

ViPLO: Vision Transformer based Pose-Conditioned Self-Loop Graph for Human-Object Interaction Detection

Jeeseung Park¹ Jin-Woo Park^{1,2} Jong-Seok Lee²
¹mAy-I Inc., Seoul, Korea, ²Yonsei University, Korea
 {jspark, jin}@may-i.io jong-seok.lee@yonsei.ac.kr

Abstract

Human-Object Interaction (HOI) detection, which localizes and infers relationships between human and objects, plays an important role in scene understanding. Although two-stage HOI detectors have advantages of high efficiency in training and inference, they suffer from lower performance than one-stage methods due to the old backbone networks and the lack of considerations for the HOI perception process of humans in the interaction classifiers. In this paper, we propose Vision Transformer based Pose-Conditioned Self-Loop Graph (ViPLO) to resolve these problems. First, we propose a novel feature extraction method suitable for the Vision Transformer backbone, called masking with overlapped area (MOA) module. The MOA module utilizes the overlapped area between each patch and the given region in the attention function, which addresses the quantization problem when using the Vision Transformer backbone. In addition, we design a graph with a pose-conditioned self-loop structure, which updates the human node encoding with local features of human joints. This allows the classifier to focus on specific human joints to effectively identify the type of interaction, which is motivated by the human perception process for HOI. As a result, ViPLO achieves the state-of-the-art results on two public benchmarks, especially obtaining a +2.07 mAP performance gain on the HICO-DET dataset.

1. Introduction

Human-Object Interaction (HOI) detection aims to localize human and objects in the image, and classify interactions between them. A detected HOI instance is represented in the form of $\langle \text{human}, \text{object}, \text{interaction} \rangle$ triplet. HOI detection is steadily attracting attention in the computer vision field, due to the high potential for use in the downstream tasks such as action recognition [46], scene understanding [44], and image captioning [47].

Existing HOI detectors can be divided into two cate-

gories: two-stage methods and one-stage methods. The former usually consists of three steps: 1) human and object detection with an off-the-shelf detector; 2) feature extraction for human and objects with ROI-Pooling; and 3) prediction of the interaction between human and objects with the extracted features [3, 8, 12, 15, 16, 40]. One-stage HOI detection framework is first proposed in PPDM [26], which directly predicts HOI triplets based on interaction points [26, 42] and union boxes [21]. Recent studies based on one-stage structure use the two-branch transformer, which predict HOI triplets with two sub-task decoders and matching processes [27, 53]. Nevertheless, these methods suffer from low training speed and large memory usage [50, 55].

In contrast, two-stage HOI detectors show high training speed using pre-trained object detectors. They are also advantageous when the object bounding box is known in advance from the image. For example, when analyzing HOI events of fixed target objects, we can efficiently detect interaction using the fixed object regions without an extra object detection process. Moreover, interaction inference can be performed only for the desired human-object pair, which is also a big advantage compared to the one-stage methods in the application situation. Despite these advantages, two-stage methods have recently been less studied due to their low performance. In this paper, we focus on improving the performance of the two-stage approach, which can also enjoy the benefit of reduced computational and memory complexity. In particular, we improve the two important parts: feature extraction and interaction prediction.

Regarding the feature extraction part, we note that most of the two-stage methods rely on ResNet [14] as backbone networks [15, 16, 49]. Recently, the Vision Transformer (ViT) has emerged as a powerful alternative yielding state-of-the-art performance in various computer vision fields [34, 36, 38], and it also has a potential as an improved feature extractor for two-stage HOI detection. However, it is not straightforward to apply a conventional feature extraction method (e.g., ROIAlign [13]) directly to a ViT backbone due to the different shapes of the output feature maps of ResNet and ViT. Therefore, we propose a novel Masking




Figure 1. The process of HOI recognition by humans. Humans first localize each person and object (Step 1), identify interactiveness using the spatial relationship and human pose (Step 2), and finally focus on specific human joints to recognize the type of interaction (Step 3).

with Overlapped Area (MOA) module for feature extraction using a ViT backbone. The MOA module resolves the spatial quantization problem using the overlapped area between each patch and the given region in the attention function. In addition, efficient computation in the MOA module keeps the inference speed similar to that of the ResNet backbone.

To improve the interaction prediction part, we refer to the human perception process of recognizing HOI. Previous studies [1, 9, 10] show that the HOI perception of humans considers information such as object identity, relative positions, reach motions, manipulation motions, and context. Following this, we assume that HOI recognition by humans consists of three steps (Fig. 1): 1) Localize each human and object; 2) Identify interactiveness using the spatial relationship between human and object and human pose information; 3) Focus on specific human joints to identify the type of interaction. Motivated by these, we design a pose-conditioned graph neural network for interaction prediction. For identifying the interactiveness part, previous studies exclude non-interactive pairs with auxiliary networks [25, 30]. However, our model represents interactiveness with edge encoding, which is obtained by using the spatial human-object relationship and the human pose information. We can expect more meaningful message passing with interactiveness-aware edge encoding. For focusing on specific joints part, early methods utilize local pose information with simple message passing or attention mechanisms [40, 54]. In contrast, we introduce pose-aware attention mechanism using query (spatial features) / key (joint features) structure, resulting in human node encodings containing richer local information by a self-loop structure.

As a result, we propose Vision Transformer based Pose-Conditioned Self-Loop Graph (ViPLO), a novel two-stage HOI detector. We demonstrate the effectiveness of our framework by achieving state-of-the-art results on two public benchmarks: HICO-DET [3] and V-COCO [11]. In addition, we further evaluate and analyze our model design through extensive ablation studies. The contributions of this paper are as follows:

- We propose an effective two-stage HOI detector with a ViT backbone by introducing a novel MOA module, which is a suitable feature extraction module for ViT.

- We design the interaction classifier with a pose-conditioned self-loop graph structure, motivated by the human perception process.
- Our model outperforms the previous state-of-the-art results on a widely-used HOI benchmark.

2. Related Work

2.1. One-stage Methods

The one-stage HOI detection framework directly predicts HOI triplets by associating human and objects with predefined anchors, then predicting their interaction. Early methods utilized interaction keypoints [26, 42] or union regions [21] as the predefined anchors. Recently, Transformer-based HOI detectors [22, 37] became the main approach for HOI detection. In particular, HOI Transformers formulated the HOI detection task as a set prediction problem, which are trained with bipartite matching and Hungarian loss. Moreover, unlike the structure of DETR [2], methods using two decoders corresponding to each sub-task have also been studied [5, 27, 48, 53]. However, these methods suffer from low training speed and large memory usage. Besides, since the detection process cannot be separated from the whole framework, it is inefficient in case of already knowing the detection results or wishing to infer interaction only for specific human-object pairs.

2.2. Two-stage Methods

The two-stage HOI detection framework detects human and object with an off-the-shelf detector [35] and then classifies the interaction label for each human-object pair. After the appearance of HO-RCNN [4], which is a widely used multi-stream framework, many recent studies use a variety of additional information to get richer contextual features for the interaction classifier, such as spatial features [7, 39, 49], pose features [12, 25, 40], and linguistic features [7, 31]. Several studies [7, 32, 39, 41, 49] attempted to encode global contextual information using a message passing mechanism in a graph structure.

However, there is a lot of room for improvement on the feature extraction step, where existing methods still use

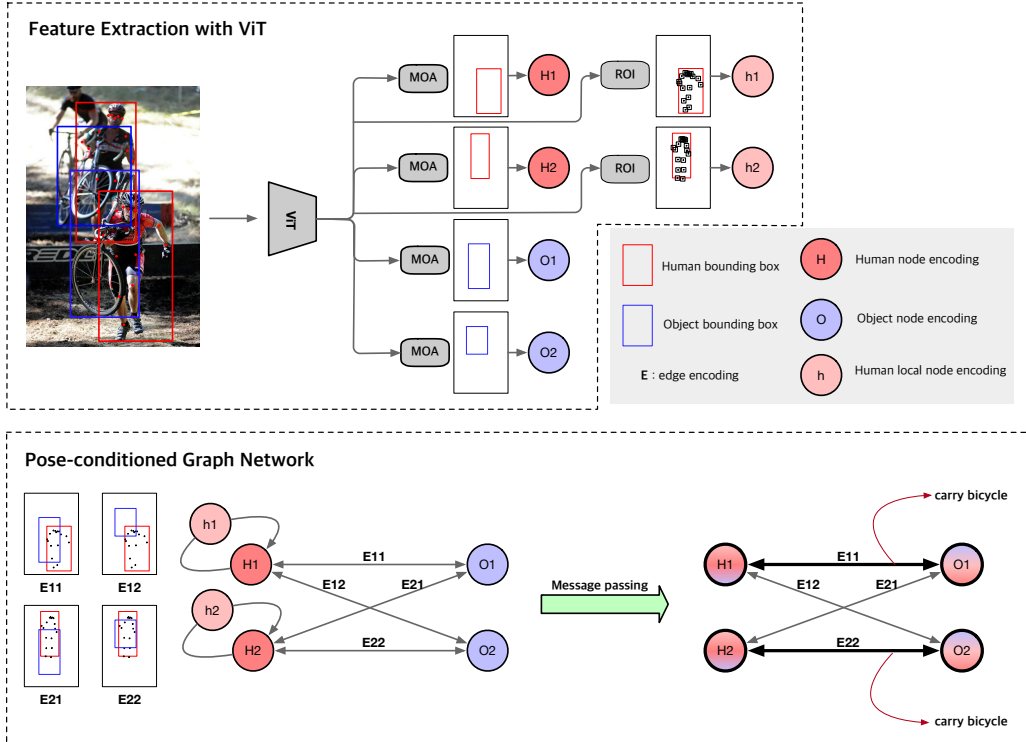


Figure 2. Overview of our ViPLO network. We first detect human and objects in a given image with Faster-RCNN [35], then estimate each human pose with an off-the-shelf pose estimator. Then, we extract features for each human and object using a ViT backbone and our novel MOA module. We also extract local features for each human with the estimated pose and ROIAlign [13]. Next, a graph neural network is used to classify interactions, where the node encoding is initialized with the extracted features and the edge encoding is obtained by spatial and human pose information. Through the message passing process, the human node encoding is also updated with the human local node encoding, helping the model focus on the necessary local part of each human. After the message passing procedure, the combination of human node, object node, and edge encoding becomes the representation of HOI triplet which is used to predict the interaction class.

ResNet backbones and ROIAlign. Since the extracted feature becomes the input of the interaction classifier, the performance of the system would vary significantly depending on the feature extraction method. DEFR [20] utilized ViT as a backbone network, but considered the HOI recognition task, not HOI detection. DEFR also did not handle the quantization problem, leading to misalignment between the given region and the extracted feature. Moreover, existing interaction classifiers do not consider the HOI perception process by humans, which focuses on the specific human joints to identify the type of interaction. In contrast, our HOI detector uses a ViT backbone and a graph neural network with pose-conditioned self-loop structure.

3. Method

In this section, we describe our ViPLO in two steps. In Sec. 3.1, we introduce a novel feature extraction module for the ViT backbone. Then in Sec. 3.2, we explain our graph neural network with a human pose-conditioned structure for HOI detection. An overview of our framework is shown in

Fig. 2.

3.1. Feature Extraction with ViT

We propose a new feature extraction method using the ViT backbone, which enables the following network to recognize interaction effectively. We first detect human and objects in a given image using Faster R-CNN [35] (off-the-shelf detector). Then, instead of ResNet, which is a commonly used backbone for feature extraction in conventional two-stage approaches, we use ViT. In the case of the ResNet backbone, visual features are extracted from the feature map of the backbone networks via ROI-Pooling or ROIAlign [13]. However, for the ViT backbone, a new feature extraction method is needed due to the different shapes of the output feature map compared to the ResNet. Therefore, we introduce a novel *Masking with Overlapped Area* (MOA) module, which is detailed below.

Masking with Overlapped Area ViT prepends a learnable embedding (i.e., CLS token) to the sequence of patch embeddings. After passing the transformer encoder, this

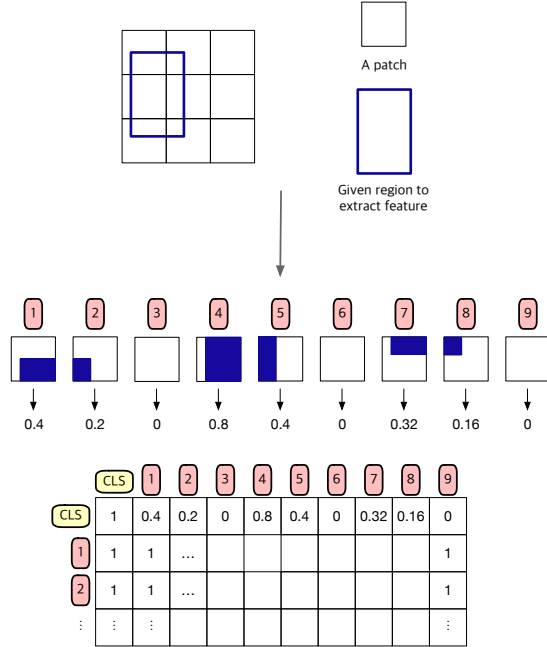


Figure 3. Overview of the MOA module. The normalized area (e.g., 0.4, 0.2, 0, ...) between each patch and the given region is calculated. Then, we use these values as the first row of the attention mask matrix (i.e., S in Eq. 1), whose logarithm is added to the attention map created by the query and key.

CLS token serves as an image representation. Therefore to extract a feature corresponding to a given region properly, we have to work with the CLS Token. Inspired by DEFR [20], we mask out all patch embeddings outside the given region for the CLS token. For example, to obtain features of an object, the CLS token can only attend to the patch embedding inside the object bounding box.

However, this approach suffers from the quantization problem as in ROI-Pooling [13]. Existing ViTs usually use 14×14 , 16×16 , and 32×32 input patch size. Since the coordinates of the bounding box are given in units of pixels, the edges of the bounding box often do not align with patch boundaries. To solve this problem, we can simply mask out or attend all patch embeddings that the bounding box edge passes through for the CLS token (i.e., quantization). However, this quantization leads to misalignment between the given region and the extracted feature.

Our MOA module address this quantization problem by utilizing the overlapped area between each patch and the given region to the attention mask in the attention function. We develop our MOA module into a quantization-free layer as shown in Fig. 3. First, we calculate the overlapped area of each patch and the given region, and normalize it with the area of the patch. Then, these normalized overlapped area become the first row of the attention mask matrix, which is added to the attention map created by the query and key of

the self-attention layer after logarithm.

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}} + \log(S)\right)V, \quad (1)$$

where Q , K , and V are the query, key, and value, respectively, and d_k is their dimension. S denotes the normalized overlapped area of the patch and the given region.

Efficient computation for MOA The MOA module leads to a large performance increment in HOI detection, as shown in the ablation studies in Sec. 4.3. To use the MOA module, however, overlapped area S has to be computed for each bounding box, which may be a computational burden. We design the entire process of computing S to be run through GPU operations, and apply techniques for reducing computations. See Appendix for more details about computational complexity and implementation of the MOA module.

3.2. Pose-conditioned Graph Neural Network

After extracting features with our MOA module, we use a graph neural network to detect interactions between human and objects. Inspired by the process of HOI detection by humans, we improve the spatially conditioned graph (SCG) [49] by using human pose information. We first review the baseline SCG, and then describe our improved network design.

SCG Baseline SCG is a bipartite graph neural network with human and object nodes. The features extracted using ResNet and ROIAlign are used to initialize the node encodings. The edge encodings are initialized with features based on the spatial information of two bounding boxes, i.e., human and object. Then, bidirectional message passing is performed between nodes conditioned on their edge encodings. After message passing, the updated node encodings and edge encodings are used to classify interaction between the human and object.

Pose-aware Edge encoding As mentioned in the introduction, the second step of the HOI perception process is identifying interactiveness using the spatial information between human and objects and human pose information. In the graph structure, the relationship between a human node and object node is expressed by the edge encoding between the two nodes. Thus, we initialize the edge encoding based on human pose information as well as spatial information. In particular, we compute pairwise spatial features (i.e., query) as in SCG. Then, we compute the pairwise joint features (i.e., key) by applying MLP to the handcrafted joint

features including the coordinates of each joint and the directional vector from the joint to the object box center. Following the attention mechanism, we compute the attention score for each human joint by the dot product of the query and key as follows:

$$\alpha_{ij} = \text{softmax}(Q_{ij}K_{ij}^T \cdot s_i), \quad (2)$$

where α_{ij} , Q_{ij} , and K_{ij} are the joint attention weight, pairwise spatial feature, and pairwise joint feature for the i th human and j th object, respectively, and s_i is the pose estimation confidence score for the i th human. The obtained attention score is used in the following message passing step. The edge encoding is initialized with a pairwise spatial feature, but becoming pose-aware due to the attention mechanism.

Message Passing with Pose As described in [40], there are cases where detailed local features are required for recognizing interaction (e.g., catch and grab). This is related to the third step of the HOI perception process by humans, focusing on specific human joints to identify the interaction. The pose-aware edge encoding alone cannot adjust the model to utilize local information. In our method, we incorporate the local information in the human node encoding, i.e., we update the human node encoding using the local features of each human joint by new message functions.

First, we extract the local features by applying ROIAlign to the ViT output for the local region box for each human joint, which is inspired by [40]. Then, we compute the weighted sum of each local feature of a human joint to get the human local feature:

$$\mathbf{x}_{ij,local} = \sum_k \alpha_{ijk} \odot \mathbf{x}_{ik,local}, \quad (3)$$

where α_{ijk} is the k th value of joint attention weight α_{ij} , $\mathbf{x}_{ik,local}$ is the local feature for the k th joint of the i th human, and $\mathbf{x}_{ij,local}^t$ is the i th human local feature for the j th object. \odot denotes the element-wise multiplication.

To use ROIAlign with the ViT output, we reshape the image patch embeddings to the form of a square feature map, i.e., Unflatten. After the unflatten process, the reshaped patch embeddings play the same role as ResNet output feature maps, so we can simply apply ROIAlign to the patch embeddings. We hypothesize that the image patch tokens contain more detailed information rather than the CLS token, choosing the ROIAlign to extract local features rather than the MOA module. When training ViT, the classification head is attached to the final CLS token output. Therefore, the CLS token has coarse information for image classification. On the other hand, patch tokens, which become the components of CLS tokens (i.e., attention mechanism),




Figure 4. Heatmaps showing joint attention. **Left:** ViPLO focuses on the local information of hand to detect interaction between cell-phone and human. **Right:** ViPLO does not focus on a specific joint for local information when detecting interaction between car and human, resulting in no-interaction.

have fine information for each region. This hypothesis is supported by the ablation study in Sec. 4.3.

Second, the extracted human local features ($\mathbf{x}_{ij,local}$) become human local node encodings, which are used to update the human node encodings (Fig. 2). In detail, the human node encodings are iteratively updated with a message from the human local features, as well as object node encodings and edge encodings:

$$M_{\mathcal{O} \rightarrow \mathcal{H}}(\mathbf{x}_i^t, \mathbf{y}_j^t, \mathbf{z}_{ij}) = \text{MBF}_o(\mathbf{x}_{ij,local} \oplus \mathbf{y}_j^t, \mathbf{z}_{ij}), \quad (4)$$

$$M_{\mathcal{H} \rightarrow \mathcal{O}}(\mathbf{x}_i^t, \mathbf{z}_{ij}) = \text{MBF}_h(\mathbf{x}_i^t, \mathbf{z}_{ij}), \quad (5)$$

where $M_{\mathcal{O} \rightarrow \mathcal{H}}$ denotes the message function from object nodes to human nodes, and vice versa. \mathbf{x}_i^t and \mathbf{y}_j^t are the i th human node encoding and j th object node encoding, respectively, at message passing step t . MBF indicates the multi-branch fusion module proposed in [49]. \oplus denotes the concatenation operation. The human nodes are updated not only with the object nodes but also with the human local features which share the similar information with the human node encodings. This algorithm resembles the self-loop structure in graph theory.

Effectiveness of Pose-Conditioned Graph Our pose-aware edge encoding scheme helps models focus only on the necessary message passing. For example, the case where a human is just standing next to the object and the case where a hand is extended to the direction of the object are different, even though the spatial relationship of the region pair is the same. We can expect more meaningful message passing in the latter case, due to the different edge encoding from each human pose information. Moreover, human local features can contain richer local information by the attention operation in Eq. 3 (hand in the latter case). Furthermore, as shown in Fig. 4, human node encodings are updated with self-loop architecture, to focus only on the necessary local part as the message passing step progresses. These human node encodings with rich local information also enrich the object node encodings with human-to-object messages, resulting in effective HOI detection.

Method	Backbone	HICO-DET						V-COCO	
		Default			Known Object			AP_{role}^{S1}	AP_{role}^{S2}
		Full	Rare	Non-Rare	Full	Rare	Non-Rare		
<i>One-stage methods</i>									
UnionDet [21]	ResNet-50-FPN	17.58	11.72	19.33	19.76	14.68	21.27	47.5	56.2
IP-Net [42]	Hourglass-104	19.56	12.79	21.58	22.05	15.77	23.92	51.0	-
PPDM [26]	Hourglass-104	21.73	13.78	21.40	24.58	16.65	26.84	-	-
GG-Net [52]	Hourglass-104	23.47	16.48	25.60	27.36	20.23	39.48	54.7	-
HOTR [22]	ResNet-50	25.10	17.34	27.42	-	-	-	55.2	64.4
HOI-Trans [56]	ResNet-101	26.61	19.15	28.84	29.13	20.98	31.57	52.9	-
AS-Net [5]	ResNet-50	28.87	24.25	30.25	31.74	27.07	33.14	53.9	-
QPIC [37]	ResNet-101	29.90	23.92	31.69	32.38	26.06	34.27	58.3	60.7
MSTR [23]	ResNet-50	31.17	25.31	32.92	34.02	28.83	35.57	62.0	65.2
SSRT [19]	ResNet-101	31.34	24.31	33.32	-	-	-	<u>65.0</u>	67.1
CDN [48]	ResNet-101	32.07	27.19	33.53	34.79	29.48	36.38	63.9	65.9
DOQ [33]	ResNet-50	33.28	29.19	34.50	-	-	-	63.5	-
IF [30]	ResNet-50	33.51	30.30	34.46	36.28	33.16	37.21	63.0	65.2
GEN-VLKT [27]	ResNet-101	34.95	31.18	<u>36.08</u>	<u>38.22</u>	34.36	<u>39.37</u>	63.6	65.9
ParMap [43]	ResNet-50	<u>35.15</u>	33.71	35.58	37.56	35.87	38.06	63.0	65.1
<i>Two-stage methods</i>									
VCL [15]	ResNet-50	23.63	17.21	25.55	25.98	19.12	28.03	48.3	-
ATL [16]	ResNet-50	23.67	17.64	25.47	26.01	19.60	27.93	-	-
IDN [24]	ResNet-50	24.58	20.33	25.86	27.89	23.64	29.16	53.3	60.3
FCL [17]	ResNet-50	24.68	20.03	26.07	26.80	21.61	28.35	52.4	-
SCG [49]	ResNet-50-FPN	29.26	24.61	30.65	32.87	27.89	34.35	54.2	60.9
STIP [51]	ResNet-50	32.22	28.15	33.43	35.29	31.43	36.45	66.0	70.7
DEFR [20]	ViT-B/16	32.35	33.45	32.02	-	-	-	-	-
UPT [50]	ResNet-101	32.62	28.62	33.81	36.08	31.41	37.47	61.3	67.1
ViPLO _s	ViT-B/32	34.95	<u>33.83</u>	35.28	38.15	<u>36.77</u>	38.56	60.9	66.6
ViPLO _l	ViT-B/16	37.22	35.45	37.75	40.61	38.82	41.15	62.2	68.0

Table 1. Performance comparison in terms of mAP on the HICO-DET and V-COCO datasets. For fair comparison between one-stage methods and two-stage methods, we report results using an object detector finetuned on the training dataset for two-stage methods. We use the finetuned object detector from [50], i.e., DETR with the ResNet-101 backbone. We do not consider the results using detections from [7] due to a problem of reproducibility. In each evaluation setting, the best result is marked with bold and the second best result is underlined.

3.3. Training and Inference

We follow the training and inference procedure of the SCG [49, 50]. From the pose-conditioned graph after the message passing, we get interaction classification scores for each pair of human and object node using the corresponding edge encoding. The focal loss [28] is used as the multi-label classification loss to train the possible interactions for each human-object pair.

4. Experiments

In this section, we first introduce our experimental settings (Sec. 4.1). We then compare our proposed model with state-of-the-art approaches (Sec. 4.2). Finally we conduct ablation studies (Sec. 4.3) demonstrating the advantages and effectiveness of our model. Qualitative results can be

found in Appendix.

4.1. Experimental Settings

Datasets We evaluate our model on two public datasets, HICO-DET [3] and V-COCO [11], which are widely used in HOI detection. HICO-DET contains 47,776 images (38,118 training images and 9,658 test images), with 80 object categories (same categories with the MS-COCO [29] dataset) and 117 verb classes constructing 600 types of HOI triplets. HICO-DET provides more than 150k annotated human-object pairs. V-COCO is a subset of MS-COCO, which is a much smaller dataset than HICO-DET. It contains 10,346 images (2,533 training images, 2,867 validation images, and 4,946 test images), with the same 80 object categories and 29 verb classes. Since HICO-DET containing various verb types is more suitable in the real world, we focus our

evaluation mainly on the HICO-DET dataset.

Evaluation Metrics We follow the standard settings in [8], reporting mean Average Precision (mAP) for evaluation. Prediction of a HOI triplet is considered as a true positive when both predicted human and object bounding boxes have IoUs larger than 0.5 compared to the ground truth boxes, and HOI category prediction is accurate. For HICO-DET, we report mAP over two evaluation settings (Default and Known Object), with three HOI category subsets: all 600 HOI triplets (Full), 138 HOI triplets with fewer than 10 training samples (Rare), and 462 HOI triplets with 10 or more training samples (Non-Rare). For V-COCO, we report mAP for two scenarios: including a prediction of the occluded object bounding box (Scenario 1), or ignoring such a prediction (Scenario 2).

Implementation Details Since neither the HOI Detection dataset nor the Faster R-CNN detector result contains a ground truth or estimated human pose, we need an off-the-shelf pose estimator. We apply Vitpose [45] as a pose estimator, which is trained on the MS-COCO Keypoint dataset [29]. As a result, each human bounding box (including the ground truth human box) has 17 keypoints.

We consider the following two variants of our method: 1) ViPLO_s: a small version of ViPLO, which uses ViT-B/32 as the backbone; and 2) ViPLO_l: large version of ViPLO, which uses ViT-B/16 as the backbone. All other hyperparameters are the same between the two models. The feature extractor (i.e., backbone) is initialized using CLIP [34] with pre-trained weights. We cannot directly apply CLIP image pre-processing, due to the center-crop process which is not suitable for object detection. Hence, we resize input images to 672 × 672 pixels using data transformation in [18]. Bounding boxes and human joints are resized accordingly. ViT-B is pre-trained on 224 × 224 pixels, but the model can be fine-tuned on 672 × 672 pixels with 2D interpolation of the pre-trained position embeddings, as described in [6].

4.2. Comparison to State-of-the-Art

Table 1 shows the performance comparison of ViPLO and other state-of-the-art methods, which are grouped into one-stage and two-stage methods. For HICO-DET, our ViPLO outperforms all existing one-stage and two-stage methods by a large margin in each evaluation setting. Specifically, ViPLO_l achieves new state-of-the-art performance of **37.22** mAP in the Default Full setting, obtaining a relative performance gain of 5.89% compared to the most recent ParMap [43]. Compared with UPT [50], which is the previous state-of-the-art two-stage method, our model achieves a significant relative performance gain of 14.10%. These results indicate the effectiveness of our method, utilizing the ViT backbone with the MOA module, and up-

Method	Backbone	Default		
		Full	Rare	Non-Rare
<i>Two-stage methods with ground truth boxes</i>				
iCAN [8]	ResNet-50	33.38	21.43	36.95
TIN [25]	ResNet-50	34.26	22.90	37.65
VCL [15]	ResNet-50	43.09	32.56	46.24
IDN [24]	ResNet-50	43.98	40.27	45.09
FCL [17]	ResNet-50	44.26	35.46	46.88
ATL [16]	ResNet-50	44.27	35.52	46.89
SCG [49]	ResNet-50-FPN	51.53	41.01	54.67
ViPLO _s	ViT-B/32	<u>58.23</u>	<u>54.08</u>	<u>59.46</u>
ViPLO _l	ViT-B/16	62.09	59.26	62.93

Table 2. Performance comparison in terms of mAP on the HICO-DET dataset with ground truth detection results. The best result is marked in bold and the second best result is underlined.

Method	Backbone	mAP ↑	GPU memory ↓	Speed ↓
SCG [49]	ResNet-50	29.26	3423 MiB	106 ms
ViPLO _s	ViT-B/32	33.92	2531 MiB	110 ms
ViPLO _l	ViT-B/16	36.97	3119 MiB	131 ms

Table 3. Comparison of mAP, memory usage, and speed for inference on the HICO-DET test dataset. We exclude the pose-conditioned graph part in our model for a fair comparison. We conduct an inference process with a batch size of 1 on a Geforce RTX 3090 GPU. “Speed” column means the elapsed time for processing one image.

dating human features with the pose-conditioned self-loop structure. For V-COCO, our ViPLO achieves comparable performance to previous state-of-the-art methods. The performance gain is not as noticeable as on HICO-DET, due to the fewer training samples which can be fatal to the ViT architecture [38].

As mentioned in the introduction, an advantage of two-stage methods is that they can also be used for HOI recognition when ground truth bounding boxes are available. Performance comparison using ground truth boxes for HICO-DET is shown in Table 2. Our method outperforms all previous state-of-the-art methods with a significant large margin of 10.56 in terms of mAP in the Default Full setting.

Meanwhile, since most previous methods use ResNet as the backbone network, memory efficiency and speed of our method using the ViT backbone are examined. As shown in Table 3, we obtain a large performance improvement in terms of mAP, while maintaining speed and memory efficiency at a similar level to the ResNet-50 backbone. Surprisingly, ViPLO consumes less GPU memory than SCG using ResNet-50 as the backbone network. These results show the efficiency and efficacy of the ViT backbone and proposed MOA module.

Methods	Full	Rare	Non-Rare	Methods	Full	Rare	Non-Rare
SCG (ResNet-50)	29.26	24.61	30.65	Base (ViT + MOA)	33.92	30.84	34.83
ViT + ROI	32.84	28.48	34.14	+ pose edge	34.51	32.74	35.04
ViT + MOA ^Q	33.20	31.26	33.78	+ pose edge + local pose (MOA)	34.44	31.21	35.41
ViT + MOA	33.92	30.84	34.83	+ pose edge + local pose (ROI)	34.95	33.83	35.28

(a) Effect of ViT backbone and MOA module.

(b) Effect of the pose-conditioned graph neural network.

Table 4. Ablation study of ViPLO components on the HICO-DET test dataset under the Default setting.

4.3. Ablation Study

In this subsection, we conduct ablation studies to evaluate and analyze the effectiveness of our proposed module and our model design. All experiments are conducted on the HICO-DET test dataset under the Default setting, with a small version of ViPLO, i.e., ViPLO_s.

MOA module Our model adopts the ViT backbone and MOA module for feature extraction, where this feature serves as the node encoding in the subsequent graph neural network. We explore the benefits of the ViT backbone and MOA module, as shown in Table 4a. We exclude the pose-conditioned graph structure for a fair comparison with the SCG baseline. ViT + ROI (Row 2) means the case extracting features with reshaped patch embeddings and ROIAlign for the bounding box, which is described in the human local feature extraction process in Sec. 3.2. ViT + MOA^Q (Row 3) means the case using the MOA module without calculating the overlapped area, and simply masking out the patch embeddings that bounding box edges pass through (i.e., quantization). Compare to the SCG baseline, the ViT backbone significantly improves the performance by 3.58 in terms of mAP for the Full setting, even without the MOA module. Further using our MOA module (ViT + MOA) yields additional performance improvement by 1.08 in mAP, demonstrating the effectiveness of the MOA module. In particular, the performance difference depending on the quantization of the MOA module supports the importance of the quantization-free method by calculating the overlapped area, which is a key part of our MOA module.

Pose-conditioned Graph Neural Network Table 4b explore the effectiveness of the components of pose-conditioned graph neural network, on top of the ViT backbone and MOA module (Row 1). We first add the human pose information to formulate the edge encoding (Row 2), which shows a performance improvement of 0.59 in mAP. Moreover, adding a self-loop structure with human local node encodings boosts the performance by 0.44 in mAP, proving the effectiveness of a self-loop structure that aligns with the perception process by humans (Row 4). However, extracting human local features by the MOA module does

Method	Backbone	Pre-trained Dataset	mAP ↑
SCG	ResNet50	MS COCO	29.26
ViPLO	ResNet50	CLIP	32.23
ViPLO _s	ViT-B/32	CLIP	33.92
ViPLO _l	ViT-B/16	CLIP	36.97

Table 5. Effect of CLIP pre-trained parameters. Experiment setting is the same as in Table 3

not boost the performance, it actually decreases the performance (Row 3).

CLIP pre-trained parameters As we mentioned, the feature extractor (i.e., ViT backbone) is initialized with CLIP pre-trained mode. However, the ResNet backbone of SCG is initialized with COCO pre-trained parameters, which is trained with much less amount of images than CLIP. Therefore, we conduct an additional experiment comparing the ResNet backbone and ViT backbones with both the same CLIP pre-trained parameters, as shown in Table 5. Please note that we cannot apply the same preprocess/feature extraction method for Row 1 and Row 2, due to the modified architecture of CLIP ResNet50, such as the attention pooling mechanism. So we adopt the preprocess/feature extraction method of Row 2 in the same way as ViT + ROI in Table 4a in the paper. We can still observe the superior performance of our model in the same pre-trained mode, showing the effectiveness of our ViT + MOA scheme.

5. Conclusion

In this paper, we have proposed the ViPLO, the state-of-the-art two-stage HOI detector assisted by the MOA module and pose-conditioned graph. With the MOA module, our detector fully utilizes the ViT as the feature extractor by addressing the quantization problem. In addition, the pose-conditioned graph inspired by the HOI perception process of humans makes our detector exploit the rich information from human poses. In addition, unlike the one-stage methods, ViPLO has the advantage of low complexity and convenience to apply in real world scenarios.

References

- [1] Christopher Baldassano, Diane M Beck, and Li Fei-Fei. Human-object interactions are more than the sum of their parts. *Cerebral Cortex*, 27(3):2276–2288, 2017. 2
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision*, pages 213–229. Springer, 2020. 2
- [3] Yu-Wei Chao, Yunfan Liu, Xieyang Liu, Huayi Zeng, and Jia Deng. Learning to detect human-object interactions. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, 2018. 1, 2, 6
- [4] Yu-Wei Chao, Yunfan Liu, Xieyang Liu, Huayi Zeng, and Jia Deng. Learning to detect human-object interactions. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 381–389. IEEE, 2018. 2
- [5] Mingfei Chen, Yue Liao, Si Liu, Zhiyuan Chen, Fei Wang, and Chen Qian. Reformulating hoi detection as adaptive set prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9004–9013, 2021. 2, 6
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceedings of the 9th International Conference on Learning Representations*, 2021. 7
- [7] Chen Gao, Jiarui Xu, Yuliang Zou, and Jia-Bin Huang. Drg: Dual relation graph for human-object interaction detection. In *Proceedings of the European Conference on Computer Vision*, pages 696–712. Springer, 2020. 2, 6
- [8] Chen Gao, Yuliang Zou, and Jia-Bin Huang. iCAN: Instance-centric attention network for human-object interaction detection. In *Proceedings of the British Machine Vision Conference*, 2018. 1, 7
- [9] Abhinav Gupta and Larry S Davis. Objects in action: An approach for combining action understanding and object perception. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. 2
- [10] Abhinav Gupta, Aniruddha Kembhavi, and Larry S Davis. Observing human-object interactions: Using spatial and functional compatibility for recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(10):1775–1789, 2009. 2
- [11] Saurabh Gupta and Jitendra Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015. 2, 6
- [12] 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. 1, 2
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2961–2969, 2017. 1, 3, 4
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 1
- [15] Zhi Hou, Xiaojiang Peng, Yu Qiao, and Dacheng Tao. Visual compositional learning for human-object interaction detection. In *Proceedings of the European Conference on Computer Vision*, pages 584–600. Springer, 2020. 1, 6, 7
- [16] Zhi Hou, Baosheng Yu, Yu Qiao, Xiaojiang Peng, and Dacheng Tao. Affordance transfer learning for human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 495–504, 2021. 1, 6, 7
- [17] Zhi Hou, Baosheng Yu, Yu Qiao, Xiaojiang Peng, and Dacheng Tao. Detecting human-object interaction via fabricated compositional learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14646–14655, 2021. 6, 7
- [18] Junjie Huang, Zheng Zhu, Feng Guo, and Guan Huang. The devil is in the details: Delving into unbiased data processing for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5700–5709, 2020. 7
- [19] ASM Iftekhar, Hao Chen, Kaustav Kundu, Xinyu Li, Joseph Tighe, and Davide Modolo. What to look at and where: Semantic and spatial refined transformer for detecting human-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5353–5363, 2022. 6
- [20] Ying Jin, Yinpeng Chen, Lijuan Wang, Jianfeng Wang, Pei Yu, Lin Liang, Jenq-Neng Hwang, and Zicheng Liu. The overlooked classifier in human-object interaction recognition. *arXiv preprint arXiv:2203.05676*, 2022. 3, 4, 6
- [21] Bumsoo Kim, Taeho Choi, Jaewoo Kang, and Hyunwoo J Kim. Uniondet: Union-level detector towards real-time human-object interaction detection. In *Proceedings of the European Conference on Computer Vision*, pages 498–514. Springer, 2020. 1, 2, 6
- [22] Bumsoo Kim, Junhyun Lee, Jaewoo Kang, Eun-Sol Kim, and Hyunwoo J Kim. Hotr: End-to-end human-object interaction detection with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 74–83, 2021. 2, 6
- [23] Bumsoo Kim, Jonghwan Mun, Kyoung-Woon On, Minchul Shin, Junhyun Lee, and Eun-Sol Kim. Mstr: Multi-scale transformer for end-to-end human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19578–19587, 2022. 6
- [24] Yong-Lu Li, Xinpeng Liu, Xiaoqian Wu, Yizhuo Li, and Cewu Lu. Hoi analysis: Integrating and decomposing human-object interaction. *Advances in Neural Information Processing Systems*, 33:5011–5022, 2020. 6, 7
- [25] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 3585–3594, 2019. 2, 7
- [26] Yue Liao, Si Liu, Fei Wang, Yanjie Chen, Chen Qian, and Jishi Feng. Ppdm: Parallel point detection and matching for real-time human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 482–490, 2020. 1, 2, 6
- [27] Yue Liao, Aixi Zhang, Miao Lu, Yongliang Wang, Xiaobo Li, and Si Liu. Gen-vlkt: Simplify association and enhance interaction understanding for hoi detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20123–20132, 2022. 1, 2, 6
- [28] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2980–2988, 2017. 6
- [29] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Proceedings of the European Conference on Computer Vision*, pages 740–755. Springer, 2014. 6, 7
- [30] Xinpeng Liu, Yong-Lu Li, Xiaoqian Wu, Yu-Wing Tai, Cewu Lu, and Chi-Keung Tang. Interactiveness field in human-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20113–20122, 2022. 2, 6
- [31] Ye Liu, Junsong Yuan, and Chang Wen Chen. Consnet: Learning consistency graph for zero-shot human-object interaction detection. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4235–4243, 2020. 2
- [32] Siyuan Qi, Wenguan Wang, Baoxiong Jia, Jianbing Shen, and Song-Chun Zhu. Learning human-object interactions by graph parsing neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 401–417, 2018. 2
- [33] Xian Qu, Changxing Ding, Xingao Li, Xubin Zhong, and Dacheng Tao. Distillation using oracle queries for transformer-based human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19558–19567, 2022. 6
- [34] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 1, 7
- [35] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems*, 28, 2015. 2, 3
- [36] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7262–7272, 2021. 1
- [37] Masato Tamura, Hiroki Ohashi, and Tomoaki Yoshinaga. Qpic: Query-based pairwise human-object interaction detection with image-wide contextual information. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10410–10419, 2021. 2, 6
- [38] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *Proceedings of the International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. 1, 7
- [39] Oytun Ulutan, ASM Iftekhar, and Bangalore S Manjunath. Vsgnet: Spatial attention network for detecting human object interactions using graph convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13617–13626, 2020. 2
- [40] 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/CVF International Conference on Computer Vision*, pages 9469–9478, 2019. 1, 2, 5
- [41] Hai Wang, Wei-shi Zheng, and Ling Yingbiao. Contextual heterogeneous graph network for human-object interaction detection. In *Proceedings of the European Conference on Computer Vision*, pages 248–264. Springer, 2020. 2
- [42] Tiancai Wang, Tong Yang, Martin Danelljan, Fahad Shahbaz Khan, Xiangyu Zhang, and Jian Sun. Learning human-object interaction detection using interaction points. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4116–4125, 2020. 1, 2, 6
- [43] Xiaoqian Wu, Yong-Lu Li, Xinpeng Liu, Junyi Zhang, Yuzhe Wu, and Cewu Lu. Mining cross-person cues for body-part interactiveness learning in hoi detection. In *Proceedings of the European Conference on Computer Vision*, pages 121–136. Springer, 2022. 6, 7
- [44] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 418–434, 2018. 1
- [45] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vit-pose: Simple vision transformer baselines for human pose estimation. *arXiv preprint arXiv:2204.12484*, 2022. 7
- [46] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the Thirty-second AAAI Conference on Artificial Intelligence*, 2018. 1
- [47] Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. Image captioning with semantic attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4651–4659, 2016. 1
- [48] Aixi Zhang, Yue Liao, Si Liu, Miao Lu, Yongliang Wang, Chen Gao, and Xiaobo Li. Mining the benefits of two-stage and one-stage hoi detection. *Advances in Neural Information Processing Systems*, 34:17209–17220, 2021. 2, 6
- [49] Frederic Z Zhang, Dylan Campbell, and Stephen Gould. Spatially conditioned graphs for detecting human-object interactions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13319–13327, 2021. 1, 2, 4, 5, 6, 7
- [50] Frederic Z Zhang, Dylan Campbell, and Stephen Gould. Efficient two-stage detection of human-object interactions

- with a novel unary-pairwise transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20104–20112, 2022. 1, 6, 7
- [51] Yong Zhang, Yingwei Pan, Ting Yao, Rui Huang, Tao Mei, and Chang-Wen Chen. Exploring structure-aware transformer over interaction proposals for human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19548–19557, 2022. 6
- [52] Xubin Zhong, Xian Qu, Changxing Ding, and Dacheng Tao. Glimpse and gaze: Inferring action-aware points for one-stage human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13234–13243, 2021. 6
- [53] Desen Zhou, Zhichao Liu, Jian Wang, Leshan Wang, Tao Hu, Errui Ding, and Jingdong Wang. Human-object interaction detection via disentangled transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19568–19577, 2022. 1, 2
- [54] Penghao Zhou and Mingmin Chi. Relation parsing neural network for human-object interaction detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 843–851, 2019. 2
- [55] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. 1
- [56] Cheng Zou, Bohan Wang, Yue Hu, Junqi Liu, Qian Wu, Yu Zhao, Boxun Li, Chenguang Zhang, Chi Zhang, Yichen Wei, et al. End-to-end human object interaction detection with hoi transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11825–11834, 2021. 6