This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# **PromptHMR: Promptable Human Mesh Recovery**

Yufu Wang<sup>1,4</sup> Yu Sun<sup>1</sup> Priyanka Patel<sup>1</sup> Kostas Daniilidis<sup>4,5</sup> Michael J. Black<sup>1,2</sup> Muhammed Kocabas<sup>1,2,3</sup>

<sup>1</sup> Meshcapade <sup>2</sup> MPI for Intelligent Systems <sup>3</sup> ETH Zürich <sup>4</sup> University of Pennsylvania <sup>5</sup> Archimedes



Figure 1. **PromptHMR** is a promptable human pose and shape (HPS) estimation method that processes images with spatial or semantic prompts. It takes "side information" readily available from vision-language models or user input to improve the accuracy and robustness of 3D HPS. PromptHMR recovers human pose and shape from spatial prompts such as (a) face bounding boxes, (b) partial or complete person detection boxes, or (c) segmentation masks. It refines its predictions using semantic prompts such as (c) person-person interaction labels for close contact scenarios, or (d) natural language descriptions of body shape to improve body shape predictions. Both image and video versions of PromptHMR achieve state-of-the-art accuracy. https://yufu-wang.github.io/phmr-page

#### Abstract

Human pose and shape (HPS) estimation presents challenges in diverse scenarios such as crowded scenes, personperson interactions, and single-view reconstruction. Existing approaches lack mechanisms to incorporate auxiliary "side information" that could enhance reconstruction accuracy in such challenging scenarios. Furthermore, the most accurate methods rely on cropped person detections and cannot exploit scene context while methods that process the whole image often fail to detect people and are less accurate than methods that use crops. While recent language-based methods explore HPS reasoning through large language or vision-language models, their metric accuracy is well below the state of the art. In contrast, we present PromptHMR, a transformer-based promptable method that reformulates HPS estimation through spatial and semantic prompts. Our method processes full images to maintain scene context and accepts multiple input modalities: spatial prompts like bounding boxes and masks, and semantic prompts like language descriptions or interaction labels. PromptHMR demonstrates robust performance across challenging scenarios: estimating people from bounding boxes as small as faces in crowded scenes, improving body shape estimation through language descriptions, modeling person-person interactions, and producing temporally coherent motions in videos. Experiments on benchmarks show that PromptHMR achieves state-of-the-art performance while offering flexible prompt-based control over the HPS estimation process.

## **1. Introduction**

The estimation of 3D human pose and shape (HPS) is classically viewed as regressing the parameters of shape and pose from pixels. In particular, most methods take a tightly cropped image of a person and output the pose and shape in camera coordinates. While the accuracy of such methods has increased rapidly, they do not address the whole problem. In particular, an HPS method should be able to take an image or video containing complex human-human and human-scene interactions, return the parameters of every person in the scene, and place these people in a consistent global coordinate frame.

Our key observation is that the classical "pixels to parameters" formulation of the problem is too narrow. Today, we have large vision-language foundation models (VLMs) that understand a great deal about images and what people are doing in them. What these models lack, however, is an understanding of 3D human pose and shape. Recent work [10, 15] has tried to bring together VLMs and 3D HPS but with 3D accuracy well below the best classical methods.

Consequently, we need to think about the problem in a different way and ask whether we can exploit readily available side information (e.g. provided by a VLM) to improve 3D HPS regression robustness, usefulness, and accuracy. To that end, we develop a novel "promptable" HPS architecture called **PromptHMR**. Consider the sample images shown in Fig. 1. In crowded scenes, existing person detection methods struggle, while face detection methods remain reliable. When people closely interact, their body parts overlap and occlude each other, introducing ambiguity in pose estimation. Moreover, 3D body shape estimation from monocular views is challenging due to perspective ambiguity. In all these cases, we can extract cues, or prompts, that provide "side information" that can help an HPS method better analyze the scene. PromptHMR formalizes this intuition by combining image evidence with different types of spatial and semantic information that can come from either humans or AI systems such as VLMs.

Specifically, our approach combines three key components: (1) a vision transformer that extracts features from high-resolution full images to preserve scene context, (2) a multi-modal prompt encoder that processes spatial and semantic inputs, and (3) a transformer decoder that attends to both prompt and image tokens to generate SMPL-X [45] body parameters. This design addresses the limitations of cropped-image HPS methods by processing full images using side information in the form of prompts. It addresses the challenges that full-image HPS methods have in detecting all people in a scene by accepting readily available bounding boxes. Last, our method incorporates auxiliary semantic information through text descriptions or interaction labels.

By combining spatial and semantic prompting, our method offers a powerful and versatile approach to 3D HPS

estimation from the whole image. At test time, we show that this promptable structure (1) can take various bounding boxes or segmentation masks to recover full body HPS in a robust way, (2) improve its body shape predictions by using textual descriptions as input, (3) is capable of modeling person-person close interaction directly in the regression process, and (4) uses full image context to reconstruct people coherently in the camera space and the world space. Our model can handle video by incorporating temporal transformer layers at the SMPL-X decoding phase, yielding temporally stable and smooth motions. Last, following TRAM [65], we combine the temporal version of our model with metric SLAM to estimate human motion in world coordinates.

We make several key design choices that make Prompt-HMR successful. To achieve robustness to different spatial inputs, we train our model by simulating noisy fullbody and face-region bounding boxes. For improved body shape estimation, we leverage SHAPY [8] to generate automatic body shape descriptions for training samples and process them with a pretrained text encoder [48]. To enhance person-person interaction reconstruction, we use segmentation masks as more precise spatial prompts and develop person-person attention layers that operate between prompted people, producing coherent reconstructions of close interactions. Through random masking of different input types during training, our model learns to work with any combination of prompts at test time.

Quantitative experiments on the EMDB [24], 3DPW [63], RICH [20], Hi4D [68], CHI3D [16] and HBW [8] benchmark datasets demonstrate that our method outperforms state-of-the-art (SOTA) approaches and strong baselines. We also provide many qualitative examples of in-the-wild images and videos that illustrate the robustness and generalization of PromptHMR.

By moving away from the pure pixels-to-parameters approach, PromptHMR not only achieves a new SOTA, it shows a new way of improving both accuracy and robustness by leveraging side information that is easily available. One can think of this as a collaboration between VLMs, which know a lot about people in images but not in 3D, and a metric regressor that knows a lot about 3D humans but not about the semantics of what they do. We show that this combination has significant upside potential to increase both generality and accuracy. Our code and model are available for research purposes.

### 2. Related Work

**Human pose and shape estimation from images**. Existing methods for human pose and shape (HPS) estimation can be broadly categorized into two main approaches. The first [6, 17, 22, 27, 28, 30–32, 39, 64, 70] uses a tightly cropped image of an individual as input, and estimates

pose and shape in camera coordinates. While effective for isolated individuals, this approach discards scene context that is essential to resolve human pose in cases of occlusion, severe overlap and close interaction in multi-person scenes [16, 68].

The second category [2, 21, 55, 57–59] build upon object detection frameworks [5, 12] to jointly detect humans and estimate their pose and shape parameters. Having access to the entire image, they can better perceive occluded individuals and infer depth relationships, but they often suffer from detection failures due to the difficulty in simultaneously performing detection and reconstruction. Our "promptable" architecture leverages detection box prompts to resolve such conflicts while having access to the entire scene context.

**Human pose and shape estimation from video.** Methods for human motion estimation from video can also be divided into two main categories. The first [7, 23, 26, 36, 56] focuses on estimating smooth human motion in camera space. These methods build upon single-person HPS estimation approaches [22, 30] by adding temporal layers during the SMPL decoding phase to introduce temporal coherence.

More recent methods estimate human motion in world coordinates from videos captured with dynamic cameras. These methods follow a two-stage approach, first estimating camera motion using SLAM techniques [18, 19, 40, 41, 60, 61], and then leveraging human motion priors to optimize the human world motion [29, 67, 69]. Others [51, 52] learn temporal models to directly regress human world motion from image and camera features. Still others [65, 71] use monocular metric depth estimation to solve for the scale of camera motion and transform human motion from camera space to world coordinates.

In our approach, we extend PromptHMR to video by taking the SMPL-X output tokens and utilizing a temporal transformer module to estimate temporally stable and smooth human motion and translation in camera space. We follow TRAM [65] to transform human motion to world co-ordinates due to its simplicity and effectiveness.

Semantic reasoning about 3D humans in images. Recent methods explore combining different types of semantic information, such as language descriptions and knowledge of person-person interactions, to improve reasoning about 3D humans from images and videos. For example, ChatPose [15] follows the common approach of visual language models (VLMs) [34] by fine-tuning a large language model (LLM) with a combination of images and tokens to estimate SMPL parameters. In a similar direction, PoseEmbroider [10] is a multi-modal framework that aligns image, 3D pose, and text representations in a shared latent space. While ChatPose focuses on combining high-level scene reasoning with 3D HPS, PoseEmbroider exploits detailed language descriptions of human pose. While promising, neither method achieves SOTA accuracy on the HPS task. Note that many other methods relate language to human pose or motion, without considering images [1, 9, 35, 46, 62], but these are outside our scope.

Additionally, several methods [8, 47, 53] focus on modeling the relationship between SMPL body shape and natural language descriptions. These methods show that language descriptions and images can provide complementary information to solve this task. Other approaches, such as BUDDI [38] and ProsePose [54], address the challenge of estimating person-person interactions. BUDDI is an optimization-based approach that leverages diffusion model as a prior over interacting people, while ProsePose queries a VLM to estimate contact points on the human body surface and uses these contact points to guide an optimization process for improving human interaction.

Overall, methods like ChatPose [15] and PoseEmbroider [10] are promising steps toward jointly learning the relationship between vision, language, and 3D humans, but their understanding of 3D humans remains limited, as indicated by their relatively low 3D pose accuracy. Meanwhile, SHAPY [8], BodyShapeGPT [47], and BodyTalk [53] focus solely on exploring the relationship between SMPL body shape and natural language. BUDDI and ProsePose are post-processing approaches for interaction that do not directly reason using image information.

Our approach addresses the limitations of these methods by training a single model capable of flexible prompting that achieves state-of-the-art (SOTA) performance, not only on standard HPS benchmarks but also on benchmarks tailored to body shape and person-person interaction.

## 3. Method

Given an image I containing N people and a set of prompts, our main goal is to recover the pose, shape, and locations of the people in the camera space to form a coherent humancentric 3D scene. Figure 2 shows an overview.

#### 3.1. Promptable mesh regression

We adopt SMPL-X [45] to represent each person *i* in the 3D space, including the orientation  $\phi_i \in \mathbb{R}^3$ , local body pose  $\theta_i \in \mathbb{R}^{22\times3}$ , shape  $\beta_i \in \mathbb{R}^{10}$ , and translation  $\tau_i \in \mathbb{R}^3$  in the camera space. We do not include face and hand parameters in this work. Each human  $H_i$  is mapped to a 3D mesh with the differentiable SMPL-X layer.

$$H_i = \{\phi_i, \theta_i, \beta_i, \tau_i\}.$$
 (1)

Each person can be prompted with spatial and semantic prompts. Spatial prompts include a bounding box  $b_i \in \mathbb{R}^{2\times 2}$  (the two corners) and a segmentation mask  $m_i \in \mathbb{R}^{h\times w}$ . Semantic prompts consist of text and two-person interaction labels. The text prompt is the CLIP embedding  $t_i$  of a sentence describing the body shape. The interaction prompt is a binary variable  $k_i$  indicating whether two



Figure 2. **Method overview.** PromptHMR estimates SMPL-X parameters for each person in an image based on various types of prompts, such as boxes, language descriptions, and person-person interaction cues. Given an image and prompts, we utilize a vision transformer to generate image embeddings and mask and prompt encoders to map different types of prompts to tokens. Optionally, camera intrinsics can be embedded along with the image embeddings. The image embeddings and prompt tokens are then fed to the SMPL-X decoder. The SMPL-X decoder is a transformer-based module that attends to both the image and prompt tokens to estimate SMPL-X parameters. Note that the language and interaction prompts are optional, but providing them enhances the accuracy of the estimated SMPL-X parameters.

people are in close contact. While semantic prompts are optional, each human needs at least one spatial prompt to be reconstructed. Overall, the input prompts are represented as  $P_i$ :

$$P_i \subseteq \{b_i, m_i, t_i, k_i\}$$
  
$$b_i \in P_i \text{ or } m_i \in P_i$$
(2)

Promptable human mesh recovery (PromptHMR) is defined as a learnable function that maps an image and a set of prompts to a set of 3D humans

$$f: (I, \{P_i\}_{i=1}^N) \to \{H_i\}_{i=1}^N.$$
(3)

This task definition integrates all available contexts to locate and reconstruct prompted humans in the image.

#### 3.2. Model

**Image encoder.** The image is first encoded as tokens by a vision transformer (ViT) encoder from DINOv2 [11, 42]:

$$F = \text{Encoder}(I), \tag{4}$$

To ensure sufficient resolution for modeling humans at both near and far distances, we use  $896 \times 896$  images. The encoder is run once per frame regardless of the number of people prompted. When camera intrinsics are provided, we add positional encoding of the camera rays to the image tokens to make them camera-aware [2, 14].

**Mask encoder.** When available, masks are first processed by an encoder consisting of strided convolutional layers that downsample the masks. The output mask features are added to the image tokens. If no mask is provided, a learned "no mask" token is added instead.

$$F_i = \text{Encoder}_{\mathsf{m}}(m_i) + F.$$
(5)

**Prompt encoder.** The prompt encoder consists of a set of transformations that map different types of prompts to token vectors of the same dimension. When a prompt is not available, it is replaced with a learned null token.

For bounding boxes, we encode  $b_i$  using positional encoding summed with learned embeddings to form the box prompt tokens  $T_{bi} = PE(b_i)$ , with  $T_{bi} \in \mathbb{R}^{2 \times d}$ . We design different box transformations during training to allow the model to use different boxes as a human identifier. In the training phase, each instance is prompted with either a whole-body bounding box, a face bounding box, or a truncated box covering part of the body. Gaussian noise is added to both corners. At inference time, the model accepts boxes without needing to know the box types.

Language is a natural way to supply semantic information, and in this paper, we use language to supplement spatial prompts with information on body shape. A sentence such as "a muscular and tall male" is encoded with the CLIP text encoder  $T_{ti} = \text{CLIP}(t_i)$ , with  $T_{ti} \in \mathbb{R}^d$ . To generate paired (image, text) data, we run SHAPY's [8] shape-toattribute method on the ground truth shape parameters to obtain shape attribute scores and randomly pick a subset of top attributes to form a sentence.

The interaction prompt  $k_i$  passes through the prompt encoder without modification and directly switches on-off the cross-person attention that is described in Sec. 3.3.

**SMPL-X decoder.** The SMPL-X decoder appends two query tokens  $T_{\text{smpl}}, T_{\text{depth}}$  with the prompt tokens  $T_{bi}, T_{ti}$  to form the person-specific prompt  $T_i \in \mathbb{R}^{5 \times d}$ .

Finally, we use a standard transformer decoder and two MLP heads to produce the final output

$$T'_{smpl}, T'_{depth} = \text{Decoder}(F_i, T_i)$$
  

$$\phi_i, \theta_i, \beta_i = \text{Head}_{smpl}(T'_{smpl})$$
  

$$\tau_i = \text{Head}_{depth}(T'_{depth}).$$
(6)

The transformer consists of three attention blocks. Each block applies self-attention on the tokens, cross-person attention (described in Sec. 3.3), and then two-way cross-attention between the tokens and the image embeddings [25]. The self-attention and cross-attention with the

image are applied to each prompted person independently. We use separate tokens  $T_{\rm smpl}$  and  $T_{\rm depth}$  to make the location representation invariant to the 3D human pose and shape representation.

Regressing the location of the human in the camera space is much more challenging than most prior work that models humans in a cropped image space. Therefore, we do not regress  $\tau$  directly. We regress focal length normalized 2D translation  $p_{xy} \in \mathbb{R}^2$  and inverse depth  $p_z \in \mathbb{R}$ , and then transform them to  $\tau$  as follows

$$t_{xy} = \frac{p_{xy}}{p_z} \quad t_z = \frac{1}{p_z} \times \frac{f}{f_c} \quad \tau = [t_{xy}, tz],$$
 (7)

where f is the ground truth or estimated focal length of the image, and  $f_c$  is the canonical focal length. Predicting the normalized inverse depth follows the recent monocular depth literature [49] and is also intuitive since the inverse depth is linearly related to the size of the human in the image. Predicting  $p_{xy}$  is equivalent to predicting the 2D location of the human in a normalized image plane.

#### 3.3. Two-person interaction

We introduce promptable layers in the decoder to model two-person interaction. We describe the case where there are two people in the image, but the implementation can extend to model an interacting pair in a larger group.

The promptability is modeled as a flow control with a residual connection (Fig. 3). Specifically, if two humans are interacting (as indicated by  $k_i$ ), their query tokens pass through an additional self-attention layer; otherwise, non-interacting humans skip this.

Applying attention to every person often creates unnecessary dependency in crowded scenes, and there is limited training data for large-group scenarios. However, there is high-quality data featuring two-person social interactions. By making the interaction layers promptable, we mitigate data diversity issues and increase flexibility, regardless of the number of people in the scene.

Our proposed interaction layer uses a standard selfattention mechanism. First, we add positional encodings to the query tokens to distinguish the two individuals. The encoded tokens then go through a self-attention layer, whose output is combined with the original tokens via a residual connection. Our experiments demonstrate that including these interaction layers significantly improves inter-person pose accuracy in two-person interaction benchmarks.

#### 3.4. PromptHMR video version

In addition to the single-image variant of PromptHMR, we train an extended version that processes videos to estimate human motion in world coordinates. To achieve this, we introduce a simple and efficient temporal transformer module. Given a monocular video sequence  $\{I^t\}_{t=0}^T$ , we first run



Figure 3. **SMPL-X decoder.** The top row shows one attention block in the decoder. The cross-person interaction module can be turned on/off. The bottom row shows the cross-person attention.

PromptHMR to obtain per-subject SMPL-X decoder output tokens  $T'_{smpl}$  and  $T'_{depth}$ , assuming that the subject identities are provided with the prompts. These tokens, along with the positional encoding of time t, are fed to a decoder-only temporal transformer module with twelve attention blocks. The output tokens are converted to SMPL-X parameters  $\phi_t, \theta_t, \beta_t$ , translation  $\tau_t$ , and joint contact probabilities  $c_t$ . The contact probabilities indicate whether a given joint is in contact with the ground plane similar to [50–52].

To obtain results in world coordinates, we adopt the approach from TRAM [65]. Specifically, we use DROID-SLAM [60] and a monocular metric depth estimation model, ZoeDepth [3], to estimate camera motion in metric world coordinates. The translation parameters  $\tau_t$  are then transformed to world coordinates using the estimated camera motion. To refine the human trajectory and mitigate foot-skating artifacts, we leverage the estimated contact probabilities and run a fast postprocessing that optimizes the contact joints to have zero velocity.

#### 3.5. Losses

PromptHMR is trained with a combination of 2D and 3D losses, following traditional HMR methods [22, 30]:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{2D} + \lambda_2 \mathcal{L}_{3D} + \lambda_3 \mathcal{L}_{SMPL} + \lambda_4 \mathcal{L}_V + \lambda_5 \mathcal{L}_t$$

with each term calculated as

$$\begin{aligned} \mathcal{L}_{2D} &= ||\hat{\mathcal{J}}_{2D} - \Pi(\mathcal{J}_{3D})||_{F}^{2} \\ \mathcal{L}_{3D} &= ||\hat{\mathcal{J}}_{3D} - \mathcal{J}_{3D}||_{F}^{2} \\ \mathcal{L}_{SMPL} &= ||\hat{\Theta} - \Theta||_{2}^{2} \\ \mathcal{L}_{V} &= ||\hat{V} - V||_{F}^{2} \\ \mathcal{L}_{t} &= ||\hat{p}_{xy} - p_{xy}||_{F}^{2} + ||\hat{p}_{z} - p_{z}||_{F}^{2} \end{aligned}$$

where  $\mathcal{J}_{3D}$  and V are the 3D joints and vertices of the SMPL-X model, with the hat operator denoting the ground truth. II is the camera reprojection operator. Additionally, on datasets with ground truth translation labels, we supervise the normalized translation  $p_{xy}$  and inverse depth  $p_z$ .

### 4. Experiments

**Datasets.** We train PromptHMR with standard datasets: BEDLAM [4], AGORA [44], 3DPW [63], COCO [33], and MPII [37]. Following 4DHumans, we add AIC [66] and InstaVariety [23] as in-the-wild data, with pseudoground truth from CamSMPLify [43]. Additionally, we add CHI3D [16] and HI4D [68] to enable learning two-person interaction following the train-test splits from BUDDI [38]. Including CHI3D and HI4D does not improve performance on other benchmarks.

**Implementation.** We train PromptHMR with AdamW with a batch size of 96 images of resolution  $896 \times 896$ . We use a learning rate of  $1e^{-5}$  for the image encoder and  $3e^{-5}$  for the prompt encoder and the SMPL-X decoder, with a weight decay of  $5e^{-5}$ . The training converges within 350K steps.

**Evaluation.** We evaluate camera space reconstruction accuracy on 3DPW [63], EMDB [24] and RICH [20], using MPJPE, Procrustes-aligned MPJPE (PA-MPJPE) and Per Vertex Error (PVE) [22]. We evaluate inter-person accuracy on HI4D and CHI3D by Pair-PA-MPJPE, which aligns the two people as a whole with the ground truth [38].

To evaluate world-grounded motion on EMDB with PromptHMR video (PromptHMR-vid), we compute Worldaligned MPJPE (WA-MPJPE<sub>100</sub>), World MPJPE (W-MPJPE<sub>100</sub>) and Root Translation Error (RTE in %) [52, 67].

#### 4.1. Reconstruction accuracy

For camera space reconstruction, as shown in Table 1, PromptHMR and PromptHMR-Vid demonstrate state-ofthe-art performance, matching crop-based methods while achieving better results than other full-image methods. PromptHMR and CameraHMR use the same training data and have similar performance, which validates that this prompt-based approach can achieve metrically accurate results. For representative results, see Fig. 7, where PromptHMR recovers coherent 3D scenes of people.

For interaction reconstruction, PromptHMR achieves good accuracy as indicated in Table 2. Compared to BUDDI which is also trained on CHI3D and HI4D, our method achieves better overall accuracy on per-person and interperson metrics. We show qualitative results in Fig. 8. As a monocular regression method, PromptHMR still cannot avoid interpenetration between closely interacting people.

PromptHMR-Vid achieves SOTA performance among methods that estimate human motion in world coordinates, as shown in Table 4. Unlike TRAM, we estimate the joint contact probabilities similar to [51, 52]. Therefore, we achieve lower foot skating than TRAM, even though we use the same metric SLAM method to transform motion in camera space to world coordinates. Please refer to our supplementary material (SupMat) for qualitative results of PromptHMR-Vid.



Figure 4. Effect of box prompts. Our method remains stable with different boxes, including noisy truncated boxes.



Figure 5. **Effect of mask prompts.** Results are from the same model with different prompt inputs. Masks are better for close interaction scenarios where boxes are ambiguous.



Figure 6. Effect of shape prompts. Compared to the baseline that does not incorporate shape description during training and testing, the model with shape prompts has better accuracy on HBW, especially in ambiguous images.

### 4.2. Effect of multimodal prompts

We conduct qualitative and quantitative evaluations of the multimodal prompts. For efficient ablation, we train models with  $448 \times 448$  input resolution and select the best model within 150K steps of training.

For box prompts, as shown in rows 3-4 of Fig. 7, our method is able to take a combination of different boxes from in-the-wild images to reconstruct crowded scenes. Figure 4 also shows an example with varying box inputs. PromptHMR remains stable when the boxes change and uses full image context to reconstruct the human even when the boxes are truncated.

The mask prompt is more effective than boxes when people closely overlap (Fig. 5), as boxes are ambiguous in such

		3DPW (14)		EMDB (24)			RICH (24)			
	Models	PA-MPJPE	MPJPE	PVE	PA-MPJPE	MPJPE	PVE	PA-MPJPE	MPJPE	PVE
cropped image	CLIFF* [31]	43.0	69.0	81.2	68.3	103.3	123.7	68.1	103.3	128.0
	HMR2.0a [17]	44.4	69.8	82.2	61.5	97.8	120.0	60.7	98.3	120.8
	TokenHMR [13]	44.3	71.0	84.6	55.6	91.7	109.4	_	-	_
	CameraHMR [43]	35.1	56.0	65.9	43.3	70.2	81.7	34.0	55.7	64.4
ge	BEV [58]	46.9	78.5	92.3	70.9	112.2	133.4	-	_	_
ima	Multi-HMR* [2]	45.9	73.1	87.1	50.1	81.6	95.7	46.3	73.8	83.0
full	PromptHMR*	36.6	58.7	69.4	41.0	71.7	84.5	37.3	56.6	65.5
video	WHAM [52]	37.5	59.8	71.5	52.0	81.6	96.9	44.3	80.0	91.2
	TRAM [65]	35.6	59.3	69.6	45.7	74.4	86.6	-	-	-
	GVHMR [51]	37.0	56.6	68.7	44.5	74.2	85.9	39.5	66.0	74.4
	PromptHMR-Vid	35.5	56.9	67.3	40.1	68.1	79.2	37.0	57.4	65.8

Table 1. Comparison of mesh reconstruction on the 3DPW, EMDB and RICH datasets, with the number of joints in parenthesis.  $\star$  denotes methods that use ground truth focal length during inference. Note that we remove the test-time flip augmentation from all of the video methods to ensure a fair comparison. All metrics are in mm.

	HI4D (14)			CHI3D (14)			
Models	PA-MPJPE	MPJPE	Pair-PA-MPJPE	PA-MPJPE	MPJPE	Pair-PA-MPJPE	
BEV* [58]	81	_	136	51	-	96	
BUDDI [38]	73	-	98	47	-	68	
Multi-HMR* [2]	49.8	67.8	80.6	31.7	54.0	100.0	
PromptHMR*	39.2	63.9	78.1	27.2	48.0	58.5	
PromptHMR	30.1	39.6	39.5	24.7	46.5	45.3	

Table 2. **Comparison on interaction reconstruction**. PromptHMR is more accurate in per-person and inter-person accuracy. \* denote a method or baseline is not trained on HI4D or CHI3D. All metrics are in mm. The impact of HI4D and the interaction prompt are evaluated in Table 5.

Train	Test w/ text	HBW					
w/ text		Height	Chest	Waist	Hip	P2P-20k	
×	×	69	51	88	63	26	
$\checkmark$	×	69	48	86	60	26	
$\checkmark$	$\checkmark$	62	43	76	58	24	

Table 3. Ablation of shape prompts using text. Training with shape prompts improves shape accuracy. Using shape prompts during inference further improves shape accuracy. The ablation study is conducted with a  $448 \times 448$  model. Errors are in mm.

cases. Ablation of HI4D (rows 1-2 in Tab. 5) shows that using masks as the spatial prompt improves accuracy.

Experiments on the HBW validation set (Tab. 3) show that text prompts effectively improve shape accuracy when used during both training and testing. Moreover, training with shape descriptions alone provides an accuracy boost even if prompts are not given at test time. As illustrated in Fig. 6, text prompts provide notable improvements, especially when large perspective effects create ambiguity.

	EMDB-2 (24)					
Models	WA-MPJPE <sub>100</sub>	$W\text{-}MPJPE_{100}$	RTE	Jitter	Foot Skating	
WHAM [52]	135.6	354.8	6.0	22.5	4.4	
TRAM [65]	76.4	222.4	1.4	18.5	23.4	
GVHMR [51]	111.0	276.5	2.0	16.7	3.5	
PromptHMR-Vid	71.0	216.5	1.3	16.3	3.5	

Table 4. **Evaluation of motion in world coordinates.** PromptHMR-Vid combined with metric SLAM from TRAM [65] surpasses SOTA methods at predicting human motion in world co-ordinates.

	Trained with	ı	HI4D (14)				
Mask	Interaction	HI4D	PA-MPJPE	MPJPE	Pair-PA-MPJPE		
×	×	×	47.0	71.4	87.2		
	×	×	43.4	60.5	83.0		
$\times$		×	43.7	61.3	73.0		
×	×		36.3	<u>49.4</u>	<u>52.6</u>		
$\checkmark$	$\checkmark$	$\checkmark$	<u>36.5</u>	47.1	47.9		

Table 5. Ablation on interaction prompt. The interaction module improves inter-person reconstruction metrics Pair-PA-MPJPE on HI4D, especially when the method does not include HI4D in training. Ablation is conducted with a  $448 \times 448$  model. All metrics are in mm.

For interaction prompts, we show an ablation in Table 5. The proposed interaction module is beneficial and largely improves inter-person accuracy on HI4D even without HI4D training, indicating out-off-domain generalization. When trained on HI4D, the interaction module does not improve per-person PA-MPJPE but still improves interperson Pair-PA-MPJPE. Please refer to our SupMat for more qualitative results on interaction prompts.



Figure 7. Qualitative comparison: Multi-HMR vs PromptHMR . Our model can recover coherent 3D scenes of people. In crowded scenes, face detection provides reliable box prompts for our model. Please zoom in to see the details.



Figure 8. **Qualitative results.** PromptHMR recovers coherent two-person close interaction. Despite suffering from some interpenetration, the relative positions of the interacting people are accurately recovered. More examples are provided in the Supplementary.

# 5. Limitations

We see PromptHMR as a step towards a holistic perception model for 3D humans, but several limitations need to be addressed in future work. Currently, the shape description and interaction prompts are not automatically generated and need to be supplied by the user. Future work should explore how to effectively integrate our promptable model with VLMs to automate prompting. We show how semantic prompts can improve reconstruction accuracy, but many other potential types of side information such as action descriptions, 3D scene context, or body measurements may provide additional benefits in different scenarios.

### 6. Conclusion

We have presented PromptHMR, a promptable HPS estimation approach that leverages full image context with spatial and semantic prompts to infer 3D humans in the scene. Our method demonstrates state-of-the-art accuracy across diverse benchmarks and generalizes well in the wild. Our experiments show that incorporating diverse input information through flexible prompting enables robustness and adaptability in challenging scenarios. Acknowledgement. The authors would like to thank Yan Zhang, Yao Feng, and Nitin Saini for their suggestions. The majority of the work was done when Yufu was an intern at Meshcapade. Yufu and Kostas thank the support of NSF NCS-FO 2124355, NSF FRR 2220868, and NSF IIS-RI 2212433.

**Disclosure.** While MJB is a co-founder and Chief Scientist at Meshcapade, his research in this project was performed solely at, and funded solely by, the Max Planck Society.

### References

- Nikos Athanasiou, Alpár Ceske, Markos Diomataris, Michael J. Black, and Gül Varol. MotionFix: Text-driven 3D human motion editing. In *SIGGRAPH Asia*, 2024. 3
- [2] Fabien Baradel, Matthieu Armando, Salma Galaaoui, Romain Brégier, Philippe Weinzaepfel, Grégory Rogez, and Thomas Lucas. Multi-HMR: Multi-person whole-body human mesh recovery in a single shot. *European Conference* on Computer Vision, 2024. 3, 4, 7
- [3] Shariq Farooq Bhat, Reiner Birkl, Diana Wofk, Peter Wonka, and Matthias Müller. ZoeDepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*, 2023. 5
- [4] Michael J Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. BEDLAM: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8726–8737, 2023. 6
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-toend object detection with transformers. In *European conference on computer vision*, pages 213–229, 2020. 3
- [6] Hongsuk Choi, Gyeongsik Moon, and Kyoung Mu Lee. Pose2Mesh: Graph convolutional network for 3D human pose and mesh recovery from a 2D human pose. In *European Conference on Computer Vision*, pages 769–787. Springer, 2020. 2
- [7] Hongsuk Choi, Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Beyond static features for temporally consistent 3D human pose and shape from a video. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1964–1973, 2021. 3
- [8] Vasileios Choutas, Lea Müller, Chun-Hao P. Huang, Siyu Tang, Dimitrios Tzionas, and Michael J. Black. Accurate 3D body shape regression using metric and semantic attributes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2718–2728, 2022. 2, 3, 4
- [9] Ginger Delmas, Philippe Weinzaepfel, Francesc Moreno-Noguer, and Grégory Rogez. PoseFix: correcting 3D human poses with natural language. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15018–15028, 2023. 3
- [10] Ginger Delmas, Philippe Weinzaepfel, Francesc Moreno-Noguer, and Grégory Rogez. PoseEmbroider: Towards a

3D, visual, semantic-aware human pose representation. In *European Conference on Computer Vision*, 2024. 2, 3

- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 4
- [12] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6569–6578, 2019. 3
- [13] Sai Kumar Dwivedi, Yu Sun, Priyanka Patel, Yao Feng, and Michael J Black. TokenHMR: Advancing human mesh recovery with a tokenized pose representation. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1323–1333, 2024. 7
- [14] Jose M. Facil, Benjamin Ummenhofer, Huizhong Zhou, Luis Montesano, Thomas Brox, and Javier Civera. CAM-Convs: Camera-aware multi-scale convolutions for singleview depth. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, page 11818–11827, 2019. 4
- [15] Yao Feng, Jing Lin, Sai Kumar Dwivedi, Yu Sun, Priyanka Patel, and Michael J. Black. ChatPose: Chatting about 3D human pose. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024. 2, 3
- [16] Mihai Fieraru, Mihai Zanfir, Elisabeta Oneata, Alin-Ionut Popa, Vlad Olaru, and Cristian Sminchisescu. Reconstructing three-dimensional models of interacting humans. arXiv preprint arXiv:2308.01854, 2023. 2, 3, 6
- [17] Shubham Goel, Georgios Pavlakos, Jathushan Rajasegaran, Angjoo Kanazawa, and Jitendra Malik. Reconstructing and tracking humans with transformers. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 2, 7
- [18] Dorian F Henning, Tristan Laidlow, and Stefan Leutenegger. BodySLAM: joint camera localisation, mapping, and human motion tracking. In *European Conference on Computer Vi*sion, pages 656–673. Springer, 2022. 3
- [19] Dorian F Henning, Christopher Choi, Simon Schaefer, and Stefan Leutenegger. BodySLAM++: Fast and tightlycoupled visual-inertial camera and human motion tracking. In *IEEE/RSJ International Conference on Intelligent Robots* and Systems, pages 3781–3788. IEEE, 2023. 3
- [20] Chun-Hao P Huang, Hongwei Yi, Markus Höschle, Matvey Safroshkin, Tsvetelina Alexiadis, Senya Polikovsky, Daniel Scharstein, and Michael J Black. Capturing and inferring dense full-body human-scene contact. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13274–13285, 2022. 2, 6
- [21] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5579–5588, 2020. 3

- [22] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 7122–7131, 2018. 2, 3, 5, 6
- [23] Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3D human dynamics from video. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5614–5623, 2019. 3, 6
- [24] Manuel Kaufmann, Jie Song, Chen Guo, Kaiyue Shen, Tianjian Jiang, Chengcheng Tang, Juan José Zárate, and Otmar Hilliges. EMDB: The electromagnetic database of global 3d human pose and shape in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14632–14643, 2023. 2, 6
- [25] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. arXiv preprint arXiv:2304.02643, 2023. 4
- [26] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. VIBE: Video inference for human body pose and shape estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5253–5263, 2020. 3
- [27] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. PARE: Part attention regressor for 3D human body estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11127– 11137, 2021. 2
- [28] Muhammed Kocabas, Chun-Hao P. Huang, Joachim Tesch, Lea Müller, Otmar Hilliges, and Michael J. Black. SPEC: Seeing people in the wild with an estimated camera. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 11035–11045, 2021. 2
- [29] Muhammed Kocabas, Ye Yuan, Pavlo Molchanov, Yunrong Guo, Michael J Black, Otmar Hilliges, Jan Kautz, and Umar Iqbal. PACE: Human and camera motion estimation from inthe-wild videos. In *International Conference on 3D Vision*, pages 397–408, 2024. 3
- [30] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2252–2261, 2019. 2, 3, 5
- [31] 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 *European Conference on Computer Vision*, pages 590–606. Springer, 2022. 7
- [32] Kevin Lin, Lijuan Wang, and Zicheng Liu. Mesh graphormer. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 12939–12948, 2021.
   2
- [33] 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

*European Conference on Computer Vision*, pages 740–755. Springer, 2014. 6

- [34] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023. 3
- [35] Thomas Lucas, Fabien Baradel, Philippe Weinzaepfel, and Grégory Rogez. Posegpt: Quantization-based 3d human motion generation and forecasting. In *European Conference on Computer Vision*, pages 417–435, 2022. 3
- [36] Zhengyi Luo, S. Alireza Golestaneh, and Kris M. Kitani. 3d human motion estimation via motion compression and refinement. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 3
- [37] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3D human pose estimation in the wild using improved cnn supervision. In *International Conference on 3D Vision*, pages 506–516. IEEE, 2017. 6
- [38] Lea Muller, Vickie Ye, Georgios Pavlakos, Michael J. Black, and Angjoo Kanazawa. Generative proxemics: A prior for 3D social interaction from images. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024. 3, 6, 7
- [39] Gyeongsik Moon and Kyoung Mu Lee. I2L-MeshNet: Image-to-lixel prediction network for accurate 3d human pose and mesh estimation from a single RGB image. In *European Conference on Computer Vision*, pages 752–768. Springer, 2020. 2
- [40] Raul Mur-Artal and Juan D Tardós. ORB-SLAM: An opensource slam system for monocular, stereo, and RGB-D cameras. *IEEE Transactions on Robotics*, 33(5):1255–1262, 2017. 3
- [41] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. ORB-SLAM: A versatile and accurate monocular SLAM system. *IEEE Transactions on Robotics*, 31(5):1147– 1163, 2015. 3
- [42] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision, 2023. 4
- [43] Priyanka Patel and Michael J. Black. Camerahmr: Aligning people with perspective. *International Conference on 3D Vision (3DV)*, 2025. 6, 7
- [44] 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 Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 13468– 13478, 2021. 6
- [45] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings of*

the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10975–10985, 2019. 2, 3

- [46] Mathis Petrovich, Michael J Black, and Gül Varol. Actionconditioned 3D human motion synthesis with transformer VAE. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 10985–10995, 2021. 3
- [47] Baldomero R. Árbol and Dan Casas. BodyShapeGPT: SMPL body shape manipulation with LLMs. In European Conference on Computer Vision Workshops, 2024. 3
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021. 2
- [49] Rene Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(3):1623–1637, 2022. 5
- [50] Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J Guibas. HUMOR: 3D human motion model for robust pose estimation. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 11488–11499, 2021. 5
- [51] Zehong Shen, Huaijin Pi, Yan Xia, Zhi Cen, Sida Peng, Zechen Hu, Hujun Bao, Ruizhen Hu, and Xiaowei Zhou. World-grounded human motion recovery via gravity-view coordinates. In SIGGRAPH Asia, 2024. 3, 6, 7
- [52] Soyong Shin, Juyong Kim, Eni Halilaj, and Michael J Black. WHAM: Reconstructing world-grounded humans with accurate 3D motion. *arXiv preprint arXiv:2312.07531*, 2023. 3, 5, 6, 7
- [53] Stephan Streuber, M Alejandra Quiros-Ramirez, Matthew Q Hill, Carina A Hahn, Silvia Zuffi, Alice O'Toole, and Michael J Black. Body talk: Crowdshaping realistic 3D avatars with words. ACM TOG, 35(4):1–14, 2016. 3
- [54] Sanjay Subramanian, Evonne Ng, Lea Müller, Dan Klein, Shiry Ginosar, and Trevor Darrell. Pose priors from language models. arXiv preprint arXiv:2405.03689, 2024. 3
- [55] Qingping Sun, Yanjun Wang, Ailing Zeng, Wanqi Yin, Chen Wei, Wenjia Wang, Haiyi Mei, Chi-Sing Leung, Ziwei Liu, Lei Yang, and Zhongang Cai. AiOS: All-in-one-stage expressive human pose and shape estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, page 1834–1843, 2024. 3
- [56] Yu Sun, Yun Ye, Wu Liu, Wenpeng Gao, Yili Fu, and Tao Mei. Human mesh recovery from monocular images via a skeleton-disentangled representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019. 3
- [57] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Michael J Black, and Tao Mei. Monocular, one-stage, regression of multiple 3D people. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11179–11188, 2021. 3
- [58] 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 *Proceedings of the IEEE/CVF Con-*

ference on Computer Vision and Pattern Recognition, pages 13243–13252, 2022. 7

- [59] 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 Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8856–8866, 2023. 3
- [60] Zachary Teed and Jia Deng. DRPOID-SLAM: Deep visual slam for monocular, stereo, and RGB-D cameras. Advances in Neural Information Processing Systems, 34:16558–16569, 2021. 3, 5
- [61] Zachary Teed, Lahav Lipson, and Jia Deng. Deep patch visual odometry. Advances in Neural Information Processing Systems, 36, 2024. 3
- [62] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *International Conference on Learning Representations*, 2023. 3
- [63] Timo Von Marcard, Roberto Henschel, Michael J Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *European Conference on Computer Vision*, pages 601–617, 2018. 2, 6
- [64] Yufu Wang and Kostas Daniilidis. ReFit: Recurrent fitting network for 3D human recovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14644–14654, 2023. 2
- [65] Yufu Wang, Ziyun Wang, Lingjie Liu, and Kostas Daniilidis. TRAM: Global trajectory and motion of 3d humans from inthe-wild videos. In *European Conference on Computer Vi*sion, 2024. 2, 3, 5, 7
- [66] Jiahong Wu, He Zheng, Bo Zhao, Yixin Li, Baoming Yan, Rui Liang, Wenjia Wang, Shipei Zhou, Guosen Lin, Yanwei Fu, et al. AI challenger: A large-scale dataset for going deeper in image understanding. arXiv preprint arXiv:1711.06475, 2017. 6
- [67] Vickie Ye, Georgios Pavlakos, Jitendra Malik, and Angjoo Kanazawa. Decoupling human and camera motion from videos in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21222–21232, 2023. 3, 6
- [68] Yifei Yin, Chen Guo, Manuel Kaufmann, Juan Jose Zarate, Jie Song, and Otmar Hilliges. Hi4D: 4D instance segmentation of close human interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17016–17027, 2023. 2, 3, 6
- [69] Ye Yuan, Umar Iqbal, Pavlo Molchanov, Kris Kitani, and Jan Kautz. GLAMR: Global occlusion-aware human mesh recovery with dynamic cameras. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 11038–11049, 2022. 3
- [70] Hongwen Zhang, Yating Tian, Xinchi Zhou, Wanli Ouyang, Yebin Liu, Limin Wang, and Zhenan Sun. PyMAF: 3D human pose and shape regression with pyramidal mesh alignment feedback loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11446– 11456, 2021. 2

[71] Yizhou Zhao, Tuanfeng Yang Wang, Bhiksha Raj, Min Xu, Jimei Yang, and Chun-Hao Paul Huang. Synergistic globalspace camera and human reconstruction from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1216–1226, 2024. 3