

Camera-Space Hand Mesh Recovery via Semantic Aggregation and Adaptive 2D-1D Registration

Xingyu Chen^{1*} Yufeng Liu^{1,3} Chongyang Ma¹ Jianlong Chang² Huayan Wang¹
Tian Chen¹ Xiaoyan Guo¹ Pengfei Wan¹ Wen Zheng¹

¹Y-tech, Kuaishou Technology ²Huawei Cloud & AI

³SEU-ALLEN Joint Center, Institute for Brain and Intelligence, Southeast University, China.

Abstract

Recent years have witnessed significant progress in 3D hand mesh recovery. Nevertheless, because of the intrinsic 2D-to-3D ambiguity, recovering camera-space 3D information from a single RGB image remains challenging. To tackle this problem, we divide camera-space mesh recovery into two sub-tasks, i.e., root-relative mesh recovery and root recovery. First, joint landmarks and silhouette are extracted from a single input image to provide 2D cues for the 3D tasks. In the root-relative mesh recovery task, we exploit semantic relations among joints to generate a 3D mesh from the extracted 2D cues. Such generated 3D mesh coordinates are expressed relative to a root position, i.e., wrist of the hand. In the root recovery task, the root position is registered to the camera space by aligning the generated 3D mesh back to 2D cues, thereby completing camera-space 3D mesh recovery. Our pipeline is novel in that (1) it explicitly makes use of known semantic relations among joints and (2) it exploits 1D projections of the silhouette and mesh to achieve robust registration. Extensive experiments on popular datasets such as FreiHAND, RHD, and Human3.6M demonstrate that our approach achieves state-of-the-art performance on both root-relative mesh recovery and root recovery. Our code is publicly available at <https://github.com/SeanChenxy/HandMesh>.

1. Introduction

Monocular 3D mesh recovery has attracted tremendous attention due to its extensive applications in AR/VR, human-machine interaction, etc. The task is to estimate 3D locations of mesh vertices from a single RGB image. It is particularly challenging owing to highly articulated structures, 2D-to-3D ambiguity, and self-occlusion. Significant efforts have been made recently for accurate 2D-to-3D

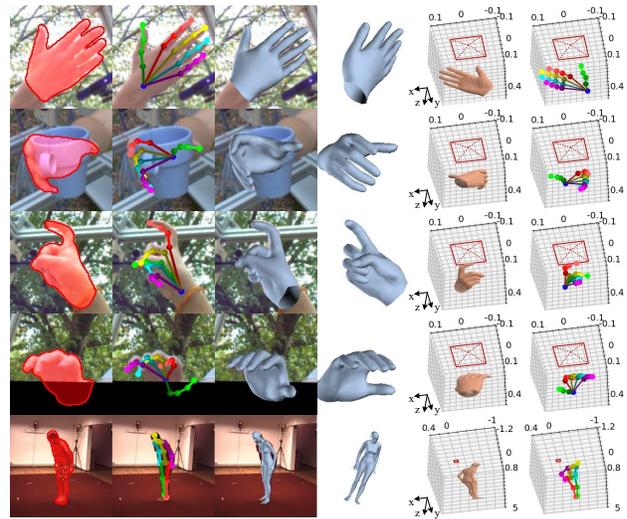


Figure 1. Qualitative results of the proposed CMR. We show silhouette, 2D pose, projection of mesh, side-view mesh, camera-space mesh and pose (unit: meter). The red rectangles indicate the camera. Our method is robust enough to handle cases of occlusion, truncation, and challenging poses.

reconstruction, including [16, 17, 20, 21, 23, 27, 30, 31, 32, 33, 36, 39], to name a few.

Most of the aforementioned methods [6, 9, 16, 20, 21, 22, 27, 40, 43] have difficulty in predicting absolute camera-space coordinates. Instead, they define a root (i.e., wrist of the hand) and estimate root-relative coordinates of the 3D mesh. In this aspect, these methods cannot be applied to many high-level tasks, e.g., hand-object interaction, that requires camera-space mesh information. To this end, we propose to jointly solve root-relative mesh recovery and root recovery by integrating these two sub-tasks into a unified framework, thereby bridging the gap between root-relative predictions and camera-space estimation.

RGB images consist of 2D patterns that are indirect cues of the underlying 3D structure. Therefore, 2D cues have long been leveraged to assist 3D tasks. For example, 2D

*Corresponding author, chenxingyu@kuaishou.com

pose and silhouette have been used to facilitate 3D pose regression [9, 20, 27, 33, 42, 43, 44].

However, the relationship between 2D cues and 3D structure remains unclear. We observe that 2D joint landmarks together with their semantic relations describe the 2D pose, while the silhouette indicates the holistic 3D-to-2D projection of the hand. They have different 2D properties and should be treated in different manners in the 3D task. Inspired by these observations, we set to explore the following aspects of the 2D-to-3D task: (1) different roles of 2D cues, (2) the reason for their different effects, and (3) how to construct more effective 2D cues.

In this paper, we propose a camera-space mesh recovery (CMR) framework to integrate the tasks of 3D hand mesh and root recovery into a unified framework. CMR consists of three phases, *i.e.*, 2D cue extraction, 3D mesh recovery, and global mesh registration. For 2D cue extraction, we predict joint landmarks and silhouette from a single RGB image. For mesh recovery, we introduce an Inception Spiral Module for robust 3D decoding. Moreover, an aggregation method is designed for composing more effective 2D cues. Specifically, instead of implicitly learning the relations among joints, we exploit their known relations by aggregating landmark heatmaps in groups, which proves to be effective for the subsequent 3D task. Finally, camera-space root location is obtained by a global mesh registration step that aligns the generated 3D mesh with the extracted 2D landmarks and silhouette. This step is carried out via an adaptive 2D-1D registration method that achieves robustness by leveraging matching objectives in different dimensions. Our full pipeline surpasses state-of-the-art methods in the 3D mesh and root recovery tasks. While our approach is mainly described for hand mesh, it can be readily applied to full-body mesh as shown in the experiments. Figure 1 demonstrates several example results of our CMR for camera-space mesh recovery.

Our main contributions are summarized as follows:

- We propose a novel aggregation method to collect effective 2D cues and exploit high-level semantic relations for root-relative mesh recovery.
- We design an adaptive 2D-1D registration method to sufficiently leverage both joint landmarks and silhouette in different dimensions for robust root recovery.
- We present a unified pipeline CMR for camera-space mesh recovery and demonstrate state-of-the-art performance on both mesh and root recovery tasks via extensive experiments on FreiHAND, RHD, and Human3.6M.

2. Related Work

Root-relative mesh/pose recovery. According to different output property, we categorize methods for single-view RGB-based mesh recovery into three types, *i.e.*,

RGB→*MANO/SMPL* [16, 40, 42, 44], *RGB*→*Voxel* [15, 27, 32, 39], and *RGB*→*Coord* (coordinate) [8, 9, 21, 22].

MANO [35] and SMPL [25] are parameterized 3D models of hand and human body, factorizing 3D human mesh into coefficients of shape and pose. Tremendous literature attempts to predict these coefficients for human/hand mesh recovery. For example, Zhou *et al.* [44] estimated MANO coefficients based on the kinematic chain and developed an inverse kinematics network to improve prediction accuracy on pose coefficients. MANO/SMPL can reconstruct 3D mesh, but they embed 3D information into a parametrized space (*e.g.*, PCA space), where the 3D structure is less straightforward (compared to 3D vertices).

Voxel is one type of Euclidean 3D representation, to which the canonical convolutional operator can be directly applied [28]. Thereby, the mesh recovery task can be explored in voxels. For example, Moon *et al.* [27] proposed an I2L-MeshNet by dividing voxels into three voxel spaces, where a 2.5D representation is leveraged for human mesh. The voxel/2.5D-based paradigm has impressive performance in terms of human mesh recovery because the merits of Euclidean space are fully leveraged. However, voxel/2.5D representations are not efficient enough in capturing 3D details (compared to 3D vertices).

Defferrard *et al.* [7] proposed a graph convolution network (GCN) based on spectral filtering to process 3D vertices in the non-Euclidean space. Based on GCN, Kolotouros *et al.* [21] developed a graph convolutional mesh regressor to directly estimate the 3D coordinates of mesh vertices. Ge *et al.* [9] also developed a graph-based method for hand mesh recovery by learning from mixed real and synthetic data. Instead of spectral filtering, Lim *et al.* [24] proposed spiral convolution (SpiralConv) to process mesh data in the spatial domain. Based on SpiralConv, Kulon *et al.* [22] developed an encoder-decoder structure for efficient hand mesh recovery. We follow the *RGB*→*Coord* paradigm and explore versatile aggregation of 2D cues.

Root recovery. In analogy to root recovery, estimation of external camera parameters has been widely studied [3, 16, 42]. For instance, Zhang *et al.* [42] designed an iterative regression method to simultaneously estimate external camera parameters and MANO coefficients. However, camera parameter estimation from RGB data is an ill-posed problem, leading to relatively low generalization performance. Moon *et al.* [26] proposed RootNet to predict the absolute 3D human root. RootNet essentially modeled object size in images, but the pixel-level object area has a relatively low correlation with 3D root. Rogez *et al.* [34] predicted both 2D and 3D pose so that the 3D root location can be obtained by aligning predicted 2D pose with projected 3D pose. We argue that this 2D-3D alignment cannot sufficiently leverage 2D information and propose an adaptive 2D-1D registration method for root recovery.

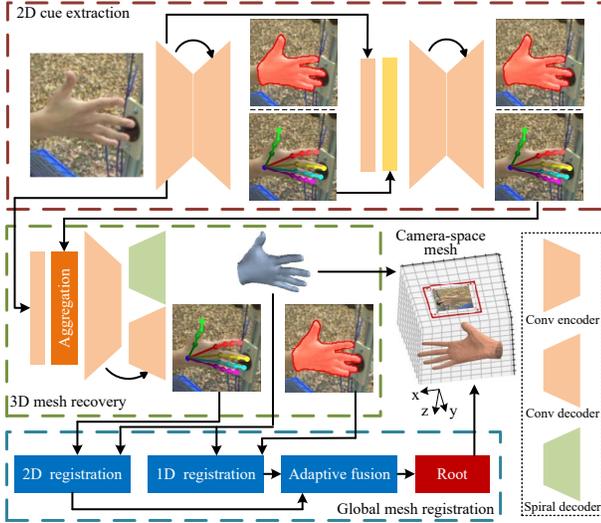


Figure 2. Overview of our CMR framework.

2D cues in 3D shape/pose recovery. Researchers have long exploited 2D cues in recovering 3D human shape and body parts. Pavlakos *et al.* [33] utilized 2D pose to regress pose coefficients and used silhouette to estimate shape coefficients. Varol *et al.* [39] first predicted 2D joint landmarks and body part segmentation, both of which were then combined to predict 3D pose. In this work, we aim to investigate how 2D cues work in 3D tasks and to leverage them effectively for hand mesh and root recovery.

3. Our Method

To represent a 3D mesh in camera-space, we divide it into the root-relative mesh and the camera-space root location. As shown in Figure 2, CMR includes three phases, *i.e.*, 2D cue extraction, 3D mesh recovery, and global mesh registration. In the step of 2D cue extraction, we predict 2D pose and silhouette, which are used later for both mesh and root recovery. In the step of 3D mesh recovery, we generate a root-relative mesh which is then registered to the camera space in the final phase of the pipeline.

3.1. Mesh Recovery by Semantic Aggregation

2D cues for 3D mesh recovery. The first phase of our pipeline extracts 2D pose and silhouette. Both 2D pose and silhouette are represented by heatmaps. To refine the 2D cues gradually, we use a multi-stack hourglass network [29].

The 3D mesh is defined by its *shape* and *pose* [35]. The silhouette is a holistic 3D-to-2D projection so it captures important shape cues. However, it can hardly describe the pose accurately. On the other hand, joint landmark locations are very informative for the pose. Given their different roles, how to better combine them becomes an interesting question. Therefore, we have an insight that to improve accuracy of the subsequent 3D tasks, it is essential not only

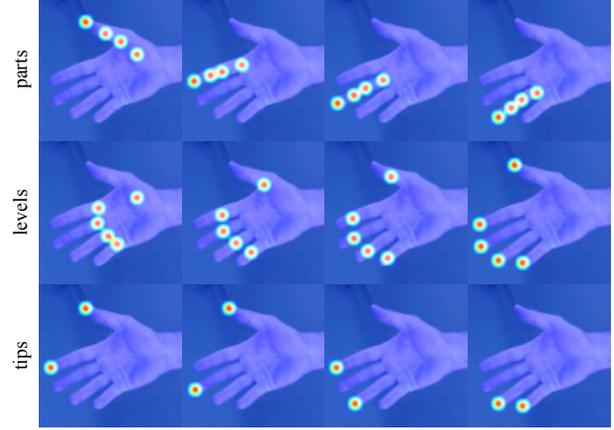


Figure 3. Examples of sub-poses for high-level semantic relations. From top to bottom: part, level, and tip grouping.

to improve the accuracy of the 2D tasks respectively, but also to better aggregate them according to their semantic relations. Specifically, we propose to aggregate a series of 2D cues denoted as follows:

- \mathbf{H}_p : N heatmaps of 2D poses. Each heatmap corresponds to a joint landmark.
- \mathbf{H}_s : a single heatmap of the silhouette.
- $cat(\mathbf{H}_p, \mathbf{H}_s)$: concatenating heatmaps of \mathbf{H}_p and \mathbf{H}_s to aggregate 2D pose and silhouette.
- $sum(\mathbf{H}_p)$: combining all the joint landmarks as a single heatmap to aggregate joint locations.
- $group(\mathbf{H}_p)$: concatenating \mathbf{H}_p and tip-, part-, or level-grouped landmarks to aggregate joint semantics for high-level semantic relations.

2D silhouette and joint landmarks represent pixel-level locations in different aspects. A straightforward way of combining them is $cat(\mathbf{H}_p, \mathbf{H}_s)$. Another simple baseline, $sum(\mathbf{H}_p)$, would discard semantics of individual joints by encoding their locations in a single heatmap. Thus, comparing \mathbf{H}_p and $sum(\mathbf{H}_p)$ would reveal the effect of joint semantics. These simple baselines, as we will show, are inferior to more semantically meaningful ways of aggregation: $group(\mathbf{H}_p)$. That is, we sum the joint heatmaps in groups. As shown in Figure 3, three ways of grouping are introduced, *i.e.*, by part, by level, and by tip grouping. Part grouping integrates joint landmarks on a finger, leg, arm, or torso, while level grouping integrates joint landmarks at the kinematic level [44]. Tip grouping integrates pairwise part tips. As a result, $group(\mathbf{H}_p)$ forms sub-poses that exploit high-level semantic relation of 2D joints.

Spiral decoder. The second phase of our pipeline generates the root-relative 3D mesh from the aggregated 2D cues using an improved spiral convolution decoder.

A 3D mesh \mathcal{M} contains vertices $\mathcal{V} = \{\mathbf{v}_i = (x_i, y_i, z_i)\}_{i=1}^M$ and faces \mathcal{F} . Convolution methods for \mathcal{M}

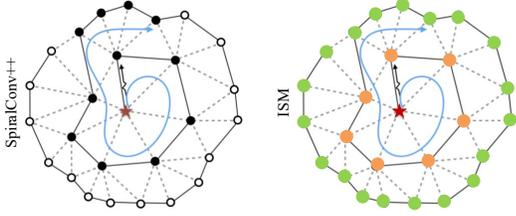


Figure 4. Neighbor selection of different spiral convolutions. **Left:** SpiralConv++ [10]. **Right:** Our method (ISM) that features hierarchical receptive fields (shown in red, orange, and green).

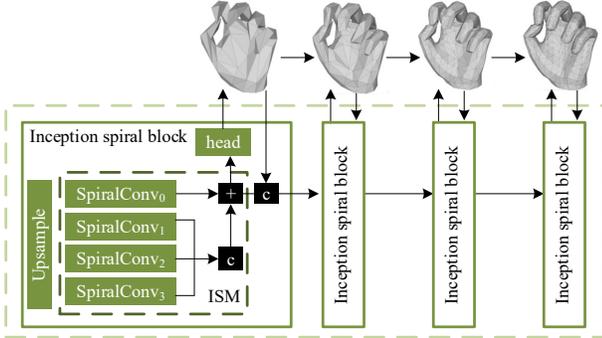


Figure 5. Spiral decoder. We improve the spiral decoder by ISM, multi-scale mechanism, and self-regression.

essentially process vertex features $f(\mathbf{v})$. SpiralConv [24] is a graph-based convolution operator, which processes vertex features in the spatial domain. By explicitly formulating the order of aggregating neighboring vertices, SpiralConv++ [10] presents an efficient version of SpiralConv. SpiralConv++ depends on a spiral manner of neighbor selection and adopts a fully-connected layer for feature fusion:

$$\begin{aligned}
 0\text{-ring}(\mathbf{v}) &= \{\mathbf{v}\} \\
 (k+1)\text{-ring}(\mathbf{v}) &= \mathbb{N}(k\text{-ring}(\mathbf{v})) \setminus k\text{-disk}(\mathbf{v}) \\
 k\text{-disk}(\mathbf{v}) &= \cup_{i=0, \dots, k} i\text{-ring}(\mathbf{v}) \\
 \text{SpiralConv}++(\mathbf{v}) &= W(f(k\text{-disk}(\mathbf{v}))) + b,
 \end{aligned} \tag{1}$$

where \mathbb{N} represents vertex neighborhood, W and b are weights and bias shared for $\forall \mathbf{v} \in \mathcal{V}$.

As shown in Figure 4(left), SpiralConv++ collects vertex neighbors (black dots) of the cell (red star) in a spiral manner. Then these neighbors are treated indiscriminately by a fully connected layer. Inspired by Inception and residual models [13, 38], we design an Inception Spiral Module (ISM) to enhance the receptive field of SpiralConv. Specifically, as shown in Figure 4(right), we distinguish neighbors according to the spiral hierarchy (red, orange, and green dots) and adopt parallel layers with diverse receptive field for 3D decoding. The ISM can be described as

$$\begin{aligned}
 o_i(\mathbf{v}) &= W_i f(i\text{-disk}(\mathbf{v})) + b_i \quad i = 0, 1, 2, 3 \\
 \text{ISM}(\mathbf{v}) &= o_0(\mathbf{v}) + [o_1(\mathbf{v}), o_2(\mathbf{v}), o_3(\mathbf{v})].
 \end{aligned} \tag{2}$$

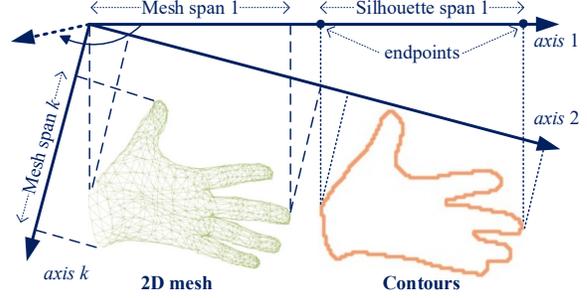


Figure 6. Multiple 1D projections for the mesh and silhouette.

where $[\cdot]$ denotes concatenating. In ISM we keep the number of parameters manageable by controlling the channel size of o . Note that the i -disk for each vertex may contain different number of elements. Similar to SpiralConv++, we truncate it to obtain a fixed-length sequence so that W_i and b_i can be shared for all the vertices.

The overall architecture of our 3D mesh decoder is shown in Figure 5. Our design improves the spiral decoder in three aspects: (1) we replace SpiralConv with ISM; (2) we leverage multi-scale prediction and coarse-to-fine fusion; and (3) we introduce a self-regression mechanism by concatenating scale-level predictions with the same-scale feature. Meanwhile, we use a convolutional decoder which runs in parallel with the spiral decoder to refine the estimation of 2D pose and silhouette.

3.2. Root Recovery by Global Mesh Registration

1D projections. The silhouette reflects holistic 3D-to-2D projection and contains strong geometric information for root recovery. Given the intrinsic matrix K of the camera, predicted 3D vertices \mathcal{V} can be projected into the 2D space by $K\mathcal{V}$, resulting in a 2D mesh which consists of 2D vertices $\mathcal{V}^{2D} = \{\mathbf{v}_i^{2D} = (x_i^{2D}, y_i^{2D})\}_{i=1}^M$ with the original connectivity. Ideally, the 2D mesh should align well with the silhouette. However, they cannot be directly aligned because (1) the large number of 2D points leads to prohibitive computational cost, *i.e.*, the silhouette contains thousands of pixels while the 2D mesh has 778 (for hand) or 6,890 (for human body) vertices; (2) 2D mesh vertices have no explicit correspondence with respect to the silhouette.

To overcome these difficulties, we design a group of 1D projections to align the silhouette with the 2D mesh. First, the silhouette is converted into contours $\mathcal{C} = \{\mathbf{c} = (x_i^c, y_i^c)\}_{i=1}^C$ by edge detection [5]. We define a set of 1D axes $\mathcal{A} = \{\mathbf{a}_j = (x_j^a, y_j^a)\}_{j=1}^A$. These axes have unit length and are uniformly distributed, *i.e.*, the angle difference between neighboring axes is π/A . As shown in Figure 6, we project contours and 2D vertices onto \mathbf{a}_j , resulting in their 1D span, *i.e.*, $\mathcal{S}_j^c = \{\mathbf{c}_i \cdot \mathbf{a}_j\}_{i=1}^C$ and $\mathcal{S}_j^{v^{2D}} = \{\mathbf{v}_i^{2D} \cdot \mathbf{a}_j\}_{i=1}^M$.

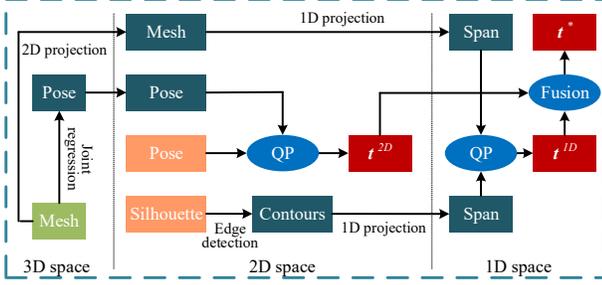


Figure 7. Adaptive 2D-1D registration. For root recovery, 2D registration aligns 3D mesh with 2D pose while 1D registration aligns 3D mesh with silhouette. “QP” indicates quadratic programming. The green and orange elements are from the network.

Adaptive 2D-1D registration. We denote the camera-space root by $\mathbf{t} = (x^r, y^r, z^r)$ and 2D joint landmark predictions by $\mathcal{P} = \{\mathbf{p}_i = (x_i^p, y_i^p)\}_{i=1}^N$. Given K and joint regressor J defined by MANO [35] or SMPL [25], the camera-space 3D vertices $\mathcal{V} + \mathbf{t}$ can be converted into 2D joints by $\mathcal{Q} = KJ(\mathcal{V} + \mathbf{t}) = \{\mathbf{q}_i = (x_i^q, y_i^q)\}_{i=1}^N$. Since \mathcal{P} and \mathcal{Q} have intrinsic correspondence, we define the energy function for 2D matching as

$$E_{2D} = \sum_{i=1}^N (\mathbf{p}_i - \mathbf{q}_i)^2, \quad (3)$$

whose solution is denoted as $\mathbf{t}^{2D} = \arg \min_{\mathbf{t}} E_{2D}(\mathbf{t})$. Further, we use $\mathcal{V} + \mathbf{t}$ to implement the aforementioned 3D-2D-1D projection. Camera-space $S^{\mathcal{V}^{2D}}$ is thereby produced. Then the 1D correspondence is obtained using two endpoints of the 1D span. The energy function for alignment of 1D spans is defined as

$$E_{1D} = \sum_{j=1}^A \{(\max(\mathcal{S}_j^c) - \max(\mathcal{S}_j^{\mathcal{V}^{2D}}))^2 + (\min(\mathcal{S}_j^c) - \min(\mathcal{S}_j^{\mathcal{V}^{2D}}))^2\}. \quad (4)$$

In 1D space, $\mathbf{t}^{1D} = \arg \min_{\mathbf{t}} E_{1D}(\mathbf{t})$. Both 2D and 1D optimizations are based on quadratic programming [2]. With \mathbf{t}^{1D} and \mathbf{t}^{2D} , we develop an adaptive method for the final root \mathbf{t}^* with their distance $d = \|\mathbf{t}^{2D} - \mathbf{t}^{1D}\|_2$:

$$\mathbf{t}^* = \begin{cases} \mathbf{t}^{1D} & \text{if } d > \delta_1 \\ \frac{\delta_1 - d}{\delta_1 - \delta_2} \mathbf{t}^{2D} + \frac{d - \delta_2}{\delta_1 - \delta_2} \mathbf{t}^{1D} & \text{else if } d > \delta_2 \\ \mathbf{t}^{2D} & \text{otherwise} \end{cases} \quad (5)$$

where $\delta_1 > \delta_2$, both of which are robust hyper-parameters according to 3D scale. Without wide search, we empirically use 0.06, 0.02 for the hand and 1.0, 0.5 for the body. This design attempts to sufficiently leverage the merits of 2D-1D registration. It is known that 2D pose is more fragile than silhouette. Hence, 2D pose is prone to be erroneous when there is a huge 2D-1D discrepancy, and silhouette is more dependable if $d > \delta_1$. From another perspective, joint correspondence is more explicit than that of 1D projection.

Thereby, the 2D process is more reliable if 2D and 1D results are similar ($d < \delta_2$). The whole process of adaptive 2D-1D registration is presented in Figure 7, from which we can see that geometrical information is sufficiently exploited from 3D to 1D spaces for root recovery.

3.3. Loss Functions

We use L1 norm for loss terms of 3D mesh/pose \mathcal{L}_{mesh} , \mathcal{L}_{pose3D} , and our 2D pose/silhouette losses \mathcal{L}_{pose2D} , \mathcal{L}_{sil} are based on binary cross entropy (BCE). We adopt normal loss \mathcal{L}_{norm} and edge length loss \mathcal{L}_{edge} for smoother reconstruction [9]. Formally, we have

$$\begin{aligned} \mathcal{L}_{mesh} &= \|\mathcal{V} - \mathcal{V}^*\|_1, \mathcal{L}_{pose3D} = \|J\mathcal{V} - \mathcal{J}^*\|_1, \\ \mathcal{L}_{pose2D} &= \text{BCE}(U, U^*), \mathcal{L}_{sil} = \text{BCE}(S, S^*), \\ \mathcal{L}_{norm} &= \sum_{\mathbf{k} \in \mathcal{F}} \sum_{(i,j) \subset \mathbf{k}} \left| \frac{\mathcal{V}_i - \mathcal{V}_j}{\|\mathcal{V}_i - \mathcal{V}_j\|_2} \cdot \mathbf{n}_{\mathbf{k}}^* \right|, \\ \mathcal{L}_{edge} &= \sum_{\mathbf{k} \in \mathcal{F}} \sum_{(i,j) \subset \mathbf{k}} \left| \|\mathcal{V}_i^p - \mathcal{V}_j^p\|_2 - \|\mathcal{V}_i^* - \mathcal{V}_j^*\|_2 \right|, \end{aligned} \quad (6)$$

where \mathcal{F}, \mathcal{V} are faces and vertices of a mesh; J, \mathcal{J} are the pose regressor and 3D joints; $\mathbf{n}_{\mathbf{k}}^*$ indicates unit normal vector of face \mathbf{k} ; U, S are heatmaps of 2D pose and silhouette; and \star denotes the ground truth. Following [27], U^* is constructed with Gaussian distribution.

Our overall loss function is $\mathcal{L}_{total} = \mathcal{L}_{mesh} + \mathcal{L}_{pose3D} + \lambda_p \mathcal{L}_{pose2D} + \lambda_s \mathcal{L}_{sil} + \lambda_n \mathcal{L}_{norm} + \mathcal{L}_{edge}$, where $\lambda_p = 10$, $\lambda_s = 0.5$, $\lambda_n = 0.1$ are used to balance different terms.

4. Experiments

4.1. Experimental Setup

We conduct experiments on several commonly-used benchmarks as listed below.

FreiHAND [46] is a 3D hand dataset with 130,240 training images and 3,960 evaluation samples. The annotations of the evaluation set are not available, so we submit our predictions to the official server for online evaluation.

Rendered Hand Pose Dataset (RHD) [45] consists of 41,258 and 2,728 virtually rendered samples for training and testing on hand pose estimation, respectively.

Human3.6M [14] is a large-scale 3D body pose benchmark containing 3.6 million video frames with annotations of 3D joint coordinates. SMPLify-X [30] is used to obtain ground-truth SMPL coefficients. We follow existing methods [6, 16, 27, 32] to use subjects $S1, S5, S6, S7, S8$ for training and subjects $S9, S11$ for testing.

COCO is a wild dataset with annotations of 2D human joints. Following previous work [6, 20, 27], we use the SMPL coefficients to produce the 3D human mesh. COCO is used for training.

We use the following metrics in quantitative evaluations.

Method	PA-MPJPE ↓	PA-MPVPE ↓
YoutubeHand [22]	8.82	9.06
+ Our spiral decoder	8.46	8.54

Table 1. Comparison of our spiral decoder with baseline.

Stack 1	Stack 2	PA-MPJPE ↓	PA-MPVPE ↓
—	—	8.46	8.54
\mathbf{H}_s	—	8.05	8.10
\mathbf{H}_p	—	7.71	7.77
$cat(\mathbf{H}_p, \mathbf{H}_s)$	—	7.83	7.95
$sum(\mathbf{H}_p)$	—	7.89	7.94
$group(\mathbf{H}_p)$	—	7.55	7.60
\mathbf{H}_s	$group(\mathbf{H}_p)$	7.47	7.55
\mathbf{H}_p	$group(\mathbf{H}_p)$	7.39	7.46

Table 2. Effects of 2D-cues aggregations after each ‘‘Stack’’.

MPJPE/MPVPE measures the mean per joint/vertex position error in terms of Euclidean distance (mm) between the root-relative prediction and ground-truth coordinates. **PA-MPJPE/MPVPE** is the MPJPE/MPVPE based on procrustes analysis [11] with global variation being ignored. **CS-MPJPE/MPVPE** measures MPJPE/MPVPE in the camera space for evaluation of the root recovery task. **AUC** is the area under the curve of PCK (percentage of correct keypoints) vs. error thresholds.

Implementation details. Our backbone is based on ResNet [13] and we use the Adam optimizer [19] to train the network with a mini-batch size of 32. Serving as network inputs, image patches are cropped and resized to resolutions of 224×224 (for the hand) or 256×256 (for the body). Because of different data amounts, the total number of iterations is set as 38 and 25 epochs for tasks on hand and body. The initial learning rate is 10^{-4} , which is divided by 10 at the 20th or 30th epoch. Data augmentation includes random box scaling/rotation, color jitter, etc.

4.2. Ablation Study

Baseline. Our ablation studies are based on ResNet18. YoutubeHand [22] serves as the baseline, but its code and models are inaccessible. Our re-implemented model obtains 9.06mm PA-MPVPE (see Table 1). With our spiral decoder, 8.54mm PA-MPVPE is achieved. Thus, our designs can strengthen the robustness of 2D-to-3D decoding.

Effects of various 2D cues. This paper explores cues of 2D pose and silhouette for 3D mesh recovery, and a two-stack network is leveraged. At first, discarding the second stack, we only use one encoder-decoder structure for exposing details of 2D-cue effects. As shown in Table 2, \mathbf{H}_s leads to 8.10mm PA-MPVPE while $sum(\mathbf{H}_p)$ reduces PA-MPVPE to 7.94mm. Both \mathbf{H}_s and $sum(\mathbf{H}_p)$ are heatmaps that contain location information, so the locations of joint landmarks are more instructive than that of a holistic shape. This phenomenon is evident, because pose is more difficult to estimate than shape, and joint locations can directly

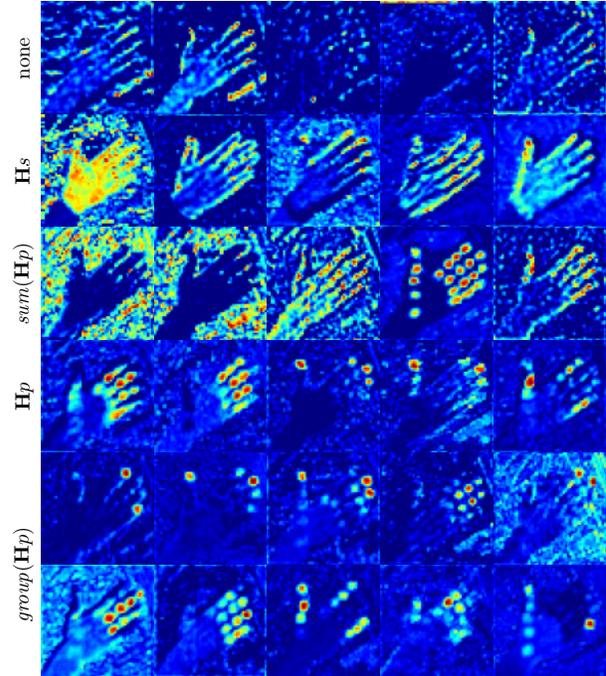


Figure 8. Comparison of feature representations after various 2D cues. \mathbf{H}_s and $sum(\mathbf{H}_p)$ invite holistic shape and pose. \mathbf{H}_p implies semantic relation while $group(\mathbf{H}_p)$ induces more exhaustive joint relations by explicit semantic aggregation. The input image is shown in the first row of Figure 1.

provide cues on poses. Furthermore, \mathbf{H}_p leads to an improved PA-MPVPE of 7.77mm by providing both joint locations and semantics. Compared to $sum(\mathbf{H}_p)$, it is seen that joint semantics is also important. In addition, the effect of $cat(\mathbf{H}_p, \mathbf{H}_s)$ is worse than that of \mathbf{H}_p . Therefore, there is no complementary benefit from 2D pose and silhouette.

To reveal potential reasons for different effects of 2D cues, we dissect feature representation acted by them. It is known that channel-specific feature usually embeds semantics, so we focus on typical channels in the first encoding block in the 3D mesh recovery phase. As shown in Figure 8, this block tends to describe trivial properties such as edges (see ‘‘none’’). If \mathbf{H}_s is employed, holistic 2D shapes emerge, but pose cues are ignored. With $sum(\mathbf{H}_p)$, holistic 2D shapes still can be learned based on joint landmark locations. Thus, this phenomenon can be the reason why there exists no complementary benefit in joint landmarks and silhouette. Moreover, holistic joint locations can also be captured with $sum(\mathbf{H}_p)$. As for \mathbf{H}_p , although joint information is provided in separate heatmaps, the features after \mathbf{H}_p have a tendency that multiple joints are simultaneously activated. Features shown in \mathbf{H}_p of Figure 8 imply semantic relation of joints, which can have significant impact on 3D tasks. However, relation representation invited by \mathbf{H}_p is not exhaustive enough. Specifically, the left two representations have similar patterns while both

Parts	Levels	Tips	PA-MPJPE ↓	PA-MPVPE ↓
			7.71	7.77
✓			7.58	7.65
	✓		7.63	7.69
		✓	7.58	7.65
✓		✓	7.55	7.60
✓	✓	✓	7.62	7.70

Table 3. Effects of semantic aggregation of joint landmarks.

Model	Root recovery	CJ ↓	CV ↓
Absolute coordinate prediction		77.6	77.4
CMR-P (2D AUC=0.762, mIoU=0.824)	PnP	98.9	99.0
	RootNet	63.0	62.9
	2D (ours)	53.7	53.7
	1D (ours, A=2)	55.8	55.8
	2D+1D (ours, A=2)	52.7	52.7
	2D+1D (ours, A=12)	52.1	52.1
CMR-PG (2D AUC= 0.798 , mIoU=0.826)	RootNet	62.7	62.7
	2D (ours)	51.5	51.5
	1D (ours, A=2)	54.3	54.3
	2D+1D (ours, A=2)	51.1	51.1
	2D+1D (ours, A=12)	50.6	50.6
CMR-SG (2D AUC=0.790, mIoU= 0.832)	RootNet	62.8	62.7
	2D (ours)	52.0	52.0
	1D (ours, A=2)	52.6	52.6
	2D+1D (ours, A=2)	50.4	50.4
	2D+1D (ours, A=12)	49.7	49.8

Table 4. Root recovery ablation/comparison. 2D AUC and mIoU indicate the accuracy of 2D pose and silhouette. Because official evaluation do not support them, 2D AUC is about RHD while mIoU is based on our silhouette annotation. PnP [18] and RootNet [26] are compared. “CJ, CV” denote CS-MPJPE/MPVPE.

the 3rd and 4th representations mainly focus on the tips of thumb and middle finger.

Effects of semantic aggregation. Instead of implied relation learning, $group(\mathbf{H}_p)$ aims to aggregate joint semantics for explicit semantic relation. Referring to Table 3, part-, level-, and tip-based relations have instructive effects, and the integration of part- and tip-based relations leads to better performance on mesh recovery. Part-based aggregation provides within-part relations while tip-based aggregation models pairwise cross-part relations, so they are complementary for creating robust semantic relations. As a result, $group(\mathbf{H}_p)$ surpasses \mathbf{H}_p by 0.17mm on PA-MPVPE, *i.e.*, 7.60mm. Referring to Figure 8, compared to \mathbf{H}_p , $group(\mathbf{H}_p)$ captures more exhaustive semantic relation, *e.g.*, finger parts and various combinations of fingertips.

Besides, the two-stack network can provide more effective 2D cues, leading to better PA-MPVPE (see Table 2). Let CMR-P, CMR-SG, and CMR-PG be models shown in the 3rd, 7th, and 8th rows of Table 2 for later analysis.

Effects of adaptive 2D-1D registration. We train a model that directly predicts absolute camera-space coordi-

Method	PJ ↓	PV ↓	CJ ↓	CV ↓
Boukhayma <i>et al.</i> [3]	35.0	13.2	–	–
ObMan [12]	13.3	13.3	85.2	85.4
MANO CNN [46]	11.0	10.9	71.3	71.5
YoutubeHand [22]	8.4	8.6	–	–
Pose2Mesh [6]	7.7	7.8	–	–
I2L-MeshNet [27]	7.4	7.6	60.3	60.4
CMR-SG (ResNet18)	7.5	7.6	49.7	49.8
CMR-PG (ResNet18)	7.4	7.5	50.6	50.6
CMR-SG (ResNet50)	7.0	7.1	48.8	48.9
CMR-PG (ResNet50)	6.9	7.0	48.9	49.0

Table 5. Results on FreiHAND. “PJ, PV, CJ, CV” denote PA-MPJPE, PA-MPVPE, CS-MPJPE, and CS-MPVPE, respectively.

nates for comparison. Although this operation easily suffers from overfitting, it can handle the single-dataset task, obtaining 77.4mm CS-MPVPE (see Table 4). Perspective-n-Point (PnP) is a standard approach that solves external parameters of a camera based on pairwise 2D-3D points [18]. Based on 2D joint landmarks and root-relative 3D joints from CMR-P, the PnP method can predict absolute root coordinates. However, it obtains 99.0mm CS-MPVPE, lagging considerably behind the baseline. With our 2D or 1D registration, CMR-P achieves 53.7mm or 55.8mm CS-MPVPE, so both our 2D and 1D designs are valid. In detail, 2D registration is better because of unambiguous 2D-3D correspondence. Furthermore, based on our proposed adaptive fusion, the adaptive 2D-1D registration achieves 52.7mm CS-MPVPE. When $A = 12$, the 2D-1D process can induce a better CS-MPVPE of 52.1mm. Thus, our adaptive 2D-1D registration can sufficiently leverage 2D cues from joint landmarks and silhouette for root recovery.

With CMR-PG and CMR-SG, the same tendency emerges, which also validates the effectiveness of our designs. Note that CMR-PG induces better 2D pose (*i.e.*, 0.798 AUC) while CMR-SG invites better silhouette (*i.e.*, 0.832 mIoU), so their 2D and 1D performances are distinct. Overall, CMR-SG achieves the best performance on root recovery and camera-space mesh recovery.

4.3. Comparisons with Existing Methods

We perform a comprehensive comparison on the FreiHAND dataset. As shown in Table 5, our proposed CMR outperforms other methods in terms of all the aforementioned metrics. Specifically, ResNet50-based CMR-PG achieves the best performance on root-relative mesh recovery, *i.e.*, 7.0mm PA-MPVPE and 6.9mm PA-MPJPE.

In Table 5, we evaluate the root recovery performance of camera-space mesh on the FreiHAND dataset. Our CMR outmatches ObMan [12] and I2L-MeshNet [27] by 36.5mm and 11.5mm on CS-MPVPE. To the best of our knowledge, CMR achieves state-of-the-art performance on camera-space mesh recovery, *i.e.*, 48.9mm CS-MPVPE and 48.8mm CS-MPJPE. From Figure 9, we can see that

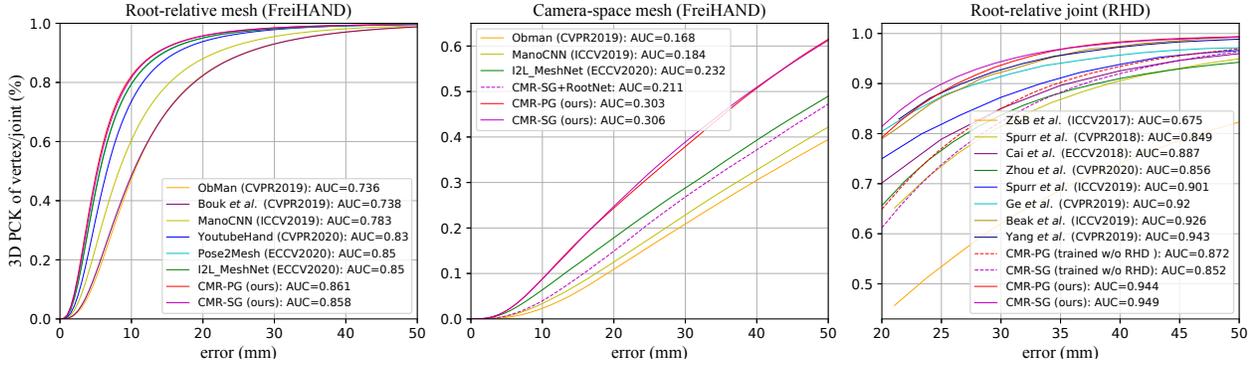


Figure 9. PCK vs. error thresholds. Our CMR significantly outperforms all the other methods.

Method	Type	<i>Human3.6M</i>		<i>Human3.6M+COCO</i>		FPS \uparrow	GPU Mem. \downarrow
		MPJPE \downarrow	PA-MPJPE \downarrow	MPJPE \downarrow	PA-MPJPE \downarrow		
HMR [16]	<i>RGB</i> \rightarrow <i>SMPL</i>	184.7	88.4	153.2	85.5	—	—
SPIN [20]	<i>RGB</i> \rightarrow <i>SMPL</i>	85.6	55.6	72.9	51.9	—	—
Pose2Mesh [6]	<i>Pose</i> \rightarrow <i>Coord</i>	64.9	<u>48.7</u>	67.9	49.9	4	6G
I2L-MeshNet [27]	<i>RGB</i> \rightarrow <i>Voxel</i>	82.6	59.8	55.7	41.7	25	4.6G
GraphCMR [21]	<i>RGB</i> \rightarrow <i>Coord</i>	148.0	104.6	78.3	59.5	—	—
CMR-PG (ours, ResNet18)	<i>RGB</i> \rightarrow <i>Coord</i>	<u>67.9</u>	49.9	59.0	44.7	59	1.5G
CMR-PG (ours, ResNet50)	<i>RGB</i> \rightarrow <i>Coord</i>	69.8	47.9	<u>57.3</u>	<u>42.6</u>	30	<u>1.9G</u>

Table 6. Results of root-relative mesh recovery on Human3.6M. In each column, the best number is highlighted in **bold**, and the second best number is highlighted with an underline.

the proposed CMR outperforms all the compared methods on 3D PCK by a large margin. With the root-relative mesh from CMR, we demonstrate that our adaptive 2D-1D registration consistently outperforms RootNet [26] as shown in Table 4.

Referring to CMR-PG’s PA-MPJPE and CS-MPJPE in Table 5, only an error of 6.9mm is incurred by relative pose with an error larger than 40mm from global translation, rotation, and scaling. Thus, compared with root-relative information, camera-space 3D reconstruction is more essential to improve the practicability of hand mesh recovery, and we advocate studying the camera-space problem.

In Figure 9, we compare our CMR with several pose estimation methods [1, 4, 9, 37, 41, 42, 44, 45] on the RHD dataset. Following the criterion of PA-MPJPE, the predicted 3D pose are processed with procrustes analysis. The AUC of CMR-PG and CMR-SG is 0.944 and 0.949, respectively, surpassing all the other methods. In addition, we directly use the FreiHAND models for RHD test, inducing comparable AUC of 0.872 and 0.852. Hence, the cross-domain generalization ability of CMR is verified.

In Table 6, we compare our method with several state-of-the-art approaches on body mesh recovery task using the Human3.6M and COCO datasets. Pose2Mesh [6] uses joint coordinates as the input and its performance downgrades slightly if the COCO dataset is added. As an RGB-to-voxel method, I2L-MeshNet performs better when both the Human3.6M and COCO datasets are used. By contrast, our CMR is more robust to different choices of training

data and we achieve comparable or even better numbers of MPJPE and PA-MPJPE. Moreover, our approach has faster inference speed and consumes less GPU memory.

In Figure 1, we show several qualitative evaluation results on the FreiHAND and Human3.6M datasets. As demonstrated in this figure, our CMR can deal with a variety of complex situations for camera-space mesh recovery, *e.g.*, challenging poses, object occlusion, and truncation.

5. Conclusions and Future Work

In this work, we aim to recover 3D hand and human mesh in camera-centered space and we present CMR to unify tasks of root-relative mesh recovery and root recovery. We first investigate 2D cues including 2D joint landmarks and silhouette for 3D tasks. Then, an aggregation method is proposed to collect effective 2D cues. Through aggregation of joint semantics, high-level semantic relations are explicitly captured, which is instructive for root-relative mesh recovery. We also explore 2D information for root recovery and design an adaptive 2D-1D registration to sufficiently leverage 2D pose and silhouette to estimate absolute camera-space information. Our CMR achieves state-of-the-art performance on both mesh and root recovery tasks when evaluated on FreiHAND, RHD, and Human3.6M datasets.

In the future, we plan to integrate 2D semantic information together with biomechanical relationship for more robust monocular 3D representation. We are also interested in extending our CMR with a human/hand detector in a top-down manner for multi-person tasks.

References

- [1] Seungryul Baek, Kwang In Kim, and Tae-Kyun Kim. Pushing the envelope for RGB-based dense 3D hand pose estimation via neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 8
- [2] Paul T Boggs and Jon W Tolle. Sequential quadratic programming. *Acta Numerica*, 4(1):1–51, 1995. 5
- [3] Adnane Boukhayma, Rodrigo de Bem, and Philip HS Torr. 3D hand shape and pose from images in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 7
- [4] Yujun Cai, Lihao Ge, Jianfei Cai, and Junsong Yuan. Weakly-supervised 3D hand pose estimation from monocular RGB images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 8
- [5] John Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, PAMI-8(6):679–698, 1986. 4
- [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 *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 5, 7, 8
- [7] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2016. 2
- [8] Bardia Doosti, Shujon Naha, Majid Mirbagheri, and David J Crandall. HOPE-Net: A graph-based model for hand-object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [9] Lihao Ge, Zhou Ren, Yuncheng Li, Zehao Xue, Yingying Wang, Jianfei Cai, and Junsong Yuan. 3D hand shape and pose estimation from a single RGB image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 5, 8
- [10] Shunwang Gong, Lei Chen, Michael Bronstein, and Stefanos Zafeiriou. SpiralNet++: A fast and highly efficient mesh convolution operator. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, 2019. 4
- [11] John C Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975. 6
- [12] Yana Hasson, Gul Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J Black, Ivan Laptev, and Cordelia Schmid. Learning joint reconstruction of hands and manipulated objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 7
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4, 6
- [14] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 36(7):1325–1339, 2013. 5
- [15] Umar Iqbal, Pavlo Molchanov, Thomas Breuel Juergen Gall, and Jan Kautz. Hand pose estimation via latent 2.5D heatmap regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [16] 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 (CVPR)*, 2018. 1, 2, 5, 8
- [17] 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 (CVPR)*, 2019. 1
- [18] Tong Ke and Stergios I Roumeliotis. An efficient algebraic solution to the perspective-three-point problem. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 7
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014. 6
- [20] 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 (ICCV)*, 2019. 1, 2, 5, 8
- [21] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 8
- [22] Dominik Kulon, Riza Alp Guler, Iasonas Kokkinos, Michael M Bronstein, and Stefanos Zafeiriou. Weakly-supervised mesh-convolutional hand reconstruction in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2, 6, 7
- [23] Dominik Kulon, Haoyang Wang, Riza Alp Güler, Michael Bronstein, and Stefanos Zafeiriou. Single image 3D hand reconstruction with mesh convolutions. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2019. 1
- [24] Isaak Lim, Alexander Dielen, Marcel Campen, and Leif Kobbelt. A simple approach to intrinsic correspondence learning on unstructured 3D meshes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2, 4
- [25] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics (TOG)*, 34(6):1–16, 2015. 2, 5
- [26] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Camera distance-aware top-down approach for 3D multi-person pose estimation from a single RGB image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 2, 7, 8

- [27] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 5, 7, 8
- [28] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. V2V-PoseNet: Voxel-to-voxel prediction network for accurate 3D hand and human pose estimation from a single depth map. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [29] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 3
- [30] 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 (CVPR)*, 2019. 1, 5
- [31] Georgios Pavlakos, Nikos Kolotouros, and Kostas Daniilidis. Texture-Pose: Supervising human mesh estimation with texture consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 1
- [32] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3D human pose. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 5
- [33] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3D human pose and shape from a single color image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 3
- [34] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. LCR-Net: Localization-classification-regression for human pose. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [35] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics (TOG)*, 36(6):245:1–245:17, 2017. 2, 3, 5
- [36] Adrian Spurr, Umar Iqbal, Pavlo Molchanov, Otmar Hilliges, and Jan Kautz. Weakly supervised 3D hand pose estimation via biomechanical constraints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1
- [37] Adrian Spurr, Jie Song, Seonwook Park, and Otmar Hilliges. Cross-modal deep variational hand pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 8
- [38] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 4
- [39] Gul Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. Bodynet: Volumetric inference of 3D human body shapes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 1, 2, 3
- [40] Lixin Yang, Jiasen Li, Wenqiang Xu, Yiqun Diao, and Cewu Lu. Bihand: Recovering hand mesh with multi-stage bisected hourglass networks. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2020. 1, 2
- [41] Linlin Yang, Shile Li, Dongheui Lee, and Angela Yao. Aligning latent spaces for 3D hand pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 8
- [42] Xiong Zhang, Qiang Li, Hong Mo, Wenbo Zhang, and Wen Zheng. End-to-end hand mesh recovery from a monocular RGB image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 2, 8
- [43] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3D human pose estimation in the wild: A weakly-supervised approach. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 1, 2
- [44] Yuxiao Zhou, Marc Habermann, Weipeng Xu, Ikhsanul Habibie, Christian Theobalt, and Feng Xu. Monocular real-time hand shape and motion capture using multi-modal data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 3, 8
- [45] Christian Zimmermann and Thomas Brox. Learning to estimate 3D hand pose from single RGB images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 5, 8
- [46] Christian Zimmermann, Duygu Ceylan, Jimei Yang, Bryan Russell, Max Argus, and Thomas Brox. FreiHAND: A dataset for markerless capture of hand pose and shape from single RGB images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 5, 7