

Detailed 2D-3D Joint Representation for Human-Object Interaction

Yong-Lu Li, Xinpeng Liu, Han Lu, Shiyi Wang, Junqi Liu, Jiefeng Li, Cewu Lu*

Shanghai Jiao Tong University

yonglu.li@sjtu.edu.cn, xinpengliu0907@gmail.com,

{sjtu_luhan, shiywang, lj435, ljf_likit, lucewu}@sjtu.edu.cn

Abstract

Human-Object Interaction (HOI) detection lies at the core of action understanding. Besides 2D information such as human/object appearance and locations, 3D pose is also usually utilized in HOI learning since its view-independence. However, rough 3D body joints just carry sparse body information and are not sufficient to understand complex interactions. Thus, we need detailed 3D body shape to go further. Meanwhile, the interacted object in 3D is also not fully studied in HOI learning. In light of these, we propose a detailed 2D-3D joint representation learning method. First, we utilize the single-view human body capture method to obtain detailed 3D body, face and hand shapes. Next, we estimate the 3D object location and size with reference to the 2D human-object spatial configuration and object category priors. Finally, a joint learning framework and cross-modal consistency tasks are proposed to learn the joint HOI representation. To better evaluate the 2D ambiguity processing capacity of models, we propose a new benchmark named Ambiguous-HOI consisting of hard ambiguous images. Extensive experiments in large-scale HOI benchmark and Ambiguous-HOI show impressive effectiveness of our method. Code and data are available at <https://github.com/DirtyHarryLYL/DJ-RN>.

1. Introduction

Human-Object Interaction (HOI) detection recently receives lots of attentions. It aims at locating the active human-object and inferring the action simultaneously. As a sub-task of visual relationship [37], it can facilitate activity understanding [6, 43, 44, 55], imitation learning [2], etc.

What do we need to understand HOI? The possible answers are human/object appearance, spatial configuration, context, pose, etc. Among them, human body information

*Cewu Lu is the corresponding author, member of Qing Yuan Research Institute and MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, China.

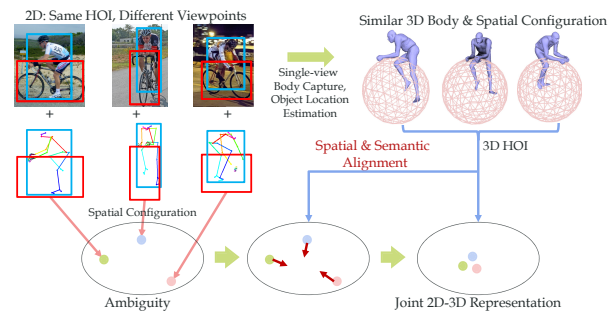


Figure 1. HOI detection based on 2D may have ambiguities under various viewpoints. HOI representation in 3D is more robust. Thus, we estimate the 3D detailed human body and interacted object location and size to represent the HOI in 3D. Then we learn a joint 2D-3D representation to combine multi-modal advantages.

often plays an important role, such as 2D pose [31, 21, 58, 13] and 3D pose [63, 39]. Because of the various viewpoints, 2D human pose [7] or segmentation [22, 61, 16] often has ambiguities, e.g. same actions may have very different 2D appearances and poses. Although 3D pose is more robust, rough 3D body joints are not enough to encode essential geometric and meaningful patterns. For example, we may need detailed hand shape to infer the action “use a knife to cut”, or facial shape for “eat and talk”. And body shape would also largely affect human posture. In light of this, we argue that *detailed 3D body* can facilitate the HOI learning. Meanwhile, the object in HOI is also important, e.g. “hold an apple” and “hold the horse” have entirely different patterns. However, few studies considered how to embed 3D interacted objects in HOI. The reasons are two-fold: first, it is hard to reconstruct objects because of the 6D pose estimation and diverse object shapes (detailed point cloud or mesh [8, 64]). Second, estimating 3D human-object spatial relationship is also difficult for single-view.

In this work, we propose a method to not only borrow essential discriminated clues from the detailed 3D body but also consider the 3D human-object spatial configuration. First, we represent the HOI in 3D. For human, we utilize the single-view human body capture [45] to obtain detailed human shape. For object, referring to the 2D human-object

spatial configuration and object category prior, we estimate its rough *location* and *size* through perspective projection, and use a hollow sphere to represent it. Then, we put the 3D detailed human and object sphere into a normalized volume as the **3D HOI spatial configuration volume**, as shown in Fig. 1. Next, we propose Detailed Joint Representation Network (DJ-RN), which consists of two feature extractors: a 2D Representation Network (2D-RN) and a 3D Representation Network (3D-RN). Finally, we adopt several consistency tasks to learn the 2D-3D joint HOI representation. In detail, we align the 2D spatial features according to more robust 3D spatial features. And we perform semantic alignment to ensure the cross-modal semantic consistency. To better embed the body posture, we estimate body part attentions in a 2D-3D joint way with consistency. That is if 2D features tell us the hands and head are important for “work on laptop”, so will the 3D features. DJ-RN is the *first* joint learning method to utilize single-view 3D recover for **HOI**. It is a novel **paradigm** instead of an ad-hoc model, and flexible to replace 2D/3D modules/extracted features. We believe it would promote not only HOI learning but also action related tasks, *e.g.*, image caption, visual reasoning.

To better evaluate the ability of processing 2D ambiguities, we propose a new benchmark named **Ambiguous-HOI**, which includes ambiguous examples selected from existing datasets like HICO-DET [9], V-COCO [20], Open-Image [28], HCVRD [65]. We conduct extensive experiments on widely-used HOI detection benchmark and Ambiguous-HOI. Our approach achieves significant improvements with 2D-3D joint learning. The main contributions are as follows: 1) We propose a 2D-3D joint representation learning paradigm to facilitate HOI detection. 2) A new benchmark Ambiguous-HOI is proposed to evaluate the disambiguation ability of models. 3) We achieve state-of-the-art results on HICO-DET [9] and Ambiguous-HOI.

2. Related Work

Human-Object Interaction Detection. Recently, great progress has been made in HOI detection. Large-scale datasets [9, 20, 28] have been released to promote this field. Meanwhile, lots of deep learning based methods [19, 17, 31, 51, 21, 58, 48] have been proposed. Chao *et al.* [9] proposed a multi-stream framework, which is proven effective and followed by subsequent works [17, 31]. Differently, GPNN [51] proposed a graph model and used message passing to address both image and video HOI detection. Gkioxari *et al.* [19] adopted an action density map to estimate the 2D location of interacted objects. iCAN [17] utilized self-attention to correlate the human-object and context. TIN [31] proposed an explicit interactiveness learning network to identify the non-interactive human-object pairs and suppress them in inference. HAKE [30] proposes a novel hierarchical paradigm based on human body

part states [38]. Previous methods mainly relied on the visual appearance and human-object spatial relative locations, some of them [31] also utilized the 2D estimated pose. But the 2D ambiguity in HOI is not well studied before.

3D Pose-based Action Recognition. Recent deep learning based 3D pose estimation methods [26, 14, 46] have achieved substantial progresses. Besides 2D pose based action understanding [31, 29, 35, 10, 18, 59], many works also utilized the 3D human pose [54, 13, 62, 63, 39, 49, 27, 33, 3, 34, 40, 57, 24]. Yao *et al.* [62] constructed a 2.5D graph with 2D appearance and 3D human pose, and selected exemplar graphs of different actions for the exemplar-based action classification. In [63], 2D pose is mapped to 3D pose and the actions are classified by comparing the 3D pose similarity. Luvizon *et al.* [39] estimated the 2D/3D pose and recognized actions in a unified model from both image and video. Wang *et al.* [60] used the RGB-D data to obtain the 3D human joints and adopted an actionlet ensemble method for HOI learning. Recently, Pham *et al.* [49] proposed a multi-task model to operate 3D pose estimation and action recognition simultaneously from RGB video. Most 3D pose based methods [3, 34, 40, 57, 33, 49, 60, 13, 27, 24] are using Recurrent Neural Network (RNN) based framework for spatio-temporal action recognition, but few studies focus on the complex HOI understanding from single RGB image.

Single-view 3D Body Recover. Recently the single-view human body capture and reconstruction methods [45, 26, 42, 47, 4] have made great progresses. With the help of deep learning and large-scale scanned 3D human database [11, 25, 1], they are able to directly recover 3D body shape and pose from single RGB images. SMPLify-X [45] is a holistic and efficient model that takes the 2D human body, face and hand poses as inputs to capture 3D body, face and hands. To obtain more accurate and realistic body shape, SMPLify-X [45] utilizes the Variational Human Body Pose prior (VPoser) trained on large-scale MoCap datasets, which carries lots of human body pose prior and knowledge. It supports us to recover 3D detailed human body from HOI images and embed more body posture knowledge.

3. Representing HOI in 3D

Our goal is to learn the 2D-3D joint HOI representation, thus we need to first represent HOI in 3D. Given a still image, we use object detection [52] and pose estimation [7] to obtain 2D instance boxes and human pose. Next, we adopt the 3D human body capture [45] to estimate the 3D human body with above 2D detection (Sec. 3.1), and estimate the object location and size in 3D to construct the 3D spatial configuration volume (Sec. 3.2).

3.1. Single-view 3D Body Capture

Rough 3D pose is not sufficient to discriminate various actions, especially the complex interactions with daily ob-

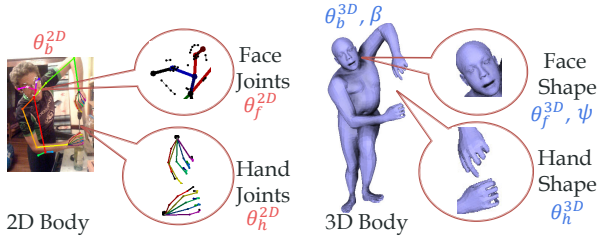


Figure 2. We adopt detailed human information. We use OpenPose [7] and SMPLify-X [45] to estimate the 2D, 3D poses and shapes of face and hands. These information would largely help the HOI learning, especially on actions related to face and hands.

jects. Thus we need holistic and fine-grained 3D body information as a clue. To this end, we adopt a holistic 3D body capture method [45] to recover detailed 3D body from single RGB images. Given the 2D detection of image I , *i.e.*, 2D human and object boxes b_h and b_o , 2D human pose $\theta^{2D} = \{\theta_b^{2D}, \theta_f^{2D}, \theta_h^{2D}\}$ (main body joints θ_b^{2D} , jaw joints θ_f^{2D} and finger joints θ_h^{2D} in Fig. 2). We input them into SMPLify-X [45] to recover 3D human estimations, *i.e.*, fitting the SMPL-X [45] model to I and θ^{2D} . Then we can obtain the optimized shape parameters $\{\theta^{3D}, \beta, \psi\}$ by minimizing the body pose, shape objective function, where θ^{3D} are pose parameters and $\theta^{3D} = \{\theta_b^{3D}, \theta_f^{3D}, \theta_h^{3D}\}$, β are body, face and hands shape parameters, ψ are facial expression parameters. The template body mesh is finally blended and deformed to fit the target body posture and shape in images. With function $M(\theta^{3D}, \beta, \psi) : \mathbb{R}^{|\theta^{3D}| \times |\beta| \times |\psi|} \rightarrow \mathbb{R}^{3N}$, we can directly generate the 3D body mesh according to the estimated $\{\theta^{3D}, \beta, \psi\}$ from images and utilize it in the next stage, some examples are shown in Fig. 2.

3.2. 3D Spatial Configuration Volume

After obtaining the 3D body, we further represent HOI in 3D, *i.e.* estimate the 3D object location and size. For robustness and efficiency, we do not reconstruct the object shape, but use a *hollow sphere* to represent it. Thus we can avoid the difficult 6D pose estimation under the circumstance of single-view and various categories. Our procedure has *two stages*: 1) locating the sphere center on a plane according to the camera perspective projection, 2) using the prior object size and human-object distance to estimate the depth of the sphere. For each image, we adopt the estimated camera parameters from SMPLify-X [45], where focal length f is set to a fixed value of 5000, and the camera distortions are not considered. For clarification, the camera optical center is expressed as $C(t_1, t_2, t_3)$ in the world coordinate system, and the object sphere center is $O(x_o, y_o, z_o)$.

Object Sphere Center. As shown in Fig. 3, we assume that O is projected to the midperpendicular of object box border-top, indicating the sphere center falls on plane P_{ABC} . And we suppose the highest and lowest visible points of the

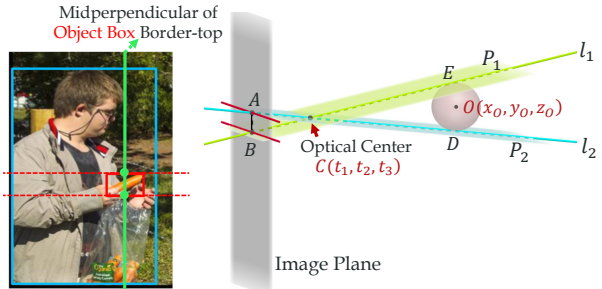


Figure 3. Object location estimation. Given prior radius r , we can get the sphere center location by solving projection equations, which restricts the sphere to be tangent to plane P_1 and P_2 , and assures the sphere center falls on plane P_{ABC} .

sphere are projected to the object box border-bottom and border-top respectively. Then we can get two tangent planes of the sphere: P_1 (contains points B, C, E) and P_2 (contains points A, C, D) as shown in Fig. 3. P_1 and P_2 intersect with P_{ABC} , restricting a region on P_{ABC} where the sphere center may locate. To get the depth of sphere center, we need to know the sphere radius, *i.e.*, $r = |\overrightarrow{OD}| = |\overrightarrow{OE}|$.

Object Sphere Radius. As for the sphere radius, we determine it by both considering the object box relative size (to the 2D human box) and the object category prior. With object detection in the first step, we can obtain the object category j . Thus, we can set a rough object size according to Wikipedia and daily life experience. In practice, we set prior sizes for COCO 80 objects [32] to suit the HICO-DET setting [9]. First, for small objects or objects with similar size along different axes (*e.g.* ball, table), we define the *prior object scale ratio* between the sphere radius and the human shoulder width. Second, for objects that are usually partly seen or whose projection is seriously affected by the 6D pose (*e.g.* boat, skis), we use the relative scale ratio of the human and object boxes as the referenced ratio. The estimated sphere center is denoted as $\hat{O}(\hat{x}_o, \hat{y}_o, \hat{z}_o)$.

The sphere depth is very sensitive to the radius and may make the sphere away from human. Thus, we regularize the estimated depth \hat{z}_c using the maximum and minimum depth z_H^{max}, z_H^{min} of the recovered human. We define prior *object depth regularization factor* $\Gamma = \{[\gamma_j^{min}, \gamma_j^{max}]\}_{j=1}^{80}$ for COCO objects [32]. Specifically, with pre-defined depth bins (very close, close, medium, far, very far), we invite fifty volunteers from different backgrounds to watch HOI images and choose the degree of the object relative depth to the human. We then use their votes to set the empirical regularization factors Γ . For estimated $\hat{O}(\hat{x}_o, \hat{y}_o, \hat{z}_o)$, if \hat{z}_o falls out of $[\gamma_j^{min} z_H^{min}, \gamma_j^{max} z_H^{max}]$, we shift \hat{O} to $(\hat{x}_o, \hat{y}_o, \gamma_j^{max} z_H^{max})$ or $(\hat{x}_o, \hat{y}_o, \gamma_j^{min} z_H^{min})$, depending on which is closer to \hat{O} . Size and depth priors can effectively restrict the error boundaries. Without them, 3D volume would have large deviation and degrade performance.

Volume Formalization. Next, we perform translations to align different configurations in 3D. First, we set the coordinate origin as the human pelvis. The direction of gravity estimated is kept same with the negative direction of the z-axis, and the line between two human shoulder joints is rotated to be parallel to the x-axis. Then, we down-sample the 3D body to 916 points and randomly sample 312 points on *spherical surface*. The hollow sphere can keep the body information of the interacted body parts within the sphere. We then normalize the whole volume by setting unit length as the distance between the pupil joints. At last, we can obtain a normalized 3D volume including 3D body and object sphere, which not only carries essential 3D action information but also 3D human-object spatial configuration.

4. 2D-3D Joint Learning

In this section, we aim to learn the joint representation. To this end, we propose Detailed Joint Representation Network (DJ-RN), as seen in Fig. 5. DJ-RN has two modules: 2D Representation Network (2D-RN) and 3D Representation Network (3D-RN). We use them to extract features from two modalities respectively (Sec. 4.1, 4.2). Then we align 2D spatial feature with 3D spatial feature (Sec. 4.3), and use body part attention consistency (Sec. 4.4) and semantic consistency (Sec. 4.5) to guide the learning.

4.1. 2D Feature Extraction

2D-RN is composed of human, object, and spatial streams following [9, 17, 31]. Within each stream, we adopt different blocks to take in 2D information with different properties and extract corresponding features (Fig. 5).

Human/Object Block. Human and object streams mainly utilize visual appearance. We use a COCO [32] pre-trained Faster-RCNN [52] to extract ROI pooling features from detected boxes. To enhance the representation ability, we adopt the iCAN block [17] which computes the self-attention via correlating the context and instances, and obtain the human feature f_H^{2D} and object feature f_O^{2D} .

Spatial Block. Although appearance carries important clues, it also imports noise and misleading patterns from various viewpoints. Thus human-object spatial configuration can be used additionally to provide discriminative features [9, 17, 31]. Spatial stream mainly considers the 2D human-object relative locations. We input the 2D pose map and spatial map [31] to the spatial block, which consists of convolution and fully-connected (FC) layers to extract the spatial feature f_{sp}^{2D} . The spatial map consists of two channels, human and object maps, which are all 64×64 and generated from the human and object boxes. The value is 1 in the box and 0 elsewhere. The pose map consists of 17 joint heatmaps of size 64×64 from OpenPose [7].

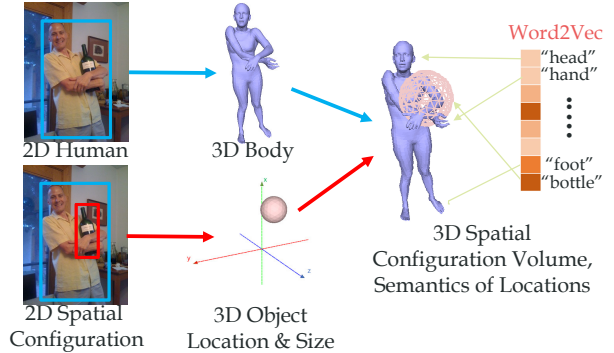


Figure 4. 3D spatial configuration volume. After 3D body capture, we use 2D boxes, estimated camera parameters and object category prior to estimate the 3D object location and size, and then put 3D human and object together in a normalized volume. We also pair the 3D location with semantic knowledge (Sec. 4.2).

4.2. 3D Feature Extraction

3D-RN contains a 3D spatial stream with volume block which takes in the 3D spatial configuration volume, and a 3D human stream with 3D body block to encode 3D body.

Volume Block. In 3D spatial stream, we adopt PointNet [50] to extract 3D spatial feature f_{sp}^{3D} . We first pre-train it to segment the human and object points in the generated 3D spatial configuration volume. Thus it can learn to discriminate the geometric difference and shape of human and object. Then we use it to extract features from 3D spatial volume point cloud. To further embed the **semantic** information of 3D **locations**, we *pair* the spatial feature with the corresponding semantics, *i.e.*, the word embedding of object or body part category. We first divide the volume point cloud into 18 sets: 17 part sets and an object sphere set. Then, for the feature of part set, we concatenate it with PCA reduced word embedding [41] of part name (*e.g.* “hand”). Similarly, for the feature of the sphere set, we concatenate it with the object category word embedding (*e.g.* “bottle”), as seen in Fig. 4. The concatenated feature is used as f_{sp}^{3D} .

3D Body Block. In 3D body block, we extract features based on SMPL-X [45] parameters: joint body, face and hands shape β , face expression ψ and pose θ^{3D} , consisting of jaw joints θ_f^{3D} , finger joints θ_h^{3D} and body joints θ_b^{3D} . For body shape and expression, we directly use their parameters. For pose, we adopt the VPoser [45] to encode the 3D body into latent representations $\{f_b^{3D}, f_f^{3D}, f_h^{3D}\}$ for body, face and hands corresponding to $\{\theta_b^{3D}, \theta_f^{3D}, \theta_h^{3D}\}$. VPoser is a variational auto-encoder trained with large-scale MoCap datasets [11, 25, 1]. Thus it learns a latent space encoding the manifold of the physically plausible pose, and effectively embeds the 3D body pose. We concatenate the latent representations, shape parameters and face expression, feed them to two 1024 sized FC layers, and get the 3D human feature $f_H^{3D} = FC_{3D}(\{\beta, \psi, f_b^{3D}, f_f^{3D}, f_h^{3D}\})$.

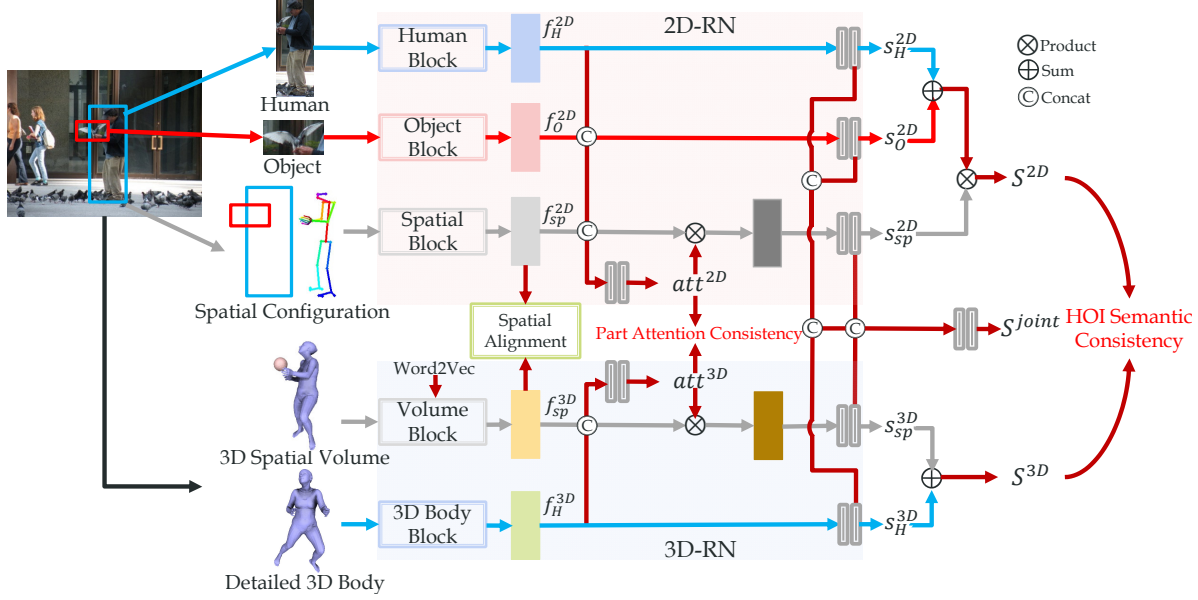


Figure 5. Overview of DJ-RN. The framework consists of two main modules, named 2D Representation Network (2D-RN) and 3D Representation Network (3D-RN). They extract HOI representations from 2D and 3D information respectively. Hence, we can use spatial alignment, part attention consistency and semantic consistency to learn a joint 2D-3D representation for HOI learning.

4.3. 2D-3D Spatial Alignment

In view of that 2D spatial features lack robustness and may bring in ambiguities, we propose the 2D spatial alignment. 3D spatial features are more robust, thus we refer them as anchors in the spatial space which describes the manifold of HOI spatial configuration. Given the 2D spatial feature f_{sp}^{2D} of a sample, from the train set we randomly sample a positive 3D spatial feature f_{sp+}^{3D} with the same HOI label and a negative feature f_{sp-}^{3D} with non-overlapping HOIs (a person may perform multiple actions at the same time). For a human-object pair, we use triplet loss [53] to align its 2D spatial feature, *i.e.*,

$$\mathcal{L}_{tri} = [d(f_{sp}^{2D}, f_{sp+}^{3D}) - d(f_{sp}^{2D}, f_{sp-}^{3D}) + \alpha]_+ \quad (1)$$

where $d(\cdot)$ indicates the Euclidean distance, and $\alpha = 0.5$ is the margin value. For 2D samples with the same HOIs but different 2D spatial configurations, this spatial alignment will gather them together in the spatial space.

4.4. Joint Body Part Attention Estimation

Body parts are important in HOI understanding, but not all parts make great contributions in inference. Thus, adopting attention mechanism is apparently a good choice. Different from previous methods [12, 15], we generate body part attention by considering both 2D and 3D clues. Specifically, we use a part attention consistency loss to conduct self-attention learning, as shown in Fig. 6. With the 2D and 3D features, we can generate two sets of body attention.

2D Attention. We concatenate the input $f_H^{2D}, f_O^{2D}, f_{sp}^{2D}$ to get f_g^{2D} , and apply global average pooling (GAP) to get

the global feature vector f_g^{2D} . Then we calculate the inner product $\langle f_g^{2D}, f^{2D} \rangle$ and generate the attention map att^{2D} by $att^{2D} = \text{Softmax}(\langle f_g^{2D}, f^{2D} \rangle)$. Because 2D pose joints can indicate the part locations, we use joint attention to represent 2D part attention. If a joint location has high attention, its *neighboring* points should have high attention too. Thus we can calculate the pose joint attention by summarizing the attentions of its neighboring points. We represent the attention of 17 pose joints as $A^{2D} = \{a_i^{2D}\}_{i=1}^{17}$,

$$\hat{a}_i^{2D} = \frac{\sum_{u,v} att_{(u,v)}^{2D} / (1 + d[(u,v), (u_i, v_i)])}{\sum_{u,v} 1 / (1 + d[(u,v), (u_i, v_i)])}, \quad (2)$$

and $a_i^{2D} = \frac{\hat{a}_i^{2D}}{\sum_{i=1}^{17} \hat{a}_i^{2D}}$, where (u, v) denotes arbitrary point on attention map att^{2D} , (u_i, v_i) indicates the coordinate of the i -th joint (calculated by scaling the joint coordinate on image). $d[\cdot]$ denotes the Euclidean distance between two points. Eq. 2 means: if point (u, v) is far from (u_i, v_i) , the attention value of (u, v) contributes less to the attention value of (u_i, v_i) ; if (u, v) is close to (u_i, v_i) , it contributes more. After the summarizing and normalizing, we finally obtain the attention of (u_i, v_i) , *i.e.* a_i^{2D} .

3D Attention. We use 3D joint attention to represent the 3D body part attention. Input f_{sp}^{3D} is $[1228 \times 384]$ and f_H^{3D} is $[1024]$. We first **tile** f_H^{3D} 1228 times to get shape $[1228 \times 1024]$, then concatenate it with f_{sp}^{3D} to get f^{3D} ($[1228 \times 1408]$). Then we apply GAP to f^{3D} to get a $[1408]$ tensor, and feed it to two 512 sized FC layers and Softmax, finally obtain the attention for 17 joints, *i.e.*, $A^{3D} = \{a_j^{3D}\}_{j=1}^{17}$.

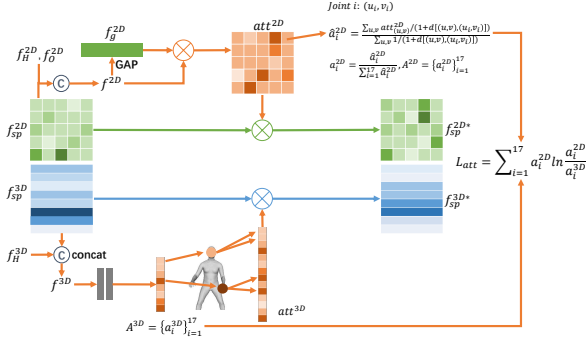


Figure 6. Body part attention alignment. For 2D, we apply self-attention on f_{sp}^{2D} to generate attention map att^{2D} and 2D part attention A^{2D} . For 3D, we use f_{sp}^{3D} to generate 3D part attention A^{3D} , and get attention map att^{3D} using the correspondence between point cloud and joints. Finally, we construct consistency loss \mathcal{L}_{att} with A^{3D} and A^{2D} . att^{2D} and att^{3D} are used to re-weight and generate f_{sp}^{2D*} and f_{sp}^{3D*} .

Attention Consistency. Whereafter, we operate the attention alignment via an attention consistency loss:

$$\mathcal{L}_{att} = \sum_i^{17} a_i^{2D} \ln \frac{a_i^{2D}}{a_i^{3D}}. \quad (3)$$

where a_i^{2D} and a_i^{3D} are 2D and 3D attentions of the i -th joint. L_{att} is the Kullback–Leibler divergence between A^{2D} and A^{3D} , which enforces two attention estimators to generate similar **part importance** and keep the consistency.

Next, in 2D-RN, we multiply f_{sp}^{2D} by att^{2D} , *i.e.* the Hadamard product $f_{sp}^{2D*} = f_{sp}^{2D} \circ att^{2D}$. In 3D-RN, we first assign attention to each 3D point in the spatial configuration volume (n points in total). For human 3D points, we divide them into different sets according to 17 joints, and each set is corresponding to a body part. Within the i -th set, we *tile* the body part attention a_i^{3D} to each point. For object 3D points, we set all their attention as one. Because each element of f_{sp}^{3D} is corresponding to a 3D point in the spatial configuration volume, we organize the attentions of both human and object 3D points as att^{3D} of size $n \times 1$, where n is the number of elements in f_{sp}^{3D} (Fig. 6). Thus we can calculate the Hadamard product $f_{sp}^{3D*} = f_{sp}^{3D} \circ att^{3D}$. After the part feature re-weighting, our model can learn to neglect the parts unimportant to the HOI inference.

4.5. 2D-3D Semantic Consistency

After the feature extraction and re-weighting, we perform the HOI classification. All classifiers in each stream are composed of two 1024 sized FC layers and Sigmoids. The HOI score of the 2D-RN is $\mathcal{S}^{2D} = (s_H^{2D} + s_O^{2D}) \circ s_{sp}^{2D}$, where s_H^{2D} , s_O^{2D} , s_{sp}^{2D} are the scores of human, object and spatial stream. $\mathcal{S}^{3D} = s_H^{3D} + s_{sp}^{3D}$ indicates the final prediction of the 3D-RN. To maintain the semantic consistency of

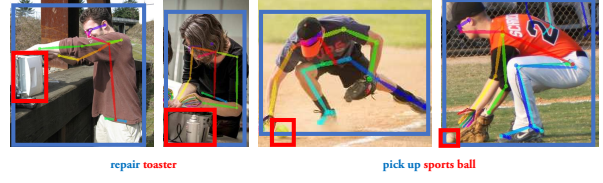


Figure 7. Ambiguous samples from Ambiguous-HOI.

2D and 3D representations, *i.e.* they should make the same prediction for the same sample, we construct:

$$\mathcal{L}_{sem} = \sum_i^m \|\mathcal{S}_i^{2D} - \mathcal{S}_i^{3D}\|_2, \quad (4)$$

where m is the number of HOIs.

Multiple HOI Inferences. Moreover, we concatenate the features from the last FC layers in 2D-RN and 3D-RN as f^{joint} (early fusion), and make the third classification to obtain the score \mathcal{S}^{joint} . The joint classifier is also composed of two 1024 sized FC layers and Sigmoids. The multi-label classification cross-entropy losses are expressed as \mathcal{L}_{cls}^{2D} , \mathcal{L}_{cls}^{3D} , $\mathcal{L}_{cls}^{joint}$. Thus, the total loss of DJ-RN is:

$$\mathcal{L}_{total} = \lambda_1 \mathcal{L}_{tri} + \lambda_2 \mathcal{L}_{att} + \lambda_3 \mathcal{L}_{sem} + \lambda_4 \mathcal{L}_{cls}, \quad (5)$$

where $\mathcal{L}_{cls} = \mathcal{L}_{cls}^{2D} + \mathcal{L}_{cls}^{3D} + \mathcal{L}_{cls}^{joint}$, and we set $\lambda_1=0.001$, $\lambda_2=0.01$, $\lambda_3=0.01$, $\lambda_4=1$ in experiments. The final score is

$$\mathcal{S} = \mathcal{S}^{2D} + \mathcal{S}^{3D} + \mathcal{S}^{joint}. \quad (6)$$

5. Experiment

In this section, we first introduce the adopted datasets and metrics. Then we describe the detailed implementation of DJ-RN. Next, we compare DJ-RN with the state-of-the-art on HICO-DET [9] and Ambiguous-HOI. At last, ablation studies are operated to evaluate modules in DJ-RN.

5.1. Ambiguous-HOI

Existing benchmarks mainly focus on evaluating generic HOIs, but not to specially examine the ability to process 2D pose and appearance ambiguities. Hence, we propose a new benchmark named Ambiguous-HOI. Ambiguous-HOI consists of hard examples collected from the test set of HICO-DET [9], and other whole datasets such as V-COCO [20], OpenImage [28], HCVRD [65] and Internet images. We choose HOI categories from HICO-DET [9] for its well-designed verbs and objects. For Internet images, we labeled the HOIs according to HICO-DET setting. The 2D pose and spatial configuration ambiguities are mainly considered in the selection. First, we put all images and corresponding labels in a candidate pool and manually choose some template 2D pose samples for each HOI. Then we use Procrustes transformation [5] to align the 2D pose of samples to the templates. Next, we cluster all samples to find the

samples far from the cluster center and repeat clustering according to different templates. The mean distance between a sample and multiple cluster centers is recorded as reference. Meanwhile, we train an MLP taking the 2D pose and spatial map as inputs on HICO-DET train set. Then we use it as an *ambiguity probe* to find the most easily misclassified samples. Combining the above two references, we finally select 8,996 images with 25,188 annotated human-object pairs. Ambiguous-HOI finally includes 87 HOI categories, consisting of 48 verbs and 40 object categories from HICO-DET [9]. Some sample are shown in Fig. 7.

5.2. Dataset and Metric

Dataset. We adopt the widely-used HOI benchmark HICO-DET [9] and our novel Ambiguous-HOI. HICO-DET [9] is an instance-level benchmark consisting of 47,776 images (38,118 for training and 9,658 for testing) and 600 HOI categories. It contains 80 object categories from COCO [32], 117 verbs and more than 150k annotated HOI pairs.

Metric. We use mAP metric from [9] for two benchmarks: true positive need to contain accurate human and object locations (box IoU with reference to the ground truth box is larger than 0.5) and accurate interaction/verb classification.

5.3. Implementation Details

For 3D body recovery, we first use OpenPose [7] to detect the 2D pose of body, face and hands. Then we feed them with the image to SMPLify-X [45] to get 3D body. Since cases with severe occlusion might fail the 3D recovery, we only recover 3D bodies for those which at least includes detected 2D head, pelvis, one shoulder and one hip joints. For the rest, we assign them the body with standard template 3D pose, *i.e.* generated by setting all SMPL-X parameters to zero. Sometimes the recovered body can be implausible, *i.e.* “monsters”. To exclude them, we use VPoser[45] to extract the latent embedding of every recovered 3D body. With the *mean latent embedding* of the generated 3D body from HICO-DET train set as a reference, we assume that the farthest 10% embeddings from the mean embedding are “monsters”. At last, 81.2% of the annotated instances are assigned with SMPLify-X [45] generated mesh, and we assign standard templates for the rest to avoid importing noise.

For feature extraction, we use COCO [32] pre-trained ResNet-50 [23] in 2D-RN. In 3D-RN, we first train a PointNet [50] to segment the human and object points in 3D volume, and then use it to extract the 3D local feature of volume. The PointNet is trained for 10K iterations, using SGD with learning rate of 0.01, momentum of 0.9 and batch size of 32. In spatial alignment, we adopt the triplet loss with semi-hard sampling, *i.e.*, for a sample, we only calculate the loss for its nearest negative and farthest positive samples in the same mini-batch with respect to their Euclidean

Method	Default			Known Object		
	Full	Rare	Non-Rare	Full	Rare	Non-Rare
Shen <i>et al.</i> [56]	6.46	4.24	7.12	-	-	-
HO-RCNN [9]	7.81	5.37	8.54	10.41	8.94	10.85
InteractNet [19]	9.94	7.16	10.77	-	-	-
GPNN [51]	13.11	9.34	14.23	-	-	-
iCAN [17]	14.84	10.45	16.15	16.26	11.33	17.73
Interactiveness [31]	17.03	13.42	18.11	19.17	15.51	20.26
No-Frills [21]	17.18	12.17	18.68	-	-	-
PMFNet [58]	17.46	15.65	18.00	20.34	17.47	21.20
Julia <i>et al.</i> [48]	19.40	14.60	20.90	-	-	-
\mathcal{S}^{2D}	19.98	16.97	20.88	22.56	19.48	23.48
\mathcal{S}^{3D}	12.41	13.08	12.21	16.95	17.74	16.72
\mathcal{S}^{Joint}	20.61	17.01	21.69	23.21	19.66	24.28
DJ-RN	21.34	18.53	22.18	23.69	20.64	24.60

Table 1. Results comparison on HICO-DET [9].

distances. In joint training, we train the whole model for 400K iterations, using SGD with momentum of 0.9, following cosine learning rate restart [36] with initial learning rate of 1e-3. For a fair comparison, we use object detection from iCAN [17]. We also adopt the Non-Interaction Suppression (NIS) and Low-grade Instance Suppression (LIS) [31] in inference. The interactiveness model from [31] is trained on HICO-DET train set only. The thresholds of NIS are 0.9 and 0.1 and LIS parameters follow [31].

5.4. Results and Comparisons

HICO-DET. We demonstrate our quantitative results in Tab. 1, compared with state-of-the-art methods [56, 9, 19, 51, 17, 31, 21, 58, 48]. The evaluation follows the settings in HICO-DET[9]: Full(600 HOIs), Rare(138 HOIs) and Non-Rare(462 HOIs) in Default and Known Object mode. We also evaluate different streams in our model, *i.e.* 2D (\mathcal{S}^{2D}), 3D (\mathcal{S}^{3D}) and Joint (\mathcal{S}^{Joint}). Our 2D-RN has a similar multi-stream structure, object detection and backbone following HO-RCNN [9], iCAN [17], Interactiveness [31] and PMFNet [58]. With **joint learning**, 2D-RN (\mathcal{S}^{2D}) directly outperforms above methods with **13.53, 6.50, 4.31, 3.88** mAP on Default Full set. This strongly proves the effectiveness of the consistency tasks in joint learning. Meanwhile, 3D-RN (\mathcal{S}^{3D}) achieves 12.41 mAP on Default Full set and shows obvious complementarity for 2D-RN. Especially, 3D performs better on Rare set than Non-Rare set. This suggests that 3D representation has much *weaker data-dependence* than 2D representation and is less affected by the long-tail data distribution. Joint learning (\mathcal{S}^{Joint}) performs better than both 2D and 3D, achieving 20.61 mAP, while unified DJ-RN (late fusion) finally achieves 21.34 mAP, which outperforms the latest state-of-the-art [48] with 1.94 mAP. Facilitated by the detailed 3D body information, we achieve **21.71** mAP on 356 *hand-related* HOIs, which is higher than the 21.34 mAP on 600 HOIs.

Ambiguous-HOI. To further evaluate our method, we conduct an experiment on the proposed Ambiguous-HOI. We choose methods [17, 31, 48] with open-sourced code as baselines. All models are trained on HICO-DET train set



Figure 8. Visualized attention. Three rows are images, 2D and 3D attentions respectively. Red indicates high attention and blue is the opposite. 2D attention is in line with 3D attention, and they both capture reasonable part attentions for various HOIs.

Method	Ambiguous-HOI
iCAN [17]	8.14
Interactiveness [31]	8.22
Julia <i>et al.</i> [48]	9.72
DJ-RN	10.37

Table 2. Results comparison on Ambiguous-HOI.

Method	Default			Known Object		
	Full	Rare	Non-Rare	Full	Rare	Non-Rare
DJ-RN	21.34	18.53	22.18	23.69	20.64	24.60
3D Pose	20.42	16.88	21.47	22.95	19.48	23.99
Point Cloud	20.05	16.52	21.10	22.61	19.11	23.66
w/o Face	21.02	17.56	22.05	23.48	19.80	24.58
w/o Hands	20.83	17.36	21.87	23.40	19.99	24.41
w/o Face & Hands	20.74	17.36	21.75	23.33	19.82	24.37
w/o Volume Block	20.34	17.19	21.28	22.97	19.94	23.87
w/o 3D Body Block	20.01	16.14	21.17	22.73	18.88	23.88
w/o \mathcal{L}_{att}	20.70	16.56	21.93	23.32	19.13	24.57
w/o \mathcal{L}_{tri}	20.83	17.66	21.77	23.50	20.31	24.45
w/o \mathcal{L}_{sem}	20.80	17.51	21.78	23.45	20.27	24.39

Table 3. Results of ablation studies.

and achieve respective best performances. To test the ability of disambiguation and generalization, we directly test all models on Ambiguous-HOI. Ambiguous-HOI is much more difficult, thus all methods get relatively low scores (Tab. 2). DJ-RN outperforms previous method by 0.65, 2.15 and 2.23 mAP. This strongly verifies the advantage of our joint representation.

Visualizations. We visualize the part attention in Fig. 8. We can find that two kinds of attention are aligned well and both capture essential parts for various HOIs. We also visualize



Figure 9. Visualized results and the corresponding 3D volumes.

HOI predictions paired with estimated 3D spatial configuration volumes in Fig. 9. Our method performs robustly in HOI inference and 3D spatial configuration estimation.

Time Complexity. 2D-RN has similar complexity with iCAN [17] and Interactiveness [31]. 3D-RN is very efficient because of the **pre-extracted features** (about 50 FPS). SMPLif-X runs with GPU acceleration is about 5 FPS.

5.5. Ablation Study

We evaluate different components of our method on HICO-DET. The results are shown in Tab. 3.

3D Formats. Using 3D pose or point cloud for 3D body block in 3D-RN performs worse than VPoser embedding.

3D Human Inputs. Without detailed face and hand shape, DJ-RN shows obvious degradation, especially DJ-RN without hand shape. Because about 70% verbs in HICO-DET are hand-related, which is consistent with daily experience.

Blocks. Without volume or body block in 3D-RN hurts the performance with 1.00 and 1.33 mAP.

Losses: Without \mathcal{L}_{att} , \mathcal{L}_{tri} and \mathcal{L}_{sem} , the performance degrades 0.64, 0.51 and 0.54 mAP.

6. Conclusion

In this paper, we propose a novel 2D-3D joint HOI representation learning paradigm, DJ-RN. We first represent the HOI in 3D with detailed 3D body and estimated object location and size. Second, a 2D Representation Network and a 3D Representation Network are proposed to extract multi-modal features. Several cross-modal consistency tasks are finally adopted to drive the joint learning. On HICO-DET and our novel benchmark Ambiguous-HOI, DJ-RN achieves state-of-the-art results.

Acknowledgement: This work is supported in part by the National Key R&D Program of China, No. 2017YFA0700800, National Natural Science Foundation of China under Grants 61772332 and Shanghai Qi Zhi Institute.

References

- [1] Ijaz Akhter and Michael J Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *CVPR*, 2015. 2, 4
- [2] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5), 2009. 1
- [3] Fabien Baradel, Christian Wolf, and Julien Mille. Pose-conditioned spatio-temporal attention for human action recognition. *arXiv preprint arXiv:1703.10106*, 2017. 2
- [4] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it simple: Automatic estimation of 3d human pose and shape from a single image. In *ECCV*, 2016. 2
- [5] I Borg and P Groenen. Modern multidimensional scaling: Theory and applications springer-verlag new york 471. 1997. 6
- [6] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015. 1
- [7] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. In *arXiv preprint arXiv:1812.08008*, 2018. 1, 2, 3, 4, 7
- [8] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 1
- [9] Yu-Wei Chao, Yunfan Liu, Xieyang Liu, Huayi Zeng, and Jia Deng. Learning to detect human-object interactions. In *WACV*, 2018. 2, 3, 4, 6, 7
- [10] Vasileios Choutas, Philippe Weinzaepfel, Jérôme Revaud, and Cordelia Schmid. Potion: Pose motion representation for action recognition. In *CVPR*, 2018. 2
- [11] CMU. Cmu mocap dataset. 2, 4
- [12] Wenbin Du, Yali Wang, and Yu Qiao. Rpan: An end-to-end recurrent pose-attention network for action recognition in videos. In *ICCV*, 2017. 5
- [13] Yong Du, Wei Wang, and Liang Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *CVPR*, 2015. 1, 2
- [14] Haoshu Fang, Yuanlu Xu, Wenguan Wang, Xiaobai Liu, and Song-Chun Zhu. Learning pose grammar to encode human body configuration for 3d pose estimation. In *AAAI*, 2018. 2
- [15] Hao-Shu Fang, Jinkun Cao, Yu-Wing Tai, and Cewu Lu. Pairwise body-part attention for recognizing human-object interactions. In *ECCV*, 2018. 5
- [16] Hao-Shu Fang, Jianhua Sun, Runzhong Wang, Minghao Gou, Yong-Lu Li, and Cewu Lu. Instaboost: Boosting instance segmentation via probability map guided copy-pasting. In *ICCV*, 2019. 1
- [17] Chen Gao, Yuliang Zou, and Jia-Bin Huang. ican: Instance-centric attention network for human-object interaction detection. *arXiv preprint arXiv:1808.10437*, 2018. 2, 4, 7, 8
- [18] Rohit Girdhar and Deva Ramanan. Attentional pooling for action recognition. In *NIPS*, 2017. 2
- [19] Georgia Gkioxari, Ross Girshick, Piotr Dollár, and Kaiming He. Detecting and recognizing human-object interactions. *arXiv preprint arXiv:1704.07333*, 2017. 2, 7
- [20] Saurabh Gupta and Jitendra Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015. 2, 6
- [21] Tanmay Gupta, Alexander Schwing, and Derek Hoiem. No-frills human-object interaction detection: Factorization, appearance and layout encodings, and training techniques. *arXiv preprint arXiv:1811.05967*, 2018. 1, 2, 7
- [22] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 1
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 7
- [24] Zhiwu Huang, Chengde Wan, Thomas Probst, and Luc Van Gool. Deep learning on lie groups for skeleton-based action recognition. In *CVPR*, 2017. 2
- [25] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *TPAMI*, 2013. 2, 4
- [26] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 2
- [27] Qihong Ke, Mohammed Bennamoun, Senjian An, Ferdous Sohel, and Farid Boussaid. A new representation of skeleton sequences for 3d action recognition. In *CVPR*, 2017. 2
- [28] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv preprint arXiv:1811.00982*, 2018. 2, 6
- [29] Inwoong Lee, Doyoung Kim, Seoungyoon Kang, and Sanghoon Lee. Ensemble deep learning for skeleton-based action recognition using temporal sliding lstm networks. In *ICCV*, 2017. 2
- [30] Yong-Lu Li, Liang Xu, Xijie Huang, Xinpeng Liu, Ze Ma, Mingyang Chen, Shiyi Wang, Hao-Shu Fang, and Cewu Lu. Hake: Human activity knowledge engine. *arXiv preprint arXiv:1904.06539*, 2019. 2
- [31] Yong-Lu Li, Siyuan Zhou, Xijie Huang, Liang Xu, Ze Ma, Hao-Shu Fang, Yanfeng Wang, and Cewu Lu. Transferable interactivity knowledge for human-object interaction detection. In *CVPR*, 2019. 1, 2, 4, 7, 8
- [32] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 3, 4, 7
- [33] Jun Liu, Amir Shahroudy, Dong Xu, and Gang Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In *ECCV*, 2016. 2
- [34] Jun Liu, Gang Wang, Ping Hu, Ling-Yu Duan, and Alex C Kot. Global context-aware attention lstm networks for 3d action recognition. In *CVPR*, 2017. 2

- [35] Mengyuan Liu and Junsong Yuan. Recognizing human actions as the evolution of pose estimation maps. In *CVPR*, 2018. 2
- [36] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 7
- [37] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *ECCV*, 2016. 1
- [38] Cewu Lu, Hao Su, Yonglu Li, Yongyi Lu, Li Yi, Chi-Keung Tang, and Leonidas J Guibas. Beyond holistic object recognition: Enriching image understanding with part states. In *CVPR*, 2018. 2
- [39] Diogo C Luvizon, David Picard, and Hedi Tabia. 2d/3d pose estimation and action recognition using multitask deep learning. In *CVPR*, 2018. 1, 2
- [40] Diogo Carbonera Luvizon, Hedi Tabia, and David Picard. Learning features combination for human action recognition from skeleton sequences. *Pattern Recognition*, 2017. 2
- [41] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013. 4
- [42] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *3DV*, 2018. 2
- [43] Bo Pang, Kaiwen Zha, Hanwen Cao, Chen Shi, and Cewu Lu. Deep rnn framework for visual sequential applications. In *CVPR*, 2019. 1
- [44] Bo Pang, Kaiwen Zha, Yifan Zhang, and Cewu Lu. Further understanding videos through adverbs: A new video task. In *AAAI*, 2020. 1
- [45] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 1, 2, 3, 4, 7
- [46] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. In *CVPR*, 2017. 2
- [47] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3d human pose and shape from a single color image. In *CVPR*, 2018. 2
- [48] Julia Peyre, Ivan Laptev, Cordelia Schmid, and Josef Sivic. Detecting rare visual relations using analogies. In *ICCV*, 2019. 2, 7, 8
- [49] Huy Hieu Pham, Houssam Salmane, Louahdi Khoudour, Alain Crouzil, Pablo Zegers, and Sergio A Velastin. A unified deep framework for joint 3d pose estimation and action recognition from a single rgb camera. *arXiv preprint arXiv:1907.06968*, 2019. 2
- [50] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 4, 7
- [51] Siyuan Qi, Wenguan Wang, Baoxiong Jia, Jianbing Shen, and Song-Chun Zhu. Learning human-object interactions by graph parsing neural networks. In *ECCV*, 2018. 2, 7
- [52] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 2, 4
- [53] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015. 5
- [54] Amir Shahrudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+d: A large scale dataset for 3d human activity analysis. In *CVPR*, 2016. 2
- [55] Dian Shao, Yu Xiong, Yue Zhao, Qingqiu Huang, Yu Qiao, and Dahua Lin. Find and focus: Retrieve and localize video events with natural language queries. In *ECCV*, 2018. 1
- [56] Liyue Shen, Serena Yeung, Judy Hoffman, Greg Mori, and Fei Fei Li. Scaling human-object interaction recognition through zero-shot learning. In *WACV*, 2018. 7
- [57] Sijie Song, Cuiling Lan, Junliang Xing, Wenjun Zeng, and Jiaying Liu. An end-to-end spatio-temporal attention model for human action recognition from skeleton data. In *AAAI*, 2017. 2
- [58] Bo Wan, Desen Zhou, Yongfei Liu, Rongjie Li, and Xuming He. Pose-aware multi-level feature network for human object interaction detection. *arXiv preprint arXiv:1909.08453*, 2019. 1, 2, 7
- [59] Chunyu Wang, Yizhou Wang, and Alan L Yuille. An approach to pose-based action recognition. In *CVPR*, 2013. 2
- [60] Jiang Wang, Zicheng Liu, Ying Wu, and Junsong Yuan. Learning actionlet ensemble for 3d human action recognition. *TPAMI*, 2013. 2
- [61] Wenqiang Xu, Yonglu Li, and Cewu Lu. Srda: Generating instance segmentation annotation via scanning, reasoning and domain adaptation. In *ECCV*, 2018. 1
- [62] Bangpeng Yao and Li Fei-Fei. Action recognition with exemplar based 2.5 d graph matching. In *ECCV*, 2012. 2
- [63] Bangpeng Yao, Jiayuan Ma, and Li Fei-Fei. Discovering object functionality. In *ICCV*, 2013. 1, 2
- [64] Yang You, Yujing Lou, Qi Liu, Yu-Wing Tai, Weiming Wang, Lizhuang Ma, and Cewu Lu. Prin: Pointwise rotation-invariant network. *arXiv preprint arXiv:1811.09361*, 2018. 1
- [65] Bohan Zhuang, Qi Wu, Chunhua Shen, Ian Reid, and Anton van den Hengel. Care about you: towards large-scale human-centric visual relationship detection. *arXiv preprint arXiv:1705.09892*, 2017. 2, 6