

Stereo-Inertial Poser: Towards Metric-Accurate Shape-Aware Motion Capture Using Sparse IMUs and a Single Stereo Camera

Tutian Tang^{1,2*}, Xingyu Ji^{1,2*}, Yutong Li¹, MingHao Liu¹, Wenqiang Xu¹, and Cewu Lu^{1,3}

Abstract—Recent advancements in visual-inertial motion capture systems have demonstrated the potential of combining monocular cameras with sparse inertial measurement units (IMUs) as cost-effective solutions, which effectively mitigate occlusion and drift issues inherent in single-modality systems. However, they are still limited by metric inaccuracies in global translations stemming from monocular depth ambiguity, and shape-agnostic local motion estimations that ignore anthropometric variations. We present Stereo-Inertial Poser, a real-time motion capture system that leverages a single stereo camera and six IMUs to estimate metric-accurate and shape-aware 3D human motion. By replacing the monocular RGB with stereo vision, our system resolves depth ambiguity through calibrated baseline geometry, enabling direct 3D keypoint extraction and body shape parameter estimation. IMU data and visual cues are fused for predicting drift-compensated joint positions and root movements, while a novel shape-aware fusion module dynamically harmonizes anthropometry variations with global translations. Our end-to-end pipeline achieves over 200 FPS without optimization-based post-processing, enabling real-time deployment. Quantitative evaluations across various datasets demonstrate state-of-the-art performance. Qualitative results show our method produces drift-free global translation under a long recording time and reduces foot-skating effects. The code, data, and supplementary materials are available at <https://sites.google.com/view/stereo-inertial-poser>.

I. INTRODUCTION

Human motion capture enables many robotic applications like human-robot interaction, teleoperation, and imitation learning [4], [5]. Current commercial solutions can be divided into optical-based, IMU-based, and hybrid ones according to the modality of sensors. They are precise but often expensive and confined to controlled environments, limiting their accessibility for real-world deployment. Therefore, developing a low-cost and easy-to-use motion capture system has been an active topic in the research community for a long time. One straightforward strategy is to reduce the number of sensors, whether cameras or IMUs. First, it is possible to estimate human motions from monocular videos [6]–[8], but they are vulnerable to occlusion, extreme light conditions, and depth ambiguities. Second, pure inertial

systems [1], [2], [9] suffer from the drifting issue due to sensor noise. Third, to avoid the inherent limitations of single-modality systems, recent hybrid solutions introduce visual-inertial fusion, combining a single RGB camera with six sparse IMUs [3], [10]. By decoupling the full human motion into *global translations* in the world coordinate and *local motions* in the body’s root coordinate, they demonstrate reduced drifting effects and improved robustness against occlusion.

However, fundamental limitations still persist in current visual-inertial fusion systems. First, the predicted translations from pure 2D images are **not metric-accurate**, due to the well-known depth ambiguity problem. Second, the estimated local motions are **not shape-aware**. The anthropometry variations (*i.e.*, the body shape) are ignored, which causes inconsistency between estimated local motions and global translations, as illustrated in Figure 1 (b). To address these challenges, we propose to lift the 2D human pose estimation to 3D perception by replacing the monocular camera with a camera of 3D capability. In this paper, we use the stereo camera, as opposed to RGB-D cameras equipped with active IR or ToF sensors, for two reasons. First, stereo cameras can be built on top of plain RGB cameras, making them potentially more cost-effective and available. Second, we can make full use of those well-established monocular methods to make initial 2D predictions and then lift them to 3D coordinates directly based on the camera model.

In this paper, we propose the Stereo-Inertial Poser, which estimates *metric-accurate* and *shape-aware* human motions from six sparse IMUs and a single stereo camera. We first introduce a stereo pose and body shape estimation module, making initial 3D estimations from the visual modality. From the detected 3D keypoints and IMU measurements, we then use state space models to predict the intermediate 3D human joint positions and body root movements. Finally, a shape-aware fusion module is used to fuse these intermediate results and the body shape parameters together, and the refinement network predicts local human motions and global translations consistent in metric 3D coordinates. Our inference pipeline can achieve very high speed (more than 200 FPS) thanks to the end-to-end design, without the optimization-based post-processing procedures commonly used in previous works. Extensive quantitative and qualitative evaluations demonstrate our method can produce accurate global translation and reduce the foot-skating effect compared with baseline methods.

Our main contributions can be summarized as:

(1) The Stereo-Inertial Poser, a metric-accurate human

¹ School of Computer Science, Shanghai Jiao Tong University

² Meta Robotics Institute, SJTU ³ Shanghai Innovation Institute

* Equal contribution. This work was supported by the Science and Technology Major Project of Jiangsu Province (No. BG2024041), by Shanghai Committee of Science and Technology (No. 24511103200), by the National Key Research and Development Project of China (No. 2022ZD0160102), by Shanghai Artificial Intelligence Laboratory, XPLOER PRIZE grants, and by Shanghai Municipal Education Commission (No. 2024AIYB010). {tttang, ji_xingyu, davidliyutong, lmh209, vinjohn, lucewu}@sjtu.edu.cn




Fig. 1: (a) Our method estimates full human motion, including the local motions in the body’s root coordinate and the global translation, with six inexpensive wearable IMUs and a stereo camera. (b) Previous methods in this track usually ignore the *body shape* (i.e., β parameters of the SMPL model), which may lead to the foot-skating effect or the incorrect global translation estimation. (c) Qualitative comparison with previous methods [1]–[3] shows the proposed method can estimate metric-accurate, drift-free global translation in the world coordinate.

motion capture pipeline with a single stereo camera and six IMUs. It accurately estimates both local motions and full 3D trajectories in the metric coordinate system.

(2) The 3D human pose and shape estimation modules. The pose module lifts pose detections from well-established monocular methods to metric 3D coordinates. The shape module optimizes the initial body shape parameters.

(3) A shape-aware visual-inertial fusion module and related loss functions that enable the correct response to different body shapes without post-processing optimizations.

II. RELATED WORK

The proposed method falls into the category of human motion capture by fusing a single camera and sparse IMUs. It is also closely related to those methods that use either visual or inertial modality.

A. Single Modal Human Motion Capture

1) *Vision-Based Methods*: The gold standard for motion capture has historically been optical systems using reflective markers, which provide high precision but are confined to specialized studios [11]. To improve accessibility, research shifted towards markerless methods, first with multi-camera setups [12], [13] and more recently to systems using only a single monocular camera. The advent of deep learning enabled motion capture from monocular videos. Early methods trained on large-scale 2D image datasets [14], [15] predicted 2D or pseudo-3D keypoints in the image coordinate system [16]–[18]. These methods lacked scale awareness and 3D spatial reasoning in world coordinates. To infer 3D structure, researchers integrated parametric body models like SMPL [19]. Early works optimized SMPL parameters against 2D keypoints using iterative fitting [6], [20] or geometric constraints [21]. Later, direct regression networks [22]–[24] realized real-time inference, but they still remained limited to camera-relative 3D coordinates. While it’s possible to use structure-from-motion techniques to get the absolute 3D predictions, it only works with a moving camera, which is usually head-mounted [25]–[27]. However, there is still

need for finding metric 3D keypoints in world coordinates with a single fixed camera. This need motivated the use of cameras with 3D capabilities. While RGB-D cameras are one option, the domain of 3D pose estimation from point clouds (i.e., RGB-D images) is less mature than its 2D counterpart [28], [29]. Stereo cameras, however, present a powerful alternative. A key advantage of the stereo approach is that we can directly use those well-established 2D pose estimators to predict on each image in the stereo pair. The resulting 2D keypoints can then be triangulated into metric 3D coordinates using the known camera geometry [30]. This design choice strategically leverages the immense progress in the 2D domain to build a robust foundation for metric 3D perception, avoiding the need to develop a new 3D estimator from scratch.

2) *Inertial-Based Methods*: Current commercial solutions use dense IMUs (typically more than 15) to fully determine the orientations of all human parts and precise kinematics. The research community focused on reducing the number of IMUs to make the system less intrusive, more convenient to set up, and more affordable. Early methods started by focusing on the local body movements, ignoring the body’s translation in space. For example, SIP [9] and DIP [31] proposed to recover human motions from 6 IMUs. This 6-IMU setting is followed by later works. TransPose [32] is the first method to predict the body’s global translation in open space, by utilizing the foot-ground contact hypothesis. TIP [33] and DynaIP [34] adopt more advanced network structures to further improve precision. PIP [1], PNP [2], and GlobalPose [35] progressively integrate advanced physics priors into the pipeline, making them work in complex settings and produce realistic motions. Despite these advances, the drifting issue is inherent in inertial systems. Therefore, it’s necessary to fuse visual and inertial modalities for precise metric root translation [3].

B. Human Motion Capture with Visual-Inertial Fusion

Fusing multi-view cameras with IMUs for human motion capture has been studied for a long time [36], [37]. These

methods require cameras to be mounted in a panoptic-like way and well-calibrated. Some works use a head-mounted camera and the underlying SLAM algorithm to capture the 6D pose of the human head [38], [39]. However, fusing a single, fixed, third-person-view camera with sparse IMUs, is a different track. In this track, early works [40]–[42] formulate it as an optimization problem to minimize energy functions of the visual and inertial signals. In the era of deep learning, researchers focused on monocular fusion [10], [43]. A recent work [43] proposes a large-scale dataset and a baseline method to be available to the public. The current state-of-the-art method, RobustCap [3], demonstrates impressive robustness but inherits the fundamental scale ambiguity of the monocular camera. While they can produce smooth and plausible trajectories, the absolute scale of the motion relative to the environment remains unknown. In this paper, we argue and demonstrate that by introducing the stereo camera, our system achieves true metric accuracy and shape-aware motion estimations while maintaining the low-cost, portable setup that makes hybrid systems attractive.

III. METHOD

A. Overview

Figure 2 shows the overview of the proposed pipeline, which aims to recover human motion with six wearable IMUs and one external stereo camera. The setup is in line with prior works [2], [3], [32], where the IMUs are mounted on the root joint (*i.e.*, the pelvis) and five leaf joints, including the head, forearms, and lower legs, while the third-person view camera is fixed. We can get 3D human keypoint detections and body shape estimations from the stereo camera (Section III-B). The shape-aware visual-inertial fusion module fuses the visual detection results along with IMU signals, and finally predicts local motions represented by the SMPL model [19] and global translations (Section III-C).

B. 3D Pose and Body Shape Estimation

1) *3D Metric Pose Estimation*: As mentioned in previous sections, a stereo camera, with known intrinsics and extrinsics, is the practical choice for fully using the well-established monocular pose estimation methods to obtain the metric 3D keypoints in the world coordinate.

At each timestamp t , the two synchronized frames from the left and right cameras, $F_{t,l}$ and $F_{t,r}$ are first processed by the widely-used MediaPipe Pose Landmarker toolbox [44]. Then we get the pose detections, $\mathbf{p}_{2d,l}$, $\mathbf{p}_{2d,r}$, $\mathbf{p}_{3d,l}$, and $\mathbf{p}_{3d,r}$, respectively, where $\mathbf{p}_{2d} \in \mathbb{R}^{2N}$ are 2D keypoints on the image plane and $\mathbf{p}_{3d} \in \mathbb{R}^{3N}$ are canonical 3D predictions relative to the human body's root coordinate. We use the layout of $N = 17$ keypoints, as defined in the widely-used COCO dataset [14]. Each detected keypoint $p_i \in \mathbf{p}$ is associated with a scalar $c_i \in [0, 1]$, representing how confident the model is, or, in other words, the visibility of that point. We denote the confidence in vector form as \mathbf{c}_l and $\mathbf{c}_r \in \mathbb{R}^N$. Then, we can fuse the 3D keypoints estimated from the left and the right camera by the confidence,

$$\mathbf{p}_R = \lambda_c \mathbf{p}_{3d,l} + (1 - \lambda_c) \mathbf{p}_{3d,r}, \quad (1)$$

where $\lambda_c = \mathbf{c}_l \cdot (\mathbf{c}_l + \mathbf{c}_r)^{-1}$. However, the 3D keypoints \mathbf{p}_R are relative to the root coordinate of the human body. To acquire the metric body pose, we need to make full use of the stereo camera. We suppose the stereo camera setup has a known baseline distance d_{base} and shared intrinsics \mathbf{K} , with the world coordinate aligned with the left lens. Then, with paired 2D keypoints $\mathbf{p}_{2d,l}$ and $\mathbf{p}_{2d,r}$, we derive the depth from each keypoint to the camera plane by

$$d_z = f_x \frac{d_{base}}{d_{disp}}, \quad (2)$$

where $f_x \in \mathbf{K}$ is the horizontal focal length, $d_{disp} = |x_r - x_l|$ is the disparity for each point pair $p_l(x_l, y_l) \in \mathbf{p}_{2d,l}$ and $p_r(x_r, y_r) \in \mathbf{p}_{2d,r}$. Then we have the following homogeneous equation

$$\mathbf{p}_C = \lambda_c \mathbf{d}_z \mathbf{K}^{-1} \mathbf{p}_{2d,l} + (1 - \lambda_c) (\mathbf{d}_z \mathbf{K}^{-1} \mathbf{p}_{2d,r} + \mathbf{t}_r), \quad (3)$$

where $\mathbf{t}_r = [d_{base}, 0, 0]^T$ translates from the right lens to the world frame. The resulting 3D keypoints \mathbf{p}_C in the world coordinate are fused by the detections according to the confidence. The corresponding confidence c_C can be derived by averaging \mathbf{c}_l and \mathbf{c}_r . Now we have the 3D keypoints \mathbf{p}_R in the root coordinate for further estimating the local motions, and metric 3D keypoints \mathbf{p}_C in the world coordinate for the global translation.

2) *Body Shape Estimation*: Body shape is important for consistency between the local motions and global translations. Therefore, we need to estimate each subject's body shape by fitting the SMPL model into the subject. We first ask the subject to stand in front of the camera and keep it in T-pose, which aligns with the rest pose of SMPL. We then have the raw point cloud \mathcal{P}_{raw} by stereo matching [45], as well as the 3D skeleton $J_{raw} = \mathbf{p}_C$ following the previous section. The point cloud is then segmented according to the subject's bounding box, and further voxel-downsampled into around 4000 points \mathcal{P}_{down} for computational efficiency. Besides, a 6D transformation, denoted by matrix $\mathbf{M} = [\mathbf{R}|\mathbf{t}]$, should be applied to \mathcal{P}_{down} and J_{raw} in order to align the coordinate between the camera and SMPL. We denote the transformed point cloud as \mathcal{P} , and the transformed skeleton as J . The body shape estimation is formulated as the minimization of a composite energy function:

$$E(\beta, \Phi, \mathbf{M}) = \lambda_{skel} E_{skel} + \lambda_{cd} E_{cd} + \lambda_\Phi E_\Phi + \lambda_\beta E_\beta, \quad (4)$$

where the body shape parameters β and the body pose parameters Φ are from the SMPL model. The skeletal alignment

$$E_{skel} = \sum_{i=1}^N \|J_i - \hat{J}_i\|^2 \quad (5)$$

penalizes the distance between observed 3D skeleton J and fitted joints $\hat{J} = f(\beta, \Phi)$, where $f(\cdot)$ is the SMPL forward kinematics function. The Chamfer distance

$$E_{cd} = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \min_{v \in V} \|p - v\|^2 + \frac{1}{|V|} \sum_{v \in V} \min_{p \in \mathcal{P}} \|v - p\|^2 \quad (6)$$

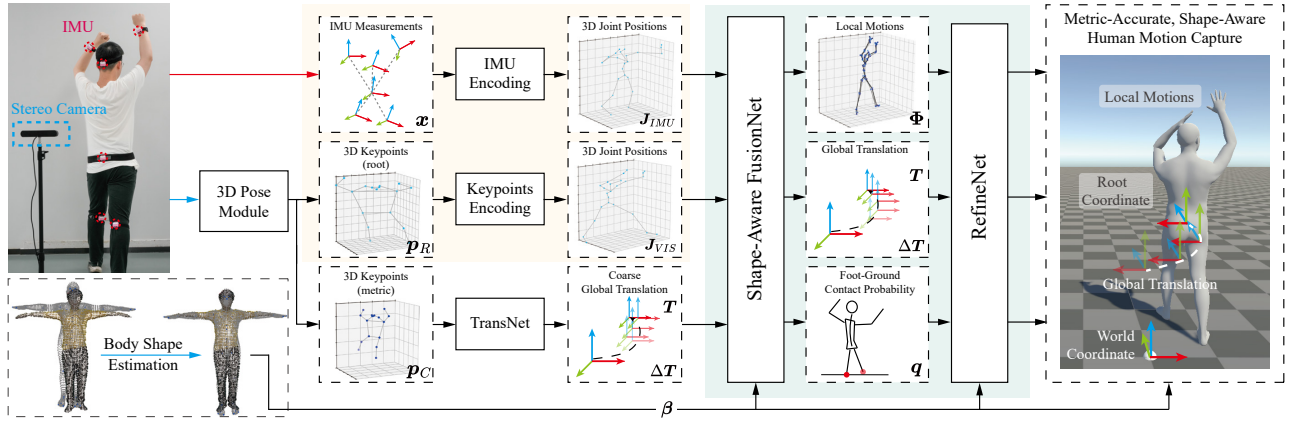


Fig. 2: Our pipeline starts with one stereo camera and six IMUs. From the stereo image pairs, the 3D Pose Module predicts 3D keypoints in the root coordinate \mathbf{p}_R and the world coordinate \mathbf{p}_C . The Body Pose Module estimates the body shape parameters β (Sec. III-B). These initial measurements and predictions are encoded by three separate networks (Sec. III-C) into initial global translations \mathbf{T} , velocities $\Delta\mathbf{T}$, and 3D joint positions \mathbf{J}_{IMU} , \mathbf{J}_{VIS} in the root coordinate. The FusionNet and RefineNet (Sec. III-C) fuse these intermediate results and further refine them into metric-accurate, shape-aware final results, including local motions Φ , global translations \mathbf{T} , and the auxiliary contact probability \mathbf{q} .

describes the distance between the point cloud \mathcal{P} and the SMPL mesh vertices V constructed by Φ and β . The pose prior $E_\Phi = \|\Phi\|^2$ and the shape prior $E_\beta = \|\beta\|^2$ regularize the pose parameters Φ and shape parameters β towards $\mathbf{0}$. We empirically set $\lambda_{skel} = 1$, $\lambda_{cd} = 15$, $\lambda_\Phi = 1$, and $\lambda_\beta = 0.01$. The energy function is solved by an SGD with momentum optimizer in a few seconds. In the supplementary, we visualize the results of four subjects.

C. Shape-Aware Visual-Inertial Fusion

Full human motions in 3D space can be decoupled into two elements: global translations and local motions. In line with prior work [3], the global translation is defined as the translation \mathbf{T}_{Cr} of the root pelvis joint in the world coordinate. The local motion is defined as the SMPL joint pose parameters $\Phi \in \mathbb{R}^{24 \times 3}$. This section describes estimating metric-accurate and shape-aware human motions from IMU measurements and detected 3D keypoints, involving the initial global translation and local motion estimation, and the final shape-aware visual-inertial fusion.

1) *Initial Global Translation Estimation*: We build the TransNet with the State Space Model (SSM) [46] building block with MLP layers to learn the inherent temporal-spatial features in the 3D keypoints and predict the global translation. The global translation is more directly related to the keypoints of the head and the body trunk than the limbs. Therefore, we select a subset of 9 keypoints \mathbf{p}'_C from \mathbf{p}_C as the network inputs, including the nose, eyes, ears, shoulders, and hips. To help the networks learn from the high-frequency signals better, the coordinates are linearly scaled to $[0, 1]$ and further processed by the positional encoding [47] with width $d = 4$. The encoded coordinates and the corresponding confidence \mathbf{c}'_C are concatenated as the network inputs $\mathbf{x} \in \mathbb{R}^n$, where $n = 9 \times (8 \times 3 + 1)$. The TransNet predicts the global translation vector in the metric camera coordinate

$\mathbf{T} \in \mathbb{R}^3$ and also the change of translation between two frames $\Delta\mathbf{T} \in \mathbb{R}^3$. We use L_2 loss to supervise \mathbf{T} ,

$$\mathcal{L}_T = \|\mathbf{T} - \mathbf{T}^{GT}\|_2^2. \quad (7)$$

We apply a spatial-temporal cycle consistency loss to supervise $\Delta\mathbf{T}_t$,

$$\mathcal{L}_{\Delta T_t} = \|\Delta\mathbf{T}_t - (\mathbf{T}_t - \mathbf{T}_{t-1})\|_2^2 + \|\Delta\mathbf{T}_t - \Delta\mathbf{T}_t^{GT}\|_2^2. \quad (8)$$

Therefore, the final loss for the global pose estimation network goes $\mathcal{L}_T = \sum_t (\mathcal{L}_{T_t} + \mathcal{L}_{\Delta T_t})$. Please note that the proposed method does not rely on specific network structures. We use the SSM mainly because it balances accuracy and performance well. It's possible to use other temporal models like plain RNNs [48], LSTMs [49], and Transformers [50].

2) *Initial Local Motion Estimation*: Directly estimating human motions from two different modalities is a difficult problem. Therefore, we use two separate networks, the IMU Encoding Network (IENet) and the Keypoints Encoding Network (KENet), to first unify these two modalities into the canonical root joint space, which is defined by the forward kinematics function from the SMPL model.

a) *IMU Encoding Network*: We use an SSM to predict the 3D joint positions from the IMU measurements. The linear accelerations $\alpha \in \mathbb{R}^{6 \times 3}$ and rotations $\mathbf{R}^{6 \times 3 \times 3}$ are first transformed into the root frame's coordinate, according to the rotation measured by the IMU mounted on the pelvis. The linear accelerations are again processed by positional encoding, and the rotations are converted to the 6D vector representation [51]. These pre-processing techniques can help the network learn better. Then the accelerations $\hat{\alpha} \in \mathbb{R}^{6 \times 3 \times 8}$ and rotations $\hat{\mathbf{R}} \in \mathbb{R}^{6 \times 6}$ are flattened and concatenated together into the input vector $\mathbf{x}_t \in \mathbb{R}^{180}$. The network continuously takes in a sequence of \mathbf{x}_t and predicts the 3D joint positions of the SMPL model $\mathbf{J}_{IMU} \in \mathbb{R}^{24 \times 3}$

in the canonical root coordinate. We use L_2 loss to supervise the network,

$$\mathcal{L}_{J_{IMU}} = \|\mathbf{J}_{IMU} - \mathbf{J}^{GT}\|_2^2. \quad (9)$$

b) *Keypoints Encoding Network*: In parallel to the IENet, another SSM is used to predict 3D joint positions from the 3D keypoints \mathbf{p}_R and confidence \mathbf{c}_C . Given a camera frame, \mathbf{p}_R is first normalized into $[0, 1]$ according to the minimum and maximum values of each axis. The normalized $\hat{\mathbf{p}}_R$ and \mathbf{c}_C are then positional-encoded similarly. The KENet predicts the 3D joint positions \mathbf{J}_{VIS} in the canonical root coordinate. We use L_2 loss to supervise it,

$$\mathcal{L}_{J_{VIS}} = \|\mathbf{J}_{VIS} - \mathbf{J}^{GT}\|_2^2. \quad (10)$$

The KENet is important, since the input 3D keypoints are defined slightly differently from the output 3D joint positions. Also, it's good for adaptability and compatibility. Since different 3D keypoint estimation models may use different definitions of keypoints, such a modular design makes it possible to switch to new state-of-the-art pose estimation methods easily, avoiding re-training the whole pipeline.

3) *Shape-Aware Visual-Inertial Fusion*: Now, we can fuse the intermediate results and raw measurements together using FusionNet to get our desired output. The inputs include the canonical joint positions \mathbf{J}_{IMU} and \mathbf{J}_{VIS} , the confidence \mathbf{c}_C , the body shape parameters β , the initial translations \mathbf{T} and $\Delta\mathbf{T}$, and the IMU measurements \mathbf{x} . We apply positional encoding on \mathbf{J}_{IMU} , \mathbf{J}_{VIS} , \mathbf{T} , $\Delta\mathbf{T}$ and \mathbf{x} . The FusionNet is shape-aware, and it predicts the refined global translation \mathbf{T} and $\Delta\mathbf{T} \in \mathbb{R}^3$, and joint rotations in axis-angle representation $\Phi \in \mathbb{R}^{24 \times 3}$ for local motions. We use the L_2 loss \mathcal{L}_T and the cycle consistency loss $\mathcal{L}_{\Delta T}$ to supervise the refined global translation, as described in Section III-C.1. The loss function for supervising Φ is defined as

$$\begin{aligned} \mathcal{L}_\Phi &= \mathcal{L}_{rot} + \lambda \mathcal{L}_{fk} \\ &= \|\Phi - \Phi^{GT}\|_2^2 + \lambda \|f(\Phi, \beta) - f(\Phi^{GT}, \beta)\|_2^2, \end{aligned} \quad (11)$$

where λ is a scalar for balancing the weights. $f(\cdot)$ is a kinematics function that takes in joint rotations and body shape, and outputs the 3D positions of each joint. We empirically set $\lambda = 2.5$.

Additionally, we design some auxiliary outputs and loss functions to help train the network. To prevent the foot-skating effect, we introduce the foot-ground contact vector $\mathbf{q} = \langle q_l, q_r \rangle \in [0, 1]^2$ to indicate if the left or the right foot stably contacts the ground. The binary cross-entropy loss is applied to supervise it,

$$\mathcal{L}_{fc} = -(\mathbf{q}^{GT} \log \mathbf{q} + ((1 - \mathbf{q}^{GT}) \log (1 - \mathbf{q})). \quad (12)$$

Then, we define the foot-skating loss at given timestamp t ,

$$\mathcal{L}_{fs,t} = \sum_{j \in \{l,r\}} q_j \|f_j(\Phi_t, \beta) - f_j(\Phi_{t-1}, \beta) + \Delta\mathbf{T}_t\|_2^2, \quad (13)$$

where $f_j(\cdot)$ indicates the forward kinematics function that maps joint rotations and body shape into the 3D position of a given joint j in the root coordinate, and j indicates the left

or the right foot. It can be intuitively interpreted as, if the foot contact does not slide with the ground, then its movements in the root coordinate should exactly reverse the body root's movement in the world coordinate. Similarly, we also define a jerk loss on each joint position at time t to stabilize the output motion,

$$\mathcal{L}_{jk,t} = \sum_j \|f_{j,t} - 3f_{j,t-1} + 3f_{j,t-2} - f_{j,t-3}\|_2^2. \quad (14)$$

Finally, we add those loss functions together,

$$\mathcal{L} = \lambda_\Phi \mathcal{L}_\Phi + \lambda_T \mathcal{L}_T + \lambda_{\Delta T} \mathcal{L}_{\Delta T} + \lambda_{fc} \mathcal{L}_{fc} + \lambda_{fs} \mathcal{L}_{fs} + \lambda_{jk} \mathcal{L}_{jk}. \quad (15)$$

We empirically set $\lambda_\Phi = 20$, $\lambda_T = 5$, $\lambda_{\Delta T} = 5$, $\lambda_{fc} = 0.001$, $\lambda_{fs} = 100$, and $\lambda_{jk} = 50$.

4) *Final Refinement Network*: In practice, we find it helpful to add an extra refinement module, RefineNet, to refine the coarse results of the fusion module. The refinement module shares a similar design with previous modules. The input includes the FusionNet predictions, including local motions Φ , global translations \mathbf{T} , velocities $\Delta\mathbf{T}$, and foot-ground contact probability \mathbf{q} , along with the body shape β . The output and loss function design are the same as the fusion module. We conduct ablation studies on the effectiveness of RefineNet in Section IV-C.

IV. EXPERIMENTS

A. Implementation Details

We run the training and evaluation processes on a desktop computer with an i7-14700 CPU, an RTX 4070 Ti Super GPU, and 64 GB RAM. We use the ZED 2 stereo camera, which features 720P resolution, 60 Hz rate, and a 120 degree diagonal field of view. The wearable IMUs are developed in-house within our laboratory, featuring high-performance IMU chips and precise synchronization mechanism. Since our method does not rely on specific stereo cameras or IMUs, users can adopt any other commercially-available cameras or IMUs. For more details, please refer to the supplementary.

1) *Datasets*: In line with prior work [3], we train our method on the AMASS dataset [52] and the AIST++ dataset [53], with synthesized keypoints, IMUs, and the ground truth for foot-ground contact. On AIST++, since it's built with multi-view RGB cameras, we use MediaPipe to generate the detections \mathbf{p}_R in the root coordinate from each camera view and synthesize 3D keypoints \mathbf{p}_C in the world coordinate by adding a random noise $\sigma = 5\text{cm}$ on the ground truth value. We follow the body shape parameters provided by the AMASS dataset. We evaluate our method on the AIST++ dataset and the TotalCapture dataset [54]. On these two datasets, for lack of ground truth shape labels, we use the proposed method (Section III-B.2) to estimate the body shape for each subject, with manually selected frames where the subject acts like the T-pose. Also, we ignore the chamfer distance item by setting $\lambda_{cd} = 0$ for lack of point clouds.

TABLE I: Quantitative comparison with other methods on the AIST++ and TotalCapture dataset.

Method	AIST++						TotalCapture					
	JPE	PVE	SIP	TE	Jerk	FS	JPE	PVE	SIP	TE	Jerk	FS
PIP [1]	87.0	116.5	28.1	45.2	1.04	0.91	49.1	66.0	12.9	43.8	0.20	0.43
PNP [2]	-	-	-	-	-	-	47.4	-	10.8	-	0.26	-
GlobalPose [35]	-	-	-	-	-	-	43.1	49.6	-	-	0.21	-
HybridCap [10]	33.3	-	-	-	-	-	-	-	-	-	-	-
RobustCap [3]	33.1	43.2	9.34	9.91	2.49	0.61	48.7	63.4	13.4	23.5	2.09	1.4
Ours (<i>ideal</i>)	29.8	39.0	8.91	1.13	1.50	0.57	46.5	54.3	10.6	5.42	0.94	0.69
Ours (<i>noise lvl. 5</i>)	30.6	39.8	9.12	4.83	1.95	0.58	47.6	55.8	11.0	8.93	1.39	0.71
Ours (<i>noise lvl. 15</i>)	35.5	46.8	10.1	13.2	2.45	0.63	51.1	60.1	11.4	15.8	1.92	0.79
Ours (<i>syn. stereo</i>)	31.7	43.1	9.89	7.76	1.74	0.60	49.0	57.1	10.9	9.41	1.73	0.72
RobustCap3D (<i>ideal</i>)	33.0	43.2	9.27	6.97	2.34	0.57	48.0	60.8	12.0	14.1	1.35	0.71
RobustCap3D (<i>noise lvl. 5</i>)	35.6	43.4	9.35	7.58	2.74	0.60	49.1	61.6	12.5	16.4	1.46	0.75
RobustCap3D (<i>noise lvl. 15</i>)	38.1	48.9	10.6	16.8	4.00	0.75	52.9	66.5	13.3	21.9	2.01	0.99
RobustCap3D (<i>syn. stereo</i>)	36.0	46.1	11.8	9.91	3.31	0.60	50.8	63.2	12.9	16.9	1.68	0.81

TABLE II: Ablation studies on the AIST++ dataset.

Method	MPJPE	PVE	SIP	TE	Jerk	FS
Ours	29.8	39.0	8.91	1.13	1.50	0.57
- cano. space	30.3	39.7	8.91	1.13	1.58	0.57
- shape-aware	31.1	40.8	8.81	1.26	1.57	0.64
- pos. encoding	29.9	39.1	8.82	1.60	1.65	0.60
- RefineNet	31.5	41.4	8.90	1.36	2.15	0.63
- jerk loss	30.2	39.9	8.80	1.49	4.56	0.79
- cycle loss	29.7	39.0	9.09	1.38	1.53	0.58
- foot skt. loss	30.4	39.8	8.91	2.84	1.65	0.66

2) *Metrics*: We evaluate our method on 6 metrics. Joint Position Error (JPE) measures the mean Euclidean distance error of all joints in the root coordinate in mm. PVE measures the per-vertex error in mm. SIP Error measures the mean global rotation error of hips and shoulders in degrees. Translation Error (TE) measures the absolute translational error between predicted root joint and the ground truth, in cm. Jerk Error measures the averaged jerk (*i.e.*, the first time derivative of acceleration) in km/s^3 . Foot-Skating Error (FS) measures the averaged predicted foot movements in mm in the global space, when ground truth foot movement equals zero (*i.e.*, when q_l or $q_r = 1$). All metrics follow the less is better principle. The first three metrics measure the accuracy of local motions, while the latter three metrics are more related to global translations.

B. Evaluations

Quantitative Results. We compare our method with the SOTA visual-inertial fusion methods HybridCap [10] and RobustCap [3], and SOTA inertial-only methods PIP [1], PNP [2], and GlobalPose [35]. For fair comparison, we enhance and retrain RobustCap by replacing its 2D keypoints inputs with 3D keypoints, denoted by RobustCap3D, which should improve its global translation estimation.

The quantitative results are shown in Table I. Please note that some numbers are missing because they are not reported by the original authors. During the evaluation, we adopt two different ways to synthesize 3D keypoints p_C , including

adding random gaussian noises to the ground-truth 3D keypoints, and building a synthetic, virtual stereo camera. We add random noises with three different levels $\sigma = 0$ (denoted as *ideal*), 5 and 15 cm. Please note that the rated depth accuracy of common RGB-D cameras are usually around several millimeters to centimeters. The synthetic stereo camera, denoted as *syn. stereo*, is built by pairing adjacent RGB cameras from the dataset to simulate stereo geometry for detecting 3D keypoints p_C . Our method shows significant improvements in TE over baseline methods, except under the impractical 15 cm noise, which is way beyond typical depth noise ranges (*i.e.*, less than 10 centimeters). On the AIST++ dataset with aggressive dance motions (*e.g.* jumps, rolls), our method achieves SOTA performance across most metrics, attributed to the shape-aware fusion module that fully exploits metric-accurate 3D keypoints. Notably, RobustCap3D fails to exploit 3D keypoint advantages even with ideal inputs, as evidenced by its suboptimal TE. On TotalCapture, visual-inertial fusion methods underperform recent inertial-only baselines due to limited translational movements and some out-of-view frames with severe occlusion. Nevertheless, our method outperforms all existing fusion approaches. Finally, while recent inertial methods exhibit strong Jerk and FS metrics through physics-based post-optimization, their body shape-agnostic design may limit translational accuracy as illustrated in Figure 1 but related metrics are not reported by the authors.

In terms of inference speed, in offline evaluation, our method runs at more than 200 FPS, thanks to the end-to-end structure without the need of post-processing. In the supplementary video, we also demonstrate online evaluations where our method is running in real time, synchronized with the camera at 60 FPS by the timestamps.

Qualitative Results. We conduct extensive real-world experiments to evaluate our method against the baselines. The evaluations took place in a $6m \times 6m$ motion capture field to assess long-term drift and accuracy of local motions. In a key experiment, a subject walked along a circular path with radius of 2 meters for 10 laps. PIP and PNP exhibit

severe drift, especially during dynamic motions like jumping. RobustCap remains drift-free but produces a geometrically distorted trajectory, a consequence of monocular scale ambiguity. Our method produces both a drift-free and a metrically accurate trajectory, correctly recovering the circular path's scale and shape. For more detailed comparisons and the live demo, please refer to the supplementary video.

C. Ablation Studies

We perform ablation studies on the AIST++ dataset with ideal p_C . Table II shows the results of our ablation studies, where we quantify the contribution of each individual component in our pipeline. We first examine the design of the canonical root space for intermediate representations. Specifically, we do not normalize the 3D keypoints per frame as the KENet inputs, and make the KENet and IENet predict joint positions in the metric coordinate. The dropped performance in local motions could be attributed to different body shapes inherent in keypoints and joint positions which confuse the neural networks. Next, we examine the shape-aware fusion module. By setting the input $\beta = \mathbf{0}$ as the input during the training and the inference process, the fusion module is then unaware of the body shape. The resulting performance degradation in all metrics demonstrates the importance of shape-awareness for motion capture. Then, we remove the positional encoding process and witness a rise in translation error and jerk, indicating the resulted networks fail to learn stably and thus proving the effectiveness of the positional encoding. After that, we test the performance without the RefineNet by directly evaluating the initial outputs of the fusion module. The performance degradations in both local motions and global translations prove the refine module can bridge the gap between the local and the world coordinate. Finally, we check the effectiveness of the loss function design, including the jerk loss, the foot-skating loss and the cycle consistency loss, which suggests we can resolve the *global translation v.s. local motions* (e.g. foot-skating) dilemma illustrated in Figure 1 (b) by introducing correct shape-awareness and proper supervision.

D. Limitations

In our pipeline, despite the IMUs supporting up to 400 Hz sampling rate, our system synchronizes them with the camera at 60 FPS. Future work could decouple the temporal visual-inertial fusion in an asynchronous manner to fully exploit the high-frequency inertial measurements for transient motion estimation. Second, our shape-aware fusion is the initial step towards improving anthropometric modeling in this track through implicit neural networks. Future work could focus on making comprehensive use of the body shape, including bone-length and soft-tissue variations (e.g. body fat or muscle mass). Third, akin to many existing approaches, we utilize IMU acceleration and orientation but under-utilize gyroscope data, which could be further utilized with physics-based constraints. Finally, the current stereo camera setup imposes limits on the workspace area and may be sensitive to poor lighting conditions.

V. CONCLUSION

We presented Stereo-Inertial Poser, a real-time motion capture system that integrates a stereo camera with six IMUs. Our method resolves depth ambiguity through calibrated stereo geometry while compensating for inertial drift via state-space fusion. The novel shape-aware fusion module ensures anthropometric consistency between global translations and local motions. Experiments demonstrate our system achieves significant improvements in global translation accuracy and reduced foot-skating artifacts compared to state-of-the-art baselines. With real-time inference performance, our approach is practical for applications in robotics. Our work bridges the gap between cost-effective hardware and professional-grade motion fidelity, advancing democratized tools for human motion capture and human-robot interaction.

REFERENCES

- [1] X. Yi, Y. Zhou, M. Habermann, S. Shimada, V. Golyanik, C. Theobalt, and F. 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.
- [2] X. Yi, Y. Zhou, and F. Xu, "Physical non-inertial poser (pnp): Modeling non-inertial effects in sparse-inertial human motion capture," in *SIGGRAPH 2024 Conference Papers*, 2024.
- [3] S. Pan, Q. Ma, X. Yi, W. Hu, X. Wang, X. Zhou, J. Li, and F. Xu, "Fusing Monocular Images and Sparse IMU Signals for Real-time Human Motion Capture," Sept. 2023, arXiv:2309.00310 [cs].
- [4] C. Nakatsuka and S. Komorita, "Stable 3d human pose estimation in low-resolution videos with a few views," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 427–433.
- [5] M. Lorenz, G. Bleser, T. Akiyama, T. Niikura, D. Stricker, and B. Taetz, "Towards artefact aware human motion capture using inertial sensors integrated into loose clothing," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 1682–1688.
- [6] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black, "Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image," in *Computer Vision – ECCV 2016*, ser. Lecture Notes in Computer Science. Springer International Publishing, Oct. 2016.
- [7] Y. Rong, T. Shiratori, and H. Joo, "Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration," in *IEEE International Conference on Computer Vision Workshops*, 2021.
- [8] J. Li, C. Xu, Z. Chen, S. Bian, L. Yang, and C. Lu, "Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 3383–3393.
- [9] T. von Marcard, B. Rosenhahn, M. J. Black, and G. Pons-Moll, "Sparse inertial poser: Automatic 3d human pose estimation from sparse imus," *Comput. Graph. Forum*, vol. 36, no. 2, p. 349–360, May 2017.
- [10] H. Liang, Y. He, C. Zhao, M. Li, J. Wang, J. Yu, and L. Xu, "Hybridcap: inertia-aid monocular capture of challenging human motions," 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*, ser. AAAI'23/IAAI'23/EAAI'23. AAAI Press, 2023.
- [11] A. Menache, *Understanding Motion Capture for Computer Animation and Video Games*, 1st ed. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1999.
- [12] N. D. Reddy, L. Guigues, L. Pishchulin, J. Eledath, and S. G. Narasimhan, "TesseTrack: End-to-End Learnable Multi-Person Articulated 3D Pose Tracking," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Nashville, TN, USA: IEEE, June 2021, pp. 15 185–15 195.
- [13] Z. Xu, S. Peng, H. Lin, G. He, J. Sun, Y. Shen, H. Bao, and X. Zhou, "4k4d: Real-time 4d view synthesis at 4k resolution," in *CVPR*, 2024.

- [14] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common Objects in Context," in *Computer Vision – ECCV 2014*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. Cham: Springer International Publishing, 2014, pp. 740–755.
- [15] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele, "2d human pose estimation: New benchmark and state of the art analysis," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014.
- [16] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh, "Convolutional pose machines," in *CVPR*, 2016.
- [17] J. Li, C. Wang, H. Zhu, Y. Mao, H.-S. Fang, and C. Lu, "Crowdpose: Efficient crowded scenes pose estimation and a new benchmark," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 10 863–10 872.
- [18] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, "Realtime multi-person 2d pose estimation using part affinity fields," in *CVPR*, 2017.
- [19] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, "SMPL: A skinned multi-person linear model," *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, vol. 34, no. 6, pp. 248:1–248:16, Oct. 2015.
- [20] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. Osman, D. Tzionas, and M. J. Black, "Expressive body capture: 3D hands, face, and body from a single image," in *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 10975–10985.
- [21] V. Tan, I. Budvytis, and R. Cipolla, "Indirect deep structured learning for 3D human body shape and pose prediction," in *Proceedings of the British Machine Vision Conference 2017*. London, UK: British Machine Vision Association, 2017, p. 15.
- [22] Y. Sun, Q. Bao, W. Liu, Y. Fu, B. Michael J., and T. Mei, "Monocular, One-stage, Regression of Multiple 3D People," in *ICCV*, 2021.
- [23] M. Kocabas, C.-H. P. Huang, O. Hilliges, and M. J. Black, "PARE: Part attention regressor for 3D human body estimation," in *Proceedings International Conference on Computer Vision (ICCV)*. IEEE, Oct. 2021, pp. 11 127–11 137.
- [24] M. Kocabas, N. Athanasiou, and M. J. Black, "Vibe: Video inference for human body pose and shape estimation," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [25] J. Li, Y. Yuan, D. Rempe, H. Zhang, P. Molchanov, C. Lu, J. Kautz, and U. Iqbal, "Coin: Control-inpainting diffusion prior for human and camera motion estimation," in *European Conference on Computer Vision (ECCV)*, 2024.
- [26] M. Kocabas, Y. Yuan, P. Molchanov, Y. Guo, M. Black, O. Hilliges, J. Kautz, and U. Iqbal, "PACE: Human and camera motion estimation from in-the-wild videos," in *International Conference on 3D Vision (3DV)*, March 2024.
- [27] Z. Shen, H. Pi, Y. Xia, Z. Cen, S. Peng, Z. Hu, H. Bao, R. Hu, and X. Zhou, "World-grounded human motion recovery via gravity-view coordinates," in *SIGGRAPH Asia Conference Proceedings*, 2024.
- [28] H. P. H. Shum, E. S. L. Ho, Y. Jiang, and S. Takagi, "Real-time posture reconstruction for microsoft kinect," *IEEE Transactions on Cybernetics*, vol. 43, no. 5, pp. 1357–1369, 2013.
- [29] R. Bashirov, A. Ianina, K. Isakov, Y. Kononenko, V. Strizhkova, V. Lempitsky, and A. Vakhtov, "Real-time rgbd-based extended body pose estimation," in *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2021, pp. 2806–2815.
- [30] H. Akada, J. Wang, V. Golyanik, and C. Theobalt, "3d human pose perception from egocentric stereo videos," in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 767–776.
- [31] Y. Huang, M. Kaufmann, E. Aksan, M. J. Black, O. Hilliges, and G. Pons-Moll, "Deep inertial poser: learning to reconstruct human pose from sparse inertial measurements in real time," *ACM Trans. Graph.*, vol. 37, no. 6, Dec. 2018.
- [32] X. Yi, Y. Zhou, and F. Xu, "Transpose: Real-time 3d human translation and pose estimation with six inertial sensors," *ACM Transactions on Graphics*, vol. 40, no. 4, 08 2021.
- [33] Y. Jiang, Y. Ye, D. Gopinath, J. Won, A. W. Winkler, and C. K. Liu, "Transformer inertial poser: Real-time human motion reconstruction from sparse imus with simultaneous terrain generation," in *SIGGRAPH Asia 2022 Conference Papers*, ser. SA '22. New York, NY, USA: Association for Computing Machinery, 2022.
- [34] Y. Zhang, S. Xia, L. Chu, J. Yang, Q. Wu, and L. Pei, "Dynamic Inertial Poser (DynaIP): Part-Based Motion Dynamics Learning for Enhanced Human Pose Estimation with Sparse Inertial Sensors," in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Seattle, WA, USA: IEEE, June 2024, pp. 1889–1899.
- [35] X. Yi, S. Pan, and F. Xu, "Improving global motion estimation in sparse imu-based motion capture with physics," 2025.
- [36] Z. Zhang, C. Wang, W. Qin, and W. Zeng, "Fusing Wearable IMUs With Multi-View Images for Human Pose Estimation: A Geometric Approach," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Seattle, WA, USA: IEEE, June 2020, pp. 2197–2206.
- [37] T. v. Marcard, G. Pons-Moll, and B. Rosenhahn, "Human pose estimation from video and imus," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 8, pp. 1533–1547, 2016.
- [38] J. Wang, R. Dabral, D. Luvizon, Z. Cao, L. Liu, T. Beeler, and C. Theobalt, "Ego4o: Egocentric Human Motion Capture and Understanding from Multi-Modal Input," Apr. 2025.
- [39] J. Lee and H. Joo, "Mocap Everyone Everywhere: Lightweight Motion Capture with Smartwatches and a Head-Mounted Camera," in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2024, pp. 1091–1100.
- [40] T. von Marcard, R. Henschel, M. J. Black, B. Rosenhahn, and G. Pons-Moll, "Recovering accurate 3d human pose in the wild using imus and a moving camera," in *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [41] Z. Zheng, T. Yu, H. Li, K. Guo, Q. Dai, L. Fang, and Y. Liu, "HybridFusion: Real-Time Performance Capture Using a Single Depth Sensor and Sparse IMUs," in *Computer Vision – ECCV 2018*, V. Ferrari, M. Hebert, C. Sminchiescu, and Y. Weiss, Eds. Cham: Springer International Publishing, 2018, vol. 11213, pp. 389–406, series Title: Lecture Notes in Computer Science.
- [42] T. Helten, M. Müller, H.-P. Seidel, and C. Theobalt, "Real-time body tracking with one depth camera and inertial sensors," in *2013 IEEE International Conference on Computer Vision*, 2013, pp. 1105–1112.
- [43] X. Chen, W. Liu, Q. Bao, X. Liu, Q. Yang, R. Dai, and T. Mei, "Motion capture from inertial and vision sensors," 2024.
- [44] C. Lugaresi, J. Tang, H. Nash, C. McClanahan, E. Uboweja, M. Hays, F. Zhang, C.-L. Chang, M. G. Yong, J. Lee, W.-T. Chang, W. Hua, M. Georg, and M. Grundmann, "Mediapipe: A framework for building perception pipelines," 2019.
- [45] X. Guo, C. Zhang, J. Lu, Y. Wang, Y. Duan, T. Yang, Z. Zhu, and L. Chen, "Openstereo: A comprehensive benchmark for stereo matching and strong baseline," *arXiv preprint arXiv:2312.00343*, 2023.
- [46] T. Dao and A. Gu, "Transformers are ssms: generalized models and efficient algorithms through structured state space duality," in *Proceedings of the 41st International Conference on Machine Learning*, ser. ICML'24. JMLR.org, 2024.
- [47] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: representing scenes as neural radiance fields for view synthesis," *Commun. ACM*, vol. 65, no. 1, p. 99–106, Dec. 2021.
- [48] J. L. Elman, "Finding structure in time," *Cognitive Science*, vol. 14, no. 2, pp. 179–211, 1990.
- [49] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, p. 1735–1780, Nov. 1997.
- [50] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017.
- [51] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, "On the Continuity of Rotation Representations in Neural Networks," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Long Beach, CA, USA: IEEE, June 2019, pp. 5738–5746.
- [52] N. Mahmood, N. Ghorbani, N. F. Troje, G. Pons-Moll, and M. Black, "Amass: Archive of motion capture as surface shapes," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 5441–5450.
- [53] R. Li, S. Yang, D. A. Ross, and A. Kanazawa, "Ai choreographer: Music conditioned 3d dance generation with aist++," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 13 381–13 392.
- [54] M. Trumble, A. Gilbert, C. Malleon, A. Hilton, and J. Collomosse, "Total capture: 3d human pose estimation fusing video and inertial sensors," in *2017 British Machine Vision Conference (BMVC)*, 2017.