

A2J-Transformer: Anchor-to-Joint Transformer Network for 3D Interacting Hand Pose Estimation from a Single RGB Image

Changlong Jiang¹, Yang Xiao^{1†}, Cunlin Wu¹, Mingyang Zhang², Jinghong Zheng¹, Zhiguo Cao¹,
 and Joey Tianyi Zhou^{3,4}

¹Key Laboratory of Image Processing and Intelligent Control, Ministry of Education, School of Artificial Intelligence and Automation, Huazhong University of Science and Technology, Wuhan 430074, China

²Alibaba Group

³Centre for Frontier AI Research, Agency for Science, Technology and Research (A*STAR), Singapore

⁴Institute of High Performance Computing, Agency for Science, Technology and Research (A*STAR), Singapore

changlongj, Yang.Xiao, cunlin.wu@hust.edu.cn, changhai.zmy@alibaba-inc.com,

deepzheng, zgcao@hust.edu.cn, zhouty@cfar.a-star.edu.sg

Abstract

3D interacting hand pose estimation from a single RGB image is a challenging task, due to serious self-occlusion and inter-occlusion towards hands, confusing similar appearance patterns between 2 hands, ill-posed joint position mapping from 2D to 3D, etc.. To address these, we propose to extend A2J-the state-of-the-art depth-based 3D single hand pose estimation method-to RGB domain under interacting hand condition. Our key idea is to equip A2J with strong local-global aware ability to well capture interacting hands' local fine details and global articulated clues among joints jointly. To this end, A2J is evolved under Transformer's non-local encoding-decoding framework to build A2J-Transformer. It holds 3 main advantages over A2J. First, self-attention across local anchor points is built to make them global spatial context aware to better capture joints' articulation clues for resisting occlusion. Secondly, each anchor point is regarded as learnable query with adaptive feature learning for facilitating pattern fitting capacity, instead of having the same local representation with the others. Last but not least, anchor point locates in 3D space instead of 2D as in A2J, to leverage 3D pose prediction. Experiments on challenging InterHand 2.6M demonstrate that, A2J-Transformer can achieve state-of-the-art model-free performance (3.38mm MPJPE advancement in 2-hand case) and can also be applied to depth domain with strong generalization. The code is available at <https://github.com/ChanglongJiangGit/A2J-Transformer>.

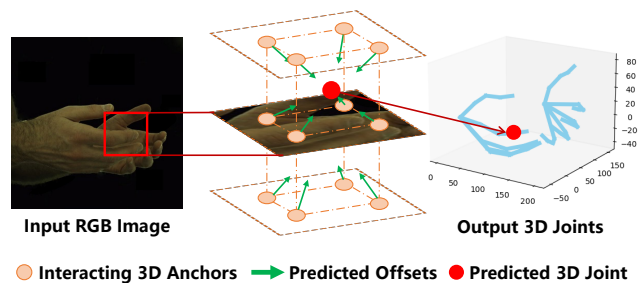


Figure 1. The main idea of A2J-Transformer. 3D anchors are uniformly set and act as local regressors to predict each hand joint. Meanwhile, they are also used as queries, and the interaction among them is established to acquire global context.

1. Introduction

3D interacting hand pose estimation from a single RGB image can be widely applied to the fields of virtual reality, augmented reality, human-computer interaction, etc.. [32, 34, 37]. Although the paid efforts, it still remains as a challenging research task due to the main issues of serious self-occlusion and inter-occlusion towards hands [7, 12, 16, 22, 27], confusing similar appearance patterns between 2 hands [12, 19, 27], and the ill-posed characteristics of estimating 3D hand pose via monocular RGB image [7, 16, 28].

The existing methods can be generally categorized into model-based [1, 2, 21, 29, 30, 35, 39, 41, 48] and model-free [5, 7, 12, 17, 19, 22, 26, 27, 29, 43] groups. Due to model's strong prior knowledge on hands, the former paradigm is overall of more promising performance. However, model-based methods generally require complex personalized model calibration, which is sensitive to initialization and susceptible

[†]Yang Xiao is corresponding author (Yang.Xiao@hust.edu.cn).

to trap in local minima [11, 12]. This is actually not preferred by the practical applications. Accordingly, we focus on model-free manner in regression way. The key idea is that, *for effective 3D interacting hand pose estimation the predictor should be well aware of joints' local fine details and global articulated context simultaneously* to resist occlusion and confusing appearance pattern issues. To this end, we propose to extend the SOTA depth-based single hand 3D pose estimation method A2J [43] to 3D interacting hand pose estimation task from a single RGB image.

Although A2J's superiority with ensemble local regression, intuitively applying it to our task cannot ensure promising performance, since it generally suffers from 3 main defects as below. First, the local anchor points for predicting offsets between them and joints lack interaction among each other. This leads to the fact that, joints' global articulated clues cannot be well captured to resist occlusion. Secondly, the anchor points within the certain spatial range share the same single-scale local convolution feature, which essentially limits the discrimination capacity on confusing visual patterns towards the interacting hands. Last, anchor points locate within 2D plane, which is not optimal for alleviating the ill-posed 2D to 3D lifting problem with single RGB image. To address these, we propose to *extend A2J under Transformer's non-local encoding-decoding framework to build A2J-Transformer, with anchor point-wise adaptive multi-scale feature learning and 3D anchor point setup*.

Particularly, the anchor point within A2J is evolved as the learnable query under Transformer framework. Each query will predict its position offsets to all the joints of the 2 hands. Joint's position is finally estimated via fusing the prediction results from all queries in a linear weighting way. That is to say, joint's position is determined by all the queries located over the whole image of global spatial perspective. Meanwhile, the setting query number is flexible, which is not strictly constrained by joint number as in [12]. Thanks to Transformer's non-local self-attention mechanism [40], during feature encoding stage the queries can interact with each other to capture joints' global articulated clues, which is essentially beneficial for resisting self-occlusion and inter-occlusion. Concerning the specific query, adaptive local feature learning will be conducted to extract query-wise multi-scale convolutional feature based Resnet-50 [14]. Compared with A2J's feature sharing strategy among the neighboring anchor points, our proposition can essentially facilitate query's pattern fitting capacity both for accurate joint localization and joint's hand identity verification. In summary, each query will be of strong local-global spatial awareness ability to better fit interacting hand appearance pattern. Meanwhile to facilitate RGB-based 2D to 3D hand pose lifting problem, the queries will be set within the 3D space instead of 2D counterpart as in A2J [43]. In this way, each query can directly predict its

3D position offset between the joints, which cannot be acquired by A2J. Overall, A2J-Transformer's main research idea is shown in Fig. 1.

Compared with the most recently proposed model-free method [12] that also addresses 3D interacting hand pose estimation using Transformer, our proposition still takes some essential advantages. First, joint-like keypoint detection is not required. Secondly, query number is not strictly constrained to be equal to joint number to facilitate pattern fitting capacity. Thirdly, our query locates within 3D space instead of 2D counterpart.

The experiments on the challenging Interhand 2.6M [29] dataset verify that, our approach can achieve the state-of-the-art model-free performance (3.38mm MPJPE advancement in 2-hand case) for 3D interacting hand pose estimation from a single RGB image. And, it significantly outperforms A2J by large margins (i.e., over 5mm on MPJPE). In addition, experiments on HANDS2017 dataset [46] demonstrate that A2J-Transformer can also be applied to depth domain with promising performance.

Overall, the main contributions of this paper include:

- For the first time, we extend A2J from depth domain to RGB domain to address 3D interacting hand pose estimation from a single RGB image with promising performance;
- A2J's anchor point is evolved with Transformer's non-local self-attention mechanism with adaptive local feature learning, to make it be aware of joints' local fine details and global articulated context simultaneously;
- Anchor point is proposed to locate within 3D space to facilitate ill-posed 2D to 3D hand pose lifting problem based on monocular RGB information.

2. Related Works

Many methods have been proposed for 3D hand pose estimation from either RGB images or depth maps. At the same time, these methods can also be divided into single hand pose estimation and interacting hand pose estimation methods based on the number of input hands. Here we categorize all these methods into model-based and model-free groups, and mainly analysis works that estimate interacting 3D hand pose from RGB images. Besides, we discuss the usage of Transformer architectures in 3D hand pose estimation field as they are highly relevant to our work.

Model-based approach. Considering that model-based methods [1, 21, 29, 30, 33, 35, 39, 41, 48] can provide strong prior knowledge, model-based 3D hand pose estimation methods could achieve relatively better results by fitting hand models. Early methods [1, 30, 39] for model-based 3D hand pose estimation used complex optimization methods to fit their proposed parameter models. However, due to the lack of a unified model paradigm, the development of model-based methods was somewhat limited at that time. After the introduction of the compatible 3D

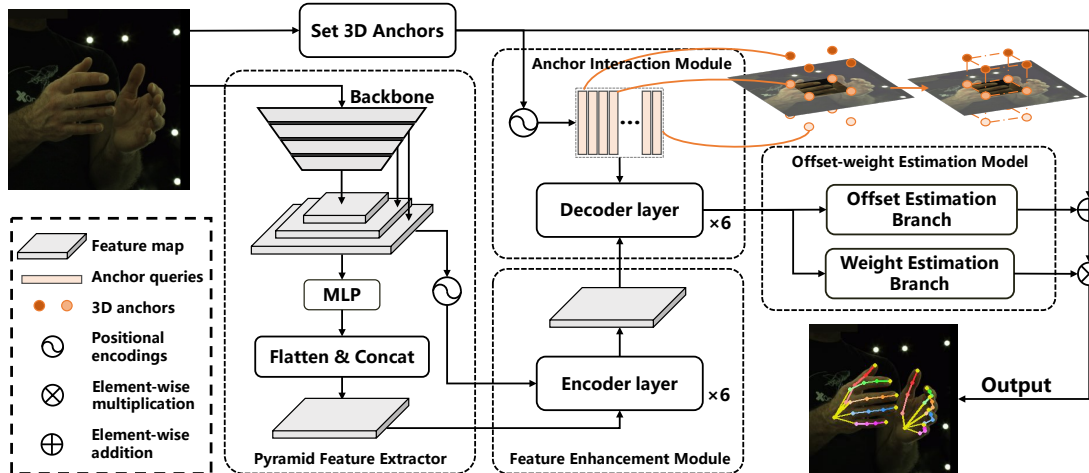


Figure 2. The main technical pipeline of A2J-Transformer. A2J-Transformer consists of 3 main models: pyramid feature extractor, anchor refinement model (containing feature enhancement module and anchor interaction module) and anchor offset-weight estimation model. The anchor interaction module aims to establish the connection (orange line) between anchors (orange dots).

hand model MANO [33], subsequent model-based methods [21, 35, 41, 48] are mostly based on it while using CNN or GCN modules. Due to the presence of a sufficiently strong prior model, model-based methods generally have good performance and are more stable than model-free methods. However, these approaches usually lose tracking when there are strong hand interactions and occlusions. At the same time, modeling the hands of different people is needed [13] in practical usage, which to a certain extent reduces the generalization ability of the model. Therefore, we turn our attention to the model-free approach, which needs no prior information and has more flexibility.

Model-free approach. Model-free approaches [2, 5, 7, 17, 19, 22, 27–29, 43] have been developed for a long time. In particular, the task of single hand pose estimation based on depth maps have been available for very mature methods [2, 17, 28, 43]. However, their extensions to two hands and RGB domain are non-trivial due to the severe occlusion and similar appearance between hand joints. After Moon *et al.* [29] propose the *InterHand2.6M* dataset, model-free approaches [5–7, 19, 22, 27] for interacting hand pose estimation has made great progress. For be better resistant to occlusion, some research [7, 22, 27] tend to separate the interacting hands and estimate the two hands separately, some methods [5, 6] perform dense modeling by using point cloud networks. However, the prediction of details of interacting hands by these methods depend heavily on the quality of the segmentation results or the point cloud generations. Some methods [12, 29] obtain the coordinate of hand joints by directly regressing the heatmap, which could be intuitive and flexible. However, the current methods are not ideal for local detail feature extraction and still have performance shortcomings for 3D interacting hand pose estimation. The proposed method of Hampali *et al.* [12] is

similar to ours, which directly regress the keypoints of two hands, but there is still a problem of poor prediction effect when having strong occlusions. In contrast, by regarding densely distributed anchor points as local regressioners and establishing interactions between them, our proposed A2J-Transformer can not only extract local detailed hand poses, but also obtain global articulated hand joints’ information.

Transformer in hand pose estimation. With the rise of the self-attention mechanism and the proposal of the transformer model [40], more and more visual fields promote their development by introducing the transformer model, like image classification, object detection, 3D mesh reconstruction and so on [18]. Since the Transformer model has a strong ability in capturing non-local features which is surely helpful for the hand pose estimation field, there has been many works [15, 23, 24] to introduce Transformer into this area. However, these architectures are all designed for single hand pose estimation. Recent methods for interacting hand pose estimation have achieved good results, but still suffer from performance shortcomings [12] or model limitations in flexibility [21].

Accordingly, our A2J-Transformer belongs to model-free region and introduced the Transformer module. Different from previous works, we integrates A2J and Transformer into an uniform model (i.e., A2J-Transformer) with end-to-end learning capacity, to reveal our key theoretical insight on addressing 3D interacting hand pose estimation (IHPE) task via concerning local and global visual context jointly. Meanwhile, 2D anchor point within A2J is evolved to 3D version adaptive to A2J-Transformer, to alleviate ill-posed 2D to 3D hand pose lifting problem using monocular RGB image. These propositions technically sound with promising performance and concern 3D IHPE’s specific characteristics deeply.

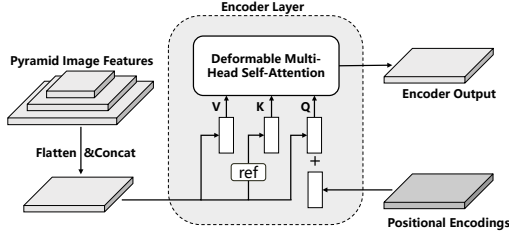


Figure 3. The first encoder layer of feature enhancement module.

3. A2J-Transformer: Anchor-to-Joint Transformer Network

As shown in Fig. 2, A2J-Transformer consists of 3 main models: pyramid feature extractor, anchor refinement model and anchor offset-weight estimation model.

3.1. Pyramid feature extractor

Since multi-scale features can obtain both global information of the input image and retain enough detailed information, feature pyramids are well suitable for the task of interacting hand pose estimation. Therefore, ResNet-50 [14] is used as backbone network to extract the pyramid features from input RGB images. In particular, we get the pyramid features by using the output layer 2-4 with $8, 16, 32 \times$ down sample rates on in-plane size. Meanwhile, 3 convolution layers are used for generating inputs of transformer model of each feature, and 1 convolution layer is used additionally for extracting the last feature layer to maintain more spatial information. Finally, these 4 feature maps are sent to the next anchor refinement model.

3.2. Anchor refinement model

Anchor refinement model aims to simultaneously focus on the non-local articulated and the local fine-grained features. It contains feature enhancement module and anchor interaction module, which can enhance image features and establish the interactions between anchors respectively.

3.2.1 Feature enhancement module

Since multi-scale features are useful for capturing global clues and recovering local details, we integrated the self-attention module [25] to enhance multi-scale features. So we refer to this module as feature enhancement module, which consists of six encoder layers. The first encoder layer of this module is shown in Fig. 3, and the input features of the rest encoder layers are the outputs of the previous layers. All dimensions of input and output features are 256.

Technically, for the input feature pyramid, convolution layer and group normalization layer [42] are firstly used to process them to a same in-plane size. After flatten and concatenation, the generated features F are added to the positional encodings P_{xy} :

$$P_{xy} = \text{PE}(x, y), \quad (1)$$

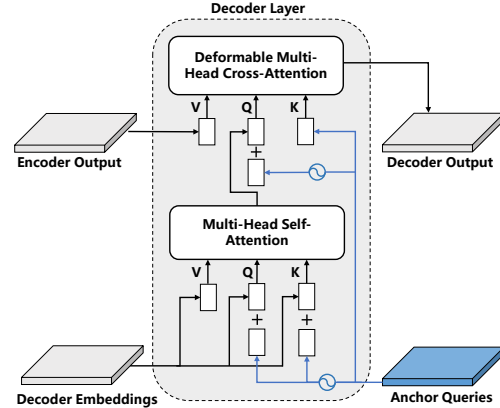


Figure 4. One decoder layer of anchor interaction module.

where PE means positional encoding to generate sinusoidal embeddings from float numbers [25], and x, y represent the in-plane positions of the feature F . Besides, we replace the self-attention module with multi-scale deformable attention module (MSDAM) [49] to mitigate issues of slow convergence and limited feature resolution.

For self-attention module, the queries Q , keys K and values V have the same content item F , and the queries contains an extra position item P_{xy} :

$$Q = F + P_{xy}, \quad K = \text{ref}(F), \quad V = F, \quad (2)$$

where $\text{ref}(\cdot)$ means sample reference keys following [49]. Then, Q, K, V are sent to the MSDAM to get the enhanced features for next encoder layers.

Finally, global-aware features are generated after 6 encoder layers and sent to the anchor interaction module.

3.2.2 Anchor interaction module

In A2J-Transformer, a uniform distribution of 3D anchor points are densely set up to perform direct estimation of hand joints through these 3D anchor points. In other words, these 3D anchor points play the role of local coordinate regressors. More details on the implementation of 3D anchor settings are described in Sec. 3.3. Estimating hand pose through local anchor points has two advantages. First, the setting of dense local 3D anchor points can effectively capture the refined local details from images, which has a good effect for estimating the detail information of strong interacting hands. Second, cross-attention module can establish interaction between local anchor points to capture global clues, which is beneficial for handling occlusion.

Based on this, anchor interaction module containing 6 decoder layers are designed to link individual anchor points, making global information available for each anchor point. One decoder layer is shown in Fig. 4 and for the first decoder layer, the Decoder Embeddings will be replaced by the Encoder Output. All dimensions of input and output features are 256.

Symbol	Definition
J & j	Joint set and joint. $j \in J$.
A & a	Anchor point set and anchor point. $a \in A$.
T_j^i	In-plane coordinate of joint j .
T_j^d	Depth coordinate of joint j .
$C^i(a)$	In-plane coordinate of anchor point a .
$C^d(a)$	Depth coordinate of anchor point a .
$W_j(a)$	Weight of anchor a towards joint j .
$O_j^i(a)$	Predicted in-plane offset towards joint j from anchor point a .
$O_j^d(a)$	Predicted depth offset towards joint j from anchor point a .

Table 1. Symbol definition within A2J-Transformer.

The symbols within A2J-Transformer are defined in Table 1 for better explaining. Different from previous Transformer-based works, we take the understanding of DAB-DETR [25] and explicitly set the coordinates of each anchor a to the queries, which we call Anchor Queries. We denote $a_q = (x_q, y_q, d_q)$ as the q -th anchor query, while $x_q, y_q, d_q \in \mathbb{R}$ denotes the coordinate of a in in-plane and depth. For a_q , the spatial encodings P_q is generated by:

$$P_q = \text{MLP}(\text{PE}(a_q)), \quad (3)$$

where parameters of MLP are shared across all layers.

For self-attention module, settings of queries, keys and values of decoder layers are similar to the setting in feature enhancement module:

$$Q = D + P_q, \quad K = D + P_q, \quad V = D, \quad (4)$$

where D denotes the decoder embeddings.

In cross-attention module, we add the positional query embeddings P_q to the output of self-attention module D to get the context aware anchor informative queries Q . Besides, anchor queries are directly set to the reference points K , and V is the encoder output E :

$$Q = D + P_q, \quad K = a_q, \quad V = E, \quad (5)$$

and MSDAM is applied for calculating cross-attention.

3.3. Anchor offset-weight estimation model

As described in Section 3.2, when each anchor point is linked to each other through the Transformer module, they have both the ability to recover local details and perceive global information. To get final output, anchor offset-weight estimation model is used to estimate the 3D offsets and weights of each anchor with respect to each hand joints. That is, each anchor acts as a local estimator. The offsets and weights are estimated separately for all hand joints. Finally, we fuse the estimation results of all anchors in a weighted summation way to get the final result of joints.

The 3D anchor structure is shown in Fig. 5. The in-plane coordinates of 3D anchors are densely distributed on the input RGB image with an in-plane stride $S_t = 16$. This could ensure that for each pixel in the extracted feature maps,

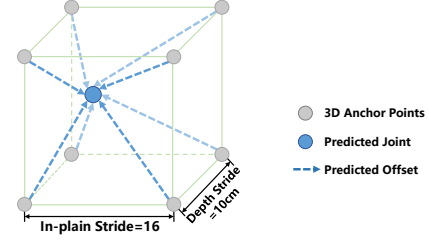


Figure 5. 3D anchors in A2J-Transformer. Joints will be estimated from anchors and offsets.

there can be at least one anchor point corresponding to it while reducing the model size. On this basis, we extend the depth value number of the anchor points. In addition to the original 0 depth value, two depth values are set at the position of ± 100 mm under the world coordinate, centered on the root joint of each hand. This is due to the data processing procedure within baseline. That is, hand joints outside the range of ± 200 mm from the root of the hand are considered as invalid joints. Therefore, for the input image size 256×256 , there are $16 \times 16 \times 3$ total anchors. This setting method extends the anchor point to the 3D space, so as to better fit the depth coordinates of the predicted joints.

Essentially, anchor points are local regressors used to estimate each joint relative to itself. As shown in Fig. 2, each anchor point returns a 3D coordinate offset from itself to all joints in **offset estimation branch**. Since different anchor points focus on different feature ranges, the contribution to each anchor point will also be different. So we predict the weight of each anchor point by **weight estimation branch**. Therefore, by these two branches, the coordinates of each joint can be calculated as the weighted sum of the prediction results of all anchor points' coordinates.

Technically, to get the offsets $O_j^i(a)$, $O_j^d(a)$ and the anchor weights $W_j(a)$, 2 MLP layers are added on the outputs of anchor interaction model. 3D offsets from each anchor point to each joint $O_j(a)$ are regressed by 1 MLP layer and then divided into $O_j^i(a)$, $O_j^d(a)$. Another MLP layer regresses each anchor weight $W_j(a)$ for each joint. Finally, the 3D coordinates of predicted joint j can be expressed as:

$$\begin{cases} \hat{T}_j^i = \sum_{a \in A} \tilde{W}_j(a) (C^i(a) + O_j^i(a)) \\ \hat{T}_j^d = \sum_{a \in A} \tilde{W}_j(a) (C^d(a) + O_j^d(a)) \end{cases}, \quad (6)$$

where \hat{T}_j^i and \hat{T}_j^d indicate the estimated in-plane and depth coordinate of target joint j , C_a^i and C_a^d denote the in-plane and depth coordinates of an anchor point a . $\tilde{W}_j(a)$ is the normalized weight from anchor point a towards joint j , which could be calculated by soft-max:

$$\tilde{W}_j(a) = \frac{e^{W_j(a)}}{\sum_{a \in A} e^{W_j(a)}}. \quad (7)$$

In this way, the estimated hand joints will adaptively select those anchor points with greater contributions to itself and give them large weights. Finally, the joint coordinates and the anchor weights are supervised through joint estimation loss and anchor point surrounding loss.

3.4. Loss functions

For training our performed A2J-Transformer model, we utilize two loss functions: (1) joint estimation loss, (2) anchor surrounding loss following [43].

Joint estimation loss. After getting the estimated 3D joint coordinates, we use the joint estimation loss to supervise the final output, which is formulated as:

$$loss_1 = \alpha \sum_{j \in J} L_{\tau_1}(\hat{T}_j^i - T_j^i) + \sum_{j \in J} L_{\tau_2}(\hat{T}_j^d - T_j^d), \quad (8)$$

where \hat{T}_j^i and \hat{T}_j^d denotes the estimated in-plane coordinate and depth coordinate of joint j from Eq.6, and T_j^i and T_j^d are the given in-plane and depth GT coordinates of joint j ; parameter α defaults to 0.5 to balance the loss between in-plane and depth offset estimation task. $L_{\tau}(\cdot)$ is the *smooth_{L1}* like loss function [31] given by:

$$L_{\tau}(x) = \begin{cases} \frac{1}{2\tau}x^2, & \text{for } |x| < \tau, \\ |x| - \frac{\tau}{2}, & \text{otherwise.} \end{cases} \quad (9)$$

τ_1, τ_2 are set to 1, 3 for better smoothing the depth value.

Anchor surrounding loss. To lead the informative anchor points locate around the hand joints, thus facilitating the generalization ability of our model, we define the anchor surrounding loss by:

$$loss_2 = \sum_{j \in J} L_{\tau_1} \left(\sum_{a \in A} \tilde{W}_j(a) C^i(a) - T_j^i \right) + \sum_{j \in J} L_{\tau_2} \left(\sum_{a \in A} \tilde{W}_j(a) C^d(a) - T_j^d \right), \quad (10)$$

where τ_1 and τ_2 are also set to 1 and 3.

Finally, the total loss function is formulated as:

$$loss = \lambda_1 loss_1 + \lambda_2 loss_2. \quad (11)$$

where λ_1 and λ_2 are set to 3 and 1 to balance two losses.

4. Experiments

4.1. Experimental setting

4.1.1 Datasets

InterHand2.6M dataset [29]. InterHand2.6M is a representative two-hand RGB image dataset with challenging hand interacting scenarios. It contains 1.36M train images and 849K test images. The ground-truth contains semi-automatically annotated 3D coordinates of 42 hand joints. For fair comparison, we choose all test frames for result evaluation following InterNet [29].

RHP dataset [50]. RHP is a synthesized dataset contains two isolated hand data. 41K training and 2.7K testing samples are contained. Since the background of this dataset is an outdoor scene, we use this dataset to approximate the generalization ability of our model on in-the-wild conditions. We also follow InterNet for fair comparison.

NYU dataset [38]. NYU is a single-hand depth image dataset which has 72K training images and 8.2K testing images with 3D annotation on 36 hand joints. Following [4, 10, 28, 43], we pick 14 of the 36 joints for evaluation.

HANDS 2017 dataset [46]. HANDS 2017 is a single-hand depth image dataset which has 957K training images and 295K testing images combined from BigHand2.2M [47] and First-Person Hand Action [46]. The ground-truth contains 3D coordinate of 21 hand joints.

4.1.2 Evaluation metrics

The **Mean Per Joint Position Error (MPJPE)** is used for evaluation on InterHand2.6M [29]. It is defined as a Euclidean distance (mm) between predicted and ground-truth 3D joint locations. Following [29], this metric is used after root joint alignment for each left and right hand separately. For RHP dataset, **end point error (EPE)** is used, which is defined as a mean Euclidean distance (mm) between the predicted and ground-truth 3D hand pose after root joint alignment. For the two depth image dataset, the **average 3D distance error** is used following [28, 43]. Besides, **FPS** is used to evaluate the inference speed, and all models are tested on single NVIDIA RTX 2080ti GPU during inference.

4.1.3 Implementation details

A2J-Transformer is implemented using PyTorch. For InterHand2.6M and RHP dataset, we directly crop the RGB images and resize them to 256×256 resolution. The data augmentations are exactly the same as InterNet [29]. For NYU and HANDS 2017 dataset, we follow [28] to crop and resize the depth image to 176×176 . We train our model using the Adam optimizer [20]. The learning rate is set to 1×10^{-4} with a weight decay of 1×10^{-4} in all cases. There are totally 42 epochs for InterHand2.6M, RHP and NYU dataset and 17 epochs for HANDS 2017 dataset.

4.2. Results

InterHand2.6M dataset: Comparison with the state-of-the-art methods on InterHand2.6M is listed in Table 2. It can be observed that:

- In general, A2J-Transformer outperforms other model-free methods by a large margin testing under all scenarios. This proves that our method has a significant advancement in extracting effective information from interacting hands. Compared with model-based methods, A2J-Transformer has a comparable result with the state-of-the-art method without using any hand prior information. Be-

Methods	MPJPE (mm)			FPS (s)	Model Size(M)
	Single	Two	All		
Model-based					
Zhang <i>et al.</i> [48]	-	13.48	-	17.02	143
Meng <i>et al.</i> [27]	8.51	13.12	10.97	15.47	55
Li <i>et al.</i> [21]	-	8.79	-	18.05	39
Model-free					
Moon <i>et al.</i> [29]	12.16	16.02	14.22	107.08	47
Kim <i>et al.</i> [19]	-	-	12.08	-	-
Fan <i>et al.</i> [7]	11.32	15.57	-	-	-
Hampali <i>et al.</i> [12]	10.99	14.34	12.78	19.66	48
Ours	8.10	10.96	9.63	25.65	42

Table 2. Comparison with state-of-the-art model-based and model-free methods on InterHand2.6M [29]. MPJPE, FPS and model size are reported.

Methods	GT S	GT H	EPE
Zimm. <i>et al.</i> [50]	✓	✓	30.42
chen <i>et al.</i> [3]	✓	✓	24.20
Yang <i>et al.</i> [44]	✓	✓	19.95
Spurr <i>et al.</i> [36]	✓	✓	19.73
Spurr <i>et al.</i> [36]	✗	✗	19.73
Moon <i>et al.</i> [29]	✗	✗	20.89
A2J-Transformer(Ours)	✗	✗	17.75

Table 3. EPE comparison with previous state-of-the-art methods on RHP. Following [29], the checkmark denotes a method use ground-truth information during inference time. S and H denote scale and handness, respectively.

sides, A2J-Transformer has fairly fast inference speed just behind baseline [29] and the smallest model size. In conclusion, our model achieves the best overall performance in terms of performance, running speed and model size.

- Specifically, compared with baseline [29], A2J-Transformer could get an improvement of 4.06, 5.06 and 4.59mm under three scenarios. Compared with the SOTA model-free method [12], the improvement of our method is 2.89, 3.38 and 3.15mm. Compared with the SOTA model-based method [21], our method could receive a comparable performance under two hands scenario without requiring any hand prior, which makes our model more flexible.

- For the running speed, A2J-Transformer has a fast inference speed with 25 FPS, surpassing all methods except baseline. Besides, A2J-Transformer also has the smallest model size with only 42M parameters. These characteristics brings our model great convenience for the future expansion and real-time 3D hand pose estimation.

RHP dataset: Comparison on RHP dataset is shown in Table 3. It shows that A2J-Transformer outperforms previous methods without relying on ground-truth information during inference time. The experiments demonstrate the effectiveness on in-the-wild images and shows the good generalization ability of A2J-Transformer.

NYU and HANDS 2017 dataset: Comparison with state-of-the-art depth based single hand estimation methods

Methods	Mean Error (mm)	FPS(s)
Moon <i>et al.</i> [28]	9.22	35
Xiong <i>et al.</i> [43]	8.61	105.06
Fang <i>et al.</i> [8]	8.29	111.20
Ours	8.43	24.81

Table 4. Performance comparison on NYU dataset [38]. Our proposed A2J-Transformer can guarantee a competitive performance for the depth image input.

Methods	Mean Error (mm)	FPS(s)
Ge <i>et al.</i> [9]	11.30	48
Yuan <i>et al.</i> [45]	9.97	-
Moon <i>et al.</i> [28]	9.95	3.5
Xiong <i>et al.</i> [43]	8.57	105.06
Ours	8.32	24.81

Table 5. Performance comparison on HANDS 2017 dataset [46]. Our method can get state-of-the-art performance on this dataset.

on NYU and HANDS 2017 dataset are given in Table 4 and Table 5. Since A2J-Transformer is not specifically designed for single hand estimation on depth image, we just changed the input channel to verify the generalization ability of our model through this experiment. We can summarize that:

- Although A2J-Transformer is based on the RGB image of interacting hands, it still achieves state-of-the-art performance on HANDS 2017 dataset and gets comparable result on NYU dataset. This relies on the strong ability of A2J-Transformer to grasp the articulated hand information and the fitting ability of 3D anchor points. Compared with A2J [43], certain performance improvement can be achieved on two datasets, which proves that A2J-Transformer has a strong generalization ability.

4.3. Ablation study

4.3.1 Component effectiveness analysis

The component effectiveness analysis within A2J-Transformer is executed on Interhand2.6M dataset. We explore the effectiveness of four parts: (1) Transformer-based model (anchor refinement model), (2) A2J (anchor-to-joint) module, (3) 3D anchor weights, (4) MSDAM. The specific implementation details are respectively set as: (1) replacing the anchor refinement model with the convolution modules in A2J, (2) directly regressing the hand joints without using anchor-to-joint module, (3) setting the weights of all anchors to all the same values and normalize them, (4) replacing the MSDAM with the origin attention module. The results are listed in Table 6. It can be observed that:

- After removing the Transformer-based model and A2J module, the performance of A2J-Transformer drops by 5 mm and 6 mm respectively, proving the effectiveness of addressing 3D interacting hand pose estimation task via concerning local and global visual context jointly.

- After removing the 3D anchor weights, the perfor-

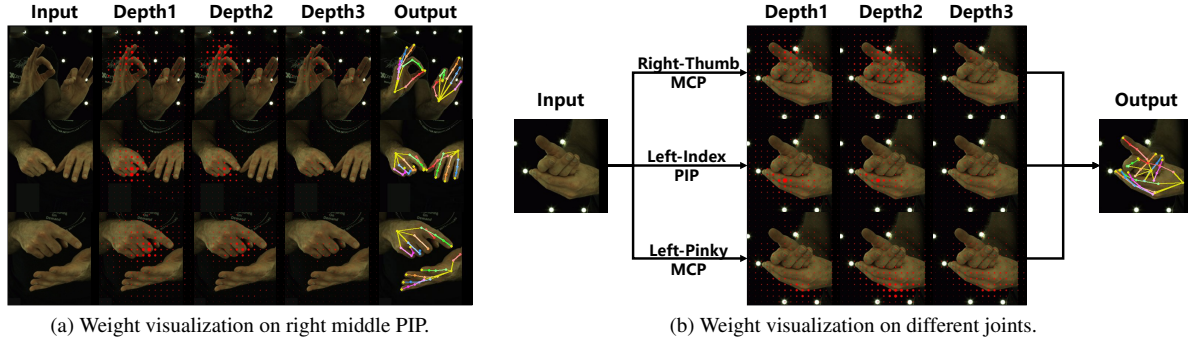


Figure 6. Qualitative results of A2J-Transformer. We show the input, output and weights of anchors on different depth value layers. Red dots in the three depth maps indicate the anchors set at depth positions +100mm, 0mm, and -100mm from the root joint respectively. The shade of red dots represent the weights assigned to these anchors as described in Sec. 3.3.

Trans.	A2J	Weights	MSDAM	MPJPE (mm)
✗	✓	✓	✓	14.44
✓	✗	✓	✓	15.36
✓	✓	✗	✓	14.04
✓	✓	✓	✗	10.69
✓	✓	✓	✓	9.63

Table 6. Component effectiveness analysis of A2J-Transformer. ‘Trans.’ means Transformer-based model (anchor refinement model) and ‘Weights’ means 3D anchor weights.

mance of A2J-Transformer drops by 4.4 mm. This proves that there is a performance difference in the regression results of each 3D anchor point, so the weights predicted by the model are crucial for the prediction of hand joints.

- After replacing the MADAM with the origin attention module, the performance of A2J-Transformer drops by 1mm, which proves the MSDAM is useful to our model.

4.3.2 Anchor setting analysis

In order to explore the impact on model performance, more in-plane and depth values are set for comparative experiments. The specific setting methods and their performance results are shown in Table 7. All depth values are uniformly selected near the hand joints, just like the selection of 3 depth values as described in Sec. 3.3. It can be noticed that, when more anchor in-plane and depth values are set, the performance of A2J-Transformer will improve while the inference speed will decrease in general. In order to strike a balance between accuracy and efficiency, the value of 256 and 3 are finally chosen.

4.4. Qualitative evaluation and limitation

We show the qualitative evaluation results in Fig. 6. We can see that, A2J-Transformer could automatically enlarge the informative anchors’ weights when different joint coordinates need to be predicted. The model achieves accurate results even with severe occlusions in the interacting hands. The major limitation of our method is when there is a large area of occlusion or missing in the hand area, the results

In-plane	Depth	MPJPE (mm)	FPS (s)
256	7	9.50	19.33
256	5	9.61	21.21
256	3	9.63	25.65
256	1	9.75	26.06
64	3	12.28	25.25
16	3	14.07	27.39
4	3	15.48	27.63

Table 7. Anchor setting analysis of A2J-Transformer. ‘In-plane’ and ‘Depth’ denotes the number of selected anchor number values for in-plane and depth direction, respectively.

predicted by our model will have deviations.

5. Conclusion

In this paper, an 3D monocular RGB interacting hand pose estimation approach termed A2J-Transformer is proposed. Equipped with Transformer’s non-local encoding-decoding framework, A2J is evolved to capture interacting hands’ local fine details and global articulated clues among joints simultaneously. Besides, 3D anchors are used to better fit the depth information and estimation of accurate 3D coordinates. Experiments on InterHand2.6M and RHP dataset demonstrate the effectiveness and superiority of A2J-Transformer and extensions on NYU and HANDS 2017 dataset show the generalization ability. In future work, we will try to represent the movement of anchor points and extend our method to model-based region.

Acknowledgment

This work is jointly supported by the National Natural Science Foundation of China (Grant No. 62271221 and U1913602). Joey Tianyi Zhou is funded by the SERC (Science and Engineering Research Council) Central Research Fund (Use-Inspired Basic Research), and the Singapore Government’s Research, and Innovation and Enterprise 2020 Plan (Advanced Manufacturing and Engineering Domain) under programmatic Grant A18A1b0045.

References

- [1] Luca Ballan, Aparna Taneja, Jürgen Gall, Luc Van Gool, and Marc Pollefeys. Motion capture of hands in action using discriminative salient points. In *European Conference on Computer Vision*, pages 640–653. Springer, 2012. 1, 2
- [2] 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. 1, 3
- [3] Liangjian Chen, Shih-Yao Lin, Yusheng Xie, Hui Tang, Yufan Xue, Xiaohui Xie, Yen-Yu Lin, and Wei Fan. Generating realistic training images based on tonality-alignment generative adversarial networks for hand pose estimation. *arXiv preprint arXiv:1811.09916*, 2018. 7
- [4] Xinghao Chen, Guijin Wang, Hengkai Guo, and Cairong Zhang. Pose guided structured region ensemble network for cascaded hand pose estimation. *Neurocomputing*, 395:138–149, 2020. 6
- [5] Wencan Cheng, Jae Hyun Park, and Jong Hwan Ko. Hand-foldingnet: A 3d hand pose estimation network using multiscale-feature guided folding of a 2d hand skeleton. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11260–11269, 2021. 1, 3
- [6] Xiaoming Deng, Dexin Zuo, Yinda Zhang, Zhaopeng Cui, Jian Cheng, Ping Tan, Liang Chang, Marc Pollefeys, Sean Fanello, and Hongan Wang. Recurrent 3d hand pose estimation using cascaded pose-guided 3d alignments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):932–945, 2022. 3
- [7] Zicong Fan, Adrian Spurr, Muhammed Kocabas, Siyu Tang, Michael J Black, and Otmar Hilliges. Learning to disambiguate strongly interacting hands via probabilistic per-pixel part segmentation. In *2021 International Conference on 3D Vision (3DV)*, pages 1–10. IEEE, 2021. 1, 3, 7
- [8] 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. In *European Conference on Computer Vision*, pages 120–137. Springer, 2020. 7
- [9] Lihao Ge, Yujun Cai, Junwu Weng, and Junsong Yuan. Hand pointnet: 3d hand pose estimation using point sets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8417–8426, 2018. 7
- [10] Hengkai Guo, Guijin Wang, Xinghao Chen, Cairong Zhang, Fei Qiao, and Huazhong Yang. Region ensemble network: Improving convolutional network for hand pose estimation. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 4512–4516. IEEE, 2017. 6
- [11] Shreyas Hampali, Mahdi Rad, Markus Oberweger, and Vincent Lepetit. Honnotate: A method for 3d annotation of hand and object poses. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3196–3206, 2020. 2
- [12] Shreyas Hampali, Sayan Deb Sarkar, Mahdi Rad, and Vincent Lepetit. Keypoint transformer: Solving joint identification in challenging hands and object interactions for accurate 3d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11090–11100, 2022. 1, 2, 3, 7
- [13] Shangchen Han, Po-chen Wu, Yubo Zhang, Beibei Liu, Linguang Zhang, Zheng Wang, Weiguang Si, Peizhao Zhang, Yujun Cai, Tomas Hodan, et al. Umetrack: Unified multi-view end-to-end hand tracking for vr. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. 3
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 4
- [15] Lin Huang, Jianchao Tan, Ji Liu, and Junsong Yuan. Hand-transformer: Non-autoregressive structured modeling for 3d hand pose estimation. In *European Conference on Computer Vision*, pages 17–33. Springer, 2020. 3
- [16] Lin Huang, Boshen Zhang, Zhilin Guo, Yang Xiao, Zhiguo Cao, and Junsong Yuan. Survey on depth and rgb image-based 3d hand shape and pose estimation. *Virtual Reality & Intelligent Hardware*, 3(3):207–234, 2021. 1
- [17] Umar Iqbal, Andreas Doering, Hashim Yasin, Björn Krüger, Andreas Weber, and Juergen Gall. A dual-source approach for 3d human pose estimation from single images. *Computer Vision and Image Understanding*, 172:37–49, 2018. 1, 3
- [18] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM computing surveys (CSUR)*, 54(10s):1–41, 2022. 3
- [19] Dong Uk Kim, Kwang In Kim, and Seungryul Baek. End-to-end detection and pose estimation of two interacting hands. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11189–11198, 2021. 1, 3, 7
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [21] Mengcheng Li, Liang An, Hongwen Zhang, Lianpeng Wu, Feng Chen, Tao Yu, and Yebin Liu. Interacting attention graph for single image two-hand reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2761–2770, 2022. 1, 2, 3, 7
- [22] Fanqing Lin, Connor Wilhelm, and Tony Martinez. Two-hand global 3d pose estimation using monocular rgb. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2373–2381, 2021. 1, 3
- [23] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1954–1963, 2021. 3
- [24] Kevin Lin, Lijuan Wang, and Zicheng Liu. Mesh graphormer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12939–12948, 2021. 3
- [25] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. *arXiv preprint arXiv:2201.12329*, 2022. 4, 5
- [26] Yang Liu, Jie Jiang, and Jiahao Sun. Hand pose estimation from rgb images based on deep learning: A survey. In

- 2021 *IEEE 7th International Conference on Virtual Reality (ICVR)*, pages 82–89. IEEE, 2021. 1
- [27] Hao Meng, Sheng Jin, Wentao Liu, Chen Qian, Mengxiang Lin, Wanli Ouyang, and Ping Luo. 3d interacting hand pose estimation by hand de-occlusion and removal. *arXiv preprint arXiv:2207.11061*, 2022. 1, 3, 7
- [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 conference on computer vision and pattern recognition*, pages 5079–5088, 2018. 1, 3, 6, 7
- [29] Gyeongsik Moon, Shoou-I Yu, He Wen, Takaaki Shiratori, and Kyoung Mu Lee. Interhand2. 6m: A dataset and baseline for 3d interacting hand pose estimation from a single rgb image. In *European Conference on Computer Vision*, pages 548–564. Springer, 2020. 1, 2, 3, 6, 7
- [30] Iasonas Oikonomidis, Nikolaos Kyriazis, and Antonis A Argyros. Tracking the articulated motion of two strongly interacting hands. In *2012 IEEE conference on computer vision and pattern recognition*, pages 1862–1869. IEEE, 2012. 1, 2
- [31] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 6
- [32] Javier Romero, Hedvig Kjellström, and Danica Kragic. Monocular real-time 3d articulated hand pose estimation. In *2009 9th IEEE-RAS International Conference on Humanoid Robots*, pages 87–92. IEEE, 2009. 1
- [33] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *arXiv preprint arXiv:2201.02610*, 2022. 2, 3
- [34] Jamie Shotton, Andrew Fitzgibbon, Mat Cook, Toby Sharp, Mark Finocchio, Richard Moore, Alex Kipman, and Andrew Blake. Real-time human pose recognition in parts from single depth images. In *CVPR 2011*, pages 1297–1304. Ieee, 2011. 1
- [35] Breannan Smith, Chenglei Wu, He Wen, Patrick Peluse, Yaser Sheikh, Jessica K Hodgins, and Takaaki Shiratori. Constraining dense hand surface tracking with elasticity. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020. 1, 2, 3
- [36] Adrian Spurr, Jie Song, Seonwook Park, and Otmar Hilliges. Cross-modal deep variational hand pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 89–98, 2018. 7
- [37] Danhang Tang, Hyung Jin Chang, Alykhan Tejani, and Tae-Kyun Kim. Latent regression forest: Structured estimation of 3d articulated hand posture. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3786–3793, 2014. 1
- [38] Jonathan Tompson, Murphy Stein, Yann Lecun, and Ken Perlin. Real-time continuous pose recovery of human hands using convolutional networks. *ACM Transactions on Graphics (ToG)*, 33(5):1–10, 2014. 6, 7
- [39] Dimitrios Tzionas, Luca Ballan, Abhilash Srikantha, Pablo Aponte, Marc Pollefeys, and Juergen Gall. Capturing hands in action using discriminative salient points and physics simulation. *International Journal of Computer Vision*, 118(2):172–193, 2016. 1, 2
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2, 3
- [41] Jiayi Wang, Franziska Mueller, Florian Bernard, Suzanne Sorli, Oleksandr Sotnychenko, Neng Qian, Miguel A Otaduy, Dan Casas, and Christian Theobalt. Rgb2hands: real-time tracking of 3d hand interactions from monocular rgb video. *ACM Transactions on Graphics (ToG)*, 39(6):1–16, 2020. 1, 2, 3
- [42] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. 4
- [43] Fu Xiong, Boshen Zhang, Yang Xiao, Zhiguo Cao, Taidong Yu, Joey Tianyi Zhou, and Junsong Yuan. A2j: Anchor-to-joint regression network for 3d articulated pose estimation from a single depth image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 793–802, 2019. 1, 2, 3, 6, 7
- [44] 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*, pages 9877–9886, 2019. 7
- [45] Shanxin Yuan, Guillermo Garcia-Hernando, Björn Stenger, Gyeongsik Moon, Ju Yong Chang, Kyoung Mu Lee, Pavlo Molchanov, Jan Kautz, Sina Honari, Liuhaog Ge, et al. Depth-based 3d hand pose estimation: From current achievements to future goals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2636–2645, 2018. 7
- [46] Shanxin Yuan, Qi Ye, Guillermo Garcia-Hernando, and Tae-Kyun Kim. The 2017 hands in the million challenge on 3d hand pose estimation. *arXiv preprint arXiv:1707.02237*, 2017. 2, 6, 7
- [47] Shanxin Yuan, Qi Ye, Bjorn Stenger, Siddhant Jain, and Tae-Kyun Kim. Bighand2. 2m benchmark: Hand pose dataset and state of the art analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4866–4874, 2017. 6
- [48] Baowen Zhang, Yangang Wang, Xiaoming Deng, Yinda Zhang, Ping Tan, Cuixia Ma, and Hongan Wang. Interacting two-hand 3d pose and shape reconstruction from single color image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11354–11363, 2021. 1, 2, 3, 7
- [49] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. 4
- [50] Christian Zimmermann and Thomas Brox. Learning to estimate 3d hand pose from single rgb images. In *Proceedings of the IEEE international conference on computer vision*, pages 4903–4911, 2017. 6, 7