

Lifting Motion to the 3D World via 2D Diffusion

Jiaman Li C. Karen Liu[†] Jiajun Wu[†]
Stanford University

{jiamanli, karenliu, jiajunwu}@cs.stanford.edu

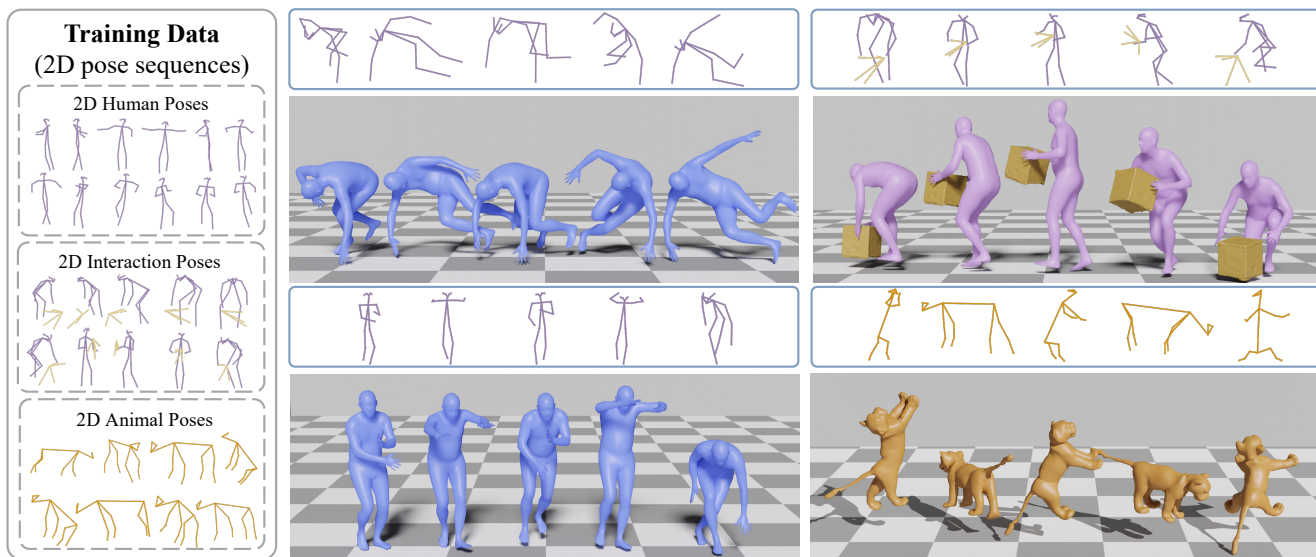


Figure 1. Our framework, MVLift, can be trained only on 2D pose sequences and generate 3D motions including joint rotations and root trajectories in the world coordinate system. The approach generalizes to the various domains of human poses, interactions, and animal poses.

Abstract

Estimating 3D motion from 2D observations is a long-standing research challenge. Prior work typically requires training on datasets containing ground truth 3D motions, limiting their applicability to activities well-represented in existing motion capture data. This dependency particularly hinders generalization to out-of-distribution scenarios or subjects where collecting 3D ground truth is challenging, such as complex athletic movements or animal motion. We introduce MVLift, a novel approach to predict global 3D motion—including both joint rotations and root trajectories in the world coordinate system—using only 2D pose sequences for training. Our multi-stage framework leverages 2D motion diffusion models to progressively generate consistent 2D pose sequences across multiple views, a key step in recovering accurate global 3D motion. MVLift generalizes across various domains, including human poses, human-object interactions, and animal poses. Despite not requiring 3D supervision, it outperforms prior work on five datasets, including those methods that require 3D supervision.

[†] indicates equal advising.

1. Introduction

Generating 3D motion is critical for applications in computer graphics, embodied AI, and robotics. High-quality 3D motion datasets [1, 11, 24, 32, 46], essential for developing generative models that synthesize human movement, are typically obtained using advanced motion capture systems. However, these systems are often confined to lab environments and are costly to scale, particularly for capturing complex motions that require specialized expertise, such as dancing, yoga, and gymnastics. Extensive work has studied monocular video pose estimation to predict 3D human motion from 2D observations without the need for elaborate mocap setups ([15, 43, 66]). Despite advancements, even state-of-the-art methods ([43]) still heavily rely on motion capture datasets and videos with paired 3D ground truth. This dependence limits their ability to generate accurate global 3D motions for out-of-distribution motions—those not well-represented in existing 3D datasets.

To overcome the challenge of generating out-of-distribution 3D motions, recent studies have explored methods to leverage in-domain 2D poses and generate 3D poses using only 2D poses during training. Some efforts train mod-

els exclusively on 2D poses of a single step, often resulting in unsmooth motions and noticeable artifacts when applied to 2D pose sequences [54]. Recent work has employed a pre-trained unconditional 2D motion diffusion model to generate 3D motions from sampled noise [14], demonstrating the ability to produce complex movements not typically covered by existing datasets, such as gymnastics, horse motion, and basketball. However, this approach does not account for the global translation of 3D motion. In this work, we address the challenge of predicting global 3D motion for out-of-distribution scenarios using only domain-specific 2D pose data (e.g., human, animal, or interaction). Unlike prior work that primarily focuses on generating 3D motion with a fixed root joint, our goal is to accurately predict realistic 3D motion with root trajectories in the world coordinate system from 2D pose sequence inputs.

The problem of estimating global 3D motion from 2D pose sequences can be effectively addressed using a single-stage network when 3D training data is available [66]. However, training a single-stage model using solely 2D poses is challenging due to the lack of direct supervision for either 3D motion or consistent multi-view 2D poses. To address this fundamental challenge, we propose a multi-stage approach, MVLift, which progressively establishes multi-view consistency through 2D motion diffusion models, enabling accurate 3D motion recovery without 3D training data.

MVLift consists of four stages. First, we train a line-conditioned diffusion model that learns to predict 2D pose sequences following epipolar lines. These epipolar lines, derived from the fundamental matrix relating two camera views, ensure corresponding points in different views lie along corresponding lines, enforcing basic geometric consistency. Second, we develop a joint optimization method for multi-view 2D motion sequences using two objectives: (1) a multi-view consistency objective to ensure geometric relationships across views, and (2) Score Distillation Sampling (SDS) based on our trained line-conditioned diffusion model to maintain motion realism. While this approach generates plausible multi-view 2D sequences that enable 3D motion recovery, the optimization process cannot guarantee perfect view consistency, affecting the accuracy of reconstructed 3D motions. To address this limitation, we leverage these optimized results to create a synthetic dataset. Although the optimized 3D motions may not perfectly align with input 2D sequences, they preserve motion realism. By reprojecting these synthetic 3D motions to different views, we obtain strictly consistent multi-view 2D sequences. Finally, we use these sequences to train an efficient diffusion model that directly generates multi-view consistent 2D sequences in a single forward pass.

Our work makes three contributions. First, we introduce a novel framework, MVLift, for estimating global 3D motion from a single-view 2D pose sequence without any 3D

training data, addressing a fundamental limitation in existing approaches. Second, we demonstrate how multi-view consistency can be progressively established through 2D motion diffusion, providing a new perspective on 3D motion estimation. Third, we show that MVLift generalizes across various domains (humans, animals, and interactions) and significantly outperforms existing methods on five datasets, even those trained with 3D supervision.

2. Related Work

Human Pose Estimation from 2D Poses. With the advent of datasets containing paired 2D videos and ground truth 3D motions [10, 34, 44], various approaches have been proposed to learn the mapping between 2D human joints and 3D human joints, including fully connected networks [33], temporal convolution networks [37, 37], GCN-based networks [3, 6, 55], and more recently, Transformer-based models [26, 39, 63, 65]. The typical paradigm involves training neural networks to regress 3D poses from 2D pose inputs. However, this is constrained by the limited availability of paired data and is often incapable of predicting accurate 3D poses outside the distribution of paired data. To address this problem, a recent work, MotionBERT [66], leverages the large-scale mocap dataset [32] to generate paired 2D and 3D data for training, showcasing superior results. Nevertheless, the issue of limited data availability persists. To overcome this data challenge, some work explores the problem of estimating 3D poses by training solely on in-domain 2D poses, without any 3D ground truth poses. For example, ElePose [54] predicts 3D pose and camera pose simultaneously with a reprojection loss and a pose prior loss. MAS [14] optimizes 3D motion using the learned 2D motion diffusion model. However, these methods do not account for the global trajectory of 3D motion. In our work, we use a 2D motion sequence as input and predict global 3D motion, including joint rotations and a root trajectory in the world coordinate system, without training on any 3D motion data.

Human Pose Estimation from Videos. Estimating SMPL [30] parameters from images [2, 12, 16, 17, 19, 22, 23, 36, 62] or videos [13, 15, 40, 43, 53, 56] has been extensively studied in recent years. Reconstructing human mesh from images has been improved by incorporating optimization techniques [2] into the learning pipeline [19] and leveraging more powerful backbones [8] such as Transformers [52]. In terms of recovering human mesh from videos, prior works [5, 15, 31] employ recurrent neural networks to integrate temporal information and produce smooth 3D motion predictions. However, these works do not predict global root trajectories. Recently, there has been growing attention on predicting global 3D motion from videos [18, 43, 45, 60, 61, 64]. Some works leverage the success of local 3D pose estimation as initializa-

tion and use motion priors to further optimize global pose and camera parameter predictions [60, 61]. A recent work, WHAM [43], developed a pipeline that generates 3D motion in the world coordinate frame with a single forward pass. However, the global trajectory prediction module in this pipeline was trained on datasets containing ground truth 3D motion [10, 32, 34], which constrains its ability to predict accurate root trajectories for motions outside of the data distribution. In this work, we generate global 3D motions from 2D pose sequences without RGB data input. Our approach does not rely on any 3D motion data or paired video and 3D data for training, leveraging only 2D pose sequences extracted from monocular videos. The approach generalizes to various domains such as animal motions and human-object interactions, lifting 2D pose sequences to global 3D motion.

Multi-View Image Generation. Based on a large-scale 3D object dataset [7], considerable research has explored generating 3D objects through novel view image synthesis. Zero-1-to-3 [27] pioneered the approach of using a single image and a specified camera pose as inputs to predict the corresponding image from a new camera view. Subsequent works [4, 28, 29, 41, 42, 47, 49, 58] have advanced the field by directly training multi-view image diffusion models that can simultaneously generate consistent multi-view images. Recent efforts have also extended this idea to the more challenging task of multi-view video generation [20, 51]. Inspired by these advancements, we reformulate the prediction of 3D motion as the generation of multi-view 2D motion sequences. Unlike prior efforts, which require synchronized multi-view images during training, our work circumvents the need for consistent multi-view 2D motion sequences. Instead, we train exclusively on single-view 2D motion sequences extracted from monocular videos.

3. MVLift

Our goal is to estimate global 3D motion, including both joint rotations and root trajectories in the world coordinate system, from a single-view 2D pose sequence. For motions well-represented in existing 3D datasets, this task can be addressed using a single-stage network trained with 3D supervision, as demonstrated in prior work [66]. However, for many out-of-distribution scenarios such as complex athletic movements or animal motions, 3D motion data and multi-view recordings are often unavailable or difficult to collect. In this case, a single-stage approach becomes infeasible due to the lack of direct 3D supervision. To address this challenge, we propose a multi-stage framework, MVLift, which only requires readily available single-view 2D data and leverages learned 2D motion diffusion priors to progressively establish multi-view consistency for 2D pose sequences across different views, enabling accurate reconstruction of realistic 3D motions without requiring any 3D supervision.

The key insight is that while a single 2D sequence provides limited 3D information, a diffusion model trained on diverse 2D motions can learn rich priors about how poses appear from different views. We use the learned 2D diffusion priors combined with geometric constraints to progressively strengthen the consistency of multi-view 2D sequences.

As illustrated in Figure 2, our framework operates in four stages, each strengthening multi-view consistency in different ways. In Stage 1, we train a line-conditioned diffusion model that generates 2D pose sequences respecting epipolar line constraints. While this ensures basic pairwise consistency between views, achieving global multi-view consistency requires jointly considering all views. In Stage 2, we establish stronger multi-view consistency through an optimization method that directly optimizes multi-view 2D sequences using explicit multi-view consistency objectives, while maintaining motion realism through Score Distillation Sampling (SDS) based on our trained diffusion model. Although this optimization approach generates plausible multi-view 2D sequences, achieving perfect geometric consistency across all views through optimization alone remains challenging. To overcome this limitation, we leverage these roughly consistent multi-view sequences in Stage 3 to recover plausible 3D motions through optimization with 2D reprojection objectives. By reprojecting these recovered 3D motions to multiple views, we obtain a synthetic dataset of strictly consistent multi-view 2D sequences. Finally, in Stage 4, we train a specialized diffusion model to directly generate multi-view consistent 2D sequences in a single forward pass.

This progression from pairwise geometric constraints to global optimization to learned multi-view generation allows us to obtain high-quality 3D motion estimation without any ground truth 3D data for training. The following sections detail each component of our framework and how they work together to achieve 3D motion reconstruction.

3.1. Line-Conditioned 2D Motion Diffusion Model

We denote a 2D motion sequence as $\mathbf{X} \in \mathbb{R}^{T \times J \times 2}$, where T is the number of time steps and J is the number of joints. Each joint position at time step t is denoted by (x_t^j, y_t^j) . We define line conditions $\mathbf{L} \in \mathbb{R}^{T \times J \times 3}$, where each joint’s line is parameterized by coefficients (a_t^j, b_t^j, c_t^j) following the equation $a_t^j x_t^j + b_t^j y_t^j + c_t^j = 0$. Given line conditions, we learn to generate a 2D pose sequence where each predicted joint position $(\hat{x}_t^j, \hat{y}_t^j)$ aligns with its input line.

Epipolar Line Generation. Epipolar geometry provides geometric constraints that help ensure multi-view consistency in our approach. During testing, given a 2D pose sequence \mathbf{X} and a specified (rather than real) camera intrinsics matrix \mathbf{K} , we compute epipolar lines across different camera views using essential matrices $\mathbf{E}_v \in \mathbb{R}^{3 \times 3}$, where $v \in 1, \dots, 5$ indexes the camera views. These essential matrices encode the relative rotations and translations between




Figure 2. Overview of our multi-stage framework. In Stage 1, we train a 2D motion diffusion model conditioned on simulated epipolar lines. Stage 2 utilizes this model to optimize multi-view 2D motion sequences, achieving only roughly consistent sequences. These sequences are used for 3D motion optimization in Stage 3, and the synthetic 3D data is then reprojected into strictly consistent multi-view 2D sequences. In Stage 4, we train a multi-view 2D motion diffusion model on these data to efficiently generate consistent 2D sequences across views.

views and are used to compute fundamental matrices as $\mathbf{F}_v = \mathbf{K}^{-T} \mathbf{E}_v \mathbf{K}^{-1}$. For a 2D point $\mathbf{x} = (x, y, 1)^T$ in the input view, its corresponding epipolar line in view v is computed as $\mathbf{l}_v = \mathbf{F}_v \mathbf{x}$, with coefficients (a_v, b_v, c_v) defining the line equation $a_v x + b_v y + c_v = 0$. During training, since only single-view 2D pose sequences are available, we simulate epipolar constraints by randomly sampling a virtual epipole for each sequence. For each joint j at timestep t , we create a line connecting it to this virtual epipole, resulting in line coefficients (a_t^j, b_t^j, c_t^j) that serve as conditions for training our line-conditioned diffusion model.


Conditional Diffusion Model. Given line conditions \mathbf{L} , we train a conditional diffusion model that learns to generate 2D motion sequences that respect epipolar geometry constraints. The diffusion model consists of a forward diffusion process and a reverse diffusion process. The forward diffusion process progressively adds noise to the data representation \mathbf{X} until, after n steps, it resembles Gaussian noise. The reverse diffusion process involves learning a neural network that denoises the data representation starting from a noisy data state \mathbf{X}_n . Each denoising step is represented as,

$$p_\theta(\mathbf{X}_{n-1} | \mathbf{X}_n, \mathbf{L}) := \mathcal{N}(\mathbf{X}_{n-1}; \boldsymbol{\mu}_\theta(\mathbf{X}_n, n, \mathbf{L}), \boldsymbol{\Sigma}_n), \quad (1)$$

where θ represents the parameters of a neural network, $\boldsymbol{\mu}_\theta$ represents the predicted mean, and $\boldsymbol{\Sigma}_n$ is a fixed variance. Predicting the mean is re-parameterized as predicting the clean data representation \mathbf{X}_0 . The training loss is defined as a reconstruction loss following prior work [9],

$$\mathcal{L} = \mathbb{E}_{\mathbf{X}_0, n} \|\hat{\mathbf{X}}_\theta(\mathbf{X}_n, n, \mathbf{L}) - \mathbf{X}_0\|_1. \quad (2)$$

We implement the denoising network using a transformer architecture [52], which has proven effective for motion synthesis [24, 48, 50]. The line conditions \mathbf{L} are incorporated by



concatenating them with the noisy pose features \mathbf{X}_n along the feature dimension before feeding into the transformer.

Line Matching Loss. We introduce an additional loss to ensure the predicted 2D joint positions conform to the input line conditions. This loss, defined as the perpendicular distance from each joint to its associated line, encourages the generated joints to satisfy the geometric constraints imposed by these lines. The loss is formulated as

$$\mathcal{L}_{\text{line}} = \sum_{t=1}^T \sum_{j=1}^J |a_t^j \hat{x}_t^j + b_t^j \hat{y}_t^j + c_t^j|, \quad (3)$$

where $(\hat{x}_t^j, \hat{y}_t^j)$ are predictions of the j -th joint at time step t .

3.2. Multi-View 2D Motion Optimization

While our line-conditioned diffusion model can generate 2D pose sequences for unobserved views using epipolar line constraints, these generated sequences may lack consistency across all views. This occurs because epipolar geometry only enforces pairwise constraints between each unobserved view and the input view, without ensuring global multi-view consistency. Such inconsistencies can lead to suboptimal 3D reconstruction results. To address this limitation, we propose a multi-view optimization strategy that jointly optimizes 2D motion sequences $\{\mathbf{X}_k\}$, $k = 1, 2, \dots, 5$ across all unobserved views. The camera views are arranged in a circular configuration around the origin, with the input view and five unobserved views separated by 60-degree intervals. Given this structured camera setup, our approach combines two optimization objectives: one objective is to ensure each 2D motion sequence is realistic, achieved using Score Distillation Sampling (SDS), and the other is to maintain consistency among different camera views, achieved

through our designed multi-view consistency loss. We detail each optimization objective as follows.

Score Distillation Sampling. We adapt Score Distillation Sampling (SDS) proposed in DreamFusion [38] to optimize multi-view 2D pose sequences. The goal is to ensure that the 2D pose sequence in each view conforms to the distribution determined by the learned line-conditioned 2D motion diffusion model. The gradient for SDS is computed as follows:

$$\nabla_{\Phi_{2D}} \mathcal{L}_{SDS} = \mathbb{E}_{n, \epsilon} [\omega(n) (\epsilon_{\theta}(\mathbf{X}_n, n, L) - \epsilon)], \quad (4)$$

where the weight $\omega(n)$ is determined by the scheduler of the diffusion model, \mathbf{X}_n represents a perturbed 2D pose sequence at noise level n , and ϵ is random noise.

Multi-View Consistency Loss. We introduce a multi-view consistency loss to enforce geometric constraints across different views. Given an input view and five unobserved views ($V = 6$ total views), we form $M = \binom{6}{2} = 15$ pairs by considering all possible combinations of two different views. The loss is formulated as

$$\mathcal{L}_{\text{multi-view}} = \frac{1}{2M} \sum_{m=1}^M \left(\mathcal{L}_{\text{line}}^{(v \rightarrow w)} + \mathcal{L}_{\text{line}}^{(w \rightarrow v)} \right), \quad (5)$$

where $\mathcal{L}_{\text{line}}^{(v \rightarrow w)}$ measures how well the joints in view w satisfy the epipolar constraints derived from view v :

$$\mathcal{L}_{\text{line}}^{(v \rightarrow w)} = \sum_{t=1}^T \sum_{j=1}^J |a(v, w)x_{j,t}^w + b(v, w)y_{j,t}^w + c(v, w)|. \quad (6)$$

Here, $(x_{j,t}^w, y_{j,t}^w)$ denotes the position of joint j at time t in view w , and $(a(v, w), b(v, w), c(v, w))$ are the coefficients of the epipolar line in view w computed from the corresponding point in view v using the fundamental matrix $\mathbf{F}_{v,w}$. By computing this loss between all pairs of views, we encourage geometric consistency across the entire multi-view setup.

3.3. Synthetic Multi-View 2D Data Generation

The prior stage produces roughly consistent multi-view 2D pose sequences that enable plausible 3D motion reconstruction, although they may deviate from the input 2D pose sequence. In this stage, we create a synthetic dataset of strictly consistent multi-view 2D pose sequences using these optimized results. The process consists of three steps. First, we recover 3D joint positions by minimizing reprojection errors across all optimized multi-view 2D sequences. Following [14], we then fit SMPL parameters to these 3D joints using VPoser [35], which provides plausible joint rotations and root translations. Finally, we reproject each fitted 3D motion sequence into four camera views (arranged with 90-degree intervals) to obtain strictly consistent multi-view 2D pose sequences. The resulting synthetic dataset serves as training data for our multi-view 2D motion diffusion model described in the next section.

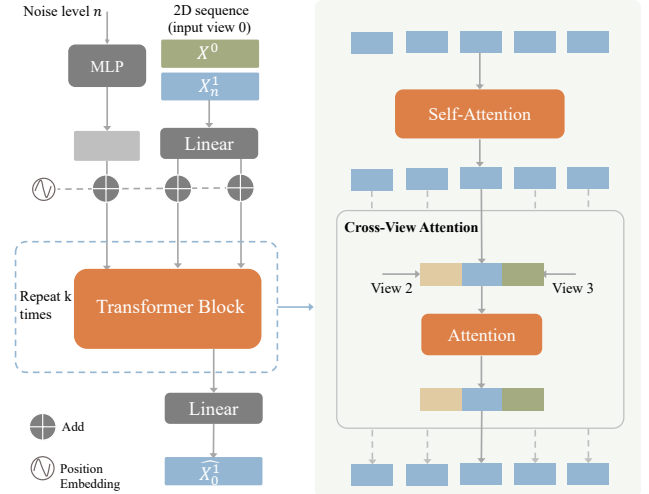


Figure 3. Denoising network of the multi-view 2D motion diffusion model, using View 1 as an example for illustration.

3.4. Multi-View 2D Motion Diffusion Model

Based on the synthetic multi-view 2D motion dataset obtained in the prior stage, we propose a multi-view 2D motion diffusion model to generate consistent multi-view 2D motion sequences simultaneously from a single 2D sequence input.

Model Architecture. We extend the conditional diffusion model from Stage 1 to handle multi-view generation. Given an input 2D pose sequence \mathbf{X}^0 of view 0, our model generates corresponding sequences for three additional views, denoted as $\mathbf{X}^k, k = 1, 2, 3$. The denoising network uses a transformer-based architecture where each transformer block is augmented with a cross-view attention layer following its self-attention layer, as illustrated in Figure 3.

4. Experiments

We first introduce the datasets and evaluation metrics used. Then, we compare our approach against prior work across various datasets, including human motion, animal motion, and human-object interaction motion data. Additionally, we conduct ablation studies to evaluate each stage.

4.1. Datasets and Evaluation Metrics

AIST++ [25] consists of multi-view RGB videos and corresponding 3D dance motions. We use this dataset to evaluate the accuracy of predicted 3D joint positions. During training, we randomly select one video from each sequence and extract 2D poses using ViTPose [57], intentionally excluding videos from other camera views to maintain a setting where only single-view videos are available.

Steezy consists of dance videos from an online dance teaching website and was originally collected by prior work to study dance motion synthesis [21]. We processed the videos using ViTPose [57] to obtain 2D pose sequences.

Method	AIST++ [25]						Steezy			NicoleMove			CatPlay		
	T_{root}	MPJPE	PA-MPJPE	J_{2D}	J_{2D}^C	FID	J_{2D}	J_{2D}^C	FID	J_{2D}	J_{2D}^C	FID	J_{2D}	J_{2D}^C	FID
MotionBERT [66]	101.6	134.0	108.6	59.6	32.3	3.6	67.5	31.3	14.1	78.8	62.5	6.4	N/A	N/A	N/A
WHAM [43]	164.3	104.8	75.1	94.3	24.1	2.6	60.8	22.5	12.1	69.4	43.0	6.9	N/A	N/A	N/A
ElePose [54]	N/A	269.4	215.1	104.9	57.1	4.8	132.7	47.6	13.3	157.4	91.1	13.3	N/A	N/A	N/A
SMPLify [2]	77.4	171.6	146.7	28.8	18.8	2.6	22.8	14.3	13.1	45.1	28.6	7.8	58.8	42.0	12.4
MAS [14]	N/A	191.1	155.6	81.5	38.1	4.4	106.4	61.6	16.1	100.1	56.4	6.3	180.6	59.7	22.0
MVLift (ours)	67.6	110.7	79.2	14.0	12.8	1.9	11.7	11.7	12.4	26.2	22.7	5.1	57.0	39.4	8.5

Table 1. Results of human pose lifting and animal pose lifting. The baselines MotionBERT [66] and WHAM [43] require training on ground truth 3D motions, while the other approaches do not rely on 3D data for training.

NicoleMove is our newly collected dataset, which consists of in-the-wild videos from a YouTube channel. We extracted 2D motion sequences using ViTPose [57]. The dataset primarily contains motions for yoga, pilates, and fitness.

CatPlay is also our newly collected dataset, which contains videos of cats playing, captured in an indoor environment. We processed these videos using an advanced approach for generic 2D keypoint estimation [59].

OMOMO [24] consists of synchronized 3D human motions and object motions. We evaluate our approach on a subset of OMOMO that includes interactions with a specific object, the *largebox*. As the dataset does not include RGB videos, for each sequence, we randomly reproject the 3D motion into a single-view 2D sequence to form a dataset containing 2D human poses and 2D object keypoints.

Evaluation Metrics. We conduct evaluations on datasets with available 3D motion (AIST++, OMOMO) and datasets with only 2D motion (Steezy, NicoleMove, CatPlay). For datasets with 3D motion, we evaluate the predicted 3D motion by computing the root translation error (T_{root}), the mean per-joint position error (**MPJPE**), and the mean per-joint position error with Procrustes Alignment (**PA-MPJPE**), following standard pose estimation protocols [43]. For datasets with only 2D motion, we assess the reprojected 2D motion from the predicted 3D. That is, we calculate the 2D joint position error (J_{2D}) and the 2D joint position error with the 2D root joint translated to the image center (J_{2D}^C). Additionally, we train a 2D motion feature extractor on each dataset and use it to compute the Fréchet Inception Distance (**FID**) for 2D motions reprojected in other camera views, excluding those from the input view. For the object interaction dataset, OMOMO [24], we include two additional metrics: T_{root}^O , which denotes the object translation error, and **O-MPJPE**, which represents the error in 3D object keypoints with the object translated to the origin.

4.2. Human Pose Lifting

Baselines. We compare our approach with two categories of baselines. The first category includes baselines that require training their model on 3D ground truth motion data, which provides them with a significant data advantage. This

category includes MotionBERT [66], which is trained solely on AMASS, and WHAM [43], which utilizes AMASS [32] along with various other datasets containing paired videos and 3D human motions. The second category of baselines does not rely on 3D motion data for training. SMPLify [2] conducts optimization using 2D reprojection objectives to obtain SMPL [30] parameters. ElePose [54], designed for single 3D pose lifting, proposes learning a 2D-to-3D lifting network using reprojection loss and a 2D pose prior loss. MAS [14] leverages a 2D motion diffusion model to perform unconditional 3D motion generation. We modified MAS to take an input 2D pose sequence as an optimization objective for comparison. We train ElePose and MAS for each dataset using their official codebases.

Results. We evaluate our approach and baselines on the test sets of AIST++, Steezy and NicoleMove, as shown in Table 1. We outperform all baselines across several metrics, including those evaluating root trajectory and reprojected 2D poses of the input and other camera views. In terms of 3D joint position errors, our results are superior to all methods that do not require 3D motion data during training, and we also outperform MotionBERT [66], which requires training on AMASS [32]. We achieve comparable results to the baseline WHAM [43] in 3D joint position errors with a significant improvement in the root translation metric.

We showcase qualitative comparisons for AIST++ in Figure 4. Note that ElePose and MAS cannot predict root trajectory. ElePose often generates unrealistic 3D poses due to training instability because its pre-trained 2D pose prior is insufficient to prevent unrealistic 2D poses. MAS predicts mean poses that poorly match the input 2D poses because its unconditional 2D motion diffusion model struggles to produce consistent multi-view 2D sequences essential for accurate 3D motion optimization. MotionBERT fails to produce highly dynamic 3D motions as it is trained on the AMASS dataset, which limits its performance on motions outside the dataset distribution. WHAM produces unrealistic root trajectories, deviating from ground truth, because it is trained on constrained 3D motion datasets and cannot generalize well beyond this data. SMPLify, optimizing only based on 2D joint positions, often results in 3D poses with sud-

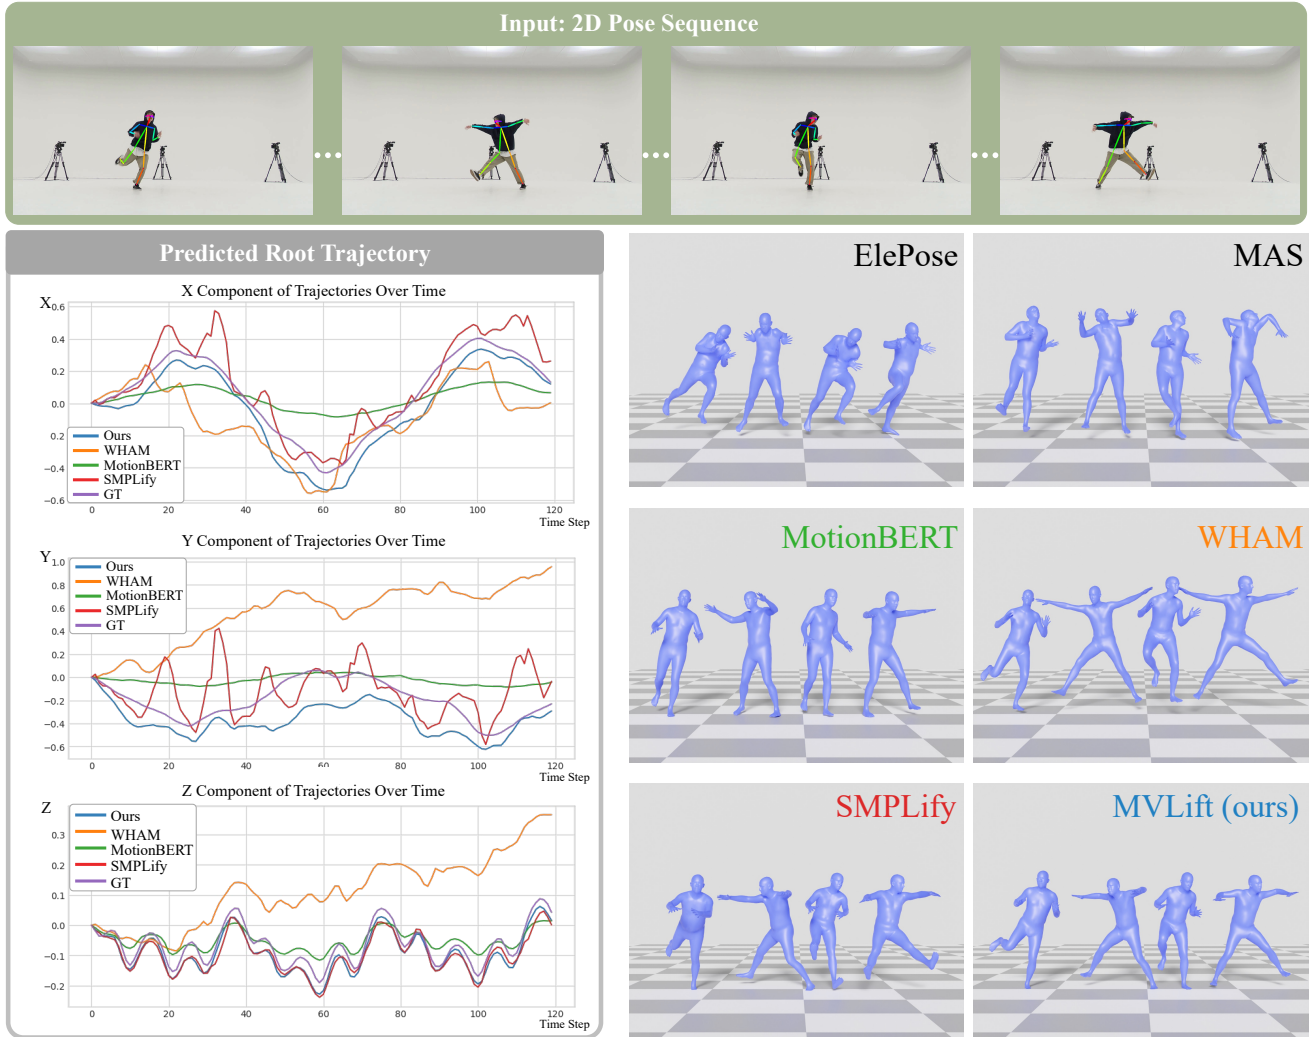


Figure 4. Qualitative results of AIST++. On the left, we show the trajectories for each component (x, y, z) of the predicted root trajectory.

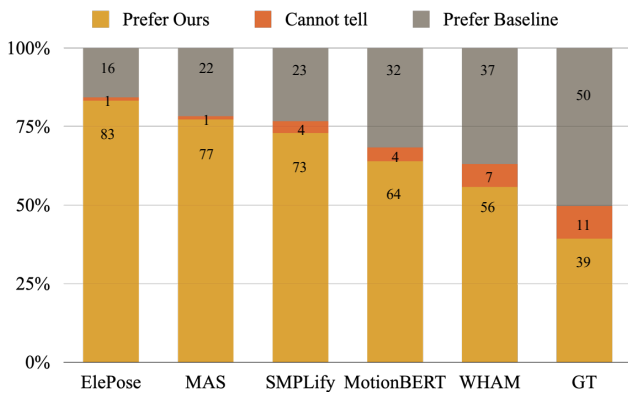


Figure 5. Results of human perceptual studies.

den and unrealistic depth changes due to unresolved depth ambiguity.

Human Perceptual Study. We conducted a human perceptual study to complement the evaluation of generated 3D motion quality. For each dataset (AIST++, Steezy, Nicole-

Move), we selected 15 generated motion sequences from each approach and created pairs comparing our approach with each baseline. The evaluation was conducted using Amazon Mechanical Turk (AMT). As shown in Figure 5, participants preferred our results over all baselines. Additionally, when comparing our AIST++ results with ground truth motions, 39% of participants preferred our results, and 11% found no noticeable difference between ours and ground truth, demonstrating the realism of our generated sequences.

4.3. Animal Pose Lifting

Baselines. Since some baselines used in evaluating human pose lifting are designed specifically for human pose predictions, we compare our approach with two baselines that are suitable for animal poses. These include MAS [14] and modified SMPLify [2], adapted to use the SMAL model [67].

Results. As shown in Table 1, we outperform all baselines. Qualitative results are showcased in Figure 6. While MAS




Figure 6. Qualitative result comparisons of CatPlay.




Figure 7. Qualitative result comparisons of OMOMO.

cannot generate 3D motions with root trajectories and SMPLify produces unrealistic depth changes due to its single-view limitation, our method generates realistic 3D motions with natural root trajectories.

4.4. Human-Object Interaction Lifting

Baselines. Since no existing work has addressed the task of predicting 3D human-object interaction motions in a world coordinate system from 2D object keypoints and 2D human poses, we adapt SMPLify as a baseline for comparison. This baseline reconstructs 3D human motion and object motion by optimizing SMPL parameters and 3D object pose to match the input 2D keypoint positions of human and object.

Results. We present comparisons in Table 2 and Figure 7. SMPLify fails to produce plausible human-object interactions due to significant depth ambiguity from single-view observations. Its predictions lack realistic hand-object contact, and the predicted object motions often contain abrupt changes along the depth direction. Our approach produces realistic human-object interactions by leveraging our generated multi-view 2D pose sequences, which are geometrically consistent across views, for 3D interaction optimization.

4.5. Ablation Study

We conduct ablations on the AIST++ dataset, which contains 3D ground truth motions for comprehensive evaluation. We present comparisons in Table 3.

Method	\mathbf{T}_{root}	MPJPE	PA-MPJPE	\mathbf{T}_{root}^O	O-MPJPE
SMPLify	97.9	142.0	99.3	751.8	106.7
MVLift (ours)	54.9	67.0	53.4	172.9	76.9

Table 2. Results of interaction lifting on OMOMO [24].

Method	\mathbf{T}_{root}	MPJPE	PA-MPJPE	\mathbf{J}_{2D}
MVLift-Stage 1	73.1	135.2	104.4	31.0
MVLift-Stage 2	65.3	127.4	96.2	19.7
SDS for 3D	72.9	137.3	103.5	25.2
SDS for 3D, w/o l_{epi}	752.3	230.4	186.2	54.9
MVLift (ours)	67.6	110.7	79.2	14.0

Table 3. Ablation study on AIST++ [25].

Evaluation of Each Stage. We compare 3D motions obtained using generated multi-view 2D sequences from different stages. MVLift-Stage 1 uses 2D sequences directly generated from the line-conditioned 2D motion diffusion model using epipolar line conditions. MVLift-Stage 2 uses multi-view 2D sequences optimized with SDS and multi-view consistency loss. As shown in Table 3, our final approach (MVLift) using the learned multi-view 2D motion diffusion model achieves more accurate 3D motion estimation than MVLift-Stage 1 and MVLift-Stage 2.

Evaluation of Stage 2 Ablations. In Stage 2, we optimize multi-view 2D pose sequences using SDS and consistency loss. One alternative approach is directly optimizing 3D motion for multi-view consistency while using SDS on the reprojected multi-view 2D sequences to maintain motion realism (SDS for 3D). However, this approach causes the reprojected 2D poses in the input camera view to deviate from the input 2D sequence, resulting in inaccurate 3D motions. Additionally, we experimented with using SDS loss based on an unconditional 2D motion diffusion model (SDS for 3D, without l_{epi}). This variant failed to generate plausible 3D motions and resulted in large joint position errors.

5. Conclusion

We addressed the challenging problem of estimating 3D motion in the world coordinate system from a 2D motion sequence input, without requiring 3D supervision. The key idea is to combine 2D motion diffusion priors and geometric constraints in a novel multi-stage framework to progressively establish consistency for 2D pose sequences across views. We first trained a line-conditioned 2D motion diffusion model to optimize roughly consistent multi-view 2D sequences, which were then used to create a synthetic dataset with strictly consistent multi-view sequences through 3D optimization and reprojection. Using this dataset, we trained a final diffusion model for efficient multi-view generation. We demonstrated our approach’s effectiveness in diverse domains, including human poses, human-object interactions, and animal poses.

Acknowledgement. We thank Zhengfei Kuang for fruitful discussions and Waffle and Mochi for CatPlay data collection. This work is in part supported by the Wu Tsai Human Performance Alliance at Stanford University and ONR MURI N00014-22-1-2740.

References

- [1] Joao Pedro Araujo, Jiaman Li, Karthik Vetrivel, Rishi Agarwal, Deepak Gopinath, Jiajun Wu, Alexander Clegg, and C Karen Liu. CIRCLE: Capture in rich contextual environments. In *CVPR*, 2023. 1
- [2] 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, 6, 7
- [3] Yujun Cai, Lihao Ge, Jun Liu, Jianfei Cai, Tat-Jen Cham, Junsong Yuan, and Nadia Magnenat Thalmann. Exploiting spatial-temporal relationships for 3d pose estimation via graph convolutional networks. In *ICCV*, 2019. 2
- [4] Eric R Chan, Koki Nagano, Matthew A Chan, Alexander W Bergman, Jeong Joon Park, Axel Levy, Miika Aittala, Shalini De Mello, Tero Karras, and Gordon Wetzstein. Generative novel view synthesis with 3d-aware diffusion models. In *ICCV*, 2023. 3
- [5] 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. 2
- [6] Hai Ci, Chunyu Wang, Xiaoxuan Ma, and Yizhou Wang. Optimizing network structure for 3d human pose estimation. In *ICCV*, 2019. 2
- [7] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *CVPR*, 2023. 3
- [8] Shubham Goel, Georgios Pavlakos, Jathushan Rajasegaran, Angjoo Kanazawa, and Jitendra Malik. Humans in 4d: Reconstructing and tracking humans with transformers. In *ICCV*, 2023. 2
- [9] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 4
- [10] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. 2, 3
- [11] Nan Jiang, Zhiyuan Zhang, Hongjie Li, Xiaoxuan Ma, Zan Wang, Yixin Chen, Tengyu Liu, Yixin Zhu, and Siyuan Huang. Scaling up dynamic human-scene interaction modeling. In *CVPR*, 2024. 1
- [12] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 2
- [13] Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. In *CVPR*, 2019. 2
- [14] Roy Kapon, Guy Tevet, Daniel Cohen-Or, and Amit H Bermano. Mas: Multi-view ancestral sampling for 3d motion generation using 2d diffusion. In *CVPR*, 2024. 2, 5, 6, 7
- [15] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. Vibe: Video inference for human body pose and shape estimation. In *CVPR*, 2020. 1, 2
- [16] 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
- [17] 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
- [18] Muhammed Kocabas, Ye Yuan, Pavlo Molchanov, Yunrong Guo, Michael J Black, Otmar Hilliges, Jan Kautz, and Umar Iqbal. Pace: Human and motion estimation from in-the-wild videos. *3DV*, 2024. 2
- [19] 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. 2
- [20] Zhengfei Kuang, Shengqu Cai, Hao He, Yinghao Xu, Hongsheng Li, Leonidas Guibas, and Gordon Wetzstein. Collaborative video diffusion: Consistent multi-video generation with camera control. *arXiv preprint arXiv:2405.17414*, 2024. 3
- [21] Jiaman Li, Yihang Yin, Hang Chu, Yi Zhou, Tingwu Wang, Sanja Fidler, and Hao Li. Learning to generate diverse dance motions with transformer. *arXiv preprint arXiv:2008.08171*, 2020. 5
- [22] 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
- [23] Jiefeng Li, Siyuan Bian, Qi Liu, Jiasheng Tang, Fan Wang, and Cewu Lu. Niki: Neural inverse kinematics with invertible neural networks for 3d human pose and shape estimation. In *CVPR*, 2023. 2
- [24] Jiaman Li, Jiajun Wu, and C Karen Liu. Object motion guided human motion synthesis. *ACM Trans. Graph.*, 42(6), 2023. 1, 4, 6, 8
- [25] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *ICCV*, 2021. 5, 6, 8
- [26] Wenhao Li, Hong Liu, Hao Tang, Pichao Wang, and Luc Van Gool. Mhformer: Multi-hypothesis transformer for 3d human pose estimation. In *CVPR*, 2022. 2
- [27] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *ICCV*, 2023. 3
- [28] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453*, 2023. 3
- [29] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. In *CVPR*, 2024. 3

- [30] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. 2, 6
- [31] Zhengyi Luo, S Alireza Golestaneh, and Kris M Kitani. 3d human motion estimation via motion compression and refinement. In *ACCV*, 2020. 2
- [32] 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. 1, 2, 3, 6
- [33] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *ICCV*, 2017. 2
- [34] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *3DV*, 2017. 2, 3
- [35] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 5
- [36] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *CVPR*, 2019. 2
- [37] Dario Pavullo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *CVPR*, 2019. 2
- [38] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 5
- [39] Wenkang Shan, Zhenhua Liu, Xinfeng Zhang, Shanshe Wang, Siwei Ma, and Wen Gao. P-stmo: Pre-trained spatial temporal many-to-one model for 3d human pose estimation. In *ECCV*, 2022. 2
- [40] Xiaolong Shen, Zongxin Yang, Xiaohan Wang, Jianxin Ma, Chang Zhou, and Yi Yang. Global-to-local modeling for video-based 3d human pose and shape estimation. In *CVPR*, 2023. 2
- [41] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110*, 2023. 3
- [42] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv preprint arXiv:2308.16512*, 2023. 3
- [43] Soyong Shin, Juyong Kim, Eni Halilaj, and Michael J Black. Wham: Reconstructing world-grounded humans with accurate 3d motion. In *CVPR*, 2024. 1, 2, 3, 6
- [44] Leonid Sigal, Alexandru O Balan, and Michael J Black. Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *International journal of computer vision*, 87(1):4–27, 2010. 2
- [45] 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
- [46] Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. GRAB: A dataset of whole-body human grasping of objects. In *ECCV*, 2020. 1
- [47] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior. In *ICCV*, 2023. 3
- [48] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Amit H Bermano, and Daniel Cohen-Or. Human motion diffusion model. In *ICLR*, 2023. 4
- [49] Ayush Tewari, Tianwei Yin, George Cazenavette, Semon Rezhikov, Josh Tenenbaum, Frédo Durand, Bill Freeman, and Vincent Sitzmann. Diffusion with forward models: Solving stochastic inverse problems without direct supervision. *NeurIPS*, 2023. 3
- [50] Jonathan Tseng, Rodrigo Castellon, and C Karen Liu. EDGE: Editable dance generation from music. In *CVPR*, 2023. 4
- [51] Basile Van Hoorick, Rundi Wu, Ege Ozguroglu, Kyle Sargent, Ruoshi Liu, Pavel Tokmakov, Achal Dave, Changxi Zheng, and Carl Vondrick. Generative camera dolly: Extreme monocular dynamic novel view synthesis. *arXiv preprint arXiv:2405.14868*, 2024. 3
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017. 2, 4
- [53] Ziniu Wan, Zhengjia Li, Maoqing Tian, Jianbo Liu, Shuai Yi, and Hongsheng Li. Encoder-decoder with multi-level attention for 3d human shape and pose estimation. In *ICCV*, 2021. 2
- [54] Bastian Wandt, James J Little, and Helge Rhodin. Elepose: Unsupervised 3d human pose estimation by predicting camera elevation and learning normalizing flows on 2d poses. In *CVPR*, 2022. 2, 6
- [55] Jingbo Wang, Sijie Yan, Yuanjun Xiong, and Dahua Lin. Motion guided 3d pose estimation from videos. In *ECCV*, 2020. 2
- [56] 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
- [57] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vitpose: Simple vision transformer baselines for human pose estimation. *NeurIPS*, 2022. 5, 6
- [58] Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli, Gordon Wetzstein, Zexiang Xu, et al. Dmv3d: Denoising multi-view diffusion using 3d large reconstruction model. *arXiv preprint arXiv:2311.09217*, 2023. 3
- [59] Jie Yang, Ailing Zeng, Ruimao Zhang, and Lei Zhang. Unipose: Detecting any keypoints. *arXiv preprint arXiv:2310.08530*, 2023. 6
- [60] Vickie Ye, Georgios Pavlakos, Jitendra Malik, and Angjoo Kanazawa. Decoupling human and camera motion from videos in the wild. In *CVPR*, 2023. 2, 3
- [61] 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, 3

- [62] Hongwen Zhang, Yating Tian, Yuxiang Zhang, Mengcheng Li, Liang An, Zhenan Sun, and Yebin Liu. Pymaf-x: Towards well-aligned full-body model regression from monocular images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(10):12287–12303, 2023. [2](#)
- [63] Jinlu Zhang, Zhigang Tu, Jianyu Yang, Yujin Chen, and Jun-song Yuan. Mixste: Seq2seq mixed spatio-temporal encoder for 3d human pose estimation in video. In *CVPR*, 2022. [2](#)
- [64] Yuxiang Zhang, Hongwen Zhang, Liangxiao Hu, Hongwei Yi, Shengping Zhang, and Yebin Liu. Real-time monocular full-body capture in world space via sequential proxy-to-motion learning. *arXiv preprint arXiv:2307.01200*, 2023. [2](#)
- [65] Ce Zheng, Sijie Zhu, Matias Mendieta, Taojiannan Yang, Chen Chen, and Zhengming Ding. 3d human pose estimation with spatial and temporal transformers. In *ICCV*, 2021. [2](#)
- [66] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang. Motionbert: A unified perspective on learning human motion representations. In *ICCV*, 2023. [1](#), [2](#), [3](#), [6](#)
- [67] Silvia Zuffi, Angjoo Kanazawa, David Jacobs, and Michael J. Black. 3D menagerie: Modeling the 3D shape and pose of animals. In *CVPR*, 2017. [7](#)