, test-time optimization is time-consuming, prone to local minima, and requires significant manual tuning [11, 75].

To overcome these limitations, we propose to leverage the iterative, generative nature of diffusion models. Diffusion models were initially proposed for 2D generation tasks [12, 24, 25, 76]; recently, they achieved compelling results in 3D human motion generation from input such as text and action labels [35, 71, 82, 95], music [83], and sparse (noise-free) keypoints [6, 13, 71]. In particular, these models proved effective in learning and modelling long-term motion dependencies over time [82], which go beyond the single-frame conditioning considered in [67, 72]. Furthermore, their iterative denoising process poses them as a data-driven alternative to the iterative minimization of optimization-based techniques. However, so far diffusion-based approaches have mostly focused on *synthesizing* motion from user input, rather than *reconstructing* motion from monocular videos exhibiting noise and occlusions – where reconstructions need to match image evidence whenever available. Here, we explore how to leverage diffusion models in such reconstruction scenarios.

Given a monocular video, we obtain initial, per-frame 3D pose estimates using off-the-shelf regressors [15, 49, 70] and/or per-frame optimization (similar to [67, 72, 100]). These estimates are inaccurate and incomplete, with implausible motions. From them, our goal is to reconstruct a smooth and complete 3D motion. This is a complex task, requiring us to address different problems (motion denoising and infilling) in different solution spaces (local and global motion). We observe that previously proposed diffusion-based motion models [71, 82] do not work well in this scenario: they expect noise-free, user-provided keypoints as input, model global and local motion together, and cannot ensure alignment between reconstruction and image evidence. Inspired by [79, 94], we therefore propose to decompose the problem. We leverage motions from the AMASS dataset [57] to learn two diffusion models conditioned on noisy input, one for global trajectory reconstruction and one for local motion prediction – showing the benefits of decoupling the two spaces. Still, this formulation ignores the correlations between global and local motion space. While [94] addresses this by predicting trajectory based only on infilled local motion, in our scenario we require the estimated trajectory to remain faithful to the input. Drawing inspiration from [98], we therefore propose a flexible conditioning strategy for trajectory reconstruction, exploiting both input data and denoised local motion. We show that this module, combined with an iterative scheme

at inference time, improves both local and global motion quality. Finally, to further encourage physically plausible motions that match image evidence, we guide sampling at test time with physics-based and image-based scores.

In summary, our contributions are: 1) **RoHM**, a novel diffusion-based approach for **Robust Human Motion** reconstruction in consistent global coordinates from monocular sequences with noise and occlusions; 2) a flexible conditioning strategy to capture inter-dependencies between root trajectory and local pose; 3) various applications enabled by RoHM, including motion reconstruction, denoising, spatial and temporal infilling. Extensive experiments on three widely used datasets show that RoHM achieves superior accuracy and realism compared to state-of-the-art optimization-based methods, while being 30 times faster at inference time.

2. Related Work

Regression-based methods. Many approaches in the literature focus on 3D human shape and pose reconstruction from a single image [9, 14, 15, 33, 41–46, 49–51, 75, 86, 90, 97], recently also considering robustness to occlusions [38, 40, 48, 53, 69, 102, 104]. Dealing with occlusions is even more challenging for monocular human motion reconstruction, requiring one to model plausible dynamics over time for non-visible body parts. Many approaches train neural networks to predict 3D motion from RGB videos and are not easily adaptable to different input modalities [8, 10, 16, 19, 34, 39, 55, 59, 63, 65, 78, 88, 92, 96]. Some methods introduce adversarial priors [34, 39], leverage multi-view cues [63] or learn denoising models [59] to achieve robustness. Still, most of them estimate only local motion in the camera frame without recovering global trajectory, thus suffering from jitter and motion artifacts. [47, 79, 94] estimate plausible global trajectories from per-frame local features, but are not robust to occlusions [91]. In contrast, our method robustly recovers both global and local motion and can be applied to different inputs.

Optimization-based methods. These methods typically fit a parametric body model [54] to observations (such as body keypoints, depth, silhouettes *etc.*) by iteratively minimizing an objective function [3, 62, 67, 72, 91, 100]. To regularize motion, such function contains one or more temporal priors. Some approaches define hand-crafted priors encouraging motion smoothness (*e.g.*, on body joint velocity or acceleration) [1, 58], or constrain motion in a low-dimensional space [30]; however, they may produce over-smooth motions and foot skating [67]. Lately, the availability of large-scale motion capture datasets such as AMASS [57] made it possible to learn powerful motion models from data [67, 72, 100]. LEMO [100] learns two separate, fully deterministic priors for motion smoothness and infilling. HuMor [67] and PhaseMP [72] propose gen-

erative priors, modelling the distribution of state transitions between frames via Conditional Variational Autoencoders (CVAEs) [4, 52]. Learning transitions only between pair of frames, these methods struggle when occlusions span long time intervals (see Sec. 5). Optimization-based methods tend to match input data more closely than regression methods, thanks to their iterative minimization; but they are in general slower, prone to local minima, and require significant manual tuning [11, 75]. In contrast, we propose to leverage the iterative nature of diffusion models.

Human motion models. A variety of approaches has been proposed for motion tracking and synthesis, including mixtures-of-Gaussians [27], Gaussian processes [84], pose embeddings [61, 64, 81, 85], VAEs [52, 103], 2D CNN [37], and normalizing flows [23]. These methods may not generalize well to unseen motions [67] and body-scene interactions. Physics-based methods [17, 18, 28, 55, 56, 66, 73, 89, 93] address these challenges by enforcing physics laws via simulation. However, simulators are computationally expensive, non-differentiable, and may introduce discrepancies between predicted motions and observed input.

Diffusion models for human motion. Diffusion models have demonstrated compelling results for human motion synthesis conditioned on input such as text and action labels [7, 32, 35, 36, 82, 87, 95, 99], music [83], environment [29, 68], and (noise-free) 3D joints [71]. Given their ability to model long-range spatio-temporal relationships, they have been applied to motion forecasting [2, 80] and infilling [71, 82]. [6, 13] use diffusion to synthesize lower body motion given head and wrist positions and rotations. These methods do not focus on scenarios with noisy input or body occlusion varying over time. Recently, diffusion has been leveraged for 3D body pose estimation [16, 20], even under severe body truncations [102], but without considering the temporal dimension. In contrast, we leverage diffusion models to achieve robustness against varying ranges of occlusion and noise over long temporal sequences.

3. Preliminaries

SMPL-X [62] is a parametric model that represents the human body as a function $\mathcal{M}(\gamma, \Phi, \theta, \beta, \theta_h, \phi)$, parameterized by global translation γ , global orientation Φ , body pose θ , body shape β , hand pose θ_h and facial expression ϕ . The function returns a triangle mesh \mathcal{M} with 10,475 vertices. In the following we do not use θ_h and ϕ and consider only the main body parameters, with a total of 22 body joints. We use SMPL-X instead of SMPL [54] to leverage the gender-neutral annotations from AMASS [57].

Conditional Diffusion Models. We adopt the Denoising Diffusion Probabilistic Models (DDPMs) formulation [25]. At the core of it are a forward diffusion process and an inverse, iterative denoising process. Let $\mathbf{X}_0 \sim q(\mathbf{X}_0)$ be the real motion distribution. The forward diffusion process is

a Markov chain adding *i.i.d.* Gaussian noise at each step $t \in \{1, \dots, T\}$:

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

where $\alpha_t \in [0, 1]$ defines the variance at each step according to a pre-defined schedule and \mathbf{I} is the identity matrix. [25] shows that \mathbf{X}_t can be directly sampled from \mathbf{X}_0 :

$$\mathbf{X}_t = \sqrt{\alpha_t} \mathbf{X}_0 + \sqrt{1 - \alpha_t} \epsilon, \text{ where } \epsilon \sim \mathcal{N}(0, \mathbf{I}). \quad (2)$$

Starting from Gaussian noise \mathbf{X}_T , the inverse diffusion process reconstructs \mathbf{X}_0 by iteratively denoising \mathbf{X}_T over T steps. In practice, we train a denoiser neural network $D(\cdot)$ to remove the added Gaussian noise based on condition signal \mathbf{c} and step t : $\hat{\mathbf{X}}_0 = D(\mathbf{X}_t, t, \mathbf{c})$ (cf. [82]). Leveraging Eq. (2), the denoiser can be trained by sampling noise step t and optimizing the simple objective [25]:

$$\mathcal{L}_{\text{simple}} = E_{\mathbf{X}_0 \sim q(\mathbf{X}_0 | \mathbf{c}), t \sim [1, T]} [\|\mathbf{X}_0 - D(\mathbf{X}_t, t, \mathbf{c})\|_2^2]. \quad (3)$$

In the following, we use the subscript t to denote the diffusion step and n to denote the frame in the motion sequence.

Note that conditioning is crucial in our setup, where we want to exploit input observations whenever available.

4. Method

Our goal is to reconstruct realistic motions from monocular RGB(-D) sequences in the presence of noise and occlusions. We use off-the-shelf, per-frame regressors [15, 49, 70] and/or per-frame optimization to obtain initial SMPL-X estimates for each frame (see Sec. 5.4 and Supp. Mat. for more details). We stack these estimates into a motion sequence $\hat{\mathbf{X}} \in \mathbb{R}^{N \times d}$, where N is the number of frames and d the dimensionality of our representation. Such estimates are noisy, inaccurate under occlusions, and temporally inconsistent. Given $\hat{\mathbf{X}}$, we aim to generate realistic motions $\hat{\mathbf{X}}_0$ in consistent global coordinates. As reconstructing simultaneously global trajectory *and* local articulated pose is challenging, we learn their dynamics with two diffusion-based models, *TrajNet* and *PoseNet* (Sec. 4.1). To capture correlations between the two and further refine motion plausibility, we introduce *TrajControl*, a flexible conditioning module (Sec. 4.2) that we leverage with an iterative schedule at inference time (Sec. 4.3). We describe training objectives in Sec. 4.4. Fig. 2 shows an overview of our approach.

Motion Representation. We represent an (input/output) motion sequence as $\mathbf{X} = (\mathbf{R}, \mathbf{P})$, where $\mathbf{R} \in \mathbb{R}^{N \times d_R}$, $\mathbf{P} \in \mathbb{R}^{N \times d_P}$ denote root trajectory and local body features, respectively. At each frame n , $\mathbf{x}^n = (\mathbf{r}^n, \mathbf{p}^n)$ includes: (1) global trajectory features \mathbf{r}^n , including root (pelvis) linear position $\mathbf{r}^l \in \mathbb{R}^2$, root angular rotation $\mathbf{r}^a \in \mathbb{R}$, root height $\mathbf{r}^z \in \mathbb{R}$, SMPL-X global translation $\gamma \in \mathbb{R}^3$ and global orientation $\Phi \in \mathbb{R}^6$; (2) local body features \mathbf{p}^n , including

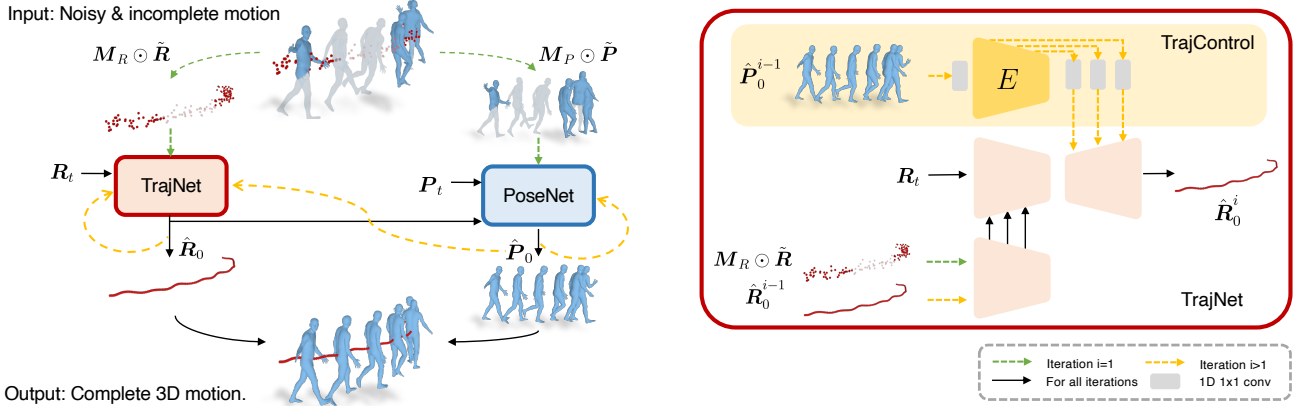


Figure 2. **Overview of our approach.** Given an initial noisy motion sequence $\tilde{X} = (\tilde{R}, \tilde{P})$ and the corresponding root/body joint occlusion masks M_R and M_P , we employ two diffusion-based models, TrajNet and PoseNet, to estimate global root trajectory \hat{R}_0 and local pose \hat{P}_0 , separately (Sec. 4.1). We leverage an additional conditioning module, TrajControl, to fine-tune TrajNet and flexibly condition it on denoised local pose \hat{P}_0 , leading to improved trajectory reconstruction (Sec. 4.2). At inference time, TrajNet, PoseNet, and TrajControl are leveraged in an iterative inference scheme to refine local and global motion (Sec. 4.3).

local joint positions $J \in \mathbb{R}^{21 \times 3}$, joint rotations $\theta \in \mathbb{R}^{21 \times 6}$, body shape $\beta \in \mathbb{R}^{10}$, and foot contact labels $f \in \{0, 1\}^4$:

$$r^n = (r^l, \dot{r}^l, r^a, \dot{r}^a, r^z, \gamma, \dot{\gamma}, \Phi, \dot{\Phi}), \quad (4)$$

$$p^n = (J, \dot{J}, \theta, \beta, f), \quad (5)$$

where the dot indicates the first derivative (velocity). For rotations, we adopt the 6D representation from [105]. See Supp. Mat for more details. For f , we select two joints per foot and set the corresponding contact label as 1 if the joint is in contact with the ground, else 0 [66, 97]. For each frame n , we define a local coordinate system such that local joint positions are relative to the current frame pelvis joint, projected on the ground [21, 26, 82, 100]. This over-parameterized representation allows explicit modelling of both 3D skeleton joints and SMPL-X meshes, such that they can benefit the one from the other. Together with R and P , we define root joint and local joint visibility masks $M_R \in \{0, 1\}^{N \times d_R}$ and $M_P \in \{0, 1\}^{N \times d_P}$, respectively (1 denotes visible, 0 otherwise). They are obtained by randomly masking joints at training time and computing joint visibility at test time (see Sec. 4.4 and Supp. Mat.).

4.1. Diffusing Global and Local Motion

We tackle the problem of denoising and infilling global and local motion by using two networks, *TrajNet* and *PoseNet*.

TrajNet. Given a noisy root trajectory \tilde{R} and root occlusion mask M_R , we train a denoiser $D_R(\cdot)$ to recover smooth and plausible global trajectory \hat{R}_0 :

$$\hat{R}_0 = D_R(R_t, t, c_R), \text{ where } c_R = M_R \odot \tilde{R}. \quad (6)$$

R_t denotes root trajectory at each diffusion denoising step t and \odot denotes element-wise matrix multiplication. The

trajectory representation $(R_t, \tilde{R}, \hat{R}_0)$ for TrajNet is parameterized as $(r^l, r^a, r^z, \gamma, \Phi)$, excluding first derivatives to avoid global drifting caused by inaccurate velocities.

PoseNet. Given denoised, infilled trajectory \hat{R}_0 from TrajNet, noisy local motion \tilde{P} and joint occlusion mask M_P , we train a denoiser $D_P(\cdot)$ to recover smooth and plausible local motion \hat{P}_0 :

$$(\hat{R}_0, \hat{P}_0) = D_P((\hat{R}_0, P_t), t, c_P), \quad (7)$$

where P_t denotes local motion at each diffusion denoising step t . The conditional signal $c_P = (\hat{R}_0, M_P \odot \tilde{P})$ includes TrajNet’s output and the corrupted initial local motion. To leverage the clean and complete information from TrajNet, we use the reconstructed trajectory \hat{R}_0 as both input and output for $D_P(\cdot)$, overwriting the global motion part with \hat{R}_0 at each diffusion step.

Architectures. TrajNet adopts a U-Net encoder-decoder structure built upon [31, 68]. An extra conditioning encoder maps the conditional trajectory signal into multi-layer features, which are concatenated with U-Net encoder features at each intermediate layer. PoseNet employs a transformer encoder structure akin to [82]. The conditioning signal c_P is encoded via an MLP (Multilayer Perceptron) and fed into the transformer. See Supp. Mat. for more details.

4.2. Controlling Global Motion Reconstruction

Learning two distinct models for global and local motion does not capture the correlations between the two. In particular, we observe that, under significant noise, TrajNet outputs might still lead to physically implausible motions with foot skating (see Sec. 5.5). [94] proposes to first infill local pose and then use that to *predict* trajectory. However, in our case we want to keep the *estimated* trajectory as close as possible to input observations. Conditioning TrajNet on the

initial corrupted local motion $\tilde{\mathbf{P}}$ is too challenging, since such initial poses are noisy and incomplete. Our solution is to adapt TrajNet such that it can be *flexibly* conditioned on local motion estimates, when they are available.

Inspired by [98], we introduce *TrajControl*, an auxiliary module for fine-tuning TrajNet with additional control signal from local body pose (Fig. 2 right). Specifically, we freeze the parameters of our pre-trained TrajNet and clone the encoder layers to a trainable duplicate, $E(\cdot)$, which serves as the *TrajControl* module. Iteratively, we feed the output of PoseNet into $E(\cdot)$, thus improving TrajNet output based on denoised, complete local motion. In turn, the improved TrajNet output can be leveraged to further refine local body motion. We detail this iterative scheme in Sec. 4.3. **Architecture.** $E(\cdot)$ is connected to the frozen pre-trained TrajNet with zero convolution layers (1D convolution layer with kernel size 1, initialized from zero bias and zero weight), to ensure a smooth start for the fine-tuning. During fine-tuning, we update only the weights of $E(\cdot)$. In this way, the fine-tuned TrajNet can still be conditioned on noisy trajectory only, when clean local motion is not available.

4.3. Inference

Iterative inference. We leverage TrajNet, PoseNet, and TrajControl to iteratively refine local and global motion at inference time. Given the initial noisy/occluded motion, we first run our ‘vanilla’ TrajNet and PoseNet sequentially. TrajNet is conditioned on the initial noisy input root trajectory $\tilde{\mathbf{R}}$ and root joint occlusion mask \mathbf{M}_R ; PoseNet is conditioned on the denoised trajectory $\hat{\mathbf{R}}_0$, noisy local motion input $\tilde{\mathbf{P}}$, and body joint visibility mask \mathbf{M}_P :

$$\mathbf{c}_R = \mathbf{M}_R \odot \tilde{\mathbf{R}}, \quad (8)$$

$$\mathbf{c}_P = (\hat{\mathbf{R}}_0^i, \mathbf{M}_P \odot \tilde{\mathbf{P}}). \quad (9)$$

For any subsequent iteration i , TrajNet is conditioned on estimated trajectory $\hat{\mathbf{R}}_0^{i-1}$ and local pose $\hat{\mathbf{P}}_0^{i-1}$, output of PoseNet at iteration $i-1$, as the additional control signal; PoseNet is conditioned on current TrajNet output $\hat{\mathbf{R}}_0^i$ and local pose $\hat{\mathbf{P}}_0^{i-1}$ predicted at iteration $i-1$:

$$\mathbf{c}_R = (\hat{\mathbf{R}}_0^{i-1}, E(t, \hat{\mathbf{P}}_0^{i-1})), \quad (10)$$

$$\mathbf{c}_P = (\hat{\mathbf{R}}_0^i, \hat{\mathbf{P}}_0^{i-1}). \quad (11)$$

Algorithm 1 summarizes the approach. Note that at each ‘iteration’ i ($i \in \{1, \dots, K\}$), we sample from TrajNet and PoseNet running all T diffusion denoising steps. T is set to 100 for TrajNet and 1000 for PoseNet. Empirically, we find that $K = 2$ iterations are sufficient to obtain accurate results. While one could still run this iterative inference between TrajNet and PoseNet without using TrajControl, we show in Sec. 5.5 that this leads to degraded results.

Algorithm 1: Iterative inference with TrajControl.

Result: Reconstructed motion $(\hat{\mathbf{R}}_0^K, \hat{\mathbf{P}}_0^K)$
Init: Noisy motion $(\tilde{\mathbf{R}}, \tilde{\mathbf{P}})$, root occlusion mask \mathbf{M}_R , body joint occlusion mask \mathbf{M}_P , TrajNet $D_R(\cdot)$, PoseNet $D_P(\cdot)$, TrajControl $E(\cdot)$;
for $i = 1 : K$ **do**
 compute $\hat{\mathbf{R}}_0^i, \hat{\mathbf{P}}_0^i$ by Eq. (6) (7) (8) (9) if $i = 1$;
 compute $\hat{\mathbf{R}}_0^i, \hat{\mathbf{P}}_0^i$ by Eq. (6) (7) (10) (11) if $i > 1$;
end

Score-guided sampling. Besides embedding condition signals in the decoder architecture, diffusion models enable also test-time conditioning via classifier-based guidance, which has been already leveraged for image generation [12, 74, 77], trajectory prediction [68], and human mesh recovery [102]. In a similar spirit, we enhance physical plausibility and accuracy of reconstructed motions by guiding PoseNet sampling with two scores, penalizing foot skating and enforcing 2D joint reprojection consistency:

$$\mathcal{J}_{\text{skate}}(\hat{\mathbf{P}}_0) = \|\hat{\mathbf{f}}_0 \dot{J}_{3D}^{\text{foot}}(\hat{\mathbf{R}}_0, \hat{\mathbf{P}}_0)\|^2, \quad (12)$$

$$\mathcal{J}_{2D}(\hat{\mathbf{P}}_0) = \|c_{\text{conf}}(\Pi_{\mathcal{K}}(J_{3D}(\hat{\mathbf{R}}_0, \hat{\mathbf{P}}_0)) - J_{\text{det}})\|^2, \quad (13)$$

where J_{3D} and \dot{J}_{3D} denote body 3D joint positions and velocities obtained via forward kinematics, respectively. J_{det} and c_{conf} denote 2D body keypoints and their confidence scores, obtained by running OpenPose [5] on input images; $\Pi_{\mathcal{K}}$ is the 3D-to-2D projection to image space with camera intrinsics \mathcal{K} . The superscript ‘foot’ identifies foot joints. $\hat{\mathbf{f}}_0$ denotes the foot contact labels predicted by PoseNet, so that $\mathcal{J}_{\text{skate}}$ penalizes foot joint velocity when the joint is predicted as touching the ground [67, 100]. The gradients $\nabla \mathcal{J}_{\cdot}(\hat{\mathbf{P}}_0)$ effectively guide the diffusion sampling to further alleviate foot skating and better align reconstructed motion to image observations (if available). At each sampling step t , PoseNet predicts $\hat{\mathbf{P}}_0$ by Eq. (7), which is then noised back to \mathbf{P}_{t-1} by sampling from the Gaussian distribution:

$$\mathbf{P}_{t-1} \sim \mathcal{N}(\mu_t(\mathbf{P}_t, \hat{\mathbf{P}}_0) + (\lambda_{\text{skate}} \nabla \mathcal{J}_{\text{skate}} + \lambda_{2D} \nabla \mathcal{J}_{2D}) \Sigma_t, \Sigma_t), \quad (14)$$

with μ_t as a linear combination of \mathbf{P}_t and $\hat{\mathbf{P}}_0$. The guidance is modulated by Σ_t , the variance of a pre-scheduled Gaussian distribution [25], and by the scaling weights λ_{skate} and λ_{2D} .

4.4. Training

We train our diffusion denoisers $D_R(\cdot)$ and $D_P(\cdot)$ using $\mathcal{L}_{\text{simple}}$ in Eq. (3), plus losses enforcing consistency with the ground truth in terms of 3D joint position ($\mathcal{L}_{J_{3D}}$) and 3D joint velocity (\mathcal{L}_{vel}), and penalizing foot skating ($\mathcal{L}_{\text{skate}}$):

$$\mathcal{L}_{J_{3D}} = \|J_{3D}(\mathbf{R}_0, \mathbf{P}_0) - J_{3D}(\mathbf{R}, \hat{\mathbf{P}}_0)\|^2, \quad (15)$$

$$\mathcal{L}_{\text{vel}} = \|\dot{J}_{3D}(\mathbf{R}_0, \mathbf{P}_0) - \dot{J}_{3D}(\mathbf{R}, \hat{\mathbf{P}}_0)\|^2, \quad (16)$$

$$\mathcal{L}_{\text{skate}} = \|\mathbf{f}_0 \dot{J}_{3D}^{\text{foot}}(\mathbf{R}_0, \hat{\mathbf{P}}_0)\|^2, \quad (17)$$

where $(\mathbf{R}_0, \mathbf{P}_0)$ is the ground-truth motion; \mathbf{R} refers to ground-truth root trajectory \mathbf{R}_0 for PoseNet, and predicted root trajectory $\hat{\mathbf{R}}_0$ for TrajNet. \mathbf{f}_0 denotes ground-truth foot contact labels, and $\hat{J}_{3D}^{\text{foot}}$ denotes the predicted foot joint velocities. The overall loss is defined as:

$$\mathcal{L} = \mathcal{L}_{\text{simple}} + \lambda_{J_{3D}} \mathcal{L}_{J_{3D}} + \lambda_{\text{vel}} \mathcal{L}_{\text{vel}} + \lambda_{\text{skate}} \mathcal{L}_{\text{skate}}, \quad (18)$$

PoseNet and TrajNet are trained separately on the AMASS dataset [57], with each sequence trimmed into short clips of $N = 144$ frames. For both training the “vanilla” TrajNet and fine-tuning TrajNet with TrajControl, we exclude $\mathcal{L}_{\text{skate}}$, and only compute $\mathcal{L}_{J_{3D}}$ and \mathcal{L}_{vel} for the root joint. We utilize ground-truth local pose \mathbf{P}_0 as the control input of TrajControl to fine-tune TrajNet.

During training, we synthesize noisy and partially occluded motion $\tilde{\mathbf{X}}$ by corrupting ground-truth motion \mathbf{X}_0 : we add Gaussian noise to SMPL-X parameters, obtaining noisy joint positions by forward kinematics, and mask out subsets of joints. We employ a curriculum training scheme by gradually increasing noise levels and occlusion rates as training progresses. See Supp. Mat. for more details.

5. Experiments

5.1. Datasets

AMASS [57] is a large-scale motion capture dataset collecting high-quality 3D human pose and shape annotations. We use the official SMPL-X neutral body annotations for training and evaluation. We downsample each sequence to 30fps. As in [100], we use TCD_handMocap, TotalCapture, and SFU for testing and the remaining subsets for training. PROX [22] collects monocular RGB-D videos of people interacting with various 3D indoor scenes. Since the dataset does not provide ground-truth annotations, we use a subset of sequences to evaluate physical plausibility as in [67, 72]. EgoBody [101] collects sequences of people interacting with each other in various 3D indoor environments, capturing multi-modal input streams with both head-mounted (first-person) and external (third-person) cameras. The dataset provides ground-truth SMPL/SMPL-X annotations. We manually select a set of third-person RGB sequences (around 24k frames) exhibiting severe human-scene occlusions, and use them for evaluation.

5.2. Evaluation Metrics

Accuracy. We adopt the Mean Per-Joint Position Error in mm to evaluate predicted body joint accuracy in the pelvis-aligned coordinate system (*MPJPE*) and in the *global* coordinate system (*GMPJPE*). We report numbers for full-body (*all*), visible (*vis*) and occluded (*occ*) body joints separately, considering the 22 SMPL-X main body joints. Furthermore, we measure foot-ground contact binary classification accuracy (*Contact*) for the 4 foot joints as in [67].

Physical Plausibility. We report additional metrics to assess motion and scene-interaction plausibility. When ground-truth motion is available (AMASS and EgoBody), we report the acceleration error (*Accel*) as the difference in acceleration between predicted and ground-truth 3D joints [39]; for PROX, we report the norm of mean per-joint acceleration ($\|Accel\|$). Both metrics are in m/s^2 . Foot skating ratio (*Skating*) measures how often feet slide on the floor. Following [100], we define sliding as happening when the velocity of all foot joints exceeds 10cm/s, toe joints’ height above the ground is lower than 10cm, and ankle joints’ height is lower than 15cm. The mean ground penetration distance (*Dist*) [67, 72] measures to what extent the toe joints penetrate into the ground, measured in mm .

5.3. Motion from 3D Observations

Experimental setup. To evaluate RoHM’s robustness to noisy and occluded data, we run two sets of experiments on AMASS: (1) motion denoising + infilling (**Occ-L.**), and (2) motion denoising + in-betweening (**Occ-10%**). Given a SMPL-X motion sequence from our AMASS test set, in (1) we mask out all lower body joint parameters (both positions and rotations), simulating scenarios commonly occurring when people move in a 3D scene; in (2), we mask out an entire subsequence of frames (10% of the original sequence), thus requiring the model to generate the in-between motion. In both setups, we add Gaussian noise to SMPL-X pose and translation parameters and use the resulting noisy and occluded 3D motion data as input for our model. We consider different, increasing noise levels: noise level k corresponds to Gaussian noise with standard deviation of $(k^\circ, k^\circ, k \text{ cm})$ for (Φ, θ, γ) . Note that the noise, defined on SMPL-X parameters, will accumulate along the kinematic tree.

Baselines. We compare RoHM with VPoser-t, HuMoR [67], and an adapted version of MDM [82] (‘MDM++’). As in [67, 72], VPoser-t is an optimization-based method using VPoser [62] and 3D joint smoothing [30]. It serves as the initialization stage for HuMoR. We cannot directly apply MDM/PriorMDM [71] to our setup, since they require noise-free visible joints for infilling and do not support explicit conditioning on noisy observations – resulting in both pose and trajectory drifting. We therefore train an adapted and improved version, which allows conditioning on the initial corrupted motion, using the same data augmentation as RoHM (see Supp. Mat. for details).

Results. Tab. 1 reports results on the AMASS test set. Our approach demonstrates significantly superior performance, in both accuracy and physical plausibility. Specifically, when compared to HuMoR, we achieve $\geq 48\%$ reduction in GMPJPE for occluded body parts in both **Occ-L.** and **Occ-10%** setups. The reduced acceleration errors suggest RoHM can recover more realistic motion dynamics. This facilitates also accurate foot contact label predictions,




Figure 3. **Model performance wrt different input noise levels** for the **Occ-L** setup on the AMASS test set.




Figure 4. **Qualitative results on AMASS.** Given noisy input with occluded lower body, we reconstruct more accurate and realistic motions (row 1-2), with fewer foot-ground penetrations (row 3) than the baseline method.

leading to $\geq 44\%$ improvement in foot skating over HuMoR and fewer foot-ground inter-penetrations (Fig. 4, row 3). MDM++ performs similarly to us in foot-ground collisions, but reconstruction accuracy and other plausibility metrics are noticeably inferior in comparison. In scenarios with larger noise (e.g., level 7), baselines struggle to recover plausible lower body motions, often leading to legs floating in the air (see Fig. 4, row 1-2). This is particularly evident for VPoser-t, which therefore exhibits a low skating ratio, as the skating score only considers frames with foot-ground contact. Fig. 3 (left, middle) compares robustness to noise of our method and HuMoR. Increasing input noise levels corresponds to a substantial decline in performance for HuMoR, while RoHM shows more robustness. Note that our method is only trained with a noise level of 2.

5.4. Motion from RGB(-D) Videos

We compare the performance of RoHM against baselines on motion reconstruction from RGB/RGB-D videos on PROX, and from RGB videos on EgoBody.

Initialization. On PROX, we obtain initial noisy motion estimates \tilde{X} by running off-the-shelf per-frame 3D pose and shape regressors [15, 49, 70], returning per-frame SMPL-X parameters. We directly feed their output to our networks in RGB scenarios. For RGB-D sequences, we roughly align the regressor estimates to depth data via optimization minimizing joint errors. On EgoBody, we follow HuMoR and use VPoser-t for initialization. This allows us to perform

Input	Noise	Method	GMPJPE \downarrow			Accel \downarrow	Cont \uparrow	Skat \downarrow	Dist \downarrow
			-vis	-occ	-all				
Occ-L.	3	VPoser-t	33.0	242.6	109.2	5.1	-	0.219	25.91
		HuMoR [67]	42.4	167.9	88.0	3.3	0.68	0.230	2.59
		MDM++	36.2	71.9	49.2	3.0	0.94	0.102	0.67
		Ours	21.8	57.4	34.8	2.3	0.95	0.078	<u>0.69</u>
	5	VPoser-t	43.0	243.1	115.7	7.2	-	0.179	22.5
		HuMoR	46.1	163.9	88.9	4.3	0.60	0.257	1.81
		MDM++	40.9	75.4	53.4	4.4	0.93	0.126	<u>0.70</u>
		Ours	31.3	66.1	44.0	3.0	0.94	0.105	0.69
	7	VPoser-t	55.1	247.6	125.1	9.4	-	0.116	18.93
		HuMoR	70.7	186.2	112.7	5.9	0.52	0.269	2.56
		MDM++	53.5	100.0	77.5	9.3	0.74	0.287	0.70
		Ours	45.6	88.9	61.3	4.1	0.87	<u>0.150</u>	<u>0.76</u>
Occ-10%	3	VPoser-t	58.9	136.4	66.4	3.4	-	0.379	3.12
		HuMoR [67]	50.0	109.0	55.7	2.6	0.88	0.192	0.66
		Ours	26.3	56.3	29.2	2.3	0.96	0.085	0.62

Table 1. **Evaluation on AMASS.** The best / second best results are in **boldface**, and underlined, respectively. ‘Cont’ denotes Contact, and ‘Skat’ denotes Skating.

Method	RGB-D			RGB		
	Skating \downarrow	$\ \text{Accel}\ \downarrow$	Dist \downarrow	Skating \downarrow	$\ \text{Accel}\ \downarrow$	Dist \downarrow
CLIFF [49]	-	-	-	0.707	49.6	61.80
VPoser-t	0.286	3.4	48.75	0.219	3.2	50.14
LEMO [100]	0.176	1.8	<u>34.22</u>	-	-	-
HuMoR [67]	<u>0.117</u>	<u>1.9</u>	54.76	<u>0.139</u>	2.3	<u>35.41</u>
PhaseMP [72]	-	-	-	0.180	1.8	46.96
Ours-init	0.565	24.4	28.70	0.758	43.7	73.83
Ours	0.038	1.8	3.36	0.116	<u>2.2</u>	9.73

Table 2. **Evaluation on PROX.** The best / second best results are in **boldface**, and underlined, respectively.

a fair quantitative comparison against baselines – factoring out the impact of the initialization strategy. Please refer to the Supp. Mat. for more implementation details.

Baselines. We compare our method against (1) a per-frame human mesh regressor, CLIFF [49] (RGB only) and (2) four optimization-based methods leveraging motion priors: VPoser-t [62, 67], HuMoR [67], LEMO [100] (RGB-D only) and PhaseMP [72] (RGB only) *. Note that methods of type (2) are currently the ones reporting the best results for robust monocular motion reconstruction. For reference, we also include as a baseline our initialization stage (‘Ours-init’).

Results. Tab. 2 reports physical plausibility results obtained on PROX. CLIFF and Ours-init (RGB/RGB-D) are per-frame methods, producing noticeable motion jitter and foot skating. VPoser-t simply enforces 3D joint smoothness and struggles to recover realistic motion dynamics. LEMO tackles noise and occlusions separately, generalizing less well to such complex scenarios. HuMoR and PhaseMP model motion transitions between frames but are less effective in capturing longer-range temporal correlations – pro-

*We compare with PhaseMP in the PROX-RGB setup using the results kindly provided by PhaseMP authors.




Figure 5. **Qualitative results on PROX** (RGB-D input, left) and **EgoBody** (RGB input, right).

Method	GMPJPE↓	MPJPE↓		Accel↓	Skating↓	Dist↓
		-vis	-occ			
VPoser-t	344.8	<u>63.8</u>	<u>126.2</u>	3.8	<u>0.143</u>	<u>13.34</u>
HuMoR [67]	340.3	74.5	164.6	<u>3.5</u>	0.147	17.44
Ours	314.7	60.0	122.9	1.6	0.010	0.96

Table 3. **Evaluation on EgoBody (RGB)**. The best / second best results are in **boldface**, and underlined, respectively.

ducing implausible results under heavy occlusions. In contrast, RoHM reconstructs smooth motions with improved foot-ground interactions (Skating and Dist), and realistic motions for occluded body parts, as shown in Fig. 5. Notably, our method starts from a much more challenging initialization (Ours-init) compared to HuMoR and PhaseMP (VPoser-t), with more severe jitter and foot skating, as can be observed by comparing rows 2 and 6 in Tab. 2.

Factoring out the impact of the initialization stage, Tab. 3 presents quantitative results on EgoBody. We start from the same initialization as HuMoR (VPoser-t) and consistently outperform the baselines across all metrics. Qualitative results are shown in Fig. 5 (right). This indicates that RoHM can generate more plausible motions, even in the highly occluded, challenging scenarios of this dataset.

Runtime. Our approach exhibits significantly reduced runtime compared to HuMoR, being **30 times faster** factoring out the initialization stage (see details in Supp. Mat.).

5.5. Ablation Study

We perform ablation studies on AMASS (Fig. 3 right, with respect to different noise levels) and PROX (Tab. 4). Our iterative inference scheme leveraging TrajControl effectively alleviates foot skating by closing the gap between PoseNet and TrajNet, particularly in the presence of large noise (see Fig. 3 right, and ‘w/o TC’ versus ‘Ours’ in Tab. 4). Iterating between TrajNet and PoseNet twice as described in

Method	RGB-D			RGB		
	Skating↓	Accel ↓	Dist↓	Skating↓	Accel ↓	Dist↓
Ours	0.038	<u>1.8</u>	3.36	0.116	2.2	9.73
w/o TC itr=2	0.046	<u>1.8</u>	4.22	<u>0.146</u>	<u>2.3</u>	10.99
w/o TC	0.056	2.1	4.62	0.165	2.7	11.51
w/o \mathcal{J} w/o TC	0.072	1.7	<u>3.42</u>	0.213	2.2	10.20

Table 4. **Ablation study on PROX**. The best / second best results are in **boldface**, and underlined, respectively. \mathcal{J} denotes the test-time guidance in Eq. (12)(13), and ‘TC’ denotes TrajControl. ‘itr=2’ denotes two iterative iterations as in Sec. 4.3.




Figure 6. **Ablation for test-time guidance \mathcal{J}_{2D}** . For each example, left/right denote without and with \mathcal{J}_{2D} (Eq. (13)), respectively.

Sec. 4.3 without TrajControl (‘w/o TC itr=2’ in Tab. 4) improves motion plausibility to some extent but is still sub-optimal. Test-time score guidance further improves result-observation alignment (see Fig. 6) and alleviates foot skating in all setups. As expected, test-time guidance slightly impacts motion smoothness – an aspect compensated by the iterative inference scheme.

6. Conclusion

We proposed RoHM, an approach for robust human motion reconstruction. Differently from previous work relying on test-time optimization [67, 72], the approach learns how to reconstruct motion from data using diffusion models. The approach decouples the problem of recovering global and local motion by learning two models and conditioning them on available image evidence; a flexible control module captures correlations between global and local dynamics, leveraged by an iterative inference scheme to refine motion plausibility. Experiments on three publicly available datasets show that the approach can reconstruct more realistic and accurate motions than state-of-the-art baselines, especially in challenging scenarios exhibiting noise and occlusions.

Limitations and Future Work. RoHM currently does not work online at real-time framerates. In the future, we plan to evaluate accuracy-efficiency tradeoffs using different architectures (e.g., [13]). Moreover, adapting RoHM to further incorporate 3D environment conditions to model human-scene interactions is an exciting avenue for future research. Finally, while here we focused on full-body reconstruction, future work should extend RoHM to model also facial expressions and articulated hand poses over time. **Acknowledgements.** Siyu Tang is partially funded by the SNSF project grant 200021 204840. We sincerely thank Korrawe Karunratanakul, Sebastian Starke, Marko Mihajlovic, Tony Tung, Yuting Ye, Artsiom Sanakoyeu, and Yuhua Chen for insightful discussions.

References

- [1] Anurag Arnab, Carl Doersch, and Andrew Zisserman. Exploiting temporal context for 3d human pose estimation in the wild. In *CVPR*, 2019. 2
- [2] German Barquero, Sergio Escalera, and Cristina Palmero. Belfusion: Latent diffusion for behavior-driven human motion prediction. In *ICCV*, 2023. 3
- [3] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J. Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *ECCV*, 2016. 2
- [4] Yujun Cai, Yiwei Wang, Yiheng Zhu, Tat-Jen Cham, Jianfei Cai, Junsong Yuan, Jun Liu, Chuanxia Zheng, Sijie Yan, Henghui Ding, et al. A unified 3d human motion synthesis model via conditional variational auto-encoder. In *ICCV*, 2021. 3
- [5] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017. 5
- [6] 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 *ICCV*, 2023. 2, 3
- [7] 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 *CVPR*, 2023. 3
- [8] Yu Cheng, Bo Yang, Bo Wang, Yan Wending, and Robby Tan. Occlusion-Aware Networks for 3D Human Pose Estimation in Video. In *ICCV*, 2019. 2
- [9] Junhyeong Cho, Kim Youwang, and Tae-Hyun Oh. Cross-attention of disentangled modalities for 3D human mesh recovery with transformers. In *ECCV*, 2022. 2
- [10] Hongsuk Choi, Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Beyond static features for temporally consistent 3d human pose and shape from a video. In *CVPR*, 2021. 1, 2
- [11] Vasileios Choutas, Federica Bogo, Jingjing Shen, and Julien Valentin. Learning to fit morphable models. In *ECCV*, 2022. 2, 3
- [12] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, 2021. 2, 5
- [13] Yuming Du, Robin Kips, Albert Pumarola, Sebastian Starke, Ali Thabet, and Arsiom Sanakoyeu. Avatars grow legs: Generating smooth human motion from sparse tracking inputs with diffusion model. In *CVPR*, 2023. 2, 3, 8
- [14] Qi Fang, Qing Shuai, Junting Dong, Hujun Bao, and Xiaowei Zhou. Reconstructing 3D human pose by watching humans in the mirror. In *CVPR*, 2021. 2
- [15] Yao Feng, Vasileios Choutas, Timo Bolkart, Dimitrios Tzionas, and Michael J Black. Collaborative regression of expressive bodies using moderation. In *2021 International Conference on 3D Vision (3DV)*, 2021. 2, 3, 7
- [16] Lin Geng Foo, Jia Gong, Hossein Rahmani, and Jun Liu. Distribution-aligned diffusion for human mesh recovery. In *ICCV*, 2023. 2, 3
- [17] Erik Gärtner, Mykhaylo Andriluka, Erwin Coumans, and Cristian Sminchisescu. Differentiable dynamics for articulated 3d human motion reconstruction. In *CVPR*, 2022. 3
- [18] Erik Gärtner, Mykhaylo Andriluka, Hongyi Xu, and Cristian Sminchisescu. Trajectory optimization for physics-based reconstruction of 3d human pose from monocular video. In *CVPR*, 2022. 3
- [19] Shubham Goel, Georgios Pavlakos, Jathushan Rajasegaran, Angjoo Kanazawa, and Jitendra Malik. Humans in 4D: Reconstructing and tracking humans with transformers. In *ICCV*, 2023. 1, 2
- [20] Jia Gong, Lin Geng Foo, Zhipeng Fan, Qihong Ke, Hossein Rahmani, and Jun Liu. Diffpose: Toward more reliable 3d pose estimation. In *CVPR*, 2023. 3
- [21] 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 *CVPR*, 2022. 4
- [22] Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J. Black. Resolving 3D human pose ambiguities with 3D scene constraints. In *ICCV*, 2019. 6
- [23] Gustav Henter, Simon Alexanderson, and Jonas Beskow. MoGlow: Probabilistic and controllable motion synthesis using normalising flows. *ACM TOG*, 39(4):1–14, 2020. 3
- [24] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. 2
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2, 3, 5
- [26] Daniel Holden, Jun Saito, and Taku Komura. A deep learning framework for character motion synthesis and editing. *ACM TOG*, 35(4):1–11, 2016. 4
- [27] Nicholas Howe, Michael Leventon, and William Freeman. Bayesian reconstruction of 3D human motion from single-camera video. In *NeurIPS*, 2000. 3
- [28] Buzhen Huang, Liang Pan, Yuan Yang, Jingyi Ju, and Yang-gang Wang. Neural mocon: Neural motion control for physically plausible human motion capture. In *CVPR*, 2022. 3
- [29] 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 *CVPR*, 2023. 3
- [30] Yinghao Huang, Federica Bogo, Christoph Lassner, Angjoo Kanazawa, Peter V Gehler, Javier Romero, Ijaz Akhter, and Michael J Black. Towards accurate marker-less human shape and pose estimation over time. In *2017 International Conference on 3D Vision (3DV)*, 2017. 2, 6
- [31] Michael Janner, Yilun Du, Joshua B Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022. 4
- [32] 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*, 36, 2024. 3
- [33] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 2

- [34] Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. In *CVPR*, 2019. 1, 2
- [35] Korrawe Karunratanakul, Konpat Preechakul, Supasorn Suwajanakorn, and Siyu Tang. Guided motion diffusion for controllable human motion synthesis. In *ICCV*, 2023. 2, 3
- [36] Korrawe Karunratanakul, Konpat Preechakul, Emre Aksan, Thabo Beeler, Supasorn Suwajanakorn, and Siyu Tang. Optimizing diffusion noise can serve as universal motion priors. In *CVPR*, 2024. 3
- [37] Manuel Kaufmann, Emre Aksan, Jie Song, Fabrizio Pece, Remo Ziegler, and Otmar Hilliges. Convolutional autoencoders for human motion infilling. In *2020 International Conference on 3D Vision (3DV)*, pages 918–927. IEEE, 2020. 3
- [38] Rawal Khirodkar, Shashank Tripathi, and Kris Kitani. Occluded human mesh recovery. In *CVPR*, 2022. 2
- [39] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *CVPR*, 2020. 1, 2, 6
- [40] Muhammed Kocabas, Chun-Hao P. Huang, Otmar Hilliges, and Michael J. Black. PARE: Part attention regressor for 3D human body estimation. In *ICCV*, 2021. 2
- [41] Muhammed Kocabas, Chun-Hao P. Huang, Joachim Tesch, Lea Müller, Otmar Hilliges, and Michael J. Black. SPEC: Seeing people in the wild with an estimated camera. In *ICCV*, 2021. 2
- [42] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *ICCV*, 2019.
- [43] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *CVPR*, 2019.
- [44] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *CVPR*, 2019.
- [45] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J Black, and Peter V Gehler. Unite the people: Closing the loop between 3D and 2D human representations. In *CVPR*, 2017.
- [46] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. HybriK: A hybrid analytical-neural inverse kinematics solution for 3D human pose and shape estimation. In *CVPR*, 2021. 2
- [47] Jiefeng Li, Siyuan Bian, Chao Xu, Gang Liu, Gang Yu, and Cewu Lu. D&d: Learning human dynamics from dynamic camera. In *ECCV*, 2022. 2
- [48] Jiahao Li, Zongxin Yang, Xiaohan Wang, Jianxin Ma, Chang Zhou, and Yi Yang. Jotr: 3d joint contrastive learning with transformers for occluded human mesh recovery. In *ICCV*, 2023. 2
- [49] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. Cliff: Carrying location information in full frames into human pose and shape estimation. In *ECCV*, 2022. 2, 3, 7
- [50] Jing Lin, Ailing Zeng, Haoqian Wang, Lei Zhang, and Yu Li. One-stage 3d whole-body mesh recovery with component aware transformer. In *CVPR*, 2023.
- [51] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *CVPR*, 2021. 2
- [52] Hung Yu Ling, Fabio Zinno, George Cheng, and Michiel van de Panne. Character controllers using motion vaes. *ACM TOG*, 39(4):1–12, 2020. 3
- [53] Qihao Liu, Yi Zhang, Song Bai, and Alan Yuille. Explicit occlusion reasoning for multi-person 3d human pose estimation. In *ECCV*, 2022. 2
- [54] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM TOG*, 34(6):1–16, 2015. 2, 3
- [55] Zhengyi Luo, S Alireza Golestaneh, and Kris M Kitani. 3d human motion estimation via motion compression and refinement. In *ACCV*, 2020. 2, 3
- [56] Zhengyi Luo, Shun Iwase, Ye Yuan, and Kris Kitani. Embodied scene-aware human pose estimation. In *NeurIPS*, 2022. 3
- [57] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *ICCV*, 2019. 2, 3, 6
- [58] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiee, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM TOG*, 36(4):1–14, 2017. 2
- [59] Hyeongjin Nam, Daniel Sungho Jung, Yeonguk Oh, and Kyoung Mu Lee. Cyclic test-time adaptation on monocular video for 3d human mesh reconstruction. In *ICCV*, 2023. 1, 2
- [60] Jorge Nocedal and Stephen Wright. *Numerical Optimization*. Springer, 2006. 2
- [61] Dirk Ormoneit, Hedvig Sidenbladh, Michael Black, and Trevor Hastie. Learning and tracking cyclic human motion. In *NeurIPS*, 2000. 3
- [62] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3D hands, face, and body from a single image. In *CVPR*, 2019. 2, 3, 6, 7
- [63] Georgios Pavlakos, Jitendra Malik, and Angjoo Kanazawa. Human mesh recovery from multiple shots. In *CVPR*, 2022. 2
- [64] Vladimir Pavlovic, James Rehg, and John MacCormick. Learning switching linear models of human motion. In *NeurIPS*, 2001. 3
- [65] Jathushan Rajasegaran, Georgios Pavlakos, Angjoo Kanazawa, and Jitendra Malik. Tracking people by predicting 3d appearance, location and pose. In *CVPR*, 2022. 2
- [66] Davis Rempe, Leonidas J Guibas, Aaron Hertzmann, Bryan Russell, Ruben Villegas, and Jimei Yang. Contact and human dynamics from monocular video. In *ECCV*, 2020. 3, 4

- [67] Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J. Guibas. HuMoR: 3D human motion model for robust pose estimation. In *ICCV*, 2021. 1, 2, 3, 5, 6, 7, 8
- [68] Davis Rempe, Zhengyi Luo, Xue Bin Peng, Ye Yuan, Kris Kitani, Karsten Kreis, Sanja Fidler, and Or Litany. Trace and pace: Controllable pedestrian animation via guided trajectory diffusion. In *CVPR*, 2023. 3, 4, 5
- [69] Chris Rockwell and David Fouhey. Full-body awareness from partial observations. In *ECCV*, 2020. 2
- [70] István Sárádi, Timm Linder, Kai O. Arras, and Bastian Leibe. MeTRAbs: metric-scale truncation-robust heatmaps for absolute 3D human pose estimation. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(1):16–30, 2021. 2, 3, 7
- [71] Yonatan Shafir, Guy Tevet, Roy Kapon, and Amit H Bermano. Human motion diffusion as a generative prior. *arXiv preprint arXiv:2303.01418*, 2023. 2, 3, 6
- [72] Mingyi Shi, Sebastian Starke, Yuting Ye, Taku Komura, and Jungdam Won. PhaseMP: Robust 3D pose estimation via phase-conditioned human motion prior. In *ICCV*, 2023. 1, 2, 6, 7, 8
- [73] Soshi Shimada, Vladislav Golyanik, Weipeng Xu, and Christian Theobalt. Physcap: Physically plausible monocular 3d motion capture in real time. *ACM TOG*, 39(6):1–16, 2020. 3
- [74] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. 5
- [75] Jie Song, Xu Chen, and Otmar Hilliges. Human body model fitting by learned gradient descent. In *ECCV*, 2020. 2, 3
- [76] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 2
- [77] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021. 5
- [78] Yu Sun, Yun Ye, Wu Liu, Wenpeng Gao, YiLi Fu, and Tao Mei. Human mesh recovery from monocular images via a skeleton-disentangled representation. In *ICCV*, 2019. 2
- [79] Yu Sun, Qian Bao, Wu Liu, Tao Mei, and Michael J. Black. TRACE: 5D Temporal Regression of Avatars with Dynamic Cameras in 3D Environments. In *CVPR*, 2023. 2
- [80] Julian Tanke, Linguang Zhang, Amy Zhao, Chengcheng Tang, Yujun Cai, Lezi Wang, Po-Chen Wu, Juergen Gall, and Cem Keskin. Social diffusion: Long-term multiple human motion anticipation. In *ICCV*, 2023. 3
- [81] Graham Taylor, Geoffrey Hinton, and Sam Roweis. Modeling human motion using binary latent variables. In *NeurIPS*, 2007. 3
- [82] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *ICLR*, 2023. 2, 3, 4, 6
- [83] Jonathan Tseng, Rodrigo Castellon, and Karen Liu. Human motion diffusion as a generative prior. In *CVPR*, 2023. 2, 3
- [84] Raquel Urtasun, David Fleet, and Pascal Fua. 3D people tracking with gaussian process dynamical models. In *CVPR*, 2006. 3
- [85] Raquel Urtasun, David Fleet, and Pascal Fua. Temporal motion models for monocular and multiview 3d human body tracking. *CVIU*, 104(2), 2006. 3
- [86] Dongkai Wang and Shiliang Zhang. 3d human mesh recovery with sequentially global rotation estimation. In *ICCV*, 2023. 2
- [87] Yin Wang, Zhiying Leng, Frederick WB Li, Shun-Cheng Wu, and Xiaohui Liang. Fg-t2m: Fine-grained text-driven human motion generation via diffusion model. In *ICCV*, 2023. 3
- [88] Wen-Li Wei, Jen-Chun Lin, Tyng-Luh Liu, and Hong-Yuan Mark Liao. Capturing humans in motion: Temporal-attentive 3d human pose and shape estimation from monocular video. In *CVPR*, 2022. 2
- [89] Kevin Xie, Tingwu Wang, Umar Iqbal, Yunrong Guo, Sanja Fidler, and Florian Shkurti. Physics-based human motion estimation and synthesis from videos. In *ICCV*, 2021. 3
- [90] Yuanlu Xu, Song-Chun Zhu, and Tony Tung. Denserac: Joint 3D pose and shape estimation by dense render-and-compare. In *ICCV*, 2019. 2
- [91] Vickie Ye, Georgios Pavlakos, Jitendra Malik, and Angjoo Kanazawa. Decoupling human and camera motion from videos in the wild. In *CVPR*, 2023. 2
- [92] Yingxuan You, Hong Liu, Ti Wang, Wenhao Li, Runwei Ding, and Xia Li. Co-evolution of pose and mesh for 3d human body estimation from video. In *ICCV*, 2023. 2
- [93] Ye Yuan, Shih-En Wei, Tomas Simon, Kris Kitani, and Jason Saragih. Simpoe: Simulated character control for 3d human pose estimation. In *CVPR*, 2021. 3
- [94] Ye Yuan, Umar Iqbal, Pavlo Molchanov, Kris Kitani, and Jan Kautz. GLAMR: Global occlusion-aware human mesh recovery with dynamic cameras. In *CVPR*, 2022. 2, 4
- [95] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. In *ICCV*, 2023. 2, 3
- [96] Andrei Zanfir, Eduard Gabriel Bazavan, Hongyi Xu, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Weakly supervised 3d human pose and shape reconstruction with normalizing flows. In *ECCV*, 2020. 2
- [97] Jianfeng Zhang, Dongdong Yu, Jun Hao Liew, Xuecheng Nie, and Jiashi Feng. Body meshes as points. In *CVPR*, 2021. 2, 4
- [98] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. 2, 5
- [99] Mingyuan Zhang, Xinying Guo, Liang Pan, Zhongang Cai, Fangzhou Hong, Huirong Li, Lei Yang, and Ziwei Liu. Remodiffuse: Retrieval-augmented motion diffusion model. *arXiv preprint arXiv:2304.01116*, 2023. 3
- [100] Siwei Zhang, Yan Zhang, Federica Bogo, Marc Pollefeys, and Siyu Tang. Learning motion priors for 4d human body capture in 3d scenes. In *ICCV*, 2021. 1, 2, 4, 5, 6, 7
- [101] Siwei Zhang, Qianli Ma, Yan Zhang, Zhiyin Qian, Taein Kwon, Marc Pollefeys, Federica Bogo, and Siyu Tang.

- Egobody: Human body shape and motion of interacting people from head-mounted devices. In *ECCV*, 2022. 6
- [102] Siwei Zhang, Qianli Ma, Yan Zhang, Sadegh Aliakbarian, Darren Cosker, and Siyu Tang. Probabilistic human mesh recovery in 3d scenes from egocentric views. In *ICCV*, 2023. 2, 3, 5
- [103] Yan Zhang and Siyu Tang. The wanderings of odysseus in 3d scenes. In *CVPR*, 2022. 3
- [104] Yi Zhang, Pengliang Ji, Angtian Wang, Jieru Mei, Adam Kortylewski, and Alan Yuille. 3d-aware neural body fitting for occlusion robust 3d human pose estimation. In *ICCV*, 2023. 2
- [105] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *CVPR*, 2019. 4