



# EgoLocate: Real-time Motion Capture, Localization, and Mapping with Sparse Body-mounted Sensors

XINYU YI, School of Software and BNRist, Tsinghua University, China

YUXIAO ZHOU, ETH Zurich, Switzerland

MARC HABERMANN, Max Planck Institute for Informatics, Saarland Informatics Campus, Germany

VLADISLAV GOLYANIK, Max Planck Institute for Informatics, Saarland Informatics Campus, Germany

SHAOHUA PAN, School of Software and BNRist, Tsinghua University, China

CHRISTIAN THEOBALT, Max Planck Institute for Informatics, Saarland Informatics Campus, Germany

FENG XU, School of Software and BNRist, Tsinghua University, China




Fig. 1. Key results of our method. *Left*: for a user moving in a 3D scene, our technique estimates the user's pose, localizes the user's positions in the scene (green moving trajectory, loop closed after 143m motion with upstairs and downstairs), and reconstructs sparse 3D points of the scene (black scene points) simultaneously in real time. *Right*: benefiting from the sensor fusion, the localization accuracy of our method is largely improved compared with the state-of-the-art techniques.

Human and environment sensing are two important topics in Computer Vision and Graphics. Human motion is often captured by inertial sensors, while the environment is mostly reconstructed using cameras. We integrate the two techniques together in EgoLocate, a system that simultaneously performs human motion capture (mocap), localization, and mapping in real time from sparse body-mounted sensors, including 6 inertial measurement units (IMUs) and a monocular phone camera. On one hand, inertial mocap suffers from

large translation drift due to the lack of the global positioning signal. EgoLocate leverages image-based simultaneous localization and mapping (SLAM) techniques to locate the human in the reconstructed scene. On the other hand, SLAM often fails when the visual feature is poor. EgoLocate involves inertial mocap to provide a strong prior for the camera motion. Experiments show that localization, a key challenge for both two fields, is largely improved by our technique, compared with the state of the art of the two fields. Our codes are available for research at <https://xinyu-yi.github.io/EgoLocate/>.

Authors' addresses: Xinyu Yi, School of Software and BNRist, Tsinghua University, Beijing, China, [yixy20@mails.tsinghua.edu.cn](mailto:yixy20@mails.tsinghua.edu.cn); Yuxiao Zhou, ETH Zurich, Zurich, Switzerland, [yuxiao.zhou@outlook.com](mailto:yuxiao.zhou@outlook.com); Marc Habermann, Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany, [mhaberma@mpi-inf.mpg.de](mailto:mhaberma@mpi-inf.mpg.de); Vladislav Golyanik, Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany, [golyanik@mpi-inf.mpg.de](mailto:golyanik@mpi-inf.mpg.de); Shaohua Pan, School of Software and BNRist, Tsinghua University, Beijing, China, [ishspan@163.com](mailto:ishspan@163.com); Christian Theobalt, Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany, [theobalt@mpi-inf.mpg.de](mailto:theobalt@mpi-inf.mpg.de); Feng Xu, School of Software and BNRist, Tsinghua University, Beijing, China, [xufeng2003@gmail.com](mailto:xufeng2003@gmail.com).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2023/8-ART76 \$15.00 <https://doi.org/10.1145/3592099>

CCS Concepts: • **Computing methodologies** → **Motion capture**.

Additional Key Words and Phrases: Human-centered, Sparse IMUs, Map-free, SLAM, Real-time

## ACM Reference Format:

Xinyu Yi, Yuxiao Zhou, Marc Habermann, Vladislav Golyanik, Shaohua Pan, Christian Theobalt, and Feng Xu. 2023. EgoLocate: Real-time Motion Capture, Localization, and Mapping with Sparse Body-mounted Sensors. *ACM Trans. Graph.* 42, 4, Article 76 (August 2023), 17 pages. <https://doi.org/10.1145/3592099>

## 1 INTRODUCTION

Humans and environments are two important components of our world. Human motion sensing, aiming to capture and understand human activities, is a widely studied topic with many applications such as VR/AR, sports, and gaming. Meanwhile, environment sensing

techniques, which get the knowledge of the surrounding environment including its 3D shape and semantics, have long been studied and are widely used in areas such as autonomous driving and robot navigation.

While humans and environments are two indispensable and interdependent parts of the world, existing techniques mostly handle them independently. We believe that the combinative sensing of human motion and environments is of great importance for scenarios where humans interact with the environment. And it may benefit new applications such as human behavior understanding, human localization, and motion planning.

As an important human sensing technique, human motion capture (mocap) aims to reconstruct kinematic body motions and is widely studied in the literature. The works that leverage multi-view external cameras (some even with multiple markers on the human body) usually achieve robust and accurate results [Chen et al. 2020; Dong et al. 2019; Reddy et al. 2021; Shao et al. 2022; Zhang et al. 2020]. However, the capture space is constrained in the camera frustums as the performers need to be visible to the cameras. While egocentric works do not constrain the recording space [Akada et al. 2022; Tome et al. 2019; Xu et al. 2019] and achieve lightweight reconstruction and real-time performance [Jiang et al. 2022b; Yi et al. 2022, 2021], they often suffer from translation drifts as there are no global positioning signals, which results in the wrong position of the human after a period of motion.

For environment sensing, simultaneous localization and mapping (SLAM) is the capstone technique that leverages a moving camera to reconstruct the 3D scene and locate the camera in the scene simultaneously. In this setting, camera localization is a key as it is required to fuse the 3D information at different time instances together to make a complete 3D scene. Despite previous efforts, current pure visual SLAM [Davison 2003; Engel et al. 2014; Forster et al. 2014; Klein and Murray 2007; Mur-Artal and Tardós 2017a] may still fail to track the camera when feature absence or motion blur happens. By incorporating an IMU into the camera, visual-inertial SLAM becomes more robust when reliable visual features are insufficient [Campos et al. 2021; Leutenegger et al. 2015; Mourikis et al. 2007; Qin et al. 2018], but still suffers from tracking losses in long-time camera occlusion and fast camera rotation, especially when the frequency of the IMU signals is low. This will also result in large localization errors or even system failure.

Localization is a key task for both mocap and SLAM and is very challenging. While inertial-sensor-based mocap explores *inner information* such as human motion signals and motion priors, SLAM majorly relies on *outer information*, *i.e.*, environment captured by the camera. The former achieves good stability, but the global position drifts accumulate in long-time motions as no correction from outside is available; the latter estimates the global position in the scene with high accuracy, but suffers from tracking losses when the environmental information is unreliable, *e.g.*, in textureless or occluded places. To this end, we propose EgoLocate, which effectively combines the two complementary techniques (mocap and SLAM) together and achieves robust and accurate localization (see Fig. 1). Our system leverages sparse body-mounted sensors including 6 IMUs mounted on the forearms, the lower legs, the head, and the pelvis of the human, and a monocular RGB camera attached to the

head looking outwards. This design is inspired by the behavior of real humans: when humans are in a new environment, they use their eyes to get knowledge of the scene with their location and plan their movements in the scene. The monocular camera acts as the "eye" that provides visual signals to our method for real-time scene reconstruction and self-localization, while the IMUs measure human motion in the scene.

We propose an organically coupled framework to exploit the complementary advantages of the sparse inertial mocap and SLAM techniques. In this framework, the human motion prior is integrated into multiple key components of SLAM, and the SLAM localization is fed back to the mocap routine. By jointly optimizing the camera pose and map point positions with the awareness of mocap, both the tracking and mapping are improved in terms of accuracy and robustness. When the visual cues are reliable, the SLAM is able to correct the drift of the mocap using the environmental information; when the visual features are sparse due to camera occlusion or extreme lighting, the mocap module can provide pose and translation estimation for the SLAM system, avoiding total failure in previous SLAM systems. Furthermore, we propose a mocap-related map point confidence, which helps to dynamically determine the importance of each map point in the bundle adjustment. By reducing the influence of potentially wrong 3D points, our method achieves a good balance between the map point-based constraints and the mocap constraints, which helps to reduce the uncertainty in the tracking and improve the localization accuracy. Combining the mocap and the SLAM system results in a win-win. On the one hand, with the help of camera localization, the translation drift in inertial mocap is largely reduced. On the other hand, with the help of mocap, the accuracy and robustness of localization and mapping are greatly improved.

To the best of our knowledge, our method is the first effort that achieves real-time mocap and simultaneous mapping from sparse body-mounted sensors. The most related works to ours are HPS [Guzov et al. 2021] and HSC4D [Dai et al. 2022]. HPS performs offline human motion capture and localization in a known scene from 17 densely placed IMUs and a body-mounted camera, and requires a pre-scanned scene with plenty of registered photos. HSC4D leverages 17 IMUs and a LiDAR sensor placed on the human body, and estimates the human motion and the scene simultaneously but also in an offline manner. Compared with these works, our system is much more *lightweight*, as we only use 6 IMUs and a body-mounted monocular camera. Importantly, we also do *not* rely on pre-scanning the scene. Moreover, our system performs motion capture, localization, and mapping in *real time*.

In summary, our contributions are:

- We propose the first real-time simultaneous human motion capture, localization, and mapping system using only 6 IMUs and a body-mount camera, which estimates drift-free human motion and sparse scene points in unconstrained 3D space.
- We propose a tightly-coupled optimization framework by incorporating the mocap prior in key SLAM modules, which improves both the localization and mapping accuracy as well as the robustness to poor visual features.

- We propose mocap-related map point confidence to dynamically adjust the influence of each map point and the mocap constraints in the bundle adjustment, which improves the localization accuracy while reducing the uncertainty in the results.

Our work is trying to merge human sensing and environment sensing together. Even though we mostly focus on localization, we believe this work is a first step toward joint motion capture and dense environment sensing and reconstruction.

## 2 RELATED WORK

### 2.1 Egocentric Human Motion Capture

Egocentric human motion capture (mocap) is widely studied due to its advantage that the recording space is not constrained in a fixed volume. Existing works leverage different body-mount sensors to capture or infer human motions. One category of works uses body-mounted cameras, either capturing the human motion from egocentric views (cameras looking at the wearer) [Akada et al. 2022; Hori et al. 2022; Rhodin et al. 2016; Tome et al. 2020, 2019; Wang et al. 2021, 2023] or estimating the human motion from first-person views (cameras looking outwards) [Jiang and Grauman 2017; Li et al. 2022; Luo et al. 2021; Shiratori et al. 2011; Yuan and Kitani 2019, 2018]. These techniques enable space-free human mocap and have yielded many impressive results. However, the former setting may encounter self-occlusions and difficulty in detecting global movements, limiting the accuracy of pose and translation estimation. While the latter is limited in its ability to reconstruct certain types of human motion as the human body is not visible in the camera view. A combination of egocentric-view and first-person-view cameras can help with translation estimation [Liu et al. 2022]. Some approaches [Henning et al. 2022; Liu et al. 2021; von Marcard et al. 2018; Yuan et al. 2022] leverage a moving external camera, *e.g.*, held by another person, to capture the human motion. They do an excellent job of decoupling human and camera motion, *e.g.*, by leveraging/reconstructing the static background scene. Overall, these camera-based approaches have made remarkable progress in enabling space-free human motion capture. Another category of works captures human motions from sparse non-vision-based body-mounted sensors, including accelerometers [Riaz et al. 2015; Slyper and Hodgins 2008; Tautges et al. 2011], inertial measurement units (IMUs) [Cha et al. 2021; Geissinger and Asbeck 2020; Huang et al. 2018; Jiang et al. 2022b; Puchert and Ropinski 2021; Vlasic et al. 2007; von Marcard et al. 2017; Xia et al. 2022; Yi et al. 2022, 2021], and electromagnetic sensors [Kaufmann et al. 2021]. They overcome common limitations of vision-based approaches like occlusions, and can estimate accurate human articulated pose and global movements within a short period after calibration. However, since there are no global positioning signals from the sensors, their translation drifts can be large, resulting in the wrong position of the human after a long period of motion. This can lead to penetration and even wrong interaction with the digital scene, which is undesirable in applications like VR and gaming. Some works leverage body-mounted VR devices [Jiang et al. 2022a; Winkler et al. 2022] to capture the human motion in a constrained space. HPS [Guzov et al. 2021] and its extension [Guzov et al. 2022] propose to locate the

human in the environment by leveraging a head-mounted camera and densely-worn IMUs. They achieve good accuracy on human localization, but they rely on the pre-scanned scene with plenty of registered photos and run in an offline manner. HSC4D [Dai et al. 2022] simultaneously reconstructs the human motion and the scene with good consistency. However, they need dense IMUs and a body-worn LiDAR sensor and also run in an offline manner. While we use *sparse* IMUs and a body-mounted camera to capture the space-free human motion in *real time*, without the need to pre-scan the scene.

### 2.2 Visual and Visual-Inertial SLAM

Simultaneous localization and mapping (SLAM) has long been studied due to its importance in scene reconstruction and localization in unknown environments. Visual SLAM [Bloesch et al. 2018; Davison 2003; Engel et al. 2017, 2014; Forster et al. 2014; Koestler et al. 2022; Mur-Artal and Tardós 2017a; Teed and Deng 2021; Yang et al. 2020; Zubizarreta et al. 2020] leverages a monocular camera and estimates the camera movements while simultaneously reconstructing the scene in 3D points or structures. However, they often fail to track the camera when the visual feature is poor, *e.g.*, due to motion blur or camera occlusion. They are also unaware of the true scale of the scene and the camera movement. Visual-inertial SLAM incorporates an IMU to the camera, which becomes more robust to fast camera motions and poor visual features, and can also determine the scale factor [Campos et al. 2021; Leutenegger et al. 2015; Mourikis et al. 2007; Mur-Artal and Tardós 2017b; Qin et al. 2018; Von Stumberg et al. 2018]. However, they still suffer from tracking failures when the camera is occluded for a longer time, and are also sensitive to the camera-IMU relative pose (IMU extrinsic) and IMU noise parameters (IMU intrinsic). Some works [Campos et al. 2021; Castle et al. 2008; Schmuck and Chli 2017] use a multi-map system to improve the robustness. [Campos et al. 2021] starts a new map when the tracking is lost. However, the online localization results will be reset in the new map, which means it is not possible to obtain the absolute position in real time until the new map is merged with the initial map. Our system focuses on situations where humans wear the camera. In this situation, the existing SLAM techniques are mostly unaware of the human prior. To this end, we propose to incorporate inertial human mocap with SLAM techniques, which not only improve the accuracy of the localization and mapping, but also enhance the robustness of the system.

## 3 METHOD

Our goal is to use body-worn sensors to capture the body motion while simultaneously reconstructing the environment and locating the human in the environment in real time. The input of our system is the synchronized signals from the sensors, including the inertial measurements from 6 IMUs and the monocular images from a head-mounted camera. The system estimates the human pose and global motion, as well as the reconstructed scene as a sparse 3D point cloud; see Fig. 2 for an overview. Our method first performs human motion capture from the inertial inputs (Sec. 3.1). Then, the camera pose is optimized by minimizing the reprojection error of the reconstructed map with mocap constraints (Sec. 3.2). The optimized camera pose is then used to update the human translation to avoid the drifting




Fig. 2. **Overview of our method.** The inputs are real-time measurements of 6 IMUs and images from a body-worn camera. We first estimate initial human poses and camera poses from sparse inertial signals (Inertial Motion Capture, Sec. 3.1). Next, we refine the camera pose by minimizing the mocap constraints and the reprojection errors of the 3D-2D feature matches (Camera Tracking, Sec. 3.2). Then, the refined camera pose and a confidence are used to correct the human’s global position and velocity (Body Translation Updater, Sec. 3.4). This results in 60-FPS drift-free human motion output. In parallel, we perform mapping and loop closing using keyframes (Mapping & Loop Closing, Sec. 3.3). We calculate another confidence for each map point, and use it in bundle adjustment (BA), where we jointly optimize the map points and the camera poses by a combined mocap and weighted reprojection error. If a loop is detected, we also perform pose graph optimization with mocap constraints. The estimated map is also outputted in real time.

artifacts caused by inertial error accumulation (Sec. 3.4). In parallel to the motion tracking, our method reconstructs the scene map and performs loop closing with mocap awareness (Sec. 3.3). The calibration of the system is discussed in Sec. 3.5, including how to initialize the inertial mocap and SLAM respectively, as well as how to calculate the extrinsic between the mocap and the SLAM system.

### 3.1 Inertial Motion Capture

From the orientation and acceleration measurements of 6 body-mounted IMUs, we first estimate an initial human pose and translation. Our method follows the sparse inertial mocap work PIP [Yi et al. 2022], but we remove its planar assumption that the whole scene is a known level ground and the human movements are on the ground because our goal is to enable unconstrained mocap in free 3D space. We also re-design the motion optimizer in PIP by removing the force calculation. This is because we do not know the dense geometry of the scene (the reconstruction is only in sparse scene points), making it impossible to detect human-environment collisions to apply forces on the contacts. More details about the modification to PIP are presented in App. C.

Our inertial mocap estimates human pose (in terms of joint angles) and the human global rotation and translation in a user-defined mocap coordinate system at 60 FPS. Then, the 6-DoF pose of the head-mounted camera can be extracted from the inertial mocap result with the calibrated camera-head extrinsic (see Sec. 3.5), which will be used as an initialization of the following camera tracking step. Note that the inertial mocap provides good articulated pose

estimation, but the global motion of the human drifts significantly over time.

### 3.2 Camera Tracking

This module aims to estimate better camera poses by using the 30-FPS images from the head-mounted camera and the estimated camera poses from the inertial mocap, which are synchronized to the same frame rate of the images. We design this module based on ORB-SLAM3 [Campos et al. 2021], where we first extract ORB feature points (keypoints) from the image and find the matches between the keypoints and the 3D map points using feature similarity. Note that the map points are reconstructed and maintained online by a mapping module (see Sec. 3.3). Then, we optimize the camera pose by minimizing the reprojection error of the matches with additional mocap constraints.

**3.2.1 Mocap-constrained Camera Tracking.** Formally, we denote the world position of the  $i$ th 3D map point as  $x_i^{3D} \in \mathbb{R}^3$  and the pixel coordinates of the matched 2D keypoint as  $x_i^{2D} \in \mathbb{R}^2$ , where  $i \in \mathcal{X}$  for all matches. We use  $\bar{R} \in SO(3)$  and  $\bar{t} \in \mathbb{R}^3$  to denote the camera pose *w.r.t* the world before the optimization, which is obtained from the inertial mocap. The mocap-constrained camera tracking is then

formulated as the nonlinear optimization of the camera pose  $\mathbf{R}, \mathbf{t}$ :

$$\begin{aligned} & \arg \min_{\mathbf{R}, \mathbf{t}} \mathcal{E}_{\text{proj}}(\mathbf{R}, \mathbf{t}) + \mathcal{E}_{\text{mocap}}(\mathbf{R}, \mathbf{t}) \\ \mathcal{E}_{\text{proj}} &= \sum_{i \in \mathcal{X}} \rho \left( \left\| \mathbf{x}_i^{2D} - \pi \left( \mathbf{R}^T (\mathbf{x}_i^{3D} - \mathbf{t}) \right) \right\|_{\Sigma_i}^2 \right) \\ \mathcal{E}_{\text{mocap}} &= \lambda_R \left\| \text{Log}(\bar{\mathbf{R}}^T \mathbf{R}) \right\|^2 + \lambda_t \|\bar{\mathbf{t}} - \mathbf{t}\|^2, \end{aligned} \quad (1)$$

where  $\rho$  is the Huber robust cost function,  $\Sigma_i$  is the covariance matrix associated with the keypoint scale<sup>1</sup>,  $\text{Log} : \text{SO}(3) \rightarrow \mathbb{R}^3$  is a mapping from the Lie group to the vector space, and  $\pi$  is the camera projection function.  $\lambda_R$  and  $\lambda_t$  are the coefficients aligning the unit and controlling the weights of the corresponding terms. The reprojection term  $\mathcal{E}_{\text{proj}}(\mathbf{R}, \mathbf{t})$  refines the camera pose by minimizing the projection error of the matched 3D-2D points, while the mocap term  $\mathcal{E}_{\text{mocap}}(\mathbf{R}, \mathbf{t})$  constrains the camera pose to be near to the mocap-based pose. Similar to [Campos et al. 2021], we perform the optimization 3 times and classify the keypoint matches as inliers/outliers after each optimization, where outliers are excluded from the next optimization. To this end, the mocap term provides a strong prior on the camera pose, which helps with the classification of the wrong 3D-2D matches and reduces the errors caused by outliers. The coefficients of the mocap terms are set to  $\lambda_R = 0.01f^2$  and  $\lambda_t = 0.5f^2s^2$  where  $f$  is the camera focal length and  $s$  is the scale factor of the SLAM coordinate system computed in the calibration module (Sec. 3.5). This automatically converts the unit of the mocap errors to pixels, which is at the same magnitude as the reprojection error. Such a strategy makes the optimization independent with the camera intrinsic and the scale factor. We leverage the Levenberg–Marquardt algorithm implemented in g2o [Kümmerle et al. 2011] to solve the nonlinear optimization. To ensure real-time performance, we calculate the Jacobian analytically, which is presented in App. D, and retrieve a local visible map (instead of using the entire map) for the keypoint matching and camera tracking, which is detailed in App. E. After the optimization, we extract the number of the inlier 3D-2D matches  $n$  as the confidence to indicate the quality of the refined camera pose, which will be used in the following module. Intuitively, we have larger confidence if more map points have been successfully matched.

**3.2.2 Camera Pose Alignment.** In practice, the absolute camera translation derived from pure inertial mocap usually contains large drifts due to inertial error accumulation. As a result, directly using it in the camera tracking optimization is not feasible for long sequences because such non-linear optimizations typically require a good initialization. Therefore, before optimization, we perform a camera pose alignment step for each frame to reduce the drifts. Specifically, we compute the *relative* camera rotation and translation from the mocap and add it to the previous SLAM-optimized camera pose. This is based on the observation that the refined camera poses in SLAM consider the visual information and the drift problem is largely reduced. We denote the SLAM-optimized camera pose as  $\mathbf{R} \in \text{SO}(3)$ ,  $\mathbf{t} \in \mathbb{R}^3$  for the camera orientation and position *w.r.t* the world respectively, and the camera pose extracted from mocap as

<sup>1</sup>The *scale* of a keypoint is determined by the level where it is detected in the image pyramid. See [Mur-Artal et al. 2015].

$\tilde{\mathbf{R}}, \tilde{\mathbf{t}}$ . Then, the alignment can be written as:

$$\begin{aligned} \bar{\mathbf{R}}_{\text{cur}} &= \mathbf{R}_{\text{last}} \tilde{\mathbf{R}}_{\text{last}}^T \tilde{\mathbf{R}}_{\text{cur}} \\ \bar{\mathbf{t}}_{\text{cur}} &= \mathbf{t}_{\text{last}} - \tilde{\mathbf{t}}_{\text{last}} + \tilde{\mathbf{t}}_{\text{cur}}, \end{aligned} \quad (2)$$

where  $\bar{\mathbf{R}}_{\text{cur}}$  and  $\bar{\mathbf{t}}_{\text{cur}}$  are the aligned camera pose at the current frame, subscript  $\cdot_{\text{cur}}$  and  $\cdot_{\text{last}}$  denote the current and last frame respectively. We use the last *keyframe* (defined later) for the alignment because its pose is more accurate as it has been optimized in the bundle adjustment (BA) (Sec. 3.3). However, if a relocalization<sup>2</sup> recently happens, when no keyframe is available after the relocalization, we choose to use the latest *frame* after the relocalization to avoid the drift. As the orientation estimation  $\tilde{\mathbf{R}}_{\text{cur}}$  does not suffer from evident drifts in inertial mocap, we finally always rotate  $\bar{\mathbf{R}}_{\text{cur}}$  towards  $\tilde{\mathbf{R}}_{\text{cur}}$  by 0.1 through spherical linear interpolation. This ensures that the aligned orientation does not deviate from the mocap estimation significantly. After the alignment, the drifts in the camera pose are mostly removed, and then we can perform the keypoint matching and the camera tracking (Eq. 1).

**3.2.3 Keyframe Decision.** This module also determines whether the current frame becomes a *keyframe*. Keyframes are a set of representative frames with minimal redundancy on visual and mocap information, and are mainly used for mapping and loop closing. This is a key to achieving real-time performance [Klein and Murray 2007]. We select the keyframes following [Campos et al. 2021], where multiple criteria are considered, such as sufficient time interval between keyframes, sufficient matched keypoints, and enough visual changes. We also include the mocap-derived camera poses  $\tilde{\mathbf{R}}, \tilde{\mathbf{t}}$  in the keyframes, as they provide mocap priors for the global optimizations detailed later.

### 3.3 Mapping and Loop Closing

In this section, we involve mocap information in the global optimization schemes in traditional SLAM, including the bundle adjustment (BA) scheme and the pose graph optimization in loop closure, which leads to more accurate mapping and localization results. Specifically, we integrate the mocap prior in the BA algorithm, where the map and keyframe poses<sup>3</sup> are optimized by considering both 3D reprojection errors and mocap constraints (Sec. 3.3.2). In BA, we calculate mocap-related map point confidences to dynamically match the weight of the projection error of each map point with the mocap constraints (Sec. 3.3.1). Besides, when a loop closure is detected (*i.e.*, the camera returns to a visited place), we perform a *pose graph optimization* to close the loop, where the mocap prior is also incorporated, which constrains the relative pose between each keyframe from the view of mocap (Sec. 3.3.3).

**3.3.1 Map Point Confidence Calculation.** In order to achieve accurate results in BA, it is essential to consider the relative weights of mocap constraints in comparison to traditional reprojection terms. To make a delicate design of the weights, we no longer treat all the map points equally but calculate their individual confidence to match their relative weights to the mocap constraints. We consider

<sup>2</sup>Relocalization refers to the process of re-estimating the camera pose in the built map when the visual tracking is lost. See [Mur-Artal et al. 2015].

<sup>3</sup>For brevity, we use *keyframe pose* to denote the camera pose (position and orientation) at the keyframe throughout the paper.




Fig. 3. Explanation on the mocap-related map point confidence. (A) The confidence is computed from the baseline length and the parallax of the keyframes that observe the point. (B) With the same baseline length  $b_1 = b_2$  and the same perturbation on the camera positions (denoted as the transparent cameras), a larger parallax (blue cameras) is often less prone to errors compared with a smaller parallax (red cameras) (*i.e.*, point error  $e_1 < e_2$ ).




Fig. 4. Factor graph for the mocap-aware bundle adjustment. Circles denote optimizable/fixed variables (*i.e.*, map points and keyframe poses). Squares denote constraints among the connected variables (*i.e.*, reprojection residual in blue for the observation constraints, mocap residual in green for the orientation constraints and pink for the position constraints).

a 3D point as accurate if 1) the frames that observe it are located in largely different positions in the real world and 2) the viewing direction in which the frames look at the point span a sufficiently large angle. This is inspired by the multi-view 3D reconstruction, where a larger baseline and parallax lead to accurate estimation (see Fig. 3).

To be specific, for each map point  $i$  that will participate in the BA optimization, we first calculate the baseline length  $b_i$  in the mocap coordinate system and the parallax angle  $\theta_i$  of multiple keyframes in which the point  $i$  is visible. Then, the map point confidence  $c_i$  is calculated as the scaled product of the baseline length and the parallax:

$$c_i = kb_i\theta_i, \quad (3)$$

where  $k$  is experimentally set to 50 to make the confidence near 1 for most of the points. The transformation to the mocap coordinate system is necessary because the confidence should be invariant to the scale factor of SLAM. This is a requirement to combine mocap and SLAM in BA.

**3.3.2 Mocap-aware Bundle Adjustment.** With the map point confidence, we perform a mocap-aware BA to refine the recent keyframe poses and the visible map. The variables to be optimized include the poses of the recent keyframes and the positions of the map points

observed by these keyframes. Other keyframes that also observe the points also participate in the BA, but remain fixed. We denote the optimizable and fixed set of keyframe poses as  $\mathcal{K}_o$  and  $\mathcal{K}_f$  respectively, and the observed map points of keyframe  $j$  as  $\mathcal{X}_j$ . Let  $\mathcal{R}_o = \{R_j | j \in \mathcal{K}_o\}$  and  $\mathcal{T}_o = \{t_j | j \in \mathcal{K}_o\}$  denote the optimizable keyframe poses *w.r.t* the world, and  $\mathcal{X}^{3D} = \bigcup_{j \in \mathcal{K}_o} \{x_i^{3D} | i \in \mathcal{X}_j\}$  denotes the map point positions in the world. The mocap-aware BA is then defined as:

$$\begin{aligned} & \arg \min_{\mathcal{R}_o, \mathcal{T}_o, \mathcal{X}^{3D}} \mathcal{E}_{\text{proj}}(\mathcal{R}_o, \mathcal{T}_o, \mathcal{X}^{3D}) + \mathcal{E}_{\text{mocap}}^{(R)}(\mathcal{R}_o) + \mathcal{E}_{\text{mocap}}^{(t)}(\mathcal{T}_o) \\ \mathcal{E}_{\text{proj}} &= \sum_{j \in \mathcal{K}_o} \sum_{i \in \mathcal{X}_j} \rho \left( c_i \left\| x_i^{2D} - \pi \left( R_j^T (x_i^{3D} - t_j) \right) \right\|_{\Sigma_i}^2 \right) \\ \mathcal{E}_{\text{mocap}}^{(R)} &= \mu_R \sum_{j \in \mathcal{K}_o} \left\| \text{Log}(\tilde{R}_j^T R_j) \right\|^2 \\ \mathcal{E}_{\text{mocap}}^{(t)} &= \mu_t \sum_{j \in \mathcal{K}_o} \left\| (\tilde{t}_j - \tilde{t}_{\text{prev}(j)}) - (t_j - t_{\text{prev}(j)}) \right\|^2, \end{aligned} \quad (4)$$

where  $\text{prev}(j)$  denotes the previous keyframe of keyframe  $j$ ,  $\tilde{R}$ ,  $\tilde{t}$  are the initial camera pose obtained from the mocap, and the coefficients of the mocap terms are set to  $\mu_R = 0.01f^2$  and  $\mu_t = 0.05f^2s^2$  to align the unit to pixels and make them scale-invariant. The mocap-aware BA optimization is easier to understand through the factor graph as shown in Fig. 4. The reprojection term  $\mathcal{E}_{\text{proj}}$  in Eq. 4 corresponds to the blue nodes in the factor graph, which computes the reprojection error using the keyframe poses and the map points (shown as the adjacent variables connected to the node). The mocap terms involve a rotation term  $\mathcal{E}_{\text{mocap}}^{(R)}$  (green nodes) that constrains the keyframe rotations to be near to the mocap results, and a translation term  $\mathcal{E}_{\text{mocap}}^{(t)}$  (pink nodes) that constrains the translations between adjacent keyframes to be similar to the mocap estimation. We apply *absolute* constraints on the keyframe rotations but *relative* constraints on the keyframe translations because the translation estimation suffers from drifts over time while the rotation estimation is relatively stable in the inertial mocap.

Note that the map point confidence  $c_i$  is used as the weights of the reprojection error, where potentially good points have larger weights compared with the mocap priors, which help to refine the keyframe poses, while probably inaccurate points (*e.g.*, recently triangulated points) have smaller weights compared with the mocap priors, which can be refined in the optimization and have a lower impact in estimating the keyframe poses. With the help of the mocap priors and the map point confidence, BA becomes more stable, and both the mapping and keyframe pose accuracy can be improved.

We solve the nonlinear problem using graph optimization implemented in g2o [Kümmerle et al. 2011]. To accelerate the algorithm, the map points are marginalized as in [Campos et al. 2021], and we calculate the Jacobian of the keyframe poses and map points analytically, which is presented in App. D. As our system targets at *online* localization, the optimized keyframe poses after the BA are not directly used in the real-time tracking system due to the delay. Specifically, it affects the real-time tracking implicitly: the mapping module updates the keyframe camera poses and the map points in parallel with the real-time modules in another thread, and

the real-time camera tracking module tracks the camera pose *w.r.t* the *updated* map. Thus, our real-time results can be improved.

**3.3.3 Mocap-aware Pose Graph Optimization.** When a loop is detected, we perform a pose graph optimization to close the loop. This optimization involves only keyframe poses but no map points, which aims at closing the loop by distributing the aligning error to all keyframes. Based on ORB-SLAM3 [Campos et al. 2021], we perform the optimization on the *essential graph* over the keyframe poses expressed as rigid transformations in SE(3) (rather than similarity transformations in Sim(3)) as we do not suffer from the scale drift due to the combination with mocap. We denote the keyframes (vertices) and connections (edges) of the essential graph as  $\mathcal{F}$  and  $C$  respectively. The pose graph optimization is then defined as:

$$\begin{aligned} & \arg \min_{\{T_j\}_{j \in \mathcal{F}}} \mathcal{E}_{\text{graph}} + \mathcal{E}_{\text{mocap}} \\ \mathcal{E}_{\text{graph}} &= \sum_{(i,j) \in C} \left\| \text{Log}(T_{ij} T_j^{-1} T_i) \right\|^2 \\ \mathcal{E}_{\text{mocap}} &= \omega_{\text{pose}} \sum_{j \in \mathcal{F}} \left\| \text{Log}(\tilde{T}_{\text{prev}(j)}^{-1} \tilde{T}_j T_j^{-1} T_{\text{prev}(j)}) \right\|^2, \end{aligned} \quad (5)$$

where  $T_j \in \text{SE}(3)$  is the pose of keyframe  $j$  *w.r.t* the world,  $T_{ij} \in \text{SE}(3)$  is the relative transformation between keyframe  $i$  and  $j$  before the optimization,  $\tilde{T}_j$  is the transformation obtained by mocap,  $\text{Log} : \text{SE}(3) \rightarrow \mathbb{R}^6$  maps a rigid transformation to the vector space, and  $\omega_{\text{pose}}$  is the coefficient of the mocap term, which we empirically set to 0.2. Intuitively, the graph term  $\mathcal{E}_{\text{graph}}$  constrains the change in the relative pose between the connected keyframes to be small before and after the optimization, while the mocap term regularizes the keyframe 6-DoF trajectory to be similar to the estimation of the mocap in a local manner.

### 3.4 Body Translation Updater

The camera tracking module considers both visual and inertial information, so the output 30-FPS camera localization should be able to refine the 60-FPS human motion from the inertial motion capture module. To make the refinement fully constrained, we only update the human translation and implement the module using a prediction-correction algorithm in Kalman Filter.

Specifically, we define the global position and velocity of the human as the state variables  $(\mathbf{p}, \mathbf{v})$ , which will be refined in this module. Using the global acceleration  $\mathbf{a}$  extracted from the mocap module, the state variables can be predicted by the following prediction equation:

$$\begin{aligned} \mathbf{p}_k &= \mathbf{p}_{k-1} + \mathbf{v}_{k-1} \Delta t \\ \mathbf{v}_k &= \mathbf{v}_{k-1} + \mathbf{a}_{k-1} \Delta t + \boldsymbol{\varphi}_{k-1}, \end{aligned} \quad (6)$$

where the subscript  $k$  indicates the  $k$ th frame,  $\Delta t = 1/60$  is the time interval, and  $\boldsymbol{\varphi} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$  models the error in the mocap prediction, where we experimentally set  $\sigma = \Delta t$  based on the assumption that the mocap-estimated acceleration follows a normal distribution with a variance of 1. With the optimized camera position  $\mathbf{p}_{\text{cam}}$  and its confidence  $n$  from the camera tracking module, the state can also be corrected in 30 FPS based on the following correction equation:

$$\mathbf{p}_k^{\text{cam}} = \mathbf{p}_k + \mathbf{p}_k^{\text{root} \rightarrow \text{cam}} + \boldsymbol{\psi}_k, \quad (7)$$

where  $\mathbf{p}_{\text{root} \rightarrow \text{cam}}$  is the position difference between the camera and the root, which can be computed by the estimated human pose using forward kinematics, and  $\boldsymbol{\psi} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\text{cam}})$  models the noise in the camera tracking, where the covariance matrix  $\Sigma_{\text{cam}}$  is computed from the camera position confidence  $n$  as:

$$\Sigma_{\text{cam}} = \frac{1000}{n + \epsilon} \mathbf{I}, \quad (8)$$

where  $\mathbf{I} \in \mathbb{R}^{3 \times 3}$  is the identity matrix and  $\epsilon = 10^{-3}$  is used to avoid division by zero. Given the prediction and correction equations (Eq. 6 and 7), we predict the state variables based on the prediction-correction algorithm (details are presented in App. F).

### 3.5 System Calibration

Here we discuss how to initialize the inertial mocap module and the SLAM module respectively, as well as how to calibrate the coordinate system of the two modules. Inertial mocap needs a T-pose calibration before capturing the human motion, where the sensor-to-bone rotations and the relative rotation between the IMU inertial frame and the user-defined global frame are determined [Huang et al. 2018; Yi et al. 2022, 2021]. The SLAM module needs a monocular initialization, where the initial map points are triangulated by computing the relative pose between two frames based on the homography or the fundamental matrix [Mur-Artal et al. 2015].

In addition to these two initialization steps, we also need to calibrate the extrinsic between the mocap coordinate frame and the SLAM coordinate frame. We propose an easy-to-perform method, where we require the user to walk in a curve for several seconds, and then the extrinsic can be automatically computed. Specifically, we express the mocap-to-SLAM extrinsic as a similarity transformation in Sim(3) as the real scale is unknown for monocular SLAM. During the curve motion, we first independently reconstruct the trajectory of the root by mocap and the camera by SLAM. Then, given the body poses and fixing the camera on the head, we can align the two trajectories and get the similarity transformation between the mocap coordinate frame and the SLAM coordinate frame.

## 4 EXPERIMENTS

In this section, we first present the implementation details including the setup of our live system and the datasets used in our experiments (Sec. 4.1). Then, we compare our method with the state-of-the-art techniques in both the research fields of motion capture (mocap) with sparse inertial sensors and SLAM (Sec. 4.2). Next, we evaluate the key designs and components in our method (Sec. 4.3). Finally, we discuss our limitations (Sec. 4.4). More results can be found in our supplemental video.

### 4.1 Implementation Details

**4.1.1 System Setup.** The system receives 60Hz IMU signals from 6 Xsens Dot [Xsens [n. d.]] IMUs via Bluetooth and 30Hz RGB images in the resolution of  $640 \times 480$  from a smartphone camera via Wi-Fi. Then, the two signals are synchronized by a user's jumping motion, which is easy to be detected from the two signals. Our method runs on a laptop with Intel(R) Core(TM) i7-12700H CPU. Note that we do not need a specific graphic card as our method runs purely on CPU. The whole system runs in real time at about 60 FPS. We implement

Table 1. Overview of the datasets. The table shows whether the datasets contain pose, translation, IMU measurements, videos from first-person cameras, and scene meshes/point clouds. "Y/N" indicates the dataset does/does not contain such data. "S" means such data is synthetic.

Datasets	Pose	Translation	IMU	FP Camera	Scene	Minutes
AMASS	Y	Y	S	N	N	1217
DIP-IMU	Y	N	Y	N	N	80
TotalCapture	Y	Y	Y	S	S	49
HPS	Y*	Y*	Y	Y	Y	180

\*We use the results of the HPS [Guzov et al. 2021] method as ground truths.

our main method using Pytorch [Pytorch [n. d.]] and Rigid Body Dynamic Library (RBDL) [Felis 2017] in python. The SLAM module is implemented in C++ to achieve real-time performance and is compiled as a shared object library that can be dynamically linked by our main method.

**4.1.2 Datasets.** The datasets used in the experiments include DIP-IMU [Huang et al. 2018], AMASS [Mahmood et al. 2019], TotalCapture [Trumble et al. 2017], and HPS [Guzov et al. 2021]. AMASS and DIP-IMU datasets are used to train our mocap networks and we follow [Yi et al. 2022, 2021] to use them for training and fine-tuning respectively. TotalCapture and HPS datasets are used to evaluate our technique. Since we combine sparse inertial mocap and SLAM, we require synchronized and calibrated visual and inertial data to run our technique. Thus, for TotalCapture, we synthesize visual data by constructing virtual scenes and putting virtual cameras on a virtual character driven by the mocap data. For HPS, we calibrate the inertial and visual signals by the first short clip of each motion sequence and discard those sequences that cannot be calibrated. Details about the dataset preprocessing can be found in App. A. An overview of the training and evaluation datasets is shown in Tab. 1.

## 4.2 Comparisons

Our work simultaneously captures human motion and performs localization and mapping from sparse IMUs and a wearable camera. Since there is no existing technique that solves these two tasks simultaneously, we compare our method with the state-of-the-art methods of both fields, *i.e.*, TransPose [Yi et al. 2021], TIP [Jiang et al. 2022b], and PIP [Yi et al. 2022] for mocap from sparse IMUs, ORB-SLAM3 [Campos et al. 2021] for monocular/monocular-inertial SLAM.

**4.2.1 Comparisons with Sparse Inertial Mocap.** We first quantitatively compare the methods using *absolute root position error*, *i.e.*, the average root position error of all frames, on TotalCapture [Trumble et al. 2017] and HPS [Guzov et al. 2021] datasets, with the root's position and orientation aligned with the ground truth at the first frame. The results are shown in Tab. 2, where we achieve as large as 41% and 38% improvement on TotalCapture and HPS datasets respectively over the state-of-the-art methods. Our method significantly outperforms previous works in different scenes including small/large indoor/outdoor environments and for different human motions. It demonstrates that our method successfully and robustly leverages the camera to reduce the translation drift in inertial mocap.

One special case is the "rom" motions in the TotalCapture dataset, which contain very few global movements of humans, making it difficult to correctly initialize the SLAM system and determine the scale factor in our method, and such errors reflect in the results. Another special case is the "EG" scene in the HPS dataset, where TIP outperforms our method on the root position error. However, TIP still suffers from large translation drifts most of the time, which is reflected in other sequences. While our method is more robust and reduces 43% errors on average on the HPS dataset compared with TIP. Note that since all the methods assume a mean body shape, we rescale the estimated translation with the ground-truth leg lengths when evaluating the position error. Also note that the three pure mocap algorithms (TransPose, TIP, and PIP) are *deterministic*, *i.e.*, there is no randomness in the results. While our method is *stochastic* (due to the RANSAC algorithm and multi-threading in the SLAM part). Thus, following previous SLAM works, we test our method 9 times for each sequence and report the median error as well as the standard deviation.

We also present qualitative mocap comparison results in Fig. 5. Due to the error accumulation in inertial mocap, previous approaches suffer from translation drifts. Thus, as shown in the figure, they estimate wrong human positions in the scene as time goes by. By incorporating SLAM in the human mocap, our method largely reduces the translation drifts during human motion, and the global position of the human is the most accurate. Note that the beginning human global position and orientation are aligned with the ground truth in all the methods, and the results are selected after a period of human motion (shown as the trajectory lines in the figure). All methods are not aware of the ground-truth scene.

Regarding the comparisons on human pose estimation, we would like to note that involving SLAM in mocap majorly contributes to determining the global positions, rather than the human poses. Thus, we present the pose estimation results in App. B, where we achieve similar pose estimation accuracy compared with the state-of-the-art mocap methods.

**4.2.2 Comparisons with SLAM.** We compare our method with ORB-SLAM3 [Campos et al. 2021], which is a state-of-the-art technique in visual and visual-inertial SLAM. We run two modes in ORB-SLAM3: *monocular (M)* mode which takes RGB stream as input, and *monocular-inertial (MI)* mode which uses synchronized RGB images and IMU measurements. In each mode, we output both the *online (On)* localization results from the SLAM front-end at running frame rate, and the *offline (Off)* localization results from the SLAM back-end after running the whole sequence. The offline results are usually better than the online results as they have been optimized in the bundle adjustment in the back end, but they cannot be obtained online in real time. For our system, we always calculate real-time camera positions online from the estimated human pose and translation.

**Localization Comparisons.** We quantitatively compare different solutions on TotalCapture [Trumble et al. 2017] and HPS [Guzov et al. 2021] datasets by *absolute camera position error*. For a fair comparison, we should calculate errors on the *full* sequences for all methods. However, due to fast human motions, the SLAM system may often fail to track the camera, which will lead to very large



Table 2. Comparisons with sparse inertial mocap approaches TransPose [Yi et al. 2021], TIP [Jiang et al. 2022b], and PIP [Yi et al. 2022] on TotalCapture [Trumble et al. 2017] (categorized by motions) and HPS [Guzov et al. 2021] (categorized by scenes) datasets. The reported numbers are absolute root position errors averaged over the frames in meters. For our method, we test it 9 times and report the median error (the top number) and the standard deviation (the bottom number, beginning with "±").

Method	TotalCapture					HPS								
	acting	freestyle	rom	walking	average	BIB_AB	BIB_EG	EG	Etag6	GEB	BIB_UG	KINO	BIB_OG	average
TransPose	0.53	0.69	0.18	0.45	0.42	6.83	5.53	2.40	2.77	1.91	4.99	2.20	4.58	4.11
TIP	0.43	0.87	0.21	0.49	0.45	2.23	3.41	<b>1.43</b>	3.87	1.38	2.92	0.89	3.89	3.00
PIP	0.61	0.51	<b>0.07</b>	0.49	0.37	1.26	2.59	1.89	1.78	1.35	2.49	1.50	4.81	2.75
Ours	<b>0.28</b>	<b>0.33</b>	0.10	<b>0.25</b>	<b>0.22</b>	<b>1.23</b>	<b>1.54</b>	1.83	<b>1.35</b>	<b>1.17</b>	<b>2.40</b>	<b>0.87</b>	<b>1.90</b>	<b>1.70</b>
	±0.06	±0.06	±0.02	±0.03	±0.04	±0.37	±0.28	±0.37	±0.18	±0.29	±0.49	±0.16	±0.41	±0.34

Table 3. Comparisons with ORB-SLAM3 [Campos et al. 2021] in [M]onocular/[M]onocular-[I]nertial [Off]line/[On]line modes. We calculate camera localization errors in meters and we present the numbers on both the [full] sequences and the [tracked] frames (excluding the tracking-failed frames) for the compared methods. We test all methods 9 times on each sequence and report the median error (the top number), the standard deviation (beginning with "±"), and the percentage of the tracked frames (ending with "%"). If a method fails too many times (>85%) during evaluation, we mark it as a failure (denoted as "-").

Method	TotalCapture					HPS									
	acting	freestyle	rom	walking	average	BIB_AB	BIB_EG	EG	Etag6	GEB	BIB_UG	KINO	BIB_OG	average	
Full	M-Off	0.82	0.89	0.25	0.42	0.54	8.58	12.57	8.87	5.62	1.62	9.29	4.79	7.61	8.18
		±0.44	±0.17	±0.16	±0.46	±0.29	±0.92	±2.30	±2.23	±1.62	±0.15	±1.52	±0.52	±2.06	±1.71
	M-On	1.34	1.01	0.25	0.80	0.76	8.59	13.94	8.95	9.04	1.84	9.91	5.78	7.77	9.38
		±0.47	±0.17	±0.16	±0.48	±0.30	±0.92	±2.02	±1.20	±1.76	±0.24	±1.43	±0.67	±1.19	±1.41
	MI-Off	10.54	4.75	-	1.08	4.87	-	78.32	-	-	-	7.23	-	-	67.64
		±5.48	±2.62	-	±1.88	±3.24	-	±136.05	-	-	-	±5.98	-	-	±116.52
Tracked	M-Off	0.35	<b>0.32</b>	0.23	0.36	0.30	2.96	5.65	6.67	4.43	<b>0.67</b>	4.20	0.88	8.73	5.19
		±0.50	±0.20	±0.28	±0.36	±0.34	±2.48	±3.50	±3.11	±1.62	±0.62	±5.57	±1.03	±3.04	±2.75
	M-On	85.9%	58.9%	85.0%	89.7%	80.7%	51.8%	62.1%	37.3%	61.1%	55.5%	22.5%	61.2%	40.4%	46.5%
		0.92	0.52	0.24	0.52	0.50	3.88	9.44	6.41	7.59	1.09	3.46	2.04	10.33	7.15
	M-On	±0.44	±0.29	±0.25	±0.54	±0.37	±1.33	±3.35	±3.22	±2.48	±0.94	±5.12	±1.34	±2.08	±2.62
		83.1%	62.3%	83.3%	87.6%	79.7%	51.2%	59.1%	43.7%	58.1%	55.6%	18.1%	62.3%	39.5%	45.2%
Ours	MI-Off	0.32	0.39	-	0.58	0.53	-	77.84	-	-	-	7.15	-	-	66.98
		±0.07	±0.11	-	±0.60	±0.48	-	±132.20	-	-	-	±5.73	-	-	±112.77
	M-Off	14.5%	5%	0%	59.7%	17.1%	0%	88.6%	0%	0%	0%	15.9%	0%	0%	18.7%
	M-On	0.41	0.37	-	0.59	0.54	-	78.35	-	-	-	7.23	-	-	67.44
		±0.08	±0.11	-	±0.62	±0.49	-	±132.23	-	-	-	±5.97	-	-	±112.84
		12.1%	7.6%	0%	57.7%	16.7%	0%	88.6%	0%	0%	0%	15.9%	0%	0%	18.7%
Ours	<b>0.29</b>	0.35	<b>0.13</b>	<b>0.25</b>	<b>0.24</b>	<b>1.25</b>	<b>1.53</b>	<b>1.81</b>	<b>1.34</b>	1.18	<b>2.39</b>	<b>0.86</b>	<b>1.90</b>	<b>1.69</b>	
	±0.06	±0.06	±0.02	±0.04	±0.04	±0.37	±0.28	±0.36	±0.18	±0.29	±0.49	±0.16	±0.41	±0.33	

Table 4. Comparisons on mapping accuracy with [M]onocular/[M]onocular-[I]nertial ORB-SLAM3 [Campos et al. 2021]. For each of the three scenes, we synthesize monocular videos based on the human motions in TotalCapture [Trumble et al. 2017] dataset. We test each method 9 times and report the median map point errors (the top number) and the standard deviations (the bottom number, beginning with "±") in meters. If a method fails too many times (>85%) during evaluation, *i.e.*, no map point is reconstructed, we mark it as a failure (denoted as "-").

Method	Japan Office					Flooded Grounds					SciFi Warehouse				
	acting	freestyle	rom	walking	average	acting	freestyle	rom	walking	average	acting	freestyle	rom	walking	average
SLAM (M)	0.88	1.10	1.20	0.60	0.95	2.89	2.06	3.10	1.36	2.31	0.40	0.43	<b>0.44</b>	0.35	0.41
	±0.26	±0.76	±1.69	±0.30	±0.76	±1.52	±1.30	±1.65	±1.31	±1.44	±0.03	±0.13	±0.05	±0.02	±0.06
SLAM (MI)	0.66	0.54	-	0.61	0.61	-	-	-	0.90	-	0.49	0.40	-	0.51	0.48
	±0.83	±1.50	-	±0.49	±0.79	-	-	-	±0.77	-	±0.10	±0.13	-	±0.41	±0.28
Ours	<b>0.32</b>	<b>0.49</b>	<b>0.72</b>	<b>0.24</b>	<b>0.45</b>	<b>0.94</b>	<b>1.49</b>	<b>1.75</b>	<b>0.80</b>	<b>1.23</b>	<b>0.24</b>	<b>0.37</b>	<b>0.44</b>	<b>0.18</b>	<b>0.31</b>
	±0.09	±0.09	±0.23	±0.06	±0.12	±0.31	±0.70	±1.13	±0.19	±0.56	±0.03	±0.10	±0.17	±0.02	±0.08

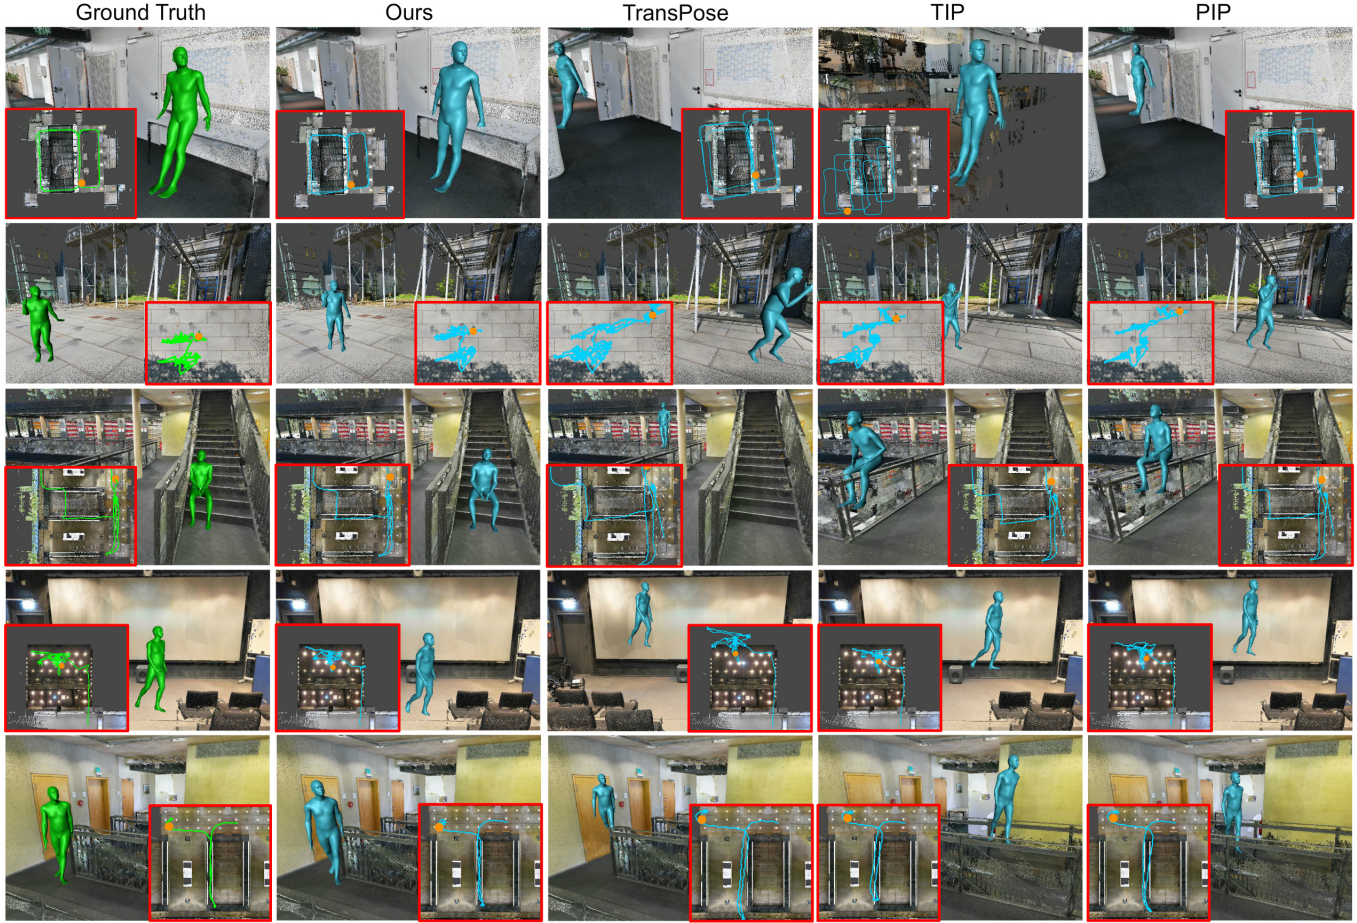


Fig. 5. Qualitative comparisons on human motion capture with TransPose [Yi et al. 2021], TIP [Jiang et al. 2022b], and PIP [Yi et al. 2022] on HPS [Guzov et al. 2021] dataset. We show the ground-truth humans in green and the estimated humans in blue. The movements of the human are also plotted in the top view. The orange point in the top view denotes the current global position of the human.

errors. To this end, we also report the errors calculated on the successfully *tracked* frames for the SLAM solutions, while for our method we consistently report the errors on full frames as we do not suffer the failure.

The comparison results are shown in Tab. 3. Among different settings, online monocular-inertial SLAM (MI-on) is the closest to our method, as both methods leverage inertial sensors and run online. Our method consistently outperforms MI-on on both full sequences and tracked frames, achieving 82% and 56% improvement respectively. Offline monocular SLAM (M-Off) occasionally achieves lower errors on tracked frames than our method. However, it runs offline, and up to 36% human motions fail to be tracked. Furthermore, our real-time method still outperforms it on average on both datasets. Overall, our method achieves lower localization errors and lower standard deviations on both datasets, which reflects that the combination of mocap and SLAM greatly improves both the accuracy and robustness of SLAM systems. Note that for monocular SLAM (M-Off, M-On), we calculate the scale of the estimated trajectory by

the ground truth as monocular SLAM has an inherent ambiguity here. Our method and MI do not suffer this ambiguity.

We would like to list two further advantages of our system over monocular/monocular-inertial ORB-SLAM3. 1) *Less sensitive to calibration errors*. Readers may find that monocular-inertial ORB-SLAM3 performs well on TotalCapture but nearly fails on HPS. This is mainly because HPS contains more calibration errors on camera-IMU extrinsic and the frequency of IMU signals is lower (see dataset preprocessing in App. A). Nevertheless, our method performs well on both datasets with the same extrinsic calibration. 2) *No failures*. Pure SLAM systems often fail to track the camera, which is reflected in the low percentage of tracked frames. This could happen when the human moves quickly or just turn around. By combining mocap and SLAM, our system always gives an estimation based on the mocap result even when the visual tracking is lost.

*Mapping Comparisons*. In addition to localization comparisons, we also compare the mapping accuracy with ORB-SLAM3 [Campos et al. 2021] by extracting the reconstructed 3D map points at the




Fig. 6. Qualitative mapping comparisons with [M]onocular/[M]onocular-[I]nertial ORB-SLAM3 [Campos et al. 2021]. We show the reconstructed map points in a walking sequence for different methods. Color encodes the map point error. For monocular SLAM, we calculate the map scale from the ground truth.

end of each sequence. We evaluate *map point error*, *i.e.*, the average distance between each reconstructed map point and the nearest scene point, on three different scenes<sup>4</sup>. The quantitative evaluation results are shown in Tab. 4. Our method consistently outperforms monocular/monocular-inertial ORB-SLAM3 in different scenes and for different human motions, demonstrating that the deep fusion of mocap and SLAM not only helps with human localization but also enhances the mapping accuracy largely. We also qualitatively show the mapping results of the three scenes in Fig. 6. Our method accurately reconstructs sparse environment points with much fewer noises compared with the SLAM methods. It also robustly handles unconstrained outdoor large environments ("Flooded Grounds"). This is attributed to the close fusion of inertial mocap and SLAM, which involves mocap priors in the mapping. Readers may notice that the reconstructed points are sparser in "Flooded Grounds" and has a larger error than the indoor scenes. This is because the outdoor environment has fewer visual feature points, and reconstructing

<sup>4</sup>The three scenes include Japan Office (<http://aec.unity3d.jp/>), Flooded Grounds (<https://assetstore.unity.com/packages/3d/environments/flooded-grounds-48529>), and SciFi Warehouse (<https://assetstore.unity.com/packages/3d/environments/sci-fi/sci-fi-construction-kit-modular-159280>). All the scenes are free assets and can be accessed from Unity Asset Store.

Table 5. Run-time performance on an Intel(R) Core(TM) i7-12700H CPU on a laptop. We show the time costs in milliseconds when processing one frame.

sync signals	motion capture	cam tracking	tran update	mapping
4.7	5.3	12.8	0.2	97.1

a far point of a large scene is more difficult and prone to errors in multi-view 3D reconstruction. Besides, the mapping errors on "rom" motions are larger compared with other human motions. This is because the "rom" motions contain very few global movements, making it difficult to reconstruct the map due to the high similarity of camera views.

### 4.3 Evaluations

**4.3.1 Performance.** We first evaluate the run-time performance of our method and the results are shown in Tab. 5. Note that our system receives IMU signals in 60 FPS and color images in 30 FPS. To process a frame with inertial signals, we need to synchronize multiple IMU signals, perform inertial mocap, and update the translation, which takes about 10.2 milliseconds. To process a frame with both inertial

Table 6. Ablation study on the methods of combining mocap and SLAM. We report the median root position errors and the standard deviations in meters in 9 tests on the TotalCapture [Trumble et al. 2017] dataset.

Method	TotalCapture				
	acting	freestyle	rom	walking	average
inertial tracking	0.72 ±0	0.64 ±0	0.12 ±0	0.57 ±0	0.46 ±0
monocular SLAM	1.34 ±0.47	1.01 ±0.17	0.25 ±0.16	0.80 ±0.48	0.76 ±0.30
fuse as init	1.90 ±0.36	0.74 ±0.17	0.13 ±0.03	0.98 ±0.23	0.82 ±0.17
fused tracking	1.13 ±0.35	0.66 ±0.07	0.12 ±0.01	0.92 ±0.26	0.63 ±0.15
fused tracking&mapping (w/o confidence)	0.29 ±0.13	0.34 ±0.08	<b>0.09</b> ±0.03	0.28 ±0.08	0.23 ±0.07
Ours	<b>0.28</b> ±0.06	<b>0.33</b> ±0.06	0.10 ±0.02	<b>0.25</b> ±0.03	<b>0.22</b> ±0.04

and visual signals, we need additional camera tracking, which takes 23 milliseconds in total. Therefore, in 1 second, we can process 30 pure-inertial frames and 30 visual-inertial frames in total, *i.e.*, our system can leverage all the available input signals and output pose and translation in 60 FPS. While the mapping (and loop closing) takes much more time than the tracking, we run it only on *keyframes* and as a background task through multi-threading to avoid additional time costs. Thus, it does not affect our real-time 60 FPS performance.

**4.3.2 Effectiveness.** We demonstrate the effectiveness of our method by conducting ablative experiments on different combining methods of inertial mocap and SLAM. Specifically, we evaluate 1) *inertial tracking*, the baseline mocap method where we estimate the translation only from 6 IMUs using the inertial motion capture module, which is similar to PIP [Yi et al. 2022] but the flat ground assumption and the force calculation are removed; 2) *monocular SLAM*, the baseline SLAM method where the translation is estimated by tracking the monocular video purely; 3) *fuse as init*, where we initialize the camera pose in monocular SLAM by the mocap-derived values, while removing all the mocap terms in the optimization, and use the body translation updater to get the root translation from the camera. In other words, it is the same as the standard monocular SLAM except that the mocap-derived camera pose is used to initialize the camera tracking optimization, replacing the original initialization scheme based on the constant-speed assumption; 4) *fused tracking*, where we fuse mocap and SLAM in the camera tracking, *i.e.*, using the proposed mocap-constrained camera tracking; 5) *fused tracking&mapping*, where we fuse mocap and SLAM in both camera tracking and mapping/loop closing, which is the same as our full method except that the mocap-related map point confidences are not used in the mocap-aware bundle adjustment during optimization. We evaluate the root position errors on the TotalCapture [Trumble et al. 2017] dataset, and the results are shown in Tab. 6 Due to the translation drifts in pure 1) *inertial tracking* and the frequent tracking failure in pure 2) *monocular SLAM*, the two baselines have relatively large errors (Row 1&2). We then show that a naive fusion of mocap and SLAM by 3) *fuse as init* will not give good results (Row 3). With the 4) *fused tracking*, the system performs slightly better

(Row 4), because the mocap constraints help to reduce the errors caused by keypoint mismatches in the camera tracking. Nevertheless, the shallow fusion of mocap and SLAM in the camera tracking still can not outperform the mocap baseline. This is because, though the camera tracking accuracy *w.r.t the map* is improved, the map itself is often poorly reconstructed (*e.g.*, in a wrong scale) due to the unawareness of mocap. Then, tracking the camera in an inaccurate map will lead to poor localization results, which further degenerate the mapping in a positive-feedback loop. To this end, we propose 5) *fused tracking&mapping* (Row 5), which shows a large enhancement in tracking accuracy and robustness compared with the baselines. We attribute this to the mocap-aware bundle adjustment algorithm, where the map and keyframe camera poses are jointly optimized with mocap prior, which greatly enhances the mapping accuracy. Furthermore, with the proposed mocap-related map point confidence (Row 6), the standard deviation of multiple experiments is reduced by 43%, reflecting that the uncertainty in the tracking is largely reduced. On the other hand, the overall accuracy is further improved slightly. This shows the effectiveness of the deep fusion of mocap and SLAM in both front-end tracking and back-end mapping, as well as the proposed mocap-related map point confidence.

#### 4.4 Limitations

**Online Loop Closure.** In our system, visual information could help to reduce the drift in inertial mocap as we did in Sec. 3.4, but there are still remaining drifts and loop closing is applied to handle them. However, loop closure requires modifying the camera pose of previous frames when a loop is detected, which is not consistent with an online mocap system whose results of previous frames have already been exported and cannot be modified. So, in our online system, even though we can leverage the loop closure component in SLAM, we have to teleport the user to its correct position to reduce the drift rather than generate a smooth trajectory.

**Body Shape and Scene Scale.** By aligning the mocap trajectory and the SLAM trajectory during initialization, the scene scale is consistent with the body shape scale. However, since we assume mean shape rather than measure the true shape scale of the user, our reconstructed scene scale may be slightly different from the real one. This problem could be handled by measuring the bone lengths and scaling the mocap output accordingly.

**Calibration Errors.** To obtain optimal results, data synchronization and system calibration should be carefully performed. In our implementation, we need a jumping motion for sensor synchronization, a still T pose for sensor-to-bone calibration, and a curve movement for global frame alignment. To ensure the calibration is well done, we check the curve trajectory alignment error to determine whether the calibration is successful. Typically, a successful calibration takes about 20 seconds. These steps are required because a bad calibration will lead to large inconsistencies between mocap and SLAM estimations, yielding poor results.

**Degenerated Cases.** Our system is robust to a few dynamic objects by leveraging the human prior (*e.g.*, playing ping-pong in the supplemented video, where the second person is always moving in front of the head-mounted camera), but the system will degenerate into

inertial mocap or become less accurate if the scene is too dynamic or sparsely textured. Besides, in extreme cases where the human shakes the head fast or performs very weird poses, our system may generate wrong scene points as the mocap fails.

## 5 CONCLUSION

This is the first work that combines inertial mocap with SLAM to achieve simultaneously human motion capture (mocap), localization, and mapping in real time. The system is still lightweight as it only requires sparse body-mounted sensors, including 6 inertial measurement units (IMUs) and a monocular phone camera. For on-line tracking, the mocap and SLAM densely exchange information by coupled optimizations and Kalman filter techniques, leading to more robust and accurate localization. For the back-end optimization, both bundle adjustment and loop closure in SLAM benefit from the inertial mocap and the localization errors can be further reduced.

## ACKNOWLEDGMENTS

This work was supported by the National Key R&D Program of China (2018YFA0704000), the NSFC (No.62021002), Beijing Natural Science Foundation (M22024), and the Key Research and Development Project of Tibet Autonomous Region (XZ202101ZY0019G). This work was also supported by THUIBCS, Tsinghua University, and BLBCI, Beijing Municipal Education Commission. This work was partially supported by the ERC consolidator grant 4DRepl (770784). The authors would like to thank Wenbin Lin, Zunjie Zhu, Guofeng Zhang, and Haoyu Hu for their extensive help on the experiments and live demos. The authors would also like to thank Ting Shu, Shuyan Han, and Kelan Liu for their help on this project. Feng Xu is the corresponding author.

## REFERENCES

- Hiroyasu Akada, Jian Wang, Soshi Shimada, Masaki Takahashi, Christian Theobalt, and Vladislav Golyanik. 2022. UnrealEgo: A New Dataset for Robust Egocentric 3D Human Motion Capture. In *European Conference on Computer Vision (ECCV)*.
- Michael Bloesch, Jan Czarowski, Ronald Clark, Stefan Leutenegger, and Andrew J Davison. 2018. CodeSLAM—learning a compact, optimisable representation for dense visual SLAM. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2560–2568.
- Carlos Campos, Richard Elvira, Juan J. Gómez, José M. M. Montiel, and Juan D. Tardós. 2021. ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual-Inertial and Multi-Map SLAM. *IEEE Transactions on Robotics* 37, 6 (2021), 1874–1890.
- Robert Castle, Georg Klein, and David W Murray. 2008. Video-rate localization in multiple maps for wearable augmented reality. In *2008 12th IEEE International Symposium on Wearable Computers*. IEEE, 15–22.
- Young-Woon Cha, Husam Shaik, Qian Zhang, Fan Feng, Andrei State, Adrian Ilie, and Henry Fuchs. 2021. Mobile. Egocentric Human Body Motion Reconstruction Using Only Eyeglasses-mounted Cameras and a Few Body-worn Inertial Sensors. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. 616–625.
- Long Chen, Haizhou Ai, Rui Chen, Zijie Zhuang, and Shuang Liu. 2020. Cross-View Tracking for Multi-Human 3D Pose Estimation at Over 100 FPS. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yudi Dai, Yitai Lin, Chenglu Wen, Siqi Shen, Lan Xu, Jingyi Yu, Yuexin Ma, and Cheng Wang. 2022. HSC4D: Human-Centered 4D Scene Capture in Large-Scale Indoor-Outdoor Space Using Wearable IMUs and LiDAR. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 6792–6802.
- Andrew J Davison. 2003. Real-time simultaneous localisation and mapping with a single camera. In *Computer Vision, IEEE International Conference on*, Vol. 3. IEEE Computer Society, 1403–1403.
- Junting Dong, Wen Jiang, Qixing Huang, Hujun Bao, and Xiaowei Zhou. 2019. Fast and Robust Multi-Person 3D Pose Estimation From Multiple Views. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jakob Engel, Vladlen Koltun, and Daniel Cremers. 2017. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence* 40, 3 (2017), 611–625.
- Jakob Engel, Thomas Schöps, and Daniel Cremers. 2014. LSD-SLAM: Large-scale direct monocular SLAM. In *European conference on computer vision*. Springer, 834–849.
- Martin Felis. 2017. RBDL: an efficient rigid-body dynamics library using recursive algorithms. *Autonomous Robots* 41 (02 2017).
- Christian Forster, Matia Pizzoli, and Davide Scaramuzza. 2014. SVO: Fast semi-direct monocular visual odometry. In *2014 IEEE international conference on robotics and automation (ICRA)*. IEEE, 15–22.
- Jack H. Geissinger and Alan T. Asbeck. 2020. Motion Inference Using Sparse Inertial Sensors, Self-Supervised Learning, and a New Dataset of Unscripted Human Motion. *Sensors* 20 (2020).
- Vladimir Guzov, Aymen Mir, Torsten Sattler, and Gerard Pons-Moll. 2021. Human POSEitioning System (HPS): 3D Human Pose Estimation and Self-Localization in Large Scenes from Body-Mounted Sensors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Vladimir Guzov, Torsten Sattler, and Gerard Pons-Moll. 2022. Visually plausible human-object interaction capture from wearable sensors. In *arXiv*.
- Dorian F. Henning, Tristan Laidlow, and Stefan Leutenegger. 2022. BodySLAM: Joint Camera Localisation, Mapping, and Human Motion Tracking. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIX*. 656–673.
- Ryosuke Hori, Ryo Hachiuma, Mariko Isogawa, Dan Mikami, and Hideo Saito. 2022. Silhouette-Based 3D Human Pose Estimation Using a Single Wrist-Mounted 360° Camera. *IEEE Access* 10 (2022), 54957–54968.
- Yinghao Huang, Manuel Kaufmann, Emre Aksan, Michael J. Black, Otmar Hilliges, and Gerard Pons-Moll. 2018. Deep Inertial Poser Learning to Reconstruct Human Pose from Sparse Inertial Measurements in Real Time. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)* 37 (nov 2018).
- Hao Jiang and Kristen Grauman. 2017. Seeing Invisible Poses: Estimating 3D Body Pose from Egocentric Video. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jiaxi Jiang, Paul Strelj, Huajian Qiu, Andreas Fender, Larissa Laich, Patrick Snape, and Christian Holz. 2022a. Avatarposer: Articulated full-body pose tracking from sparse motion sensing. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V*. Springer, 443–460.
- Yifeng Jiang, Yuting Ye, Deepak Gopinath, Jungdam Won, Alexander W. Winkler, and C. Karen Liu. 2022b. Transformer Inertial Poser: Real-Time Human Motion Reconstruction from Sparse IMUs with Simultaneous Terrain Generation. In *SIGGRAPH Asia 2022 Conference Papers*.
- Manuel Kaufmann, Yi Zhao, Chengcheng Tang, Lingling Tao, Christopher Twigg, Jie Song, Robert Wang, and Otmar Hilliges. 2021. Em-pose: 3d human pose estimation from sparse electromagnetic trackers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 11510–11520.
- Georg Klein and David Murray. 2007. Parallel Tracking and Mapping for Small AR Workspaces. In *2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality*. 225–234. <https://doi.org/10.1109/ISMAR.2007.4538852>
- Lukas Koestler, Nan Yang, Niclas Zeller, and Daniel Cremers. 2022. Tandem: Tracking and dense mapping in real-time using deep multi-view stereo. In *Conference on Robot Learning*. PMLR, 34–45.
- Rainer Kümmerle, Giorgio Grisetti, Hauke Strasdat, Kurt Konolige, and Wolfram Burgard. 2011. G2o: A general framework for graph optimization. In *2011 IEEE International Conference on Robotics and Automation*.
- Stefan Leutenegger, Simon Lynen, Michael Bosse, Roland Siegwart, and Paul Furgale. 2015. Keyframe-based visual-inertial odometry using nonlinear optimization. *The International Journal of Robotics Research* 34, 3 (2015), 314–334.
- Jiaman Li, C Karen Liu, and Jiajun Wu. 2022. Ego-Body Pose Estimation via Ego-Head Pose Estimation. *arXiv preprint arXiv:2212.04636* (2022).
- Miao Liu, Dexin Yang, Yan Zhang, Zhaopeng Cui, James M. Rehg, and Siyu Tang. 2021. 4D Human Body Capture from Egocentric Video via 3D Scene Grounding. In *2021 International Conference on 3D Vision (3DV)*.
- Yuxuan Liu, Jianxin Yang, Xiao Gu, Yao Guo, and Guang-Zhong Yang. 2022. Ego+X: An Egocentric Vision System for Global 3D Human Pose Estimation and Social Interaction Characterization. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 5271–5277.
- Matthew Loper, Nareen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. 2015. SMPL: A Skinned Multi-Person Linear Model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)* 34 (oct 2015).
- Zhengyi Luo, Ryo Hachiuma, Ye Yuan, and Kris Kitani. 2021. Dynamics-Regulated Kinematic Policy for Egocentric Pose Estimation. In *Advances in Neural Information Processing Systems*.
- Nareen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. 2019. AMASS: Archive of Motion Capture as Surface Shapes. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Anastasios I Mourikis, Stergios I Roumeliotis, et al. 2007. A Multi-State Constraint Kalman Filter for Vision-aided Inertial Navigation. In *ICRA*, Vol. 2. 6.

- Raul Mur-Artal, J. M. M. Montiel, and Juan D. Tardos. 2015. ORB-SLAM: A Versatile and Accurate Monocular SLAM System. *IEEE Transactions on Robotics* 31, 5 (oct 2015), 1147–1163.
- Raul Mur-Artal and Juan D Tardós. 2017a. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE transactions on robotics* 33, 5 (2017), 1255–1262.
- Raúl Mur-Artal and Juan D Tardós. 2017b. Visual-inertial monocular SLAM with map reuse. *IEEE Robotics and Automation Letters* 2, 2 (2017), 796–803.
- Patrik Puchert and Timo Ropinski. 2021. Human Pose Estimation from Sparse Inertial Measurements through Recurrent Graph Convolution. *CoRR* abs/2107.11214 (2021). arXiv:2107.11214
- Pytorch. [n. d.]. Pytorch. Website. <https://pytorch.org/>.
- Tong Qin, Peiliang Li, and Shaojie Shen. 2018. Vins-mono: A robust and versatile monocular visual-inertial state estimator. *IEEE Transactions on Robotics* 34, 4 (2018), 1004–1020.
- N Dinesh Reddy, Laurent Guigues, Leonid Pishchulin, Jayan Eledath, and Srinivasa G Narasimhan. 2021. Tesseract: End-to-end learnable multi-person articulated 3d pose tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15190–15200.
- Helge Rhodin, Christian Richardt, Dan Casas, Eldar Insafutdinov, Mohammad Shafiei, Hans-Peter Seidel, Bernt Schiele, and Christian Theobalt. 2016. EgoCap: Egocentric Marker-Less Motion Capture with Two Fisheye Cameras. 35 (2016).
- Qaiser Riaz, Guan hong Tao, Björn Krüger, and Andreas Weber. 2015. Motion Reconstruction Using Very Few Accelerometers and Ground Contacts. *Graph. Models* 79 (may 2015).
- Patrik Schmuck and Margarita Chli. 2017. Multi-uav collaborative monocular slam. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 3863–3870.
- Ruizhi Shao, Zerong Zheng, Hongwen Zhang, Jingxiang Sun, and Yebin Liu. 2022. Diffustereo: High quality human reconstruction via diffusion-based stereo using sparse cameras. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*. Springer, 702–720.
- Takaaki Shiratori, Hyun Soo Park, Leonid Sigal, Yaser Sheikh, and Jessica K. Hodgins. 2011. Motion Capture from Body-Mounted Cameras. *ACM Trans. Graph.* 30, 4, Article 31 (jul 2011), 10 pages.
- Ronit Slycer and Jessica Hodgins. 2008. Action Capture with Accelerometers. *ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (01 2008).
- Jochen Tautges, Arno Zinke, Björn Krüger, Jan Baumann, Andreas Weber, Thomas Helten, Meinard Müller, Hans-Peter Seidel, and Bernhard Eberhardt. 2011. Motion Reconstruction Using Sparse Accelerometer Data. *ACM Transactions on Graphics* 30 (05 2011).
- Zachary Teed and Jia Deng. 2021. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in Neural Information Processing Systems* 34 (2021), 16558–16569.
- Denis Tome, Thiemo Alldieck, Patrick Peluse, Gerard Pons-Moll, Lourdes Agapito, Hernan Badino, and Fernando de la Torre. 2020. SelfPose: 3D Egocentric Pose Estimation from a Headset Mounted Camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (Oct 2020).
- Denis Tome, Patrick Peluse, Lourdes Agapito, and Hernan Badino. 2019. xr-egopose: Egocentric 3d human pose from an hmd camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Matthew Trumble, Andrew Gilbert, Charles Malleon, Adrian Hilton, and John Collo-mosse. 2017. Total Capture: 3D Human Pose Estimation Fusing Video and Inertial Sensors. In *2017 British Machine Vision Conference (BMVC)*.
- Daniel Vlasic, Rolf Adelsberger, Giovanni Vannucci, John Barnwell, Markus Gross, Wojciech Matusik, and Jovan Popović. 2007. Practical Motion Capture in Everyday Surroundings. *ACM Trans. Graph.* 26 (jul 2007).
- Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. 2018. Recovering Accurate 3D Human Pose in The Wild Using IMUs and a Moving Camera. In *European Conference on Computer Vision (ECCV)*.
- Timo von Marcard, Bodo Rosenhahn, Michael Black, and Gerard Pons-Moll. 2017. Sparse Inertial Poser: Automatic 3D Human Pose Estimation from Sparse IMUs. *Computer Graphics Forum* 36(2), *Proceedings of the 38th Annual Conference of the European Association for Computer Graphics (Eurographics)* (2017).
- Lukas Von Stumberg, Vladyslav Usenko, and Daniel Cremers. 2018. Direct sparse visual-inertial odometry using dynamic marginalization. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2510–2517.
- Jian Wang, Lingjie Liu, Weipeng Xu, Kripasindhu Sarkar, and Christian Theobalt. 2021. Estimating Egocentric 3D Human Pose in Global Space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Jian Wang, Diogo Luvizon, Weipeng Xu, Lingjie Liu, Kripasindhu Sarkar, and Christian Theobalt. 2023. Scene-aware Egocentric 3D Human Pose Estimation. *CVPR* (2023).
- Alexander Winkler, Jungdam Won, and Yuting Ye. 2022. QuestSim: Human Motion Tracking from Sparse Sensors with Simulated Avatars. In *SIGGRAPH Asia 2022 Conference Papers*. 1–8.
- Di Xia, Yeqing Zhu, and Heng Zhang. 2022. Faster Deep Inertial Pose Estimation with Six Inertial Sensors. *Sensors* 22, 19 (2022).
- Xsens. [n. d.]. Xsens 3D motion tracking. Website. <https://www.xsens.com/>.
- Weipeng Xu, Avishek Chatterjee, Michael Zollhoefer, Helge Rhodin, Pascal Fua, Hans-Peter Seidel, and Christian Theobalt. 2019. Mo<sup>2</sup>Cap<sup>2</sup> : Real-time Mobile 3D Motion Capture with a Cap-mounted Fisheye Camera. *IEEE Transactions on Visualization and Computer Graphics* (2019).
- Nan Yang, Lukas von Stumberg, Rui Wang, and Daniel Cremers. 2020. D3vo: Deep depth, deep pose and deep uncertainty for monocular visual odometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1281–1292.
- Xinyu Yi, Yuxiao Zhou, Marc Habermann, Soshi Shimada, Vladislav Golyanik, Christian Theobalt, and Feng Xu. 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)*.
- Xinyu Yi, Yuxiao Zhou, and Feng Xu. 2021. TransPose: Real-time 3D Human Translation and Pose Estimation with Six Inertial Sensors. *ACM Transactions on Graphics* 40 (08 2021).
- Ye Yuan, Umar Iqbal, Pavlo Molchanov, Kris Kitani, and Jan Kautz. 2022. GLAMR: Global Occlusion-Aware Human Mesh Recovery with Dynamic Cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ye Yuan and Kris Kitani. 2019. Ego-Pose Estimation and Forecasting As Real-Time PD Control. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Ye Yuan and Kris M. Kitani. 2018. 3D Ego-Pose Estimation via Imitation Learning. In *ECCV*.
- Yuxiang Zhang, Liang An, Tao Yu, Xiu Li, Kun Li, and Yebin Liu. 2020. 4D Association Graph for Realtime Multi-Person Motion Capture Using Multiple Video Cameras. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jon Zubizarreta, Iker Aguinaga, and Jose Maria Martinez Montiel. 2020. Direct sparse mapping. *IEEE Transactions on Robotics* 36, 4 (2020), 1363–1370.

## A DATASET PREPROCESSING

In this section, we explain the details of our preprocessing for the datasets, including AMASS [Mahmood et al. 2019], DIP-IMU [Huang et al. 2018], TotalCapture [Trumble et al. 2017], and HPS [Guzov et al. 2021]. Since these datasets are recorded for different purposes and have different usages in our experiments, we process each dataset differently. 1) *AMASS*. This dataset is used only in training. Thus, we use ground-truth pose and translation, and synthesize 60Hz acceleration and orientation measurements from the human motion, following [Yi et al. 2022, 2021]. 2) *DIP-IMU*. This dataset contains real inertia measurements but no human translation. Thus, we use it in the network fine-tuning step following [Yi et al. 2022, 2021]. 3) *Total-Capture*. This dataset is used in the evaluation and contains character pose, translation, and IMUs, but no first-person videos. Thus, we transfer the ground-truth pose and translation to a SMPL [Loper et al. 2015] human model and put the model in virtual scenes. With a virtual camera on the model’s head, first-person videos are synthesized for this dataset (see Fig. 7). Such videos are used to run our method during evaluation, and the virtual scenes are used to validate the mapping accuracy. The camera-head relative pose is fixed when synthesizing the first-person videos, and is assumed known in the experiments. Besides, visual-inertial SLAM assumes known IMU noise parameters (IMU intrinsic), which we assign experimentally in the values that produce the best results, and known camera-IMU relative pose (camera-IMU extrinsic), which we calculate from T-pose calibration. Note that the intrinsic and extrinsic are only used by visual-inertial SLAM, while our method does not need them. 4) *HPS*. This dataset is recorded in large scenes and provides both real inertial measurements and first-person videos. HPS also estimates character pose and translation by combining dense inertial motion capture and scene-based localization, which are used as the ground truths for our evaluations. To evaluate different methods on the dataset, we need three extrinsics: imu-to-bone rotation,



Fig. 7. We synthesize first-person videos for the TotalCapture dataset. The pictures at the bottom right corner are the synthesized first-person views.

Table 7. Comparisons on pose accuracy with TransPose [Yi et al. 2021], TIP [Jiang et al. 2022b], and PIP [Yi et al. 2022] on TotalCapture [Trumble et al. 2017] dataset. We report the mean per joint position error (MPJPE) in millimeters with the root position aligned.

Method	TotalCapture				
	acting	freestyle	rom	walking	average
TransPose	62.9	89.3	55.6	55.0	63.6
TIP	<b>56.9</b>	<b>79.4</b>	56.0	49.2	59.3
PIP	60.1	82.1	45.7	<b>46.6</b>	56.2
Ours	59.6	82.5	<b>44.9</b>	47.6	<b>56.1</b>

camera-to-head rotation, and camera-to-imu rotation, which are not provided in this dataset. We first use the beginning 150 frames of each sequence to perform IMU-to-bone calibration, which determines imu-to-bone rotations. Then, the camera-to-head rotation is calculated as the mean rotation between the head orientations measured by the IMU and the camera orientations measured by scene-based localization in the first 60 frames. Note that sequences with too much false camera localization in the beginning 60 frames or have evidently wrong IMU-camera synchronization are discarded. Finally, the camera-to-imu rotation is computed by multiplying the camera-to-head rotation and the head-to-imu rotation. To perform visual-inertial SLAM on this dataset, we synthesize the IMU angular velocities and local accelerations from the sensor orientation and free acceleration signals as they are not provided in this dataset. Mathematically, the sensor-local angular velocity  $\omega_S$  and acceleration  $\mathbf{a}_S$  are calculated as:

$$\omega_S(t) = \frac{1}{\Delta t} \text{Log}(\mathbf{R}_{IS}(t)^{-1} \mathbf{R}_{IS}(t + 1)), \quad (9)$$

$$\mathbf{a}_S(t) = \mathbf{R}_{IS}(t)^{-1} (\mathbf{a}_I(t) - \mathbf{g}_I), \quad (10)$$

where  $\mathbf{R}_{IS}(t)$  and  $\mathbf{a}_I(t)$  are the sensor orientation and free acceleration in the global inertial frame at frame  $t$  respectively,  $\Delta t$  is the time interval between two frames,  $\mathbf{g}_I$  is the gravity in the global inertial frame, and the logarithm map  $\text{Log} : \text{SO}(3) \rightarrow \mathbb{R}^3$  maps from the Lie group to the vector space. This results in 30Hz inertia measurements, the highest frequency of signals we can acquire from the provided data. The visual-inertial SLAM also needs the IMU intrinsic, which is assigned experimentally, and the camera-IMU extrinsic, which is assigned with the camera-to-imu rotation and zero relative position.

## B POSE COMPARISONS WITH INERTIAL MOCAP

We conduct pose comparisons with sparse inertial mocap methods TransPose [Yi et al. 2021], TIP [Jiang et al. 2022b], and PIP [Yi et al. 2022] using mean per joint position error (MPJPE) with the root positions aligned with the ground truth. The results are shown in Tab. 7. We achieve comparable pose estimation accuracy with the state-of-the-art mocap methods. We would like to note that introducing SLAM in the inertial mocap majorly helps with improving the translation estimation accuracy rather than the pose in our algorithm.

## C MODIFICATION OVER PIP IN MOCAP

The motion capture module is designed following [Yi et al. 2022]. We use the same neural kinematics estimator but a different motion optimizer as discussed in Sec. 3.1. Following the notations in [Yi et al. 2022], our new motion optimizer is defined as:

$$\begin{aligned} \arg \min_{\tilde{\mathbf{q}}} \quad & \|\tilde{\mathbf{q}}_3 - \tilde{\boldsymbol{\theta}}_{\text{des}}\|^2 + \|\mathbf{J}\tilde{\mathbf{q}} + \dot{\mathbf{J}}\tilde{\mathbf{q}} - \dot{\mathbf{r}}_{\text{des}}\|^2 \\ \text{s.t.} \quad & \dot{\mathbf{r}}_j(\tilde{\mathbf{q}}) \in \mathcal{C}. \end{aligned} \quad (11)$$

The only optimizable variable is the generalized acceleration  $\tilde{\mathbf{q}}$  of the human pose and translation, while we do not consider the forces. The optimization solves the acceleration that best reproduces the kinematically estimated human pose and joint velocity leveraging the dual PD controller. The contact constraints  $\mathcal{C}$  is defined by:

$$\begin{aligned} \mathcal{C}_j &= \{\dot{\mathbf{r}}_j \in \mathbb{R}^3 \mid |\dot{r}_j^x| \leq \sigma, |\dot{r}_j^z| \leq \sigma\}, \\ \mathcal{C} &= \{[\dot{\mathbf{r}}_1 \cdots \dot{\mathbf{r}}_{n_j}] \in \mathbb{R}^{3n_j} \mid \dot{\mathbf{r}}_j \in \mathcal{C}_j, j = 1, 2, \dots, n_j\}, \end{aligned} \quad (12)$$

where we remove the vertical velocity constraints of the original method as we do not assume a known ground (we allow the capture in the free 3D space). Besides, we only consider the kinematically estimated contact probabilities during the contact determination, where a foot with a probability larger than 0.5 is classified as in contact and vice versa. Our new optimization can be efficiently solved using quadratic programming algorithms. Readers are referred to [Yi et al. 2022] for the notations and more details.

## D NONLINEAR OPTIMIZATIONS IN SLAM

To accelerate our algorithm, we compute the partial Jacobians analytically for the key optimizations, including the mocap-constrained camera tracking (Eq. 1) and the mocap-aware bundle adjustment (Eq. 4). We present the details in the following.

*Parameterization.* As we express the camera orientation  $\mathbf{R} \in \text{SO}(3)$  in the over-parameterized matrix representation that lies in a non-Euclidean space, we optimize it by adding a perturbation  $\Delta\boldsymbol{\theta} \in \mathbb{R}^3$  around the current value, which uses a minimal representation for 3D rotations. We define the new "adding" operator  $\oplus$  between a rotation matrix and a perturbation as:

$$\mathbf{R} \oplus \Delta\boldsymbol{\theta} = \mathbf{R} \text{Exp}(\Delta\boldsymbol{\theta}), \quad (13)$$

where the exponential map  $\text{Exp} : \mathbb{R}^3 \rightarrow \text{SO}(3)$  maps from the Lie algebra to the Lie group. Similarly, we redefine the "adding" operator for the camera translation  $\mathbf{t}$  to make the perturbation  $\Delta\mathbf{t} \in \mathbb{R}^3$  acts in the camera-local space:

$$\mathbf{t} \oplus \Delta\mathbf{t} = \mathbf{t} + \mathbf{R}\Delta\mathbf{t}, \quad (14)$$

where  $\mathbf{R}$  is the camera orientation. For the map point positions, we express the increment in the same space of the point (*i.e.*, the world).

*Mocap-constrained Camera Tracking.* According to Eq. 1, without loss of generality, we respectively denote the reprojection error of point  $i$  in keyframe  $j$ , the orientation error of keyframe  $j$ , and the position error of keyframe  $j$  as:

$$\begin{aligned} \mathbf{e}_{ij}^{(\text{proj})} &= \mathbf{x}_i^{2\text{D}} - \pi \left( \mathbf{R}_j^T (\mathbf{x}_i^{3\text{D}} - \mathbf{t}_j) \right) \\ \mathbf{e}_j^{(\text{ori})} &= \text{Log}(\tilde{\mathbf{R}}_j^T \mathbf{R}_j) \\ \mathbf{e}_j^{(\text{pos})} &= \tilde{\mathbf{t}}_j - \mathbf{t}_j. \end{aligned} \quad (15)$$

For brevity, we denote the map point  $i$ 's position in the camera  $j$ 's coordinate frame as  $\mathbf{y}_{ij} = \mathbf{R}_j^T (\mathbf{x}_i^{3\text{D}} - \mathbf{t}_j)$ . Then, the Jacobians of the errors to the pose perturbations can be computed as:

$$\begin{aligned} \begin{pmatrix} \frac{\partial \mathbf{e}_{ij}^{(\text{proj})}}{\partial \Delta \theta_j} & \frac{\partial \mathbf{e}_{ij}^{(\text{proj})}}{\partial \Delta \mathbf{t}_j} \end{pmatrix} &= \begin{pmatrix} -\mathbf{J}_p(\mathbf{y}_{ij}) \mathbf{y}_{ij}^\wedge & \mathbf{J}_p(\mathbf{y}_{ij}) \end{pmatrix} \\ \begin{pmatrix} \frac{\partial \mathbf{e}_j^{(\text{ori})}}{\partial \Delta \theta_j} & \frac{\partial \mathbf{e}_j^{(\text{ori})}}{\partial \Delta \mathbf{t}_j} \end{pmatrix} &= \begin{pmatrix} \mathbf{J}_r \left( \text{Log}(\tilde{\mathbf{R}}_j^T \mathbf{R}_j) \right)^{-1} & \mathbf{O} \end{pmatrix} \\ \begin{pmatrix} \frac{\partial \mathbf{e}_j^{(\text{pos})}}{\partial \Delta \theta_j} & \frac{\partial \mathbf{e}_j^{(\text{pos})}}{\partial \Delta \mathbf{t}_j} \end{pmatrix} &= \begin{pmatrix} \mathbf{O} & -\mathbf{R}_j \end{pmatrix}, \end{aligned} \quad (16)$$

where  $\cdot^\wedge$  converts the 3D vector to the skew symmetric matrix,  $\mathbf{J}_p(\cdot)$  is calculated as:

$$\mathbf{J}_p(\mathbf{y}) = \begin{pmatrix} \frac{f_x}{y_z} & 0 & -\frac{f_x y_x}{y_z^2} \\ 0 & \frac{f_y}{y_z} & -\frac{f_y y_y}{y_z^2} \end{pmatrix}, \quad (17)$$

where we assume  $\mathbf{y} = (y_x, y_y, y_z)^T$  and omit its subscript  $i$  and  $j$ , and  $f_x, f_y$  are the camera focal lengths.  $\mathbf{J}_r(\cdot)^{-1}$  in Eq. 16 is calculated as:

$$\mathbf{J}_r(\theta \mathbf{a})^{-1} = \frac{\theta}{2} \cot \frac{\theta}{2} \mathbf{I} + \left(1 - \frac{\theta}{2} \cot \frac{\theta}{2}\right) \mathbf{a} \mathbf{a}^T + \frac{\theta}{2} \mathbf{a}^\wedge, \quad (18)$$

where we assume  $\text{Log}(\tilde{\mathbf{R}}_j^T \mathbf{R}_j) = \theta \mathbf{a}$  in Eq. 16, where  $\theta \in \mathbb{R}$ ,  $\mathbf{a} \in \mathbb{R}^3$  corresponds to the angle and axis respectively. Note that we omit the subscript  $j$  for brevity.

*Mocap-aware Bundle Adjustment.* The mocap-aware bundle adjustment (Eq. 4) shares the same reprojection error with the camera tracking. The absolute orientation error and the relative position error in the BA are defined as:

$$\begin{aligned} \mathbf{e}_j^{(\text{aori})} &= \text{Log}(\tilde{\mathbf{R}}_j^T \mathbf{R}_j) \\ \mathbf{e}_j^{(\text{rpos})} &= (\tilde{\mathbf{t}}_j - \tilde{\mathbf{t}}_{\text{prev}(j)}) - (\mathbf{t}_j - \mathbf{t}_{\text{prev}(j)}). \end{aligned} \quad (19)$$

The non-zero jacobians of the two errors can be computed as:

$$\begin{aligned} \frac{\partial \mathbf{e}_j^{(\text{aori})}}{\partial \Delta \theta_j} &= \mathbf{J}_r \left( \text{Log}(\tilde{\mathbf{R}}_j^T \mathbf{R}_j) \right)^{-1} \\ \frac{\partial \mathbf{e}_j^{(\text{rpos})}}{\partial \Delta \mathbf{t}_j} &= -\mathbf{R}_j \\ \frac{\partial \mathbf{e}_j^{(\text{rpos})}}{\partial \Delta \mathbf{t}_{\text{prev}(j)}} &= \mathbf{R}_{\text{prev}(j)}. \end{aligned} \quad (20)$$

As the map point positions are also optimized in the BA (through reprojection errors), the Jacobian of the reprojection error to the map point positions can be computed as:

$$\frac{\partial \mathbf{e}_{ij}^{(\text{proj})}}{\partial \mathbf{x}_i^{3\text{D}}} = -\mathbf{J}_p(\mathbf{y}_{ij}) \mathbf{R}_j^T. \quad (21)$$

---

#### ALGORITHM 1: Body Translation Updater

---

**Input:** Global acceleration  $\mathbf{a}_{k-1}$ , camera-root position difference  $\mathbf{p}_{\text{root} \rightarrow \text{cam}}$  (computed from the human pose), camera localization  $\mathbf{p}_{\text{cam}}$  and its confidence  $\Sigma_{\text{cam}}$  (only if the camera image is available in the current frame).

**Output:** The human's global position  $\mathbf{p}_k$  for the current frame.

```

// initialization
if first frame then
    |  $\hat{\mathbf{x}}_{k-1} \leftarrow \mathbf{0}; \mathbf{P}_{k-1} \leftarrow \mathbf{O};$ 
    | initialize  $\mathbf{A}, \mathbf{B}, \mathbf{H}$ , and  $\mathbf{Q}$  according to Eq. 23;
end

// prediction
 $\mathbf{u}_{k-1} \leftarrow \mathbf{a}_{k-1};$ 
 $\hat{\mathbf{x}}_k^- \leftarrow \mathbf{A} \hat{\mathbf{x}}_{k-1} + \mathbf{B} \mathbf{u}_{k-1};$ 
 $\mathbf{P}_k^- \leftarrow \mathbf{A} \mathbf{P}_{k-1} \mathbf{A}^T + \mathbf{Q};$ 

// correction
if camera not available in the current frame then
    |  $\hat{\mathbf{x}}_k \leftarrow \hat{\mathbf{x}}_k^-; \mathbf{P}_k \leftarrow \mathbf{P}_k^-;$ 
else
    |  $\mathbf{R}_k \leftarrow \Sigma_{\text{cam}}; \mathbf{z}_k \leftarrow \mathbf{p}_{\text{cam}} - \mathbf{p}_{\text{root} \rightarrow \text{cam}};$ 
    | if relocalization detected then
    | |  $\hat{\mathbf{x}}_k \leftarrow \hat{\mathbf{x}}_k^-; \mathbf{P}_k \leftarrow \mathbf{P}_k^-;$ 
    | | replace the position in  $\hat{\mathbf{x}}_k$  with  $\mathbf{z}_k$ ;
    | else
    | |  $\mathbf{K}_k \leftarrow \mathbf{P}_k^- \mathbf{H}^T (\mathbf{H} \mathbf{P}_k^- \mathbf{H}^T + \mathbf{R}_k)^{-1};$ 
    | |  $\hat{\mathbf{x}}_k \leftarrow \hat{\mathbf{x}}_k^- + \mathbf{K}_k (\mathbf{z}_k - \mathbf{H} \hat{\mathbf{x}}_k^-);$ 
    | |  $\mathbf{P}_k \leftarrow (\mathbf{I} - \mathbf{K}_k \mathbf{H}) \mathbf{P}_k^-;$ 
    | end
end

// output
 $\mathbf{p}_k \leftarrow \mathbf{H} \hat{\mathbf{x}}_k;$ 

```

---

## E LOCAL MAP CONSTRUCTION FOR REAL-TIME TRACKING

As the reconstructed map can grow very large during tracking and finding matches between the image keypoints and 3D map points can be very time-consuming, we only use a *local* visible map in the camera tracking module to ensure its real-time performance. Following [Campos et al. 2021], we construct the local map from the observed points in recent and neighboring frames. One difference is that, when the visual tracking is lost, we additionally add to the local map the 3D points observed by any of 1) the candidate keyframes selected for relocalization and 2) the covisible<sup>5</sup> keyframes of the candidates. This effectively accelerates the recovery from visual tracking losses because the camera tracking can often succeed ahead of the strict place-recognition-based relocalization since we still have an accurate camera pose due to the combination with mocap. Therefore, we do not give up the camera tracking during visual losses as the original system does. Instead, we actively construct the local map from similar keyframes and try tracking the camera with respect to it.

<sup>5</sup>If two keyframes observe some shared map points, we say they are *covisible*.



## F DETAILS OF BODY TRANSLATION UPDATER

The task of the translation updater is to estimate the drift-free human global position from the global acceleration  $\mathbf{a}$ , camera localization  $\mathbf{p}_{\text{cam}}$ , and the camera confidence  $\Sigma_{\text{cam}}$ . We formulate the prediction-correction algorithm in the Kalman filter framework, where the state and observation equations are defined as:

$$\begin{aligned} \mathbf{x}_k &= \mathbf{A}\mathbf{x}_{k-1} + \mathbf{B}\mathbf{u}_{k-1} + \boldsymbol{\varphi}_{k-1} \\ \mathbf{z}_k &= \mathbf{H}\mathbf{x}_k + \boldsymbol{\psi}_k, \end{aligned} \quad (22)$$

where, in our setting, the state  $\mathbf{x} = [\mathbf{p}^T \ \mathbf{v}^T]^T$  is the concatenation of the human's global position and velocity,  $\mathbf{u}_{k-1} = \mathbf{a}_{k-1}$  is the acceleration input,  $\mathbf{z}_k$  is the root position observation calculated from the camera localization and the estimated human pose,  $\boldsymbol{\varphi}_{k-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$  and  $\boldsymbol{\psi}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$  denote the process and measurement

noise respectively. Then, we have:

$$\begin{aligned} \mathbf{A} &= \begin{pmatrix} \mathbf{I} & \Delta t \mathbf{I} \\ \mathbf{O} & \mathbf{I} \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} \mathbf{O} \\ \Delta t \mathbf{I} \end{pmatrix}, \quad \mathbf{H} = (\mathbf{I} \ \mathbf{O}), \\ \mathbf{Q} &= \begin{pmatrix} \mathbf{O} & \\ & \sigma^2 \mathbf{I} \end{pmatrix}, \quad \mathbf{R} = \Sigma_{\text{cam}}. \end{aligned} \quad (23)$$

Readers can refer to Sec. 3.4 as this equation is the matrix version of Eq. 6 and 7. Notice that we overload the notation  $\boldsymbol{\varphi}$  here to model the full-state transition noise. The body translation updater algorithm is presented in Alg. 1. We run the updater at 60 FPS. As the camera tracking module runs at the frame rate of the monocular camera (*i.e.*, 30 FPS), we distinguish the frames *with/without* the camera image in the algorithm. The drift-free human global positions can finally be outputted at 60 FPS.