

3D Human Pose Perception from Egocentric Stereo Videos

Hiroyasu Akada

Jian Wang

Vladislav Golyanik

Christian Theobalt

Max Planck Institute for Informatics, SIC




Figure 1. **3D human pose estimation results of our proposed method from egocentric stereo fisheye videos.** **Left:** results on synthetic images; (a) reference RGB view of the scene; (b) 3D-to-2D pose re-projections, and (c) a 3D pose in a scene mesh reconstructed by our framework. **Right:** results on real-world images; (d) reference view; (e) 3D-to-2D pose re-projections; (f) a 3D pose in the reconstructed scene, and (g) 3D virtual character animation (possible future application of our method).

Abstract

While head-mounted devices are becoming more compact, they provide egocentric views with significant self-occlusions of the device user. Hence, existing methods often fail to accurately estimate complex 3D poses from egocentric views. In this work, we propose a new transformer-based framework to improve egocentric stereo 3D human pose estimation, which leverages the scene information and temporal context of egocentric stereo videos. Specifically, we utilize 1) depth features from our 3D scene reconstruction module with uniformly sampled windows of egocentric stereo frames, and 2) human joint queries enhanced by temporal features of the video inputs. Our method is able to accurately estimate human poses even in challenging scenarios, such as crouching and sitting. Furthermore, we introduce two new benchmark datasets, i.e., *UnrealEgo2* and *UnrealEgo-RW* (RealWorld). The proposed datasets offer a much larger number of egocentric stereo views with a wider variety of human motions than the existing datasets, allowing comprehensive evaluation of existing and upcoming methods. Our extensive experiments show that the proposed approach significantly outperforms previous methods. *UnrealEgo2*, *UnrealEgo-RW*, and trained models are available on our project page¹ and Benchmark Challenge².

¹<https://4dqv.mpi-inf.mpg.de/UnrealEgo2/>

²<https://unrealego.mpi-inf.mpg.de/>

1. Introduction

Egocentric 3D human motion capture using wearable devices has received increased attention recently [1, 11, 22, 31, 37, 38, 40–42, 45, 48, 52, 53]. Different from traditional vision-based motion capture setups that require a fixed recording space, egocentric systems allow flexible motion capture in less constrained situations. Therefore, the egocentric setups offer various applications, such as motion analysis and XR technologies (Fig. 1-(g)).

Previous works proposed various egocentric methods to capture device users. On the one hand, the vast majority of existing methods—which use a monocular camera—would fail for complex human poses due to depth ambiguity and self-occlusion. On the other hand, the methods designed for stereo devices do not yet realize the full potential of their stereo settings, especially with the most recent compact eyeglasses-based setups [1, 53]. Specifically, they do not deliver high 3D reconstruction accuracy across different scenarios. Moreover, these approaches do not consider scene information, which further limits their accuracy.

To address the challenges outlined above, we propose a new transformer-based framework for egocentric 3D human motion capture from compact eyeglasses-based devices; see Fig. 1. The first step of our framework is to estimate 2D joint heatmaps from egocentric stereo fisheye RGB videos (Sec. 4.1). These 2D joint heatmaps are then processed with human joint queries in our transformer-based 3D mod-

ule to estimate 3D poses. Here, we leverage the scene information and temporal context of the input videos in the 3D module to improve estimation accuracy. Firstly, we use uniformly sampled windows of egocentric stereo frames to reconstruct a 3D background scene using Structure from Motion (SfM) [33], obtaining scene depth as additional information for the 3D module (Sec. 4.2 and 4.3). In our challenging eyeglasses-based setup, however, the 3D scene and camera poses can not always be estimated due to severe self-occlusion in the egocentric images. This results in depth maps with zero (invalid) values and undesired computation of network gradients during training. To mitigate this issue, we propose to use depth padding masks that prevent processing such invalid depth values in the 3D module. Additionally, we propose video-dependent query augmentation that enhances the joint queries with the temporal context of stereo video inputs to effectively capture the temporal relation of human motions at a joint level (Sec. 4.4).

We also introduce two new benchmark datasets: *UnrealEgo2* and *UnrealEgo-RW*. *UnrealEgo2* is an extended version of *UnrealEgo* [1] and the largest eyeglasses-based synthetic data with various new motions, offering $2.8\times$ larger data (2.5M images) than the existing dataset [1]. *UnrealEgo-RW* is a real-world dataset recorded with our newly developed device that resembles the virtual eyeglasses-based setup [1], offering 260k images with various motions and 3D poses. The proposed datasets make it possible to evaluate existing and upcoming methods on a variety of motions, not only in synthetic scenes but also in real-world cases.

In short, the contributions of this paper are as follows:

- The transformer-based framework for egocentric stereo 3D human pose estimation that accounts for temporal context in egocentric stereo views.
- 3D pose estimation is enhanced via the utilization of scene information from our video-based 3D scene reconstruction module as well as joint queries obtained from our video-dependent query augmentation policy.
- A new portable device for egocentric stereo view capture with its specification and two new benchmark datasets: *UnrealEgo2* and *UnrealEgo-RW* recorded with our device. The proposed datasets allow for a comprehensive evaluation of methods for egocentric 3D human pose estimation from stereo views.

Our experiments demonstrate that the proposed method outperforms the previous state-of-the-art approaches by a substantial margin, *i.e.*, $>15\%$ on *UnrealEgo* [1], $\geq 40\%$ on *UnrealEgo2*, and $\geq 10\%$ on *UnrealEgo-RW* (on MPJPE). We release *UnrealEgo2*, *UnrealEgo-RW*, and our trained models on our project page³ and Benchmark Challenge⁴ to foster the area of egocentric 3D vision.

³<https://4dqv.mpi-inf.mpg.de/UnrealEgo2/>

⁴<https://unrealego.mpi-inf.mpg.de/>

2. Related Work

Egocentric 3D Human Motion Capture. Recent years witnessed significant innovations in egocentric 3D human pose estimation. To capture device users, many existing works use downward-facing cameras and the existing methods can be categorized into two groups. The first group are monocular approaches [11, 21, 22, 27, 37, 38, 40, 41, 43, 45, 48, 52]. For example, Wang *et al.* [43] uses a diffusion-based [10] motion prior to tackle self-occlusions. Due to the depth ambiguity, monocular methods often fail to estimate accurate 3D poses. Wang *et al.* [42] tackled this issue by projecting depth and 2D pose features into a pre-defined voxel space. This method requires additional training with ground-truth depths and human body segmentation; it cannot easily be extended for multi-view or temporal inputs. Zhang *et al.* [51] utilized a diffusion model [10] conditioned on a 3D scene to generate poses. They require pre-scanned scene mesh as an input and cannot capture a device user.

The second group, including our work, focuses on the multi-view (often stereo) setting. Rhodin *et al.* [31] proposed an optimization approach whereas Cha *et al.* [3] used eight cameras to estimate a 3D body and reconstruct a 3D scene separately. Other works [1, 53] used the multi-branch autoencoder [37] to the stereo setup. Kang *et al.* [12] (arXiv pre-print at the time of submission) leveraged a stereo-matching mechanism and perspective embedding heatmaps. In contrast to the existing methods, we propose a new transformer-based method that effectively utilizes egocentric stereo videos via our video-based 3D scene reconstruction module and video-dependent query augmentation policy. Our method considers the scene information without the supervision of the scene data.

Transformers in 3D Human Pose Estimation from External Cameras. 3D pose estimation from external cameras has shown significant progress due to the advances in transformer architectures [39]. Some works [20, 47] predict 3D human pose and mesh from monocular views. Other works [5, 18, 19, 28, 29, 36, 46, 49, 54–58] present a 2D-to-3D lifting module that estimates 3D poses from monocular 2D joints obtained with off-the-shelf 2D joint detectors. Although their lifting modules show impressive results, those monocular methods cannot be easily applied to our stereo setting. On the other hand, some works utilize transformers in multi-view settings. He *et al.* [9] and Ma *et al.* [23] aggregate stereo information on epipolar lines of stereo images, which are difficult to obtain from fisheye images. Recent work [44] regresses multi-person 3D poses from multi-view inputs, powered by projective attention and query adaptation. However, no existing works explored the potential of transformers along with 2D joint heatmaps or explicit scene information in stereo 3D pose estimation. In this paper, we propose a transformer-based framework that accounts for the temporal relation of human motion at a joint level via

intermediate 2D joint heatmap and depth maps even with inaccurate depth values mixed in the framework.

Datasets for Egocentric 3D Human Pose Estimation.

Several works proposed unique setups to create datasets, using a monocular camera [11, 17, 22, 37, 40, 41, 45, 48] and forward-facing cameras [11, 14, 17, 22, 26, 48, 50, 51]. There also exist datasets captured with stereo devices [3, 7, 14, 26, 31, 53]. However, they are small [31] with limited motion types [31, 53], not publicly available [3, 53], or do not provide ground truth 3D poses of device users [7, 14, 26]. Recently, Akada *et al.* [1] introduced UnrealEgo, a synthetic dataset based on virtual eyeglasses with two fisheye cameras. However, they provide only synthetic images. Meanwhile, more glasses-based stereo datasets that offer a wider variety of motions or real-world footage are required nowadays for an extensive evaluation of existing and upcoming methods. Hence, we introduce two new benchmark datasets that in their characteristics go beyond the existing data: *UnrealEgo2* and *UnrealEgo-RW*. We describe the proposed datasets in the following section.

3. Mobile Device and Datasets

We present two new datasets for egocentric stereo 3D motion capture: *UnrealEgo2* and *UnrealEgo-RW*; see Fig. 1. Please watch our supplementary video for visualizations.

UnrealEgo2 Dataset. To create *UnrealEgo2* (an extension of *UnrealEgo* [1]), we adapt the publicly available setup with a virtual eyeglasses device [1]. This setup comes with two downward-facing fisheye cameras attached 12cm apart from each other on the glasses frames. The camera’s field of view is 170°. With this device, we capture 17 realistic 3D human models [30] animated by the Mixamo [25] dataset in various 3D environments. We record simple to highly complex motions such as crouching and crawling, for 14 hours.

Overall, *UnrealEgo2* offers 15,207 motions and >1.25M stereo views (2.5M images) as well as depth maps with a resolution 1024×1024 pixel rendered at 25 frames per second. Each frame is annotated with 32 body and 40 hand joints. Note that *UnrealEgo2* is the largest glasses-based dataset and 2.8× larger than *UnrealEgo*. Also, it does not share the same motions with *UnrealEgo*, providing a larger motion variety for a comprehensive evaluation.

Design of Our Mobile Device. Evaluation with real-world datasets plays a pivotal role in computer vision research. Therefore, we build a new portable device; see Fig. 2. Our device is based on a helmet with two RIBCAGE RX0 II cameras [32] and two FUJINON FE185C057HA-1 fisheye lenses [6]. We placed the cameras 12cm away from each other and 2cm away from user’s face. We cropped the margins of the egocentric images to resemble the field of view of 170° of the *UnrealEgo* and *UnrealEgo2* setups. Note that our setup is more compact than *EgoCap* [31] that placed cameras 25cm away from user’s face.



Figure 2. Our portable setup to acquire *UnrealEgo-RW*.

UnrealEgo-RW (Real-World) Dataset. With our device, we record various motions of 16 identities in a multi-view motion capture studio (Fig. 1-(d)). We capture simple and challenging activities, *e.g.*, crawling and dancing, for 1.5 hours. This is in strong contrast to the existing real-world stereo dataset [53] (not publicly available) that records only three simple actions, *i.e.*, sitting, standing, and walking.

In total, we obtained 591 motion segments from 16 identities with various textured clothing. This results in more than 130k stereo views (260k images) of a resolution 872×872 pixel rendered at 25 frames per second with ground-truth 3D poses of 16 joints. Note that *UnrealEgo-RW* offers 4.3× larger data with a wider variety of motions than the publicly available real-world stereo data [31].

4. Method

We propose a new framework for egocentric stereo 3D human pose estimation as shown in Fig. 3. Our framework first estimates the 2D joint heatmaps from egocentric stereo fisheye videos in our 2D module (Sec. 4.1). The heatmaps and input videos are then processed in our segmentation module to obtain 2D human body masks (Sec. 4.2). Next, we use uniformly sampled windows of input frames and human body masks to reconstruct 3D scenes (Sec. 4.3). Here, we render depth maps and depth region masks from the reconstructed mesh. Finally, our transformer-based 3D module processes the joint heatmaps, depth information, and joint queries to estimate 3D poses (Sec. 4.4). Here, the 3D module leverages depth padding masks based on the availability of the depth maps as well as joint queries enhanced by the stereo video features from the 2D module.

4.1. 2D Pose Estimation

Given egocentric stereo videos with T frames $\{\mathbf{I}_{\text{Left}}^t, \mathbf{I}_{\text{Right}}^t \in \mathbb{R}^{H \times W \times 3} | t = 1, 2, \dots, T\}$, we use the existing stereo 2D joint heatmap estimator [1] to obtain a sequence of corresponding 2D heatmaps of 15 joints $\{\mathbf{H}_{\text{Left}}^t, \mathbf{H}_{\text{Right}}^t \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 15}\}$, including the neck, upper arms, lower arms, hands, thighs, calves, feet, and balls of the feet. We also extract intermediate feature maps $\{\mathbf{F}_{\text{Left}}^t, \mathbf{F}_{\text{Right}}^t \in \mathbb{R}^{\frac{H}{32} \times \frac{W}{32} \times C}\}$ where $C = 512$, which are used later in the 3D module.

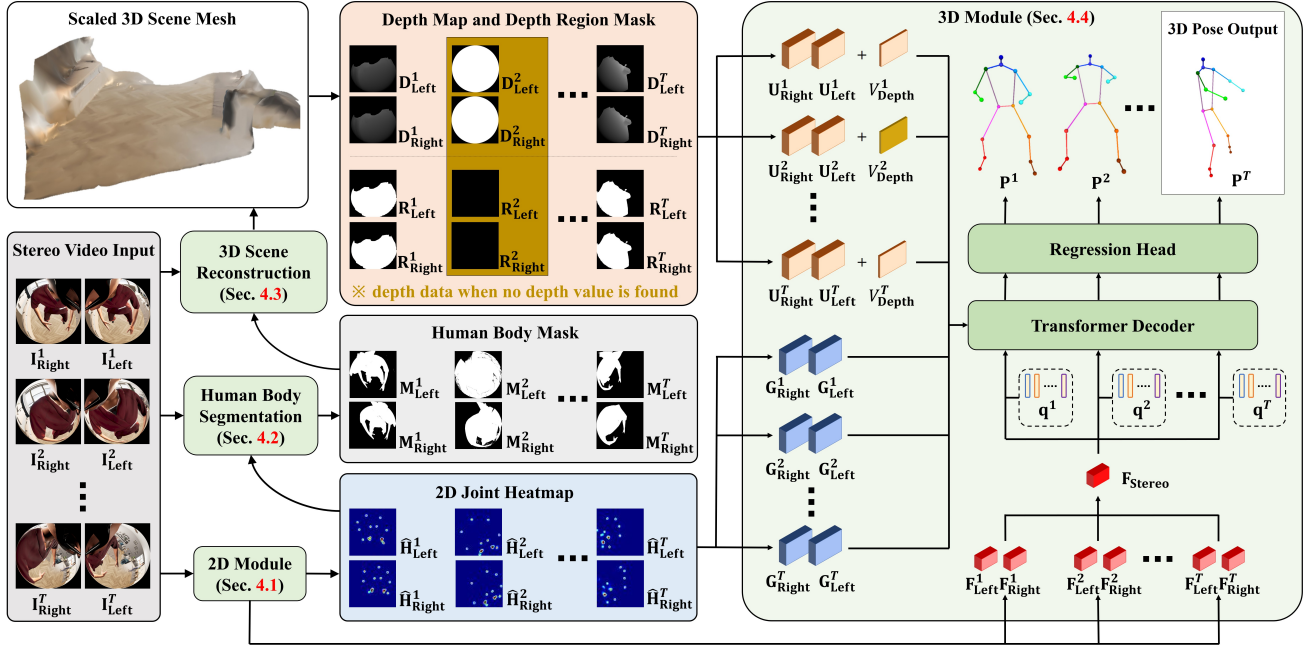


Figure 3. **Overview of our framework.** Our method takes egocentric stereo videos $\{\mathbf{I}_{\text{Left}}^t, \mathbf{I}_{\text{Right}}^t\}$ as inputs. We first apply the 2D module to obtain 2D joint heatmaps $\{\mathbf{H}_{\text{Left}}^t, \mathbf{H}_{\text{Right}}^t\}$ and video features $\{\mathbf{F}_{\text{Left}}^t, \mathbf{F}_{\text{Right}}^t\}$ (Sec. 4.1). The heatmaps are used with input videos to create human body masks $\{\mathbf{M}_{\text{Left}}^t, \mathbf{M}_{\text{Right}}^t\}$ (Sec. 4.2). Next, we use uniformly sampled windows of input frames and human body masks to reconstruct a 3D scene mesh (Sec. 4.3). From the mesh, we generate depth maps $\{\mathbf{D}_{\text{Left}}^t, \mathbf{D}_{\text{Right}}^t\}$ and depth region masks $\{\mathbf{R}_{\text{Left}}^t, \mathbf{R}_{\text{Right}}^t\}$. Note that this diagram shows an example case of missing depth values for the second input frame. Lastly, the depth data, 2D joint heatmaps, video features, joint queries q^t and the padding masks V_{Depth}^t are processed in the 3D module to estimate 3D poses \mathbf{P}^t (Sec. 4.4).

4.2. Human Body Segmentation

To reconstruct 3D scenes from egocentric videos, it is necessary to identify the pixels corresponding to the background environment. Therefore, we integrate an existing segmentation method, *i.e.*, ViT-H SAM model [16], as our segmentation network \mathcal{F}_{SAM} . In this module, we firstly obtain 2D joint locations from the 2D joint heatmap $\{\hat{\mathbf{H}}_{\text{Left}}^t, \hat{\mathbf{H}}_{\text{Right}}^t\}$. Then, we use the input video frames $\{\mathbf{I}_{\text{Left}}^t, \mathbf{I}_{\text{Right}}^t\}$ and its corresponding 2D joints to extract a human body mask $\{\mathbf{M}_{\text{Left}}^t, \mathbf{M}_{\text{Right}}^t \in \mathbb{R}^{H \times W \times 1}\}$:

$$\mathbf{M}_{\text{Left}}^t = \mathcal{F}_{\text{SAM}}(\mathbf{I}_{\text{Left}}^t, \hat{\mathbf{H}}_{\text{Left}}^t). \quad (1)$$

The same process can be applied to obtain $\mathbf{M}_{\text{Right}}^t$. Note that we use the SAM model without re-training on ground-truth human body masks. Instead, we guide the predictions of SAM using joint positions extracted from the 2D heatmaps.

4.3. 3D Scene Reconstruction

We aim to reconstruct 3D environments from uniformly sampled windows of input frames $\{\mathbf{I}_{\text{Left}}^t, \mathbf{I}_{\text{Right}}^t\}$ and human body masks $\{\mathbf{M}_{\text{Left}}^t, \mathbf{M}_{\text{Right}}^t\}$ with a fixed length. The length is set to 4 seconds (some motion data contains shorter sequences). Given these data, we use Metashape [24] to perform SfM to obtain camera poses and a 3D scene

mesh. Here, as the baseline length between stereo cameras is known, *i.e.*, 12cm, we can obtain the mesh in the real-world scale. Next, we render down-sampled depth maps $\{\mathbf{D}_{\text{Left}}^t, \mathbf{D}_{\text{Right}}^t \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 1}\}$ and depth region masks $\{\mathbf{R}_{\text{Left}}^t, \mathbf{R}_{\text{Right}}^t \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 1}\}$ from the reconstructed 3D scene mesh. The depth region masks show the regions where the depth values are obtained from the 3D scene. This depth information will be used later in the 3D module as additional cues for pose estimation. However, there are some cases where the egocentric RGB videos are largely occupied by a human body. In such scenarios, the 3D scene can not be reconstructed or camera poses can not be estimated. This results in missing (invalid) depth values and undesired computation of network gradients during training. Therefore, we tackle this issue in our 3D module.

4.4. 3D Pose Estimation

In the 3D module, we aim to estimate a sequence of 3D poses by considering scene information and the temporal context of the egocentric stereo videos. Specifically, given the 2D joint heatmaps, depth maps, depth region masks, and T sets of joint queries $q^t \in \mathbb{R}^{16 \times \frac{C}{2}}$, we use a transformer decoder to estimate a sequence of 3D poses $\{\mathbf{P}^t \in \mathbb{R}^{16 \times 3} | t = 1, 2, \dots, T\}$. Our pose output is the 3D pose at the last time step \mathbf{P}^T . We follow the existing

Method	Task	MPJPE(↓)	PA-MPJPE(↓)	3D PCA(↑)	AUC(↑)
Zhao <i>et al.</i> [53]	Pelvis relative	86.45	63.71	85.97	50.50
Akada <i>et al.</i> [1]		78.98	59.30	88.81	54.31
Kang <i>et al.</i> [12]		60.82	<u>48.47</u>	-	-
Baseline		<u>59.85</u>	49.14	<u>92.07</u>	<u>63.88</u>
Ours		50.55	40.50	93.83	70.61
Zhao <i>et al.</i> [53]	Device relative	88.12	65.36	85.10	50.37
Akada <i>et al.</i> [1]		84.53	63.92	87.05	52.76
Baseline		<u>63.44</u>	<u>50.97</u>	<u>92.30</u>	<u>64.54</u>
Ours		46.20	40.19	94.02	73.53

Table 1. Quantitative results on UnrealEgo [1] with mm-scale.

Method	MPJPE(↓)	PA-MPJPE(↓)	3D PCA(↑)	AUC(↑)
Zhao <i>et al.</i> [53]	79.64	58.22	88.50	53.82
Akada <i>et al.</i> [1]	72.80	52.88	91.32	55.81
Baseline	<u>52.23</u>	<u>39.78</u>	<u>95.72</u>	<u>68.13</u>
Ours	30.53	26.72	97.22	80.75

Table 2. Quantitative results of device-relative pose estimation on UnrealEgo2 with mm-scale.

works [1, 37, 38] to estimate 16 joints including the head. **Depth and Heatmap Features.** We use the sequence of the depth maps, depth region masks, and the 2D joint heatmaps as the memory of a cross-attention operation in the transformer decoder. For this purpose, we extract depth features $\{\mathbf{U}_{\text{Left}}^t, \mathbf{U}_{\text{Right}}^t \in \mathbb{R}^{\frac{H}{32} \times \frac{W}{32} \times \frac{C}{2}}\}$ from the depth data:

$$\mathbf{U}_{\text{Left}}^t = \mathcal{F}_{\text{Depth}}(\mathbf{D}_{\text{Left}}^t \oplus \widehat{\mathbf{R}}_{\text{Left}}^t), \quad (2)$$

where “ \oplus ” is a concatenation operation along the channel axis and $\mathcal{F}_{\text{Depth}}$ represents a feature extractor. The same process can be applied to obtain $\mathbf{U}_{\text{Right}}^t$.

Similarly, we extract heatmap features $\{\mathbf{G}_{\text{Left}}^t, \mathbf{G}_{\text{Right}}^t \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times C}\}$ from the 2D heatmaps:

$$\mathbf{G}_{\text{Left}}^t = \mathcal{F}_{\text{HM}}(\widehat{\mathbf{H}}_{\text{Left}}^t), \quad (3)$$

where \mathcal{F}_{HM} represents another feature extractor. The same process can be applied to obtain $\mathbf{G}_{\text{Right}}^t$.

These features are forwarded with positional embeddings into the transformer. However, as mentioned in Sec. 4.3, depth values can be missing in some frames. To prevent processing features of such depth data and let the network focus only on valid frames, we propose to add padding masks $V_{\text{Depth}}^t \in \mathcal{R}$ to all the elements of $\{\mathbf{U}_{\text{Left}}^t, \mathbf{U}_{\text{Right}}^t\}$:

$$V_{\text{Depth}}^t = \begin{cases} -\text{inf}, & \text{if depth values are missing} \\ 0, & \text{otherwise} \end{cases}. \quad (4)$$

When $V_{\text{Depth}}^t = -\text{inf}$, the depth features $\{\mathbf{U}_{\text{Left}}^t, \mathbf{U}_{\text{Right}}^t\}$ after the softmax function in self-attention layers of the transformer will have zero effect on the network training.

Stereo-Video-Dependent Joint Query Adaptation. The existing work [44] represents human joints as learnable positional embeddings called joint queries that encode prior

Method	MPJPE(↓)	PA-MPJPE(↓)	3D PCA(↑)	AUC(↑)
Zhao <i>et al.</i> [53]	117.57	88.01	73.12	38.94
Akada <i>et al.</i> [1]	122.64	86.55	72.51	38.67
Baseline	<u>115.95</u>	<u>85.00</u>	<u>74.13</u>	<u>40.11</u>
Ours	104.14	82.18	80.20	46.22

Table 3. Quantitative results of device-relative pose estimation on UnrealEgo-RW with mm-scale.

knowledge about the skeleton joints. In our problem setting, the simplest way to design such joint queries is to set queries for each pose in a motion sequence. However, this can not capture the temporal context in video inputs, *e.g.*, human motions and background changes. Therefore, we extend the multi-view joint query augmentation technique [44] for our stereo video setting to account for sequential information. Specifically, we enhance the joint queries with the temporal intermediate features of stereo RGB frames $\{\mathbf{F}_{\text{Left}}^t, \mathbf{F}_{\text{Right}}^t\}$. Firstly, from the sequence of the intermediate features, we create a sequence of combined features $\mathbf{F}^t \in \mathbb{R}^{\frac{H}{32} \times \frac{W}{32} \times \frac{C}{2}}$:

$$\mathbf{F}^t = \text{conv}(\mathbf{F}_{\text{Left}}^t \oplus \mathbf{F}_{\text{Right}}^t), \quad (5)$$

where “ $\text{conv}(\cdot)$ ” is a convolution operation with a kernel size of 1×1 .

Next, we fuse the sequence of the combined features \mathbf{F}^t to obtain a fused stereo features $\mathbf{F}_{\text{Stereo}} \in \mathbb{R}^{\frac{C}{2}}$:

$$\mathbf{F}_{\text{Stereo}} = \mathbf{F}_{\text{p}}^1 \oplus \dots \oplus \mathbf{F}_{\text{p}}^T, \text{ where } \mathbf{F}_{\text{p}}^i = p(\mathbf{F}^i), \quad (6)$$

where “ $p(\cdot)$ ” is an operation of adaptive average pooling. Now, the feature $\mathbf{F}_{\text{Stereo}}$ contains stereo video information.

Lastly, with $\mathbf{F}_{\text{Stereo}}$ and a fully connected layer “ $\text{fc}(\cdot)$ ”, we augment each query q^t to obtain $q_{\text{Aug}}^t \in \mathbb{R}^{16 \times \frac{C}{2}}$:

$$\mathbf{q}_{\text{Aug}}^t = \text{fc}(\mathbf{F}_{\text{Stereo}}) + q^t. \quad (7)$$

Transformer Decoder. We adopt a DETR [2]-based transformer decoder and a pose regression head. In decoder layers, all of the augmented joint queries q_{Aug}^t first interact with each other on a self-attention layer. Then, the queries extract all of the temporal stereo features from the memory $\{\mathbf{U}_{\text{Left}}^t, \mathbf{U}_{\text{Right}}^t, \mathbf{G}_{\text{Left}}^t, \mathbf{G}_{\text{Right}}^t\}$ with the padding masks V_{Depth}^t on a cross-attention layer. Lastly, the pose regression head estimates a sequence of 3D poses $\{\hat{\mathbf{P}}^t \in \mathbb{R}^{16 \times 3} | t = 1, 2, \dots, T\}$, yielding the final pose output \mathbf{P}^T .

Similar to the previous works [5, 49], we train the 3D module with the pose supervision of the current and past frames:

$$L_{3\text{D}} = L_{\text{pose}}(\mathbf{P}^T, \hat{\mathbf{P}}^T) + \frac{\lambda_{\text{past}}}{(T-1)} \sum_{t=1}^{T-1} L_{\text{pose}}(\mathbf{P}^t, \hat{\mathbf{P}}^t), \quad (8)$$

$$L_{\text{pose}}(\mathbf{P}, \hat{\mathbf{P}}) = \lambda_{\text{pose}} (\text{mpjpe}(\mathbf{P}, \hat{\mathbf{P}}) + \lambda_{\text{cos}} \cos(\text{bone}(\mathbf{P}), \text{bone}(\hat{\mathbf{P}}))), \quad (9)$$

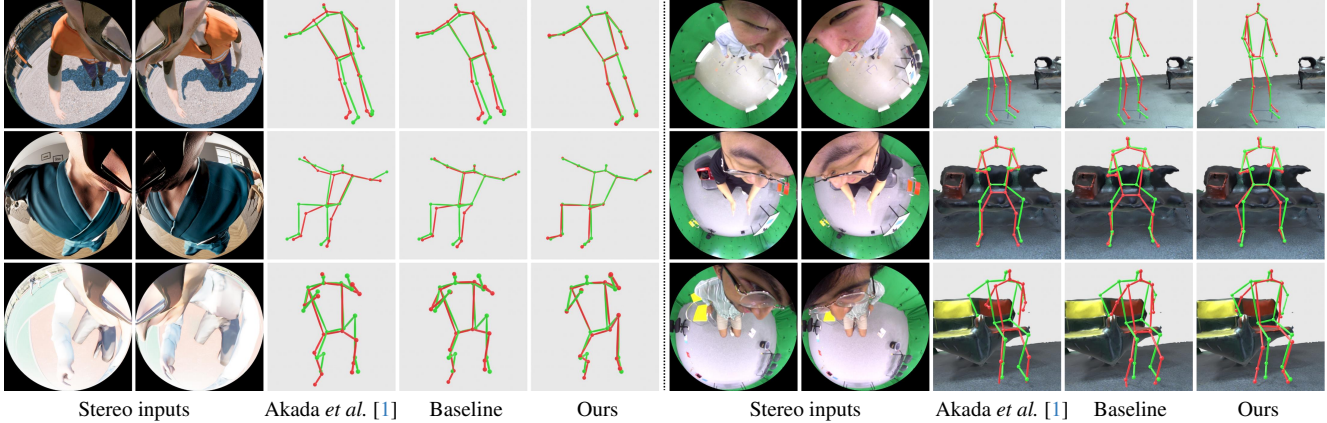


Figure 4. Qualitative results of device-relative pose estimation. **Left:** UnrealEgo2. **Right:** UnrealEgo-RW. 3D pose prediction and ground truth are displayed in red and green, respectively. For UnrealEgo-RW, we show ground-truth scene meshes for visualization.

Method	MPJPE(\downarrow)	PA-MPJPE(\downarrow)
(a) Baseline with depth information	120.39	86.23
Baseline	115.36	84.80
(b) Ours w/o query adaptation	108.33	86.69
(c) Ours w/o depth information	112.56	84.37
(d) Ours w/o depth padding mask	108.70	84.26
(e) Ours with latest pose supervision only	105.67	83.46
(f) Ours with a single set of queries	105.58	85.68
Ours	104.14	82.18

Table 4. Ablation study of our model for device-relative pose estimation on UnrealEgo-RW with mm-scale.

where \mathbf{P} is a ground-truth 3D pose, $\text{mpjpe}(\cdot)$ is the mean per joint position error, $\cos(\cdot)$ is a negative cosine similarity, and $\text{bone}(\cdot)$ is an operation of obtaining bones of the 3D poses as used in the previous work [1]:

$$\text{mpjpe}(\mathbf{P}, \hat{\mathbf{P}}) = \frac{1}{NJ} \sum_{n=1}^N \sum_{j=1}^J \|\mathbf{P}_{n,j} - \hat{\mathbf{P}}_{n,j}\|_2, \quad (10)$$

$$\cos(\mathbf{B}, \hat{\mathbf{B}}) = -\frac{1}{N} \sum_{n=1}^N \sum_{m=1}^M \frac{\mathbf{B}_{n,m} \cdot \hat{\mathbf{B}}_{n,m}}{\|\mathbf{B}_{n,m}\| \|\hat{\mathbf{B}}_{n,m}\|}, \quad (11)$$

where N is batch size, J is the number of joints, M is the number of bones, and $\mathbf{B}_{n,m} \in \mathbb{R}^3$ is a vector of m -th bone.

5. Experiments

5.1. Datasets for Evaluation

We use three datasets for our experiments: UnrealEgo [1], UnrealEgo2, and UnrealEgo-RW. For UnrealEgo, we use their proposed data splits. Also, we divide UnrealEgo2 into 12,139 motions (1,002,656 stereo views) for training, 1,545 motions (127,968 stereo views) for validation, and 1523 motions (123,488 stereo views) for testing. Similarly, we

Method	Upper body MPJPE(\downarrow)	Lower body MPJPE(\downarrow)	Foot MPJPE(\downarrow)	Foot MPE(\downarrow)
Ours w/o depth information	80.82	144.31	174.45	6.39
Ours w/o depth padding masks	77.29	140.10	169.95	5.02
Ours	<u>77.85</u>	130.97	155.86	4.83

Table 5. The effect of scene information (depth) per body part on UnrealEgo-RW. The numbers are in *mm*.

split UnrealEgo-RW into 547 motions (51,936 stereo views) for training, 77 motions (7,616 stereo views) for validation, and 86 motions (7,936 stereo views) for testing. we follow the existing works [1, 11, 37, 38, 40–42, 45, 52, 53] to report the results of device-relative 3D pose estimation. For UnrealEgo, we also follow the existing works [1, 12] to include the results of pelvis-relative 3D pose estimation.

5.2. Training Details

We resize the input RGB images and ground-truth 2D keypoint heatmaps to 256×256 and 64×64 pixels, respectively. For the training of the 2D module, we follow the previous work [1] to use the ResNet18 [8] pre-trained on ImageNet [4] as an encoder and train the module with a batch size of 16 and an initial learning rate of 10^{-3} . Then, we train the 3D module with a batch size of 32 and an initial learning rate of $2 \cdot 10^{-4}$. The modules are trained with Adam optimizer [15] for ten epochs, starting with the initial learning rate for the first half epochs and applying a linearly decaying rate for the next half. Also, we set the hyperparameters as $\lambda_{\text{pose}} = 0.1$, $\lambda_{\text{cos}} = 0.01$, and $\lambda_{\text{past}} = 0.1$. We use five sequential stereo views as inputs to our model, *i.e.*, $T = 5$, with a skip size of 3. See our supplement for more details on the network architecture.

5.3. Evaluation

We compare our method with existing stereo-based ego-centric pose estimation methods [1, 53]. We use the of-

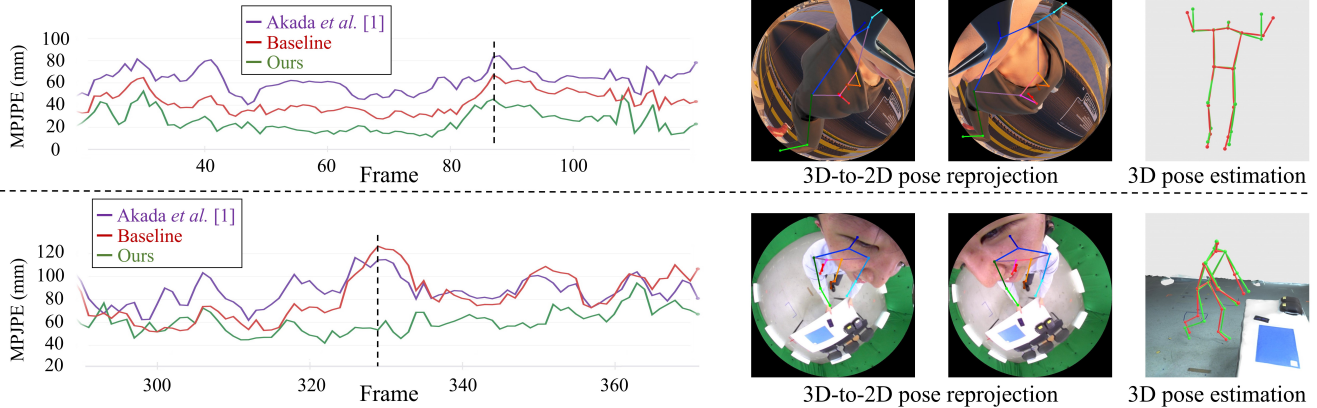


Figure 5. Results of our framework and comparison methods on example sequences from UnrealEgo2 (above) and UnrealEgo-RW (below). **Left:** MPJPE curves. **Right:** Outputs of our method at frame 87 and 329 of the sequences, respectively. 3D pose estimation and ground truth are colored in red and green, respectively.

ficial source code of Akada *et al.* [1] and re-implement the framework of Zhao *et al.* [53] as its source code is not available. Note that the comparison methods are trained on the same datasets as our model. Kang *et al.* [12] (arXiv preprint at the time of submission) only shows results of the pelvis-relative estimation on UnrealEgo. Therefore, we include them for reference. Furthermore, we are interested in the performance of the publicly available state-of-the-art method [1] with temporal inputs. Thus, we modify their 3D module such that it can take as an input a sequence of stereo 2D keypoint heatmaps with the same time step as ours, *i.e.*, $T = 5$. Here, we replace the first and the last fully connected layers in the encoder, the pose decoder, and the heatmap reconstruction decoder of their autoencoder-based 3D module [1] by those with T times the size of the original hidden dimension. We denote this model as Baseline and train it with the same training procedure as Akada *et al.* [1]. Note that Akada *et al.* [1], Baseline, and our model use the same 2D module.

We follow the existing works [1, 11, 37, 38, 40–42, 45, 52, 53] to report Mean Per Joint Position Error (MPJPE) and Mean Per Joint Position Error with Procrustes Alignment [13] (PA-MPJPE). We additionally report 3D Percentage of Correct Keypoints (3D PCK) and Area Under the Curve (AUC) for UnrealEgo2 and UnrealEgo-RW.

Results on Synthetic Datasets. Tables 1 and 2 report the results with UnrealEgo [1] and UnrealEgo2. Our method outperforms the existing methods [1, 12, 53] and Baseline across all metrics by a significant margin, *e.g.*, $>15\%$ on UnrealEgo [1] and $\geq 40\%$ on UnrealEgo2 (on MPJPE). The qualitative results on UnrealEgo2 in Fig. 4-(left part) show that existing methods and Baseline fail to estimate lower bodies of complex poses with severe self-occlusions, such as crouching. Even under such challenging scenarios, however, our approach yields accurate 3D poses. See Fig. 5-(above part) for a MPJPE curve and visual outputs of our

Method	MPJPE(\downarrow)	PA-MPJPE(\downarrow)	3D PCA(\uparrow)	AUC(\uparrow)
T = 1	108.63	84.69	77.98	44.15
T = 3	108.23	85.28	78.35	44.54
T = 5	<u>104.14</u>	82.18	<u>80.20</u>	46.22
T = 7	104.01	<u>82.43</u>	80.52	<u>46.10</u>

Table 6. Ablation study of our model with different sequence lengths on UnrealEgo-RW. The numbers are in *mm*.

framework on UnrealEgo2. Our method is able to constantly estimate accurate 3D poses compared to the existing methods. As evidenced by these results, our method demonstrates superiority and becomes a strong benchmark method in the egocentric stereo 3D pose estimation tasks. See our supplementary material and video for more results.

Results on the Real-World Dataset. Table 3 shows quantitative results on UnrealEgo-RW. Again, our method outperforms the existing methods [1, 53] and Baseline across all metrics, *e.g.*, by more than 10% on MPJPE. See Fig. 4-(right) for qualitative results. The current state-of-the-art methods [1, 53] or Baseline show floating feet, inaccurate pelvis position, and penetration to the floor ground. However, our method is able to estimate accurate 3D poses. See Fig. 5-(below part) for a MPJPE curve and visual outputs on an example motion of UnrealEgo-RW. The curve indicates that our method constantly shows lower 3D errors than the comparison methods. All of the results indicate the effectiveness of our proposed framework compared to the existing methods. We also visualize 2D heatmaps, 3D-to-2D pose reprojection, and 3D pose prediction from our method in Fig. 6. Even when the joint locations of the lower body are estimated closely in the 2D heatmaps, our approach predicts accurate lower body poses. These results suggest that the proposed method with our portable device can open up the possibility of many future applications, including animating virtual humans (Fig. 1-(g)). For the virtual human animation, we applied inverse kinematics with estimated 3D joint locations and ground-truth camera poses to drive the




Figure 6. Visualization of outputs from our model on UnrealEgo-RW. 3D-to-2D pose reprojection is visualized in the same colors as in Fig. 1-(e). 3D pose estimation and ground truth are displayed in red and green, respectively.

character in a world coordinate system.

Ablation Study. In Table 4, we first ablate (a) the CNN-based 3D module (Baseline) with depth data concatenated to the heatmap inputs. However, naively adding this extra scene information to this 3D module does not help probably because the CNN layers can be affected by invalid depth values even with the depth region masks.

Next, we test our transformer-based 3D module (b) without query augmentation (c) without depth data. They perform worse than our full framework. We also ablate our method (d) without the padding mask. The result indicates that adding depth padding masks helps because the padding mask can filter out the invalid values in depth maps in the attention module. These results validate that our video-based 3D scene reconstruction module and video-dependent query augmentation policy boost 3D joint localization accuracy. Next, we ablate our model (e) with 3D pose supervision of the latest frame only. Note that this ablation uses the same sets of input data and joint queries as the original model, *i.e.*, $T = 5$. This model estimates less accurate poses due to the loss of supervision from past 3D poses. We also test (f) a single set of joint queries, *i.e.*, q^1 , instead of T sets to predict the latest 3D pose. Similar to (e), this model cannot benefit from the supervision of past 3D poses.

We further investigate the effect of the scene information. Table 5 shows the MPJPE per body part and Mean Penetration Error (MPE) [34, 35] between feet and floor ground. The results reveal that depth features with the padding masks reduce the errors in the lower body while maintaining the performance in the upper body.

In Table 6, we ablate the effect of the sequence length of input frames for our method. It is worth noting that our model with $T=1$ yields better results than the best existing method [1] and Baseline that utilizes temporal information (see Table 3). Since our model uses the same 2D module as Akada *et al.* [1] and Baseline, the difference comes only from the 3D module. This suggests that their autoencoder-based 3D modules with the heatmap reconstruction component are, very likely, not the most suitable solution for estimating 3D poses from 2D joint heatmaps, highlighting the potential of our transformer-based framework. The result also indicates that although the longer sequence can bring performance improvement to some extent, the se-

Method	Initial training data	MPJPE(↓)	PA-MPJPE(↓)	3D PCA(↑)	AUC(↑)
Zhao <i>et al.</i> [53]	UnrealEgo [1]	99.09	72.47	79.82	43.55
Akada <i>et al.</i> [1]		94.87	69.79	82.78	46.80
Baseline		83.89	64.30	86.20	51.63
Ours		75.34	57.29	89.43	55.77

Zhao <i>et al.</i> [53]	UnrealEgo2	97.86	69.92	81.53	46.32
Akada <i>et al.</i> [1]		92.48	67.15	84.25	48.04
Baseline		82.16	61.60	87.07	52.72
Ours		72.89	56.19	90.29	57.19

Table 7. Fine-tuning results of device-relative 3D pose estimation on UnrealEgo-RW with mm-scale.

quence lengths of five and seven show comparable results. **Synthetic Data for Pre-training.** No existing works explored the efficacy of synthetic data for pre-training in egocentric 3D pose estimation. Thus, we further conduct experiments with models pre-trained on the synthetic datasets and fine-tuned on the real-world data. Tables 3 and 7 show that all methods benefit from the training with the large-scale synthetic data even with the differences in the synthetic and real-world setups, *e.g.*, fisheye distortion and syn-to-real domain gaps. Note that the gain of our method from UnrealEgo to UnrealEgo2 is significant, *i.e.*, 3.3% on MPJPE (75.34mm to 72.89mm). This suggests that it is helpful to develop not only new models but also large-scale synthetic datasets even with different distortion and domain gaps.

6. Conclusion

In this paper, we proposed a new transformer-based framework that significantly boosts the accuracy of egocentric stereo 3D human pose estimation. The proposed framework leverages the scene information and temporal context of egocentric stereo video inputs via our video-based 3D scene reconstruction module and video-based joint query augmentation policy. Our extensive experiments on the new synthetic and real-world datasets with challenging human motions validate the effectiveness of our approach compared to the existing methods. We hope that our proposed benchmark datasets and trained models will foster the further development of methods for egocentric 3D vision.

Acknowledgment. The work was supported by the ERC Consolidator Grant 4DReply (770784) and the Nakajima Foundation. We thank Silicon Studio Corp. for providing the fisheye plug-in for Unreal Engine.

References

- [1] 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 (ECCV)*, 2022. 1, 2, 3, 5, 6, 7, 8
- [2] 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 (ECCV)*, 2020. 5
- [3] Young-Woon Cha, True Price, Zhen Wei, Xinran Lu, Nicholas Rewkowski, Rohan Chabra, Zihe Qin, Hyounghun Kim, Zhaoqi Su, Yebin Liu, Adrian Ilie, Andrei State, Zhenlin Xu, Jan-Michael Frahm, and Henry Fuchs. Towards fully mobile 3d face, body, and environment capture using only head-worn cameras. *IEEE Transactions on Visualization and Computer Graphics*, 24(11):2993–3004, 2018. 2, 3
- [4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *Computer Vision and Pattern Recognition (CVPR)*, 2009. 6
- [5] Moritz Einfalt, Katja Ludwig, and Rainer Lienhart. Uplift and upsample: Efficient 3d human pose estimation with uplifting transformers. In *Winter Conference on Applications of Computer Vision (WACV)*, 2023. 2, 5
- [6] FUJINON FE185C057HA-1 fisheye lens, 2023. <https://www.fujifilm.com/de/de/business/optical-devices/mvlens/fe185.3>
- [7] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [8] Kaiping He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [9] Yihui He, Rui Yan, Katerina Fragkiadaki, and Shoou-I Yu. Epipolar transformers. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. 2020. 2
- [11] Hao Jiang and Vamsi Krishna Ithapu. Egocentric pose estimation from human vision span. In *International Conference on Computer Vision (ICCV)*, 2021. 1, 2, 3, 6, 7
- [12] Taeho Kang, Kyungjin Lee, Jinrui Zhang, and Youngki Lee. Ego3dpose: Capturing 3d cues from binocular egocentric views. In *SIGGRAPH Asia Conference*, 2023. 2, 5, 6, 7
- [13] David G. Kendall. A Survey of the Statistical Theory of Shape. *Statistical Science*, 4(2):87–99, 1989. 7
- [14] Rawal Khirodkar, Aayush Bansal, Lingni Ma, Richard Newcombe, Minh Vo, and Kris Kitani. Ego-humans: An egocentric 3d multi-human benchmark. In *International Conference on Computer Vision (ICCV)*, 2023. 3
- [15] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 6
- [16] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023. 4
- [17] Jiaman Li, Karen Liu, and Jiajun Wu. Ego-body pose estimation via ego-head pose estimation. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [18] Wenhao Li, Hong Liu, Runwei Ding, Mengyuan Liu, Pichao Wang, and Wenming Yang. Exploiting temporal contexts with strided transformer for 3d human pose estimation. *IEEE Transactions on Multimedia (TMM)*, 2022. 2
- [19] Wenhao Li, Hong Liu, Hao Tang, Pichao Wang, and Luc Van Gool. Mhformer: Multi-hypothesis transformer for 3d human pose estimation. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [20] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [21] Yuxuan Liu, Jianxin Yang, Xiao Gu, Yijun Chen, Yao Guo, and Guang-Zhong Yang. EgoFish3d: Egocentric 3d pose estimation from a fisheye camera via self-supervised learning. *IEEE Transactions on Multimedia (TMM)*, pages 1–12, 2023. 2
- [22] Zhengyi Luo, Ryo Hachiuma, Ye Yuan, and Kris Kitani. Dynamics-regulated kinematic policy for egocentric pose estimation. 2021. 1, 2, 3
- [23] Haoyu Ma, Liangjian Chen, Deying Kong, Zhe Wang, Xingwei Liu, Hao Tang, Xiangyi Yan, Yusheng Xie, Shih-Yao Lin, and Xiaohui Xie. Transfusion: Cross-view fusion with transformer for 3d human pose estimation. In *British Machine Vision Conference (BMVC)*, 2021. 2
- [24] Metashape, 2023. <https://www.agisoft.com/>. 4
- [25] Mixamo, 2022. <https://www.mixamo.com>. 3
- [26] Xiaqing Pan, Nicholas Charron, Yongqian Yang, Scott Peters, Thomas Whelan, Chen Kong, Omkar Parkhi, Richard Newcombe, and Yuheng (Carl) Ren. Aria digital twin: A new benchmark dataset for egocentric 3d machine perception. In *International Conference on Computer Vision (ICCV)*, 2023. 3
- [27] Jinman Park, Kimathi Kaai, Saad Hossain, Norikatsu Sumi, Sirisha Rambhatla, and Paul Fieguth. Domain-guided spatio-temporal self-attention for egocentric 3d pose estimation. In *Conference on Knowledge Discovery and Data Mining (KDD)*, 2023. 2
- [28] Sungchan Park, Eunyi You, Inho Lee, and Joonseok Lee. Towards robust and smooth 3d multi-person pose estimation from monocular videos in the wild. In *International Conference on Computer Vision (ICCV)*, 2023. 2
- [29] Dario Pavllo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [30] RenderPeople, 2022. <https://renderpeople.com>. 3
- [31] 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. 1, 2, 3

- [32] RIBCAGE RX0 II camera, 2023. <https://www.backbone.ca/product/ribcage-rx0-2/>. 3
- [33] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [34] Soshi Shimada, Vladislav Golyanik, Weipeng Xu, and Christian Theobalt. Physcap: Physically plausible monocular 3d motion capture in real time. *ACM Transactions on Graphics (TOG)*, 39(6), 2020. 8
- [35] Soshi Shimada, Vladislav Golyanik, Weipeng Xu, Patrick Pérez, and Christian Theobalt. Neural monocular 3d human motion capture with physical awareness. *ACM Transactions on Graphics (TOG)*, 40(4), 2021. 8
- [36] Zhenhua Tang, Zhaofan Qiu, Yanbin Hao, Richang Hong, and Ting Yao. 3d human pose estimation with spatio-temporal criss-cross attention. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [37] Denis Tome, Patrick Peluse, Lourdes Agapito, and Hernan Badino. xr-egopose: Egocentric 3d human pose from an hmd camera. In *International Conference on Computer Vision (ICCV)*, 2019. 1, 2, 3, 5, 6, 7
- [38] 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 (PAMI)*, 45(6):6794–6806, 2023. 1, 2, 5, 6, 7
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems (NeurIPS)*, 2017. 2
- [40] Jian Wang, Lingjie Liu, Weipeng Xu, Kripasindhu Sarkar, and Christian Theobalt. Estimating egocentric 3d human pose in global space. In *International Conference on Computer Vision (ICCV)*, 2021. 1, 2, 3, 6, 7
- [41] Jian Wang, Lingjie Liu, Weipeng Xu, Kripasindhu Sarkar, Diogo Luvizon, and Christian Theobalt. Estimating egocentric 3d human pose in the wild with external weak supervision. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3
- [42] Jian Wang, Diogo Luvizon, Weipeng Xu, Lingjie Liu, Kripasindhu Sarkar, and Christian Theobalt. Scene-aware egocentric 3d human pose estimation. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2, 6, 7
- [43] 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 *Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [44] Tao Wang, Jianfeng Zhang, Yujun Cai, Shuicheng Yan, and Jiashi Feng. Direct multi-view multi-person 3d human pose estimation. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2, 5
- [45] Weipeng Xu, Avishek Chatterjee, Michael Zollhoefer, Helge Rhodin, Pascal Fua, Hans-Peter Seidel, and Christian Theobalt. Mo²Cap²: Real-time mobile 3d motion capture with a cap-mounted fisheye camera. *IEEE Transactions on Visualization and Computer Graphics*, 2019. 1, 2, 3, 6, 7
- [46] Honghong Yang, Longfei Guo, Yumei Zhang, and Xiaojun Wu. U-shaped spatial-temporal transformer network for 3d human pose estimation. *Machine Vision and Applications*, 33(6):82, 2022. 2
- [47] Yingxuan You, Hong Liu, Ti Wang, Wenhao Li, Runwei Ding, and Xia Li. Co-evolution of pose and mesh for 3d human body estimation from video. In *International Conference on Computer Vision (ICCV)*, 2023. 2
- [48] Ye Yuan and Kris Kitani. Ego-pose estimation and forecasting as real-time pd control. In *International Conference on Computer Vision (ICCV)*, 2019. 1, 2, 3
- [49] Jinlu Zhang, Zhigang Tu, Jianyu Yang, Yujin Chen, and Jun-song Yuan. Mixste: Seq2seq mixed spatio-temporal encoder for 3d human pose estimation in video. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 5
- [50] Siwei Zhang, Qianli Ma, Yan Zhang, Zhiyin Qian, Taein Kwon, Marc Pollefeys, Federica Bogo, and Siyu Tang. Ego-body: Human body shape and motion of interacting people from head-mounted devices. In *European conference on computer vision (ECCV)*, 2022. 3
- [51] Siwei Zhang, Qianli Ma, Yan Zhang, Sadegh Aliakbarian, Darren Cosker, and Siyu Tang. Probabilistic human mesh recovery in 3d scenes from egocentric views. In *International Conference on Computer Vision (ICCV)*, 2023. 2, 3
- [52] Yahui Zhang, Shaodi You, and Theo Gevers. Automatic calibration of the fisheye camera for egocentric 3d human pose estimation from a single image. In *Winter Conference on Applications of Computer Vision (WACV)*, 2021. 1, 2, 6, 7
- [53] Dongxu Zhao, Zhen Wei, Jisan Mahmud, and Jan-Michael Frahm. Egoglass: Egocentric-view human pose estimation from an eyeglass frame. In *International Conference on 3D Vision (3DV)*, 2021. 1, 2, 3, 5, 6, 7, 8
- [54] Qitao Zhao, Ce Zheng, Mengyuan Liu, Pichao Wang, and Chen Chen. Poseformerv2: Exploring frequency domain for efficient and robust 3d human pose estimation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [55] Ce Zheng, Sijie Zhu, Matias Mendieta, Taojiannan Yang, Chen Chen, and Zhengming Ding. 3d human pose estimation with spatial and temporal transformers. In *International Conference on Computer Vision (ICCV)*, 2021.
- [56] Jieming Zhou, Tong Zhang, Zeeshan Hayder, Lars Petersson, and Mehrtash Harandi. Diff3dhpe: A diffusion model for 3d human pose estimation. In *International Conference on Computer Vision (ICCV) Workshops*, 2023.
- [57] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang. Motionbert: A unified perspective on learning human motion representations. In *International Conference on Computer Vision (ICCV)*, 2023.
- [58] Yiran Zhu, Xing Xu, Fumin Shen, Yanli Ji, Lianli Gao, and Heng Tao Shen. Posegtac: Graph transformer encoder-decoder with atrous convolution for 3d human pose estimation. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2021. 2