

DiffusionPoser: Real-time Human Motion Reconstruction From Arbitrary Sparse Sensors Using Autoregressive Diffusion

Tom Van Wouwe Seunghwan Lee Antoine Falisse Scott Delp C. Karen Liu
 Stanford University

[tvwouwe, lsw9021, afalisse, delp, ckliu38]@stanford.edu

Abstract

Motion capture from a limited number of body-worn sensors, such as inertial measurement units (IMUs) and pressure insoles, has important applications in health, human performance, and entertainment. Recent work has focused on accurately reconstructing whole-body motion from a specific sensor configuration using six IMUs. While a common goal across applications is to use the minimal number of sensors to achieve required accuracy, the optimal arrangement of the sensors might differ from application to application. We propose a single diffusion model, DiffusionPoser, which reconstructs human motion in real-time from an arbitrary combination of sensors, including IMUs placed at specified locations, and, pressure insoles. Unlike existing methods, our model grants users the flexibility to determine the number and arrangement of sensors tailored to the specific activity of interest, without the need for retraining. A novel autoregressive inferencing scheme ensures real-time motion reconstruction that closely aligns with measured sensor signals. The generative nature of DiffusionPoser ensures realistic behavior, even for degrees-of-freedom not directly measured. Qualitative results can be found on our website: <https://diffusionposer.github.io/>.

1. Introduction

Portable and minimally intrusive tools for estimating human motion have important applications in health and human performance as they allow continuous monitoring of the motions of the body and the forces in the musculoskeletal system [19, 45]. Mobile sensing technologies also have applications in virtual and augmented reality, and video games [32, 33]. Truly wearable sensors such as inertial measurement units (IMUs) (e.g. [2]), pressure insoles (e.g. [6]) or electromagnetic sensors (e.g. [15]), have an advantage compared to video-based motion capture, as they are egocentric, do not suffer from occlusion or poor lighting, and offer an infinite measurement volume. While egocentric




Figure 1. (Left) Examples of live reconstruction using DiffusionPoser. (Right) Subject instrumented with IMUs. We assume IMUs may be attached at 13 specific locations: pelvis, thighs, shanks, feet, arms, wrists, torso, head.

video represents an alternative or supplementary modality, it introduces potential privacy concerns. With IMU sensors being the cheapest alternative of these mobile sensing technologies, much research has focused on motion reconstruction from these. Configurations with few sensors, i.e. sparse configurations, are desirable for user convenience and adherence. However, a sparse configuration as opposed to a dense configuration, does not allow direct measurement of the full body motion. Besides sparsity, sensor noise is an important challenge. To tackle the problems of sparsity and noise in the case of IMUs, both optimization [48] and data-driven [11, 36] methods have been used.

Prior data-driven methods that reconstruct motion from sparse IMU sensor configurations focused on a specific number of sensors in a specific configuration [11, 13, 36, 48, 52, 53]. In practice, applications would benefit from a

real-time system that is more flexible and where the user can optimize the number of sensors and configuration to the activity of interest.

Here we present a single diffusion generative model, DiffusionPoser, that reconstructs human motion in real-time from ad-hoc sensor configurations, i.e. any number (between one and thirteen) and combination of IMUs and/or pressure insoles can be used.

The use of a diffusion generative model is key to allow any combination of sensors. It results in realistic motion reconstruction in correspondence with measured signals, even for very sparse configurations. Our generative model is also robust to cases where sensor signals are corrupted or lost.

Two developments, that both utilize the inpainting denoising technique [22], differentiate our model from previous motion diffusion models such as [7, 20, 42, 44] and enable our real-time sensor-based motion reconstruction application:

1. We use an **autoregressive inference scheme** (Section 3.4.1) where at every time step the motion history is used to generate a full motion sequence, including the current frame of interest.
2. We use a **tailored motion representation and inpainting denoising**, rather than a conditional model, to control motion generation. The motion representation consists of contact labels, root translation, global segment orientations and global linear accelerations of segment-specific sites.

We demonstrate the capability of DiffusionPoser to optimize sensors configurations for applications that might require different levels of accuracy for specific body regions. We show that, in the case of using three sensors, the best configuration instruments, the pelvis and wrists when optimizing whole-body reconstruction, the pelvis and thighs when optimizing reconstruction of leg motion, and the upper arms and pelvis when optimizing reconstruction of the back, neck and shoulder joints. All these configurations can be used without retraining. We show that such flexibility does not come at the cost of accuracy as DiffusionPoser achieves accuracy that is on par with state-of-the-art regressive models that focus on a single specific six sensor IMU setup [11, 13, 52, 53].

Our method is agnostic to the underlying skeleton models and we present two versions of DiffusionPoser for different skeletons. First, we use the SMPL body model [21]. Second, because we aim to enable biomedical research applications, we implemented DiffusionPoser for a different and more physiologically realistic skeleton: the OpenSim musculoskeletal model [34]. While DiffIP for the SMPL model was trained on the available AMASS dataset we created a new dataset specific to OpenSim models by combining three existing motion capture datasets [14, 40, 43].

DiffusionPoser has interesting features for health and biomedical research applications. It reconstructs motion from wearable IMU sensors that have extended battery lives and are portable. It allows users to select a sensor configuration while trading off comfort and specific accuracy that suits their goal best, on-the-fly, without retraining. Finally, DiffusionPoser runs at 20Hz making it useful for health and performance interventions such as guidance of rehabilitation exercises or training using biofeedback [46].

2. Related Work

Motion Capture using sparse IMUs. An IMU measures 3D linear acceleration, 3D angular velocity, and the direction of the magnetic North. These measurements are noisy and expressed in the local IMU frame. Commercially available IMU systems come with proprietary sensor fusion algorithms providing orientation and acceleration for use in downstream applications (e.g., [31]). Linear accelerations remain noisy and are therefore sometimes smoothed before being used as reconstruction input (e.g. [13]).

Early work performing human motion reconstruction from sparse IMUs performed database search based on four [41] or five accelerometer signals [36]. In recent years, several works focused on reconstructing whole-body motion from six IMU sensors placed on the wrists (2), on the shanks (2), on the pelvis, and on the head. Sparse Inertial Poser [48] relies on this configuration and does offline motion reconstruction by optimizing a sequence of poses to match measured signals.

Deep Inertial Poser [11] introduced the use of a human motion prior: a bi-directional RNN architecture is trained on a motion dataset. IMU signals for training were synthesized from the AMASS motion dataset [23]. Other data-driven algorithms have built on this work to improve accuracy and address root translation estimation. Transpose [53] combines several RNNs to predict the full pose, including the root motion. An extension of this work, Physics Inertial Poser (PIP) [52], adds a physics layer to refine the joint orientations and root motion predictions. Transformer Inertial Poser (TIP) [13], addresses root motion estimation by predicting stationary points on feet, hands and pelvis and applying a correction accordingly.

A downside of the prior work, that we address with DiffusionPoser, is that only one specific IMU configuration is allowed. One other prior work performs online reconstruction from a flexible and sparse IMU configuration: IMUPoser [25] uses an LSTM to predict motion from different IMU configurations including up to three sensors embedded in a phone, watch, and earbuds. Compared to our work, IMUPoser has a limited number of possible locations for the sensors and does not reconstruct root motion.

Finally, several systems complement inertial measurements with another sensory modality to improve reconstruction

tion (e.g.: optical markers [3], lidar [30], third-person video [24], first-person video [51], and depth camera [55]).

Diffusion for human motion reconstruction. Diffusion models [37] are a class of deep generative models. During training, a neural net learns the inverse mapping of samples from the target distribution that are gradually noised (i.e. denoising). At inference, the neural net is used to perform a denoising process starting from pure noise, thus generating a clean sample. A recent overview of diffusion models for human motion generation can be found in Section 7 of [28].

A diffusion generative model can be controlled in different manners. One approach that has been employed in motion generative systems is to utilize a conditional generative model. In this model, a specific condition, such as text (e.g., [42]) or music (e.g., [44]), is provided during both the training and inference phases. Text-to-motion generative systems that rely on diffusion [16, 42, 54] have been shown to improve expressiveness and robustness compared to prior generative models [1, 27] based on variational autoencoders [17].

Similar to our work, human motion diffusion models have been used to reconstruct motion from different sensory modalities such as monocular camera [8, 10], egocentric camera [20], or position and orientation sensors [7, 12, 50]. There are two important distinctions between our work and the systems mentioned here. First, the mentioned systems have access to a good estimate of position as input, whereas we start from noisy linear acceleration signals. Second, the prior works are conditional models for which the sensor measurement is the condition of the generative process. We use the sensor measurement for guidance instead of using it as a condition. This allows to maintain flexibility with respect to what sensor signals we have access to.

Indeed, when using conditional models, we cannot use a different condition during inference than during training. Guidance on the other hand typically only interferes during inference. There are several ways to do guidance [28] and we take inspiration from [42, 44] that use inpainting denoising as a flexible approach to control (part of) the features of a generated motion. Different from how inpainting denoising is used in [44] to generate long sequence motions by having small overlapping temporal patches, we generate a continuous motion sequence in a fully autoregressive manner. Each newly generate window of motion has an overlapping temporal patch with the previous window that covers all but the last frame.

3. DiffusionPoser

DiffusionPoser reconstructs whole-body human motion in real-time based on measurements from sparse IMUs in arbitrary configurations and/or pressure insoles.

3.1. Skeleton model and IMU instrumentation

We implemented DiffusionPoser for two body models: SMPL [21] and OpenSim [18]. The system works identical for both models. Training is slightly different because of different definitions of the underlying kinematic tree. We provide further detail on the OpenSim model in Appendix A and focus on SMPL here.

The SMPL skeleton model is a 75 dof kinematic tree consisting of a root segment (6 dof) and 23 additional segments connected by ball-and-socket joints (3 dof). We assume that IMUs can only be attached at specific locations to 13 of the 23 body segment (Figure 1). These specific attachment sites should be respected at inference as they are used to synthesize the training data (Section 3.3).

For the upper and lower limbs we opted to attach the sensors in positions as distal as possible without hampering joint motion. As location within a body segment does not influence the orientation estimate of a body, we choose this placement to maximize the signal-to-noise ratio for the measured linear acceleration. The root and torso IMU were attached at locations that provide a good articulation between sensor and body segment with few soft tissue artifacts. Similar to prior work [13, 52, 53] we will ignore the toes, wrist and finger joints in our evaluations as in several motion training datasets these dofs are not articulated.

3.2. Diffusion Model

3.2.1 Features

Our diffusion model generates sequences of $N=61$ feature vectors. The feature vector in each frame represents the combination of whole-body pose and IMU information in a compact manner:

$$\mathbf{x}_{frame} = (\mathbf{R}, \mathbf{a}, \Delta\mathbf{p}, \mathbf{p}_y, \mathbf{b}). \quad (1)$$

$\mathbf{R} \in \mathbf{R}^{24 \times 6}$ are the global orientations of the body segments, parameterized by 6-DOF representation [56]. The orientations of the 13 body segments that are potentially instrumented also represent IMU orientation estimates as there is no relative orientation between IMU and body segment after calibration. $\mathbf{a} \in \mathbf{R}^{13 \times 3}$ are the linear accelerations, expressed in the world frame, of the IMU locations on the potentially instrumented body segments. To complete our whole-body motion representation we have $\Delta\mathbf{p}$, the 2D change in root position from one frame to the next, and \mathbf{p}_y the root vertical position. For the heel and toe of each foot, we add a binary contact feature $\mathbf{b} \in \mathbf{R}^{2 \times 2}$ that will be used to regularize generated motion (Section 3.2.3). The full feature vector is $\mathbf{x} \in \mathbf{R}^{61 \times 190}$.

3.2.2 Diffusion Framework

Training and architecture of our diffusion framework are similar to [42, 44]. Diffusion is modeled as a Markov noise-




Figure 2. **DiffIP transformer decoder network.** Architecture of the denoiser $f_\theta(\hat{z}_t, t, h)$ that predicts the sample (\hat{x}_0) given the noised sample z_t , denoising step t and the body height h . We use the transformer decoder architecture from [47] and use the step embedding and height embedding for cross attention as well as self attention by concatenating them to the input embedding.

ing process with latents $\{x_t\}_{t=0:T}$, with $T = 1000$. The forward noising process is defined as:

$$q(z_t|x) \sim \mathcal{N}(\sqrt{\bar{\alpha}_t}x, (1 - \bar{\alpha}_t)\mathbf{I}) \quad (2)$$

where $\bar{\alpha}_t \in [0, 1]$ are constants which follow a monotonically decreasing (cosine) schedule for increasing t , such that $x_T \sim \mathcal{N}(0, \mathbf{I})$. We learn an approximation of the reverse diffusion process, i.e., the denoising process, by training a transformer neural network f_θ (Figure 2), with parameters θ , which takes a noised version \hat{z}_t of the ground truth motion x , the noise step t and condition h and generates a denoised version \hat{x}_0 that aims to match x :

$$f_\theta(\hat{z}_t, t, h) = \hat{x}_0. \quad (3)$$

The condition h is subject height. Height is associated with body segment lengths, which mathematically underlie the relation between body orientations and body accelerations and, body orientations and root motion.

3.2.3 Training losses

For training f_θ we follow DDPM [9] by sampling a diffusion step t from a uniform distribution $U \in [0, T]$, randomly sampling a motion sequence from our training dataset, noising this motion x to \hat{z}_t , predicting \hat{x}_0 and performing gradient descent on our loss \mathcal{L}

$$\mathbb{E}_{x,t} \mathcal{L}(\hat{x}_0, x). \quad (4)$$

Similar to [44], our loss is composed of a simple loss [9, 42] and several auxiliary losses:

$$\mathcal{L} = \mathcal{L}_{\text{simple}} + \mathcal{L}_{\text{vel}} + \mathcal{L}_{\text{FK}} + \mathcal{L}_{\text{drift}} + \mathcal{L}_{\text{slide}} \quad (5)$$

with

$$\mathcal{L}_{\text{simple}} = \sum_{i=1}^N \|\hat{x}_0^{(i)} - x^{(i)}\|^2, \quad (6)$$

$$\mathcal{L}_{\text{vel}} = \sum_{i=1}^{N-1} \|(\hat{R}_0^{(i+1)} - \hat{R}_0^{(i)}) - (R^{(i+1)} - R^{(i)})\|^2, \quad (7)$$

$$\mathcal{L}_{\text{FK}} = \sum_{i=1}^N \|\text{FK}(\hat{R}_0^{(i)}) - \text{FK}(R^{(i)})\|^2, \quad (8)$$

$$\mathcal{L}_{\text{drift}} = \sum_{i=1}^N \|\hat{p}_0^{(i)} - p^{(i)}\|^2 \quad \text{where } p^{(i)} = \sum_{j=1}^i \Delta p^{(j)}, \quad (9)$$

$$\mathcal{L}_{\text{slide}} = \sum_{i=1}^{N-1} \|\hat{b}_0^{(i)} \cdot [\text{FK}_{\text{ft}}(\hat{R}_0^{(i+1)}) - \text{FK}_{\text{ft}}(\hat{R}_0^{(i)}) + \Delta p_0^{(i)}]\|^2. \quad (10)$$

$\mathcal{L}_{\text{simple}}$ supervises the features, including contact labels $\hat{b}_0^{(i)}$, directly. \mathcal{L}_{vel} encourages smooth motion [26]. \mathcal{L}_{FK} is a kinematic loss to encourage realistic joint positions [35], with FK_{ft} the forward kinematics functions mapping global orientations to foot positions relative to the root. Compared to prior work we explicitly add $\mathcal{L}_{\text{drift}}$, which penalizes drift by accumulating the absolute root translation error across frames. $\mathcal{L}_{\text{drift}}$ is required because we choose to predict Δp , rather than p ; a design choice to improve performance when generative motion models are used autoregressively [29]. $\mathcal{L}_{\text{slide}}$ encourages the model to generate motion where foot motion is consistent with foot contact prediction (minimizing foot velocity when in contact) [44].

3.3. Training data

We used AMASS [23] for training. IMU orientations were synthesized as the global orientations of the body segments. IMU linear accelerations were derived by double differentiation of selected vertex positions that were obtained by a forward kinematic pass. To better match measured and synthesized linear accelerations we averaged these at the original sampling rate using a moving average sliding window of 166ms (e.g., 11 frames at 60Hz) [13]. Finally, we resampled at 20Hz. Contact labels were annotated by applying a velocity threshold (0.3m/s) for each of the four potential contact points (heels and toes) [53].

Sampling from training data was done following a probability that was assigned to each trial based on a simple energy metric. The sampling strategy encouraged the model to predict more diverse motion when deployed purely generative and improved reconstruction accuracy (Section 4.7). For the energy metric, we estimated center-of-mass (COM) position of each segment as the midpoint between joints and

calculated COM velocity. Then we calculated linear kinetic energy for each segment: $0.5 \cdot m_{segment} \cdot \dot{COM}^2$, with $m_{segment}$ in Appendix D.

3.4. Inference

3.4.1 Autoregressive inference

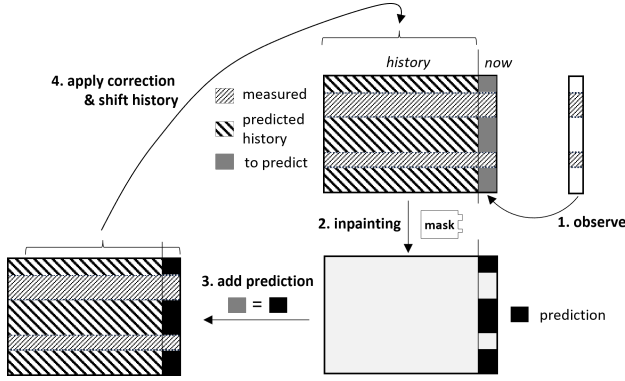


Figure 3. **Four step autoregressive inference including denoising inpainting.** Motion is reconstructed frame-by-frame in real-time following a four step process. New predictions are shifted into history and serve as input for the reconstruction at the consecutive timesteps.

Inference is run at 20Hz and consists of a four step process (Figure 3). At each timestep a new observation is made of the features measured by the sensors. The new observation, together with the history of measured features and the history of reconstructed motion are input to the inpainting denoising process. The output from the inpainting denoising [38] process is a prediction of unobserved features at the current frame. The predicted features are concatenated with the measured features. Finally, we apply a correction of the root motion 3.4.2 and shift the new frame into the history before a new measurement is acquired.

Our inpainting denoising algorithm 1 takes as input $\mathbf{x}^{\text{input}}$ which consists for frames 0 to $N - 1$ of the measured and reconstructed feature history. For the final frame, frame N , the observed features are set to the sensor measurements. The unobserved features of frame N , that we will predict, are initialized from the reconstructions at frame $N - 1$. The inpainting mask (\mathbf{m}) covers all features for frames 0 to $N - 1$ and the observed features for frame N . $\mathbf{x}^{\text{input}}$ are inpainted using our denoiser, which leads to $\mathbf{x}^{\text{output}}$. From $\mathbf{x}^{\text{output}}$ we take the unobserved features at frame N and concatenate these with the measured features.

3.4.2 Root correction

We exploit the predicted contact information to correct foot sliding artifacts and reduce root drift. The motion of points on the feet that are predicted to be in contact during the

ALGORITHM 1: Inpainting denoising with DiffusionPoser

```

Given :  $c, \mathbf{m}, \mathbf{x}^{\text{input}}$ 
 $\mathbf{x}_T = \mathbf{x}^{\text{input}}$ 
for  $t = T-1, T-2, \dots, 1$  do
    noise  $\hat{\mathbf{z}}_t \sim \mathcal{N}(\sqrt{\bar{\alpha}_t} \mathbf{x}_{t+1}, (1 - \bar{\alpha}_t) \mathbf{I})$ 
    predict  $\hat{\mathbf{x}}_0 = p_\theta(\hat{\mathbf{z}}_t, t, c)$ 
    edit  $\mathbf{x}_t = \mathbf{m} \odot \mathbf{x}^{\text{input}} + (1 - \mathbf{m}) \odot \hat{\mathbf{x}}_0$ 
end
 $\mathbf{x}^{\text{output}} = \mathbf{x}_t$ 

```

transition from frame $N - 1$ to N is corrected by changing the root motion. We do this by first calculating the mean horizontal position change between frames $N - 1$ and N across the points that are predicted to be in contact. We then subtract this mean position change from the predicted horizontal root position change between frames $N - 1$ and N . Note that this heuristic does not guarantee predicted contact points to be static, as it only makes the mean of the contact points to be static.

4. Evaluations

We performed several experiments and evaluations to quantify reconstruction performance for different IMU configurations, to compare DiffusionPoser to Transpose [53], PIP [52] and TIP [13] and to demonstrate multimodality. We show that DiffusionPoser has a comparable reconstruction accuracy to other methods, while offering users the flexibility to select sensor setting on-the-fly. Next, we performed experiments to show robustness of DiffusionPoser to sensor signal corruption and to examine the effect of different denoising schemes. Finally we did several ablations.

Following [13, 52, 53] we performed evaluations using the TotalCapture dataset with real IMU data for sensors on the pelvis, head, wrists and shanks (‘TotalCaptureReal’) [11, 43]. Next we also used ‘TotalCaptureSynth’, a version of the TotalCapture dataset where we synthesized all IMU signals. ‘TotalCaptureSynth’ does not allow to evaluate sim-to-real error but is useful to understand the effect of different configurations on reconstruction accuracy. We also used DIP-IMU for additional baseline comparison [11]. Evaluations are done using 30 denoising steps (Section 4.6).

We report the following evaluation metrics:

Local Angular Error, LA [°]: Rotation difference between ground-truth and reconstructed local joint angles. We use specific LA metrics to quantify reconstruction accuracy of the back (neck, shoulder and spine joints) and legs (hip, knee and ankle joints).

Global Angular Error, GA [°]: Rotation difference between ground-truth and reconstructed local global segment orientations.

Joint Position Error, JPE [cm]: Difference between corresponding joint positions, expressed in the root frame, of reconstructed and ground-truth motions.

Jitter [-]: Ratio of the global jitter averaged across joints of the reconstructed over the jitter of the ground truth motion. Global jitter of a joint was calculated as the third derivative of absolute position using finite differences.

Root Translation Error, RE [m]: Distance between root position of ground-truth and reconstructed motion at 2s, 5s and 10s into the motion.

The metrics are based on Transpose, TIP and PIP. Because there is no exact correspondence between metrics across papers, values do not exactly correspond to the original papers. This is further explained in Appendix B.




Figure 4. **Motion reconstructions with PIP and DiffusionPoser (Ours) for different IMU configurations of a TotalCaptureReal sequence.** Yellow: PIP with pelvis, head, wrists and shanks. Grey: ground truth. Purple: Ours with pelvis, head, wrists and shanks. Orange: Ours with wrists, shanks. Blue: Ours with pelvis and wrists. Green: Ours with shanks.

4.1. Optimizing IMU configurations across tasks

We show the capability of DiffusionPoser to optimize sensor configurations for specific activities by evaluating on different metrics to quantify accuracy for: whole-body (GA), leg kinematics (legsLA), back kinematics (backLA) and global translation (RE10). We optimized sensor configurations for two, three and four IMU sensors by evaluating

the performance for a range of selected configurations on TotalCaptureSynth. This was the only experiment where we used TotalCaptureSynth because we did not have real IMU data for all the potentially instrumented segments that we wanted to test. A full table with results can be found in Appendix C, Table 4. Table 1 shows the optimal configuration when using four, three and two IMU sensors for the different evaluation metrics. Note how instrumenting thighs is best to reconstruct leg kinematics, whereas shanks are preferred to improve root translation error.

Table 1. **Optimal IMU configurations** for four, three and two sensors for different tasks.

#	GA [°]	legsLA [°]	backLA [°]	RE10 [m]
4	<i>plvs, hd, wr_r, wr_l</i>	<i>plvs, hd th_r, th_l</i>	<i>plvs, hd arm_r, arm_l</i>	<i>th_r, th_l sh_r, sh_l</i>
3	<i>plvs, wr_r, wr_l</i>	<i>plvs, th_r, th_l</i>	<i>plvs, arm_r, arm_l</i>	<i>plvs, sh_r, sh_l</i>
2	<i>arm_r, arm_l</i>	<i>plvs, sh_r</i>	<i>arm_r, arm_l</i>	<i>sh_r, sh_l</i>

4.2. Sim-to-real error and reconstruction quality

We evaluated real vs synthetic reconstruction quality. We performed exhaustive evaluation for all possible IMU configurations with six or fewer sensors with the potentially instrumented segments limited to wrists, shanks, pelvis and head. Qualitative results of reconstructions from real data and a comparison to PIP [52] are shown in Figure 4. Table 2 reports sim-to-real comparisons for selected configurations. Quantitative results for all configurations we tested with real IMU data are in Appendix C, Table 5.

Table 2. **Sim-to-real error** for selected IMU configurations.

config	IMU	GA [°]	Jitter[-]	RE2[m]	RE10[m]
<i>plvs, hd, wr_r, wr_l, sh_r, sh_l</i>	real	14.4	2.8	0.25	0.25
	synth	7.0	2.8	0.09	0.17
<i>plvs, hd, sh_r, sh_l</i>	real	24.9	2.7	0.26	0.26
	synth	22.3	2.5	0.17	0.17
<i>plvs, sh_r, sh_l</i>	real	36.4	2.8	0.30	0.30
	synth	28.1	2.5	0.21	0.21
<i>sh_r, sh_l</i>	real	39.2	4.3	0.47	0.47
	synth	29.9	3.1	0.26	0.26

Despite the many sources of sim-to-real-error our model performs well on real data (see also Section 4.3). The following are sources that contribute to sim-to-real error: (1) imperfect sensor-to-bone orientation calibration, (2) relative motion between sensor and bone, (3) imperfect placement of the IMU with respect to vertex positions used during data synthesis, (4) IMU orientation estimation error by the proprietary algorithm and (5) differences synthetic and real

Table 3. **Comparison of DiffusionPoser to different baselines for TotalCaptureReal.** Numbers are averages over all trials, bracketed numbers are the metrics from the trial with the highest error.

system	LA [°]	GA [°]	JPE [cm]	Jitter [-]	RE2s [m]	RE5s [m]	RE10s [m]
Transpose	13.9(24.3)	16.1(25.2)	6.4(12.7)	5.0(9.7)	0.19(0.98)	0.26(1.38)	0.29(1.4)
PIP	11.9(20.1)	14.4(22.7)	5.3(10.7)	1.1(1.5)	0.12(0.45)	0.17(0.60)	0.27(1.07)
TIP	12.0(20.0)	14.3(21.4)	6.2(12.1)	5.3(13.0)	0.17(0.46)	0.33(0.53)	0.32(0.98)
DiffusionPoser	13.0(21.5)	14.4(23.4)	6.1(12.3)	2.8(4.5)	0.14(0.41)	0.20(0.6)	0.25(0.75)

data noise properties. We noted that (1) and (2) are substantial in ‘TotalCaptureReal’.

4.3. Comparison to baselines for six IMUs

We compare DiffusionPoser to Transpose, PIP and TIP on ‘TotalCaptureReal’ for the six IMU sensor configuration: pelvis, head, wrists, shanks (Table 3). We reran and evaluated Transpose, PIP and TIP to ensure we were making a fair comparison starting from the exact same input data and reporting the same evaluation metrics.

Evaluation metrics for all four systems are close. DiffusionPoser is within 1.1 degrees and 1cm of the system with the best scores for angular error and joint position error. Optical motion capture will result in variations that are similar or even larger than 1 degree for the exact same underlying motion performed by the same subject due to marker noise, differences in marker placement and parameter settings of the motion capture processing pipeline [5, 49]. The jitter metric is significantly lower in PIP than in DiffusionPoser. Several trials reconstructed by PIP have less jitter than the ground truth motion, indicating that PIP could be oversmoothing in such cases. From our videos it is clear that our reconstructed motion smoothness is reasonable.

We performed an additional validation with real IMU data using the DIP-IMU dataset [11] with comparison to Transpose, PIP and TIP. Conclusions are the same with evaluation metrics being even closer as for TotalCaptureReal. Results are in Appendix Table 3.

4.4. Multimodality: Adding shoe insoles

Because we use contact information while reconstructing motion we can exploit ground truth information when pressure insoles are worn. We performed a quantitative analysis on ‘TotalCaptureReal’ by using ground truth contact labels as part of the measured signals. Using ground truth contact labels improved accuracy (Table 4). When wearing IMUs on the shank the improvement from wearing insoles is less because the shank IMUs provide good information to estimate the contact label.

4.5. Dealing with signal corruption and loss

DiffusionPoser can deal with corrupted or lost signal by relying on its generative nature. Signal loss and corrup-

Table 4. **Adding insoles for ground truth contact labels improves reconstruction accuracy.**

config	insole	GA[°]	Jitter[-]	RE2[m]	RE10[m]
<i>plvs, hd, sh_r, sh_l</i>	No	29.4	2.7	0.14	0.33
	Yes	29.1	2.6	0.12	0.26
<i>plvs, hd, wr_r, wr_l</i>	No	18.0	3.2	0.34	0.96
	Yes	16.9	3.0	0.3	0.67

tion occurs regularly in practical settings (e.g. packet loss, out-of-range). We provide examples on our project website where we drop the signal from all sensors for a couple of seconds. DiffusionPoser continues to generate realistic motion in real-time. Once the signal is back, a natural transition to the actual motion is generated. Regressive models, such as PIP [52] are not capable of such online infilling.

4.6. Reducing denoising steps and optimal spread

For DiffusionPoser to run online we sped up the denoising process by reducing the number of denoising steps following DDIM [7, 38]. For our large base model, a NVIDIA A4000 GPU can achieve real-time at 30 denoising steps, but for weaker GPUs it is required to reduce to 10 or even 5 denoising steps. Here we analyzed how accuracy degrades with reducing the number of denoising steps (Table 5) and how these steps are best spread (Table 6).

Accuracy improved with increasing number of denoising steps but saturated around 30 steps. Scaling up to 100 denoising steps only gave a marginal improvement. When reducing the number of denoising steps to 5 we sacrificed the jitter metric mostly.

Concerning different denoising schemes (Table 6), for 10A we only used the last 10 denoising steps. For 10B we used the last 20 denoising steps with a spread of 2. For 10C we started at step 100 and decayed exponentially to zero. 10D has the following irregular spread: 1000/850/700/550/400/250/100/10/2/0. We observed a trade-off where starting from high noise levels resulted in slightly worse angular errors and more jitter but improved root estimation error. This could be explained as we initialized the unknown features of the last frame as a copy of the prior frame. Since orientations are smooth, a few denoising steps are enough for the model to make a good prediction

Table 5. **Evaluation of decreasing numbers of denoising steps for the 6 IMU configuration.**

#steps	GA[°]	Jitter[-]	RE2[m]	RE10[m]
5	16.0	4.4	0.17	0.32
10	15.8	3.3	0.15	0.25
30	14.4	2.8	0.14	0.25
100	14.3	2.8	0.14	0.24

Table 6. **Evaluation of differently spreading denoising steps.**

spread	GA[°]	Jitter[-]	RE2[m]	RE10[m]
10A	15.3	2.0	0.42	0.97
10B	15.4	2.1	0.42	0.89
10C	15.4	2.4	0.45	0.67
10D	15.8	3.3	0.15	0.25

of the next frame. The change in root position (Δp) is not smooth and adding more noise allows the model to change the initialization of the last frame more. For all other evaluations in this paper we used a spacing similar to 10D as the increase in angular error and jitter was small relative to the gain in root estimation accuracy.

4.7. Additional experiments and ablations

We performed several additional experiments and ablations (Table 7). We found that reducing the number of transformer layers is a more effective way to reduce model size than reducing the feature dimension. Interestingly, and we do not know why, root translation error was slightly better in smaller models. We found that not using the energy metric for sampling during training increased the GA by 1°. Next, we found that ablating height as a condition did not change much. This could in part be explained by the subjects in TotalCapture having an average stature. Finally, it is clear that using the contact label predictions to perform root motion correction improves root translation error. Root correction slightly increased jitter.

We performed a comparison of DiffusionPoser with a regressive model that takes in an input mask to encode sensor configuration. This was only implemented for the OpenSim version of DiffusionPoser and we show that DiffusionPoser is more accurate and robust than the regressive model (Appendix A, Table 2).

4.8. Live demonstration

We tested our system live with different configurations in an indoor and outdoor setting (tennis court). Using our system requires a quick sensor-to-world calibration. Next, the sensors are attached at the correct body locations and a sensor-to-bone calibration is done by standing in T-pose.

Table 7. **Ablations.** Transformer model size is reported between brackets: layers/feature dimension/feedforward dimension

Ablation	GA[°]	Jitter[-]	RE2[m]	RE10[m]
Ours (8/512/2048)	14.4	2.8	0.14	0.25
model size (4/512/2048)	14.9	2.7	0.11	0.22
model size (8/256/1024)	15.3	2.7	0.11	0.25
Ours w/o energy metric	15.4	2.6	0.14	0.28
Ours w/o height	14.5	2.8	0.14	0.27
Ours w/o root correction	14.4	2.15	0.32	0.54

We use XSens IMU sensors and stream processed (MTw Awinda) orientation and linear acceleration at 60Hz. We apply a moving average filter with current frame and five past and five future frames on the acceleration and downsample both signals to 20Hz to match our synthetic training data.

We perform indoor experiments with a mixture of motions including walking, jogging, jumping, lunges and kicking and some specific gait deviations that are observed in patient populations. Next we also perform a live demonstration of our system with a tennis player outdoors.

The latency between the real and visualized motion results from (1) XSens to system communication, (2) moving average filter (83ms) and (3) our reconstruction algorithm (45ms). We noticed that the latency of the XSens system increased when the distance to the receiver station increased. The latency of our reconstruction algorithm is similar to TIP [13], and significantly more than PIP (16ms) [52].

5. Conclusion

DiffusionPoser allows immediate use of arbitrary sensor configurations and thus optimizing these configurations for specific activities. This generative model provides robustness in motion prediction against sparsity, noisy and corrupted sensor measurements. Our evaluations show state-of-the-art performance when using six sensors and little degradation in performance when using fewer sensors.

The latency of DiffusionPoser is a limitation for applications that require short latencies such as visual illusion (<50ms according to [39]) or control of an assistive device (e.g., 40-60ms to assist balance with an exoskeleton according to [4]). Other future work is to extend our system to estimate joint torques and muscle forces. This could lead to a breakthrough system for biomechanists and users interested in health and performance.

References

- [1] Chaitanya Ahuja and Louis-Philippe Morency. Language2pose: Natural language grounded pose forecasting. *2019 International Conference on 3D Vision (3DV)*, pages 719–728, 2019. [3](#)
- [2] Mazen Al Borno, Johanna O’Day, Vanessa Ibarra, James Dunne, Ajay Seth, Ayman Habib, Carmichael Ong, Jennifer Hicks, Scott Uhrlich, and Scott Delp. Opensense: An open-source toolbox for inertial-measurement-unit-based measurement of lower extremity kinematics over long durations. *Journal of neuroengineering and rehabilitation*, 19(1):1–11, 2022. [1](#)
- [3] Sheldon Andrews, Ivan Huerta, Taku Komura, Leonid Sigal, and Kenny Mitchell. Real-time physics-based motion capture with sparse sensors. In *Proceedings of the 13th European Conference on Visual Media Production (CVMP 2016)*, CVMP 2016, New York, NY, USA, 2016. Association for Computing Machinery. [3](#)
- [4] Owen N Beck, Max K Shepherd, Rish Rastogi, Giovanni Martino, Lena H Ting, and Gregory S Sawicki. Exoskeletons need to react faster than physiological responses to improve standing balance. *Science robotics*, 8(75):eadf1080, 2023. [8](#)
- [5] A Cappozzo, F Catani, A Leardini, MG Benedetti, and U Della Croce. Position and orientation in space of bones during movement: experimental artefacts. *Clinical Biomechanics*, 11(2):90–100, 1996. [7](#)
- [6] Chariklia Chatzaki, Vasileios Skaramagkas, Nikolaos Tachos, Georgios Christodoulakis, Evangelia Maniadi, Zinovia Kefalopoulou, Dimitrios I Fotiadis, and Manolis Tsiknakis. The smart-insole dataset: Gait analysis using wearable sensors with a focus on elderly and parkinson’s patients. *Sensors*, 21(8):2821, 2021. [1](#)
- [7] Yuming Du, Robin Kips, Albert Pumarola, Sebastian Starke, Ali Thabet, and Artsiom Sanakoyeu. Avatars grow legs: Generating smooth human motion from sparse tracking inputs with diffusion model. In *CVPR*, 2023. [2](#), [3](#), [7](#)
- [8] Jia Gong, Lin Geng Foo, Zhipeng Fan, Qihong Ke, Hossein Rahmani, and Jun Liu. Diffpose: Toward more reliable 3d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023. [3](#)
- [9] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. [4](#)
- [10] Karl Holmquist and Bastian Wandt. Diffpose: Multi-hypothesis human pose estimation using diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15977–15987, October 2023. [3](#)
- [11] Yinghao Huang, Manuel Kaufmann, Emre Aksan, Michael J. Black, Otmar Hilliges, and Gerard Pons-Moll. Deep inertial poser: Learning to reconstruct human pose from sparse inertial measurements in real time. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 37:185:1–185:15, Nov. 2018. Two first authors contributed equally. [1](#), [2](#), [5](#), [7](#)
- [12] Jiaxi Jiang, Paul Strelci, Huajian Qiu, Andreas Fender, Larissa Laich, Patrick Snape, and Christian Holz. Avatarposer: Articulated full-body pose tracking from sparse motion sensing. In *Proceedings of European Conference on Computer Vision*. Springer, 2022. [3](#)
- [13] 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*, SA ’22, New York, NY, USA, 2022. Association for Computing Machinery. [1](#), [2](#), [3](#), [4](#), [5](#), [8](#)
- [14] Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *The IEEE International Conference on Computer Vision (ICCV)*, 2015. [2](#)
- [15] Manuel Kaufmann, Yi Zhao, Chengcheng Tang, Lingling Tao, Christopher Twigg, Jie Song, Robert Wang, and Otmar Hilliges. Em-pose: 3d human pose estimation from sparse electromagnetic trackers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11510–11520, 2021. [1](#)
- [16] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. Flame: free-form language-based motion synthesis & editing. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI’23/IAAI’23/EAAI’23. AAAI Press, 2023. [3](#)
- [17] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [3](#)
- [18] Adrian KM Lai, Allison S Arnold, and James M Wakeling. Why are antagonist muscles co-activated in my simulation? a musculoskeletal model for analysing human locomotor tasks. *Annals of biomedical engineering*, 45:2762–2774, 2017. [3](#)
- [19] Chang June Lee and Jung Keun Lee. Inertial motion capture-based wearable systems for estimation of joint kinetics: A systematic review. *Sensors*, 22(7):2507, 2022. [1](#)
- [20] Jiaman Li, C. Karen Liu, and Jiajun Wu. Ego-body pose estimation via ego-head pose estimation, 2023. [2](#), [3](#)
- [21] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Smpl: A skinned multi-person linear model. *ACM Trans. Graph.*, 34(6), nov 2015. [2](#), [3](#)
- [22] Andreas Lugmayr, Martin Danelljan, Andrés Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11451–11461, 2022. [2](#)
- [23] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *International Conference on Computer Vision*, pages 5442–5451, Oct. 2019. [2](#), [4](#)
- [24] Charles Malleon, Andrew Gilbert, Matthew Trumble, John Collomosse, Adrian Hilton, and Marco Volino. Real-time

- full-body motion capture from video and imus. In *2017 International Conference on 3D Vision (3DV)*, pages 449–457, 2017. [3](#)
- [25] Vimal Molyn, Riku Arakawa, Mayank Goel, Chris Harrison, and Karan Ahuja. Imuposer: Full-body pose estimation using imus in phones, watches, and earbuds. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New York, NY, USA, 2023. Association for Computing Machinery. [2](#)
- [26] Mathis Petrovich, Michael J. Black, and Gül Varol. Action-conditioned 3D human motion synthesis with transformer VAE. In *International Conference on Computer Vision (ICCV)*, 2021. [4](#)
- [27] Mathis Petrovich, Michael J. Black, and Gül Varol. TEMOS: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision (ECCV)*, 2022. [3](#)
- [28] Ryan Po, Wang Yifan, and Vladislav Golyanik et al. State of the art on diffusion models for visual computing. In *ArXiv*, 2023. [3](#)
- [29] Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J. Guibas. Humor: 3d human motion model for robust pose estimation. In *International Conference on Computer Vision (ICCV)*, 2021. [4](#)
- [30] Yiming Ren, Chengfeng Zhao, Yannan He, Peishan Cong, Han Liang, Jingyi Yu, Lan Xu, and Yuexin Ma. Lidar-aid inertial poser: Large-scale human motion capture by sparse inertial and lidar sensors. *IEEE Transactions on Visualization and Computer Graphics*, 29(5):2337–2347, 2023. [3](#)
- [31] Daniel Roetenberg, Henk Luinge, Per Slycke, et al. Xsens mvn: Full 6dof human motion tracking using miniature inertial sensors. *Xsens Motion Technologies BV, Tech. Rep.*, 1:1–7, 2009. [2](#)
- [32] Martin Schepers, Matteo Giuberti, Giovanni Bellusci, et al. Xsens mvn: Consistent tracking of human motion using inertial sensing. *Xsens Technol.*, 1(8), 2018. [1](#)
- [33] Paul Schreiner, Maksym Perepichka, Hayden Lewis, Sune Darkner, Paul G. Kry, Kenny Erleben, and Victor B. Zordan. Global position prediction for interactive motion capture. *Proc. ACM Comput. Graph. Interact. Tech.*, 4(3), sep 2021. [1](#)
- [34] Ajay Seth, Jennifer L Hicks, Thomas K Uchida, Ayman Habib, Christopher L Dembia, James J Dunne, Carmichael F Ong, Matthew S DeMers, Apoorva Rajagopal, Matthew Millard, et al. Opensim: Simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement. *PLoS computational biology*, 14(7):e1006223, 2018. [2](#)
- [35] Mingyi Shi, Kfir Aberman, Andreas Aristidou, Taku Komura, Dani Lischinski, Daniel Cohen-Or, and Baoquan Chen. Motionet: 3d human motion reconstruction from monocular video with skeleton consistency. *ACM Trans. Graph.*, 40(1), sep 2020. [4](#)
- [36] Ronit Slyper and Jessica K. Hodgins. Action capture with accelerometers. SCA '08, page 193–199, Goslar, DEU, 2008. Eurographics Association. [1](#), [2](#)
- [37] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. [3](#)
- [38] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [5](#), [7](#)
- [39] Jan-Philipp Stauffert, Florian Niebling, and Marc Erich Latoschik. Latency and cybersickness: Impact, causes, and measures. a review. *Frontiers in Virtual Reality*, 1:582204, 2020. [8](#)
- [40] Agnieszka Szczesna, Monika Błaszczyszyn, and Magdalena Pawlyta. Optical motion capture dataset of selected techniques in beginner and advanced kyokushin karate athletes. *Scientific Data*, 8(1):13, 2021. [2](#)
- [41] Jochen Tautges, Arno Zinke, Björn Krüger, Jan Baumann, Andreas Weber, Thomas Helten, Meinard Müller, Hans-Peter Seidel, and Bernd Eberhardt. Motion reconstruction using sparse accelerometer data. *ACM Trans. Graph.*, 30(3), may 2011. [2](#)
- [42] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*, 2023. [2](#), [3](#), [4](#)
- [43] Matt Trumble, Andrew Gilbert, Charles Malleon, Adrian Hilton, and John Collomosse. Total capture: 3d human pose estimation fusing video and inertial sensors. In *2017 British Machine Vision Conference (BMVC)*, 2017. [2](#), [5](#)
- [44] Jonathan Tseng, Rodrigo Castellon, and Karen Liu. Edge: Editable dance generation from music. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 448–458, June 2023. [2](#), [3](#), [4](#)
- [45] Scott D. Uhlich, Antoine Falisse, Łukasz Kidziński, Julie Muccini, Michael Ko, Akshay S. Chaudhari, Jennifer L. Hicks, and Scott L. Delp. Opencap: Human movement dynamics from smartphone videos. *PLOS Computational Biology*, 19(10):1–26, 10 2023. [1](#)
- [46] Scott D. Uhlich, Rachel W. Jackson, Ajay Seth, Julie Kolesar, and Scott L. Delp. Muscle coordination retraining inspired by musculoskeletal simulations reduces knee contact force. *Scientific reports*, 2022. [2](#)
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, page 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc. [4](#)
- [48] Timo Von Marcard, Bodo Rosenhahn, Michael J Black, and Gerard Pons-Moll. Sparse inertial poser: Automatic 3d human pose estimation from sparse imus. In *Computer graphics forum*, volume 36, pages 349–360. Wiley Online Library, 2017. [1](#), [2](#)
- [49] Jason M. Wilken, Kelly M. Rodriguez, Melissa Brawner, and Benjamin J. Darter. Reliability and minimal detectable change values for gait kinematics and kinetics in healthy adults. *Gait and Posture*, 35(2):301–307, 2012. [7](#)

- [50] Alexander Winkler, Jungdam Won, and Yuting Ye. Questsim: Human motion tracking from sparse sensors with simulated avatars. In *SIGGRAPH Asia 2022 Conference Papers*, SA '22, New York, NY, USA, 2022. Association for Computing Machinery. 3
- [51] Xinyu Yi, Yuxiao Zhou, Marc Habermann, Vladislav Golyanik, Shaohua Pan, Christian Theobalt, and Feng Xu. Egolocate: Real-time motion capture, localization, and mapping with sparse body-mounted sensors. *ACM Trans. Graph.*, 42(4), jul 2023. 3
- [52] 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)*, June 2022. 1, 2, 3, 5, 6, 7, 8
- [53] 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), jul 2021. 1, 2, 3, 4, 5
- [54] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. 3
- [55] Zerong Zheng, Tao Yu, Hao Li, Kaiwen Guo, Qionghai Dai, Lu Fang, and Yebin Liu. Hybridfusion: Real-time performance capture using a single depth sensor and sparse imus. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 3
- [56] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5738–5746, 2019. 3