

Hand Image Understanding via Deep Multi-Task Learning

Xiong Zhang^{1*}, Hongsheng Huang², Jianchao Tan³, Hongmin Xu⁴,
Cheng Yang¹, Guozhu Peng¹, Lei Wang¹, Ji Liu³

¹YY Live, Baidu Inc., ²Joyy Inc., ³AI Platform, Kwai Inc., ⁴OPPO Inc.,

Abstract

Analyzing and understanding hand information from multimedia materials like images or videos is important for many real world applications and remains active in research community. There are various works focusing on recovering hand information from single image, however, they usually solve a single task, for example, hand mask segmentation, 2D/3D hand pose estimation, or hand mesh reconstruction and perform not well in challenging scenarios. To further improve the performance of these tasks, we propose a novel Hand Image Understanding (HIU) framework to extract comprehensive information of the hand object from a single RGB image, by jointly considering the relationships between these tasks. To achieve this goal, a cascaded multi-task learning (MTL) backbone is designed to estimate the 2D heat maps, to learn the segmentation mask, and to generate the intermediate 3D information encoding, followed by a coarse-to-fine learning paradigm and a self-supervised learning strategy. Qualitative experiments demonstrate that our approach can recover reasonable mesh representations even in challenging situations. Quantitatively, our method significantly outperforms the state-of-the-art approaches on various widely-used datasets, in terms of diverse evaluation metrics <https://github.com/MandyMo/HIU-DMTL>.

1. Introduction

Hand image understanding (HIU) keeps active in both computer vision and graphics communities, aiming to recover the spatial configurations from RGB/depth images, including 2D/3D hand pose estimation, hand mask segmentation and hand mesh reconstruction, which have been employed in various domains [23, 25, 30, 33, 49, 52]. Recovering the spatial configurations is still challenging, due to the inherent depth and scale ambiguity, diverse appearance variation, and complicated articulations. While a bunch of existing works have considered markerless HIU, most of whom require depth camera [37, 62, 51, 41, 26, 15,

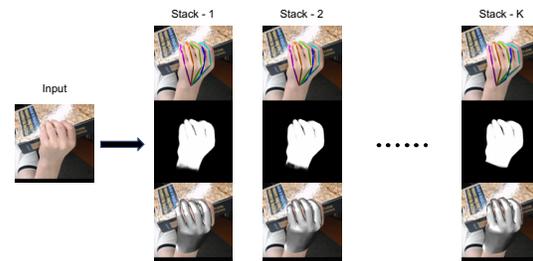


Figure 1. **Hand Image Understanding.** The figure illustrates the fundamental idea of this work. We derive the 2D hand pose, hand mask, hand mesh (including 3D hand pose) representation simultaneously from a monocular RGB image of the hand object in a coarse-to-fine manner.

67, 28, 32, 14, 22, 21] or synchronized multi-view images [3, 18, 46, 51, 59] to deal with the aforementioned challenges. As a consequence, most of those methods are impractical in real-world situations where only monocular RGB images are available.

For the monocular RGB scenario, the main obstacles reside in three-folds. Firstly, the lack of high-quality large-scale datasets with precise annotations. Existing datasets are either synthesized using software [73, 19, 35], or labeled in a semi-automated manner [74, 27], or collected in a controlled experimental environment [66, 68, 50]. Secondly, the incapability of the current datasets makes the trained models not generalize well to various wild images, especially under self occlusions and complex configurations, which may hinder its applications. Thirdly, contemporary approaches fail to exploit unlabeled images, which is more widely distributed than those with annotations.

The above obstacles motivate us to bring out two questions: *Can the existing multi-modality data be harnessed comprehensively to tackle the aforementioned difficulties?* *Can the tremendous wild images without labels be exploited well enough to favor the HIU?*

In this work, we demonstrate the answers to be **yes**, and the fundamental idea is illustrated in Figure 1. Specifically, an innovative multi-task learning (MTL) framework is designed to tackle the HIU problem, which follows the cascade coarse-to-fine design manner. Concretely, the frame-

*corresponding author, zhangxiong@yy.com

work consists of a backbone and several regressor heads corresponding to different tasks. The backbone aims to learn the various elementary representations from hand images, including 2D pose estimation, hand mask segmentation, and 3D part orientation field (POF) encoding. To reconstruct the whole hand mesh efficiently, we exploit the generative hand model MANO [44], and employ the regressor heads to regress MANO’s parameters based on the semantic features of the elementary tasks. To efficiently fuse beneficial semantic-features among various tasks, we conceive the task attention module (TAM) to aggregate the semantic features across individual tasks and derive compact high-level representations by removing redundant features. Note that the 3D hand joints can be achieved as a side output in MANO. With these designs together, one can obtain the 2D/3D hand pose, hand mask, and hand mesh from a RGB image simultaneously.

It is clear that the whole framework can be trained with generic supervised learning by leveraging existing multi-modality datasets. The self-supervised learning strategies can be adopted by leveraging the implicit relationship constraints maintained among the reasonable predictions from each task. For instance, the mask rendered from the hand mesh with proper camera parameters, shall match the one that is estimated by the backbone; the coordinates of the re-projected 2D hand pose shall be close to the integral of locations encoded in the heat-maps. The self-supervised learning make it possible to exploit enormous wild images, which can improve the accuracy of the framework, and enhance the generalization capability. Additionally, considering the absence of a large-scale hand dataset with well-labeled hand mask and 2D pose, we collect a high-quality dataset with manually labeled 2D hand pose and hand mask.

To summarize, our main contributions are as follows:

- We design an innovative cascade multi task learning (MTL) framework, dubbed HIU-DMTL, for hand image understanding, which can exploit existing multi-modality hand datasets effectively.
- The self-supervised learning (SSL) was firstly introduced to alleviate the insufficient data problem for HIU, and the effectiveness of which has been verified through extensive experiments comprehensively.
- We propose a simple while effective task attention module (TAM), targeting to aggregate semantic features across various tasks, which proves to be instrumental for MTL on the HIU tasks.
- Our HIU-DMTL framework outperforms contemporary hand mesh recovery approaches [19, 16, 4, 70, 72, 34], and demonstrates new state-of-the-art performances on widely used benchmarks [68, 73, 50, 74, 46], in terms of various evaluation metrics.

2. Related Work

Due to the extensive scope of related works, it is hard to summarize all of them comprehensively. We only discuss works that are strongly related with our framework settings.

3D Hand Pose Estimation. The pioneering work [73] firstly applies the deep learning technique to estimate the 3D hand pose from a single RGB image. Since then, 3D hand pose estimation draws a great deal of attention from the community [38, 35, 5, 24, 48, 65, 6, 55, 71, 47, 13]. Those methods either target to solve the perspective ambiguity problem by introducing geometry-constraints of hand articulation [73, 38, 35, 47], to investigate intricate learning strategies to achieve better performance [48, 24, 5, 65, 6, 55, 13], or to tackle the challenge of lacking sufficient high-quality hand data [73, 35, 65, 71]. Despite the significant progress have been achieved, estimating 3D hand pose from a monocular RGB image remains to be challenging, and the lacking of sufficient well-labeled data is still one of the main obstacles.

Hand Mesh Recovery. Besides the hand pose estimation, hand mesh recovery is another essential and active research topic. One line of works focus on reconstructing the hand mesh representation in a general situation. For instance, some works [4, 1, 70, 74, 27, 72, 63] capture the hand mesh representation by regressing the generative hand model MANO’s parameters, while others [16, 34, 8] instead turn to estimate the 3D mesh vertices directly in order to learn nonlinear hand shape variations. Another line of works [19, 2, 22] attempt to reconstruct the hand mesh in the hand-object interaction environment, in which the hand and object are reconstructed jointly by introducing interactive relationship constraints. Though the above approaches can recover reasonable hand mesh in experimental benchmarks, in practice, we find that approaches [4, 19, 34, 16, 70, 72, 10], with the publicly released code and pre-trained model, do not work well across different datasets or generalize well to the real-world situations.

Multi-Task Learning. Multi-task learning (MTL) is one methodology to improve the performance of tasks by utilizing the limited training samples and sharing the beneficial information among all tasks, which has been successfully adopted in many domains [11, 17, 45]. One broadly employed MTL method is hard-parameter sharing, which trains one shared encoder, followed by multiple task-specific decoders for different tasks. Some of them also further design a decoder-fusing module to distill information of different tasks to refine the final tasks’ prediction. Recent works [43, 39, 31, 56, 58, 12, 9] have applied such multi-task framework into pose estimation tasks and achieved state-of-the-art performance. Our approach follows general settings of previous multi-task learning methods, which makes use of an encoder, several certain task-specific decoders, and a features fusing module. Solving

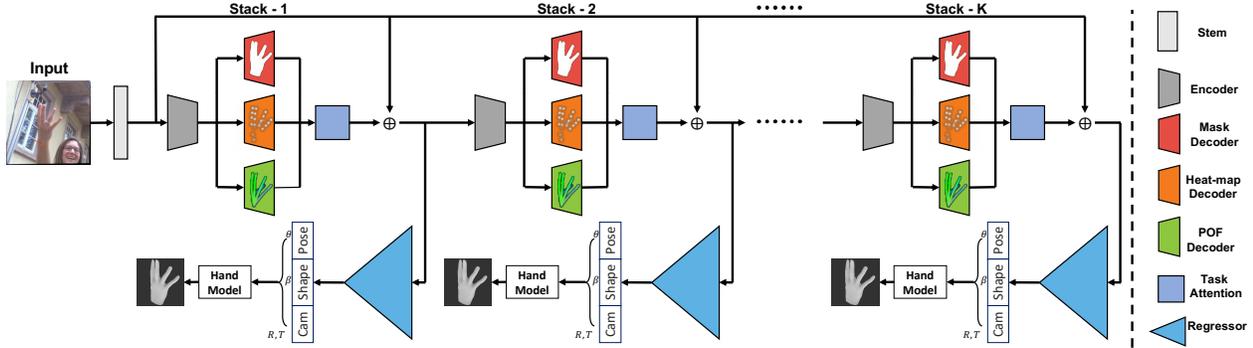


Figure 2. **Framework Architecture.** The whole pipeline follows the classical cascade coarse-to-fine design paradigm, which consists of two main components: (1) a novel multi-task learning backbone aiming to estimate heat-maps that encode the 2D hand pose, to learn the hand segmentation mask, and to generate the POF encoding that covers 3D information, (2) the regressor heads aiming to regress 3D parameters $\Theta = \{\theta, \beta, R, T\}$, based on the multi-task features, of the parametric hand model MANO [44] and the perspective camera.

HIU tasks under cascade multi-stack MTL framework still remains less explored in community, our work demonstrates such combination can achieve SOTA performance.

3. Framework

The primary goal of this work is to design a unified framework to provide comprehensive information of the hand object from a single RGB image, including 2D hand pose, 3D hand joints, hand mask, and hand mesh representation. To achieve this goal, a multi-stacked multi-branch backbone is designed to learn various elementary representations of hand object, then the mesh reconstruction is achieved by estimating parameters of the parametric model MANO [44] based on the elementary representations. The whole framework is illustrated in Figure 2.

3.1. Backbone

The backbone consists of a stem module and several MTL blocks sharing the same structure. Each MTL block consists of an encoder that is shared by all task-specified decoders in that block, several certain dedicated decoders that aim to tackle individual primary task, and a TAM that aggregating features across various tasks. In practice, we take the 2D hand pose, the hand mask, and the POF encoding as the learning target of the intermediate elementary tasks.

Stem Module. The stem module aims to extract low-level semantic features that are shared by succeeding modules. To keep the stem module as simple as possible while be capable of covering sufficient information, we implement the stem module with two 7×7 convolution layers with stride 2, which may quickly downsample the feature-maps to reduce the computation while obtaining wide enough receptive field.

Encoder. The encoder targets at generating high-level features to support various individual tasks, which takes the low-level features from the stem module and the aggregating

high-level representations from the preceding TAM together as inputs. For the first MTL block (denoted as stack 1 in Figure 2), since no preceding MTL block exists, the high-level semantic features remain $\mathbf{0}$ with a proper size.

Heat-Map Decoder. The purpose of heat-map decoder branch is to perform 2D hand pose estimation. Similar to recent methods [5, 16, 4, 70, 72, 34], we employ the 2D Gaussian-like heat-maps $\mathbf{H} \in \mathbb{R}^{K \times H \times W}$ to encode the 2D hand pose, where K indicates the number of joints and $\{H, W\}$ are the resolution size. Each key point corresponds to a heat-map. The pixel value at (x, y) is defined as $\exp\{-\frac{(x-\hat{x})^2+(y-\hat{y})^2}{2\sigma^2}\}$, where the (\hat{x}, \hat{y}) refers to the ground-truth location, corresponds to the confidence score of the key point locating in this 2D position (x, y) . With this heat-map, one may derive the k_{th} key point with $\arg \max_{\{h,w\}} \mathbf{H}_{(k,h,w)}$, or in a differentiable form,

$$\sum_{h=1}^H \sum_{w=1}^W \mathbf{H}_{(k,h,w)}(h, w) / \sum_{h=1}^H \sum_{w=1}^W \mathbf{H}_{(k,h,w)}. \quad (1)$$

Mask Decoder. The hand mask branch proves to be indispensable in HIU tasks, because the segmentation mask may further boost the performance of key point detection and vice versa [57, 69, 20]. More importantly, the hand mesh may deform to best fit joint locations and may ignore hand geometrical properties, leading to an unreasonable hand mesh representation when imposing supervision over 2D/3D hand pose only [4, 70, 72]. Fortunately, the mask branch not only exploits samples that are labeled with masks which lead to better performance, but also refines the hand mesh and camera parameters by penalizing the misalignment errors between the rendered mask and the estimated segmentation mask via self-supervised learning.

POF Decoder. To bridge the gap between the 2D semantic features and the 3D information, we introduce the part orientation field (POF) to encode the 3D orientation of the articulated structure in 2D image space. In practice,

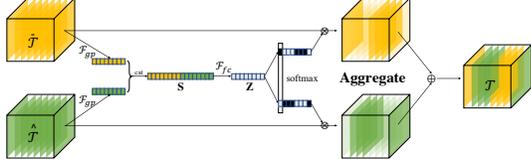


Figure 3. **Task Attention Module.** The figure illustrates the data flow of the task attention module, where \mathcal{F}_{gp} and \mathcal{F}_{fc} represents the global average pooling and the fully connected layers.

we introduce the standard hand skeleton hierarchy \mathbb{S} data structure, which consists of a set of ‘(parent, child)’ pairs. For a specific bone $(A, B) \in \mathbb{S}$, where A and B are two joints, and denote \overrightarrow{AB}_{3d} and \overrightarrow{AB}_{2d} as the normalized orientation from joint A to joint B in 3D and 2D domain respectively. Then, for the bone (A, B) , its POF, represented as $\mathbf{L}_{(a,b)} \in \mathbb{R}^{3 \times H \times W}$, encodes the 3D semantic information as a 3-channel feature-map, and the value of $\mathbf{L}_{a,b}$ at location x is defined as,

$$\mathbf{L}_{(a,b)}(x) = \begin{cases} (\overrightarrow{AB}_{2d}, \overrightarrow{AB}_{3d.z}) & x \in \text{bone} \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

We shall point out that, POF values are non-zero only for the pixels belonging to the current target bone part, and we employ a different but more appropriate definition, comparing with [61], because our POF encoding can exploit numerous wild training samples with only 2D labels.

Task Attention Module. The TAM aims to bring together semantic features across individual tasks, which can be formalized as a transformation, $\mathbb{R}^{N \times C \times H \times W} \mapsto \mathbb{R}^{C \times H \times W}$, where N refers to the number of tasks, and $\{C, H, W\}$ denotes the spatial resolution of the feature-maps. Figure 3 demonstrates the structure of TAM (with $N = 2$ for simplicity). Frankly speaking, our design is motivated by SKNet [29], but with several meaningful and reasonable modifications from original setting. The beginning element-wise addition step is replaced by a global average pooling and a feature concatenation. Such modest but necessary adjustments make the TAM more suitable to select critical semantic features among various tasks at the expense of additional but negligible computation burden.

3.2. Regressor Heads

The goal of the regressor heads is to reconstruct the hand mesh surfaces. To achieve this goal, we exploit the generative hand model MANO[44], and estimate the parameters of MANO that governs the mesh representation.

The mesh surface of MANO can be fully deformed and posed by the shape parameters $\beta \in \mathbb{R}^{10}$ and pose parameters $\theta \in \mathbb{R}^{15 \times 3}$, where β models hand shape and θ represents joints rotation. Given a pair of parameter $\{\beta, \theta\}$, the shape deformation $B_S(\beta) : \mathbb{R}^{10} \mapsto \mathbb{R}^{N \times 3}$ outputs the blended shape to characterize the identity subject and pose

deformation $B_P(\theta) : \mathbb{R}^{15 \times 3} \mapsto \mathbb{R}^{N \times 3}$ is applied to the mean template \bar{T} . Then we can obtain final mesh by rotating each bone part around joints $J(\beta)$ with the standard blend skinning function $W(\cdot)$:

$$M(\beta, \theta) = W(T(\beta, \theta), J(\beta), \theta, \mathcal{W}), \quad (3)$$

$$T(\beta, \theta) = \bar{T} + B_S(\beta) + B_P(\theta), \quad (4)$$

where \mathcal{W} refers to the skinning weights.

In the skinning procedure, not only the hand mesh can be obtained, but also the 3D hand joints can be achieved by rotating joints $J(\beta)$ with the pose parameters θ . Thus, an alternative way to estimate 2D hand pose is to project 3D hand joints with proper camera parameters. Particularly, in this work, we assume an ideal pinhole camera setting, with the projection matrix represented as,

$$\mathbf{Q} = \begin{pmatrix} f & 0 & p_0 \\ 0 & f & q_0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (5)$$

where f is the focal length and (p_0, q_0) is at the image center position, making f the only unknown variable.

Note that, instead of replicate K regressor heads with stand-alone training parameters, we make the K regressors share a global group of training parameters, similar to the strategy adopted in [60].

3.3. Training Objective

The differentiable property of the backbone, regressor heads and MANO makes our HIU-DMTL framework end-to-end trainable. The overall loss function can be divided into three categories, targeting at backbone training, the regressor heads optimizing, and self-supervised learning (*loss weights are ignored in the following discussions*).

Training the Backbone. The target of the backbone is to derive certain kinds of instrumental representations from the hand image, including the 2D heat-maps, hand mask, and the POF encoding. To train the backbone, the outputs of three branches in every MTL block are directly supervisory trained. Specifically, the training objective is defined as: $\mathcal{L}_{backbone} = \mathcal{L}_{hm} + \mathcal{L}_{pof} + \mathcal{L}_{seg}$, where \mathcal{L}_{hm} makes the estimated heat-maps close to the ground truth, the \mathcal{L}_{pof} as well. And the \mathcal{L}_{seg} makes use of the classical cross-entropy loss often used in semantic image segmentation task.

Training the Regressor Heads. The regression module aims to regress the camera parameters $\{R, T\}$ and mesh parameters $\{\beta, \theta\}$ of MANO. However, it is impracticable to obtain the ground-truth labels. Fortunately, the regressors can be trained with widely-available samples with 3D/2D annotations through weakly-supervised learning. Concretely, the loss function comprises three terms, $\mathcal{L}_{regressor} = \mathcal{L}_{3d} + \mathcal{L}_{2d} + \mathcal{L}_{mask}$, where \mathcal{L}_{3d} measures the orientation similarity between the estimated bones and the



Figure 4. **Qualitative Evaluation.** The first column exhibits the RGB inputs that cover widely used benchmarks [68, 73, 74, 46] (2nd - 5th rows) and wild situations (1st row). The following columns demonstrate the reconstruction results of ours, [34], [72], [70], [19], [4], and [16] respectively. To make a fair comparison, all methods are evaluated on the public available pre-trained checkpoints.

ground-truth bones in aforementioned skeleton hierarchy \mathbb{S} , while \mathcal{L}_{2d} and \mathcal{L}_{mask} refer to the re-projection loss of 2D joints and hand mask, as in [4, 70, 16].

Achieving Self-Supervised Learning. For a well reconstructed hand mesh, the projected mask shall be consistent with the silhouette, and such constraint has been exploited when training the regressor heads. However, more implicit constraints shall be maintained among the reasonable predictions. For instance, the rendered hand mask shall match the mask estimated by the backbone; the coordinates of the re-projected hand pose shall be close to the one inferred from the heat-maps of the backbone. Such consistencies enable us to exploit unlabeled hand images to achieve self-supervised learning. In practice, the re-projected differentiable hand mask is obtained via differentiable renderers contained in Pytorch3D [42], and Equation 1 is adopted to derive the differentiable 2D pose from the heat-maps.

4. Experiments

To demonstrate the effectiveness of our HIU-DMTL framework, we first present qualitative comparison over our recovered hand mesh with the results from recent works. We then quantitatively evaluate the superiority of HIU-DMTL on 2D/3D hand pose estimation, hand mask segmen-

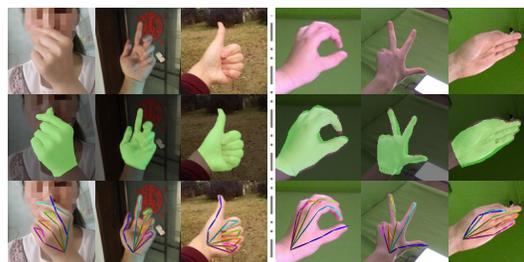


Figure 5. **Comparisons between HIU-Data and FreiHAND dataset.** The first three columns demonstrate samples from the HIU-Data, and the last three columns present examples in FreiHAND [74]. For each sample, the center cropped hand image, the segmentation mask, and the 2D hand pose are visualized.

tation, and hand mesh reconstruction tasks over several publicly available datasets. Finally, we perform several ablation studies to evaluate the importance of different design strategies and analyze the impact of the hyper-parameters. Due to the limited space, the implementation details of experiments are provided in the supplemental materials, please refer it for more details.

4.1. Experiment Settings

Datasets. We mainly leverage two kinds of datasets, *i.e.*, publicly available benchmarks and our newly annotated

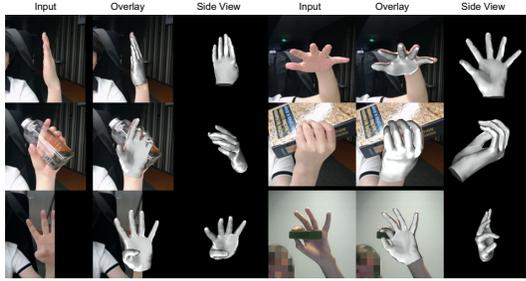


Figure 6. **Mesh Recovery Results.** We demonstrate the performance of our method under several typical challenging scenarios. For each sample, the overlay results between reconstructed hand and input image, and the side-view rendering results are presented.

Method	DeepLab v3 [7]	Fast-SCNN [40]	HIU-DMTL
mIoU \uparrow	96.6%	96.2%	97.5%

Table 1. **Hand Segmentation.** The table presents the mIoU of [7], [40], and our HIU-DMTL on HIU-Data.

dataset. For public datasets, we evaluate our method on CMU Panoptic Dataset (CMU) [46]; the Rendered Hand dataset (RHD) [73]; the Stereo Hand Pose Tracking Benchmark (STB) [68]; the FreiHAND [74]; and the Dexter Object (Dexter) [50] dataset. Since no accessible datasets contain high-quality hand masks, making it difficult to perform training directly on aforementioned HIU tasks.

HIU-Data. Regarding above concern, we manually collect 33,000 hand images in challenging scenarios to be our new dataset, namely HIU-data. For each sample, both the 2D hand pose and hand mask are manually annotated rather than generating approximate labels automatically as did in [74, 61, 66]. As Figure 5 demonstrates several samples with the corresponding labels, one may observe that our dataset achieves better annotation accuracy than the newest released automatically annotated FreiHAND [74] dataset.

4.2. Qualitative Evaluation

To visually evaluate the quality of the recovered hand mesh representation and the robustness of HIU-DMTL in various cases, we illustrate several representative samples in Figure 4 and Figure 6. As shown in Figure 4, the recovered hand mesh of HIU-DMTL have better overlays with the input image than the those reconstructed by other contemporary approaches [34, 72, 70, 19, 4, 16]. Furthermore, HIU-DMTL demonstrates superior robustness and generalization capability on in-the-wild unseen images, which is infeasible for previous existing methods. For example, Figure 6 shows that our method is robust enough to conduct hand meshes recovery in several challenging situations, such as exaggerated hand articulations, extremely unconstrained camera views, the existence of image truncations, heavy self-occlusions, and hand object interaction.

4.3. Quantitative Evaluation

We quantitatively evaluate the superiority of HIU-DMTL on 2D/3D hand pose estimation, hand mask segmentation, and hand mesh reconstruction tasks respectively.

Comparisons of 3D Hand Pose. To be consistent with [4, 70, 19, 16, 72], 3D PCK is adopted to evaluate the performance of 3D hand pose estimation. All the comparison methods are evaluated on STB, RHD, Dexter, and FreiHAND datasets, and Figure 7 reports the overall experimental results. On the **STB** dataset, the 3D PCK curves remain intertwined because the STB dataset is relatively small and lacks diversity. Our method achieves competitively results among newly launched approaches [4, 70, 16, 65, 1, 61, 71], which is reasonable when considering the saturated performance on this dataset [70, 72, 61, 64, 24]. On the **RHD** dataset, which is relatively complex and more diverse, our method outperforms all existing approaches [27, 70, 1, 16, 65, 71] and achieves state-of-the-art performance. On the **Dexter Object** dataset, our method dramatically outperforms existing works [2, 61, 1, 24, 35] and remains comparable with [50], which exploits additional depth information to calibrate the joint locations. On the **FreiHAND** benchmark, our approach outperforms contemporary methods [19, 4, 74] by a large margin, and obtains slightly better performance comparing with [34], while [34] exploits the additional ground truth hand mesh to fulfill mesh recovery.

Comparisons of Hand Mesh. We perform apple to apple comparison with these approaches [34, 72, 70, 16, 4, 19], containing accessible source code/checkpoint, which are focused on hand mesh recovery task under various evaluation metrics. Specifically, we take the IoU PCK of the re-projected mask, the 2D PCK of re-projected hand pose, and the 3D PCK of per-vertex reconstruction error as the evaluation metrics, and evaluate the aforementioned methods on the HIU-Data, CMU, and FreiHAND, as shown in Figure 8. Regarding the **2D PCK** of the re-projected 2D hand pose, our method outperforms [34, 72, 70, 16, 4, 19] by a wide margin. Regarding the **IOU PCK** of the re-projected hand mask, our approach retains a significantly higher PCK score, which is also already demonstrated in Figure 4 that our recovered hand mesh overlays with the input images better. Regarding the **3D PCK** of per-vertex reconstruction error, our framework obtains the minimum misalignment, comparing to [34, 74, 4, 19]. Note that the method [34] additionally exploits the ground truth hand mesh in FreiHAND dataset for training.

Comparisons of Hand Mask. For the hand mask segmentation, we compare our approach with general SOTA segmentation methods [7, 40]. Table 1 reveals that our method is more suitable than general sophisticated approaches for the task of hand mask segmentation. As the certain representative samples shown in Figure 9, our approach can estimate a precise mask even in the texture area

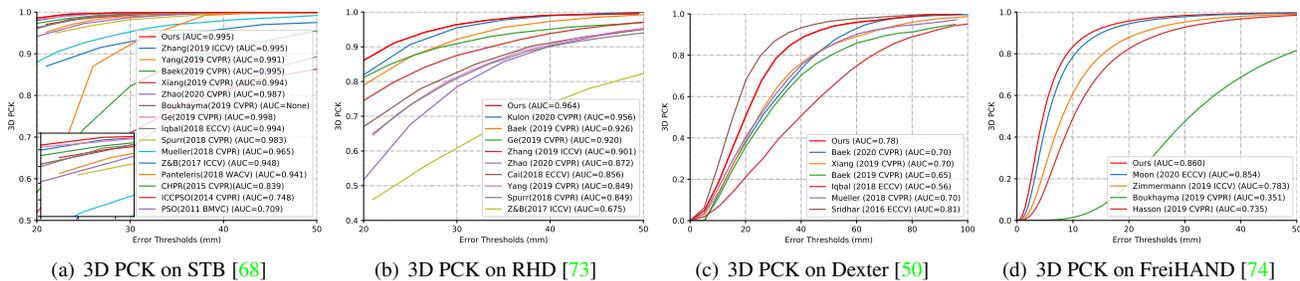


Figure 7. **Quantitative Evaluation.** The plots present the 3D PCK on the STB, RHD, Dexter, and FreiHAND datasets, respectively.

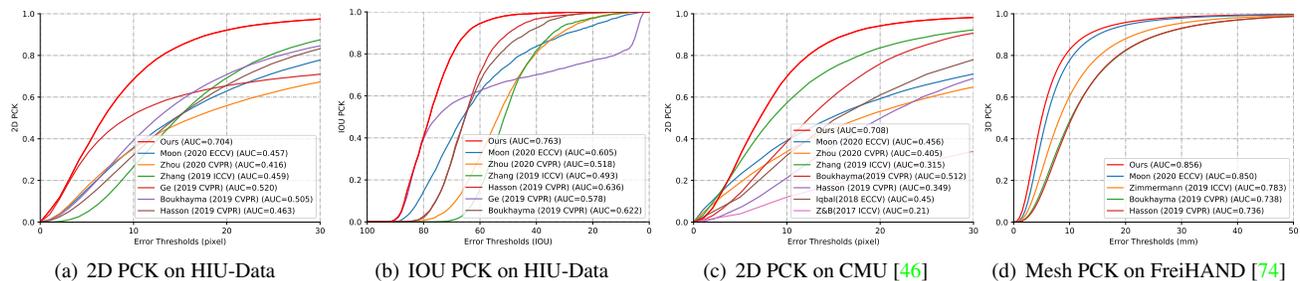


Figure 8. **Quantitative Evaluation.** The figures demonstrate the 2D PCK of the re-projected hand pose on HIU-data, IOU PCK of the re-projected hand mask on HIU-data, 2D PCK of the re-projected hand pose on CMU, 3D per-vertex PCK on FreiHAND, respectively.

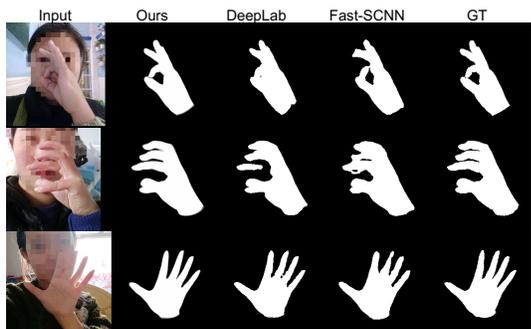


Figure 9. **Hand Segmentation.** The figure presents several results of mask segmentation task from our HIU-DMTL, [7], and [40].

with high-frequency, which is incapable in [7, 40]. Besides, [7, 40] may generate unreasonable hand masks that violates the biomedical appearance of hand. However, our approach can avoid such a dilemma by fully exploiting SSL and using priors that are comprised in hand mesh.

4.4. Ablation Study

The superior performance of HIU-DMTL mainly attributes to the cascaded design (CD) paradigm, the multi-task learning (MTL) setup, the task-attention module (TAM), and the self-supervised learning (SSL) strategy. In this section, we perform ablation studies to better understand the importance of these different design choices. Also, the ablation studies deeply analyze the impact of some parameters related with network structure, such as the num-

	2D Pose	2D Pose [†]	Hand Mask	Hand Mask [†]	3D Pose	Hand Mesh
4-Stack	0.866	0.704	0.974	0.770	0.860	0.856
2-Stack	0.857	0.701	0.969	0.765	0.857	0.853
1-Stack	0.837	0.686	0.948	0.752	0.852	0.842

Table 2. **Ablation of CD Paradigm.** The table presents the ablation results across different stacking arrangements under various evaluation metrics, where [†] indicates inferring the pose/mask by projecting the 3D pose/mesh with proper camera parameters. The 3D hand pose/mesh are quantified on FreiHAND benchmark, while the 2D pose/mask are evaluated on the HIU-Data, since the quality of masks in FreiHAND benchmark is not good enough.

ber of the network stack in our framework.

Cascaded Design Paradigm. The CD paradigm has been widely adopted in 2D pose estimation task [36, 60, 53, 54], while **the multi-branched cascaded MTL design for HIU tasks remains less explored.** Similar to [36], we investigate the impact of the number of MTL blocks. Specifically, we construct three models, which contain 1, 2, and 4 MTL blocks respectively. Meanwhile, the three models are designed to have similar FLOPs for a fair comparison. As shown in Table 2, changing from 1 stack to 4 stacks, the performance of 2D hand pose and hand mask tasks can improve substantially. We also investigate how performance improves, according to the increases of MTL block. Table 3 presents the performance of each intermediate stack in a 8-stack HIU-DMTL framework, where we take same evaluation metrics as in Table 2. We can conclude that from stack 1 to stack 4, the performance increases rapidly, and then this growth trend gradually flattened in later stacks.

Stack	2D Pose	2D Pose [†]	Hand Mask	Hand Mask [†]	3D Pose	Hand Mesh
Stack-1	0.818	0.661	0.915	0.712	0.824	0.819
Stack-2	0.855	0.685	0.954	0.744	0.845	0.834
Stack-3	0.865	0.700	0.968	0.764	0.854	0.847
Stack-4	0.867	0.704	0.975	0.769	0.857	0.853
Stack-5	0.867	0.706	0.975	0.771	0.859	0.855
Stack-6	0.868	0.706	0.976	0.773	0.860	0.857
Stack-7	0.869	0.705	0.976	0.771	0.861	0.857
Stack-8	0.870	0.707	0.977	0.773	0.861	0.859

Table 3. **Ablation of Intermediate Stacks.** The table presents the performance of the intermediate stacks under various evaluation metrics, where [†] shares similar definition as in Table 2.

2D Pose	Hand Mask	Hand Mesh	2D Pose	Hand Mask	3D Pose	Hand Mesh
✓	✗	✗	0.752	-	-	-
✗	✓	✗	-	0.863	-	-
✗	✗	✓	-	-	0.802	0.801
✓	✓	✗	0.791	0.875	-	-
✗	✓	✓	-	0.891	0.816	0.813
✓	✗	✓	0.795	-	0.812	0.808
✓	✓	✓	0.802	0.907	0.821	0.817

Table 4. **Ablation of MTL Setup.** The table presents the comparisons of different task combinations on FreiHAND dataset.

Multi-Task Learning Setup. Table 4 reports the influences of jointly learning the hand mask segmentation, 2D hand pose estimation, and hand mesh recovery tasks. One may observe two conclusions: first, when jointly training of any two tasks together, each task’s performance is better than training any stand-alone task; second, when jointly learning the whole three tasks, each task’s performance is better than when training in any other configurations. This is because the aforementioned relationship constraints among elementary tasks can improve performance with each other. For instance, the priors comprised in the hand mesh can correct the unreasonable estimations of hand mask/pose that violate the hand’s biomedical appearance.

Task Attention Module. To investigate the effectiveness of the TAM, we conduct ablation experiments on 2D hand pose estimation, hand mask segmentation tasks over the HIU-Data, as well as 3D hand pose regression, hand mesh recovery tasks on the FreiHAND benchmark. For the baseline model without the TAM, we employ a concatenation followed by a 3×3 convolution to aggregate the feature-maps from individual tasks. As shown in Table 5, we compare the baseline with and without TAM module in 1-stack and 4-stacks setting. As we can see, the TAM significantly improves the performance of individual tasks in both single-stack mode and multi-stacks mode, which means TAM is not only beneficial to the aggregation of feature representations for regressor branch but also helpful to the transition of feature representation between cascaded stacks.

Self-Supervised Learning Strategy. To verify the effectiveness of the SSL paradigm, we take the AUC of 2D hand pose and hand mask as the evaluation metric and compare the performance of framework with and without SSL. As plotted in Figure 10 (left), the 2D hand pose retrieved from the backbone and 2D hand pose that derived from pro-

	2D Pose	2D Pose [†]	Hand Mask	Hand Mask [†]	3D Pose	Hand Mesh
4-stack	0.866	0.704	0.974	0.770	0.860	0.856
4-stack [§]	0.793	0.662	0.957	0.739	0.813	0.808
1-stack	0.837	0.686	0.948	0.752	0.852	0.842
1-stack [§]	0.821	0.648	0.935	0.739	0.841	0.826

Table 5. **Ablation of TAM.** The table presents the ablation results of the TAM under various evaluation metrics, where [§] denotes to not employ the TAM, and [†] shares similar definition as in Table 2.

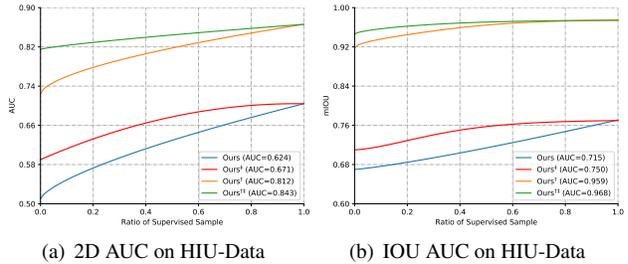


Figure 10. **Ablation of SSL Strategy.** The plots present the comparison of with and without applying self-supervised learning strategy over the various portions of samples that are used for supervised learning, [‡] denotes to apply self-supervised training over the remaining samples that are not chosen to do supervised learning, [†] shares the similar definition as in Table 2.

jecting the 3D hand joints are two factors evaluated in this comparison. One may observe that: when no samples are used for supervised training ($x=0$), the self-learning strategy can improve baseline performance dramatically. As all samples are under supervised training mode ($x=1$), the vacancy of self-supervised learning results in no influence over final performance at all. Similarly, the right side of Figure 10 plots the comparisons of hand mask, which also demonstrates that SSL strategy can significantly improve the performance of hand mask. The above observations also explain the generalization capability of HIU-DMTL on unseen in the wild images (Figure 6) from another aspect.

5. Discussion

In this work, we proposed a novel cascaded multi-task learning (MTL) framework, namely **HIU-DMTL**, to tackle the hand image understanding tasks in the coarse-to-fine manner. Our approach can efficiently exploit existing multi-modality datasets due to the MTL tactics and can harness enormous in-the-wild images through the self-supervised learning strategy, making it more practicable than previous methods. The performance of our **HIU-DMTL** framework have been verified by extensive quantitative and qualitative experiments. Besides, most of the components are well explored and discussed in various ablation studies, which make it straightforward to explain why our framework can achieve good performance over HIU tasks. For further work, we hope the core designs of our HIU-DMTL framework can be adopted generally to other research domains.

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)*, pages 1067–1076, 2019. 2, 6
- [2] Seungryul Baek, Kwang In Kim, and Tae-Kyun Kim. Weakly-supervised domain adaptation via gan and mesh model for estimating 3d hand poses interacting objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6121–6131, 2020. 2, 6
- [3] Luca Ballan, Aparna Taneja, Jürgen Gall, Luc Van Gool, and Marc Pollefeys. Motion capture of hands in action using discriminative salient points. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 640–653, 2012. 1
- [4] 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, 3, 5, 6
- [5] 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)*, pages 666–682, 2018. 2, 3
- [6] Yujun Cai, Lihao Ge, Jun Liu, Jianfei Cai, Tat-Jen Cham, Junsong Yuan, and Nadia Magnenat Thalmann. Exploiting spatial-temporal relationships for 3d pose estimation via graph convolutional networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2272–2281, 2019. 2
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 40(4):834–848, 2017. 6, 7
- [8] Xingyu Chen, Yufeng Liu, Chongyang Ma, Jianlong Chang, Huayan Wang, Tian Chen, Xiaoyan Guo, Pengfei Wan, and Wen Zheng. Camera-space hand mesh recovery via semantic aggregation and adaptive 2d-1d registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13274–13283, 2021. 2
- [9] Yifei Chen, Haoyu Ma, Deying Kong, Xiangyi Yan, Jianbao Wu, Wei Fan, and Xiaohui Xie. Nonparametric structure regularization machine for 2d hand pose estimation. In *The IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 381–390, 2020. 2
- [10] Yujin Chen, Zhigang Tu, Di Kang, Linchao Bao, Ying Zhang, Xuefei Zhe, Ruizhi Chen, and Junsong Yuan. Model-based 3d hand reconstruction via self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10451–10460, 2021. 2
- [11] Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the ACM International Conference on Machine Learning (ICML)*, pages 160–167, 2008. 2
- [12] Kuo Du, Xiangbo Lin, Yi Sun, and Xiaohong Ma. Cross-infonet: Multi-task information sharing based hand pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9896–9905, 2019. 2
- [13] Zhipeng Fan and YW Jun Liu. Adaptive computationally efficient network for monocular 3d hand pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 127–144, 2020. 2
- [14] Linpu Fang, Xingyan Liu, Li Liu, Hang Xu, and Wenxiong Kang. Jgr-p2o: Joint graph reasoning based pixel-to-offset prediction network for 3d hand pose estimation from a single depth image. pages 120–137, 2020. 1
- [15] Lihao Ge, Hui Liang, Junsong Yuan, and Daniel Thalmann. Robust 3d hand pose estimation in single depth images: from single-view cnn to multi-view cnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3593–3601, 2016. 1
- [16] 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)*, pages 10833–10842, 2019. 2, 3, 5, 6
- [17] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1440–1448, 2015. 2
- [18] Francisco Gomez-Donoso, Sergio Orts-Escolano, and Miguel Cazorla. Large-scale multiview 3d hand pose dataset. *Image and Vision Computing*, 81:25–33, 2019. 1
- [19] 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)*, pages 11807–11816, 2019. 1, 2, 5, 6
- [20] Kaiping He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2961–2969, 2017. 3
- [21] Lin Huang, Jianchao Tan, Ji Liu, and Junsong Yuan. Hand-transformer: Non-autoregressive structured modeling for 3d hand pose estimation. pages 17–33, 2020. 1
- [22] Lin Huang, Jianchao Tan, Jingjing Meng, Ji Liu, and Junsong Yuan. Hot-net: Non-autoregressive transformer for 3d hand-object pose estimation. In *Proceedings of the International Conference on Multimedia (MM)*, pages 3136–3145, 2020. 1, 2
- [23] Wolfgang Hürst and Casper Van Wezel. Gesture-based interaction via finger tracking for mobile augmented reality. *Multimedia Tools and Applications*, 62(1):233–258, 2013. 1
- [24] Umar Iqbal, Pavlo Molchanov, Thomas Breuel, Jürgen Gall, and Jan Kautz. Hand pose estimation via latent 2.5 d heatmap regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 118–134, 2018. 2, 6
- [25] Youngkyoon Jang, Seung-Tak Noh, Hyung Jin Chang, Tae-Kyun Kim, and Woontack Woo. 3d finger cape: Clicking

- action and position estimation under self-occlusions in ego-centric viewpoint. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 21(4):501–510, 2015. [1](#)
- [26] Sameh Khamis, Jonathan Taylor, Jamie Shotton, Cem Keskin, Shahram Izadi, and Andrew Fitzgibbon. Learning an efficient model of hand shape variation from depth images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2540–2548, 2015. [1](#)
- [27] 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)*, pages 4990–5000, 2020. [1](#), [2](#), [6](#)
- [28] Shile Li and Dongheui Lee. Point-to-pose voting based hand pose estimation using residual permutation equivariant layer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11927–11936, 2019. [1](#)
- [29] Xiang Li, Wenhai Wang, Xiaolin Hu, and Jian Yang. Selective kernel networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 510–519, 2019. [4](#)
- [30] R Liang and M Ouhyoung. A real-time continuous gesture recognition system for sign language. In *Proceedings Third IEEE International Conference on Automatic Face and Gesture Recognition*, pages 558–558, 1998. [1](#)
- [31] Mude Lin, Liang Lin, Xiaodan Liang, Keze Wang, and Hui Cheng. Recurrent 3d pose sequence machines. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 810–819, 2017. [2](#)
- [32] Jameel Malik, Ibrahim Abdelaziz, Ahmed Elhayek, Soshi Shimada, Sk Aziz Ali, Vladislav Golyanik, Christian Theobalt, and Didier Stricker. Handvoxnnet: Deep voxel-based network for 3d hand shape and pose estimation from a single depth map. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7113–7122, 2020. [1](#)
- [33] Anders Markussen, Mikkel Rønne Jakobsen, and Kasper Hornbæk. Vulture: a mid-air word-gesture keyboard. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1073–1082, 2014. [1](#)
- [34] 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. 2020. [2](#), [3](#), [5](#), [6](#)
- [35] Franziska Mueller, Florian Bernard, Oleksandr Sotnychenko, Dushyant Mehta, Srinath Sridhar, Dan Casas, and Christian Theobalt. Gnerated hands for real-time 3d hand tracking from monocular rgb. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 49–59, 2018. [1](#), [2](#), [6](#)
- [36] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 483–499, 2016. [7](#)
- [37] Iason Oikonomidis, Nikolaos Kyriazis, and Antonis A Argyros. Efficient model-based 3d tracking of hand articulations using kinect. In *The British Machine Vision Conference (BMVC)*, volume 1, page 3, 2011. [1](#)
- [38] Paschalis Panteleris, Iason Oikonomidis, and Antonis Argyros. Using a single rgb frame for real time 3d hand pose estimation in the wild. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 436–445, 2018. [2](#)
- [39] Alin-Ionut Popa, Mihai Zanfir, and Cristian Sminchisescu. Deep multitask architecture for integrated 2d and 3d human sensing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6289–6298, 2017. [2](#)
- [40] Rudra PK Poudel, Stephan Liwicki, and Roberto Cipolla. Fast-scnn: Fast semantic segmentation network. In *The British Machine Vision Conference (BMVC)*, 2019. [6](#), [7](#)
- [41] Chen Qian, Xiao Sun, Yichen Wei, Xiaoou Tang, and Jian Sun. Realtime and robust hand tracking from depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1106–1113, 2014. [1](#)
- [42] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501*, 2020. [5](#)
- [43] 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)*, pages 3433–3441, 2017. [2](#)
- [44] 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, 2017. [2](#), [3](#), [4](#)
- [45] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. [2](#)
- [46] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1145–1153, 2017. [1](#), [2](#), [5](#), [6](#), [7](#)
- [47] Adrian Spurr, Umar Iqbal, Pavlo Molchanov, Otmar Hilliges, and Jan Kautz. Weakly supervised 3d hand pose estimation via biomechanical constraints. pages 211–288, 2020. [2](#)
- [48] 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)*, pages 89–98, 2018. [2](#)
- [49] Srinath Sridhar, Anna Maria Feit, Christian Theobalt, and Antti Oulasvirta. Investigating the dexterity of multi-finger input for mid-air text entry. In *Proceedings of the Annual ACM Conference on Human Factors in Computing Systems*, pages 3643–3652, 2015. [1](#)
- [50] Srinath Sridhar, Franziska Mueller, Michael Zollhöfer, Dan Casas, Antti Oulasvirta, and Christian Theobalt. Real-time joint tracking of a hand manipulating an object from rgb-d

- input. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 294–310, 2016. 1, 2, 6, 7
- [51] Srinath Sridhar, Antti Oulasvirta, and Christian Theobalt. Interactive markerless articulated hand motion tracking using rgb and depth data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2456–2463, 2013. 1
- [52] Thad Starner, Joshua Weaver, and Alex Pentland. Real-time american sign language recognition using desk and wearable computer based video. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 20(12):1371–1375, 1998. 1
- [53] Xiao Sun, Yichen Wei, Shuang Liang, Xiaoou Tang, and Jian Sun. Cascaded hand pose regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 824–832, 2015. 7
- [54] Danhang Tang, Jonathan Taylor, Pushmeet Kohli, Cem Keskin, Tae-Kyun Kim, and Jamie Shotton. Opening the black box: Hierarchical sampling optimization for estimating human hand pose. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3325–3333, 2015. 7
- [55] Thomas Theodoridis, Theocharis Chatzis, Vassilios Sotlachidis, Kosmas Dimitropoulos, and Petros Daras. Cross-modal variational alignment of latent spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 960–961, 2020. 2
- [56] Denis Tome, Chris Russell, and Lourdes Agapito. Lifting from the deep: Convolutional 3d pose estimation from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2500–2509, 2017. 2
- [57] Subarna Tripathi, Maxwell Collins, Matthew Brown, and Serge Belongie. Pose2instance: Harnessing keypoints for person instance segmentation. *arXiv preprint arXiv:1704.01152*, 2017. 3
- [58] Chengde Wan, Thomas Probst, Luc Van Gool, and Angela Yao. Dense 3d regression for hand pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5147–5156, 2018. 2
- [59] Robert Wang, Sylvain Paris, and Jovan Popović. 6d hands: markerless hand-tracking for computer aided design. In *Proceedings of the Annual ACM Symposium on User Interface Software and Technology*, pages 549–558, 2011. 1
- [60] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4732, 2016. 4, 7
- [61] Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. Monocular total capture: Posing face, body, and hands in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10965–10974, 2019. 4, 6
- [62] Chi Xu and Li Cheng. Efficient hand pose estimation from a single depth image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3456–3462, 2013. 1
- [63] John Yang, Hyung Jin Chang, Seungeui Lee, and Nojun Kwak. Seqhand: Rgb-sequence-based 3d hand pose and shape estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 122–139, 2020. 2
- [64] 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)*, pages 2335–2343, 2019. 6
- [65] Linlin Yang and Angela Yao. Disentangling latent hands for image synthesis and pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9877–9886, 2019. 2, 6
- [66] Zhixuan Yu, Jae Shin Yoon, In Kyu Lee, Prashanth Venkatesh, Jaesik Park, Jihun Yu, and Hyun Soo Park. Humbi: A large multiview dataset of human body expressions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2990–3000, 2020. 1, 6
- [67] Shanxin Yuan, Guillermo Garcia-Hernando, Björn Stenger, Gyeongsik Moon, Ju Yong Chang, Kyoung Mu Lee, Pavlo Molchanov, Jan Kautz, Sina Honari, Lihua Ge, et al. Depth-based 3d hand pose estimation: From current achievements to future goals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2636–2645, 2018. 1
- [68] Jiawei Zhang, Jianbo Jiao, Mingliang Chen, Liangqiong Qu, Xiaobin Xu, and Qingxiong Yang. A hand pose tracking benchmark from stereo matching. In *International Conference on Image Processing (ICIP)*, pages 982–986, 2017. 1, 2, 5, 6, 7
- [69] Song-Hai Zhang, Ruilong Li, Xin Dong, Paul Rosin, Zixi Cai, Xi Han, Dingcheng Yang, Haozhi Huang, and Shi-Min Hu. Pose2seg: Detection free human instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 889–898, 2019. 3
- [70] 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)*, pages 2354–2364, 2019. 2, 3, 5, 6
- [71] Long Zhao, Xi Peng, Yuxiao Chen, Mubbasir Kapadia, and Dimitris N Metaxas. Knowledge as priors: Cross-modal knowledge generalization for datasets without superior knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6528–6537, 2020. 2, 6
- [72] 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)*, pages 5346–5355, 2020. 2, 3, 5, 6
- [73] 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)*, pages 4903–4911, 2017. 1, 2, 5, 6, 7

- [74] 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 International Conference on Computer Vision*, pages 813–822, 2019. [1](#), [2](#), [5](#), [6](#), [7](#)