

# H+O: Unified Egocentric Recognition of 3D Hand-Object Poses and Interactions

Bugra Tekin<sup>1</sup>

Federica Bogo<sup>1</sup>

Marc Pollefeys<sup>1,2</sup>

<sup>1</sup> Microsoft

<sup>2</sup> ETH Zürich

## Abstract

We present a unified framework for understanding 3D hand and object interactions in raw image sequences from egocentric RGB cameras. Given a single RGB image, our model jointly estimates the 3D hand and object poses, models their interactions, and recognizes the object and action classes with a single feed-forward pass through a neural network. We propose a single architecture that does not rely on external detection algorithms but rather is trained end-to-end on single images. We further merge and propagate information in the temporal domain to infer interactions between hand and object trajectories and recognize actions. The complete model takes as input a sequence of frames and outputs per-frame 3D hand and object pose predictions along with the estimates of object and action categories for the entire sequence. We demonstrate state-of-the-art performance of our algorithm even in comparison to the approaches that work on depth data and ground-truth annotations.

## 1. Introduction

Human behavior can be characterized by the individual actions a person takes in interaction with the surrounding objects and the environment. A significant amount of research has focused on visual understanding of humans [8, 25, 36, 40, 59, 65, 70, 74] and objects [4, 28, 64, 66], in isolation from each other. However, the problem of jointly understanding humans and objects, although crucial for a semantically meaningful interpretation of the visual scene, has received far less attention. In this work, we propose, for the first time, a unified method to jointly recognize 3D hand and object poses, and their interactions from egocentric monocular color images. Our method jointly estimates the hand and object poses in 3D, models their interactions and recognizes the object and activity classes. An example result is shown in Fig. 1. Our unified framework is highly relevant for augmented and virtual reality [62], fine-grained recognition of people’s actions, robotics and telepresence.



**Figure 1:** Unified reasoning on first-person views. Our method takes as input color images and produces a comprehensive egocentric scene interpretation. We estimate simultaneously 3D hand and object poses (shown as skeletons and 3D bounding boxes), object class (e.g. juice bottle) and action category (e.g. pouring).

Capturing hands in action while taking into account the objects in contact is an extremely challenging problem. Jointly reasoning about hands and objects from moving, egocentric cameras is even more challenging as this would require understanding of the complex and often subtle interactions that take place in cluttered real-world scenes where the hand is often occluded by the object or the viewpoint. Recent research in computer vision has successfully addressed some of the challenges in joint understanding of hands and objects for depth and multi-camera input. Sridhar et al. [60] have demonstrated that accounting jointly for hands and objects helps predict the 3D pose more accurately than models that ignore interaction. Pioneering works by [21, 37, 41, 68] have proposed ways to model hand-object interactions to increase robustness and accuracy in recovering the hand motion.

Most of these works, however, are limited by the following factors: Firstly, they either rely on active depth sensors or multi-camera systems. Depth sensors are power hungry and less prevalent than regular RGB cameras. On the other hand, multi-camera systems are impractical due to the cost and effort in setting up a calibrated and synchronous system of sensors. Secondly, they do not reason about the action the subject is performing. While estimating the 3D hand pose is crucial for many applications in robotics and graphics, the sole knowledge of the pose lacks semantic meaning about the actions of the subject. Thirdly, they focus mostly

on only capturing hand motion without recovering the object pose in 3D [21, 37], and, therefore, lack environmental understanding.

Our method aims at tackling all these issues. To this end, we propose an approach for predicting simultaneously 3D hand and object poses, object classes and action categories through a novel data-driven architecture. Our model jointly produces 3D hand and object pose estimates, action and object classes from a single image and requires neither external region proposals nor pre-computed detections. Our method further models the temporal nature of 3D hand and object motions to recognize actions and infer interactions.

Our contributions can be summarized as follows:

- We propose a unified framework for recognizing 3D hand and object interactions by simultaneously solving four tasks in a feed-forward pass through a neural network: 3D hand pose estimation, object pose estimation, object recognition and activity classification. Our method operates on monocular color images and relies on joint features that are shared among all the tasks.
- We introduce a novel single-shot neural network framework that jointly solves for the 3D articulated and rigid pose estimation problems within the same architecture. Our scheme relies on a common output representation for both hands and objects that parametrizes their pose with 3D control points. Our network directly predicts the control points in 3D rather than in 2D, in contrast to the common single-shot neural network paradigm [48, 66], dispenses with the need to solve a 2D-to-3D correspondence problem [32] and yields large improvements in accuracy.
- We present a temporal model to merge and propagate information in the temporal domain, explicitly model interactions and infer relations between hand and objects, directly in 3D.

In Section 4, we show quantitatively that these contributions allow us to achieve better overall performance in targeted tasks, while running at real-time speed and not requiring detailed 3D hand and object models. Our approach, which we call *Hand + Object (H+O)*, achieves state-of-the-art results on challenging sequences, and outperforms existing approaches that rely on the ground-truth pose annotations and depth data.

## 2. Related Work

We now review existing work on 3D hand and object pose estimation – both jointly and in isolation – and action recognition, with a focus on egocentric scenarios.

**Hands and Objects.** Many approaches in the literature tackle the problem of estimating either hand or object pose in isolation.

Brachman et al. [4] recover 6D object pose from single RGB images using a multi-stage approach, based on regression forests. More recent approaches [28, 46] rely on Convolutional Neural Networks (CNNs). BB8 [46] uses CNNs to roughly segment the object and then predict the 2D locations of the object’s 3D bounding box, projected in image space. 6D pose is then computed from these estimates via PnP [32]. SSD-6D [28] predicts 2D bounding boxes together with an estimate of the object pose. These methods need a detailed textured 3D object model as input, and require a further pose refinement step to improve accuracy. Tekin et al. [66] overcome these limitations by introducing a single-shot architecture which predicts 2D projections of the object’s 3D bounding box in a single forward pass, at real-time speed. All these approaches do not address the problem of estimating object pose in hand-object interaction scenarios, where objects might be largely occluded.

3D hand pose and shape estimation in egocentric views has started receiving attention recently [7, 22, 37, 51, 70, 72, 73]. The problem is challenging with respect to third-person scenarios [58, 74], due to self-occlusions and the limited amount of training data available [37, 39, 51, 70]. Mueller et al. [37] train CNNs on synthetic data and combine them with a generative hand model to track hands interacting with objects from egocentric RGB-D videos. This hybrid approach has then been extended to work with RGB videos [36]. Iqbal et al. [25] estimate 3D hand pose from single RGB images, from both first- and third-person views, regressing 2.5D heatmaps via CNNs. These methods focus on hand pose estimation and try to be robust in the presence of objects, but do not attempt to model hand and objects together.

While reasoning about hands in action, object interactions can be exploited as additional constraints [41, 49, 51, 53]. By observing that different object shapes induce different hand grasps, the approaches in [7, 51] discriminatively estimate 3D hand pose from depth input. Model-based approaches [54] have also been proposed to jointly estimate hand and object parameters at a finer level of detail. However, most approaches focus on third-view scenarios, taking depth as input [41, 43, 67, 68].

To our knowledge, no approach in the literature jointly estimates 3D hand and object pose from RGB video only.

**Action Recognition.** While action recognition is a long-standing problem in computer vision [3, 10, 18, 26, 31, 38, 69], first-person action recognition started to emerge as an active field only recently, largely due to the advent of consumer-level wearable sensors and large egocentric datasets [9, 20, 44, 52]. First-person views bring upon

unique challenges to action recognition due to fast camera motion, large occlusions and background clutter [33].

Early approaches for egocentric action recognition rely on motion cues [29, 50, 55]. In particular, [29] uses optical flow-based global motion descriptors to categorize “ego-actions” across sports genres. [45] feeds sparse optical flow to 3D CNNs to index egocentric videos. Motion and appearance cues are used in conjunction with depth in [63]. In addition to motion, features based on gaze information [13], head motion [33], and, recently, CNN-learned features [35] have been proposed by a number of studies. Another line of work has focused specifically on hand and object cues for first person action recognition [12, 16, 27, 57, 61]. Pirsivash and Ramanan [44] explore active object detection as an auxiliary task for activity recognition. Koppula et al. [30] learn a model of object affordances to understand activities from RGB-D videos. EgoNet [2] detects “action objects”, *i.e.* objects linked to visual or tactile interactions, from first-person RGB-D images. Fathi et al. [12, 14] use motion cues to segment hands and objects, and then extract features from these foreground regions. All these approaches, however, focus on 2D without explicitly modeling hand-object interactions in 3D. Recently, Garcia-Hernando et al. [17] demonstrate that 3D heuristics are beneficial to first person action recognition. However, they work on depth input and rely on ground-truth object poses.

Similarly to us, Cai et al. [34] propose a structured approach where grasp types, object attributes and their contextual relationships are analyzed together. However, their single-image framework does not consider the temporal dimension. Object Relation Network [1] models contextual relationships between detected semantic object instances, through space and time. All these approaches aim at understanding the scene only in 2D. Here, we model more complex hand and object attributes in 3D, as well as their temporal interaction.

### 3. Method

Our goal is to construct comprehensive interpretations of egocentric scenes from raw image sequences to understand human activities. To this end, we propose a unified framework to jointly estimate 3D hand and object poses and recognize object and action classes.

#### 3.1. Overview

The general overview of our *Hand+Object* model is given in Fig. 2. Our model takes as input a sequence of color frames  $\mathbf{I}^t$  ( $1 \leq t \leq N$ ) and predicts *per-frame* 3D hand and object poses, object classes and action categories, along with *per-sequence* interaction classes. Here, we define actions as *verbs*, *e.g.* “pour”, and interactions as (*verb, noun*) pairs, *e.g.* “pour juice”. We represent hand and object poses with  $N_c$  3D control points. Details about control

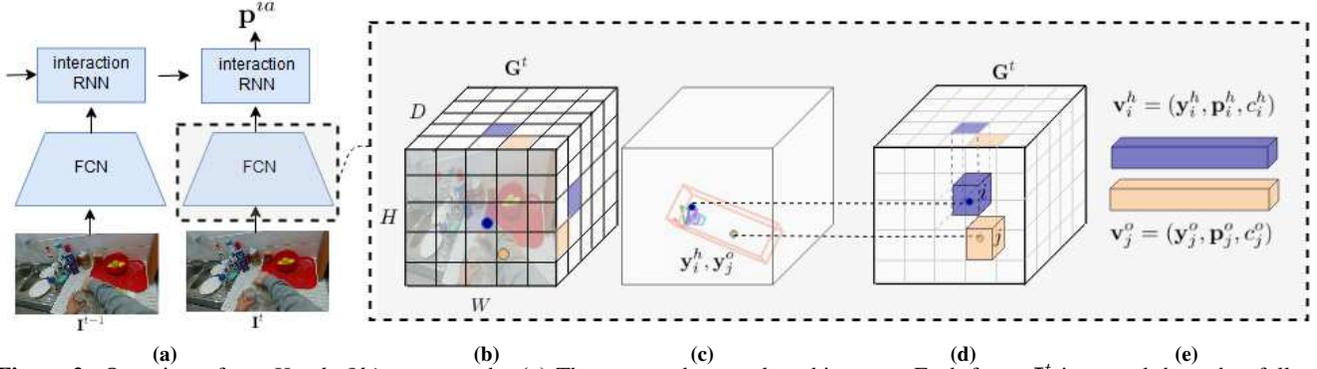
point definitions are provided in Sec. 3.2. We denote the number of object classes by  $N_o$ , the number of actions by  $N_a$  and the number of interactions by  $N_{ia}$ .

Our model first processes each frame,  $\mathbf{I}^t$ , of a sequence with a fully convolutional network (Fig. 2a) and divides the input image into a regular grid  $\mathbf{G}^t$  containing  $H \times W \times D$  cells that span the 3D scene in front of the camera (Fig. 2b). We keep the target values of our network for hands and objects in tensor  $\mathbf{G}^t$  (Fig. 2c-d). Namely, the target values for a hand or an object at a specific cell location,  $i \in H \times W \times D$ , are placed in the  $i$ -th cell of  $\mathbf{G}^t$  in the form of a multi-dimensional vector,  $\mathbf{v}_i$ . To be able to jointly estimate the pose of a hand and an object potentially occluding each other, we allow each cell to store two separate sets of values, one for the hand, denoted by  $\mathbf{v}_i^h$ , and one for the object, denoted by  $\mathbf{v}_i^o$  (Fig. 2e). Vector  $\mathbf{v}_i^h$  stores the control points for hand pose,  $\mathbf{y}_i^h \in \mathbb{R}^{3N_c}$ , action probabilities,  $\mathbf{p}_i^a \in \mathbb{R}^{N_a}$ , and an overall confidence value for the hand pose estimation,  $c_i^h \in [0, 1]$ . Vector  $\mathbf{v}_i^o$  stores the control points for object pose,  $\mathbf{y}_i^o \in \mathbb{R}^{3N_c}$ , object class probabilities,  $\mathbf{p}_i^o \in \mathbb{R}^{N_o}$ , and an overall confidence value for the object pose estimation,  $c_i^o \in [0, 1]$ .

We train our single pass network based on [48] to be able to predict these target values. At test time, predictions at cells with low confidence values, *i.e.* where the hands or objects of interest are not present, are pruned. All these predictions are obtained with a single forward pass in the network. While very efficient, this step works on each frame independently, thus ignoring the temporal dimension. Therefore, we add a recurrent module to integrate information across frames and model the interaction between hands and objects (Fig. 2a). This module takes as input hand and object predictions with high confidence values, and outputs a probability vector,  $\mathbf{p}^{ia} \in \mathbb{R}^{N_{ia}}$ , over interaction classes. In the following sections, we describe each of these components in more detail.

#### 3.2. Joint 3D Hand-Object Pose Estimation

In the context of rigid object pose estimation, [46, 66] regress the 2D locations of 8 keypoints – namely, the projections of the 8 corners of the object’s 3D bounding box. The object’s 6D pose is then estimated using a PnP algorithm [32]. Adopting a similar approach would not work in our case, since we aim at estimating also the articulated 3D pose of the hand. To tackle this problem and jointly estimate 3D articulated and rigid pose within the same architecture, we propose to use a common output representation for both hands and objects. To this end, we parametrize both hand and object poses jointly with 3D control points, corresponding to 21 skeleton joints for the hand pose and 3D locations of object keypoints, corresponding to the locations on the 3D object bounding box. For simplicity, we choose  $N_c = 21$ , and define 8 keypoints for the objects as



**Figure 2:** Overview of our *Hand+Object* approach. (a) The proposed network architecture. Each frame  $\mathbf{I}^t$  is passed through a fully-convolutional network to produce a 3D regular fixed grid  $\mathbf{G}^t$ . (b) The  $H \times W \times D$  grid showing cells responsible for recognizing hands and objects. (c) Each cell predicts the 3D hand pose and object bounding box coordinates in a 3D grid. (d) The output tensor from our network, in which the target values for hands and objects are stored. (e) Cells are associated with a vector that contains target values for hand and object pose, object and action categories and an overall confidence value. Predictions with high confidence values are then passed through the *interaction RNN* to propagate information in the temporal domain and model interactions in 3D between hands and objects.

proposed in [46, 66], along with the 12 edge midpoints and the centroid of the 3D bounding box.

Adopting a coherent parameterization for hand and object simplifies the regression task. We subdivide the input image into a grid of  $H \times W$  cells, and further discretize depth into  $D$  cells. Note that discretization is defined in pixel space for the first two dimensions, and in metric space for depth. Therefore each cell has a size of  $C_u \times C_v$  pixels,  $\times C_z$  meters. Within each cell, we predict offsets  $\Delta u$ ,  $\Delta v$ ,  $\Delta z$  for the locations corresponding to the control points with respect to the top-left corner of the cell that is closer to the camera,  $(u, v, z)$ . For the hand root joint and the object centroid, we constrain the offset to lie between 0 and 1, where a size of 1 corresponds to the full extent of an individual cell within grid dimensions. For other control points, we do not constrain the network’s output as those points should be allowed to fall outside the cell. The predicted location of the control point  $(\hat{w}_u, \hat{w}_v, \hat{w}_z)$  is then computed as:

$$\hat{w}_u = g(\Delta u) + u \quad (1)$$

$$\hat{w}_v = g(\Delta v) + v \quad (2)$$

$$\hat{w}_z = g(\Delta z) + z \quad (3)$$

where  $g(\cdot)$  is chosen to be a 1D sigmoid function for the root joint and the object centroid, and the identity function for other control points. Here,  $(u, v, z)$  are indices for the cell in grid dimensions. Given the camera intrinsics matrix  $K$ , and the prediction for the grid location,  $(\hat{w}_u, \hat{w}_v, \hat{w}_z)$ , the 3D location  $\hat{\mathbf{y}}$  of the control point in the camera coordinate system is then computed as:

$$\hat{\mathbf{y}} = \hat{w}_z \cdot C_z \cdot K^{-1}[\hat{w}_u \cdot C_u, \hat{w}_v \cdot C_v, 1]^T. \quad (4)$$

3D joint predictions already define the hand pose. Given the control point predictions on the 3D bounding box, 6D object pose could also be computed efficiently by aligning the pre-

diction to the reference 3D bounding box with a rigid transformation. This dispenses with the need to solve for a 2D-to-3D correspondence problem using PnP as in [46, 66] and recovers the 6D pose via Procrustes transformation [19]. Such a formulation also reduces depth ambiguities caused by projection from 3D to 2D. We show in Sec. 4 that this results in an improved object pose estimation accuracy.

In addition to hand and object control point locations, our network also predicts high confidence values for cells where the hand (or the object) is present, and low confidence where they are not present. Computing reliable confidence values is key for obtaining accurate predictions at test time. We define the confidence of a prediction as a function of the distance of the prediction to the ground truth, inspired by [66]. Namely, given a predicted 2D location  $(\hat{w}_u, \hat{w}_v)$  and its Euclidean distance  $D_T(\hat{w}_u, \hat{w}_v)$  from the ground truth, measured in image space, the confidence value  $c_{uv}(\hat{w}_u, \hat{w}_v)$  is computed as an exponential function with cut-off value  $d_{th}$  and sharpness parameter  $\alpha$ :

$$c_{uv}(\hat{w}_u, \hat{w}_v) = e^{\alpha \left(1 - \frac{D_T(\hat{w}_u, \hat{w}_v)}{d_{th}}\right)} \quad (5)$$

if  $D_T(\hat{w}_u, \hat{w}_v) < d_{th}$ , and  $c_{uv}(\hat{w}_u, \hat{w}_v) = 0$  otherwise. However, in contrast to [66] that computes a confidence value for 2D prediction, we need to also consider the depth dimension. Therefore, for a given depth prediction,  $\hat{w}_z$ , we compute an additional confidence value,  $c_z(\hat{w}_z)$ , measuring the distance (in metric space) between prediction and ground truth as in Eq. 5. We then compute the final confidence value  $c = 0.5 \cdot c_{uv}(\hat{w}_u, \hat{w}_v) + 0.5 \cdot c_z(\hat{w}_z)$ .

### 3.3. Object and Action Recognition

In addition to the 3D control locations and the confidence value, our model also predicts the object and action classes (*i.e. nouns* and *verbs*, as defined in Sec. 3.1). Intuitively, features learned to predict hand and object pose could also

help recognize actions and object classes. In order to predict actions, at each cell  $i$  we store into vector  $\mathbf{v}_i^h$ , together with hand pose  $\mathbf{y}_i^h$  and the corresponding confidence value  $c_i^h$ , the target action class probabilities  $\mathbf{p}_i^a$ . Similarly, to be able to predict object classes, we store into vector  $\mathbf{v}_i^o$ , together with object pose  $\mathbf{y}_i^o$  and the corresponding confidence value  $c_i^o$ , the target object class probabilities  $\mathbf{p}_i^o$ .

In total,  $\mathbf{v}_i^h$  stores  $3 \cdot N_c + N_a + 1$  values and  $\mathbf{v}_i^o$  stores  $3 \cdot N_c + N_o + 1$  values. We train our network to be able to predict  $\mathbf{v}_i^h$  and  $\mathbf{v}_i^o$  for each cell  $i$  of the 3D grid  $\mathbf{G}^t$ , for each  $t$ . Given *verb* and *noun* predictions, our model is able to recognize interactions from only a single image. Ultimately, our network learns to jointly predict 3D hand and object poses along with object, action and interaction classes – with a single forward pass.

### 3.4. Temporal Reasoning and Interaction Modeling

While simple actions can be recognized by looking at single frames, more complex activities require to model longer-term dependencies across sequential frames. To reason along the temporal dimension, we add a RNN module to our architecture. We choose to use a Long Short-Term Memory (LSTM) [23], given its popularity in the action recognition literature [17]. We also experimented with different modules (*e.g.* Gated Recurrent Units [6]), without observing substantial differences in our results.

A straightforward approach would be to consider the highest confidence predictions for hand and object poses at each frame, and give them as input to the RNN module, as in [17]. The RNN output is then processed by a softmax layer to predict the activity class. However, we can improve upon this baseline by explicitly reasoning about *interactions* between hands and objects. For example, in the context of visual reasoning, [1] proposes to model relational dependencies between objects in the scene and demonstrates that such object-level reasoning improves accuracy in targeted visual recognition tasks. In a similar manner, we aim to explicitly model interactions – in this case, however, of hands and objects, and directly in 3D. Co-training of 3D hand and object pose estimation networks already implicitly accounts for interactions in a data-driven manner. We further propose to model hand-object interactions at the structured output level with an *interaction RNN*. To do so, instead of directly feeding the hand and object poses as input to the temporal module, inspired by [1, 56], we model dependencies between hands and objects with a composite learned function and only then give the resulting mapping as input to RNN:

$$f_\phi(g_\theta(\hat{\mathbf{Y}}^h, \hat{\mathbf{Y}}^o)) \quad (6)$$

where  $f_\phi$  is an LSTM and  $g_\theta$  is an MLP, parameterized by  $\phi$  and  $\theta$ , respectively. Ultimately this mapping learns the explicit dependencies between hand and object poses and models interactions. We analyze and quantitatively evaluate the benefits of this approach in Sec. 4.

### 3.5. Training

The final layer of our single-shot network produces, for each cell  $i$ , hand and object pose predictions,  $\hat{\mathbf{y}}_i^h$  and  $\hat{\mathbf{y}}_i^o$ , with their overall confidence values,  $\hat{c}_i^h$  and  $\hat{c}_i^o$ , as well as probabilities for actions,  $\hat{\mathbf{p}}_i^a$ , and object classes,  $\hat{\mathbf{p}}_i^o$ . For each frame  $t$ , the loss function to train our network is defined as follows:

$$\mathcal{L} = \lambda_{pose} \sum_{i \in \mathbf{G}^t} (\|\hat{\mathbf{y}}_i^h - \mathbf{y}_i^h\|_2^2 + \|\hat{\mathbf{y}}_i^o - \mathbf{y}_i^o\|_2^2) + \quad (7)$$

$$\lambda_{conf} \sum_{i \in \mathbf{G}^t} ((\hat{c}_i^h - c_i^h)^2 + (\hat{c}_i^o - c_i^o)^2) - \quad (8)$$

$$\lambda_{actcls} \sum_{i \in \mathbf{G}^t} \hat{\mathbf{p}}_i^a \log \mathbf{p}_i^a - \quad (9)$$

$$\lambda_{objcls} \sum_{i \in \mathbf{G}^t} \hat{\mathbf{p}}_i^o \log \mathbf{p}_i^o. \quad (10)$$

Here, the regularization parameters for the pose and classification losses,  $\lambda_{pose}$ ,  $\lambda_{actcls}$  and  $\lambda_{objcls}$  are simply set to 1. As suggested by [48], for cells that contain a hand or an object, we set  $\lambda_{conf}$  to 5 and for cells that do not contain any of them, we set it to 0.1 to increase model stability. We feed the highest confidence predictions of our network to the recurrent module and define an additional loss based on cross entropy for the action recognition over the entire sequence. In principle, it is straightforward to merge and train the single-image and temporal models jointly with a softargmax operation [5, 71]. However, as backpropagation requires to keep all the activations in memory, which is not possible for a batch of image sequences, we found it effective to train our complete model in two stages. We first train on single frames to jointly predict 3D hand and object poses, object classes and action categories. We then keep the weights of the initial model fixed and train our recurrent network to propagate information in the temporal domain and model interactions. The complete model takes as input a sequence of images and outputs per-frame 3D hand-object pose predictions, object and action classes along with the estimates of interactions for the entire sequence.

## 4. Evaluation

In this section, we first describe the datasets and the corresponding evaluation protocols. We then compare our *Hand+Object* approach against the state-of-the-art methods and provide a detailed analysis of our general framework.

### 4.1. Datasets

We evaluate our framework for recognizing 3D hand-object poses and interactions on the recently introduced First-Person Hand Action (FPHA) dataset [17]. It is the only publicly available dataset for 3D hand-object interaction recognition that contains labels for 3D hand pose, 6D object pose and action categories. FPHA is a large and diverse dataset including 1175 videos belonging to 45 dif-

ferent activity categories performed by 6 actors. A total of 105,459 frames are annotated with accurate hand poses and action categories. The subjects carry out complex motions corresponding to daily human activities. A subset of the dataset contains annotations for objects’ 6-dimensional poses along with corresponding mesh models for 4 objects involving 10 different action categories. We denote this subset of the dataset as FPHA-HO.

As there are no other egocentric datasets containing labels for both 3D hand pose and 6D object pose, we further annotate 6D object poses on a part of the EgoDexter hand pose estimation dataset [37], on which we validate our joint hand-object pose estimation framework. To this end, we annotate the *Desk* sequence between frames 350 and 500 which features a cuboid object. We use training data from the SynthHands 3D hand pose estimation dataset [37]. We augment this dataset by randomly superimposing on the image synthetic cuboid objects that we generate and object segmentation masks from the third-person view dataset of [60] which features the same object. To gain robustness against background changes, we further replace the backgrounds using random images from [11].

## 4.2. Evaluation Metrics

We evaluate our unified framework on a diverse set of tasks: egocentric activity recognition, 3D hand pose estimation and 6D object pose estimation, and use standard metrics and official train/test splits to evaluate our performance in comparison to the state of the art. We use the percentage of correct video classifications, percentage of correct key-point estimates (3D PCK) and percentage of correct poses to measure accuracy on activity recognition, 3D hand pose estimation and 6D object pose estimation, respectively. When using the 3D PCK metric for hand pose estimation, we consider a pose estimate to be correct when the mean distance between the predicted and ground-truth joint positions is less than a certain threshold. When using the percentage of correct poses to evaluate 6D object pose estimation accuracy, we take a pose estimate to be correct if the 2D projection error or the average 3D distance of model vertices is less than a certain threshold (the latter being also referred to as the ADD metric).

## 4.3. Implementation Details

We initialize the parameters of our single-image network based on [48] with weights pretrained on ImageNet. The input to our model is a  $416 \times 416$  image. The output grid,  $\mathbf{G}^t$ , has the following dimensions:  $H = 13$ ,  $W = 13$  and  $D = 5$ . We set the grid size in image dimensions to  $C_u = C_v = 32$  pixels as in [48] and in depth dimension to  $C_z = 15$  cm. Further details about the architecture can be found in the supplemental material. We set the sharpness of the confidence function,  $\alpha$ , to 2, and the distance thresh-

old to 75 pixels for the spatial dimension and 75 mm for the depth dimension. We use a 2-layer LSTM with a hidden layer size of 512. The nonlinearity,  $g_\theta$ , is implemented as an MLP with 1 hidden layer with ReLU activation consisting of 512 units. We use stochastic gradient descent for optimization. We start with a learning rate of 0.0001 and divide the learning rate by 10 at the 80<sup>th</sup> and 160<sup>th</sup> epoch. All models are trained with a batch size of 16 for 200 epochs. We use extensive data augmentation to prevent overfitting. We randomly change the hue, saturation and exposure of the image by up to a factor of 50%. We also randomly translate the image by up to a factor of 10%. Our implementation is based on PyTorch.

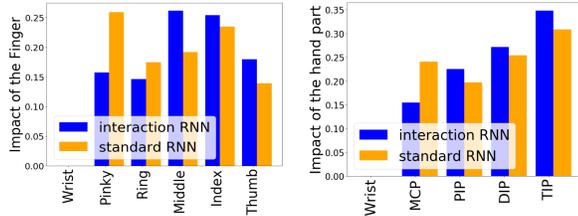
## 4.4. Experimental Results

We first report activity recognition accuracy on the FPHA-HO dataset and compare our results to the state-of-the-art results of [17] in Table 1. We further use the following baselines and versions of our approach in the evaluation:

- **SINGLE-IMAGE** - Our single pass network that predicts the action (*i.e. verb*), and object class (*i.e. noun*). The individual predictions for action ( $\mathbf{p}_i^a$ , *e.g.* open) and object ( $\mathbf{p}_i^o$ , *e.g.* bottle) class are combined to predict the interaction type (*e.g.* open bottle), *i.e.* (*verb, noun*) pair. This version of our model does not use a temporal model.
- **HAND POSE** - A temporal model that uses the hand pose predictions of our approach as input to the RNN to recognize activities.
- **OBJECT POSE** - similar to the previous baseline, but trained to predict activities based on the predicted keypoints on the 3D object bounding box.
- **HAND + OBJECT POSE** - A version of our model that combines 3D hand-object pose predictions to feed them as input to the RNN.
- **HAND POSE + OBJECT POSE + INTERACT** - A complete version of our model with temporal reasoning and interaction modeling.

Method	Model	Action Accuracy (%)
[17]	Ground-truth Hand Pose	87.45
	Ground-truth Object Pose	74.45
	Ground-truth Hand + Object Pose	91.97
OURS	SINGLE-IMAGE	85.56
	HAND POSE	89.47
	OBJECT POSE	85.71
	HAND + OBJECT POSE	94.73
	HAND + OBJECT POSE + INTERACT	<b>96.99</b>

**Table 1:** Action recognition results on FPHA-HO. We evaluate the impact of hand and object poses for action recognition. We demonstrate that hand-object predictions of our unified network yield more accurate results than [17] which relies on ground-truth pose as input. Furthermore, explicitly modeling interactions of hand-object poses results in a clear improvement in accuracy.

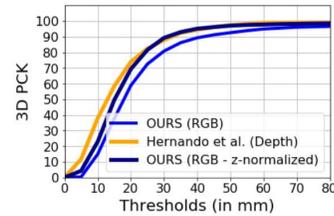


**Figure 3:** The impact of each finger (left) and each hand part (right) in recognizing interactions. The impact is measured by the normalized magnitude of neural network weights tied to the corresponding hand joint positions and compared for a standard RNN and an interaction RNN. MCP, PIP and DIP denote the 3 consecutive joints located inbetween wrist and fingertip (TIP) on each finger, in their respective order.

**Recognizing Interactions.** Our model achieves state-of-the-art performance for recognizing egocentric activities and 3D hand-object interactions even without ground-truth 3D hand and object poses. Our SINGLE-IMAGE baseline already achieves close results to the state of the art. We demonstrate that temporal reasoning on 3D hand and object poses individually improves activity recognition accuracy. The combination of hand and object poses further boosts the overall scores. With interaction modeling, performance improvements are even more noticeable. To analyze the importance of interaction modeling further, in Fig. 3, we quantify the importance of each input to the RNN by measuring the magnitude of the network weights tied to the inputs, both for a standard RNN and our interaction RNN. We demonstrate that, in contrast to a standard RNN, our temporal model attributes more importance to the index fingers and fingertips that interact more commonly with objects and learns which joints are relevant in interaction. In Fig. 5, we provide some visual results, 3D hand and object poses along with action and object labels that are produced by our HAND + OBJECT + INTERACT model.

We further evaluate the performance of our approach for the task of egocentric activity and interaction recognition on the full FPFA dataset. On the full dataset, object poses are not available for all the action categories, therefore, we train our models only using 3D hand poses. Here, however, to increase the descriptive power, we further leverage the object class and action category predictions produced by our single pass network in our temporal model. To this end we augment the 3D hand pose input (HP) with the output object class (OC) and action category (AC) probabilities. In Table 2, we demonstrate that these additional features yield improved action recognition accuracy. Overall, our method consistently outperforms the baselines by a large margin, including the ones that rely on depth data.

**3D Hand Pose Prediction.** We further compare the accuracy of our 3D hand pose predictions to the state-of-the-art results on FPFA in Fig. 4. Even though we only use *color* images, in contrast to [17] that uses *depth* input, we achieve



**Figure 4:** Comparison of the hand pose estimation results of our approach with those of Garcia-Hernando et al. [17] using different thresholds for the 3D pose error.

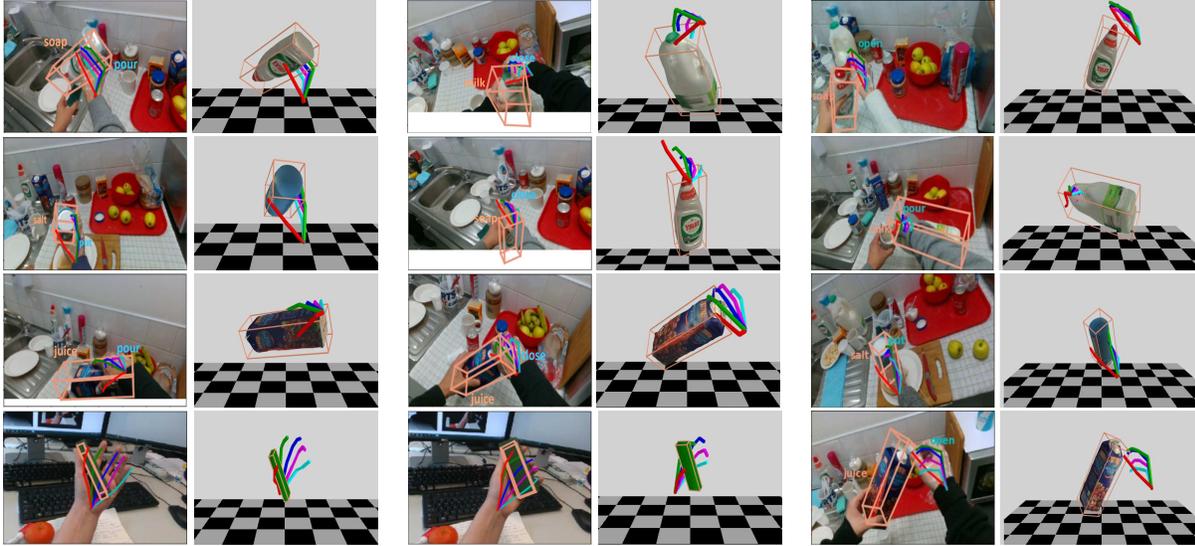
Model	Input modality	Accuracy
Two-stream-color [15]	Color	61.56
Two-stream-flow [15]	Color	69.91
Two-stream-all [15]	Color	75.30
Joule-color [24]	Color	66.78
HON4D [42]	Depth	70.61
Novel View [47]	Depth	69.21
Joule-depth [24]	Depth	60.17
[17] + Gram Matrix	Depth	32.22
[17] + Lie Group	Depth	69.22
[17] + LSTM	Depth	72.06
Ours - HP	Color	62.54
Ours - HP + AC	Color	74.20
Ours - HP + AC + OC	Color	<b>82.43</b>

**Table 2:** Action recognition results on the full FPFA dataset [17]. Our method significantly improves upon the baselines, including the ones that rely on depth data. We further provide action-specific recognition accuracies in the supplemental material.

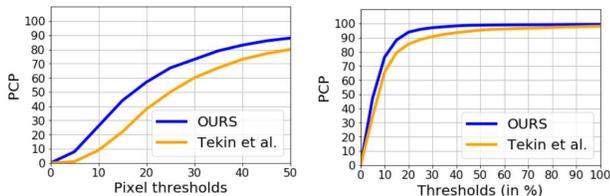
competitive 3D hand pose estimation accuracy. Furthermore, we do not assume knowledge of the hand bounding box as in [17]; our model takes as input a single full color image. Note also that the method of [17] is specifically trained for 3D hand pose estimation, whereas this is a sub-task of our method which simultaneously tackles multiple tasks within a unified architecture.

**6D Object Pose Prediction.** To evaluate our object pose accuracy, we compare our approach to the state-of-the-art results of [66] in Fig. 6. To this end, we run their approach on FPFA with their publicly available code. We demonstrate that explicitly reasoning about 6D object pose in 3D, in contrast to [66] that relies on solving 2D-to-3D correspondences, yields improved pose estimation accuracy. We conjecture that posing the 6D pose estimation problem in 2D is prone to depth ambiguities and our approach brings in robustness against it by directly reasoning in 3D.

**Unified Framework.** We analyze the influence of simultaneously training hand and object pose estimation tasks within the same single pass network on individual pose estimation accuracies. We compare the results of our *Hand + Object* network to those of the networks trained only for hand pose estimation and only for object pose estimation in Table 3. While we obtain similar accuracies for hand pose estimation with co-training and individual training, object pose estimation accuracy for the unified network is significantly better than that of the individual pose estimation net-



**Figure 5:** Qualitative results on the FPHA and EgoDexter dataset. We visualize the 3D hand pose estimates, 3D object bounding boxes which are transformed with the learned 6D object poses, and activity labels. The proposed approach can handle motion blur, self-occlusions, occlusions by viewpoints, clutter and complex articulations. We provide further qualitative results in our supplemental material.



**Figure 6:** Comparison of the object pose estimation results of our approach with those of Tekin et al. [66] using different thresholds for the 2D projection (left) and ADD metric (right).

work by a large margin of 9.65%. This indicates that having a joint representation which is shared across multiple tasks leads to an improvement in combined pose estimation accuracy. The results suggest that 3D hand pose highly constrains the 6D object pose, while the effect of rigid object pose on the articulated hand pose is not as pronounced. We further show that the prediction of the interaction class improves the hand and object pose estimation accuracy. This further validates that our unified framework allows us to achieve better overall performance in targeted tasks.

Network	HP error	OP error
HAND ONLY	16.15	-
OBJECT ONLY	-	28.27
HAND + OBJECT	16.87	25.54
HAND + OBJECT + INTERACT	<b>15.81</b>	<b>24.89</b>

**Table 3:** Comparison of the pose estimation results of our unified network to those of the networks trained only for hand and object pose estimation. Error metric is the mean 3D distance in mm.

**Generalization.** To demonstrate the generalization power of our joint hand-object pose estimation framework, we annotate a part of the EgoDexter hand pose estimation dataset [37] with 6D object poses, as explained in Sec. 4.1, and report quantitative results in Table 4. We demonstrate

Part	Error (in cm)
Fingertips	4.84
Object coordinates	2.37

**Table 4:** Results on the augmented EgoDexter dataset. Even when trained on synthetic data, our approach yields accurate poses.

that even when trained on synthetic data, our approach generalizes well to unconstrained environments and results in reliable and accurate joint 3D hand-object pose estimates. We provide visual results on this dataset in Fig. 5.

**Runtime.** Our single pass network that produces per-frame predictions simultaneously for 3D hand poses, 6D object poses, object classes and action categories runs at real-time speed of 25 fps on an NVIDIA Tesla M40. Without action and object recognition, when estimating only the poses of hands and objects, our method runs at a greater speed of 33 fps. Given hand and object poses, the interaction RNN module further processes a sequence with virtually no time overhead, at an average of 0.003 seconds.

## 5. Conclusion

In this paper, we propose the first method to jointly recognize 3D hand and object poses from monocular color image sequences. Our unified *Hand+Object* model simultaneously predicts per-frame 3D hand poses, 6D object poses, object classes and action categories, while being able to run at real-time speeds. Our framework jointly solves 3D articulated and rigid pose estimation problems within the same single-pass architecture and models the interactions between hands and objects in 3D to recognize actions from first-person views. Future work will apply the proposed framework to explicitly capture interactions between two hands and with other people in the scene.

## References

- [1] F. Baradel, N. Neverova, C. Wolf, J. Mille, and G. Mori. Object level Visual Reasoning in Videos. In *ECCV*, 2018.
- [2] G. Bertasius, H. S. Park, S. X. Yu, and J. Shi. First-Person Action-Object Detection with EgoNet. In *Robotics: Science and Systems*, 2017.
- [3] A. Bobick and J. Davis. The Recognition of Human Movement Using Temporal Templates. *PAMI*, 23(3):257–267, 2001.
- [4] E. Brachmann, F. Michel, A. Krull, M. Y. Yang, S. Gumhold, and C. Rother. Uncertainty-Driven 6D Pose Estimation of Objects and Scenes from a Single RGB Image. In *CVPR*, 2016.
- [5] O. Chapelle and M. Wu. Gradient Descent Optimization of Smoothed Information Retrieval Metrics. *Information Retrieval*, 13(3):216–235, 2009.
- [6] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In *EMNLP*, 2014.
- [7] C. Choi, S. H. Yoon, C. Chen, and K. Ramani. Robust Hand Pose Estimation during the Interaction with an Unknown Object. In *ICCV*, 2017.
- [8] V. Choutas, P. Weinzaepfel, J. Revaud, and C. Schmid. PoTion: Pose MoTion Representation for Action Recognition. In *CVPR*, 2018.
- [9] D. Damen, H. Doughy, G. M. Farinella, S. Fidler, A. Furnari, E. K. D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray. Scaling Egocentric Vision: The EPIC-KITCHENS Dataset. In *ECCV*, 2018.
- [10] P. Dollar, V. Rabaud, G. Cottrell, and S. Belongie. Behavior Recognition via Sparse Spatio-temporal Features. In *VS-PETS*, 2005.
- [11] M. Everingham, L. V. Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes (VOC) Challenge. *IJCV*, 88(2):303–338, 2010.
- [12] A. Fathi, A. Farhadi, and J. Rehg. Understanding Egocentric Activities. In *ICCV*, 2011.
- [13] A. Fathi, Y. Li, and J. Rehg. Learning to Recognize Daily Actions Using Gaze. In *ECCV*, 2012.
- [14] A. Fathi, X. Ren, and J. Rehg. Learning to Recognize Objects in Egocentric Activities. In *CVPR*, 2011.
- [15] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional Two-Stream Network Fusion for Video Action Recognition. In *CVPR*, 2016.
- [16] D. Fouhey, W. Kuo, A. Efros, and J. Malik. From Lifestyle VLOGs to Everyday Interaction. In *CVPR*, 2018.
- [17] G. Garcia-Hernando, S. Yuan, S. Baek, and T.-K. Kim. First-Person Hand Action Benchmark with RGB-D Videos and 3D Hand Pose Annotations. In *CVPR*, 2018.
- [18] G. Gkioxari, R. Girshick, P. Dollar, and K. He. Detecting and Recognizing Human-Object Interactions. In *CVPR*, 2018.
- [19] J. C. Gower. Generalized Procrustes Analysis. *Psychometrika*, 40(1):33–51, 1975.
- [20] R. Goyal, S. E. Kahou, V. Michalski, J. Materzynska, S. Westphal, H. Kim, V. Haenel, I. Freund, P. Yianilos, M. Mueller-Freitag, F. Hoppe, C. Thureau, I. Bax, and R. Memisevic. The “Something Something” Video Database for Learning and Evaluating Visual Common Sense. In *ICCV*, 2017.
- [21] H. Hamer, K. Schindler, E. Koller-Meier, and L. V. Gool. Tracking a Hand Manipulating an Object. In *ICCV*, 2009.
- [22] Y. Hasson, G. Varol, D. Tzionas, I. Kalevatykh, M. J. Black, I. Laptev, and C. Schmid. Learning Joint Reconstruction of Hands and Manipulated Objects. In *CVPR*, 2019.
- [23] S. Hochreiter and J. Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [24] J.-F. Hu, W.-S. Zheng, J. Lai, and J. Zhang. Jointly Learning Heterogeneous Features for RGB-D Activity Recognition. In *CVPR*, 2015.
- [25] U. Iqbal, P. Molchanov, T. Breuel, J. Gall, and J. Kautz. Hand Pose Estimation via Latent 2.5D Heatmap Regression. In *ECCV*, 2018.
- [26] H. Jhuang, T. Serre, L. Wolf, and T. Poggio. A Biologically Inspired System for Action Recognition. In *ICCV*, 2007.
- [27] V. Kalogeiton, P. Weinzaepfel, V. Ferrari, and C. Schmid. Joint Learning of Object and Action Detectors. In *ICCV*, 2017.
- [28] W. Kehl, F. Manhardt, F. Tombari, S. Ilic, and N. Navab. SSD-6D: Making RGB-based 3D Detection and 6D Pose Estimation Great Again. In *ICCV*, 2017.
- [29] K. Kitani, T. Okabe, Y. Sato, and A. Sugimoto. Fast Unsupervised Ego-action Learning for First-person Sports Videos. In *CVPR*, 2011.
- [30] H. Koppula, R. Gupta, and A. Saxena. Learning Human Activities and Object Affordances from RGB-D Videos. *IJRR*, 32(8):951–970, 2013.
- [31] I. Laptev and T. Lindeberg. Space-time Interest Points. In *ICCV*, 2003.
- [32] V. Lepetit, F. Moreno-Noguer, and P. Fua. EPnP: An Accurate O(n) Solution to the PnP Problem. *IJCV*, 81(2):155–166, 2009.
- [33] Y. Li, Z. Ye, and J. Rehg. Delving into Egocentric Actions. In *CVPR*, 2015.
- [34] Y. S. M. Cai, K. M. Kitani. Understanding Hand-Object Manipulation with Grasp Types and Object Attributes. In *Robotics: Science and Systems*, 2016.
- [35] M. Ma, H. Fan, and K. Kitani. Going Deeper into First-Person Activity Recognition. In *CVPR*, 2016.
- [36] F. Mueller, F. Bernard, O. Sotnychenko, D. Mehta, S. Sridhar, D. Casas, and C. Theobalt. GANerated Hands for Real-Time 3D Hand Tracking from Monocular RGB. In *CVPR*, 2018.
- [37] F. Mueller, D. Mehta, O. Sotnychenko, S. Sridhar, D. Casas, and C. Theobalt. Real-time Hand Tracking under Occlusion from an Egocentric RGB-D Sensor. In *ICCV*, 2017.
- [38] J. C. Nibbles, H. Wang, and L. Fei-Fei. Unsupervised Learning of Human Action Categories Using Spatio-temporal Words. In *BMVC*, 2006.
- [39] M. Oberweger, G. Riegler, P. Wohlhart, and V. Lepetit. Efficiently Creating 3D Training Data for Fine Hand Pose Estimation. In *CVPR*, 2016.

- [40] M. Oberweger, P. Wohlhart, and V. Lepetit. Training a Feedback Loop for Hand Pose Estimation. In *ICCV*, 2015.
- [41] I. Oikonomidis, N. Kyriazis, and A. Argyros. Full DOF Tracking of a Hand Interacting with an Object by Modeling Occlusions and Physical Constraints. In *ICCV*, 2011.
- [42] O. Oreifej and Z. Liu. HON4D: Histogram of Oriented 4D Normals for Activity Recognition from Depth Sequences. In *CVPR*, 2013.
- [43] P. Panteleris, N. Kyriazis, and A. Argyros. 3D Tracking of Human Hands in Interaction with Unknown Objects. In *BMVC*, 2015.
- [44] H. Pirsiavash and D. Ramanan. Detecting Activities of Daily Living in First-Person Camera Views. In *CVPR*, 2012.
- [45] Y. Poleg, E. Phrat, S. Peleg, and C. Arora. Compact CNN for Indexing Egocentric Videos. In *WACV*, 2016.
- [46] M. Rad and V. Lepetit. BB8: A Scalable, Accurate, Robust to Partial Occlusion Method for Predicting the 3D Poses of Challenging Objects without Using Depth. In *ICCV*, 2017.
- [47] H. Rahmani and A. Mian. 3D Action Recognition from Novel Viewpoints. In *CVPR*, 2016.
- [48] J. Redmon and A. Farhadi. YOLO9000: Better, Faster, Stronger. In *CVPR*, 2017.
- [49] X. Ren and C. Gu. Figure-ground Segmentation Improves Handled Object Recognition in Egocentric Video. In *CVPR*, 2010.
- [50] X. Ren and C. Gu. First-Person Activity Recognition: What Are They Doing to Me? In *CVPR*, 2013.
- [51] G. Rogez, J. Supancic, and D. Ramanan. First-Person Pose Recognition Using Egocentric Workspaces. In *CVPR*, 2015.
- [52] G. Rogez, J. Supancic, and D. Ramanan. Understanding Everyday Hands in Action from RGB-D Images. In *ICCV*, 2015.
- [53] J. Romero, H. Kjellström, and D. Kragic. Hands in Action: Real-Time 3D Reconstruction of Hands in Interaction with Objects. In *ICRA*, 2010.
- [54] J. Romero, D. Tzionas, and M. J. Black. Embodied Hands: Modeling and Capturing Hands and Bodies Together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6), 2017.
- [55] M. Ryoo, B. Rothrock, and L. Matthies. Pooled Motion Features for First-Person Videos. In *CVPR*, 2015.
- [56] A. Santoro, D. Raposo, D. G. Barrett, M. Malinowski, R. Pascanu, P. Battaglia, and T. Lillicrap. A Simple Neural Network Model for Relational Reasoning. In *NIPS*, 2017.
- [57] G. Sigurdsson, A. Gupta, C. Schmid, A. Farhadi, and K. Alahari. Actor and Observer: Joint Modeling of First and Third-Person Videos. In *CVPR*, 2018.
- [58] A. Spurr, J. Song, S. Park, and O. Hilliges. Cross-modal Deep Variational Hand Pose Estimation. In *CVPR*, 2018.
- [59] S. Sridhar, F. Mueller, A. Oulasvirta, and C. Theobalt. Fast and Robust Hand Tracking Using Detection-Guided Optimization. In *CVPR*, 2015.
- [60] S. Sridhar, F. Mueller, M. Zollhoefer, D. Casas, A. Oulasvirta, and C. Theobalt. Real-time Joint Tracking of a Hand Manipulating an Object from RGB-D Input. In *ECCV*, 2016.
- [61] S. Sundaram and W. Mayol-Cuevas. High Level Action Recognition Using Low Resolution Wearable Vision. In *CVPR Workshops*, 2009.
- [62] D. Surie, T. Pederson, F. Lagriffoul, L. Janlert, and D. Sjölie. Activity Recognition Using an Egocentric Perspective of Everyday Objects. In *UIC*, 2007.
- [63] Y. Tang, Y. Tian, J. Lu, J. Feng, and J. Zhou. Action Recognition in RGB-D Egocentric Videos. In *ICIP*, 2017.
- [64] A. Tejani, D. Tang, R. Kouskouridas, and T.-K. Kim. Latent-Class Hough Forests for 3D Object Detection and Pose Estimation. In *ECCV*, 2014.
- [65] B. Tekin, P. Marquez-Neila, M. Salzmann, and P. Fua. Learning to Fuse 2D and 3D Image Cues for Monocular Body Pose Estimation. In *ICCV*, 2017.
- [66] B. Tekin, S. N. Sinha, and P. Fua. Real-Time Seamless Single Shot 6D Object Pose Prediction. In *CVPR*, 2018.
- [67] A. Tsoli and A. Argyros. Joint 3D Tracking of a Deformable Object in Interaction with a Hand. In *ECCV*, 2018.
- [68] D. Tzionas, L. Ballan, A. Srikantha, P. Aponte, M. Pollefeys, and J. Gall. Capturing Hands in Action Using Discriminative Salient Points and Physics Simulation. *IJCV*, 118(2):172–193, 2016.
- [69] S.-F. Wong, T. K. Kim, and R. Cipolla. Learning Motion Categories Using Both Semantic and Structural Information. In *CVPR*, 2007.
- [70] Q. Ye and T.-K. Kim. Occlusion-aware Hand Pose Estimation Using Hierarchical Mixture Density Network. In *ECCV*, 2018.
- [71] K. M. Yi, E. Trulls, V. Lepetit, and P. Fua. LIFT: Learned Invariant Feature Transform. In *ECCV*, 2016.
- [72] S. Yuan, G. Garcia-Hernando, B. Stenger, G. Moon, J. Y. Chang, K. M. Lee, P. Molchanov, J. Kautz, S. Honari, L. Ge, J. Yuan, X. Chen, G. Wang, F. Yang, K. Akiyama, Y. Wu, Q. Wan, M. Madadi, S. Escalera, S. Li, D. Lee, I. Oikonomidis, A. Argyros, and T.-K. Kim. Depth-Based 3D Hand Pose Estimation: From Current Achievements to Future Goals. In *CVPR*, 2018.
- [73] S. Yuan, Q. Ye, B. Stenger, S. Jain, and T.-K. Kim. Big-Hand2.2M Benchmark: Hand Pose Dataset and State of the Art Analysis. In *CVPR*, 2017.
- [74] C. Zimmermann and T. Brox. Learning to Estimate 3D Hand Pose from Single RGB Images. In *ICCV*, 2017.