

HybridCap: Inertia-Aid Monocular Capture of Challenging Human Motions

Han Liang¹, Yannan He¹, Chengfeng Zhao¹, Mutian Li¹, Jingya Wang^{1,2}, Jingyi Yu^{1,2}, Lan Xu^{1,2*}

¹School of Information Science and Technology, ShanghaiTech University

²Shanghai Frontiers Science Center of Human-centered Artificial Intelligence
{lianghan, heyn, zhaochf, limt1, wangjingya, yujingyi, xulan1}@shanghaitech.edu.cn

Abstract

Monocular 3D motion capture (mocap) is beneficial to many applications. The use of a single camera, however, often fails to handle occlusions of different body parts and hence it is limited to capture relatively simple movements. We present a light-weight, hybrid mocap technique called HybridCap that augments the camera with only 4 Inertial Measurement Units (IMUs) in a learning-and-optimization framework. We first employ a weakly-supervised and hierarchical motion inference module based on cooperative pure residual recurrent blocks that serve as limb, body and root trackers as well as an inverse kinematics solver. Our network effectively narrows the search space of plausible motions via coarse-to-fine pose estimation and manages to tackle challenging movements with high efficiency. We further develop a hybrid optimization scheme that combines inertial feedback and visual cues to improve tracking accuracy. Extensive experiments on various datasets demonstrate HybridCap can robustly handle challenging movements ranging from fitness actions to Latin dance. It also achieves real-time performance up to 60 fps with state-of-the-art accuracy.

Introduction

The past ten years have witnessed a rapid development of human motion capture (Davison, Deutscher, and Reid 2001; Hasler et al. 2009; Stoll et al. 2011; Wang et al. 2017), which benefits broad applications like VR/AR, gaming, sports and movies. However, capturing challenging human motions in a light-weight and convenient manner remains unsolved.

The high-end vision-based solutions require attaching dense optical markers (Vicon 2022) or expensive multi-camera setup (Stoll et al. 2011; Joo et al. 2015; Collet et al. 2015; Joo, Simon, and Sheikh 2018) to capture professional motions, which are undesirable for consumer-level usage. Recent learning-based methods enable robust human capture from monocular RGB video (Kanazawa et al. 2019; Kocabas, Athanasiou, and Black 2020; Zheng et al. 2019; Xiang, Joo, and Sheikh 2019). They require specific human templates for space-time coherent capture (Habermann et al. 2019; Xu et al. 2018b, 2020; Habermann et al. 2020), or utilize parametric human model (Kanazawa et al. 2018;




Figure 1: Our HybridCap achieves robust 3D capture of challenging human motions from a single RGB camera and only 4 IMUs.

Mahmood et al. 2019; Kolotouros et al. 2019; Kocabas, Athanasiou, and Black 2020; Loper et al. 2015). However, these monocular methods are fragile to capture specific challenging motions such as fitness actions or Latin dance, which suffer from complex motion patterns and severe self-occlusion. Recent advances compensate occlusions and challenging motions using probabilistic or attention-based partial occlusion modeling (Kolotouros et al. 2021; Kocabas et al. 2021), or using the generative or weakly-supervised prior (Rempe et al. 2021; He et al. 2021). But the above methods still suffer from the inherent monocular ambiguity due to the lack of reliable observation for self-occluded regions.

In stark contrast, combining occlusion-unaware body-worn sensors for robust motion capture has been widely explored (Henschel, Von Marcard, and Rosenhahn 2020; Pons-Moll et al. 2011; Gilbert et al. 2019; Zhang et al. 2020; Kaufmann et al. 2021). Utilizing Inertial Measurement Units (IMUs) for motion inertia recording is a very popular trend. However, most previous works (Pons-Moll et al. 2010, 2011; Von Marcard, Pons-Moll, and Rosenhahn 2016; Malleison et al. 2017; von Marcard et al. 2018; Gilbert et al. 2019; Malleison, Collomosse, and Hilton 2019; Zhang et al. 2020) utilize multi-view video or a relatively large amount of IMUs (from 8 to 17), which is undesirable for daily usage.

*Corresponding author

Recently, learning-based approaches (Huang et al. 2018; Yi, Zhou, and Xu 2021) enable a real-time motion capture with only 6 IMUs. But the lack of visual cue leads to inherent drifts and overlay artifacts.

In this paper, we tackle the above challenges and present *HybridCap* – a high-quality inertia-aid monocular approach for capturing challenging human motions, as shown in Fig. 1. We revisit the light-weight hybrid setting using a single RGB camera and sparse IMUs (only 4 in our setting), and reframe such hybrid motion capture into the neural data-driven realm. It enables practical and robust human motion capture under challenging scenarios at 60 fps with state-of-the-art accuracy.

To this end, we introduce a learning-and-optimization framework, consisting of a hybrid motion inference module and a robust hybrid optimization scheme. The former is formulated in a weakly supervised and hierarchical multi-stage framework. Specifically, we adopt a fully differentiable architecture in an analysis-by-synthesis fashion without explicit 3D ground truth annotation. Thus, for dense hybrid weak supervision, we propose a new largest multi-modal human motion dataset called Hybrid Challenging Motions (HCM) consisting of 176-minute challenging motions recorded from 14 RGB cameras and 17 IMUs with a total of 8.9M images.

Our network takes the sequential 2D pose estimation and motion inertia from our light-weight hybrid setting as input, and learns to infer 3D human motion by comparing predicted tracking results against the dense hybrid observations. To handle the challenging motions, we further formulate our network into a series of cooperative subtasks, including a limb tracker, a body tracker, a root tracker and a hybrid IK (inverse kinematics) solver. For all these four subtasks, we also propose a novel bone tracking block design consisting of light-weight pure sequential recurrent layers with skip connection for temporal motion state modeling, which is trained to consider bone length while tracking or solving. Such hierarchical multi-stage design makes full use of our hybrid weak supervision by progressively narrowing the search space of plausible motions in a coarse-to-fine manner, so as to enable efficient and effective inference of challenging motions.

Finally, besides the data-driven motion characteristics from previous module, current multi-modal input also encodes reliable visual and inertial hints, especially for the non-occluded regions. Thus, we further propose a robust hybrid optimization scheme with inferential prior to refine the inferred results. It directly fits 2D and inertial observations to improve the accuracy while keeping multi-modal inferential characteristics to preserve motion plausibility.

To summarize, our main contributions include:

- A real-time and accurate motion capture approach augmenting the light-weight monocular setting using only 4 IMUs aiding, achieving significant superiority to state-of-the-arts.
- A novel network with a hierarchical framework containing cooperative bone tracking blocks with bone length awareness, and a robust optimization scheme combining

inferential prior with input observations, improving capture results effectively.

- A new largest multi-modal human motion dataset containing a wide range of challenging motions along with abundant records of RGB cameras and IMUs.

Related Work

Based on our inertia-aid monocular setup, we focus on human motion capture from optical and inertial solutions.

Optical Motion Capture. Marker-based methods (Vicon 2022; Vlasic et al. 2007) have achieved success in capturing professional human motions which are widely utilized in industry, but they are inapplicable for daily usage due to expensive and tedious setup. The exploration of markerless mocap (Bregler and Malik 1998; De Aguiar et al. 2008; Theobalt et al. 2010) has made great progress in order to get rid of body-worn markers and pursue efficiency. Benefiting from researches on parametric human models (Anguelov et al. 2005; Loper et al. 2015; Pavlakos et al. 2019; ?) in the last decade, various data-driven approaches are proposed to estimate 3D human pose and shape by optimizing (Huang et al. 2017; Lassner et al. 2017; Bogo et al. 2016; Kolotouros, Pavlakos, and Daniilidis 2019) or directly regressing (Kanazawa et al. 2018, 2019; Kocabas, Athanasiou, and Black 2020; Zanfir et al. 2021) human model parameters. Taking specific template mesh as prior, multi-view (Gall et al. 2010; Stoll et al. 2011; Liu et al. 2013; Robertini et al. 2016; Pavlakos et al. 2017; Simon et al. 2017; Xu et al. 2018a) and monocular (Xu et al. 2018b; Habermann et al. 2019; Xu et al. 2020; Habermann et al. 2020) template-based approaches combine free-form and parametric methods, which produce high quality skeletal and surface motions. Besides, to alleviate the inherent estimation ambiguity from 2D input to 3D motion, recent approaches (Kolotouros et al. 2021; Kocabas et al. 2021; Liang et al. 2023) handle complex patterns using probabilistic or attention-based semantic modeling, (Rempe et al. 2021; He et al. 2021) learn to model generative or weakly-supervised prior to solve unseen and non-periodic motions. However, these methods still suffer from challenging motions, especially for rare pose patterns and extreme self-occlusion.

Inertial Motion Capture. To overcome the limitations of vision cues only, another category of works propose to use IMUs. Previously, purely inertial methods using large amounts of sensors like Xsens MVN (Movella 2022) has been commercially used. However, intrusive capture system prompts researchers forward to sparse-sensor setup. SIP (Von Marcard et al. 2017), which uses only 6 IMUs, presents a pioneering exploration. However, limitations of its traditional optimization framework make real-time application impractical. Recent data-driven works (Huang et al. 2018; Yi, Zhou, and Xu 2021; Yi et al. 2022) achieve great improvements on accuracy and efficiency with sparse sensors, but substantial drift is still unsolved for challenging motions. Preceding sensor-aid solutions propose to combine IMUs with videos (Gilbert et al. 2019; Henschel, Von Marcard, and Rosenhahn 2020; Malleson, Collomosse, and Hilton 2019; Malleson et al. 2017), RGB-D cameras (Hel-

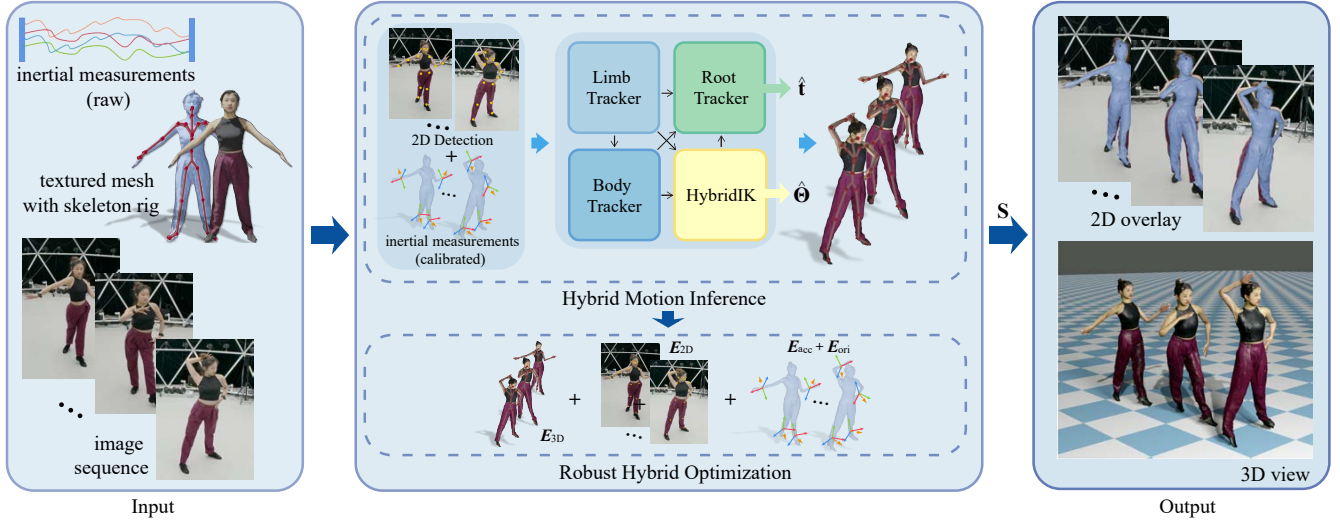


Figure 2: The pipeline of HybridCap with multi-modal input. Our approach combines a hybrid motion inference stage with a robust hybrid motion optimization stage to capture 3D challenging motions.

ten et al. 2013; Zheng et al. 2018), optical markers (Andrews et al. 2016) or even LiDAR (Ren et al. 2023). Although these approaches partially solve scene-occlusion problem and correct drift effectively, they are restricted from either undesirable system complexity or implausible estimation for challenging motions.

Preliminary and Overview

The goal of our work is to capture challenging 3D human motions using a single camera with few inertial sensors aiding, which suffers from complex motion patterns and extreme self-occlusion. Fig. 2 provides an overview of HybridCap, which relies on a template mesh of the actor and makes full usage of multi-modal input in a learning-and-optimization framework. In the inference stage, our hierarchical design extracts different characteristics from multi-modal observations and learns to estimate plausible motions in a weakly supervised manner. Then, a robust optimization stage is introduced to refine the skeletal motions to increase the tracking accuracy and overlay performance.

Template and Motion Representation. We first scan the actor with a 3D body scanner to generate the textured template mesh of the actor. Then, we rig it automatically by fitting the Skinned Multi-Person Linear Model (SMPL) (Loper et al. 2015) to the template mesh and transferring the SMPL skinning weights to our scanned mesh. The kinematic skeleton is parameterized as $\mathbf{S} = [\theta, \mathbf{R}, \mathbf{t}]$, including the Euler angles $\theta \in \mathbb{R}^{N_J \times 3}$ of the N_J joints, the global rotation $\mathbf{R} \in \mathbb{R}^3$ and translation $\mathbf{t} \in \mathbb{R}^3$. Furthermore, let Θ denotes the 6D representation (Zhou et al. 2019) of the global rotation and joint rotations. Then, we can formulate $\mathbf{S} = \mathcal{M}(\Theta, \mathbf{t})$ where \mathcal{M} denotes the motion transformation between various representations.

Input Preprocessing. Our system takes the RGB video, inertial measurements and a well pre-scanned template of the actor as the overall input. Given an image frame, we extract

N_M 2D keypoints $\mathbf{p} \in \mathbb{R}^{N_M \times 2}$ and corresponding confidence $\sigma \in \mathbb{R}^{N_M}$ using OpenPose (Cao et al. 2017). To generalize to various camera settings during inference time, we refer (Shimada et al. 2021) to use canonicalized 2D keypoints \mathbf{p}_c by projecting them onto $Z = 1$ plane. Then we transform inertial measurements from inertial frame \mathcal{F}_I into camera frame \mathcal{F}_C and obtain IMU accelerations $\mathbf{A}_n \in \mathbb{R}^3$ and orientations $\mathbf{R}_{\mathbf{b}_n} \in \mathbb{R}^{3 \times 3}$ of N_i corresponding bones \mathbf{b}_n with calibrated \mathbf{R}_{I2C} and $\mathbf{R}_{S2B,n}$:

$$\mathbf{A}_n = \mathbf{R}_{I2C} \mathbf{A}_{I,n} \quad (1)$$

$$\mathbf{R}_{\mathbf{b}_n} = \mathbf{R}_{I2C} \tilde{\mathbf{R}}_n \mathbf{R}_{S2B,n}^T, \quad (2)$$

where \mathbf{R}_{I2C} is the transformation from \mathcal{F}_I to \mathcal{F}_C , and $\mathbf{R}_{S2B,n}$ is the transformation from the n -th IMU sensor \mathcal{F}_{s_n} to $\mathcal{F}_{\mathbf{b}_n}$ of its corresponding bone \mathbf{b}_n . Besides, to provide prior knowledge on anthropometry, we heuristically calculate $N_b = 7$ key bone lengths \mathbf{L}_k (uparm, lowarm, upleg, lowleg, foot, clavicle, and spine) from the rigging skeleton \mathbf{S} and concatenate them into the input. Thus the overall input of the network in a single frame is $[\mathbf{p}_c, \sigma, \mathbf{R}_b, \mathbf{A}, \mathbf{L}_k] \in \mathbb{R}^{2N_M + N_M + 9N_i + 3N_i + N_b}$.

Approach

Hybrid Motion Inference

As illustrated in Fig. 3, we adopt a hierarchical multi-stage motion inference scheme. It takes the sequential 2D pose estimation and motion inertia as input and predicts 3D human motions in a weakly supervised manner without explicit 3D ground truth annotation. To tackle challenging motions, we divide the motion inference task into a series of cooperative subtasks. Specifically, we introduce a limb tracker, a body tracker, a root tracker, and a hybrid IK (inverse kinematics) solver, so as to model the hierarchical knowledge of articulated human body structure.

Pure Residual Recurrent Blocks. For temporal motion state modeling in these subtasks, we propose a novel effi-

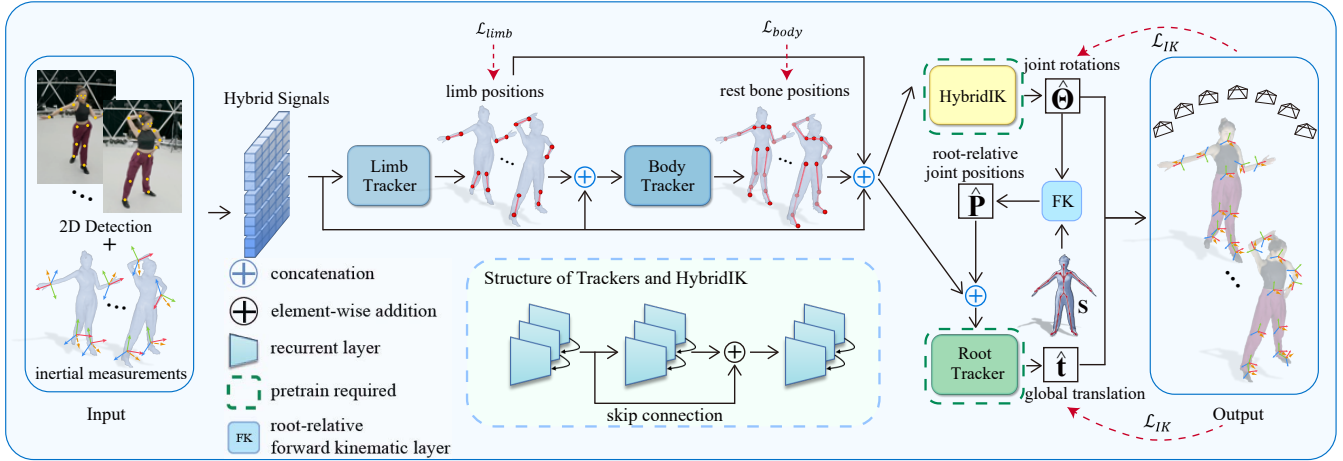


Figure 3: Illustration of our hybrid motion inference module, which is based on cooperative pure residual recurrent blocks that serve as limb, body and root trackers as well as a hybrid inverse kinematics (IK) solver. The limb tracker focuses on accurately tracking for 4 limbs while the body tracker estimates the rest bone positions from accurate limb positions. Then, the hybrid IK solver and root tracker are employed to combine initial input with well-estimated bone positions to solve rotations $\hat{\Theta}$ and global translation $\hat{\mathbf{t}}$.

cient block design using pure recurrent layers with skip connection. In contrast to block design in TransPose (Yi, Zhou, and Xu 2021) adopting both recurrent (LSTM) and fully-connected (FC) layers, we find that such FC layers are the key bottleneck, which encode non-temporal representation and tend to overfit. Thus we simply remove the FC layers and improve the accuracy.

We further introduce skip connection. Our key intuition is that motion inference or position tracking requires less high-level features. We adopt skip connection enabling easier identity function learning to retain more information from the input layer, where the layer close to the output is responsible for adding low-level details. Thus the low-level 2D and inertial features could be selected and passed to the output layer directly, which helps more accurate and detailed motion inference.

Hierarchical Bone Trackers In our hierarchical design, our trackers track bones rather than joints. For a bone \mathbf{b}_n , we track root-relative positions of its two endpoint $\hat{\mathbf{J}}$ with distance constraint using the known bone lengths $\mathbf{L}_{\mathbf{b}_n}$. Specifically, the limb tracker focuses on accurately tracking four limbs with reliable inertial measurements. Here, the loss of limb tracker is formulated as:

$$\mathcal{L}_{limb} = \mathcal{L}_{joint}^{(limb)} + \mathcal{L}_{bone}^{(limb)}, \quad (3)$$

where $\mathcal{L}_{joint}^{(limb)}$ is the 3D joint position loss and $\mathcal{L}_{bone}^{(limb)}$ is the limb bone length loss. Next, the body tracker estimates the rest rigid body parts concatenating the initial input and well-estimated limb positions as the input. Bone length constraints are also utilized to reduce depth ambiguity. Furthermore, the bone orientations help to narrow the search space to the target joint positions. Similar to the limb tracker loss, the body tracker loss is formulated as:

$$\mathcal{L}_{body} = \mathcal{L}_{joint}^{(body)} + \mathcal{L}_{bone}^{(body)}. \quad (4)$$

Note that \mathcal{L}_{joint} is formulated as the reprojection error to all camera views, which guides the corresponding tracker to predict joint positions in a weakly supervised manner:

$$\mathcal{L}_{joint} = \sum_{t=1}^T \sum_{c=1}^{N_C} \sum_{j=1}^{N_J} \sigma_{c,j}^{(t)} \|\Pi_c(\hat{\mathbf{J}}_j^{(t)} + \mathbf{t}) - \mathbf{p}_{c,j}^{(t)}\|_2^2, \quad (5)$$

where $\sigma_{c,j}^{(t)}$ denotes the confidence of 2D joint $\mathbf{p}_{c,j}^{(t)}$; Π_c denotes the projection function of camera c ; \mathbf{t} denotes global translation from hybrid optimization of full observations; N_J denotes the number of bone endpoints (8 for limb tracker and 7 for body tracker). Then, the bone length loss is formulated as the L_2 loss between the predicted bone length and ground-truth bone length $\mathbf{L}_{\mathbf{b}_n}$:

$$\mathcal{L}_{bone} = \sum_{t=1}^T \sum_{n=1}^{N_B} (\|\hat{\mathbf{J}}_{\mathbf{b}_n,0}^{(t)} - \hat{\mathbf{J}}_{\mathbf{b}_n,1}^{(t)}\|_2 - \mathbf{L}_{\mathbf{b}_n})^2, \quad (6)$$

where the predicted bone length can be calculated by the distance of two output endpoint positions $\hat{\mathbf{J}}_{\mathbf{b}_n,i}^{(t)}$ ($i = 0, 1$) of bone \mathbf{b}_n . Note that N_B is the number of target bones which is 4 for the limb tracker and 8 for the body tracker.

Hybrid Inverse Kinematics Solver. Based on the accurate tracking of bones, we introduce our hybrid IK solver and root tracker to solve rotations and translation respectively. The initial input and well-estimated root-relative 3D joints are concatenated and fed into our hybrid IK solver, which outputs global rotation and local joint rotations $\hat{\Theta}$ in the 6D representation. Then we perform forward kinematics using predicted rotations $\hat{\Theta}$ and bone lengths to obtain refined root-relative 3D joints. Next we send them with the initial input into the root tracker which predicts root position $\hat{\mathbf{t}}$ (i.e. global translation) in the camera frame.

To utilize our dense hybrid weak supervision, we further calculate N_M 3D marker positions $\hat{\mathbf{P}}_m(\hat{\Theta}, \hat{\mathbf{t}})$ attached to the




Figure 4: Our qualitative results. Each picture shows the input image, 2D overlay and motion in 3D space from left to right respectively. The results show that our approach produces good 2D overlay and plausible 3D motion.

skeleton (corresponding to body joints and face landmarks of OpenPose), N_I bone orientations $\hat{\mathbf{R}}_{\mathbf{b}_n}(\hat{\Theta})$, and N_I simulated IMU sensor positions $\hat{\mathbf{P}}_n(\hat{\Theta}, \hat{\mathbf{t}})$ respectively. The IK loss is formulated as:

$$\mathcal{L}_{IK} = \lambda_{2D}\mathcal{L}_{2D} + \lambda_{acc}\mathcal{L}_{acc} + \lambda_{ori}\mathcal{L}_{ori} + \lambda_{prior}\mathcal{L}_{prior} + \lambda_{trans}\mathcal{L}_{trans}. \quad (7)$$

The 2D reprojection loss \mathcal{L}_{2D} ensures each estimated 3D marker $\hat{\mathbf{P}}_m$ projects onto the corresponding 2D keypoint $\mathbf{p}_{c,m}$ in all camera views, formulated as:

$$\mathcal{L}_{2D} = \sum_{t=1}^T \sum_{c=1}^{N_C} \sum_{m=1}^{N_M} \sigma_{c,m}^{(t)} \|\Pi_c(\hat{\mathbf{P}}_m^{(t)}(\hat{\Theta}, \hat{\mathbf{t}})) - \mathbf{p}_{c,m}^{(t)}\|_2^2, \quad (8)$$

where $\sigma_{c,m}^{(t)}$ denotes the confidence of 2D keypoint $\mathbf{p}_{c,m}$ and Π_c is the projection function of camera c . Then, we introduce the acceleration loss \mathcal{L}_{acc} to encourage the network to learn the implicit physical constraints and generate plausible motions:

$$\mathcal{L}_{acc} = \sum_{t=2}^{T-1} \sum_{n=1}^{N_I} \|\hat{\mathbf{A}}_n^{(t)}(\hat{\Theta}, \hat{\mathbf{t}}) - \mathbf{A}_n^{(t)}\|_2^2, \quad (9)$$

where $\hat{\mathbf{A}}_n^{(t)}$ is the estimated acceleration calculated from predicted IMU position $\hat{\mathbf{P}}_n$, formulated as below where st is the sampling time:

$$\hat{\mathbf{A}}_n^{(t)} = (\hat{\mathbf{P}}_n^{(t+1)} - 2\hat{\mathbf{P}}_n^{(t)} + \hat{\mathbf{P}}_n^{(t-1)})/st^2. \quad (10)$$

We further propose the orientation loss \mathcal{L}_{ori} , which ensures the predicted orientation $\hat{\mathbf{R}}_{\mathbf{b}_n}$ of each bone bound with IMU sensor fits the observation $\mathbf{R}_{\mathbf{b}_n}$ of the corresponding IMU measurement.

$$\mathcal{L}_{ori} = \sum_{t=1}^T \sum_{n=1}^{N_I} \|\hat{\mathbf{R}}_{\mathbf{b}_n}^{(t)}(\hat{\Theta}) - \mathbf{R}_{\mathbf{b}_n}^{(t)}\|_2^2. \quad (11)$$

Furthermore, the prior loss \mathcal{L}_{prior} is the L_2 loss between predicted rotations $\hat{\Theta}$ and reference Θ , while \mathcal{L}_{trans} the one related to the predicted root position $\hat{\mathbf{t}}$ and reference position \mathbf{t} . Both Θ and \mathbf{t} are dense ‘‘pseudo ground-truth’’ references.

Hybrid Motion Optimization

Despite the data-driven motion inference stage learns the mapping from multi-modal observations to 3D motions, the generalization error is non-negligible due to noisy 2D detection and inertial measurements.

We further introduce a hybrid motion optimization stage to refine the skeletal motions to increase the tracking accuracy and overlay performance. It jointly utilizes the learned 3D prior from the network of multi-modal weak supervision, the 2D keypoints in the visible regions as well as inertial measurements.

In this stage, we first initialize the skeletal motion sequence \mathbf{S} using network output by representation transformation $\mathcal{M}(\hat{\Theta}, \hat{\mathbf{t}})$ and then perform the optimization procedure. We adopt the Euler angle representation so that the joint angles θ of \mathbf{S} locate in the pre-defined range $[\theta_{min}, \theta_{max}]$ of physically plausible joint angles to prevent unnatural poses. Our energy function is formulated as:

$$\mathbf{E}(\mathbf{S}) = \lambda_{3D}\mathbf{E}_{3D} + \lambda_{2D}\mathbf{E}_{2D} + \lambda_{acc}\mathbf{E}_{acc} + \lambda_{ori}\mathbf{E}_{ori}. \quad (12)$$

Here, \mathbf{E}_{3D} enforces the final motion sequence close to the predicted one; \mathbf{E}_{2D} ensures that each final 3D marker reprojects onto the corresponding 2D keypoint.

Besides, we adopt the acceleration energy \mathbf{E}_{acc} to enforce the final motion to be temporally consistent with the network estimating accelerations $\hat{\mathbf{A}}_n$ supervised by N_I IMUs and the measured ground-truth accelerations \mathbf{A}_n from N_i input IMUs. Specifically, the acceleration term \mathbf{E}_{acc} is formulated as:

$$\mathbf{E}_{acc} = \sum_{t=2}^{T-1} \sum_{n=1}^{N_i} \gamma_n^{(t)} \|\mathbf{A}_n^{(t)}(\mathbf{S}) - \mathbf{A}_n^{(t)}\|_2^2 + \sum_{t=2}^{T-1} \sum_{n=N_i+1}^{N_I} \|\mathbf{A}_n^{(t)}(\mathbf{S}) - \hat{\mathbf{A}}_n^{(t)}\|_2^2, \quad (13)$$

where the first term means that we directly use the acceleration observation from the input IMUs and the second term

Method	HCM			3DPW			AIST++		
	MPJPE↓	PCK@0.2↑	AE↓	MPJPE ↓	PCK@0.2↑	AE ↓	MPJPE ↓	PCK@0.2↑	Accel
VIBE	103.4	61.5	96.8	82.9	70.3	21.1	73.5	69.1	94.6
HuMoR	81.9	72.7	31.8	77.0	75.8	13.2	59.4	81.3	33.1
TransPose	73.1	77.4	65.6	76.5	76.9	19.3	61.6	78.2	57.0
ChallenCap	69.5	79.5	98.1	78.2	74.2	48.4	53.2	86.8	91.4
Ours	43.3	90.1	17.9	72.1	80.5	5.4	33.3	95.4	14.4

Table 1: Quantitative comparison of several previous state-of-the-art methods in terms of tracking accuracy and plausibility.




Figure 5: Our method outperforms the previous state-of-the-arts on the performance of both overlay and limb orientation.

implies that we trust the network for those unobserved parts. The final acceleration $\mathbf{A}_n^{(t)}(\mathbf{S})$ is obtained from skeletal motion sequence \mathbf{S} as same as Eqn. 10. Then, we adopt the orientation energy \mathbf{E}_{ori} to enforce the final bone orientations to be consistent with the observations of N_i input IMUs, which is formulated as:

$$\mathbf{E}_{\text{ori}} = \sum_{t=1}^T \sum_{n=1}^{N_i} \|\mathbf{R}_{\mathbf{b}_n}^{(t)}(\mathbf{S}_t) - \mathbf{R}_{\mathbf{b}_n}^{(t)}\|_2^2, \quad (14)$$

where $\mathbf{R}_{\mathbf{b}_n}^{(t)}(\mathbf{S}_t)$ denotes the final orientation of bone \mathbf{b}_n , which is obtained from skeleton \mathbf{S}_t .

Optimization. The constrained optimization problem to minimize the Eqn. 12 is solved using the Levenberg-Marquardt (LM) algorithm of ceres (Agarwal, Mierle, and The Ceres Solver Team 2022) with 4 iterations. In our experiments, we empirically set the parameters $\lambda_{3D} = 10$, $\lambda_{2D} = 1$, $\lambda_{\text{acc}} = 10$ and $\lambda_{\text{ori}} = 30$.

Experiments

We run our pipeline on a PC with an i7-10700k CPU and RTX3070 GPU, where the inference module and optimization module take $2.8(\pm 0.4)$ ms and $6.2(\pm 3.5)$ ms respectively, achieving 60 fps. In this section, first, we describe the datasets used for training and evaluation. Next, we further qualitatively (Fig. 5) and quantitatively (Tab. 1) illustrate that our method outperforms previous state-of-the-arts. We also provide extensive ablation studies (Tab. 2, Tab. 3 and Fig. 6, Fig. 7) to evaluate our technical design and input setting. Finally, we present more qualitative results in Fig. 4.

Training Dataset. We use mixed multi-modal datasets to train our inference module. To provide sufficient 2D and inertial observations, we build and propose a new dataset Hybrid Challenging Motions (HCM), which consists of abundant records of RGB cameras and IMUs. Please refer to the

supplementary to appreciate more details. We further utilize AIST++ (Li et al. 2021; Tsuchida et al. 2019) which consists of various challenging dance sequences. Note that when using AIST++ we simulate N_i IMU measurements only as the input rather than supervision.

Evaluation Dataset. For thorough evaluation, we use 3DPW (von Marcard et al. 2018), AIST++, and our HCM, which contain complex in-the-wild scenarios or challenging motions. We follow the standard protocols to split out test sets for all the datasets. Note that we use 3DPW that only provides monocular data for testing rather than training. Our code and HCM dataset will be made publicly available to stimulate future research.

Metrics. To evaluate capture accuracy, we report the Mean Per Joint Position Error (MPJPE) in terms of *mm*, and the Percentage of Correct Keypoints (PCK) with 0.2 torso size as the threshold. For HCM and 3DPW with ground-truth acceleration measured by IMUs, we report acceleration error (AE) (*m/s²*), calculated as the difference between the ground-truth and estimated acceleration. For AIST++, we report mean per-joint accelerations (Accel) (Kanazawa et al. 2019; Rempe et al. 2021) to evaluate motion plausibility.

Comparison

We compare HybridCap with various representative methods using a single RGB camera or sparse IMUs. Specifically, we apply ChallenCap (He et al. 2021) and HuMoR (Rempe et al. 2021) which are also based on a learning-and-optimization framework, VIBE (Kocabas, Athanasiou, and Black 2020) using adversarial learning, as well as TransPose (Yi, Zhou, and Xu 2021) using pure 6 IMUs. For fair comparisons, we fine-tune them with the same training dataset. As shown in Fig. 5, our method gets better overlays of the captured body and obtains more accurate limb orientation results. Tab. 1 provides the corresponding quantitative comparison results under various metrics and datasets.

Ablation Study

Evaluation on inference module. We first evaluate our inference module by comparing to the variants of our approach and previous state-of-the-arts. The quantitative results are provided in Rows 1-5 in Tab. 2. **VIBE structure** (Row 1) shows VIBE variant with 4 IMUs and RGB input, where we concatenate inertial input and key bone lengths with image features obtained from pre-trained feature extractor used in the vanilla VIBE. **TransPose structure** (Row 2) shows TransPose variant with 4 IMUs and RGB input, where we concatenate 2D keypoints and key bone lengths with inertial




Figure 6: Qualitative evaluation on our inference module and optimization module.

Ablation	MPJPE↓	PCK↑	AE ↓
1) VIBE net	75.4	76.5	34.2
2) TransPose net	61.0	82.7	19.2
3) w/o inference	136.5	51.1	82.1
4) w/o skip connection	54.5	86.9	18.7
5) w/o bone length	65.8	81.3	19.6
6) w/o optimization	57.5	85.3	18.6
7) w/o inertial terms	51.2	87.4	62.2
8) Ours (complete)	43.3	90.1	17.9

Table 2: Quantitative evaluation on our network structure, bone length utilization and optimization configurations.

input. We train both two networks using the same training set and losses as ours. The results show our network with nuanced designs to progressively fuse multi-modal input is superior to both two previous state-of-the-arts. Row 3-5 corresponds to the variants without our entire inference module, without using the skip connection in our block design, without using bone length, respectively. **Purely recurrent design** (Row 4) without FC layers (adopted in Row 2), effectively alleviates overfitting. **Skip connection** (Row 4) operation enables to preserve more low-level features from the input layer, and the **bone length** (Row 5) awareness effectively narrows the search space, which are the two key designs to improve the inference results. Corresponding qualitative results are provided in Fig. 6 (a), which demonstrates our two nuanced designs also contribute to improving the overlay performance. These results illustrate the effectiveness of inference module and highlight the contribution of our algorithmic component designs.

Evaluation on optimization module. To evaluate the effectiveness of our optimization module design, we further compare to the two variants without the entire optimization module and without using the inertial terms in Eqn. 13 and Eqn. 14, respectively. The quantitative results are provided in Rows 6-7 of Tab. 2. Our inertial terms improve position accuracy while preserving motion plausibility (acceleration) from the inference module. The qualitative results provided in Fig. 6 (b) show our robust optimization improves the overlay performance effectively. Note that these variants suffer from tracking loss especially for the limbs with fast and challenging motions (generalization error). In contrast, our full pipeline achieves more robust capture.




Figure 7: Qualitative evaluation on our multi-modal input setting and IMU number configurations.

Input setting	MPJPE↓	PCK↑	AE↓
a) 4 IMUs Only	104.7	57.4	18.5
b) RGB Only	77.6	74.2	28.1
c) 2-IMU-aid	65.0	81.0	23.4
d) 4-IMU-aid	43.3	90.1	17.9
e) 12-IMU-aid	36.1	93.2	15.1

Table 3: Quantitative evaluation on our input configuration.

Evaluation on 4-IMU-aid setting. Finally, to verify the necessity and rationality of our multi-modal input setting with only 4 IMUs aiding, we compare several variants of our approach with various network input settings. As shown in Tab. 3, even using 2 IMUs aiding, our multi-modal input outperforms pure IMU or RGB input under single modality. Besides, our approach with 4 IMUs significantly outperforms the one with only two IMUs (one for left arm and one for right leg) and closes to the one with tedious 12 IMUs, which serves as a good compromise of acceptable performance and light-weight convenient capture setting. As shown in Fig. 7, our inertia-aid setting augments the camera by alleviating its inherent defects in terms of depth ambiguity and occlusion. As shown in Fig. 7, 4-IMU-aid setting enables to constantly track frequently occluded extremities while minimizing the intrusive body-worn sensors.

Conclusion

We present a practical approach HybridCap to capture challenging 3D human motions using only a single camera and 4 IMUs and achieve superior results compared to previous state-of-the-art methods. HybridCap uses a learning-and-optimization framework through novel cooperative bone tracking blocks with bone length awareness and an optimization-based refinement module with inferential prior. We also propose a new largest multi-modal human motion dataset, called HCM dataset, to evaluate their approach thoroughly. The experimental results demonstrate the robustness of HybridCap in capturing challenging human motions in various scenarios. We believe that it is a significant step for convenient and robust capture of human motions, with many potential applications in VR/AR and motion evaluations for gymnastics and dancing.

Acknowledgements

This work was supported by Shanghai YangFan Program (21YF1429500), Shanghai Local college capacity building program (22010502800), NSFC programs (61976138, 61977047), the National Key Research and Development Program (2018YFB2100500), STCSM (2015F0203-000-06), and SHMEC (2019-01-07-00-01-E00003).

References

- Agarwal, S.; Mierle, K.; and The Ceres Solver Team. 2022. Ceres Solver. <https://github.com/ceres-solver/ceres-solver>. Accessed: 2023-03-26.
- Andrews, S.; Huerta, I.; Komura, T.; Sigal, L.; and Mitchell, K. 2016. Real-time physics-based motion capture with sparse sensors. In *Proceedings of the 13th European conference on visual media production (CVMP 2016)*, 1–10.
- Anguelov, D.; Srinivasan, P.; Koller, D.; Thrun, S.; Rodgers, J.; and Davis, J. 2005. SCAPE: Shape Completion and Animation of People. In *ACM SIGGRAPH 2005 Papers*, SIGGRAPH '05, 408–416. New York, NY, USA: Association for Computing Machinery. ISBN 9781450378253.
- Bogo, F.; Kanazawa, A.; Lassner, C.; Gehler, P.; Romero, J.; and Black, M. J. 2016. Keep It SMPL: Automatic Estimation of 3D Human Pose and Shape from a Single Image. In Leibe, B.; Matas, J.; Sebe, N.; and Welling, M., eds., *Computer Vision – ECCV 2016*, 561–578. Cham: Springer International Publishing.
- Bregler, C.; and Malik, J. 1998. Tracking people with twists and exponential maps. In *Computer Vision and Pattern Recognition (CVPR)*.
- Cao, Z.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2017. Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields. In *Computer Vision and Pattern Recognition (CVPR)*.
- Collet, A.; Chuang, M.; Sweeney, P.; Gillett, D.; Evseev, D.; Calabrese, D.; Hoppe, H.; Kirk, A.; and Sullivan, S. 2015. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics (TOG)*, 34(4): 69.
- Davison, A. J.; Deutscher, J.; and Reid, I. D. 2001. Markerless Motion Capture of Complex Full-body Movement for Character Animation. In *Eurographics Workshop on Computer Animation and Simulation*.
- De Aguiar, E.; Stoll, C.; Theobalt, C.; Ahmed, N.; Seidel, H.-P.; and Thrun, S. 2008. Performance capture from sparse multi-view video. In *ACM SIGGRAPH 2008 papers*, 1–10.
- Gall, J.; Rosenhahn, B.; Brox, T.; and Seidel, H.-P. 2010. Optimization and Filtering for Human Motion Capture. *International Journal of Computer Vision (IJCV)*, 87(1–2): 75–92.
- Gilbert, A.; Trumble, M.; Malleson, C.; Hilton, A.; and Colloso, J. 2019. Fusing visual and inertial sensors with semantics for 3d human pose estimation. *International Journal of Computer Vision*, 127(4): 381–397.
- Habermann, M.; Xu, W.; Zollhöfer, M.; Pons-Moll, G.; and Theobalt, C. 2019. LiveCap: Real-Time Human Performance Capture From Monocular Video. *ACM Transactions on Graphics (TOG)*, 38(2): 14:1–14:17.
- Habermann, M.; Xu, W.; Zollhofer, M.; Pons-Moll, G.; and Theobalt, C. 2020. DeepCap: Monocular Human Performance Capture Using Weak Supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hasler, N.; Rosenhahn, B.; Thormahlen, T.; Wand, M.; Gall, J.; and Seidel, H.-P. 2009. Markerless motion capture with unsynchronized moving cameras. In *Computer Vision and Pattern Recognition (CVPR)*, 224–231.
- He, Y.; Pang, A.; Chen, X.; Liang, H.; Wu, M.; Ma, Y.; and Xu, L. 2021. Challengcap: Monocular 3d capture of challenging human performances using multi-modal references. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11400–11411.
- Helten, T.; Müller, M.; Seidel, H.-P.; and Theobalt, C. 2013. Real-time body tracking with one depth camera and inertial sensors. In *Proceedings of the IEEE international conference on computer vision*, 1105–1112.
- Henschel, R.; Von Marcard, T.; and Rosenhahn, B. 2020. Accurate long-term multiple people tracking using video and body-worn IMUs. *IEEE Transactions on Image Processing*, 29: 8476–8489.
- Huang, Y.; Bogo, F.; Lassner, C.; Kanazawa, A.; Gehler, P. V.; Romero, J.; Akhter, I.; and Black, M. J. 2017. Towards Accurate Marker-Less Human Shape and Pose Estimation over Time. In *2017 International Conference on 3D Vision (3DV)*, 421–430.
- Huang, Y.; Kaufmann, M.; Aksan, E.; Black, M. J.; Hilliges, O.; and Pons-Moll, G. 2018. Deep inertial poser: Learning to reconstruct human pose from sparse inertial measurements in real time. *ACM Transactions on Graphics (TOG)*, 37(6): 1–15.
- Joo, H.; Liu, H.; Tan, L.; Gui, L.; Nabbe, B.; Matthews, I.; Kanade, T.; Nobuhara, S.; and Sheikh, Y. 2015. Panoptic Studio: A Massively Multiview System for Social Motion Capture. In *Proceedings of the IEEE International Conference on Computer Vision*, 3334–3342.
- Joo, H.; Simon, T.; and Sheikh, Y. 2018. Total Capture: A 3D Deformation Model for Tracking Faces, Hands, and Bodies. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kanazawa, A.; Black, M. J.; Jacobs, D. W.; and Malik, J. 2018. End-to-end Recovery of Human Shape and Pose. In *Computer Vision and Pattern Recognition (CVPR)*.
- Kanazawa, A.; Zhang, J. Y.; Felsen, P.; and Malik, J. 2019. Learning 3D Human Dynamics From Video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kaufmann, M.; Zhao, Y.; Tang, C.; Tao, L.; Twigg, C.; Song, J.; Wang, R.; and Hilliges, O. 2021. EM-POSE: 3D Human Pose Estimation From Sparse Electromagnetic Trackers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11510–11520.
- Kocabas, M.; Athanasiou, N.; and Black, M. J. 2020. VIBE: Video Inference for Human Body Pose and Shape Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

- Kocabas, M.; Huang, C.-H. P.; Hilliges, O.; and Black, M. J. 2021. PARE: Part Attention Regressor for 3D Human Body Estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 11127–11137.
- Kolotouros, N.; Pavlakos, G.; Black, M. J.; and Daniilidis, K. 2019. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2252–2261.
- Kolotouros, N.; Pavlakos, G.; and Daniilidis, K. 2019. Convolutional Mesh Regression for Single-Image Human Shape Reconstruction. In *Computer Vision and Pattern Recognition (CVPR)*.
- Kolotouros, N.; Pavlakos, G.; Jayaraman, D.; and Daniilidis, K. 2021. Probabilistic Modeling for Human Mesh Recovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11605–11614.
- Lassner, C.; Romero, J.; Kiefel, M.; Bogo, F.; Black, M. J.; and Gehler, P. V. 2017. Unite the people: Closing the loop between 3d and 2d human representations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6050–6059.
- Li, R.; Yang, S.; Ross, D. A.; and Kanazawa, A. 2021. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13401–13412.
- Liang, H.; Zhang, W.; Li, W.; Yu, J.; and Xu, L. 2023. InterGen: Diffusion-based Multi-human Motion Generation under Complex Interactions. *arXiv preprint arXiv:2304.05684*.
- Liu, Y.; Gall, J.; Stoll, C.; Dai, Q.; Seidel, H.-P.; and Theobalt, C. 2013. Markerless motion capture of multiple characters using multiview image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(11): 2720–2735.
- Loper, M.; Mahmood, N.; Romero, J.; Pons-Moll, G.; and Black, M. J. 2015. SMPL: A Skinned Multi-person Linear Model. *ACM Trans. Graph.*, 34(6): 248:1–248:16.
- Mahmood, N.; Ghorbani, N.; Troje, N. F.; Pons-Moll, G.; and Black, M. J. 2019. AMASS: Archive of Motion Capture As Surface Shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Malleson, C.; Collomosse, J.; and Hilton, A. 2019. Real-time multi-person motion capture from multi-view video and IMUs. *International Journal of Computer Vision*, 1–18.
- Malleson, C.; Gilbert, A.; Trumble, M.; Collomosse, J.; Hilton, A.; and Volino, M. 2017. Real-time full-body motion capture from video and imus. In *2017 International Conference on 3D Vision (3DV)*, 449–457. IEEE.
- Movella. 2022. Movella Xsens Products. <https://www.movella.com/products/xsens>. Accessed: 2023-03-26.
- Pavlakos, G.; Choutas, V.; Ghorbani, N.; Bolkart, T.; Osman, A. A. A.; Tzionas, D.; and Black, M. J. 2019. Expressive Body Capture: 3D Hands, Face, and Body from a Single Image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 10975–10985.
- Pavlakos, G.; Zhou, X.; Derpanis, K. G.; and Daniilidis, K. 2017. Harvesting Multiple Views for Marker-less 3D Human Pose Annotations. In *Computer Vision and Pattern Recognition (CVPR)*.
- Pons-Moll, G.; Baak, A.; Gall, J.; Leal-Taixe, L.; Mueller, M.; Seidel, H.-P.; and Rosenhahn, B. 2011. Outdoor human motion capture using inverse kinematics and von mises-fisher sampling. In *2011 International Conference on Computer Vision*, 1243–1250. IEEE.
- Pons-Moll, G.; Baak, A.; Helten, T.; Müller, M.; Seidel, H.-P.; and Rosenhahn, B. 2010. Multisensor-fusion for 3d full-body human motion capture. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 663–670. IEEE.
- Rempe, D.; Birdal, T.; Hertzmann, A.; Yang, J.; Sridhar, S.; and Guibas, L. J. 2021. HuMoR: 3D Human Motion Model for Robust Pose Estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 11488–11499.
- Ren, Y.; Zhao, C.; He, Y.; Cong, P.; Liang, H.; Yu, J.; Xu, L.; and Ma, Y. 2023. LiDAR-aid Inertial Poser: Large-scale Human Motion Capture by Sparse Inertial and LiDAR Sensors. *IEEE Transactions on Visualization and Computer Graphics*, 1–11.
- Robertini, N.; Casas, D.; Rhodin, H.; Seidel, H.-P.; and Theobalt, C. 2016. Model-based Outdoor Performance Capture. In *International Conference on 3D Vision (3DV)*.
- Shimada, S.; Golyanik, V.; Xu, W.; Pérez, P.; and Theobalt, C. 2021. Neural Monocular 3D Human Motion Capture with Physical Awareness. *ACM Transactions on Graphics*, 40(4).
- Simon, T.; Joo, H.; Matthews, I.; and Sheikh, Y. 2017. Hand Keypoint Detection in Single Images using Multiview Bootstrapping. In *Computer Vision and Pattern Recognition (CVPR)*.
- Stoll, C.; Hasler, N.; Gall, J.; Seidel, H.-P.; and Theobalt, C. 2011. Fast articulated motion tracking using a sums of Gaussians body model. In *International Conference on Computer Vision (ICCV)*.
- Theobalt, C.; de Aguiar, E.; Stoll, C.; Seidel, H.-P.; and Thrun, S. 2010. Performance capture from multi-view video. In *Image and Geometry Processing for 3-D Cinematography*, 127–149. Springer.
- Tsuchida, S.; Fukayama, S.; Hamasaki, M.; and Goto, M. 2019. AIST Dance Video Database: Multi-genre, Multi-dancer, and Multi-camera Database for Dance Information Processing. In *Proceedings of the 20th International Society for Music Information Retrieval Conference, ISMIR 2019*, 501–510. Delft, Netherlands.
- Vicon. 2022. Vicon Mocap Systems. <https://www.vicon.com>. Accessed: 2023-03-26.
- Vlasic, D.; Adelsberger, R.; Vannucci, G.; Barnwell, J.; Gross, M.; Matusik, W.; and Popović, J. 2007. Practical motion capture in everyday surroundings. *ACM transactions on graphics (TOG)*, 26(3): 35–es.
- von Marcard, T.; Henschel, R.; Black, M.; Rosenhahn, B.; and Pons-Moll, G. 2018. Recovering Accurate 3D Human

- Pose in The Wild Using IMUs and a Moving Camera. In *European Conference on Computer Vision (ECCV)*.
- Von Marcard, T.; Pons-Moll, G.; and Rosenhahn, B. 2016. Human pose estimation from video and imus. *IEEE transactions on pattern analysis and machine intelligence*, 38(8): 1533–1547.
- Von Marcard, T.; Rosenhahn, B.; Black, M. J.; and Pons-Moll, G. 2017. Sparse inertial poser: Automatic 3d human pose estimation from sparse imus. In *Computer Graphics Forum*, volume 36, 349–360. Wiley Online Library.
- Wang, Y.; Liu, Y.; Tong, X.; Dai, Q.; and Tan, P. 2017. Outdoor Markerless Motion Capture With Sparse Handheld Video Cameras. *Transactions on Visualization and Computer Graphics (TVCG)*.
- Xiang, D.; Joo, H.; and Sheikh, Y. 2019. Monocular Total Capture: Posing Face, Body, and Hands in the Wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xu, L.; Liu, Y.; Cheng, W.; Guo, K.; Zhou, G.; Dai, Q.; and Fang, L. 2018a. FlyCap: Markerless Motion Capture Using Multiple Autonomous Flying Cameras. *IEEE Transactions on Visualization and Computer Graphics*, 24(8): 2284–2297.
- Xu, L.; Xu, W.; Golyanik, V.; Habermann, M.; Fang, L.; and Theobalt, C. 2020. Eventcap: Monocular 3d capture of high-speed human motions using an event camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4968–4978.
- Xu, W.; Chatterjee, A.; Zollhöfer, M.; Rhodin, H.; Mehta, D.; Seidel, H.-P.; and Theobalt, C. 2018b. MonoPerfCap: Human Performance Capture From Monocular Video. *ACM Transactions on Graphics (TOG)*, 37(2): 27:1–27:15.
- Yi, X.; Zhou, Y.; Habermann, M.; Shimada, S.; Golyanik, V.; Theobalt, C.; and Xu, F. 2022. 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)*.
- Yi, X.; Zhou, Y.; and Xu, F. 2021. Transpose: Real-time 3d human translation and pose estimation with six inertial sensors. *ACM Transactions on Graphics (TOG)*, 40(4): 1–13.
- Zanfir, A.; Bazavan, E. G.; Zanfir, M.; Freeman, W. T.; Sukthankar, R.; and Sminchisescu, C. 2021. Neural descent for visual 3d human pose and shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14484–14493.
- Zhang, Z.; Wang, C.; Qin, W.; and Zeng, W. 2020. Fusing wearable imus with multi-view images for human pose estimation: A geometric approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2200–2209.
- Zheng, Z.; Yu, T.; Li, H.; Guo, K.; Dai, Q.; Fang, L.; and Liu, Y. 2018. HybridFusion: Real-time Performance Capture Using a Single Depth Sensor and Sparse IMUs. In *European Conference on Computer Vision (ECCV)*.
- Zheng, Z.; Yu, T.; Wei, Y.; Dai, Q.; and Liu, Y. 2019. DeepHuman: 3D Human Reconstruction From a Single Image. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Zhou, Y.; Barnes, C.; Lu, J.; Yang, J.; and Li, H. 2019. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5745–5753.