

RR-Net: Injecting Interactive Semantics in Human-Object Interaction Detection

Dongming Yang¹, Yuexian Zou^{1,2*}, Can Zhang¹, Meng Cao¹ and Jie Chen^{1,2}

¹School of ECE, Peking University, Shenzhen, China, 518055

²Peng Cheng Laboratory, Shenzhen, China, 518055

{yangdongming, zouyx, zhangcan, mengcao}@pku.edu.cn, chenjie@pcl.ac.cn

Abstract

Human-Object Interaction (HOI) detection devotes to learn how humans interact with surrounding objects. Latest end-to-end HOI detectors are short of relation reasoning, which leads to inability to learn HOI-specific interactive semantics for predictions. In this paper, we therefore propose novel relation reasoning for HOI detection. We first present a progressive Relation-aware Frame, which brings a new structure and parameter sharing pattern for interaction inference. Upon the frame, an Interaction Intensifier Module and a Correlation Parsing Module are carefully designed, where: a) interactive semantics from humans can be exploited and passed to objects to intensify interactions, b) interactive correlations among humans, objects and interactions are integrated to promote predictions. Based on modules above, we construct an end-to-end trainable framework named Relation Reasoning Network (abbr. RR-Net). Extensive experiments show that our proposed RR-Net sets a new state-of-the-art on both V-COCO and HICO-DET benchmarks and improves the baseline about 5.5% and 9.8% relatively, validating that this first effort in exploring relation reasoning and integrating interactive semantics has brought obvious improvement for end-to-end HOI detection.

1 Introduction

Fine-grained understanding of visual contents is one of the fundamental problems in computer vision. The task of Human-Object Interaction (HOI) detection aims to localize and infer triplets of $\langle \text{human, verb, object} \rangle$ from a still image. Beyond comprehending instances, e.g., object detection [Ren *et al.*, 2017; Yang *et al.*, 2019], human pose estimation [Fang *et al.*, 2017] and action recognition [Sharma *et al.*, 2015], detecting HOIs requires a deeper understanding of visual semantics to depict complex and finer-grained relationships between $\langle \text{human, object} \rangle$ pairs. Therefore, HOI detection is challenging and usually needs to mine high-level and HOI-specific representations in solutions.

*Contact Author

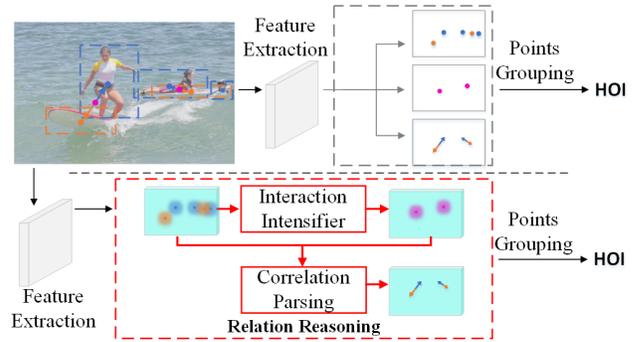


Figure 1: Visualization of how our method facilitates HOI detection. The baseline framework (above) lacks learning relative information during inference, while we present a framework (bottom) that performs relation reasoning to learn HOI-specific interactive semantics.

The conventional two-stage HOI detection methods [Gkioxari *et al.*, 2018; Gao *et al.*, 2018; Xu *et al.*, 2019] first obtain human and object region proposals from a pre-trained object detector [Ren *et al.*, 2017] and then generate $\langle \text{human, object} \rangle$ pairs to make interaction classifications. The related techniques of HOI feature learning are all region-based where instance (e.g., human and object) attention features [Gkioxari *et al.*, 2018; Gao *et al.*, 2018; Xu *et al.*, 2019], spatial configuration [Gao *et al.*, 2018; Gupta *et al.*, 2019; Ulutan *et al.*, 2020] and human pose [Wan *et al.*, 2019; Zhou and Chi, 2019] are exploited. Inspired by the point-based object detectors [Law and Deng, 2018; Zhou *et al.*, 2019], as illustrated in Figure 1 (above), latest end-to-end HOI detectors pose HOI detection as a keypoint detection and grouping problem, which effectively improve the speed of HOI detection.

Previous region-based feature enhancements are hardly applicable to end-to-end detectors due to these end-to-end frameworks do not provide proposed instance regions during inference. Besides, although previous feature learning techniques generally benefit HOI representations, they bring several drawbacks: 1) Firstly, dominant methods focus on region-based feature enhancement but lack relation reasoning, however, comprehending HOI-specific interactive semantics from relation dependency among humans, objects and interactions is significant for HOI inference. 2) Sec-

only, region-based feature learning is complicated, and employing human pose or body-part features brings additional workload. Therefore, studying relation reasoning to integrate region-independent interactive semantics has become the key to improve end-to-end HOI detection.

In this work, as illustrated in Figure 1 (bottom), we propose relation reasoning for HOI detection, which captures and transmits HOI-specific interactive semantics among humans, objects and interactions where no human and object proposals are needed. The contributions are summarized as follows:

- A progressive Relation-aware Frame is first presented to build a coherent structure and parameter sharing pattern for interaction inference.
- Upon the frame, an Interaction Intensifier Module (abbr. IIM) is proposed, where interactive semantics from humans can be exploited and easily passed to objects.
- A Correlation Parsing Module (abbr. CPM) is designed subsequently, in which interactive correlations among visual targets can be exploited and integrated.
- We offer a framework named Relation Reasoning Network (RR-Net), which is the first to carry out region-independent relation reasoning for HOI detection.

We perform extensive experiments on two public benchmarks (i.e., V-COCO [Yatskar *et al.*, 2016] and HICO-DET [Chao *et al.*, 2018] datasets). Our method with relation reasoning provides obvious performance gain compared with the baseline and outperforms the state-of-the-art methods by a sizable margin. Detailed ablation studies of our framework are also provided to facilitate future research.

2 Related Work

2.1 Frameworks in HOI Detection

Since human and object instances are basic elements in detecting HOIs, conventional HOI detectors are mostly based on object detection frameworks. InteractNet [Gkioxari *et al.*, 2018] extends the object detector Faster R-CNN [Ren *et al.*, 2017] with an additional branch for HOI detection. iCAN [Gao *et al.*, 2018] then introduces a three-branch architecture with one branch each for a human candidate, an object candidate, and an interaction pattern, which has become the template framework for most later HOI detectors. TIN [Li *et al.*, 2019] further proposes a multiple network to perform feature extraction, interactivity discrimination and HOI classification, respectively. All above frameworks are two-stage and region-based, where region proposals obtained from pre-trained object detectors are indispensable.

More recently, it has become popular to treat object detection as a keypoint estimation task [Law and Deng, 2018; Zhou *et al.*, 2019] to improve detection efficiency. Following the same idea, IPNet [Wang *et al.*, 2020] and PPDM [Liao *et al.*, 2020] pose HOI detection as a keypoint detection and grouping problem, which propose to build end-to-end HOI detectors. Specifically, PPDM [Liao *et al.*, 2020] contains two parallel branches to predict three points (i.e., human, object and interaction points) and two displacements from the interaction point to its corresponding human and

object points. Then, human with its corresponding object can be paired using the interaction point and displacements as bridge. However, without relation reasoning, these end-to-end HOI detectors lack learning HOI-specific interactive semantics in their methods, thus may be deficient in modelling HOIs (e.g., bring problem of miss-pairing between human and object).

2.2 Feature Learning in HOI Detection

Contextual Cues. Several contextual cues have been explored to improve HOI detection. Given a \langle human, object \rangle candidate with bounding boxes, the interaction pattern [Chao *et al.*, 2018] or spatial configuration [Gao *et al.*, 2018; Ulutan *et al.*, 2020] is a coarse region layout encoded by a binary image with two channels. Tanmay [Gupta *et al.*, 2019] encodes both absolute and relative position of the human and object boxes. Additionally, human pose obtained from pre-trained pose estimators [Fang *et al.*, 2017] is unitized as a valid cue to support HOI detection. With human pose vector, human body-parts [Fang *et al.*, 2018; Wan *et al.*, 2019] are explored as a set of local regions centered at human keypoints with certain size. Although extracting region-based spatial encodings or body-parts improve detection results, these methods are not favored since additional computation and premade annotations are indispensable.

Appearance Attention and Reasoning. Contextual Attention [Wang *et al.*, 2019] and GID-Net [Yang *et al.*, 2020] propose to capture instance-centric attention features and context-aware appearance features for human and object. PMFNet [Wan *et al.*, 2019] focuses on pose-aware attention by employing attentional human parts. VSGNet [Ulutan *et al.*, 2020] proposes to refine visual features by the spatial configuration of human-object pair. Moreover, GPNN [Qi *et al.*, 2018] and in-GraphNet [Yang and Zou, 2020] introduce learnable graph-based structures, in which HOIs are represented with graphs. RPNN [Zhou and Chi, 2019] introduces an object-body graph and a human-body graph to capture contexts between body-parts and surrounding instances, where fine-grained human body-part regions are required as prior information.

Overall, methods employing appearance attention mechanisms or graphs are enlightening but complicated, e.g., implementing attention mechanisms upon M human regions and N object regions needs $M + N$ standalone calculations. Analogously, constructing graphs with $M * N$ possible \langle human, object \rangle pairs needs to calculate $M * N$ candidate samples. More importantly, above methods neglect relation reasoning to explore implicit interactive semantics among humans, objects and interactions, whereas comprehending relation is essential to improve HOI detection. In this paper, unlike the previous techniques of HOI feature learning, we bring region-independent relation reasoning to integrate and transmit HOI-specific interactive semantics for HOI inference.

3 Proposed Method

3.1 Overview of RR-Net

An overview of our proposed RR-Net is shown in Figure 2. DLA-34 [Yu *et al.*, 2018] is employed as the backbone net-

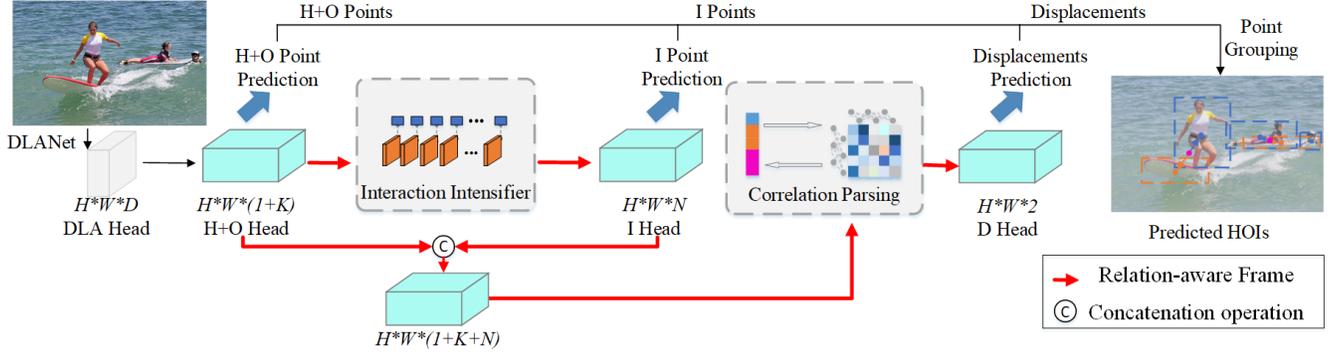


Figure 2: An overview of RR-Net. The framework is coherent, where instance points prediction, interaction points prediction and displacements prediction are treated as highly correlated processes. We also develop an Interaction Intensifier Module and a Correlation Parsing Module to achieve efficient relation reasoning to learn region-independent interactive semantics for HOI inference.

work in our implementation to extract the global $128 \times 128 \times 64$ DLA feature from an image. Our framework formulates HOI detection as a point detection and grouping task [Wang *et al.*, 2020; Liao *et al.*, 2020], where predicting HOI triplets consists of three steps: 1) predicting human points (H points), K -class object points (O points) and N -class interaction points (I points). Here, I point is defined as the midpoint of a corresponding $\langle \text{human, object} \rangle$ pair. 2) predicting two displacements from the I point to the H point and O point respectively, where each displacement contains an abscissa value and an ordinate value. 3) grouping each I point with the H point and O point respectively to generate a set of triplets based on the predicted points and displacements.

We have noticed that most methods focus on HOI feature learning but none of them study the association between HOI cognitive processes or explore interactive semantics in these processes. In order to explore above information, a Relation-aware Frame (See Section 3.2) is first proposed to provide a coherent structure and parameter sharing pattern for interaction inference. Upon the frame, an Interaction Intensifier Module (See Section 3.3) and a Correlation Parsing Module (See Section 3.4) are then developed to achieve region-independent relation reasoning so that interactive semantics and correlations among visual targets can be integrated. Next, we will introduce the proposed framework in detail.

3.2 Relation-aware Frame

As illustrated in Figure 3 (a), to predict H points, O points, I points and two displacements, the baseline build different prediction feature heads independently by the function group $G() = [G_{ho}(), G_i(), G_{dh}(), G_{do}()]$, and head parameters are optimized severally in backpropagation. It should be mentioned that we have merged $G_{dh}()$ and $G_{do}()$ as $G_d()$ in Figure 3 for simplicity. We argue that such a structure neglects the implicit relation among these predictions. It is noted that the human visual system usually detects instances and groups related ones coherently from the scene to recognize a HOI. Based on the above inspiration, we propose a correlated and coherent structure for interaction inference, which is termed as a Relation-aware Frame.

For predicting humans and objects, we first generate H and

O head feature $f_{ho} \in \mathbb{R}^{H \times W \times (1+K)}$ from global DLA head feature $f_{dla} \in \mathbb{R}^{H \times W \times D}$ by the function $G_{ho}()$, where $H \times W$ denotes range, D denotes feature dimension and $(1 + K)$ is the category synthesis of human and K -class objects. The two stages of our Relation-aware Frame are then established to break the independent predictions of baseline, which are: 1) Part A: building I head feature $f_i \in \mathbb{R}^{H \times W \times N}$ upon feature f_{ho} by the function $G_i()$; 2) Part B: building two D head features $f_{dh}, f_{do} \in \mathbb{R}^{H \times W \times 2}$ upon all point head features $f_p = \text{concat}(f_{ho}, f_i)$ by the function $G_d()$. As illustrated in Figure 3 (b), in the fully Relation-aware Frame, output head features of each previous stage are used as input in the next stage. The function group $G()$ are computed by:

$$G(X) = \text{Conv2d}(\text{ReLU}(\text{Conv2d}(X))), \quad (1)$$

where the first 3×3 convolution transforms the features from input dimension to hidden dimension, and the second 1×1 convolution acts as a fully connected layer to reduce the feature from hidden dimension to specified dimension to match the predictions. The next predictions of H points, O points and I points can be regarded as heatmap estimation processes by sigmoid functions. In this way, instance points prediction loss L_{ho} , interaction point prediction loss L_i and two displacements prediction loss $L_d = L_{dh} + L_{do}$ are correlated in backpropagation and parameters are optimized connectedly.

3.3 Interaction Intensifier Module

Human contexts often imply information of taking interactions with different objects. For example, we can clearly distinguish between kickers and riders by their postures and distribution. Meanwhile, kickers can easily remind us of football, but riders remind us of bicycles or motorcycles. To exploit interactive semantics above to intensify interactions, we propose an Interaction Intensifier Module (abbr. IIM) upon the Relation-aware Frame Part A.

After running $G_{ho}()$, each obtained feature map of H and O heads contains characteristic information for prediction of corresponding instances. We innovatively employ these features to perform relation reasoning, where no human and object regions are needed. As mentioned above, the total number $(1 + K)$ of categories contains person and K other

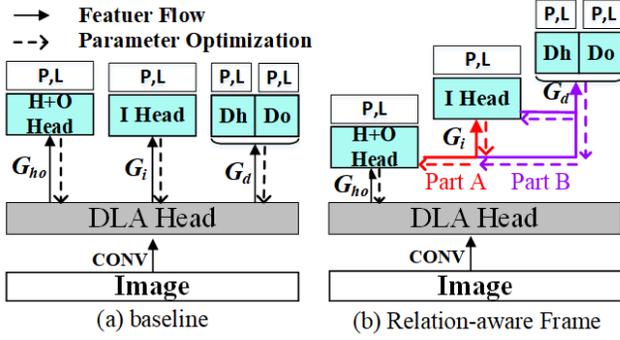


Figure 3: Relation-aware Frame visualization. (a) In baseline, different heads are built independently, and head parameters are optimized severally in backward propagation. (b) Relation-aware Frame: building I head upon H and O heads. Meanwhile, building two D heads upon all points heads (i.e., H, O and I heads). Here, P means prediction operation and L means loss function.

object categories. As illustrated in Figure 4 (a), We can further cut $f_{ho} \in \mathbb{R}^{H*W*(1+K)}$ into $f_h \in \mathbb{R}^{H*W*1}$ and $f_o \in \mathbb{R}^{H*W*K}$. Then, a multilayer perceptron (MLP) with two fully connected layers followed by a sigmoid operation is adopted on feature f_h to generate the embedded interactive vector $\beta = [\beta_1, \beta_2, \dots, \beta_K]$, denoting relevance between humans and different objects. Afterwards, O head features f_o are weighted with vector β by a multiplication operation, generating the valued O head features f'_o .

$$\beta = \text{sigmoid}(\text{MLP}(f_h)), f'_o = \beta * f_o, \quad (2)$$

where $\beta_k \in [0, 1]$ is the k -th interactive value for f_o . Finally, IIM outputs $f'_{ho} = \text{concat}(f_h, f'_o)$, which is in same form as inputs. Differing from treating human and object features separately, with the formulation above, interactive semantics from humans can be dynamically passed to objects to intensify interactions and promote I points prediction.

3.4 Correlation Parsing Module

Interactive correlations among humans, objects, and interactions provide crucial information for HOI inference. We propose Correlation Parsing Module (abbr. CPM) upon the Relation-aware Frame Part B to break away from the previous idea of region-based feature enhancement and exploit the interactive correlations among these targets. The detailed design of CPM is provided in Figure 4 (b). The CPM takes all point head features $f_p \in \mathbb{R}^{H*W*(1+K+N)}$ as input, where the human class, K object classes and N interaction classes are considered as $(1 + K + N)$ visual targets.

In particular, a project function first provides a pattern to convert head feature $X = f_p$ into a dynamic adjacency matrix A . Specifically, given the input $X \in \mathbb{R}^{H*W*(1+K+N)}$, we embed two feature tensors (X^θ, X^ϕ) by $\theta(X)$ and $\phi(X)$ modeled by two $1 * 1$ convolutions. The obtained tensors are then reshaped from $H * W * (1 + K + N)$ to $L * (1 + K + N)$ by the *Planar* operation. Finally, a linear combination transforms visual targets into the unified adjacency matrix $A \in \mathbb{R}^{(1+K+N)*(1+K+N)}$, where $(1 + K + N)$ denotes the length

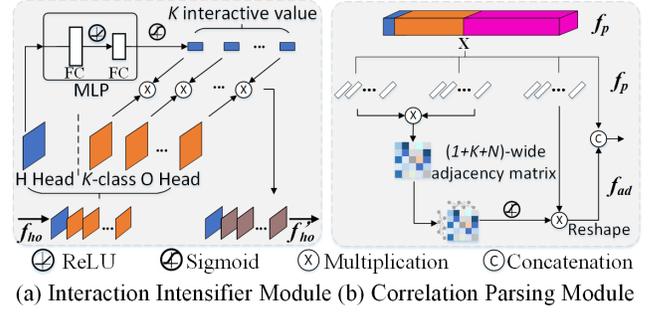


Figure 4: The detailed design of proposed IIM and CPM. IIM takes H and O head features as inputs and passes interactive semantics from humans to objects to intensify interactions. CPM exploits and integrates interactive correlations among visual targets and helps group HOI triplets without region proposals.

and width of the adjacency matrix. Each element $a_{ij} \in A$ denotes the correlation between two visual targets.

$$A = \text{Planar}[\theta(X)]^T * \text{Planar}[\phi(X)]. \quad (3)$$

Based on the adjacency matrix, a message spreading process is further adopted to broadcast and integrate interactive correlations from all elements, which can be computed by:

$$A' = \text{Conv1D}((\text{Conv1D}(A)^T)^T \oplus A), \quad (4)$$

where two operations of *Conv1D* are single-layer 1D convolutions to build communication among all elements. Between the convolutions, the \oplus implements addition point to point which updates the hidden element states according to incoming information. The A' is the updated adjacency matrix.

The following step provides a reverse projection of A' to convolution space. Firstly, a $g(X)$ function and a reshape operation act as linear embedding to compute the representation of input signal as visual targets. Then a weighted linear combination is employed as the reverse projection:

$$Y = \text{sigmoid}(A') * \text{Planar}[g(X)]. \quad (5)$$

In this way, A' plays a role of dynamical weights so that interactive correlations among visual targets can be weighted aggregated. Finally, we reshape the planar tensor Y to three-dimensional f_{ad} . For displacements prediction, f_{ad} as the correlation encoding is concatenated with appearance feature f_p to provide correlative information among all visual targets.

3.5 Training and Inference

Training. In addition to predicting categories of instances and interactions, we also need to regress the size and offset of instance boxes by $G_{wh}()$ and $G_{off}()$, which have the same computation with $G()$ in Section 3.2. In summary, the final loss $L = L_p + \lambda L_d + L_r$ can be calculated as the weighted sum of point prediction loss L_p , displacement prediction loss L_d and instance box regression loss L_r , where λ is set as 0.1 [Zhou *et al.*, 2019]. As mentioned in Section 3.2, point detection can be transformed into a heatmap estimation problem, where $L_p = L_{ho} + L_i$ are optimized by element-wise focal loss. The $L_d = L_{dh} + L_{do}$ is trained by L1 loss, where

Method	Backbone Network	mAP _{role} (%)
InteractNet [Gkioxari <i>et al.</i> , 2018]	ResNet-50-FPN	40.0
GPNN [Qi <i>et al.</i> , 2018]	Deformable CNN	44.0
iCAN [Gao <i>et al.</i> , 2018]	ResNet-50	45.3
Contextual Att [Wang <i>et al.</i> , 2019]	ResNet-50	47.3
RPNN [Zhou and Chi, 2019]	ResNet-50	47.5
TIN(R _d C _d) [Li <i>et al.</i> , 2019]	ResNet-50	47.8
in-GraphNet [Yang and Zou, 2020]	ResNet-50	48.9
IPNet [Wang <i>et al.</i> , 2020]	Hourglass-104	51.0
PMFNet [Wan <i>et al.</i> , 2019]	ResNet-50	52.0
our baseline	DLA-34	51.4
RR-Net	DLA-34	54.2

Table 1: Performance comparison on V-COCO.

Method	Backbone Network	Default		
		full	rare	non-rare
InteractNet [Gkioxari <i>et al.</i> , 2018]	ResNet-50-FPN	9.94	7.16	10.77
GPNN [Qi <i>et al.</i> , 2018]	Deformable CNN	13.11	9.34	14.23
iCAN [Gao <i>et al.</i> , 2018]	ResNet-50	14.84	10.45	16.15
Contextual Att [Wang <i>et al.</i> , 2019]	ResNet-50	16.24	11.16	17.75
RPNN [Zhou and Chi, 2019]	ResNet-50	17.35	12.78	18.71
TIN(R _d C _d) [Li <i>et al.</i> , 2019]	ResNet-50	17.03	13.42	18.11
No-Frills [Gupta <i>et al.</i> , 2019]	ResNet-50	17.18	12.17	18.68
PMFNet [Wan <i>et al.</i> , 2019]	ResNet-50	17.46	15.65	18.00
IPNet [Wang <i>et al.</i> , 2020]	Hourglass-104	19.56	12.79	21.58
PPDM [Liao <i>et al.</i> , 2020]	DLA-34	20.29	13.06	22.45
our baseline	DLA-34	18.87	11.32	21.13
RR-Net	DLA-34	20.72	13.21	22.97

Table 2: Performance comparison on HICO-DET.

L_{dh} and L_{do} are losses for predicted displacement dp^{ih} to H point and displacement dp^{io} to O point respectively for an I point. The $L_r = L_{wh} + L_{off}$ is optimized by L1 loss, which is in charge of regressing box size (h, w) and local offset $d = (d_x, d_y)$ for location of ground truth instance points.

Inference. After points detection, we select top T human, object and interaction points according to the corresponding confidence scores respectively. Then, point grouping process will determine the optimal triplets. Specifically, for each detected I point (x^i, y^i) and predicted displacements dp^{ih} and dp^{io} , several candidate instance points will be ranked according to two criteria, which are 1) have high confidence score, 2) be close to the human/object location that generated by interaction point location plus the human/object displacement.

For a grouped regressive H point (x^h, y^h) and an O point (x^o, y^o) , their bounding boxes are determined as $B_h = (x^h, y^h, h^h, w^h)$ and $B_o = (x^o, y^o, h^o, w^o)$, where (h^h, w^h) and (h^o, w^o) denote the width and height of the boxes. With a set of detected HOI triplets, we use a fusion of scores to predict the final score S_{hoi} for each of them:

$$S_{hoi} = S_{(x^h, y^h)}^h * S_{(x^i, y^i)}^i * S_{(x^o, y^o)}^o, \quad (6)$$

where $S_{(x^h, y^h)}^h$, $S_{(x^i, y^i)}^i$ and $S_{(x^o, y^o)}^o$ are confidence scores of candidate H point, I point and O point.

4 Experiments and Evaluations

4.1 Experimental Setup

Datasets and Evaluation Metrics. We evaluate our method on two large-scale benchmarks, including V-COCO [Yatskar *et al.*, 2016] and HICO-DET [Chao *et al.*, 2018] datasets. V-COCO includes 10,346 images, which contains 16,199 human instances in total and provides 26 common verb categories. HICO-DET contains 47,776 images, where 80 object

Method	mAP _{role} (%)
baseline	51.4
baseline+Relation-aware Frame Part A	52.4
baseline+Relation-aware Frame Part B	52.2
baseline+Fully Relation-aware Frame	52.8

Table 3: Impact of adopting progressive Relation-aware Frame.

Baseline	✓					
Relation-aware Frame Part A		✓		✓	✓	✓
Relation-aware Frame Part B			✓	✓	✓	✓
IIM		✓		✓		✓
CPM			✓		✓	✓
mAP _{role} (%)	51.4	52.9	52.8	53.6	53.4	54.2

Table 4: Impact of employing IIM and CPM.

categories and 117 verb categories compose of 600 HOI categories. There are three different HOI category sets in HICO-DET, which are: (a) all 600 HOI categories (Full), (b) 138 HOI categories with less than 10 training instances (Rare), and (c) 462 HOI categories with 10 or more training instances (Non-Rare). A detected \langle human, verb, object \rangle triplet is considered as a true positive if: 1) it has correct interaction label; 2) both the predicted human and object bounding boxes have IoU of 0.5 or higher with the ground-truth boxes. Following the standard protocols, we use role mean average precision (mAP_{role}) [Yatskar *et al.*, 2016] to report evaluation results.

Implementation details. During training, input images have the resolution of $512 * 512$, yielding a resolution of $128 * 128$ for all output head features. We employ standard data augmentation following [Zhou *et al.*, 2019]. Our model is optimized with Adam. The batch-size is set as 15 for V-COCO and 20 for HICO-DET. We train the model for 140 epochs, with the initial learning rate of $5e-4$ which drops 10x at 90 and 120 epochs respectively. The top predictions T is set as 100. Our experiments are conducted by Pytorch on a single GPU of NVIDIA Tesla P100.

4.2 Overall Performance

We compare our method with state-of-the-arts in this subsection. Meanwhile, we strip all components related to the Relation-aware Frame structure, IIM and CPM as baseline.

Performance on V-COCO. Comparison results in terms of mAP_{role} are shown in Table 1. It can be seen that our proposed RR-Net has a mAP(%) of 54.2, obtaining the best performance among all methods. Although we do not adopt previous region-based feature learning (e.g., RPNN [Zhou and Chi, 2019], Contextual Att [Wang *et al.*, 2019]), or employ human pose (e.g., PMFNet [Wan *et al.*, 2019], TIN [Li *et al.*, 2019]), our method outperforms these methods with sizable gains. Besides, our method achieves an absolute gain of 2.8 points, which is a relative improvement of 5.5% compared with the baseline, validating its efficacy in HOI detection task.

Performance on HICO-DET. Table 2 shows the comparison results on HICO-DET. Firstly, the detection results of our RR-Net are the best among all methods under the Full and Non-Rare settings, demonstrating that our method is more competitive than the others in detecting most common HOIs. It is noted that RR-Net is not preminent in detecting rare

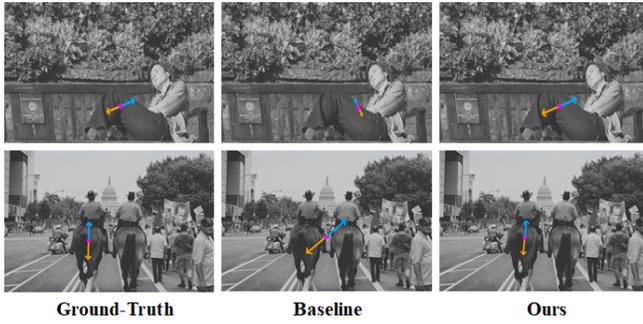


Figure 5: Visualized effects of IIM and CPM. The first column shows examples of ground-truth annotation. The second column visualizes interaction point and displacements predictions before point grouping using baseline method. The third column represents above predictions with IIM and CPM. Best viewed in color.

HOIs (HOI categories with less than 10 training instances), because our relation reasoning modules need enough training samples to exert their effects. Besides, our RR-Net obtains 20.72 mAP on HICO-DET (Default Full), which achieves relative gain of 9.8% compared with the baseline. These results quantitatively show the efficacy of our method.

4.3 Ablation Studies

Performance Impact of Proposed Components. V-COCO [Yatskar *et al.*, 2016] serves as the testbed in this subsection. We first study the performance impact of our proposed progressive Relation-aware Frame, Interaction Intensifier Module (IIM) and Correlation Parsing Module (CPM). As shown in Table 3, when we adopt the Relation-aware Frame Part A or Part B (See Figure 3) severally, the mAP are 52.4 and 52.2, indicating the respective effects of the two parts. Adopting the fully Relation-aware Frame obtains the mAP of 52.8, which boosts the performance by 1.4 points compared with the baseline. Meanwhile, we investigate the impact of employing IIM and CPM in Table 4. Compared with the baseline, employing IIM (upon Relation-aware Frame Part A) and CPM (upon Relation-aware Frame Part B) brings gains of 1.5 mAP and 1.4 mAP, respectively. Furthermore, improvements of 2.2 mAP and 2.0 mAP can be obtained while employing IIM and CPM respectively upon our fully Relation-aware Frame. Finally, RR-Net achieves an absolute gain of 2.8 mAP compared with the baseline. These evidences show that each proposed component of our region-independent relation reasoning indeed contributes to the final performance.

Effects visualization of IIM and CPM. Moreover, we take several examples of interaction point and displacements predictions before point grouping in Figure 5 to visualize the effects of IIM and CPM. Intuitively, with IIM and CPM, our framework makes better predictions of interaction point location and displacements compared with baseline. It is worth noting that due to the inaccurate prediction of interaction point location and displacements, the baseline framework in the second row may lead to miss-pairing between the human and horse. However, our framework provides predictions similar to the ground-truth annotation, which helps produce



Figure 6: Visualization of HOI detections. The first two rows show our results compared with baseline, where texts below indicate the detected interaction (verb), and two numbers in turn represent scores predicted by baseline and our method. The third and fourth rows show multiple people take interactions with various objects concurrently detected by our method. Best viewed in color.

correct HOI detection in point grouping procedure.

4.4 Qualitative Examples

In Figure 6, we first compare our results with baseline in the first two rows to demonstrate our improvements. Each subplot displays one detected $\langle \text{human, verb, object} \rangle$ triplet for easy observation. We can see that our method is capable of detecting various HOIs with higher scores. In addition, the third and fourth rows show that our method is able to detect multiple people taking different interactions with diversified objects. Summarizing all above results validates that our method can learn HOI-specific semantics well from complex environments to provide high quality HOI detections.

5 Conclusion

Recent end-to-end HOI detectors are limited in modelling HOIs due to they are incapable to perform relation reasoning in their methods. Besides, previous region-based HOI feature learning is hardly applicable to end-to-end detectors. Therefore, we make a first try to develop a RR-Net to explore novel techniques of region-independent relation reasoning, which opens up a new direction to improve HOI detection. Specifically, a Relation-aware Frame is proposed, upon which an Interaction Intensifier Module and a Correlation Parsing Module are designed successively to reason HOI-specific interactive semantics without instance regions. Extensive experiments show that our RR-Net outperforms existing methods with sizable gaps, validating its efficacy in detecting HOIs.

Acknowledgments

Thanks to IER foundation (No. HT-JD-CXY-201904).

References

- [Chao *et al.*, 2018] Yuwei Chao, Yunfan Liu, Xieyang Liu, Huayi Zeng, and Jia Deng. Learning to detect human-object interactions. *workshop on applications of computer vision*, pages 381–389, 2018.
- [Fang *et al.*, 2017] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. Rmpe: Regional multi-person pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2334–2343, 2017.
- [Fang *et al.*, 2018] Hao-Shu Fang, Jinkun Cao, Yu-Wing Tai, and Cewu Lu. Pairwise body-part attention for recognizing human-object interactions. In *Proceedings of the European Conference on Computer Vision*, pages 51–67, 2018.
- [Gao *et al.*, 2018] Chen Gao, Yuliang Zou, and Jiabin Huang. ican: Instance-centric attention network for human-object interaction detection. *british machine vision conference*, page 41, 2018.
- [Gkioxari *et al.*, 2018] Georgia Gkioxari, Ross B Girshick, Piotr Dollar, and Kaiming He. Detecting and recognizing human-object interactions. *computer vision and pattern recognition*, pages 8359–8367, 2018.
- [Gupta *et al.*, 2019] Tanmay Gupta, Alexander Schwing, and Derek Hoiem. No-frills human-object interaction detection: Factorization, layout encodings, and training techniques. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9677–9685, 2019.
- [Law and Deng, 2018] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision*, pages 734–750, 2018.
- [Li *et al.*, 2019] Yong-Lu Li, Siyuan Zhou, Xijie Huang, Liang Xu, Ze Ma, Hao-Shu Fang, Yanfeng Wang, and Cewu Lu. Transferable interactiveness knowledge for human-object interaction detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3585–3594, 2019.
- [Liao *et al.*, 2020] Yue Liao, Si Liu, Fei Wang, Yanjie Chen, Chen Qian, and Jiashi Feng. Ppdm: Parallel point detection and matching for real-time human-object interaction detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 482–490, 2020.
- [Qi *et al.*, 2018] Siyuan Qi, Wenguan Wang, Baoxiong Jia, Jianbing Shen, and Songchun Zhu. Learning human-object interactions by graph parsing neural networks. *european conference on computer vision*, pages 407–423, 2018.
- [Ren *et al.*, 2017] Shaoqing Ren, Kaiming He, Ross B Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1137–1149, 2017.
- [Sharma *et al.*, 2015] Shikhar Sharma, Ryan Kiros, and Ruslan Salakhutdinov. Action recognition using visual attention. *arXiv preprint arXiv:1511.04119*, 2015.
- [Ulutan *et al.*, 2020] Oytun Ulutan, A S M Iftekhar, and B. S. Manjunath. Vsgnet: Spatial attention network for detecting human object interactions using graph convolutions. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13617–13626, 2020.
- [Wan *et al.*, 2019] Bo Wan, Desen Zhou, Yongfei Liu, Rongjie Li, and Xuming He. Pose-aware multi-level feature network for human object interaction detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9469–9478, 2019.
- [Wang *et al.*, 2019] Tiancai Wang, Rao Muhammad Anwer, Muhammad Haris Khan, Fahad Shahbaz Khan, Yanwei Pang, Ling Shao, and Jorma Laaksonen. Deep contextual attention for human-object interaction detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5694–5702, 2019.
- [Wang *et al.*, 2020] Tiancai Wang, Tong Yang, Martin Danelljan, Fahad Shahbaz Khan, Xiangyu Zhang, and Jian Sun. Learning human-object interaction detection using interaction points. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4116–4125, 2020.
- [Xu *et al.*, 2019] Bingjie Xu, Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S Kankanhalli. Interact as you intend: Intention-driven human-object interaction detection. *IEEE Transactions on Multimedia*, 22(6):1423–1432, 2019.
- [Yang and Zou, 2020] Dongming Yang and Yuexian Zou. A graph-based interactive reasoning for human-object interaction detection. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 1111–1117, 2020.
- [Yang *et al.*, 2019] Dongming Yang, YueXian Zou, Jian Zhang, and Ge Li. C-rpns: Promoting object detection in real world via a cascade structure of region proposal networks. *Neurocomputing*, 367:20 – 30, 2019.
- [Yang *et al.*, 2020] Dongming Yang, YueXian Zou, Jian Zhang, and Ge Li. Gid-net: Detecting human-object interaction with global and instance dependency. *Neurocomputing*, pages 366–377, 2020.
- [Yatskar *et al.*, 2016] Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. Situation recognition: Visual semantic role labeling for image understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5534–5542, 2016.
- [Yu *et al.*, 2018] Fisher Yu, Dequan Wang, Evan Shelhamer, and Trevor Darrell. Deep layer aggregation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2403–2412, 2018.
- [Zhou and Chi, 2019] Penghao Zhou and Mingmin Chi. Relation parsing neural network for human-object interaction detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 843–851, 2019.
- [Zhou *et al.*, 2019] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. In *arXiv preprint arXiv:1904.07850*, 2019.