

Depth-Based 3D Hand Pose Estimation: From Current Achievements to Future Goals

Shanxin Yuan¹ Guillermo Garcia-Hernando¹ Björn Stenger²
Gyeongsik Moon⁴ Ju Yong Chang⁵ Kyoung Mu Lee⁴ Pavlo Molchanov⁶
Jan Kautz⁶ Sina Honari⁷ Lihao Ge⁸ Junsong Yuan⁹ Xinghao Chen¹⁰ Guijin Wang¹⁰
Fan Yang¹¹ Kai Akiyama¹¹ Yang Wu¹¹ Qingfu Wan¹² Meysam Madadi¹³ Sergio Escalera^{13,14}
Shile Li¹⁵ Dongheui Lee^{15,16} Iason Oikonomidis³ Antonis Argyros³ Tae-Kyun Kim¹

Abstract

In this paper, we strive to answer two questions: What is the current state of 3D hand pose estimation from depth images? And, what are the next challenges that need to be tackled? Following the successful Hands In the Million Challenge (HIM2017), we investigate the top 10 state-of-the-art methods on three tasks: single frame 3D pose estimation, 3D hand tracking, and hand pose estimation during object interaction. We analyze the performance of different CNN structures with regard to hand shape, joint visibility, view point and articulation distributions. Our findings include: (1) isolated 3D hand pose estimation achieves low mean errors (10 mm) in the view point range of [70, 120] degrees, but it is far from being solved for extreme view points; (2) 3D volumetric representations outperform 2D CNNs, better capturing the spatial structure of the depth data; (3) Discriminative methods still generalize poorly to unseen hand shapes; (4) While joint occlusions pose a challenge for most methods, explicit modeling of structure constraints can significantly narrow the gap between errors on visible and occluded joints.

1. Introduction

The field of 3D hand pose estimation has advanced rapidly, both in terms of accuracy [2, 4, 5, 6, 7, 26, 28, 32, 38, 45, 48, 49, 53, 54, 57, 60, 62] and dataset quality [10, 22, 43, 47, 52, 59]. Most successful methods treat the estimation task as a learning problem, using random

forests or convolutional neural networks (CNNs). However, a review from 2015 [44] surprisingly concluded that a simple nearest-neighbor baseline outperforms most existing systems. It concluded that *most systems do not generalize beyond their training sets* [44], highlighting the need for more and better data. Manually labeled datasets such as [31, 41] contain just a few thousand examples, making them unsuitable for large-scale training. Semi-automatic annotation methods, which combine manual annotation with tracking, help scaling the dataset size [43, 46, 52], but in the case of [46] the annotation errors are close to the lowest estimation errors. Synthetic data generation solves the scaling issue, but has not yet closed the realism gap, leading to some kinematically implausible poses [37].

A recent study confirmed that cross-benchmark testing is poor due to different capture set-ups and annotation methods [59]. It showed that training a standard CNN on a million-scale dataset achieves state-of-the-art results. However, the estimation accuracy is not uniform, highlighting the well-known challenges of the task: variations in view point and hand shape, self-occlusion, and occlusion caused by objects being handled.

In this paper we analyze the top methods of the *HIM2017* challenge [58]. The benchmark dataset includes data from *BigHand2.2M* [59] and the *First-Person Hand Action dataset (FHAD)* [10], allowing the comparison of different algorithms in a variety of settings. The challenge considers three different tasks: single-frame pose estimation, tracking, and hand-object interaction. In the evaluation we consider different network architectures, preprocessing strategies, and data representations. Over the course of the challenge the lowest mean 3D estimation error could be reduced from 20 mm to less than 10 mm. This paper analyzes the errors with regard to seen and unseen subjects, joint visibility, and view point distribution. We conclude by providing insights for designing the next generation of methods.

¹Imperial College London, ²Rakuten Institute of Technology, ³University of Crete and FORTH, ⁴Seoul National University, ⁵Kwangwoon University, ⁶NVIDIA, ⁷University of Montreal, ⁸Nanyang Technological University, ⁹State University of New York at Buffalo, ¹⁰Tsinghua University, ¹¹Nara Institute of Science and Technology, ¹²Fudan University, ¹³Computer Vision Center, ¹⁴University of Barcelona, ¹⁵Technical University of Munich, ¹⁶German Aerospace Center.

Corresponding author's email: s.yuan14@imperial.ac.uk



Figure 1: Evaluated tasks. For each scenario the goal is to infer the 3D locations of the 21 hand joints from a depth image. In **Single frame pose estimation** (left) and the **Interaction task** (right), each frame is annotated with a bounding box. In the **Tracking task** (middle), only the first frame of each sequence is fully annotated.

Related work. Public benchmarks and challenges in other areas such as ImageNet [35] for scene classification and object detection, PASCAL [9] for semantic and object segmentation, and the VOT challenge [19] for object tracking, have been instrumental in driving progress in their respective field. In the area of hand tracking, the review from 2007 by Erol *et al.* [8] proposed a taxonomy of approaches. Learning-based approaches have been found effective for solving single-frame pose estimation, optionally in combination with hand model fitting for higher precision, *e.g.*, [50]. The review by Supancic *et al.* [44] compared 13 methods on a new dataset and concluded that deep models are well-suited to pose estimation [44]. It also highlighted the need for large-scale training sets in order to train models that generalize well. In this paper we extend the scope of previous analyses by comparing deep learning methods on a large-scale dataset, carrying out a fine-grained analysis of error sources and different design choices.

2. Evaluation tasks

We evaluate three different tasks on a dataset containing over a million annotated images using standardized evaluation protocols. Benchmark images are sampled from two datasets: *BigHand2.2M* [59] and *First-Person Hand Action dataset (FHAD)* [10]. Images from *BigHand2.2M* cover a large range of hand view points (including third-person and first-person views), articulated poses, and hand shapes. Sequences from the *FHAD* dataset are used to evaluate pose estimation during hand-object interaction. Both datasets contain 640×480 -pixel depth maps with 21 joint annotations, obtained from magnetic sensors and inverse kinematics. The 2D bounding boxes have an average diagonal length of 162.4 pixels with a standard deviation of 40.7 pixels. The evaluation tasks are 3D single hand pose estimation, *i.e.*, estimating the 3D locations of 21 joints, from (1) individual frames, (2) video sequences, given the pose in the first frame, and (3) frames with object interaction, *e.g.*, with a juice bottle, a salt shaker, or a milk carton. See Figure 1 for an overview. Bounding boxes are provided as input for tasks (1) and (3). The training data is sampled from the *BigHand2.2M* dataset and only the interaction task uses test data from the *FHAD* dataset. See Table 1 for dataset sizes and the number of total and unseen subjects for each task.

Number of	Train	Test single	Test track	Test interact
frames	957K	295K	294K	2,965
subjects (unseen)	5	10 (5)	10 (5)	2 (0)

Table 1: Data set sizes and number of subjects.

3. Evaluated methods

We evaluate the top 10 among 17 participating methods [58]. Table 2 lists the methods with some of their key properties. We also indirectly evaluate *DeepPrior* [29] and *REN* [15], which are components of *rvhand* [1], as well as *DeepModel* [61], which is the backbone of *LSL* [20]. We group methods based on different design choices.

2D CNN vs. 3D CNN. 2D CNNs have been popular for 3D hand pose estimation [1, 3, 14, 15, 20, 21, 23, 29, 57, 61]. Common pre-processing steps include cropping and resizing the hand volume by normalizing the depth values to $[-1, 1]$. Recently, several methods have used a 3D CNN [12, 24, 56], where the input can be a 3D voxel grid [24, 56], or a projective D-TSDF volume [12]. Ge *et al.* [13] project the depth image onto three orthogonal planes and train a 2D CNN for each projection, then fusing the results. In [12] they propose a 3D CNN by replacing 2D projections with a 3D volumetric representation (projective D-TSDF volumes [40]). In the *HIM2017* challenge [58], they apply a 3D deep learning method [11], where the inputs are 3D points and surface normals. Moon *et al.* [24] propose a 3D CNN to estimate per-voxel likelihoods for each hand joint. *NAIST_RV* [56] proposes a 3D CNN with a hierarchical branch structure, where the input is a 50^3 -voxel grid.

Detection-based vs. Regression-based. Detection-based methods [23, 24] produce a probability density map for each joint. The method of *RCN-3D* [23] is an RCN+ network [17], based on Recombinator Networks (RCN) [18] with 17 layers and 64 output feature maps for all layers except the last one, which outputs a probability density map for each of the 21 joints. *V2V-PoseNet* [24] uses a 3D CNN to estimate per-voxel likelihood of each joint, and a CNN to estimate the center of mass from the cropped depth map. For training, 3D likelihood volumes are generated by placing normal distributions at the locations of hand joints. Regression-based methods [1, 3, 11, 14, 20, 21, 29, 56]

Method	Model	Input	Aug. range (s, θ ,t)	3D	De	Hi	St	M	R
<i>V2V-PoseNet</i> [24]	3D CNN, per-voxel likelihood of each joint	88×88×88 voxels	[0.8, 1.2] [-40,40] [-8,8]	✓	✓	✗	✗	✗	✓
<i>RCN-3D</i> [23]	RCN+ network [17] with 17 convolutional layers	80×80	[0.7, 1.1] [0, 360] [-8,8]	✗	✓	✗	✗	✓	✓
<i>oasis</i> [11]	Hierarchical PointNet with three set abstraction levels and three full-connected layers	1024 3D points	random scaling*	✗	✗	✗	✓	✗	✗
<i>THU_VCLab</i> [3]	Pose-REN [3]: <i>REN</i> [15] + cascaded + hierarchical.	96×96	[0.9,1.1] [-45,45] [-5,5]	✗	✗	✓	✗	✓	✓
<i>NAIST_RV</i> [56]	3D CNN with 5 branches, one for each finger	50×50×50 3D grid	[0.9,1.1] [-90, 90] [-15,15]	✓	✗	✓	✗	✗	✗
<i>Vanora</i> [14]	shallow CNN trained end-to-end	resized 2D	random scaling*	✗	✗	✗	✗	✗	✗
<i>strawberryfg</i> [55]	ResNet-152 + [42]	224×224	None	✗	✗	✗	✓	✗	✓
<i>rvhand</i> [1]	ResNet [16] + <i>REN</i> [15] + Deep Prior [29]	192×192	[0.9,1.1] [-90, 90] [-15,15]	✗	✗	✓	✓	✗	✓
<i>mmadadi</i> [21]	Hierarchical tree-like structured CNN [21]	192×192	[random] [-30, 30] [-10,10]	✗	✗	✓	✓	✗	✗
<i>LSL</i> [20]	ScaleNet to estimate hand scale + <i>DeepModel</i> [61]	128×128	[0.85,1.15] [0,360] [-20,20]	✗	✗	✗	✓	✗	✗

Table 2: Methods evaluated in the hand pose estimation challenge. Methods are ordered by average error on the leader-board.

* in both methods, hand segmentation is performed considering different hand arm lengths. **3D**, **De**, **Hi**, **St**, **M**, and **R** denote **3D CNN**, **Detection-based method**, **Hierarchical model**, **Structure model**, **Multistage model**, and **Residual net**, respectively.

directly map the depth image to the joint locations or the joint angles of a hand model [39, 61]. *rvhand* [1] combines ResNet [16], Region Ensemble Network (*REN*) [15], and *DeepPrior* [29] to directly estimate the joint locations. *LSL* [20] uses one network to estimate a global scale factor and a second network [61] to estimate all joint angles, which are fed into a forward kinematic layer to estimate the hand joints.

Hierarchical models divide the pose estimation problem into sub-tasks [1, 3, 15, 21, 56]. The evaluated methods divide the hand joints either by finger [21, 56], or by joint type [1, 3, 15]. *mmadadi* [21] designs a hierarchically structured CNN, dividing the convolution+ReLU+pooling blocks into six branches (one per finger with palm and one for palm orientation), each of which is then followed by a fully connected layer. The final layers of all branches are concatenated into one layer to predict all joints. *NAIST_RV* [56] chooses a similar hierarchical structure of a 3D CNN, but uses five branches, each to predict one finger and the palm. *THU_VCLab* [3], *rvhand* [1], and *REN* [15] apply constraints per finger and joint-type (across fingers) in their multiple regions extraction step, each region containing a subset of joints. All regions are concatenated in the last fully connected layers to estimate the hand pose.

Structured methods embed physical hand motion constraints into the model [11, 20, 21, 29, 55, 61]. Structural constraints are included in the CNN model [20, 27, 29] or in the loss function [21, 55]. *DeepPrior* [29] learns a prior model and integrates it into the network by introducing a *bottleneck* in the last CNN layer. *LSL* [20] uses prior knowledge in *DeepModel* [61] by embedding a kinematic model layer into the CNN and using a fixed hand model. *mmadadi* [21] includes the structure constraints in the loss

function, which incorporates physical constraints about natural hand motion and deformation. *strawberryfg* [55] applies a structure-aware regression approach, Compositional Pose Regression [42], and replaces the original ResNet-50 with ResNet-152. It uses phalanges instead of joints for representing pose, and defines a loss function that encodes long-range interaction between the phalanges.

Multi-stage methods propagate results from each stage to enhance the training of the subsequent stages [3, 23]. *THU_VCLab* [3] uses *REN* [15] to predict an initial hand pose. In the following stages, feature maps are computed with the guidance of the hand pose estimate in the previous stage. *RCN-3D* [23] has five stages: (1) 2D landmark estimation using an RCN+ network [17], (2) estimation of corresponding depth values by multiplying probability density maps with the input depth image, (3) inverse perspective projection of the depth map to 3D, (4) error compensation for occlusions and depth errors (a 3-layer network of residual blocks) and, (5) error compensation for noise (another 3-layer network of residual blocks).

Residual networks. ResNet [16] is adopted by several methods [1, 3, 15, 23, 24, 42]. *V2V-PoseNet* [24] uses residual blocks as main building blocks. *strawberryfg* [55] implements the Compositional Pose Regression method [42] by using ResNet-152 as basic network. *RCN-3D* [23] uses two small residual blocks in its fourth and fifth stage.

4. Results

The aim of this evaluation is to identify success cases and failure modes. We use both standard error metrics [30, 37, 51] and new proposed metrics to provide further insights. We consider joint visibility, seen vs. unseen subjects, hand view point distribution, articulation distribution, and per-joint accuracy.

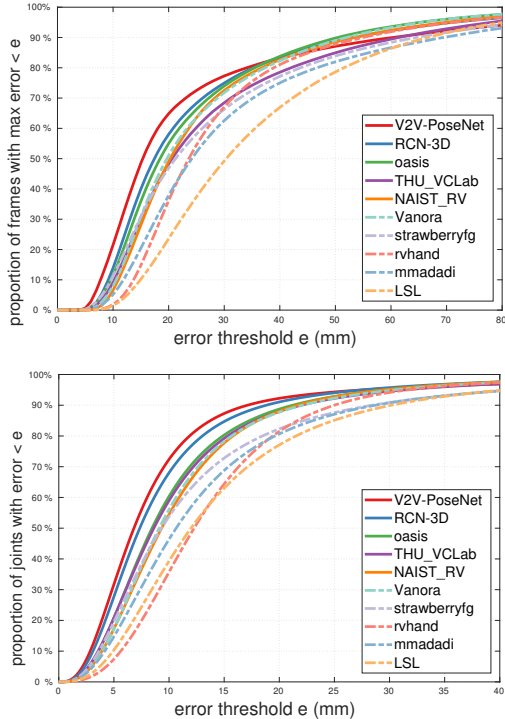


Figure 2: Estimation errors. (top) proportion of frames within maximum error threshold [51], (bottom) proportion of joints within an error threshold [37].

Method \ Case	Seen Visible	Seen Occ	Unseen Visible	Unseen Occ
V2V-PoseNet	6.2	8.0	11.1	14.6
RCN-3D	6.9	9.0	10.6	14.8
oasis	8.2	9.8	12.4	14.9
THU_VCLab	8.4	10.2	12.5	16.1
NAIIST_RV	8.8	10.1	13.1	15.6
Vanora	8.8	10.5	12.9	15.5
strawberryfg	9.3	10.7	16.4	18.8
rvhand	12.2	11.9	16.1	17.6
mmadadi	10.6	13.6	15.6	19.7
LSL	11.8	13.1	18.1	19.2
Top5	7.7	9.4	11.9	15.2
All	9.1	10.7	13.9	16.7

Table 3: Mean errors (in mm) for single frame pose estimation, divided by cases. ‘Seen’ and ‘Unseen’ refers to whether or not the hand shape was in the training set, and ‘Occ’ denotes ‘Occluded joints’.

4.1. Single frame pose estimation

Over the 6-week period of the challenge the lowest mean error could be reduced from 19.7 mm to 10.0 mm by exploring new model types and improving data augmentation, optimization and initialization. For hand shapes seen during training, the mean error was reduced from 14.6 mm




Figure 3: Estimating annotation errors. Two examples of *Easy Poses* overlaid with estimates by the top five methods (shown in different colors). Poses vary slightly, but are close to the ground truth (black).

to 7.0 mm, and for unseen hand shapes from 24.0 mm to 12.2 mm. Considering typical finger widths of 10-20 mm, these methods are becoming applicable to scenarios like pointing or motion capture, but may still lack sufficient accuracy for fine manipulation that is critical in some UI interactions.

We evaluate ten state-of-the-art methods (Table 2) directly and three methods indirectly, which were used as components of others, *DeepPrior* [29], *REN* [15], and *DeepModel* [61]. Figure 2 shows the results in terms of two metrics: (top) the proportion of frames in which all joint errors are below a threshold [51] and (bottom) the total proportion of joints below an error threshold [37].

Figure 4 (top-left) shows the success rates based on per-frame average joint errors [30] for a varying threshold. The top performer, *V2V-PoseNet*, estimates 70% of frames with a mean error of less than 10 mm, and 20% of frames with mean errors under 5 mm. All evaluated methods achieve a success rate greater than 80% with an average error of less than 20 mm.

As has been noted by [15, 27], data augmentation is beneficial, especially for small datasets. However, note that even though the top performing methods employ data augmentation, it is still difficult to generalize to hands from unseen subjects, see Table 3, with an error gap of around 6 mm between seen and unseen subjects. Some methods generalize better than others, in particular *RCN-3D* is the top performer on unseen subjects, even though it is not the best on seen subjects.

4.1.1 Annotation error

The annotation error takes into account inaccuracies due to small differences of 6D sensor placement for different subjects during the annotation, and uncertainty in the wrist joint location. To quantify this error we selected poses for which all methods achieved a maximum error [51] of less than 10 mm. We denote these as *Easy Poses*. The pose estimation task for these can be considered solved, as shown in Figure 3. The outputs of the top five methods are visu-

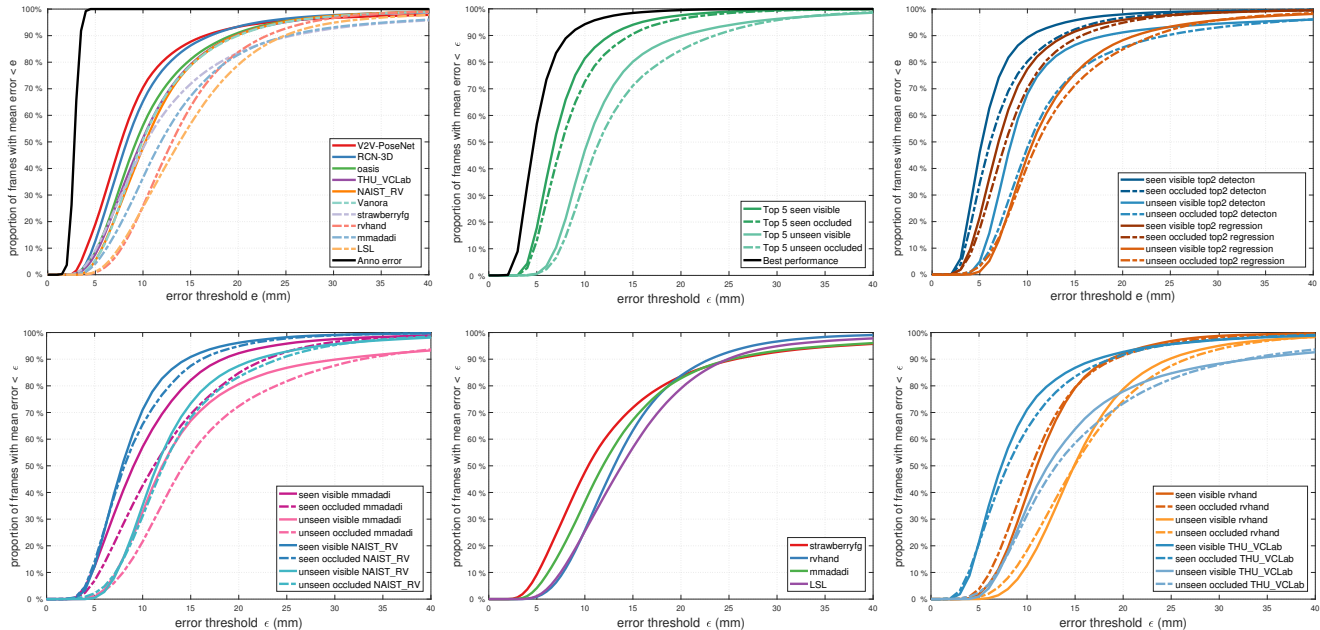


Figure 4: Success rates for different methods. **Top-left:** all evaluated methods using all test data. **Top-middle:** the average of the top five methods for four cases. **Top-right:** the average of top-two detection-based and regression-based methods in four cases. **Bottom-left:** direct comparison with 2D CNN and 3D CNN, *mmadadi* is a 2D CNN, *NAIST_RV* has the same structure but replaced 2D CNN with 3D CNN. **Bottom-middle:** comparison among structured methods. **Bottom-right:** comparison of a cascaded multistage method (*THU_VCLab*) and a one-off method (*rvhand*). Both of them use *REN* [15] as backbone.

ally accurate and close to the ground truth. We estimate the *Annotation Error* as the error on these poses, which has a mean value of 2.8 mm and a standard deviation of 0.5 mm.

4.1.2 Analysis by occlusion and unknown subject

Average error for four cases: To analyze the results with respect to joint visibility and hand shape, we partition the joints into four groups, based on whether or not they are visible, and whether or not the subject was seen at training time. Different hand shapes and joint occlusions are responsible for a large proportion of errors, see Table 3. The error for unseen subjects is significantly larger than for seen subjects. Moreover, the error for visible joints is smaller than for occluded joints. Based on the first group (visible, seen), we carry out a best-case performance estimate for the current state-of-the-art. For each frame of seen subjects, we first choose the best result from all methods, and calculate the success rate based on the average error for each frame, see the black curve in Figure 4 (top-middle).

2D vs. 3D CNNs: We compare two hierarchical methods with similar structure but different representation. The bottom-left plot of Figure 4 shows *mmadadi* [21], which employs a 2D CNN, and *NAIST_RV* [56], using a 3D CNN. *mmadadi* and *NAIST_RV* have almost the same structure, but *NAIST_RV* [56] uses a 3D CNN, while *mmadadi* [21] uses a 2D CNN. *NAIST_RV* [56] outperforms *mmadadi* [21] in all four cases.

Detection-based vs. regression-based methods: We compare the average of the top two detection-based methods with the average of the top two regression-based methods. In all four cases, detection-based methods outperform regression-based ones, see the top-right plot of Figure 4. In the challenge, the top two methods are detection-based methods, see Table 2. Note that a similar trend can be seen in the field of full human pose estimation, where only one method in a recent challenge was regression-based [25].

Hierarchical methods: Hierarchical constraints can help in the case of occlusion. The hierarchical model in *rvhand* [1] shows similar performance on visible and occluded joints. *rvhand* [1] has better performance on occluded joints when the error threshold is smaller than 15 mm, see the bottom-right plot of Figure 4. The underlying *REN* module [15], which includes finger and joint-type constraints seems to be critical. Methods using only per-finger constraints, e.g., *mmadadi* [21] and *NAIST_RV* [56], generalize less well to occluded joints, see the bottom-left plot of Figure 4.

Structural methods: We compare four structured methods *LSL* [20], *mmadadi* [21], *rvhand* [1], and *strawberryfg* [55], see the bottom-middle plot of Figure 4. *strawberryfg* [55] and *mmadadi* [21] have higher success rates when the error threshold is below 15 mm, while *LSL* [20] and *rvhand* [1] perform better for thresholds larger than 25 mm. Embedding structural constraints in the loss func-




Figure 5: Joint visibility. **Top:** Joint visibility distributions for training set and testing sets. **Bottom:** Average error (mm) for different numbers of visible joints and different methods.

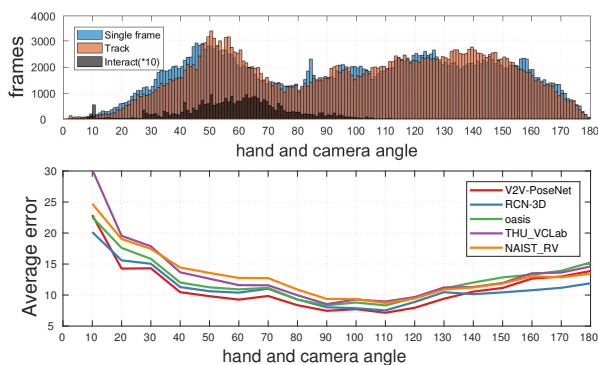


Figure 6: View point distributions. The error is significantly higher for small angles between hand and camera orientations.

tion has been more successful than including them within the CNN layers. *strawberryfg* [55] performs the best, using constraints on phalanges rather than on joints.

Single- vs. multi-stage methods: Cascaded methods work better than single-stage methods, see the bottom-right plot of Figure 4. Compared to other methods, *rvhand* [1] and *THU_VCLab* [3] both embed structural constraints, employing *REN* as their basic structure. *THU_VCLab* [3] takes a cascaded approach to iteratively update results from previous stages, outperforming *rvhand* [1].

4.1.3 Analysis by number of occluded joints

Most frames contain joint occlusions, see Figure 5 (top). We assume that a visible joint lies within a small range of the 3D point cloud. We therefore detect joint occlusion by thresholding the distance between the joint’s depth annotation value and its re-projected depth value. As shown in Figure 5 (bottom), the average error decreases nearly monotonously for increasing numbers of visible joints.

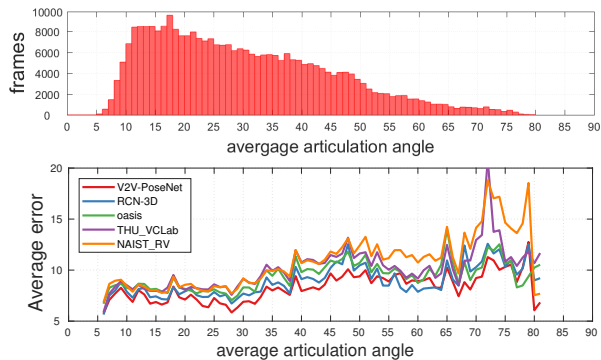


Figure 7: Articulation distribution. Errors increase for larger finger articulation angles.

4.1.4 Analysis based on view point

The view point is defined as the angle between the palm and camera directions. The test data covers a wide range of view points for the *Single frame pose estimation* task, see Figure 6 (top). View points in the [70, 120] range have a low mean error of below 10 mm. View points in the [0, 10] range have a significantly larger error due to the amount of self occlusion. View points in the [10, 30] range have an average error of 15–20 mm. View point ranges of [30, 70] and [120, 180] show errors of 10–15 mm. Third-person and egocentric views are typically defined by the hand facing toward or away from the camera, respectively. However, as shown in Figure 6, there is no clear separation by view point, suggesting a uniform treatment of both cases is sensible. Note that *RCN-3D* [23] outperforms others with a margin of 2-3 mm on extreme view points in the range of [150, 180] degrees due to their depth prediction stage.

4.1.5 Analysis based on articulation

We evaluate the effect of hand articulation on estimation accuracy, measured as the average of 15 finger flexion angles, see Figure 7. To reduce the influence from other factors such as view point, we select frames with view point angles within the range of [70, 120]. We evaluate the top five performers, see Figure 7 (bottom). For articulation angles smaller than 30 degrees, the mean error is 7 mm, when the average articulation angle increases to the range of [35, 70], errors increase to 9-10 mm. When the articulation angle is larger than 70 degrees, close to a fist pose, the mean error increases to over 12 mm.

4.1.6 Analysis by joint type

As before we group joints according to their visibility and the presence of the subject in the training set. We report the top five performers, see Figure 10. For the easiest case (visible joints of seen subjects), all 21 joints have a similar average error of 6–10 mm. For seen subjects, along the kinematic hand structure from the wrist to finger tips, occluded joints have increasingly larger errors, reaching 14 mm in the




Figure 8: Error curves for hand tracking. **Left:** number of frames with IoU below threshold. **Middle:** success rate for the detection part of each method in two cases. **Right:** success rate for three methods in four cases.




Figure 9: Error curves for hand-object interaction. **Left:** success rate for each method using average error per-frame. **Middle:** success rate for visible and occluded joints. **Right:** average error for each joint.




Figure 10: Average error of the top five methods for each joint in the Single frame pose estimation task. Finger tips have larger errors than other joints. For non-tip joints, joints on ring finger and middle finger have lower average errors than other fingers. ‘T’, ‘I’, ‘M’, ‘R’, ‘P’ denotes ‘Thumb’, ‘Index’, ‘Middle’, ‘Ring’, and ‘Pinky’ finger, respectively.

finger tips. Visible joints of unseen subjects have larger errors (10-13 mm) than that of seen subjects. Occluded joints of unseen subjects have the largest errors, with a relatively smaller error for the palm, and larger errors for finger tips (24 – 27 mm). We draw two conclusions: (1) all the top performers have difficulty in generalizing to hands from unseen subjects, (2) occlusions have more effect on finger tips than other joints. An interesting observation is that middle and ring fingers tend to have smaller errors in *MCP* and *PIP* joints than other fingers. One reason may be that the motion of these fingers is more restricted. The thumb’s *MCP* joint has a larger error than for other fingers, because it tends to have more discrepancy among different subjects.

Method	Model	AVG
<i>RCN-3D_track</i> [23]	scanning window + post-processing + pose estimator [23]	10.5
<i>NAIST_RV_track</i> [56]	hand detector [34] + hand verifier + pose estimator [56]	12.6
<i>THU_VCLab_track</i> [3]	Estimation [3] with the aid of tracking + re-initialization [33]	13.7

Table 4: Methods evaluated on 3D hand pose tracking. The last column is the average error in mm for all frames.

4.2. Hand pose tracking

In this task we evaluate three state-of-the-art methods, see Table 4 and Figure 8. Discriminative methods [3, 23, 56] break tracking into two sub-tasks: detection and hand pose estimation, sometimes merging the sub-tasks [3]. Based on the detection methods, 3D hand pose estimation can be grouped into pure tracking [23], tracking-by-detection [56], and a combination of tracking and re-initialization [3], see Table 4.

Pure tracking: *RCN-3D_track* estimate the bounding box location by scanning windows based on the result in the previous frame, including a motion estimate. Hand pose within the bounding box is estimated using *RCN-3D* [23].

Tracking-by-detection: *NAIST_RV_track* is a tracking-by-detection method with three components: hand detector, hand verifier, and pose estimator. The hand detector is built on U-net [34] to predict a binary hand-mask, which, after

Method	Model	AVG
NAIST_RV_obj [56]	Hand-object segmentation CNN + pose estimation [56]	25.0
THU_VCLab_obj [3]	Hand-object segmentation (intuitive) + pose estimation [3]	29.2
rvhand_obj [1]	Hand-object segmentation CNN + pose estimation [1]	31.3
RCN-3D_obj [23]	Two RCNs: Feature maps of first are used in the second RCN's stage 2.	32.4

Table 5: Methods evaluated on hand pose estimation during hand-object interaction. The last column is the average error (mm) for all frames.

verification, is passed to the pose estimator *NAIST_RV* [56]. If verification fails, the result from the previous frame is chosen.

Hybrid tracking and detection: *THU_VCLab_track* [3] makes use of the previous tracking result and the current frame’s scanning window. The hand pose of the previous frame is used as a guide to predict the hand pose in the current frame. The previous frame’s bounding box is used for the current frame. During fast hand motion, Faster R-CNN [33] is used for re-initialization.

Detection accuracy: We first evaluate the detection accuracy by evaluating the bounding box overlap, *i.e.*, the intersection over union (IoU) of the detection and ground truth bounding boxes, see Figure 8 (middle). Overall, *RCN-3D_track* is more accurate than *THU_VCLab_track*, which itself outperforms *NAIST_RV_track*. Pure detection methods have a larger number of false negatives, especially when multiple hands appear in the scene, see Figure 8 (left). There are 72 and 174 missed detections (IoU of zero), for *NAIST_RV_track* and *THU_VCLab_track*, respectively. By tracking and re-initializing, *THU_VCLab_track* achieves better detection accuracy overall. *RCN-3D_track*, using motion estimation and single-frame hand pose estimation, shows the lowest error.

Tracking accuracy is shown in Figure 8 (right). Even through *THU_VCLab* performs better than *NAIST_RV* in the *Single frame pose estimation* task, *NAIST_RV_track* performs better on the tracking tasks due to per-frame hand detection.

4.3. Hand object interaction

For this task, we evaluate four state-of-the-art methods, see Table 5 and Figure 9. Compared to the other two tasks there is significantly more occlusion, see Figure 5 (top). Methods explicitly handling occlusion achieve higher accuracy with errors in the range of 25–29 mm: (1) *NAIST_RV_obj* [56] and *rvhand_obj* [1] segment the hand area from the object using a network. (2) *THU_VCLab_obj* [3] removes the object region from cropped hand images with image processing operations [36]. (3) *RCN-3D_obj* [23] modify their original network to infer the depth values of 2D keypoint locations.

Current state-of-the-art methods have difficulty generalizing to the hand-object interaction scenario. However, *NAIST_RV_obj* [56] and *rvhand_obj* [1] show similar performance for visible joints and occluded joints, indicating that CNN-based segmentation can better preserve structure than image processing operations, see the middle plot of Figure 9.

5. Discussion and conclusions

The analysis of the top 10 among 17 participating methods from the *HIM2017* challenge [58] provides insights into the current state of 3D hand pose estimation.

(1) 3D volumetric representations used with a 3D CNN show high performance, possibly by better capturing the spatial structure of the input depth data.

(2) Detection-based methods tend to outperform regression-based methods, however, regression-based methods can achieve good performance using explicit spatial constraints. Making use of richer spatial models, *e.g.*, bone structure [42], helps further. Regression-based methods perform better in extreme view point cases [23], where severe occlusion occurs.

(3) While joint occlusions pose a challenge for most methods, explicit modeling of structure constraints and spatial relation between joints can significantly narrow the gap between errors on visible and occluded joints [1].

(4) Discriminative methods still generalize poorly to unseen hand shapes. Data augmentation and scale estimation methods model only global shape changes, but not local variations. Integrating hand models with better generative capability may be a promising direction.

(5) Isolated 3D hand pose estimation achieves low mean errors (10 mm) in the view point range of [70, 120] degrees. However, errors remain large for extreme view points, *e.g.*, view point range of [0,10], where the hand is facing away from the camera. Multi-stage methods [23] tend to perform better in these cases.

(6) In hand tracking, current discriminative methods divide the problem into two sub-tasks: detection and pose estimation, without using the hand shape provided in the first frame. Hybrid methods may work better by using the provided hand shape.

(7) Current methods perform well on single hand pose estimation when trained on a million-scale dataset, but have difficulty in generalizing to hand-object interaction. Two directions seem promising, (a) designing better hand segmentation methods, and (b) training the model with large datasets containing hand-object interaction.

Acknowledgement: This work was partially supported by Huawei Technologies.

References

- [1] K. Akiyama, F. Yang, and Y. Wu. Naist rvlab g2's solution for 2017 hand challenge. *Hands 2017*. 2, 3, 5, 6, 8
- [2] S. Baek, K. I. Kim, and T.-K. Kim. Augmented skeleton space transfer for depth-based hand pose estimation. In *CVPR, 2018*. 1
- [3] X. Chen, G. Wang, H. Guo, and C. Zhang. Pose guided structured region ensemble network for cascaded hand pose estimation. *arXiv preprint arXiv:1708.03416*, 2017. 2, 3, 6, 7, 8
- [4] C. Choi, S. Kim, and K. Ramani. Learning hand articulations by hallucinating heat distribution. In *ICCV, 2017*. 1
- [5] C. Choi, S. H. Yoon, C.-N. Chen, and K. Ramani. Robust hand pose estimation during the interaction with an unknown object. In *ICCV, 2017*. 1
- [6] N. Cihan Camgoz, S. Hadfield, O. Koller, and R. Bowden. Subunets: End-to-end hand shape and continuous sign language recognition. In *ICCV, 2017*. 1
- [7] X. Deng, S. Yang, Y. Zhang, P. Tan, L. Chang, and H. Wang. Hand3d: Hand pose estimation using 3d neural network. *arXiv preprint arXiv:1704.02224*, 2017. 1
- [8] A. Erol, G. Bebis, M. Nicolescu, R. D. Boyle, and X. Twombly. Vision-based hand pose estimation: A review. *CVIU*, 2007. 2
- [9] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *IJCV*, 2015. 2
- [10] G. Garcia-Hernando, S. Yuan, S. Baek, and T.-K. Kim. First-person hand action benchmark with rgb-d videos and 3d hand pose annotations. In *CVPR, 2018*. 1, 2
- [11] L. Ge, Y. Cai, J. Weng, and J. Yuan. Hand PointNet: 3d hand pose estimation using point sets. In *CVPR, 2018*. 2, 3
- [12] L. Ge, H. Liang, J. Yuan, and D. Thalmann. 3d convolutional neural networks for efficient and robust hand pose estimation from single depth images. In *CVPR, 2017*. 2
- [13] L. Ge, H. Liang, J. Yuan, and D. Thalmann. Robust 3D hand pose estimation in single depth images: from single-view cnn to multi-view cnns. In *CVPR, 2016*. 2
- [14] L. Ge and J. Yuan. 2017 hand challenge imi ntu team. *Hands 2017*. 2, 3
- [15] H. Guo, G. Wang, X. Chen, and C. Zhang. Towards good practices for deep 3d hand pose estimation. *arXiv preprint arXiv:1707.07248*, 2017. 2, 3, 4, 5
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR, 2016*. 3
- [17] S. Honari, P. Molchanov, S. Tyree, P. Vincent, C. Pal, and J. Kautz. Improving landmark localization with semi-supervised learning. In *CVPR, 2018*. 2, 3
- [18] S. Honari, J. Yosinski, P. Vincent, and C. Pal. Recombinator networks: Learning coarse-to-fine feature aggregation. In *CVPR, 2016*. 2
- [19] M. Kristan, J. Matas, A. Leonardis, M. Felsberg, L. Cehovin, G. Fernández, T. Vojir, G. Hager, G. Nebehay, and R. Pflugfelder. The visual object tracking vot2015 challenge results. In *ICCV Workshop 2015*. 2
- [20] S. Li and D. Lee. 2017 hand challenge/frame-based/team hcr: Method description. *Hands 2017*. 2, 3, 5
- [21] M. Madadi, S. Escalera, X. Baro, and J. Gonzalez. End-to-end global to local cnn learning for hand pose recovery in depth data. *arXiv preprint arXiv:1705.09606*, 2017. 2, 3, 5
- [22] M. Madadi, S. Escalera, A. Carruesco, C. Andujar, X. Baró, and J. González. Occlusion aware hand pose recovery from sequences of depth images. In *FG, 2017*. 1
- [23] P. Molchanov, J. Kautz, and S. Honari. 2017 hand challenge nvresearch and umontreal team. *Hands 2017*. 2, 3, 6, 7, 8
- [24] G. Moon, J. Y. Chang, and K. M. Lee. V2v-posesnet: Voxel-to-voxel prediction network for accurate 3d hand and human pose estimation from a single depth map. In *CVPR, 2018*. 2, 3
- [25] MPII Leader Board. <http://human-pose.mpi-inf.mpg.de>. 5
- [26] F. Mueller, D. Mehta, O. Sotnychenko, S. Sridhar, D. Casas, and C. Theobalt. Real-time hand tracking under occlusion from an egocentric rgb-d sensor. In *ICCV, 2017*. 1
- [27] M. Oberweger and V. Lepetit. Deeprior++: Improving fast and accurate 3d hand pose estimation. In *ICCV Workshop, 2017*. 3, 4
- [28] M. Oberweger, G. Riegler, P. Wohlhart, and V. Lepetit. Efficiently creating 3D training data for fine hand pose estimation. In *CVPR, 2016*. 1
- [29] M. Oberweger, P. Wohlhart, and V. Lepetit. Hands deep in deep learning for hand pose estimation. In *CVWW, 2015*. 2, 3, 4
- [30] I. Oikonomidis, N. Kyriazis, and A. A. Argyros. Efficient model-based 3D tracking of hand articulations using kinect. In *BMVC*, 2011. 3, 4
- [31] C. Qian, X. Sun, Y. Wei, X. Tang, and J. Sun. Realtime and robust hand tracking from depth. In *CVPR, 2014*. 1
- [32] E. Remelli, A. Tkach, A. Tagliasacchi, and M. Pauly. Low-dimensionality calibration through local anisotropic scaling for robust hand model personalization. In *ICCV, 2017*. 1
- [33] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS, 2015*. 7, 8
- [34] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI, 2015*. 7
- [35] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 2015. 2
- [36] J. Serra. *Image analysis and mathematical morphology*, v. 1. Academic press, 1982. 8
- [37] T. Sharp, C. Keskin, D. Robertson, J. Taylor, J. Shotton, D. Kim, C. Rhemann, I. Leichter, A. Vinnikov, Y. Wei, D. Freedman, P. Kohli, E. Krupka, A. Fitzgibbon, and S. Izadi. Accurate, robust, and flexible real-time hand tracking. In *CHI, 2015*. 1, 3, 4
- [38] T. Simon, H. Joo, I. Matthews, and Y. Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *CVPR, 2017*. 1
- [39] A. Sinha, C. Choi, and K. Ramani. Deephand: Robust hand pose estimation by completing a matrix imputed with deep features. In *CVPR, 2016*. 3

- [40] S. Song and J. Xiao. Deep sliding shapes for amodal 3d object detection in rgb-d images. In *CVPR, 2016*. 2
- [41] S. Sridhar, A. Oulasvirta, and C. Theobalt. Interactive markerless articulated hand motion tracking using rgb and depth data. In *ICCV, 2013*. 1
- [42] X. Sun, J. Shang, S. Liang, and Y. Wei. Compositional human pose regression. In *ICCV, 2017*. 3, 8
- [43] X. Sun, Y. Wei, S. Liang, X. Tang, and J. Sun. Cascaded hand pose regression. In *CVPR, 2015*. 1
- [44] J. S. Supancic, G. Rogez, Y. Yang, J. Shotton, and D. Ramanan. Depth-based hand pose estimation: methods, data, and challenges. In *ICCV, 2015*. 1, 2
- [45] A. Tagliasacchi, M. Schröder, A. Tkach, S. Bouaziz, M. Botsch, and M. Pauly. Robust articulated-icp for real-time hand tracking. *Computer Graphics Forum*, 2015. 1
- [46] D. Tang, H. J. Chang, A. Tejani, and T.-K. Kim. Latent regression forest: Structured estimation of 3D articulated hand posture. In *CVPR, 2014*. 1
- [47] D. Tang, H. J. Chang, A. Tejani, and T.-K. Kim. Latent regression forest: structured estimation of 3d hand poses. *TPAMI*, 2017. 1
- [48] D. Tang, J. Taylor, P. Kohli, C. Keskin, T.-K. Kim, and J. Shotton. Opening the black box: Hierarchical sampling optimization for estimating human hand pose. In *ICCV, 2015*. 1
- [49] D. Tang, T.-H. Yu, and T.-K. Kim. Real-time articulated hand pose estimation using semi-supervised transductive regression forests. In *ICCV, 2013*. 1
- [50] J. Taylor, L. Bordeaux, T. Cashman, B. Corish, C. Keskin, E. Soto, D. Sweeney, J. Valentin, B. Luff, A. Topalian, E. Wood, S. Khamis, P. Kohli, T. Sharp, S. Izadi, R. Banks, A. Fitzgibbon, and J. Shotton. Efficient and precise interactive hand tracking through joint, continuous optimization of pose and correspondences. *TOG*, 2016. 2
- [51] J. Taylor, J. Shotton, T. Sharp, and A. Fitzgibbon. The vitruvian manifold: Inferring dense correspondences for one-shot human pose estimation. In *CVPR, 2012*. 3, 4
- [52] J. Tompson, M. Stein, Y. Lecun, and K. Perlin. Real-time continuous pose recovery of human hands using convolutional networks. *TOG*, 2014. 1
- [53] C. Wan, T. Probst, L. Van Gool, and A. Yao. Crossing nets: Combining gans and vaes with a shared latent space for hand pose estimation. In *CVPR, 2017*. 1
- [54] C. Wan, A. Yao, and L. Van Gool. Hand pose estimation from local surface normals. In *ECCV, 2016*. 1
- [55] Q. Wan. 2017 hand challenge fudan university team. *Hands 2017*. 3, 5, 6
- [56] F. Yang, K. Akiyama, and Y. Wu. Naist rv’s solution for 2017 hand challenge. *Hands 2017*. 2, 3, 5, 7, 8
- [57] Q. Ye, S. Yuan, and T.-K. Kim. Spatial attention deep net with partial PSO for hierarchical hybrid hand pose estimation. In *ECCV, 2016*. 1, 2
- [58] S. Yuan, Q. Ye, G. Garcia-Hernando, and T.-K. Kim. The 2017 hands in the million challenge on 3d hand pose estimation. In *arXiv preprint arXiv:1707.02237*, 2017. 1, 2, 8
- [59] S. Yuan, Q. Ye, B. Stenger, S. Jain, and T.-K. Kim. Big-hand2.2m benchmark: Hand pose dataset and state of the art analysis. In *CVPR, 2017*. 1, 2
- [60] J. Zhang, J. Jiao, M. Chen, L. Qu, X. Xu, and Q. Yang. 3d hand pose tracking and estimation using stereo matching. In *arXiv, 2017*. 1
- [61] X. Zhou, Q. Wan, W. Zhang, X. Xue, and Y. Wei. Model-based deep hand pose estimation. In *IJCAI, 2016*. 2, 3, 4
- [62] C. Zimmermann and T. Brox. Learning to estimate 3d hand pose from single rgb images. In *ICCV, 2017*. 1