

Measuring Generalisation to Unseen Viewpoints, Articulations, Shapes and Objects for 3D Hand Pose Estimation under Hand-Object Interaction

Anil Armagan¹, Guillermo Garcia-Hernando^{1,2}, Seungryul Baek^{1,20}, Shreyas Hampali³, Mahdi Rad³, Zhaohui Zhang⁴, Shipeng Xie⁴, MingXiu Chen⁴, Boshen Zhang⁵, Fu Xiong⁶, Yang Xiao⁶, Zhiguo Cao⁶, Junsong Yuan⁷, Pengfei Ren⁸, Weiting Huang⁸, Haifeng Sun⁸, Marek Hrúz⁹, Jakub Kanis⁹, Zdeněk Krňoul⁹, Qingfu Wan¹⁰, Shile Li¹¹, Linlin Yang¹², Dongheui Lee¹¹, Angela Yao¹³, Weiguo Zhou¹⁴, Sijia Mei¹⁴, Yunhui Liu¹⁵, Adrian Spurr¹⁶, Umar Iqbal¹⁷, Pavlo Molchanov¹⁷, Philippe Weinzaepfel¹⁸, Romain Brégier¹⁸, Gregory Rogez¹⁸, Vincent Lepetit^{3,19}, and Tae-Kyun Kim^{1,21}

Abstract. We study how well different types of approaches generalise in the task of 3D hand pose estimation under single hand scenarios and hand-object interaction. We show that the accuracy of state-of-the-art methods can drop, and that they fail mostly on poses absent from the training set. Unfortunately, since the space of hand poses is highly dimensional, it is inherently not feasible to cover the whole space densely, despite recent efforts in collecting large-scale training datasets. This sampling problem is even more severe when hands are interacting with objects and/or inputs are RGB rather than depth images, as RGB images also vary with lighting conditions and colors. To address these issues, we designed a public challenge (HANDS’19) to evaluate the abilities of current 3D hand pose estimators (HPEs) to interpolate and extrapolate the poses of a training set. More exactly, HANDS’19 is designed (a) to evaluate the influence of both depth and color modalities on 3D hand pose estimation, under the presence or absence of objects; (b) to assess the generalisation abilities *w.r.t.* four main axes: shapes, articulations, viewpoints, and objects; (c) to explore the use of a synthetic hand models to fill the gaps of current datasets. Through the challenge, the overall accuracy has dramatically improved over the baseline, especially on extrapolation tasks, from 27mm to 13mm mean joint error. Our analyses highlight the impacts of: Data pre-processing, ensemble approaches, the use of the MANO model, and different HPE methods/backbones.

¹Imperial College London, ²Niantic, Inc., ³Graz Uni. of Technology, ⁴Rokid Corp. Ltd., ⁵HUST, ⁶Megvii Research Nanjing, ⁷SUNY Buffalo, ⁸BUPT, ⁹Uni. of West Bohemia, ¹⁰Fudan Uni., ¹¹TUM, ¹²Uni. of Bonn, ¹³NUS, ¹⁴Harbin Inst. of Technology, ¹⁵CUHK, ¹⁶ETH Zurich, ¹⁷NVIDIA Research, ¹⁸NAVER LABS Europe, ¹⁹ENPC ParisTech, ²⁰UNIST, ²¹KAIST.

1 Introduction

3D hand pose estimation is crucial to many applications including natural user-interaction in AR/VR, robotics, teleoperation, and healthcare. The recent successes primarily come from large-scale training sets [40], deep convolutional neural networks [10, 22], and fast optimisation for model fitting [15, 23]. State-of-the-art methods now deliver satisfactory performance for viewpoints seen at training time and single hand scenarios. However, as we will show, these methods substantially drop accuracy when applied to egocentric viewpoints for example, and in the presence of significant foreground occlusions. These cases are not well represented on the training sets of existing benchmarks [5, 20, 21]. The challenges become even more severe when we consider RGB images and hand-object interaction scenarios. These issues are well aligned with the observations from the former public challenge HANDS'17 [39]: The state-of-the-art methods dropped accuracy from frontal to egocentric views, and from open to closure hand postures. The average accuracy was also significantly lower under hand-object interaction [5].

Given the difficulty to interpolate and extrapolate poses from the training set, one may opt for creating even larger training sets. Unfortunately, an inherent challenge in 3D hand pose estimation is the very high dimensionality of the problem, as hand poses and shapes and camera viewpoints have a large number of degrees-of-freedom that can vary independently. This complexity increases even more when we consider the case of a hand manipulating an object. Despite

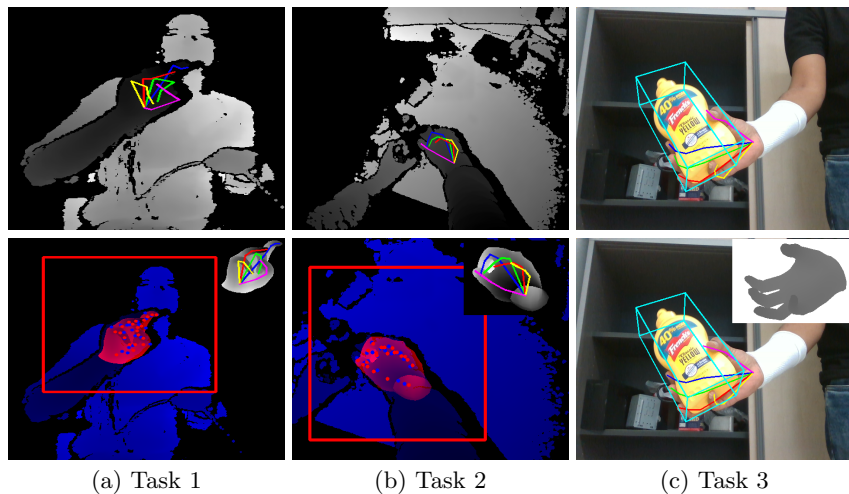


Fig. 1: Frames from the three tasks of our challenge. For each task, we show the input depth or RGB image with the ground-truth hand skeleton (top) and a rendering of the fitted 3D hand model as well as a depth rendering of the model (bottom). The ground-truth and estimated joint locations are shown in blue and red respectively.

the recent availability of large-scale datasets [40], and the development of complex calibrated multi-view camera systems to help the annotation of synthetic data [13, 28, 44], capturing a training set that covers completely the domain of the problem remains extremely challenging.

In this work, we therefore study in depth the ability of current methods to interpolate and extrapolate the training set, and how this ability can be improved. To evaluate this ability, we consider the three tasks depicted in Fig. 1, which vary the input (depth and RGB images) or the camera viewpoints, and introduce the possible manipulation of an object by the hand. We carefully designed training and testing sets in order to evaluate the generalisation performance to unseen viewpoints, articulations, and shapes of the submitted methods.

HANDS’19 fostered dramatic accuracy improvement compared to a provided baseline, which is a ResNet-50 [10]-based 3D joint regressor trained on our training set, from **27mm** to **13mm**. This paper provides an in-depth analysis of the different factors that made this improvement possible.

2 HANDS 2019 Challenge Overview

The challenge consists of three different tasks, in which the goal is to predict the 3D locations of the hand joints given an image. For training, images, hand pose annotations, and a 3D parametric hand model [26] for synthesizing data are provided. For inference, only the images and bounding boxes of the hands are given to the participants. These tasks are defined as follows:

Task 1: Depth-Based 3D Hand Pose Estimation: This task builds on Big-Hand2.2M [40] dataset, as for the HANDS 2017 challenge [38]. No objects appear in this task. Hands appear in both third person and egocentric viewpoints.

Task 2: Depth-Based 3D Hand Pose Estimation while Interacting with Objects: This task builds on the F-PHAB dataset [5]. The subject manipulates objects with their hand, as captured from an egocentric viewpoint. Some object models are provided by [5].

Task 3: RGB-Based 3D Hand Pose Estimation while Interacting with Objects: This task builds on the HO-3D [8] dataset. The subject manipulates objects with their hand, as captured from a third person viewpoint. The objects are used from the YCB dataset [35]. The ground truth wrist position of the test images is also provided in this task.

The BigHand2.2M [40] and F-PHAB [5] datasets have been used by 116 and 123 unique institutions to date. HANDS’19 received 80 requests to access the datasets with the designed partitions, and 17, 10 and 9 participants have evaluated their methods on Task 1, Task 2 and Task 3, respectively.

3 Evaluation Criteria

We evaluate the generalisation power of HPEs in terms of 4 "axes": Viewpoint, Articulation, Shape, and Object. For each axis, frames within a dataset are automatically annotated by using the ground-truth 3D joint locations and the object

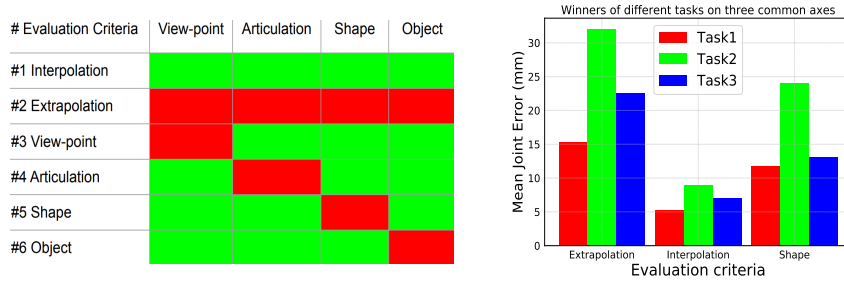


Fig. 2: Left: The 6 evaluation criteria used in the challenge. For each axis (View-point, Articulation, Shape, Object), we indicate if hand poses from the frames in an evaluation criterion are also available (green) in the training set or not (red). Right: MJE comparison of the best methods for the Extrapolation, Interpolation, and Shape criteria, on each task.

information to annotate each frame in each axis. The annotation distribution of the dataset for each axis are used to create a training and a test set. Using the frame annotations on each axis, the sets are sampled in a structured way to have the test frames that are similar to the frames in the training data (for interpolation) and also the test frames where axes’ annotations are never seen in the training data (for extrapolation). More details on the dataset are given in Section 4. To measure the generalisation of HPEs, 6 evaluation criteria are further defined with the 4 main axes: **Viewpoint**, **Articulation**, **Shape** and **Object** are respectively used for measuring the extrapolation performance of HPEs on the frames with articulation cluster, viewpoint angle, hand shape and object type (axis annotations) that are not present in the training set. **Extrapolation** is used to measure the performance on the frames with axis annotations that do not overlap/present in the training set. Lastly, **Interpolation** is defined to measure the performance on the frames with the axis annotations present in the training set.

The challenge uses the mean joint error (MJE) [23] as the main evaluation metric. Results are ranked according to the **Extrapolation** criterion which measures the total extrapolation power of the approaches with MJE on all axes. We also consider success rates based on maximum allowed distance errors for each frame and each joint for further analysis.

Fig. 2 (left) summarises the 6 evaluation strategies, and Fig. 2 (right) shows the accuracies obtained by the best approaches, measured for the three evaluation criteria that could be evaluated for all 3 tasks. Articulation and viewpoint criteria are only considered for Task 1 since the joint angles are mostly fixed during object interaction and hence the Articulation criteria is not as meaningful as in Task 1 for the other tasks. The Viewpoint criteria is not meaningful for Task 2 which is for egocentric views since the task’s dataset constrains the relative palm-camera angle to a small range. For Task 3, the data scarcity is not helping to sample enough diverse viewpoints. The extrapolation errors tend to be three

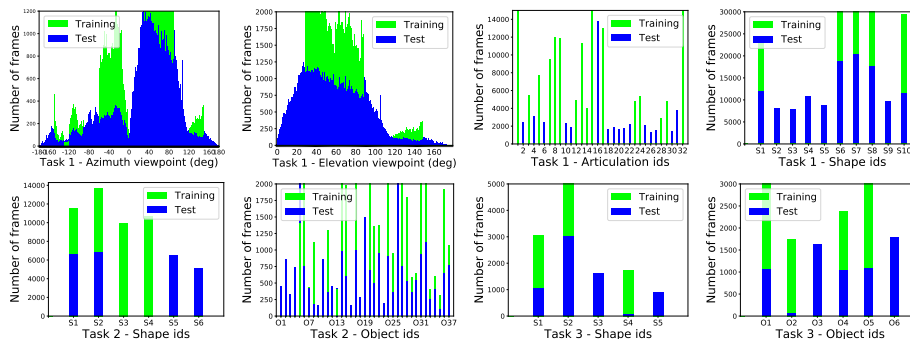


Fig. 3: Distributions of the training and test datasets for Task 1 (top), Task 2 (bottom left), and Task 3 (bottom right). The splits are used to evaluate the extrapolation power of the approaches and decided based on the viewpoints, the articulation clusters of the hand pose, the hand shape, and the type of the object present.

times larger than the interpolation errors while the shape is a bottleneck among the other attributes. Lower errors on Task 3 compared to Task 2 are likely due to the fact that the ground truth wrist position is provided for Task 3.

4 Datasets

Given a task, the training set is the same and the test frames used to evaluate each criterion can be different or overlapped. The number of training frames are 175k, 45k and 10k for Task 1, 2 and 3 respectively. The sizes of the test sets for each evaluation criterion are shown in Table 1.

Fig. 3 shows the distributions of the training and test data for each task. The viewpoints are defined as elevation and azimuth angles of the hand *w.r.t.* the camera using the ground-

Table 1: Detailed analytics on the number of frames provided on the training and test sets for the different tasks.

Dataset	Task id	#Frames							#Subjects	#Objects	#Actions	#Seq.
		Total	Ext.	Int.	Art.	View.	Sha.	Obj.				
Test	1	125K	20%	16%	16%	32%	16%	✗	10	✗	✗	✗
	2	25K	14%	32%	✗	✗	37%	17%	4	37	71	292
	3	6.6K	24%	35%	✗	✗	14	27%	5	5	1	5
Training	1								5	✗	✗	✗
	2								4	26	45	539
	3								3	4	1	12

truth joint annotations. The articulation of the hand is defined and obtained by clustering on the ground-truth joint angles in a fashion similar to [18], by using binary representations (open/closed) of each finger *e.g.* '00010' represents a hand articulation cluster with frames with the index finger closed and the rest of the fingers open, which ends up with $2^5 = 32$ clusters. A visualization of the articulation clusters is provided in the supplementary document. Note that the use of low-dimensional embedding such as PCA or t-SNE is not adequate

here to compare the two data distributions, because the dimensionality of the distributions is very high and a low-dimensional embedding would not be very representative. Fig. 3 further shows the repartition in terms of subjects/shapes, where five seen subjects and five unseen subjects are present. Similarly, the data partition was done on objects. This way we can control the data to define the evaluation metrics.

Use of 3D Hand Models for HPEs. A series of methods [1, 3, 7, 9, 41] have been proposed in the literature to make use of 3D hand models for supervision of HPEs. Ge et al. [7] proposed to use Graph CNNs for mapping RGB images to infer the vertices of 3D meshes. Hasson et al. [9] jointly infers both hands and object meshes and investigated the effect of the 3D contact loss penalizing the penetration of object and hand surfaces. Others [1, 3, 41] attempted to make use of MANO [26], a parametric 3D hand model by learning to estimate low-dimensional PCA parameters of the model and using it together with differentiable model renderers for 3D supervision. All the previous works on the use of 3D models in learning frameworks have shown to help improving performance on the given task. Recently, [16] showed that fitting a 3D body model during the estimation process can be accelerated by using better initialization of the model parameters however, our goal is slightly different since we aim to explore the use of 3D models for better generalisation from the methods. Since the hand pose space is huge, we make use of a 3D hand model to fill the gaps in the training data distribution to help approaches to improve their extrapolation capabilities. In this study, we make use of the MANO [26] hand model by providing the model’s parameters for each training image. We fit the 3D model for each image in an optimization-based framework which is described in more details below.

Gradient-based Optimization for Model Fitting. We fit the MANO [26] models’ shape $\mathbf{s} = \{s_j\}_{j=1}^{10}$, camera pose $\mathbf{c} = \{c_j\}_{j=1}^8$, and articulation $\mathbf{a} = \{a_j\}_{j=1}^{45}$ parameters to the i -th raw skeletons of selected articulations $\mathbf{z} = \{z_i\}_{i=1}^K$, by solving the following equation:

$$(\mathbf{s}^{i*}, \mathbf{c}^{i*}, \mathbf{a}^{i*}) = \arg \min_{(\mathbf{s}, \mathbf{c}, \mathbf{a})} O(\mathbf{s}, \mathbf{c}, \mathbf{a}, \mathbf{z}^i), \forall i \in [1, K], \quad (1)$$

where our proposed objective function $O(\mathbf{s}, \mathbf{c}, \mathbf{a}, \mathbf{z}^i)$ for the sample i is defined as follows:

$$O(\mathbf{s}, \mathbf{c}, \mathbf{a}, \mathbf{z}^i) = \|f^{reg}(V(\mathbf{s}, \mathbf{c}, \mathbf{a})) - \mathbf{z}^i\|_2^2 + \sum_{j=1}^{10} \|\mathbf{s}_j\|_2^2 + R_{Lap}(V(\mathbf{s}, \mathbf{c}, \mathbf{a})). \quad (2)$$

$V(\mathbf{s}, \mathbf{c}, \mathbf{a})$ denotes the 3D mesh as a function of the three parameters $\mathbf{s}, \mathbf{c}, \mathbf{a}$. Eq. (2) is composed of the following terms: *i*) the Euclidean distance between 3D skeleton ground-truths \mathbf{z}^i and the current MANO mesh model’s 3D skeleton values $f^{reg}(V(\mathbf{s}, \mathbf{c}, \mathbf{a}))$ ¹; *ii*) A shape regularizer enforcing the shape parameters \mathbf{s} to be close to their MANO model’s mean values, normalized to 0 as in [26], to maximize the shape likelihood; and *iii*) A Laplacian regularizer $R_{Lap}(V(\mathbf{s}, \mathbf{c}, \mathbf{a}))$ to obtain the smooth mesh surfaces as in [14]. Eq. (1) is solved iteratively by

¹ f^{reg} geometrically regresses the skeleton from the mesh vertex coordinates. It is provided with the MANO model and the weights are fixed during the process.




Fig. 4: Depth renderings of the hand model for different iterations in gradient-based optimization fitting. Target image (joints) (a), optimization iterations 0, 100, 300, 400, 600, 700 (b), final fitted hand pose at iteration 3000 (c).

using the gradients from Eq. (2) as follows:

$$(\mathbf{s}_{t+1}, \mathbf{c}_{t+1}, \mathbf{a}_{t+1}) = (\mathbf{s}_t, \mathbf{c}_t, \mathbf{a}_t) - \gamma \cdot \nabla O(\mathbf{s}_t, \mathbf{c}_t, \mathbf{a}_t, \mathbf{z}^i), \forall t \in [1, T], \quad (3)$$

where $\gamma = 10^{-3}$ and $T = 3000$ are empirically set. This process is similar to the refinement step of [32, 1], which refines estimated meshes by using the gradients from the loss. In Fig. 4, both the target and the fitted depth images during the process described by Eq. (3) are depicted. Minor errors of the fitting are not a problem for our purpose given that we will generate input and output pairs of the fitted model by exploiting fitted meshes’ self-data generation capability while ignoring original depth and skeletons. Here the aim of fitting the hand model is to obtain a plausible and a complete articulation space. The model is fitted without optimizing over depth information from the model and the input depth image since we did not observe an improvement on the parameter estimation. Moreover, the optimization needs to be constrained to produce plausible hand shapes and noise and other inconsistencies may appear in the depth image.

5 Evaluated Methods

In this section, we present the gist of selected 14 methods among 36 participants (17 for Task 1, 10 for Task 2, 9 for Task 3) to further analyze their results in Section 6. Methods are categorized based on their main components and properties. See Tables 1, 2 and 3 of the supplementary document provided with this work for a glance of the properties of the methods in HANDS’19.

2D and 3D supervision for HPEs. Approaches that embed and process 3D data obtain high accuracies but less efficient [39] in terms of their complexity compared to 2D-based approaches. 3D-based methods use 3D convolutional layers for point-clouds input similar to *NTIS* which uses an efficient voxel-based representation V2V-PoseNet [19] with a deeper architecture and weighted sub-voxel predictions on quarter of each voxel representations for robustness. Some other approaches adopts 3D as a way of supervision similar to *Strawberryfg* [33] which employs a render-and-compare stage to enforce voxel-wise supervision for

model training and adopts a 3D skeleton volume renderer to re-parameterize an initial pose estimate obtained similar to [30]. *BT* uses a permutation invariant feature extraction layer [17] to extract point-cloud features and uses a two branch framework for point-to-pose voting and point-to-latent voting. 3D supervision is employed by point-cloud reconstruction from a latent embedding in Task 1 whereas 3D hand model parameters are estimated and used in a differentiable model renderer for 3D supervision for the other tasks.

2D CNN-based approaches has been a standard way for learning regression models as used by *Rokid* [42] where they adopt a two stage regression models. The first regression model is used to predict an initial pose and the second model built on top of the first model. *A2J* [36] uses a 2D supervised method based on 2D offset and depth estimations with anchor points. Anchor points are densely set on the input image to behave as local regressors for the joints and able to capture global-local spatial context information. *AWR* [11] adopts a learnable and adaptive weighting operation that is used to aggregate spatial information of different regions in dense representations with 2D convolutional CNNs. The weighting operation adds direct supervision on joint coordinates and draw consensus between the training and inference as well as enhancing the model’s accuracy and generalisation ability by adaptively aggregating spatial information from related regions. *CrazyHand* uses a hierarchically structured regression network by following the joints’ distribution on the hand morphology. *ETH_NVIDIA* adopts the latent 2.5D heatmap regression [12]; additionally an MLP is adopted for denoising the absolute root depth. Absolute 3D pose in scale-normalized space is obtained with the pinhole camera equations. *NLE* [25] first performs a classification of the hand into a set of canonical hand poses (obtained by clustering on the poses in the training set), followed by a fine class-specific regression of the hand joints in 2D and 3D. *NLE* adopts the only approach proposing multiple hand poses in a single stage with a Region Proposal Network (RPN) [24] integration.

Detection, regression and combined HPEs. Detection methods are based on hand key-points and producing a probability density maps for each joint. *NTIS* uses a 3D CNN [19] to estimate per-voxel likelihood of each joint. Regression-based methods estimate the joint locations by learning a direct mapping from the input image to hand joint locations or the joint angles of a hand model [29, 43]. *Rokid* uses joint regression models within two stages to estimate an initial hand pose for hand cropping and estimates the final pose from the cleaned hand image. *A2J* adopts regression framework by regressing offsets from anchors to final joint location. *BT*’s point-wise features are used in a voting scheme which behaves as a regressor to estimate the pose.

Some approaches take advantage of both detection-based and regression-based methods. Similarly, *AWR*, *Strawberryfg* estimates probability maps to estimate joint locations with a differentiable *soft-argmax* operation [30]. A hierarchical approach proposed by *CrazyHand* regresses the joint locations from joint probability maps. *ETH_NVIDIA* estimates 2D joint locations from estimated probability maps and regresses relative depth distance of the hand joints *w.r.t.* a root joint.

NLE first localizes the hands and classifies them to anchor poses and the final pose is regressed from the anchors.

Method-wise ensembles. *A2J* uses densely set anchor points in a voting stage which helps to predict location of the joints in an ensemble way for better generalisation leveraging the uncertainty in reference point detection. In a similar essence, *AWR* adaptively aggregates the predictions from different regions and *Strawberryfg* adopts local patch refinement [34] where refinement models are adopted to refine bone orientations. *BT* uses the permutation equivariant features extracted from the point-cloud in a point-to-pose voting scheme where the votes are ensembled to estimate the pose. *NLE* ensembles anchor poses to estimate the final pose.

Ensembles in post-processing. Rather than a single pose estimator, an ensemble approach was adopted by multiple entries by randomly replicating the methods and fusing the predictions in the post-prediction stage, *e.g.* *A2J*, *AWR*, *NTIS*, *NLE* and *Strawberryfg*.

A2J ensembles predictions from 10 different backbone architectures in Task 1 like *AWR* (5 backbones) and augments test images to ensemble the predictions with different scales and rotations as similar to rotation augmentation adopted by *NLE*. *NTIS* uses predictions obtained from the same model at 6 different training epochs.

A similar ensembling is also adopted by *A2J* in Task 2. *NTIS* adopts a different strategy where N most confident sub-voxel predictions are ensembled to further use them in a refinement stage with Truncated SVDs together with temporal smoothing (Task 2). *NLE* takes advantage of ensembles from multiple pose proposals [25]. *Strawberryfg* employs a different strategy and ensembles the predictions from models that are trained with various input modalities.

Real + synthetic data usage. The methods *Rokid* in Task 1 and *BT* in Tasks 2 and 3 make use of the provided MANO [26] model parameters to synthesize more training samples. *Rokid* leverages the synthesized images and combines them the real images—see Fig. 5—to train their initial pose regression network which effectively boosts accuracies—see Table 5. However, the amount of synthetic data created is limited to 570K for *Rokid* and 32K in Task 2, 100K in Task 3 for *BT*. Considering the continuous high-dimensional hand pose space with or without objects, if we sub-sample uniformly and at minimum, for instance, 10^2 (azimuth/elevation angles) \times 2^5 (articulation) \times 10^1 (shape) \times 10^1 (object) = 320K, the number is already very large, causing a huge compromise issue for memory and training GPU hours. Random sampling was applied without a prior on the data distribution or smart sampling techniques [4, 2]. *BT* generates synthetic images with objects and hands similar to [20] by randomly placing the objects from [8] to nearby hand locations without taking into account the hand




Fig. 5: Visualization of synthetic depth images by *Rokid* [42]: (a) input depth image, (b) rendered depth image using 3D hand model, (c) the mixed by using the pixels with the closest depth values from real and synthetic images.

and object interaction. The rest of the methods use the provided real training data only.

Multi-modal inputs for HPEs. *BT* adopts [37] in Task 3 to align latent spaces from depth and RGB input modalities and to embed the inherit depth information in depth images during learning. makes use of multi-inputs where each is obtained from different representations of the depth image, *e.g.* point-cloud, 3D point projection [6], multi-layer depth map [27], depth voxel [19].

Dominating HPE backbones. ResNet [10] architectures with residual connections have been a popular backbone choice among many HPEs *e.g.* *A2J*, *AWR*, *NTIS*, *Strawberryfg*, *CrazyHand*, *ETH_NVIDIA*, *NLE* or implicitly by *BT* within the ResPEL [17] architecture. *Rokid* adopts EfficientNet-b0 [31] as a backbone which uniformly scales the architecture’s depth, width, and resolution.

6 Results and Discussion

We share our insights and analysis of the results obtained by the participants’ approaches: 6 in Task 1, 4 in Task 2, and 3 in Task 3. Our analyses highlight the impacts of data pre-processing, the use of an ensemble approach, the use of MANO model, different HPE methods, and backbones and post-processing strategies for the pose refinement.

Analysis of Submitted Methods for Task 1. We consider the main properties of the selected methods and the evaluation criteria for comparisons. Table 2 provides the errors for the MJE metric and Fig. 6 show that high success rates are easier to achieve in absence of an object for low distance d thresholds. 2D-based approaches such as *Rokid*, with the advantage of additional data synthesizing, or *A2J*, with cleverly designed local regressors, can be considered to be best when the MJE score is evaluated for the Extrapolation criterion. *AWR* performs comparable to the other 2D-based approaches by obtaining the lowest MJE errors on the Interpolation and Articulation criteria. *AWR* performs best for the distances less than 50mm on Extrapolation as well as showing better generalisation to unseen Viewpoints and Articulations, while excelling to interpolate well. A similar trend is observed with the 3D-voxel-based approach *NTIS*. However, the other 3D supervised methods, *Strawberryfg* and *BT* show lower generalisation capability compared to other approaches while performing reasonably well on the Articulation, Shape, and Interpolation criteria but not being able to show a similar performance for the Extrapolation and Viewpoint criteria.

Table 2: Task 1 - MJE (mm) and ranking of the methods on five evaluation criteria. Best results on each evaluation criteria are highlighted.

Username	Extrapolation	Interpolation	Shape	Articulation	Viewpoint
<i>Rokid</i>	13.66 (1)	4.10 (2)	10.27 (1)	4.74 (3)	7.44 (1)
<i>A2J</i>	13.74 (2)	6.33 (6)	11.23 (4)	6.05 (6)	8.78 (6)
<i>AWR</i>	13.76 (4)	3.93 (1)	11.75 (5)	3.65 (1)	7.50 (2)
<i>NTIS</i>	15.57 (7)	4.54 (3)	12.05 (6)	4.21 (2)	8.47 (4)
<i>Strawberryfg</i>	19.63 (12)	8.42 (10)	14.21 (10)	7.50 (9)	14.16 (12)
<i>BT</i>	23.62 (14)	18.78 (16)	21.84 (16)	16.73 (16)	19.48 (14)

Table 3: Task 2 - MJE (mm) and ranking of the methods on four evaluation criteria.

Username	Extrapolation	Interpolation	Object	Shape
<i>NTIS</i>	33.48 (1)	17.42 (1)	29.07 (2)	23.62 (2)
<i>A2J</i>	33.66 (2)	17.45 (2)	27.76 (1)	23.39 (1)
<i>CrazyHand</i>	38.33 (4)	19.71 (4)	32.60 (4)	26.26 (4)
<i>BT</i>	47.18 (5)	24.95 (6)	38.76 (5)	32.36 (5)

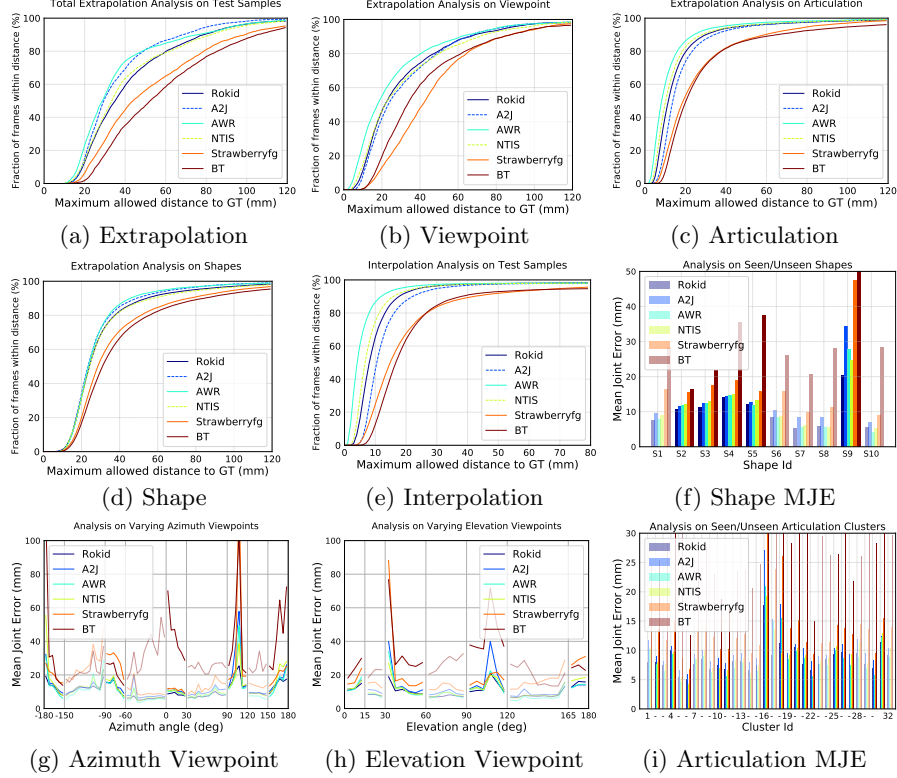


Fig. 6: Task 1 - Success rate analysis (a-e) and MJE analysis on extrapolation and interpolation using shapes (f), viewpoints (g, h) and articulations (i). Solid colors depict samples of extrapolation and transparent colors depict interpolation samples in plots (f-i).

Analysis of Submitted Methods for Task 2. We selected four submitted methods to compare on Task 2, where a hand interacts with an object in an egocentric viewpoint. Success rates illustrated in Fig. 7 highlight the difficulty of extrapolation. All methods struggle to show good performance on estimating frames with joint errors less than $15mm$. On the other hand, all methods can estimate 20% to 30% of the joints correctly with less than $15mm$ error for the other criteria in this task.

NTIS (a voxel-based) and *A2J* (weighted local regressors with anchor points) perform similarly when MJEs for all joints are considered. However, *NTIS* obtains higher success rates on the frame-based evaluation for all evaluation criteria with low distance error thresholds (*d*), see Fig. 7. Its performance is relatively much higher when Extrapolation is considered, especially for the frames with unseen objects, see Fig. 7. This can be explained by having a better embedding of the

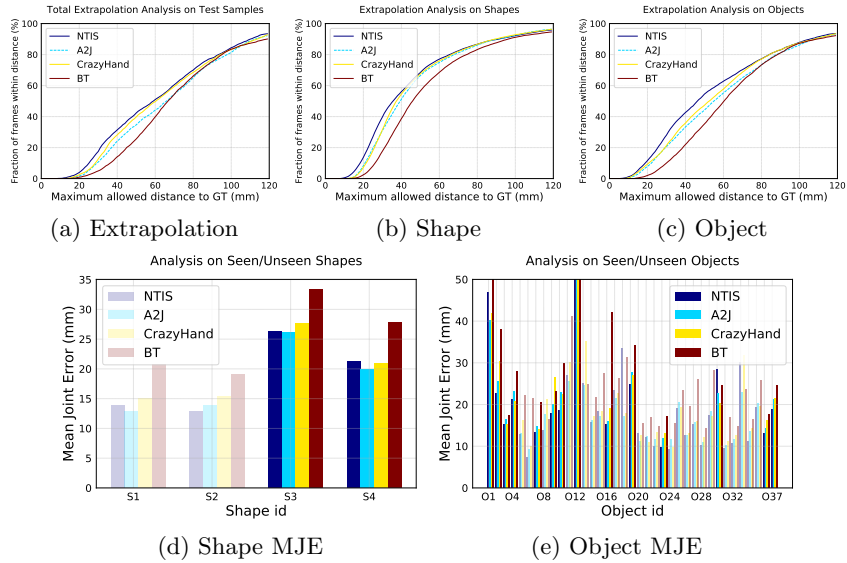


Fig. 7: Task 2 - Success rate analysis (a,b,c) and interpolation (seen, transparent) and extrapolation (unseen, solid) errors for subject (d) and object (e).

occluded hand structure with the voxels in the existence of seen/unseen objects. *NTIS* interpolates well under low distance thresholds.

Note that the first three methods, *NTIS*, *A2J*, and *CrazyHand* perform very similar for high error thresholds *e.g.* $d > 30mm$. *CrazyHand* uses a structured detection-regression-based HPE where a heatmap regression is employed for the joints from palm to tips in a sequential manner which is highly valuable for egocentric viewpoints, helps to obtain comparable results with *A2J* where the structure is implicitly refined by the local anchor regressors.

Analysis of Submitted Methods for Task 3. We selected 3 entries, with different key properties for this analysis. It is definitively harder for the participants to provide accurate poses compared to the previous tasks. None of the methods can estimate frames that have all joints estimated with less than $25mm$ error, see Fig. 8. The $25mm$ distance threshold shows the difficulty of estimating a hand pose accurately from RGB input modality even though the participants of this task were provided with the ground-truth wrist joint location.

The task is based on hand-object interaction in RGB modality. Therefore, the problem raises the importance of multi-modal data and learn from different modalities. Only *BT* uses the MANO parameters provided by the organizers to synthesize 100K images and adds random objects near the hand. This approach supports the claim on the importance of multi-modality and filling the real data gaps with synthetic data with its close performance to the two higher ranked methods in MJE.

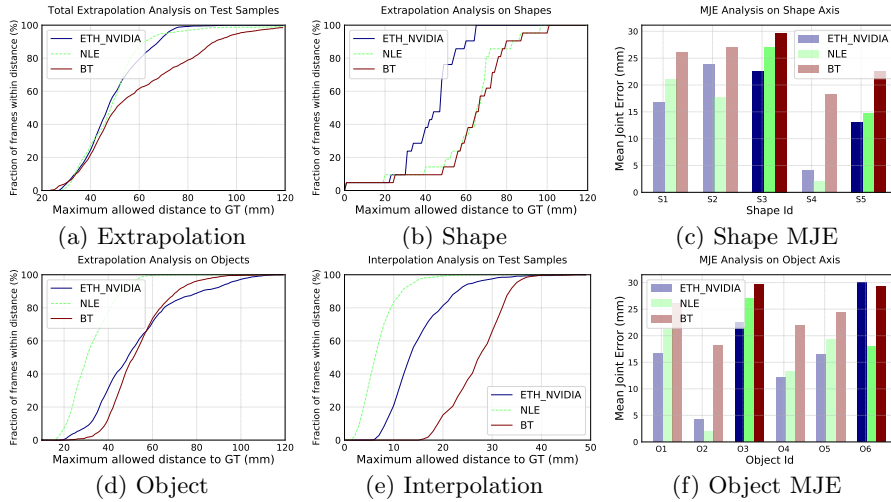


Fig. 8: Task 3 - Success rate analysis on the evaluation criteria (a,b,d,e) and MJE error analysis on the seen/unseen subjects (c) and objects (f). For (c) and (f), solid and transparent colors are used to depict extrapolation and interpolation.

The generalisation performance of *BT* in Task 3 compared to the team’s approaches with similar gist in Tasks 1 and 2 supports the importance of multi-model learning and synthetic data augmentation. The close performance of the method to generalise to unseen objects compared to *ETH_NVIDIA* and to generalise to unseen shapes compared to *NLE* also supports the argument with the data augmentation. The approach is still outperformed in MJE for this task although it performs close to the other methods.

NLE’s approach shows the impact of learning to estimate 2D joints+3D joints ($28.45mm$) compared to learning 3D joints alone ($37.31mm$) on the Object as well as improvements for the Interpolation. Object performance is further improved to $23.22mm$ with PPI integration. Further insights put by *NLE*’s own experiments on the number rotation augmentations (n) in post-processing helps to better extrapolate for unseen shapes ($17.35mm$, $16.77mm$, $15.79mm$ where $n = 1, 4, 12$, respectively).

Analysis on the Usage of Synthetic Images. The best performing method of Task 1 (*Rokid*) in MJE uses the 3D hand model parameters to create 570K

Table 4: Task 3 - MJE (mm) and ranking of the methods on four evaluation criteria.

Username	Extrapolation	Interpolation	Object	Shape
<i>ETH_NVIDIA</i>	24.74 (1)	6.70 (3)	27.36 (2)	13.21 (1)
<i>NLE</i>	29.19 (2)	4.06 (1)	18.39 (1)	15.79 (3)
<i>BT</i>	31.51 (3)	19.15 (5)	30.59 (3)	23.47 (4)

Table 5: Impact of synthetic data reported by *Rokid* [42] with learning from different ratios of synthetic data and the Task 1 training set. 100% = 570K.

Synthetic Data %	-	10%	30%	70%	100%
Extrapolation MJE (mm)	30.11	16.70	16.11	15.81	15.73

synthetic images by either perturbing (first stage) the model parameters or not (second stage). Synthetic data usage significantly helps in training the initial model (see Fig. 5). Table 5 shows the impact of different proportions of the 570K synthetic data usage to train the model together with the real training images. Using synthetic data can boost such a simple 3D joint regressor’s performance from MJE of $30.11mm$ to $15.73mm$, a $\sim 50\%$ improvement. Moreover, *Rokid*’s experiments with a regression model trained for 10 epochs shows the impact of the mixed depth inputs, Fig. 5, to lower the total extrapolation error ($26.16mm$) compared to the use of raw depth renderings ($30.13mm$) or the renderings averaged ($31.92mm$) with the real input images. *BT* uses synthetic images in a very small amount of $32K$ and $100K$ in Tasks 2 and 3 since 3D reconstruction is difficult to train at a larger scale. However, favorable impact the data can be observed by comparing performances in Tasks 1 and 2.

Please refer to the supplementary document provided with this work for qualitative results and also for more results and discussions on the performance of the approaches on each axis, backbone architectures, ensembling and post-processing techniques and joint success rates obtained by the participants.

7 Conclusion

We carefully designed structured training and test sets for 3D HPEs and organized a challenge for the hand pose community to show state-of-the-art methods still tend to fail to extrapolate on large pose spaces. Our analyses highlight the impacts of using ensembles, the use of synthetic images, different type of HPEs *e.g.* 2D, 3D or local-estimators and post-processing. Ensemble techniques, both methodologically in 2D and 3D HPEs and in post-processing, help many approaches to boost their performance on extrapolation. The submitted HPEs were proven to be successful while interpolating in all the tasks, but their extrapolation capabilities vary significantly. Scenarios such as hands interacting with objects present the biggest challenges to extrapolate by most of the evaluated methods both in depth and RGB modalities.

Given the limited extrapolation capabilities of the methods, usage of synthetic data is appealing. Only a few methods actually were making use of synthetic data to improve extrapolation. 570K synthetic images used by the winner of Task 1 is still a very small number compared to how large, potentially infinite, it could be. We believe that investigating these possibilities, jointly with data sub-sampling strategies and real-synthetic domain adaptation is a promising and interesting line of work. The question of what would be the outcome if we sample ‘dense enough’ in the continuous and infinite pose space and how ‘dense enough’ is defined when we are limited by hardware and time is significant to answer.

8 Acknowledgements

This work is partially supported by Huawei Technologies Co. Ltd. and Samsung Electronics. S. Baek was supported by IITP funds from MSIT of Korea (No. 2020-0-01336) for AIGS of UNIST.

References

1. Baek, S., Kim, K.I., Kim, T.K.: Pushing the envelope for RGB-based dense 3D hand pose estimation via neural rendering. In: CVPR (2019)
2. Bhattarai, B., Baek, S., Bodur, R., Kim, T.: Sampling strategies for GAN synthetic data. In: ICASSP (2020)
3. Boukhayma, A., de Bem, R., Torr, P.H.: 3D hand shape and pose from images in the wild. In: CVPR (2019)
4. Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: AutoAugment: Learning augmentation strategies from data. In: CVPR (2019)
5. Garcia-Hernando, G., Yuan, S., Baek, S., Kim, T.K.: First-person hand action benchmark with RGB-D videos and 3D hand pose annotations. In: CVPR (2018)
6. Ge, L., Liang, H., Yuan, J., Thalmann, D.: Robust 3D hand pose estimation in single depth images: from single-view CNN to multi-view CNNs. In: CVPR (2016)
7. Ge, L., Ren, Z., Li, Y., Xue, Z., Wang, Y., Cai, J., Yuan, J.: 3D hand shape and pose estimation from a single RGB image. In: CVPR (2019)
8. Hampali, S., Oberweger, M., Rad, M., Lepetit, V.: HO-3D: A multi-user, multi-object dataset for joint 3D hand-object pose estimation. In: arXiv preprint arXiv:1907.01481v1 (2019)
9. Hasson, Y., Varol, G., Tzionas, D., Kalevatykh, I., Black, M.J., Laptev, I., Schmid, C.: 3D hand shape and pose estimation from a single RGB image. In: CVPR (2019)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
11. Huang, W., Ren, P., Wang, J., Qi, Q., Sun, H.: AWR: Adaptive weighting regression for 3D hand pose. In: AAAI (2020)
12. Iqbal, U., Molchanov, P., Breuel, T., Gall, J., Kautz, J.: Hand pose estimation via latent 2.5D heatmap regression. In: ECCV (2018)
13. Joo, H., Liu, H., Tan, L., Gui, L., Nabbe, B., Matthews, I., Kanade, T., Nobuhara, S., Sheikh, Y.: Panoptic studio: a massively multiview system for social motion capture. In: ICCV (2015)
14. Kanazawa, A., Tulsiani, S., Efros, A.A., Malik, J.: Learning category-specific mesh reconstruction from image collections. In: ECCV (2018)
15. Kennedy, J., Eberhart, R.: Particle Swarm Optimization. In: ICNN (1995)
16. Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In: ICCV (2019)
17. Li, S., Lee, D.: Point-to-pose voting based hand pose estimation using residual permutation equivariant layer. In: CVPR (2019)
18. Lin, J., Wu, Y., Huang, T.S.: Modeling the constraints of human hand motion. In: HUMO (2000)
19. Moon, G., Chang, J.Y., Lee, K.M.: V2V-PoseNet: Voxel-to-voxel prediction network for accurate 3D hand and human pose estimation from a single depth map. In: CVPR (2018)
20. Mueller, F., Bernard, F., Sotnychenko, O., Mehta, D., Sridhar, S., Casas, D., Theobalt, C.: GANerated hands for real-time 3D hand tracking from monocular RGB. In: CVPR (2018)
21. Mueller, F., Mehta, D., Sotnychenko, O., Sridhar, S., Casas, D., Theobalt, C.: Real-time hand tracking under occlusion from an egocentric RGB-D sensor. In: ICCV (2017)
22. Oberweger, M., Lepetit, V.: DeepPrior++: Improving fast and accurate 3D hand pose estimation. In: ICCV Workshop on HANDS (2017)

23. Oikonomidis, I., Kyriazis, N., Argyros, A.A.: Efficient model-based 3D tracking of hand articulations using kinect. In: BMVC (2011)
24. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: NIPS (2015)
25. Rogez, G., Weinzaepfel, P., Schmid, C.: LCR-Net++: Multi-person 2D and 3D pose detection in natural images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **42**(5), 1146–1161 (2019)
26. Romero, J., Tzionas, D., Black, M.J.: Embodied Hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* **36**(6), 245:1–245:17 (2017)
27. Shin, D., Ren, Z., Sudderth, E.B., Fowlkes, C.C.: Multi-layer depth and epipolar feature transformers for 3D scene reconstruction. In: CVPR Workshops (2019)
28. Simon, T., Joo, H., Matthews, I., Sheikh, Y.: Hand keypoint detection in single images using multiview bootstrapping. In: CVPR (2017)
29. Sinha, A., Choi, C., Ramani, K.: DeepHand: Robust hand pose estimation by completing a matrix imputed with deep features. In: CVPR (2016)
30. Sun, X., Xiao, B., Wei, F., Liang, S., Wei, Y.: Integral human pose regression. In: ECCV (2018)
31. Tan, M., Le, Q.V.: EfficientNet: Rethinking model scaling for convolutional neural networks. In: ICML (2019)
32. Tung, H.Y.F., Tung, H.W., Yumer, E., Fragkiadaki, K.: Self-supervised learning of motion capture. In: NIPS (2017)
33. Wan, Q.: SenoritaHand: Analytical 3D skeleton renderer and patch-based refinement for HANDS19 challenge Task 1 - Depth-based 3D hand pose estimation (Dec 2019), <https://github.com/strawberryfg/Senorita-HANDS19-Pose>
34. Wan, Q., Qiu, W., Yuille, A.L.: Patch-based 3D human pose refinement. In: CVPRW (2019)
35. Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: A convolutional neural network for 6D object pose estimation in cluttered scenes. In: RSS (2018)
36. Xiong, F., Zhang, B., Xiao, Y., Cao, Z., Yu, T., Zhou, J.T., Yuan, J.: A2J: Anchor-to-joint regression network for 3D articulated pose estimation from a single depth image. In: ICCV (2019)
37. Yang, L., Li, S., Lee, D., Yao, A.: Aligning latent spaces for 3D hand pose estimation. In: ICCV (2019)
38. Yuan, S., Ye, Q., Garcia-Hernando, G., Kim, T.K.: The 2017 hands in the million challenge on 3d hand pose estimation. *arXiv preprint arXiv:1707.02237* (2017)
39. Yuan, S., Garcia-Hernando, G., Stenger, B., Moon, G., Chang, J.Y., Lee, K.M., Molchanov, P., Kautz, J., Honari, S., Ge, L., Yuan, J., Chen, X., Wang, G., Yang, F., Akiyama, K., Wu, Y., Wan, Q., Madadi, M., Escalera, S., Li, S., Lee, D., Oikonomidis, I., Argyros, A., Kim, T.K.: Depth-based 3D hand pose estimation: From current achievements to future goals. In: CVPR (2018)
40. Yuan, S., Ye, Q., Stenger, B., Jain, S., Kim, T.K.: BigHand 2.2M Benchmark: hand pose data set and state of the art analysis. In: CVPR (2017)
41. Zhang, X., Li, Q., Mo, H., Zhang, W., Zheng, W.: End-to-end hand mesh recovery from a monocular RGB image. In: ICCV (2019)
42. Zhang, Z., Xie, S., Chen, M., Zhu, H.: HandAugment: A simple data augmentation method for depth-based 3D hand pose estimation. *arXiv:2001.00702* (2020)
43. Zhou, X., Wan, Q., Zhang, W., Xue, X., Wei, Y.: Model-based deep hand pose estimation. In: IJCAI (2016)
44. Zimmermann, C., Brox, T.: Learning to estimate 3D hand pose from single RGB images. In: ICCV (2017)