

# Egocentric Visibility-Aware Human Pose Estimation

Peng Dai, Yu Zhang, Yiqiang Feng, Zhen Fan, Yang Zhang  
PICO

{daipeng.2022, zhangyu.118, fengyiqiang, fanzhen.0315, zhangyang.0621}@bytedance.com

## Abstract

*Egocentric human pose estimation (HPE) using a head-mounted device is crucial for various VR and AR applications, but it faces significant challenges due to keypoint invisibility. Nevertheless, none of the existing egocentric HPE datasets provide keypoint visibility annotations, and the existing methods often overlook the invisibility problem, treating visible and invisible keypoints indiscriminately during estimation. As a result, their capacity to accurately predict visible keypoints is compromised. In this paper, we first present Eva-3M, a large-scale egocentric visibility-aware HPE dataset comprising over 3.0M frames, with 435K of them annotated with keypoint visibility labels. Additionally, we augment the existing EMHI dataset with keypoint visibility annotations to further facilitate the research in this direction. Furthermore, we propose EvaPose, a novel egocentric visibility-aware HPE method that explicitly incorporates visibility information to enhance pose estimation accuracy. Extensive experiments validate the significant value of ground-truth visibility labels in egocentric HPE settings, and demonstrate that our EvaPose achieves state-of-the-art performance in both Eva-3M and EMHI datasets.*

## 1. Introduction

Egocentric human pose estimation (HPE) using a head-mounted device (HMD) has gained great attention in recent years for its numerous applications in virtual reality (VR), augmented reality (AR), robotic control, etc [6]. Different from outside-in HPE with external cameras, a key challenge of egocentric HPE lies in the keypoint invisibility problem [7, 26, 48], which usually results from two causes. As shown in Fig. 1 (e), the first cause arises from severe self-occlusions of body parts, particularly the lower body. Another cause is the limited field-of-view (FoV) of head-mounted cameras, making it difficult to fully capture the body especially when hands and legs are stretched out.

Annotating the visibility of each keypoint is a labor-intensive and time-consuming process. To the best of our knowledge, none of the existing egocentric HPE datasets

include visibility annotations for keypoints. Consequently, current egocentric HPE methods [15, 43, 48] often neglect the invisibility issue and treat visible and invisible keypoints indiscriminately in pose estimation. For invisible keypoints, the absence of direct visual evidence introduces inherent ambiguity in predicting their 3D positions. As our experiments demonstrate, this undifferentiated treatment compromises the estimation accuracy of visible keypoints.

We introduce Eva-3M, a large-scale, visibility-aware egocentric HPE dataset comprising over 3.0M frames, including 435K with detailed keypoint visibility labels. As shown in Fig. 1, Eva-3M was collected from 31 subjects performing 24 types of daily activities using a commercial Pico4 Ultra VR-MR headset, offering realistic insights into HPE with real VR devices. It stands as the first and largest real egocentric dataset to provide both ground-truth SMPL [29] poses and visibility labels. To foster broader research, we also extend our contribution by providing keypoint visibility labels for the existing EMHI [15] dataset.

To effectively leverage keypoint visibility and mitigate the interference of invisible keypoints on visible ones, we propose EvaPose, a novel visibility-aware framework for egocentric HPE. As depicted in Fig. 2, EvaPose comprises three key components: (1) To ensure the generation of plausible 3D poses particularly for invisible keypoints, we employ a Vector Quantized-VAE (VQ-VAE) [41] model pre-trained on extensive motion capture (mocap) datasets to embed realistic human pose priors. (2) EvaPose presents a visibility-aware 3D estimation network to predict per-frame 3D keypoints and their visibility states for the first time. This is complemented by a visibility-based loss-weighting scheme that applies distinct supervision to visible and invisible keypoints during training. (3) An iterative intra- and inter-frame attention module is designed to refine the initial predictions. Within this module, per-frame 3D keypoints and their visibility scores iteratively interact with stereo image features and with each other across the whole sequence, achieving robust multi-view and temporal fusion. Finally, the fused features are processed by the pre-trained VQ-VAE decoder to reconstruct the final, high-fidelity 3D poses.

In summary, our main contributions are as follows:

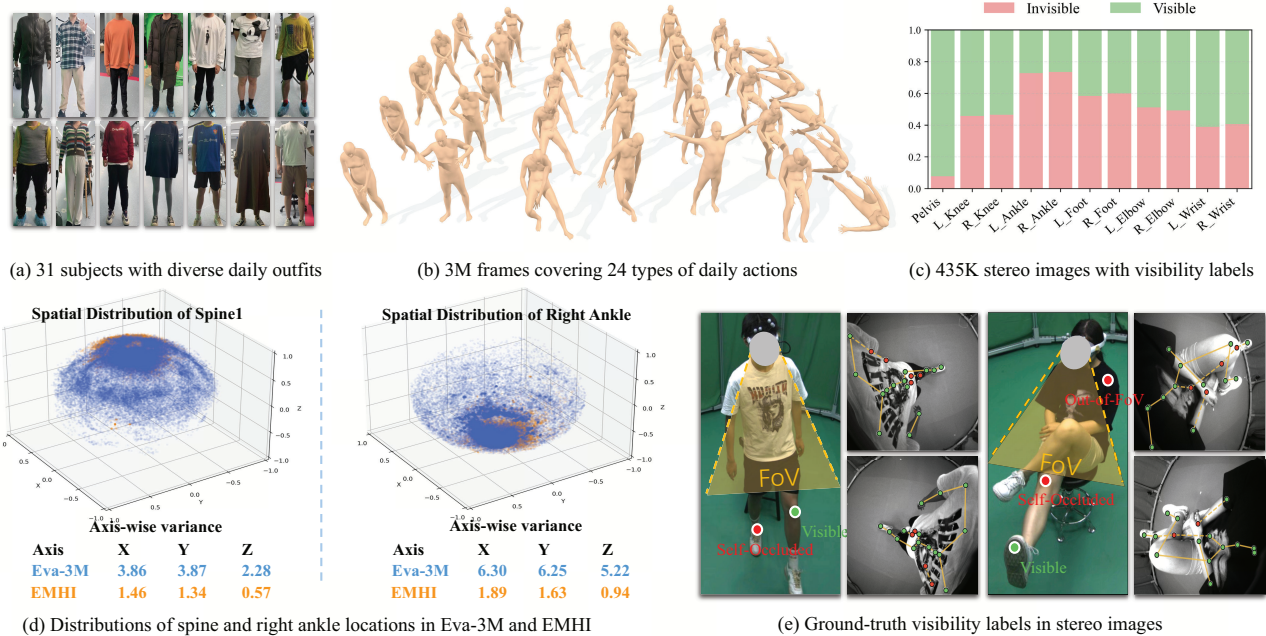


Figure 1. We introduce Eva-3M, a large-scale egocentric visibility-aware dataset comprising over 3.0M frames from (a) 31 subjects in daily outfits (b) performing 24 types of daily actions, and of which (c) 435K are annotated with keypoint visibility labels. (d) shows the normalized spatial distribution comparison between Eva-3M and EMHI, indicating that Eva-3M has a wider range of motion diversity than EMHI. (e) illustrates a few representative examples of keypoint invisibility due to self-occlusion and out-of-FoV in Eva-3M dataset.

- **Eva-3M Dataset:** A large-scale, real egocentric HPE dataset with keypoint visibility labels for the first time.
- **EMHI Enrichment:** We augment the existing EMHI dataset [15] by contributing keypoint visibility labels.
- **EvaPose Method:** A novel visibility-aware framework that explicitly utilizes keypoint visibility information to significantly enhance the 3D pose estimation accuracy.
- **Superior Performance:** Extensive experiments validate the superiority of our method over previous methods.

## 2. Related Work

### 2.1. Egocentric HPE Datasets

Recent years have seen the release of large-scale datasets for egocentric video understanding [50], egocentric action recognition [16, 17], and egocentric life assistant [49]. However, in the field of egocentric HPE, the availability of large-scale datasets still remains limited. Collecting a real-world egocentric HPE dataset is challenging due to specialized devices and complex setups are required for accurate markerless motion capture. Consequently, prior work has predominantly relied on synthetic datasets [2, 3, 10, 38, 43, 44, 46]. Although synthetic data enables scalable generation with perfect ground-truth annotations, the inherent domain gap between synthetic and real-captured data inevitably limits the performance of the models when deployed in real-world scenarios.

For real-world egocentric HPE datasets, most existing works use custom-made wearable capture rigs, such as EgoCap [35], EgoPW-Scene [43], UE-RW [3], Ego4View-RW [4], and SELF [7] dataset. However, these capture rigs differ significantly from commercial VR/AR devices in both design and imaging characteristics. For instance, custom setups often employ protruding camera placements to minimize self-occlusions, while commercial VR/AR devices follow a slim and lightweight design, which typically induces more severe self-occlusions. In response to this discrepancy, several recent datasets have been constructed using actual VR/AR devices. Nymeria [31] uses Project Aria glasses [36] to collect 300 hours of in-the-wild daily motion data. EgoBody3M [52] introduces the first egocentric HPE dataset captured with a realistic VR headset configuration, while EMHI [15] provides a multi-modal egocentric human motion dataset combining real VR devices with body-worn inertial measurement units (IMUs). Despite these advances, neither EgoBody3M [52] nor EMHI [15] addresses the issue of keypoint invisibility, and neither dataset provides visibility annotations for keypoints.

### 2.2. Egocentric HPE Methods

Existing methods for estimating the wearer’s body pose vary by input modality, including sparse IMUs [5, 11, 24, 45, 54], front-facing egocentric cameras [19, 22, 27, 51], and down-facing egocentric cameras [2, 3, 7, 15, 25, 26, 38,

39, 46, 48, 52, 53]. Among these, front-facing camera approaches often treat pose estimation as a motion generation or inpainting task, since the wearer’s body is largely out of view. In this work, we focus primarily on methods using down-facing egocentric cameras, which are most relevant to our setting.

Early efforts in this line of work mainly addressed 3D pose estimation from visual-only inputs. EgoGlass [53] first predicts 2D heatmaps from each view and then applies a 3D lifting module to reconstruct 3D keypoints from multi-view 2D heatmaps. UnrealEgo [2] follows a similar pipeline but introduces a dual weight-sharing encoder architecture to better leverage stereo information for heatmap estimation. EgoPoseFormer [48] further advances this direction with a Deformable Self-Attention (DSA) mechanism for effective multi-view feature fusion, and adopts a DETR-style [9] Transformer for coarse-to-fine pose estimation.

More recent studies incorporate camera or device poses from SLAM systems to improve both 3D pose estimation accuracy and physical plausibility. FRAME [7] first estimates per-frame 3D poses in local coordinates, aligns them into a global coordinate system using device poses, and refines the sequence of 3D poses with a temporal fusion module. REWIND [26] proposes a diffusion-based approach [21] that generates global whole-body motion conditioned on stereo egocentric images and camera poses.

Despite these advances, existing egocentric HPE methods overlook the adverse effect of invisible keypoints on the estimation accuracy of visible ones. In contrast, we propose a visibility-aware 3D pose estimation network along with a visibility-based loss-weighting scheme, which explicitly differentiates between visible and invisible keypoints during training to mitigate such interference.

### 3. Eva-3M Dataset

The primary challenge in egocentric HPE lies in the prevalent issue of keypoint invisibility. As shown in Fig. 1 (a), which illustrates the visibility statistics of 11 selected keypoints in our Eva-3M dataset, limb keypoints are invisible nearly half of the time. Despite its significant impact, none of the existing datasets provide visibility annotations. To address this gap, we introduce Eva-3M, the first and largest real-world egocentric dataset offering both ground-truth SMPL poses and comprehensive visibility labels, with substantially greater motion diversity. This combination supports more reliable training and evaluation of egocentric pose estimation models under realistic conditions.

A detailed comparison in Tab. 1 highlights the advantages of our Eva-3M across several dimensions. First, unlike prior datasets built with custom capture rigs, Eva-3M is captured using a commercial Pico4 Ultra VR-MR headset, providing direct insights into HPE performance in real VR devices. Second, as depicted in Fig. 1(b) and (c), Eva-3M

Dataset	R/S	Cams	Frames	Vis	SMPL	Subj	Act
Mo <sup>2</sup> Cap <sup>2</sup> [46]	Syn.	1	530K	×	×	700	3K
xR-EgoPose [38]	Syn.	1	383K	×	×	46	-
UnrealEgo [2]	Syn.	2	450K	×	×	17	30
EgoGTA [43]	Syn.	1	320K	×	×	5	-
SynthEgo [10]	Syn.	2	60K	×	✓	6K	-
EgoCap [35]	Real	2	75K	×	×	8	-
EgoGlobal [42]	Real	1	12K	×	×	2	13
EgoGlass [53]	Real	2	173K	×	×	10	-
UE-RW [3]	Real	2	260K	×	×	16	-
SELF [7]	Real	2	1.6M	×	×	14	-
EMHI [15]	Real	2	3.1M	×	✓	58	39
EgoBody3M [52]	Real	4	3.4M	×	×	120	30
<b>Eva-3M</b>	Real	2	3.0M	✓	✓	31	24

Table 1. Comparison of egocentric motion datasets. R/S denotes whether the dataset is collected in real or synthetic setting. Cams indicates the number of egocentric cameras. Vis denotes the availability of keypoint visibility labels. Subj is the number of subjects, and Act is the number of action categories.

includes 3.0 million synchronized frames from 31 subjects performing 24 categories of common VR motions, offering richer motion diversities than existing benchmarks such as EMHI [15]. To quantify this, we compare the motion diversity between Eva-3M and EMHI by randomly sampling 70K root-relative 3D keypoint coordinates from both. For better visualization, the root-relative coordinates are normalized by their maximum distance. The spatial distributions of an internal keypoint (spine1) and an extremity keypoint (right ankle) are visualized in Fig. 1(d), where Eva-3M covers a noticeably broader spatial range. We further compute the variance of the original (non-normalized) 3D coordinates in decimeters, and observe consistently larger variances for both keypoints in Eva-3M, confirming its wider range of motion diversities. Third, Eva-3M provides the first large-scale keypoint visibility annotations across 435K frames. Representative examples of these annotations are visualized in Fig. 1(e).

Eva-3M consists of 1,353 motion sequences recorded at 30 fps. Each frame includes paired stereo grayscale images (640×480), ground-truth SMPL parameters in both camera and world coordinates, and corresponding ground-truth 2D/3D keypoints. Further details regarding the data capture system, ground-truth acquisition, and visibility annotation process are provided in the supplementary materials.

## 4. EvaPose

### 4.1. Overview

Our task is to estimate the body poses of the wearer using egocentric sensors from a head-mounted device (HMD). Following recent works [7, 26], our method takes as in-

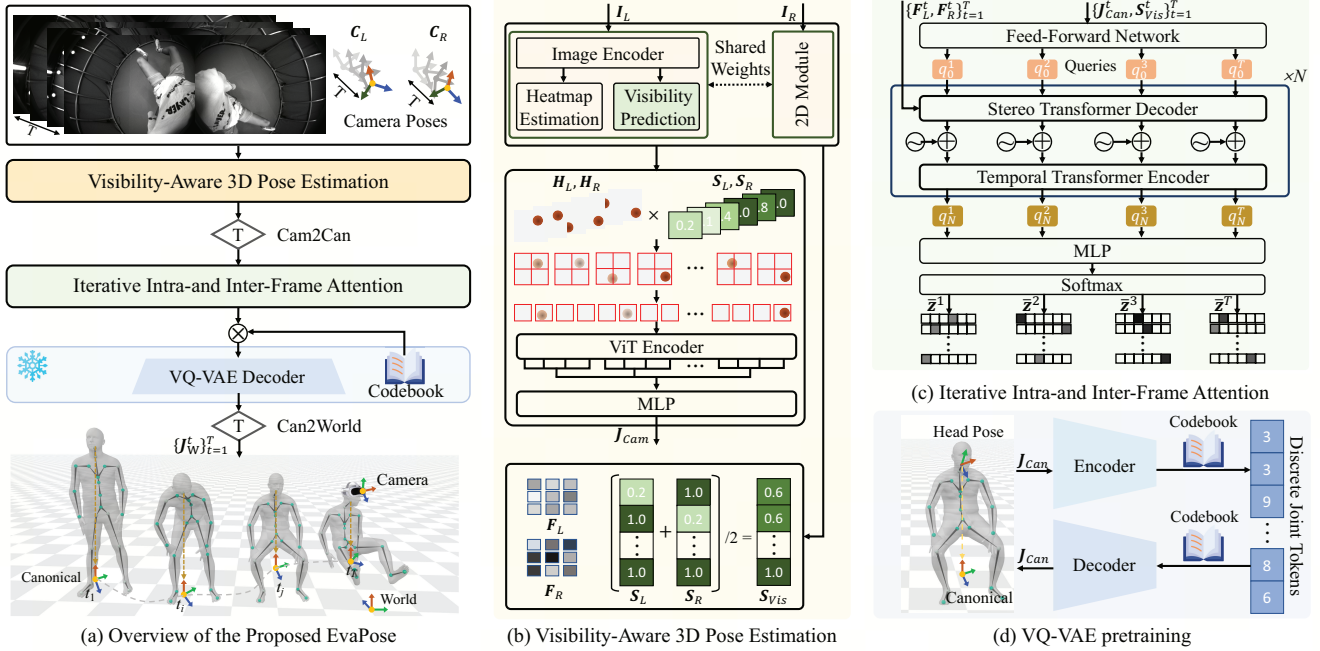


Figure 2. Overview of the proposed EvaPose. Given a sequence of egocentric observations, we first propose a visibility-aware 3D pose estimation network to extract stereo image features, and predict per-frame 3D keypoints in the camera coordinate system (defined as the left camera coordinate system in this paper) and their corresponding visibility confidence scores. Then, the predicted 3D keypoints are transformed to the canonical coordinate system with the help of camera poses from SLAM system. Next, an iterative intra-and inter-frame attention network is used for temporal feature fusion. Finally, we estimate the 3D poses with the pre-trained VQ-VAE decoder.

put both stereo egocentric videos and camera poses from the HMD’s built-in SLAM system. Formally, our proposed EvaPose models the human pose estimation problem conditioned on a sequence of egocentric observations over a window of  $T$  frames:

$$f_{\phi}(\mathbf{J}_W^{1:T} | \mathbf{I}_L^{1:T}, \mathbf{I}_R^{1:T}, \mathbf{C}_L^{1:T}, \mathbf{C}_R^{1:T}) \quad (1)$$

where  $\phi$  denotes the learnable network parameters, and  $\mathbf{J}_W^t$  represents the set of  $N_J$  SMPL [30] keypoints in the world coordinate system at frame  $t$ . The inputs consist of stereo egocentric images  $\mathbf{I}_{v \in \{L,R\}}^t \in \mathbb{R}^{C \times H \times W}$  from the left ( $L$ ) and right ( $R$ ) views, along with their corresponding camera poses  $\mathbf{C}_{v \in \{L,R\}}^t = [\mathbf{R}_v^t | \mathbf{T}_v^t] \in \mathbb{R}^{3 \times 4}$ , which include camera rotation  $\mathbf{R}_v^t \in \mathbb{R}^{3 \times 3}$  and translation  $\mathbf{T}_v^t \in \mathbb{R}^{3 \times 1}$ . The camera poses are provided by the HMD’s SLAM system.

As illustrated in Fig. 2, EvaPose mainly consists of three components. (1) A VQ-VAE [41] pretrained on extensive mocap datasets encodes plausible human poses into a discrete codebook, serving as a strong pose prior (Sec. 4.2). (2) A visibility-aware 3D pose estimation network extracts stereo image features  $[\mathbf{F}_L^t, \mathbf{F}_R^t]$ , and jointly predicts per-frame 3D keypoints  $\mathbf{J}_{Cam}^t$  along with their visibility confidence scores  $\mathbf{S}_{Vis}^t$  (Sec. 4.3). (3) An iterative intra-and inter-frame attention network fuses the above data over the whole time window for 3D pose refinement (Sec. 4.4).

## 4.2. Learning the Pose Prior

This section describes how we encode 3D human poses into a discrete representation using VQ-VAE [41]. Previous methods typically adopt either the HumanML3D [18] representation or local pose parameters [14] as the pose representation. However, both approaches overlook shape parameters and present limitations in our context. Since our goal is to estimate accurate 3D keypoint positions for different users with varying body proportions and heights, we represent human poses using 3D keypoints defined in a canonical coordinate system. As illustrated in Fig. 2 (d), this coordinate system is constructed by projecting the head joint onto the floor plane, aligning the vertical ( $y$ ) axis accordingly, and using the projection of the head’s horizontal orientation as the horizontal axis. A key advantage of this representation is its invariance to arbitrary translations in the floor plane and rotations around the vertical axis.

The overall VQ-VAE architecture, illustrated in Fig. 2(d), consists of an encoder  $\mathbf{E}$ , a decoder  $\mathbf{D}$ , and a learnable codebook  $\mathbf{CB} = \{c_k\}_{k=1}^K$ , where each code  $c_k \in \mathbb{R}^D$  and  $D$  is the code dimension. Given 3D keypoints  $\mathbf{J}_{Can}$  in the canonical coordinate system as input, the encoder produces a latent feature sequence  $\mathbf{z} = \mathbf{E}(\mathbf{J}_{Can}) = [z_1, z_2, \dots, z_M]$ , where each  $z_i \in \mathbb{R}^D$  and  $M$  is the number of features. Each  $z_i$  is quantized

to tokens by finding the most similar code element in the codebook. During training, we adopt the strategy from prior work [14], utilizing exponential moving average updates and periodic codebook resets to prevent codebook collapse. Meanwhile, a combination of three loss terms including the reconstruction loss, the embedding loss, and the commitment loss are used to optimize the network.

### 4.3. Visibility-Aware 3D Pose Estimation

In contrast to prior egocentric HPE methods [7, 44, 48], we introduce a visibility-aware 3D pose estimation network which, for the first time, explicitly predicts the visibility state of each keypoint. By leveraging these predicted states, EvaPose differentiates between visible and invisible keypoints during training, leading to improved HPE accuracy. For notational simplicity, we omit frame index  $t$  in this subsection, as all operations are applied identically per frame.

As illustrated in Fig. 2 (b), given a stereo image pair  $\mathbf{I}_{v \in \{L, R\}} \in \mathbb{R}^{C \times H \times W}$  as input, an image encoder first extracts visual features  $\mathbf{F}_L$  and  $\mathbf{F}_R$ . These features are then processed by two separate lightweight decoders: one estimates 2D heatmaps  $\mathbf{H}_v \in \mathbb{R}^{N_J \times H' \times W'}$ , and the other predicts visibility scores  $\mathbf{S}_v \in \mathbb{R}^{N_J}$ , where  $H'$  and  $W'$  denote the spatial dimensions of the heatmaps. Following the design of ViTPose [47], the heatmap decoder comprises a series of deconvolution layers followed by a final prediction layer. The visibility decoder, for simplicity, is implemented as a convolution layer followed by a Multi-Layer Perceptron (MLP). Subsequently, a ViT encoder integrates the multi-view heatmaps and visibility scores for 3D pose estimation. Let  $\{\mathbf{H}_{i,v}\}$  and  $\{s_{i,v}\}$  denote the sets of  $2N_J$  joint heatmaps and visibility scores, respectively. We first compute visibility-aware heatmaps as  $\mathbf{H}'_{i,v} = s_{i,v} \cdot \mathbf{H}_{i,v}$ . These heatmaps are then embedded into tokens via a patch embedding layer and augmented with learnable positional encodings before being fed into the transformer. The tokens are processed by three successive ViT encoder layers, which effectively model cross-joint and cross-view dependencies. Finally, for the  $i$ -th keypoint, all tokens corresponding to its multi-view heatmaps (i.e.,  $\mathbf{H}'_{i,L}$  and  $\mathbf{H}'_{i,R}$ ) are concatenated and passed through an MLP to regress its 3D position  $\mathbf{J}_{Cam}^i$  in the camera coordinate system. Meanwhile, the visibility scores from the left and right views are averaged to produce the final 3D visibility score:  $\mathbf{S}_{Vis} = (\mathbf{S}_L + \mathbf{S}_R)/2$ .

### 4.4. Iterative Intra-and Inter-Frame Attention

We propose an iterative intra-and inter-frame attention network to further improve both the accuracy and temporal smoothness of pose estimation. First, the 3D keypoints  $\mathbf{J}_{Cam}^{1:T}$  in the camera coordinate system are transformed into the canonical coordinate system as  $\mathbf{J}_{Can}^{1:T}$  using the camera poses. The detailed derivation of coordinate transformations is provided in the supplementary material. The

canonical 3D keypoints  $\mathbf{J}_{Can}^t$  and visibility scores  $\mathbf{S}_{Vis}^t$  at each frame are then concatenated and processed by a feed-forward network to generate frame-level queries  $\{\mathbf{q}_0^t\}_{t=1}^T$ . As shown in Fig. 2 (c), these queries undergo alternating processing between a Stereo Transformer Decoder (STD) and a Temporal Transformer Encoder (TTE). The STD fuses stereo visual features within each frame, while the TTE aggregates data across the entire temporal window. Formally, at the  $n$ -th iteration, given features  $\{\mathbf{q}_{n-1}^t\}_{t=1}^T$  from the previous iteration, the STD employs two transformer decoders to enable each  $\mathbf{q}_{n-1}^t$  to interact with the visual features from the left and right views independently:

$$\mathbf{f}_v^t = \text{Decoder}(\mathbf{q}_{n-1}^t, \mathbf{F}_v), \quad v \in \{L, R\} \quad (2)$$

The resulting features  $\mathbf{f}_L^t$  and  $\mathbf{f}_R^t$  are concatenated and passed through an MLP to produce a multi-view fused feature  $\mathbf{f}_n^t$ . All multi-view fused features  $\{\mathbf{f}_n^t\}_{t=1}^T$  within the temporal window are then fed into the TTE, implemented as a transformer encoder, for temporal fusion:

$$[\mathbf{q}_n^1, \mathbf{q}_n^2, \dots, \mathbf{q}_n^T] = \text{Encoder}([\mathbf{f}_n^1, \mathbf{f}_n^2, \dots, \mathbf{f}_n^T]) \quad (3)$$

After  $N$  iterations, we obtain multi-view and temporally fused features  $\{\mathbf{q}_N^t\}_{t=1}^T$ , which are subsequently processed by an MLP and a softmax operation to estimate the logits:

$$\bar{\mathbf{z}}^t = \text{Softmax}(\text{MLP}(\mathbf{q}_N^t)) \quad (4)$$

where  $\bar{\mathbf{z}}^t \in \mathbb{R}^{M \times K}$  represents the predicted logits for frame  $t$ . Following prior work [14], we avoid the non-differentiable operation of directly selecting an embedding from the codebook by instead estimating logits and multiplying them with the pretrained codebook, yielding approximated quantized features:

$$\mathbf{z}^t = \bar{\mathbf{z}}_{M \times K}^t \times \mathbf{CB}_{K \times D} \quad (5)$$

Finally, the features  $\{\mathbf{z}^t\}_{t=1}^T$  are passed through the VQ-VAE decoder to reconstruct the refined 3D poses  $\mathbf{J}_{Can}^{1:T}$  in the canonical coordinate system, which are then transformed back to the world coordinate system to produce the final 3D poses  $\mathbf{J}_W^{1:T}$ .

### 4.5. Training EvaPose

The training process is divided into two stages. In the first stage, we introduce a visibility-based loss-weighting scheme to train the visibility-aware 3D pose estimation network. This strategy mitigates interference from invisible keypoints by assigning different loss weights according to their visibility states. The overall loss is defined as:

$$\mathcal{L}_{\text{stage1}} = \lambda_{\text{vis}} \mathcal{L}_{\text{vis}} + \lambda_{\text{heatmap}} \mathcal{L}_{\text{heatmap}} + \lambda_{3D} \mathcal{L}_{3D} \quad (6)$$

where  $\lambda_{\text{vis}}$ ,  $\lambda_{\text{heatmap}}$ , and  $\lambda_{3D}$  balance the individual loss terms.  $\mathcal{L}_{\text{vis}}$  is the binary cross-entropy loss between the predicted and ground-truth visibility labels. The heatmap loss

$\mathcal{L}_{\text{heatmap}}$  and the 3D pose loss  $\mathcal{L}_{3D}$  are formulated as:

$$\mathcal{L}_{\text{heatmap}} = \frac{1}{2N_J} \sum_{j=1}^2 \sum_{i=1}^{N_J} w(s_{i,j}) \cdot \text{MSE}(\mathbf{H}_{i,j}, \bar{\mathbf{H}}_{i,j}) \quad (7)$$

$$\mathcal{L}_{3D} = \frac{1}{N_J} \sum_{i=1}^{N_J} \frac{w(s_{i,1}) + w(s_{i,2})}{2} \cdot \text{MSE}(\mathbf{J}_{Cam}^i, \bar{\mathbf{J}}_{Cam}^i) \quad (8)$$

where  $\mathbf{H}_{i,j}$  and  $\bar{\mathbf{H}}_{i,j}$  denote the predicted and ground-truth heatmaps for the  $i$ -th keypoint in the  $j$ -th camera view.  $\mathbf{J}_{Cam}^i$  and  $\bar{\mathbf{J}}_{Cam}^i$  represent the predicted and ground-truth 3D positions of the  $i$ -th keypoint in the camera coordinate system. The terms  $s_{i,1}$  and  $s_{i,2}$  indicate the visibility states of the  $i$ -th keypoint in the left and right image, respectively. The weighting function  $w(\cdot)$  assigns a value of 1.0 to visible keypoints and 0.1 to invisible ones, thereby reducing the influence of unreliable invisible keypoints during training.

In the second stage, we train the iterative intra-and inter-frame attention network using a combination of a joint position loss  $\mathcal{L}_{\text{joint}}$  and a smooth loss  $\mathcal{L}_{\text{smooth}}$ . The joint loss  $\mathcal{L}_{\text{joint}}$  is computed as the mean squared error between the predicted and ground-truth 3D joint positions. Following [11],  $\mathcal{L}_{\text{smooth}}$  is defined as the mean absolute error between the predicted and ground-truth joint accelerations.

## 5. Experiments

### 5.1. Dataset

**Eva-3M.** Our dataset comprises 1,353 motion sequences from 31 subjects performing 24 categories of actions common in VR scenarios. We reserve all sequences from two randomly chosen subjects (one male and one female, 98 sequences total) for testing, using the remainder for training. **EMHI [15].** We adhere to the official split, dividing the data into three subsets: 70% for training, and two separate test sets, i.e., P1 (16%) and P2 (14%). Notably, P2 contains unseen action categories and is designed specifically to evaluate the model’s generalization capability to out-of-distribution motions. We further annotate 488K frames in this dataset with keypoint visibility labels to support visibility-aware model training and evaluation.

### 5.2. Implementation Details and Metrics

**Implementation Details.** For the VQ-VAE module, we set the number of tokens to  $M = 160$ , and the codebook size to  $2048 \times 256$ . The VQ-VAE network is pre-trained on the AMASS [32], MOYO [40] and AIST++ [28] datasets, and remains frozen during the training of EvaPose. For EvaPose, we employ two image backbones: ResNet50 [20] and ViT-Large [13, 33], with the corresponding models denoted as EvaPose-ResNet50 and EvaPose-ViT-L. Input images are resized to  $448 \times 336$  for EvaPose-ViT-L, and kept at the original  $640 \times 480$  resolution for EvaPose-ResNet50. Training

proceeds in two stages. First, the visibility-aware 3D pose estimation module is trained with a batch size of 24 and learning rate of  $1 \times 10^{-5}$  for 20 epochs. The weights are set as  $\lambda_{\text{vis}} = 5 \times 10^{-3}$ ,  $\lambda_{\text{heatmap}} = 0.1$ , and  $\lambda_{3D} = 1.0$ . Then, the temporal 3D pose refinement network is trained with a batch size of 4 and a learning rate of  $1 \times 10^{-5}$  for 40 epochs, using a time window of  $T = 24$  frames. Further architectural details are provided in the supplementary material.

**Metrics.** As in prior works, we evaluate pose estimation accuracy using six metrics: mean per joint positional error (MPJPE), Procrustes-aligned MPJPE (PA-MPJPE), upper-body MPJPE (U-PE), lower-body MPJPE (L-PE), foot MPJPE (FootPE), and hand MPJPE (HandPE). All values are reported in millimeters. To assess motion smoothness, we adopt Jitter metric ( $10^2 m/s^3$ ) as in [11]. Inference speed is measured in frames per second (FPS), evaluated on a single NVIDIA V100 GPU under consistent settings.

### 5.3. Comparison with State-of-the-Art Methods

To ensure a rigorous comparison, we retrain three leading egocentric HPE methods, i.e., UnrealEgo [2], EgoPoseFormer [48], and FRAME [7], on both the EMHI and our Eva-3M datasets, adhering to their original experimental settings. As summarized in Tab. 2, our approach consistently outperforms prior methods across most key metrics on all test sets. Notably, on the EMHI P2 test set, which contains actions unseen during training, both EvaPose-ResNet50 and EvaPose-ViT-L achieve substantial improvements over the previous state-of-the-art, FRAME [7], in both MPJPE and Jitter metrics. These experimental results validate the superiority of our method not only in terms of accuracy and smoothness of pose estimation, but also in generalization capability to unseen actions when applying to real VR scenarios.

To further understand how EvaPose-ResNet50 and EvaPose-ViT-L outperform previous methods, we compare the detailed accuracy on visible and invisible keypoints using the Eva-3M dataset. Tab. 3 presents a breakdown for representative limb keypoints, from which we draw the following observations: (1) Previous methods yield comparable accuracy on visible limb keypoints, yet perform considerably worse than our approach. This suggests that indiscriminately treating visible and invisible keypoints degrades their estimation accuracy on visible ones. (2) Compared to FRAME [7], our EvaPose-ResNet50 and EvaPose-ViT-L bring more pronounced accuracy improvements on visible keypoints than on invisible ones, indicating that our EvaPose can effectively utilize visibility information to mitigate interference from invisible joints. (3) Ground-truth visibility annotations are highly valuable in egocentric HPE, where self-occlusion and out-of-FoV are common.

In terms of inference speed, EvaPose-ResNet50 runs at 48.0 FPS on an NVIDIA V100 GPU, enabling real-

Dataset	Methods	MPJPE↓	PA-MPJPE↓	U-PE↓	L-PE↓	FootPE↓	HandPE↓	Jitter↓	FPS↑
<b>Eva-3M</b>	UnrealEgo [2]	60.7	38.8	46.5	81.2	104.9	86.6	7.6	58.9
	EgoPoseFormer [48]	52.9	33.7	40.6	70.5	93.4	79.2	15.4	<u>67.0</u>
	FRAME [7]	49.8	35.1	42.4	60.5	77.4	80.2	<u>3.1</u>	<b>108.7</b>
	EvaPose-ResNet50 (Ours)	<u>35.6</u>	<u>24.7</u>	<u>28.3</u>	46.1	<u>58.4</u>	<u>49.6</u>	<b>3.0</b>	48.0
	EvaPose-ViT-L (Ours)	<b>34.2</b>	<b>24.0</b>	<b>27.2</b>	<b>44.5</b>	<b>56.4</b>	<b>49.2</b>	3.2	9.4
<b>EMHI: P1</b>	UnrealEgo [2]	55.2	38.6	39.9	77.3	100.2	85.2	5.9	58.9
	EgoPoseFormer [48]	50.7	29.9	32.9	76.5	96.6	51.0	7.3	<u>67.0</u>
	FRAME [7]	37.4	28.2	<u>28.0</u>	51.1	62.1	45.1	<b>2.4</b>	<b>108.7</b>
	EvaPose-ResNet50 (Ours)	36.2	28.0	30.0	45.3	<u>55.7</u>	<u>44.8</u>	<u>2.6</u>	48.0
	EvaPose-ViT-L (Ours)	<b>31.7</b>	<b>25.2</b>	<b>24.9</b>	<b>41.9</b>	<b>54.5</b>	<b>39.9</b>	3.0	9.4
<b>EMHI: P2</b>	UnrealEgo [2]	63.7	42.8	46.4	88.7	115.5	100.8	6.1	58.9
	EgoPoseFormer [48]	62.6	32.7	42.0	92.3	118.1	66.8	7.6	<u>67.0</u>
	FRAME [7]	60.5	44.3	55.8	67.4	78.6	75.5	6.4	<b>108.7</b>
	EvaPose-ResNet50 (Ours)	38.5	<u>29.5</u>	<u>31.4</u>	48.8	<u>61.9</u>	<u>57.8</u>	<b>3.1</b>	48.0
	EvaPose-ViT-L (Ours)	<b>33.3</b>	<b>26.2</b>	<b>25.4</b>	<b>44.7</b>	<b>58.9</b>	<b>47.5</b>	<u>3.4</u>	9.4

Table 2. Comparison with state-of-the-art methods on the Eva-3M and EMHI datasets. The best and second best results are highlighted in **boldface** and underlined, respectively.

Methods	Ankle		Foot		Elbow		Wrist		Mean	
	Visible	Invisible	Visible	Invisible	Visible	Invisible	Visible	Invisible	Visible	Invisible
UnrealEgo [2]	83.0	87.1	91.3	106.8	52.0	63.0	67.0	85.8	73.3	85.7
EgoPoseFormer [48]	78.2	94.2	83.9	96.4	49.1	63.2	57.9	101.0	67.3	88.7
FRAME [7]	77.8	79.1	81.2	92.5	53.0	59.5	67.9	86.5	70.0	79.4
EvaPose-ResNet50	<u>52.2</u>	<u>66.4</u>	<u>60.8</u>	<u>82.5</u>	<u>32.2</u>	48.0	38.0	<u>65.7</u>	45.8	<u>65.6</u>
EvaPose-ViT-L	<b>49.2</b>	<b>63.8</b>	<b>57.2</b>	<b>79.4</b>	<b>29.5</b>	<b>45.6</b>	<b>34.2</b>	<b>63.3</b>	<b>42.5</b>	<b>63.0</b>

Table 3. Performance comparisons of visible and invisible keypoints on the Eva-3M dataset.

Method	MPJPE↓	PA-MPJPE↓	VLK-PE↓	ILK-PE↓
w/o visibility	40.6	27.9	53.5	<b>65.3</b>
with visibility	<b>35.6</b>	<b>24.7</b>	<b>45.8</b>	65.6

Table 4. Impact of the visibility-aware model and loss design. VLK-PE is the MPJPE for visible limb keypoints, including ankle, foot, elbow and wrist. ILK-PE denotes the MPJPE for invisible limb keypoints.

TTE	STD	VQ-VAE	MPJPE↓	PA-MPJPE↓	Jitter↓
<b>x</b>	<b>x</b>	<b>x</b>	46.0	33.3	4.7
✓	<b>x</b>	<b>x</b>	40.3	27.0	<b>2.5</b>
✓	✓	<b>x</b>	39.1	25.9	2.7
✓	✓	✓	<b>35.6</b>	<b>24.7</b>	3.0

Table 5. Ablation study for key components in our model.

time performance despite being slower than prior methods. EvaPose-ViT-L, which employs a large ViT-Large backbone with 0.3B parameters, achieves the highest accuracy across all datasets at the cost of lower FPS, as expected.

The qualitative results in Fig. 3 further highlight our

method’s superior capability in estimating accurate and physically plausible 3D poses, including those of invisible keypoints. A significant enhancement is evident in the limb regions, which is consistent with the quantitative results in Tab. 2. We also provide more quantitative and qualitative comparison results in our supplementary materials.

## 5.4. Ablation Study

We conduct comprehensive ablation studies on the Eva-3M dataset using the ResNet50 image backbone to evaluate the contribution of each proposed module.

**Impact of the visibility-aware model and loss design.** As one of our key contributions, we introduce a visibility-aware 3D pose estimation network along with a visibility-based loss-weighting scheme to explicitly model keypoint visibility and leverage it for improved pose estimation accuracy. To validate their effectiveness, we perform an ablation by removing all visibility-related components. As reported in Tab. 4, incorporating the visibility-aware model and loss design significantly reduces MPJPE for visible limb keypoints, leading to notable improvements in both overall MPJPE and PA-MPJPE metrics. These results jointly



Figure 3. Qualitative comparisons on both EMHI (the first two rows) and Eva-3M (the last two rows) datasets. For better illustration, the predicted (red) and ground-truth (green) 3D poses are re-projected onto external reference views which are not used for pose estimation.

demonstrate the importance of visibility annotations and the efficacy of our visibility-aware modeling approach.

**Impact of the key components.** As described in Section 4, there are three key components in our model: the VQ-VAE pose prior, the Stereo Transformer Decoder (STD), and the Temporal Transformer Encoder (TTE). Tab. 5 summarizes our ablation study by removing these components sequentially to isolate their contributions. First, introducing the pose prior via VA-VAE decreases MPJPE from 39.1mm to 35.6mm, demonstrating that 3D pose priors learned from large-scale mocap datasets can help resolve ambiguities, particularly for invisible keypoints. Second, incorporating stereo visual features via the STD network can contribute to additional accuracy improvement. Third, the TTE module plays an important role in reducing both MPJPE and Jitter metrics, which confirms the importance of temporal fusion in improving smoothness and accuracy.

## 6. Conclusion

We introduced Eva-3M, a large-scale real egocentric visibility-aware HPE dataset. It provides visibility labels for the first time and surpasses existing datasets in motion diversity. Captured using a Pico4 Ultra VR-MR device, Eva-3M offers realistic insights into HPE in real VR-MR devices. We also proposed EvaPose, a novel egocentric visibility-aware HPE method explicitly leveraging the visibility information to improve the HPE accuracy. Extensive experiments validate that EvaPose achieves state-of-the-art performance on both the EMHI and Eva-3M benchmarks.

**Limitations.** As a data-driven approach, our method relies on large-scale, high-quality GT labels, which are challenging to acquire in in-the-wild scenarios. Exploring weakly-supervised or self-supervised training methods to improve generalization will be a promising direction for future work.

## References

- [1] Abien Fred Agarap. Deep learning using rectified linear units (relu). *arXiv preprint arXiv:1803.08375*, 2018. 3
- [2] Hiroyasu Akada, Jian Wang, Soshi Shimada, Masaki Takahashi, Christian Theobalt, and Vladislav Golyanik. Unrealego: A new dataset for robust egocentric 3d human motion capture. In *European Conference on Computer Vision*, pages 1–17. Springer, 2022. 2, 3, 6, 7
- [3] Hiroyasu Akada, Jian Wang, Vladislav Golyanik, and Christian Theobalt. 3d human pose perception from egocentric stereo videos. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 767–776, 2024. 2, 3
- [4] Hiroyasu Akada, Jian Wang, Vladislav Golyanik, and Christian Theobalt. Bring your rear cameras for egocentric 3d human pose estimation. In *International Conference on Computer Vision (ICCV)*, 2025. 2
- [5] Sadegh Aliakbarian, Fatemeh Saleh, David Collier, Pashmina Cameron, and Darren Cosker. Hmd-nemo: Online 3d avatar motion generation from sparse observations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9622–9631, 2023. 2
- [6] Md Mushfiqur Azam and Kevin Desai. A survey on 3d egocentric human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1643–1654, 2024. 1
- [7] Andrea Boscolo Camilletto, Jian Wang, Eduardo Alvarado, Rishabh Dabral, Thabo Beeler, Marc Habermann, and Christian Theobalt. Frame: Floor-aligned representation for avatar motion from egocentric video. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 17497–17507, 2025. 1, 2, 3, 5, 6, 7
- [8] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. 1
- [9] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 3
- [10] Hanz Cuevas-Velasquez, Charlie Hewitt, Sadegh Aliakbarian, and Tadas Baltrušaitis. Simpleego: Predicting probabilistic body pose from egocentric cameras. In *2024 International Conference on 3D Vision*, pages 1446–1455. IEEE, 2024. 2, 3
- [11] Peng Dai, Yang Zhang, Tao Liu, Zhen Fan, Tianyuan Du, Zhuo Su, Xiaozheng Zheng, and Zeming Li. Hmd-pose: On-device real-time human motion tracking from scalable sparse observations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 874–884, 2024. 2, 6
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 3
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 6, 3
- [14] Sai Kumar Dwivedi, Yu Sun, Priyanka Patel, Yao Feng, and Michael J Black. Tokenhmr: Advancing human mesh recovery with a tokenized pose representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1323–1333, 2024. 4, 5
- [15] Zhen Fan, Peng Dai, Zhuo Su, Xu Gao, Zheng Lv, Jiarui Zhang, Tianyuan Du, Guidong Wang, and Yang Zhang. Emhi: A multimodal egocentric human motion dataset with hmd and body-worn imus. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2879–2887, 2025. 1, 2, 3, 6
- [16] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, 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*, pages 18995–19012, 2022. 2
- [17] Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, 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*, pages 19383–19400, 2024. 2
- [18] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5152–5161, 2022. 4
- [19] Vladimir Guzov, Yifeng Jiang, Fangzhou Hong, Gerard Pons-Moll, Richard Newcombe, C Karen Liu, Yuting Ye, and Lingni Ma. Hmd 2: Environment-aware motion generation from single egocentric head-mounted device. In *2025 International Conference on 3D Vision (3DV)*, pages 1394–1405. IEEE, 2025. 2
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6, 3
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3
- [22] Fangzhou Hong, Vladimir Guzov, Hyo Jin Kim, Yuting Ye, Richard Newcombe, Ziwei Liu, and Lingni Ma. EgoLM: Multi-modal language model of egocentric motions. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5344–5354, 2025. 2
- [23] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015. 3
- [24] Jiayi Jiang, Paul Strelci, Huajian Qiu, Andreas Fender, Larissa Laich, Patrick Snape, and Christian Holz. Avatarposer: Ar-

- ticated full-body pose tracking from sparse motion sensing. In *European Conference on Computer Vision*, pages 443–460, 2022. 2
- [25] Taeho Kang and Youngki Lee. Attention-propagation network for egocentric heatmap to 3d pose lifting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 842–851, 2024. 2
- [26] Jihyun Lee, Weipeng Xu, Alexander Richard, Shih-En Wei, Shunsuke Saito, Shaojie Bai, Te-Li Wang, Minhyuk Sung, Tae-Kyun Kim, and Jason Saragih. Rewind: Real-time egocentric whole-body motion diffusion with exemplar-based identity conditioning. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 7095–7104, 2025. 1, 2, 3
- [27] Jiaman Li, Karen Liu, and Jiajun Wu. Ego-body pose estimation via ego-head pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17142–17151, 2023. 2
- [28] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 13401–13412, 2021. 6
- [29] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015. 1
- [30] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 851–866. 2023. 4
- [31] Lingni Ma, Yuting Ye, Fangzhou Hong, Vladimir Guzov, Yifeng Jiang, Rowan Postyeni, Luis Pesqueira, Alexander Gamino, Vijay Baiyya, Hyo Jin Kim, et al. Nymeria: A massive collection of multimodal egocentric daily motion in the wild. *arXiv preprint arXiv:2406.09905*, 2024. 2
- [32] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. 6
- [33] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 6, 3, 5
- [34] Kejie Qiu, Tong Qin, Jie Pan, Siqi Liu, and Shaojie Shen. Real-time temporal and rotational calibration of heterogeneous sensors using motion correlation analysis. *IEEE Transactions on Robotics*, 37(2):587–602, 2021. 1
- [35] Helge Rhodin, Christian Richardt, Dan Casas, Eldar Insafutdinov, Mohammad Shafiei, Hans-Peter Seidel, Bernt Schiele, and Christian Theobalt. Egocap: egocentric marker-less motion capture with two fisheye cameras. *ACM Transactions on Graphics (TOG)*, 35(6):1–11, 2016. 2, 3
- [36] Kiran Somasundaram, Jing Dong, Huixuan Tang, Julian Straub, Mingfei Yan, Michael Goesele, Jakob Julian Engel, Renzo De Nardi, and Richard Newcombe. Project aria: A new tool for egocentric multi-modal ai research. *arXiv preprint arXiv:2308.13561*, 2023. 2
- [37] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5693–5703, 2019. 1
- [38] Denis Tome, Patrick Peluse, Lourdes Agapito, and Hernan Badino. xr-egopose: Egocentric 3d human pose from an hmd camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7728–7738, 2019. 2, 3
- [39] Denis Tome, Thiemo Alldieck, Patrick Peluse, Gerard Pons-Moll, Lourdes Agapito, Hernan Badino, and Fernando De la Torre. Selfpose: 3d egocentric pose estimation from a head-set mounted camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6):6794–6806, 2020. 3
- [40] Shashank Tripathi, Lea Müller, Chun-Hao P Huang, Omid Taheri, Michael J Black, and Dimitrios Tzionas. 3d human pose estimation via intuitive physics. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4713–4725, 2023. 6
- [41] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 1, 4
- [42] Jian Wang, Lingjie Liu, Weipeng Xu, Kripasindhu Sarkar, and Christian Theobalt. Estimating egocentric 3d human pose in global space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11500–11509, 2021. 3
- [43] Jian Wang, Diogo Luvizon, Weipeng Xu, Lingjie Liu, Kripasindhu Sarkar, and Christian Theobalt. Scene-aware egocentric 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13031–13040, 2023. 1, 2, 3
- [44] Jian Wang, Zhe Cao, Diogo Luvizon, Lingjie Liu, Kripasindhu Sarkar, Danhang Tang, Thabo Beeler, and Christian Theobalt. Egocentric whole-body motion capture with fisheyevit and diffusion-based motion refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 777–787, 2024. 2, 5
- [45] Songpengcheng Xia, Yu Zhang, Zhuo Su, Xiaozheng Zheng, Zheng Lv, Guidong Wang, Yongjie Zhang, Qi Wu, Lei Chu, and Ling Pei. Envposer: Environment-aware realistic human motion estimation from sparse observations with uncertainty modeling. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1839–1849, 2025. 2
- [46] Weipeng Xu, Avishek Chatterjee, Michael Zollhoefer, Helge Rhodin, Pascal Fua, Hans-Peter Seidel, and Christian Theobalt. Mo 2 cap 2: Real-time mobile 3d motion capture with a cap-mounted fisheye camera. *IEEE transactions on visualization and computer graphics*, 25(5):2093–2101, 2019. 2, 3
- [47] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vitpose: Simple vision transformer baselines for human pose estimation. *Advances in neural information processing systems*, 35:38571–38584, 2022. 5, 3
- [48] Chenhongyi Yang, Anastasia Tkach, Shreyas Hampali, Linguang Zhang, Elliot J Crowley, and Cem Keskin. Egopose-

- former: A simple baseline for stereo egocentric 3d human pose estimation. In *European conference on computer vision*, pages 401–417. Springer, 2024. [1](#), [3](#), [5](#), [6](#), [7](#), [2](#)
- [49] Jingkang Yang, Shuai Liu, Hongming Guo, Yuhao Dong, Xiamengwei Zhang, Sicheng Zhang, Pengyun Wang, Zitang Zhou, Binzhu Xie, Ziyue Wang, et al. Egolife: Towards egocentric life assistant. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 28885–28900, 2025. [2](#)
- [50] Hanrong Ye, Haotian Zhang, Erik Daxberger, Lin Chen, Zongyu Lin, Yanghao Li, Bowen Zhang, Haoxuan You, Dan Xu, Zhe Gan, et al. Mm-ego: Towards building egocentric multimodal llms for video qa. *arXiv preprint arXiv:2410.07177*, 2024. [2](#)
- [51] Brent Yi, Vickie Ye, Maya Zheng, Yunqi Li, Lea Müller, Georgios Pavlakos, Yi Ma, Jitendra Malik, and Angjoo Kanazawa. Estimating body and hand motion in an ego-sensed world. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 7072–7084, 2025. [2](#)
- [52] Amy Zhao, Chengcheng Tang, Lezi Wang, Yijing Li, Mihika Dave, Lingling Tao, Christopher D Twigg, and Robert Y Wang. Egobody3m: Egocentric body tracking on a vr headset using a diverse dataset. In *European Conference on Computer Vision*, pages 375–392. Springer, 2024. [2](#), [3](#)
- [53] Dongxu Zhao, Zhen Wei, Jisan Mahmud, and Jan-Michael Frahm. Egoglass: Egocentric-view human pose estimation from an eyeglass frame. In *2021 International Conference on 3D Vision (3DV)*, pages 32–41. IEEE, 2021. [3](#)
- [54] Xiaozheng Zheng, Zhuo Su, Chao Wen, Zhou Xue, and Xiaojie Jin. Realistic full-body tracking from sparse observations via joint-level modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14678–14688, 2023. [2](#)