

# AiOS: All-in-One-Stage Expressive Human Pose and Shape Estimation

Qingping Sun<sup>\*,1,2</sup>, Yanjun Wang<sup>\*,1</sup>, Ailing Zeng<sup>3</sup>, Wanqi Yin<sup>1</sup>, Chen Wei<sup>1</sup>,  
Wenjia Wang<sup>5</sup>, Haiyi Mei<sup>1</sup>, Chi-Sing Leung<sup>2</sup>, Lei Yang<sup>1,5</sup>, Ziwei Liu<sup>4</sup>, Zhongang Cai<sup>†,1,4,5</sup>

<sup>1</sup> SenseTime Research <sup>2</sup> City University of Hong Kong

<sup>3</sup> International Digital Economy Academy (IDEA)

<sup>4</sup> S-Lab, Nanyang Technological University <sup>5</sup> Shanghai AI Laboratory

\* Equal Contributions, † Corresponding Author

<https://ttxskk.github.io/AiOS/>



Figure 1. A comparison of existing methods in EHPS. (a) Top-down, multi-stage methods typically use detectors to detect humans, then use different networks to regress body parts on cropped images. (b) Top-down, one-stage methods use only one network for regression but still require detectors and rely on the cropped image. (c) Our all-in-one-stage pipeline, end-to-end human detection, and regression on full frame.

## Abstract

Expressive human pose and shape estimation (a.k.a. 3D whole-body mesh recovery) involves the human body, hand, and expression estimation. Most existing methods have tackled this task in a two-stage manner, first detecting the human body part with an off-the-shelf detection model and then inferring the different human body parts individually. Despite the impressive results achieved, these methods suffer from 1) loss of valuable contextual information via cropping, 2) introducing distractions, and 3) lacking inter-association among different persons and body parts, inevitably causing performance degradation, especially for crowded scenes. To address these issues, we introduce a novel all-in-one-stage framework, AiOS, for multiple expressive human pose and shape recovery without an additional human detection step. Specifically, our method is built upon DETR, which treats multi-person whole-body mesh recovery task as a progressive set prediction problem with various sequential detection. We devise the decoder tokens and extend them to our task. Specifically, we first employ a human token to probe a human location in the image and encode global features for each instance, which provides a coarse location for the later transformer block. Then, we introduce a joint-related token

to probe the human joint in the image and encoder a fine-grained local feature, which collaborates with the global feature to regress the whole-body mesh. This straightforward but effective model outperforms previous state-of-the-art methods by a 9% reduction in NMVE on AGORA, a 30% reduction in PVE on EHF, a 10% reduction in PVE on ARCTIC, and a 3% reduction in PVE on EgoBody.

## 1. Introduction

Expressive human pose and shape estimation (EHPS)<sup>1</sup> is a rapidly developing area. It plays an important role in human understanding and has broad applications in the animation, gaming, and streaming industries. Unlike human pose and shape estimation (HPS), which focuses solely on the human body, EHPS is designed to jointly estimate human body poses, hand gestures, and facial expressions from the image.

In mainstream studies, the common approaches involve utilizing parametric human models, such as SMPL-X [30], to represent the articulated mesh model of a human and to regress the parameters for each body part. Drawing from research experience in single-part estimation, such as body

<sup>1</sup>EHPS is used interchangeably with 3D whole-body human mesh recovery in this work

pose and shape estimation [8, 13–15, 19, 20, 34, 39, 40, 51], existing methods [3, 7, 11, 18, 26, 27, 31] employ a multi-stage paradigm. As shown in Fig. 1a), the process begins by cropping the body parts using bounding boxes detected either by off-the-shelf detection models or provided via ground truth annotations. Following this, distinct models are utilized for the separate reconstruction of each individual body part.

Obviously, this design compromises both complexity and accuracy. The images are processed multiple times with each model. The separate parts model blocks the inter-part, inter-human connection and brings inconsistent poses and unnatural artifacts at the connected joints. Recently, OSX [21] and SMPLer-X [3] discard part experts and regress the model in a holistic manner, which alleviates the artifacts. Their paradigm can be abstract to Fig. 1b), however, they still need to be given a bounding box to crop the image. While their benchmarks show promising results, the accurate ground truth bounding boxes are not attainable in real-world scenarios. RoboSMPLX [27] has demonstrated that the performance drops significantly under noisy boxes. Moreover, CLIFF [20] points out that the cropping operation discards the location information, which degrades the performance.

A direct solution to address the challenges posed by the multi-stage paradigm is to utilize a one-stage framework that directly recovers EHPS from the entire image without requiring additional boxes for cropping. However, current one-stage methods [35–37] are proposed for HPS. Both of them use a body center heatmap and mesh parameter map to represent the potential human location and corresponding features. Relying solely on these human-centered global features is insufficient for achieving accurate part-wise regression. Although numerous two-stage HPS methods [14, 47] that extract local features in various ways, it is non-trivial to extend to a one-stage model, as most of the representations, like part-attention maps, are designed for a single person.

In order to tackle the above challenges, we have proposed the first All-in-One-Stage (AiOS) EHPS method. This novel approach is capable of predicting every individual present in an image solely based on a single image input without any additional requirements. Inspired by the achievement of DETR-based [5] methods in various vision tasks [41, 42, 46, 48, 54], we designed our pipeline in a DETR [5] style with image feature encoder and various location-aware decoders. We tailored different queries and association strategies to progressively guide the decoder to perceive global and local human features from the entire image.

Three key design features distinguish the AiOS model. **First**, it is built upon DETR structure, with a CNN backbone, transformer encoders, and decoders, and progressively detects human and decodes person features in an end-to-end manner. **Second**, we introduced the "Human-as-Tokens" design, where humans are conceptualized as a collection of box tokens and joint tokens. With different supervision and loca-

tion cues, these tokens aggregate both global and local feature representations with cross-attention for enhanced model accuracy in diverse scenarios. **Third**, using self-attention and cross-attention mechanisms in our model allows for an in-depth analysis of inter-human and intra-human relationships, enhancing performance in crowded and occlusion-heavy environments.

Extensive experiments show that our proposed model has overpass state-of-the-art (SOTA) methods that utilize ground truth bounding boxes and also SOTA methods when the bounding box is not given. Further, our bounding box is accurate enough to improve the other two-stage methods on the AGORA benchmark.

In summary, our contributions are i) The first one-stage method for EHPS that eliminates the need for extra detection networks; ii) A unified framework to integrate local and global features for whole-body regression; iii) SOTA performance on mainstream benchmarks without ground truth bounding boxes.

## 2. Related Work

### 2.1. Expressive Human Mesh Recovery Methods

EHPS focuses on reconstructing the mesh of the human body, hands, and face from monocular images. Pioneering research in this domain introduced whole-body parametric models such as SMPL-X [30]. With advancements in regression techniques for the human body, hands, and face, early studies adopted multi-stage solutions [11, 26, 29, 31]. They independently recover body pose, hand pose, and facial expressions from cropped images before integration. However, these multi-stage methods often produce artifacts at joint intersections and present complex network designs.

Given the recent surge in whole-body datasets [1, 2, 4, 28, 43], many approaches have transitioned to a holistic paradigm. OSX [21] presents a groundbreaking one-stage method, eliminating part-specific experts and cropped image regression. SMPLer-X [3] further amplifies one-stage methods utilizing large vision models and extensive datasets. However, they still rely on bounding boxes for image cropping. Despite their precision with ground truth bounding boxes, performance degrades under detected boxes [27].

### 2.2. One-Stage Human Mesh Recovery Methods

Most of the existing HPS methods [6, 8, 13, 16, 22, 39, 44, 45, 47] are multi-staged. Although these methods preserve relatively high-resolution images and generally have higher accuracy, they neglect other information in the full frame, including inter-person occlusions and individual positions [20]. To address these limitations, ROMP [35] first proposed to recover humans from an entire frame. It locates the human locations from a body center heatmap and indexes the corresponding features from the feature map to regress all human

meshes. Furthermore, BEV [37] extends the 2D heatmap to 3D by incorporating a bird-eye-view. It enables the model to discern 3D relative positions within the frame. TRACE [36] further achieved simultaneously tracking humans and predicting camera motions with added motion maps. However, these center-map-based methods often distill the human into a single vector on the feature map and recover the human pose and shape based on this global feature. We reckon that this representation is insufficient for the EHPS task, particularly given that hand pose and expression require more fine-grained local features for accurate regression.

### 3. Method

#### 3.1. Motivation

For EHPS, using cropped images presents significant problems. The cropping discards the location information [20], and inaccurate bounding boxes may lead to missing body parts, negatively impacting performance. In crowded scenes, cropping struggles to distinguish individual humans, with parts from others intruding into the frame, leading to errors in human part detection and regression. Especially when people overlap significantly, the model struggles to differentiate them due to unclear bounding boxes. Furthermore, the detectors used are typically trained on general object detection datasets and are not specifically designed for human detection, adding to these difficulties.

To tackle these problems, we introduced the AiOS, the first fully end-to-end network for EHPS. Abandoning the uncertain assumption of box-as-subject, our model leverages feature tokens and position queries for more precise human localization. We’ve developed a cohesive approach that combines global and local feature representations for accurate regression. To handle crowded scenarios and enhance the separation of human figures, our model employs attention mechanisms to establish intricate relationships between different body parts and between multiple individuals.

#### 3.2. Preliminaries

**SMPL-X.** We use 3D parametric model SMPL-X [30] to study EHPS. It utilizes a set of parameters to model body, face, and hands geometries. Specifically, our model estimates pose parameters  $\theta \in \mathbb{R}^{53 \times 3}$ , which include body poses  $\theta_{body} \in \mathbb{R}^{22 \times 3}$ , left hand poses  $\theta_{lhand} \in \mathbb{R}^{15 \times 3}$ , right hand poses  $\theta_{rhand} \in \mathbb{R}^{15 \times 3}$ , and jaw poses  $\theta_{jaw} \in \mathbb{R}^{1 \times 3}$ . Additionally, it estimates shape parameters  $\beta \in \mathbb{R}^{10}$ , and facial expression parameters  $\psi \in \mathbb{R}^{10}$ . We use the joint regressor  $J$  to obtain the 3D joint from the parameters by  $J(\mathcal{M}(\beta, \theta, \psi))$ , where  $\mathcal{M}$  is the SMPL-X function.

#### 3.3. Overview

AiOS includes the backbone and transformer encoder-decoder structures. It has three steps, 1) localize coarse body

location and extract global features of the body; 2) refine body location, extract body local features, localize coarse hands and face locations and extract global features of the hands and face; 3) refine hand and face location, extract local features for whole body.

**Backbone.** AiOS utilizes the ResNet-50 [12] to extract a multi-scale feature maps  $F_{img}$ , which provide features from detailed to holistic.

**Encoder.** As our task needs more than local associations, we utilize a standard transformer encoder [38] for long-distance relations. To transform the CNN-based feature map into a transformer-compatible feature vector, we flatten the multi-layer feature maps along their spatial dimensions and concatenate them. The flattened feature is added with position encodings  $PE \in \mathbb{R}^{M \times D}$  to derive the image feature token  $T_{img} \in \mathbb{R}^{M \times D}$ , where  $M$  represents the total length of the image feature token. We fed  $T_{img}$  a transformer encoder, which produces the refined image feature tokens  $T'_{img}$ , serving as a reference for cross-attention in the decoder. Utilizing  $T'_{img}$ , a feed-forward network (FFN) is applied to classify each token as a human token. Following the approach in DINO [49] and ED-Pose [42], we filter based on the classification score and retain the top  $M_h = 900$  tokens. These tokens serve as candidate human body localization tokens  $T_{body} \in \mathbb{R}^{M_h \times D}$ , and they also function as the input for the subsequent decoders.

**Generic Decoder.** Similar to PETR [33] and ED-Pose [42], which extend the deformable decoder [53] to 2D human body-only pose estimation, AiOS extends the deformable decoder to 3D whole-body mesh recovery. It mainly has three inputs, image content tokens  $T'_{img} \in \mathbb{R}^{M \times D}$ , object content tokens  $T \in \mathbb{R}^{M_h \times D}$  and object position queries  $Q \in \mathbb{R}^{M_h \times 4}$ . Utilizing this decoder, our model can automatically probe the suitable global and local features around the body parts for each human conditioned by various queries. We will introduce our key decoder designs in the following sections.

#### 3.4. Naive AiOS

Drawing inspiration from ROMP [35], we extend the DETR structure [5] to EHPS and progressively regress SMPL-X parameters. Specifically, we follow DAB-DETR [24] and introduce the location queries to probe the body, face, and hands-related features, guided by bounding boxes  $(x, y, w, h)$ , that considers both the location and size of each body part boxes. The model first extracts features related to the body using body box location queries and refines them through the body-location decoder. Subsequently, they are expanded to include hands and face queries and leverage the whole-body-location decoder to extract whole-body features.

**Body-location Decoder.** The first two decoders are body-centric, and the input object content tokens  $T$  are the body location tokens  $T_{bl}$ . We derive body location query  $Q_{bl}$  with FFN from the corresponding  $T_{bl}$ . The decoder first

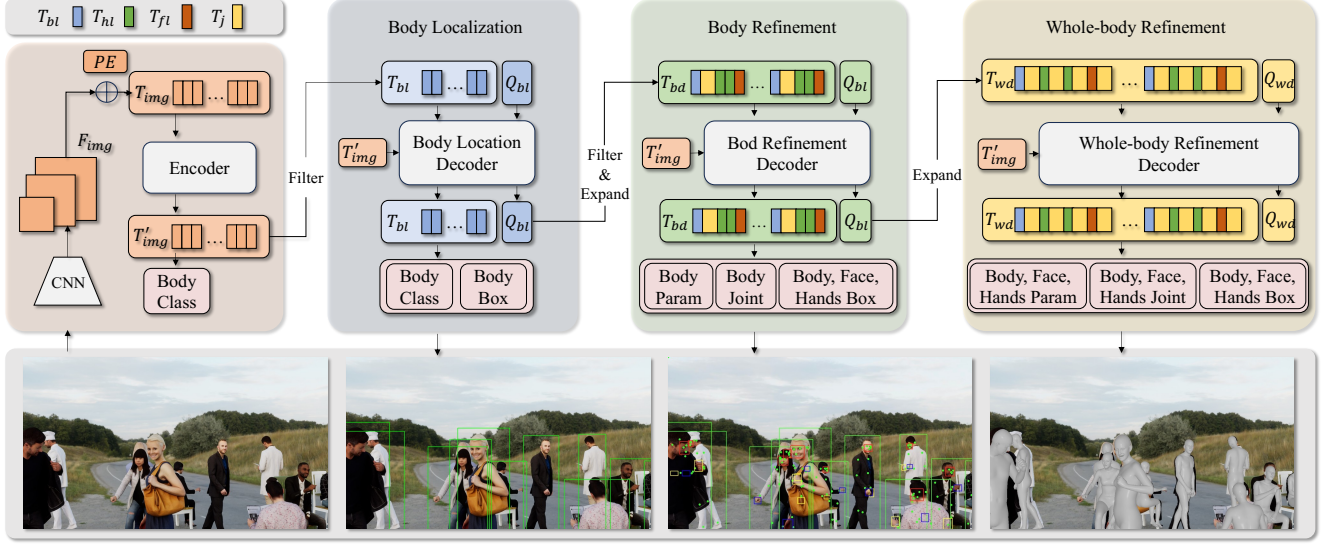


Figure 2. **Pipeline overview.** AiOS performs human localization and SMPL-X estimation in a progressive manner. It is composed of (1) the body localization stage that predicts coarse human location; (2) the Body refinement stage that refines body features and produces face and hand locations; (3) the Whole-body Refinement stage that refines whole-body features and regress SMPL-X parameters.

associates the body location tokens and updates them by the self-attention mechanism. Then, the decoder takes image tokens  $T'_{img}$  as the value and the updated body location tokens as the query for cross-attention, and the  $Q_{bl}$  acts as an indicator, which is used to aggregate the information focusing on the corresponding body area. After that, the body location tokens  $T_{bl}$  and body location queries  $Q_{bl}$  are refined with the decoder.

We estimate the body bounding box with an FFN from  $T_{bl}$ , which is supervised by  $L_{box}$ . This supervision makes sure the tokens aggregate global information of the human. Similar to the encoder, we classify the output  $T_{bl}$  with an FFN on whether it is a token representing a human. The classification results from  $T_{bl}$  are supervised with classification loss  $L_{cls}$ . At the end of the second decoder, we downsample the body tokens again to  $M_b = 100$  to further distill potential human tokens and lower the computational complexity.

**Whole-body-location Decoder.** The latter four decoders of naive AiOS jointly consider whole-body information and their association. With the body location tokens from the previous step, we expand them to hands and face location tokens with learnable embedding. We first broadcast the given embedding  $E_{bl} \in \mathbb{R}^D$  and add it to the body location token  $T_{bl} \in \mathbb{R}^{M_b \times D}$ . After that, we obtained hand location tokens  $T_{lhl}$ ,  $T_{rhl}$ , and face location tokens  $T_{fl}$ , which have the same shape as  $T_{bl}$ . Then we concat them into a whole-body token  $T_{full} = [T_{bl}, T_{lhl}, T_{rhl}, T_{fl}]$ . Similarly, the whole-body location queries  $Q_{full}$  are expanded from  $Q_{bl}$  with learnable embeddings.

The decoders use a self-attention module to explore inter-part and inter-human relations and then extract each part's features around their bounding boxes with a conditioned

cross-attention module. We utilize an attention mask to ensure that the bounding boxes for each person's hands and face are associated only with their own and others' body bounding boxes. As our model is already capable of recognizing each person's body in the first two stages, this specific attention mechanism allows for more accurate identification of body parts in crowded scenes. We provide an illustration of the attention mechanism in the Supplementary Material.

We regress body bounding boxes from  $T_{bl}$ , face boxes from  $T_{fl}$  and hand boxes from  $T_{rhl}$ ,  $T_{lhl}$ , and supervise them with  $L_{box}$ . We regress different part's parameters from the refined whole body  $T_{full}$  tokens. The parameters are supervised with SMPL-X loss  $L_{smplx}$ , which includes parameter loss  $L_{param}$ , 3d keypoints loss  $L_{kp3d}$ , and the 2d keypoints reprojection loss  $L_{kp2d}$ .

### 3.5. AiOS

Previous methods [35, 37] have shown that regressing multi-person body meshes from global features alone can achieve impressive results, but in EHPS, relying on global information alone is insufficient. The model should also consider local information to obtain a detailed context of the whole-body regression. Therefore, to elevate the model's ability, we introduce joint-related tokens and their corresponding queries to our model. Combined with location tokens, the AiOS expresses human context in multilevel. We will further regress the SMPL-X parameter on this well-rounded feature group. Specifically, we adopt a progressive detection and decoding strategy. The first two layers are body-location decoders same as our naive design, which outputs coarse human location. Further, two layers of body-refinement decoders utilize body joint tokens to enrich local body features



and estimate rough hand and face location simultaneously on the basis of human location. At last, two layers of whole-body-refinement decoders extract whole-body local features with extra hands and face joint tokens.

**Body-refinement Decoder.** This decoder is built on body-location decoders in naive AiOS. In detail, we expand body joints tokens, hands location tokens, and face location tokens. We adopt the learnable-embedding  $E_{bj} \in \mathbb{R}^{17 \times D}$  to expand body joint tokens  $T_{bj} \in \mathbb{R}^{M_b \times 17 \times D}$  from box location tokens, and then we obtain detailed body token set  $T_{bd} = [T_{bl}, T_{bj}, T_{lhl}, T_{rhl}, T_{fl}]$ . Note that we use an attention mask to limit the joint attention within its subject as inter-joint attention among different subjects brings no incremental but much higher computation complexity.

The  $T_{bd}$  are refined with layers of decoders. Within each layer, similar to naive AiOS, we regress bounding boxes of body parts from their location tokens and supervise them with  $L_{box}$ . Further, we regress body joint location from  $T_{bj}$  and supervise them with  $L_{j2d}$ , helping these joint tokens learn the local human features. Different from Naive AiOS, in this stage, we regress SMPL-X body parameters based on  $T_{bl}, T_{bj}$ . We use  $L_{smplx}$  to supervise the body parameter, helping to refine the body-related tokens representing more accurate body features.

**Whole-body-refinement Decoder.** This decoder further expands the face and hand joint tokens. Similarly, we use embedding  $E_{lhj}, E_{rhj},$  and  $E_{fj}$  to expand  $T_{lhl}, T_{rhl},$  and  $T_{fl}$  to  $T_{lhj}, T_{rhj},$  and  $T_{fj}$ , respectively. At this stage, the model forms the complete tokens that represent a human  $T_{wd} = [T_{bl}, T_{bj}, T_{lhl}, T_{lhj}, T_{rhl}, T_{rhj}, T_{fl}, T_{fj}]$ .

Based on  $T_{wd}$ , we utilize FFN to regress box location from  $T_{bl}, T_{lhl}, T_{rhl}, T_{fl}$  and supervised with  $L_{box}$ . We also regress whole-body joint location from  $T_{bj}, T_{lhj}, T_{rhj},$  and  $T_{fj}$ , and supervise them with  $L_{j2d}$ . Finally, we estimate SMPL-X body, hands, and face parameters from body, hand, and face-related tokens, respectively, and supervise whole-body parameters with  $L_{smplx}$ .

**Overall Loss Functions.** The overall loss function is the sum of all the losses at each stage. Please refer to the Supplementary Material for the details.

## 4. Experiment

### 4.1. Experimental Setup

Due to the page limit, we put the detailed experiment setup, implementation, and partial quantitative and qualitative comparison with SOTA methods in the Supplementary Material.

**Datasets.** AiOS is trained on the multi-person datasets AGORA [28], BEDLAM [1], and COCO [23], and single-person datasets UBody [21], ARCTIC [9], and EgoBody [52]. We evaluate it on AGORA, UBody, EHF [30], ARCTIC [9], EgoBody [52], and BEDLAM [1].

**Implementation.** The training is conducted on 16 V100

GPUs, with a total batch size of 32. We first train our model for 60 epochs on AGORA, BEDLAM, and COCO. We fine-tune it for 50 epochs on all train datasets.

**Evaluation metrics.** Following the previous EHPS methods [3, 21, 26], we report Procrustes Aligned per-vertex position error (PA-MPVPE) and the mean per-vertex position error (MPVPE) across all benchmarks. In AGORA Leaderboard, we report mean vertex error (MVE), mean per-joint position error (MPJPE) for pure reconstruction accuracy; F Score, precision, recall for detection accuracy; Normalized mean vertex error (NMVE) and normalized mean joint error (NMJE) that considered regression accuracy with detection accuracy. All metrics are reported in millimeters (mm).

### 4.2. Quantitative comparison with SOTA

In Table 1, we compare AiOS with the SOTA methods on the AGORA test set. The results are provided by the leaderboard<sup>2</sup> with their bounding boxes on the upper part of the table. We also feed our estimated bounding boxes to OSX [21] and SMPLer-X [3] on the lower part, which helps to verify our model’s localization quality.

For a fair comparison with the SOTA methods, we utilize a threshold of 0.5 to filter the detected samples with lower confidence, which generally have severe occlusions. As shown on the upper part of Table 1, our model’s NMVE and NMJE greatly surpass the current SOTA method SMPLer-X. This observation proves that our one-stage pipeline achieves the best overall quality, combining localization and reconstruction. In terms of pure reconstruction quality, our model also achieves SOTA performance with a relatively accurate detection result on MVE and MPJPE. While BEDLAM [1] excels in face and hand reconstruction, its recall performance is comparatively low, omitting some instances for evaluation.

On the lower-part comparison, we lower the detection threshold to 0.3, which has higher recall than any current results, allowing more hard cases to be detected. We feed the same bounding boxes to the OSX and SMPLer-X, and their performance on whole-body MVE improves compared with the results reported in the original paper (122.8 to 121.3 for OSX, 99.7 to 98.3 for SMPLer-X) even with a higher recall. This indicates that improvement is achieved not by filtering out hard cases but by providing high-quality bounding boxes. This finding proves that the current two-stage method is sensitive to bounding box quality, and using the ground truth box to crop images in other benchmarks is biased from real use cases. Notably, under this bounding box setting, AiOS is still much higher than the current SOTA OSX and comparable with the foundation model SMPLer-X L20.

As the first one-stage method in EHPS, we cannot find relevant one-stage methods for a fair comparison. Therefore, similar to H4W [26], we compare the results of our body part with existing body-only methods, which can be broadly

<sup>2</sup><https://agora-evaluation.is.tuebingen.mpg.de/>

Methods	F Score $\uparrow$	Precision $\uparrow$	Recall $\uparrow$	NMVE $\downarrow$ (mm)		NMJE $\downarrow$ (mm)		MVE $\downarrow$ (mm)					MPJPE $\downarrow$ (mm)				
				All	Body	All	Body	All	Body	Face	LHand	RHand	All	Body	Face	LHand	RHand
BEDLAM [1]	0.73	<b>0.98</b>	0.59	179.5	132.2	177.5	131.4	131.0	96.5	25.8	38.8	<b>39.0</b>	129.6	95.9	27.8	36.6	<b>36.7</b>
H4W [26] $\dagger$	<b>0.94</b>	<b>0.96</b>	<b>0.92</b>	144.1	96.0	141.1	92.7	135.5	90.2	41.6	46.3	48.1	132.6	87.1	46.1	44.3	46.2
BEDLAM [1] $\dagger$	0.73	<b>0.98</b>	0.59	142.2	102.1	141.0	101.8	103.8	74.5	<b>23.1</b>	<b>31.7</b>	<b>33.2</b>	102.9	74.3	<b>24.7</b>	<b>29.9</b>	<b>31.3</b>
PyMaF-X [50] $\dagger$	0.89	0.90	0.89	141.2	94.4	140.0	93.5	125.7	84.0	35.0	44.6	45.6	124.6	83.2	37.9	42.5	43.7
OSX [21] *	<b>0.94</b>	<b>0.96</b>	<b>0.93</b>	130.6	85.3	127.6	83.3	122.8	80.2	36.2	45.4	46.1	119.9	78.3	37.9	43.0	43.9
HybrIK-X [18]	<b>0.93</b>	0.95	<b>0.92</b>	120.5	73.7	115.7	72.3	112.1	68.5	37.0	46.7	47.0	107.6	67.2	38.5	41.2	41.4
SMPLer-X [3]	<b>0.93</b>	<b>0.96</b>	0.90	133.1	88.1	128.9	84.6	123.8	81.9	37.4	43.6	44.8	119.9	78.7	39.5	41.4	44.8
SMPLer-X [3] $\dagger$	<b>0.93</b>	<b>0.96</b>	0.90	107.2	68.3	104.1	66.3	99.7	63.5	29.9	39.1	39.5	96.8	61.7	31.4	36.7	37.2
Native AiOS	<b>0.93</b>	<b>0.98</b>	0.89	<b>105.7</b>	<b>66.5</b>	<b>103.9</b>	<b>65.8</b>	<b>98.3</b>	<b>61.8</b>	27.2	40.7	41.7	<b>96.6</b>	<b>61.2</b>	28.4	38.4	39.4
AiOS	<b>0.94</b>	<b>0.98</b>	0.90	<b>97.8</b>	<b>61.3</b>	<b>96.0</b>	<b>60.7</b>	<b>91.9</b>	<b>57.6</b>	<b>24.6</b>	<b>38.7</b>	39.6	<b>90.2</b>	<b>57.1</b>	<b>25.7</b>	<b>36.4</b>	37.3
OSX [21] $\diamond$	<b>0.96</b>	<b>0.97</b>	<b>0.95</b>	126.4	81.8	123.4	80.0	121.3	78.5	36.1	45.9	46.3	118.5	76.8	37.6	43.5	44.0
SMPLer-X [3] $\dagger$ $\diamond$	<b>0.96</b>	<b>0.97</b>	<b>0.95</b>	<b>102.4</b>	<b>63.8</b>	<b>99.5</b>	<b>62.1</b>	<b>98.3</b>	<b>61.2</b>	<b>30.3</b>	<b>40.4</b>	<b>40.7</b>	<b>95.5</b>	<b>59.6</b>	<b>31.7</b>	<b>37.9</b>	<b>38.2</b>
AiOS	<b>0.96</b>	<b>0.97</b>	<b>0.95</b>	<b>103.0</b>	<b>63.5</b>	<b>100.8</b>	<b>62.6</b>	<b>98.9</b>	<b>61.0</b>	<b>27.7</b>	<b>42.5</b>	<b>43.4</b>	<b>96.8</b>	<b>60.1</b>	<b>29.2</b>	<b>40.1</b>	<b>40.9</b>

Table 1. **AGORA SMPL-X test set**.  $\dagger$  denotes the methods finetuned on the AGORA training set. \* denotes the methods trained on the AGORA training set only.  $\diamond$  denotes the methods that use the AiOS’s bounding box to crop the image. The **best results** are colored with **red**, and the **second-best results** are colored with **blue** for the upper and lower parts of the table, respectively.

Methods	F1-score $\uparrow$	Precision $\uparrow$	Recall $\uparrow$	NMVE $\downarrow$	NMJE $\downarrow$	MVE $\downarrow$	MPJPE $\downarrow$
Top-down Methods							
HMR [13]	0.80	0.93	0.70	217.0	226.0	173.6	180.5
PyMAF [51]	0.84	0.86	0.82	200.2	207.4	168.2	174.2
PARE [14]	0.84	0.96	0.75	167.7	174.0	140.9	146.2
H4W [26] $\dagger$	<b>0.94</b>	<b>0.96</b>	<b>0.93</b>	90.2	95.5	84.8	89.8
CLIFF [20] $\dagger$	0.91	0.96	0.87	83.5	89.0	76.0	81.0
HybrIK [19] $\dagger$	0.91	0.92	0.90	81.2	84.6	73.9	77.0
ProPose [10] $\dagger$	0.90	0.91	0.89	78.8	82.7	70.9	74.4
PLIKS [32] $\dagger$	<b>0.94</b>	0.95	<b>0.93</b>	71.6	76.1	67.3	71.5
NIKI [17] $\dagger$	0.91	0.92	0.90	70.2	74.0	63.9	67.3
One-stage Methods							
ROMP [35] $\dagger$	0.91	0.95	0.88	113.6	118.8	103.4	108.1
BEV [37] $\dagger$	0.93	0.96	0.90	108.3	113.2	100.7	105.3
AiOS <sub>0.5</sub>	<b>0.94</b>	<b>0.98</b>	0.90	<b>61.2</b>	<b>68.0</b>	<b>57.5</b>	<b>63.9</b>
AiOS <sub>0.3</sub>	<b>0.96</b>	<b>0.97</b>	<b>0.95</b>	<b>63.4</b>	<b>70.1</b>	<b>60.9</b>	67.3

Table 2. **AGORA SMPL test set**.  $\dagger$  indicates that this method is fine-tuned on the AGORA training set. AiOS<sub>0.5</sub> and AiOS<sub>0.3</sub>, representing the use of a 0.5 score threshold and a 0.3 score threshold to filter the data, respectively.

categorized into top-down methods [10, 13, 14, 17, 19, 20, 32, 51] and one-stage methods [35, 37], on the AGORA SMPL test set. Specifically, we downsample the SMPL-X mesh estimated by AiOS to the SMPL [25] mesh using official tools [30] and then measure MVE and NMVE. We use the J-regressor to regress joints from the downsampled SMPL mesh to measure NMJE and MPJPE.

As shown in Table 2, even though AiOS is designed for EHPS, it still outperforms ROMP [35] and BEV [37], with a notable improvement in NMVE of 43% (from 108.3 mm to 61.2 mm) and an NMJE enhancement of 40% (from 113.2 mm to 68.0 mm). It is worth noting that we do not deliberately fine-tune our model exclusively on AGORA.

**Single datasets.** We compare UBody in Table 3, EHF in Table 4. Note that the other methods utilize ground-truth bounding boxes. Without any given bounding boxes, our model achieves SOTA performance on real-life datasets.

### 4.3. Qualitative comparison with SOTA

We perform a qualitative comparison with current SOTA methods on AGORA and EHF. To overlay the results onto the image, we apply an affine transformation for the two-stage methods that use images cropped by ground truth boxes.

Method	PA-PVE $\downarrow$ (mm)			PVE $\downarrow$ (mm)		
	All	Hands	Face	All	Hands	Face
PIXIE [11]	61.7	12.2	4.2	168.4	55.6	45.2
H4W [26]	44.8	8.9	2.8	104.1	45.7	27.0
OSX [21]	42.4	10.8	<b>2.4</b>	92.4	47.7	24.9
OSX [21] $\dagger$	42.2	<b>8.6</b>	<b>2.0</b>	81.9	41.5	21.2
SMPLer-X [3]	33.2	10.6	2.8	61.5	43.3	23.1
SMPLer-X [3] $\dagger$	<b>31.9</b>	10.3	2.8	<b>57.4</b>	<b>40.2</b>	21.6
Native AiOS	35.6	<b>8.6</b>	2.9	62.7	41.3	<b>20.8</b>
AiOS	<b>32.5</b>	<b>7.3</b>	2.8	<b>58.6</b>	<b>39.0</b>	<b>19.6</b>

Table 3. **UBody**.  $\dagger$  indicates the model is finetuned with the UBody training set.

Method	PA-PVE $\downarrow$ (mm)			PVE $\downarrow$ (mm)		
	All	Hands	Face	All	Hands	Face
H4W [26]	50.3	<b>10.8</b>	5.8	76.8	<b>39.8</b>	26.1
OSX [21]	48.7	15.9	6.0	70.8	53.7	26.4
SMPLer-X [3]	<b>37.8</b>	15.0	5.1	65.4	49.4	17.4
Native AiOS	38.8	13.8	<b>4.0</b>	<b>50.2</b>	49.8	<b>17.3</b>
AiOS	<b>34.0</b>	<b>12.8</b>	<b>3.8</b>	<b>45.4</b>	<b>44.1</b>	<b>16.9</b>

Table 4. **EHF**. As EHF is absent from our training data, it serves as a valuable tool to assess the generalization ability of our models.

In contrast, our method can be directly overlaid on the image. Further, with accurate betas estimation, we are able to recover the depth order, as shown in the Fig. 3. We achieve comparable visual quality in both scenes, proving our model’s accuracy.

We further perform a qualitative comparison with SOTA one-stage methods [35, 37]. As shown in Fig. 4, while ROMP and BEV can achieve decent results for body reconstruction in multi-person scenarios, they are limited by the constraints of the SMPL [25] model, preventing them from reconstructing detailed hand gestures and facial expressions.

### 4.4. Ablation Study

In this subsection, we analyze the effectiveness of the proposed components in detail. All experiments are conducted






Figure 3. Comparison of current SOTA methods [3, 21, 26] with our AiOS model. The upper part is visualization results on AGORA [28], and the lower is EHF test [7].




Figure 4. Visual comparisons with SOTA one-stage HPS methods [35, 37] on the Internet data<sup>3</sup>.

on the AGORA validation set.

**Analysis of the naive AiOS and full AiOS.** Whole-body mesh recovery requires attention to both small-scale gestures, expression details, and large-scale pose details. To validate the effectiveness of our joint-guided local feature query, we compared naive AiOS and full AiOS models across the benchmarks. The table shows that even the naive setting

<sup>3</sup><https://www.pexels.com/>

achieves comparable performance with SOTA methods, indicating our one-stage pipeline, which treats EHPS as a progressive set prediction problem with various sequential detections following the DETR, is ideal for SMPL-X parameter regression. On this solid base, the full AiOS consistently achieves higher accuracy on all parts of the human, and the increment on the whole-body aspect is especially outstanding. Since the body tends to have a relatively higher area on




Figure 5. **Attention Visualization.** The green dots represent the location of the reference point, and the red dots are the sampling points.

the image, adding joint queries to the body provides a large number of local features for reference, while for smaller areas like face and gestures, the difference between global and local features is not that obvious. However, adding the local joints feature overall brings more comprehensive features.

**The Scheme of the SMPL-X supervision.** In this part, we investigate how to supervise different tokens. For our original AiOS, we don’t supervise the SMPL-X parameter in the first stage, as we want the model to focus on body localization. In the second stage, we don’t supervise hands and face for the same reason, but supervise SMPL-X body parameters as we have detailed body feature tokens. And we supervise the whole body parameter at the third stage. In the first ablation setting, we add body parameter supervision in every stage and hand and face supervision in the second stage, meaning every stage has SMPL-X supervision. In the second setting, we remove the SMPL-X body supervision in the second stage so that the model will be only supervised by SMPL-X in the last stage. As shown in Table 5, a comparison between AiOS and all stage settings shows adding SMPL-X parameters when the location is not properly refined will hinder the model’s performance. Comparing the AiOS and 3rd stage setting shows the design of gradually whole-body estimation from body to whole-body increased performance.

**The association between the human body, hands, and face.** We focus on the self-attention relations on this part. Our stock design allows free attention among body, face, and hand location tokens, but limits joint tokens to only attend with tokens belonging to the same human. In the full attention setting, we allow tokens to any other tokens. The inter-person setting will further limit the hand and face location tokens to attend with only its subject. As shown in Table 5, the unlimited setting is the worst, as the complicated attention mechanism is not properly learned. And the limited setting is also not ideal compared to our original attention mechanism. Furthermore, we visualize the cross-attention of our model. As shown in Fig. 5, our model is able to

Ablation Studies	PA-PVE↓ (mm)			PVE↓ (mm)		
	All	Hands	Face	All	Hands	Face
Attention Format						
Full	42.5	7.2	4.2	54.8	39.0	25.8
Inter-human Only	41.7	7.3	4.2	52.8	38.9	24.5
Ours	<b>39.9</b>	<b>7.2</b>	<b>4.1</b>	<b>50.5</b>	<b>37.4</b>	<b>23.3</b>
SMPL-X Supervision Manners						
All stages	42.7	7.4	4.2	55.7	39.8	25.1
3rd stage only	40.3	7.2	4.2	51.8	38.0	23.8
Ours (2,3 stage)	<b>39.9</b>	<b>7.2</b>	<b>4.1</b>	<b>50.5</b>	<b>37.4</b>	<b>23.3</b>

Table 5. **Ablation Studies.** The upper part studies the attention format, and the bottom part studies the SMPL-X supervision manners.

localize global features with body location tokens and local features with joint tokens. The lower part shows the attention map under occlusion, and it shows that our model will take reference from other body parts.

## 5. Conclusion

In this work, we propose the first all-in-one-stage model for expressive human pose and shape estimation. We explored the incorporation of body-, face-, and hand-related tokens, as well as the aggregation of local and global features with various supervision. Moreover, we carefully designed a self-attention mechanism to establish the associations between inter- and intra-human body and body parts, which helps us to achieve its best performance. The SOTA results indicate our one-stage pipeline, which treats EHPS as a progressive set prediction problem with various sequential detections following the DETR, is a crucial factor contributing to the overall performance. This can be further proved by the performance of our naive AiOS baseline. We hope this work can contribute new insights to the EHPS research community.

**Limitations.** First, our model achieves SOTA, but there is still a large room for improvement if we add more datasets for training, particularly those containing multi-person real data. Second, the versatile design can be further extended with more dimensions of human perception tasks such as tracking and 3D localization. Exploring the estimation of hands under limited resolution is also worth investigating.

**Acknowledgement.** This project is supported by the Hong Kong Innovation and Technology Commission (InnoHK Project CIMDA). It is also supported by the Ministry of Education, Singapore, under its MOE AcRF Tier 2 (MOET2EP20221- 0012), NTU NAP, and under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).



## References

- [1] Michael J Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. Bedlam: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8726–8737, 2023. 2, 5, 6
- [2] Zhongang Cai, Daxuan Ren, Ailing Zeng, Zhengyu Lin, Tao Yu, Wenjia Wang, Xiangyu Fan, Yang Gao, Yifan Yu, Liang Pan, et al. Humman: Multi-modal 4d human dataset for versatile sensing and modeling. In *Eur. Conf. Comput. Vis.*, pages 557–577. Springer, 2022. 2
- [3] Zhongang Cai, Wanqi Yin, Ailing Zeng, Chen Wei, Qingping Sun, Yanjun Wang, Hui En Pang, Haiyi Mei, Mingyuan Zhang, Lei Zhang, et al. Smpler-x: Scaling up expressive human pose and shape estimation. *arXiv preprint arXiv:2309.17448*, 2023. 2, 5, 6, 7
- [4] Zhongang Cai, Mingyuan Zhang, Jiawei Ren, Chen Wei, Daxuan Ren, Zhengyu Lin, Haiyu Zhao, Lei Yang, and Ziwei Liu. Playing for 3d human recovery. *arXiv preprint arXiv:2110.07588*, 2021. 2
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Eur. Conf. Comput. Vis.*, pages 213–229. Springer, 2020. 2, 3
- [6] Junhyeong Cho, Kim Youwang, and Tae-Hyun Oh. Cross-attention of disentangled modalities for 3d human mesh recovery with transformers. In *Eur. Conf. Comput. Vis.*, pages 342–359. Springer, 2022. 2
- [7] Vasileios Choutas, Georgios Pavlakos, Timo Bolkart, Dimitrios Tzionas, and Michael J Black. Monocular expressive body regression through body-driven attention. In *Eur. Conf. Comput. Vis.*, pages 20–40. Springer, 2020. 2, 7
- [8] Zhiyang Dou, Qingxuan Wu, Cheng Lin, Zeyu Cao, Qiangqiang Wu, Weilin Wan, Taku Komura, and Wenping Wang. Tore: Token reduction for efficient human mesh recovery with transformer. *arXiv preprint arXiv:2211.10705*, 2022. 2
- [9] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J Black, and Otmar Hilliges. Arctic: A dataset for dexterous bimanual hand-object manipulation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12943–12954, 2023. 5
- [10] Qi Fang, Kang Chen, Yinghui Fan, Qing Shuai, Jiefeng Li, and Weidong Zhang. Learning analytical posterior probability for human mesh recovery. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2023. 6
- [11] Yao Feng, Vasileios Choutas, Timo Bolkart, Dimitrios Tzionas, and Michael Black. Collaborative regression of expressive bodies using moderation. In *International Conference on 3D Vision (3DV)*, pages 792–804, Dec. 2021. 2, 6
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. 3
- [13] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7122–7131, 2018. 2, 6
- [14] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *Int. Conf. Comput. Vis.*, pages 11127–11137, 2021. 2, 6
- [15] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Int. Conf. Comput. Vis.*, pages 2252–2261, 2019. 2
- [16] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4501–4510, 2019. 2
- [17] Jiefeng Li, Siyuan Bian, Qi Liu, Jiasheng Tang, Fan Wang, and Cewu Lu. NIKI: Neural inverse kinematics with invertible neural networks for 3d human pose and shape estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2023. 6
- [18] Jiefeng Li, Siyuan Bian, Chao Xu, Zhicun Chen, Lixin Yang, and Cewu Lu. Hybrik-x: Hybrid analytical-neural inverse kinematics for whole-body mesh recovery. *arXiv preprint arXiv:2304.05690*, 2023. 2, 6
- [19] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3383–3393, 2021. 2, 6
- [20] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. Cliff: Carrying location information in full frames into human pose and shape estimation. In *Eur. Conf. Comput. Vis.*, pages 590–606. Springer, 2022. 2, 3, 6
- [21] Jing Lin, Ailing Zeng, Haoqian Wang, Lei Zhang, and Yu Li. One-stage 3d whole-body mesh recovery with component aware transformer. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 21159–21168, 2023. 2, 5, 6, 7
- [22] Kevin Lin, Lijuan Wang, and Zicheng Liu. Mesh graphormer. In *Int. Conf. Comput. Vis.*, pages 12939–12948, 2021. 2
- [23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Eur. Conf. Comput. Vis.*, pages 740–755. Springer, 2014. 5
- [24] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. DAB-DETR: Dynamic anchor boxes are better queries for DETR. In *Int. Conf. Learn. Represent.*, 2022. 3
- [25] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM Trans. Graph.*, 34(6):1–16, 2015. 6
- [26] Gyeongsik Moon, Hongsuk Choi, and Kyoung Mu Lee. Accurate 3d hand pose estimation for whole-body 3d human mesh estimation. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, 2022. 2, 5, 6, 7
- [27] Hui En Pang, Zhongang Cai, Lei Yang, Qingyi Tao, Zhonghua Wu, Tianwei Zhang, and Ziwei Liu. Towards robust and expressive whole-body human pose and shape estimation. In *Adv. Neural Inform. Process. Syst.*, 2023. 2
- [28] Priyanka Patel, Chun-Hao P Huang, Joachim Tesch, David T Hoffmann, Shashank Tripathi, and Michael J Black. AGORA: Avatars in geography optimized for regression analysis. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 13468–13478, 2021. 2, 5, 7

- [29] Georgios Pavlakos, Vasileios Choutas, Timo Bolkart, Dimitrios Tzionas, Michael J. Black, Vasileios Choutas, Georgios Pavlakos, Timo Bolkart, Dimitrios Tzionas, and Michael J. Black. Monocular expressive body regression through body-driven attention. *Eur. Conf. Comput. Vis.*, 2020. [2](#)
- [30] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. [1](#), [2](#), [3](#), [5](#), [6](#)
- [31] Yu Rong, Takaaki Shiratori, and Hanbyul Joo. Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration. In *Int. Conf. Comput. Vis. Worksh.*, 2021. [2](#)
- [32] Karthik Shetty, Annette Birkhold, Srikrishna Jaganathan, Norbert Strobel, Markus Kowarschik, Andreas Maier, and Bernhard Egger. Pliks: A pseudo-linear inverse kinematic solver for 3d human body estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 574–584, 2023. [6](#)
- [33] Dahu Shi, Xing Wei, Liangqi Li, Ye Ren, and Wenming Tan. End-to-end multi-person pose estimation with transformers. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11059–11068, 2022. [3](#)
- [34] Qingping Sun, Yi Xiao, Jie Zhang, Shizhe Zhou, Chi-Sing Leung, and Xin Su. A local correspondence-aware hybrid cnn-gcn model for single-image human body reconstruction. *IEEE Transactions on Multimedia*, 25:4679–4690, 2023. [2](#)
- [35] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Michael J Black, and Tao Mei. Monocular, one-stage, regression of multiple 3d people. In *Int. Conf. Comput. Vis.*, pages 11179–11188, 2021. [2](#), [3](#), [4](#), [6](#), [7](#)
- [36] Yu Sun, Qian Bao, Wu Liu, Tao Mei, and Michael J Black. Trace: 5d temporal regression of avatars with dynamic cameras in 3d environments. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8856–8866, 2023. [3](#)
- [37] Yu Sun, Wu Liu, Qian Bao, Yili Fu, Tao Mei, and Michael J Black. Putting people in their place: Monocular regression of 3d people in depth. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 13243–13252, 2022. [2](#), [3](#), [4](#), [6](#), [7](#)
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Adv. Neural Inform. Process. Syst.*, 30, 2017. [3](#)
- [39] Wenjia Wang, Yongtao Ge, Haiyi Mei, Zhongang Cai, Qingping Sun, Yanjun Wang, Chunhua Shen, Lei Yang, and Taku Komura. Zolly: Zoom focal length correctly for perspective-distorted human mesh reconstruction. *arXiv preprint arXiv:2303.13796*, 2023. [2](#)
- [40] Yanjun Wang, Qingping Sun, Wenjia Wang, Jun Ling, Zhongang Cai, Rong Xie, and Li Song. Learning dense uv completion for human mesh recovery. *arXiv preprint arXiv:2307.11074*, 2023. [2](#)
- [41] Jie Yang, Ailing Zeng, Feng Li, Shilong Liu, Ruimao Zhang, and Lei Zhang. Neural interactive keypoint detection. In *Int. Conf. Comput. Vis.*, pages 15122–15132, 2023. [2](#)
- [42] Jie Yang, Ailing Zeng, Shilong Liu, Feng Li, Ruimao Zhang, and Lei Zhang. Explicit box detection unifies end-to-end multi-person pose estimation. *arXiv preprint arXiv:2302.01593*, 2023. [2](#), [3](#)
- [43] Zhitao Yang, Zhongang Cai, Haiyi Mei, Shuai Liu, Zhaoxi Chen, Weiye Xiao, Yukun Wei, Zhongfei Qing, Chen Wei, Bo Dai, et al. Synbody: Synthetic dataset with layered human models for 3d human perception and modeling. *arXiv preprint arXiv:2303.17368*, 2023. [2](#)
- [44] Ailing Zeng, Xuan Ju, Lei Yang, Ruiyuan Gao, Xizhou Zhu, Bo Dai, and Qiang Xu. Deciwatc: A simple baseline for 10× efficient 2d and 3d pose estimation. In *Eur. Conf. Comput. Vis.*, pages 607–624. Springer, 2022. [2](#)
- [45] Ailing Zeng, Lei Yang, Xuan Ju, Jiefeng Li, Jianyi Wang, and Qiang Xu. Smoothnet: A plug-and-play network for refining human poses in videos. In *Eur. Conf. Comput. Vis.*, pages 625–642. Springer, 2022. [2](#)
- [46] Fangao Zeng, Bin Dong, Yuang Zhang, Tiancai Wang, Xiangyu Zhang, and Yichen Wei. Motr: End-to-end multiple-object tracking with transformer. In *Eur. Conf. Comput. Vis.*, 2022. [2](#)
- [47] Wang Zeng, Wanli Ouyang, Ping Luo, Wentao Liu, and Xiaogang Wang. 3d human mesh regression with dense correspondence. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7054–7063, 2020. [2](#)
- [48] Aixi Zhang, Yue Liao, Si Liu, Miao Lu, Yongliang Wang, Chen Gao, and Xiaobo Li. Mining the benefits of two-stage and one-stage hoi detection. *Adv. Neural Inform. Process. Syst.*, 34:17209–17220, 2021. [2](#)
- [49] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. [3](#)
- [50] Hongwen Zhang, Yating Tian, Yuxiang Zhang, Mengcheng Li, Liang An, Zhenan Sun, and Yebin Liu. Pymaf-x: Towards well-aligned full-body model regression from monocular images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2023. [6](#)
- [51] Hongwen Zhang, Yating Tian, Xinchu Zhou, Wanli Ouyang, Yebin Liu, Limin Wang, and Zhenan Sun. Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In *Int. Conf. Comput. Vis.*, 2021. [2](#), [6](#)
- [52] 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 *Eur. Conf. Comput. Vis.*, pages 180–200. Springer, 2022. [5](#)
- [53] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. [3](#)
- [54] Cheng Zou, Bohan Wang, Yue Hu, Junqi Liu, Qian Wu, Yu Zhao, Boxun Li, Chenguang Zhang, Chi Zhang, Yichen Wei, et al. End-to-end human object interaction detection with hoi transformer. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11825–11834, 2021. [2](#)

# AiOS: All-in-One-Stage Expressive Human Pose and Shape Estimation

## Supplementary Material

### 1. Overview

Due to the limited space on the main paper, we present more additional details in this supplementary material, covering the following aspects:

- Additional details on datasets, experimental parameters related to data augmentation, and implementation details in Sec. 2.
- Supplementary experiments investigating the impact of bounding boxes on algorithm performance in Sec. 3.
- Extra SOTA quantitative and qualitative comparison experiments on EgoBody-EgoSet [22], BEDLAM [1], and ARCTIC [6], and extra visualization comparison on single person, multiperson, and synthetic images. in Sec. 4.

### 2. Experiment Setup

#### 2.1. Datasets

This subsection primarily describes the characteristics of the datasets utilized in our main paper and provides details on how we use these datasets in the training and testing stages. **AGORA** [16] is a synthetic image dataset that encompasses lots of complex scenarios, including severe occlusion and truncations. It has 14K training images with 122K instances and 1K validation images with around 8K instances. Recently, AGORA has become an essential benchmark for tasks related to SMPL [12] and SMPL-X [17], primarily for its effectiveness in assessing algorithm performance in occlusion-heavy scenes. Our approach utilizes the complete dataset for training, validation, and testing purposes.

**BEDLAM** [1] is a synthetic video dataset offering a wide variety of data, including diverse body shapes, motions, skin tones, hairstyles, and clothing. The clothing is notably rendered using a professional physics simulator, enhancing realism, particularly in depicting character movement. Each image in the dataset features between 1 to 10 individuals. Originally comprising 286K images, we downsampled it by a factor of 5, yielding 57L images with 190K instances, which we then utilized for training.

**MSCOCO** [10] is a large-scale real-world dataset designed for object detection, segmentation, keypoint detection, and captioning. On the basis of the keypoint subset, we use the pseudo SMPL-X [17] annotations from NeuralAnnot [14] for training. It contains 56K multi-person images, featuring a total of 147K instances.

**EHF** [17] is an indoor dataset consisting of only 100 images along with corresponding SMPL-X labels. Due to the absence of a corresponding training set, current algorithms often utilize it to assess algorithm generalization. It only

contains test sets with 100 single-person images.

**UBody** [8] is a dataset containing diverse real-life scenarios, including movies, TV shows, talk shows, vlogs, sign language, online classes, and more. UBody contains an extensive collection of rich gestures and expressions that are not present in other real-world datasets. We down-sample the training set to 54K images with 66K instances. We utilize a downsampled test set, as used in SMPLer-X[2] and OSX [8], which has 2642 images with 2642 instances.

**ARCTIC** [6] is an indoor dataset which mainly focuses on the hand-object interaction. We downsample the original training set 1000 times for a train set of 50K images with 50K instances. We discard the egocentric view, which only contains hands, in our training. For the test set, following [2], we keep the original test set with 207K images to evaluate our method.

**Egobody** [22] is a large-scale dataset for egocentric views. The egocentric view datasets are collected using Microsoft HoloLens 2, including RGB, depth, eye gaze, head tracking, and hand tracking. The SMPL-X data is obtained by fitting the multi-Kinect rig data. We downsample 2 times to get 45K images with 45K instances in the egocentric-view split for training. The test set has 62K images.

We pre-train our model by randomly sampling data from AGORA, BEDLAM, and COCO. We choose these three datasets because their multi-person images satisfy our method of perceiving the positions of individuals in the images. Additionally, AGORA and BEDLAM have the advantage of accurate ground truth, while COCO provides the diversity of real-world scenes, preventing our method from overfitting to synthetic data.

SMPLer-X [2] indicates that scaling up the data can enhance algorithm performance. Hence, we incorporate Egobody, ARCTIC, and UBody data into our training. The sampling probabilities are 0.2, 0.2, 0.2, 0.2, 0.1, and 0.1, respectively, for AGORA, BEDLAM, COCO, Egobody, ARCTIC, and UBody. ARCTIC and Egobody provide more accurate real-world whole-body data compared to COCO, while UBody contributes abundant and diverse gestures and expressions. All the datasets are standardized into the HumanData [5] format. Some visual demonstrations of AiOS are provided in Fig. 1 and Fig. 2.

#### 2.2. Implementation details

The training is conducted on 16 V100 GPUs, with a total batch size of 32. We initialize AiOS’s backbone, encoder, and part of the decoder from a weight trained on COCO human pose estimation, provided by ED-Pose [21]. We use Adam optimizer with a step learning rate for training. We





Figure 1. **Illustration of AiOS in indoor datasets.** The first two columns are the qualitative results on ARCTIC [6], while the last two columns are the qualitative results on Egobody [22].

first train our model for 60 epochs with an initial learning rate  $1 \times 10^{-4}$  and drop at the 50th epoch on AGORA, BEDLAM, and COCO. We finetune it for 50 epochs with  $1 \times 10^{-5}$  initial learning rate and drop at the 25th epoch on all train datasets.

During the training stage, we adopt color jittering, random horizontal flipping, random image resizing, and random instance cropping. For the color jittering, we randomly apply a variation of  $\pm 0.2$  in the RGB channels. For the random horizontal flipping, we augment images and their corresponding annotations by horizontally flipping them with a probability of 0.5. For random image resizing, during the training process, we resize the images in proportion, ensuring that the shorter side is kept between 480 and 800 pixels and the longer side does not exceed 1333 pixels. During testing, we set the shorter side to 800 pixels, and the longer side is scaled proportionally, with the constraint that it does not exceed 1333 pixels. For random instance cropping, we first randomly enable the cropping operation with a probability of 0.5. When performing the cropping operation on a multi-person dataset, such as AGORA [16], BEDLAM [1], COCO [10], We randomly sample 1 to  $N$  instances from the image with probability 0.5, where  $N$  is the total number of people in the image. The image is then cropped according

to their collective bounding box. Alternatively, with a 0.5 probability, the cropping operation is applied to the collective bounding box of all instances. We maintain the original aspect ratio during the cropping process to avoid cropping unusual aspect ratios.

### 2.3. Loss Functions

**Body-location Decoder.** The losses for supervising Body-location Decoders are  $L_{box}$  and  $L_{cls}$ .  $L_{box}$  contains L1 loss and the GIOU loss [18] for the body location.  $L_{cls}$  is focal loss [9] for classify body tokens.

**Body-refinement Decoder.** The losses for supervising Body-refinement Decoders are  $L_{box}$ ,  $L_{j2d}$  and  $L_{smplx}$ .  $L_{box}$  contains L1 and loss GIOU loss for the body location, hands location, and face location.  $L_{j2d}$  is the L1 loss and OKS loss supervising the body joints location and  $L_{smplx}$  contains L1 loss with ground truth SMPL-X body parameters  $L_{param}$ , L1 loss  $L_{kp3d}$  for 3D body keypoints regressed by SMPL-X J-regressor, and L1 loss  $L_{kp2d}$  for projected 2D keypoints.

**Wholebody-refinement Decoder.** The losses for supervising Wholebody-refinement Decoders are  $L_{j2d}$  and  $L_{smplx}$ .  $L_{j2d}$  is the L1 loss and OKS loss supervising the whole-body joints location and  $L_{smplx}$  contains L1 loss with ground truth





Figure 2. **Illustration of AiOS in outdoor datasets.** The first three rows are qualitative results on UBody [8], and the last row is qualitative results on COCO [10]

SMPL-X parameters  $L_{param}$ , L1 loss  $L_{kp3d}$  for 3D whole-body keypoints regressed by SMPL-X J-regressor, and L1 loss  $L_{kp2d}$  for projected 2D whole-body keypoints.

**Loss Weights** We weighted-sum all the losses at all stages as the final loss. The loss weights of the same type of loss in different stages are the same, which are shown as follows:  $L_{cls}$ : 2.0,  $L_{smplx}$ : 1.0 (pose), 0.01 (shape, expression),  $L_{kp3d}$ : 1.0 (body), 0.5 (face, hand),  $L_{kp2d}$ : 1.0 (body), 0.5 (face, hand),  $L_{j2d}^1$ : 10,  $L_{oks}$ : 4.0 (body), 0.5 (face), 0.5 (hands),  $L_{giou}$ : 2.0,  $L_{box}^1$ : 5.0.

### 3. Sensitivity Analysis of Performance to Bounding Box Accuracy

In this section, we present that current methodologies exhibit a significant sensitivity to bounding boxes. Our experiments are carried out on the AGORA [16] validation set, utilizing official evaluation tools<sup>1</sup> for a thorough assessment. We report several key metrics, including Normalized Mean Vertex Error (NMVE), Normalized Mean Joint Error (NMJE), Mean Vertex Error (MVE), and Mean Per-Joint Position Error (MPJPE), which focus on the reconstruction accuracy of body, hands, and face. Additionally, we include F-Score, precision, and recall to evaluate the accuracy of detection.

<sup>1</sup>[https://github.com/pixelite1201/agora\\_evaluation](https://github.com/pixelite1201/agora_evaluation)




Figure 3. **Illustration of Inter-human, Intra-human, and Cross Attention.** Inter-human attention is conducted between the body tokens and body, hand, and face location tokens of different persons'. Intra-human attention is conducted with the location and joint tokens of the same person. Cross-attention is conducted between image features and all the tokens.

Firstly, we utilize the ground truth (GT) bounding box to crop the image and evaluate the performance of OSX [8] and SMPLer-X [2], denoted by **GT Box** in Table 1. Following this, we employ DAB-DETR[3, 11], an off-the-shelf detection model, to identify human bounding boxes, replacing the GT boxes for image cropping. We present detection results under three bounding box score thresholds: 0.1, 0.2, and 0.3, labeled as detected boxes with scores 0.1, 0.2, and 0.3, respectively. A higher score indicates greater confidence in the detection results, but it may lead to missing instances with severe occlusion. This is reflected in the metrics as high accuracy but low recall. Conversely, a lower score threshold retains results with lower confidence, capturing

Methods	Box Type	F Score $\uparrow$	Precision $\uparrow$	Recall $\uparrow$	NMVE $\downarrow$ (mm)		NMJE $\downarrow$ (mm)		MVE $\downarrow$ (mm)					MPIPE $\downarrow$ (mm)				
					All	Body	All	Body	All	Body	Face	LHand	RHand	All	Body	Face	LHand	RHand
OSX [8]	GT box	<b>1.0</b>	<b>1.0</b>	<b>1.0</b>	<b>120.4</b>	<b>75.2</b>	<b>116.1</b>	<b>72.1</b>	120.4	75.2	39.4	47.5	48.8	116.1	72.1	40.5	45.1	46.4
	Detected box w score 0.1	0.58	0.42	0.92	237.8	155.9	228.3	147.4	137.9	90.4	43.8	48.3	50.4	132.4	85.5	46.4	46.1	48.1
	Detected box w score 0.2	0.8	0.74	0.87	165.9	110.6	159.2	104.7	132.7	88.5	40.7	<b>44.9</b>	<b>46.9</b>	127.4	83.8	43.1	42.9	44.8
	Detected box w score 0.3	0.86	0.89	0.83	148.1	100.3	142.1	95.1	127.4	86.3	<b>37.8</b>	<b>41.8</b>	<b>43.6</b>	122.2	81.8	39.8	<b>39.8</b>	<b>41.6</b>
	GT box $\times$	<b>0.95</b>	<b>0.93</b>	<b>0.97</b>	<b>123.6</b>	<b>77.3</b>	<b>119.3</b>	<b>74.2</b>	<b>117.4</b>	<b>73.4</b>	38.2	46.3	47.7	<b>113.3</b>	<b>70.5</b>	<b>39.0</b>	43.9	45.3
	AiOS box	<b>0.95</b>	<b>0.93</b>	<b>0.97</b>	124.3	78.7	120.0	75.7	<b>118.1</b>	<b>74.8</b>	<b>37.4</b>	45.5	47.0	<b>114.0</b>	<b>71.9</b>	<b>38.3</b>	<b>43.2</b>	<b>44.7</b>
SMPLer-X [2]	GT box	<b>1.0</b>	<b>1.0</b>	<b>1.0</b>	<b>99.5</b>	<b>60.2</b>	<b>95.5</b>	<b>57.5</b>	99.5	<b>60.2</b>	32.9	41.9	43.0	95.5	<b>57.5</b>	34.1	39.5	40.5
	Detected box w score 0.1	0.58	0.43	0.92	203.4	131.6	195.3	124.7	118.0	76.3	37.1	43.6	44.4	113.3	72.3	39.5	41.3	42.2
	Detected box w score 0.2	0.81	0.75	0.88	140.7	92.6	135.1	87.8	114.0	75.0	34.7	40.7	<b>41.7</b>	109.4	71.1	36.9	38.5	39.5
	Detected box w score 0.3	0.86	0.9	0.83	126.4	84.4	121.3	80.1	108.7	72.6	31.8	<b>37.8</b>	<b>38.7</b>	104.3	68.9	33.7	<b>35.7</b>	<b>36.7</b>
	GT box $\times$	<b>0.95</b>	<b>0.93</b>	<b>0.97</b>	<b>101.5</b>	<b>61.4</b>	<b>97.6</b>	<b>58.8</b>	<b>96.4</b>	<b>58.3</b>	<b>31.7</b>	40.8	41.9	<b>92.7</b>	<b>55.9</b>	<b>32.6</b>	38.3	<b>39.4</b>
	AiOS box	<b>0.95</b>	<b>0.93</b>	<b>0.97</b>	103.3	63.5	99.6	61.1	<b>98.1</b>	60.3	<b>31.3</b>	<b>40.4</b>	41.8	<b>94.6</b>	58.0	<b>32.3</b>	<b>38.0</b>	<b>39.4</b>
AiOS	-	0.95	0.93	0.97	106.4	64.2	103.4	62.1	101.1	61.0	30.7	43.9	45.7	98.2	59.0	32.8	41.5	43.4

Table 1. **AGORA validation set**. OSX [8] and SMPLer-X [2] are finetuned on the AGORA training set. However, AiOS is not intentionally fine-tuned exclusively on AGORA. **GT Box** means that this method uses the ground truth bounding box to crop the image. **GT box $\times$**  means that this method uses the ground truth bounding box to crop the image but filters the instances that AiOS fails to detect. **AiOS box** means that this method uses the bounding box provided by AiOS. The **best results** are colored with **red**, and the **second-best results** are colored with **blue** for OSX and SMPLer-X, respectively.

Methods	Box Type	F Score $\uparrow$	Precision $\uparrow$	Recall $\uparrow$	NMVE $\downarrow$ (mm)		NMJE $\downarrow$ (mm)		MVE $\downarrow$ (mm)					MPIPE $\downarrow$ (mm)				
					All	Body	All	Body	All	Body	Face	LHand	RHand	All	Body	Face	LHand	RHand
OSX [8]	GT box $\times$	0.95	0.93	0.97	<b>123.6</b>	<b>77.3</b>	<b>119.3</b>	<b>74.2</b>	<b>117.4</b>	<b>73.4</b>	<b>38.2</b>	<b>46.3</b>	<b>47.7</b>	<b>113.3</b>	<b>70.5</b>	<b>39.0</b>	<b>43.9</b>	<b>45.3</b>
	GT box with noise $\times$	0.95	0.93	0.97	126.1	78.8	121.8	<b>75.6</b>	119.8	74.9	38.7	46.9	49.1	115.7	72.1	39.6	44.5	46.6
	AiOS box	0.95	0.93	0.97	<b>124.3</b>	<b>78.7</b>	<b>120.0</b>	75.7	<b>118.1</b>	<b>74.8</b>	<b>37.4</b>	<b>45.5</b>	<b>47.0</b>	<b>114.0</b>	<b>71.9</b>	<b>38.3</b>	<b>43.2</b>	<b>44.7</b>
SMPLer-X [2]	GT box $\times$	0.95	0.93	0.97	<b>101.5</b>	<b>61.4</b>	<b>97.6</b>	<b>58.8</b>	<b>96.4</b>	<b>58.3</b>	<b>31.7</b>	<b>40.8</b>	<b>41.9</b>	<b>92.7</b>	<b>55.9</b>	<b>32.6</b>	<b>38.3</b>	<b>39.4</b>
	GT box with noise $\times$	0.95	0.93	0.97	105.6	64.0	101.6	61.5	100.3	60.8	32.5	42.4	43.6	96.5	58.4	33.5	39.9	<b>41.0</b>
	AiOS box	0.95	0.93	0.97	<b>103.3</b>	63.5	<b>99.6</b>	<b>61.1</b>	<b>98.1</b>	<b>60.3</b>	<b>31.3</b>	<b>40.4</b>	<b>41.8</b>	<b>94.6</b>	<b>58.0</b>	<b>32.3</b>	<b>38.0</b>	<b>39.4</b>

Table 2. **AGORA validation set with noise bounding box**. OSX [8] and SMPLer-X [2] are finetuned on the AGORA training set. **GT box** means that this method uses the ground truth bounding box to crop the image. **GT box with noise** means translating the ground truth bounding box by **10%** of the image size in the horizontal direction, causing the human to deviate from the image center. This is a very small noise that ensures the person is not removed from the image plane, avoiding truncation.  $\times$  means that this method filters the instances that AiOS fails to detect. The **best results** are colored with **red**, and the **second-best results** are colored with **blue** for OSX and SMPLer-X, respectively.

more instances with severe occlusion but resulting in many redundant detections. This is reflected in the metrics as high recall but lower accuracy.

From the data presented in Table 1, it’s evident that regardless of using a low or high threshold for filtering detection results, both OSX [8] and SMPLer-X [2] show a marked decrease in performance in terms of NMVE and NMJE when compared to results achieved using GT bounding boxes. This performance gap is largely due to the way NMVE and NMJE are normalized using the F1 score. The F1 score, being the harmonic mean of recall and precision, penalizes methods for both missed detections and false positives, which explains the observed discrepancy in performance.

Notably, when we filter the detection results with a score of 0.1, denoted by **Detected box w score 0.1**, the recall is 0.92, indicating that we detect the majority of instances, although there were many redundant detections. In this case, if we concentrate on reconstruction accuracy, represented by MVE and MPIPE, we observe that the results using detected bounding boxes OSX and SMPLer-X are significantly worse than using GT bounding box for cropping images.

On the lower part of Table 1, we present the AiOS results. To allow more hard cases to be detected, we use a threshold of 0.1 to filter the result. Similar to the main paper, providing AiOS’s bounding boxes to OSX and SMPLer-X denoted by **AiOS box**, leads to a significant performance increase com-

pared to using the detected boxes (**Detected box w score 0.3**), even though AiOS’s box includes more challenging cases. Specifically, there is a significant improvement in NMVE with a 16% reduction from 148.1 to 124.3 for OSX and an 18% decrease from 126.4 to 103.3 for SMPLer-X. Additionally, there is an improvement in NMJE, with a 15% reduction from 142.1 to 120.0 for OSX and a 17% decrease from 121.3 to 99.6 for SMPLer-X. These results substantiate that one-stage methods demonstrate superior performance in real-world scenarios compared to existing two-stage methods.

For a fair comparison with OSX and SMPLer-X, which use the GT bounding boxes to crop images, we filter out the same instances that AiOS failed to detect for both OSX and SMPLer-X. Interestingly, we found that the results obtained using AiOS bounding boxes (denoted with **AiOS box** in Table 1) are comparable to those obtained using GT bounding boxes (denoted with **GT box $\times$** ). Notably, under this bounding box setting, AiOS still outperforms the current state-of-the-art OSX, even though it utilizes the GT bounding box and is on par with the foundational model SMPLer-X L20. To further demonstrate the superiority of AiOS, we follow the procedure in RoboSMPLX [15] to translate the image by 0.1 of the image size horizontally, introducing a small noise to the GT bounding boxes (denoted by **GT box with noise $\times$**  in Table 2). It’s worth noting that this noise





Figure 4. **Additional visual comparisons with existing methods.** **Upper part:** The first column is the input images, and they are downloaded from the internet. The second column is the visualization results of ROMP [19], the third column shows the visualization results of BEV [20], and the last column illustrates our visualization results; **Middle part:** Comparison of current SOTA methods [2, 8] with our AiOS model on AGORA [16]. **Lower part:** Comparison of current SOTA methods [2, 8, 13] with our AiOS model on EHF test [4].

Methods	F Score $\uparrow$	Precision $\uparrow$	Recall $\uparrow$	NMVE $\downarrow$ (mm)		NMJE $\downarrow$ (mm)		MVE $\downarrow$ (mm)					MPIPE $\downarrow$ (mm)				
				All	Body	All	Body	All	Body	Face	LHand	RHand	All	Body	Face	LHand	RHand
PIXIE [7]	0.94	0.99	<b>0.90</b>	158.7	107.2	153.7	103.5	149.2	100.8	51.4	44.8	48.9	144.5	97.3	55.4	41.3	44.8
BEDLAM-CLIFF	0.94	0.99	<b>0.90</b>	100.6	65.2	98.0	64.3	94.6	61.3	29.8	34.7	35.5	92.1	60.4	30.4	32.2	32.6
BEDLAM-CLIFF++ $\dagger$	0.94	0.99	<b>0.90</b>	93.2	61.2	90.9	60.4	87.6	57.5	27.3	30.3	32.6	85.4	56.8	28.0	28.0	29.9
AiOS $\dagger$	<b>0.95</b>	<b>1</b>	<b>0.90</b>	<b>87.6</b>	<b>57.7</b>	<b>85.8</b>	<b>57.7</b>	<b>83.2</b>	<b>54.8</b>	<b>26.4</b>	<b>28.1</b>	<b>30.8</b>	<b>81.5</b>	<b>54.8</b>	<b>26.2</b>	<b>25.9</b>	<b>28.1</b>

Table 3. **BEDLAM test set.** The best results are in **bold**.  $\dagger$  denotes methods that include the AGROA training set.

only shifts the person away from the image center and does not remove the person from the image plane, avoiding trun-

cation and occlusion. Firstly, compare it with the results cropping with the GT bounding box denoted by **GT box $\times$** ,

Method	PA-PVE↓ (mm)			PVE↓ (mm)		
	All	Hands	Face	All	Hands	Face
H4W [13]	63.4	18.1	4.0	136.8	54.8	59.2
OSX [8]	56.9	<b>17.5</b>	3.9	102.6	56.5	44.6
OSX [8]†	33.0	18.8	3.3	58.4	39.4	30.4
SMPLer-X [2]	31.9	18.9	2.5	52.2	<u>39.3</u>	27.0
Native AiOS	<u>31.8</u>	19.7	<u>2.3</u>	<u>51.6</u>	40.6	<u>26.5</u>
AiOS	<b>30.2</b>	19.2	<b>2.1</b>	<b>47.1</b>	<b>38.3</b>	<b>26.1</b>

Table 4. **ARCTIC**. † denotes the method finetuned on the ARCTIC training set.

Method	PA-PVE↓ (mm)			PVE↓ (mm)		
	All	Hands	Face	All	Hands	Face
H4W [13]	58.8	9.7	3.7	121.9	50.0	42.5
OSX [8]	54.6	11.6	3.7	115.7	50.6	41.1
OSX [8]†	45.3	10.0	<u>3.0</u>	82.3	46.8	35.2
SMPLer-X [2]	38.9	9.9	<u>3.0</u>	66.6	42.7	31.8
SMPLer-X [2]†	<b>37.8</b>	9.9	<b>2.9</b>	<u>63.6</u>	46.3	<u>32.3</u>
Native AiOS	40.8	<u>9.1</u>	<u>3.0</u>	64.6	<u>42.3</u>	<b>26.3</b>
AiOS	<u>38.0</u>	<b>9.0</b>	<b>2.9</b>	<b>61.6</b>	<b>40.0</b>	<u>26.7</u>

Table 5. **EgoBody-EgoSet**. † denotes the methods that are finetuned on the EgoBody-EgoSet training set.

we can observe a drop in performance. Specifically, for NMVE, OSX increased from 123.6 to 126.1, and SMPLer-X increased from 101.5 to 105.6. Regarding NMJE, OSX increased from 119.3 to 121.8, and SMPLer-X increased from 97.6 to 101.6. This observation indicates that current two-stage methods are highly sensitive to the accuracy of the bounding box, as even a slight noise introduced, causing the person to be off-center, results in a performance drop, even with GT bounding boxes.

When comparing the results obtained by cropping images with bounding boxes provided by AiOS to those obtained by cropping with GT bounding boxes and GT bounding boxes with added noise, we observe that the results of cropping with AiOS-provided bounding boxes are slightly inferior to those obtained with GT bounding boxes in body-related metrics but better than GT bounding boxes with added noise. However, for the face and hands, the results of cropping with AiOS-provided bounding boxes can even be better than those obtained with GT bounding boxes. We attribute this improvement to AiOS’s attention to not only the body but also the hands and face.

#### 4. Extra SOTA comparison experiments

In this section, we show the extra single datasets on BED-LAM, ARCTIC, and EgoBody-EgoSet in Table 3, 4, 5. Our proposed AiOS achieved SOTA performance across all these datasets. Also, we provide extra visualization comparison in Fig. 4.



## References

- [1] Michael J Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. Bedlam: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8726–8737, 2023. 1, 2
- [2] Zhongang Cai, Wanqi Yin, Ailing Zeng, Chen Wei, Qingping Sun, Yanjun Wang, Hui En Pang, Haiyi Mei, Mingyuan Zhang, Lei Zhang, et al. Smpler-x: Scaling up expressive human pose and shape estimation. *arXiv preprint arXiv:2309.17448*, 2023. 1, 3, 4, 5, 6
- [3] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 3
- [4] Vasileios Choutas, Georgios Pavlakos, Timo Bolkart, Dimitrios Tzionas, and Michael J Black. Monocular expressive body regression through body-driven attention. In *Eur. Conf. Comput. Vis.*, pages 20–40. Springer, 2020. 5
- [5] MMHuman3D Contributors. Openmmlab 3d human parametric model toolbox and benchmark. <https://github.com/open-mmlab/mmhuman3d>, 2021. 1
- [6] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J Black, and Otmar Hilliges. Arctic: A dataset for dexterous bimanual hand-object manipulation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12943–12954, 2023. 1, 2
- [7] Yao Feng, Vasileios Choutas, Timo Bolkart, Dimitrios Tzionas, and Michael Black. Collaborative regression of expressive bodies using moderation. In *International Conference on 3D Vision (3DV)*, pages 792–804, Dec. 2021. 5
- [8] Jing Lin, Ailing Zeng, Haoqian Wang, Lei Zhang, and Yu Li. One-stage 3d whole-body mesh recovery with component aware transformer. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 21159–21168, 2023. 1, 3, 4, 5, 6
- [9] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 2
- [10] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Eur. Conf. Comput. Vis.*, pages 740–755. Springer, 2014. 1, 2, 3
- [11] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. DAB-DETR: Dynamic anchor boxes are better queries for DETR. In *Int. Conf. Learn. Represent.*, 2022. 3
- [12] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM Trans. Graph.*, 34(6):1–16, 2015. 1
- [13] Gyeongsik Moon, Hongsuk Choi, and Kyoung Mu Lee. Accurate 3d hand pose estimation for whole-body 3d human mesh estimation. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, 2022. 5, 6
- [14] Gyeongsik Moon, Hongsuk Choi, and Kyoung Mu Lee. Neuralannot: Neural annotator for 3d human mesh training sets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2299–2307, 2022. 1
- [15] Hui En Pang, Zhongang Cai, Lei Yang, Qingyi Tao, Zhonghua Wu, Tianwei Zhang, and Ziwei Liu. Towards robust and expressive whole-body human pose and shape estimation. In *Adv. Neural Inform. Process. Syst.*, 2023. 4
- [16] Priyanka Patel, Chun-Hao P Huang, Joachim Tesch, David T Hoffmann, Shashank Tripathi, and Michael J Black. AGORA: Avatars in geography optimized for regression analysis. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 13468–13478, 2021. 1, 2, 3, 5
- [17] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. 1
- [18] Hamid Rezaatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2019. 2
- [19] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Michael J Black, and Tao Mei. Monocular, one-stage, regression of multiple 3d people. In *Int. Conf. Comput. Vis.*, pages 11179–11188, 2021. 5
- [20] Yu Sun, Wu Liu, Qian Bao, Yili Fu, Tao Mei, and Michael J Black. Putting people in their place: Monocular regression of 3d people in depth. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 13243–13252, 2022. 5
- [21] Jie Yang, Ailing Zeng, Shilong Liu, Feng Li, Ruimao Zhang, and Lei Zhang. Explicit box detection unifies end-to-end multi-person pose estimation. *arXiv preprint arXiv:2302.01593*, 2023. 1
- [22] 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 *Eur. Conf. Comput. Vis.*, pages 180–200. Springer, 2022. 1, 2