

HOnnotate: A method for 3D Annotation of Hand and Object Poses

Shreyas Hampali¹, Mahdi Rad¹, Markus Oberweger¹, and Vincent Lepetit^{2,1}

¹Institute for Computer Graphics and Vision, Graz University of Technology, Austria

²LIGM, Ecole des Ponts, Univ Gustave Eiffel, CNRS, Marne-la-Vallée, France

{hampali, rad, oberweger, lepetit}@icg.tugraz.at

Project Page: <https://www.tugraz.at/index.php?id=40231>

Abstract

We propose a method for annotating images of a hand manipulating an object with the 3D poses of both the hand and the object, together with a dataset created using this method. Our motivation is the current lack of annotated real images for this problem, as estimating the 3D poses is challenging, mostly because of the mutual occlusions between the hand and the object. To tackle this challenge, we capture sequences with one or several RGB-D cameras and jointly optimize the 3D hand and object poses over all the frames simultaneously. This method allows us to automatically annotate each frame with accurate estimates of the poses, despite large mutual occlusions. With this method, we created HO-3D, the first markerless dataset of color images with 3D annotations for both the hand and object. This dataset is currently made of 77,558 frames, 68 sequences, 10 persons, and 10 objects. Using our dataset, we develop a single RGB image-based method to predict the hand pose when interacting with objects under severe occlusions and show it generalizes to objects not seen in the dataset.

1. Introduction

Methods for 3D pose estimation of rigid objects and hands from monocular images have made significant progress recently, thanks to the development of Deep Learning, and the creation of large datasets or the use of synthetic images for training [32, 35, 46, 55, 71, 73]. However, these recent methods still fail when a hand interacts with an object, mostly because of large mutual occlusions, and of the absence of datasets specific to 3D pose estimation for hand+object interaction. Breaking this limit is highly desirable though, as 3D hand and object poses would be very useful in augmented reality applications, or for learning-by-imitation in robotics, for example.

Several pioneer works have already considered this problem, sometimes with impressive success [26, 52, 61]. These works typically rely on tracking algorithms to exploit temporal constraints, often also considering physical constraints between the hand and the object to improve the pose estimates. While these temporal and physical constraints remain relevant, we would like to also benefit from the power of data-driven methods for 3D hand+object pose estimation from a single image. Being able to estimate these poses from a single frame would avoid manual initialization and drift of tracking algorithms. A data-driven approach, however, requires real or synthetic images annotated with the 3D poses of the object and the hand. Unfortunately, creating annotated data for the hand+object problem is very challenging. Both common options for creating 3D annotations, annotating real images and generating synthetic images, raise challenging problems.

Annotating real images. One can rely on some algorithm for automated annotation, as was done for current benchmarks in 3D hand pose estimation [45, 54, 57, 71, 74], where the “ground truth” annotations are obtained automatically with a tracking algorithm. Though these annotations are noisy, they are usually taken for granted and used for training and evaluation [37]. Another approach is to use sensors attached to the hand as in [14] (bottom right image of Fig. 1). This directly provides the 3D poses, however, the sensors are visible in the images, and thus bias learning methods. Significant effort is still required in developing algorithms for automated annotation of real images.

Generating synthetic images. Relying on synthetic images is appealing, as the 3D poses are known perfectly. Realistic rendering and domain transfer can be used to train 3D pose estimation on synthetic images [31, 47, 73]. Generating physically correct grasps is possible [29], as shown in [18], but complex manipulation is difficult to simulate. However, real images with accurate 3D annotations would still be needed to evaluate the generalizability of the method

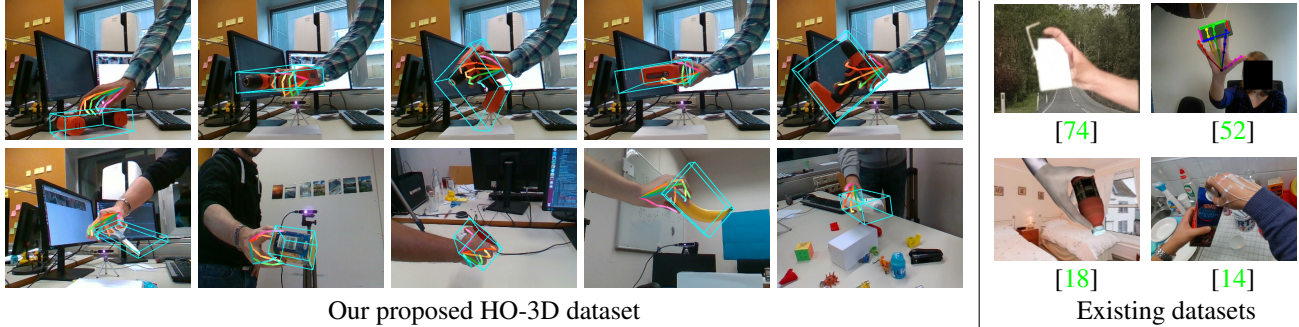


Figure 1: We introduce a method for labelling real images of hand+object interaction with the 3D poses of the hand and of the object. With this method, we automatically created a dataset made of more than 75,000 frames, 10 different objects and 10 different users. In comparison, existing datasets have several limitations: The 3D objects are very simple, the interaction is not realistic, the images are synthetic, corrupted by sensors, and/or the number of samples is limited. More illustrations of annotations in our dataset are shown in the supplementary material.

to real data.

We therefore propose a method to automatically annotate real images of hands grasping objects with their 3D poses. Our method works with a single RGB-D camera, but can exploit more cameras if available for better robustness and accuracy. The single-camera setup works under the assumption that the grasp pose varies marginally over the sequence; the multi-camera setup can handle complex hand+object interaction scenarios. Instead of tracking the poses frame-by-frame, our method optimizes jointly all the 3D poses of the hand and the object over the sequence. As our evaluations show, this allows us to exploit temporal consistency in a stronger way than a tracking algorithm. Using differentiable rendering, we can optimize a complex objective function by exploiting the new powerful gradient descent methods originally developed for Deep Learning [24]. We see this approach as the equivalent of bundle adjustment for SLAM algorithms, where we track objects instead of points.

We rely on the MANO hand model [50], and the 3D model of the objects. We use objects from the YCB-Video dataset [68], as they have various shapes and materials, and can be bought online [1] by researchers interested in performing their own experiments. Being able to use a single camera also enables easier expansion of the dataset by other researchers with a larger variety of objects and grasping poses as multi-camera capture is often complex to setup.

Using our method, we created a dataset, depicted in Fig. 1, which we call HO-3D. In addition, we used this dataset to learn to predict from a single RGB image the 3D pose of a hand manipulating an object. More exactly, we train a Deep Network to predict the 2D joint locations of the hand along with the joint direction vectors and lift them to 3D by fitting a MANO model to these predictions. This validates the fact that the 3D poses estimated by our annotation method can actually be used in a data-driven method for hand pose estimation. By comparing with an existing

method for hand+object pose estimation [18] that directly estimates MANO parameters, we show that predicting 2D keypoints and lifting them to 3D performs more accurately.

2. Related Work

The literature on hand and object pose estimation is extremely broad, and we review only some works here.

2.1. 3D Object Pose Estimation

Estimating the 3D pose of an object from a single frame is still one of the fundamental problems of Computer Vision. Some methods are now robust to partial occlusions [20, 36, 41], but many works rely on RGB-D data to handle this problem [5, 8, 22, 30], by fitting the 3D object model to depth data. This can fail when a hand grasps the object, since the surface of the hand can be mistaken for the surface of the object.

2.2. 3D Hand Pose Estimation

Single image hand pose estimation is also a very popular problem in Computer Vision, and approaches can be divided into discriminative and generative methods. Discriminative approaches directly predict the joint locations from RGB or RGB-D images. Recent works based on Deep Networks [15, 32, 34, 35, 59, 69, 73] show remarkable performance, compared to early works based on Random Forests such as [23]. However, discriminative methods perform poorly in case of partial occlusion.

Generative approaches take advantage of a hand model and its kinematic structure to generate hand pose hypotheses that are physically plausible [12, 28, 45, 51, 53, 63, 70]. [31, 40] predict 2D joint locations and then lift them to 3D. Generative approaches are usually accurate and can be made robust to partial occlusions. They typically rely on some pose prior, which may require manual initialization or result in drift when tracking.

Our work is related to both discriminative and generative approaches: we use a generative approach within a global optimization framework to generate the pose annotations, and use a discriminative method to initialize this complex optimization. [64] also combines generative and discriminative approaches to train a network in a self-supervised setting. However, they only consider hands. We also train a discriminative method using our dataset, to predict the hand poses which are robust to occlusions from interacting objects.

2.3. Synthetic Images for 3D Pose Estimation

Being able to train discriminative methods on synthetic data is valuable as it is difficult to acquire annotations for real images [73]. [18, 47] show that because of the domain gap between synthetic and real images, training on synthetic images only results in sub-optimal performance. A GAN method was used in [31] to make synthetic images of hands more realistic. While using synthetic images remains appealing for many problems, creating the virtual scenes can be expensive and time-consuming. Generating animated realistic hand grasps of various objects, as it would be required to solve the problem considered in this paper remains challenging. Being able to use real sequences for training has thus also its advantages. Moreover, evaluation should be performed on real images.

2.4. Joint Hand+Object Pose Estimation

Early approaches for joint hand+object pose estimation [2, 38, 65] typically relied on multi-view camera setups, and frame-by-frame tracking methods, which may require careful initialization and drift over time. [39, 62] propose generative methods to track finger contact points for in-hand RGB-D object shape scanning. [43, 44] consider sensing from vision to estimate contact forces during hand+object interactions using a single RGB-D camera, and then estimate the hand and the object pose. However, these methods are limited to small occlusions.

[26, 61] propose to use a physics simulator and a 3D renderer for frame-to-frame tracking of hands and objects from RGB-D. [27] uses an ensemble of Collaborative Trackers for multi-object and multiple hand tracking from RGB-D images. The accuracy of these methods seems to be qualitatively high, but as the ground truth acquisition in a real-world is known to be hard, they evaluate the proposed method on synthetic datasets, or by measuring the standard deviation of the difference in hand/object poses during a grasping scenario.

[60] considers the problem of tracking a deformable object in interaction with a hand, by optimizing an energy function on the appearance and the kinematics of the hand, together with hand+object contact configurations. However, it is evaluated quantitatively only on synthetic images,

which points to the difficulty of evaluation on real data. In addition, they only consider scenarios where the hand is visible from a top view, restricting the range of the hand poses and not allowing occlusions.

Very recently, [25] uses a coarse hand pose estimation to retrieve the 3D pose and shape of hand-held objects. However, they only consider a specific type of object and do not estimate the object pose. [18] presents a model with contact loss that considers physically feasible hand+object interaction to improve grasp quality. However, to estimate 3D hand pose, they predict PCA components for the pose, which results in lower accuracy compared to ours, as our experiments show. [58] proposes a deep model to jointly predict 3D hand and object poses from egocentric views, but the absence of physical constraints might result in infeasible grasps.

2.5. Hand+Object Datasets

Several datasets for hand+object interactions have already been proposed. Many works provide egocentric RGB or RGB-D sequences for action recognition [3, 6, 7, 13, 48]. However, they focus on grasp and action labels and do not provide 3D poses. [10, 32, 49, 60] generate synthetic datasets with 3D hand pose annotations, but fine interaction between a hand and an object remains difficult to generate accurately.

[61, 63] captured sequences in the context of hand+hand and hand+object interaction, with 2D hand annotations only. [33] collected a dataset of real RGB images of hands holding objects. They also provide 2D joint annotations of pairs of non-occluded and occluded hands, by removing the object from the grasp of the subject, while maintaining their hand in the same pose. [16] proposes two datasets, a hand+object segmentation dataset, and a hand+object pose estimation dataset. However, for both datasets, the background pixels have been set to zero, and the training images only consist of a hand interacting with a tennis ball. They provide hand pose annotations and object positions, by manually labelling the joints and using a generative method to refine the joint positions. [21] generate a large scale dataset with full body pose and hand pose annotations in a multi-view setup. They use a generative approach to fit the body and hand models to 3D keypoints and point cloud. However, their dataset focuses on total body pose annotation and not hand+object interactions exclusively and does not provide object pose annotations.

[52] proposed an RGB-D dataset of a hand manipulating a cube, which contains manual ground truth for both fingertip positions and 3D poses of the cube. [42] collected a dataset where they measure motion and force under different object-grasp configurations using sensors, but do not provide 3D poses. In contrast to these previous works, [14] provides a dataset of hand and object interactions with 3D annotations for both hand joints and object pose. They used

Dataset	No. of Frames	3D Object Pose	Marker-less	Real Images	Labels	No. of Objects	No. of Subjects
PAN [21]	675K	-	+	+	automatic	-	70
GAN [31]	300K	-	+	-	synthetic	-	-
FPFA [14]	100K	+(23K frames)	-	+	automatic	26 (4 models)	6
ObMan [18]	150K	+	+	-	synthetic	2.7K	20
FreiHAND [74]	37K	-	+	+	hybrid	27	35
HO-3D (ours)	78K	+	+	+	automatic	10	10

Table 1: Comparison of hand+object datasets.

a motion capture system made of magnetic sensors attached to the user’s hand and to the object in order to obtain hand 3D pose annotations in RGB-D video sequences. However, this changes the appearance of the hand in the color images as the sensors and the tape attaching them are visible.

Very recently, [18] introduced ObMan, a large dataset of images of hands grasping objects. The images in ObMan dataset are synthetic and the grasps are generated using an algorithm from robotics. Even more recently, [74] proposed a multi-view RGB dataset, FreiHAND, which includes hand-object interactions. However, the annotations are limited to the 3D poses and shapes of the hand. Further, [74] uses a human-in-the-loop method to obtain annotations from multiple RGB cameras in a green-screen background environment. Our method, on the other hand is fully automatic, capable of working even on a single RGBD-camera setup and does not make any assumption on the background. The objects in our dataset are also larger than those in FreiHAND, thus resulting in a more challenging scenario as the occlusions are larger. The annotation accuracy of our method is comparable to [74] as described in Section 6.1.

As illustrated in Fig. 1 and Table 1, our HO-3D dataset is the first markerless dataset providing both 3D hand joints and 3D object pose annotations for real images, while the hand and the object are heavily occluded by each other.

3. 3D Annotation Method

We describe below our method for annotating a sequence $\mathcal{T} = \{ \{ (I_c^t, D_c^t) \}_{c=1}^{N_C} \}_{t=1}^{N_F}$ of $N_C \times N_F$ of RGB-D frames, captured by N_C cameras. The sequence captures a hand interacting with an object. Each RGB-D frame is made of a color image I_c^t and a depth map D_c^t .

We define the 3D hand and object poses in Section 3.1, and our general cost function in Section 3.2. We initialize the poses automatically and optimize the cost function in multiple stages as described in Sections 4.1 and 4.2.

3.1. 3D Hand and Object Poses

We aim to estimate the 3D poses $\mathcal{P} = \{ (p_h^t, p_o^t) \}_{t=1}^{N_F}$ for both the hand and the object in all the images of the sequence. We adopt the MANO hand model [50] and use the objects from the YCB-Video dataset [68] as their corresponding 3D models are available and of good quality. The MANO hand pose $p_h^t \in \mathbb{R}^{51}$ consists of 45 DoF (3 DoF for each of the 15 finger joints) plus 6 DoF for rotation and

translation of the wrist joint. The 15 joints together with the wrist joint form a kinematic tree with the wrist joint node as the first parent node. In addition to the pose parameters p_h^t , the hand model has shape parameters $\beta \in \mathbb{R}^{10}$ that are fixed for a given person and we follow a method similar to [56] to estimate these parameters. More details about the shape parameter estimation are provided in the supplementary material. The object pose $p_o^t \in \text{SE}(3)$ consists of 6 DoF for global rotation and translation.

3.2. Cost Function

We formulate the hand+object pose estimation as an energy minimization problem:

$$\hat{\mathcal{P}} = \arg \min_{\mathcal{P}} \sum_{t=1}^{N_F} (E_{\mathcal{D}}(p_h^t, p_o^t) + E_{\mathcal{C}}(p_h^t, p_o^t)), \quad (1)$$

where $E_{\mathcal{D}}$ and $E_{\mathcal{C}}$ represent the energy from data terms and constraints, respectively. We define $E_{\mathcal{D}}$ as,

$$E_{\mathcal{D}}(p_h^t, p_o^t) = \sum_{c=1}^{N_C} \left(\alpha E_{\text{mask}}(I_c^t, p_h^t, p_o^t) + \beta E_{\text{dpt}}(D_c^t, p_h^t, p_o^t) + \gamma E_{\text{j2D}}(I_c^t, p_h^t) \right) + \delta E_{\text{3D}}(\{D_c^t\}_{c=1..N_C}, p_h^t, p_o^t), \quad (2)$$

where $E_{\text{mask}}(\cdot)$ is a silhouette discrepancy term, $E_{\text{dpt}}(\cdot)$ a depth residual term, $E_{\text{j2D}}(\cdot)$ a 2D error in hand joint locations, and $E_{\text{3D}}(\cdot)$ a 3D distance term. This last term is not absolutely necessary, however, we observed that it significantly speeds up convergence. $\alpha, \beta, \gamma, \delta$ are weights.

The constraints energy $E_{\mathcal{C}}$ is defined as,

$$E_{\mathcal{C}}(p_h^t, p_o^t) = \epsilon E_{\text{joint}}(p_h^t) + \zeta E_{\text{phy}}(p_h^t, p_o^t) + \eta E_{\text{tc}}(p_h^t, p_o^t, p_h^{t-1}, p_o^{t-1}, p_h^{t-2}, p_o^{t-2}), \quad (3)$$

where $E_{\text{joint}}(\cdot)$ denotes a prior on the hand pose to prevent unnatural poses, $E_{\text{phy}}(\cdot)$ is a physical plausibility term ensuring the hand and the object do not interpenetrate, and $E_{\text{tc}}(\cdot)$ is a temporal consistency term. The terms are weighted by parameters ϵ, ζ and η .

We detail each of the terms in $E_{\mathcal{D}}$ and $E_{\mathcal{C}}$ below. For simplicity, we omit the frame index t from our above notation except when necessary.

Silhouette discrepancy term E_{mask} . The $E_{\text{mask}}(\cdot)$ term compares the silhouettes of the hand and the object models rendered with the current estimated poses and their segmentation masks. We obtain a segmentation $S_c(I)$ of the hand and the object in the color image I of camera c using DeepLabv3 [9] trained on images created by synthetically over-laying and under-laying images of hands on YCB objects. More details about this step are given in the supplementary material. The hand and object models are rendered on the camera plane using a differentiable renderer [19],

which enables computing the derivatives of E_{mask} with respect to the pose parameters. The silhouette of the hand and object rendered on camera c is denoted by $\text{RS}_c(\mathbf{p}_h, \mathbf{p}_o)$ and the silhouette discrepancy is defined as,

$$E_{\text{mask}}(\mathbf{I}_c, \mathbf{p}_h, \mathbf{p}_o) = \|\text{RS}_c(\mathbf{p}_h, \mathbf{p}_o) - S(\mathbf{I}_c)\|^2. \quad (4)$$

Depth residual term E_{dpt} . The depth residual term is similar to the segmentation discrepancy term:

$$E_{\text{dpt}}(\mathbf{D}_c, \mathbf{p}_h, \mathbf{p}_o) = \text{Tukey}(\|\text{RD}_c(\mathbf{p}_h, \mathbf{p}_o) - \mathbf{D}_c\|), \quad (5)$$

where $\text{RD}_c(\mathbf{p}_h, \mathbf{p}_o)$ is the depth rendering of the hand and the object under their current estimated poses \mathbf{p}_h and \mathbf{p}_o . The Tukey function is a robust estimator that is similar to the ℓ_2 loss close to 0, and constant after a threshold. It is useful to be robust to small deviations in the scale and shape of the hand and object models and also noise in the captured depth maps. E_{dpt} is differentiable as we employ a differentiable renderer [19] for rendering the depth maps.

2D Joint error term E_{j2D} . The 2D joint error term is defined as,

$$E_{\text{j2D}}(\mathbf{I}_c, \mathbf{p}_h) = \sum_{i=1}^{21} h[i] \left\| \text{proj}_c(\mathbf{J}_{\mathbf{p}_h}[i]) - \mathbf{K}_c[i] \right\|^2, \quad (6)$$

where $\mathbf{J}_{\mathbf{p}_h}[i]$ denotes the i^{th} 3D hand joint location under pose \mathbf{p}_h , the $\text{proj}_c(\cdot)$ operator projects it onto camera c , $\mathbf{K}_c[i]$ is its predicted 2D location, and $h[i]$ its confidence. The 21 hand joints in $E_{\text{j2D}}(\cdot)$ consist of 15 finger joints, 5 finger tips, and the wrist joint.

In practice, we take the $\mathbf{K}_c[i]$ as the locations of the maximum values of heatmaps, and the $h[i]$ as the maximum values themselves. To predict these heatmaps, we trained a CNN based on the architecture of [66]. Training data come from an initial dataset [17] we created using a semi-automatic method. This dataset is made of 15,000 frames from 15 sequences in a single camera setup. We manually initialized the grasp pose and object pose for the first frame of each sequence. The manipulators were asked to maintain their grasp poses as rigid as possible to make the registration easier. We then ran the optimization stages for the single camera case described in Section 4.2. After optimization, we augmented the resulting dataset by scaling and rotating the images, and adding images from the Panoptic Studio dataset [67], which contain 3D annotations for hands.

3D error term E_{3D} . This term is not absolutely necessary as the depth information from all the cameras is already exploited by E_{dpt} , however it accelerates the convergence by guiding the optimization towards the minimum even from far away. We build a point cloud \mathbf{P} by merging the depth maps from the RGB-D cameras after transforming them to

a common reference frame. More details on the point cloud reconstruction can be found in the supplementary material.

We segment \mathbf{P} into an object point cloud \mathbf{P}_o and a hand point cloud \mathbf{P}_h using the segmentation mask $S_c(\mathbf{I})$ in each camera image. At each iteration of the optimization, for each point $\mathbf{P}_o[j]$ of the object point cloud, we look for the closest vertex $\mathbf{V}_o[j^*]$ on the object mesh, and for each point $\mathbf{P}_h[k]$ of the hand point cloud, we look for the closest vertex $\mathbf{V}_h[k^*]$ on the hand mesh. $E_{\text{3D}}(\mathbf{P}, \mathbf{p}_h, \mathbf{p}_o)$ is then defined as,

$$\sum_j \|\mathbf{P}_o[j] - \mathbf{V}_o[j^*]\|^2 + \sum_k \|\mathbf{P}_h[k] - \mathbf{V}_h[k^*]\|^2. \quad (7)$$

Joint angle constraint E_{joint} . This term imposes restrictions on the 15 joints of the hand to ensure the resulting pose is natural. The three-dimensional rotation of a joint is parameterized using the axis-angle representation in MANO model, resulting in 45 joint angle parameters. A common solution when using MANO model is to optimize in the PCA space of 3D joint angles with an ℓ_2 regularizer [4, 74] for pose coefficients. However, we observed in practice that optimizing Eq. 1 in PCA space had less expressibility: some of the complex grasp poses in our dataset could not be accurately expressed in the PCA space, which was constructed with relatively simpler grasp poses and free-hand gestures. Instead, we optimize on joint angles directly and derive our own limits for each of the 45 joint parameters (please refer to supplementary material for these limits). As in [72], the joint angle constraint term $E_{\text{joint}}(\mathbf{p}_h^t)$ is given by,

$$\sum_{i=1}^{45} \max(\underline{a}_i - a[i], 0) + \max(a[i] - \bar{a}_i, 0), \quad (8)$$

where $a[i]$ denotes the i^{th} joint angle parameter for pose \mathbf{p}_h , and \underline{a}_i and \bar{a}_i correspond to its lower and upper limits.

Physical plausibility term E_{phy} . During optimization, the hand model might penetrate the object model, which is physically not possible. To avoid this, we add a repulsion term that pushes the object and the hand apart if they interpenetrate each other. For each hand vertex $\mathbf{V}_h[m]$, the amount of penetration $\Gamma[m]$ is taken as,

$$\Gamma[m] = \max(-\mathbf{n}_o(\mathbf{V}_o[m^*])^T (\mathbf{V}_h[m] - \mathbf{V}_o[m^*]), 0), \quad (9)$$

where $\mathbf{V}_o[m^*]$ is the vertex on object closest to hand vertex $\mathbf{V}_h[m]$, and the $\mathbf{n}_o(\cdot)$ operator provides the normal vector for a vertex. In words, the amount of penetration is estimated by projecting the vector joining a hand vertex and its nearest object vertex onto the normal vector at the object vertex location. The physical plausibility term is then defined as,

$$E_{\text{phy}}(\mathbf{p}_h^t, \mathbf{p}_o^t) = \sum_m \exp(w \Gamma[m]). \quad (10)$$

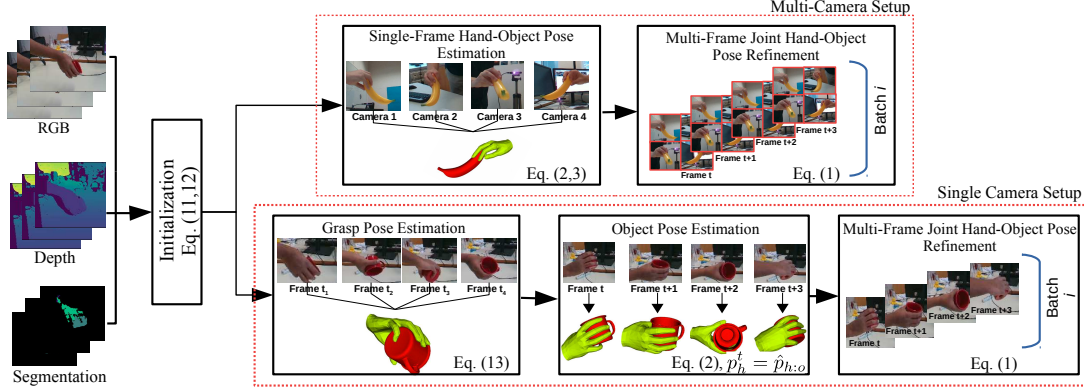


Figure 2: The different stages of the multi-camera and single camera setups. See Section 4 for more details.

We use $w = 5$ in practice, and only a subsampled set of vertices of the hand to compute E_{phy} efficiently.

Temporal consistency term E_{tc} . The previous terms are all applied to each frame independently. The temporal consistency term E_{tc} allows us to constrain together the poses for all the frames. We apply a 0-th and 1-st order motion model on both the hand and object poses:

$$E_{\text{tc}}(\mathbf{p}_h^t, \mathbf{p}_o^t, \mathbf{p}_h^{t-1}, \mathbf{p}_o^{t-1}, \mathbf{p}_h^{t-2}, \mathbf{p}_o^{t-2}) = \|\Delta_h^t\|^2 + \|\Delta_o^t\|^2 + \|\Delta_h^t - \Delta_h^{t-1}\|^2 + \|\Delta_o^t - \Delta_o^{t-1}\|^2,$$

where $\Delta_h^t = \mathbf{p}_h^t - \mathbf{p}_h^{t-1}$ and $\Delta_o^t = \mathbf{p}_o^t - \mathbf{p}_o^{t-1}$. Since we optimize a sum of these terms over the sequence, this effectively constrains all the poses together.

4. Optimization

Optimizing Eq. (1) is a challenging task, as it is a highly non-convex problem with many parameters to estimate. We therefore perform the optimization in multiple stages as shown in Fig. 2. These stages are different for multi-camera and single camera scenarios, and we detail them below.

4.1. Multi-Camera Setup

Initialization. In the multi-camera setup, we obtain a first estimate $\tilde{\mathbf{p}}_h^0$ for the hand pose in the first frame ($t = 0$) as,

$$\tilde{\mathbf{p}}_h^0 = \arg \min_{\mathbf{p}_h} \sum_{c=1}^{N_C} E_{\text{j2D}}(\mathbf{I}_c^0, \mathbf{p}_h) + \nu E_{\text{joint}}(\mathbf{p}_h). \quad (11)$$

We use the Dogleg optimizer [11] to perform this optimization. A first estimate $\tilde{\mathbf{p}}_o^0$ for the object pose in this frame is obtained using [46] trained by synthetically over-laying hands on YCB objects as explained in Section 3.2.

Single-frame joint pose optimization. We then obtain estimates $\tilde{\mathbf{p}}_h^t$ and $\tilde{\mathbf{p}}_o^t$ for all the other frames ($t = 1..N_F$)

by tracking. We minimize $(E_{\mathcal{D}}(\mathbf{p}_h^t, \mathbf{p}_o^t) + E_{\mathcal{C}}(\mathbf{p}_h^t, \mathbf{p}_o^t))$ w.r.t. \mathbf{p}_h^t and \mathbf{p}_o^t , using $\tilde{\mathbf{p}}_h^{t-1}$ and $\tilde{\mathbf{p}}_o^{t-1}$ for initialization.

Multi-frame joint pose optimization. We finally perform a full optimization of Eq. (1) w.r.t. \mathbf{p}_h^t and \mathbf{p}_o^t for $t = 0..N_F$ over all the frames simultaneously using estimates $\tilde{\mathbf{p}}_h^t$ and $\tilde{\mathbf{p}}_o^t$ for initialization. Due to memory constraints, we optimize Eq. (1) in batches instead of considering all the frames in sequence. We use a batch size of 20 frames with $\alpha = 20$, $\beta = 20$, $\gamma = 5 \times 10^{-5}$, $\delta = 50$, $\epsilon = 100$, $\zeta = 50$, and $\eta = 100$, and the Adam optimizer with learning rate of 0.01 for 100 iterations.

4.2. Single-Camera Setup

Initialization. In the single camera setup, as we assume that the grasp pose varies marginally across the sequence, we initially make the assumption that it remains constant throughout the sequence. In order to account for the minor changes which occur in practice, we relax this assumption in the latter stages of the optimization. We thus obtain initial estimates $\tilde{\mathbf{p}}_h^t$ for the hand poses as,

$$\{\tilde{\mathbf{p}}_h^t\}_t = \arg \min_{\{\mathbf{p}_h^t\}_t} \sum_t E_{\text{j2D}}(\mathbf{I}^t, \mathbf{p}_h^t) + \nu E_{\text{joint}}(\mathbf{p}_h^t), \quad (12)$$

where the joint angle parameters are constrained to be the same over all the frames, and only the rotation and translation parameters for the wrist joint can be different. In practice, we perform this optimization only over a random subset Ω of the frames to save time. We set $\nu = 50$, size of Ω to 20 and use the Dogleg optimizer [11]. First estimates $\tilde{\mathbf{p}}_o^t$ for the object poses in Ω are obtained as for the multi-camera setup.

From the $\tilde{\mathbf{p}}_h^t$ and the $\tilde{\mathbf{p}}_o^t$, we can compute the grasp pose in the *object coordinate system* $\tilde{\mathbf{p}}_{h:o}^t$, which is assumed to be constant at this stage. The initial estimate of the constant grasp pose $\tilde{\mathbf{p}}_{h:o}$ is taken as the average of $\{\tilde{\mathbf{p}}_{h:o}^t\}_{t \in \Omega}$.

Grasp pose estimation. We obtain a better estimate of the grasp pose $\hat{p}_{h:o}$ and object poses \hat{p}_o^t under fixed grasp pose assumption as,

$$\hat{p}_{h:o}, \{\hat{p}_o^t\}_{t \in \Omega} = \arg \min_{p_{h:o}, \{p_o^t\}_{t \in \Omega}} \sum_{t \in \Omega} E_D(f_{p_o^t}(p_{h:o}), p_o^t) + \zeta E_{\text{phy}}(f_{p_o^t}(p_{h:o}), p_o^t) + \epsilon E_{\text{joint}}(p_{h:o}), \quad (13)$$

wrt $p_{h:o}$ and p_o^t over the frames in Ω , using $\tilde{p}_{h:o}$ and \tilde{p}_o^t for initialization. $f_{p_o^t}(\cdot)$ converts the grasp pose to the hand pose in the world coordinate system given object pose p_o^t . This optimization accounts for the mutual occlusions between hand and object.

Object pose estimation. Having a good estimate $\hat{p}_{h:o}$ for the grasp pose, we obtain object pose estimates \hat{p}_o^t for all the frames by minimizing $E_D(f_{p_o^t}(\hat{p}_{h:o}), p_o^t)$ wrt p_o^t over each frame independently. We use \tilde{p}_o^{t-1} to initialize the optimization over p_o^t at frame t , except for \hat{p}_o^0 where \tilde{p}_o^0 is used. Note that the hand pose is not optimized in this stage.

Multi-frame joint hand+object pose refinement. In this final stage, we allow variations in the grasp pose across frames and introduce temporal constraints. We thus optimize Eq. (1) w.r.t. $\{(p_h^t, p_o^t)\}_{t=1}^{N_F}$ over all the frames simultaneously, using pose parameters $\hat{p}_o^t, \hat{p}_h^t = f_{\hat{p}_o^t}(\hat{p}_{h:o})$ estimated in the previous stages as initialization for p_o^t and p_h^t .

5. Monocular RGB based 3D Hand Pose

For establishing a baseline on our proposed dataset for single RGB image based hand pose prediction, we use a CNN architecture based on a Convolutional Pose Machine (CPM) [66] to predict the 2D hand joint locations $\{k_i\}_{i=1..21}$. In addition, we also predict the root-relative hand joint directions $\{d_i\}_{i=1..20}$, by adding an additional stage at the end of the CPM and replacing the last layer with a fully connected layer. More details on the architecture are provided in the supplementary material. The 3D joint locations and shape parameters of the hand are then obtained by fitting a MANO model to these predictions. The loss function for this fitting procedure is:

$$\sum_{i=1}^{21} \|\hat{k}_i - k_i\|^2 + \rho \sum_{i=1}^{20} (1 - \hat{d}_i \cdot d_i) + \sigma E_{\text{joint}}(p_h) + \tau \|\beta\|^2, \quad (14)$$

where $\hat{d}_i = \frac{p_h[i+1] - p_h[1]}{\|p_h[i+1] - p_h[1]\|}$, $\hat{k}_i = \text{proj}(J_{p_h}[i])$ and E_{joint} is defined in Eq. (8). We use $\rho = 10$, $\sigma = 5$, and $\tau = 1$.

6. Benchmarking HO-3D

In this section, we evaluate both our annotation method and our baseline for hand pose prediction from a single color image in hand+object interaction scenarios. We used

our 3D pose annotation method to annotate 68 sequences, totalling 77,558 frames of 10 different users manipulating one among 10 different objects from the YCB dataset. The image sizes are 640×480 pixels for both the color and depth cameras, and we used 5 synchronized cameras in our multi-camera setup. The cameras were synchronized with an accuracy of 5ms.

6.1. Evaluation of the Annotation Method

For validating the accuracy of our annotation method, we manually annotated the 3D locations of the 3D joints in randomly selected frames of a sequence, by relying on the consolidated point cloud from the 5 cameras. We then compared these locations to the ones predicted with our method (explained in Section 4.1) using the multi-camera setup.

As shown in the last column of Table 2, our method achieves an average joint error accuracy of lower than 8mm on average, with an Area Under the Curve metric (AUC) of 0.79. This metric is comparable with the results reported for the recent FreiHAND dataset [74] (AUC=0.791). Note that the occlusions in our dataset are higher due to larger objects and that we do not use green screens.

To analyze the influence of the different terms in Eq. (1), we run the optimization of Eq. (1) by enabling only a subset of these terms, and report the results in Table 2. While E_{silh} and E_{dpt} terms alone cannot provide good pose estimates, together they provide better estimates as it leads to a loss function with less local minima. The E_{3D} term provides a minor improvement in estimates but speeds up the convergence. Though the physical plausibility term E_{phy} does not help in improving the pose estimates, it results in more natural grasps. The last two columns show the effect of multi-frame based joint optimization compared to single-frame based optimization when all the terms are considered. The multi-frame multi-camera based optimization over all the terms improves the accuracy by about 15%.

The accuracy of the single camera based annotation method is calculated by considering the annotations from the multi-camera method as ground truth for a given sequence. More specifically, for a sequence of 1000 frames, we compute the average difference between hand+object mesh vertices obtained from single and multi-camera setups. Further, we calculate the accuracy after each stage of the single-camera setup. The results are given in Table 3. The estimated poses with these two methods are consistent with each other with an average mesh error of 0.77cm and 0.45cm for hand and object, respectively. The final refinement stage yields a 15% improvement in accuracy.

6.2. Evaluation of Hand Pose Prediction Method

We trained our single frame hand pose prediction method explained in Section 5 on 66,034 frames from our HO-3D dataset. We evaluated it on a test set of 13 sequences cap-

Terms	Initialization	Single-frame Optimization						Multi-frame Opt. (Eq. 1)
		E_{silh}	E_{dpt}	$E_{silh} + E_{dpt}$	$E_{silh} + E_{dpt} + E_{3D}$	$E_{silh} + E_{dpt} + E_{3D} + E_{phy}$	$E_{silh} + E_{dpt} + E_{3D} + E_{phy} + E_{tc}$	
mean (std)	4.20 (± 3.32)	1.17 (± 1.12)	2.22 (± 1.22)	1.04 (± 0.43)	0.98 (± 0.40)	0.99 (± 0.40)	0.92 (± 0.34)	0.77 (± 0.29)

Table 2: Evaluation of the accuracy for the multi-camera setup. We report the average hand-joint errors (in cm) for different combinations of the terms in Eq. (1). The final error is comparable to the recent FreiHAND dataset [74].

Stages	Init.	Grasp Pose Est.	Object Pose Est.	Refinement
Hand	5.40	3.60	0.91	0.77
Object	4.02	4.02	0.52	0.45

Table 3: Evaluation of the accuracy for the single-camera setup. The accuracy (average mesh error in cm) is measured at each stage of optimization by comparing with the annotations from multi-camera setup. The results show that the annotation quality of our single camera method is similar to that of the multi-camera setup.

Method	Mesh Error↓	F@5mm↑	F@15mm↑	Joint Error↓
Joints2D	1.14	0.49	0.93	3.14
Joints2D + Dir. Vec.	1.06	0.51	0.94	3.04
[18]	1.30	0.42	0.90	8.31

Table 4: Evaluation of different methods for single frame hand pose prediction. The Mesh Error (in cm) and F -score are obtained after aligning the predicted meshes with ground truth meshes. The Mean joint error (in cm) is obtained after aligning the position of the root joint and overall scale with the ground truth. Hand pose prediction using joint direction predictions along with 2D joint predictions provides better accuracy than directly predicting the MANO parameters as in [18].

tured from different viewpoints and totaling 11,524 frames. The test set sequences also contain subjects and objects not present in the training set.

We report three different metrics from previous works: Mean joint position error after aligning the position of the root joint and global scale with ground truth [73]; Mesh error measuring the average Euclidean distance between predicted and ground truth mesh vertices [74]; and the F -score [74], defined as the harmonic mean between recall and precision between two meshes given a distance threshold. The mesh error and F -score are obtained after aligning the predicted meshes using Procrustes alignment with the ground truth meshes and hence does not measure the accuracy of wrist joint rotation. The mean joint error on the other hand considers wrist joint location as the 3D points are not rotated before evaluation.

To understand the effect of joint direction predictions on the overall accuracy, we evaluate the results of the MANO fitting by dropping the second term in Eq. (14). We also compare our results with the hand branch of [18], a very recent work that predicts the MANO pose and shape parameters directly from a single RGB image, retrained on our dataset. As shown in Table 4, predicting joint directions along with 2D joint locations significantly improves




Figure 3: Qualitative results of our single color image hand pose estimation method. It can recover hand poses even when the hand is heavily occluded by objects and in cluttered scenes. The last row shows it can handle unseen objects.

the hand pose estimation accuracy. It can also be inferred that predicting 2D hand joint locations and fitting MANO model to them is more accurate than direct MANO parameter predictions as in [18]. Qualitative results are shown in Fig. 3. The last row shows that our method robustly predicts hand poses even when interacting with unknown objects.

7. Conclusion

We introduced a fully automatic method to annotate images of a hand manipulating an object with their 3D poses, even under large occlusions, by exploiting temporal consistency. We also introduced the first markerless dataset of color images for benchmarking 3D hand+object pose estimation. To demonstrate the usefulness of our dataset, we proposed a method for predicting the 3D pose of the hand from a single color image. Another future application is the joint estimation of hand+object poses from a single RGB frame.

The lack of high quality segmentations (we had to use a synthetic dataset as explained in Section 3) of the hand sometimes affects accuracy. Improving these segmentations and/or introducing an attraction term and physics constraints as in [52, 61] would further improve our annotations.

8. Acknowledgment

This work was supported by the Christian Doppler Laboratory for Semantic 3D Computer Vision, funded in part by Qualcomm Inc.

References

- [1] YCB Benchmarks – Object and Model Set. <http://ycbbenchmarks.org/>. 2
- [2] 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 (ECCV)*, pages 640–653, Firenze, October 2012. 3
- [3] Sven Bambach, Stefan Lee, David J. Crandall, and Chen Yu. Lending a Hand: Detecting Hands and Recognizing Activities in Complex Egocentric Interactions. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 1949–1957, 2015. 3
- [4] Adnane Boukhayma, Rodrigo de Bem, and Philip HS Torr. 3d hand shape and pose from images in the wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10843–10852, 2019. 5
- [5] Anders Glent Buch, Lilita Kiforenko, and Dirk Kraft. Rotational subgroup voting and pose clustering for robust 3d object recognition. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 4137–4145, United States, 2017. 2
- [6] Ian M. Bullock, Thomas Feix, and Aaron M. Dollar. The Yale Human Grasping Dataset: Grasp, Object, and Task Data in Household and Machine Shop Environments. *The International Journal of Robotics Research*, 34(3):251–255, 2015. 3
- [7] Minjie Cai, Kris M. Kitani, and Yoichi Sato. A Scalable Approach for Understanding the Visual Structures of Hand Grasps. In *The IEEE International Conference on Robotics and Automation (ICRA)*, 2015. 3
- [8] Qixin Cao and Haoruo Zhang. Combined Holistic and Local Patches for Recovering 6D Object Pose. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 2219–2227, 2017. 2
- [9] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking Atrous Convolution for Semantic Image Segmentation. *CoRR*, abs/1706.05587, 2017. 4
- [10] Chiho Choi, Sang Ho Yoon, Chin-Ning Chen, and Karthik Ramani. Robust Hand Pose Estimation During the Interaction with an Unknown Object. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 3142–3151, 2017. 3
- [11] Andrew R. Conn, Nicholas I. M. Gould, and Philippe L. Toint. *Trust-Region Methods*. SIAM, 2000. 6
- [12] Martin de La Gorce, David J. Fleet, and Nikos Paragios. Model-Based 3D Hand Pose Estimation from Monocular Video. *IEEE Trans. Pattern Anal. Mach. Intell.(PAMI)*, 33(9):1793–1805, 2011. 2
- [13] Alireza Fathi, Xiaofeng Ren, and James M. Rehg. Learning to Recognize Objects in Egocentric Activities. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3281–3288, 2011. 3
- [14] Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, and Tae-Kyun Kim. First-Person Hand Action Benchmark with RGB-D Videos and 3D Hand Pose Annotations. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 409–419, 2018. 1, 2, 3, 4
- [15] Lihao Ge, Yujun Cai, Junwu Weng, and Junsong Yuan. Hand PointNet: 3D Hand Pose Estimation Using Point Sets. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8417–8426, 2018. 2
- [16] Duncan Goudie and Aphrodite Galata. 3D Hand-Object Pose Estimation from Depth with Convolutional Neural Networks. In *IEEE International Conference on Automatic Face & Gesture Recognition*, 2017. 3
- [17] Shreyas Hampali, Markus Oberweger, Mahdi Rad, and Vincent Lepetit. HO-3D: a Multi-User, Multi-Object Dataset for Joint 3D Hand-Object Pose Estimation. *CoRR*, abs/1907.01481, 2019. 5
- [18] Yana Hasson, Gül Varol, Dimitrios Tzionas, Igor Kalyvayk, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning Joint Reconstruction of Hands and Manipulated Objects. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11807–11816, 2019. 1, 2, 3, 4, 8
- [19] Paul Henderson and Vittorio Ferrari. Learning Single-Image 3D Reconstruction by Generative Modelling of Shape, Pose and Shading. *International Journal of Computer Vision (IJCV)*, pages 1573–1405, 2019. 4, 5
- [20] Yinlin Hu, Joachim Hugonot, Pascal Fua, and Mathieu Salzmann. Segmentation-Driven 6d Object Pose Estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3385–3394, 2019. 2
- [21] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total Capture: A 3D Deformation Model for Tracking Faces, Hands, and Bodies. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8320–8329, 2018. 3, 4
- [22] Wadim Kehl, Fausto Milletari, Federico Tombari, Slobodan Ilic, and Nassir Navab. Deep Learning of Local RGB-D Patches for 3D Object Detection and 6D Pose Estimation. In *European Conference on Computer Vision (ECCV)*, pages 205–220, 2016. 2
- [23] Cem Keskin, Furkan Kiraç, Yunus Emre Kara, and Lale Akarun. Hand Pose Estimation and Hand Shape Classification Using Multi-Layered Randomized Decision Forests. In *European Conference on Computer Vision (ECCV)*, pages 852–863, 2012. 2
- [24] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *The International Conference on Learning Representations (ICLR)*, 2015. 2
- [25] Mia Kokic, Danica Kragic, and Jeannette Bohg. Learning to Estimate Pose and Shape of Hand-Held Objects from RGB Images. In *The IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3980–3987, 11 2019. 3
- [26] Nikolaos Kyriazis and Antonis A. Argyros. Physically Plausible 3D Scene Tracking: The Single Actor Hypothesis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013. 1, 3
- [27] Nikolaos Kyriazis and Antonis A. Argyros. Scalable 3D Tracking of Multiple Interacting Objects. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 3

- [28] Stan Melax, Leonid Keselman, and Sterling Orsten. Dynamics Based 3D Skeletal Hand Tracking. In *Proceedings of Graphics Interface*, 2013. 2
- [29] Andrew T. Miller and Peter K. Allen. Graspit! a Versatile Simulator for Robotic Grasping. *Robotics Automation Magazine*, pages 110–122, 2004. 1
- [30] Chaitanya Mitash, Abdeslam Boularias, and Kostas E. Bekris. Improving 6D Pose Estimation of Objects in Clutter via Physics-Aware Monte Carlo Tree Search. In *The IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8, 2018. 2
- [31] Franziska Mueller, Florian Bernard, Oleksandr Sotnychenko, Dushyant Mehta, Srinath Sridhar, Dan Casas, and Christian Theobalt. Generated Hands for Real-Time 3D Hand Tracking from Monocular RGB. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 49–59, 2018. 1, 2, 3, 4
- [32] Franziska Mueller, Dushyant Mehta, Oleksandr Sotnychenko, Srinath Sridhar, Dan Casas, and Christian Theobalt. Real-Time Hand Tracking Under Occlusion from an Egocentric Rgb-D Sensor. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 2, 3
- [33] Battushig Myanganbayar, Cristina Mata, Gil Dekel, Boris Katz, Guy Ben-Yosef, and Andrei Barbu. Partially Occluded Hands: A Challenging New Dataset for Single-Image Hand Pose Estimation. In *The Asian Conference on Computer Vision (ACCV)*, pages 85–98, 2018. 3
- [34] Natalia Neverova, Christian Wolf, Florian Nebout, and Graham W. Taylor. Hand Pose Estimation through Semi-Supervised and Weakly-Supervised Learning. *Comput. Vis. Image Underst. (CVIU)*, 164:56–67, 2017. 2
- [35] Markus Oberweger and Vincent Lepetit. DeepPrior++: Improving Fast and Accurate 3D Hand Pose Estimation. In *The IEEE International Conference on Computer Vision Workshops (ICCV)*, pages 585–594, 2017. 1, 2
- [36] Markus Oberweger, Mahdi Rad, and Vincent Lepetit. Making Deep Heatmaps Robust to Partial Occlusions for 3D Object Pose Estimation. In *European Conference on Computer Vision (ECCV)*, pages 125–141, 2018. 2
- [37] Markus Oberweger, Gernot Riegler, Paul Wohlhart, and Vincent Lepetit. Efficiently Creating 3D Training Data for Fine Hand Pose Estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4957–4965, 2016. 1
- [38] Iasonas Oikonomidis, Nikolaos Kyriazis, and Antonis A. Argyros. Full DoF Tracking of a Hand Interacting with an Object by Modeling Occlusions and Physical Constraints. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 2088–2095, 2011. 3
- [39] Paschalis Panteleris, Nikolaos Kyriazis, and Antonis A. Argyros. 3D Tracking of Human Hands in Interaction with Unknown Objects. In *Proceedings of the British Machine Vision Conference 2015 (BMVC)*, pages 123.1–123.12, 2015. 3
- [40] Paschalis Panteleris, Iason Oikonomidis, and Antonis A. Argyros. Using a Single RGB Frame for Real Time 3D Hand Pose Estimation in the Wild. In *The IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 436–445, 2018. 2
- [41] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pvnnet: Pixel-Wise Voting Network for 6DoF Pose Estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4561–4570, 2019. 2
- [42] Tu-Hoa Pham, Nikolaos Kyriazis, Antonis A. Argyros, and Abderrahmane Kheddar. Hand-Object Contact Force Estimation from Markerless Visual Tracking. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 40(12):2883–2896, 2018. 3
- [43] Tu-Hoa Pham, Abderrahmane Kheddar, Ammar Qammaz, and Antonis A. Argyros. Capturing and Reproducing Hand-Object Interactions through Vision-Based Force Sensing. In *Object Understanding for Interaction*, 2015. 3
- [44] Tu-Hoa Pham, Abderrahmane Kheddar, Ammar Qammaz, and Antonis A. Argyros. Towards Force Sensing from Vision: Observing Hand-Object Interactions to Infer Manipulation Forces. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 3
- [45] Chen Qian, Xiao Sun, Yichen Wei, Xiaou Tang, and Jian Sun. Realtime and Robust Hand Tracking from Depth. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 1, 2
- [46] Mahdi Rad and Vincent Lepetit. BB8: A Scalable, Accurate, Robust to Partial Occlusion Method for Predicting the 3D Poses of Challenging Objects Without Using Depth. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 6
- [47] Mahdi Rad, Markus Oberweger, and Vincent Lepetit. Domain Transfer for 3D Pose Estimation from Color Images Without Manual Annotations. In *The Asian Conference on Computer Vision (ACCV)*, 2018. 1, 3
- [48] Grégory Rogez, James Steven Supancic III, and Deva Ramanan. Understanding Everyday Hands in Action from RGB-D Images. In *The IEEE International Conference on Computer Vision (ICCV)*, 2015. 3
- [49] Grégory Rogez, Maryam Khademi, James Steven Supancic III, J. M. M. Montiel, and Deva Ramanan. 3D Hand Pose Detection in Egocentric RGB-D Images. In *European Conference on Computer Vision (ECCV)*, 2014. 3
- [50] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied Hands: Modeling and Capturing Hands and Bodies Together. *ACM Trans. Graph.*, 36(6):245, 2017. 2, 4
- [51] Toby Sharp, Cem Keskin, Duncan P. Robertson, Jonathan Taylor, Jamie Shotton, David Kim, Christoph Rhemann, Ido Leichter, Alon Vinnikov, Yichen Wei, Daniel Freedman, Pushmeet Kohli, Eyal Krupka, Andrew W. Fitzgibbon, and Shahram Izadi. Accurate, Robust, and Flexible Real-Time Hand Tracking. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI*, pages 3633–3642, 2015. 2
- [52] Srinath Sridhar, Franziska Mueller, Michael Zollhöfer, Dan Casas, Antti Oulasvirta, and Christian Theobalt. Real-Time Joint Tracking of a Hand Manipulating an Object from RGB-D Input. In *European Conference on Computer Vision (ECCV)*, volume 9906, pages 294–310, 2016. 1, 2, 3, 8
- [53] Srinath Sridhar, Antti Oulasvirta, and Christian Theobalt. Interactive Markerless Articulated Hand Motion Tracking Using RGB and Depth Data. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 2456–2463, 2013. 2

- [54] Xiao Sun, Yichen Wei, Shuang Liang, Xiaoou Tang, and Jian Sun. Cascaded Hand Pose Regression. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 824–832, 2015. 1
- [55] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3D Orientation Learning for 6D Object Detection from RGB Images. In *European Conference on Computer Vision (ECCV)*, pages 712–729, 2018. 1
- [56] David Joseph Tan, Thomas J. Cashman, Jonathan Taylor, Andrew W. Fitzgibbon, Daniel Tarlow, Sameh Khamis, Shahram Izadi, and Jamie Shotton. Fits Like a Glove: Rapid and Reliable Hand Shape Personalization. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5610–5619, June 2016. 4
- [57] Danhang Tang, Hyung Jin Chang, Alykhan Tejani, and Tae-Kyun Kim. Latent Regression Forest: Structured Estimation of 3D Articulated Hand Posture. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3786–3793, 2014. 1
- [58] Bugra Tekin, Federica Bogo, and Marc Pollefeys. H+O: Unified Egocentric Recognition of 3D Hand-Object Poses and Interactions. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4511–4520, 2019. 3
- [59] Jonathan Tompson, Murphy Stein, Yann LeCun, and Ken Perlin. Real-Time Continuous Pose Recovery of Human Hands Using Convolutional Networks. *ACM Trans. Graph.*, 33:169:1–169:10, 2014. 2
- [60] Aggeliki Tsoli and Antonis A. Argyros. Joint 3D Tracking of a Deformable Object in Interaction with a Hand. In *European Conference on Computer Vision (ECCV)*, pages 504–520, 2018. 3
- [61] 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 (IJCV)*, 118(2):172–193, 2016. 1, 3, 8
- [62] Dimitrios Tzionas and Juergen Gall. 3D Object Reconstruction from Hand-Object Interactions. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 729–737, 2015. 3
- [63] Dimitrios Tzionas, Abhilash Srikantha, Pablo Aponte, and Juergen Gall. Capturing Hand Motion with an RGB-D Sensor, Fusing a Generative Model with Salient Points. In *German Conference on Pattern Recognition (GCPR)*, pages 277–289, 2014. 2, 3
- [64] Chengde Wan, Thomas Probst, Luc Van Gool, and Angela Yao. Self-Supervised 3D Hand Pose Estimation through Training by Fitting. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10853–10862, June 2019. 3
- [65] Robert Y. Wang, Sylvain Paris, and Jovan Popovic. 6D Hands: Markerless Hand-Tracking for Computer Aided Design. In *ACM Symposium on User Interface Software and Technology*, pages 549–558, 2011. 3
- [66] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional Pose Machines. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4732, 2016. 5, 7
- [67] Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. Monocular Total Capture: Posing Face, Body, and Hands in the Wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10965–10974, 2019. 5
- [68] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes. In *Robotics: Science and Systems XIV (RSS)*, 2018. 2, 4
- [69] Chi Xu, Lakshmi Narasimhan Govindarajan, Yu Zhang, and Li Cheng. Lie-X: Depth Image Based Articulated Object Pose Estimation, Tracking, and Action Recognition on Lie Groups. *International Journal of Computer Vision (IJCV)*, 123:454–478, 2017. 2
- [70] Qi Ye, Shanxin Yuan, and Tae-Kyun Kim. Spatial Attention Deep Net with Partial PSO for Hierarchical Hybrid Hand Pose Estimation. In *European Conference on Computer Vision (ECCV)*, pages 346–361, 2016. 2
- [71] Shanxin Yuan, Qi Ye, Bjorn Stenger, Siddhant Jain, and Tae-Kyun Kim. Big Hand 2.2m Benchmark: Hand Pose Data Set and State of the Art Analysis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [72] Xingyi Zhou, Qingfu Wan, Wei Zhang, Xiangyang Xue, and Yichen Wei. Model-based deep hand pose estimation. In *The International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2421–2427, 2016. 5
- [73] Christian Zimmermann and Thomas Brox. Learning to Estimate 3D Hand Pose from Single RGB Images. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 4913–4921, 2017. 1, 2, 3, 8
- [74] Christian Zimmermann, Duygu Ceylan, Jimei Yang, Bryan C. Russell, Max J. Argus, and Thomas Brox. FreiHAND: A Dataset for Markerless Capture of Hand Pose and Shape from Single RGB Images. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 813–822, 2019. 1, 2, 4, 5, 7, 8