




# Ego3DPose: Capturing 3D Cues from Binocular Egocentric Views

Taeho Kang  
Seoul National University  
Seoul, South Korea  
taeho.kang@hcs.snu.ac.kr

Jinrui Zhang  
Central South University  
Changsha, China  
zhangjinruicsu@gmail.com

Kyungjin Lee  
Seoul National University  
Seoul, South Korea  
jin11542@snu.ac.kr

Youngki Lee  
Seoul National University  
Seoul, South Korea  
youngkilee@snu.ac.kr



**Figure 1: The binocular egocentric input and a qualitative result of 3D pose estimation of previous methods and our system. The ground truth pose is in red, and the gray one corresponds to [Zhao et al. 2021], [Akada et al. 2022], and our system from left to right.**

## ABSTRACT

We present Ego3DPose, a highly accurate binocular egocentric 3D pose reconstruction system. The binocular egocentric setup offers practicality and usefulness in various applications, however, it remains largely under-explored. It has been suffering from low pose estimation accuracy due to viewing distortion, severe self-occlusion, and limited field-of-view of the joints in egocentric 2D images. Here, we notice that two important 3D cues, stereo correspondences, and perspective, contained in the egocentric binocular input are neglected. Current methods heavily rely on 2D image features, implicitly learning 3D information, which introduces biases towards commonly observed motions and leads to low overall accuracy. We observe that they not only fail in challenging occlusion cases but also in estimating visible joint positions. To address these challenges, we propose two novel approaches. First, we design a two-path network architecture with a path that estimates pose per limb independently with its binocular heatmaps. Without full-body information provided, it alleviates bias toward trained full-body distribution. Second, we leverage the egocentric view of body limbs, which exhibits strong perspective variance (e.g., a

significantly large-size hand when it is close to the camera). We propose a new perspective-aware representation using trigonometry, enabling the network to estimate the 3D orientation of limbs. Finally, we develop an end-to-end pose reconstruction network that synergizes both techniques. Our comprehensive evaluations demonstrate that Ego3DPose outperforms state-of-the-art models by a pose estimation error (i.e., MPJPE) reduction of 23.1% in the UnrealEgo dataset. Our qualitative results highlight the superiority of our approach across a range of scenarios and challenges.

## CCS CONCEPTS

• Computing methodologies;

## KEYWORDS

Egocentric, 3D Human Pose Estimation, Stereo vision, Heatmap

### ACM Reference Format:

Taeho Kang, Kyungjin Lee, Jinrui Zhang, and Youngki Lee. 2023. Ego3DPose: Capturing 3D Cues from Binocular Egocentric Views. In *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers '23)*, December 12–15, 2023, Sydney, NSW, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3610548.3618147>

## 1 INTRODUCTION

Egocentric 3D pose reconstruction brings mobility to many promising applications such as avatar creation [Ahuja et al. 2019], motion analysis for sports [Hwang et al. 2017; Zecha et al. 2018] and health [Maskeliūnas et al. 2023], and action recognition [Núñez-Marcos et al. 2022]. Among various egocentric setups [Akada et al.

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

*SA Conference Papers '23, December 12–15, 2023, Sydney, NSW, Australia*

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0315-7/23/12...\$15.00

<https://doi.org/10.1145/3610548.3618147>

2022; Rhodin et al. 2016; Tome et al. 2019; Xu et al. 2019; Zhao et al. 2021], several works show the potential of eyeglasses with two binocular fisheye cameras [Zhao et al. 2021]. EgoCap [Rhodin et al. 2016] is one of the first works that adopt a helmet form factor with two mounted cameras on an extended stick. EgoGlass [Zhao et al. 2021] extends the work to handle self-occlusion cases and UnrealEgo [Akada et al. 2022] improves the accuracy by using a large-scale synthetic dataset. The accuracy, however, is still far below third-person-view-based 3D pose estimation techniques.


In this work, we introduce Ego3DPose, a novel system designed with an in-depth investigation of the binocular and egocentric setting. Compared to third-person view inputs, egocentric binocular inputs inherently possess limited information due to their restricted field of view and significant self-occlusions among the joints. Yet, existing approaches have not fully harnessed the potential of these already constrained inputs. Figure 2 shows that existing work fails on joints clearly shown in the inputs with sufficient visual information. This is especially common in motions with highly independent body parts movement, as shown in the supplementary video, where full-body information barely helps estimate specific limb poses.

We identify two under-utilized sources of 3D information within binocular and egocentric inputs: **stereo correspondence** and **perspective**. While these concepts are not new, we recognize that effectively incorporating them into the network poses unique challenges and requires careful investigation.

Utilizing stereo correspondence is challenging in two-fold. Firstly, existing methods in the domain of multiview pose estimation are highly data-dependent. Their performance lacks generalizability across poses or camera configurations [Bartol et al. 2022]. This is, in particular, challenging in egocentric settings with limited datasets. Secondly, 2D-to-3D lifting is inevitably ambiguous unless both 2D joint positions and camera calibrations are accurate [Chen and Ramanan 2017a]. Existing methods focus on optimizing these two components working sequentially. In contrast to third-person-view settings, where high accuracy can be achieved in 2D pose estimation, egocentric settings suffer from severe self-occlusions and view distortions propagating the error to stereo correspondence estimation. Thus, tighter cooperation between stereo information and the 2D pose estimation holds greater significance.

The amplified perspective effect of fisheye camera views, especially with the egocentric setup, has been considered a challenge rather than an opportunity. Specifically, the body parts further from the camera, such as the lower limb, are shown much smaller than the real size, while the shoulder and upper limb views are enlarged. Existing methods aim to tackle the challenge with reprojection methods to alleviate the distortion effect [Liu et al. 2023; Xu et al. 2019; Zhang et al. 2021] or extensively train the network with large synthetic datasets [Akada et al. 2022; Tome et al. 2019]. However, these approaches overlook the potential value of the perspective factor, which can be useful information conveying the intensity of the 3D effect.

We propose two novel modules integrated into an end-to-end neural network to tackle the challenges. First, we devise a Stereo Matcher network to serve two purposes: it i) focuses on learning the stereo correspondences independent from the full body pose distribution in the dataset and ii) generates an explicit estimation of the 3D orientation instead of directly lifting 2D joint estimations



**Figure 2: Heatmaps (Left) and 3D pose reprojected to the input images (Right) of [Akada et al. 2022] and our method. The heatmap estimation indicates that the visual information is sufficient.**

to a final 3D position. Specifically, the Stereo Matcher takes the optical features of each limb independently and predicts per-limb 3D orientation with a weight-shared network. Using the stereo correspondences and perspective effect, the architecture enables the module to focus on estimating the limb’s orientation without reliance on the full-body information. Furthermore, the estimated 3D orientation is concatenated with the 2D heatmaps’ encoded pose features to assist lifting from 2D to 3D poses.

We also introduce the Perspective Embedding Heatmap representation. This representation allows the 2D module to explicitly extract and utilize the 3D information of the perspective effect. We represent the perspective effect as the 3D orientation of the limbs. For example, as shown in Figure 2, the width of the top part of the upper limb is much larger than the width of the lower part. This visual cue suggests that the upper limb is nearly perpendicular to the camera view plane. Then, we devise a trigonometry-based representation of the 3D orientation. It allows embedding the information into a 2D heatmap to utilize its uncertainty estimates.

Our end-to-end network, which integrates the two modules, significantly outperforms prior works. The quantitative measure of 3D pose estimation errors, MPJPE, was improved by 27.0% and 23.1% compared to the baselines EgoGlass and UnrealEgo in the UnrealEgo dataset, respectively. Furthermore, the qualitative results shown in Figure 1 and Figure 7 also show noticeable improvements made by our work.

In summary, the main contributions are as follows:

- We identify two under-utilized but crucial 3D information from binocular egocentric 3D pose estimation.
- We propose a novel two-path network that separates stereo correspondence features from the pose features.
- We introduce a new heatmap representation for 3D orientation to leverage the perspective effect of egocentric views.
- We achieve significant quantitative and qualitative performance improvements compared to state-of-the-art methods.

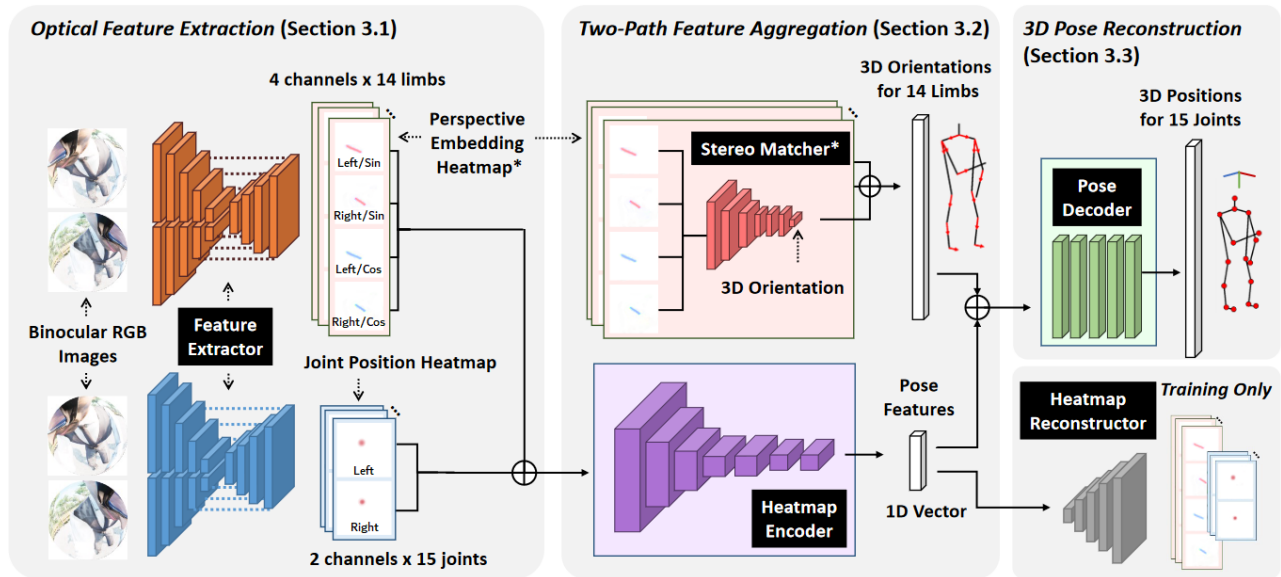


Figure 3: Our system consists of three stages: optical feature extraction, two-path feature aggregation, and 3D pose reconstruction with our two key components Perspective Embedding Heatmap and Stereo Matcher (marked with asterisk\*). The optical feature extraction uses two feature extractors to generate Joint Position Heatmaps and Perspective Embedding Heatmaps. The Stereo Matcher network estimates the 3D orientation and learns the stereo correspondence from Perspective Embedding Heatmaps. The heatmap encoder takes all heatmaps as input and generates pose features. Finally, the 3D pose decoder takes the estimated 3D orientations and pose features to predict the 3D pose of the joints. The heatmap reconstructor is only used during training.

## 2 RELATED WORKS

### 2.1 Binocular Egocentric 3D Pose Reconstruction

There has been a growing body of literature on egocentric 3D pose reconstruction, particularly in the context of binocular vision-based setups. Compared to monocular vision-based methods that use a single fisheye camera with a downward view towards the full body [Tome et al. 2019; Xu et al. 2019], binocular vision allows higher accuracy in 3D pose reconstruction via the wider view with cameras in proximity to the face [Zhao et al. 2021].

EgoCap [Rhodin et al. 2016] proposes the first binocular fisheye camera-based prototype mounted on a helmet or a headset. It uses the prototype to collect data with 8 subjects in a multiview studio setting for ground truth collection and creates a CNN-based 3D pose estimation model. EgoGlass [Zhao et al. 2021] proposes compact binocular cameras embedded in eyeglasses, which are located in closer proximity to the user’s body. However, the compact form factor leads to frequent self-occlusions of partial body parts and limited body coverage for various motions. The work partially solves the problem by explicitly forcing the model to learn the body part information with a segmentation mask. UnrealEgo [Akada et al. 2022] recently released a large-scale synthesized dataset for a similar setup of EgoGlass. It also shows that the overall 3D pose estimation accuracy improves when the weights between the two heatmap estimators, one for each of the binocular inputs, are shared. Despite the potential of such a setting, the overall pose estimation accuracy still requires significant improvement.

### 2.2 Methods for Robust 3D Pose Estimation

3D pose estimation solutions can be broadly divided into two types: i) direct estimation of 3D poses [Li and Chan 2014; Li et al. 2015; Pavlakos et al. 2017, 2018; Sun et al. 2017; Tekin et al. 2016a] and ii) leveraging robust 2D pose estimation as an intermediate result and lifting to 3D [Chen and Ramanan 2017b; Jain et al. 2013; Li and Lee 2019; Moreno-Noguer 2016; Qiu et al. 2019; Tekin et al. 2016b; Tompson et al. 2015; Zhou et al. 2022]. The 2D-to-3D lifting method is known to be more robust in various situations [Chen and Ramanan 2017b]. In particular, maintaining the probabilistic nature of the intermediate 2D representation significantly contributed to performance improvement. Earlier work designed the 2D module to make multiple pose predictions to have ambiguity [Jahangiri and Yuille 2017; Moreno-Noguer 2017]. Currently, a common approach is representing the estimated 2D joint positions as heatmaps [Tekin et al. 2016b; Tompson et al. 2015; Zhou et al. 2022]. This approach has been explored in various ways. In particular, using skeletal heatmaps instead of joint heatmaps was proposed to handle occlusion cases in multi-person 3D pose estimation [Chen and Yang 2018]. The key idea was to represent the skeleton as lines connecting two joints. We also leverage the line-based representation but propose a trigonometry-based embedding showing better performance. The heatmap-based approach was explored in multiview settings as well [Iskakov et al. 2019; Qiu et al. 2019; Zhang et al. 2020]. We also leverage the heatmap representations but propose methods to integrate 3D supervision from earlier stages of the 2D module.



Most of the recent works in egocentric 3D pose reconstruction also adopt the heatmap-based two-staged approach. xR-EgoPose [Tome et al. 2019], which relies on a monocular fisheye view, uses a 2D module to generate Joint Position Heatmaps. Then, the 3D pose estimator module uses the heatmap as input to estimate the 3D pose. EgoGlass [Zhao et al. 2021] uses the same architecture but extends the 2D module to take binocular inputs and estimates an additional segmentation mask for body part information. UnrealEgo [Akada et al. 2022] improves this work by using a weight-shared encoder of the U-Net for the two binocular inputs and a shared decoder for the heatmap estimation. Compared to the large body of work in third-person view 3D pose estimation, more studies are required to increase the robustness in egocentric settings.

### 3 METHOD


To better highlight the effectiveness of our two modules, incorporating stereo correspondence and perspective effect, we integrate them into widely used network architecture instead of redesigning the end-to-end network. In this paper, the overall structure follows the encoder-decoder network, where the 2D encoder module extracts optical features from the binocular RGB inputs as 2D Joint Position Heatmaps. The heatmaps are encoded into pose features and then reconstructed into 3D joint positions with the final 3D decoder. Here, we add the Perspective Embedding Heatmap generator in the optical feature extraction stage and the Stereo Matcher to the encoder stage.

Figure 3 shows the overall network architecture. The input to the network is two  $256 \times 256$  RGB images. The first stage, optical feature extraction, generates two sets of heatmaps: i) Perspective Embedding Heatmaps based on our novel perspective-aware trigonometry representation and ii) Joint Position Heatmaps similar to previous work. Then, the two-path feature aggregation network separately learns i) the stereo correspondences from the binocular Perspective Embedding Heatmaps and ii) pose feature vector extracted from both Joint Position Heatmaps and Perspective Embedding Heatmaps. Our network then concatenates the two disentangled information and generates the final 3D local pose. The local pose is defined as a 3D position of the joints relative to the pelvis in the camera coordinate system. A total of 15 joints are estimated, including the neck, shoulders, elbows, hands, thighs, calves, ankles, and feet. We additionally train a heatmap reconstruction network which is not used for inferences.

#### 3.1 Optical Feature Extraction

This module consists of two networks. Both networks take the binocular RGB images as input but generate different types of optical features: Perspective Embedding Heatmap and Joint Position Heatmap. They use the same network architecture, a U-Net-based model with a ResNet18 backbone. We first explain the Joint Position Heatmap as a preliminary for our Perspective Embedding Heatmap.

**3.1.1 Joint Position Heatmap.** This heatmap is a single-channel image where each pixel represents the confidence level associated with the 2D positions of the joints. It has been widely adopted in various pose estimation methods [Bulat and Tzimiropoulos 2016] due to its effectiveness in capturing uncertainty. To align with previous work, we adopt a per-joint heatmap generation approach



**Figure 4: Visualization of a joint position, limb position, limb depth, and our Perspective Embedding Heatmap, where the color represents the cosine and sine of the limb-view angle shown in the right image.**

instead of using a single heatmap for all joints. A total of 30 Joint Position Heatmaps are generated, corresponding to the left and right input images for the 15 individual joints.


**3.1.2 Perspective Embedding Heatmap.** The goal of the Perspective Embedding Heatmap is to incorporate the probabilistic information of the perspective effect instead of the joint positions. The perspective effect can be expressed in different ways, such as the depth of each joint or the angle between two adjacent joints (Refer to Figure 4 for the visualization of different types and the specific explanations in Section 5.3.3). We, in particular, choose the 3D orientation of the limbs defined with the angle between the viewing plane and the limb (referred to as limb-view angle hereinafter). The limb-view angle  $\theta_l$  for each limb is computed as follows:

$$\theta_l = \text{atan} \left( \frac{z_l}{\sqrt{x_l^2 + y_l^2}} \right)$$


, where  $x_l, y_l, z_l$  denotes the relative 3D coordinate of the child joint from the parent joint in the camera coordinate space, and  $xy$ -plane is equivalent to the viewing plane of the camera. Limbs are represented as a line rather than a pair of two adjacent joints to enhance the robustness against occlusions because the visible parts of the limbs can be utilized.

Embedding the 3D orientation to a heatmap is not straightforward. We have to embed three pieces of information, 2D position, confidence, and 3D orientation, into a 2D heatmap. One approach would be simply multiplying the magnitude of the orientation with the probabilistic output of the heatmap. However, this can confuse the network as the pixel values contain information on confidence and orientation. Adding the perspective value as a separate channel is also not an effective representation because the confidence only refers to the 2D joint position estimation rather than the orientation value.

We use the trigonometric function to embed the limb-view angle into a 2D heatmap. This choice is motivated by the properties of trigonometric functions. The sine and cosine value represents an angle with the constant norm as a vector. Our heatmap has two channels where the pixel values are scaled sine and cosine of the



**Figure 5: Our Stereo Matcher module without Perspective Embedding Heatmap.**



**Figure 6: Pose feature extraction module from prior work.**

limb-view angle. The scale of the resulting vector represents the confidence of the orientation estimation, while the ratio of the sine and cosine values can represent the estimated limb-view angle. In this way, using the trigonometric functions ensures that the perspective information is entangled with the probability. Perspective Embedding Heatmaps are defined for 14 limbs, excluding the head-neck connection from the 15 limbs connecting the tree hierarchy of 16 joints. Each set of Perspective Embedding Heatmaps contains four channels of heatmaps for sine and cosine for the left and right camera views.

## 3.2 Two-Path Feature Aggregation

The key to our high 3D pose estimation accuracy is devising the network to learn the stereo feature matching separately from the distribution of full body pose in the train dataset. Thus, we replace the one-path encoder stage of the previous work, which generates pose features from the 2D joint heatmaps, with a two-path feature aggregation stage. We first explain the heatmap encoder network, which extends prior work estimating the pose features. Then, we provide the details of the new Stereo Matcher network.

**3.2.1 Heatmap Encoder.** The input of the heatmap encoder is all the heatmaps generated from the previous optical feature extraction stage. The heatmaps include the joint positions of all 15 joints and Perspective Embedding Heatmaps for 14 limbs concatenated channel-wise. The heatmap encoder is a CNN-based network, and the output is a 1D vector of 20 pose feature values.

**3.2.2 Stereo Matcher Network.** Our Stereo Matcher network is a simple yet effective technique that takes the Perspective Embedding Heatmaps of each limb as input and estimates the 3D orientation of the limb individually. The 3D orientation of each limb is defined as,

$$\mathbf{o}_l = \frac{[x_l, y_l, z_l]}{\sqrt{x_l^2 + y_l^2 + z_l^2}}$$

, a normalized vector of relative positions of a child from the parent joint in a camera’s coordinate system.  $x_l, y_l, z_l$  denotes the relative 3D coordinate.

The Stereo Matcher takes the heatmaps of one limb at a time and estimates the specific limb’s 3D orientation. Figure 5 visualizes the architecture, with the conventional heatmaps used for comparison. Each limb is represented as a pair of position heatmaps of adjacent joints (i.e., a total of 4 channel heatmaps, including binocular Joint Position Heatmaps of two joints). The per-limb process is clearly distinguished from the heatmap encoder shown in Figure 6,

which takes heatmaps for all body parts to encode them as one pose feature vector. The estimation without full body information is possible due to stereo information and 3D cues from Perspective Embedding Heatmaps, and it prevents reliance on the trained full-body pose distribution. The process does not require specific knowledge of the limb. Thus, the weights of the Stereo Matcher network are shared across different limbs for sample efficiency. Our final network architecture integrates the Stereo Matcher with the Perspective Embedding Heatmaps shown in Figure 3. The four channels input is the Perspective Embedding Heatmaps of one limb, including two channels each for sine and cosine. Our empirical analysis shows that the integration has a synergistic effect (Table 3).

## 3.3 3D Pose Reconstruction

In the final 3D pose reconstruction stage, the 3D pose decoder takes the concatenated vector of the pose features and the 3D orientations of 14 limbs as input. Then, it estimates the final 3D local pose of the 15 joints. The 3D pose decoder consists of only fully connected layers. Lastly, our system adopts a heatmap reconstructor of [Tome et al. 2019], only used during training, to constrain the heatmap encoder to preserve the uncertainty represented in the heatmap.

## 4 TRAINING DETAILS

Our Ego3DPose is trained and evaluated on a cloud server with Nvidia RTX A6000 48GB GPU, and two AMD EPYC 7313 CPUs. The entire process of training takes approximately 20 hours on the UnrealEgo dataset. The network is trained in 2 stages, first for only an optical feature extractor, and second for only a 3D pose reconstructor. For the training of the 3D pose reconstructor, pre-trained and frozen optical feature extractor weights are used to estimate heatmaps.

The ResNet-18 in the U-Nets of the optical feature extractors are initialized with pre-trained weights. Other networks are initialized with standard normal distribution. The 3D orientation estimated from Stereo Matcher (Section 3.2.2) is detached from the gradient graph before being used as input to a 3D pose estimator to ensure that the module is learned independently.

The network is trained for 10 epochs, with linear learning rate decay from epochs 5 to 10 on the UnrealEgo dataset. For the EgoCap dataset, due to its volume, the network is trained for 100 epochs. Batch size 16 is used. Adam [Kingma and Ba 2015] optimizer with learning rate  $10^{-3}$  is used for the optical feature extractor. The learning rate-free D-Adaptation optimizer [Defazio and Mishchenko 2023], D-Adapt SGD is used for the 3D pose reconstruction to facilitate the tuning process with a growth rate of 1.02. Empirically, the SGD variant is more effective with the intermediate output detached during the training. More details of the network architecture are described in the supplementary material.

## 4.1 Loss

We introduce the details of the loss term used for training our system. We train modules in 2 stages, and multiple losses are defined for each stage.

**4.1.1 Optical Feature Extractor.** Denoting the estimated Joint Position Heatmaps as  $\hat{JH}$  and the ground truth heatmaps as  $JH$ , the loss for the Joint Position Heatmap estimation is computed as follows:

$$L_{JH} = \text{mse}(\hat{JH}, JH)$$

To encourage learning the heatmap evenly for limbs with different pixel lengths, the Perspective Embedding Heatmap estimation loss for each limb is divided by its pixel length, accounting for the estimation error dependency on the length of the limb in the pixel space. We observed that this helps to get stable test accuracy. For the Perspective Embedding Heatmaps, denoting the estimated heatmaps as  $\hat{PH}$  and the ground truth heatmaps as  $PH$ , the loss is computed as follows:

$$L_{PH} = \frac{1}{|S|} \sum_{l \in S} \frac{1}{|\ell_l|} \text{mse}(\hat{PH}_l, PH_l)$$

$PH_l$  denotes a set of heatmaps for limbs,  $\ell_l$  is the pixel length of the line in the ground truth heatmap, connecting two projected coordinates of the ends of the limb, and  $S$  is the set of limbs used for the Perspective Embedding Heatmaps.

Loss weights for the optical feature extractor are set to  $\lambda_{JH} = 1$  and  $\lambda_{PH} = 1$ . The total loss ( $L_{2D}$ ) for the optical feature extractor is the following:

$$L_{2D} = \lambda_{JH} * L_{JH} + \lambda_{PH} * L_{PH}$$

**4.1.2 3D Pose Reconstructor.** The loss for the Stereo Matcher is defined as follows,

$$L_{Trans} = \text{mse}(\hat{\mathbf{o}}, \mathbf{o})$$

, where  $\mathbf{o}$  is the ground truth of all limbs' 3D orientation and  $\hat{\mathbf{o}}$  is an estimated orientation of all limbs.

The pose estimator is trained with 3D pose reconstruction loss, heatmaps reconstruction loss, and cosine-similarity loss [Akada et al. 2022] as follows:

$$\begin{aligned} L_{Pose} &= \sum_{j \in J} \|\hat{p}_j - p_j\| \\ L_{Recon} &= (\text{mse}(\hat{JH}, JH) + \text{mse}(\hat{PH}, PH)) \\ L_{Cos} &= \sum_{l \in S} \frac{f(l, \hat{p}_j) \cdot f(l, p_j)}{\|f(l, \hat{p}_j)\| \|f(l, p_j)\|} \end{aligned}$$

, where  $J$  is the set of joints to estimate 3D position.  $f(l, p_j)$  is a limb  $l$ 's translation, the child joint's 3D coordinate relative to the parent joint, computed from given joint positions  $p_j$ .  $S$  is the set of limbs used by our systems. We set loss weights of  $\lambda_{Trans} = 1$ ,  $\lambda_{Pose} = 10^{-1}$ ,  $\lambda_{Recon} = 10^{-3}$ , and  $\lambda_{Cos} = -10^{-2}$  for the 3D pose reconstructor. The total loss ( $L_{3D}$ ) for the 3D pose reconstructor is the following:

$$L_{3D} = \lambda_{Trans} * L_{Trans} + \lambda_{Pose} * L_{Pose} + \lambda_{Recon} * L_{Recon} + \lambda_{Cos} * L_{Cos}$$

## 5 EVALUATION

### 5.1 Experimental Setup

**5.1.1 Datasets.** We evaluate using the publicly available synthetic UnrealEgo dataset, generated through the Unreal Engine<sup>1</sup>. Moreover, we further demonstrate the practicality of our approach on

the real-world EgoCap [Rhodin et al. 2016] dataset. For our ablation studies, we primarily use the UnrealEgo dataset. It provides much more comprehensive coverage of diverse human movements across various settings, which surpasses that of the EgoCap dataset. Furthermore, it encompasses more challenging scenarios, particularly those involving severe self-occlusions. Lastly, the dataset covers various ethnicities and body types with 17 different characters.

**5.1.2 Baselines.** We evaluate our system against two state-of-the-art binocular egocentric pose estimation models, e.g., UnrealEgo [Akada et al. 2022] and EgoGlass [Zhao et al. 2021], using the evaluation dataset. The results for baselines are based on the provided open-source code [Akada et al. 2022]. The evaluation of UnrealEgo on its dataset is based on the open-source weight. To perform a comprehensive evaluation of EgoGlass, we re-implement its full architecture as the body part branch, which is the primary factor for its achieved performance, is not included in the source code. As the body part branch requires segmentation data not provided in the original UnrealEgo dataset, we utilize the pseudo-segmentation technique proposed by EgoGlass.

**5.1.3 Metrics.** We utilize the widely recognized Mean Per Joint Position Error (**MPJPE**) as a metric, which calculates the mean Euclidean distance between the predicted joint positions and their respective ground truth counterparts. Additionally, we also employ Procrustes Analysis MPJPE (**PA-MPJPE**) as a complementary metric to further quantify the precision and robustness of our approach.

### 5.2 Overall Performance

**5.2.1 Evaluations on UnrealEgo Dataset.** Table 1 shows the 3D pose estimation errors of our proposed system, Ego3DPose, and two baseline systems, e.g., EgoGlass and UnrealEgo, on the UnrealEgo dataset. Tests are conducted for all motion categories within the dataset, and our proposed system achieves the best results across the board. Specifically, our system demonstrates an MPJPE of only 60.82, which represents a significant improvement of 23.1% and 27.0% over the UnrealEgo and EgoGlass systems, respectively, in terms of overall error. Our system outperforms previous methods in every motion category by a large margin, including challenging motions such as "Crawling" and "Sitting on the Ground".

**5.2.2 Evaluations on EgoCap Dataset.** As the motions are not categorized in this dataset, we present the overall results. It should be noted that the EgoCap dataset solely provides 3D pose annotations for its test dataset publicly. We sorted the 3D test dataset by subject and frame, trained on 80% of it, and reserved the remaining 20% for testing purposes. Table 2 shows that our method performs well in different settings and in real-world images. Our approach manifests augmentations, as evidenced by 11.2% improvements in MPJPE and 12.6% enhancements in PA-MPJPE, in contrast to EgoGlass. Moreover, it outperforms UnrealEgo with 8.0% advances in MPJPE and 16.5% progressions in PA-MPJPE. It is worth noting that the EgoCap dataset exhibits fewer instances of severe self-occlusions and a narrower range of pose variations in comparison to the UnrealEgo dataset. The baseline methods achieve modest accuracy resulting in smaller performance improvement compared to the evaluation on the UnrealEgo dataset.

<sup>1</sup><https://www.unrealengine.com>

**Table 1: Comparison with the state-of-the-art binocular egocentric pose reconstruction methods on the UnrealEgo dataset.**

MPJPE(PA-MPJPE)								
Method	Overall	Jumping	Falling Down	Exercising	Pulling	Singing	Rolling	Crawling
EgoGlass	83.33(61.56)	81.33(63.42)	131.64(94.70)	100.17(74.11)	81.96(62.07)	70.59(52.57)	103.29(92.66)	182.36(113.65)
UnrealEgo	79.08(59.26)	76.63(61.13)	126.82(95.99)	90.27(68.36)	78.20(61.39)	67.34(49.85)	86.26(71.61)	181.50(116.42)
Ours	<b>60.82(48.47)</b>	<b>60.84(49.41)</b>	<b>95.78(79.17)</b>	<b>76.15(63.31)</b>	<b>58.08(45.69)</b>	<b>51.75(41.05)</b>	<b>81.94(71.16)</b>	<b>148.72(105.14)</b>
MPJPE(PA-MPJPE)								
Method	Laying	Sitting on the Ground	Crouching	Crouching and turning	Crouching to Standing	Crouching-Forward	Crouching-Backward	Crouching-Sideways
EgoGlass	109.64(81.75)	207.31(144.72)	132.74(110.83)	128.50(100.83)	82.30(60.48)	79.23(65.41)	94.77(77.98)	93.89(77.91)
UnrealEgo	97.84(76.62)	194.34(150.33)	116.77(97.88)	128.23(104.32)	76.80(55.18)	73.60(61.89)	78.19(62.01)	84.91(71.41)
Ours	<b>83.36(69.85)</b>	<b>153.81(133.37)</b>	<b>93.59(79.64)</b>	<b>109.05(89.97)</b>	<b>65.34(48.70)</b>	<b>53.99(45.94)</b>	<b>58.13(46.90)</b>	<b>67.67(56.69)</b>
MPJPE(PA-MPJPE)								
Method	Standing-Whole Body	Standing-Upper Body	Standing-Turning	Standing to Crouching	Standing-Forward	Standing-Backward	Standing-Sideways	
EgoGlass	70.17(51.55)	70.18(46.32)	77.72(59.29)	77.82(70.02)	77.34(62.27)	75.13(58.47)	85.19(67.76)	
UnrealEgo	68.47(50.10)	66.54(45.44)	74.33(57.92)	74.09(61.25)	70.79(56.61)	68.13(55.17)	78.92(62.27)	
Ours	<b>50.37(39.90)</b>	<b>50.51(37.19)</b>	<b>58.52(48.08)</b>	<b>59.64(52.78)</b>	<b>55.71(48.72)</b>	<b>51.96(42.58)</b>	<b>60.93(51.22)</b>	
MPJPE(PA-MPJPE)								
Method	Dancing	Boxing	Wrestling	Soccer	Baseball	Basketball	Football	Golf
EgoGlass	83.72(65.50)	71.38(54.25)	86.09(65.28)	82.29(58.24)	77.69(61.04)	59.68(47.73)	102.15(80.56)	69.37(48.22)
UnrealEgo	79.99(63.20)	68.69(51.39)	83.73(64.47)	78.32(56.24)	72.81(56.73)	60.48(45.41)	99.56(81.81)	73.86(49.01)
Ours	<b>61.30(50.82)</b>	<b>49.91(40.25)</b>	<b>65.48(51.76)</b>	<b>54.46(43.70)</b>	<b>69.25(55.58)</b>	<b>48.39(35.83)</b>	<b>83.28(72.96)</b>	<b>54.52(38.10)</b>

**Table 2: Overall error comparison of Ego3DPose, EgoGlass and UnrealEgo on real-world EgoCap dataset.**

Methods \ Error	EgoGlass	UnrealEgo	Ours
MPJPE	61.78	59.16	<b>54.41</b>
PA-MPJPE	46.06	48.22	<b>40.24</b>

**Table 3: Ablation results of Perspective Embedding Heatmap (PH), Stereo Matcher (SM) on UnrealEgo Dataset. The input for B+SM evaluation is two adjacent joint heatmaps.**

Methods \ Error	Baseline(B)	B+PH	B+SM	B+PH+SM
MPJPE	79.08	75.82	66.72	<b>60.82</b>
PA-MPJPE	59.26	58.52	52.29	<b>48.47</b>

**Table 4: Ablation of two-path feature aggregation on UnrealEgo dataset. The Heatmap Encoder (HE) module is used exclusively, and then with our Stereo Matcher (SM), Rotation Decoder [Tome et al. 2020] (RD), and Orientation Decoder (OD).**

Paths \ Error	HE	HE+RD	HE+OD	HE+SM	SM
MPJPE	75.82	73.48	70.77	<b>60.82</b>	73.40
PA-MPJPE	58.52	57.32	55.83	<b>48.47</b>	51.62

**5.2.3 Inference Time.** The total time was 18ms making our network suitable for real-time applications. Our latency breakdown shows that the per-limb Stereo Matcher network incurs minimal

**Table 5: Comparison of results on UnrealEgo dataset with different representations used for Perspective Embedding Heatmaps on the UnrealEgo dataset. Tri. is the result with the trigonometric function.**

Heatmaps \ Error	Joints	Line(L)	L&Depth	L&Angle	L&Tri.
MPJPE	66.72	65.66	82.36	65.68	<b>60.82</b>
PA-MPJPE	52.29	51.82	63.41	52.27	<b>48.47</b>

computation overhead (e.g., Optical Feature Extraction: 13.88ms, Heatmap Encoder: 1.07ms, Stereo Matcher: 1.51ms, Pose Reconstruction: 0.32ms). [Akada et al. 2022] has shown that naively using a larger computational capacity, replacing the ResNet-18 backbone with up to the ResNet-101 backbone, does not improve accuracy. Our method effectively takes advantage of additional computation.

**5.2.4 Qualitative Comparison.** We demonstrate our system’s accuracy qualitatively. Fig. 7 shows the comparison to the previous methods in both the UnrealEgo dataset and the EgoCap dataset. The results on the UnrealEgo dataset show that our method is significantly more accurate on challenging poses. The EgoCap dataset evaluation shows that our method performs much better compared to previous methods in the real-world setting. Refer to the video attached as supplementary material for more results.

### 5.3 Ablation Studies

**5.3.1 Breakdown of Accuracy Gains.** Table 3 shows the effectiveness of our two newly proposed techniques, Perspective Embedding Heatmap (PH) and Stereo Matcher (SM). The reduction of MPJPE is 4.1% when only the Perspective Embedding Heatmap is used without the Stereo Matcher. Alternatively, with only the Stereo

Matcher, using the joint position heatmaps as shown in Figure 5, a reduction of 15.6% in MPJPE is achieved. Lastly, the combination of techniques results in a synergistic effect, leading to a substantial reduction of 23.1% in MPJPE.

**5.3.2 Effectiveness of the Two-Path Feature Aggregation.** We now compare the effectiveness of our Stereo Matcher with design alternatives in Table 4. The Perspective Embedding Heatmap is used as an input to the Heatmap Encoder for all experiments. The two key features are i) per-limb 3D orientation estimation and ii) utilization of 3D orientation as an intermediate output to assist the Heatmap Encoder for the final 3d pose reconstruction. Regarding the first feature, we evaluate it against a single network that predicts all the 3D orientations simultaneously from the concatenated heatmaps (referred to as Orientation Decoder in Table 4). This alternative only reduced the MPJPE by 6.6%, while our network achieved an impressive 19.8% improvement. We also show the effectiveness of the second feature by incorporating the rotation decoder proposed by [Tome et al. 2020]. The decoder estimates 3D joint rotation from encoded pose features and guides the heatmap encoder to learn rotational features. This method is not as effective as using the 3D estimation as an input to the pose decoder. The experimental result indicates that the inclusion of a rotation decoder yields only a minimal increase in accuracy for 3D pose recognition by only 3%. Lastly, we evaluate using only the Stereo Matcher path in the last column. Stereo Matcher alone outperforms the one with only the Heatmap Encoder, which shows the importance of the per-limb independent feature extraction.

**5.3.3 Effectiveness of Perspective Embedding Heatmap.** We experimented with alternative representations to embed perspective information in the heatmaps. The result is shown in Table 5. We first compare joint and line-based representation. The line-based representation is a raw limb heatmap without perspective information. The heatmap value represents the position of the limb rather than the orientation, as in ours. The results show that limb representation has a slightly better performance compared to the two joint-based one. We also test the alternatives of our trigonometric-based representation of the perspective effect. Specifically, we compare two other alternatives: depth and raw angle. The pixel value of the depth heatmap is the depth of each point on the limb. We normalized and added an offset to the depth value to keep it larger than zero. The raw angle representation refers to the limb-view angle without applying the trigonometric functions. The results show that our trigonometric-based representation achieves the highest accuracy.

## 6 CONCLUSION

In this work, we push the performance boundaries of binocular egocentric 3D pose estimation. We identify an under-utilized source of information in the highly constrained setting, stereo correspondence, and perspective effect. Our system incorporates a novel Stereo Matcher network, which learns to estimate 3D pose independently using stereo correspondence of heatmaps for each limb. It alleviates the issue of overfitting to the distribution of full body pose in the dataset. Furthermore, we introduce a 3D Perspective Embedding Heatmap representation. It effectively captures the perspective effect while preserving the inherent uncertainty associated

with heatmaps. Our system demonstrates substantial improvements compared to state-of-the-art methods. It is noteworthy that our proposed techniques are not exclusively tailored to the egocentric and binocular setting. They hold potential for direct application in various contexts of human 3D pose reconstruction, providing benefits beyond the scope of this work.

## ACKNOWLEDGMENTS

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MIST) (No. 2022RIA2C3008495, No. RS-2023-00218601).

## REFERENCES

- Karan Ahuja, Chris Harrison, Mayank Goel, and Robert Xiao. 2019. MeCap: Whole-Body Digitization for Low-Cost VR/AR Headsets. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 453–462. <https://doi.org/10.1145/3332165.3347889>
- Hiroyasu Akada, Jian Wang, Soshi Shimada, Masaki Takahashi, Christian Theobalt, and Vladislav Golyanik. 2022. UnrealEgo: A New Dataset for Robust Egocentric 3D Human Motion Capture. In *European Conference on Computer Vision (ECCV)*.
- Kristijan Bartol, David Bojanić, Tomislav Petković, and Tomislav Pribanić. 2022. Generalizable Human Pose Triangulation. In *Proceedings of IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Adrian Bulat and Georgios Tzimiropoulos. 2016. Human Pose Estimation via Convolutional Part Heatmap Regression. In *Computer Vision – ECCV 2016*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer International Publishing, Cham, 717–732.
- Ching-Hang Chen and Deva Ramanan. 2017a. 3D Human Pose Estimation = 2D Pose Estimation + Matching. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21–26, 2017*. IEEE Computer Society, 5759–5767. <https://doi.org/10.1109/CVPR.2017.610>
- C. Chen and D. Ramanan. 2017b. 3D Human Pose Estimation = 2D Pose Estimation + Matching. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 5759–5767. <https://doi.org/10.1109/CVPR.2017.610>
- Xiao Chen and Genke Yang. 2018. Multi-Person Pose Estimation with LIMB Detection Heatmaps. In *2018 25th IEEE International Conference on Image Processing (ICIP)*. 4078–4082. <https://doi.org/10.1109/ICIP.2018.8451055>
- Aaron Defazio and Konstantin Mishchenko. 2023. Learning-Rate-Free Learning by D-Adaptation. arXiv:2301.07733 [cs.LG]
- Jihye Hwang, Sungheon Park, and Nojun Kwak. 2017. Athlete Pose Estimation by a Global-Local Network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Karim Iskakov, Egor Burkov, Victor Lempitsky, and Yury Malkov. 2019. Learnable Triangulation of Human Pose. In *International Conference on Computer Vision (ICCV)*.
- Ehsan Jahangiri and Alan L. Yuille. 2017. Generating Multiple Diverse Hypotheses for Human 3D Pose Consistent with 2D Joint Detections. In *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*. 805–814. <https://doi.org/10.1109/ICCVW.2017.100>
- Arjun Jain, Jonathan Tompson, Mykhaylo Andriluka, Graham W. Taylor, and Christoph Bregler. 2013. Learning Human Pose Estimation Features with Convolutional Networks. *CoRR* abs/1312.7302 (2013).
- Diederik Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*. San Diego, CA, USA.
- Chen Li and Gim Hee Lee. 2019. Generating multiple hypotheses for 3d human pose estimation with mixture density network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 9887–9895.
- Sijin Li and Antoni B. Chan. 2014. 3D Human Pose Estimation from Monocular Images with Deep Convolutional Neural Network. In *Asian Conference on Computer Vision*.
- S. Li, W. Zhang, and A. B. Chan. 2015. Maximum-Margin Structured Learning with Deep Networks for 3D Human Pose Estimation. In *2015 IEEE International Conference on Computer Vision (ICCV)*. IEEE Computer Society, Los Alamitos, CA, USA, 2848–2856. <https://doi.org/10.1109/ICCV.2015.326>
- Yuxuan Liu, Jianxin Yang, Xiao Gu, Yijun Chen, Yao Guo, and Guang-Zhong Yang. 2023. EgoFish3D: Egocentric 3D Pose Estimation from a Fisheye Camera via Self-Supervised Learning. *IEEE Transactions on Multimedia* (2023), 1–12. <https://doi.org/10.1109/TMM.2023.3242551>
- Rytis Maskeliūnas, Audrius Kulikajėvas, Robertas Damaševičius, Julius Griškevičius, and Aušra Adomavičienė. 2023. Biomac3D: 2D-to-3D Human Pose Analysis Model for Tele-Rehabilitation Based on Pareto Optimized Deep-Learning Architecture. *Applied Sciences* 13, 2 (2023). <https://doi.org/10.3390/app13021116>



- Francesc Moreno-Noguer. 2016. 3D Human Pose Estimation from a Single Image via Distance Matrix Regression. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016), 1561–1570.
- F. Moreno-Noguer. 2017. 3D Human Pose Estimation from a Single Image via Distance Matrix Regression. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 1561–1570. <https://doi.org/10.1109/CVPR.2017.170>
- Adrián Núñez-Marcos, Gorka Azkune, and Ignacio Arganda-Carreras. 2022. Egocentric Vision-based Action Recognition: A survey. *Neurocomputing* 472 (2022), 175–197. <https://doi.org/10.1016/j.neucom.2021.11.081>
- Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. 2017. Coarse-to-Fine Volumetric Prediction for Single-Image 3D Human Pose. In *Computer Vision and Pattern Recognition (CVPR)*.
- G. Pavlakos, L. Zhu, X. Zhou, and K. Daniilidis. 2018. Learning to Estimate 3D Human Pose and Shape from a Single Color Image. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 459–468. <https://doi.org/10.1109/CVPR.2018.00055>
- Haibo Qiu, Chunyu Wang, Jingdong Wang, Naiyan Wang, and Wenjun Zeng. 2019. Cross View Fusion for 3D Human Pose Estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Helge Rhodin, Christian Richardt, Dan Casas, Eldar Insafutdinov, Mohammad Shafiei, Hans-Peter Seidel, Bernt Schiele, and Christian Theobalt. 2016. EgoCap: Egocentric Marker-Less Motion Capture with Two Fisheye Cameras. *ACM Trans. Graph.* 35, 6, Article 162 (dec 2016), 11 pages. <https://doi.org/10.1145/2980179.2980235>
- Xiao Sun, Jiaxiang Shang, Shuang Liang, and Yichen Wei. 2017. Compositional Human Pose Regression. *2017 IEEE International Conference on Computer Vision (ICCV)* (2017), 2621–2630.
- Bugra Tekin, Isinsu Katircioglu, Mathieu Salzmann, Vincent Lepetit, and Pascal Fua. 2016a. Structured Prediction of 3D Human Pose with Deep Neural Networks. In *Proceedings of the British Machine Vision Conference (BMVC)*, Edwin R. Hancock Richard C. Wilson and William A. P. Smith (Eds.), BMVA Press, Article 130, 11 pages. <https://doi.org/10.5244/C.30.130>
- Bugra Tekin, Pablo Márquez-Neila, Mathieu Salzmann, and P. Fua. 2016b. Learning to Fuse 2D and 3D Image Cues for Monocular Body Pose Estimation. *2017 IEEE International Conference on Computer Vision (ICCV)* (2016), 3961–3970.
- D. Tome, T. Alldieck, P. Peluse, G. Pons-Moll, L. Agapito, H. Badino, and F. De la Torre. 2020. SelfPose: 3D Egocentric Pose Estimation from a Headset Mounted Camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020), 1–1. <https://doi.org/10.1109/TPAMI.2020.3029700>
- Denis Tome, Patrick Peluse, Lourdes Agapito, and Hernan Badino. 2019. xR-EgoPose: Egocentric 3D Human Pose from an HMD Camera. In *Proceedings of the IEEE International Conference on Computer Vision*. 7728–7738.
- Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christoph Bregler. 2015. Efficient object localization using Convolutional Networks. In *CVPR*. IEEE Computer Society, 648–656. <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2015.html#TompsonGJLB15>
- Weipeng Xu, Avishek Chatterjee, Michael Zollhoefer, Helge Rhodin, Pascal Fua, Hans-Peter Seidel, and Christian Theobalt. 2019. Mo<sup>2</sup>Cap<sup>2</sup>: Real-time Mobile 3D Motion Capture with a Cap-mounted Fisheye Camera. *IEEE Transactions on Visualization and Computer Graphics* (2019), 1–1. <https://doi.org/10.1109/TVCG.2019.2898650>
- Dan Zecha, Moritz Einfalt, Christian Eggert, and Rainer Lienhart. 2018. Kinematic Pose Rectification for Performance Analysis and Retrieval in Sports. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Yahui Zhang, Shaodi You, and Theo Gevers. 2021. Automatic Calibration of the Fisheye Camera for Egocentric 3D Human Pose Estimation from a Single Image. In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*. 1771–1780. <https://doi.org/10.1109/WACV48630.2021.00181>
- Zhe Zhang, Chunyu Wang, Weichao Qiu, Wenhui Qin, and Wenjun Zeng. 2020. AdaFuse: Adaptive Multiview Fusion for Accurate Human Pose Estimation in the Wild. *IJCV* (2020), 1–16.
- Dongxu Zhao, Zhen Wei, Jisan Mahmud, and Jan-Michael Frahm. 2021. EgoGlass: Egocentric-View Human Pose Estimation From an Eyeglass Frame. In *2021 International Conference on 3D Vision (3DV)*. 32–41. <https://doi.org/10.1109/3DV53792.2021.00014>
- K. Zhou, X. Han, N. Jiang, K. Jia, and J. Lu. 2022. HEMlets PoSh: Learning Part-Centric Heatmap Triplets for 3D Human Pose and Shape Estimation. *IEEE Transactions on Pattern Analysis & Machine Intelligence* 44, 06 (jun 2022), 3000–3014. <https://doi.org/10.1109/TPAMI.2021.3051173>

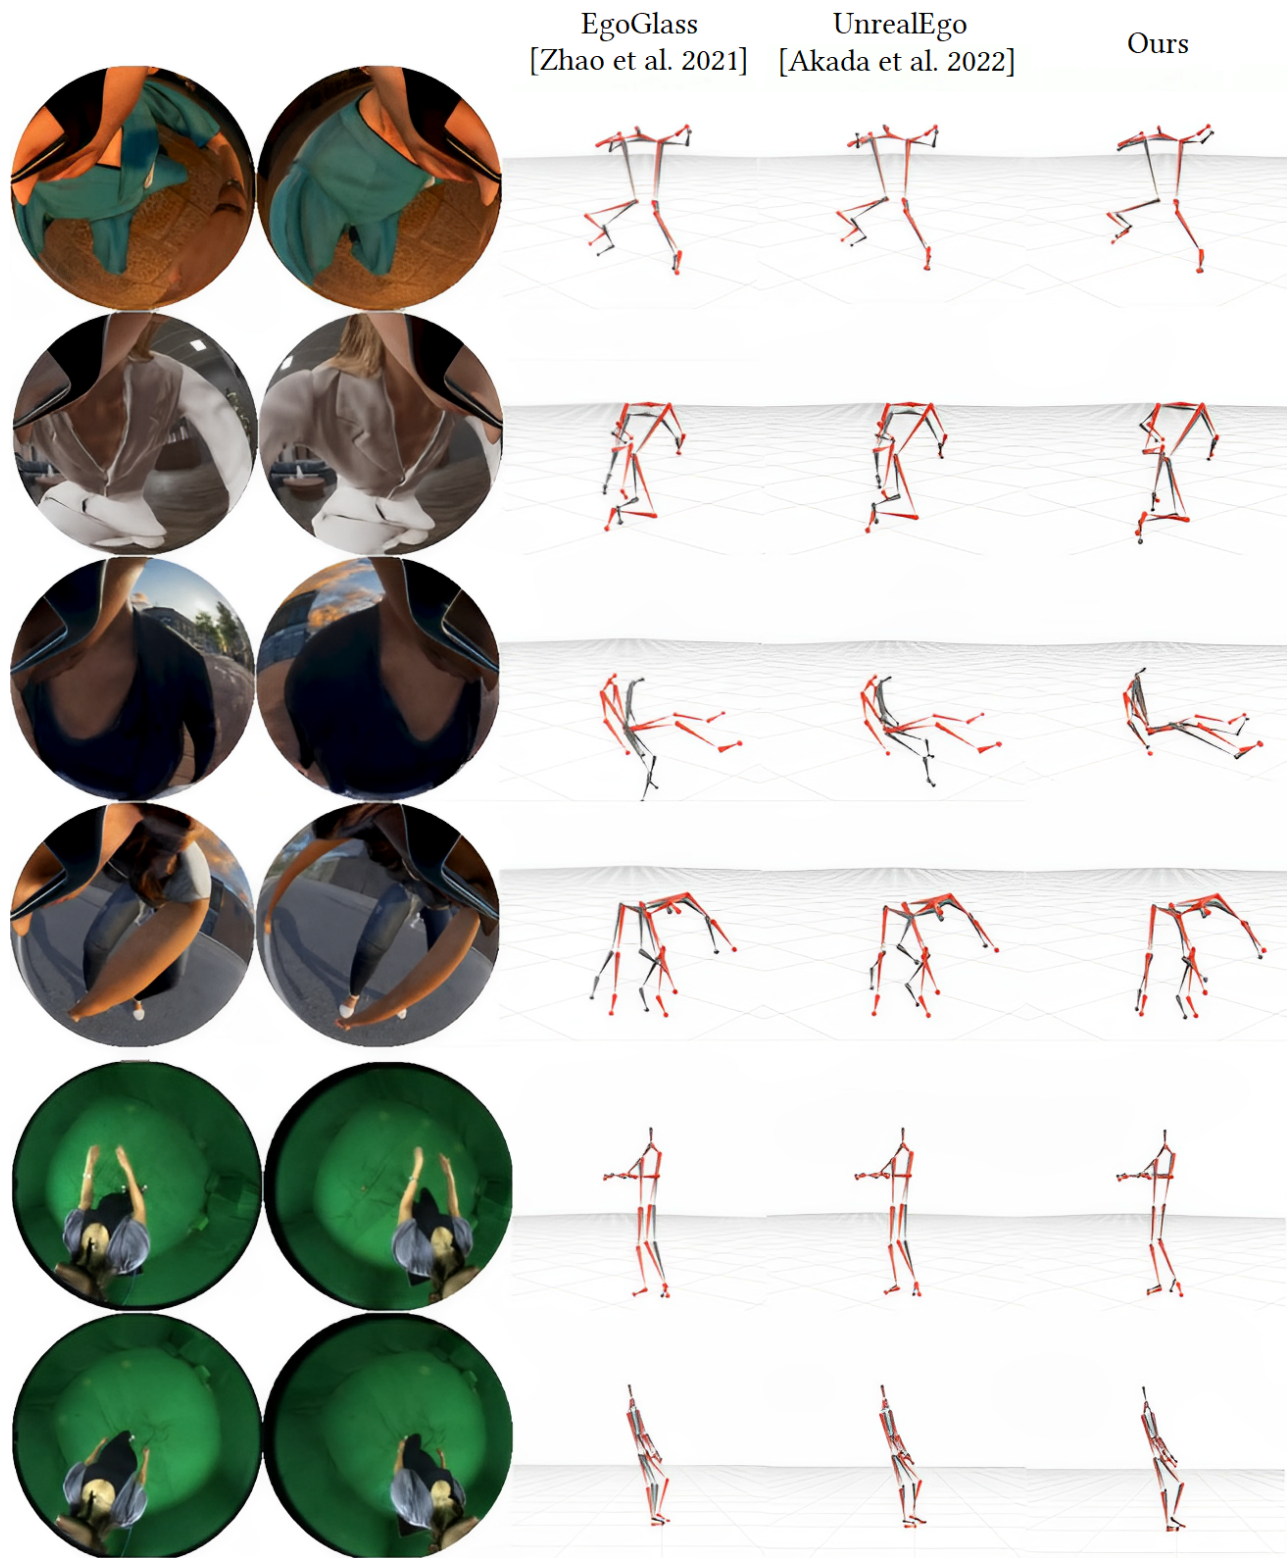


Figure 7: Qualitative evaluation of our method. The first four poses are from the synthetic UnrealEgo dataset, and the last two poses are from the real-world EgoCap dataset. The EgoCap validation 3D set has two sequences after our train/test set splits, so we show two samples from each. The red pose is the ground truth and the grey pose is the predicted pose from methods.