

Exploring Self- and Cross-Triplet Correlations for Human-Object Interaction Detection

Weibo Jiang¹, Weihong Ren^{1*}, Jiandong Tian², Liangqiong Qu³, Zhiyong Wang¹, Honghai Liu¹

¹ State Key Laboratory of Robotics and System, School of Mechanical Engineering and Automation, Harbin Institute of Technology, Shenzhen

² State Key Laboratory of Robotics, Shenyang Institute of Automation, Chinese Academy of Science

³ Department of Statistics and Actuarial Science, The University of Hong Kong

jiangweibo@stu.hit.edu.cn, {renweihong, wangzhiyong, honghai.liu}@hit.edu.cn, tianjd@sia.cn, liangqqu@hku.hk

Abstract

Human-Object Interaction (HOI) detection plays a vital role in scene understanding, which aims to predict the HOI triplet in the form of $\langle \text{human, object, action} \rangle$. Existing methods mainly extract multi-modal features (e.g., appearance, object semantics, human pose) and then fuse them together to directly predict HOI triplets. However, most of these methods focus on seeking for self-triplet aggregation, but ignore the potential cross-triplet dependencies, resulting in ambiguity of action prediction. In this work, we propose to explore Self- and Cross-Triplet Correlations (SCTC) for HOI detection. Specifically, we regard each triplet proposal as a graph where Human, Object represent nodes and Action indicates edge, to aggregate self-triplet correlation. Also, we try to explore cross-triplet dependencies by jointly considering instance-level, semantic-level, and layout-level relations. Besides, we leverage the CLIP model to assist our SCTC obtain interaction-aware feature by knowledge distillation, which provides useful action clues for HOI detection. Extensive experiments on HICO-DET and V-COCO datasets verify the effectiveness of our proposed SCTC.

Introduction

Human-object interaction detection plays a core role in human-centric behavior analysis and provides deeply understanding of human intentions (Chao et al. 2015; Gupta and Malik 2015). HOI detection approaches need to extract the semantic relationships between humans and objects and finally predict a set of $\langle \text{human, object, action} \rangle$ triplets within an image. The challenges in the HOI detection task mainly lie in the modeling the relationships between human and object and leveraging these relationships to perform human-object paring as well as predict their interactions.

The HOI detection approaches can be roughly divided into two categories: one-stage detectors and two-stage ones. One-stage methods implicitly build the relationship between human and object, and try to simultaneously perform object detection and interaction prediction. E.g., previous CNN-based work (Wang et al. 2020b) localizes and classifies the interaction by detecting additional interaction point, while

*Corresponding author

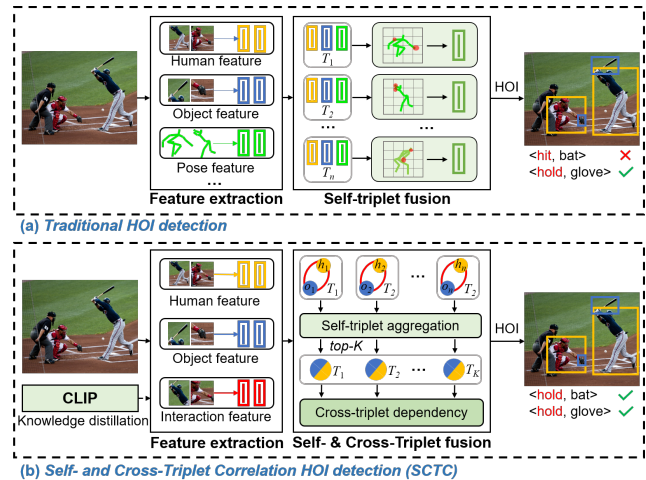


Figure 1: Architecture comparison of different HOI detection pipelines. (a) Traditional HOI detection methods focus on exploring multiple features and fuse them by self-triplet fusion (e.g. body-part attention); (b) our SCTC jointly explores Self- and Cross-triplet Correlations, aiming to eliminate action ambiguity and promote scene understanding.

other method (Kim et al. 2020a) directly captures the interaction region using a union-level box. Though the CNN-based methods can effectively model human-object relationships, they usually need post-grouping strategies to form a complete triplet. Recently, to capture long range contexts, the transformer-based works (Tamura, Ohashi, and Yoshinaga 2021; Zou et al. 2021; Dong et al. 2021; Kim et al. 2022; Liao et al. 2022; Zhou et al. 2022; Kim, Jung, and Cho 2023) have greatly advanced the HOI detection using self-attention and cross-attention mechanisms. They directly predict a HOI triplet without extracting instance-level priors or exploring their dependencies, but this end-to-end way may suffer insufficient exchange between contextual clues, potentially leading to a sub-optimal solution.

The two-stage methods firstly perform object detection, and then explicitly build the relationships between human and object. As shown in Fig. 1 (a), some works (Xu et al. 2019; Kim et al. 2020b) extract multi-modal features based

on object detection, and then fuse them through the Multi-Layer Perceptron (MLP) for final action prediction. However, such simple fusion fails to delve deep correlations between different instance-level features. The recent works (Li et al. 2022b; Park, Park, and Lee 2023) jointly use instance level and body part features to mine interactiveness in HOI. However, they only focus on seeking for the self-triplet attentions, but ignore the potential cross-triplet dependencies, which may lead to ambiguous action prediction. In contrast, as shown in Fig. 1(b), we propose to explore Self- and Cross-Triplet Correlations (SCTC) for HOI detection. Using object detections, we first extract instance-level and interaction-level features simultaneously. Then, we regard each candidate triplet as a graph (“human” and “object” are nodes, “interaction” is edge) to aggregate self-triplet features. After that, we model a triplet-level graph where each node represents a triplet proposal, to explore cross-triplet dependencies. Here, the instance-level, semantic-level, and layout-level relations are used to formulate adjacency matrix, thereby eliminating ambiguous action prediction.

Interaction features play a vital role in self-triplet aggregation, and they are utilized to associate different humans and objects, which helps to predict the action of a human-object pair. To enhance this representation, we leverage the CLIP model (Radford et al. 2021) to transfer its vision-language knowledge to the interaction feature. This is different from other methods (Liao et al. 2022; Yuan et al. 2022a; Ning et al. 2023; Wan et al. 2023), which adopt CLIP to distill the HOI classification in the final stage for few-shot detection. Experimental results on public datasets show the superiority of our proposed model. Also, the ablation studies prove the effectiveness of the Self- and Cross-Correlation mechanism.

To summarize, our contributions are three-fold:

- In order to obtain self-triplet aggregation, we formulate a graph for each triplet where human and object are represented as nodes, and the interaction prior is innovatively employed as the graph edge that models relationships between different instances.
- To build the contextual dependencies across different candidate triplets, we also create a triplet-level graph, and utilize the instance-level, semantic-level, and layout-level relations to formulate adjacency matrix, thereby eliminating ambiguous action prediction.
- We adopt CLIP model to transfer vision-language knowledge to the interaction feature, which serves as the key element in exploring the semantic relationships from different instances.

Related Works

One-Stage HOI Detection Methods

The one-stage HOI detection methods predict the HOI triplet directly, and build the dependencies between human and object using geometric or structural constrains. Previous CNN-based works (Kim et al. 2020a; Liao et al. 2020; Wang et al. 2020b; Zhong et al. 2021) predict the triplet using two parallel branches: in the first branch, the image feature map is used to find the interaction point or the interaction area and

accordingly the interaction. In the second branch, object detection is performed to find the human-object pair. Finally, the two branches are used to regress the offsets and match the interaction class with the detected instances. These methods capture the contextual information by pairing the human and object from an early stage in feature extraction. However, when the interacting human and object are far away from each other, the limited global modeling capability of CNNs may lead to ambiguity in the semantic features.

Recent transformer-based works (Zou et al. 2021; Tamura, Ohashi, and Yoshinaga 2021; Kim et al. 2021; Zhou et al. 2022; Liao et al. 2022; Lim et al. 2023) leverage the encoder-decoder architecture to jointly represent human, object, and interaction features, and build their relationships implicitly. To learn the representations that focus on different feature regions, Kim *et al.* (2023) propose three decoder branches to represent human, object, and interaction, respectively. The relationships are built through large amount of self-attention and cross-attention modules, leading to the increase of computation cost. These methods made great exploration of applying transformer to HOI detection, however, they only build associations of HOI in instance level, and also suffer from the low training speed and large memory usage (Park, Park, and Lee 2023).

Two-Stage HOI Detection Methods

The two-stage HOI detection approaches initially detect all possible interacting pairs, and then extract the corresponding features (e.g., appearance, spatial, pose, semantics). Finally, these features are fused to predict the interactiveness and the action for each pair. Chao *et al.* (2018) propose the multi-stream framework to obtain contextual information for interaction classification, but it only considers the spatial relations. The subsequent works encode a variety of additional priors for robust representation, including pose features (Kim et al. 2020b; Li et al. 2022b; Park, Park, and Lee 2023) and semantic features (Bansal et al. 2020; Liu, Li, and Lu 2022; Zhang et al. 2022b). E.g., Zhang *et al.* (2022b) and Kim *et al.* (2020b) first concatenate multi-modal features and then use the MLP to encode them for interaction classification, but such simple fusion strategy fails to mine the relationships between these features. Li *et al.* (2022b) explore the interactiveness using human pose, and build the dependencies between human and object based on the human-part attentions. However, it only considers the self-triplet relationships, which may lead to action ambiguity.

Recent studies (Gao et al. 2020; Zhou et al. 2021; Zhang, Campbell, and Gould 2021b; Park, Park, and Lee 2023) leverage graph network to explore relationships between human and object. E.g., Park *et al.* (2023) employ human joints to focus on specific human regions, and then use a graph network to aggregate local features within each human-object pair. The above graph-assisted methods have made contributions to explore the potential relationships between human and object. However, they only model the instance-level relations within a HOI triplet, and don’t consider the dependencies across different triplets. Thus, in this work, we propose to explore self- and cross-triplet correlations in HOI detection, and encode the interaction feature as the funda-

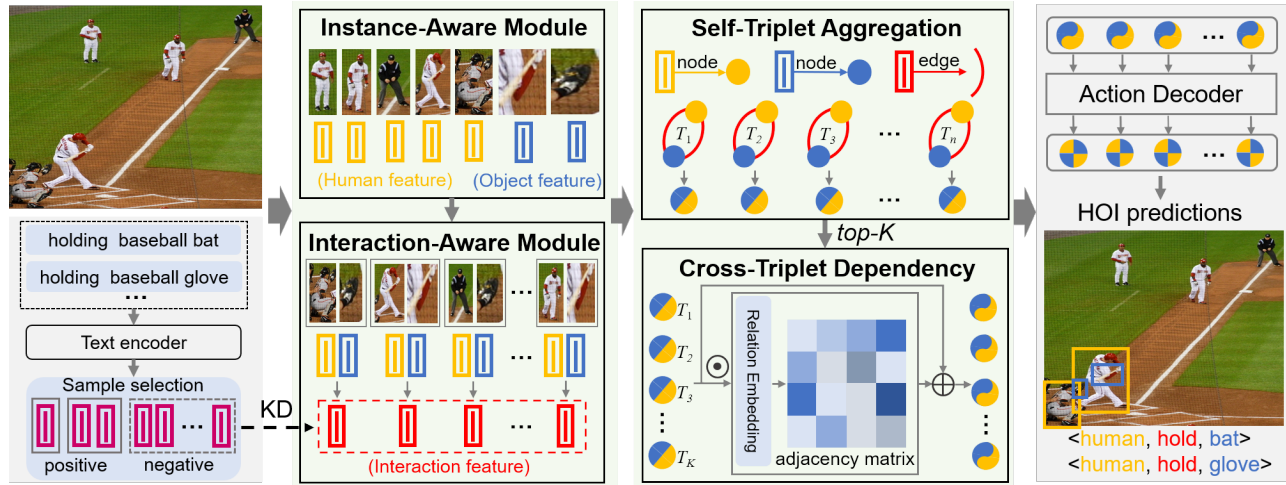


Figure 2: The overall pipeline of SCTC. For an input image, the instance-aware module is first used to extract instance-level features (appearance and semantics). Then, the interaction-aware module match human-object pairs and fuse them together to generate interaction feature. Next, Self-Triplet Aggregation (STA) is employed to explore self-triplet attentions, while the Cross-Triplet Dependency (CTD) is used to build connections across different HOI triplets. Finally, the action decoder is utilized to predict HOI triplets. Also, Knowledge Distillation (KD) is used to transfer text embeddings from CLIP to interaction feature.

mental element for the reasoning process, which is a new attempt in HOI relationship exploration.

Exploiting Vision-Language Models

Recently, the CLIP model (Radford et al. 2021) has demonstrated strong generalization to various downstream tasks (Gu et al. 2022; Li et al. 2022a; Esmaeilpour et al. 2022), and several studies have attempted to leverage the knowledge from CLIP to HOI detection. E.g., Dong et al. (2022), Yuan et al. (2022a), Wan et al. (2023) and Zhang et al. (2023) integrate text embeddings generated from the CLIP to enhance the representation of semantic features, and have achieved promising results. In contrast, Liao et al. (2022) leverage the CLIP knowledge for interaction classification and visual feature distillation, but such a kind of distillation may introduce ambiguity when multiple HOI triplets exist in the scenario. Ning et al. (2023) also try to transfer visual-linguistic knowledge to HOI detection, utilizing the prior knowledge in CLIP to generate a classifier by encoding HOI descriptions. Though this method can effectively integrate CLIP knowledge, it still requires the participation of CLIP in the inference stage, which results in a high computation cost. Different from the methods above, we leverage the CLIP model to distill the interaction feature existing in each <human, object, action> triplet, which is a fine-grained knowledge distillation method.

Method

We propose the Self- and Cross-Triplet Correlation HOI model to explore relationships among humans, objects, and interactions for better understanding the human-object pairing and the following action prediction. As shown in Fig. 2, we first leverage Instance-Aware Module to detect all the objects and obtain the instance-level features. Then, we employ

Interaction-Aware Module to perform human-object pairing and project the fused features into interaction feature. After that, the Self-Triplet Aggregation (STA) is adopted to explore self-triplet attentions, while the following Cross-Triplet Dependency (CTD) is used to build connections across different HOI triplets. Finally, the actions for human-object pairs are predicted through the action decoder.

Instance-Aware Module

Given an input image, we first leverage the DETR (Carion et al. 2020), which is composed of ResNet (He et al. 2016) and the following Transformer (Vaswani et al. 2017), to detect all the humans and objects. Then, we extract the features of them from two aspects: appearance and semantics. The appearance feature is directly derived from the output of DETR decoder, which is highly related with the detected object. The semantic feature is obtained by embedding the object category into a feature vector. Finally, both the human feature \mathbf{F}_h and the object feature \mathbf{F}_o can be represented as the concatenation of appearance and semantic feature.

Interaction-Aware Module

The interaction-aware module aims to learn the interaction feature from the instance-level features (\mathbf{F}_h and \mathbf{F}_o). Specifically, we pair all the detected humans and objects to construct candidate human-object pairs. The hard mining strategy (Shrivastava, Gupta, and Girshick 2016) is employed to sample negative pairs to facilitate the learning of interactivity prediction following Zhang et al. (2022b). Then, we extract the spatial feature of each candidate pair, which can be denoted as $[\Delta x, \Delta y, \Delta s, \arctan(\frac{\Delta y}{\Delta x}), A_h, A_o, A_I, A_U]$, where Δx , Δy , Δs denote the distance between human and object in the x -axis, y -axis, and the coordinate system, respectively; A_h and A_o denote the area of human and object,

respectively; A_I and A_U denote the intersection area and union area of the human-object pair, respectively. By concatenating the human feature \mathbf{F}_h , object feature \mathbf{F}_o , and the spatial feature together, we obtain the pair-level feature \mathbf{F}_{ho} , which contains the appearance, semantic, and spatial features of a human-object pair. Finally, the interaction feature \mathbf{F} of each pair is obtained by projecting its corresponding \mathbf{F}_{ho} with the MLP.

Visual-Linguistic Knowledge Transfer

To enhance the interaction feature, we transfer the visual-linguistic knowledge from CLIP (Radford et al. 2021) to our interaction-aware module. Specifically, we first adopt the CLIP model to generate text embeddings for the ground truth HOIs. Next, we utilize the text embeddings to distill our interaction feature \mathbf{F} .

To generate the text embeddings, we first convert the HOI triplet labels into textual descriptions. Given an object and a verb, we use the template of ‘‘A photo of a person [verb-ing] a/an [object]’’ to generate a sentence. Then we generate the text embedding \mathbf{E} for each candidate human-object pair through the pre-trained CLIP text encoder offline.

After that, we leverage the text embeddings to distill our interaction feature. For positive human-object pairs that involve a single interaction, we directly take their corresponding text embeddings as supervision, which promotes the interaction feature learning from the text encoder of CLIP. While for other positive pairs with multiple interactions, we calculate the average representation of the text embeddings for each pair to minimize semantic ambiguity. Besides, for the negative pairs, we take the average representation of all the text embeddings, which helps to push interaction feature away from non-interactive cluster. Formally, we use the \mathcal{L}_1 loss to supervise the distance between the interaction feature \mathbf{F} and the text embedding \mathbf{E} :

$$\mathcal{L}_{kd} = \frac{1}{N_p} \sum_i |\mathbf{E}_i - \mathbf{F}_i|, \quad (1)$$

where i denotes the i -th candidate human-object pair, and N_p is the number of candidate human-object pairs.

Self-Triplet Aggregation

Previous works explore self-triplet relations solely on the basis of instance-level features, but ignore the interaction feature, which is the fundamental element of a HOI triplet. In our Self-Triplet Aggregation (STA), we propose to aggregate features from both instance level and interaction level. Specifically, we design a graph neural network (Kipf and Welling 2017) for each candidate triplet, where the human and object are viewed as nodes and the interaction feature is treated as edge. The human node ν_h and object node ν_o are directly derived from the human feature \mathbf{F}_h and object feature \mathbf{F}_o , respectively. Each node (ν_h or ν_o) contains the appearance feature and semantic feature for a specific instance, which only focuses on the instance itself. To aggregate features from ν_h and ν_o , we initialize the edge ε as the interaction feature \mathbf{F} . The edge contains the interaction semantics, spatial feature, and instance features, forging a robust connection between human and object in a candidate HOI pair.

We aim to build correlation between human and object via the interaction-aware edge, empowering STA to capture both individual attributes and the potential relationship between them. For the positive HOI triplets, the interaction feature distilled by CLIP tends to encompass specific action semantics, steering the human node ν_h and object node ν_o to acquire action clues through the edge ε . Meanwhile, for the negative triplets, the interaction feature learns the semantics of non-interactiveness, and the implicit information that the human and object is not a interactive pair is transmitted to the corresponding nodes.

The human node and the object node are updated with the message propagation from the each other and the edge:

$$\nu_{h \rightarrow o} = f_i(\varepsilon) \odot f_h(\nu_h), \quad (2)$$

$$\nu_{o \rightarrow h} = f_i(\varepsilon) \odot f_o(\nu_o), \quad (3)$$

where f_i , f_h , f_o are projection functions consist of a linear function and a ReLU activation function, and \odot is the Hadamard product.

The updated nodes can be obtained by:

$$\hat{\nu}_o = \nu_{h \rightarrow o} + \nu_o, \quad (4)$$

$$\hat{\nu}_h = \nu_{o \rightarrow h} + \nu_h. \quad (5)$$

Here, we use skip connections to enable the graph network to learn the additional correlations between human node and object node. After message propagating, each node has a comprehensive understanding of its potential interactions.

Afterwards, we concatenate the updated human and object nodes together (denoted as $\nu_{h,oi}$) to calculate the interactiveness scores followed by two-layer linear projections. We select the top-K human-object proposals according to the interactiveness scores, and the corresponding features are also abstracted. We optimize the human-object matching process by focal loss (FL) (Lin et al. 2017):

$$\mathcal{L}_{pair} = \frac{1}{\sum_{k=1}^K p_k} \sum_{k=1}^K \mathbf{FL}(\hat{p}_k, p_k), \quad (6)$$

where $p_k \in \{0, 1\}$ indicates whether an interaction exists in ground truth and \hat{p}_k is the predicted interactiveness score.

Cross-Triplet Dependency

The current two-stage works focus on exploring relationships in the self-triplet level, but ignore the implicit dependencies in the cross-triplet level. Actually, the relations between humans and objects are often context-dependent and can affect the reasoning of individual interactions. Thus, we create the Cross-Triplet Dependency (CTD), a triplet-level graph to exploit the correlations across different triplets. Different from STA, each node in CTD denotes a HOI triplet proposal and is initialized by the $\nu_{h,oi}$ obtained in STA. Here, we mainly introduce the formulation of the edges, i.e., the adjacency matrix.

As shown in Fig. 3, we embed the instance relation, semantic relation, and layout relation into an adjacency matrix to link all the nodes. The instance relation means whether different HOI triplets share the same human or object. We create a matrix $\mathbf{M}_{ins} \in \mathbb{R}^{K \times K \times 2}$ to represent the instance

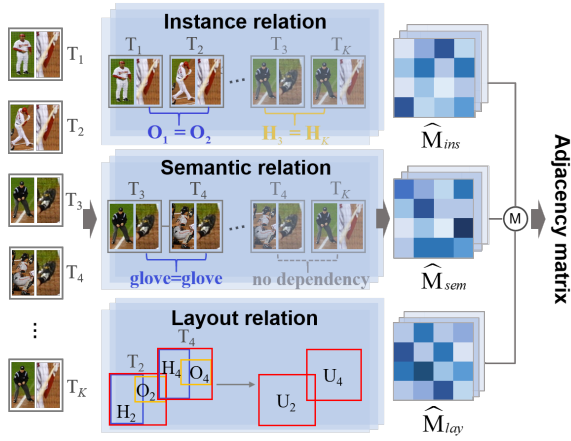


Figure 3: Generation of the edges in CTD, where the \textcircled{M} indicates the MLP.

relation, where K means the selected K triplets. If the j -th triplet ν_{hoi}^j and the k -th triplet ν_{hoi}^k share the same human or object, the corresponding value $M_{ins}^{(j,k,0)}$ or $M_{ins}^{(j,k,1)}$ is set to 1 (channel “0” and “1” represent human and object, respectively). Otherwise, the value is set to 0. If there exists the instance-level connection between two different triplets, there maybe exists the interaction relationship, which is helpful for the HOI detection with multiple actions. For example, if we find a triplet of $\langle \text{human, hold, fork} \rangle$, it is high likely that the mentioned human is also associated with another triplet of $\langle \text{human, hold, knife} \rangle$. The instance relation aims to exploit such kind of implied relations, and its embedding $\widehat{M}_{ins} \in \mathbb{R}^{K \times K \times d_{ins}}$ ($d_{ins} = 64$) is obtained by the linear embedding on M_{ins} .

The instance relation only considers whether the instances in different triplets are identical, neglecting the dependencies among different triplets with the same object category, and may cause a lack of semantic information. Thus, the semantic relation is introduced to measure whether different HOI triplets share the same object semantics. We design a matrix $M_{sem} \in \mathbb{R}^{K \times K \times C_o}$ to represent the semantic relation, where C_o denotes the number of object categories. If ν_{hoi}^j and ν_{hoi}^k share the same object category c , the corresponding value $M_{sem}^{(j,k,c)}$ is set to 1. Otherwise, the value is set to 0. For example, if we find a triplet of $\langle \text{human, ride, motorbike} \rangle$, it is likely that another triplet with the same interaction also exists in the scenario but hard to be detected. The semantic relation aims to inject the semantic information to the adjacency matrix and find the hard positive samples, and its embedding $\widehat{M}_{sem} \in \mathbb{R}^{K \times K \times d_{sem}}$ ($d_{sem} = 256$) is obtained by the linear embedding on M_{sem} .

The layout relation aims to build spatial relationships between two different triplets. We first calculate the union box for each triplet, and then extract the spatial feature between the two union boxes, which can be represented as $[\Delta x, \Delta y, \Delta s, \arctan(\frac{\Delta y}{\Delta x}), A_h, A_o, A_I, A_U]$ following the interaction-aware module. Thus, the layout relation $M_{lay} \in$

$\mathbb{R}^{K \times K \times 8}$ is formulated by calculating the spatial feature between every two triplets. The corresponding embedding $\widehat{M}_{lay} \in \mathbb{R}^{K \times K \times d_{lay}}$ ($d_{lay} = 64$) is also obtained by the linear embedding on M_{lay} .

After all the relations are built and embedded, the final adjacency matrix is derived by:

$$M_{adj} = \text{MLP}(\text{Concat}(\widehat{M}_{ins}, \widehat{M}_{sem}, \widehat{M}_{lay})), \quad (7)$$

where Concat is the concatenation, and $M_{adj} \in \mathbb{R}^{K \times K}$. Then, we update the HOI triplet nodes by:

$$\widehat{\nu}_{hoi} = (M_{adj} \cdot \nu_{hoi}) \oplus \nu_{hoi} \quad (8)$$

where \cdot is the matrix multiplication, and \oplus is the element-wise addition. Similar to the STA, we still adopt the skip connections to enable the CTD to learn additional cross-triplet dependencies.

Human-Object Interaction Prediction

As shown in Fig. 2, we predict the final actions for each human-object pair using the action decoder. The decoder query is initialized by projecting the $\widehat{\nu}_{hoi}$ with the MLP, and the key and value come from the features extracted by the ResNet50 in DETR. The final actions \widehat{y} are obtained by linear projection on the output of action decoder, and the classification is optimized by focal loss:

$$\mathcal{L}_a = \frac{1}{\sum_{k=1}^K \sum_{c=1}^C \mathbf{y}_{k,c}} \sum_{k=1}^K \sum_{c=1}^C \text{FL}(\widehat{\mathbf{y}}_{k,c}, \mathbf{y}_{k,c}), \quad (9)$$

where $\mathbf{y}_{k,c} \in \{0, 1\}$ indicates whether the k^{th} proposal contains the c^{th} action class, and $\widehat{\mathbf{y}}_{k,c}$ is the predicted probability of c^{th} action occurs in the k^{th} proposal.

Finally, the HOI predictions are achieved by combining the human-object pairs and the predicted actions. The overall objective of our SCTC contains the knowledge distillation in (1), the human-object proposal prediction in (6), and the action classification in (9):

$$\mathcal{L}_{SCTC} = \alpha \cdot \mathcal{L}_{kd} + \beta \cdot \mathcal{L}_{pair} + \gamma \cdot \mathcal{L}_a, \quad (10)$$

where α, β, γ are hyper-parameters.

Experiment

Experimental Setting

Datasets. HICO-DET (Chao et al. 2015) consists of 38,118 images in training set and 9,658 in test set, and has 600 HOI categories over 117 actions and 80 objects. The 600 HOI categories are split into 138 Rare and 462 Non-Rare based on the number of instances. V-COCO (Gupta and Malik 2015) consists of 5,400 images in the trainval set and 4,946 images in the test set. It has 259 HOI categories over 29 actions and 80 objects.

Evaluation metric. Following the standard metric, we use the commonly used mean Average Precision (mAP) to evaluate the model performance for the two benchmarks. A HOI triplet is considered as true positive if it localizes the human and object accurately (i.e., the Interaction-over-Union (IOU) between the predicted bounding boxes and ground truth is greater than 0.5) and also predicts the action correctly.

Method	Backbone	Full / Rare / Non-Rare
<i>One-stage methods</i>		
UnionDet (Kim et al. 2020a)	R50	14.25 / 10.23 / 15.46
IPNet (Wang et al. 2020b)	HG-104	19.56 / 12.79 / 21.58
HOTR (Kim et al. 2021)	R50	23.46 / 16.21 / 25.60
AS-Net (Chen et al. 2021)	R50	24.40 / 22.39 / 25.01
HOITrans (Zou et al. 2021)	R101	26.61 / 19.15 / 28.84
QPIC (Tamura et al. 2021)	R101	29.90 / 23.92 / 31.69
PhraseHOI (Li et al. 2022c)	R101	30.03 / 23.48 / 31.99
DT (Zhou et al. 2022)	R50	31.75 / 27.45 / 33.03
CATN (Dong et al. 2022)	R50	31.86 / 25.15 / 33.84
CDN (Zhang et al. 2021a)	R101	32.07 / 27.19 / 33.53
GEN-VLKT (Liao et al. 2022)	R50	33.75 / 29.25 / 35.10
HOICLIP (Ning et al. 2023)	R50	34.69 / 31.12 / 35.74
ParMap (Wu et al. 2022)	R50	35.15 / 33.71 / 35.58
ERNet (Lim et al. 2023)	ENetV2-L	35.92 / 30.12 / 38.29
<i>Two-stage methods</i>		
GPNN (Zhou et al. 2021)	R50	19.42 / 13.98 / 20.91
ACP (Kim et al. 2020b)	R152	20.59 / 15.92 / 21.98
OCN (Yuan et al. 2022b)	R101	31.43 / 25.80 / 33.11
STIP (Zhang et al. 2022b)	R50	32.22 / 28.15 / 33.43
UPT (Zhang et al. 2022a)	R101	32.62 / 28.62 / 33.81
OpenCat (Zheng et al. 2023)	R101	32.68 / 28.42 / 33.75
CQL (Xie et al. 2023)	R101	36.03 / 33.16 / 36.89
ViPLO (Park et al. 2023)	ViT-B/16	37.22 / 35.45 / 37.75
RmLR (Cao et al. 2023)	R101	37.41 / 28.81 / 39.97
SCTC (Ours)	R50	<u>37.92 / 34.78 / 38.86</u>
SCTC (Ours)	R101	39.12 / 36.09 / 39.87

Table 1: Performance comparison on HICO-DET dataset. The best result is marked with bold and the second best result is underlined. For fair comparison, we report results using an object detector finetuned on the training dataset of HICO-DET for the two-stage methods.

Implementation details. We take the DETR for object detection. All the learnable parameters in DETR are frozen during training. We leverage the pre-trained ViT-16/B CLIP on its official data to generate text embeddings for knowledge distillation, and also keep its parameters fixed during training. We set the number of action decoder layers to 6 and the number of queries to 32. The parameters α , β , δ are all set to 1.0. The entire model is trained on two Nvidia 3090 GPUs with a batch size 8. The AdamW (Loshchilov and Hutter 2017) optimizer is used for training with 30 epochs, where the starting learning rate is 5×10^{-5} , and then decays with the Cosine annealing training strategy.

Comparisons with the State-of-the-Arts

Tab. 1 shows the performance comparison of SCTC and other state-of-the-art methods on HICO-DET dataset. It can be observed that SCTC outperforms all the existing methods by a large margin. Specifically, SCTC achieves **39.12** mAP in the default full setting, obtaining a performance gain of 1.71 mAP (relatively 4.57%) compared to the most recent approach RmLR. Also, compared with ERNet, which is the state-of-the-art one-stage method, our model achieves a performance gain of 3.20 mAP (relatively 8.90 %). The results validate the superiority of our proposed SCTC. ACP and STIP directly concatenate the human and object fea-

Method	Backbone	$AP_{role}^{\#1}$ / $AP_{role}^{\#2}$
<i>One-stage methods</i>		
UnionDet (Kim et al. 2020a)	R50	47.5 / 56.2
IPNet (Wang et al. 2020b)	HG-104	51.0 / -
HOITrans (Zou et al. 2021)	R101	52.9 / -
AS-Net (Chen et al. 2021)	R50	53.9 / -
HOTR (Kim et al. 2021)	R50	55.2 / 64.4
PhraseHOI (Li et al. 2022c)	R50	57.4 / -
QPIC (Tamura et al. 2021)	R101	58.3 / 60.7
CATN (Dong et al. 2022)	R50	60.1 / -
GEN-VLKT (Liao et al. 2022)	R50	62.4 / 64.5
ParMap (Wu et al. 2022)	R50	63.0 / 65.1
HOICLIP (Ning et al. 2023)	R50	63.5 / 64.8
CDN (Zhang et al. 2021a)	R101	63.9 / 65.9
ERNet (Lim et al. 2023)	ENetV2-L	64.2 / -
DT (Zhou et al. 2022)	R50	66.2 / 68.5
<i>Two-stage methods</i>		
GPNN (Zhou et al. 2021)	R50	50.4 / -
ACP (Kim et al. 2020b)	R152	53.0 / -
UPT (Zhang et al. 2022a)	R101	61.3 / 67.1
OpenCat (Zheng et al. 2023)	R101	61.9 / 63.2
ViPLO (Park et al. 2023)	ViT-B/16	62.2 / 68.0
RmLR (Cao et al. 2023)	R101	64.2 / 70.2
OCN (Yuan et al. 2022b)	R101	65.3 / 67.1
STIP (Zhang et al. 2022b)	R50	66.0 / 70.7
CQL (Xie et al. 2023)	R101	66.8 / 69.8
SCTC (Ours)	R50	<u>67.1 / 71.7</u>
SCTC (Ours)	R101	68.2 / 72.5

Table 2: Performance comparison on V-COCO dataset.

tures together and use MLP to fuse them for action prediction, which can't well explore the complex relations between human and object, resulting in relatively worse results. ViPLO and ParMap explore the potential interactivity by leveraging body-part features, which guides the detection model to focus on crucial regions of human and object, and have achieved relatively better results. However, they only consider the relations in the self-triplet level, and ignore the dependencies among different triplets. For our SCTC, it utilizes the STA to aggregate features within each triplet, and then build correlations across different triplets by CTD, which can further promote HOI detection.

For V-COCO, as shown in Tab. 2, SCTC also performs the best among all the state-of-the-art methods, only using R50 as backbone. E.g., SCTC works better than the three recent HOI detectors ERNet, RmLR and CQL ($AP_{role}^{\#1}$ of 67.1 vs 64.2, 64.2 and 66.8). It is noted that $AP_{role}^{\#2}$ is significantly improved by the SCTC (about 1.9 higher than CQL), and the reason may be that our model can make fully use of scene semantics to infer missed or occluded information.

Ablation Study

In this subsection, we explore how Self-Triplet Aggregation (STA), Cross-Triplet Dependency (CTD), and the CLIP-based Knowledge Distillation (KD) affect the HOI detection performance using R50 as backbone on HICO-DET dataset. All training settings are consistent with those mentioned above.

KD	STA	CTD	MLP	Full	Rare	Non-Rare
-	-	-	✓	33.54	30.62	34.14
✓	-	-	✓	35.63	32.82	36.35
✓	✓	-	-	36.41	32.95	37.45
✓	-	✓	-	37.25	34.29	38.17
✓	✓	✓	-	37.92	34.78	38.86

Table 3: Effect of STA, CTD, and KD in our model.

IF	SF	LE	Full	Rare	Non-Rare
-	-	✓	35.42	32.24	36.31
-	✓	-	35.91	32.45	36.77
✓	-	-	37.76	34.83	38.71
✓	✓	-	37.92	34.78	38.86

Table 4: Effect of embedding content in STA. IF: interaction feature, SF: spatial feature, LE: learnable edge.

Self-Triplet Aggregation. SCTC adopts STA to aggregate features from human, object, and interaction. As shown in Tab. 3, STA with KD can improve MLP with KD by 0.78 mAP in Full category, validating that our proposed STA is an effective choice to aggregate triplet features. The nodes in the graph of STA are initialized with human and object features. Innovatively, we propose to leverage the interaction feature as the edge to connect the nodes. As shown in Tab. 4, we explore the effect of applying different features to serve as edges to connect the graph nodes. ‘Learnable Edge’ (LE) means that we directly initialize the edges with learnable parameters without any priors. It can be observed that the interaction feature (IF) exceeds the LE by 2.34 mAP in the Full category, which proves the effectiveness of IF to connect nodes. Besides, the spatial feature (SF) also contributes to HOI detection. E.g., IF+SF achieves the best results, improving IF from 37.76 to 37.92 in the Full category.

IR.	SR	LR	LE	Full	Rare	Non-Rare
-	-	-	✓	36.55	33.49	37.40
✓	-	-	-	37.14	34.38	38.04
-	✓	-	-	37.49	34.62	38.21
-	-	✓	-	37.26	34.12	38.22
✓	✓	✓	-	37.92	34.78	38.86

Table 5: Effect of adjacency matrix in CTD. IR: instance relation, SR: semantic relation, LR: layout relation, LE: learnable edge.

Cross-Triplet Dependency. SCTC adopts CTD to build correlations among different HOI triplets. From Tab. 3, KD+STA+CTD can further improve KD+STA by 1.51 mAP in Full category, which proves the importance to explore relationships across different triplets. The nodes in CTD are derived from the STA directly, representing the triplet-level feature. For the graph edges (i.e., the adjacency matrix), we initialize them by making full use of instance relation (IR),

semantic relation (SR), and layout relation (LR). In Tab. 5, we analyze the effect of different relations in the CTD. ‘Learnable edge’ (LE) also means that the edges are initialized with learnable parameters without any priors. It can be found that any one of the IR, SR, and LR works better than the LE by 0.59 mAP, 0.94 mAP, and 0.71 mAP in the Full category, respectively. Besides, when jointly considering IR, SR and LR, the HOI detection can be further improved.

CLIP	Bert	CE	Full	Rare	Non-Rare
-	-	-	35.95	32.57	37.01
✓	-	-	37.92	34.78	38.86
-	✓	-	35.78	32.25	36.69
-	-	✓	37.36	34.21	38.41
✓	-	✓	37.79	34.61	38.66

Table 6: Effect of Knowledge Distillation (KD) strategies.

Knowledge distillation. We transfer the knowledge from CLIP to our SCTC for the refinement of interaction feature. Specifically, we use the text encoder to generate text embeddings to distill our interaction feature. From Tab. 3, KD+MLP improves MLP by 2.09 mAP, which validates the significance of knowledge distillation to enhance interaction feature. Here, we explore three strategies to distill the interaction feature as shown in Tab. 6. ‘Bert’ means that we leverage Bert model (Devlin et al. 2019) to generate text embeddings, and ‘CE’ means that we directly use interaction feature to generate action classification predictions, supervised by ground truth action labels. It can be observed that CLIP achieves the better results than those of the Bert and CE. Bert model lacks the alignment of visual feature and linguistic feature, which results in the incompatibility between the interaction feature and its text embeddings. Besides, the interaction feature contains integrated information of human, object, and interaction, so it is inappropriate to distill it only with the action labels. This is also the reason why CLIP+CE performs worse than CLIP. Differently, CLIP model can fully encode the <human, object, action> triplet and its textual embedding is compatible with the visual feature, so it is suitable to distill the interaction feature.

Conclusion

In this paper, we propose the SCTC, the state-of-the-art two-stage HOI detector, to jointly explore Self- and Cross-Triplet Correlations for HOI detection. For the Self-Triplet Aggregation, SCTC creates a graph for each candidate triplet to jointly aggregate instance-level features from human, object and the potential action semantics. As for the Cross-triplet Dependency, it builds the contextual dependencies across different triplet proposals, where the adjacency matrix is formulated as the combination of instance relation, semantic relation and layout relation. Besides, we transfer the text knowledge from CLIP to SCTC, which helps to enhance the representation of interaction features. Experimental results on HICO-DET and VCOCO datasets verify the effectiveness of the proposed SCTC.

Acknowledgments

This work was supported in part by the National Key Research and Development Program of China under Grant 2022YFB4703201, in part by the National Natural Science Foundation of China under Grants 62206075, 61733011, and 62261160652, in part by the Guangdong Basic and Applied Basic Research Foundation under Grant 2021A1515110438, and in part by the Shenzhen Science and Technology Program under Grant RCBS20221008093220004.

References

- Bansal, A.; Rambhatla, S. S.; Shrivastava, A.; and Chellappa, R. 2020. Detecting human-object interactions via functional generalization. In *AAAI*, 10460–10469.
- Cao, Y.; Su, X.; Tang, Q.; Yang, F.; You, S.; Lu, X.; and Xu, C. 2023. Re-mine, Learn and Reason: Exploring the Cross-modal Semantic Correlations for Language-guided HOI detection. In *ICCV*.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *ECCV*, 213–229.
- Chao, Y.-W.; Liu, Y.; Liu, X.; Zeng, H.; and Deng, J. 2018. Learning to detect human-object interactions. In *WACV*, 381–389.
- Chao, Y.-W.; Wang, Z.; He, Y.; Wang, J.; and Deng, J. 2015. Hico: A benchmark for recognizing human-object interactions in images. In *ICCV*, 1017–1025.
- Chen, M.; Liao, Y.; Liu, S.; Chen, Z.; Wang, F.; and Qian, C. 2021. Reformulating hoi detection as adaptive set prediction. In *CVPR*, 9004–9013.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*.
- Dong, L.; Li, Z.; Xu, K.; Zhang, Z.; Yan, L.; Zhong, S.; and Zou, X. 2022. Category-aware transformer network for better human-object interaction detection. In *CVPR*, 19538–19547.
- Dong, Q.; Tu, Z.; Liao, H.; Zhang, Y.; Mahadevan, V.; and Soatto, S. 2021. Visual relationship detection using part-and-sum transformers with composite queries. In *ICCV*, 3550–3559.
- Esmailpour, S.; Liu, B.; Robertson, E.; and Shu, L. 2022. Zero-shot out-of-distribution detection based on the pre-trained model clip. In *AAAI*, 6568–6576.
- Gao, C.; Xu, J.; Zou, Y.; and Huang, J.-B. 2020. Drg: Dual relation graph for human-object interaction detection. In *ECCV*, 696–712.
- Gu, X.; Lin, T.-Y.; Kuo, W.; and Cui, Y. 2022. Open-vocabulary object detection via vision and language knowledge distillation. In *ICLR*.
- Gupta, S.; and Malik, J. 2015. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Kim, B.; Choi, T.; Kang, J.; and Kim, H. J. 2020a. Uniondet: Union-level detector towards real-time human-object interaction detection. In *ECCV*, 498–514.
- Kim, B.; Lee, J.; Kang, J.; Kim, E.-S.; and Kim, H. J. 2021. Hotr: End-to-end human-object interaction detection with transformers. In *CVPR*, 74–83.
- Kim, B.; Mun, J.; On, K.-W.; Shin, M.; Lee, J.; and Kim, E.-S. 2022. MSTR: Multi-Scale Transformer for End-to-End Human-Object Interaction Detection. In *CVPR*, 19578–19587.
- Kim, D.-J.; Sun, X.; Choi, J.; Lin, S.; and Kweon, I. S. 2020b. Detecting human-object interactions with action co-occurrence priors. In *ECCV*, 718–736.
- Kim, S.; Jung, D.; and Cho, M. 2023. Relational Context Learning for Human-Object Interaction Detection. In *CVPR*, 2925–2934.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Li, B.; Weinberger, K. Q.; Belongie, S.; Koltun, V.; and Ranftl, R. 2022a. Language-driven semantic segmentation. In *ICLR*.
- Li, Y.-L.; Liu, X.; Wu, X.; Huang, X.; Xu, L.; and Lu, C. 2022b. Transferable interactiveness knowledge for human-object interaction detection. *PAMI*, 44(7): 3870–3882.
- Li, Z.; Zou, C.; Zhao, Y.; Li, B.; and Zhong, S. 2022c. Improving human-object interaction detection via phrase learning and label composition. In *AAAI*, 1509–1517.
- Liao, Y.; Liu, S.; Wang, F.; Chen, Y.; Qian, C.; and Feng, J. 2020. Ppdm: Parallel point detection and matching for real-time human-object interaction detection. In *CVPR*, 482–490.
- Liao, Y.; Zhang, A.; Lu, M.; Wang, Y.; Li, X.; and Liu, S. 2022. GEN-VLKT: Simplify Association and Enhance Interaction Understanding for HOI Detection. In *CVPR*, 20123–20132.
- Lim, J.; Baskaran, V. M.; Lim, J. M.-Y.; Wong, K.; See, J.; and Tistarelli, M. 2023. ERNet: An Efficient and Reliable Human-Object Interaction Detection Network. *TIP*, 32: 964–979.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *ICCV*, 2980–2988.
- Liu, X.; Li, Y.-L.; and Lu, C. 2022. Highlighting object category immunity for the generalization of human-object interaction detection. In *AAAI*, 1819–1827.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Ning, S.; Qiu, L.; Liu, Y.; and He, X. 2023. HOICLIP: Efficient Knowledge Transfer for HOI Detection with Vision-Language Models. In *CVPR*, 23507–23517.
- Park, J.; Park, J.-W.; and Lee, J.-S. 2023. ViPLO: Vision Transformer based Pose-Conditioned Self-Loop Graph for Human-Object Interaction Detection. In *CVPR*, 17152–17162.

- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763.
- Shrivastava, A.; Gupta, A.; and Girshick, R. 2016. Training region-based object detectors with online hard example mining. In *CVPR*, 761–769.
- Tamura, M.; Ohashi, H.; and Yoshinaga, T. 2021. Qpic: Query-based pairwise human-object interaction detection with image-wide contextual information. In *CVPR*, 10410–10419.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *NeurIPS*, 5998–6008.
- Wan, B.; Liu, Y.; Zhou, D.; Tuytelaars, T.; and He, X. 2023. Weakly-supervised HOI Detection via Prior-guided Bi-level Representation Learning. In *ICLR*.
- Wang, T.; Yang, T.; Danelljan, M.; Khan, F. S.; Zhang, X.; and Sun, J. 2020b. Learning human-object interaction detection using interaction points. In *CVPR*, 4116–4125.
- Wu, X.; Li, Y.-L.; Liu, X.; Zhang, J.; Wu, Y.; and Lu, C. 2022. Mining cross-person cues for body-part interactivity learning in hoi detection. In *ECCV*, 121–136.
- Xie, C.; Zeng, F.; Hu, Y.; Liang, S.; and Wei, Y. 2023. Category Query Learning for Human-Object Interaction Classification. In *CVPR*, 15275–15284.
- Xu, B.; Wong, Y.; Li, J.; Zhao, Q.; and Kankanhalli, M. S. 2019. Learning to detect human-object interactions with knowledge. In *CVPR*, 2019–2028.
- Yuan, H.; Jiang, J.; Albanie, S.; Feng, T.; Huang, Z.; Ni, D.; and Tang, M. 2022a. RLIP: Relational Language-Image Pre-training for Human-Object Interaction Detection. *NeurIPS*.
- Yuan, H.; Wang, M.; Ni, D.; and Xu, L. 2022b. Detecting human-object interactions with object-guided cross-modal calibrated semantics. In *AAAI*, 3206–3214.
- Zhang, A.; Liao, Y.; Liu, S.; Lu, M.; Wang, Y.; Gao, C.; and Li, X. 2021a. Mining the benefits of two-stage and one-stage hoi detection. In *NeurIPS*, 17209–17220.
- Zhang, F.; Sheng, L.; Guo, B.; Chen, R.; and Chen, J. 2023. SQA: Strong Guidance Query with Self-Selected Attention for Human-Object Interaction Detection. In *ICASSP*, 1–5.
- Zhang, F. Z.; Campbell, D.; and Gould, S. 2021b. Spatially conditioned graphs for detecting human-object interactions. In *ICCV*, 13319–13327.
- Zhang, Y.; Pan, Y.; Yao, T.; Huang, R.; Mei, T.; and Chen, C.-W. 2022b. Exploring structure-aware transformer over interaction proposals for human-object interaction detection. In *CVPR*, 19548–19557.
- Zhong, X.; Qu, X.; Ding, C.; and Tao, D. 2021. Glance and gaze: Inferring action-aware points for one-stage human-object interaction detection. In *CVPR*, 13234–13243.
- Zhou, D.; Liu, Z.; Wang, J.; Wang, L.; Hu, T.; Ding, E.; and Wang, J. 2022. Human-Object Interaction Detection via Disentangled Transformer. In *CVPR*, 19568–19577.
- Zhou, T.; Qi, S.; Wang, W.; Shen, J.; and Zhu, S.-C. 2021. Cascaded parsing of human-object interaction recognition. *PAMI*, 44(6): 2827–2840.
- Zou, C.; Wang, B.; Hu, Y.; Liu, J.; Wu, Q.; Zhao, Y.; Li, B.; Zhang, C.; Zhang, C.; Wei, Y.; et al. 2021. End-to-end human object interaction detection with hoi transformer. In *CVPR*, 11825–11834.