

End-to-End Human Object Interaction Detection with HOI Transformer

Cheng Zou*, Bohan Wang*, Yue Hu, Junqi Liu, Qian Wu, Yu Zhao, Boxun Li,
Chenguang Zhang, Chi Zhang, Yichen Wei, Jian Sun
MEGVII Technology

{zoucheng, wangbohan, huyue, liujunqi, wuqian, zhaoyu03, liboxun,
zhangchenguang, zhangchi, weiyichen, sunjian}@megvii.com

Abstract

We propose HOI Transformer to tackle human object interaction (HOI) detection in an end-to-end manner. Current approaches either decouple HOI task into separated stages of object detection and interaction classification or introduce surrogate interaction problem. In contrast, our method, named HOI Transformer, streamlines the HOI pipeline by eliminating the need for many hand-designed components. HOI Transformer reasons about the relations of objects and humans from global image context and directly predicts HOI instances in parallel. A quintuple matching loss is introduced to force HOI predictions in a unified way. Our method is conceptually much simpler and demonstrates improved accuracy. Without bells and whistles, HOI Transformer achieves 26.61% AP on HICO-DET and 52.9% AP_{role} on V-COCO, surpassing previous methods with the advantage of being much simpler. We hope our approach will serve as a simple and effective alternative for HOI tasks. Code is available at <https://github.com/bbepoch/HoiTransformer>.

1. Introduction

Human-Object Interaction (HOI) detection plays an important role in the high level human-centric scene understanding, and has attracted considerable research interest recently. The HOI research can also contribute to other tasks, such as action analysis, weakly-supervised object detection, and visual question answering, etc.

The goal of HOI detection aims at localizing human and object, as well as recognizing the interaction between them. Previous studies [7, 6, 9, 22, 19, 3] present promising results on HOI detection by decouple this task into object detection and interaction classification (Fig. 1(a)). More specifically, human and object detection results is first obtained by pre-trained object detector, then interaction classification is

*Equal contribution.

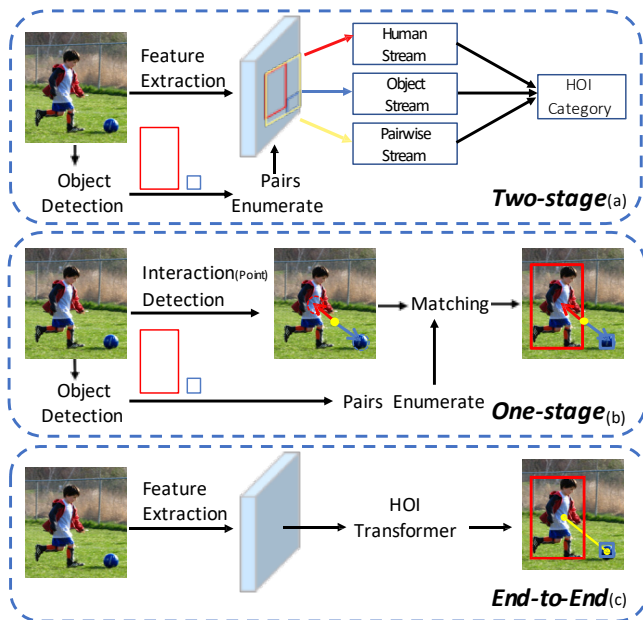


Figure 1: Comparison of recent approaches on HOI detection. (a) two-stage methods, typically use pre-trained detectors to generate human, object proposal and enumerate all (human, object) pairs, followed by a multi-stream architecture to classify interactions. (b) one-stage methods, detect interaction point/box and object proposals simultaneously, followed by complex matching process to assign interactions to object pairs. (c) our end-to-end method, input an image and predict HOI instances directly.

conducted on the pair-wisely combined human-object proposals. The limitations of these methods are mainly caused by the *separated* two stages. The independent optimization on two sub-problems may lead to sub-optimal solution. The generated human-object proposals have relatively low quality for interaction classification [19], because only object-level confidence has been taken into account. Moreover, all pair-wise proposals need to be processed, which brings

large redundant computation cost.

More recent approaches [35, 3, 19] have introduced a surrogate interaction detection problem to optimize HOI detection indirectly (Fig. 1(b)). Firstly, the interaction proposal has been pre-defined based on human priors. For example, UnionDet [3] defines the interaction proposal as union box of the human and object box. PPDM [19] uses the center point between human and object as interaction point. Secondly, the human, object and interaction proposals are detected in parallel. Finally, each interaction result is assigned to one (human, object) pair based on pre-defined matching strategy in post processing. However, such definition of interaction proposal are not always valid under different circumstance and make the pipeline more complex and costly in computation.

For HOI detection, how to capture the dependencies, especially long range, between human and object in the image space is the main problem. The above methods used complex but sub-optimal strategies, i.e. decouple into two-stages or introduce surrogate proposals to empower models the ability of capturing dependencies. However, the transformer network [32] is designed to exhaustively capture the long range dependencies, which inspire us to address the problem with transformer.

In this paper, we propose a new architecture to directly predict the HOI instance, i.e. (human, object, interaction), in an *end-to-end* manner. Our method consists of two parts, a transformer encoder-decoder architecture and a quintuple HOI matching loss. The architecture first use CNN backbone to extract high-level image features, then the encoder is leveraged to generate global memory feature, which models the relation between the image feature explicitly. Next the global memory from encoder and the HOI queries are send to decoder to generate the output embeddings. Finally, a multi-layer perception is used to predict HOI instances based on the output embeddings of decoder. Meanwhile, a quintuple HOI matching loss is proposed to supervise the learning of HOI instance prediction. Our method achieves state-of-the-art results on different challenging HOI benchmarks.

2. Related work

2.1. Two-Stage HOI Detection

Modern two-stage HOI detection methods usually consists of an object detector in the first stage and an interaction classifier in the second stage. More specifically, In the first stage, a fine-tuned object detector is used to get the human and object bounding boxes and class labels. In the second stage, a multi-stream architecture is used to predict the interactions for each human-object pair.

Typically there are three streams in the mentioned multi-stream interaction classifier: human stream, object stream,

and pairwise stream. Both human stream and object stream usually encode visual features for human and object boxes respectively [7]. In FCMNet [22], object visual feature is replaced by word embedding for the reason that detailed visual appearance of the object is often not crucial for the interaction category. Besides visual features, Bansal et. al [1] introduced word embedding in human stream for feature augmentation. PDNet [38] introduced word embedding for all the streams to get language prior-guided channel attention and feature augmentation. Plenty of researches have been done on the pairwise stream. This stream usually encodes the relationship between the human and object. A two-channel binary image representation is first advocated in iCAN [7] to encode the spatial relation, but in FCMNet [22], a fine-grained version from human parsing is proposed to amplify the key cues. Apart from spatial relation, graph neural networks in DRG [6], CHG [34], RPNN [39] were proposed to explicitly model the interactions between human and objects, which sure improved the model's representation capability.

Auxiliary models can be easily introduced to two-stage pipeline to help improving HOI, e.g. human pose feature, human body-part [17], language model [26] and graph model [37], etc. Interestingly Bansal et.al [1] and Hou et.al [14] introduced feature level augmentation, which is proved to be effective to HOI. However, these methods suffer from heavy complexity and low efficiency due to the sequential and separated two-stage architecture.

2.2. One-Stage HOI Detection

InteractNet [9] may be one of the earliest detection-based methods but it needs cascaded inference. PPDM [19] and IPNet [35] treated HOI as a point detection problem and directly detect interactions in a one stage manner by introducing a novel definition of interaction point. Further, PPDM predicts both object detection and HOI detection in a unified CenterNet-based [41] model. In UnionDet [3], HOI detection is regarded as a union box detection problem and based on the popular RetinaNet [20], another unified one stage HOI detection model is proposed, an extra union branch for detecting union box is added parallel to the conventional object detection branch.

Compared to two-stage methods, the pipeline becomes simpler, faster, more efficient and easier to deploy for real world applications. However, one-stage methods still need complex post processing to group object detection results and the interaction predictions.

2.3. End-to-End Object Detection

Russell [30] proposed an end-to-end people detection method by LSTM-based encoder-decoder, which is an autoregressive model that predicts the output sequence one element at a time. DETR [4] improved it by replacing

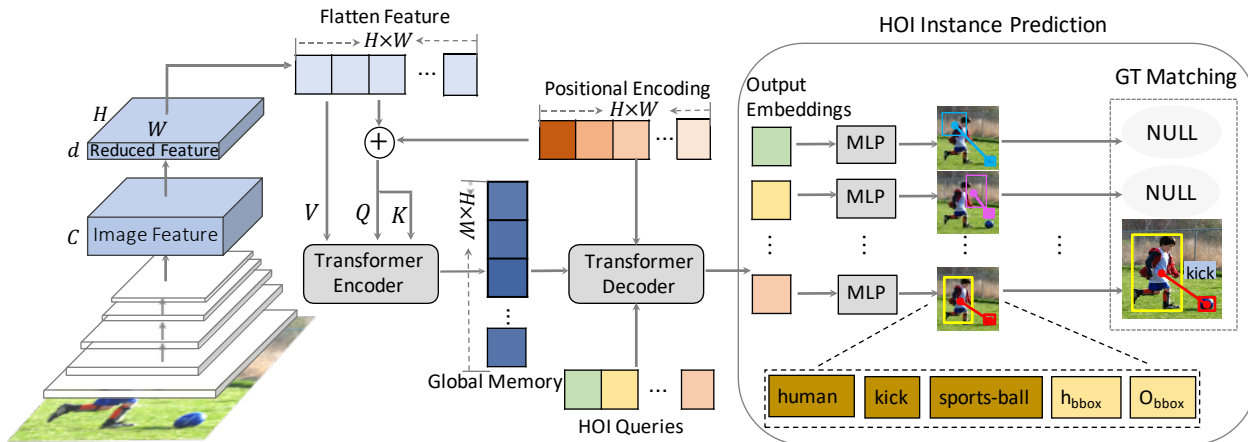


Figure 2: Overall architecture. Our model first use a CNN backbone to extract visual feature from input image. Then the feature is reduced in channel-dimension, flatten in spatial-dimension and complemented by positional encoding. The transformer encoder generates the global memory feature based on K, Q, V . The transformer decoder transforms N learnt positional embeddings (denoted as HOI queries) into N output embeddings. Finally, the multi-layer perceptron (MLP) predict the quintuple HOI instance based on the embeddings. The HOI instances are directly output simultaneously.

LSTM with transformer, which decodes N objects in parallel by leveraging the recent transformers with parallel decoding [24, 10, 8]. Both methods use Hungarian algorithm to match the ground truths and predictions though different matching costs are used. Unlike traditional object detectors, the end-to-end methods, usually have an NMS free architecture, and to make this reality, a good one-to-one matching strategy for duplicates reduction is important, and Hungarian matching seems to be a better choice so far.

3. Method

3.1. Overview

Different from previous methods, we solve human-object interaction detection in an end-to-end manner both in training and inference: input an image and then output the HOI relations directly, without any post processing. The proposed method consists of two main parts, an end-to-end transformer encoder-decoder architecture and a quintuple HOI instance matching loss.

3.2. Network Architecture

The proposed architecture illustrated in Fig. 2 consists of three main parts: (i) a backbone to extract visual feature from the input image, (ii) a transformer encoder-decoder to digest backbone feature and produce output embeddings, and (iii) a multi-layer perceptron (MLP) to predict HOI instances.

Backbone: A CNN backbone is used to extract visual feature from the input image. First, a color image is fed into the backbone and generate a feature map of shape (H, W, C)

which contains high level semantic concepts. A 1×1 convolution layer is used to reduce the channel dimension from C to d . A flatten operator is used to collapse the spatial dimension into one dimension. After that, a feature map of shape $[H \times W, d]$ is obtained, denoted as *flatten feature* in Fig. 2. The spatial dimension transformation is important because the following transformer encoder requires a sequence as input, thus the feature map can be interpreted as a sequence of length $H \times W$, and the value at each time step is a vector of size d . We use ResNet [13] as our backbone and reduce the dimension of feature conv-5 from $C = 2048$ to $d = 256$.

Encoder: The encoder layer is built upon standard transformer architecture with a multi-head self-attention module and a feed-forward network (FFN). Theoretically the transformer architecture is permutation invariant. To enable it distinguish relative position in the sequence, position encoding [25, 2] is added to the input of each attention layer. The sum of flatten feature and *positional encoding* is fed into the transformer encoder to summarize global information. The output of the encoder is denoted as *global memory* in Fig. 2.

Decoder: The decoder layer is also built upon the transformer architecture. Different from encoder layer, it contains an additional multi-head cross attention layer. The decoder transforms N learnt positional embeddings (denoted as *HOI queries* in Fig. 2) into N output embeddings. They are then decoded into HOI instances by the following MLP, which will be detailed in next section. In general, the decoder has three inputs, one is the global memory from encoder, one is HOI queries, and one is positional encoding.

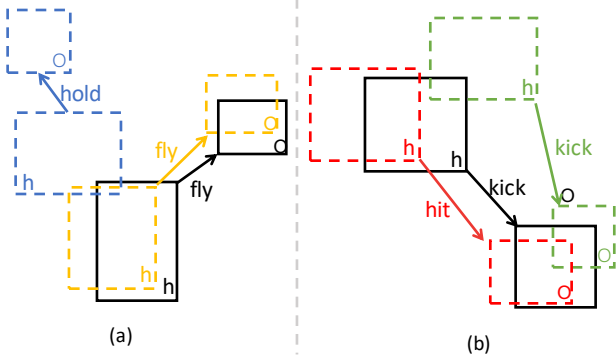


Figure 3: Illustration of the matching strategy between HOI ground-truth (black) and prediction (other colors). An HOI instance is represented by a pair of boxes in the same color. h and o represent human and object respectively.

For multi-head cross attention layer, the *value* comes from global memory directly. The *key* is the sum of global memory and the input position encoding. The *query* is the sum of input position encoding and the input HOI queries. For self-attention layer, all of the Query, Key, Value come from the HOI queries or the output of previous decoder layer. The output of the decoder is denoted as *output embeddings* in Fig. 2. This architecture design follows [4].

MLP for HOI Prediction: We define each HOI instance as a quintuple of (human class, interaction class, object class, human box, object box). The output embedding for each HOI query is decoded into one HOI instance by several multi-layer perceptron (MLP) branches. Specifically, there are three one-layer MLP branches to predict the human confidence, object confidence and interaction confidence respectively, and two three-layer MLP branches to predict human box and object box. All one-layer MLP branches for predicting confidence use a softmax function. For human confidence branch, the output size is 2, implies the confidences for foreground and background. For object confidence branch and interaction confidence branch, the output size is $C + 1$, which implies the confidences for all C kinds of objects or verbs defined in the dataset plus one for background. For both human and object box branches, the output size is 4, implies the normalized center coordinates (x_c, y_c) , height and width of the box.

3.3. HOI Instance Matching

The HOI instance is a quintuple of $(c_h, c_r, c_o, b_h, b_o)$, where (c_h, c_r, c_o) denotes human, interaction and object class confidence, (b_h, b_o) is the bounding box of the human and object. Two-stage HOI detectors first predict the object proposals $(c_h, b_h), (c_o, b_o)$ with an object detector, then enumerate the detected (human, object) pairs to predict the c_r by interaction classification. In other words, they are

trying to approximate the following probability in a given dataset,

$$p(h, r, o) = p(h, o)p(r|h, o) \approx p(h)p(o)p(r|h, o) \quad (1)$$

where $p(h)$ and $p(o)$ indicate the confidence of human and object bounding box, respectively. $p(r|h, o)$ denotes the probability of interaction r given human box h and object box o , often implemented by a multi-stream interaction recognition model. In this method, the object detector and the interaction classifier are *separately* optimized.

On the contrary, we treat HOI detection as a set prediction problem of bipartite matching between predictions and ground truth. Our method directly predicts the elements in HOI set and optimizes the proposed HOI matching loss in a unified way.

As shown in Fig. 3(a), suppose a ground truth (human, fly, object) is in the image, and the model predicts two HOI instances: the yellow one (human, fly, object), and the blue one (human, hold, object). The yellow one not only predict the interaction correctly but localize the human and object more accurately as well. To minimize the matching cost, it is more suitable to assign the black one to the yellow one, and assign \emptyset (implies nothing) to the blue one. A precise and complete matching strategy is formulated in the following.

Assume the model outputs a fixed-size set of N predictions, where N is chosen to be larger than the number of HOI relations in one image. Let us denote the set of predicted HOIs as $P = p^i, i = 1, 2, \dots, N$, the set of ground truth HOIs as $G = g^i, i = 1, 2, \dots, M, \emptyset, \dots, \emptyset$, where $M \leq N$. M represents the number of ground truth in an image. By padding \emptyset to the ground truth set, we make the length of two sets equal.

We denote the matching as an injective function: $\sigma_{G \rightarrow P}$, where $\sigma(i)$ is the index of predicted HOI assigned to the i -th groundtruth. The matching cost function is defined as:

$$\mathcal{L}_{cost} = \sum_i^N \mathcal{L}_{match}(g^i, p^{\sigma(i)}) \quad (2)$$

where $\mathcal{L}_{match}(g^i, p^{\sigma(i)})$ is a matching cost between ground truth g^i and prediction $p^{\sigma(i)}$.

In each step of training, we should first find an optimal one-to-one matching between the ground truth set and the current prediction set. We design the following matching cost for HOI:

$$\mathcal{L}_{match}(g^i, p^{\sigma(i)}) = \beta_1 \sum_{j \in h, o, r} \alpha_j \mathcal{L}_{cls}^j + \beta_2 \sum_{k \in h, o} \mathcal{L}_{box}^k \quad (3)$$

where $\mathcal{L}_{cls}^j = \mathcal{L}_{cls}(g_j^i, p_j^{\sigma(i)})$, $j \in h, o, r$ represents human, object, and interactions, g_j^i denotes the category label

of j on ground-truth g^i . We use standard softmax cross entropy loss in the paper. $\mathcal{L}_{\text{box}}^k$ is box regression loss for human box and object box, the weighted sum of GIoU [28] loss and L_1 loss are used. α and β are hyper-parameters of loss weights, which will be discussed later in ablation study.

We use Hungarian algorithm [16, 4] to solve the following problem to find a bipartite matching.

$$\hat{\sigma} = \arg \min_{\sigma \in \mathfrak{S}_N} \mathcal{L}_{\text{cost}} \quad (4)$$

where \mathfrak{S}_N denotes the one-to-one matching solution space.

After the optimal one-to-one matching between the ground truth and predictions is found, the network loss is calculated between the matched pairs, using the same loss function as Eq. 3. Although these two processes share the same formulation, the hyper-parameters of them are different theoretically and may have different optimal values. However, in practice, due to the considerable computation cost brought by large hyper-parameter search space, we made them the same, just as DETR [4] does.

Different from conventional HOI detection methods which optimize object detector and interaction classifier separately, the proposed HOI matching loss takes both the classification and localization into account. Human and object boxes will be produced simultaneously with their interactions.

4. Experiments

4.1. Experimental Setting

Datasets: We conduct experiments on HICO-DET [5] and V-COCO [11] benchmark to evaluate the proposed methods. HICO-DET consists of 47,776 images with more than 150K human-object pairs (38,118 images in training set and 9,658 in test set). It has 600 HOI categories over 117 interactions and 80 objects. Further, 600 HOI categories has been split into 138 Rare and 462 Non-Rare based on the number of training instances. V-COCO is a subset of MSCOCO [21], consists of 5,400 images in the trainval dataset and 4946 images in test set. Each human is annotated with binary labels for 29 different action categories (five of them do not involve associated objects).

Evaluation Metric: Following the standard evaluation, we use the commonly used role *mean average precision* (mAP) to examine the model performance for both datasets. An HOI detection is considered as true positive if and only if it localizes the human and object accurately (i.e. the *Interaction-over-Union* (IOU) ratio between the predicted box and ground-truth is greater than 0.5) and predict the interaction correctly.

Implementation Details

Data Augmentation: First, we adjust the brightness and contrast with a probability of 0.5 as image level augmenta-

tion. Specifically, for both brightness and contract, a parameter is randomly chosen from the range [0.8, 1.2], meaning only slight change is performed to original image. Next, we use scale augmentation, scaling the input image such that the shortest side is at least 480 and at most 800 pixels while the longest at most 1333 [36]. And also, we use random flip with a probability of 0.5. Finally, we apply random crop augmentations: an image is cropped with probability 0.5 to a random rectangular patch followed by another scale augmentation to ensure its shape, it is noteworthy that if any box in a given ground truth human-object pair is outside the cropped patch, its label will be removed.

Training Settings: The input image to the model is first scaled to [0, 1] and then normalized by channel-wise mean and std. The experiments are conducted on two popular backbone, ResNet-50 and ResNet-101. The models are trained with AdamW [23] setting the transformer’s learning rate to 1e-4, the backbone’s to 1e-5, and weight decay to 1e-4. The number of encoder layer and decoder layer are both set to 6, the number of HOI query is set to 100, and we use a COCO pre-trained DETR [4] model to initialize the weights of both backbone and transformer encoder-decoder. The batch size for ResNet-50 is set to 16 while 8 for ResNet-101. All the models are trained for 250 epochs with once learning rate decay at epoch 200. Training our network takes 7 hours on 8 NVIDIA 2080TI GPU on V-COCO and 70 hours on HICO-DET. At test times, our model runs at 24 fps on a single 2080TI GPU.

4.2. Comparisons with State-of-the-Art methods

We report the main quantitative results in terms of AP on HICO-DET in Table 1 and AP_{role} on V-COCO in Table 2.

For the HICO-DET dataset, our method compares against state-of-the-art algorithms. We achieve 4.88% point gain over one-stage methods [19] on Full categories, especially 5.37% point on Rare categories. Compared with those two-stage methods, our method achieves best performance on Full categories without introducing additional human pose feature and languages prior, which shows great potential of our method. Meanwhile, we humbly regard how to improve performance on Rare due to the long-tail distribution of HOI detection as a significant future work in our framework.

For the V-COCO dataset, our method also achieves the competitive performance compared with state-of-the-art methods without introducing external data, e.g. human pose, HICO-DET data, language prior knowledge. We obtain 1.9% point gain over previous one-stage method [35].

4.3. Ablation Study

In our ablation study, we explore how the matching strategy, loss weight, and data augmentation influence the final performance. The ablation experiments are conducted with

Methods	Feature Backbone	Detector	Pose	Language	Default			Known Object		
					Full↑	Rare↑	NonRare↑	Full↑	Rare↑	NonRare↑
<i>Two-stage methods</i>										
Shen et al. [29]	VGG-19	COCO			6.46	4.24	7.12	-	-	-
HO-RCNN [5]	CaffeNet	COCO			7.81	5.37	8.54	10.41	8.94	10.85
InteractNet [9]	ResNet-50-FPN	COCO			9.94	7.16	10.77	-	-	-
GPNN [27]	ResNet-101	COCO			13.11	9.34	14.23	-	-	-
iCAN [7]	ResNet-50	COCO			14.84	10.45	16.15	16.26	11.33	17.73
PMFNet-Base [33]	ResNet-50-FPN	COCO			14.92	11.42	15.96	18.83	15.30	19.89
PMFNet [33]	ResNet-50-FPN	COCO	✓		17.46	15.65	18.00	20.34	17.47	21.20
No-Frills [12]	ResNet-152	COCO		✓	17.18	12.17	18.68	-	-	-
TIN [18]	ResNet-50	COCO	✓		17.22	13.51	18.32	19.38	15.38	20.57
CHG [34]	ResNet-50	COCO			17.57	16.85	17.78	21.00	20.74	21.08
Peyre et al. [26]	ResNet-50-FPN	COCO		✓	19.40	14.63	20.87	-	-	-
VSGNet [31]	ResNet152	COCO			19.80	16.05	20.91	-	-	-
FCMNet [22]	ResNet-50	COCO	✓	✓	20.41	17.34	21.56	22.04	18.97	23.12
ACP [15]	ResNet-152	COCO	✓	✓	20.59	15.92	21.98	-	-	-
Bansal et al. [1]	ResNet-50-FPN	HICO-DET		✓	21.96	16.43	23.62	-	-	-
PD-Net [38]	ResNet-152	COCO		✓	20.81	15.90	22.28	24.78	18.88	26.54
PastaNet [17]	ResNet-50	COCO	✓	✓	22.65	21.17	23.09	24.53	23.00	24.99
VCL [14]	ResNet101	HICO-DET			23.63	17.21	25.55	25.98	19.12	28.03
DRG [6]	ResNet-50-FPN	HICO-DET		✓	24.53	19.47	26.04	27.98	23.11	29.43
<i>One-stage methods</i>										
UnionDet [3]	ResNet-50-FPN	HICO-DET			17.58	11.52	19.33	19.76	14.68	21.27
IPNet [35]	Hourglass	COCO			19.56	12.79	21.58	22.05	15.77	23.92
PPDM [19]	Hourglass	HICO-DET			21.73	13.78	24.10	24.58	16.65	26.84
Ours	ResNet-50	-			23.46	16.91	25.41	26.15	19.24	28.22
Ours	ResNet-101	-			26.61	19.15	28.84	29.13	20.98	31.57

Table 1: Comparison with the state-of-the-art methods on HICO-DET test set. For the Detector, COCO means that the detector is trained on COCO, while HICO-DET means that the detector is first trained on COCO and then fine-tuned on HICO-DET. Pose means that human pose feature extracted by pre-trained skeleton model. Language means that languages prior, e.g. words2vec. Our methods achieve best performance on Full categories without introducing additional human pose feature, external language prior, and object detector.

ResNet-50 backbone models, and the models are trained for 250 epochs with once learning rate decay at epoch 200. The number of encoder layer and decoder layer are both set to 6, the number of HOI query is set to 100, and the batch size is set to 16. We use a COCO pre-trained DETR model to initialize the weights of both backbone and transformer encoder-decoder.

Our baseline set all the loss weight hyper-parameters to 1.0, and uses only brightness, contrast and random flip augmentation.

Matching Strategy: Consider the situation in Fig. 3(b), black boxes indicate a ground truth pair, green boxes indicate a predicted pair, and red boxes indicate another predicted pair. For simplicity, we assume that both human and object boxes are in the right place and the object class category is correct. Then we can see, for the red pair, both human and object boxes have higher overlap with the black

ones, but its interaction prediction ‘hit’ is wrong, note that the interaction label of ground truth is ‘kick’. For the green pair, the interaction prediction ‘kick’ is right, while its human box is obviously far from the ground truth. So in this case, which pair to match with the ground truth is confused, location first or category first? We conduct ablation study to find the relative importance in matching. In Eq. 3, β_1 , β_2 dominates the weight of classification and localization respectively. As shown in Table. 3b, the best result is obtained under $\beta_1 = 2.0$, $\beta_2 = 1.0$, which reflects that classification plays a more important role than localization during the matching process.

Loss Ablation: The human/object/interaction are 2/81/117-category classification problems in HICO-DET. ‘human’ has sufficient training data because it appears in each HOI instance. Therefore, we assume human classification as the simplest one and set $\alpha_h = 1$.

Methods	Backbone	Pose	Language	AProle
<i>Two-stage methods</i>				
VSRL [11]	ResNet-50-FPN			31.8
InteractNet [9]	ResNet-50-FPN			40.0
GPNN [27]	ResNet-101			44.0
RPNN [39]	ResNet50	✓		47.5
VCL [14]	ResNet101			48.3
TIN* [18]	ResNet-50	✓		48.7
Zhou et al. [40]	ResNet-50	✓		48.9
PastaNet [17]	ResNet-50	✓	✓	51.0
DRG [6]	ResNet-50-FPN		✓	51.0
VSGNet [31]	ResNet-152			51.8
CHG [34]	ResNet-50			52.7
PMFNet [33]	ResNet-50-FPN	✓		52.0
PD-Net [38]	ResNet-152	✓		52.6
FCMNet [22]	ResNet-50	✓	✓	53.1
ACP* [15]	ResNet-152	✓	✓	53.2
<i>One-stage methods</i>				
UnionDet [3]	ResNet-50-FPN			47.5
IPNet [35]	Hourglass-104			51.0
IPNet* [35]	Hourglass-104			52.3
Ours	ResNet-101			52.9

Table 2: Comparisons of the state-of-the-art on V-COCO test set. Pose denotes whether human skeleton feature has been introduced and Language denotes external-language prior. Character * indicates that HICO-DET training data was incorporated into training data.

We conduct experiments to evaluate the relative importance of ‘object’ and ‘interaction’ in our experiments. In Eq. 3, α_o and α_r dominate the weight of object and interaction respectively in training loss. As shown in Table. 3a, our method obtains best result when $\alpha_r = 2.0$ and $\alpha_o = 1.0$, indicating that interaction tends to be more important than object in our framework.

Data Augmentation: We mainly study two kinds of data augmentation in our experiments: multi-scale training and random crop. We conduct ablation experiments on combination of them, results can be found in Table. 3c. Considerable improvements have been made, multi-scale training attains 4.29% point gain on Full categories and random crop achieves 5.08%, and the combination of them gets even better results, mainly because these two augmentations help the attention layers to learn scale-invariant and shift-invariant features much easier on a small dataset.

4.4. Discussion

We formulate HOI detection as a set prediction problem and directly predict instances in an end-to-end manner. So, compared with previous methods, which are trained on HOI instance level, treating human-object pair as data point, our

α_r	α_o	Full↑	Rare↑	NonRare↑
1.0	1.0	16.97	11.65	18.55
1.0	2.0	17.49	11.47	19.29
2.0	1.0	17.85	13.54	19.14

(a) Varying α for classification loss

β_1	β_2	Full↑	Rare↑	NonRare↑
1.0	1.0	16.97	11.65	18.55
1.0	2.0	16.20	9.54	18.19
2.0	1.0	18.46	13.41	19.97

(b) Varying β for match cost

scale	crop	Full↑	Rare↑	NonRare↑
		16.97	11.65	18.55
	✓	22.05	15.20	24.09
✓		21.26	14.61	23.25
✓	✓	22.36	15.47	24.42

(c) Data augmentation during training process

Table 3: Ablation experiments for HOI Transformer. All models use ResNet-50 backbone to extract feature; number of queries is set to 100; batch size is set for 16; trained for 250 epochs; learning rate decay from 1e-4 to 1e-5 at epoch 200.



Figure 4: Out of distribution generalization for HOI prediction. Even though no images in the training set has more than 3 ‘drink with bottle’ HOIs, our method generalize well on the synthetic image with 6 of them.

method regards a whole image as data point. There seems some differences. Here we briefly discuss our understanding on two aspects.

Image-level data point. First, the training loss for interaction in our method is the sum of classification loss between each matched (pred, GT) pair. It is relevant to the number of positive samples, which is consistent with prior instance-level work. Second, random crop augmentation provides sufficient training instances. For example, let’s think random crop in an extreme way, each time only one positive sample is cropped, then it is close to prior work.

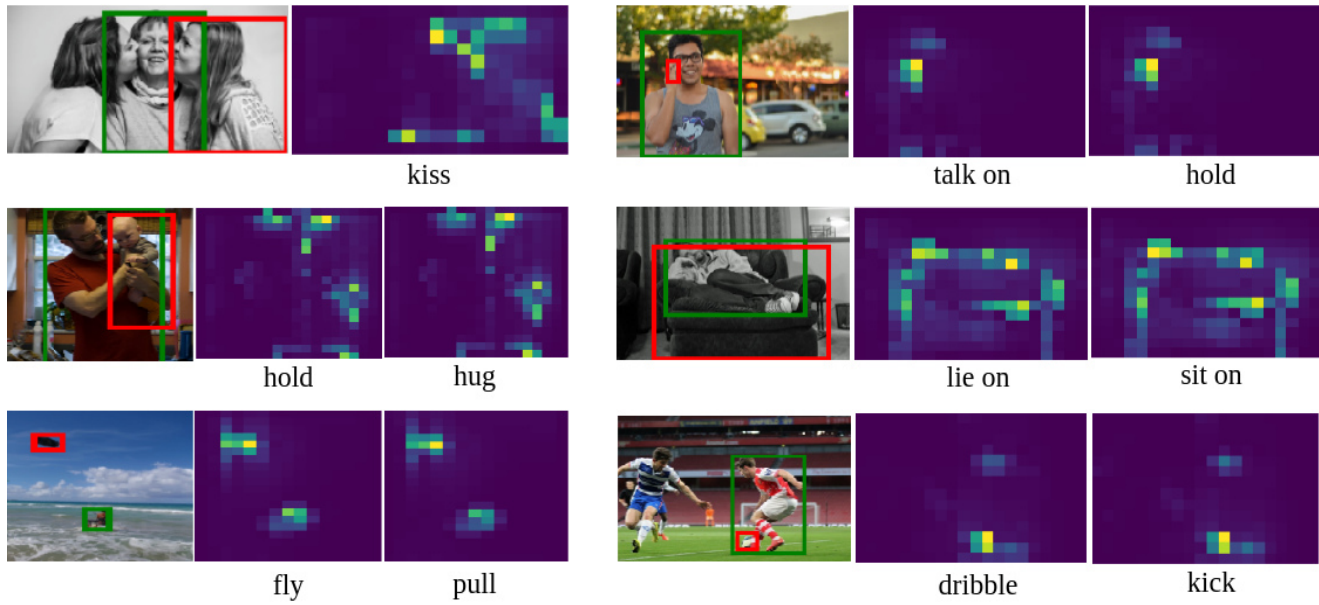


Figure 5: Visualization of attention map in decoder for predicted HOI instance (images from HICO-DET dataset). As can be seen from the figure, our method attends to the discriminative part (e.g. telephone for *talk on*, human face for *kiss*) and can capture long-distance interaction (e.g. *fly* and *pull*). Moreover, it can be seen from the figure that different interaction categories share little common pattern, which suggests that a unified definition for all interaction proposal may be sub-optimal.

Out of distribution test. Noting that in training data there are at most 3 ‘drink with bottle’ HOIs in a single image, we create a synthetic image containing 6 ‘drink with bottle’ HOIs, which is out of distribution. As shown in Fig. 4, the result indicates that the model can learn to recognize HOIs with good generalization.

4.5. Qualitative Analysis

As can be seen from Fig. 5, we visualize the decoder attention map for predicted HOI instances. The interaction heatmap highlights both the human and object area, meaning that our model reasons about the relations between human and object from a more global image context, not focusing on human or object only. It is obvious that decoder has ability to find the discriminative part for the interaction category. The model can predict different instances based on similar attention heatmap, which implies that, the MLP of the model have the ability to tell from fine-grained interaction features. Meanwhile, some local area with relatively higher attention may indicate the localization (boundaries) of human or object, because the visualized attention map is immediately followed by the MLP head for classification as well as regression. Moreover, it can be seen from the figure that different interaction categories share little common pattern, which suggests that the empirically unified definition of interaction proposal in one-stage methods, i.e. interaction point/box, is sub-optimal. And thanks to the large re-

ceptive field of attention layers, our model can easily handle long distance interaction, e.g. fly kite.

5. Conclusion

In this paper, we propose a novel HOI Transformer to directly predict the HOI instances in an *end-to-end* manner. Our core idea is to build a transformer encoder-decoder architecture to directly predict HOI instances, and a quintuple matching loss for HOI to enable supervision in a unified way. We validate the proposed method on two challenging HOI benchmark and achieve a considerable performance boost over state-of-the-art results. It is worth noting that our method not only abandon the additional features, but desert the complex post processing as well. Moreover, based on the attention map of the decoder, we found that our model has ability to dynamically attains the discerning feature for different HOI queries. We hope our method will be useful for human activity understanding research.

Acknowledgement

This research is supported by China’s ‘scientific and technological innovation 2030 - major projects’ (No. 2020AAA0104400).

References

- [1] Ankan Bansal, Sai Saketh Rambhatla, Abhinav Shrivastava, and Rama Chellappa. Detecting human-object interactions via functional generalization. In *AAAI*, pages 10460–10469, 2020. [2](#), [6](#)
- [2] Irwan Bello, Barret Zoph, Ashish Vaswani, Jonathon Shlens, and Quoc V Le. Attention augmented convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3286–3295, 2019. [3](#)
- [3] Kim Bumsu, Choi Taeho, Kang Jaewoo, and J. Kim Hyunwoo. Uniondet: Union-level detector towards real-time human-object interaction detection. In *European Conference on Computer Vision*. Springer, 2020. [1](#), [2](#), [6](#), [7](#)
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. *arXiv preprint arXiv:2005.12872*, 2020. [2](#), [4](#), [5](#)
- [5] Yu-Wei Chao, Yunfan Liu, Xiayang 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. [5](#), [6](#)
- [6] Chen Gao, Jiarui Xu, Yuliang Zou, and Jia-Bin Huang. Drg: Dual relation graph for human-object interaction detection. In *European Conference on Computer Vision*, pages 696–712. Springer, 2020. [1](#), [2](#), [6](#), [7](#)
- [7] Chen Gao, Yuliang Zou, and Jia-Bin Huang. ican: Instance-centric attention network for human-object interaction detection. *arXiv preprint arXiv:1808.10437*, 2018. [1](#), [2](#), [6](#)
- [8] Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. Mask-predict: Parallel decoding of conditional masked language models. *arXiv preprint arXiv:1904.09324*, 2019. [3](#)
- [9] Georgia Gkioxari, Ross Girshick, Piotr Dollár, and Kaiming He. Detecting and recognizing human-object interactions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8359–8367, 2018. [1](#), [2](#), [6](#), [7](#)
- [10] Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. Non-autoregressive neural machine translation. *arXiv preprint arXiv:1711.02281*, 2017. [3](#)
- [11] Saurabh Gupta and Jitendra Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015. [5](#), [7](#)
- [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 International Conference on Computer Vision*, pages 9677–9685, 2019. [6](#)
- [13] 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. [3](#)
- [14] Zhi Hou, Xiaojiang Peng, Yu Qiao, and Dacheng Tao. Visual compositional learning for human-object interaction detection. *arXiv preprint arXiv:2007.12407*, 2020. [2](#), [6](#), [7](#)
- [15] Dong-Jin Kim, Xiao Sun, Jinsoo Choi, Stephen Lin, and In So Kweon. Detecting human-object interactions with action co-occurrence priors. *arXiv preprint arXiv:2007.08728*, 2020. [6](#), [7](#)
- [16] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. [5](#)
- [17] Yong-Lu Li, Liang Xu, Xinpeng Liu, Xijie Huang, Yue Xu, Shiyi Wang, Hao-Shu Fang, Ze Ma, Mingyang Chen, and Cewu Lu. Pastanet: Toward human activity knowledge engine. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 382–391, 2020. [2](#), [6](#), [7](#)
- [18] Yong-Lu Li, Siyuan Zhou, Xijie Huang, Liang Xu, Ze Ma, Hao-Shu Fang, Yanfeng Wang, and Cewu Lu. Transferable interactivity knowledge for human-object interaction detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3585–3594, 2019. [6](#), [7](#)
- [19] Yue Liao, Si Liu, Fei Wang, Yanjie Chen, Chen Qian, and Jia-ashi 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](#), [5](#), [6](#)
- [20] 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. [2](#)
- [21] 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 *European conference on computer vision*, pages 740–755. Springer, 2014. [5](#)
- [22] Yang Liu, Qingchao Chen, and Andrew Zisserman. Amplifying key cues for human-object-interaction detection. In *European Conference on Computer Vision*, pages 248–265. Springer, 2020. [1](#), [2](#), [6](#), [7](#)
- [23] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [5](#)
- [24] Aaron Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, George Driessche, Edward Lockhart, Luis Cobo, Florian Stimberg, et al. Parallel wavenet: Fast high-fidelity speech synthesis. In *International conference on machine learning*, pages 3918–3926. PMLR, 2018. [3](#)
- [25] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Łukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. *arXiv preprint arXiv:1802.05751*, 2018. [3](#)
- [26] Julia Peyre, Ivan Laptev, Cordelia Schmid, and Josef Sivic. Detecting unseen visual relations using analogies. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1981–1990, 2019. [2](#), [6](#)
- [27] 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. [6](#), [7](#)
- [28] Hamid Rezaatofghi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE Conference on Com-*

- puter Vision and Pattern Recognition*, pages 658–666, 2019. 5
- [29] Liyue Shen, Serena Yeung, Judy Hoffman, Greg Mori, and Li Fei-Fei. Scaling human-object interaction recognition through zero-shot learning. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1568–1576. IEEE, 2018. 6
 - [30] Russell Stewart, Mykhaylo Andriluka, and Andrew Y Ng. End-to-end people detection in crowded scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2325–2333, 2016. 2
 - [31] Ulutan and Iftekhar. Vsgnet: Spatial attention network for detecting human object interactions using graph convolutions. *arXiv: Computer Vision and Pattern Recognition*, 2020. 6, 7
 - [32] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 2
 - [33] Bo Wan, Desen Zhou, Yongfei Liu, Rongjie Li, and Xuming He. Pose-aware multi-level feature network for human object interaction detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9469–9478, 2019. 6, 7
 - [34] Hai Wang, Wei-shi Zheng, and Ling Yingbiao. Contextual heterogeneous graph network for human-object interaction detection. *arXiv preprint arXiv:2010.10001*, 2020. 2, 6, 7
 - [35] 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. 2, 5, 6, 7
 - [36] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 5
 - [37] Bingjie Xu, Yongkang Wong, Junnan Li, Qi Zhao, and Mohan S Kankanhalli. Learning to detect human-object interactions with knowledge. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
 - [38] Xubin Zhong, Changxing Ding, Xian Qu, and Dacheng Tao. Polysemy deciphering network for human-object interaction detection. 2, 6, 7
 - [39] Penghao Zhou and Mingmin Chi. Relation parsing neural network for human-object interaction detection. pages 843–851, 2019. 2, 7
 - [40] Tianfei Zhou, Wenguan Wang, Siyuan Qi, Haibin Ling, and Jianbing Shen. Cascaded human-object interaction recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4263–4272, 2020. 7
 - [41] Xingyi Zhou and Dequan Wang. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 2