

EgoPoser: Robust Real-Time Egocentric Pose Estimation from Sparse and Intermittent Observations Everywhere

Jiayi Jiang, Paul Strel, Manuel Meier, Christian Holz

Department of Computer Science, ETH Zürich, Switzerland

{firstname.lastname}@inf.ethz.ch

<https://siplab.org/projects/EgoPoser>

Abstract. Full-body egocentric pose estimation from head and hand poses alone has become an active area of research to power articulate avatar representations on headset-based platforms. However, existing methods over-rely on the indoor motion-capture spaces in which datasets were recorded, while simultaneously assuming continuous joint motion capture and uniform body dimensions. We propose EgoPoser to overcome these limitations with four main contributions. 1) EgoPoser robustly models body pose from intermittent hand position and orientation tracking only when inside a headset’s field of view. 2) We rethink input representations for headset-based ego-pose estimation and introduce a novel global motion decomposition method that predicts full-body pose independent of global positions. 3) We enhance pose estimation by capturing longer motion time series through an efficient SlowFast module design that maintains computational efficiency. 4) EgoPoser generalizes across various body shapes for different users. We experimentally evaluate our method and show that it outperforms state-of-the-art methods both qualitatively and quantitatively while maintaining a high inference speed of over 600 fps. EgoPoser establishes a robust baseline for future work where full-body pose estimation no longer needs to rely on outside-in capture and can scale to large-scale and unseen environments.

Keywords: Human Pose Estimation · Egocentric Vision · Mixed Reality

1 Introduction

Current Mixed Reality (MR) systems such as Microsoft HoloLens, Meta Quest, and Apple Vision Pro derive tracking cues and user input mainly from a head-worn platform. The cameras inside these devices observe the environment as well as the user’s hand motions when inside the field of view (FoV) [17, 18]. This enables the system to track the device’s position inside the world and derive input commands from the user’s actions. Due to the sparse nature of the input signals, whose capture relies on data from the user’s head and hands, current MR systems are limited in their ability to generate comprehensive virtual avatar

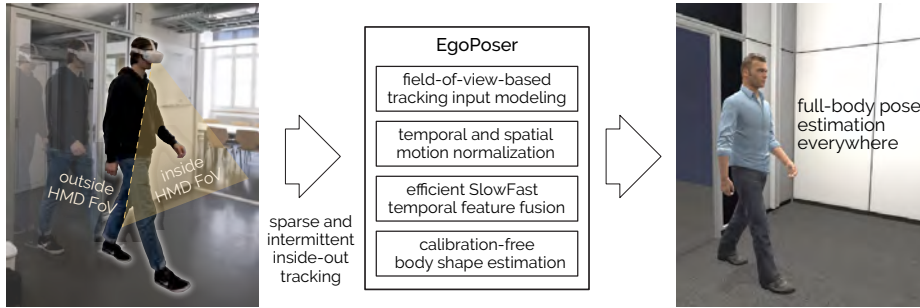


Fig. 1: Today’s Mixed Reality systems integrate all tracking inside the headset, supporting mobile use in everyday environments. This sacrifices much of the user’s body and surroundings for input, when body parts leave the cameras’ field of view. Accounting for these constraints, our novel method EgoPoser robustly estimates full-body poses that are spatially and temporally coherent, even from the sparse and intermittent inside-out tracking input available on today’s headsets.

representations, often limited to only their upper body. This reduces the fidelity of the user experience and may also affect the user’s sense of immersion.

Motivated by the goal to holistically embody users as full-body avatars in MR, several recent methods have been developed to estimate full-body poses from the sparse tracking cues current systems provide [7, 10, 20, 21, 34, 44]. These efforts all rely on large motion-capture datasets to estimate realistic body poses and animations, leveraging the robust, continuous, and high-fidelity recordings across a large variety of environments.

However, existing methods exhibit several limitations in real-world applications. (1) Prior approaches directly use the global pose in world space as the network input, causing the trained model to overfit to motions and poses that are specific to the environment and typically concentrated near the origin. Our paper reveals that using global input representation results in significantly worse predictions, even for slight meter-scale offsets. (2) Current methods assume that the pose of users’ hands is always available, which would require them to always remain within the field of view of the headset’s cameras. However, portable inside-out tracking systems obtain only intermittent information about the hands, as they occasionally move out of the field of view. (3) Existing methods only account for a mean body shape and disregard the natural variations in body shapes across different people. This limitation prevents the model from adapting to real-world inputs and accurately representing the user’s body. Often, motion artifacts such as floating and ground penetration arise from this.

To address these problems, we propose EgoPoser, an exclusively headset-centered estimation method for full-body poses that robustly performs on the sparse and intermittent tracking cues provided by today’s inside-out tracking systems. As shown in Fig. 1, EgoPoser comprises four main components that jointly enable its robust performance on real-world data and live motions outside motion-capture datasets. (1) EgoPoser’s realistic field-of-view modeling captures

both spatial and temporal information to smoothly estimate accurate full-body poses even when the user’s hands leave the camera’s view frustum. (2) Our novel global motion decomposition retains the critical relative global information using local representation, making it robust to position changes by encoding motion priors from sparse inputs. (3) We sample the original signals at different rates to capture longer motion time series through a SlowFast module as the input to the Transformer encoder, thus improving prediction accuracy without increasing the computational burden. (4) To support personalized use, we predict individual body shapes to accurately anchor each user’s representation within the virtual environment.

Taken together, we make the following contributions in this paper:

(1) We propose EgoPoser, a novel systematic approach to full-body pose estimation from the signals HMD devices provide. EgoPoser remains robust even when hands leave the field of view, and it generalizes well to various body shapes.

(2) We have identified a notable issue with existing methods: they tend to overfit to the training data due to the global input representation of the neural network. We emphasize the significance of position-invariant prediction and present an effective global motion decomposition strategy in EgoPoser.

(3) EgoPoser effectively accommodates to different body shapes unlike previous methods. We demonstrate our method’s input adaptability and the accurate output avatar representation it produces. In addition, EgoPoser significantly reduces motion artifacts such as floating and ground penetration.

(4) We demonstrate superior numerical and visual performance compared to state-of-the-art methods on the public datasets AMASS and HPS. Our demo also shows that EgoPoser can operate with real-world MR systems, making it a practical and effective solution for use with end-user devices.

2 Related Work

Egocentric human pose estimation for Mixed Reality. The task of estimating full-body poses from head and hand poses alone has gained significant attention with the growing popularity of MR-based egocentric vision [14, 15, 29]. Our previous work AvatarPoser [21] trained a single model for various motion types, combining a Transformer-based neural network with inverse kinematics (IK) optimization for realistic predictions that match the observations. QuestSim [41] and QuestEnvSim [25] combined reinforcement learning with physics simulation to ensure physically plausible predictions. While diffusion model-based methods such as AGRoL [11] and EgoEgo [26] synthesized smooth predictions, they both relied on future input signals to predict the current frame, and supported only slow sampling speeds—two factors that pose significant challenges for real-time applications. Recently, a temporal-spatial Transformer-based method [51] was proposed that can produce realistic results by joint-level modeling, but its network design is still computationally expensive. Our recent work MANIKIN [20] introduces a biomechanically accurate, differentiable full-body IK framework that matches observations and reduces ground penetration. However,

as we demonstrate in this paper, these previous projects [11, 20, 21, 26, 51] take the global pose as the network input, which is prone to overfitting to the training data and cannot generalize well to unseen positions. The design of EgoPoser thus aims to robustly perform independent of these positions.

Human pose estimation from inertial sensors. In addition to MR devices, there has also been research on full-body pose estimation from body-worn inertial sensors [8, 19, 22, 31, 38, 39, 45–48]. Because these sensors are distributed across the body, motion capture (MoCap) can become inflexible and obtrusive. Besides, inertial pose estimation usually struggles with accurate global position tracking since only relative position and orientation changes can be observed as input. In contrast, our task uses the three-point tracking of today’s MR devices without additional instrumentation. Without trackers on the lower body, we infer the complete body pose, including leg motions, from the global movements of the head and hands.

Pose estimation under field-of-view constraints. Estimating human pose when parts of the body are outside the cameras’ field of view is challenging [3, 9, 33, 35, 42, 43]. To alleviate the visibility issue, previous work has investigated alternative configurations with custom sensors such as wrist-worn cameras [27] or IMUs [36], cameras embedded inside controllers [5], and hat-mounted downward-facing fisheye cameras [6, 35, 40] or glasses [23, 50]. In terms of software solutions, FLAG [7] retained the original constraints of headset-only capture and augmented the training data by randomly masking the hands with a certain probability instead. While promising, it did not consider the spatial relative pose between the hand and headset, as hands can be masked out even if they are actually inside the FoV. In this paper, we realistically model the cameras’ FoV by considering the spatial relative pose and the temporal continuity.

3 Proposed Method: EgoPoser

We now describe our method EgoPoser for the real-time estimation of the global full-body pose based on head and hand poses tracked by an HMD device from egocentric vision. We first outline the specific task we investigate and then describe our method.

3.1 Overview

While MR systems differ in the tracking technology they implement, they all typically provide the global positions \mathbf{p} and orientations Θ of the headset, along with those of the user’s two hands. Following AvatarPoser [21], we aim to find a mapping f from the 3-point pose input to the user’s full-body pose, and thus the positions of all J full-body joints:

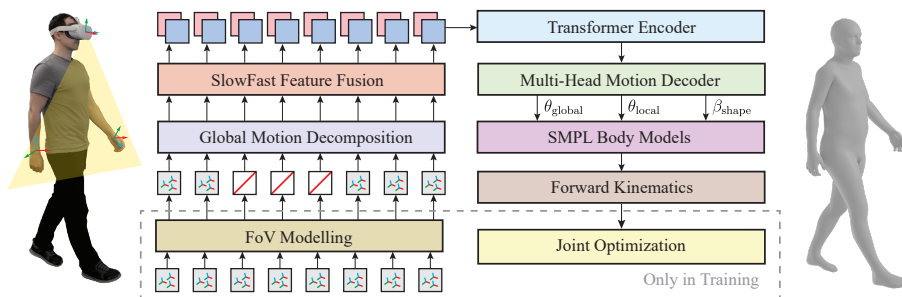


Fig. 2: The architecture of EgoPoser for full-body pose estimation from an MR device. Given $N=80$ frames as input, we generate the last frame as the full-body representation for each timestamp, facilitating real-time applications.

$$\{\mathbf{p}_t^j\}_{j=1:J} = f(\{\mathbf{p}_t^j, \Theta_t^j\}_{j=1:3}) \quad (1)$$

This is a challenging under-determined problem because the same input may correspond to multiple possible outputs. To ensure the prediction of consistent human skeletons, prior work used the first 22 joints defined in the kinematic tree of the SMPL-H [28, 30] skeleton model as the output representation of the full body, ignoring the pose of the fingers. Additionally, the SMPL-H model parameterizes the shape of the corresponding 3D human mesh with 16 shape parameters, $\beta \in \mathbb{R}^{16}$.

We show an overview of our approach in Fig. 2. The core components of our method include our proposed realistic FoV modeling, global motion decomposition, SlowFast feature fusion module, Transformer Encoder, and a human motion decoder. The output of the human motion decoder consists of the global root orientation θ_{global} , the local joint rotations θ_{local} , and the shape parameters β . The pose and shape parameters are then jointly optimized using the SMPL body model and forward kinematics. As proposed in AvatarPoser [21], the global root position is calculated using forward kinematics based on the tracked head position and the predicted joint angles.

3.2 Realistic Field of View Modeling

Previous work does not adequately address the inherent limitations of the inside-out hand tracking on today’s state-of-the-art headsets such as Apple Vision Pro, Meta Quest 2/3/Pro, and HoloLens 2. They typically model tracking failures through random frame drops uniformly sampled over the complete tracked motion but fail to account for the fact that hand tracking generally fails in regions outside the headset cameras’ field of view.

Based on the head pose, which determines the viewing angle of the cameras mounted on the headset, and the relative position of the hands, we simulate hand tracking failures for headsets with varying FoVs.




Fig. 3: An illustration of an HMD’s field of view and in-FoV conditions.

Fig. 3 shows an HMD’s field of view and the corresponding in-FoV conditions, where α_h is the horizontal FoV and α_v is the vertical FoV. Here, x_{hand} , y_{hand} , and z_{hand} are the x, y, and z coordinates of the hand position in a head-centered coordinate system with the x-axis pointing through the eyes. We train our method to robustly handle inputs with continuous tracking gaps by setting the input features of joints outside the field of view to 0.

3.3 Global Motion Decomposition

Today’s MR headsets track their own global pose as well as the user’s hands or controllers in three-dimensional space. Based on that, previous methods [11, 21, 51] use the global poses in world space as input to their networks. However, most existing datasets, like AMASS, are recorded in a limited physical space near the origin, so it remains unclear whether these methods can generalize well to different locations. Data augmentation could be a remedy, yet it makes the training process less efficient and it is impossible to cover the entire infinite 3-space. One common strategy is to decompose the global motion into a rigid body motion in global world space and a local motion relative to a root capturing the current body pose. However, in 3-point tracking problems, simply converting the reference frame to the head, the root frame, makes the prediction sensitive to head rotation. Additionally, removing the global information can lead to information loss, making the ill-posed problem even more challenging.

To combine the advantages of both global and local representations, we introduce a global motion decomposition strategy designed to be position-invariant for pose estimation across large-scale environments. Our approach contains two key operations, temporal and spatial normalization, detailed as follows:

(1) Temporal Normalization (TN). We perform temporal normalization by subtracting the translation of each joint at the first frame from the corresponding joint positions over the temporal window. This operation extracts the relative global trajectory of each joint across the temporal window.

(2) Spatial Normalization (SN). Instead of subtracting the head’s 3D translation from both the head and hand poses to obtain head-relative hand positions,




Fig. 4: An illustration of the temporal and spatial normalizations for robust position-invariant pose estimation.




Fig. 5: SlowFast feature fusion module. Original signals are sparsely and densely sampled and then concatenated.

we normalize only the horizontal translations relative to the head. The global vertical translation is retained as a crucial feature to encode motion priors.

Fig. 4 illustrates the spatial and temporal normalizations. These normalizations translate the positions of the headset and hands from a world reference frame R_{World} anchored at a static ground location to a reference frame R_{SN}^t anchored at the head’s projected ground location, and a reference frame $R_{\text{TN}}^{t,j}$ anchored at the j -th joint’s initial position for a given window. These transformed positions are then used as input to our pose estimation network. The positions provided by SN and TN at time t_i are written as:

$$\begin{aligned} \mathbf{p}_{\text{SN},h}^{t_i,\text{hand}} &= \mathbf{p}_{W,h}^{t_i,\text{hand}} - \mathbf{p}_{W,h}^{t_i,\text{head}} \\ \mathbf{p}_{\text{TN}}^{t_i,j} &= \mathbf{p}_W^{t_i,j} - \mathbf{p}_W^{t_0,j} \end{aligned} \quad (2)$$

In addition to the orientation and decomposed position information, we calculate the corresponding linear and angular velocities of the head and hands to enrich the input data. For the rotations, we use 6D representations [52] to ensure continuity. Finally, a total of 59 input features are provided to the network.

3.4 SlowFast Feature Fusion

Based on the information from a single frame, a multitude of plausible body poses exist that would fit a given set of head and hand poses. However, the problem converges toward a more unique solution as we consider the head and hand motions over a longer temporal context. Yet, simply adding more frames to the input would significantly increase computational overhead. For example, the computational complexity of a Transformer’s self-attention module scales quadratically with input sequence length. Thus, inspired by SlowFast networks originally proposed for video recognition [12], we propose a SlowFast feature fusion module that increases the context of considered past tracking frames in a more efficient manner.

As shown in Fig. 5, given an input window of τ past frames, the SlowFast module concatenates the linear embeddings for the last $\frac{\tau}{2}$ frames (FAST) with $\frac{\tau}{2}$ frames sampled with a stride of 2 over the complete window (SLOW). In this

way, we reduce the length of the input sequence by a factor of 2 while maintaining the temporal context over the whole window. Additionally, we still capture the temporal information contained within the higher temporal resolution of the FAST input frames.

3.5 Shape-aware Pose Optimization

One common limitation of existing methods [10, 11, 21] is that they assume a mean-shaped body skeleton and mesh, ignoring differences in body sizes across users. This assumption introduces two problems: First, since all the training and testing data share the same body shape, it is unclear whether a trained model will generalize well to input data from users with diverse body shapes in real-world scenarios. Second, the animation based on the mean-shaped skeleton does not accurately reflect the user’s real body size. Even with perfect joint angles, this can result in ground penetration, self-penetration, or floating artifacts.

Solution 1: Data augmentation + calibration. To address this issue, we introduce an approach that combines data augmentation with T-pose calibration. We augment the training data by incorporating ground truth shape parameters. For evaluation on the test set, we measure the body height and arm length and compute the ratio between the measured and the corresponding mean shape values. The average of these ratios is then used as the scaling factor for the entire skeleton and applied to adjust the output representation accordingly.

Solution 2: Joint pose and shape estimation. The previous solution does not account for variations in body ratios and requires an extra calibration step to approximate a user’s body size, adding to the overall effort involved. To address this, we introduce a calibration-free method that estimates the user’s shape from the tracking input.

Instead of directly supervising the shape parameters β in the loss function, we implicitly optimize the estimated β through the error in the joint positions generated by the shape-aware differentiable SMPL body model, which takes β and the estimated joint rotations as input. Additionally, we apply L1 regularization to β to encourage sparsity, ensuring that shape parameters which do not impact joint locations and bone lengths default to zero.

As our method estimates β for each frame, we can apply the median shape parameters from an initial sequence of frames to enforce consistency. However, we did not observe frequent or sudden deviations in β for a given input sequence.

Loss functions. The loss function for body shape estimation is written as:

$$\mathcal{L}_{\text{shape}} = \lambda_{\text{pos}} \mathcal{L}_{\text{pos}} + \lambda_{\beta} \|\beta\|_1 \quad (3)$$

where the shape-guided positional loss is calculated through forward kinematics:

$$\mathcal{L}_{\text{pos}} = \|\text{FK}(\theta, \beta) - \text{FK}(\theta_{GT}, \beta_{GT})\|_1 \quad (4)$$

The final loss function is composed of an L1 global orientation loss, an L1 local rotational loss, an L1 positional loss, and an L1 regularization of β denoted by:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{ori}}\mathcal{L}_{\text{ori}} + \lambda_{\text{rot}}\mathcal{L}_{\text{rot}} + \lambda_{\text{pos}}\mathcal{L}_{\text{pos}} + \lambda_{\beta} \|\beta\|_1 \quad (5)$$

We set the weights λ_{ori} , λ_{rot} , λ_{pos} , and λ_{β} to 0.05, 1, 1, and 0.01 respectively.

4 Experiments

4.1 Datasets and Training Details

Following prior work [21], we used three subsets of the AMASS [30] dataset, namely CMU [1], BMLrub [37], and HDM05 [32], for both training and testing. We applied the data split provided by AvatarPoser [21], which randomly allocates 90% of the sequences to the training set and 10% to the test set. To evaluate performance in the wild, we also used the HPS dataset [16] for testing, which captures subjects within large-scale scenes. For ground truth, we relied on the high-quality joint optimization results described by Guzov et al. [16], who consider camera localization, IMU pose estimates, and scene constraints.

To train EgoPoser, we adopted the Adam solver [24] with a batch size of 256. We considered the latest 80 frames as input ($\tau = 80$), resulting in an input window with 40 frames after SlowFast fusion. The learning rate starts from 1×10^{-4} and decays by a factor of 0.5 every 2×10^4 iterations. We trained EgoPoser using PyTorch on one NVIDIA GeForce GTX 3090 GPU.


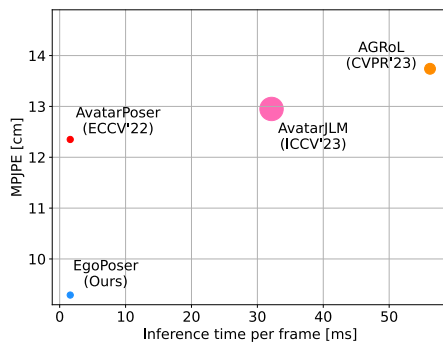
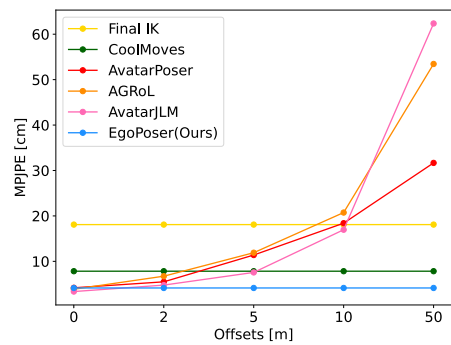
4.2 Evaluation Metrics

We use Mean Per Joint Position Error (MPJPE [cm]) and Mean Per Joint Velocity Error (MPJVE [cm/s]) as our main evaluation metrics to determine estimation accuracy and smoothness. To evaluate the shape-aware pose estimation, we use Mean Vertex Error (MVE [cm]), as well as the mean errors of the predicted heights (in [cm]) and bone lengths (in [cm]). In addition, we compute the average distance to the ground for mesh vertices below the ground to evaluate ground penetration [49]. To analyze foot floating artifacts, we calculate the mean distance (in [cm]) between the ground and the lowest vertex of the mesh across all frames where all vertices of the mesh are above the ground.

To ensure a fair comparison with state-of-the-art methods and to demonstrate the impact of each proposed component, we assume full hand visibility and use the mean body shape for comparisons with prior work making similar assumptions (Tab. 1 and Fig. 8). We evaluate the hand partial visibility problems and size-aware pose estimation independently in Tab. 3 and Tab. 2, respectively. Our motion decomposition method is employed in all evaluations of EgoPoser to ensure robust performance. This results in a slight increase in positional error when trained and tested on the three AMASS subsets, CMU, BMLrub, and HDM05, which capture motion sequences close to a fixed origin.

Table 1: Comparisons to state-of-the-art methods on HPS dataset, which is captured in large scenes.

Methods	BIB_EG_Tour		MPI_EG		Working_Standing		UG_Computers		Go_Around	
	MPJPE	MPJVE	MPJPE	MPJVE	MPJPE	MPJVE	MPJPE	MPJVE	MPJPE	MPJVE
AvatarPoser [21]	22.53	60.25	16.54	36.39	19.08	52.95	23.24	40.65	19.50	59.54
AvatarPoser-Improved	11.48	82.70	13.86	59.66	12.42	77.83	11.42	50.46	12.56	82.42
AGRoL [11]	28.95	166.34	19.41	55.52	17.67	53.97	20.90	109.12	14.16	98.34
AGRoL-Improved	15.04	124.12	13.94	89.42	13.86	89.42	12.71	106.43	13.13	128.42
AvatarJLM [51]	41.27	82.92	12.91	50.44	17.26	69.08	21.31	55.42	11.57	62.18
AvatarJLM-Improved	14.80	79.66	14.72	45.57	13.75	68.98	10.28	45.74	11.19	68.87
EgoPoser (Ours)	9.55	49.39	11.05	35.60	8.70	46.49	10.25	38.29	6.90	45.10

**Fig. 6:** Visual comparisons on the HPS dataset.**Fig. 7:** Comparisons of computational complexity. Marker size indicates parameter counts. We achieved the smallest error on the HPS dataset with fast inference.**Fig. 8:** Position error relative to users' distance from the origin. Our method remains robust with MPJPE of 4.14 cm and MPJVE of 25.95 cm/s.

4.3 Evaluation Results

Results of pose estimation in the wild. We compare our method to state-of-the-art approaches, including AvatarPoser [21], AGRoL [11], and AvatarJLM [51], on the real-world large-scene MoCap dataset HPS [16]. All methods are trained on the training sequences of the CMU, BMLrub, and HDM05 subsets of AMASS. Tab. 1 and Fig. 6 present the numerical and visual results. Our method achieves significantly better performance than existing methods. The results on the HPS dataset, encompassing motion data both in close proximity to the origin and at more distant points, demonstrate the robustness of our method across various spatial contexts. This is particularly noteworthy given that our model was exclusively trained on indoor MoCap data.

By adapting local representations, we also improved the performance of state-of-the-art methods on large-scale scenes. For each input window, we normalize the position of the head and hands to their positions in the first frame. This ensures that the network’s input range is consistently observed during training, even when the user is far from the origin. This method significantly enhances the overall performance of state-of-the-art methods, particularly when the position is far from the origin. However, our method still outperforms these methods by a substantial margin.

We also show the comparisons of the number of parameters and inference time in Fig. 7. Our method achieves significantly better performance with a much smaller model size and lower computational cost than recent methods.

Pose estimation robustness relative to the distance from the origin.

We also analyze the robustness to location variations on the CMU, BMLrub, and HDM05 subsets. To do this, we add a constant positional offset of 0, 2, 5, 10, and 50 m to each sequence to simulate a user performing a motion at different locations. We compare our method to AvatarPoser [21], AGRoL [11], and AvatarJLM [51], which use global input representations, and the classical KNN-based method CoolMoves [4], and FinalIK [2], which use local input representations. We plot the position error against different offsets in Fig. 8. While AvatarJLM achieved the best performance when no offset was applied, all prior learning-based methods, namely AvatarPoser, AGRoL, and AvatarJLM, taking global pose as input, experience a significant decrease in performance as the offset from the origin increases. In contrast, methods using local representations maintain stable performance.

Results of outside-the-FoV pose estimation. We evaluated various strategies for scenarios where hands are tracked by a headset and may go out of the cameras’ FoV. To simulate real-world scenarios, we considered various angles of available FoV: 180°(in fisheye cameras), 120°(in Quest 2), 90°(in Hololens 2). We tested the results on a model trained on hands with full visibility, denoted as ‘Full visibility’, on fine-tuned models with random hand masking using a probability $p = 0.2$ as proposed in FLAG [7] (denoted as ‘Random Masking’ or ‘RM’), and on our realistic FoV modeling (denoted as ‘Ours’).

Table 2: Evaluation of shape-aware pose estimation on AMASS dataset. All metrics are measured in [cm].

Strategies	MPJPE	Vertex Height	Arm	GP	FF
Mean Shape	6.36	6.74	7.67	7.42	3.87 5.38
Ours 1 - DA + Calib.	5.26	4.69	1.36	1.24	2.06 1.67
Ours 2 - Shape Est.	4.79	4.08	1.78	1.66	2.31 1.64

Table 4: Ablation study of different global motion decomposition methods using the AMASS dataset.

Strategies	MPJPE	MPJVE
Mean Norm. (all features)	6.25	42.69
Mean Norm. (horiz. + vert. pos.)	6.24	42.75
Mean Norm. (horiz. pos.)	6.25	42.87
Spatial Norm. (horiz. + vert. pos.)	4.96	29.59
Spatial Norm. (horiz. pos.)	4.45	27.56
Temporal Norm.	4.58	28.01
Ours	4.14	25.95

Tab. 3 and Fig. 9 present the numerical and visual results of models trained with different strategies. When testing the performance on various FoVs using the default model that assumes hands are always visible during training, we observe two main trends. First, as the FoV becomes smaller, the position error MPJPE increases. This is intuitive since a smaller FoV means there is more chance that hands are outside the FoV, rendering accurate predictions more challenging. Second, with a smaller FoV, the velocity error MPJVE initially increases and then decreases. This trend can be explained with the strong discontinuity in predictions for FoVs of 180° or 120°, switching between hands leaving and reentering the FoV. With even smaller FoVs, hands are mostly or even always outside the FoV, leading to smoother but less accurate predictions.

We also improve the random masking strategy. While random hand masking can improve results, our realistic FoV modeling strategy sets the visibility status based on the actual position of the hand relative to the head and captures the real temporal dependencies of hand visibility. It thus achieves the best performance for both position accuracy and smoothness.

Results of shape-aware pose estimation. We evaluate the performance of the shape-aware pose estimation on the same test data as in Fig. 8 from AMASS but with the true shape parameters β . This test set includes over 175 subjects with heights ranging from 145 to 207 cm. As listed in Tab. 2, the model trained using the mean body shape achieved a mere 6.36 cm in MPJPE. The average error for body dimensions such as height and arm length exceeds 7 cm. These discrepancies arise from an inaccurate shape representation, which leads to issues like ground penetration (GP) and floating feet (FF) as shown in Fig. 10. Metrics are in centimeters [cm].

Table 3: Results of different methods under various field of views on AMASS dataset.

Strategies	FoV = 180°		FoV = 120°		FoV = 90°	
	MPJPE	MPJVE	MPJPE	MPJVE	MPJPE	MPJVE
Full Visibility [21]	24.75	183.84	38.99	144.42	41.24	95.66
Random Masking [7, 11]	7.09	49.91	13.29	64.09	14.84	58.33
Improved RM	6.52	47.50	11.88	57.44	12.83	52.98
Ours	5.31	39.69	6.07	46.01	6.60	48.25

Table 5: Ablation study of the SlowFast design on AMASS dataset. Our design choice captures longer time series without introducing additional computational costs.

Strategies	MPJPE	MPJVE	FLOPs	#Parameters
length 40	4.36	28.12	0.33G	4.12M
length 80	4.11	29.27	0.65G	4.12M
length 80, s=2	4.13	30.02	0.33G	4.12M
Ours	4.14	25.95	0.33G	4.12M



Fig. 9: Visual comparisons between methods for the scenario where the hands can go out of FoV. Hands outside FoV are rendered in ●.




Fig. 10: Visual comparison between results using the mean shape and estimated shape. Our method reduces penetration and floating.

Conversely, as our first solution, data augmentation (DA) with ground truth body shapes and subsequently re-scaling the standardized model output by a body size factor obtained via T-pose calibration reduced the MPJPE to 5.26 cm. It also considerably enhanced performance across various metrics for body sizes and motion artifacts. Our calibration-free shape prediction approach showed further improvements in both MPJPE and mean vertex error, while delivering comparable outcomes in metrics for body size and motion artifacts. Calibration works slightly better for arm length and height as they are directly measured.

4.4 Ablation Studies

We conduct thorough ablation studies to show the effectiveness of each proposed component. The different approaches for shape-aware and FoV-aware pose estimation have already been discussed in Tab. 2 and Tab. 3.

Tab. 4 lists our ablation studies on global motion decomposition. Mean normalization refers to removing the mean value of each feature across the temporal window. Alternatively, we subtract only the mean of the vertical or horizontal joint translations. Spatial normalization refers to subtracting the horizontal and/or vertical head translation from all joint positions per frame. The results indicate that retaining vertical position information leads to better predictions. We also experiment with applying only temporal normalization. Combining our proposed temporal and spatial normalization achieves the best performance.

Tab. 5 shows ablation studies for the SlowFast design. We compare our approach with methods that utilize signals of varying frame lengths: 40, 80, or the original 80 samples downsampled by a factor of 2. The results demonstrate that our design benefits from extended temporal context without significantly increasing the model size or computational cost.



Fig. 11: EgoPoser’s visual demos on real data from Quest 2 with different participants.

4.5 Demos on Commodity Mixed Reality Devices

To assess EgoPoser’s robustness on actual real-world data from an end-user device, we ran it on live recordings from participants with diverse body shapes, wearing Meta Quest 2 headsets, recorded by Velt [13]. Fig. 11 shows the results.

4.6 Limitations and Discussions

While our method has demonstrated generalization ability for the challenging problems in HMD-based egocentric pose estimation, several limitations remain. First, since we encode vertical global position as an input feature, we assume that the user is moving on the same floor. Therefore, the origin system needs to be redefined when moving to another floor. Second, we address real-world inside-out tracking problems in this paper using a straightforward Transformer as our backbone model. Future work could investigate more elaborate models for this task. Third, our method has no post-processing steps. This could however further improve estimation accuracy or enhance physical plausibility.

5 Conclusion

We have proposed EgoPoser, a novel systematic approach for 3D full-body pose estimation based solely on the tracking information available on contemporary Mixed Reality head-mounted devices. We address the challenges faced by existing efforts using such platforms, specifically scaling robust estimations to arbitrary real-world settings, handling hands as input even when they are outside the cameras’ field of view, and robustness to varying body dimensions. Our experiments showed that EgoPoser achieves new state-of-the-art performance for accurate motion estimation under these challenging circumstances by combining our novel global motion decomposition, SlowFast fusion strategy, robust field-of-view modeling, and shape-aware pose estimation method. We believe that our proposed strategies can significantly contribute to the advancement of 3D full-body pose estimation and its integration into various AR/VR/MR applications.

Acknowledgments

We sincerely thank Andreas Fender for his help with data recording, testing, and manuscript proofreading.

References

1. CMU MoCap Dataset. <http://mocap.cs.cmu.edu/> (2004) 9
2. RootMotion Final IK. <https://assetstore.unity.com/packages/tools/animation/final-ik-14290> (2018) 11
3. Ahuja, K., Harrison, C., Goel, M., Xiao, R.: Mecap: Whole-body digitization for low-cost vr/ar headsets. In: Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology. pp. 453–462 (2019) 4
4. Ahuja, K., Ofek, E., Gonzalez-Franco, M., Holz, C., Wilson, A.D.: Coolmoves: User motion accentuation in virtual reality. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 5(2), 1–23 (2021) 11
5. Ahuja, K., Shen, V., Fang, C.M., Riopelle, N., Kong, A., Harrison, C.: Controller-pose: Inside-out body capture with vr controller cameras. In: CHI Conference on Human Factors in Computing Systems. pp. 1–13 (2022) 4
6. Akada, H., Wang, J., Shimada, S., Takahashi, M., Theobalt, C., Golyanik, V.: Unrealego: A new dataset for robust egocentric 3d human motion capture. In: European Conference on Computer Vision. pp. 1–17. Springer (2022) 4
7. Aliakbarian, S., Cameron, P., Bogo, F., Fitzgibbon, A., Cashman, T.J.: Flag: Flow-based 3d avatar generation from sparse observations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13253–13262 (2022) 2, 4, 11, 12
8. Armani, R., Qian, C., Jiang, J., Holz, C.: Ultra Inertial Poser: Scalable Motion Capture and Tracking from Sparse Inertial Sensors and Ultra-Wideband Ranging. In: ACM SIGGRAPH 2024 Conference Papers. SIGGRAPH ’24, Association for Computing Machinery, New York, NY, USA (2024) 4
9. Bailly, G., Müller, J., Rohs, M., Wigdor, D., Kratz, S.: Shoesense: a new perspective on gestural interaction and wearable applications. In: Proceedings of the SIGCHI conference on human factors in computing systems. pp. 1239–1248 (2012) 4
10. Dittadi, A., Dziadzio, S., Cosker, D., Lundell, B., Cashman, T.J., Shotton, J.: Full-body motion from a single head-mounted device: Generating smpl poses from partial observations. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11687–11697 (2021) 2, 8
11. Du, Y., Kips, R., Pumarola, A., Starke, S., Thabet, A., Sanakoyeu, A.: Avatars grow legs: Generating smooth human motion from sparse tracking inputs with diffusion model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2023) 3, 4, 6, 8, 10, 11, 12
12. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6202–6211 (2019) 7
13. Fender, A., Müller, J.: Velt: a framework for multi rgb-d camera systems. In: Proceedings of the 2018 ACM International Conference on Interactive Surfaces and Spaces. pp. 73–83 (2018) 14
14. Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., et al.: Ego4d: Around the world in 3,000 hours of egocentric video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18995–19012 (2022) 3

15. Grauman, K., Westbury, A., Torresani, L., Kitani, K., Malik, J., Afouras, T., Ashutosh, K., Baiyya, V., Bansal, S., Boote, B., et al.: Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19383–19400 (2024) [3](#)
16. Guzov, V., Mir, A., Sattler, T., Pons-Moll, G.: Human poseitioning system (hps): 3d human pose estimation and self-localization in large scenes from body-mounted sensors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4318–4329 (2021) [9](#), [11](#)
17. Han, S., Liu, B., Cabezas, R., Twigg, C.D., Zhang, P., Petkau, J., Yu, T.H., Tai, C.J., Akbay, M., Wang, Z., et al.: Megatrack: monochrome egocentric articulated hand-tracking for virtual reality. *ACM Transactions on Graphics (ToG)* **39**(4), 87–1 (2020) [1](#)
18. Han, S., Wu, P.c., Zhang, Y., Liu, B., Zhang, L., Wang, Z., Si, W., Zhang, P., Cai, Y., Hodan, T., et al.: Umetrack: Unified multi-view end-to-end hand tracking for vr. In: SIGGRAPH Asia 2022 Conference Papers. pp. 1–9 (2022) [1](#)
19. Huang, Y., Kaufmann, M., Aksan, E., Black, M.J., Hilliges, O., Pons-Moll, G.: Deep inertial poser: Learning to reconstruct human pose from sparse inertial measurements in real time. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)* **37**, 185:1–185:15 (Nov 2018) [4](#)
20. Jiang, J., Strel, P., Luo, X., Gebhardt, C., Holz, C.: MANIKIN: Biomechanically Accurate Neural Inverse Kinematics for Human Motion Estimation. In: European Conference on Computer Vision. Springer (2024) [2](#), [3](#), [4](#)
21. Jiang, J., Strel, P., Qiu, H., Fender, A., Laich, L., Snape, P., Holz, C.: Avatar-poser: 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. pp. 443–460. Springer (2022) [2](#), [3](#), [4](#), [5](#), [6](#), [8](#), [9](#), [10](#), [11](#), [12](#)
22. Jiang, Y., Ye, Y., Gopinath, D., Won, J., Winkler, A.W., Liu, C.K.: Transformer inertial poser: Real-time human motion reconstruction from sparse imu with simultaneous terrain generation. In: SIGGRAPH Asia 2022 Conference Papers. pp. 1–9 (2022) [4](#)
23. Kang, T., Lee, K., Zhang, J., Lee, Y.: Ego3dpose: Capturing 3d cues from binocular egocentric views. In: SIGGRAPH Asia 2023 Conference Papers. pp. 1–10 (2023) [4](#)
24. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: International Conference on Learning Representations (2015) [9](#)
25. Lee, S., Starke, S., Ye, Y., Won, J., Winkler, A.: Questenvsim: Environment-aware simulated motion tracking from sparse sensors. arXiv preprint arXiv:2306.05666 (2023) [3](#)
26. Li, J., Liu, K., Wu, J.: Ego-body pose estimation via ego-head pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17142–17151 (2023) [3](#), [4](#)
27. Li, S., Jiang, J., Ruppel, P., Liang, H., Ma, X., Hendrich, N., Sun, F., Zhang, J.: A mobile robot hand-arm teleoperation system by vision and imu. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 10900–10906. IEEE (2020) [4](#)
28. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)* **34**(6), 1–16 (2015) [5](#)

29. Ma, L., Ye, Y., Hong, F., Guzov, V., Jiang, Y., Postyeni, R., Pesqueira, L., Gamino, A., Baiyya, V., Kim, H.J., et al.: Nymeria: A massive collection of multimodal egocentric daily motion in the wild. arXiv preprint arXiv:2406.09905 (2024) [3](#)
30. Mahmood, N., Ghorbani, N., Troje, N.F., Pons-Moll, G., Black, M.J.: AMASS: Archive of motion capture as surface shapes. In: International Conference on Computer Vision. pp. 5442–5451 (Oct 2019) [5](#), [9](#)
31. Mollyn, V., Arakawa, R., Goel, M., Harrison, C., Ahuja, K.: Imuposer: Full-body pose estimation using imus in phones, watches, and earbuds. In: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. pp. 1–12 (2023) [4](#)
32. Müller, M., Röder, T., Clausen, M., Eberhardt, B., Krüger, B., Weber, A.: Documentation mocap database hdm05. Tech. Rep. CG-2007-2, Universität Bonn (June 2007) [9](#)
33. Parger, M., Tang, C., Xu, Y., Twigg, C.D., Tao, L., Li, Y., Wang, R., Steinberger, M.: Unoc: Understanding occlusion for embodied presence in virtual reality. IEEE Transactions on Visualization and Computer Graphics **28**(12), 4240–4251 (2021) [4](#)
34. Ponton, J.L., Yun, H., Aristidou, A., Andujar, C., Pelechano, N.: Sparseposer: Real-time full-body motion reconstruction from sparse data. ACM Transactions on Graphics **43**(1), 1–14 (2023) [2](#)
35. Rhodin, H., Richardt, C., Casas, D., Insafutdinov, E., Shafiei, M., Seidel, H.P., Schiele, B., Theobalt, C.: Egocap: egocentric marker-less motion capture with two fisheye cameras. ACM Transactions on Graphics (TOG) **35**(6), 1–11 (2016) [4](#)
36. Strel, P., Armani, R., Cheng, Y.F., Holz, C.: HOOV: Hand Out-Of-View Tracking for Proprioceptive Interaction using Inertial Sensing. In: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. pp. 1–16 (2023) [4](#)
37. Troje, N.F.: Decomposing biological motion: A framework for analysis and synthesis of human gait patterns. Journal of vision **2**(5), 2–2 (2002) [9](#)
38. Van Wouwe, T., Lee, S., Falisse, A., Delp, S., Liu, C.K.: Diffusionposer: Real-time human motion reconstruction from arbitrary sparse sensors using autoregressive diffusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2513–2523 (2024) [4](#)
39. Von Marcard, T., Rosenhahn, B., Black, M.J., Pons-Moll, G.: Sparse inertial poser: Automatic 3d human pose estimation from sparse imus. In: Computer graphics forum. vol. 36, pp. 349–360. Wiley Online Library (2017) [4](#)
40. Wang, J., Liu, L., Xu, W., Sarkar, K., Theobalt, C.: Estimating egocentric 3d human pose in global space. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11500–11509 (2021) [4](#)
41. Winkler, A., Won, J., Ye, Y.: Questsim: Human motion tracking from sparse sensors with simulated avatars. In: SIGGRAPH Asia 2022 Conference Papers. pp. 1–8 (2022) [3](#)
42. Wu, E., Yuan, Y., Yeo, H.S., Quigley, A., Koike, H., Kitani, K.M.: Back-hand-pose: 3d hand pose estimation for a wrist-worn camera via dorsum deformation network. In: Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology. pp. 1147–1160 (2020) [4](#)
43. Xie, X., Bhatnagar, B.L., Pons-Moll, G.: Visibility aware human-object interaction tracking from single rgb camera. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4757–4768 (2023) [4](#)
44. Yang, D., Kim, D., Lee, S.H.: Lobstr: Real-time lower-body pose prediction from sparse upper-body tracking signals. In: Computer Graphics Forum. vol. 40, pp. 265–275. Wiley Online Library (2021) [2](#)

45. Yi, X., Zhou, Y., Habermann, M., Golyanik, V., Pan, S., Theobalt, C., Xu, F.: Egolocate: Real-time motion capture, localization, and mapping with sparse body-mounted sensors. *ACM Transactions on Graphics (TOG)* **42**(4), 1–17 (2023) [4](#)
46. Yi, X., Zhou, Y., Habermann, M., Shimada, S., Golyanik, V., Theobalt, C., Xu, F.: Physical inertial poser (pip): Physics-aware real-time human motion tracking from sparse inertial sensors. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 13167–13178 (2022) [4](#)
47. Yi, X., Zhou, Y., Xu, F.: Transpose: real-time 3d human translation and pose estimation with six inertial sensors. *ACM Transactions on Graphics (TOG)* **40**(4), 1–13 (2021) [4](#)
48. Yi, X., Zhou, Y., Xu, F.: Physical non-inertial poser (pnp): Modeling non-inertial effects in sparse-inertial human motion capture. In: *ACM SIGGRAPH 2024 Conference Papers*. pp. 1–11 (2024) [4](#)
49. Yuan, Y., Wei, S.E., Simon, T., Kitani, K., Saragih, J.: Simpoe: Simulated character control for 3d human pose estimation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 7159–7169 (2021) [9](#)
50. Zhao, D., Wei, Z., Mahmud, J., Frahm, J.M.: Egoglass: Egocentric-view human pose estimation from an eyeglass frame. In: *2021 International Conference on 3D Vision (3DV)*. pp. 32–41. *IEEE* (2021) [4](#)
51. Zheng, X., Su, Z., Wen, C., Xue, Z., Jin, X.: Realistic full-body tracking from sparse observations via joint-level modeling. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 14678–14688 (2023) [3](#), [4](#), [6](#), [10](#), [11](#)
52. Zhou, Y., Barnes, C., Lu, J., Yang, J., Li, H.: On the continuity of rotation representations in neural networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5745–5753 (2019) [7](#)