

# Mining Cross-Person Cues for Body-Part Interactiveness Learning in HOI Detection

Xiaoqian Wu<sup>1</sup><sup>\*</sup>, Yong-Lu Li<sup>1,2</sup><sup>\*</sup>, Xinpeng Liu<sup>1</sup>, Junyi Zhang<sup>1</sup>, Yuzhe Wu<sup>3</sup>, and Cewu Lu<sup>1</sup><sup>\*\*</sup>

<sup>1</sup> Shanghai Jiao Tong University

<sup>2</sup> Hong Kong University of Science and Technology

<sup>3</sup> DongHua University

{enlighten,yonglu\_li,junyizhang,lucewu}@sjtu.edu.cn,  
{xinpengliu0907,wuyuzhe486}@gmail.com

**Abstract.** Human-Object Interaction (HOI) detection plays a crucial role in activity understanding. Though significant progress has been made, interactiveness learning remains a challenging problem in HOI detection: existing methods usually generate redundant negative H-O pair proposals and fail to effectively extract interactive pairs. Though interactiveness has been studied in both whole body- and part- level and facilitates the H-O pairing, previous works only focus on the *target person* once (*i.e.*, in a **local** perspective) and overlook the information of the other persons. In this paper, we argue that comparing body-parts of multi-person simultaneously can afford us more useful and supplementary interactiveness cues. That said, to learn body-part interactiveness from a **global** perspective: when classifying a target person’s body-part interactiveness, visual cues are explored not only from herself/himself but also from *other persons in the image*. We construct body-part saliency maps based on self-attention to mine cross-person informative cues and learn the holistic relationships between *all* the body-parts. We evaluate the proposed method on widely-used benchmarks HICO-DET and V-COCO. With our new perspective, the holistic global-local body-part interactiveness learning achieves significant improvements over state-of-the-art. Our code is available at <https://github.com/enlighten0707/Body-Part-Map-for-Interactiveness>.

**Keywords:** Human-Object Interaction, Interactiveness Learning, Body-Part Correlations

## 1 Introduction

Human-Object Interaction (HOI) detection retrieves human and object locations and infers the interactions simultaneously from still images. In practice, an HOI instance is represented as a triplet: (*human, verb, object*). As a sub-task of visual relationship [12,28], it is crucial for activity understanding, embodied AI, etc.

<sup>\*</sup> The first two authors contribute equally.

<sup>\*\*</sup> Cewu Lu is corresponding author, member of Qing Yuan Research Institute and Shanghai Qi Zhi institute.

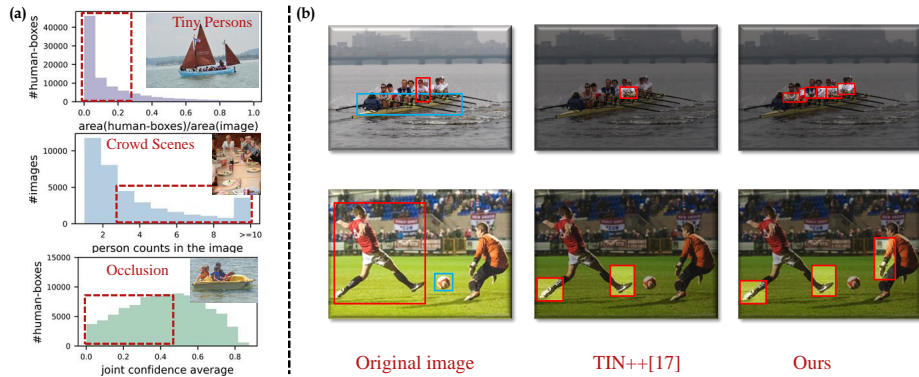


Fig. 1: a) Statistics of hard cases in HOI datasets HICO-DET [1]. The images containing tiny persons, crowded scenes, and occlusion are considered. b) Our idea of learning body-part interactiveness from the same/different body-parts of other persons in the image.

Though significant progress has been made, HOI detection is still bottlenecked by interactiveness learning [22]: they fail to effectively extract interacted pairs while generating redundant negative pairs. The problem is first raised in TIN [22], where a pairwise interactiveness classifier is inserted into the HOI detection framework and the interactiveness prediction is used for non-interaction suppression. The decent gain on HOI detection performance verifies the great potential and importance of interactiveness learning.

Recently, TIN++ [17] is further proposed to utilize body-part level features to improve instance-level interactiveness learning via a hierarchical diagram. Despite the improvement in detecting positive pairs, it focuses on *local* body-part features from the targeted person only, which is not enough. We argue that when classifying the target person’s body-part interactiveness, visual cues can be explored not only from himself/herself but also from *other persons in the image*. First, aside from the targeted person and object, it is also important to exploit contextual cues from the whole image [5]. Existing methods [5,32] have made efforts to learn relevant context to facilitate HOI learning. However, it is hard and **unstable** to retrieve useful information from various backgrounds without restriction. Therefore, we argue to better utilize contextual information by highlighting *all the persons* in the image because regions containing persons are usually more informative according to our prior. For instance, when recognizing a speaker giving a lecture, the audience is evidence more obviously than the stage decoration. Furthermore, as is shown in Fig. 1a, there are a large proportion of **hard cases** in HOI datasets [1], *e.g.*, tiny interactive persons, crowded scenes, and occlusion. In these cases, the cue carried by body parts of the local targeted person is very limited, while mining cues from a *global* multi-person perspective with other persons as a reference would be a good choice to alleviate the difficulty.

Following this insight, we propose to learn body-part interactiveness from the same/different body-parts of other persons in the image, which is illustrated in Fig. 1b. In the upper image, a crowd of tiny persons is rowing boats, and their lower bodies are occluded. When classifying a targeted H-O pair, aside from the targeted object, TIN++ [17] only focuses on the targeted person from the body-part level (*e.g.*, hand), which is highlighted in the image. Nevertheless, we also emphasize other persons’ hands even when classifying the targeted person’s hand interactiveness, which provides a **supplementary** and **contrastive** viewpoint for interactiveness learning. Suppose another person B’s hands are given as a reference which is easier to identify interactiveness, then the similarity between B’s and targeted person A’s hands will lead to the same prediction (both interactive or non-interactive), while discernable difference will lead to an opposite conclusion. Further, attention to different body-parts of multi-person also matters. In the bottom image, the left person is kicking a football while the right person is defending. When classifying the left person’s *feet* interactiveness, the right person’s *arms* would provide useful cues, since he is stretching out his arms and trying to intercept the ball. Thus, the relationship between *different* body-parts of different persons overlooked by previous works also offers supplementary information to interactiveness inference.

In light of these, we utilize a transformer for body-part interactiveness detection, where self-attention helps to capture informative cross-person visual cues. **First**, *body-part saliency maps* are constructed via image patches (*i.e.*, transformer tokens) masking, where patches not containing interested body-parts are masked and blocked from the computation. **Second**, to encode diverse visual patterns more flexibly, body-parts are *progressively masked*, where different attention mask is applied in successive transformer layers and more tokens are dropped in the late layers. **Third**, motivated by the sparsity property [21,17] of body-part interactiveness [17], the model classifies interactiveness of different body-parts via *one-time passing* to improve computation efficiency. An early filter is inserted to drop unimportant body-parts, and then the remaining saliency maps are merged. **Fourth**, we also propose a *sparsity adaptive sampling strategy* on the train set to put more emphasis on crowded scenes and guide better interactiveness knowledge. In extensive experiments, the proposed method achieves state-of-the-art. We firstly achieve all **33+** mAP on three sets of HICO-DET [1], especially the impressive improvement on Rare HOIs (**6.16** mAP improvement upon the SOTA CDN [36]) thanks to our holistic global-local interactiveness learning. Meanwhile, on the HOI hard cases, we also show our significant superiority.

Our contribution includes: 1) We propose to learn body-part interactiveness from a global perspective as an effective supplement for existing local-based methods, thus boosting interactiveness learning; 2) To mine cross-person cues, we construct body-part saliency maps based on self-attention computation and propose the progressively mask and one-time passing strategies to improve flexibility and efficiency; 3) With our proposed interactiveness detector, we achieve state-of-the-art on widely-used benchmarks HICO-DET [1] and V-COCO [8].

## 2 Related Work

**Human-Object Interaction.** Human-Object Interaction is essential to understand human-centric interaction with objects. Rapid progress has been made in HOI learning. Many large datasets [1,8,21,13] and deep learning based methods [1,30,6,5,22,17,16,9,18,23,35,32,36,25,26] have been proposed. They usually followed the two-stage pipeline, *i.e.*, first H-O pair detection, and then HOI classification. Chao *et al.* [1] proposed the widely-used multi-stream framework combining visual features and spatial locations, while GPNN [30] incorporated DNN and graph model to model the HOI relationship. InteractNet [6] utilized an action-specific density map to estimate the interacted object locations. DJ-RN [16] introduced 3D information and proposed a 2D-3D joint representation learning method. PaStaNet [21] inferred human part states [29] first and then reason out the activities based on part-level semantics. IDN [18] analyzed HOI from an integration and decomposition perspective. Recently, several one-stage methods have been proposed [23,35,32,36], where HOIs triplets are directly detected by parallel HOI detectors. PPDM [23] and IP-Net [35] adopted a variant of one-stage object detector for HOI detection. QPIC [32] utilized the recent transformer-based object detector DETR [38] to aggregate image-wide contextual information and facilitate HOI learning.

**Interactiveness Learning.** Though significant progress has been made, interactiveness learning remains challenging in HOI detection. Existing methods fail to pair interactive human and object effectively but generate redundant negative pairs. TIN and TIN++ [22,17] first raised this problem and tried to address it via an inserted pairwise interactiveness classifier. In TIN++ [17], the framework is extended to a hierarchical paradigm with jointly learning instance-level and body part-level interactiveness. Recently, CDN [36] proposed to accurately locate interactive pairs via a one-stage transformer framework disentangling human-object pair detection and interaction classification. However, it still performs not well on interactiveness detection (Sec. 5.3). Here, we point out the previously overlooked global perspective and utilize it to improve interactiveness learning.

**Part-based Action Recognition.** The part-level human feature provides finer-grained visual cues to improve HOI detection. Based on the whole person and part boxes, Gkioxari *et al.* [7] developed a part-based model to make fine-grained action recognition. Fang *et al.* [3] proposed a pairwise body-part attention model which can learn to focus on crucial parts and their correlations, while TIN++ [17] utilized the human instance and body-part features together to learn interactiveness in a hierarchical paradigm and extract deeper interactive visual clues.

## 3 Method

### 3.1 Overview

As is shown in Fig. 2, our pipeline consists of three main modules: a *box detector*, an *interactiveness classifier*, and a *verb classifier*. They are all implemented as stacked transformed decoder.

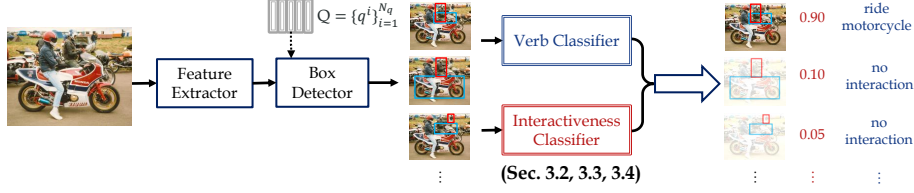


Fig. 2: The overall framework of our proposed method.

Given an input image  $x \in \mathcal{R}^{3 \times H_0 \times W_0}$ , we first adopt a ResNet-50 followed by a transformer encoder as our visual feature extractor. The output feature map is obtained as  $z \in \mathcal{R}^{D_c \times H \times W}$ , where  $D_c$  is the number of channels and  $H, W$  is the size. A fixed positional embedding  $pos \in \mathcal{R}^{D_c \times H \times W}$  is additionally input to the transformer to supplement the positional information. Then based on the feature map  $z$ , the three main components are used for HOI detection.

An interactive human-object pair is mathematically defined following [32]: **1)** a human bounding box vector  $b(h) \in [0, 1]^4$  normalized by the image size, **2)** an object bounding box vector  $b(o) \in [0, 1]^4$  normalized by the image size, **3)** an object class vector  $c \in \{0, 1\}^{N_{obj}}$ , where  $N_{obj}$  is the number of object classes, **4)** an interactiveness prediction  $p_{int} \in [0, 1]$ , and **5)** a verb prediction  $p_{verb} \in \{0, 1\}^{N_v}$ , where  $N_v$  is the number of verb classes. For **box detection**, a transformer decoder  $f_{dec1}(\cdot, \cdot, \cdot)$  transforms a set of learnable query vectors  $Q = \{q^i | q^i \in \mathcal{R}^{D_c}\}_{i=1}^{N_q}$  into a set of decoded embeddings  $D = \{d^i | d^i \in \mathcal{R}^{D_c}\}_{i=1}^{N_q}$ , which is obtained as  $D = f_{dec1}(z, Q, pos)$ . The subsequent three small feed-forward networks (FFNs): human bounding box FFN  $f_h$ , object bounding box FFN  $f_o$ , and object class FFN  $f_c$  further process the decoded embeddings  $D$  to produce  $N_q$  prediction results  $\{b(h)^i\}_{i=1}^{N_q}$ ,  $\{b(o)^i\}_{i=1}^{N_q}$ ,  $\{c^i\}_{i=1}^{N_q}$ , respectively. The decoded embeddings  $D$  are then fed into the interactiveness classifier and verb classifier as their query embeddings. For **verb classification**, another transformer decoder  $f_{dec3}(\cdot, \cdot, \cdot)$  takes  $D$  as input and outputs  $V = \{v^i | v^i \in \mathcal{R}^{D_c}\}_{i=1}^{N_q}$ . With the verb class FFN  $f_v$ , the classification results is obtained as  $p_{verb}^i = Sigmoid(f_v(v^i))$ .

Our main contribution is global-local interactiveness learning based on the proposed body-part saliency map. Thus, we focus on the design of the new proposed **interactiveness classifier** (Fig. 3). In Sec. 3.2, we introduce the construction of body-part saliency map based on self-attention computation, and provide an intuitive scheme to validate its effectiveness (Fig. 3a). Then, we improve the intuitive scheme from two aspects: progressively masking to encode diverse visual patterns flexibly (Sec. 3.3), and one-time passing to improve efficiency (Sec. 3.4). The final improved model is shown in Fig. 3b.

### 3.2 Constructing Body-Part Saliency Map

We divide a person into six body-parts: feet, legs, hip, hands, arms, and head following HAKE [20]. To construct the body-part saliency map, *e.g.*, for hands,

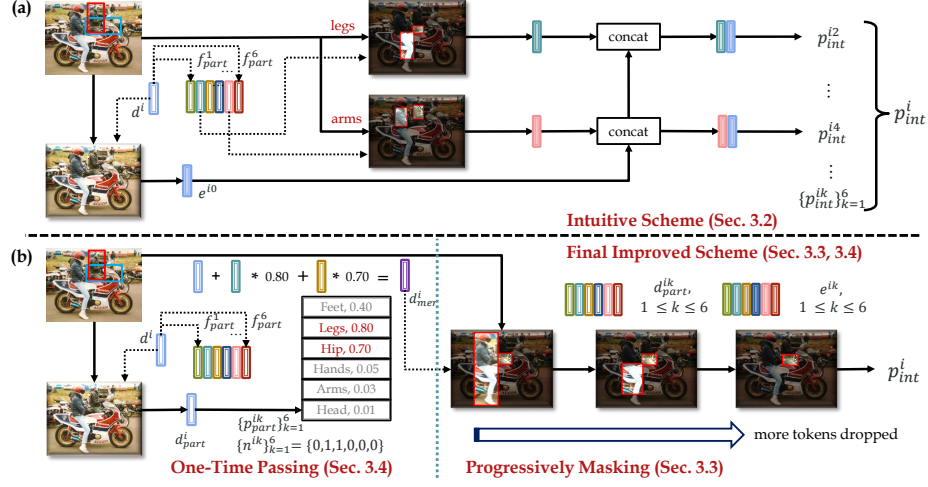


Fig. 3: The design of interactivenss classifier. a) Body-part saliency map construction and the intuitive scheme. We use legs and arms as examples. b) The final improved model. Here, we use the images to show the masking process for clarity instead of masking real images. The decoded embedding  $d^i$  is generated from the box detector and fed into the interactivenss classifier as query embeddings.

we utilize the attention mask matrix in the transformer. Only image patches (or equally, transformer tokens) containing hands are remained for attention calculation, while other patches are masked. The attention mask is pre-calculated from body-part boxes detection results.

**An Intuitive Scheme** In implementation, a transformer decoder  $f_{dec2}(\cdot, \cdot, \cdot, \cdot)$  is used to transform input query embeddings  $D$  into  $E = \{e^i | e^i \in \mathcal{R}^{D_e}\}_{i=1}^{N_q}$ . Body-part saliency map is applied in  $f_{dec2}$  by masking partial of the feature map  $z$  via a masking matrix  $M = \{m^i | m^i \in \{0, 1\}^{H \times W}\}_{i=1}^{N_q}$  (1 for effective tokens and 0 for masked tokens). First, to integrate image-wide contextual information, a set of decoded embeddings are obtained as  $E^0 = f_{dec2}(z, D, pos, M^0)$ , where  $M^0$  is an all-one matrix, *i.e.*, applying no mask. Next, the finer-grained body-part level features are used. The input query embeddings  $D_{part}^k$  for the  $k$ -th ( $1 \leq k \leq 6$ ) body-part is transformed from the original query embeddings  $D$  via FFNs  $\{f_{part}^k\}_{k=1}^6$ , *i.e.*,  $D_{part}^k = f_{part}^k(D)$ . Then the decoded embeddings are calculated via  $E^k = f_{dec2}(z, D_{part}^k, pos, M^k)$ , where  $M^k$  is the corresponding mask matrix. In  $f_{dec2}$ , the attention masks are different for each body-part, while the learned parameters are shared. Finally, based on decoded embeddings  $E^0$  and  $\{E^k\}_{k=1}^6$ , an interactivenss FFN  $f_{int}$  is used to generate body-part interactivenss prediction. For the  $i$ -th proposal and the  $k$ -th body-part, its

interactiveness score is obtained via  $p_{int}^{ik} = \text{Sigmoid}(f_{int}(\text{concat}(e^{ik}, e^{i0})))$ . The instance-level interactiveness score is then obtained as  $p_{int}^i = \max_k \{p_{int}^{ik}\}$ .

**Attention Mask** To calculate the attention mask of all persons in the image, we first obtain the body-part boxes from the pose estimation [4,15,14] results following [17]. For the failure cases of pose estimation, we use the whole-body detection boxes from [31] as body-part boxes. With the  $k$ -th body-part boxes as  $b(p)^k = \{b(p)^{kl} | b(p)^{kl} = [w_1^{kl}, h_1^{kl}, w_2^{kl}, h_2^{kl}], 1 \leq l \leq L\}$  ( $L$  is the persons count in the image), the global body-parts mask matrix  $m_{part}^{ik}$  ( $\forall i, m_{part}^{ik} = m_{part}^k$ ) is calculated as:

$$m_{part}^{k(xy)} = \begin{cases} 1 & \exists l, h_1^{kl} \leq x * (H_0/H) \leq h_2^{kl}, w_1^{kl} \leq y * (W_0/W) \leq w_2^{kl} \\ 0 & \text{otherwise} \end{cases}, \quad (1)$$

where  $x, y$  is the index of the matrix. The scaling factor  $H_0/H$  and  $W_0/W$  are used here because the size of the feature map  $z$  is scaled down from that of the original image, *e.g.*,  $H_0/H = W_0/W = 32$  with ResNet-50 backbone.

Notably, although one proposal corresponds to only *one* targeted H-O pair, its body-part saliency map contains the body-parts of *all* persons in the image. Thus, the relationship between body-parts of all persons is learned from a global perspective. Here  $m_{part}^{ik}$  is a core component of the final mask  $m^{ik}$ . We briefly set  $m^{ik} = m_{part}^{ik}$  for the intuitive scheme in this section, and will further modify  $m^{ik}$  in the next sections.

### 3.3 Progressively Body-Part Masking for Flexibility

In Sec. 3.2, an intuitive scheme with body-part saliency map is introduced, and its effectiveness is validated by experiments (Sec. 5). However, it lacks flexibility to *simply* construct body-part saliency maps to highlight all the **same** body-parts. As aforementioned in Fig. 1b, it also matters to learn attention from **different** body-parts of **different** persons. Thus, we develop a *progressively masking* strategy, where different attention masks are applied in successive transformer layers. Considering that the token representations are encoded more and more sufficiently throughout the whole network, fewer tokens are dropped in the early layers while more tokens are dropped in the late layers. In other words, the “**receptive field**” gradually becomes more focused as the computation proceeds. Attention to different body-parts of different persons is learned in the early layer, then comes the same body-parts of different persons. After encoding useful information from a global perspective, our model focuses on the targeted person in the final layer, which is similar to TIN++ [17].

**Implementation** We construct the progressive masks from three components: body-part mask  $m_{part}^k$  ( $1 \leq k \leq 6$ , Sec. 3.2) shared by proposals, detected human mask  $m_{hum}^i$  ( $1 \leq i \leq N_q$ ) dynamically calculated from the box detection of the  $i$ -th proposal, and detected object mask  $m_{obj}^i$  ( $1 \leq i \leq N_q$ ) similar to  $m_{hum}^i$ . For

the  $i$ -th proposal and its detected human bounding box  $b(h)^i \in [0, 1]^4$ , the box is first unnormalized by the image size as  $b(h)^i = [w_1^i, h_1^i, w_2^i, h_2^i]$ . Then

$$m_{hum}^{i(xy)} = \begin{cases} 1 & h_1^i \leq x * (H_0/H) \leq h_2^i, w_1^i \leq y * (W_0/W) \leq w_2^i \\ 0 & otherwise \end{cases}, \quad (2)$$

The detected object mask  $m_{obj}^i$  is obtained from  $b(o)^i$  in a similar way.

The transformer decoder  $f_{dec2}$  for interactivens inference has three stacked layers. The attention mask  $m_j^{ik} \in \{0, 1\}^{H \times W}$  for the  $i$ -th proposal, the  $j$ -th ( $j = 1, 2, 3$ , larger for later layer) layer and the  $k$ -th body-part is:

$$m_1^{ik} = \max(m_{hum'}^i, m_{part}^k, m_{obj}^i), m_{hum'}^i = \max(\max_k(m_{part}^k) - m_{hum}^i, 0), \quad (3)$$

$$m_2^{ik} = \max(m_{part}^k, m_{obj}^i), \quad (4)$$

$$m_3^{ik} = \max(\min(m_{part}^k, m_{hum}^i), m_{obj}^i). \quad (5)$$

An example is given in Fig. 3b. The targeted object is highlighted in all layers. Specifically, the 1-st layer highlights the  $k$ -th body-part of the detected human, and the whole body of the other humans in the image ( $m_{hum'}^i$ ), which allows attention computation from different body-parts of different persons. The 2-nd layer emphasizes the  $k$ -th body-part from all persons in the image, while the 3-rd layer focuses on the  $k$ -th body-part of the targeted person. With the progressive mask throughout transformer layers, different visual patterns are flexibly encoded to facilitate body-part interactivens learning.

### 3.4 One-Time Passing via Body-Parts Filtering and Merging

In the intuitive scheme, six repeated calculations are needed for each body-part because their saliency maps are not shared. However, this is computationally redundant since body-part interactivens has a notable property: **sparsity**. Namely, only several body-parts will get involved when people are interacting with daily objects. For instance, in “eat apple”, the head and hands have stronger relationships with the apple than the lower body. Therefore, we can select the most important body-parts and only feed their visual feature into the network to classify interactivens, *i.e.*, the *filtering* process. Then, for the selected body-parts, there are still repeated calculations although fewer times. Thus, we spatially *merge* the saliency maps of the selected body-parts, feed them into the network, and obtain the whole-body interactivens prediction via a one-time calculation. The rationality of merging is validated from the following two aspects. First, it seems that the spatial merging leads to “mixed-up” body-part feature, since merging all the six body-parts is equal to instance-level learning instead of body-part level. However, we emphasize that body-part level finer-grained feature remains in our model because only *several* body-parts are merged. Second, here the whole-body interactivens prediction is directly calculated, while in the intuitive scheme we obtain it from six body-part interactivens predictions. This does not impede interactivens learning because most *important* body-parts are used for calculation.



**Implementation** The implementation is illustrated in Fig. 3b. First, the importance of body-parts is calculated via a one-layer transformer decoder layer  $h(\cdot, \cdot, \cdot)$  without mask. For the  $i$ -th proposal, it transforms  $d^i$  into  $d_{part}^i = h(z, d^i, pos)$ , and a subsequent FFN is applied to get body-part attention score  $\{p_{part}^{ik}\}_{k=1}^6 = Sigmoid(FFN(d_{part}^i))$ . Then, body-parts with relatively higher scores (top 25% of  $\{p_{part}^{ik} | 1 \leq i \leq N_q, 1 \leq k \leq 6\}$ ) are chosen for the following computation and others are filtered out. The result is represented as  $\{n^{ik} \in \{0, 1\}\}$  (1 for chosen and 0 for dropped). After the filtering, both the mask matrix and query embeddings are merged as:

$$m_j^i = \max_k(m_j^{ik} * n^{ik}), \quad (6)$$

$$d_{mer}^i = d^i + \sum_k (d_{part}^{ik} * n^{ik} * p_{part}^{ik}). \quad (7)$$

For the mask matrix, the spatial merge is equivalent to the elementwise maximum operation. For the query embeddings, the weighted sum of query embeddings  $d_{part}^{ik}$  of the selected body-parts is added to the whole-image query embeddings  $d^i$  to generate the merged embeddings  $d_{mer}^i$ . Then the decoder  $f_{dec2}$  with progressively masking transforms  $d_{mer}^i$  into embedding  $e_{mer}^i$ . Finally, the interactiveness score is obtained by a FFN via  $p_{int}^i = Sigmoid(FFN(e_{mer}^i))$ . Compared with the intuitive scheme, the one-time passing strategy drops unnecessary computation and improves computation efficiency.

### 3.5 Training and Inference

When training our model, following the set-based training process of QPIC [32], we first match predictions and ground-truths via bipartite matching, then calculate the loss for the matched pairs. For box detection, the loss is composed by three parts: box regression loss  $L_b$ , intersection-over-union loss  $L_u$  and object-class loss  $L_c$ . The loss  $L_{det}$  is obtained via  $L_{det} = \lambda_1 * L_b + \lambda_2 * L_u + \lambda_3 * L_c$ , where  $\lambda_1, \lambda_2, \lambda_3$  are hyperparameters for adjusting the weights. The interactiveness classifier and verb classifier are supervised with classification loss  $L_{int}$  and  $L_{verb}$  respectively.

The training is divided into two stages. First, we train the box detector and interactiveness classifier along with the visual feature extractor with loss  $L_1 = L_{det} + L_{int}$ . Then, the box detector and verb classifier are trained along with the visual feature extractor and the loss is  $L_2 = L_{det} + L_{verb}$ . In inference, we use the interactiveness results to improve verb classification via non-interaction suppression (NIS) [17], where H-O pairs with lower interactiveness scores are filtered out.

## 4 Discussion: Sparse vs. Crowded Scene

In this section, we discuss a naturally raised question: our method focuses on crowded scenes with multi interactive pairs or multi persons/objects, then what about the sparse scenes?

First, our model is adapted to both crowded and sparse scenes. Under crowded scenes, mining cross-person cues provides more useful information of interactiveness. While for sparse scenes, our method would be operated similar to TIN++ [17]. However, our model is still superior to TIN++ [17] thanks to the integration body-part feature and DETR [38] backbone.

Next, we want to re-emphasize the importance of HOI detection in crowded scenes, especially for interactiveness learning and interactive human-object pair proposal detection. We find that crowded images occupy a large proportion of the HOI dataset, validating the effects brought by our method. Meanwhile, the performance of interactiveness learning under crowded scenes is inferior to sparse ones. We split HICO-DET [1] test set into *sparse/crowded* scenes respectively and evaluate the interactiveness AP: 16.96/9.64 AP (TIN++ [17]) and 43.62/33.10 AP (ours, Sec. 5). From the large performance gap between sparse and crowded scenes (7.32 AP for TIN++ [17] and 10.52 AP for ours, Sec. 5), we can see that interactiveness learning is mainly bottlenecked by crowded scenes, where it is usually harder to effectively extract interactive pairs. Therefore, it matters to focus on crowded scenes for interactiveness learning. Statistics are detailed in the supplementary.

Then, we further propose a novel **sparsity adaptive sampling strategy** on train set to put more emphasis on crowded scenes and facilitate interactiveness learning. The sampling probability is modified as  $1:\alpha$  ( $\alpha > 1$ , in practice  $\alpha = 3$ ) for sparse vs. crowded images, compared with the original 1:1, which guides better interactiveness knowledge to identify interactive pairs under complex scenes. Finally, the experiment (Sec. 5.3) proves that the proposed global perspective indeed boosts interactiveness learning, especially for crowded scenes.

## 5 Experiment

In this section, we first introduce the datasets and metrics in Sec. 5.1, and describe implementation details in Sec. 5.2. Next, we report the results on HICO-DET [1] and V-COCO [8] in Sec. 5.3. Some visualization results are given in Sec. 5.4. Finally, ablation studies are conducted in Sec. 5.5.

### 5.1 Dataset and Metric

**Datasets.** We adopt two datasets HICO-DET [1] and V-COCO [8]. HICO-DET [1] includes 47,776 images (38,118 in train set and 9658 in test set), 600 HOI categories on 80 object categories (same with [24]) and 117 verbs, and provides more than 150k annotated human-object pairs. V-COCO [8] provides 10,346 images (2,533 for training, 2,867 for validating, and 4,946 for testing) and 16,199 person instances. Each person has labels for 29 action categories (five of them have no paired object).

**Metrics.** We follow the settings adopted in [1], i.e., a prediction is a true positive only when the human and object bounding boxes both have IoUs larger than 0.5 with reference to ground truth, and the HOI classification result is accurate. The

role mean average precision [8] is used to measure the performance. Additionally, we measure interactiveness detection in a similar setting.

## 5.2 Implementation Details

We adopt ResNet-50 followed by a six-layer transformer encoder as our visual feature extractor. The box detector and the verb classifier are both implemented as a six-layer transformer decoder. The interactiveness classifier is implemented as a three-layer transformer decoder, where selected tokens are masked. During training, AdamW [27] with the weight decay of  $1e-4$  is used. The visual feature extractor and box decoder are initialized from COCO [24] pretrained DETR [38]. The query size is set as 64 for HICO-DET [1] and 100 for V-COCO [8] following CDN [36]. The loss weight coefficients  $\lambda_1, \lambda_2, \lambda_3$  are respectively set as 1, 2.5, 1, exactly the same as QPIC [32]. In 1st stage, the model is trained for 90 epochs with a learning rate of  $1e-4$  which is decreased by 10 times at the 60th epoch. In 2nd stage, the model is fine-tuned for 60 epochs. All experiments are conducted on four NVIDIA GeForce RTX 3090 GPUs with a batch size of 16. In inference, a pairwise NMS with a threshold of 0.6 is conducted.

## 5.3 Results

**Interactiveness Detection.** We evaluate our interactiveness detection on HICO-DET [1] and V-COCO [8]. On HICO-DET [1], we adopt the interactiveness detection AP proposed in TIN++ [17], while on V-COCO [8] we construct the benchmark in a similar way. Tab. 1 shows our interactiveness detection results compared with open-source state-of-the-art methods [17,23,32,36]. For TIN++ [17], the output interactiveness score is used. For PPDM [23], QPIC [32], and CDN [36], the mean of HOI scores (520 HOI categories for HICO-DET [1], and 24 for V-COCO [8]) is used as an approximation.

Table 1: Interactiveness detection results on HICO-DET [1] and V-COCO [8].

Method	HICO-DET [1]	V-COCO [8]
TIN++ [17]	14.35	29.36
PPDM [23]	27.34	-
QPIC [32]	32.96	38.33
CDN [36]	33.55	40.13
ours	<b>38.74 (+5.19)</b>	<b>43.61 (+3.48)</b>

As is shown in Tab. 1, for the two-stage method TIN++ [17], the interactiveness AP is unsatisfactory. We claim that it suffers from exhaustively generated negative H-O pairs from the detector in the first stage, despite the non-interaction suppression [17]. Instead, the one-stage methods PPDM [23], QPIC [32], and CDN [36] benefit from the avoidance of exhaustive pairing and achieve better performance on interactiveness detection. With the insight of

holistically modeling body-part level interactiveness, our method achieves a even better interactiveness AP of **38.74** on HICO-DET [1] and **43.61** on V-COCO [8].

Next, we evaluate the effectiveness of our method in hard cases, as is discussed in Sec. 1. We split HICO-DET [1] test set and compare the interactiveness detection performance of TIN++ [17] and our method. Compared with the local-based method TIN++ [17], our method improves an interactiveness AP gain of 26.66 (157%) / 23.46 (243%) for sparse/crowded scenes, 23.74 (147%) / 23.53(263%) for normal/tiny-persons scenes, and 22.11 (134%) / 14.69 (182%) for less/more-occlusion scenes. We can see that the proposed global perspective indeed boosts interactiveness learning, especially for hard cases where interactive persons are more difficult to identify. For detailed results, please refer to our supplementary.

**Body-Part Interactiveness Detection.** To detail the analysis of body-part level interactiveness learning, we evaluate body-part interactiveness AP from the output body-part attention score. When trained without body-part interactiveness supervision, our method still learns body-part interactiveness well. The interactiveness APs on HICO-DET [1] are: 38.74 (whole body), 11.66 (feet), 5.31 (legs) , 23.83 (hip), 23.11 (hands), 1.38 (arms), 5.51 (head). Further, we utilize annotations provided by HAKE [21,19] to apply body-part supervision, *i.e.*,  $\{p_{part}^{ik}\}_{k=1}^6$  is bound with body-part level interactiveness labels, and the loss is added to  $L_{int}$ . The results are: 38.83 (whole body), 34.89 (feet), 31.02 (legs) , 38.11 (hip), 34.18 (hands), 31.32 (arms), 24.94 (head). We can see that the performance is further improved with body-part supervision, especially for “legs”, “arms”, and “head”. Without supervision, the performance of “arms” and “head” are inferior to other body-parts. One possible reason is that “arms” suffer from relatively ambiguous definitions, and can sometimes be confused with “torso” or “hands” due to occlusion. Additionally, “head” is related to HOIs harder to identify such as “look”, “smell”.

**HOI Detection Boosting.** In Tab. 2 and Tab. 3, we evaluate how HOI learning can benefit from the interactiveness detection results of our method. Here we use instance-level supervision without annotations from HAKE [21,19] for interactiveness learning to compare with TIN++ [17]. In Tab. 2, the first part adopted COCO pre-trained detector. HICO-DET fine-tuned or one-stage detector is used in the second part. All the results are with ResNet-50.

Our method outperforms state-of-the-arts with **35.15/37.56** mAP (Default Full/ Known Object Full) on HICO-DET [1] and **63.0/65.1** mAP (Scenario 1/2) on V-COCO [8], verifying its effectiveness. For two-stage HOI methods, we feed the representative method iCAN [5] (human-object pairs are exhaustively paired) with our detected pairs. Tab. 4 reports the performance comparison on HICO-DET [1] with different pair detection results. We find that with high-quality detected interactive pairs, the performance of iCAN [5] is significantly boosted, especially from the results of our method. We leave the detailed settings in supplementary.

Table 2: Results on HICO-DET [1].

Method	Default			Known Object		
	Full	Rare	Non-Rare	Full	Rare	Non-Rare
iCAN [5]	14.84	10.45	16.15	16.26	11.33	17.73
TIN [22]	17.03	13.42	18.11	19.17	15.51	20.26
PMFNet [34]	17.46	15.65	18.00	20.34	17.47	21.20
DJ-RN [16]	21.34	18.53	22.18	23.69	20.64	24.60
PPDM [23]	21.73	13.78	24.10	24.58	16.65	26.84
VCL [9]	23.63	17.21	25.55	25.98	19.12	28.03
IDN [18]	26.29	22.61	27.39	28.24	24.47	29.37
Zou <i>et al.</i> [39]	26.61	19.15	28.84	29.13	20.98	31.57
AS-Net [2]	28.87	24.25	30.25	31.74	27.07	33.14
QPIC [32]	29.07	21.85	31.23	31.68	24.14	33.93
FCL [10]	29.12	23.67	30.75	31.31	25.62	33.02
GGNet [37]	29.17	22.13	30.84	33.50	26.67	34.89
CDN [36]	31.78	27.55	33.05	34.53	29.73	35.96
Ours	<b>35.15</b>	<b>33.71</b>	<b>35.58</b>	<b>37.56</b>	<b>35.87</b>	<b>38.06</b>

Table 4: The performance comparison on HICO-DET [1] with different pair detection results.

Method	Full	Rare	Non-Rare
<i>iCAN</i> [5]	14.84	10.45	16.15
<i>iCAN</i> [5] <sup>QPIC</sup>	20.36	11.14	23.11
<i>iCAN</i> [5] <sup>CDN</sup>	21.09	11.20	24.04
<i>iCAN</i> [5] <sup>Ours</sup>	<b>24.38</b>	<b>16.27</b>	<b>26.80</b>

Table 3: Results on V-COCO [8].

Method	$AP_{role}(S1)$	$AP_{role}(S2)$
iCAN [5]	45.3	52.4
TIN [22]	47.8	54.2
VSGNet [33]	51.8	57.0
PMFNet [34]	52.0	-
IDN [18]	53.3	60.3
AS-Net [2]	53.9	-
GGNet [37]	54.7	-
HOTR [11]	55.2	64.4
QPIC [32]	58.8	61.0
CDN [36]	62.3	64.4
Ours	<b>63.0</b>	<b>65.1</b>

Table 5: Results of ablation studies on HICO-DET [1].

Method	int AP	HOI mAP
Ours	<b>38.74</b>	<b>35.15</b>
w/o body-part	36.46	32.16
w/o body-part saliency map	37.43	32.60
intuitive scheme	37.91	33.12
w/o progressive mask	38.06	34.05
w/o sparsity adaptive sampling	38.29	34.37
w/o one-time passing	38.51	34.90

## 5.4 Visualization

Fig. 4 shows some visualization results of the learned attention. Our model can effectively learn body-part attention (Fig. 4b) and extract informative cues from other persons in the image, either from the same (Fig. 4d, f) or the different (Fig. 4a, c) body-parts. Learning holistic relationship between body-parts from different persons alleviates the difficulties of interactiveness learning in hard cases, *e.g.*, tiny interactive persons (Fig. 4e, g), crowded scenes (Fig. 4e), and occlusion (Fig. 4f). Also, our method benefits both interactive pairs and non-interactive pairs (Fig. 4d, h). For more please refer to our supplementary.

## 5.5 Ablation Studies

We conduct ablation studies on HICO-DET [1] and list interactiveness AP and corresponding HOI mAP results in Tab. 5.

We first validate the effectiveness of the body-part saliency map. In the intuitive scheme, the body-part saliency map is applied via an attention mask and interactiveness is calculated for each body-part. The intuitive scheme achieves 37.91 interactiveness AP and 33.12 HOI mAP. In contrast, the interactiveness AP falls to 37.43 when body-part saliency map is removed, *i.e.*, for each proposal, body-part and layer, the mask matrix is set as  $\mathbf{1}^{H \times W}$ . It validates the effectiveness of introducing global-level visual cues and learning holistic body-part relationship. Further, if trained with only instance-level visual features without

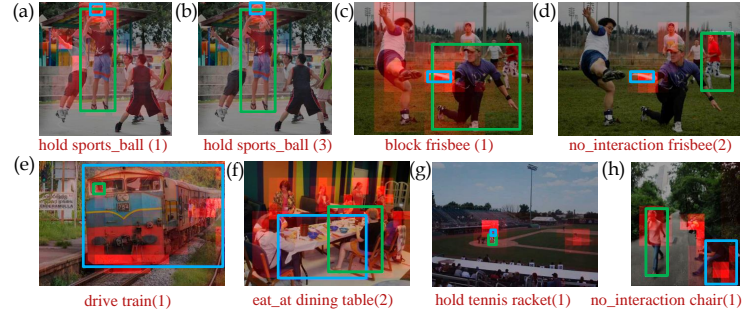


Fig. 4: Visualization results of learned attention. Number  $j$  ( $j = 1, 2, 3$ ) in the bracket represents which layer the attention results are obtained from.

emphasis on body-parts, the interactiveness AP falls to 36.46. Thus, learning body-part interactiveness benefits instance-level interactiveness learning by introducing fine-grained visual features.

Then we evaluate the proposed modules. (1) When progressively masking is removed, all decoder layers are applied with the same attention mask of body-part saliency map, which leads to an interactiveness AP drop to 38.06. The result validates that our model indeed benefits from the progressively masking strategy, where diverse body-part oriented visual patterns are combined and integrated to facilitate interactiveness learning. (2) Without sparsity adaptive sampling, the interactiveness AP falls to 38.29. We find that the sampling augmentation on crowded-scene images helps to extract interactive pairs, especially in complex scenes. (3) Finally, we evaluate the one-time passing strategy. If removed, the interactiveness AP falls to 38.51, and computation speed is reduced by approximately 20% from 28 min/epoch to 33 min/epoch. We can see that our model benefits from it to improve performance as well as efficiency.

## 6 Conclusions

Currently, HOI detection is still bottlenecked by interactiveness learning. In this paper, we focus on learning human body-part interactiveness from a previously overlooked global perspective. We construct body-part saliency maps to mine informative cues from not only the targeted person, but also other persons in the image. Our method provides an effective supplement for existing local-based methods and achieves significant improvements on widely used HOI benchmarks. Despite the improvement, we believe there is still much room left to make further progress on interactiveness learning.

**Acknowledgements.** This work was supported by the National Key R&D Program of China (No. 2021ZD0110700), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), Shanghai Qi Zhi Institute, and SHEITC (2018-RGZN-02046).

## References

1. Chao, Y.W., Liu, Y., Liu, X., Zeng, H., Deng, J.: Learning to detect human-object interactions. In: WACV (2018)
2. Chen, M., Liao, Y., Liu, S., Chen, Z., Wang, F., Qian, C.: Reformulating hoi detection as adaptive set prediction. In: CVPR (2021)
3. Fang, H.S., Cao, J., Tai, Y.W., Lu, C.: Pairwise body-part attention for recognizing human-object interactions. In: ECCV (2018)
4. Fang, H.S., Xie, S., Tai, Y.W., Lu, C.: Rmpe: Regional multi-person pose estimation. In: ICCV (2017)
5. Gao, C., Zou, Y., Huang, J.B.: ican: Instance-centric attention network for human-object interaction detection. In: BMVC (2018)
6. Gkioxari, G., Girshick, R., Dollár, P., He, K.: Detecting and recognizing human-object interactions. In: CVPR (2018)
7. Gkioxari, G., Girshick, R., Malik, J.: Actions and attributes from wholes and parts. In: ICCV (2015)
8. Gupta, S., Malik, J.: Visual semantic role labeling. arXiv preprint arXiv:1505.04474 (2015)
9. Hou, Z., Peng, X., Qiao, Y., Tao, D.: Visual compositional learning for human-object interaction detection. arXiv preprint arXiv:2007.12407 (2020)
10. Hou, Z., Yu, B., Qiao, Y., Peng, X., Tao, D.: Detecting human-object interaction via fabricated compositional learning. In: CVPR (2021)
11. Kim, B., Lee, J., Kang, J., Kim, E.S., Kim, H.J.: Hotr: End-to-end human-object interaction detection with transformers. In: CVPR (2021)
12. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. IJCV (2017)
13. Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Duerig, T., Ferrari, V.: The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. arXiv:1811.00982 (2018)
14. Li, J., Bian, S., Zeng, A., Wang, C., Pang, B., Liu, W., Lu, C.: Human pose regression with residual log-likelihood estimation. In: ICCV (2021)
15. Li, J., Wang, C., Zhu, H., Mao, Y., Fang, H.S., Lu, C.: Crowdpose: Efficient crowded scenes pose estimation and a new benchmark. In: CVPR (2019)
16. Li, Y.L., Liu, X., Lu, H., Wang, S., Liu, J., Li, J., Lu, C.: Detailed 2d-3d joint representation for human-object interaction. In: CVPR (2020)
17. Li, Y.L., Liu, X., Wu, X., Huang, X., Xu, L., Lu, C.: Transferable interactiveness knowledge for human-object interaction detection. TPAMI (2022)
18. Li, Y.L., Liu, X., Wu, X., Li, Y., Lu, C.: Hoi analysis: Integrating and decomposing human-object interaction. NeurIPS (2020)
19. Li, Y.L., Liu, X., Wu, X., Li, Y., Qiu, Z., Xu, L., Xu, Y., Fang, H.S., Lu, C.: Hake: A knowledge engine foundation for human activity understanding. arXiv preprint arXiv:2202.06851 (2022)
20. Li, Y.L., Xu, L., Liu, X., Huang, X., Xu, Y., Chen, M., Ma, Z., Wang, S., Fang, H.S., Lu, C.: Hake: Human activity knowledge engine. arXiv preprint arXiv:1904.06539 (2019)
21. Li, Y.L., Xu, L., Liu, X., Huang, X., Xu, Y., Wang, S., Fang, H.S., Ma, Z., Chen, M., Lu, C.: Pastanet: Toward human activity knowledge engine. In: CVPR (2020)

22. Li, Y.L., Zhou, S., Huang, X., Xu, L., Ma, Z., Fang, H.S., Wang, Y., Lu, C.: Transferable interactiveness knowledge for human-object interaction detection. In: CVPR (2019)
23. Liao, Y., Liu, S., Wang, F., Chen, Y., Feng, J.: Ppdm: Parallel point detection and matching for real-time human-object interaction detection. In: CVPR (2020)
24. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
25. Liu, X., Li, Y.L., Lu, C.: Highlighting object category immunity for the generalization of human-object interaction detection. In: AAAI (2022)
26. Liu, X., Li, Y.L., Wu, X., Tai, Y.W., Lu, C., Tang, C.K.: Interactiveness field in human-object interactions. In: CVPR (2022)
27. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
28. Lu, C., Krishna, R., Bernstein, M., Fei-Fei, L.: Visual relationship detection with language priors. In: ECCV (2016)
29. Lu, C., Su, H., Li, Y., Lu, Y., Yi, L., Tang, C.K., Guibas, L.J.: Beyond holistic object recognition: Enriching image understanding with part states. In: CVPR (2018)
30. Qi, S., Wang, W., Jia, B., Shen, J., Zhu, S.C.: Learning human-object interactions by graph parsing neural networks. In: ECCV (2018)
31. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. NeurIPS (2015)
32. Tamura, M., Ohashi, H., Yoshinaga, T.: QPIC: Query-based pairwise human-object interaction detection with image-wide contextual information. In: CVPR (2021)
33. Ulutan, O., Iftekhhar, A., Manjunath, B.: Vsgnet: Spatial attention network for detecting human object interactions using graph convolutions. In: CVPR (2020)
34. Wan, B., Zhou, D., Liu, Y., Li, R., He, X.: Pose-aware multi-level feature network for human object interaction detection. In: ICCV (2019)
35. Wang, T., Yang, T., Danelljan, M., Khan, F.S., Zhang, X., Sun, J.: Learning human-object interaction detection using interaction points. In: CVPR (2020)
36. Zhang, A., Liao, Y., Liu, S., Lu, M., Wang, Y., Gao, C., Li, X.: Mining the benefits of two-stage and one-stage hoi detection. arXiv preprint arXiv:2108.05077 (2021)
37. Zhong, X., Qu, X., Ding, C., Tao, D.: Glance and gaze: Inferring action-aware points for one-stage human-object interaction detection. In: CVPR (2021)
38. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)
39. Zou, C., Wang, B., Hu, Y., Liu, J., Wu, Q., Zhao, Y., Li, B., Zhang, C., Zhang, C., Wei, Y., et al.: End-to-end human object interaction detection with hoi transformer. In: CVPR (2021)